Research Square

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

AI, Occupational Entry Barriers, and Upward Mobility

Ting Zhang

tzhang@ubalt.edu

University of Baltimore

TyLee Armstrong

Thomas Jefferson High School for Science and Technology

Research Article

Keywords: occupational entry barriers, rent-seeking, artificial intelligence, occupational mobility, O*NET Job Zones

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11425793/v1

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

Additional Declarations: No competing interests reported.

AI, Occupational Entry Barriers, and Upward Mobility

Ting Zhang
Professor and Harry Y. Wright Chair
Merrick School of Business, University of Baltimore, Baltimore, Maryland, USA
tzhang@ubalt.edu
ORCID: 0000-0001-8596-7488, corresponding author.

TyLee Armstrong
Thomas Jefferson High School for Science and Technology, Alexandria, Virginia, USA
TyLeeArmstrong@gmail.com.

Abstract

Public choice has long analyzed occupational entry barriers and is increasingly concerned with artificial intelligence (AI), but it has not examined how the two relate. We develop a public choice theory of tolls to explain it. An entry barrier is a toll set by the incumbents who profit from it, charged on credentials rather than demonstrated skill, and slow to adjust when more workers need to pass through it. Because the toll is fixed while AI lowers the value of some experience built up in the occupations it exposes, the toll claims a larger share of the cost of moving up for workers leaving those occupations, so the upward-mobility premium from entering a high-barrier occupation should erode, most sharply for them. A high-barrier occupation is one restricted by a license or comparable requirement in the worker's state, measured by a positive State Occupational Licensing Index Barrier Score. Mobility is the change in O*NET Job Zone, a five-tier scale of required preparation, between the occupation a worker leaves and the one entered. Using linked Current Population Survey records for 2010–2025, we compare moves into high-barrier and low-barrier occupations before and after 2023. Entering a high-barrier occupation was associated with 0.029 more Job Zone steps of upward mobility before 2023; afterward the premium falls by 0.018 to 0.041 steps, to roughly zero or below, and the decline is concentrated among workers leaving AI-exposed occupations. The evidence is associational, and the exposure gradient is less precisely estimated than the timing result.

Keywords: occupational entry barriers; rent-seeking; artificial intelligence; occupational mobility; O*NET Job Zones

Statements and Declarations

- Competing interests. The author has no relevant financial or non-financial interests to disclose.

- Author contributions. The first author led the study and was responsible for the conceptualization, analysis and writing. The second author contributed to the ideation, data cleaning and editing. Both authors approved the final manuscript.
- Funding. The author received no specific funding for this work.
- Ethics approval and consent. This study uses only de-identified, publicly available secondary data and involved no contact with human participants. Ethics approval and informed consent were therefore not required.
- Data and code availability. All data are publicly available: Current Population Survey microdata from IPUMS CPS (Flood et al. 2025), the State Occupational Licensing Index (Trudeau and Timmons 2023, 2024), O*NET Job Zones, and generative-AI exposure scores (Felten et al. 2023). The code used to link the CPS records, merge the data sets and produce all results is available from the corresponding author upon request.

JEL Classification: D72; D73; J24; J44; J62; L51; O33

Acknowledgments. The author thanks Professor Ilia Murtazashvili and Professor Henry Thompson for their guidance on this paper, and in particular Professor Ilia Murtazashvili for detailed and constructive advice over several rounds of revision.

AI, Occupational Entry Barriers, and Upward Mobility

Abstract

Public choice has long analyzed occupational entry barriers and is increasingly concerned with artificial intelligence (AI), but it has not examined how the two relate. We develop a public choice theory of tolls to explain it. An entry barrier is a toll set by the incumbents who profit from it, charged on credentials rather than demonstrated skill, and slow to adjust when more workers need to pass through it. Because the toll is fixed while AI lowers the value of some experience built up in the occupations it exposes, the toll claims a larger share of the cost of moving up for workers leaving those occupations, so the upward-mobility premium from entering a high-barrier occupation should erode, most sharply for them. A high-barrier occupation is one restricted by a license or comparable requirement in the worker's state, measured by a positive State Occupational Licensing Index Barrier Score. Mobility is the change in O*NET Job Zone, a five-tier scale of required preparation, between the occupation a worker leaves and the one entered. Using linked Current Population Survey records for 2010–2025, we compare moves into high-barrier and low-barrier occupations before and after 2023. Entering a high-barrier occupation was associated with 0.029 more Job Zone steps of upward mobility before 2023; afterward the premium falls by 0.018 to 0.041 steps, to roughly zero or below, and the decline is concentrated among workers leaving AI-exposed occupations. The evidence is associational, and the exposure gradient is less precisely estimated than the timing result.

Keywords: occupational entry barriers; artificial intelligence; occupational mobility; Job Zones

JEL Classification: D72; J44; J62; O33

1. Introduction

When technology changes the jobs people do, the institutions that govern work determine whether workers can follow. In the 1980s and 1990s employment shifted away from routine middle-skill work (Autor et al. 2003; Goos et al. 2014), and for non-college workers the erosion of specialized middle-skill jobs was associated primarily with reallocation into lower-skill, lower-wage occupations rather than into high-skill ones (Autor 2019). Generative AI poses the question again, to different workers. The most exposed occupations are cognitive, analytical, and communication-heavy, the many credentialed white-collar jobs that sat safely above the last wave (Felten et al. 2023; Eloundou et al. 2024), so the rungs just above them are often licensed

professions such as healthcare, law, accounting, and teaching. Roughly three-quarters of healthcare practitioners and two-thirds of legal workers hold an active license, against fewer than one in ten in food service or administrative support (Cunningham 2019). The workers who most need to move up are therefore standing in front of some of the most heavily gated doors in the American labor market.

Public choice has studied such doors for decades. It explains why entry restrictions arise, who gains from them, and what resources are spent in obtaining them (Stigler 1971; Peltzman 1976; Tullock 1967; Krueger 1974), and it is increasingly concerned with the consequences of AI. Work on earlier automation links licensing to mobility (Geloso et al. 2026), but public choice has not asked whether entry barriers shape which workers can adjust to AI. The economics of AI and labor has the complementary gap. It measures which occupations are exposed (Felten et al. 2023; Eloundou et al. 2024) and who gains from the tools (Noy and Zhang 2023; Brynjolfsson et al. 2025), but it has not fully embraced the emphasis on barriers in labor markets, and on how technology interacts with them, that this tradition supplies. AI exposure tells us who will need to move, not where they can go, and the economics of labor markets and AI needs that emphasis to answer the second question.

We develop a public choice theory of tolls that links the two. A worker choosing a new occupation pays a learning cost, which falls as the old and new jobs overlap in content, and a political toll, the schooling, examinations, fees, and supervised practice a state requires for entry. The toll is set by boards of practitioners who profit from its height, and it has three properties: it is charged on credentials held rather than skill shown, it rises with the rank of the occupation entered, and it is politically sticky, because loosening it requires the consent of those who benefit from scarcity (Olson 1965; Stigler 1971). A shock that lowers the value of the experience workers have built in the occupations they leave raises the toll's share of the cost of moving up. Generative AI is such a shock, and it should be strongest for workers leaving the occupations it exposes most. The theory predicts that the effect appears in the interaction between the barrier and the post-shock period rather than in the barrier's level, and that it is larger for workers leaving the occupations AI exposes most. Because the barrier itself does not change, what changes is the pressure against it.

The theory becomes testable once two terms are defined. A high-barrier occupation is one in which practitioners already inside can restrict entry by license or a comparable requirement. Empirically, a destination is high-barrier when the Barrier Score of the State Occupational Licensing Index (Trudeau and Timmons 2023, 2024) is 1 (versus 0) for that occupation in the worker's state, that is, when the occupation's tasks are restricted by any license, including cases where the occupation carries no explicit license but is barred under another licensed title. About one move in five lands in such a destination (Section 4.1). Upward mobility is the change in

O*NET Job Zone, a five-tier scale of required preparation, between the occupation a worker leaves and the one entered. We follow workers across two consecutive Current Population Survey interviews, keep those who changed occupations, and compare moves into high-barrier and low-barrier destinations before and after 2023, the first full year of widespread public access to generative AI. Model A uses the 2023 index with a post-period of 2023–2025; Model B uses the 2024 index, whose only post-period year is 2025. We then ask whether the comparison varies with the AI exposure of the occupation the worker left (Felten et al. 2023).

Before the benchmark year, entering a high-barrier occupation was associated with 0.0287 additional Job Zone steps in the preferred 2023-index model ($p = .023$); the 2024-index level term is positive but statistically indistinguishable from zero (0.0071). After the benchmark year the premium erodes: the interaction is −0.0179 under the 2023 index ($p = .070$) and −0.0409 under the 2024 index ($p = .002$), which puts the implied post-period relationship at +0.0108 and −0.0338 steps. These magnitudes are 2.6 and 5.9 percent of the standard deviation of the outcome, and 19 and 43 percent of the bachelor's-degree coefficient. The erosion sorts by exposure. For a worker leaving an occupation one standard deviation below mean AI exposure, the post-period change in the premium is essentially zero (+0.0022 under the 2023 index); for one leaving an occupation one standard deviation above it, the change is −0.0275. The triple interaction is significant at ten percent under the 2023 index ($p = .078$) and has a similar sign and size, but lacks statistical significance, under the 2024 index ($p = .115$). The design identifies a change in association between a stable barrier grouping and the direction of mobility, not the effect of a policy change, and AI exposure correlates 0.61 with the Job Zone of the occupation a worker leaves, so we read the exposure result as suggestive rather than settled.

