The Intensive and Extensive Margins of the Decline in US Manufacturing

Bumsoo Kim* Williams College

Jaeeun Seo † UIUC

March 26, 2026

Click here for most updated version.

Abstract

We show that the decline of U.S. manufacturing is driven by the extensive margin: young workers disproportionately avoid entry into declining sectors, while incumbents exit through retirement. We document that the labor supply elasticity with respect to wages is 7 times higher for new entrants than for incumbents. We develop a model that shows this high-elasticity entrant margin creates quantity amplification and wage attenuation in the labor market in response to sectoral shocks, and a smaller underlying demand shock is sufficient to explain the observed decline compared to models lacking this margin. A multi-region quantitative application of our framework replicates key empirical facts, including local-level recoveries driven by young workers entering service sectors. Our results suggest adjustment policies may be more effective when targeting the high-elasticity extensive margin of new entrants rather than the low-elasticity intensive margin of incumbents.

*Williams College Department of Economics. Email: bk12@williams.edu

†UIUC Department of Economics. Email: sje424@gmail.com

1 Introduction

U.S. manufacturing has been in retreat for generations. Between 2000 and 2017, the sector shed roughly six million jobs, even as real output held steady. A vast literature links that contraction to a combination of import competition, particularly the "China shock" (Autor et al., 2013, 2021), rapid automation (Acemoglu and Restrepo, 2020), and other structural factors (Charles et al., 2019). But does this adjustment come mainly from incumbent workers switching out of declining sectors (the intensive margin), or from retirement of experienced workers and new entrants choosing not to step onto the shop-floor career ladder (the extensive margin)?

The distinction between these two margins is important not only to describe worker flows per se, but because it helps identify the *elasticity of sectoral reallocation* that governs the economy's response to shocks and policy. If manufacturing contracts mainly through incumbents switching sectors, then the key object is the reallocation elasticity of experienced workers, and counterfactuals, such as the employment and wage effects of tariffs or industrial policy, depend on how readily those workers move when relative wages change.

If instead adjustment occurs primarily through young workers choosing different first jobs and older workers retiring, then aggregate adjustment is driven by the elasticity of entry and initial sector choice. These alternative elasticities imply different general-equilibrium paths for sectoral wages and employment since they jointly adjust. This equilibrium adjustment will also change which policies are most effective, from retraining and adjustment assistance targeted at incumbents to interventions that directly affect job creation and entry.

In this paper, we study the extensive margin in the decline of U.S. manufacturing by combining new stylized facts with a life-cycle model of sectoral choice. We begin with an accounting decomposition of the decline in the manufacturing employment share between 1980 and 2019 and find that approximately 80% of this decline is due to the extensive margin, with older workers exiting manufacturing jobs and a corresponding lack of entry by young workers. We then quantify the elasticities of adjustment and find that young entrants are far more responsive to sectoral wage differentials than old incumbents, with an elasticity 7 times larger.

This heterogeneity changes equilibrium adjustment: negative shocks to manufacturing generate smaller wage movements but larger manufacturing employment responses than models that treat mobility as uniform across workers. In a multi-region quantitative exercise, incorporating the entry margin reduces the size of the inferred manufacturing demand shock needed to match observed import exposure by 15.7 %.

A large literature following Artuç et al. (2010) provides a tractable framework for estimating the mobility elasticity governing sectoral reallocation. We build on this approach by introducing overlapping generations that allow incumbents and entrants to differ in their mobility elasticity. While heterogeneity in reallocation costs sets baseline mobility levels, it is the heterogeneity in elasticities that governs marginal responses to counterfactual shocks and policy

experiments. Our estimation separately identifies these two objects: in our preferred specification, incumbents respond to a 1% decrease in wages by leaving 0.38 % more, while entrants respond by entering 2.71 % less—an elasticity ratio of approximately 7 that is stable across data sources, estimation methods, and assumptions about the cost structure.

In general equilibrium, we derive two theoretical implications of the elasticity gap. First, because entry and retirement decisions are far more elastic than incumbent reallocation, the aggregate labor supply to a sector is more responsive to shocks than incumbent switching alone would suggest. Negative demand shocks therefore produce smaller wage declines and larger employment responses than models with a uniform mobility elasticity. This creates a calibration wedge: attributing a given employment decline entirely to incumbent responses requires a larger underlying demand shock, overstating both the shift in labor demand and the welfare costs borne by affected workers.

Second, the composition of adjustment is nonlinear in the size of the shock. For small shocks, the more elastic entry margin absorbs a disproportionate share of the response, limiting the need for costly incumbent reallocation. As shocks accumulate, entry into the declining sector falls toward zero and additional adjustment increasingly falls on the less elastic incumbent margin. This mechanism generates the pattern documented in Section 2: young workers exit first through reduced entry, while incumbents bear the remaining burden of a protracted decline.

The elasticity gap also has direct implications for policy. In Section 5, we study a policymaker who wants to raise employment in a target sector at minimum fiscal cost. The logit structure implies that the cost of moving one additional worker into a sector through a margin is $\nu_m/(1-p)$, where ν_m is the dispersion parameter and p is the baseline probability on that margin. Because entrants have an elasticity seven times higher than incumbents, entry-directed instruments—hiring credits, apprenticeship pipelines, vocational training subsidies—are correspondingly cheaper per additional worker than incumbent-targeted instruments such as retraining programs or retention subsidies. Retention subsidies are especially costly: the high baseline probability of staying (typically 80-90%) means that nearly every recipient would have remained without the subsidy.

This cost ranking is invariant to general-equilibrium wage feedback. Subsidies that raise sectoral employment depress the sectoral wage, which partially offsets the supply response, but this markup applies equally to all instruments and so preserves the ratio. Finally, we note that expanding sectoral employment and compensating workers displaced by sectoral decline are distinct objectives that operate on different margins: the cheapest way to maintain manufacturing jobs is not the same policy as the most direct way to support the workers who lose them.

Our quantitative simulations in Section 6 discipline the mechanism with data and quantify its general-equilibrium and distributional implications in a multi-region setting. The model

reproduces the key empirical finding that manufacturing adjustment is driven primarily by cohort replacement, and that employment recoveries in trade-exposed places occur mainly through new cohorts entering low-wage service jobs. Quantitatively, the entrant margin acts as a local buffer in high-manufacturing states, muting manufacturing wage losses by shifting adjustment toward quantities. It also generates a sizable calibration wedge: relative to a canonical model without entry, matching the same import-exposure facts requires a 15.7 % smaller underlying trade shock.

Related Literature

Recent micro-evidence supports our emphasis on the extensive margin. Autor et al. (2025b) show that commuting zones that lose manufacturing jobs regain employment largely through young workers entering service-sector jobs, while incumbent manufacturing workers rarely switch sectors. We provide a structural interpretation of this pattern and show how it helps reconcile the continued decline of manufacturing employment after the China shock (Autor et al., 2021) with cohort-driven local recoveries and limited incumbent reallocation.

Our paper contributes to the literature embedding forward-looking labor supply into quantitative trade models, pioneered by Artuç et al. (2010) and applied to import competition (Caliendo et al., 2019; Dix-Carneiro et al., 2023; Kim et al., 2024; Rodriguez-Clare et al., 2025), trade policy (Caliendo and Parro, 2022), place-based policies (Donald et al., 2025; O'Connor, 2025), capital accumulation (Kleinman et al., 2023), refugee inflows (Brinatti and Guo, 2024; Caliendo et al., 2023), and environmental impact (Balboni, 2025). In these models, the dispersion of idiosyncratic preference shocks is a key parameter governing equilibrium adjustment and welfare. We show that this elasticity differs sharply between entrants and incumbents, and that recognizing this heterogeneity materially changes counterfactual predictions for wages, employment dynamics, and the distributional consequences of trade shocks.

Our framework also connects to the literature on long-run structural transformation (Pierce and Schott, 2016; Adão et al., 2024). Recent work shows that the decline of agriculture was similarly driven by demographic turnover rather than incumbent reallocation (Porzio et al., 2022; Budı-Ors, 2025), and analogous patterns hold for broader structural change (Cosar, 2013; Hobijn et al., 2018). We extend this line of inquiry by providing quantitative estimates of the cohort gap in mobility elasticities and tracing its implications for the incidence of trade shocks.

Finally, our policy analysis in Section 5 connects to a growing literature on welfare and optimal policy in discrete-choice labor economies (Mongey and Waugh, 2024; Donald et al., 2025). Rather than characterizing welfare-maximizing policy—which depends on the missing insurance markets—we ask which instruments achieve a given employment target at minimum fiscal cost, exploiting the fact that the same dispersion parameters governing equilibrium adjustment also determine the relative cost of policy instruments across margins.

2 Stylized Facts on the Margins of Manufacturing Decline

This section documents *how* U.S. manufacturing shrank: through incumbent reallocation across sectors (the intensive margin) or through older workers leaving and young workers not entering (the extensive margin). Using linked CPS microdata, we decompose annual changes in manufacturing employment into these margins and find that the extensive margin accounts for four-fifths of the long-run decline. The pattern is strongly life-cycle dependent: the margin that falls the most is entry by young workers. This margin is central to our model because first-job sector choice is more wage-elastic than sector switching by incumbents.

We begin with the linked CPS panel and our state classification, then use raw transition matrices to characterize life-cycle mobility, and finally document three facts: a post-2000 break in manufacturing employment, a decline dominated by the extensive margin, and the role of generational replacement.

2.1 Data and Measurement

We use CPS Basic Monthly microdata from IPUMS (Flood et al., 2024), 1982–2019, and construct one-year links using the CPS month-in-sample structure (canonical MIS 4→16 match) and standard validation rules (consistent sex, race, quarter-of-birth/age, and state), dropping ambiguous matches; details are in Appendix B.1. Earnings are taken from the outgoing rotation groups and linked to subsequent one-year transitions using the same structure. The resulting linked sample contains large annual panels (90,000 matched individuals per year; Appendix B.1 reports match rates and sample sizes), which is sufficient to estimate aggregate transition rates and their time-series evolution. We restrict attention to men aged 20–64 to abstract from secular changes in women's labor-force participation; the Appendix reports results for the full sample. All moments use CPS survey weights.

We classify each individual into one of four mutually exclusive states: manufacturing (M), agriculture (A), service (S), and non-employment (N). Unemployed individuals who report an industry are treated as attached to that sector. Thus, switches between M and $\{A, S\}$ are intensive-margin reallocations, while movements between M and N are extensive-margin transitions.[1] To isolate life-cycle entry and exit, we report results by age at t: young (20–29), prime-age (30–54), and old (55–64).

Table 1 highlights two features of life-cycle mobility. First, mobility is front-loaded: young workers are the most likely to move across sectors and between employment and non-employment. Second, near retirement non-employment is close to absorbing: older workers who leave employment rarely re-enter, and older workers in manufacturing are much more likely to move to

[1] Across our sample, 97% of unemployed workers report an industry, consistent with temporary separation or continued attachment, whereas about 95% of those out of the labor force report no industry.

Table 1: Gross flows between four sector groups by age group. Means with standard deviations in parentheses.

