



2026 REPORT

The State of the H-1B Program

Who files, what they pay, and who bears the cost.

Findings from 4.2 million labor condition applications, fiscal years 2020 to 2026.



Josh C. Simmons, PhD

Founder, Meridian Strategic Systems · Queen Creek, Arizona
hello@meridianstrategic.systems

Published September 21, 2026 · Data snapshot
2026-09-19

DOL OFLC and PERM disclosure files · USCIS H-1B
Employer Data Hub · BLS OEWS · SEC filings

Prepared for the agencies with jurisdiction over national-origin and citizenship-status discrimination and over labor condition application wage compliance. Every number in this report was produced by a query against the public record and reproduced by an independent second query. Queries, code, and the pre-registration record are in the appendices.

The State of the H-1B Program: 2026 Report

Who files, what they pay, and who bears the cost.

Findings from 4.2 million labor condition applications, FY2020 to FY2026.

Josh C. Simmons, PhD

Meridian Strategic Systems, Queen Creek, Arizona

hello@meridianstrategic.systems

Working paper, September 21, 2026. Data snapshot 2026-09-19.

Abstract. The H-1B program admits foreign workers on an employer’s attestation that their employment will not adversely affect the wages and working conditions of U.S. workers. This report asks what the program’s own administrative record says about how it is used. We analyze every labor condition application disclosed by the Department of Labor for fiscal years 2020 through 2026 (4.56 million decision records, 4.22 million certified), the PERM green-card certification file, USCIS employer-level petition outcomes, and BLS wage benchmarks, joined by employer and corporate family. We then score the job-title histogram of each of 361,353 employer-months against 18 semantic predicates with an evaluation model and enter the scores into two-way-clustered generalized linear models with a fixed structured baseline, corrected for multiple testing, replicated on a chronological validation split, and confirmed on a sealed, pre-registered holdout. Filing is concentrated: 100 of 205,032 employers account for 32.9% of certified applications, and the largest single filer is an IT outsourcing firm. The program is priced at the bottom of its own wage scale: 56.6% of certified applications are at the two lowest of four prevailing-wage levels, and the median Level I offer for software developers is \$89,253 against a \$130,160 national median. The outsourcing model is a temporary-labor model: Cognizant files 35 applications per approved green-card certification against 4.2 at Microsoft. The composition is shifting: employer-months that read as staffing operations fell from 67% of applications in 2019 to 32% in 2026 while AI-dominated employer-months rose nine-fold, and 39 pre-registered title-based effects replicate with the same sign in all three periods. We also document a prior machine-generated screening design that tested 1,342 hypotheses on 300-observation samples and found nothing, and diagnose why. The report closes with a targeting framework for agencies enforcing citizenship-status and national-origin discrimination law and LCA wage compliance.

Keywords: H-1B, labor condition applications, prevailing wage, outsourcing, employer concentration, citizenship-status discrimination, text-as-data, pre-registration.

1 Introduction

The Immigration and Nationality Act conditions H-1B admission on an employer attestation, filed with the Department of Labor as a labor condition application (LCA), that the foreign worker will be paid at least the prevailing wage for the occupation and area of employment and that the employment will not adversely affect the working conditions of U.S. workers similarly employed. The attestation is the program’s only wage-protection mechanism. It is not audited before certification; the Department of Labor certifies an LCA unless it is incomplete or obviously inaccurate, and the file it later discloses is the only public record of what employers said they would pay, for what job, and where.

That disclosure file is large, current, and rarely read as a whole. This paper reads it as a whole. The question is descriptive before it is causal: who uses the program, at what wage level, in what occupations, at what scale, with what persistence, and how has that changed since fiscal year 2020. Where the record permits inference beyond description, the paper tests it with a design built to fail honestly: chronological splits, corrections for multiple testing, a negative control, an out-of-time predictive test, and a holdout that was opened once, for a list of hypotheses registered before it was scored.

The statutory structure is short. Section 212(n) of the Immigration and Nationality Act requires the employer to attest that it will pay the greater of the actual wage paid to similarly qualified employees or the prevailing wage for the occupational

classification in the area of employment, and leaves “prevailing wage” undefined (*20 C.F.R. Part 655, Subpart H: Labor Condition Applications and Requirements for Employers Seeking to Employ Nonimmigrants on H-1b Visas in Specialty Occupations and as Fashion Models*, n.d.; Congressional Research Service, 2025; *Immigration and Nationality Act Section 212(N)*, 8 U.S.C. 1182(N): *Labor Condition Application*, 2026). The Department of Labor fills the gap with four wage levels drawn from the Occupational Employment and Wage Statistics survey, the lowest two of which sit below the occupational median by construction (Prevailing Wage Determination Policy Guidance: Nonagricultural Immigration Programs, 2009; U.S. Bureau of Labor Statistics, 2026). The Department’s Inspector General described the certification process in 2020 as “simply rubber-stamping”: employers file no supporting documentation, the agency lacks authority to validate attestations, and 96% of fiscal 2019 applications required no recruitment of U.S. workers at all (U.S. Department of Labor, Office of Inspector General, 2020). The Government Accountability Office reached the same conclusion a decade earlier, and noted that fewer than 1% of employers accounted for more than a quarter of approvals and that no provision holds an employer accountable when it obtains H-1B workers through a staffing company (U.S. Government Accountability Office, 2011). An attempt to raise the wage levels by interim final rule in October 2020 was vacated in 2021 and the pre-2020 methodology restored (*Strengthening Wage Protections for the Temporary and Per-*

manent Employment of Certain Aliens in the United States, 2020; *Strengthening Wage Protections for the Temporary and Permanent Employment of Certain Immigrants and Non-Immigrants in the United States, Implementation of Vacatur*, 2021); the Department’s own economic analysis scored the abandoned change as a transfer of roughly \$155.7 billion over ten years (U.S. Government Accountability Office, 2021), which is a direct estimate of what the current levels leave with employers. In March 2026 the Department proposed revising the wage-level computation again, stating in the notice that the current levels create an incentive to use the program “to replace rather than supplement” U.S. workers by employing lower-paid foreign workers (*Improving Wage Protections for the Temporary and Permanent Employment of Certain Foreign Nationals in the United States*, 2026). A September 2025 proclamation attached a one-time \$100,000 payment to new petitions for beneficiaries abroad (*Restriction on Entry of Certain Nonimmigrant Workers*, 2025); the implementing guidance was vacated by the District of Massachusetts in June 2026 (*State of California v. Mullin*, 2026), and the wage-weighted cap lottery that gives a Level IV registration four entries and a Level I registration one took effect for the fiscal 2027 season (*Weighted Selection Process for Registrants and Petitioners Seeking to File Cap-Subject H-1b Petitions*, 2025).

The empirical literature on what the program does to U.S. workers is contested, and the paper cites both sides. Audits of the disclosure data find that 60% of fiscal 2019 certifications were at the two below-median wage levels, with the same split at the thirty largest employers (Costa & Hira, 2020), and that the largest cap-subject employers hired tens of thousands of new H-1B workers in the same year they announced tens of thousands of layoffs (Costa & Hira, 2023). Lottery-identified firm-level evidence finds that an additional H-1B visa crowds out roughly 1.5 other workers at the firm with at most modest effects on patenting (Doran et al., 2022); a Richmond Fed working paper using matched employer-employee records finds no net displacement of native college graduates at the firm, concentrated in small high-productivity firms (Mahajan et al., 2025). Skill-group designs find that immigration lowered relative wages of native STEM graduates (Turner, 2022) and that the 1990s IT inflow held computer-science wages 2.6% to 5.1% below what they would otherwise have been, with the gains going to firm profits and consumers (Bound et al., 2018); city-level designs find wage gains for native college graduates from STEM inflows (Peri et al., 2015), and cap restrictions are followed by offshoring rather than native hiring at the most globalized firms (Glennon, 2024; Mayda et al., 2018). A 2026 NBER working paper estimates that H-1B workers earn 16% less than comparable natives (Borjas, 2026); an earlier compensation-survey study found a visa premium in the pre-outsourcing era (Mithas & Lucas, 2010). The replacement episodes at Southern California Edison and Disney, in which U.S. employees trained the contractor-supplied H-1B workers who replaced them at a 40% to 50% wage discount, were found by the Department of Labor to involve no violation (Hira, 2016; Preston, 2015; Thibodeau, 2015), which is the clearest statement in the record that the pattern this paper measures is a design feature rather than a compliance failure (Hira, 2010).

Three features distinguish this analysis from the compliance snapshots and advocacy reports that usually draw on the same files. First, it uses the entire disclosure record for seven fiscal years, joined across the LCA, PERM, and USCIS files by normalized employer name and by corporate family, rather than a single year or a single employer. Second, it reads the text of the record as well as its numeric fields. The job-title histogram an employer files each month is a statement about what it is hiring for, and it turns out to be a stable, machine-readable signal of how the employer uses the program. We score 361,353 employer-months against 18 title-based predicates with an evaluation model and treat those scores as measured covariates in otherwise conventional panel regressions. Third, it reports its own failed design. An earlier version of this project, described in Section 4.6, generated and tested 1,342 hypotheses by machine and found nothing, for reasons that are fully diagnosable and that we think are worth stating plainly because they are the reasons most automated “AI discovery” pipelines fail.

The audience is deliberately narrow. The findings are written for the agencies that can act on them: the Equal Employment Opportunity Commission, which enforces Title VII’s prohibition on national-origin discrimination (U.S. Equal Employment Opportunity Commission, 2016; *Unlawful Employment Practices*, 1964); the Department of Justice’s Immigrant and Employee Rights Section, which enforces the citizenship-status discrimination provisions of 8 U.S.C. § 1324b (*Unfair Immigration-Related Employment Practices*, 1986) and has already reached settlements with large employers over recruitment practices that favored visa holders over U.S. workers (U.S. Department of Justice, Office of Public Affairs, 2021; 2023); and the Department of Labor’s Office of Foreign Labor Certification and Wage and Hour Division, which certify and enforce the LCA wage attestation. Section 8 translates the findings into targeting rules those offices could apply to the same public data.

A note on what the record measures. An LCA is an employer’s application for permission to employ one or more workers in H-1B, H-1B1, or E-3 status. It is not a visa, not a hire, and not a layoff. A certified LCA may never be used. PERM records are applications for permanent (green-card) labor certification, a separate pipeline with its own recruitment test. USCIS employer records count petitions approved and denied. Every count in this paper is a count of filings, and the text says so where it matters.

2 Summary of findings

Every number below was produced by a query against the snapshot and reproduced by an independent second query. Section references give the supporting detail.

1. A third of H-1B filing comes from 100 employers.

Of 4,215,630 certified LCAs, the 10 largest filers account for 15.6%, the 100 largest for 32.9%, and the 500 largest for 46.8%, out of 205,032 filers. Concentration by corporate family is higher still. (Section 5)

2. The largest single filer is an IT outsourcing firm.

Cognizant Technology Solutions US Corp: 109,360 certified LCAs, ahead of Amazon.com Services (106,408), Google (77,182), EY (71,800), and Microsoft (69,103). By corporate family, the Amazon entities lead with 158,795 and

Cognizant is second with 110,470. Six of the top twelve families are labor-supply firms.

3. **The largest filer files internal grade codes as job titles, ten thousand times a year.** “Senior Systems Analyst JC60” was certified 10,032 times in FY2020 and 6,889 in FY2021; “Manager JC50” 8,285 and 6,487 times. No other employer-title-year cell exceeds 6,800.
4. **The outsourcing model is a temporary-labor model.** Certified LCAs per certified PERM green-card labor certification, FY2020 to FY2026: Tech Mahindra 45.4, Cognizant 35.1, Ernst & Young 26.2, Wipro 18.5, Deloitte 15.5, Accenture 13.8, HCL 13.4, Infosys 12.0; versus Oracle 2.4, Intel 4.0, Microsoft 4.2, Nvidia 4.3, Cisco 4.8, Google 5.8, Apple 5.9, Amazon 6.4.
5. **The program is priced at the bottom of its own wage scale.** 56.6% of certified applications (2,384,880) are at DOL wage Levels I or II, both defined below the occupational median. The share was 60.9% in FY2020 and 57.3% in FY2025; the Level I share alone rose from 13.3% to 18.8%.
6. **Software developers are one in three filings, and the Level I median offer is 31% below the national median.** SOC 15-1252 and 15-1132: 1,373,338 certified LCAs (32.6%). For 15-1252 in FY2023 to FY2024: Level I median \$89,253 (n = 43,611; 87.8% below the \$130,160 BLS median), Level II \$108,992 (n = 172,627; 75.3% below), Level III \$140,000, Level IV \$161,080.
7. **Employers with staffing-agency names file at the bottom levels two-thirds of the time.** 22,455 name-matched employers, 755,299 certified LCAs (17.9%), 66.8% at Level I or II versus 54.3% for all others; mean offered wage \$186,790 versus \$220,313.
8. **The named outsourcers’ share of the program has fallen by more than half; the program has not shrunk.** The eight large IT outsourcers filed 13.2% of certified LCAs in FY2020 and 5.3% in FY2026; sixteen outsourcing and consulting families, 19.6% and 9.7%. Certified volume rose from 545,621 (FY2020) to 666,228 (first three quarters of FY2026).
9. **The green-card wait has more than tripled.** Median days from PERM receipt to decision by fiscal year of decision: 137, 194, 189, 287, 388, 482 (FY2020 to FY2025), decided cases only.
10. **Denial rates are near zero, with named exceptions.** LCA denials run 0.47% to 0.79% a year and PERM denials 1.8% to 5.5%. Wipro Limited’s PERM decisions were 30.9% denied (574 of 1,856) while JPMorgan, SAP, Bank of America, Mastercard, and Rivian show 0% at comparable volume; HCL America Solutions Inc had 19.5% of 1,443 LCA decisions denied; USCIS petition denials (through FY2023) were 42.1% at Syntel and 15.2% at Cognizant on 27,900 decisions.
11. **FY2026 Q3 was the largest quarter on record.** 401,412 certified LCAs against 224,127 in FY2025 Q3 and the prior record of 381,037 in FY2021 Q3, with growth across Amazon (5,237 to 13,922), Google (1,279 to 7,423), Microsoft, EY, Apple, and Meta.
12. **Filing is moving to Texas.** Texas gained 28,498 certified worksite filings between FY2020 and FY2025 (55,863 to

84,361), the largest gain of any state; California lost 6,262. New York City is the largest worksite city (245,882).

13. **Read from the titles alone, the outsourcing share of applications fell from 67% to 32%.** Employer-months whose title histograms read as staffing, consulting, or outsourcing operations: 67.0% of applications in late 2019, 56.6% (2020), 49.4%, 43.7%, 41.6%, 37.2%, 34.6%, and 31.8% in the first half of 2026; median offered wage \$98,332 in those months versus \$112,185. The decline holds unweighted (45% to 27%), excluding the ten largest filers (57% to 37%), and excluding the two known false positives (66% to 37%). (Section 6)
14. **AI titles rose nine-fold, and Walmart is the second-largest AI filer.** AI-dominated employer-months went from 0.9% of applications (2019) to 8.2% (2026), at a median wage of \$127,213 versus \$103,074. By raw title count: Apple 3,279, Wal-Mart Associates 2,309 (9.4% of its filings, median \$160,000), Amazon.com Services 2,242, LinkedIn 1,837, Google 1,235, Microsoft 1,097.
15. **Filers stopped hiring one kind of worker.** Single-job-family employer-months fell from 53.1% of applications to 26.2%; months spanning three or more job families rose from 26.5% to 51.8%; generic titles fell from 20.8% to 10.4%; management-heavy months from 22.0% to 13.9%; academic and medical months rose from 6.7% to 11.3%.
16. **Financial firms hire H-1Bs at the bottom of the ladder.** Share of certified titles carrying an entry-level token: Goldman Sachs 67.0%, Citibank 67.0%, BlackRock 64.5%, McKinsey 54.0%, JPMorgan Chase 43.5% (top title “associate, software engineer,” 2,743), versus Google 4.9%.
17. **“Senior title, modest wage” is measurable and has names.** 3.3% of applications fall in employer-months where senior or specialist titles dominate at a modest wage (median \$82,000 versus \$105,040): Cognizant 14,606 such applications, TCS 10,508, Wipro 7,711 (58% of its filings), Deloitte & Touche 3,962. Those months are followed by the largest subsequent wage increases of any flag, in all three time periods.
18. **Universities pay H-1B holders half of what technology pays.** Median certified wage: University of Michigan \$66,837, Johns Hopkins \$71,448, Yale \$75,000, Mayo Clinic \$75,684, Columbia \$78,000, Stanford \$82,000; Google \$167,000. Academic employer-months are the steadiest filers and the least likely to contract.
19. **The title-based effects replicate across three time periods.** Of 270 full-panel regressions, 117 clear $q < 0.05$ and 78 clear $q < 0.005$ in discovery; 112 of 117 keep their sign in validation, 74 replicate at $p < 0.05$, and 39 were pre-registered for the sealed 2025-2026 holdout, where 39 of 39 keep their sign and 32 are significant at $p < 0.05$. The prior design (1,342 hypotheses on 300-observation samples) replicated none. (Section 4.6, Section 6)

3 Data

3.1 Sources and snapshot

All data come from an immutable snapshot taken on September 19, 2026 of eleven public tables loaded into Google BigQuery (dataset `h1b_research`, 8.34 GiB logical). The tables and their coverage are summarized in Table 1. The LCA disclosure file

Table	Rows	What a row is
lca_disclosure_rows	4,562,922	One LCA decision record (a case may appear in more than one quarterly file; 3,754,962 distinct case numbers)
lca_worksite_rows		One worksite line per case, with worker count
lca_appendix_a_rows	2,714	Exempt-worker appendix lines
perm_disclosure_rows	684,916	One PERM labor certification decision
uscis_h1b_employer_records		Employer-year approvals and denials, FY2020-FY2023
bls_oes_wages		Occupation by area median wage, May 2023
sec_employee_headcount		Company-year total employees, 8 companies
companies		Corporate-family key (company_slug) for name variants

Table 1: Snapshot tables used. Employer identity in the LCA and PERM files is an exact normalized legal name; the **company_slug** key groups name variants and subsidiaries where a mapping exists (for example, 77 Amazon filing entities).

is the Department of Labor’s quarterly release of LCA decision records built from Form ETA-9035, keyed to the determination date (*Labor Condition Application for Nonimmigrant Workers, Form ETA-9035 and 9035e*, 2026; U.S. Department of Labor, Employment and Training Administration, Office of Foreign Labor Certification, 2026a; 2026b); the PERM file is built from Form ETA-9089 (U.S. Department of Labor, Employment and Training Administration, Office of Foreign Labor Certification, 2026c); the USCIS Employer Data Hub reports first decisions by employer and fiscal year (U.S. Citizenship and Immigration Services, 2025; 2026); BLS wages follow the OEWS methodology (U.S. Bureau of Labor Statistics, 2026); and SEC headcounts come from the post-2020 human-capital disclosure, which prescribes no standard metric and is therefore used only illustratively (*Modernization of Regulation s-K Items 101, 103, And 105*, 2020). The LCA snapshot the snapshot spans fiscal years 2020 through the third quarter of fiscal 2026 (decisions through June 30, 2026). The PERM disclosure file covers fiscal 2020 through 2025. The USCIS H-1B Employer Data Hub covers petition approvals and denials by employer and fiscal year and ends at FY2023 in this snapshot. BLS Occupational Employment and Wage Statistics (May 2023) provide occupational median wages for three metropolitan areas and the nation for six SOC codes. SEC 10-K human-capital disclosures provide total employee counts for eight companies.

Employer identity deserves emphasis. The disclosure files identify employers by legal name as typed on the form. “Meta Platforms, Inc” and “Meta Platforms, Inc.” are two employers in the raw file (18,275 and 13,516 certified records). Amazon files through at least 77 legal entities. Every ranking by legal entity in this paper is therefore a floor for large companies, and we report corporate-family totals alongside entity totals wherever the distinction changes a conclusion.