The paper's main contribution is the theory. Licensing scholars have largely measured the level of a barrier, or its change over time, and asked what it does on average (Kleiner 2000; Kleiner and Krueger 2013; Blair and Chung 2019; Johnson and Kleiner 2020; Jackson 2023). We ask how the bite of a fixed barrier changes when a technology shock changes the pressure against it, and we derive a testable prediction: the upward-mobility premium from entering a high-barrier occupation should erode after the shock, and by more for workers leaving more-exposed occupations (Section 3). The closest related work, Geloso et al. (2026), asks how changes in licensing condition the link between exposure to industrial robots and commuting-zone intergenerational mobility; we hold the barrier fixed and ask how it governs the direction of a worker's own move when generative AI, which exposes credentialed white-collar work, is the shock. Empirically, we examine a feature of licensing that has received less attention, the direction of occupational moves, meaning whether a worker moves up, laterally, or down in required preparation, and ask whether the link between entry barriers and upward moves changed when generative AI arrived. We then test a mechanism: whether any change is concentrated

among workers leaving AI-exposed occupations, which macroeconomic explanations sharing the same dates would not predict (Section 5.4). On design, we show why a study of licensing and mobility must anchor treatment, outcome, and destination-side fixed effects to the job entered while absorbing the job left at a coarse resolution (Section 4.2).

The pattern has significant policy implications, which we develop in Section 7. If licensing simply raised the average cost of entering a profession, the remedy would be the familiar one of lowering it. If a fixed barrier fails to adjust when a shock changes how many workers want to enter or how much entry is worth, licensing may hinder adaptation; the policy question is then whether entry regulation can respond to shocks at all.

Section 2 reviews the literature, Section 3 develops the toll theory and its hypotheses, Section 4 describes the data and design, and Section 5 reports the findings. Section 6 summarizes robustness checks, Section 7 discusses interpretation, limits, and policy, and Section 8 concludes.

2. Literature Review

Regulation is not simply imposed for the public good; Stigler (1971) argued that organized industries may seek regulation that benefits them, while Peltzman (1976) modeled regulators as balancing competing interests to maximize political support. Small groups with a lot at stake per person organize easily; large groups with little at stake per person do not (Olson 1965). Because winners are organized and losers are not, competing for these protections wastes real resources on top of whatever it redistributes (Tullock 1967; Krueger 1974; Tollison 1982). Buchanan and Tullock (1962) trace this to the rules under which regulation is written, not to regulators behaving badly; McChesney (1987) adds that politicians can profit from creating protections and from threatening to take them away. Occupational licensing is where this theory is least abstract. Authority over who may practice is typically delegated to a board of practitioners drawn from the occupation being licensed, so the people who set the height of the barrier are the people who profit from its height. Capture theory describes how regulators come to serve the regulated (Dal Bó 2006); with licensing boards, no capture may even be necessary, because the regulated are the regulator. Friedman (1962) made this criticism sixty years ago, and the institutional form has scarcely changed.

Coverage has grown from roughly five percent of the workforce in the 1950s to about twenty-nine percent in 2008 (Kleiner and Krueger 2013). Licensed workers earn twelve to eighteen percent more, but caution that the estimates are descriptive rather than strictly causal (Kleiner and Krueger 2013). Licensing reduces equilibrium labor supply (Blair and Chung 2019) and suppresses interstate migration, since credentials do not travel across state lines (Johnson and Kleiner 2020). Welfare analyses find meaningful aggregate costs (Kleiner and Soltas 2023), and

federal policy analysis recommends portability, reciprocity, and periodic board review (U.S. Department of the Treasury et al. 2015). Closest to our question, Jackson (2023) links state–occupation licensing policies to Current Population Survey records from 2015–2022. He finds that licensing is associated with a 2.2-percentage-point lower probability of switching into a licensed final occupation—about 24% relative to the unlicensed group's mean—and no statistically significant effect on switching out of the baseline occupation. His results also associate lower entry with fees and minimum age and education requirements. Jackson studies whether licensing is associated with occupational switching; we ask whether workers move up or down in required Job Zone preparation, and whether those patterns vary with generative-AI exposure. This literature has largely studied the level of licensing and its effects on wages, labor supply, migration and entry, rather than the direction of moves, and it has typically taken the barrier's effect as constant over time.

The task framework showed that technology replaces activities, not whole jobs (Autor et al. 2003; Acemoglu and Autor 2011), producing labor-market polarization across rich economies (Goos et al. 2014; Autor 2019), though that account has been challenged (Hunt and Nunn 2022). Generative AI lands differently. AI exposure measures put cognitive and communication-heavy work at the center of the distribution, not the tail (Felten et al. 2023; Eloundou et al. 2024), and productivity gains concentrate among less-experienced workers (Brynjolfsson et al. 2025) or weaker performers (Noy and Zhang 2023), narrowing the gap that experience used to open between novice and veteran output. A narrower gap means a smaller pay premium for experience: expertise depreciates not because it stopped mattering, but because AI now supplies part of what it used to. If experience is worth less while the licensing toll stays fixed, the toll claims a larger share of what moving up is worth, with no one voting on it. The economics of AI and labor, however, has not fully embraced the public choice emphasis on barriers in labor markets and on how new technology interacts with them. It tells us who is exposed and who gains from the tools, but it has given less attention to the institutions that govern entry, so it cannot say where workers leaving exposed occupations are able to go. For that reason the economics of labor markets and AI needs public choice.

Workers move most easily between occupations with similar content (Poletaev and Robinson 2008; Gathmann and Schönberg 2010; Neffke et al. 2017). One prior study complicates that: similarity predicts a lower chance that a move is upward (Zhang and Schintler 2025). Closeness in skill space buys passage through ungated doors that mostly lead to lateral moves. That result motivates this paper: if the knowledge penalty reflects a credential barrier, the barrier itself has to be measured. A second political mechanism runs through the government agencies that place workers in new jobs. Like any bureau (Niskanen 1971), they are funded and judged on how many people they place and how fast — often not necessarily emphasizing whether the new job

is a step up. That incentive gets gamed: training centers have strategic timing of trainee terminations or graduations around reporting periods to flatter statistics (Courty and Marschke 2004), and consequently, the measures agencies are somewhat rewarded for showing little relation to workers' long-run earnings (Heckman et al. 2002). Our data cannot test this mechanism directly, but it shapes the policy discussion in Section 7.

Three gaps follow. The licensing literature measures barriers and their average effects but gives less attention to whether their effect changes with conditions; the AI and labor literature measures exposure but not the institutions workers must pass through; and the occupational mobility literature measures skill distance but has given less attention to barrier-related regulation. This paper measures the entry barrier directly, tracks the direction of a worker's move rather than just its level, and tests whether the barrier's association with that direction changed once generative AI began reshaping work. The closest bridge between the licensing and technology literatures is Geloso et al. (2026), who examine how changes in licensing condition the relationship between industrial robots and commuting-zone intergenerational mobility; our question is adjacent but distinct: we observe individual moves rather than area averages, and our exposed workers are white-collar workers whose next rung could be a licensed profession.

3. Theory and Hypotheses

3.1 A toll model of which way people move

The model is simple. A worker choosing where to go next faces two costs: the cost of learning the new job, which shrinks the more it resembles the old one, and the cost of being allowed to do it — the license. The first is economic and falls with skill overlap; the second is political and does not necessarily. Everything that follows is a consequence of that difference.

Formally: a worker in occupation o chooses among destinations d, each at some level of required preparation, and what we observe is the distance travelled:

$$\Delta J = J_{\mathrm{d}} - J_{\mathrm{o}}$$

Getting into d costs the worker two different things. The first is a **learning cost**, $c(s, o, d)$, which falls as the competencies of o and d overlap. This is the standard human capital story (Becker 1964; Gathmann and Schönberg 2010; Poletaev and Robinson 2008). The second is a **political toll**, $\tau(d, s)$ — the mandated schooling hours, examinations, fees, supervised practice, and state-specific credentials required to enter occupation d in state s. Conditional on moving, the worker picks the destination that maximizes

$$V_{\mathrm{d}} = v_{\mathrm{d}} - c(s, o, d) - \tau(d, s)$$

The toll τ has three properties that do all the analytical work below, and each one follows from the political logic of Section 2, not from anything about technology. First, it is competence-invariant: the toll is charged based on the credentials a worker holds, not on the skill they can actually demonstrate. A worker who already has the skills a licensed occupation requires still pays close to the full price to get the credential. Second, it is rank-increasing: occupations with more to protect build higher walls, so the toll rises with the level of the destination job. Licensed occupations cluster at the top of the required-skill scale; the least-regulated occupations sit at the bottom. Third, it is politically sticky: the toll is the result of a political bargain between organized incumbents and unorganized outsiders, and it does not automatically adjust just because more or fewer people want to get through. Figure 1 renders the setup. Because τ taxes upward moves but barely touches lateral or downward ones, it tilts realized transitions toward $\Delta J \leq 0$ even for workers whose learning cost is low.