	Nonemployment	Agriculture	Manufacturing	Service
Panel A. Young (20-29)				
Nonemployment	0.838	0.007	0.020	0.136
	(0.028)	(0.002)	(0.009)	(0.022)
Agriculture	0.231	0.642	0.023	0.104
	(0.039)	(0.033)	(0.009)	(0.024)
Manufacturing	0.135	0.005	0.759	0.101
	(0.034)	(0.003)	(0.029)	(0.025)
Service	0.165	0.004	0.020	0.810
	(0.023)	(0.001)	(0.006)	(0.022)
Panel B. Prime-age (30-54)				
Nonemployment	0.915	0.005	0.012	0.068
	(0.010)	(0.002)	(0.003)	(0.008)
Agriculture	0.148	0.774	0.016	0.062
	(0.023)	(0.021)	(0.006)	(0.017)
Manufacturing	0.080	0.003	0.847	0.069
	(0.019)	(0.002)	(0.021)	(0.024)
Service	0.091	0.003	0.015	0.891
	(0.014)	(0.001)	(0.004)	(0.012)
Panel C. Old (55-64)				
Nonemployment	0.978	0.003	0.002	0.017
	(0.006)	(0.002)	(0.001)	(0.005)
Agriculture	0.193	0.753	0.012	0.043
	(0.030)	(0.030)	(0.009)	(0.021)
Manufacturing	0.147	0.003	0.794	0.057
	(0.032)	(0.002)	(0.023)	(0.024)
Service	0.154	0.002	0.012	0.832
	(0.027)	(0.001)	(0.005)	(0.025)

N than to another sector. Manufacturing can therefore shrink mainly through exits to N even when cross-sector job-to-job switching is limited. The next subsection quantifies this point using net-flow decompositions.

2.2 Flow decomposition

Let L_t^i denote the number of workers in state $i \in \{M, A, S, N\}$ at time t, and let F_t^{ij} denote the number moving from i to j between t and $t+1$. To account for entry to and exit from the CPS sample, as well as linkage failure, we add an outside state R. Then changes in sectoral

employment admit the exact decomposition

$$L^i_{t+1} - L^i_t = \underbrace{\sum_{j \in \{M,A,S\}} \left(F^{ji}_t - F^{ij}_t\right)}_{\text{intensive (between employed sectors)}} + \underbrace{\left(F^{Ni}_t - F^{iN}_t\right)}_{\text{extensive (employment} \leftrightarrow \text{non-employment)}} + \underbrace{\left(F^{Ri}_t - F^{iR}_t\right)}_{\text{outside / residual}}. \tag{1}$$

The first term is net reallocation across employed sectors, the second is net flow between the sector and non-employment, and the third absorbs entry to and exit from the 20–64 CPS sample and any remaining stock-flow discrepancy.

We report results both in headcounts and employment shares. The corresponding share decomposition (and the age-aggregation identity used to separate within-age flows from between-age composition) are provided in Appendix B.1. We also report age-specific versions of (1); in the life-cycle interpretation, we summarize "exit" as $F^{MN}_t(\text{old})$ and "entry" as $F^{NM}_t(\text{young})$.

Because linked CPS flows need not sum exactly to observed origin and destination stocks, we reconcile the flow matrix by biproportional fitting, requiring row sums to match L_t and column sums to match L_{t+1} while leaving the R margins free. Appendix B.1 provides details of this procedure and the broader data construction.

2.3 Stylized Facts

We document three facts that pin down the timing, margin, and age structure of manufacturing decline, with particular attention to the early 2000s. First, manufacturing headcount is roughly flat through the late 1990s and then falls sharply after 2000, even though its employment share declines throughout. Second, this contraction is driven mainly by flows between manufacturing and non-employment rather than by job-to-job reallocation across sectors. Third, those extensive-margin flows reflect generational turnover: older workers exit at high rates, while young workers enter manufacturing less often. Together, these facts motivate a model in which labor-supply adjustment operates primarily through entry and exit rather than incumbent sector switching.

Fact 1. *Manufacturing's employment share declines throughout the sample, but headcount falls sharply only after 2000.*

Figure 1 shows a clear divergence between manufacturing employment shares and levels. The employment share declines smoothly throughout the sample, while headcount is roughly flat through the late 1990s and then drops sharply after 2000. The distinction matters because a falling share can arise mechanically if employment grows faster elsewhere, whereas a decline in headcount requires persistent net worker losses. Fact 1 therefore isolates the timing of the contraction; Facts 2 and 3 show which flows generate it.

Fact 2. *The decline in manufacturing is dominated by the extensive margin.*

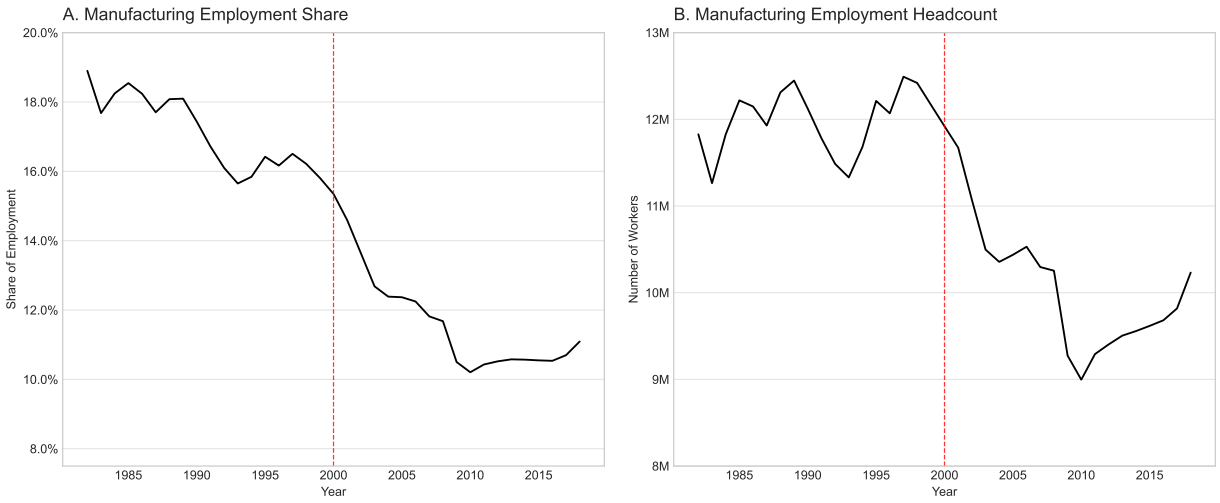


Figure 1: Divergence of Manufacturing Employment Share and Headcount.

Note. This figure is computed from the linked CPS sample to align with the flow-based decomposition. Appendix B.1 reproduces the same series using BLS employment totals and shows the same pattern.

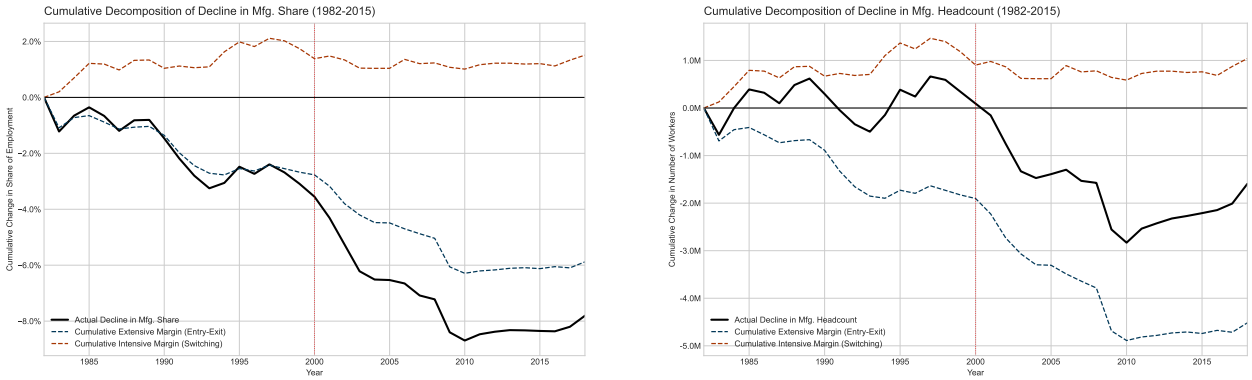
Figure 2 decomposes annual changes in manufacturing employment into net job-to-job flows across employed sectors (the intensive margin) and net flows between manufacturing and non-employment (the extensive margin), using Equation (1) and the share analogue in Appendix B.2. In both shares and levels, the intensive component is small relative to the total decline and close to zero on average; if anything, it is mildly positive before 2000. This does not contradict deindustrialization: manufacturing can lose headcount mainly through exits to non-employment even when net switching between employed sectors is limited or positive.

The contraction instead comes from the extensive margin. After 2000, net flows from manufacturing to non-employment rise sharply and are not offset by inflows from non-employment into manufacturing, accounting for most of the cumulative decline in headcount. The next fact shows that this extensive-margin adjustment reflects life-cycle turnover (older workers exit and young workers increasingly do not enter manufacturing) rather than temporary non-employment during sector-to-sector reallocation.

Fact 3. *The extensive margin reflects generational churn: high exit by older workers and low entry by young workers.*

Figure 3 disaggregates the net-flow decomposition by age at t and reveals the life-cycle structure behind the aggregate patterns. Panel A shows that the intensive margin is small for all age groups: net job-to-job switching is noisy and centered near zero, so direct cross-sector reallocation is not the main source of manufacturing decline.

The adjustment instead occurs along the extensive margin. Panel B shows that exits from manufacturing to non-employment are much larger for older workers (55–64) than for prime-age or young workers, consistent with retirement-like transitions. Panel C shows that these



(a) Change in share

(b) Change in headcount

Figure 2: Extensive vs Intensive Margin of Manufacturing Decline.

Note. The left panel reports the cumulative decomposition for employment shares; the right panel reports the corresponding decomposition for headcounts. The intensive margin is net job-to-job reallocation across employed sectors. The extensive margin is net flow between manufacturing and non-employment.

exits are not offset by young entrants (20–29): net flows from non-employment into manufacturing are small and trend downward over time.

Panel D shows the same pattern from the entry side. Manufacturing's share among new job entrants falls sharply for young cohorts, from about 20% in the early 1980s to below 10% by the 2000s. Together, these age patterns imply that manufacturing contracts mainly through cohort replacement—older workers leave, and younger workers increasingly start elsewhere—rather than through incumbent sector switching.

2.4 Discussions

Summary. Facts 1–3 suggest that manufacturing decline in the U.S. is primarily a cohort-replacement process: older workers exit to non-employment and are not replaced by young entrants, while cross-sector switching by incumbents plays only a limited role. This motivates a model that distinguishes entrants from incumbents and allows their mobility elasticities to differ.

Relation to prior evidence. These national flow patterns echo and extend the commuting-zone evidence in Autor et al. (2025b). In trade-exposed places, manufacturing employment falls, total employment later recovers in services, and incumbent reallocation explains only a small share of the trade-induced contraction; recovery is driven largely by young and foreign-born entrants. Our accounting decomposition shows that the same mechanism—limited incumbent switching and adjustment through cohort replacement—also appears in linked national microdata and over the broader secular decline, not only in China-shock-exposed local labor markets.