3.2 The employer-by-month panel

For the panel analyses we aggregate the LCA record into employer-months: for each normalized employer and each calendar month of decision, the number of applications, the median offered wage, the distribution of job titles, the distribution of SOC codes, the distribution of worksite states, and derived quantities (log application counts at one- and six-month windows, three-month growth, Shannon entropy of the title and location distributions (Shannon, 1948), unique title and location counts, and the share of titles in technical SOC families). Employers enter the panel if they have at least 24 active months in the discovery period, which yields 4,536 employers and 361,353 employer-months from October 2019 through June

2026. This is the population of sustained filers, not the full universe of 210,209 employer names; the descriptive analyses in Section 5 use the full universe.

Five outcomes are defined at horizons of one, two, and three months ahead: the application count (a count), an indicator that applications contracted relative to the prior window (binary), the change in log median wage, the change in title-distribution entropy (role diversity), and the change in location-distribution entropy (location diversity). The two diversity outcomes are defined only for months with positive subsequent activity.

The panel is partitioned chronologically, never randomly (Bergmeir & Benítez, 2012; Cerqueira et al., 2020; Hyndman & Athanasopoulos, 2021): discovery, October 2019 through September 2023 (211,665 employer-months); validation, October 2023 through January 2025 (72,576); holdout, February 2025 through June 2026 (77,112). Outcome windows that cross a partition boundary are purged from the earlier partition so that no label information leaks backward (Kaufman et al., 2012; López de Prado, 2018). The holdout was sealed at construction and was opened exactly once, in the step described in Section 4.4.

4 Methods

Nothing in this section is new. The contribution is the assembly and the discipline, not the parts. Each part is cited to its canonical source so that a reader can check that we did the standard thing.

4.1 Descriptive analysis of the raw record

The descriptive findings are aggregating SQL queries executed in DuckDB (Raasveldt & Mühleisen, 2019) directly against the Parquet snapshot; every query is reproduced in Section G. Concentration is measured with top-*k* concentration ratios (Tirole, 1988) rather than a Herfindahl index (Hirschman, 1964; Rhoades, 1993), because the policy question is how many employers one would need to audit to cover a given share of filings, not the abstract competitiveness of a market. Wage comparisons use the DOL prevailing-wage level recorded on the LCA and, for software developers, the BLS occupational median for the same SOC code and area. Each descriptive number was produced by one analyst’s query and then reproduced by a second, independent query before it was retained.

4.2 Semantic measurement of job-title histograms

The numeric fields of an LCA say how many, where, and at what wage. The job-title field says what. We measure the title

histogram with an evaluation model, TypeSafe’s Jev, accessed through the Vercel AI Gateway. Jev is not a chat model: it accepts a JSON state and a set of typed questions and returns, for each boolean question, a probability rounded to two decimal places. The use of large language models as zero-shot annotators of text for social-science measurement is by now well established (Gilardi et al., 2023; Ornstein et al., 2025; Rathje et al., 2024; Törnberg, 2024; Ziems et al., 2024), as are the cautions that such annotations must be validated and treated as measurements with error rather than as ground truth (Egami et al., 2023; Pangakis et al., 2023; Reiss, 2023).

Each employer-month is presented to the model as its structured record (application counts, wage, SOC and title histograms, location map, calendar month) with 18 boolean questions, each written as an instruction plus explicit true and false criteria and a guard forbidding inference about hiring, nationality, intent, future periods, or the employer’s identity (the employer name is not in the state). Ten predicates were written for this study and read the meaning of titles: whether the histogram reads as an IT staffing, consulting, or outsourcing operation; as a university, hospital, or research institution; as direct clinical care; whether titles are generic; whether senior-level tokens hold a majority of title mass; whether entry-level tokens hold at least a third; whether AI, machine-learning, or data-science titles hold at least a quarter; whether management titles hold at least a third; whether all titles sit in one job family; and whether senior or specialist titles dominate while the offered wage is modest for such titles. Eight are concise rewrites of predicates from the earlier campaign (Section 4.6) that had come closest to replicating. Full text of all 18 is in Section B.

One call per employer-month, 361,338 of 361,353 states scored (15 unrecoverable after retries), 1.06 billion input tokens, median latency 296 ms; the measurement cost \$0 under the gateway’s promotional rate for this model. Measurement reliability was assessed by test-retest on 80 employer-months and six predicates scored twice in fresh requests: 480 paired answers, Pearson $r = 0.998$, mean absolute difference 0.014, 1.5% of answers crossing 0.5 between passes, and measurement noise equal to 0.25% of between-state variance. The measurement is, for practical purposes, deterministic; Section E gives the per-predicate table. Face validity was checked after scoring by mapping employer-months back to names: the ten employers with the most applications in months flagged as staffing operations are Cognizant, Tata Consultancy Services, Infosys, Deloitte Consulting, Capgemini, Accenture, Wipro, IBM, Compunnel, and Tech Mahindra, each at 100% of its months; the academic predicate’s top ten are all universities and academic medical centers. Two known false positives are documented in Section 6.

The 18 scores are used in two ways. Descriptively, an employer-month is “flagged” when its probability exceeds 0.5, and we report shares of employer-months and of applications flagged by year. Inferentially, the raw probability enters the regressions as a continuous predictor, standardized per split, so that no thresholding decision is made on the outcome data.

4.3 Regression design

For each predicate j , outcome y , and horizon $h \in \{1, 2, 3\}$, we fit a generalized linear model (McCullagh & Nelder, 1989; Nelder

& Wedderburn, 1972) of $y_{i,t+h}$ on the standardized predicate score and a structured baseline that is identical across all 270 specifications: the ten structured panel variables (log application counts at one and six months, three-month growth, median wage and a wage-observed indicator, title entropy, location entropy, technical share, unique titles, unique locations), their squares, a linear time trend, calendar-month and two-digit-industry indicators, and sixteen hashed bag-of-words bins of the title text (Manning et al., 2008; Weinberger et al., 2009). The hashed title bins are the point: they give the model a deterministic, model-free representation of the same title text the semantic score is read from, so that a significant coefficient on the semantic score is evidence of information beyond a bag of words. The baseline is fixed in advance and is the same in every specification; there is no variable selection.

The family follows the outcome: Poisson quasi-likelihood for counts (Cameron & Trivedi, 2013; Gourieroux et al., 1984; Santos Silva & Tenreiro, 2006; Wedderburn, 1974), logistic for the binary contraction indicator, Gaussian for the three continuous changes. Standard errors are two-way cluster-robust by employer and by calendar month (Cameron et al., 2011; Petersen, 2009; Thompson, 2011), generalizing the sandwich estimator (White, 1980); discovery-split fits have 45 month clusters and roughly 4,000 employer clusters. Effects are reported per standard deviation of the predictor, and for Gaussian outcomes divided by the outcome’s standard deviation, so that they are comparable across specifications. Models are fit with statsmodels (Seabold & Perktold, 2010) on pandas and NumPy (Harris et al., 2020; McKinney, 2010). Employer-months with fewer than five applications are excluded from the regressions (a minimum-support rule for reading a title histogram), which leaves 47,144 to 53,457 observations per discovery fit.

4.4 Multiple testing, replication, and the sealed holdout

The 270 discovery p-values are corrected as one family with the Benjamini-Hochberg procedure (Benjamini & Hochberg, 1995; Benjamini & Yekutieli, 2001), and we report q-values (Storey, 2002; Storey & Tibshirani, 2003). Every specification is then refit on the validation split; a discovery hit is counted as replicated when the validation coefficient has the same sign and $p < 0.05$. For each validation fit we also compute the out-of-time incremental predictive gain of the model with the semantic score over the same model without it, in the proper scoring rule appropriate to the family (mean log loss for binary and Poisson deviance for counts, squared error for Gaussian) (Gneiting & Raftery, 2007); a positive gain means the score improves prediction on data the model never saw.

Thirty-nine specifications cleared a strict bar, discovery $q < 0.005$ and validation $p < 0.005$ with the same sign. Their identifiers were written to a pre-registration record with a SHA-256 digest (587314db33b635d7...) at 05:07 UTC on September 21, 2026, before the holdout had finished scoring, and the holdout was then fit once, for those 39 and no others (Dwork et al., 2015; Nosek et al., 2018; Wagenmakers et al., 2012). Every other use of the holdout in this paper is descriptive (prevalence by year) and touches no outcome variable.

4.5 Negative controls and the garden of forking paths

Each discovery fit is accompanied by a negative-control fit in which the same predictor is regressed on the outcome measured before the predictor's month rather than after (Arnold et al., 2016; Lipsitch et al., 2010). A predicate that is a stable employer trait will "predict" the past as well as the future; one that is a leading indicator will not. We report the negative-control p-value for every specification in Section A and interpret the results accordingly: most of the replicated effects are traits, which is the substantively interesting finding, and a minority are dynamics.

The whole design exists because of the garden of forking paths (Gelman & Loken, 2014; Simmons et al., 2011). The predicate list, the outcome list, the horizons, the baseline, the family, the clustering, and the thresholds were fixed before any full-panel regression was run, and all 270 results are reported.

4.6 The previously explored design and why it failed

This paper's panel, outcomes, structured baseline, and regression code were inherited from an earlier design that ran from September 19 to September 21, 2026 and produced no findings. It is described here because its failure is instructive and because the diagnosis is what motivated the design above.

The earlier design was a hypothesis-screening loop in the style of recent automated-discovery systems (Gu et al., 2024; Lu et al., 2024; Majumder et al., 2024). A large language model (Grok 4.6) proposed falsifiable hypotheses in batches of twenty, each naming a semantic predicate, an outcome, a horizon, an expected sign, controls, and a mechanism; the same model wrote the predicate definitions; the evaluation model measured the predicates; and deterministic code tested each hypothesis on a sample of 300 employer-months per split at a per-round discovery FDR of 0.005 and a validation α of 0.005, with the negative control and predictive-gain tests described above. Four campaigns ran in sequence: a screen of 20 hypotheses, then 10, 10, and 48 rounds, 1,342 hypotheses in all. None was validated. The final campaign of 927 hypotheses classified 888 as rejected and 39 as inconclusive.

The reasons are five, and all are measurable from the campaign's own output (Figure 1).

Power. With 300 observations against a design of roughly 65 columns, the median standard error of the standardized effect was 0.198 for application counts, 0.082 to 0.088 for the diversity outcomes, 0.043 for wage change, and 0.498 for contraction. The minimum detectable effect at the campaign's own threshold with 80% power (Bloom, 1995; Cohen, 1988) was therefore 0.72, 0.30 to 0.32, 0.16, and 1.82 standard deviations respectively, against a configured effect-size floor of 0.10. No application-count or contraction test, and 3% to 4% of diversity tests, could have detected an effect at the floor. The campaign was a screen for effects of a size that does not exist in administrative filing data, and it found none (Button et al., 2013; Gelman & Carlin, 2014; Ioannidis, 2005).

Miscalibration masquerading as signal. The discovery p-values were not uniform: 163 of 908 fell in the first decile against 91 expected, and 27 were below 0.005 against 4.5 expected. But the negative-control p-values showed the same excess (156 in the first decile, 26 below 0.005). Excess small p-values

that appear equally when predicting the past are a property of the test, not of the hypotheses; with 300 rows, 45 month clusters, and a near-saturated design, the two-way cluster-robust variance is unreliable (Thompson, 2011). Discovery and validation p-values were uncorrelated ($r = 0.014$), effect sizes barely correlated ($r = 0.098$), the sign matched the generator's expected direction in 47% of tests, and the sign agreed between splits in 51%. Ten hypotheses were nominally significant with the same sign in both splits, against about one expected under the null; none improved out-of-time prediction.

Degenerate measurements. Of 89 predicates the generator wrote, 25 were flagged on fewer than about 30 of 300 sampled employer-months and 5 had a standard deviation below 0.01 across 242,020 cached measurements; they were constants, and the 232 hypotheses built on them could not have been estimated. Fifteen predicates referenced no title text at all and were threshold rules over numeric fields already in the control set. The generator's guardrails, which forbade inference about the future in 79 of 85 predicate texts and about nationality in 73, had stripped the semantic content out of the semantic layer, leaving arithmetic.

The generator did not learn. Sixty-four percent of the hypotheses were interaction terms, the least powered specification available. When the 927 hypotheses of the final campaign were themselves scored by the evaluation model on four questions (whether the predicate was computable from the controls; whether it required reading title text; how well the stated mechanism supported the claim, on a 0 to 3 scale; and whether the expected sign followed from the mechanism), the mechanism score averaged 2.02 with two thirds of hypotheses in a single bin and 3% at or above 2.5, and the model judged the expected sign to follow from the stated reason with mean probability 0.57. Across 48 rounds of feedback the mechanism score drifted from 2.07 to 1.97; the only thing that improved was the rate of predicates redundant with the controls (0.27 to 0.09), the one rejection reason the loop reported back verbatim. This is the failure mode that the literature on machine-generated research ideas describes: fluent, plausible, and interchangeable (Si et al., 2025; Wang et al., 2024).

The critic could not rank noise. The evaluation model's mechanism score was uncorrelated with the observed test statistic ($r = -0.008$). When outcome variance at $n = 300$ is noise, no prior can rank it, and the useful role of a model-based critic begins only once the screen has power.

The remedy was not a better generator. It was to measure the full panel, which the measurement budget permitted at zero marginal cost, to write predicates that read titles, to drop interactions and constants, to fix a single baseline, and to replace 1,342 free-form hypotheses with a 270-cell grid corrected as one family. That is the design of Section 4.3. The chronology, design decisions, and per-campaign counts of the earlier effort are documented in Section F, reconstructed from the project's own logs and session transcripts.

4.7 Disclosure

The earlier hypothesis generation used Grok 4.6; all semantic measurement used TypeSafe Jev via the Vercel AI Gateway with zero data retention; analysis orchestration, code, and drafting were done with Anthropic's Claude under the author's




Figure 1: The earlier design and why it failed. (A) Discovery and negative-control p-value deciles across 927 hypotheses tested on 300-observation samples; the excess at small p appears equally when predicting the past. (B) The same test family on the full panel: 117 of 270 tests below $q = 0.05$. (C) Minimum detectable standardized effect at $p < 0.005$ and 80% power in the earlier design, by outcome, against its configured floor of 0.10.




Figure 2: Share of certified LCAs, FY2020 to FY2026, filed by the top k employers by legal name.

direction, with every reported number produced by deterministic queries or regressions and reproduced by an independent query before inclusion. No model approved a statistical finding.

5 Findings from the raw record

5.1 Filing is concentrated in a few hundred employers

Of 4,215,630 certified LCA records in fiscal years 2020 through 2026, the ten largest filing entities account for 659,503 (15.6%), the fifty largest for 1,159,041 (27.5%), the hundred largest for 1,384,729 (32.9%), and the five hundred largest for 1,972,266 (46.8%), out of 205,032 employers with at least one certified filing (Figure 2). Because large companies file under multiple legal names, true concentration is higher.

The largest single filing entity is Cognizant Technology Solutions US Corp, 109,360 certified records, ahead of Amazon.com Services LLC (106,408), Google LLC (77,182), Ernst & Young U.S. LLP (71,800), Microsoft Corporation (69,103), Tata Consultancy Services Limited (68,230), Infosys Limited (53,823), Apple Inc. (40,316), Deloitte Consulting LLP (34,985), and Accenture LLP (28,296). By corporate family the Amazon entities together lead with 158,795, followed by Cognizant (110,470), Google (80,406), EY (72,008), Microsoft (69,103),

TCS (68,270), Infosys (54,299), Deloitte (54,086), Meta (46,783), and Apple (40,433) (Figure 3). Six of the top twelve families are IT services, consulting, or professional-services firms whose business is supplying labor to other companies.

5.2 The largest filer files the same title by the thousand

The disclosed job title on Cognizant’s records is its internal job code. “Senior Systems Analyst JC60” was certified 10,032 times in FY2020 and 6,889 times in FY2021; “Manager JC50” 8,285 and 6,487 times. No other employer-title-year cell in the file exceeds 6,800 (Google, “Software Engineer,” 6,769 in FY2023 and 6,666 in FY2026). The prevailing-wage attestation is made per LCA, for a stated occupation, at a stated level. A filing practice in which an internal grade code stands in for the occupation, ten thousand times a year, is a practice in which the specialty-occupation and wage-level determinations are made once, in bulk, by the employer.

5.3 The program is priced at the bottom of its own wage scale

The Department of Labor’s prevailing-wage system has four levels; Level I is defined as entry-level and Level II as “qualified,” and both are set below the occupational median by construction. Of certified LCAs, 15.3% are at Level I, 41.3% at Level II, 20.5% at Level III, 15.6% at Level IV, and 7.4% record no level. Levels I and II together are 2,384,880 records, 56.6% (Figure 4). The share has been stable: 60.9% in FY2020, 55.8%, 56.7%, 56.9%, 58.6%, and 57.3% in FY2025, with the Level I share alone rising from 13.3% to 18.8% over the same years. A program defined in statute as one for specialty occupations is, by the government’s own wage taxonomy, majority-filled at the two rungs defined as below-average pay.

The occupation that dominates the file makes the point concrete. Software developers (SOC 15-1252 and its predecessor 15-1132) are 1,373,338 certified records, 32.6% of everything; the next occupation, computer systems analysts, is 2.9%. For SOC 15-1252 in FY2023 and FY2024, the median offered wage was \$89,253 at Level I ($n = 43,611$; 87.8% of these offers below the May 2023 BLS national median of \$130,160), \$108,992 at




Figure 3: Certified LCAs by (A) legal entity and (B) corporate family, FY2020 to FY2026.




Figure 4: (A) Distribution of certified LCAs across DOL prevailing-wage levels, FY2020 to FY2026. (B) Share at Levels I and II by fiscal year.

Level II ($n = 172,627$; 75.3% below), \$140,000 at Level III, and \$161,080 at Level IV. There were 216,238 Level I and II software-developer certifications in those two years. In the Washington, DC metropolitan area, 69.2% of 2,033 certified software-developer LCAs offered less than the local BLS median (\$126,533 offered against \$139,850); in the New York metropolitan area the figure was 34.7% ($n = 12,768$). The BLS coverage in this snapshot is limited to three metros and one year, so these are illustrations, not a national estimate.

5.4 Staffing-pattern employers file at the bottom more often

A crude name match (employers whose normalized names contain “consulting,” “technologies,” “infotech,” “solutions,” “systems,” “staffing,” “IT services,” or “software”) captures 22,455 employers and 755,299 certified records, 17.9% of the file. Those employers filed at Level I or II 66.8% of the time, against 54.3% for all others, at a mean offered wage of \$186,790 against \$220,313. The name match is a blunt instrument; the semantic measurement in Section 6 is the sharp one, and it agrees.

5.5 The outsourcing model is a temporary-labor model

The clearest single statistic on how an employer uses the program is the ratio of its H-1B filings to its green-card filings. An employer that sponsors the workers it brings in will show a low ratio; one that cycles temporary workers will show a high one. Over FY2020 to FY2026, by corporate family (Figure 5): Tech Mahindra 45.4 certified LCAs per certified PERM labor certification; Cognizant 35.1 (110,470 against 3,144); Ernst & Young 26.2; Wipro 18.5; Deloitte 15.5; Accenture 13.8; HCL 13.4; Infosys 12.0; TCS 8.4. Against that: Oracle 2.4; Intel 4.0; Microsoft 4.2 (69,103 against 16,407); Nvidia 4.3; Cisco 4.8; Google 5.8; Apple 5.9; Amazon 6.4. A low ratio is not proof of virtue, and the two pipelines have different timing. But a 35-to-1 ratio, sustained over seven years at a firm filing over a hundred thousand LCAs, means that the firm’s H-1B workforce is overwhelmingly not on a path to permanence, which is the condition under which an employer’s leverage over a worker’s presence in the country is greatest.