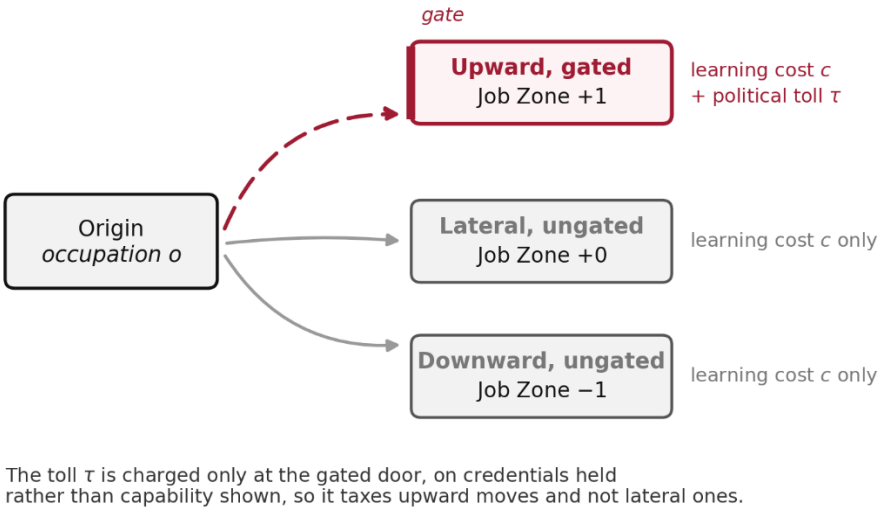


Figure 1. *Three doors out of an occupation, one of them gated.*

Note. The learning cost applies to every destination and falls as competencies overlap. The political toll applies only to the gated door and is charged on credentials held rather than capability shown, which is why competency proximity cannot substitute for it.

3.2 What generative AI does to a fixed toll

Now suppose generative AI lowers what a worker's accumulated experience is worth. A move up into a licensed occupation is worth making when its gain, the value of the destination net of the learning cost, exceeds the toll. Productivity gains from AI concentrate among less-experienced workers (Noy and Zhang 2023; Brynjolfsson et al. 2025), which narrows the gap that experience used to open, so we take the gain from moving up to shrink for workers leaving exposed occupations. Because the toll is charged on credentials rather than skill and cannot be loosened without the consent of those who benefit from scarcity, it does not shrink with the gain.

It therefore claims a larger share of a smaller gain, and some upward moves into gated occupations stop paying; moves into ungated occupations, where there is no toll to claim a share, are less affected. The adjustment is one-sided: unregulated training markets can expand to meet demand, while licensed entry cannot. A second channel is possible. If the shock also sends more workers toward the same licensed occupations, entry is rationed by seats, exam cycles, and supervised-practice slots administered by incumbents, and the marginal upward mover is diverted. We do not test this channel, and the predictions below do not depend on it.

3.2.1 Whose gain shrinks: the exposure margin

The devaluation channel operates through how much AI has devalued a worker's accumulated experience. That is occupation-specific, not a national aggregate, because generative AI does not reorganize all work at once. Exposure concentrates in cognitive, analytical, and communication-heavy occupations (Felten et al. 2023; Eloundou et al. 2024), the occupations that sit just below the licensed professions. This yields a more specific prediction than "the toll binds harder after the benchmark year," one a rival explanation cannot easily mimic. If the erosion reflects a fixed toll claiming a larger share of a shrinking gain as experience is devalued, it should be largest among workers leaving the most AI-exposed occupations. If it instead reflects macroeconomic conditions that happen to share our dates, there is no reason for it to sort by occupational exposure to a technology.

This is the central identifying idea of the paper's second half: a claim resting only on timing cannot rule out coincidence, since several macroeconomic shifts share our benchmark date; a claim resting on heterogeneity can, because no generic macroeconomic account gives a reason for that sorting.

3.3 Hypotheses

Together, the toll's three properties and the devaluation mechanism they interact with generate five testable predictions. **H1 (Gate level).** *Moves into high-barrier destination cells are associated with larger positive ΔJ than moves into low-barrier cells.* This is a fact about who ends up where, not about how people behave. Licensing tracks occupational rank closely (Cunningham 2019), so entering a high-barrier cell usually means entering a higher Job Zone. H1 is a validity check: if the level relationship were negative or zero, the barrier index would not be measuring what we claim.

H2 (Toll binding). *In the post-benchmark period, the relationship between destination entry barriers and upward movement weakens.* This is the paper's central prediction and follows from Section 3.2. Note what it does *not* say: it does not predict that the entry barrier pushes workers

downward. It predicts erosion of an upward premium, which is a different and more demanding claim.

H3 (Political stickiness). *Before the benchmark year, high-barrier and low-barrier destination cells show no differential trend in ΔJ.* This is the identifying assumption, but also a substantive claim. If τ is politically sticky, it is stable, and stability means parallel pre-trends. A violation would be theoretically informative, not merely inconvenient — it would say high-barrier cells were already on a different path.

H4 (Exposure heterogeneity). *The post-benchmark erosion of the high-barrier-destination premium is larger for workers leaving occupations more exposed to generative AI.* H4 follows from the devaluation channel (Section 3.2): as AI devalues the experience of workers in exposed occupations (Section 2), the gain from a high-barrier move shrinks for them, and the fixed toll τ claims a larger share of it.

4. Data and Method

4.1 Data and measures

Worker data come from Current Population Survey records for 2010–2025 (U.S. Census Bureau 2025; Flood et al. 2025). The CPS rotation pattern permits one-year matches for a subset of respondents, subject to successful person-level linkage and attrition. Longer panels such as the PSID and NLSY can track workers over longer periods, but occupational coding and classification changes require careful harmonization. Kambourov and Manovskii (2008) document important measurement problems in historical occupational data. We measure occupational preparation using ONET Job Zones (National Center for O*NET Development n.d.; O*NET Online n.d.), which group occupations by typical education, related experience, and on-the-job training requirements. Job Zone is a preparation scale, not a direct measure of cognitive complexity: five tiers of required preparation from Zone 1 (little or none) to Zone 5 (extensive, typically graduate training) that combine schooling, experience, and on-the-job training; higher zones involve more cognitive complexity.

This scale's usual weakness is our advantage: much of the “preparation” defining Zones 4 and 5 is related to entry barriers— the degrees, exams, and supervised hours placing surgeons, lawyers, accountants, and teachers at the top are licensing requirements as much as job requirements (Kleiner 2000). An upward Job Zone move is often a passage through a political gate, which is exactly what makes it useful here.

Occupational entry barriers are measured with the Barrier Score from the State Occupational Licensing Index (SOLI) for 2023 and 2024 (Trudeau and Timmons 2023, 2024), published by

the Archbridge Institute from the Knee Regulatory Research Center's Occupational Regulation Database and cross-walked to CPS occupations. The Barrier Score counts, for each state-occupation pair, whether the tasks associated with that occupation are restricted by any license, including cases where the occupation itself carries no explicit license but is barred under another licensed title — broader than the Index's separate License Score, which counts only explicit, title-specific licenses. We use the Barrier Score because it captures entry restriction as such. Technology exposure uses occupation-level exposure to generative AI (Felten et al. 2023), normalized to the unit interval and merged onto both origin and destination occupations; Section 4.2 explains why only the origin-side measure is estimable.

We keep workers employed at both interviews who report a different occupation at the second — people who changed jobs rather than being pushed out of employment, isolating the margin where a destination is actually selected. The merged file holds **214,148 linked switchers** with destination years from 2011 to 2025. After dropping records missing the licensing index or covariates, and after the automatic removal of twelve singleton observations, the baseline estimation sample is **142,491**, and the exposure-moderated models, which require non-missing AI exposure, use 123,717. Figure 2 shows the year composition, with a warning: CPS linkage yield declines steadily (17,696 matched switchers in 2011 to 11,676 in 2024), so both post-periods fall in the thinnest years of the panel. Destination years 2023 onward (Model A's post-period) comprise 16.52 percent of the merged file; 2025 alone (Model B's entire post-period) comprises 5.55 percent.