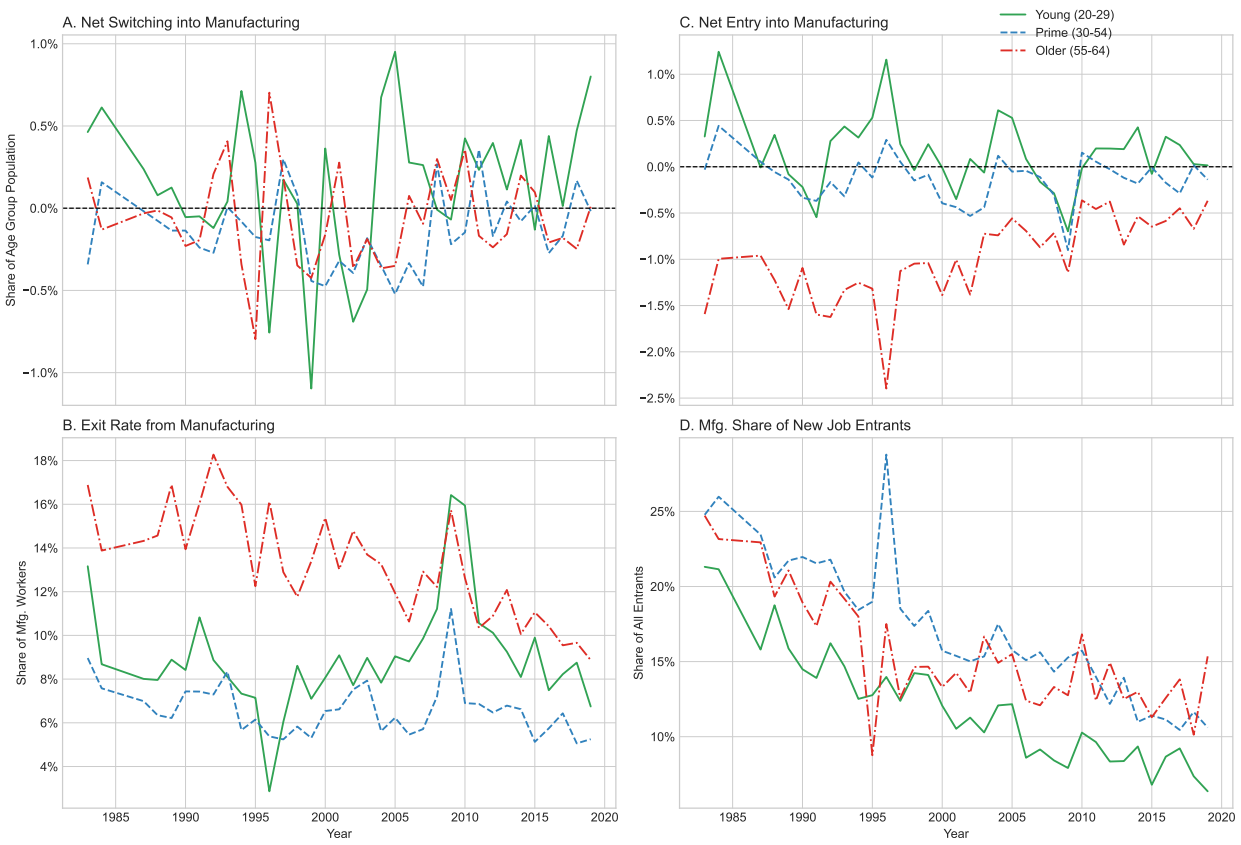


Figure 3: Extensive margin driven by old exit and lack of young entry.

Robustness. Appendix B.3 validates the CPS-based measures using two alternative sources. First, we replicate the key entry and margin decompositions in administrative LEHD Job-to-Job (J2J) data from 2001 onward. The CPS object most central to our mechanism—young inflows into manufacturing from non-employment—maps naturally to LEHD hires from persistent non-employment (`NEPersist`). The resulting entry-share series closely tracks the CPS measure and exhibits the same secular decline. LEHD job-to-job flows (`EE`) sharpen the interpretation: switching into manufacturing is flat or rising for young workers, while the contraction operates through the extensive inflow margin. Second, to avoid reliance on month-in-sample linking, we replicate the decomposition in the CPS ASEC ("March CPS"), which records both current industry and reported industry worked most in the previous calendar year. The results are similar.

Overall, the evidence points to a central role for entry and retirement. The next section develops a model with entrants and incumbents that captures this life-cycle structure and lets us study mobility elasticities, welfare, and policy during transition.

3 A Model of Labor Supply with Entry and Exit

Section 2 shows that manufacturing decline operates mainly through cohort replacement rather than incumbent sector switching: older workers exit to non-employment, and young cohorts increasingly do not enter manufacturing, even though gross worker churn remains substantial. The relevant object is therefore not the mobility of the average incumbent, but the *elasticity of the entry margin*—how strongly first-sector choices respond to expected wage paths.

We build a tractable labor-supply block that separates incumbent switching from entry and retirement, allows these margins to have different wage elasticities, and yields implications for wage pass-through, employment adjustment, and policy in general equilibrium. The model provides a structural interpretation of the flow evidence and a disciplined way to calibrate sectoral shocks.

We first define the problems of incumbents and entrants and the law of motion for sectoral employment. We then show how observed flows and wages identify two mobility elasticities, one for incumbents and one for entrants, which we estimate in Section 3.2. Section 4 shows that an elastic entry margin attenuates wage responses and amplifies employment adjustment, so models without entry require larger implied demand shocks to match the same quantity responses.

Environment. We adapt the dynamic discrete-choice framework of Artuç et al. (2010) to a perpetual-youth overlapping-generations setting as in Blanchard (1985). Time is discrete. Workers choose among S productive sectors $s \in \{1, \ldots, S\}$ and the non-employment state $s = 0$. Each period, a mass N_t of new entrants arrives and chooses an initial productive sector. Incumbents earn sectoral wages $\{w_t^s\}$, may switch sectors subject to mobility costs, and retire permanently with probability $1 - \theta$.

3.1 Incumbents and Entrants

Incumbents. An incumbent who starts period t in sector s earns wage w_t^s and chooses sector n for period $t + 1$. Moving from s to n entails a utility cost χ^{sn}, with $\chi^{ss} = 0$, and an idiosyncratic taste shock ϵ_t^n. Incumbents retire with probability $1 - \theta$ each period. Let $V_{t+1}^n = \mathbb{E}[\mathcal{V}_{t+1}^n(\epsilon)]$ denote the ex-ante value of starting period $t + 1$ in sector n. The realized lifetime utility for an incumbent in sector s is

$$\mathcal{V}_t^s(\epsilon) = u(w_t^s) + \max_n [\beta\theta V_{t+1}^n - \chi^{sn} + \epsilon_t^n]$$

Assuming ϵ_t^n is i.i.d. Type I Extreme Value with dispersion parameter $\nu > 0$, the probability that an incumbent in sector s chooses sector n is

$$\mu_t^{sn} = \frac{\exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^n - \chi^{sn}\right)\right)}{\sum_{n'}\exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^{n'} - \chi^{sn'}\right)\right)} \tag{2}$$

and the ex-ante value function V_t^s satisfies

$$V_t^s = u(w_t^s) + \nu \log\left(\sum_{n'}\exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^{n'} - \chi^{sn'}\right)\right)\right) \tag{3}$$

Entrants. At the end of period t, an entrant chooses an initial productive sector $n \in \{1, \ldots, S\}$. Choosing n yields the discounted continuation value of becoming an incumbent in that sector next period, net of an entry cost C^n, plus a taste shock $\epsilon_{t,ent}^n$:

$$\mathcal{V}_t^{ent}(\epsilon) = \max_n\left[\beta V_{t+1}^n - C^n + \epsilon_{t,ent}^n\right]$$

Assuming $\epsilon_{t,ent}^n$ is Type I Extreme Value with dispersion parameter ν^E, the share of entrants choosing sector n is

$$E_t^n = \frac{\exp\left(\frac{1}{\nu^E}\left(\beta V_{t+1}^n - C^n\right)\right)}{\sum_k \exp\left(\frac{1}{\nu^E}\left(\beta V_{t+1}^k - C^k\right)\right)} \tag{4}$$

Sectoral Labor Supply. Let L_t^s denote the number of incumbents in sector s at time t, and let N_t be the mass of new entrants. Sectoral employment then evolves according to

$$L_{t+1}^n = \sum_s \mu_t^{sn}\theta L_t^s + E_t^n N_t \tag{5}$$

Equation (5) nests the workhorse model without cohort replacement as the case $\theta = 1$ and $N_t = 0$. In the data, μ_t^{sn} corresponds to intensive-margin job-to-job transitions across employed sectors, while E_t^n corresponds to the entry margin proxied by young inflows from non-employment (Section 2).

Interpretation and use. Equations (2)–(5) characterize the labor-supply side of the model. Given a path of sectoral wages $\{w_t^s\}$, they determine sectoral values $\{V_t^s\}$, transition probabilities $\{\mu_t^{sn}\}$, entrant shares $\{E_t^n\}$, and employment dynamics $\{L_t^s\}$. The key sufficient statistics are the two dispersion parameters ν and ν^E, which govern the wage responsiveness of incumbent switching and entry. Because this labor block can be paired with any demand system mapping $\{L_t^s\}$ into $\{w_t^s\}$, it can be embedded in quantitative models of trade or structural

transformation. Section 3.2 estimates ν and ν^E, and Section 4 derives the resulting equilibrium implications.

3.2 Estimating mobility elasticities

The dispersion parameters ν and ν^E govern incumbent switching and entrant sector choice, so they are central for welfare and counterfactual analysis. We estimate them by adapting the approach in Artuç et al. (2010), which turns the model's optimality conditions into moment conditions.

Moment conditions. Rearranging the logit choice probabilities yields two Euler-style equations linking current choices to next-period wages and continuation values. For incumbents, comparing staying in sector s with moving to sector n implies

$$\log(\mu_t^{sn}/\mu_t^{ss}) = \frac{\beta\theta}{\nu}[u(w_{t+1}^n) - u(w_{t+1}^s)] + \beta\theta \log(\mu_{t+1}^{sn}/\mu_{t+1}^{nn}) - \frac{1-\beta\theta}{\nu}\chi^{sn} + \Omega_{t+1}^{sn} \tag{6}$$

For entrants,

$$\log E_t^n = \frac{\beta}{\nu^E}u(w_{t+1}^n) + \beta\theta \log(E_{t+1}^n) - \beta\theta \log \mu_{t+1}^{nn} + \alpha_t - \frac{1-\beta}{\nu^E}C^n + \Omega_{t+1}^{n,E} \tag{7}$$

Appendix A.1 derives both expressions. As in Artuç et al. (2010), the equations are recursive: current mobility depends on expected future wage gains and on future mobility objects that proxy continuation values. In (6), continuation values are proxied by future relative mobility. In (7), they are proxied by future entry E_{t+1} and future retention μ_{t+1}^{nn}.

Notice that the coefficient on the retention rate $\log \mu_{t+1}^{nn}$ in Equation (7) is $\beta\theta$. As we show in Appendix A.1, an unconstrained derivation gives an additional coefficient of $\frac{\nu}{\nu^E}$, where ν is the average elasticity of the entire incumbent pool. We instead assume that the relevant incumbent elasticity for a marginal entrant is the elasticity of a first-year incumbent and set it equal to the entrant elasticity, $\nu^{\text{1-year incumbent}} = \nu^E$, which yields the coefficient $\beta\theta$.[2] This restriction targets the continuation value relevant for the marginal entry decision and avoids estimating an unstable aggregate ratio.

The error terms Ω_{t+1}^{sn} and $\Omega_{t+1}^{n,E}$ are expectation errors and are therefore orthogonal to information available at time t. OLS is nevertheless inconsistent because shocks realized at $t+1$ affect both the realized right-hand-side variables (e.g., w_{t+1} and the continuation-value proxies μ_{t+1}, E_{t+1}) and the expectation errors. We therefore instrument the $t+1$ regressors with predetermined time-t information, using lagged sectoral wages and flows (for example, w_{t-1}, μ_{t-1}, and E_{t-1}) together with sector fixed effects, following Artuç et al. (2010). Appendix A.1 reports

[2] Appendix A.1 gives the full derivation and robustness exercises.

the exact instrument set and robustness to alternative lag structures.