Figure 6: The outsourcing share of H-1B filing. (A) Raw record: share of certified LCAs filed by the eight large IT outsourcers and by sixteen outsourcing and consulting families, by fiscal year. (B) Semantic measurement: share of applications, and of employer-months, from employer-months whose title histogram reads as a staffing or consulting operation, by calendar year, 2019 (October to December only) through the first half of 2026.




Figure 5: Certified LCAs per certified PERM labor certification by corporate family, FY2020 to FY2026. Orange: 12 or more.

5.6 The outsourcers' share is falling, and the program is not

The eight large IT outsourcing families (Cognizant, TCS, Infosys, Wipro, HCL, Tech Mahindra, Capgemini, Accenture) filed 13.2% of certified LCAs in FY2020 and 5.3% in FY2026; with EY, Deloitte, IBM, LTIMindtree, Compunnel, Mphasis, UST, and Hexaware added, 19.6% and 9.7% (Figure 6, panel A). The program did not shrink: 545,621 certified records in FY2020, 666,228 in the first three quarters of FY2026. The composition changed. Who absorbed the share is the subject of Section 6.




Figure 7: Median days from PERM receipt to decision, by fiscal year of decision.

5.7 Processing and denial anomalies

PERM processing time has more than tripled. The median interval from receipt to decision on PERM applications, by fiscal year of decision, was 137 days in FY2020, 194 in FY2021, 189 in FY2022, 287 in FY2023, 388 in FY2024, and 482 in FY2025 (Figure 7). These are decided cases only; the live backlog is longer. For a worker whose right to remain depends on an employer's willingness to wait, the wait is the leverage.

Denial rates are near zero, with named exceptions. The LCA denial rate is 0.47% to 0.79% in every fiscal year; the PERM denial rate was 4.5% in FY2020, peaked at 5.5% in FY2023, and was 1.8% in FY2025 (a figure that reflects which cases had been decided, not a policy change). Against that background: Wipro Limited's PERM decisions in FY2020 to FY2025 were 1,856, of which 574 were denied (30.9%; 35.6% of decided cases), 299 certified outright, 738 certified and expired, and 245 withdrawn, while JPMorgan Chase (802 decisions), Deloitte & Touche (641), SAP America (597), Bank of America (595), Mastercard (551), and Rivian (711) show no denials at comparable volume. One HCL entity, HCL America Solutions Inc, had 282 of 1,443 LCA decisions denied (19.5%), against 0.84% at Ford and




Figure 8: Certified LCAs per fiscal quarter, FY2020 Q1 to FY2026 Q3. Third quarters (April to June, the cap-petition season) in orange.

0.59% at Wells Fargo; the larger HCL America Inc entity has an ordinary denial rate, and this figure should be confirmed against the DOL public file before use. USCIS petition denials (data through FY2023) were 42.1% at Syntel (254 of 603) and 15.2% at Cognizant on 27,900 decisions, the largest-volume outlier; the national rate fell from 8.5% in FY2020 to 2.0% in FY2022. Denial reasons are not in the data, and none of this is evidence of intent. It is evidence of where an auditor would look first.

5.8 Timing and geography

The H-1B cap season (petitions filed April through June) produces a third-fiscal-quarter spike every year. FY2026 Q3 was the largest quarter in the series, 401,412 certified records against 224,127 in FY2025 Q3 and the prior record of 381,037 in FY2021 Q3 (Figure 8). The year-over-year growth was broad: Amazon.com Services 5,237 to 13,922; Google 1,279 to 7,423; Microsoft 2,715 to 5,159; EY 1,885 to 5,158; Apple 1,570 to 4,870; Meta 2,241 to 5,278; JPMorgan 1,301 to 2,801. TCS fell (3,922 to 2,880). The snapshot ends June 30, 2026, so the quarter is complete; the cause is outside this record and should be established against DOL's release notes before it is asserted.

By worksite state, Texas gained the most certified filings between FY2020 and FY2025 (55,863 to 84,361) and California lost the most (109,716 to 103,454). New York City is the largest single worksite city (245,882), followed by Seattle (111,127), San Francisco (102,891), Austin (93,581), and San Jose (81,077); four Texas cities (Irving, Houston, Plano, Dallas) each exceed 60,000. Five single LCA cases in FY2025 listed 1,250 to 1,325 workers at one worksite (one in the Bronx, four in Manassas Park, Virginia); the worksite file carries no employer name and the join to the case file was not completed, so these are leads.

6 What the titles say: findings from the semantic measurement

The analyses in this section use the 84,371 employer-months with at least five applications and a Jev score, from 4,346 employers, October 2019 through June 2026. "Flagged" means a predicate probability above 0.5. Prevalence is reported as the share of applications that fall in flagged employer-months (so that large filers count in proportion to their filing), with the unweighted share of employer-months in Section C.

6.1 The outsourcing share of applications has more than halved

Employer-months whose title histograms read as IT staffing, consulting, or outsourcing operations accounted for 67.0% of applications in the last quarter of 2019, 56.6% in 2020, 49.4% in 2021, 43.7% in 2022, 41.6% in 2023, 37.2% in 2024, 34.6% in 2025, and 31.8% in the first half of 2026 (Figure 6, panel B); 44.1% of applications and 33.7% of employer-months over the whole period. The decline survives every robustness cut we tried: unweighted by volume, 45% to 27%; excluding the ten largest filers, 57% to 37%; excluding Google and Amazon Web Services, the two product companies the predicate over-flags, 66% to 37%. Flagged months offer a median wage of \$98,332 against \$112,185 in other months. This is the same decline the raw record shows for the named outsourcers (Figure 6, panel A), measured without any employer name.

6.2 AI titles rose nine-fold, and Walmart is the second-largest AI filer

Certified LCAs whose title names machine learning, artificial intelligence, data science, deep learning, natural language processing, computer vision, or large language models rose from 1.4% of the file in FY2020 to 2.8% in FY2025 and 3.5% in FY2026 (Figure 9, panel A). Employer-months in which such titles are at least a quarter of the histogram rose from 0.9% of applications in late 2019 to 5.2% in 2024, 5.9% in 2025, and 8.2% in the first half of 2026 (panel B), at a median offered wage of \$127,213 against \$103,074 in other months. By employer, the largest AI-title filers over the period are Apple (3,279 certified records, 8.1% of its filings, median \$183,800), Wal-Mart Associates, Inc. (2,309; 9.4%; \$160,000), Amazon.com Services (2,242; 2.1%), LinkedIn (1,837; 18.6%), Meta (1,595 and 1,110 under two name variants), Google (1,235; 1.6%), Adobe (1,110; 14.3%), and Microsoft (1,097; 1.6%). Walmart's AI titles are data-science titles at every grade (senior data scientist 743, data scientist 407, staff data scientist 330).

6.3 Filers stopped hiring one kind of worker

Employer-months whose titles all sit in one job family fell from 53.1% of applications in late 2019 to 26.2% in the first half of 2026; months spanning three or more job families rose from 26.5% to 51.8%; generic, unqualified titles fell from 20.8% to 10.4%; management-heavy months fell from 22.0% to 13.9%;




Figure 9: AI, machine-learning, and data-science hiring in the H-1B record. (A) Share of certified LCAs with an AI title, by fiscal year. (B) Share of applications from employer-months in which AI titles are at least a quarter of the title histogram, by calendar year.




Figure 10: Share of applications from flagged employer-months, six predicates, by calendar year. 2019 covers October to December; 2026 covers January to June.

academic and medical months rose from 6.7% to 11.3% and clinical months from 0.5% to 3.3% (Figure 10). The program is moving from single-function labor pools toward broad corporate hiring, which is the same shift as the outsourcers' declining share seen from inside the filer.

6.4 Financial firms hire H-1Bs at the bottom of the ladder

Among certified LCA titles, the share carrying an entry-level token (analyst, associate, junior, intern, level I) is 67.0% at Goldman Sachs & Co. (10,211 records; median \$125,000), 67.0% at Citibank, N.A. (9,095; \$150,000), 64.5% at BlackRock Financial Management (4,217; \$105,622), 54.0% at McKinsey & Company (5,180; \$160,400), and 43.5% at JPMorgan Chase (26,067; \$144,800), against 4.9% at Google. JPMorgan's most

frequent title is "associate, software engineer" (2,743). The semantic predicate agrees: 95% of Goldman's applications, 92% of BlackRock's, and 88% of JPMorgan's fall in entry-level-heavy months. The financial sector's H-1B program is an analyst-class pipeline.

6.5 "Senior title, modest wage" is measurable and has names

Two percent of employer-months (3.3% of applications) have senior or specialist titles dominating the histogram while the offered median wage is modest for such titles. Their median wage is \$82,000 against \$105,040 elsewhere. The employers with the most applications in such months are Cognizant (14,606; 21% of its applications), TCS (10,508; 20%), Wipro (7,711; 58% of its applications; median \$78,499), Deloitte & Touche




Figure 11: The 39 pre-registered effects across three time periods. Each row is one predicate, outcome, and horizon; markers are the standardized coefficient in discovery (2019-2023), validation (2023-2025), and the sealed holdout (2025-2026).

(3,962; 48%), Deloitte Tax (1,030; 40%), UST Global (1,078), Cummins (1,077), and Persistent Systems (737). Wipro has the highest mean flag probability of any employer with 2,000 or more applications (0.52). The pattern was far more common early in the period (12.3% of applications in late 2019, 6.1% in 2020) and is now rare (0.3% in the first half of 2026).

6.6 Universities pay half of what technology pays

Academic and medical employer-months are 9.0% of applications and 17.1% of employer-months. Their median certified wage is \$66,837 at the University of Michigan (4,013 records), \$71,448 at Johns Hopkins, \$75,000 at Yale and at Harvard, \$75,684 at the Mayo Clinic, \$78,000 at Columbia, and \$82,000 at Stanford, against \$167,000 at Google and \$183,800 for AI roles at Apple. In the regressions below, the academic predicate is the strongest and most stable of all 18: academic employer-months are the least likely to contract in the following months and the most likely to diversify their roles. The postdoctoral economy runs on H-1B at postdoctoral pay. Whether that is an adverse effect on similarly employed U.S. workers is a question the wage-level system, which classifies most of these positions at Level I or II, is designed not to ask.

6.7 Two known false positives

Google LLC is flagged as a staffing operation in 53% of its months and Amazon Web Services in 92%, because both file “Software Engineer” by the thousand and the predicate reads title genericness at volume. We report them, exclude them in the robustness cut above, and note that the same limitation is inherent in any measurement of an employer from its titles alone: a product company that files generic titles is, on this record, indistinguishable from a body shop.

6.8 Replicated dynamics

Of the 270 specifications in the full-panel grid, 117 clear $q < 0.05$ in discovery and 78 clear $q < 0.005$. Of the 117, 112 keep their sign in the validation split (96%; the earlier design managed 51%), 74 are significant at $p < 0.05$ with the same sign, and 73 of those improve out-of-time prediction. Thirty-nine clear the strict bar and were pre-registered. In the sealed holdout (February 2025 through June 2026; 10,635 to 13,327 employer-months per fit), all 39 keep their sign, 32 are significant at $p < 0.05$, and 23 at $p < 0.005$ (Figure 11; Section A gives all 270).

The effects are small per standard deviation, as real effects in administrative data are, and they are stable: academic-medical on role-diversity change, +0.188, +0.173, +0.180 across the three periods; senior-title-modest-wage on subsequent wage change, +0.048, +0.081, +0.081; entry-level share on wage change, +0.055, +0.068, +0.075; AI titles on wage change, +0.042, +0.035, +0.047; the staffing signature on role-diversity change, -0.080, -0.072, -0.065; the staffing signature on location-diversity change, +0.065, +0.075, +0.009 (the weakest, and the one holdout miss that goes to zero). Which predicates replicate: academic-medical (12 of 15 tests), entry-level share (9), staffing (8), senior-title-modest-wage (7), wage-title discord (7), technical-to-business mix (6), AI (5), multi-seniority span (5). Which outcomes: role-diversity change (26), wage change (25), location-diversity change (17), contraction (6). Application volume itself: nothing replicates, which is as it should be; next month's filing count is dominated by this month's, and titles add nothing to that.

Three of the replicated dynamics bear directly on wage setting. Employer-months flagged for senior titles at a modest wage, and employer-months heavy in entry-level titles, both show the largest subsequent wage increases of any flag, in all three periods, with negative-control p-values above 0.05 in most horizons. The underpayment does not persist at the employer level; it is corrected in the following months. That is consistent with an employer filing at the lowest defensible level and adjusting, and it is the pattern an LCA wage audit is designed to catch. Whether the correction reaches the worker whose LCA was certified at the lower wage is not observable in this record.

7 Discussion

7.1 What the record supports

The statute's protective condition is a wage attestation. The record shows a program in which a majority of certified applications are made at the two wage levels the Department of Labor itself defines as below the occupational median; in which the single largest occupation is certified at a Level I median 31% below the national median for that occupation; in which the largest filer is a labor-supply firm that files internal grade codes as job titles ten thousand times a year and sponsors permanent residence for one worker in thirty-five; in which the employers whose titles read as staffing operations file at the bottom levels two-thirds of the time and pay \$14,000 less at the median; and in which the surviving dynamics show that employer-months priced at the bottom are followed by wage corrections, not by stability.

None of this measures displacement of a specific U.S. worker, and the paper does not claim it does. What the record measures is the price at which employers are attesting that foreign labor does not adversely affect U.S. workers, and that price is, for most of the program, the lowest the wage-level system allows. The economics of a wage floor set below the market median are not subtle: an employer who can hire at Level I for a job the market pays at the median has been given a discount, and the discount is paid for by whoever would otherwise have been hired at the median. That is a labor-cost instrument, and the concentration, the LCA-to-PERM ratios, and the repeated-

title batch filings are what a labor-cost instrument looks like in administrative data.

This reading is consistent with the parts of the literature that measure the same margin. The below-median wage-level share we find (56.6% over FY2020 to FY2026) is the same quantity Costa and Hira found for FY2019 (60%) (Costa & Hira, 2020), now extended over seven years and shown to be stable while the Level I share rises. The 35-to-1 and 45-to-1 LCA-to-PERM ratios at the outsourcing families are the filing-level signature of the business model Hira described as importing ordinary-skill workers who substitute for incumbents (Hira, 2010), and they are what the Edison and Disney contracts looked like from the vendor's side (Thibodeau, 2015). The firm-level crowd-out estimate of Doran, Gelber, and Isen (Doran et al., 2022) and the skill-group wage estimates of Turner and of Bound, Khanna, and Morales (Bound et al., 2018; Turner, 2022) are the causal evidence that the price effect described above reaches U.S. workers; the city-level complementarity results (Peri et al., 2015) and the Richmond Fed null (Mahajan et al., 2025) are the evidence that it does not reach all of them, or reaches them through channels the LCA record cannot see. What this paper adds is not a causal estimate but a map: which employers, at which wage levels, in which occupations, with what permanence, and how that map is moving.

The composition shift is the finding most likely to be misread. The outsourcers' share has halved; the program has not shrunk. The share moved to product companies, financial firms, and universities that file more diverse titles at higher nominal wages. Whether that is an improvement for U.S. workers depends on whether the wage-level distribution moved with it, and it did not: the Level I and II share was 60.9% in FY2020 and 57.3% in FY2025, and the Level I share alone rose. The financial sector's entry-level pipeline and the universities' postdoctoral pipeline are the two places where the new composition and the old wage levels meet.

7.2 Limitations

The unit of observation is a filing, not a person. A certified LCA may be unused, a worker may be paid above the attested wage, and none of the outcomes in the panel are employment outcomes for U.S. workers. Employer identity is a legal name; corporate-family rollups depend on a mapping that is incomplete for small firms. BLS wage benchmarks cover six occupations, three metros, and one year in this snapshot. The USCIS file ends at FY2023. The semantic predicates are one model's reading of a title histogram; the two false positives show the failure mode, and a title-only signature cannot see an employer's business model when its titles are generic. The regressions describe association at the employer-month level within a fixed baseline; they are not causal estimates, and the negative-control results show that most replicated predicates are traits rather than leading indicators. The panel is restricted to sustained filers, which excludes most small and intermittent employers from the semantic analyses (but not from the descriptive ones). The pre-registration is internal to this project and time-stamped, not lodged with a third-party registry.

8 Implications for enforcement targeting

The same public files support a targeting model. The variables that separate employers in this record are computable for every filer, monthly, from data the agencies already hold.

The enforcement record already contains the pattern. A unanimous jury found in October 2024 that Cognizant, the largest filer in this record, engaged in a pattern or practice of intentional discrimination against non-South-Asian and non-Indian employees terminated from its bench, and in December 2025 the court found that its “Visa Readiness” and “Visa Utilization” policies had a disparate impact on the class (Fan & Jones, 2024; *Palmer v. Cognizant Technology Solutions Corp.*, 2024); a jury reached the opposite verdict on the same theory against Tata Consultancy Services (*Slaight v. Tata Consultancy Services, Ltd.*, 2021), and the Seventh Circuit affirmed summary judgment for Infosys in 2026 after the plaintiffs’ statistical expert was excluded (*Koehler v. Infosys Technologies Limited, Inc.*, 2026), which is a warning about method as much as about merits. On the citizenship-status side, the Immigrant and Employee Rights Section has found a pattern or practice of reserving PERM positions for temporary visa holders at Facebook (\$14.25 million, 2021) (U.S. Department of Justice, Office of Public Affairs, 2021), Apple (\$25 million, 2023) (U.S. Department of Justice, Office of Public Affairs, 2023), and OpenAI (\$3.2 million, 2026) (U.S. Department of Justice, Office of Public Affairs, 2026a), and penalized an IT staffing intermediary, Compunnel, for job advertisements restricted to visa holders (U.S. Department of Justice, Office of Public Affairs, 2026b); Infosys paid \$34 million in 2013 over visitor-visa misuse (U.S. Department of Justice, U.S. Attorney’s Office, Eastern District of Texas, 2013). The Commission’s own 2024-2028 strategic enforcement plan lists workers on temporary visas first among vulnerable populations and targets practices that confine workers to temporary positions on a prohibited basis (U.S. Equal Employment Opportunity Commission, 2016; 2023; *Unfair Immigration-Related Employment Practices*, 1986; *Unlawful Employment Practices*, 1964). What the agencies lack is not a theory but a ranking of where to look.

For citizenship-status and national-origin enforcement (DOJ IER, EEOC). Rank employers by (1) the ratio of certified LCAs to certified PERM labor certifications over a rolling three-year window, (2) the share of filings at Levels I and II, (3) the share of an employer’s title histogram occupied by a single title string, and (4) the staffing-signature probability from a title-reading model of the kind used here. Employers in the top decile on three of the four are the population in which a recruitment or retention practice that disadvantages U.S. workers is most likely to be found and cheapest to investigate, because the same pattern that produces the statistics (bulk filing of interchangeable titles at the lowest level, without a path to permanence) is the pattern the Cognizant jury and the Facebook settlement described.

For LCA wage compliance (DOL WHD and OFLC). The “senior title, modest wage” and “entry-level heavy” flags identify employer-months where the attested wage level is most likely to be inconsistent with the attested job, and the replicated dynamics show those months are followed by upward wage revisions. An audit sample drawn from flagged months,

stratified by employer, would test the attestation at the point it is most likely to be wrong. The repeated-title batch statistic (the share of an employer’s annual filings under one title) identifies employers whose specialty-occupation determinations are made once per grade code rather than per position.

For all three. The processing-time series turns a compliance question into a leverage question: a worker whose PERM decision takes 482 days is a worker whose employer holds his presence in the country for 482 days. Enforcement priorities that weight employer leverage would weight the LCA-to-PERM ratio and the processing backlog together.

The model that produced the rankings in this paper runs on a laptop against the public files in under an hour once the semantic scores exist, and the semantic scores for the entire seven-year record cost less than \$50 at list price. A monthly refresh is a matter of scheduling.