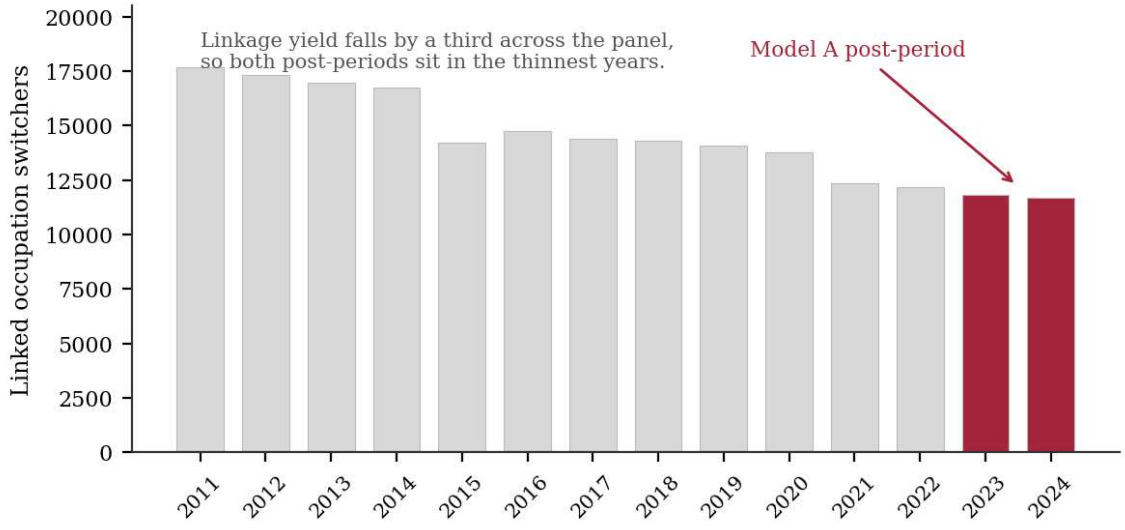


Figure 2. *Sample composition by destination year.*

Note. Linked occupation switchers by destination year, 2011–2024. Linkage yield declines by roughly a third across the panel, so both post-periods fall in the thinnest years of the data.

In plain terms, the outcome answers one question about each worker: did they move up, laterally, or down, and by how much? We measure that as the change in Job Zone across the move, on a five-tier scale that runs from −4 to +4 (Section 4.3 reports its distribution). We use the signed measure rather than a simple up-or-down indicator for three reasons: it keeps the size of the move, treats direction as a single continuum (as the toll model predicts), and permits linear estimation with 319 occupation fixed effects, where a binary model would be fragile. We flag a binary specification as a future extension in Section 8.

The treatment is a binary indicator, HighBarriers_{d}, equal to one when the destination state-by-occupation cell carries a high entry-barrier score:

$$\text{HighBarriers}_{d} = 1\{\text{HighBarriers}_{d} > 0\}$$

In plain terms, a high-barrier occupation is one where the practitioners already inside get to set the terms on which anyone else may enter — the mechanism Section 3.1 describes. The extensive-margin definition is deliberate. Most occupations carry a zero barrier score, so splitting the continuous score at some percentile would slice through structural zeros arbitrarily. What matters theoretically is whether a toll is charged at all — the extensive margin captures exactly that. In the analytic sample, 20.45 percent of moves land in a high-barrier destination cell.

4.2 Research design and specification

Treatment, outcome, and the four principal fixed effects are defined on the destination side (Figure 3), for two reasons. Theoretically, an entry barrier cannot affect a worker until they try to enter the regulated occupation: the toll is charged at the destination's gate, not the origin's exit, so measuring the barrier on the origin side would misattribute it to the occupation being left, where it was already paid or never owed. Econometrically, controlling for origin occupation is safe because the treatment varies by destination. Workers leaving the same two-digit origin occupation go on to different destinations, some high-barrier and some low-barrier. The origin fixed effect therefore removes differences in where workers start, without removing the destination differences that identify the barrier effect. Anchoring on the destination does not excuse ignoring the origin: since the outcome is a difference (destination Job Zone minus origin Job Zone) and destination fixed effects absorb the first term, origin variation passes straight into the residual. Workers who start high have little room to climb and far to fall, and vice versa, so left unabsorbed this asymmetry contaminates every covariate.

Detailed origin occupation can't be used — 319 categories interacted with the destination-side structure causes severe singleton deletion and non-convergence — but grouping to the two-digit

SOC level (22 categories) estimates cleanly on the full sample[1]. We report both specifications throughout: the destination-only model is the design's minimal form, and the model adding coarse origin occupation is our preferred specification, the one on which the paper's claims rest.

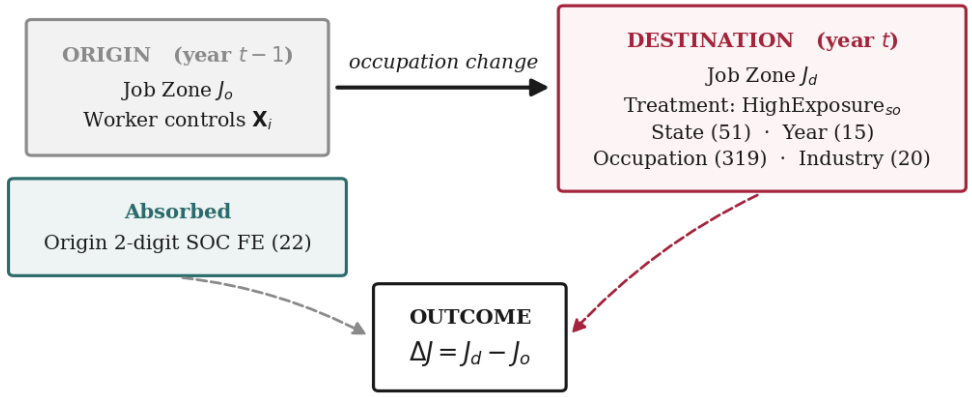


Figure 3. *The design: destination-anchored treatment, origin-absorbed composition.*

Note. Treatment, outcome, and the four principal fixed effects are defined on the destination side. Origin occupation is absorbed at the two-digit SOC level (22 categories), which removes the bulk of headroom contamination; detailed origin occupation does not converge.

The comparison we want is straightforward: among workers moving into high-barrier jobs versus low-barrier ones, did the direction of the move change after the benchmark year, the first year of the post-period? Because the barrier index exists only for 2023 and 2024, we estimate two independently anchored models rather than pooling them. Each defines its post-period by the calendar year in which the worker is observed in the destination job:

$$\Delta J_{ist} = \beta\,(HighBarriers_d \times Post_t) + \varphi\, HighBarriers_d + X_i'\gamma + \delta_s + \mu_o + \nu_n + \tau_t + \rho_g + \varepsilon_{ist}$$

where δ_s, μ_o, ν_n, and τ_t are fixed effects for destination state (51 categories), destination detailed occupation at the SOC-2018 level (319 categories), destination two-digit NAICS industry (20 categories), and destination-observed year (15 categories); and ρ_g denotes origin two-digit SOC

[1] The gain is substantial: adding these 22 categories raises the model's R^2 from 0.2161 to 0.6280 without materially altering the treatment coefficients, the signature of absorbing variation that was contaminating covariates rather than treatment. It also cuts the education gradient by roughly seventy percent and dissolves several implausible covariate patterns (Section 5.3).

occupation group (22 categories), included in our preferred specification and omitted in the destination-only comparison.

Model A uses the 2023 index with $\text{Post}_t = 1\{\text{year} \geq 2023\}$. The 2023 index was published in March 2023 from information predating the 2023 CPS collection, so treating 2023 onward as post is consistent with the informational timing. **Model B** uses the 2024 index with $\text{Post}_t = 1\{\text{year} \geq 2025\}$. The 2024 index appeared in August 2024, after that year's CPS collection, so only 2025 qualifies. That leaves a single post-year. The main effect of Post_t is absorbed by the destination-year fixed effects and reported as omitted.

The AI-exposure-moderated specification. To test H4 we interact the difference-in-differences term with a continuous measure of the worker's exposure to generative AI in the occupation they left:

$$\Delta J_{ist} = \beta 1\ (HighBarriers_d \times Post_t) + \beta 2\ (HighBarriers_d \times Post_t \times \mathbf{AIExposure_o})$$
$$+ \text{(lower-order terms)} + X'_i\,\gamma + \text{(fixed effects)} + \varepsilon_{ist}$$

It is origin-anchored: AIExposure records the exposure of the occupation a worker departs, not the one they enter — the theoretically correct choice given Section 3.2.1, since the mechanism runs through how much AI devalued the experience a worker accumulated in the occupation they are leaving.

X_i contains firm size (nine categories),[2] class of worker (seven), full-time status, age and its square, sex, race, Hispanic ethnicity, years since immigration, education (three), marital status (three), and disability status. All models are weighted by the IPUMS CPS one-year longitudinal linking weight. For the event-study version used in Section 6, the single Post_t dummy is replaced by a full set of year interactions, each in turn interacted with AI exposure.

4.3 Inference and descriptive statistics

Standard errors are clustered by destination state to account for state-level treatment variation and correlated shocks among workers entering the same state labor market— 51 clusters, 50 denominator degrees of freedom[3].

---

[2] The firm size variable contains overlapping value labels because the underlying Census Bureau survey questions changed their size categories across survey years, and the composite coding scheme retains mutually exclusive groupings meant for distinct time periods.

[3] Fifty-one clusters may be adequate for asymptotic cluster-robust inference, but MacKinnon et al. (2023) reject any universal threshold: validity depends on cluster-size imbalance, leverage, and heterogeneity across clusters, not on cluster count alone, and with only 51 clusters a reader may still want a bootstrap alternative. We attempted wild

One clarification: this is a two-group, two-period comparison across a fixed entry-barrier grouping, not a staggered-adoption study, so the known problems from units treated at different times (Goodman-Bacon 2021; de Chaisemartin and D'Haultfœuille 2020) don't arise in their usual form.