Identification and Interpretation

Equations (6)–(7) give ν and ν^E a direct (semi-)elasticity interpretation:

$$\frac{1}{\nu} = \frac{1}{\beta}\frac{d\log\mu}{d\,u(w)}, \qquad \frac{1}{\nu^E} = \frac{1}{\beta}\frac{d\log E}{d\,u(w)}.$$

As we show in Section 4, these elasticities are sufficient statistics for a range of counterfactual and welfare exercises. For example, in steady state the lifetime utility of an incumbent in sector s can be written as $V^s = \frac{1}{1-\beta\theta}(u(w^s) - \nu\log\mu^{ii})$, which links welfare directly to ν and the observed stayer share.[3]

The estimation also separates mobility costs from mobility elasticities. Average flow levels identify cost terms relative to dispersion—that is, χ^{sn}/ν and C^n/ν^E—while the elasticities ν and ν^E are identified by how flows respond over time to wage variation. To allow for origin- and destination-specific mobility costs, we estimate stacked versions of the equations with origin-destination fixed effects. For comparability with the literature, we also report estimates under more restrictive cost structures, such as constant or destination-specific costs.

Our descriptive evidence suggests that entrants should be more wage-responsive than incumbents because they have accumulated less sector-specific human capital. In the model, this implies the testable restriction $\nu^E < \nu$.

Estimation requires annual sectoral transition rates and sectoral wages. We construct one-year linked CPS panels using the month-in-sample structure (canonical MIS 4→16 match), as described in Section 2 and Appendix B.1. The sample spans 1983–2019 and contains 2,991,559 unique individuals, or about 80,000 observations per year, weighted using standard CPS weights. For the stylized-facts accounting we use the four-state partition in Section 2; for estimation, we match the model's choice set by aggregating CPS industries into the six sectors in Artuç et al. (2010).[4]

Under this classification, incumbent transition probabilities μ_t^{sn} are computed from observed job-to-job reallocations. Sectoral wages w_t^s are constructed from outgoing rotation group weekly earnings, deflated and cleaned as in Section 2. We proxy entrant shares E_t^n using transitions of young workers (ages 20–29) from non-employment into sector n, which captures school-to-work and re-entry flows in the absence of true first-job histories. Throughout, we set $u(w) = \log w$, so the estimated coefficients map directly into the wage elasticities $\beta\theta/\nu$ and β/ν^E.

[3]This sufficient statistics approach follows (Caliendo et al., 2019; Rodriguez-Clare et al., 2025).

[4]The six sectors are: (1) agriculture and mining; (2) construction; (3) manufacturing; (4) transportation, communication and utilities; (5) trade; (6) other services.

Estimators. We estimate both the incumbent equation (6) and the entrant equation (7) by IV, using $t-1$ lags and sector dummies as instruments for the $t+1$ regressors. For the incumbent equation, we also report a two-stage PPML gravity estimator in levels following Artuç and McLaren (2015). Because PPML identifies ν but not the full bilateral cost matrix χ^{ij}, we report only the implied ν from that specification; Appendix A.1 provides details.

Table 2: Incumbent mobility parameter ν

	(1)	(2)	(3)
χ specification	Dyadic (χ^{ij})	Entry–exit ($\bar{\chi}_i + \tilde{\chi}_j$)	Constant (χ)
Panel A. IV			
	2.44^{*}	2.22^{*}	1.71^{***}
	(1.86)	(1.48)	(0.11)
Panel B. Two-stage PPML			
	1.14^{***}	2.48^{***}	1.31^{***}
	(0.07)	(0.24)	(0.10)
Observations		1050	

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$ (one-tailed). Switching cost estimates χ^{ij} reported in the appendix.

Table 3: Entry parameter ν^E

	(1)	(2)
C^i specification	C^i	$C^i = C$
ν^E	0.350^{***}	0.070^{***}
	(0.094)	(0.005)
Observations	175	

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$ (one-tailed). Entry costs are identified only up to a normalization and are reported in the appendix.

Tables 2 and 3 report estimates of the mobility dispersion parameters: ν for incumbents and ν^E for entrants. Because we set $u(w) = \log w$, these parameters map directly into wage elasticities, with $\varepsilon^{\text{inc}} = \beta\theta/\nu$ for incumbents and $\varepsilon^{\text{ent}} = \beta/\nu^E$ for entrants. Smaller dispersion parameters therefore imply greater wage responsiveness.

Incumbent mobility (ν). Table 2 shows that incumbent mobility is relatively inelastic across specifications. Under the benchmark IV specification with dyadic switching costs (column (1)), we estimate $\nu = 2.44$, implying $\varepsilon^{\text{inc}} \approx 0.38$. Alternative cost restrictions (columns (2)–(3)) and

the PPML estimator deliver similar magnitudes, indicating that incumbent sector switching responds only weakly to wage differentials.

Entry (ν^E). Table 3 shows much higher responsiveness on the entry margin. Under the preferred IV specification with sector-specific entry costs (column (1)), we estimate $\nu^E = 0.350$, implying $\varepsilon^{\text{ent}} \approx 2.71$. Imposing constant entry costs (column (2)) yields an even smaller ν^E, but sector-specific entry costs are more plausible, so we use 0.350 as the benchmark.

The implied elasticity gap between entrants and incumbents is

$$\frac{\varepsilon^{\text{ent}}}{\varepsilon^{\text{inc}}} = \frac{\beta/\nu^E}{\beta\theta/\nu} = \frac{\nu}{\theta\nu^E}.$$

Using the benchmark IV estimates and $\theta \simeq 0.97$, entrants are about seven times as wage-responsive as incumbents. Specifically, a 1% decline in a sector's expected wage path reduces entrant inflows by about 2.71 %, but raises incumbent outflows by only about 0.38 %.

This entrant–incumbent elasticity gap provides a structural interpretation of the generational-churn patterns in Section 2. Manufacturing can keep shrinking even when incumbent switching is limited because retirements generate steady outflows, while young cohorts disproportionately avoid entering sectors with deteriorating expected wage paths.

Robustness. Table 4 shows that the entrant–incumbent elasticity ratio is stable across specifications even when the levels of ν and ν^E move. Across Monthly CPS and March CPS (ASEC), alternative IV lags, different values of β, and samples excluding recession years, ν/ν^E ranges from 4.2 to 9.7 and usually lies between 6 and 7. By contrast, the levels of ν and ν^E vary more substantially across datasets—most notably between Monthly and March CPS—but they move together, leaving the ratio largely unchanged.[5] The entrant–incumbent gap therefore appears to be a robust feature of the data rather than an artifact of a particular dataset or timing assumption.

3.3 Discussions

Relation to Empirical Findings. Our finding that adjustment occurs primarily on the entry margin connects naturally to the commuting-zone evidence on trade exposure. It helps explain why manufacturing employment continues to decline even after trade shocks plateau around 2010, and why exposed places eventually recover overall employment through entrants

[5] One source of the level discrepancy is the timing problem in March CPS mobility documented by Kambourov and Manovskii (2013). Although the survey appears annual, respondents likely report jobs from a shorter reference window, which biases measured transition rates and therefore the level of ν. Artuç et al. (2010) address this by annualizing the observed transition matrix as $\tilde{\mu}^{12/K}$, with $K = 5$ months best matching NLSY flows. In our data, these corrections matter for levels but not materially for the ratio ν/ν^E.

Table 4: Robustness of entrant–incumbent elasticity ratio ν/ν^E

	ν	ν^E	ν/ν^E
Monthly CPS			
Baseline	2.44 (1.86)	0.35 (0.09)	7.02 (5.69)
Excl. recessions	2.59 (2.62)	0.61 (0.26)	4.22 (4.62)
IV lag $= 3$	1.77 (0.20)	0.27 (0.08)	6.45 (1.96)
$\beta = 0.90$	1.65 (0.89)	0.34 (0.09)	4.87 (2.94)
$\beta = 0.97$	2.97 (2.72)	0.35 (0.09)	8.47 (8.09)
March CPS (ASEC)			
Baseline	4.67 (2.77)	0.72 (0.23)	6.52 (4.40)
Excl. recessions	3.69 (1.99)	0.58 (0.16)	6.38 (3.84)
IV lag $= 3$	9.85 (3.84)	1.02 (0.53)	9.68 (6.33)
$\beta = 0.90$	4.26 (2.39)	0.70 (0.22)	6.11 (3.95)
$\beta = 0.97$	4.86 (2.97)	0.72 (0.23)	6.72 (4.64)
FE specification			
ν: Dyadic FE (Dest FE for IV lag $= 3$)			
ν^E: Sector FE			

Note: Standard errors in parentheses (delta method for ν/ν^E). Benchmark specification uses Monthly CPS with dyadic switching costs, sector-specific entry costs, IV lag $= 2$, and $\beta = 0.95$. Full results across all FE combinations are reported in the Appendix.

into services while incumbent employment remains depressed and manufacturing continues to shrink, as documented by Autor et al. (2021, 2025b). Our estimates suggest a mechanism for this pattern: incumbents, because of sector-specific human capital, respond only weakly to wage signals, so structural reallocation occurs mainly through new cohorts sorting into service-sector jobs. In the next section, we use counterfactuals with and without entry to show how allowing for this extensive margin changes the shocks required to match import-penetration paths and the distribution of incidence across cohorts.

Non-employment. Our estimation focuses on mobility across productive sectors. Following Artuç et al. (2010), we treat non-employment as an involuntary outcome rather than an active choice for incumbents. This keeps the model tractable and matches the descriptive evidence in Section 2, where flows into and out of non-employment look distinct from intersectoral switching. While this approach maintains a tractable choice set, we acknowledge it as a limitation, as it abstracts from potentially important dynamics at the extensive margin of employment.

4 The Elasticity Gap and Adjustment Dynamics

Section 3.2 yields two central estimates: incumbents are relatively immobile ($\nu \approx 2.44$), whereas entrants are about seven times more wage-responsive to expected wages ($\nu/\theta\nu^E \approx 7$). This section studies the equilibrium implications. We first show that baseline flows $\{\mu_t^{sn}, E_t^n\}$ together with the two elasticity composites are sufficient statistics for counterfactual dynamics, so the levels of mobility costs are not needed (Section 4.1). We then characterize how the entry margin shapes equilibrium adjustment. Locally, elastic entry amplifies employment responses and attenuates wage pass-through (Section 4.2), with the allocation of adjustment across margins governed by a slack-elasticity product (Section 4.3). For permanent shocks, it front-loads the transition and reduces the persistence of employment deviations (Section 4.4).

To study general equilibrium, we pair the labor-supply block with a generic inverse labor-demand system for each sector s:

$$w_t^s = f^s(L_t^s, \{\xi_t^n\}) \tag{8}$$

where L_t^s is employment in sector s at time t and $\{\xi_t^n\}$ denotes exogenous sector-specific demand shifters (for example, technology, trade prices, or preferences). We assume standard properties: $df^s/dL^s < 0$, $df^s/d\xi^s > 0$, and $df^s/d\xi^n < 0$. Equilibrium combines labor supply, given by Equations (2)–(5), with labor demand in Equation (8).

4.1 Elasticities as Sufficient Statistics

Following the exact hat-algebra approach of Dekle et al. (2007), we express counterfactual dynamics in changes relative to an observed equilibrium. For incumbents, this is standard: Caliendo et al. (2019) show that baseline transition probabilities absorb all information about switching costs, so counterfactual incumbent flows depend only on observed $\{\mu_t^{sn}\}$, changes in continuation values, and the elasticity composite $\beta\theta/\nu$.

We extend the same logic to entrants. Because entrant choices have the same logit structure, baseline entry shares $\{E_t^n\}$ absorb entry-cost levels, and the only additional supply parameter is the elasticity composite β/ν^E. The joint entrant–incumbent system can therefore be evaluated using observed baseline flows and the two elasticity composites alone.