9 About the author and availability

Josh C. Simmons holds a PhD from the University of California, Irvine (2019) (Simmons, 2019) and has built and led engineering teams at Anduril, Shopify, and American Express, where he was an Engineering Director; his account of visa-linked substitution at that employer (Simmons, 2026a; 2026b) is one of the primary-source reports that motivated this study, and his prior writing on labor arbitrage (Simmons, 2025a; 2025b; 2026c) and on evaluation-model measurement (Simmons, 2026d) is cited where relevant above. He is the founder of Meridian Strategic Systems, an independent research and media company in Queen Creek, Arizona, and the host of The Dr. J Show.

Meridian Strategic Systems performs this kind of analysis on contract: administrative-record audits, employer-level targeting models, semantic measurement of text fields at scale, and pre-registered replication designs, delivered with reproducible queries and code. Inquiries: hello@meridianstrategic.systems. The queries, predicate texts, pre-registration record, and all 270 regression results are in the appendices; the snapshot itself is a deterministic extract of the public DOL, USCIS, BLS, and SEC files and can be rebuilt from their current releases.

References

- 20 C.F.R. Part 655, Subpart H: Labor Condition Applications and Requirements for Employers Seeking to Employ Non-immigrants on H-1b Visas in Specialty Occupations and as Fashion Models. <https://www.ecfr.gov/current/title-20/chapter-V/part-655/subpart-H>
- Arnold, B. F., Ercumen, A., Benjamin-Chung, J., & Colford, J., John M. (2016). Brief Report: Negative Controls to Detect Selection Bias and Measurement Bias in Epidemiologic Studies. *Epidemiology*, 27(5), 637–641. <https://doi.org/10.1097/EDE.0000000000000504>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4), 1165–1188. <https://doi.org/10.1214/aos/1013699998>

- Bergmeir, C., & Benítez, J. M. (2012). On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191, 192–213. <https://doi.org/10.1016/j.ins.2011.12.028>
- Bloom, H. S. (1995). Minimum Detectable Effects: A Simple Way to Report the Statistical Power of Experimental Designs. *Evaluation Review*, 19(5), 547–556. <https://doi.org/10.1177/0193841X9501900504>
- Borjas, G. J. (2026). *The H-1B Wage Gap, Visa Fees, and Employer Demand* (No. NBER Working Paper 34793). National Bureau of Economic Research. <https://www.nber.org/papers/w34793>
- Bound, J., Khanna, G., & Morales, N. (2018). Understanding the Economic Impact of the H-1B Program on the U.S. In G. H. Hanson, W. R. Kerr, & S. Turner (Eds.), *High-Skilled Migration to the United States and Its Economic Consequences* (pp. 109–176). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226525662.003.0004>
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365–376. <https://doi.org/10.1038/nrn3475>
- Cameron, A. C., & Trivedi, P. K. (2013). *Regression Analysis of Count Data* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9781139013567>
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust Inference With Multiway Clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249. <https://doi.org/10.1198/jbes.2010.07136>
- Cerqueira, V., Torgo, L., & Mozetič, I. (2020). Evaluating time series forecasting models: an empirical study on performance estimation methods. *Machine Learning*, 109(11), 1997–2028. <https://doi.org/10.1007/s10994-020-05910-7>
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Congressional Research Service. (2025). *The H-1B Visa for Specialty Occupation Workers* (No. IF12912). Congressional Research Service. <https://www.congress.gov/crs-product/IF12912>
- Costa, D., & Hira, R. (2020). *H-1B visas and prevailing wage levels*. Economic Policy Institute. <https://www.epi.org/publication/h-1b-visas-and-prevailing-wage-levels/>
- Costa, D., & Hira, R. (2023). *Tech and outsourcing companies continue to exploit the H-1B visa program at a time of mass layoffs*. Economic Policy Institute. <https://www.epi.org/blog/tech-and-outsourcing-companies-continue-to-exploit-the-h-1b-visa-program-at-a-time-of-mass-layoffs-the-top-30-h-1b-employers-hired-34000-new-h-1b-workers-in-2022-and-laid-off-at-least-85000-workers/>
- Doran, K., Gelber, A. M., & Isen, A. (2022). The Effects of High-Skilled Immigration Policy on Firms: Evidence from Visa Lotteries. *Journal of Political Economy*, 130(10), 2501–2533. <https://doi.org/10.1086/720467>
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., & Roth, A. (2015). The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248), 636–638. <https://doi.org/10.1126/science.aaa9375>
- Egami, N., Hinck, M., Stewart, B. M., & Wei, H. (2023). Using Imperfect Surrogates for Downstream Inference: Design-based Supervised Learning for Social Science Applications of Large Language Models. *Advances in Neural Information Processing Systems 36 (Neurips 2023)*, 68589–68601. https://proceedings.neurips.cc/paper_files/paper/2023/hash/d862f7f5445255090de13b825b880d59-Abstract-Conference.html
- Fan, E., & Jones, C. (2024, December 9). Insiders Tell How IT Giant Favored Indian H-1B Workers Over US Employees. *Bloomberg News*. <https://www.bloomberg.com/graphics/2024-cognizant-h1b-visas-discriminates-us-workers/>
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science*, 9(6), 641–651. <https://doi.org/10.1177/1745691614551642>
- Gelman, A., & Loken, E. (2014). The Statistical Crisis in Science. *American Scientist*, 102(6), 460. <https://doi.org/10.1511/2014.111.460>
- Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30). <https://doi.org/10.1073/pnas.2305016120>
- Glennon, B. (2024). How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program. *Management Science*, 70(2), 907–930. <https://doi.org/10.1287/mnsc.2023.4715>
- Gneiting, T., & Raftery, A. E. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477), 359–378. <https://doi.org/10.1198/016214506000001437>
- Gourieroux, C., Monfort, A., & Trognon, A. (1984). Pseudo Maximum Likelihood Methods: Applications to Poisson Models. *Econometrica*, 52(3), 701–720. <https://doi.org/10.2307/1913472>
- Gu, K., Shang, R., Jiang, R., Kuang, K., Lin, R.-J., Lyu, D., Mao, Y., Pan, Y., Wu, T., Yu, J., Zhang, Y., Zhang, T. M., Zhu, L., Merrill, M. A., Heer, J., & Althoff, T. (2024). BLADE: Benchmarking Language Model Agents for Data-Driven Science. *Findings of the Association for Computational Linguistics: EMNLP 2024*, 13936–13971. <https://doi.org/10.18653/v1/2024.findings-emnlp.815>
- Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M. H. van, Brett, M., Haldane, A., Río, J. Fernández del, Wiebe, M., Peterson, P., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Hira, R. (2010). *The H-1B and L-1 Visa Programs: Out of Control* (No. EPI Briefing Paper 280). Economic Policy Institute. <https://www.epi.org/publication/bp280/>
- Hira, R. (2016). *Testimony of Ronil Hira before the Subcommittee on Immigration and the National Interest, Hearing on the Impact of High-Skilled Immigration on U.S. Workers*. U.S. Senate Committee on the Judiciary. <https://www.>

- judiciary.senate.gov/imo/media/doc/02-25-16%20Hira%20Testimony.pdf
- Hirschman, A. O. (1964). The Paternity of an Index. *The American Economic Review*, 54(5), 761–762. <https://doi.org/10.2307/1818582>
- Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts. <https://otexts.com/fpp3/>
- Immigration and Nationality Act Section 212(N), 8 U.S.C. 1182(N): Labor Condition Application (2026). <https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title8-section1182&num=0&edition=prelim>
- Improving Wage Protections for the Temporary and Permanent Employment of Certain Foreign Nationals in the United States (2026). <https://www.federalregister.gov/documents/2026/03/27/2026-06017/improving-wage-protections-for-the-temporary-and-permanent-employment-of-certain-foreign-nationals>
- Ioannidis, J. P. A. (2005). Why Most Published Research Findings Are False. *Plos Medicine*, 2(8). <https://doi.org/10.1371/journal.pmed.0020124>
- Kaufman, S., Rosset, S., Perlich, C., & Stitelman, O. (2012). Leakage in data mining: Formulation, detection, and avoidance. *ACM Transactions on Knowledge Discovery from Data*, 6(4), 1–21. <https://doi.org/10.1145/2382577.2382579>
- Koehler V. Infosys Technologies Limited, Inc. (July 13, 2026). <https://law.justia.com/cases/federal/appellate-courts/ca7/25-2272/25-2272-2026-07-13.html>
- Labor Condition Application for Nonimmigrant Workers, Form ETA-9035 and 9035e (2026). <https://flag.dol.gov/programs/LCA>
- Lipsitch, M., Tchetgen Tchetgen, E., & Cohen, T. (2010). Negative Controls: A Tool for Detecting Confounding and Bias in Observational Studies. *Epidemiology*, 21(3), 383–388. <https://doi.org/10.1097/EDE.0b013e3181d61eeb>
- Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., & Ha, D. (2024). *The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery*. arXiv. <https://doi.org/10.48550/arXiv.2408.06292>
- López de Prado, M. (2018). *Advances in Financial Machine Learning*. John Wiley & Sons.
- Mahajan, P., Morales, N., Shih, K., Chen, M., & Brinatti, A. (2025). *The Impact of Immigration on Firms and Workers: Insights from the H-1B Lottery* (No. Working Paper 24-04R). Federal Reserve Bank of Richmond. <https://doi.org/10.21144/wp24-04>
- Majumder, B. P., Surana, H., Agarwal, D., Hazra, S., Sabharwal, A., & Clark, P. (2024). Position: Data-driven Discovery with Large Generative Models. *Proceedings of the 41st International Conference on Machine Learning (ICML 2024)*, 235, 34350–34382. <https://proceedings.mlr.press/v235/majumder24a.html>
- Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press. <https://nlp.stanford.edu/IR-book/information-retrieval-book.html>
- Mayda, A. M., Ortega, F., Peri, G., Shih, K., & Sparber, C. (2018). The effect of the H-1B quota on the employment and selection of foreign-born labor. *European Economic Review*, 108, 105–128. <https://doi.org/10.1016/j.euroecorev.2018.06.010>
- McCullagh, P., & Nelder, J. A. (1989). *Generalized Linear Models* (2nd ed.). Chapman and Hall.
- McKinney, W. (2010). Data Structures for Statistical Computing in Python. *Proceedings of the 9th Python in Science Conference*, 56–61. <https://doi.org/10.25080/Majora-92bf1922-00a>
- Mithas, S., & Lucas, H. C. (2010). Are Foreign IT Workers Cheaper? U.S. Visa Policies and Compensation of Information Technology Professionals. *Management Science*, 56(5), 745–765. <https://doi.org/10.1287/mnsc.1100.1149>
- Modernization of Regulation s-K Items 101, 103, And 105 (2020). <https://www.sec.gov/files/rules/final/2020/33-10825.pdf>
- Nelder, J. A., & Wedderburn, R. W. M. (1972). Generalized Linear Models. *Journal of the Royal Statistical Society. Series a (General)*, 135(3), 370–384. <https://doi.org/10.2307/2344614>
- Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115(11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>
- Ornstein, J. T., Blasingame, E. N., & Truscott, J. S. (2025). How to train your stochastic parrot: large language models for political texts. *Political Science Research and Methods*, 13(2), 264–281. <https://doi.org/10.1017/psrm.2024.64>
- Palmer V. Cognizant Technology Solutions Corp. (October 4, 2024). <https://www.courtlistener.com/docket/8410282/christy-palmer-v-cognizant-technology-solutions-corporation/>
- Pangakis, N., Wolken, S., & Fasching, N. (2023). *Automated Annotation with Generative AI Requires Validation*. arXiv. <https://doi.org/10.48550/arXiv.2306.00176>
- Peri, G., Shih, K., & Sparber, C. (2015). STEM Workers, H-1B Visas, and Productivity in US Cities. *Journal of Labor Economics*, 33(S1), S225–S255. <https://doi.org/10.1086/679061>
- Petersen, M. A. (2009). Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. *The Review of Financial Studies*, 22(1), 435–480. <https://doi.org/10.1093/rfs/hhn053>
- Preston, J. (2015, June 3). Pink Slips at Disney. But First, Training Foreign Replacements. *The New York Times*. <https://www.nytimes.com/2015/06/04/us/last-task-after-layoff-at-disney-train-foreign-replacements.html>
- Prevailing Wage Determination Policy Guidance: Nonagricultural Immigration Programs (2009). https://www.dol.gov/sites/dolgov/files/ETA/offc/pdfs/NPWHC_Guidance_Revised_11_2009.pdf
- Raasveldt, M., & Mühleisen, H. (2019). DuckDB: an Embeddable Analytical Database. *SIGMOD '19: Proceedings of the 2019 International Conference on Management of Data*, 1981–1984. <https://doi.org/10.1145/3299869.3320212>

- Rathje, S., Mirea, D.-M., Sucholutsky, I., Marjeh, R., Robertson, C. E., & Van Bavel, J. J. (2024). GPT is an effective tool for multilingual psychological text analysis. *Proceedings of the National Academy of Sciences*, 121(34). <https://doi.org/10.1073/pnas.2308950121>
- Reiss, M. V. (2023). *Testing the Reliability of ChatGPT for Text Annotation and Classification: A Cautionary Remark*. arXiv. <https://doi.org/10.48550/arXiv.2304.11085>
- Restriction on Entry of Certain Nonimmigrant Workers (2025). <https://www.federalregister.gov/documents/2025/09/24/2025-18601/restriction-on-entry-of-certain-nonimmigrant-workers>
- Rhoades, S. A. (1993). The Herfindahl-Hirschman Index. *Federal Reserve Bulletin*, 79(3), 188–189. https://fraser.stlouisfed.org/files/docs/publications/FRB/pages/1990-1994/33101_1990-1994.pdf
- Santos Silva, J. M. C., & Tenreiro, S. (2006). The Log of Gravity. *The Review of Economics and Statistics*, 88(4), 641–658. <https://doi.org/10.1162/rest.88.4.641>
- Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and Statistical Modeling with Python. *Proceedings of the 9th Python in Science Conference*, 92–96. <https://doi.org/10.25080/Majora-92bf1922-011>
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- Si, C., Yang, D., & Hashimoto, T. (2025,). Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers. *The Thirteenth International Conference on Learning Representations (ICLR 2025)*. https://proceedings.iclr.cc/paper_files/paper/2025/hash/ea94957d81b1c1caf87ef5319fa6b467-Abstract-Conference.html
- Simmons, J. C. (2019). *Towards a Postmodal Movement for Computer Musicians* [University of California, Irvine]. <https://escholarship.org/uc/item/72b184qj>
- Simmons, J. C. (2025a, December 23). *Get In Loser: We're Exposing the AI Bubble*. <https://www.drjoshcsimmons.com/writing/get-in-loser-were-exposing-the-ai>
- Simmons, J. C. (2025b, December 27). *Nvidia Just Paid 20 Billion for a Company That Missed Its Revenue Target by 75%*. <https://www.drjoshcsimmons.com/writing/nvidia-just-paid-20-billion-for-a>
- Simmons, J. C. (2026a, February 13). *Hundreds of American Express Employees Reached Out to Me This Week. They're All Saying the Same Thing*. <https://www.drjoshcsimmons.com/writing/hundreds-of-amex-employees-reached-out>
- Simmons, J. C. (2026b, June 8). *Fake and Expensive: The Future Is Here*. <https://www.drjoshcsimmons.com/writing/fake-and-expensive-the-future-is-here>
- Simmons, J. C. (2026c, June 15). *You have a job. Why are you still broke?.* <https://www.drjoshcsimmons.com/writing/you-have-a-job-why-are-you-still-broke>
- Simmons, J. C. (2026d, September 19). *It Said No 6,336 Times. Then It Picked White*. <https://www.drjoshcsimmons.com/writing/it-refused-then-i-made-it-choose>
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant. *Psychological Science*, 22(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>
- Slaight V. Tata Consultancy Services, Ltd. (January 8, 2021). <https://cdn.ca9.uscourts.gov/datastore/memoranda/2021/01/08/19-16806.pdf>
- State of California V. Mullin (June 8, 2026). <https://www.courtlistener.com/docket/72031571/state-of-california-v-noem/>
- Storey, J. D. (2002). A Direct Approach to False Discovery Rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3), 479–498. <https://doi.org/10.1111/1467-9868.00346>
- Storey, J. D., & Tibshirani, R. (2003). Statistical significance for genome wide studies. *Proceedings of the National Academy of Sciences*, 100(16), 9440–9445. <https://doi.org/10.1073/pnas.1530509100>
- Strengthening Wage Protections for the Temporary and Permanent Employment of Certain Aliens in the United States (2020). <https://www.federalregister.gov/documents/2020/10/08/2020-22132/strengthening-wage-protections-for-the-temporary-and-permanent-employment-of-certain-alien-in-the>
- Strengthening Wage Protections for the Temporary and Permanent Employment of Certain Immigrants and Non-Immigrants in the United States, Implementation of Vacatur (2021). <https://www.federalregister.gov/documents/2021/12/13/2021-26660/strengthening-wage-protections-for-the-temporary-and-permanent-employment-of-certain-immigrants-and>
- Thibodeau, P. (2015). Southern California Edison IT Workers 'Beyond Furious' Over H-1B Replacements. *Computerworld*. <https://www.computerworld.com/article/1628728/southern-california-edison-it-workers-beyond-furious-over-h-1b-replacements-2.html>
- Thompson, S. B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics*, 99(1), 1–10. <https://doi.org/10.1016/j.jfineco.2010.08.016>
- Tirole, J. (1988). *The Theory of Industrial Organization*. The MIT Press.
- Turner, P. S. (2022). High-Skilled Immigration and the Labor Market: Evidence from the H-1B Visa Program. *Journal of Policy Analysis and Management*, 41(1), 92–130. <https://doi.org/10.1002/pam.22347>
- Törnberg, P. (2024). Best Practices for Text Annotation with Large Language Models. *Sociologica*, 18(2), 67–85. <https://doi.org/10.6092/issn.1971-8853/19461>
- U.S. Bureau of Labor Statistics. (2026). *Handbook of Methods: Occupational Employment and Wage Statistics*. <https://www.bls.gov/opub/hom/oews/>
- U.S. Citizenship and Immigration Services. (2025, July 7). *Understanding Our H-1B Employer Data Hub*. <https://www.uscis.gov/tools/reports-and-studies/h-1b-employer-data-hub/understanding-our-h-1b-employer-data-hub>