The analytic sample (N = 123,717) skews older than a typical group of job changers (mean age 45), reflecting the linkage and employed-to-employed restrictions. The outcome is tightly bunched around zero, with a mean of 0.0081 and a standard deviation of 0.6962 on a −4-to-+4 scale, so most moves span at most one Job Zone step and any institutional effect will look small in absolute terms. About one move in five lands in a high-barrier destination, the margin the whole design turns on. Full statistics are in Appendix A.7.

5. Findings

Three numbers carry this section. The first says whether entering a high-barrier job means moving up. The second says whether that changed after the benchmark year. The third says for whom it changed — and that is where the technology enters the argument.

Table 1. Effect of destination entry barriers on the direction of occupational movement

Dependent variable: Δ Job Zone (destination minus origin)

	(1)	**(2)**	**(3)**	**(4)**
	Model A	Model B	Model A	Model B
	2023 index	2024 index	2023 index	2024 index
	destination FE only	destination FE only	**+ origin 2-digit SOC FE**	**+ origin 2-digit SOC FE**
HighBarrier × Post	−0.0187	−0.0321	**−0.0179***	**−0.0409****
	(0.0112)	(0.0199)	(0.0097)	(0.0122)
HighBarrier (level)	0.0278	0.0099	**0.0287****	0.0071
	(0.0186)	(0.0173)	(0.0123)	(0.0118)
Post	omitted	omitted	omitted	omitted

cluster bootstrap on both headline coefficients but could not implement it in our current specification and software: the packages available to us do not support the multi-way fixed-effect structure every specification here uses. Obtaining it would require a different estimator whose results would not be directly comparable to those below.

	(1)	(2)	(3)	(4)
Worker controls	yes	yes	yes	yes
Destination state FE	yes	yes	yes	yes
Destination year FE	yes	yes	yes	yes
Destination occupation FE	yes	yes	yes	yes
Destination industry FE	yes	yes	yes	yes
Origin 2-digit SOC FE	no	no	**yes**	**yes**
Observations	142,491	142,491	142,491	142,491
R^2	0.2161	0.2161	0.6280	0.6280
Adjusted R^2	0.2137	0.2137	0.6269	0.6269
Within R^2	0.0570	0.0570	0.0167	0.0167
Root MSE	0.6334	0.6334	0.4363	0.4363
$F(29, 50)$	1,012.38	1,068.51	291.41	247.86
Prob > F	0.0000	0.0000	0.0000	0.0000
Clusters (destination state)	51	51	51	51
Singletons dropped	12	12	12	12

Note. Weighted linear regressions with high-dimensional fixed effects. Standard errors clustered on destination state in parentheses. Weights are the IPUMS CPS one-year longitudinal linking weight. Fixed effects: destination state (51), destination year (15), destination detailed SOC-2018 occupation (319), destination two-digit NAICS industry (20), and, in columns 3–4, origin two-digit SOC occupation group (22). **** p < .001, *** p < .01, ** p < .05, * p < .10.

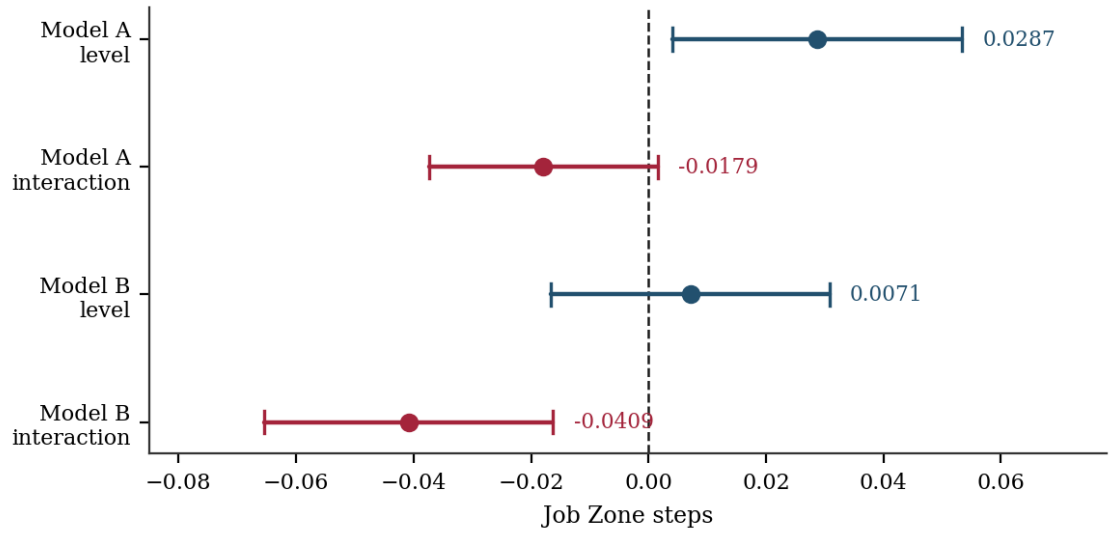


Figure 4. *Main estimates with 95 percent confidence intervals.*

Note. Level terms in blue, interaction terms in red. Estimates from Table 1, columns 3 and 4. Model B's interval is the only one that excludes zero.

5.1 Before the benchmark: the gate led up

If incumbents use their control over entry to their own advantage, the occupations they gate most tightly should also be the ones workers had to climb hardest to reach — and that is what the data show before the benchmark year. The level coefficient on the high-barrier indicator is positive in every specification, and in our preferred model (Table 1) for the 2023 index it is statistically significant: 0.0287 Job Zone steps (SE 0.0123, p = .023). The positive sign is what H1 predicts and what the institutional structure requires. Licensing clusters in healthcare, legal, and education occupations near the top of the Job Zone scale, while the least-regulated families sit at the bottom. Entering a high-barrier cell usually means climbing.

Absorbing the occupation a worker left improves the statistical significance compared to the estimate in the destination-only specification in Table 1 returns the same coefficient, 0.0278 (SE 0.0186, p = .141). The coefficient survives a demanding comparison — with 319 destination-occupation categories absorbed, it asks whether workers entering the same occupation travel a different distance up the Job Zone scale in high-barrier versus low-barrier states. Under the 2024 index the level term is 0.0071 (SE 0.0118, p = .549) — the same sign, but indistinguishable from zero. We take H1 as supported by Model A and uninformative in Model B.

5.2 After the benchmark: the premium disappears

The interactions are the paper's central result, and both are negative. Under the 2023 index the interaction (Table 1, column 3) is −0.0179 (SE 0.0097, p = .070). Under the 2024 index (column 4) it is −0.0409 (SE 0.0122, p = .002). Model A clears the ten percent threshold; Model B clears one percent, and its confidence interval excludes zero.

The comparison that matters is not to zero but to the level term above it. Model A's interaction cancels about sixty-two percent of the pre-period level relationship (0.0179 against 0.0287); Model B's exceeds its level term by roughly five and three-quarters times (0.0409 against 0.0071). Adding them gives the implied post-period relationship: +0.0108 in Model A and −0.0338 in Model B. In both cases, the post-period association between entering a high-barrier occupation and moving up is at or below zero. In Model A the gate stopped being a route upward; in Model B the association turns slightly negative.

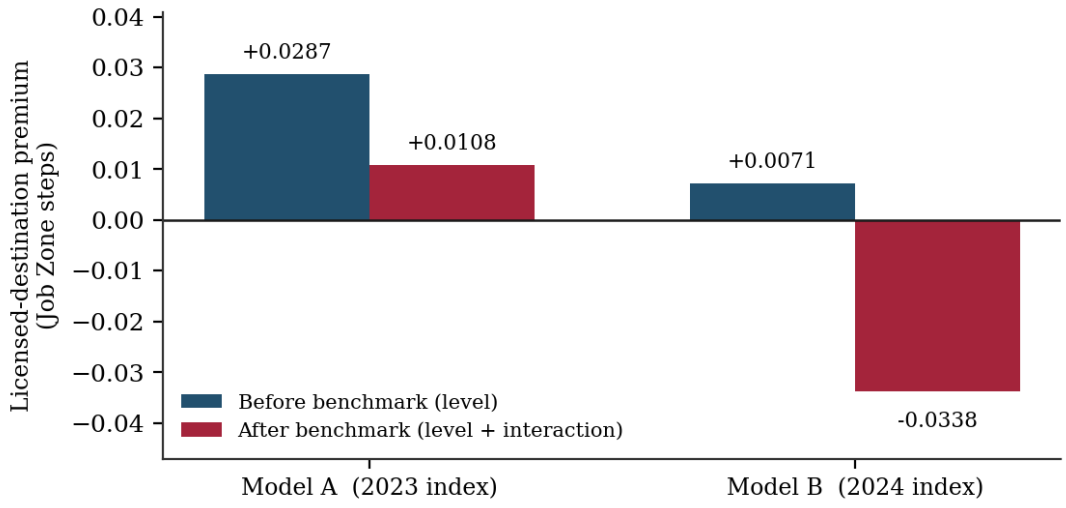


Figure 5. *The implied relationship before and after the benchmark year.*

Note. Post-period bars are the sum of the level and interaction terms from Table 1, columns 3 and 4.