Consider a counterfactual economy that shares the primitives $(\beta, \theta, \nu, \nu^E)$ and choice set of the baseline but allows different wages $\tilde{w}_t^s$ and possibly different costs $\tilde{\chi}^{sn} = \chi^{sn} + \Delta\chi^{sn}$ and $\tilde{C}^n = C^n + \Delta C^n$. Let $\Delta V_t^s := \tilde{V}_t^s - V_t^s$. Since $u(w) = \log w$, define wage changes in hat notation and policy wedges as

$$\hat{w}_t^s := \frac{\tilde{w}_t^s}{w_t^s}, \qquad \hat{\kappa}^{sn} := \exp\big(-\Delta\chi^{sn}/\nu\big), \qquad \hat{\kappa}_E^n := \exp\big(-\Delta C^n/\nu^E\big),$$

so that $\Delta u_t^s = \log \hat{w}_t^s$. We have the following:

Proposition 1 (Exact hat algebra for the entrant–incumbent block). *Suppose* $\{\log \hat{w}_t^s\}$ *is bounded. Then the system admits a unique bounded solution* $\{\Delta V_t^s\}$*, and the counterfactual path satisfies:*

$$\Delta V_t^s = \log \hat{w}_t^s + \nu \log\left(\sum_n \mu_t^{sn} \hat{\kappa}^{sn} \exp\left(\tfrac{\beta\theta}{\nu} \Delta V_{t+1}^n\right)\right), \tag{9}$$

$$\tilde{\mu}_t^{sn} = \frac{\mu_t^{sn} \hat{\kappa}^{sn} \exp\left(\tfrac{\beta\theta}{\nu} \Delta V_{t+1}^n\right)}{\sum_k \mu_t^{sk} \hat{\kappa}^{sk} \exp\left(\tfrac{\beta\theta}{\nu} \Delta V_{t+1}^k\right)}, \tag{10}$$

$$\tilde{E}_t^n = \frac{E_t^n \hat{\kappa}_E^n \exp\left(\tfrac{\beta}{\nu^E} \Delta V_{t+1}^n\right)}{\sum_k E_t^k \hat{\kappa}_E^k \exp\left(\tfrac{\beta}{\nu^E} \Delta V_{t+1}^k\right)}, \tag{11}$$

$$\tilde{L}_{t+1}^n = \theta \sum_s \tilde{L}_t^s \tilde{\mu}_t^{sn} + \tilde{N}_t \tilde{E}_t^n, \tag{12}$$

where $\tilde{N}_t$ *follows a population rule. The levels* $\{\chi^{sn}, C^n\}$ *matter only to the extent that they determine the baseline flows* $\{\mu_t^{sn}, E_t^n\}$*. Counterfactual dynamics depend on structural parameters through the elasticity composites* $(\beta\theta/\nu,\ \beta/\nu^E)$*; policy interventions that alter costs enter only through the wedges* $\hat{\kappa}^{sn}$ *and* $\hat{\kappa}_E^n$*.*

Equations (9)–(10) is exactly the dynamic hat algebra in Caliendo et al. (2019); equations (11)–(12) extend the system to entrants and the overlapping-generations law of motion. The implication is that the labor-supply side of the counterfactual never requires the levels of $\{\chi^{sn}, C^n\}$. These are absorbed entirely by the baseline flows $\{\mu_t^{sn}, E_t^n\}$. Thus the supply block contributes to counterfactual analysis only through the elasticity composites $\beta\theta/\nu$ and β/ν^E and, when policy changes costs directly, the wedges $\hat{\kappa}^{sn}$ and $\hat{\kappa}_E^n$. Estimating the full matrix of switching and entry costs is therefore unnecessary for counterfactual analysis; observed flows and the elasticities estimated in Section 3.2 are sufficient.

We now turn to local results that show how these elasticities shape equilibrium adjustment near the observed equilibrium.

4.2 Wage Pass-Through and Employment Adjustment

Section 2 showed that manufacturing decline operates mainly through cohort replacement rather than incumbent switching, and Section 3.2 showed that entrants are much more wage-responsive than incumbents. We now characterize the equilibrium implications. An active entry margin raises sectoral labor-supply elasticity, attenuates wage pass-through, and implies that models without entry require larger shocks to match a given employment decline.

To formalize this, consider a partial-equilibrium experiment for sector s: vary the sector-s wage at $t+1$, holding fixed other wages and all continuation values from $t+2$ onward. As

shown in Appendix A.2 and summarized in Lemma 1, the resulting labor-supply response combines incumbent and entrant adjustment:

Lemma 1 (Labor supply response to wages)**.** *The marginal response of labor supply to future wages is given by*

$$x_t^s := \frac{\partial L_{t+1}^s}{\partial w_{t+1}^s} = \frac{\beta}{w_{t+1}^s}[\underbrace{\frac{\theta^2}{\nu}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns})}_{\textit{incumbent margin}} + \underbrace{\frac{1}{\nu^E}E_t^s(1-E_t^s)N_t}_{\textit{entrant margin}}] > 0. \tag{13}$$

The total response depends on both mobility elasticities and on the slack in each margin, captured by $\mu_t^{ns}(1-\mu_t^{ns})$ for incumbents and $E_t^s(1-E_t^s)$ for entrants.

Quantitative importance in manufacturing. To gauge magnitudes, define the normalized labor-supply slope

$$\bar{x}_t^s \equiv \frac{w_{t+1}^s}{\beta}x_t^s = \underbrace{\frac{\theta^2}{\nu}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns})}_{\bar{x}_{t,\text{inc}}^s} + \underbrace{\frac{1}{\nu^E}E_t^s(1-E_t^s)N_t}_{\bar{x}_{t,\text{ent}}^s}.$$

Table 5 evaluates this expression for manufacturing using our benchmark estimates and observed average flows in three windows. Two facts stand out. First, entrants account for a substantial share of the local manufacturing supply slope—about one-quarter before 2000 and still about one-fifth during the China-shock period—despite being a small population margin. Second, their contribution falls in both share and level terms during the 2000s because manufacturing entry declines, shrinking the slack term $E_t^M(1-E_t^M)$. The entrant margin therefore remains quantitatively important, but its buffering role weakens as young workers increasingly avoid manufacturing.

This has immediate equilibrium implications:

Proposition 2 (Entry, Wage Pass-Through, and Calibration Wedge)**.** *Consider a negative labor demand shock $\Delta\xi_{t+1}^s < 0$. Compared to a model without an active (E_t^s not near 0, 1) and more elastic entry margin ($\nu^E < \nu$),*

(a) Wage pass-through is attenuated: The equilibrium wage response $|\partial w_{t+1}^s/\partial\xi_{t+1}^s|$ is smaller.

(b) Quantity response is amplified: The equilibrium employment response $|\partial L_{t+1}^s/\partial\xi_{t+1}^s|$ is larger.

(c) Calibration implications: To match a given observed employment change ΔL_{t+1}^s, the required shock $|\Delta\xi_{t+1}^s|$ is smaller.

Table 5: Quantitative importance of the entrant margin in manufacturing

Panel A. Manufacturing employment response to a 1% wage decrease

Window	Incumbent margin (000 jobs)	Entrant margin (000 jobs)	Total (000 jobs)	Entrant share
Pre-China shock (1987–1999)	12.0	4.1	16.1	25.3%
China shock (2000–2010)	17.0	3.5	20.5	17.3%
Full sample (1987–2019)	14.9	3.8	18.8	20.4%

Panel B. Drivers of the entrant contribution

Window	E^M	$E^M(1-E^M)$	N (millions)	μ^{MM}	Stayer share of inc. slope
Pre-China shock (1987–1999)	0.066	0.061	2.4	0.844	47.6%
China shock (2000–2010)	0.041	0.040	3.4	0.772	46.2%
Full sample (1987–2019)	0.048	0.045	3.4	0.797	46.0%

Note: Panel A reports $0.01 \times \beta \times \bar{x}$, where $\bar{x}_t^M \equiv (w_{t+1}^M/\beta)x_t^M$, so entries give the local change in manufacturing employment (in thousands of CPS-weighted jobs) implied by a 1% increase in manufacturing wages. Panel B reports the primitives that govern the entrant contribution. Parameters are set to the benchmark estimates $\nu = 2.44$, $\nu^E = 0.350$, $\theta = 0.97$, and $\beta = 0.95$.

Proof. (Sketch, see Appendix for full proof) Because entrants are more responsive ($\nu^E < \nu$), a negative shock causes them to avoid the affected sector more strongly. This larger quantity adjustment (increase in x_t^s) dampens the fall in the equilibrium wage needed to clear the market. Consequently, attributing an observed employment decline solely to incumbent responses (ignoring the more elastic entry margin) would lead to overestimating the magnitude of the underlying demand shock ξ_{t+1}^s. □

Proposition 2 shows that omitting the entry margin biases equilibrium adjustment. For a given negative labor-demand shock, a model without entry overstates wage pass-through and understates quantity adjustment. The size of this bias is state-dependent. By Lemma 1, the entrant contribution to aggregate labor-supply elasticity is proportional to the slack term $E_t^s(1-E_t^s)$, so it becomes small only when entry rates are already near 0 or 1. Table 5 shows that manufacturing was not close to that boundary: even during the China-shock period, entrants still account for 17% of the local supply slope. The entry margin therefore remained quantitatively relevant, dampening wage responses and amplifying employment adjustment.[6]

This matters for quantitative work. Because models that shut down entry understate labor-supply elasticity, they require larger demand shocks to match a given path of manufacturing employment. A common approach to studying the general equilibrium effects of the China shock (e.g. Caliendo et al. (2019)) is to find the trade-cost reductions that replicate observed

[6]Formally, the entry margin increases aggregate labor-supply elasticity relative to an incumbents-only model whenever $\theta \leq \frac{\nu}{2\nu^E}\frac{E_t^s(1-E_t^s)}{q_t^s(1-q_t^s)}$, where $q_t^s := L_t^{-1}\sum_n L_t^n \mu_t^{ns}$ is the employment-weighted share of incumbents flowing into s. Using our estimates ($\nu \approx 2.44$, $\nu^E \approx 0.350$, $\theta = 0.97$) and observed entry-to-incumbent flow ratios $E_t^s/q_t^s \in (0.5, 2)$ over 1983–2019, this condition holds throughout the sample. See Appendix A.2 for the derivation.

U.S. import penetration from China: viewed through the "reduced exchange model" lens of Adao et al. (2017), any such import shock is equivalent to a negative sectoral labor demand shock ($\Delta\xi^s_{t+1} < 0$). Proposition 2 therefore implies that a model with an active entry margin can match a given path of import penetration and manufacturing employment with a smaller underlying shock. Section 6 quantifies this wedge.