- U.S. Citizenship and Immigration Services. (2026,). *H-1B Employer Data Hub*. <https://www.uscis.gov/tools/reports-and-studies/h-1b-employer-data-hub>
- U.S. Department of Justice, Office of Public Affairs. (2021). *Justice, Labor Departments Reach Settlements with Facebook Resolving Claims of Discrimination Against U.S. Workers and Potential Regulatory Recruitment Violations*. <https://www.justice.gov/archives/opa/pr/justice-labor-departments-reach-settlements-facebook-resolving-claims-discrimination-against>
- U.S. Department of Justice, Office of Public Affairs. (2023). *Justice Department Secures 25 Million Landmark Agreement with Apple to Resolve Employment Discrimination Allegations Based on Citizenship Status*. <https://www.justice.gov/archives/opa/pr/justice-department-secures-25-million-landmark-agreement-apple-resolve-employment>
- U.S. Department of Justice, Office of Public Affairs. (2026b). *Civil Rights Division Obtains Settlement with Company that Discouraged U.S. Workers from Applying for Jobs*. <https://www.justice.gov/opa/pr/civil-rights-division-obtains-settlement-company-discouraged-us-workers-applying-jobs>
- U.S. Department of Justice, Office of Public Affairs. (2026a). *Civil Rights Division Secures Settlement with OpenAI for Discriminating Against U.S. Workers*. <https://www.justice.gov/opa/pr/civil-rights-division-secures-settlement-openai-discriminating-against-us-workers>
- U.S. Department of Justice, U.S. Attorney's Office, Eastern District of Texas. (2013). *Indian Corporation Pays Record Amount to Settle Allegations of Systemic Visa Fraud and Abuse of Immigration Processes*. <https://www.justice.gov/usao-edtx/pr/indian-corporation-pays-record-amount-settle-allegations-systemic-visa-fraud-and-abuse>
- U.S. Department of Labor, Employment and Training Administration, Office of Foreign Labor Certification. (2026a). *Public Disclosure File: LCA (H-1B, H-1B1 and E-3), Form ETA-9035*. https://www.dol.gov/sites/dolgov/files/ETA/oflc/pdfs/FY26Q3/LCA_Record_Layout_FY2026_Q3.pdf
- U.S. Department of Labor, Employment and Training Administration, Office of Foreign Labor Certification. (2026c). *Public Disclosure File: PERM, Form ETA-9089*. https://www.dol.gov/sites/dolgov/files/ETA/oflc/pdfs/FY26Q3/PERM_Record_Layout_FY2026_Q3.pdf
- U.S. Department of Labor, Employment and Training Administration, Office of Foreign Labor Certification. (2026b,). *Performance Data*. <https://www.dol.gov/agencies/eta/foreign-labor/performance>
- U.S. Department of Labor, Office of Inspector General. (2020). *Overview of Vulnerabilities and Challenges in Foreign Labor Certification Programs* (Nos. 6–21–1–3–321). U.S. Department of Labor, Office of Inspector General. <https://www.oig.dol.gov/public/reports/oa/2021/06-21-001-03-321.pdf>
- U.S. Equal Employment Opportunity Commission. (2016). *EEOC Enforcement Guidance on National Origin Discrimination*. <https://www.eeoc.gov/laws/guidance/eeoc-enforcement-guidance-national-origin-discrimination>
- U.S. Equal Employment Opportunity Commission. (2023). *Strategic Enforcement Plan Fiscal Years 2024-2028*. <https://www.eeoc.gov/strategic-enforcement-plan-fiscal-years-2024-2028>
- U.S. Government Accountability Office. (2011). *H-1B Visa Program: Reforms Are Needed to Minimize the Risks and Costs of Current Program* (Nos. GAO-11–26). U.S. Government Accountability Office. <https://www.gao.gov/products/gao-11-26>
- U.S. Government Accountability Office. (2021). *Department of Labor, Employment and Training Administration: Strengthening Wage Protections for the Temporary and Permanent Employment of Certain Aliens in the United States* (No. B-332870). U.S. Government Accountability Office. <https://www.gao.gov/products/b-332870>
- Unfair Immigration-Related Employment Practices (1986). <https://www.law.cornell.edu/uscode/text/8/1324b>
- Unlawful Employment Practices (1964). <https://www.law.cornell.edu/uscode/text/42/2000e-2>
- Wagenmakers, E.-J., Wetzels, R., Borsboom, D., Maas, H. L. J. van der, & Kievit, R. A. (2012). An Agenda for Purely Confirmatory Research. *Perspectives on Psychological Science*, 7(6), 632–638. <https://doi.org/10.1177/1745691612463078>
- Wang, Q., Downey, D., Ji, H., & Hope, T. (2024). SciMON: Scientific Inspiration Machines Optimized for Novelty. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 279–299. <https://doi.org/10.18653/v1/2024.acl-long.18>
- Wedderburn, R. W. M. (1974). Quasi-likelihood functions, generalized linear models, and the Gauss-Newton method. *Biometrika*, 61(3), 439–447. <https://doi.org/10.1093/biomet/61.3.439>
- Weighted Selection Process for Registrants and Petitioners Seeking to File Cap-Subject H-1b Petitions (2025). <https://www.federalregister.gov/documents/2025/12/29/2025-23853/weighted-selection-process-for-registrants-and-petitioners-seeking-to-file-cap-subject-h-1b>
- Weinberger, K., Dasgupta, A., Langford, J., Smola, A., & Attenberg, J. (2009). Feature hashing for large scale multi-task learning. *ICML '09: Proceedings of the 26th Annual International Conference on Machine Learning*, 1113–1120. <https://doi.org/10.1145/1553374.1553516>
- White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica*, 48(4), 817–838. <https://doi.org/10.2307/1912934>
- Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., & Yang, D. (2024). Can Large Language Models Transform Computational Social Science?. *Computational Linguistics*, 50(1), 237–291. https://doi.org/10.1162/coli_a_00502

A All 270 full-panel regressions

Effects are standardized per SD of the predictor (and per SD of the outcome for Gaussian models; log rate for counts; log odds for contraction). D: discovery (Oct 2019 to Sep 2023); V: validation (Oct 2023 to Jan 2025); H: sealed holdout (Feb 2025 to Jun 2026), fit only for the 39 pre-registered specifications (marked ✓). NC p: negative-control p-value (same predictor on the past outcome). q: Benjamini-Hochberg q-value across all 270 discovery tests.

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
academic medical signature	application contraction	1	53,457	−0.198	<0.001	0.142	−0.207	0.014	-	-
academic medical signature	application contraction	2	52,389	−0.270	<0.001	<0.001	−0.217	0.027	-	-
✓ academic medical signature	application contraction	3	51,548	−0.305	<0.001	0.011	−0.417	<0.001	−0.281	0.026
academic medical signature	application count	1	53,457	+0.030	0.271	0.595	−0.003	0.961	-	-
academic medical signature	application count	2	52,389	+0.038	0.343	0.167	+0.059	0.373	-	-
academic medical signature	application count	3	51,548	+0.025	0.575	0.003	+0.037	0.625	-	-
✓ academic medical signature	location diversity change	1	50,329	−0.110	<0.001	0.704	−0.141	<0.001	−0.116	0.002
✓ academic medical signature	location diversity change	2	48,365	−0.098	<0.001	<0.001	−0.141	<0.001	−0.111	0.025
✓ academic medical signature	location diversity change	3	47,147	−0.099	<0.001	<0.001	−0.124	0.003	−0.097	0.041
✓ academic medical signature	role diversity change	1	50,329	+0.188	<0.001	0.513	+0.173	<0.001	+0.180	<0.001
✓ academic medical signature	role diversity change	2	48,365	+0.215	<0.001	0.065	+0.174	<0.001	+0.186	<0.001
✓ academic medical signature	role diversity change	3	47,147	+0.206	<0.001	0.006	+0.208	<0.001	+0.210	<0.001
academic medical signature	wage change	1	50,328	−0.106	<0.001	0.007	−0.087	0.024	-	-
✓ academic medical signature	wage change	2	48,363	−0.124	<0.001	0.379	−0.142	0.004	−0.063	0.173
✓ academic medical signature	wage change	3	47,144	−0.121	0.001	0.858	−0.164	0.003	−0.063	0.241
ai data specialization	application contraction	1	53,457	+0.042	0.005	0.686	+0.017	0.574	-	-
ai data specialization	application contraction	2	52,389	+0.022	0.245	0.817	−0.017	0.655	-	-
ai data specialization	application contraction	3	51,548	+0.005	0.848	0.439	+0.025	0.421	-	-
ai data specialization	application count	1	53,457	−0.019	0.082	0.362	−0.015	0.142	-	-
ai data specialization	application count	2	52,389	−0.030	0.007	0.022	−0.028	0.169	-	-
ai data specialization	application count	3	51,548	−0.044	0.004	0.914	−0.022	0.221	-	-
ai data specialization	location diversity change	1	50,329	−0.028	<0.001	0.040	−0.030	0.016	-	-
ai data specialization	location diversity change	2	48,365	−0.019	0.018	0.568	−0.019	0.114	-	-
ai data specialization	location diversity change	3	47,147	−0.020	0.017	0.820	−0.031	0.005	-	-
ai data specialization	role diversity change	1	50,329	−0.005	0.512	0.237	−0.032	0.006	-	-
ai data specialization	role diversity change	2	48,365	−0.013	0.058	0.608	−0.019	0.074	-	-
ai data specialization	role diversity change	3	47,147	−0.009	0.254	0.619	−0.021	0.107	-	-
✓ ai data specialization	wage change	1	50,328	+0.042	<0.001	0.075	+0.035	<0.001	+0.047	<0.001
✓ ai data specialization	wage change	2	48,363	+0.045	<0.001	0.255	+0.042	<0.001	+0.047	<0.001
✓ ai data specialization	wage change	3	47,144	+0.046	<0.001	0.603	+0.060	<0.001	+0.049	<0.001
calendar novdec repeat batch	application contraction	1	53,457	+0.022	0.136	0.772	+0.024	0.273	-	-
calendar novdec repeat batch	application contraction	2	52,389	−0.008	0.754	0.567	−0.005	0.749	-	-
calendar novdec repeat batch	application contraction	3	51,548	+0.045	0.047	0.830	+0.017	0.040	-	-
calendar novdec repeat batch	application count	1	53,457	−0.011	0.425	-	−0.026	0.008	-	-
calendar novdec repeat batch	application count	2	52,389	−0.011	0.505	0.176	+0.020	0.534	-	-
calendar novdec repeat batch	application count	3	51,548	+0.010	0.711	0.236	+0.043	0.263	-	-
calendar novdec repeat batch	location diversity change	1	50,329	−0.003	0.693	0.565	−0.014	0.024	-	-
calendar novdec repeat batch	location diversity change	2	48,365	+0.008	0.238	0.258	−0.002	0.807	-	-
calendar novdec repeat batch	location diversity change	3	47,147	−0.001	0.890	0.702	−0.004	0.598	-	-
calendar novdec repeat batch	role diversity change	1	50,329	−0.011	0.159	0.468	−0.009	0.135	-	-
calendar novdec repeat batch	role diversity change	2	48,365	−0.005	0.434	0.082	−0.014	0.019	-	-
calendar novdec repeat batch	role diversity change	3	47,147	−0.027	0.007	0.278	−0.022	0.043	-	-
calendar novdec repeat batch	wage change	1	50,328	−0.009	0.116	0.036	+0.003	0.619	-	-
calendar novdec repeat batch	wage change	2	48,363	−0.009	0.178	0.905	+0.011	0.018	-	-
calendar novdec repeat batch	wage change	3	47,144	+0.004	0.421	0.472	+0.005	0.085	-	-

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
clinical healthcare signature	application contraction	1	53,457	+0.043	0.070	0.648	+0.045	0.328	-	-
clinical healthcare signature	application contraction	2	52,389	+0.062	0.015	0.539	+0.057	0.213	-	-
clinical healthcare signature	application contraction	3	51,548	+0.075	0.004	0.042	+0.071	0.254	-	-
clinical healthcare signature	application count	1	53,457	-0.000	0.994	0.196	-0.026	0.210	-	-
clinical healthcare signature	application count	2	52,389	-0.016	0.263	0.508	+0.007	0.735	-	-
clinical healthcare signature	application count	3	51,548	-0.019	0.269	0.472	-0.049	0.013	-	-
clinical healthcare signature	location diversity change	1	50,329	+0.006	0.586	0.750	-0.013	0.238	-	-
clinical healthcare signature	location diversity change	2	48,365	+0.013	0.267	0.028	-0.002	0.909	-	-
clinical healthcare signature	location diversity change	3	47,147	+0.017	0.151	0.195	-0.005	0.752	-	-
✓ clinical healthcare signature	role diversity change	1	50,329	-0.049	<0.001	<0.001	-0.050	<0.001	-0.045	0.005
clinical healthcare signature	role diversity change	2	48,365	-0.052	<0.001	0.009	-0.055	0.020	-	-
clinical healthcare signature	role diversity change	3	47,147	-0.055	<0.001	0.612	-0.038	0.105	-	-
clinical healthcare signature	wage change	1	50,328	+0.033	0.266	0.460	+0.070	0.070	-	-
clinical healthcare signature	wage change	2	48,363	+0.038	0.215	0.132	+0.022	0.440	-	-
clinical healthcare signature	wage change	3	47,144	+0.039	0.266	0.095	+0.006	0.829	-	-
entry level share	application contraction	1	53,457	+0.001	0.975	0.308	-0.027	0.436	-	-
entry level share	application contraction	2	52,389	+0.003	0.940	0.529	-0.022	0.557	-	-
entry level share	application contraction	3	51,548	-0.028	0.192	0.503	+0.019	0.597	-	-
entry level share	application count	1	53,457	-0.007	0.540	0.008	-0.021	0.195	-	-
entry level share	application count	2	52,389	-0.012	0.457	0.066	-0.032	0.186	-	-
entry level share	application count	3	51,548	-0.004	0.870	0.286	-0.032	0.161	-	-
entry level share	location diversity change	1	50,329	-0.030	0.003	0.158	-0.032	0.009	-	-
entry level share	location diversity change	2	48,365	-0.030	0.002	0.915	-0.033	0.029	-	-
✓ entry level share	location diversity change	3	47,147	-0.029	0.004	0.755	-0.043	0.001	-0.027	0.044
✓ entry level share	role diversity change	1	50,329	+0.057	<0.001	0.623	+0.066	<0.001	+0.059	<0.001
✓ entry level share	role diversity change	2	48,365	+0.051	<0.001	0.736	+0.048	<0.001	+0.057	<0.001
entry level share	role diversity change	3	47,147	+0.053	<0.001	0.782	+0.043	0.005	-	-
✓ entry level share	wage change	1	50,328	+0.054	<0.001	0.081	+0.067	0.001	+0.075	<0.001
✓ entry level share	wage change	2	48,363	+0.069	<0.001	0.887	+0.066	0.002	+0.081	<0.001
✓ entry level share	wage change	3	47,144	+0.060	<0.001	0.933	+0.071	0.002	+0.087	0.002
generic title mass	application contraction	1	53,457	-0.029	0.159	0.109	-0.000	0.993	-	-
generic title mass	application contraction	2	52,389	-0.026	0.266	0.443	-0.038	0.210	-	-
generic title mass	application contraction	3	51,548	-0.030	0.194	0.702	-0.007	0.862	-	-
generic title mass	application count	1	53,457	+0.039	0.009	0.520	+0.027	0.133	-	-
generic title mass	application count	2	52,389	+0.056	0.015	0.002	+0.049	0.051	-	-
generic title mass	application count	3	51,548	+0.057	0.018	0.154	+0.060	0.055	-	-
generic title mass	location diversity change	1	50,329	+0.001	0.940	0.345	+0.016	0.213	-	-
generic title mass	location diversity change	2	48,365	-0.002	0.906	0.247	+0.024	0.102	-	-
generic title mass	location diversity change	3	47,147	+0.001	0.963	0.213	+0.022	0.209	-	-
generic title mass	role diversity change	1	50,329	+0.028	0.007	0.004	+0.015	0.206	-	-
generic title mass	role diversity change	2	48,365	+0.030	0.009	0.935	+0.025	0.021	-	-
generic title mass	role diversity change	3	47,147	+0.025	0.025	0.250	+0.021	0.099	-	-
generic title mass	wage change	1	50,328	-0.013	0.242	0.360	-0.013	0.322	-	-
generic title mass	wage change	2	48,363	-0.013	0.310	0.604	-0.005	0.735	-	-
generic title mass	wage change	3	47,144	-0.003	0.870	0.381	-0.006	0.709	-	-
hybrid title presence	application contraction	1	53,457	+0.024	0.159	0.451	-0.001	0.967	-	-
hybrid title presence	application contraction	2	52,389	+0.022	0.192	0.844	+0.030	0.242	-	-
hybrid title presence	application contraction	3	51,548	+0.023	0.159	0.995	+0.064	0.064	-	-
hybrid title presence	application count	1	53,457	-0.005	0.671	0.800	-0.007	0.529	-	-
hybrid title presence	application count	2	52,389	-0.014	0.468	0.206	-	-	-	-

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
hybrid title presence	application count	3	51,548	−0.013	0.501	0.033	−0.038	0.021	-	-
hybrid title presence	location diversity change	1	50,329	+0.002	0.874	0.995	−0.011	0.338	-	-
hybrid title presence	location diversity change	2	48,365	−0.002	0.848	0.969	−0.017	0.091	-	-
hybrid title presence	location diversity change	3	47,147	+0.006	0.457	0.214	−0.024	0.015	-	-
hybrid title presence	role diversity change	1	50,329	−0.001	0.951	0.367	+0.004	0.666	-	-
hybrid title presence	role diversity change	2	48,365	+0.001	0.890	0.384	−0.001	0.896	-	-
hybrid title presence	role diversity change	3	47,147	+0.008	0.296	0.865	−0.011	0.387	-	-
hybrid title presence	wage change	1	50,328	−0.004	0.633	0.337	−0.006	0.631	-	-
hybrid title presence	wage change	2	48,363	−0.005	0.594	0.815	−0.008	0.467	-	-
hybrid title presence	wage change	3	47,144	−0.011	0.184	0.690	+0.006	0.645	-	-
management heavy mix	application contraction	1	53,457	−0.017	0.441	0.163	+0.072	0.058	-	-
management heavy mix	application contraction	2	52,389	−0.009	0.696	0.953	+0.034	0.407	-	-
management heavy mix	application contraction	3	51,548	+0.002	0.913	0.555	+0.081	0.054	-	-
management heavy mix	application count	1	53,457	+0.027	0.042	0.018	+0.005	0.812	-	-
management heavy mix	application count	2	52,389	+0.023	0.338	0.704	+0.025	0.304	-	-
management heavy mix	application count	3	51,548	+0.014	0.584	<0.001	+0.029	0.162	-	-
management heavy mix	location diversity change	1	50,329	+0.035	<0.001	0.770	+0.025	0.067	-	-
management heavy mix	location diversity change	2	48,365	+0.028	0.002	0.278	+0.039	0.020	-	-
management heavy mix	location diversity change	3	47,147	+0.023	0.010	0.927	+0.043	0.005	-	-
management heavy mix	role diversity change	1	50,329	+0.006	0.547	0.878	−0.026	0.046	-	-
management heavy mix	role diversity change	2	48,365	−0.002	0.872	0.812	−0.012	0.381	-	-
management heavy mix	role diversity change	3	47,147	+0.000	0.981	0.261	−0.022	0.043	-	-
management heavy mix	wage change	1	50,328	−0.029	0.002	0.592	−0.020	0.287	-	-
management heavy mix	wage change	2	48,363	−0.045	<0.001	0.826	−0.028	0.185	-	-
management heavy mix	wage change	3	47,144	−0.033	0.001	0.142	−0.019	0.136	-	-
minority second major group	application contraction	1	53,457	+0.017	0.526	0.300	−0.017	0.509	-	-
minority second major group	application contraction	2	52,389	−0.004	0.890	0.430	−0.038	0.183	-	-
minority second major group	application contraction	3	51,548	+0.024	0.345	0.656	−0.027	0.353	-	-
minority second major group	application count	1	53,457	−0.011	0.651	0.487	+0.028	0.336	-	-
minority second major group	application count	2	52,389	−0.012	0.540	0.371	+0.010	0.721	-	-
minority second major group	application count	3	51,548	−0.022	0.266	0.671	−0.020	0.442	-	-
minority second major group	location diversity change	1	50,329	−0.016	0.032	0.849	+0.007	0.540	-	-
minority second major group	location diversity change	2	48,365	−0.010	0.325	0.689	+0.004	0.779	-	-
minority second major group	location diversity change	3	47,147	−0.012	0.218	0.224	+0.011	0.535	-	-
✓ minority second major group	role diversity change	1	50,329	+0.027	0.002	0.178	+0.045	0.003	+0.045	0.009
✓ minority second major group	role diversity change	2	48,365	+0.029	0.002	0.365	+0.053	<0.001	+0.040	0.002
minority second major group	role diversity change	3	47,147	+0.024	0.008	0.555	+0.027	0.037	-	-
minority second major group	wage change	1	50,328	+0.013	0.155	0.953	+0.009	0.455	-	-
minority second major group	wage change	2	48,363	+0.018	0.030	0.670	+0.000	0.995	-	-
minority second major group	wage change	3	47,144	+0.012	0.269	0.042	+0.009	0.560	-	-
multi seniority title span	application contraction	1	53,457	−0.029	0.184	0.545	−0.040	0.214	-	-
multi seniority title span	application contraction	2	52,389	−0.012	0.564	0.185	−0.104	0.002	-	-
multi seniority title span	application contraction	3	51,548	+0.004	0.870	0.055	−0.042	0.270	-	-
multi seniority title span	application count	1	53,457	−0.022	0.230	0.835	−0.015	0.372	-	-
multi seniority title span	application count	2	52,389	−0.058	0.047	0.320	−0.024	0.128	-	-
multi seniority title span	application count	3	51,548	−0.057	0.049	0.025	+0.004	0.886	-	-
multi seniority title span	location diversity change	1	50,329	+0.003	0.842	0.968	+0.017	0.182	-	-