This is precisely the pattern H2 predicts, and it distinguishes the toll account from two rivals. It is not what a simple "entry barriers block mobility" story predicts, which would put the negative sign on the level term — ours is positive and, in Model A, significantly so. Nor is it what a "licensing got stricter" story predicts, since the index is a fixed cell characteristic that does not change across the boundary. What changed is the world the toll operates in, not the toll.

5.3 How big is this?

The effect is small but not negligible. Model A's interaction equals 2.6 percent of the outcome's standard deviation (0.0179 ÷ 0.6962); Model B's equals 5.9 percent. On the Job Zone scale, Model A's estimate amounts to one fewer step of upward movement for roughly every fifty-six moves into high-barrier destinations after 2023, and Model B's to one fewer for roughly every twenty-four. Set against the covariates (Appendix A.8), the barrier effect is a fifth to two-fifths the size of the education gradient: Model A's interaction is 18.8 percent of the bachelor's-degree coefficient (0.0179 against 0.0952) and Model B's is 42.9 percent. Against the full-time coefficient, the shares are 31.9 and 72.6 percent.

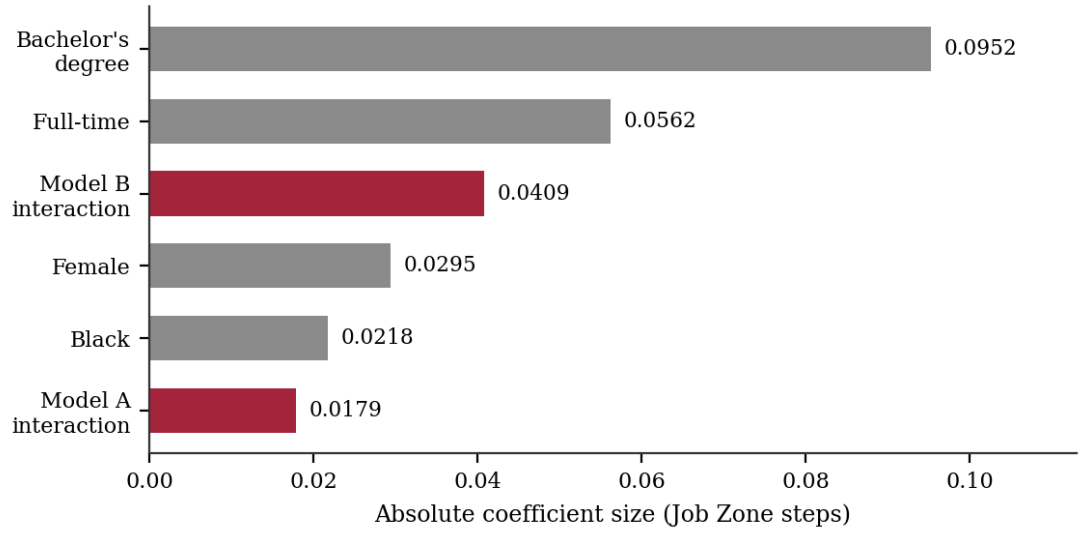


Figure 6. How large is the barrier effect?

Note. Absolute coefficient sizes in Job Zone steps, from Table 1 and Appendix A.8, using the origin-absorbed specification. Interaction terms in red.

Once origin composition is absorbed (Table A-4), the education gradient shrinks by roughly seventy percent while the barrier interaction barely moves, and several implausible covariate patterns dissolve (Appendix A.8). The barrier and human capital effects are of similar magnitude: human capital is larger, but by well under a factor of ten. Occupational entry barriers do not govern the direction of American occupational mobility, since workforce composition and human capital do far more, but the barrier is not negligible, shifting mobility at a magnitude comparable to well-established labor market fundamentals, in the direction and at the moment the toll theory anticipates. The models absorb a great deal of variation (R^2 of 0.6280 in the preferred specification), but the within-fixed-effects R^2 of 0.0167 shows how little residual variation the covariates explain: where a worker goes next is overwhelmingly idiosyncratic, and any institutional variable is a small term in that equation.

5.4 Where the erosion falls: the technology margin

Everything so far establishes when the high-barrier-destination premium eroded, not why. This section asks a different question — not when, but to whom — because a toll that incumbents control and do not adjust should bind hardest on workers whose experience AI has devalued most, not on workers in general. We interact the difference-in-differences term with each worker's exposure to generative AI in the occupation they left. If the premium erodes because AI has devalued the experience of workers leaving exposed occupations while the toll stayed fixed, the erosion should concentrate among workers whose occupations the technology actually disrupted — not fall evenly across occupations, as a macroeconomic story would predict.

Table 2. Licensing, the post-benchmark period, and exposure to generative AI

Dependent variable: Δ Job Zone. Origin two-digit SOC fixed effects included. N = 123,717.

	(1) Model A (2023)	**(2)** Model B (2024)
HighBarrier (level)	0.0150 (0.0201)	0.0425* (0.0218)
HighBarrier × Post	0.0201 (0.0168)	0.0037 (0.0210)
AI exposure, origin occupation	−2.6882**** (0.0274)	−2.6723**** (0.0261)
HighBarrier × AI exposure	0.0423 (0.0403)	−0.0783 (0.0494)
Post × AI exposure	0.0406** (0.0174)	0.0411* (0.0239)
HighBarrier × Post × AI exposure	**−0.0599*** (0.0333)	−0.0686 (0.0427)
R^2	0.7296	0.7297
Adjusted R^2	0.7286	0.7287
Within R^2	0.2963	0.2963
Root MSE	0.3668	0.3668
$F(33, 50)$	3,588.57	2,914.69
Clusters	51	51
Singletons dropped	8	8

Note. AI exposure is measured on the occupation the worker left, normalized to the unit interval (mean 0.5472, SD 0.2477). Destination-anchored exposure is collinear with the 319 destination-occupation fixed effects and cannot be estimated; see Section 4.2. Standard errors clustered on destination state in parentheses. **** p < .001, *** p < .01, ** p < .05, * p < .10.

The triple interaction (Table 2) is negative in both models: −0.0599 (SE 0.0333, p = .078) under the 2023 index, and −0.0686 (SE 0.0427, p = .115) under the 2024 index. Model A clears ten percent; Model B agrees in sign and magnitude but not statistically significant. The coefficients are easier to read as fitted effects at different levels of exposure. Under Model A, the post-benchmark change in the high-barrier-destination premium is +0.0022 at low exposure (one standard deviation below the mean), −0.0127 at average exposure, and −0.0275 at high exposure. Under Model B the same calculation gives −0.0169, −0.0339, and −0.0509.

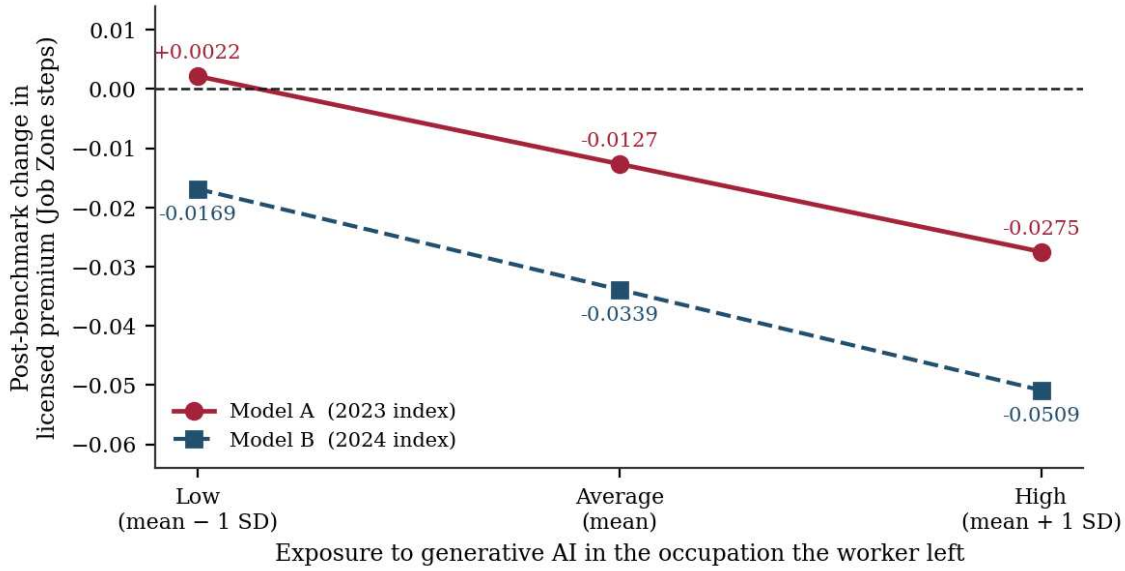


Figure 7. *The erosion by exposure to artificial intelligence.*

Note. Each point is the fitted post-benchmark change in the high-barrier premium, $\beta_1 + \beta_2 \times$ exposure, from Table 2, where β_1 is HighBarrier × Post (Model A: 0.0201; Model B: 0.0037) and β_2 is the triple interaction (−0.0599; −0.0686). Exposure is that of the occupation the worker left (mean 0.5472, SD 0.2477): low = 0.2995, average = 0.5472, high = 0.7949. Example, Model A at high exposure: 0.0201 + (−0.0599 × 0.7949) = −0.0275. Values use unrounded coefficients. Only Model A's triple interaction is statistically significant (p = .078).