4.3 Allocation of Adjustment Across Margins

Proposition 2 established that the elastic entry margin amplifies the aggregate quantity response. A natural next question is: how is the adjustment allocated between incumbents and entrants? The decomposition follows directly from Lemma 1. Each margin's share of the total employment response $\partial L^s_{t+1}/\partial\xi^s_{t+1}$ is proportional to the product of its size, its elasticity ($1/\nu$ or $1/\nu^E$), and its choice slack $p(1-p)$. For small shocks, this allocation is approximately constant. But for discrete shocks, the allocation is nonlinear in the shock size:

Proposition 3 (Nonlinear allocation of adjustment)**.** *Consider a sequence of negative demand shocks to sector s.*

(a) Small shocks *(ΔV^s_{t+1} small): Adjustment is dominated by the margin with the highest slack-weighted elasticity. Given $\nu^E < \nu$, entrants drive a disproportionate share of the employment response whenever E^s_t is not near 0 or 1.*

(b) Large cumulative shocks *(V^s_{t+1} very low): Because entrants are more elastic, the entry share E^s_t approaches zero faster than incumbent inflow shares μ^{ns}_t, and entrant slack $E^s_t(1-E^s_t)$ vanishes. The marginal adjustment to further shocks is therefore increasingly borne by the less elastic incumbent stock.*

Proposition 3 implies that the burden of adjustment shifts endogenously over the course of a protracted sectoral decline. Early on, forward-looking entrants respond strongly, absorbing most of the shock through reduced entry. But once entrants have largely abandoned the sector, the remaining adjustment falls on incumbents who are less elastic and adjust primarily through retirement rather than switching. This is precisely the pattern documented in Section 2: young workers disproportionately avoid manufacturing entry (Fact 3), while incumbent manufacturing workers face earnings losses and low mobility, eventually exiting through retirement.

The proposition also implies that the combined impact of multiple shocks (e.g., trade and automation) is non-additive: the slack available on each margin depends on the total accumulated shock, not on each shock individually. Finally, it yields a testable cross-sectional prediction: local labor markets with a historically larger entrant share in manufacturing (higher baseline E^s_t) should exhibit more attenuated wage responses and larger employment responses to subsequent negative shocks.

4.4 Permanent shocks and the timing of adjustment

The previous results characterize one-period responses. We now ask how the entry margin shapes the *speed* of adjustment to a permanent shock, the setting most relevant for episodes like the China shock, automation, and the emergence of AI.

Consider an economy in steady state at $t = -1$. At $t = 0$, an unanticipated permanent negative demand shock $\Delta\xi^s < 0$ hits sector s for all $t \geq 0$. We linearize the supply and demand systems around the pre-shock steady state. Let $a \equiv \partial w^s / \partial L^s < 0$ be the slope of labor demand and $b \equiv \partial w^s / \partial \xi^s > 0$ the impact of the shifter, so that

$$\Delta w_t^s = a\, \Delta L_t^s + b\, \Delta \xi^s.$$

On the supply side, two forces govern the evolution of employment: the wage-responsive slope x from Lemma 1 (evaluated at steady-state values) and the persistence of the incumbent stock, $\rho = \theta \mu^{ss}$, where μ^{ss} is the steady-state stayer share. The linearized law of motion is $\Delta L_{t+1}^s = \rho\, \Delta L_t^s + x\, \Delta w_{t+1}^s$. Substituting the demand equation and rearranging yields the transition equation

$$\Delta L_{t+1}^s = \underbrace{\frac{\rho}{1 - ax}}_{\equiv \kappa}\, \Delta L_t^s + \underbrace{\frac{xb}{1 - ax}}_{\equiv \gamma}\, \Delta \xi^s,$$

where $0 \leq \kappa < 1$ since $a < 0$ and $x > 0$. The persistence parameter κ governs the speed of transition: a smaller κ means faster convergence to the new steady state.

Proposition 4 (Front-loaded transition with entry)**.** *Let x denote the supply slope with an active entry margin ($E^s \in (0,1)$) and $x_{\text{noE}} < x$ the slope without it. Following a permanent shock $\Delta\xi^s < 0$, we have:*

(a) Faster Convergence (Front-Loading). The persistence parameter $\kappa = \rho/(1 - ax)$ is strictly decreasing in x, so $\kappa < \kappa_{\text{noE}}$. The fraction of the long-run adjustment realized by horizon h, given by $1 - \kappa^h$, is strictly larger for all $h \geq 1$.

(b) Muted Wage Pass-Through: The impact wage response is $\Delta w_1^s = b\, \Delta \xi^s / (1 - ax)$, and the long-run response is $\Delta w_\infty^s = b(1 - \rho)\, \Delta \xi^s / (1 - ax - \rho)$. Both are strictly decreasing in $|\cdot|$ with respect to x.

(c) Dynamic Decomposition and Front-loading: Entrants at $t = 0$ anticipate the full path of lower wages and immediately reduce entry into sector s, absorbing a large share of the employment adjustment on impact. As L_t^s falls and wages partially recover, the entry share E_t^s gradually rises from its trough toward a new, lower steady state E_∞^s.

Proof. (Sketch, see Appendix for full proof.) The results follow from the derivatives of κ and the wage-response term with respect to x: since $a < 0$ and $\rho > 0$, the persistence parameter

ρ satisfies $\frac{\partial \rho}{\partial x} = \frac{a\rho}{(1-ax)^2} < 0$. Moreover, since $a < 0$ and $b > 0$, the derivative of the wage response magnitude is $\partial |\Delta w_1^s / \Delta \xi_s| / \partial x = -ab/(1-ax)^2 < 0$. From Lemma 1, an active entry margin ($x_{\text{ent}} > 0$) strictly increases x, proving the first two parts. The fraction of adjustment is $\Delta L_h^s / \Delta L_\infty^s = (1 - \kappa^h)$, which is decreasing in κ. Result (3) follows from the decomposition of x in Lemma 1 and the evolution of $E_s(1 - E_s)$ as E_s approaches a boundary. □

The economic intuition is straightforward. The entry margin provides a large, immediate source of supply elasticity. Forward-looking entrants respond to the full anticipated wage path, producing a sharp quantity adjustment at the onset of the shock and reducing the initial wage decline needed to clear the market. Without this margin, the entire burden falls on less-elastic incumbents, requiring a larger wage drop and a more protracted employment adjustment governed by ρ alone.

Finally, Proposition 4 connects to the nonlinear logic of Proposition 3: if the shock is sufficiently large, entry collapses and slack vanishes on the extensive margin. The remaining adjustment is then gradually borne by incumbents, producing the slow tail of decline observed in the data.

Summary. The theory section delivers three results. First, baseline flows and the two estimated elasticities are sufficient statistics for counterfactual dynamics (Proposition 1). Second, the elastic entry margin amplifies employment responses, attenuates wage responses, and implies that models without entry require larger calibrated shocks to match the same employment paths (Proposition 2). Third, the entry margin front-loads adjustment and endogenously shifts the burden from entrants to incumbents as shocks accumulate (Propositions 3–4).

Proposition 1 expresses counterfactual dynamics in terms of wedges $\hat{\kappa}^{sn}$ and $\hat{\kappa}_E^n$ on switching costs and entry costs. These wedges need not arise from trade or technology; they can equally represent deliberate policy interventions. The same dispersion parameters that govern equilibrium wage pass-through therefore also govern the cost of implementing sectoral employment targets. Section 5 develops this implication.

5 Policy as an Implementation Problem

A policymaker considering sectoral employment targets faces a practical question: which instruments achieve the target at the lowest cost? We show that the answer is governed by the dispersion parameters ν and ν^E, and that the instrument ranking is robust to general-equilibrium

wage feedback.[7]

We make three points. First, conditional on wanting more employment in a sector, the entry margin is seven times cheaper per additional worker than the incumbent margin, and more so when incumbent retention subsidies are considered. Second, this cost ratio is invariant to the demand slope: general-equilibrium wage feedback raises the absolute cost of every instrument but preserves the ranking. Third, if a policymaker assigns a shadow value to employment in a target sector, the cheapest implementation loads first on entry incentives and activates incumbent instruments only as the entry margin saturates. These results concern the cost of expanding or protecting sectoral employment, not the welfare value of compensating workers displaced by sectoral decline; the two objectives operate on different margins and should be evaluated against different benchmarks.

5.1 Policy instruments

Consider a government that offers destination-specific subsidies in utility units: $a_t^n \geq 0$ to entrants who choose sector n, and $b_t^{sn} \geq 0$ to incumbents in sector s who move to n. These enter the choice problem additively, reducing effective entry costs to $C^n - a_t^n$ and switching costs to $\chi^{sn} - b_t^{sn}$. In the framework of Proposition 1, they correspond to the wedges $\hat{\kappa}_E^n = \exp(a_t^n/\nu^E)$ and $\hat{\kappa}^{sn} = \exp(b_t^{sn}/\nu)$. Substituting into equations (10)–(11) yields the subsidized allocation rules:

$$\tilde{\mu}_t^{sn}(b_t^{s\cdot}) = \frac{\mu_t^{sn}\exp(b_t^{sn}/\nu)}{\sum_k \mu_t^{sk}\exp(b_t^{sk}/\nu)}, \tag{14}$$

$$\tilde{E}_t^n(a_t) = \frac{E_t^n \exp(a_t^n/\nu^E)}{\sum_k E_t^k \exp(a_t^k/\nu^E)}, \tag{15}$$

where μ_t^{sn} and E_t^n are the baseline rules ($a = b = 0$).

The subsidies enter the logit index in utility units, which is the natural unit for the logit model. We conduct the cost analysis in these units throughout. Under $u(w) = \log w$, a utility-unit subsidy a is equivalent to raising the effective wage from w^s to $w^s e^a$, at a dollar cost of $w^s(e^a - 1)$ per recipient. At the baseline ($a = b = 0$), the marginal dollar cost per utility unit is w^s on every margin targeting sector s, so the cost ratios derived below are identical in dollars and in utility units.

[7] We study cost minimization rather than optimal policy because the fiscal cost of moving a worker through each margin depends only on the dispersion parameters ν and ν^E and observed flows, not on the completeness of insurance markets. Optimal policy in discrete-choice labor economies, by contrast, depends on the structure of missing insurance markets, which remains an active area of research; see Mongey and Waugh (2024) and Donald et al. (2025).

The law of motion under policy is

$$\tilde{L}^n_{t+1} \;=\; \theta \sum_s L^s_t \, \tilde{\mu}^{sn}_t(b^{s\cdot}_t) \;+\; N_t \, \tilde{E}^n_t(a_t). \tag{16}$$

Given a desired change in $\tilde{\mathbf{L}}_{t+1}$, the next subsection asks: what is the minimum fiscal cost of implementing it, and which margin should be used?

5.2 The minimum fiscal cost of raising employment in a sector

We study a one-period implementation problem for a single destination sector s. The dynamic case is recursive: the same problem applies at each date conditional on continuation values.[8] In the notation of Proposition 1, an entrant subsidy a^n_t for choosing sector n corresponds to the wedge $\hat{\kappa}^n_E = \exp(a^n_t/\nu^E)$, and an incumbent subsidy b^{sn}_t for moving from s to n corresponds to $\hat{\kappa}^{sn} = \exp(b^{sn}_t/\nu)$. The question is which wedge delivers a target employment change at minimum fiscal cost.

Let baseline next-period employment in sector s be

$$L^{s,0}_{t+1} \;=\; \theta \sum_n L^n_t \mu^{ns}_t \;+\; N_t E^s_t.$$

Suppose a policymaker wants to raise sector-s employment by $\Delta^s_{t+1} > 0$, so the target is $\tilde{L}^s_{t+1} = L^{s,0}_{t+1} + \Delta^s_{t+1}$. We restrict attention to destination-specific subsidies: an entrant subsidy for choosing s, and incumbent subsidies for moving into s from specific origins. These correspond to the actual policy tools used in practice such as hiring credits, training subsidies, relocation assistance, and are sufficient for the cost comparisons below.