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
multi seniority title span	location diversity change	2	48,365	−0.001	0.906	0.478	+0.005	0.602	-	-
multi seniority title span	location diversity change	3	47,147	+0.001	0.940	0.056	+0.008	0.561	-	-
multi seniority title span	role diversity change	1	50,329	+0.040	<0.001	0.482	+0.032	0.005	-	-
✓ multi seniority title span	role diversity change	2	48,365	+0.031	<0.001	0.294	+0.045	0.002	+0.026	0.059
multi seniority title span	role diversity change	3	47,147	+0.033	<0.001	0.344	+0.023	0.132	-	-
multi seniority title span	wage change	1	50,328	+0.040	<0.001	0.074	+0.030	0.035	-	-
multi seniority title span	wage change	2	48,363	+0.040	<0.001	0.119	+0.041	0.011	-	-
✓ multi seniority title span	wage change	3	47,144	+0.045	<0.001	0.086	+0.053	0.003	+0.045	<0.001
repeated title batching	application contraction	1	53,457	+0.018	0.266	0.636	+0.003	0.911	-	-
repeated title batching	application contraction	2	52,389	−0.004	0.876	0.156	−0.016	0.512	-	-
repeated title batching	application contraction	3	51,548	+0.001	0.961	0.157	−0.025	0.372	-	-
repeated title batching	application count	1	53,457	+0.004	0.842	0.051	+0.044	0.088	-	-
repeated title batching	application count	2	52,389	−0.003	0.906	0.149	+0.069	0.015	-	-
repeated title batching	application count	3	51,548	−0.004	0.816	0.044	+0.047	0.065	-	-
repeated title batching	location diversity change	1	50,329	−0.010	0.153	0.424	−0.001	0.916	-	-
repeated title batching	location diversity change	2	48,365	−0.002	0.874	0.453	+0.003	0.765	-	-
repeated title batching	location diversity change	3	47,147	−0.004	0.749	0.410	+0.002	0.846	-	-
repeated title batching	role diversity change	1	50,329	−0.028	0.001	0.097	−0.011	0.249	-	-
repeated title batching	role diversity change	2	48,365	−0.020	0.013	0.610	−0.010	0.350	-	-
repeated title batching	role diversity change	3	47,147	−0.026	<0.001	0.557	−0.011	0.354	-	-
repeated title batching	wage change	1	50,328	−0.015	0.039	0.043	−0.015	0.113	-	-
repeated title batching	wage change	2	48,363	−0.009	0.304	0.786	−0.010	0.502	-	-
repeated title batching	wage change	3	47,144	−0.013	0.136	0.668	−0.005	0.677	-	-
senior title share	application contraction	1	53,457	−0.030	0.151	0.115	−0.014	0.631	-	-
senior title share	application contraction	2	52,389	−0.049	0.009	0.767	−0.019	0.569	-	-
senior title share	application contraction	3	51,548	−0.042	0.032	0.656	−0.055	0.109	-	-
senior title share	application count	1	53,457	+0.005	0.788	0.895	−0.001	0.862	-	-
senior title share	application count	2	52,389	+0.008	0.764	0.685	−0.008	0.668	-	-
senior title share	application count	3	51,548	+0.006	0.861	0.575	+0.008	0.619	-	-
senior title share	location diversity change	1	50,329	+0.012	0.245	0.095	+0.044	0.006	-	-
senior title share	location diversity change	2	48,365	+0.019	0.072	0.070	+0.041	0.009	-	-
senior title share	location diversity change	3	47,147	+0.013	0.248	0.078	+0.038	0.035	-	-
senior title share	role diversity change	1	50,329	+0.008	0.540	0.225	+0.010	0.441	-	-
senior title share	role diversity change	2	48,365	+0.010	0.343	0.005	+0.005	0.705	-	-
senior title share	role diversity change	3	47,147	+0.011	0.390	0.575	+0.012	0.357	-	-
senior title share	wage change	1	50,328	+0.000	0.994	0.858	+0.024	0.064	-	-
senior title share	wage change	2	48,363	−0.001	0.913	0.763	+0.021	0.154	-	-
senior title share	wage change	3	47,144	+0.004	0.788	0.554	+0.031	0.024	-	-
single job family	application contraction	1	53,457	−0.000	0.985	0.852	+0.020	0.643	-	-
single job family	application contraction	2	52,389	−0.006	0.861	0.844	+0.011	0.794	-	-
single job family	application contraction	3	51,548	−0.009	0.792	0.845	+0.024	0.517	-	-
single job family	application count	1	53,457	+0.019	0.164	0.412	+0.011	0.521	-	-
single job family	application count	2	52,389	+0.019	0.447	0.073	−0.009	0.709	-	-
single job family	application count	3	51,548	+0.015	0.674	0.101	−0.011	0.600	-	-
single job family	location diversity change	1	50,329	+0.016	0.099	0.682	+0.034	0.007	-	-
single job family	location diversity change	2	48,365	+0.017	0.118	0.020	+0.030	0.067	-	-
single job family	location diversity change	3	47,147	+0.027	0.011	0.300	+0.023	0.110	-	-
single job family	role diversity change	1	50,329	−0.035	0.004	0.632	−0.016	0.349	-	-
single job family	role diversity change	2	48,365	−0.035	0.007	0.651	−0.020	0.321	-	-

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
single job family	role diversity change	3	47,147	−0.025	0.044	0.480	−0.022	0.174	-	-
single job family	wage change	1	50,328	−0.057	<0.001	0.106	−0.033	0.028	-	-
single job family	wage change	2	48,363	−0.051	<0.001	0.295	−0.027	0.081	-	-
single job family	wage change	3	47,144	−0.048	<0.001	0.393	−0.057	0.007	-	-
staffing consultancy signature	application contraction	1	53,457	−0.017	0.625	0.515	−0.026	0.555	-	-
staffing consultancy signature	application contraction	2	52,389	+0.020	0.649	0.997	−0.070	0.164	-	-
staffing consultancy signature	application contraction	3	51,548	−0.027	0.457	0.253	−0.027	0.645	-	-
staffing consultancy signature	application count	1	53,457	+0.002	0.940	0.139	+0.002	0.935	-	-
staffing consultancy signature	application count	2	52,389	+0.025	0.619	0.568	+0.063	0.198	-	-
staffing consultancy signature	application count	3	51,548	−0.007	0.838	0.563	+0.044	0.282	-	-
✓ staffing consultancy signature	location diversity change	1	50,329	+0.065	<0.001	0.232	+0.075	0.005	+0.009	0.706
✓ staffing consultancy signature	location diversity change	2	48,365	+0.068	<0.001	0.608	+0.105	0.001	+0.064	0.009
staffing consultancy signature	location diversity change	3	47,147	+0.075	<0.001	0.355	+0.077	0.008	-	-
✓ staffing consultancy signature	role diversity change	1	50,329	−0.080	<0.001	0.729	−0.072	<0.001	−0.065	0.004
staffing consultancy signature	role diversity change	2	48,365	−0.069	<0.001	0.236	−0.047	0.024	-	-
staffing consultancy signature	role diversity change	3	47,147	−0.065	0.001	0.018	−0.046	0.033	-	-
staffing consultancy signature	wage change	1	50,328	−0.040	0.011	0.288	−0.033	0.055	-	-
staffing consultancy signature	wage change	2	48,363	−0.049	0.007	0.647	−0.075	<0.001	-	-
staffing consultancy signature	wage change	3	47,144	−0.044	0.011	0.056	−0.058	0.007	-	-
technical to business transition	application contraction	1	53,457	+0.026	0.065	0.235	+0.048	0.149	-	-
technical to business transition	application contraction	2	52,389	+0.044	0.012	0.539	+0.083	0.017	-	-
technical to business transition	application contraction	3	51,548	+0.047	0.019	0.970	+0.087	0.006	-	-
technical to business transition	application count	1	53,457	−0.013	0.202	0.822	−0.033	0.077	-	-
technical to business transition	application count	2	52,389	−0.031	0.106	0.335	+0.007	0.738	-	-
technical to business transition	application count	3	51,548	−0.018	0.397	0.114	+0.006	0.848	-	-
technical to business transition	location diversity change	1	50,329	+0.026	0.004	0.007	−0.000	0.979	-	-
technical to business transition	location diversity change	2	48,365	+0.020	0.016	0.084	−0.003	0.787	-	-
technical to business transition	location diversity change	3	47,147	+0.012	0.159	0.814	+0.021	0.033	-	-
technical to business transition	role diversity change	1	50,329	−0.014	0.124	0.247	−0.021	0.072	-	-
technical to business transition	role diversity change	2	48,365	−0.020	0.019	0.350	−0.035	0.024	-	-
technical to business transition	role diversity change	3	47,147	−0.023	0.013	0.420	−0.025	0.037	-	-
✓ technical to business transition	wage change	1	50,328	+0.044	<0.001	0.275	+0.031	<0.001	+0.020	0.028
technical to business transition	wage change	2	48,363	+0.038	<0.001	0.991	+0.022	0.118	-	-
✓ technical to business transition	wage change	3	47,144	+0.040	<0.001	0.036	+0.038	<0.001	+0.048	0.001
title family heterogeneity	application contraction	1	53,457	+0.033	0.108	0.164	+0.025	0.563	-	-
title family heterogeneity	application contraction	2	52,389	+0.037	0.116	0.359	+0.033	0.459	-	-
title family heterogeneity	application contraction	3	51,548	+0.021	0.468	0.438	+0.044	0.273	-	-
title family heterogeneity	application count	1	53,457	−0.022	0.025	0.682	+0.009	0.034	-	-
title family heterogeneity	application count	2	52,389	−0.024	0.070	0.042	+0.014	0.278	-	-
title family heterogeneity	application count	3	51,548	−0.020	0.432	0.020	−0.016	0.433	-	-
title family heterogeneity	location diversity change	1	50,329	+0.001	0.906	0.728	−0.008	0.520	-	-

Predicate	Outcome	Lag	n	Effect (D)	q (D)	NC p	Effect (V)	p (V)	Effect (H)	p (H)
title family heterogeneity	location diversity change	2	48,365	−0.014	0.172	0.146	−0.012	0.287	-	-
title family heterogeneity	location diversity change	3	47,147	−0.012	0.238	0.721	−0.013	0.296	-	-
title family heterogeneity	role diversity change	1	50,329	+0.001	0.961	0.889	+0.005	0.698	-	-
title family heterogeneity	role diversity change	2	48,365	−0.004	0.842	0.841	−0.010	0.520	-	-
title family heterogeneity	role diversity change	3	47,147	−0.007	0.619	0.126	−0.009	0.572	-	-
✓ title family heterogeneity	wage change	1	50,328	+0.070	<0.001	0.021	+0.042	0.003	+0.024	0.151
✓ title family heterogeneity	wage change	2	48,363	+0.073	<0.001	0.385	+0.046	<0.001	+0.056	0.006
✓ title family heterogeneity	wage change	3	47,144	+0.074	<0.001	0.076	+0.064	<0.001	+0.073	<0.001
title wage inflation	application contraction	1	53,457	−0.010	0.636	0.451	−0.022	0.404	-	-
title wage inflation	application contraction	2	52,389	−0.034	0.026	0.986	−0.024	0.347	-	-
title wage inflation	application contraction	3	51,548	−0.024	0.153	0.467	−0.006	0.842	-	-
title wage inflation	application count	1	53,457	+0.012	0.065	0.872	−0.006	0.071	-	-
title wage inflation	application count	2	52,389	+0.030	0.070	0.706	+0.010	0.578	-	-
title wage inflation	application count	3	51,548	+0.029	0.153	0.138	+0.045	0.113	-	-
title wage inflation	location diversity change	1	50,329	−0.009	0.266	0.016	−0.016	0.096	-	-
title wage inflation	location diversity change	2	48,365	−0.021	0.015	0.069	−0.020	0.032	-	-
title wage inflation	location diversity change	3	47,147	−0.015	0.126	0.505	−0.023	0.022	-	-
✓ title wage inflation	role diversity change	1	50,329	+0.030	<0.001	0.697	+0.035	0.002	+0.026	0.020
title wage inflation	role diversity change	2	48,365	+0.024	0.004	0.554	+0.036	0.018	-	-
title wage inflation	role diversity change	3	47,147	+0.024	0.002	0.547	+0.046	0.009	-	-
✓ title wage inflation	wage change	1	50,328	+0.044	<0.001	0.278	+0.075	<0.001	+0.064	<0.001
✓ title wage inflation	wage change	2	48,363	+0.048	<0.001	0.736	+0.081	<0.001	+0.081	<0.001
✓ title wage inflation	wage change	3	47,144	+0.050	<0.001	0.874	+0.070	<0.001	+0.090	<0.001
wage title grade discord	application contraction	1	53,457	+0.003	0.932	0.084	−0.037	0.074	-	-
wage title grade discord	application contraction	2	52,389	+0.005	0.874	0.878	−0.033	0.380	-	-
wage title grade discord	application contraction	3	51,548	−0.006	0.861	0.379	−0.050	0.231	-	-
wage title grade discord	application count	1	53,457	−0.018	0.230	0.177	−0.029	0.030	-	-
wage title grade discord	application count	2	52,389	−0.038	0.043	0.196	−0.016	0.461	-	-
wage title grade discord	application count	3	51,548	−0.043	0.028	0.055	−0.030	0.061	-	-
✓ wage title grade discord	location diversity change	1	50,329	−0.036	<0.001	0.316	−0.050	0.001	−0.024	0.061
✓ wage title grade discord	location diversity change	2	48,365	−0.036	<0.001	0.026	−0.047	0.002	−0.052	0.001
✓ wage title grade discord	location diversity change	3	47,147	−0.035	<0.001	0.008	−0.052	<0.001	−0.048	0.003
wage title grade discord	role diversity change	1	50,329	+0.048	<0.001	0.025	+0.051	0.013	-	-
wage title grade discord	role diversity change	2	48,365	+0.047	<0.001	0.554	+0.038	0.008	-	-
✓ wage title grade discord	role diversity change	3	47,147	+0.047	<0.001	0.387	+0.055	0.003	+0.032	0.077
wage title grade discord	wage change	1	50,328	+0.055	<0.001	0.050	+0.029	0.143	-	-
wage title grade discord	wage change	2	48,363	+0.078	<0.001	0.671	+0.042	0.205	-	-
wage title grade discord	wage change	3	47,144	+0.078	<0.001	0.409	+0.064	0.016	-	-

B Predicate definitions

The exact instruction and criteria text sent to the evaluation model with each employer-month. The state presented to the model contained the fields `period`, `applications_1m`, `applications_6m`, `application_growth_3m`, `median_wage`, `wage_observed`, `industry`, `role_titles` (title histogram), `role_distribution` (SOC histogram), `location_distribution`, `role_entropy`, `location_entropy`, `technical_share`, `unique_roles`, `unique_locations`, and a clock note; it did not contain the employer name.

B.1 semantic.staffing_consultancy_signature.v2

Origin: new. Judge whether the `role_titles` map reads like an IT staffing, consulting, or outsourcing operation: generic interchangeable technology titles (for example programmer analyst, software developer, systems analyst, consultant, technical lead, QA analyst) repeated at volume, often across several `location_distribution` keys, with few or no product, research, or management titles. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Most title mass is generic, interchangeable technology labor titles typical of staffing or consulting placements. **False:** Titles are specific to a product, research, clinical, academic, or corporate function, or too few to judge.

B.2 semantic.academic_medical_signature.v2

Origin: new. Judge whether the `role_titles` map reads like a university, hospital, or research institution: professor, lecturer, postdoctoral, research scientist, physician, resident, fellow, clinical, nurse, or similar titles. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Most title mass is academic, medical, or research-institution titles. **False:** Titles are predominantly corporate, technology, or commercial, or too few to judge.

B.3 semantic.generic_title_mass.v2

Origin: new. Judge whether the `role_titles` are generic and uninformative about the actual work (for example engineer, analyst, consultant, associate, specialist, developer, manager with no qualifier) rather than specific. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Most title mass carries no domain, product, or specialty qualifier. **False:** Most titles name a specialty, domain, product, or level.

B.4 semantic.senior_title_share.v2

Origin: new. Judge whether senior-level tokens (senior, sr, lead, principal, staff, architect, director, head, chief, vp) account for the majority of `role_titles` mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: More than half of title mass carries a senior-level token. **False:** Half or less of title mass carries a senior-level token.

B.5 semantic.entry_level_share.v2

Origin: new. Judge whether entry-level tokens (junior, jr, associate, intern, trainee, entry, level i or 1, analyst i, graduate) account for a substantial share (at least a third) of `role_titles` mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: At least a third of title mass is entry-level titles. **False:** Less than a third of title mass is entry-level titles.

B.6 semantic.ai_data_specialization.v2

Origin: new. Judge whether titles in `role_titles` referencing machine learning, artificial intelligence, data science, deep learning, NLP, computer vision, or LLMs account for at least a quarter of title mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: At least a quarter of title mass is AI, ML, or data-science titles. **False:** Less than a quarter of title mass is AI, ML, or data-science titles.

B.7 semantic.management_heavy_mix.v2

Origin: new. Judge whether management titles (manager, director, head, vp, vice president, chief, officer, supervisor) account for at least a third of `role_titles` mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: At least a third of title mass is management titles. **False:** Less than a third of title mass is management titles.

B.8 semantic.single_job_family.v2

Origin: new. Judge whether essentially all `role_titles` belong to one job family (for example all software engineering, all finance, all nursing) as opposed to spanning two or more distinct families. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: All or nearly all title mass sits in one job family. **False:** Title mass spans two or more distinct job families.

B.9 semantic.title_wage_inflation.v2

Origin: new. Judge whether senior or specialist titles dominate role_titles while median_wage is modest for such titles (a rough guide: below 90000 for senior technology or finance titles, below 70000 for other senior titles). If wage_observed is 0 or median_wage is missing, mark false. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Senior or specialist titles dominate and median_wage is modest for those titles. **False:** Titles are not senior-dominated, wage is commensurate, or wage is missing.

B.10 semantic.clinical_healthcare_signature.v2

Origin: new. Judge whether the role_titles map reads like direct clinical healthcare delivery: physician, nurse, therapist, pharmacist, dentist, technologist, physical therapist, medical resident, or similar. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Most title mass is clinical healthcare delivery titles. **False:** Titles are predominantly non-clinical, or too few to judge.

B.11 semantic.title_family_heterogeneity.v2

Origin: rewrite. Judge whether role_titles span several distinct job families (for example engineering plus finance plus operations) rather than one family with minor variants. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Title mass is spread across three or more distinct job families. **False:** Title mass sits in one or two closely related families.

B.12 semantic.technical_to_business_transition.v2

Origin: rewrite. Judge whether role_titles mix technical titles (engineer, developer, scientist, analyst with a technical qualifier) with business titles (sales, marketing, product, finance, operations, HR) such that business titles hold at least a third of mass while technical titles still hold at least a third. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Both technical and business titles each hold at least a third of title mass. **False:** One side holds under a third of title mass, or titles are too few to judge.

B.13 semantic.wage_title_grade_discord.v2

Origin: rewrite. Judge whether the seniority implied by role_titles disagrees with median_wage: senior-heavy titles with a low median_wage, or junior-heavy titles with a high median_wage. If wage_observed is 0 mark false. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Title seniority and median_wage clearly disagree. **False:** Title seniority and median_wage are consistent, or wage is missing.

B.14 semantic.minority_second_major_group.v2

Origin: rewrite. Judge whether role_distribution (SOC codes) has one dominant major group (first two digits) plus a second major group holding between 10% and 40% of mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: A second SOC major group holds between 10% and 40% of mass beside a dominant group. **False:** Mass is in one major group, or spread with no dominant group, or too few codes.