The pattern is clearest under the 2023 index. For a worker leaving a lightly exposed occupation (+0.0022), the high-barrier route upward is essentially unchanged after 2023; for one leaving a heavily exposed occupation it has closed, with a fitted change of −0.0275, roughly three hundredths of a Job Zone step. Under the 2024 index the erosion is present at every exposure level (−0.0169 at low exposure) but also grows with exposure (−0.0509 at high exposure), although the triple interaction is not statistically significant. This is what H4 predicts, and it is hard to reconcile with the leading alternative. Post-pandemic normalization, the end of the Great Resignation's churn, and 2023 monetary tightening all share our benchmark's timing, but none has any reason to sort workers by their occupation's exposure to a language model. The interaction converts an argument about coincidence into an argument about mechanism.

We read this result as suggestive; Appendix A.5 explains why.

6. Robustness and Validity Checks

Four checks ask whether the result reflects a real institutional response or an artifact of the data. The **pre-trend test** asks whether high-barrier and low-barrier destinations were already drifting apart before the benchmark, which would undercut the parallel-trends assumption. Model B passes cleanly; Model A passes with a qualification (two of eleven pre-period coefficients

significant, no pattern suggesting a real pre-existing trend) (Appendix A.1–A.2). The **placebo test** asks whether an invented benchmark year, one where nothing happened, would produce the same break. It does not, in either model, which is evidence the true benchmark is doing real work (Appendix A.3).

The **overlap-sample check** asks whether the result depends on which destination cells each index vintage happens to observe. It does not: restricting both models to the same cells leaves the estimates essentially unchanged (Appendix A.4). The **headroom test** asks the more consequential question: whether AI exposure captures technological displacement at all, or just how much room a worker's origin occupation leaves to climb (Appendix A.5). Full results and tables for all four checks are in Appendix A.

7. Discussion

The central finding can be stated simply: moves into high-barrier occupations were associated with upward mobility, but in the most recent data that association has disappeared, most clearly for workers leaving occupations highly exposed to generative AI. Before the benchmark year, moving into a high-barrier occupation was associated with moving up (Section 5.1); after it, that association is cancelled, unevenly, with the largest decline for workers leaving heavily AI-exposed occupations (Sections 5.2, 5.4).

The interpretation turns on one asymmetry: **the toll is a price that cannot adjust, and a price that cannot adjust cannot clear the market for entry.** The barrier itself did not change within either model's own vintage — the index is a stable cell characteristic across each model's before-and-after boundary (Section 4.2). What changed is what moving up is worth: AI narrows the productivity gap between novice and experienced workers (Noy and Zhang 2023; Brynjolfsson et al. 2025), and its exposure is concentrated in cognitive occupations (Felten et al. 2023; Eloundou et al. 2024) that sit next to licensed work. We take this to lower the value of experience for workers leaving those occupations, so the gain from a high-barrier move shrinks. The fixed toll does not adjust, for the reasons in Stigler (1971) and Olson (1965), and so claims a larger share of the smaller gain. The upward-mobility premium from entering a high-barrier occupation eroded after the post-period began, most among workers leaving the most AI-exposed occupations. The licensing literature has largely studied the average effect of the barrier on wages, labor supply and entry (Kleiner 2000; Kleiner and Krueger 2013; Blair and Chung 2019; Jackson 2023); we ask whether the barrier's effect changed when generative AI arrived. The argument is dynamic: a static barrier can have a time-varying effect, which may rise when a technology shock such as generative AI increases the number of workers who need to move. It concerns institutional adaptability rather than how high the barrier is on average, and public

choice theory is well placed to make it, because the barrier does not adapt: adapting would require its beneficiaries' consent.

Several rival explanations deserve attention. The most obvious is that the post-period is confounded: 2023 also brought post-pandemic normalization and monetary tightening. But the erosion sorts on occupational AI exposure, which no macroeconomic rival predicts, and that heterogeneity carries the interpretive weight (Section 5.4). A more serious threat is that composition, not the entry barrier, is changing. We test this directly: origin fixed effects absorb the composition of the occupations workers leave (Section 5.3), and the treatment coefficients barely move, the signature of a result not driven by composition. A different channel is also plausible: the destination job itself may be getting automated. Workers' common transition destinations concentrate in a narrow band of adjacent occupations (Plaisance et al. 2026), and if those destinations are AI-exposed, part of the erosion could run through automation of the job entered. We cannot test this channel, since destination-side exposure is collinear with our destination fixed effects (Section 4.2), so it remains open. A related channel is rationing: if AI sends more workers toward the same licensed occupations, seats, exam cycles, and supervised-practice slots may bind. Our data cannot separate this from the devaluation channel, and neither the timing result nor the exposure result requires it.

Four limits qualify these results. First, no barrier law switches on at the benchmark year, so the design identifies a change in association across a fixed grouping, which is weaker than a policy-shock difference-in-differences. Second, the barrier measure is binary: it tells us whether a barrier exists, not how burdensome it is. Third, AI exposure correlates 0.6068 with a worker's starting position; the exposure gradient survives a placebo and a headroom-absorption check but not a joint test against a destination-anchored measure (Appendix A.5), so we treat it as suggestive. Fourth, Model B rests on a single post-year, Model A's result is significant only at ten percent, and wild cluster bootstrap inference is unavailable under our multi-way fixed-effect structure (Section 4.3); reporting both models together keeps either from being overstated alone. Covariates absorbed at the two-digit level should be read as descriptive rather than causal, and because the analysis is entirely American, cross-national variation in licensing regimes is a natural next test of whether the mechanism generalizes.

If the pattern survives further tests, the implications concern not whether to abolish licensing but whether entry regulation can respond to shocks. If the toll is what claims a larger share of a shrinking gain, the clearest policy options are those that reduce the toll: universal recognition of out-of-state licenses, interstate compacts, exam-based alternatives to seat-time mandates, and periodic board review (Johnson and Kleiner 2020; U.S. Department of the Treasury et al. 2015). Our data do not say how much of the post-2023 pattern such reforms would recover, because they cannot separate a toll that weighs more heavily from a destination that is worth less; if the

second dominates, easing entry would help less. Section 5.4 suggests where reform would matter most, among workers leaving AI-exposed occupations, but that gradient is suggestive rather than established.

The politics are the familiar politics of concentrated benefits and diffuse costs (Olson 1965): practitioner-dominated boards are unlikely to vote to relax their own entry requirements, so reform would have to come from legislatures rather than the boards, which is also a public choice reason to expect the toll to persist. Workforce agencies deserve the same scrutiny: rewarded on placement rates and speed, they have reason to supply the fast match rather than the slower licensed one (Niskanen 1971; Courty and Marschke 2004; Heckman et al. 2002). Metrics crediting skill-level gains would align the bureau's incentives with the worker's, though gaming remains a risk.

8. Conclusion and Future Research

Our data show that, among workers who change occupations, American occupational mobility is on net lateral. Across fifteen years of linked occupation changes, the average movement up or down the preparation scale is 0.0081 Job Zone steps, eight thousandths of one step. Workers who change occupations therefore show almost no net movement in required preparation. This paper asked whether the political architecture of occupational entry contributes to this absence of net upward movement, and whether its contribution intensified as artificial intelligence began reorganizing work. The answer is limited in scope but informative.

Under the 2023 index, entering a high-barrier occupation carried a measurable upward premium before the post-period (0.0287 Job Zone steps) and carries almost none after (0.0108); under the 2024 index the post-period association is slightly negative (−0.0338). The erosion is 0.0179 to 0.0409 steps, a fifth to two-fifths the size of the education gradient, surviving pre-trend, placebo, and overlap checks. Most importantly, the erosion is not uniform: it is larger among workers leaving occupations more exposed to generative AI and, under the 2023 index, close to zero among those leaving the least exposed. That gradient helps separate a political-toll account from the macroeconomic explanations that share our dates. This is not a causal finding. It is a pattern consistent with a public choice account in which a toll that is charged on credentials and slow to adjust claims a larger share of a shrinking gain when a shock lowers the value of moving up, and it is inconsistent with the simpler account in which the entry barrier merely raises the average cost of moving up.

Four extensions would sharpen the analysis: (1) an AI-exposure measure less correlated with occupational rank; (2) a licensing-burden measure and policy variation from states' staggered adoption of universal-recognition reforms (Goodman-Bacon 2021; de Chaisemartin and

D'Haultfœuille 2020), with the analysis focused on eligible interstate moves; (3) finer origin-occupation resolution; and (4) a binary upward-move outcome with wild-cluster-bootstrap inference as a robustness check.