Entry margin. Let $a^s_t \geq 0$ be a per-recipient subsidy paid to entrants who choose sector s. Inverting (15), the fiscal cost of raising the entry share from E^s_t to x is $\mathcal{F}^s_{E,t}(x) = N_t \, x \, \nu^E \log\!\left(\frac{x/(1-x)}{E^s_t/(1-E^s_t)}\right)$, which is strictly convex with local marginal cost per additional worker $\kappa^s_{E,t} = \nu^E/(1-E^s_t)$. This cost is zero at $x = E^s_t$, strictly increasing for $x > E^s_t$, and strictly convex. The fiscal cost per additional worker moved through the entry margin is

$$\kappa^s_{E,t} \;:=\; \left.\frac{d\mathcal{F}^s_{E,t}}{d(N_t x)}\right|_{x=E^s_t} \;=\; \frac{\nu^E}{1-E^s_t}. \tag{17}$$

Incumbent inflow margin. Let $b^{ns}_t \geq 0$ be a per-recipient subsidy paid to incumbents in origin sector n who move to s. Inverting (14), the fiscal cost of raising the transition share from

[8] Under commitment, the policymaker solves a sequence of convex programs linked through the law of motion (16). Each period's cost-minimization problem has the same structure as the one analyzed here, with continuation values determined by future policy.

μ_t^{ns} to y is $\mathcal{F}_{M,t}^{ns}(y) = \theta L_t^n \, y \, \nu \log\!\left(\frac{y/(1-y)}{\mu_t^{ns}/(1-\mu_t^{ns})}\right)$, which is strictly convex with local marginal cost per additional worker

$$\kappa_{M,t}^{ns} \;=\; \frac{\nu}{1-\mu_t^{ns}}. \tag{18}$$

Incumbent retention margin. A retention subsidy paid to workers currently in s who stay in s is the same instrument with $n = s$, so the local cost per additional retained worker is $\kappa_{R,t}^{s} = \nu/(1-\mu_t^{ss})$. Because the stayer share μ_t^{ss} is typically large (above 0.8 in most sectors), almost every recipient would have remained in s without the subsidy. The cost per marginal worker retained is therefore high, making retention the most expensive of the three margins.

The minimum-cost problem. Stacking these together, the minimum fiscal cost of raising sector-s employment by Δ_{t+1}^{s} is

$$\mathcal{C}_t^s(\Delta_{t+1}^s) \;=\; \min_{x,\,\{y_n\}} \left\{ \mathcal{F}_{E,t}^{s}(x) + \sum_n \mathcal{F}_{M,t}^{ns}(y_n) \right\} \tag{19}$$

subject to $N_t(x - E_t^s) + \theta \sum_n L_t^n (y_n - \mu_t^{ns}) = \Delta_{t+1}^s$, with $x \in [E_t^s, 1)$ and $y_n \in [\mu_t^{ns}, 1)$. The sum over n includes $n = s$ (retention). The objective is convex and the constraint is linear, so the solution is unique and the first-order conditions equalize the marginal cost per additional worker across all active margins.

Two results follow. First, for small targets, the policymaker uses only the margin with the lowest baseline marginal cost. Entry dominates whenever

$$\frac{\kappa_{M,t}^{ns}}{\kappa_{E,t}^{s}} \;=\; \frac{\nu}{\nu^E} \cdot \frac{1-E_t^s}{1-\mu_t^{ns}} \;>\; 1. \tag{20}$$

With $\nu/\nu^E \approx 7$ from Section 3.2 and comparable slack terms, the entry margin is roughly seven times cheaper per additional worker. Second, for larger targets, convexity of the cost functions implies that the policymaker loads first on entry and activates incumbent inflow margins only as the entry margin saturates. Retention is the last margin to be used.

The marginal costs κ_E and κ_M are expressed in utility units. In dollars, the marginal cost on any margin targeting sector s is $\kappa_m^{\$} = \kappa_m \cdot w^s$, so the ratio in (20) is the same in dollars as in utils.

The structure of the cost ratio (20) reveals the connection between the equilibrium results in Section 4 and the policy cost comparison. On any logit margin with baseline probability p, population M, and dispersion ν_m, an additive subsidy z shifts the choice probability with semi-elasticity $Mp(1-p)/\nu_m$, so the local subsidy cost per additional worker is $\nu_m/(1-p)$. This is the same ν_m that governs the equilibrium wage response: the margin that absorbs more adjustment per unit wage change is also the cheaper margin to target with subsidies. Since

$\beta\theta \approx 0.92$, the equilibrium wage-elasticity ratio $\nu/(\theta\nu^E) \approx 7$ and the cost ratio $\nu/\nu^E \approx 7$ are numerically close but conceptually distinct.

General-equilibrium policy cost. The marginal costs κ_E and κ_M^{ns} are partial-equilibrium objects that hold wages fixed. In general equilibrium, a subsidy that raises sector-s employment by ΔL depresses the sectoral wage through the demand schedule (8), which in turn reduces supply from both margins. To first order, the equilibrium employment gain from a direct addition of Δ_{dir} additional workers is

$$\Delta L \;=\; \Delta_{\text{dir}} + x_t^s \cdot \alpha^s \cdot \Delta L, \qquad \text{so} \qquad \Delta L \;=\; \frac{\Delta_{\text{dir}}}{1 - x_t^s\,\alpha^s},$$

where $\alpha^s := \partial f^s/\partial L^s < 0$ is the demand slope and $x_t^s := \partial L_{t+1}^s/\partial w_{t+1}^s > 0$ is the supply slope from Lemma 1. Since $x_t^s\alpha^s < 0$, the denominator exceeds one: the equilibrium employment gain is strictly smaller than the partial-equilibrium gain. To move one equilibrium worker into sector s through margin m, the policymaker must move $1 - x_t^s\alpha^s > 1$ workers in PE terms, so the general-equilibrium marginal cost on margin m is

$$\kappa_m^{GE} \;=\; \kappa_m^{PE} \cdot (1 - x_t^s\,\alpha^s). \tag{21}$$

The GE markup $(1 - x_t^s\alpha^s)$ depends on the total supply elasticity x_t^s and the demand slope α^s, but not on which instrument was used. It follows that the cost ratio across margins is invariant to general-equilibrium feedback:

$$\frac{\kappa_M^{GE}}{\kappa_E^{GE}} \;=\; \frac{\kappa_M^{PE}}{\kappa_E^{PE}} \;=\; \frac{\nu}{\nu^E} \cdot \frac{1 - E_t^s}{1 - \mu_t^{ns}}. \tag{22}$$

Entry remains roughly seven times cheaper per equilibrium worker. The GE markup raises the absolute cost of every instrument, and its magnitude connects to Lemma 1: the supply slope x_t^s is larger when the entry margin is active, which is the same mechanism that attenuates wage pass-through in Proposition 2. The elastic entry margin that dampens equilibrium wage responses also makes all policy instruments more expensive in absolute terms, but does so uniformly, preserving the ranking.

The demand slope α^s is determined by the production structure in Section 6. Together with the estimated x_t^s, it pins down the level of GE policy costs, not just their ratios. The linearization in (21) is valid for small subsidies; for large interventions, the GE feedback becomes nonlinear as subsidies interact with the saturation dynamics of Proposition 3. The convex PE cost functions in (19) already capture the supply-side nonlinearity; the linear demand approximation is the remaining simplification.

5.3 Application: preserving or expanding a target sector

Suppose a policymaker assigns a shadow value $\psi_t^{s^\star} \geq 0$ to one additional worker in a target sector $s^\star$, reflecting national security, agglomeration externalities, political economy, or any other motive outside the baseline model. We do not take a stand on whether such motives are justified. Given $\psi_t^{s^\star}$, the policymaker chooses the scale of intervention by solving

$$\max_{\Delta \geq 0} \left\{ \psi_t^{s^\star} \Delta \; - \; \mathcal{C}_t^{s^\star}(\Delta) \right\}, \tag{23}$$

where $\mathcal{C}_t^{s^\star}$ is the minimum-cost function from (19). At an interior optimum,

$$\mathcal{C}_t^{s^\star\prime}(\Delta_{t+1}^{s^\star}) \; = \; \psi_t^{s^\star}, \tag{24}$$

so the policymaker expands the target sector until the marginal implementation cost equals the shadow value.

The instrument ranking follows directly from Section 5.2. For small Δ, the policymaker uses only the cheapest margin, which is entry. Apprenticeship pipelines, vocational training, and hiring credits aimed at new cohorts all operate on this margin. Incumbent inflow subsidies are activated only once the entry margin becomes sufficiently expensive. Retention subsidies are usually last, since the high stayer share $\mu_t^{s^\star s^\star}$ makes them the most expensive per marginal worker.

This distinction matters for interpreting observed policy. Programs like Trade Adjustment Assistance and sector-specific retraining subsidies are designed primarily to support incumbent workers exposed to sectoral decline. That may be desirable as compensation, but it is a different objective from expanding sectoral employment. If the goal is employment in the target sector, incumbent-focused instruments purchase roughly $\nu/\nu^E \approx 7$ times fewer workers per fiscal dollar than entrant-directed instruments.

Summary. We do not ask which policy maximizes welfare; that question is inseparable from missing insurance markets that lie outside this paper. We ask a simpler question: conditional on wanting a different employment allocation, how expensive is it to get there? The answer is governed by the entrant–incumbent elasticity gap. The logit structure implies that policy must work through the payoff indices (Section 5.1), the cost per additional worker on any margin is $\nu_m/(1-p)$ (Section 5.2), and the resulting ranking is invariant to general-equilibrium wage feedback (Equation (22)).

The cost comparison extends to place-based policies that boost local labor demand through investment credits, infrastructure spending, or public-sector hiring. These instruments raise the sectoral wage w_t^s through (8), shifting the logit index for all worker types. Because entrants are more elastic (Lemma 1), such policies disproportionately attract new cohorts into the tar-

geted region or sector, consistent with the cohort-driven local recoveries documented by Autor et al. (2025b).

The framework also clarifies a distinction often blurred in policy debates. Instruments aimed at expanding sectoral employment, such as hiring credits, apprenticeships, and entry subsidies, should be evaluated by their cost per additional worker, where the entry margin dominates. Instruments aimed at compensating workers who bear the cost of sectoral decline, such as adjustment assistance, retraining, and income support, serve a different objective that our framework does not address. These are not substitutes: the cheapest way to maintain employment and the most direct way to support displaced workers are different policies aimed at different margins.

6 Quantitative Analysis

[This section is under construction.]

This section quantifies the mechanisms developed in Sections 4–5. We proceed in two parts. First, a calibrated model of the aggregate U.S. economy isolates the role of the entry margin by comparing the baseline model (with $\nu^E < \nu$ and cohort replacement) to the canonical specification without entrants ($\theta = 1$, $N_t = 0$). Second, a 50-region model embeds the framework into observed cross-state heterogeneity to study how the entry margin interacts with local labor market conditions.

6.1 Quantitative Model Setup

The labor supply block from Section 3 is embedded in a small open economy where labor is mobile across sectors $s \in \{1, \ldots, S\}$ within each region $i \in \{1, \ldots, I\}$ but immobile across regions.[9]

Labor Demand. We follow the demand-side structure of Artuç et al. (2010). As small open economies, all regions face exogenous world prices P_{it}^s, which we normalize to 1 in the initial steady state. Firms in each region i and sector s use a CES production function combining labor L_{it}^s and a fixed factor K_i^s (e.g., capital/land), normalized to $K_i^s = 1$, with productivity A_i^s:

$$y_{it}^s = A_i^s(\alpha_i^s(L_{it}^s)^{\rho^s} + (1-\alpha_i^s)(K_i^s)^{\rho^s})^{1/\rho^s}. \tag{25}$$

[9] Inter-regional migration provides an additional adjustment margin in principle, but we abstract from it. Regional migration flows are orders of magnitude smaller than sectoral reallocation flows, and empirically, outmigration from regions exposed to the China shock was modest (Autor et al., 2021, 2025a). Rodriguez-Clare et al. (2025) find that regional migration is highly inelastic and that shutting it down leaves counterfactual outcomes largely unchanged. Olney and Thompson (2024) document that the decline of internal migration is a recent phenomenon driven by housing costs.