B.15 semantic.calendar_novdec_repeat_batch.v2

Origin: rewrite. Judge whether period is November or December and role_titles is dominated (over half of mass) by one repeated title string. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Period is 11 or 12 and one title string holds over half of title mass. **False:** Period is another month, or no single title holds over half of mass.

B.16 semantic.repeated_title_batching.v2

Origin: rewrite. Judge whether one exact title string holds a clear majority (over 60%) of role_titles mass while applications_1m is at least 5. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: One title string holds over 60% of title mass with at least 5 applications. **False:** No single title string holds over 60%, or fewer than 5 applications.

B.17 semantic.multi_seniority_title_span.v2

Origin: rewrite. Judge whether role_titles contain the same base title at two or more seniority levels (for example software engineer, senior software engineer, staff software engineer). Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: At least one base title appears at two or more seniority levels. **False:** No base title appears at multiple seniority levels.

B.18 semantic.hybrid_title_presence.v2

Origin: rewrite. Judge whether role_titles include hybrid titles that fuse two functions (for example software engineer / data scientist, product manager / engineer, developer analyst, engineering manager) holding at least a fifth of title mass. Use only the fields supplied on this t-state. Do not infer hiring, layoffs, nationality, intent, future periods, or real-company identity. Applications are administrative counts, not workers or visas.

True: Hybrid two-function titles hold at least a fifth of title mass. **False:** Hybrid titles are absent or hold under a fifth of title mass.

C Prevalence of every predicate by year

Share of applications (percent) from employer-months flagged (probability above 0.5), employer-months with five or more applications. 2019 covers October to December; 2026 covers January to June.

Predicate	2019	2020	2021	2022	2023	2024	2025	2026 H1
staffing consultancy signature	67.0	56.6	49.4	43.7	41.6	37.2	34.6	31.8
academic medical signature	6.7	6.9	7.3	8.2	10.0	10.2	11.5	11.3
generic title mass	20.8	21.1	21.4	18.4	17.8	15.9	13.4	10.4
senior title share	24.8	21.7	25.2	25.1	25.0	28.5	28.0	26.1
entry level share	12.3	13.4	11.0	10.7	12.1	11.6	12.9	15.3
ai data specialization	0.9	2.3	3.8	3.7	3.6	5.2	5.9	8.2
management heavy mix	22.0	14.7	14.4	14.2	16.1	15.6	16.3	13.9
single job family	53.1	45.5	36.5	33.0	34.2	32.2	29.6	26.2
title wage inflation	12.3	6.1	2.5	3.1	1.9	2.7	2.5	0.3
clinical healthcare signature	0.5	1.0	1.3	1.3	1.8	1.8	2.8	3.3
title family heterogeneity	26.5	33.3	36.9	41.5	43.1	44.7	45.9	51.8
technical to business transition	3.9	4.3	5.1	3.7	5.4	4.8	5.6	7.3
wage title grade discord	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
minority second major group	56.6	60.3	67.0	67.3	68.3	67.5	70.8	73.4
calendar novdec repeat batch	8.2	1.3	1.3	1.6	1.7	1.8	2.0	0.0
repeated title batching	7.1	8.3	7.1	6.9	10.5	9.5	7.1	8.6
multi seniority title span	69.7	69.8	68.8	72.2	70.9	70.3	72.5	75.7
hybrid title presence	6.8	7.4	8.0	7.0	7.2	7.0	7.4	7.1

D Employers behind each predicate

Employers ranked by applications falling in flagged months, with the share of the employer's own applications so flagged, the employer's mean predicate probability across all its months, and its median offered wage. Employers with at least 12 panel months.

D.1 Staffing or consulting signature

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
cognizant technology solutions us corp	96,543	100%	0.83	\$101,358
tata consultancy services limited	62,001	100%	0.89	\$80,912
infosys limited	46,792	100%	0.92	\$91,666
deloitte consulting llp	29,687	100%	0.79	\$125,550
accenture llp	23,567	97%	0.77	\$116,055
google llc	23,560	36%	0.50	\$173,000
capgemini america inc	22,348	100%	0.92	\$102,461
amazon web services, inc.	20,124	91%	0.61	\$144,500
wipro limited	16,710	100%	0.88	\$80,808
hcl america inc	15,264	100%	0.80	\$110,240
ibm corporation	14,802	100%	0.80	\$114,920
compunnel software group, inc	14,513	100%	0.91	\$106,000
ernst & young u.s. llp	11,949	20%	0.42	\$136,472
mphasis corporation	9,825	100%	0.84	\$114,975
tech mahindra (americas), inc	9,593	100%	0.87	\$92,061

D.2 Academic or medical signature

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
grandison management, inc.	5,207	100%	0.62	\$69,670
university of michigan	3,682	100%	0.90	\$65,000
mayo clinic	3,415	100%	0.91	\$74,816
the leland stanford, jr university	2,720	100%	0.92	\$80,170
johns hopkins university	2,349	100%	0.94	\$70,000
trustees of the university of pennsylvania	2,286	100%	0.92	\$73,125
columbia university	2,175	100%	0.94	\$75,002
washington university	2,165	100%	0.93	\$65,000
university of california, san francisco	2,044	100%	0.93	\$77,327
icahn school of medicine at mount sinai	1,997	100%	0.93	\$87,720
management health systems, llc	1,993	100%	0.81	\$56,576
university of pittsburgh	1,984	100%	0.94	\$65,063
st. jude children's research hospital	1,940	100%	0.87	\$75,174
emory university	1,913	100%	0.94	\$70,268
national institutes of health, hhs	1,862	100%	0.93	\$93,625

D.3 AI or data specialization

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
wal-mart associates, inc.	14,891	67%	0.49	\$142,176
capital one services, llc	3,507	69%	0.56	\$110,240
facebook, inc.	3,383	40%	0.42	\$170,000
adobe inc.	3,012	48%	0.38	\$162,088
linkedin corporation	2,692	35%	0.37	\$157,477
bytedance inc.	2,405	49%	0.41	\$193,875
paypal, inc.	2,251	32%	0.36	\$159,546
aetna resources llc	2,012	86%	0.67	\$135,000
tiktok inc.	1,947	68%	0.49	\$200,905
capital one, national association	1,706	97%	0.79	\$103,584
caremark llc	1,402	59%	0.53	\$133,376
snap inc.	1,288	58%	0.54	\$151,549
expedia, inc.	1,174	37%	0.40	\$133,598
federal express corporation	1,123	62%	0.39	\$112,374
discover products inc.	1,096	98%	0.85	\$95,000

D.4 Senior or specialist titles at a modest wage

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
tata consultancy services limited	16,378	26%	0.41	\$80,912
cognizant technology solutions us corp	14,606	15%	0.15	\$101,358
wipro limited	7,711	46%	0.44	\$80,808
deloitte & touche llp	4,206	42%	0.30	\$91,150
hcl america inc	1,428	9%	0.09	\$110,240
deloitte tax llp	1,230	25%	0.32	\$87,558
ust global inc	1,078	21%	0.20	\$95,781
cummins inc.	1,077	20%	0.31	\$88,800
capital one services, llc	970	19%	0.15	\$110,240
persistent systems, inc.	737	38%	0.30	\$92,560
discover products inc.	732	66%	0.35	\$95,000
mindtree limited	683	17%	0.14	\$94,746
american airlines, inc.	349	17%	0.20	\$94,912
genpact llc	308	16%	0.09	\$117,980
slk america, inc.	283	47%	0.34	\$85,250

D.5 Entry-level heavy

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
jpmorgan chase & co.	17,033	80%	0.63	\$145,000
amazon.com services llc	13,743	15%	0.27	\$147,350
cognizant technology solutions us corp	11,832	12%	0.31	\$101,358
wipro limited	9,433	56%	0.47	\$80,808
goldman sachs & co. llc	8,101	94%	0.74	\$127,000
accenture llp	6,179	25%	0.41	\$116,055
amazon development center u.s., inc.	4,992	41%	0.42	\$153,600
amazon web services, inc.	4,306	19%	0.31	\$144,500
goldman sachs services llc	3,996	89%	0.66	\$140,000
ust global inc	3,828	76%	0.62	\$95,781
blackrock financial management, inc.	3,246	90%	0.79	\$115,000
mckinsey & company, inc. united states	2,650	62%	0.49	\$165,000
deloitte & touche llp	2,476	25%	0.19	\$91,150
johns hopkins university	2,349	100%	0.70	\$70,000
ebay inc.	2,164	42%	0.45	\$171,064

D.6 Generic titles

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
tata consultancy services limited	58,424	94%	0.61	\$80,912
deloitte consulting llp	24,370	82%	0.60	\$125,550
infosys limited	22,050	47%	0.46	\$91,666
google llc	21,834	33%	0.48	\$173,000
compunnel software group, inc	9,721	67%	0.55	\$106,000
ernst & young u.s. llp	9,666	16%	0.41	\$136,472
pricewaterhousecoopers advisory services llc	6,517	97%	0.82	\$173,750
capgemini america inc	5,574	25%	0.47	\$102,461
goldman sachs & co. llc	5,258	61%	0.61	\$127,000
cgi technologies and solutions inc.	3,871	75%	0.57	\$110,000
blackrock financial management, inc.	3,559	99%	0.87	\$115,000
kpmg llp	3,454	100%	0.85	\$131,300
mindtree limited	3,169	77%	0.60	\$94,746
ibm corporation	3,125	21%	0.35	\$114,920
goldman sachs services llc	2,899	64%	0.64	\$140,000

D.7 Management-heavy

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
cognizant technology solutions us corp	94,149	98%	0.77	\$101,358

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
deloitte consulting llp	15,949	54%	0.56	\$125,550
infosys limited	15,299	33%	0.41	\$91,666
jpmorgan chase & co.	15,025	71%	0.58	\$145,000
accenture llp	9,234	38%	0.52	\$116,055
hcl america inc	6,253	41%	0.51	\$110,240
pricewaterhousecoopers advisory services llc	6,010	90%	0.84	\$173,750
deloitte & touche llp	5,092	51%	0.60	\$91,150
bank of america n.a.	4,732	100%	0.83	\$143,610
hcl america, inc.	4,584	64%	0.55	\$97,508
capital one services, llc	3,794	74%	0.67	\$110,240
kpmg llp	3,307	96%	0.86	\$131,300
goldman sachs services llc	3,272	72%	0.74	\$140,000
goldman sachs & co. llc	3,258	38%	0.60	\$127,000
wal-mart associates, inc.	3,183	14%	0.38	\$142,176

D.8 Senior-title majority

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
infosys limited	24,769	53%	0.54	\$91,666
deloitte consulting llp	20,026	67%	0.55	\$125,550
wal-mart associates, inc.	19,724	89%	0.69	\$142,176
cognizant technology solutions us corp	14,714	15%	0.38	\$101,358
wipro limited	12,414	74%	0.57	\$80,808
qualcomm technologies, inc.	10,324	76%	0.76	\$127,829
jpmorgan chase & co.	9,098	43%	0.49	\$145,000
tata consultancy services limited	8,465	14%	0.37	\$80,912
citibank, n.a.	7,616	100%	0.74	\$147,972
fidelity technology group, llc	6,696	100%	0.88	\$106,943
linkedin corporation	6,538	84%	0.70	\$157,477
pricewaterhousecoopers advisory services llc	5,898	88%	0.66	\$173,750
hcl america inc	5,404	35%	0.37	\$110,240
capital one services, llc	5,029	98%	0.79	\$110,240
deloitte & touche llp	4,833	48%	0.52	\$91,150

D.9 Clinical healthcare

Employer (normalized legal name)	Apps in flagged months	Share of own apps	Mean p	Median wage
grandison management, inc.	5,207	100%	0.96	\$69,670
mayo clinic	2,224	65%	0.52	\$74,816
management health systems, llc	1,993	100%	0.85	\$56,576
montefiore medical center	944	99%	0.86	\$99,178
northwell health	815	89%	0.71	\$104,999
university of pittsburgh physicians	689	100%	0.87	\$260,082
the university of iowa	613	35%	0.32	\$74,999
icahn school of medicine at mount sinai	564	28%	0.31	\$87,720
ochsner clinic foundation	480	94%	0.83	\$67,392
arup laboratories	477	78%	0.58	\$66,648
cleveland clinic foundation	449	81%	0.63	\$97,500
henry ford health system	445	42%	0.38	\$94,078
mass general brigham incorporated	426	99%	0.76	\$76,900
cvs rx services inc.	396	100%	0.93	\$136,739
adventist health system/sunbelt, inc.	386	90%	0.73	\$69,347

E Measurement reliability

Test-retest of the evaluation model on 80 randomly sampled employer-months (five or more applications), six predicates, each state scored in two independent requests.

Predicate	n	mean	sd	exact match	mean abs diff	r	flips across 0.5
isolated_activity_pulse.v1	80	0.042	0.102	80%	0.002	0.998	0%
repeated_title_batching.v1	80	0.184	0.185	41%	0.01	0.996	0%
location_volume_discord.v1	80	0.42	0.164	10%	0.026	0.98	5%
title_family_heterogeneity.v1	80	0.85	0.2	48%	0.009	0.997	0%
senior_only_title_mix.v1	80	0.601	0.435	65%	0.005	1	0%
title_soc_cardinality_mismatch.v1	80	0.693	0.276	20%	0.031	0.984	3.8%
All	480			44%	0.014	0.998	1.5%

Measurement noise (half the mean squared paired difference) is 0.25% of the variance across states.

F The earlier design: chronology and diagnostics

F.1 Chronology

Reconstructed from the project’s git history, run configurations, report files, and the session transcripts of the coding agent (OpenAI Codex CLI) that built and operated it. Times are Phoenix local.

When	Event
2026-09-18, 22:55	Written specification for a Grok-proposes, Jev-measures, Python-tests pipeline over the H-1B record; private repository created. Founding rule, carried into the README: models invent and criticize hypotheses, the evaluation model measures semantic variables, deterministic code performs statistics.
09-18 to 09-19 09:05	Architecture built: BigQuery snapshot, immutable state builder (361,353 employer-months, chronological discovery/ validation/sealed-holdout split), gateway client, credential sandbox, resumable driver, retry and circuit-breaker logic.
09-19, 06:06 to 06:43	Snapshot taken (8.34 GiB, 11 tables). Unattended spend authorized under \$20. Acceptance run: 100 states, 5 predicates, 500 measurements, 20 hypotheses; an integration test, reported Jev cost \$0.
09-19, 10:32 to 18:14	Screen: 20 hypotheses on 120 employer-months per split. 0 promising, 15 rejected, 5 inconclusive (perfect-separation fit failures).
09-19, 18:53 to 09-20 02:29	Ten-round campaign: 300 discovery and 300 validation employer-months per round on fresh employer panels, per-round discovery FDR and validation alpha of 0.005 (the original 5% budget divided by ten). Paused at 8 of 10 rounds against a spending guard: 160 hypotheses, 0 validated.
09-20, 02:32 to 03:12	Threshold audit: relaxing 0.005 to 0.05 still yields zero survivors; a typical fit used about 262 observations against about 65 design terms. A separate, model-free audit of low-volume filers in the raw record is started and produces two investigative leads (an education-sector denial-rate gap; 69 filings with a wage range below the stated prevailing wage).
09-20, 06:17 to 06:42	Retry policy set to immediate first retry then progressive backoff, ignoring Retry-After. Budget accounting corrected (the model calls were free or prepaid; only source-transfer cost was cash). Rounds 9 and 10 resumed and ten follow-up rounds queued unattended.
09-20, 15:01	Overnight result: 20 rounds, 395 hypotheses, 0 validated; 189 reversed sign between discovery and validation, 233 failed to improve out-of-time prediction.
09-20, 15:46	Forty-eight-round expansion launched on the now-exhausted frozen panels, reusing cached measurements; results labeled exploratory by construction.
09-21, 04:57	Expansion result: 48 rounds, 927 hypotheses, 888 rejected, 39 inconclusive, 0 qualifying. Agent’s verdict: more iterations of the same design are not paying off.
09-21, 08:02 onward	Redesign described in the Methods section: full-panel measurement of 18 title-reading predicates, fixed 270-cell grid, one FDR family, validation replication, pre-registered holdout.

Run	Rounds	Hypotheses	Promising	Rejected	Inconclusive
screen-20260919-001	1	20	0	15	5
campaign-20260919-001	10	200	0	193	7
campaign-followup-20260920-001	10	195	0	182	13
campaign-expansion-20260920-001	48	927	0	888	39
Total	69	1,342	0	1,278	64

The design decisions that produced the result, and who made them: the 300-observation per-split samples (agent, to preserve employer-disjoint panels across rounds); the per-round 0.005 thresholds (agent, to avoid ten independent 5% budgets); interaction and rare-combination hypotheses (agent, in response to an instruction to widen the search); the 48-round escalation on reused panels (author); predicates compiled into literal threshold rules with guardrails against inference (carried from the original specification’s rule against using the evaluation model as a chatbot); and the sealed holdout (built into the state builder from the first session and never opened). The measurement cost was \$0 throughout under the gateway’s promotion; the binding constraint, as the author noted at the time, was not money but the design.

F.2 Final campaign statistics

927 hypotheses in 48 rounds; 888 rejected, 39 inconclusive (unstable coefficient covariance), 0 promising. Discovery p-value deciles: 163, 99, 90, 73, 83, 73, 78, 82, 78, 89. Negative-control p-value deciles: 156, 108, 94, 75, 92, 87, 68, 70, 73, 65. Validation p-value deciles: 148, 103, 81, 85, 91, 75, 82, 77, 72, 74. Discovery p below 0.005: 27 negative control below 0.005: 26. Sign agreed with the generator’s expected direction in 426 of 908 discovery and validation signs agreed in 457 of 888. Interaction hypotheses: 64%. Distinct semantic predicates: 85. Out-of-time predictive gain positive in 316 of 888.

Outcome	tests	median SE	MDE (p<0.005, 80%)	share able to detect 0.1 SD
role diversity change	168	0.082	0.3	3%
application count	330	0.198	0.72	0%
application contraction	147	0.498	1.82	0%
wage change	124	0.043	0.16	39%
location diversity change	139	0.088	0.32	4%

Failure reasons recorded (a hypothesis can carry several): direction mismatch (665); discovery fdr failed (878); effect too small (504); no incremental out of time prediction (572); validation multiplicity failed (884); validation p failed (863); negative control failed (88); Unstable coefficient covariance (39).

F.3 Predicate degeneracy in the earlier campaign

Moments of every predicate the generator wrote, across all 242,020 cached measurements.