The larger question this paper poses and does not answer is whether institutions built by incumbents to ration entry can accommodate a technology that changes who needs to enter. Public choice theory suggests they will find it hard to, since accommodation requires the consent of those who profit from its absence, and the evidence here is consistent with that: the upward premium erodes most for those leaving the occupations the technology exposes most.

References

Acemoglu, Daron, and David H. Autor. 2011. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, ed. Orley Ashenfelter and David Card, vol. 4B, 1043–1171. Amsterdam: Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

Autor, David H. 2019. Work of the past, work of the future. *AEA Papers and Proceedings* 109: 1–32. https://doi.org/10.1257/pandp.20191110

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118: 1279–1333. https://doi.org/10.1162/003355303322552801

Becker, Gary S. 1964. *Human capital: A theoretical and empirical analysis, with special reference to education*. New York: National Bureau of Economic Research, distributed by Columbia University Press.

Blair, Peter Q., and Bobby W. Chung. 2019. How much of a barrier to entry is occupational licensing? *British Journal of Industrial Relations* 57: 919–943. https://doi.org/10.1111/bjir.12470

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. Generative AI at work. *Quarterly Journal of Economics* 140: 889–942. https://doi.org/10.1093/qje/qjae044

Buchanan, James M., and Gordon Tullock. 1962. *The calculus of consent: Logical foundations of constitutional democracy*. Ann Arbor: University of Michigan Press.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2024. Difference-in-differences with a continuous treatment. NBER Working Paper no. 32117. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w32117

Courty, Pascal, and Gerald Marschke. 2004. An empirical investigation of gaming responses to explicit performance incentives. *Journal of Labor Economics* 22: 23–56. https://doi.org/10.1086/380402

Cunningham, Evan. 2019. Professional certifications and occupational licenses: Evidence from the Current Population Survey. *Monthly Labor Review*, June. U.S. Bureau of Labor Statistics. https://doi.org/10.21916/MLR.2019.15

Dal Bó, Ernesto. 2006. Regulatory capture: A review. *Oxford Review of Economic Policy* 22: 203–225. https://doi.org/10.1093/oxrep/grj013

de Chaisemartin, Clément, and Xavier d'Haultfœuille. 2020. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110: 2964–2996. https://doi.org/10.1257/aer.20181169

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

Felten, Edward W., Manav Raj, and Robert C. Seamans. 2023. Occupational heterogeneity in exposure to generative AI. SSRN working paper. https://doi.org/10.2139/ssrn.4414065

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. 2025. Integrated Public Use Microdata Series, Current Population Survey: Version 13.0. Dataset. Minneapolis: IPUMS. doi:10.18128/D030.V13.0.

Friedman, Milton. 1962. *Capitalism and freedom*. Chicago: University of Chicago Press.

Gathmann, Christina, and Uta Schönberg. 2010. How general is human capital? A task-based approach. *Journal of Labor Economics* 28: 1–49. https://doi.org/10.1086/649786

Geloso, Vincent, Alicia Plemmons, and Pradyot Sharma. 2026. Income mobility, automation, and occupational licensing. *Southern Economic Journal*, published online May 13. https://doi.org/10.1002/soej.70037

Goodman-Bacon, Andrew. 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225: 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review* 104: 2509–2526. https://doi.org/10.1257/aer.104.8.2509

Heckman, James J., Carolyn Heinrich, and Jeffrey Smith. 2002. The performance of performance standards. *Journal of Human Resources* 37: 778–811. https://doi.org/10.2307/3069617

Hunt, Jennifer, and Ryan Nunn. 2022. Has U.S. employment really polarized? A critical reappraisal. *Labour Economics* 75: 102117. https://doi.org/10.1016/j.labeco.2022.102117

Jackson, Osborne. 2023. "Occupational Licensing and Occupational Mobility." Federal Reserve Bank of Boston Research Department Working Papers No. 23-18. https://doi.org/10.29412/res.wp.2023.18.

Johnson, Janna E., and Morris M. Kleiner. 2020. Is occupational licensing a barrier to interstate migration? *American Economic Journal: Economic Policy* 12: 347–373. https://doi.org/10.1257/pol.20170704

Kambourov, Gueorgui, and Iourii Manovskii. 2008. Rising occupational and industry mobility in the United States: 1968–97. *International Economic Review* 49: 41–79. https://doi.org/10.1111/j.1468-2354.2008.00473.x

Kleiner, Morris M. 2000. Occupational licensing. *Journal of Economic Perspectives* 14: 189–202. https://doi.org/10.1257/jep.14.4.189

Kleiner, Morris M., and Alan B. Krueger. 2013. Analyzing the extent and influence of occupational licensing on the labor market. *Journal of Labor Economics* 31: S173–S202. https://doi.org/10.1086/669060

Kleiner, Morris M., and Evan J. Soltas. 2023. A welfare analysis of occupational licensing in U.S. states. *Review of Economic Studies* 90: 2481–2516. https://doi.org/10.1093/restud/rdad015

Krueger, Anne O. 1974. The political economy of the rent-seeking society. *American Economic Review* 64: 291–303.

MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb. 2023. Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics* 232: 272–299. https://doi.org/10.1016/j.jeconom.2022.04.001

Madrian, Brigitte C., and Lars John Lefgren. 2000. An approach to longitudinally matching Current Population Survey respondents. *Journal of Economic and Social Measurement* 26: 31–62. https://doi.org/10.3233/JEM-2000-0165

McChesney, Fred S. 1987. Rent extraction and rent creation in the economic theory of regulation. *Journal of Legal Studies* 16: 101–118. https://doi.org/10.1086/467825

National Center for O*NET Development. n.d. O*NET OnLine help: Job zones. U.S. Department of Labor, Employment and Training Administration. Accessed September 1, 2026. https://www.onetonline.org/help/online/zones

Neffke, Frank M. H., Anne Otto, and Antje Weyh. 2017. Inter-industry labor flows. *Journal of Economic Behavior & Organization* 142: 275–292. https://doi.org/10.1016/j.jebo.2017.07.003

Niskanen, William A., Jr. 1971. *Bureaucracy and representative government*. Chicago: Aldine-Atherton.

Noy, Shakked, and Whitney Zhang. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381: 187–192. https://doi.org/10.1126/science.adh2586

O*NET OnLine. n.d. O*NET OnLine. U.S. Department of Labor, Employment and Training Administration. Accessed September 1, 2026. https://www.onetonline.org

Olson, Mancur. 1965. *The logic of collective action: Public goods and the theory of groups*. Cambridge, MA: Harvard University Press.

Peltzman, Sam. 1976. Toward a more general theory of regulation. *Journal of Law and Economics* 19: 211–240. https://doi.org/10.1086/466865

Plaisance, Daniel, Aleister Montfort, Liat Krawczyk, and Melissa Pumphrey. 2026. Trapped workers: A network analysis of worker mobility in the AI economy. Issue brief, Bipartisan Policy Center, July 23. https://bipartisanpolicy.org/issue-brief/trapped-workers-a-network-analysis-of-worker-mobility-in-the-ai-economy/

Poletaev, Maxim, and Chris Robinson. 2008. Human capital specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000. *Journal of Labor Economics* 26: 387–420. https://doi.org/10.1086/588180

Roth, Jonathan. 2026. Interpreting event-studies from recent difference-in-differences methods. *Japanese Economic Review* 77: 275–288. https://doi.org/10.1007/s42973-026-00235-x

Stigler, George J. 1971. The theory of economic regulation. *Bell Journal of Economics and Management Science* 2: 3–21. https://doi.org/10.2307/3003160

Tollison, Robert D. 1982. Rent seeking: A survey. *Kyklos* 35: 575–602. https://doi.org/10.1111/j.1467-6435.1982.tb00174.x

Trudeau, Noah, and Edward Timmons. 2023. *State occupational licensing index 2023*. Washington, DC: Archbridge Institute. https://www.archbridgeinstitute.org/state-occupational-licensing-index-2023/

Trudeau, Noah, and Edward Timmons. 2024. *State occupational licensing index 2024*. Washington, DC: Archbridge Institute. https://www.archbridgeinstitute.org/state-occupational-licensing-index-2024/

Tullock, Gordon. 1967. The welfare costs of tariffs, monopolies, and theft. *Western Economic Journal* 5: 224–232. https://doi.org/10.1111/j.1465-7295.1967.tb01923.x

U.S. Census Bureau and U.S. Bureau of Labor Statistics. 2025. *Current Population Survey*. Accessed September 1, 2026. https://www.census.gov/programs-surveys/cps.html

U.S. Department of the Treasury, Office of Economic Policy, Council of Economic Advisers, and U.S. Department of Labor. 2015. *Occupational licensing: A framework for policymakers*. Washington, DC: The White House. https://obamawhitehouse.archives.gov/sites/default/files/docs/licensing_report_final_nonembargo.pdf

Zhang, Ting, and Laurie Schintler. 2025. The evolving role of knowledge, skill, and ability similarity in occupational mobility: A test of the 'similarity trap' in the AI era. SSRN working paper. https://doi.org/10.2139/ssrn.5934894

Supplementary Files

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

- 0.Appendix.docx