Predicate (original campaign)	n	mean	sd	share below 0.1	share above 0.9
wage_blank_on_multi_title.v1	1,800	0.031	0.003	100%	0%
senior_token_sparse_batch.v1	1,800	0.019	0.004	100%	0%
sparse_geo_split_batch.v1	1,200	0.016	0.005	100%	0%
wage_level_discontinuity.v1	1,200	0.079	0.007	99%	0%
midyear_isolated_batch.v1	1,200	0.020	0.007	100%	0%
compressed_wage_distribution.v1	1,052	0.390	0.011	0%	0%
majority_state_token_remainder.v1	600	0.266	0.019	0%	0%
location_map_unique_count_gap.v1	1,200	0.051	0.021	98%	0%
joint_location_encoding_gap.v1	3,000	0.052	0.024	96%	0%
geo_entropy_inconsistency.v1	3,445	0.043	0.028	100%	0%
mix_change_without_volume.v1	1,800	0.076	0.030	91%	0%
role_map_entropy_discord.v1	2,400	0.074	0.037	84%	0%
elevated_wage_early_narrow.v1	600	0.016	0.055	100%	0%
technical_composition_shift.v1	600	0.193	0.056	3%	0%
collapsed_single_title_geo.v1	3,600	0.026	0.065	99%	0%
entropy_cardinality_feasibility_gap.v1	1,800	0.150	0.069	16%	0%
two_state_single_occupation.v1	3,000	0.028	0.076	98%	0%
volume_window_mismatch.v1	3,600	0.271	0.087	0%	0%
volume_without_mix_change.v1	2,400	0.301	0.088	0%	0%
growth_volume_inconsistency.v1	4,647	0.176	0.088	1%	0%
elevated_median_wage_band.v1	2,254	0.496	0.094	0%	0%
early_stage_narrow_footprint.v1	2,849	0.047	0.102	97%	1%
technical_label_share_gap.v1	2,251	0.134	0.108	45%	0%
administrative_activity_plateau.v1	2,254	0.167	0.110	12%	0%
role_entropy_share_gap.v1	2,400	0.253	0.118	15%	0%
activity_burst_after_quiet.v1	3,444	0.216	0.120	5%	0%
equivalent_titles_elevated_wage.v1	2,400	0.068	0.121	86%	0%
synonym_title_multi_soc.v1	4,200	0.096	0.128	78%	0%
title_soc_family_discord.v1	1,800	0.175	0.137	37%	0%
fy_boundary_compressed_batch.v1	600	0.072	0.140	83%	1%
isolated_activity_pulse.v1	8,236	0.060	0.147	95%	2%
location_entropy_share_gap.v1	3,600	0.230	0.147	32%	0%
academic_year_start_alignment.v1	1,056	0.054	0.147	97%	2%
soc_mass_volume_gap.v1	1,200	0.182	0.155	42%	0%
calendar_novdec_repeat_batch.v1	1,800	0.068	0.159	91%	2%
calendar_q4_narrow_window.v1	600	0.047	0.163	97%	3%
location_volume_discord.v1	4,636	0.392	0.164	0%	0%
location_mass_volume_gap.v1	4,200	0.165	0.170	51%	0%
elevated_wage_nontechnical_titles.v1	1,200	0.070	0.179	92%	2%
soc_minor_group_fork.v1	4,800	0.305	0.180	19%	0%
fy_boundary_specialized_batch.v1	600	0.068	0.187	94%	3%
title_histogram_volume_gap.v1	1,800	0.267	0.187	18%	0%
balanced_two_state_split.v1	2,400	0.076	0.201	88%	4%
joint_title_location_concentration.v1	5,400	0.114	0.201	79%	1%
qualifier_only_title_ladder.v1	2,400	0.236	0.203	37%	0%
application_activity_acceleration.v1	1,991	0.360	0.210	28%	0%
q1_compressed_calendar_batch.v1	1,800	0.127	0.213	76%	2%
role_title_equivalence_cluster.v1	5,245	0.377	0.216	10%	3%
calendar_q2_narrow_window.v1	2,400	0.075	0.217	94%	5%
multi_state_same_title_echo.v1	4,200	0.140	0.218	74%	1%
hybrid_title_elevated_wage.v1	600	0.129	0.225	76%	1%
role_location_fragmentation.v1	3,781	0.372	0.231	23%	0%
geographic_expansion_burst.v1	2,245	0.373	0.233	27%	0%
adjacent_occupation_expansion.v1	1,646	0.756	0.247	7%	26%
role_portfolio_specialization.v1	3,787	0.544	0.250	1%	10%
repeated_title_batching.v1	6,444	0.238	0.250	34%	4%

Predicate (original campaign)	n	mean	sd	share below 0.1	share above 0.9
compressed_decision_month.v1	600	0.322	0.251	13%	2%
role_mix_reconfiguration.v1	3,000	0.820	0.252	6%	63%
midband_wage_repeated_title.v1	600	0.124	0.256	85%	6%
single_market_technical_batch.v1	600	0.111	0.259	89%	7%
wage_title_grade_discord.v1	3,000	0.237	0.261	48%	0%
hybrid_title_presence.v1	1,800	0.339	0.262	9%	4%
title_family_heterogeneity.v1	5,400	0.799	0.266	6%	58%
nonhub_cross_region_split.v1	600	0.252	0.268	44%	0%
elevated_wage_narrow_geo.v1	1,200	0.104	0.274	91%	9%
elevated_wage_volume_taper.v1	2,400	0.128	0.279	87%	2%
multi_seniority_title_span.v1	3,600	0.195	0.321	82%	13%
administrative_volume_taper.v1	6,000	0.223	0.327	71%	0%
geographic_concentration.v1	4,382	0.611	0.329	2%	32%
title_soc_cardinality_mismatch.v1	5,400	0.629	0.329	13%	25%
location_singleton_tail_spread.v1	4,200	0.229	0.330	64%	10%
technical_role_concentration.v1	3,186	0.726	0.344	10%	59%
adjacent_occupation_spread.v1	6,436	0.363	0.351	34%	4%
minority_second_major_group.v1	1,200	0.444	0.357	38%	2%
senior_title_share_elevation.v1	600	0.401	0.357	35%	12%
engineering_minor_group_divergence.v1	3,000	0.552	0.358	27%	10%
nontechnical_title_majority.v1	1,800	0.263	0.359	64%	15%
cross_region_state_split.v1	1,800	0.446	0.368	37%	7%
hub_state_mass_majority.v1	3,600	0.742	0.369	18%	66%
calendar_q1_spread_window.v1	1,200	0.225	0.372	77%	16%
occupation_family_shift.v1	5,231	0.536	0.393	26%	27%
even_multi_state_spread.v1	3,000	0.342	0.395	53%	17%
single_state_filing_dominance.v1	2,844	0.417	0.396	38%	27%
cross_major_group_spread.v1	5,238	0.523	0.397	36%	27%
technical_to_business_transition.v1	4,640	0.433	0.404	34%	31%
manager_plus_ic_title_mix.v1	1,800	0.369	0.417	60%	24%
location_footprint_shift.v1	2,400	0.686	0.427	30%	70%
senior_only_title_mix.v1	6,000	0.502	0.445	48%	46%
seniority_title_wage_alignment.v1	1,800	0.439	0.449	54%	37%

G Queries for the raw-record findings

Every query is verbatim as executed against the snapshot; results are in the text. Table paths are the snapshot's Parquet directories.

G.1 01_total_universe

```
select count(*) as total_rows,
       count(distinct case_number) as distinct_cases,
       count(distinct employer_name_norm) as distinct_employers,
       sum(is_certified) as certified_rows,
       sum(is_denied) as denied_rows,
       sum(is_withdrawn) as withdrawn_rows,
       sum(is_certified_withdrawn) as cert_withdrawn_rows
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
```

G.2 02_top10_50_100_concentration

```
with emp as (
  select employer_name_norm, count(*) as n
  from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/
*.parquet')
  where is_certified = 1
  group by 1
),
ranked as (
  select *, row_number() over (order by n desc) as rnk, sum(n) over () as total
  from emp
)
select
  sum(case when rnk <= 10 then n else 0 end) as top10,
  sum(case when rnk <= 50 then n else 0 end) as top50,
  sum(case when rnk <= 100 then n else 0 end) as top100,
  sum(case when rnk <= 500 then n else 0 end) as top500,
  max(total) as total_certified,
  count(*) as n_employers
from ranked
```

G.3 03_top25_employers_alltime_certified

```
select employer_name_norm, count(*) as certified_lcas,
       sum(workers_num) as total_workers_requested,
       round(avg(annual_wage_from),0) as avg_annual_wage_from
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1
group by 1
order by 2 desc
limit 25
```

G.4 04_top_employers_by_year_2025

```
select fiscal_year, employer_name_norm, count(*) as certified_lcas
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1 and fiscal_year in (2020,2023,2025,2026)
group by 1,2
qualify row_number() over (partition by fiscal_year order by count(*) desc) <= 15
order by fiscal_year, certified_lcas desc
```

G.5 05_wagelevel_overall_share

```
select raw_pw_wage_level, count(*) n,
       round(100.0*count(*)/sum(count(*) over(),1) as pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1
group by 1 order by 2 desc
```

G.6 06_wagelevel_trend_by_year

```
select fiscal_year,
       round(100.0*sum(case when raw_pw_wage_level in ('I','II') then 1 else 0 end)/count(*),1) as pct_level_1_2,
       round(100.0*sum(case when raw_pw_wage_level = 'I' then 1 else 0 end)/count(*),1) as pct_level_1,
       count(*) as n
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1 and fiscal_year between 2020 and 2025
group by 1 order by 1
```

G.7 07_staffing_pattern_employers

```
with flagged as (
  select *,
  case when
```

```

        regexp_matches(lower(employer_name_norm), '(consulting|technologies|infotech|info tech|solutions|systems|staffing|
outsourc|it services|software)')
    then 1 else 0 end as looks_staffing
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/
*.parquet')
    where is_certified = 1
)
select looks_staffing,
    count(*) as certified_lcas,
    count(distinct employer_name_norm) as n_employers,
    round(100.0*sum(case when raw_pw_wage_level in ('I','II') then 1 else 0 end)/count(*),1) as pct_level_1_2,
    round(avg(annual_wage_from),0) as avg_wage
from flagged
group by 1

```

G.8 08_same_title_mega_filers

```

select employer_name_norm, job_title_norm, fiscal_year, count(*) as n_filings, sum(workers_num) as total_workers
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1
group by 1,2,3
order by n_filings desc
limit 25

```

G.9 09_worksite_mega_batches

```

select case_number, worksite_city_norm, worksite_state_norm, worksite_workers_num, fiscal_year
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_worksite_rows/*.parquet')
where worksite_workers_num is not null
order by worksite_workers_num desc
limit 20

```

G.10 10_worksite_case_batches

```

select fiscal_year, worksite_city_norm, worksite_state_norm, count(*) as n_case_rows, sum(worksite_workers_num) as
total_workers
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_worksite_rows/*.parquet')
group by 1,2,3
order by total_workers desc nulls last
limit 20

```

G.11 11_quarterly_filing_trend

```

select fiscal_year, source_quarter, count(*) as certified_lcas
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1
group by 1,2
order by 1,2

```

G.12 12_yearly_totals_all_status

```

select fiscal_year, is_partial_year,
    count(*) as total_rows,
    sum(is_certified) as certified,
    sum(is_denied) as denied,
    sum(is_withdrawn) as withdrawn,
    sum(is_certified_withdrawn) as cert_withdrawn
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
group by 1,2
order by 1

```

G.13 13_geography_shift_state

```

select worksite_state_norm,
    sum(case when fiscal_year=2020 then 1 else 0 end) as fy2020,
    sum(case when fiscal_year=2025 then 1 else 0 end) as fy2025,
    sum(case when fiscal_year=2025 then 1 else 0 end) - sum(case when fiscal_year=2020 then 1 else 0 end) as delta
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified = 1 and worksite_state_norm is not null
group by 1
order by delta desc

```

G.14 14_FIXED_perm_vs_lca_by_company

```

with lca_agg as (
    select company_slug, count(*) as lca_certified
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/
*.parquet')
    where is_certified = 1 and company_slug is not null

```

```

    group by 1
),
perm_agg as (
    select company_slug, count(*) as perm_total, sum(is_certified) as perm_certified
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/
*.parquet')
    where company_slug is not null
    group by 1
)
select l.company_slug, l.lca_certified, coalesce(p.perm_certified,0) as perm_certified,
    round(l.lca_certified * 1.0 / nullif(coalesce(p.perm_certified,0),0),1) as lca_per_perm_ratio
from lca_agg l
left join perm_agg p on l.company_slug = p.company_slug
order by l.lca_certified desc
limit 40

```

G.15 14__perm_vs_lca_by_company

```

with lca_agg as (
    select company_slug, count(*) as lca_certified
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/
*.parquet')
    where is_certified = 1 and company_slug is not null
    group by 1
),
perm_agg as (
    select company_slug, count(*) as perm_certified
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_source_rows/*.parquet')
    where is_canonical_source = 1 and is_certified = 1 and company_slug is not null
    group by 1
)
select l.company_slug, l.lca_certified, coalesce(p.perm_certified,0) as perm_certified,
    round(l.lca_certified * 1.0 / nullif(coalesce(p.perm_certified,0),0),1) as lca_per_perm_ratio
from lca_agg l
left join perm_agg p on l.company_slug = p.company_slug
order by l.lca_certified desc
limit 40

```

G.16 15__FIXED_perm_cert_denial_rate_by_employer

```

select employer_name_norm,
    count(*) as total,
    sum(is_certified) as certified,
    sum(is_denied) as denied,
    sum(is_withdrawn) as withdrawn,
    round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct,
    round(100.0*sum(is_certified)/count(*),1) as cert_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/*.parquet')
group by 1
having count(*) >= 200
order by denial_rate_pct desc
limit 25

```

G.17 15_perm_cert_denial_rate_by_employer

```

with agg as (
    select employer_name_norm,
        count(*) as total,
        sum(is_certified) as certified,
        sum(is_denied) as denied,
        sum(is_withdrawn) as withdrawn
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_source_rows/*.parquet')
    where is_canonical_source = 1
    group by 1
    having count(*) >= 100
)
select *, round(100.0*denied/total,1) as denial_rate_pct, round(100.0*certified/total,1) as cert_rate_pct
from agg
order by denial_rate_pct desc
limit 25

```

G.18 15b__FIXED_perm_cert_denial_rate_lowest

```

select employer_name_norm,
    count(*) as total,
    sum(is_certified) as certified,

```

```

    sum(is_denied) as denied,
    round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/*.parquet')
group by 1
having count(*) >= 500
order by denial_rate_pct asc
limit 15

```

G.19 16_uscis_denial_rate_by_employer

```

select employer_name_norm,
    sum(total_approvals_num) as total_approvals,
    sum(total_denials_num) as total_denials,
    sum(total_approvals_num)+sum(total_denials_num) as total_decisions,
    round(100.0*sum(total_denials_num)/nullif(sum(total_approvals_num)+sum(total_denials_num),0),1) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/uscis_hlb_employer_records/*.parquet')
group by 1
having sum(total_approvals_num)+sum(total_denials_num) >= 200
order by denial_rate_pct desc
limit 25

```

G.20 17_uscis_denial_rate_trend_overall

```

select fiscal_year,
    sum(total_approvals_num) as total_approvals,
    sum(total_denials_num) as total_denials,
    round(100.0*sum(total_denials_num)/nullif(sum(total_approvals_num)+sum(total_denials_num),0),1) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/uscis_hlb_employer_records/*.parquet')
group by 1 order by 1

```

G.21 18_headcount_vs_lca_trend

```

with lca_agg as (
    select company_slug, fiscal_year, count(*) as lca_certified
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
    where is_certified=1 and company_slug is not null
    group by 1,2
)
select s.company_slug, s.raw_fiscal_year::int as fiscal_year, s.raw_total_employees, s.raw_us_employees, l.lca_certified
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/sec_employee_headcount/*.parquet') s
left join lca_agg l on s.company_slug = l.company_slug and s.raw_fiscal_year::int = l.fiscal_year
order by s.company_slug, fiscal_year

```

G.22 19_bls_wage_gap_ny_swe

```

with lca_ny_swe as (
    select annual_wage_from
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
    where is_certified=1 and soc_code_norm = '15-1252'
    and worksite_state_norm = 'ny'
    and worksite_city_norm in ('new york','newark','jersey city')
    and fiscal_year in (2023,2024)
    and annual_wage_from is not null
)
select count(*) as n_lca_rows,
    round(avg(annual_wage_from),0) as avg_lca_offered,
    round(median(annual_wage_from),0) as median_lca_offered,
    round(100.0*sum(case when annual_wage_from < 138400 then 1 else 0 end)/count(*),1) as pct_below_bls_median
from lca_ny_swe

```

G.23 19b_bls_wage_gap_dc_swe

```

with lca_dc_swe as (
    select annual_wage_from
    from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
    where is_certified=1 and soc_code_norm = '15-1252'
    and worksite_state_norm in ('dc','va','md')
    and worksite_city_norm in ('washington','arlington','alexandria')
    and fiscal_year in (2023,2024)
    and annual_wage_from is not null
)

```

```

select count(*) as n_lca_rows,
       round(avg(annual_wage_from),0) as avg_lca_offered,
       round(median(annual_wage_from),0) as median_lca_offered,
       round(100.0*sum(case when annual_wage_from < 139850 then 1 else 0 end)/count(*),1) as pct_below_bls_median
from lca_dc_swe

```

G.24 19c_bls_wage_gap_national_swe_by_level

```

select raw_pw_wage_level, count(*) n, round(avg(annual_wage_from),0) avg_wage, round(median(annual_wage_from),0)
median_wage,
       round(100.0*sum(case when annual_wage_from < 130160 then 1 else 0 end)/count(*),1) as pct_below_national_median
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and soc_code_norm = '15-1252' and fiscal_year in (2023,2024) and annual_wage_from is not null
group by 1 order by 1

```

G.25 20_2026q3_spike_drivers

```

select employer_name_norm, count(*) as n
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and fiscal_year=2026 and source_quarter=3
group by 1 order by n desc limit 20

```

G.26 20b_2026q3_vs_priorQ3_same_employers

```

select employer_name_norm,
       sum(case when fiscal_year=2025 and source_quarter=3 then 1 else 0 end) as fy2025_q3,
       sum(case when fiscal_year=2026 and source_quarter=3 then 1 else 0 end) as fy2026_q3
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and fiscal_year in (2025,2026) and source_quarter=3
group by 1
order by fy2026_q3 desc
limit 20

```

G.27 21_FIXED_perm_cert_rate_trend

```

select fiscal_year, count(*) total,
       sum(is_certified) certified, sum(is_denied) denied, sum(is_withdrawn) withdrawn,
       round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct,
       round(100.0*sum(is_certified)/count(*),1) as cert_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/*.parquet')
group by 1 order by 1

```

G.28 21_perm_cert_rate_trend

```

select fiscal_year, count(*) total,
       sum(is_certified) certified, sum(is_denied) denied, sum(is_withdrawn) withdrawn,
       round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_source_rows/*.parquet')
where is_canonical_source=1
group by 1 order by 1

```

G.29 22_lca_denial_rate_top_employers

```

select employer_name_norm, count(*) as total,
       sum(is_denied) as denied,
       round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
group by 1
having count(*) >= 1000
order by denial_rate_pct desc
limit 20

```

G.30 23_meta_name_variants

```

select employer_name_norm, count(*) as certified_lcas
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and lower(employer_name_norm) like '%meta platforms%'
group by 1

```

G.31 24_perm_processing_days_trend

```

select fiscal_year,
       round(avg(datediff('day', try_cast(raw_received_date as date), try_cast(raw_decision_date as date))),0) as avg_days,
       round(median(datediff('day', try_cast(raw_received_date as date), try_cast(raw_decision_date as date))),0) as
median_days,
       count(*) as n
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/*.parquet')
where raw_received_date is not null and raw_decision_date is not null
group by 1 order by 1

```

G.32 25_hcl_denial_check

```
select case_status_norm, count(*) from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where employer_name_norm = 'hcl america solutions inc' group by 1 order by 2 desc
```

G.33 26_fy2021q3_spike_drivers

```
select employer_name_norm, count(*) as n
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and fiscal_year=2021 and source_quarter=3
group by 1 order by n desc limit 15
```

G.34 27_top_soc_codes_overall

```
select soc_code_norm, soc_title_norm, count(*) as n,
       round(100.0*count(*)/sum(count(*) over(),1) as pct_of_total
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1
group by 1,2 order by n desc limit 15
```

G.35 28_wipro_perm_denial_detail

```
select case_status_norm, count(*) from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/perm_disclosure_rows/*.parquet')
where employer_name_norm = 'wipro limited' group by 1 order by 2 desc
```

G.36 29_top_worksite_cities_overall

```
select worksite_city_norm, worksite_state_norm, count(*) as n
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where is_certified=1 and worksite_city_norm is not null
group by 1,2 order by n desc limit 15
```

G.37 30_lca_denial_rate_trend_overall

```
select fiscal_year, count(*) total, sum(is_denied) denied, round(100.0*sum(is_denied)/count(*),2) as denial_rate_pct
from read_parquet('/Users/simsies/github/semantic-backtester/data/snapshots/20260919-001/lca_disclosure_rows/*.parquet')
where fiscal_year between 2020 and 2025
group by 1 order by 1
```