Firms are price-takers in goods and labor markets. Profit maximization implies that labor earns its marginal product, yielding the inverse labor demand function:

$$w_t^s = P_t^s \alpha^s A^s (L_t^s)^{\rho^s - 1} (\alpha^s (L_t^s)^{\rho^s} + (1 - \alpha^s))^{\frac{1}{\rho^s} - 1} \tag{26}$$

Labor Supply. Labor supply is governed by the entrant-incumbent model in Equations (2)-(5). The framework features incumbents with mobility parameter ν and survival probability θ, and entrants with mobility parameter ν^E and mass N_t. We use our estimated values for ν and ν^E from Section 3.2 (e.g., $\nu \approx 4.75, \nu^E \approx 0.71$). In our baseline, we assume the entrant mass N_t perfectly replaces the retiring worker mass, $N_t = (1 - \theta) L_t$, keeping a constant aggregate labor force in each state.[10]

Equilibrium and Transition Dynamics. An equilibrium is a sequence of wages $\{w_{it}^s\}$, labor allocations $\{L_{it}^s\}$, and continuation values $\{V_{it}^s\}$ that satisfies firm optimization, worker optimization, and labor market clearing in each region i and sector s for all t.

6.2 Testing the Mechanism: Stylized Shock in Aggregate Economy

To cleanly test the key predictions from Section 4, we first simulate an unexpected, permanent drop in the world price of manufacturing goods in a one-region model ($I = 1$) calibrated to the U.S. aggregate economy.

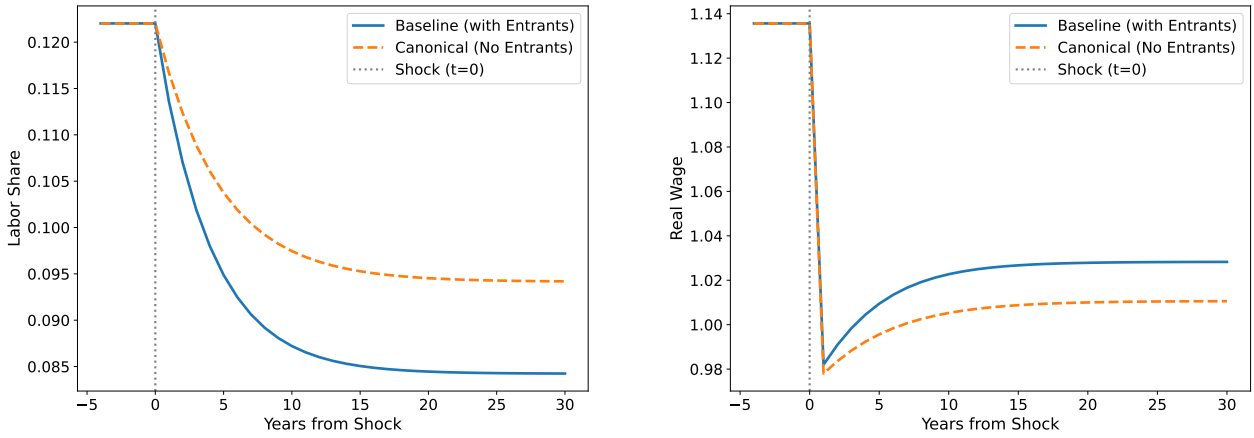
Calibration. We compare two models, both calibrating the labor demand-side parameters $(A_i^s, \alpha_i^s, \rho^s)$ to match aggregate sectoral wages, labor shares, and employment shares for the US in the year 2000. The first is the *baseline model* which uses our full labor supply block with the empirically estimated elasticities for incumbents ($\nu \approx 4.75$) and entrants ($\nu^E \approx 0.71$) and cohort replacement ($\theta < 1, N_t = (1 - \theta) L_t$). The second model is the *canonical model* which abstracts from entry by setting $\theta = 1$ and $N_t = 0$. To isolate the adjustment dynamics and ensure a comparable starting point, we simulate both models from the identical initial labor allocation L_0, which is set to the observed 2000 employment shares. This experiment allows us to isolate the precise effect of the entry margin.

One shock. We solve for the perfect-foresight transition path for both models assuming a 20% drop in the world price of manufacturing goods (P^{mfg} falls from 1 to 0.8).

First, we examine the aggregate wage and employment dynamics, which provide a quan-

[10] We are working on extensions where (1) we match the realized entrant mass, and (2) entry into regions is endogenous – following empirical evidences by Autor et al. (2025b).

Figure 4: Employment and Wage Dynamics following shock



(a) Manufacturing employment (share)

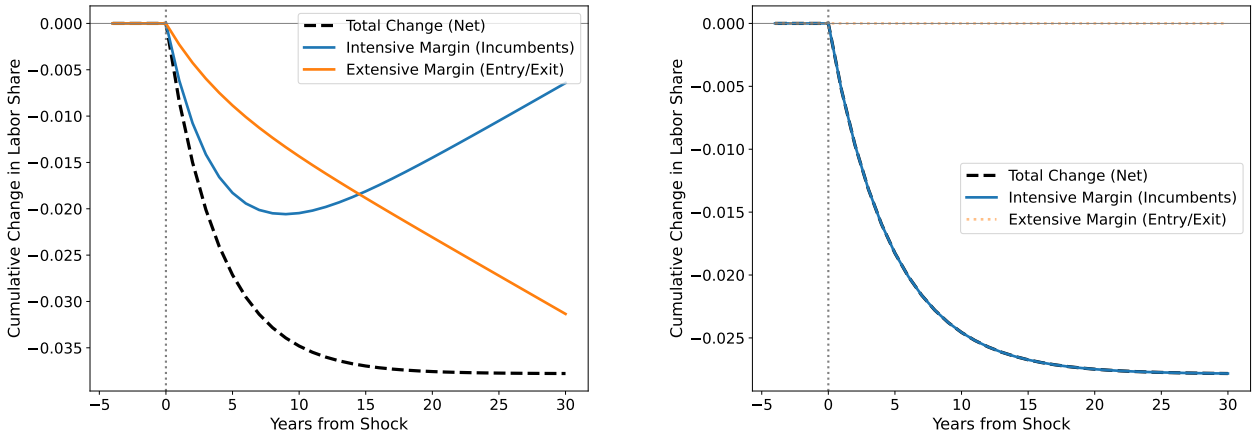
(b) Manufacturing real wage

titative test of Proposition 2. Figure 4 plots the transition paths for both models. The results confirm the predicted trade-off: incorporating entrants feature both quantity amplification and wage attenuation. The left panel shows quantity amplification: the decline in manufacturing employment is not only faster and deeper in the baseline model, but the total long-run employment drop is larger for the *same* 20% price shock. The right panel shows the wage attenuation: the manufacturing wage falls less in the baseline model than in the canonical model, both on impact and in the new steady state. The intuition is precisely as the theory suggests: because the highly elastic entry margin provides a large quantity response, the market can clear with a smaller equilibrium price (wage) adjustment. Models omitting this margin will overstate wage declines and understate the quantity adjustment.

Second, we decompose the decline to identify who adjusts, mirroring the empirical facts from Section 2. Figure 5 illustrates the cumulative change attributable to the intensive margin (incumbent switching) versus the extensive margin (net entry/exit). The results provide validation of our motivating facts. In the baseline model, the extensive margin – a collapse in new entry combined with incumbent retirements – drives the vast majority of the adjustment. The intensive margin plays some role but diminishes, despite the fact that only a small fraction $(1-\theta)$ enter, and the bulk of the population are incumbents (θ). This is consistent with the low estimated elasticity of incumbents.

Nonlinearity in Shock Magnitude. To test the *nonlinear* predictions of Corollary **??**, we now simulate both models over a grid of 40 permanent manufacturing price shocks, ranging from a 1% to 40% decline. For each shock, we decompose the employment adjustment in manufacturing into the shares attributable to the extensive (entrant) and intensive (incumbent) margins at

Figure 5: Decomposition of Manufacturing Decline (Cumulative)



(a) Baseline Model (with Entrants)

(b) Canonical Model (No Entrants)

different horizons.

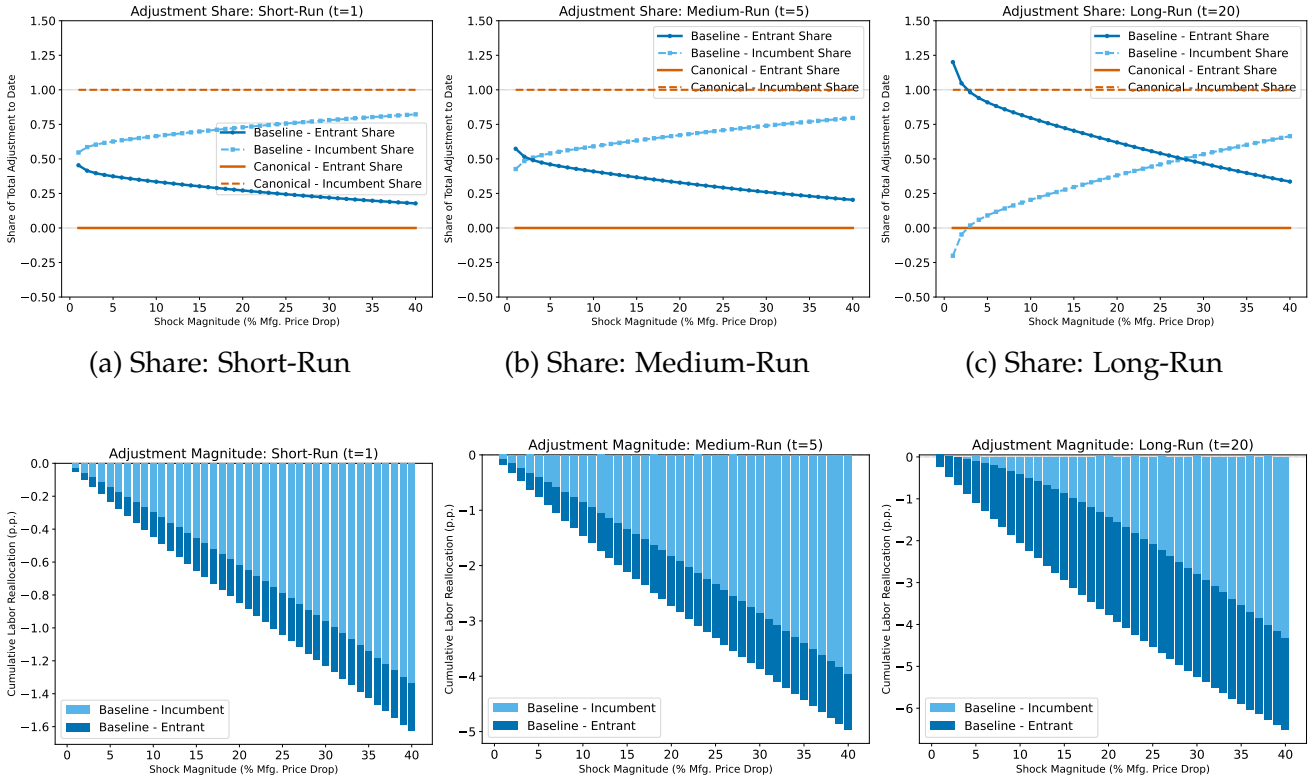
Figure 6 plots the plots the decomposition of adjustment shares (top row, panels a-c) and total magnitudes (bottom row, panels d-f). The canonical model provides a clear benchmark. In this case, with $\theta = 1$, the entrant margin is nonexistent, and 100% of the adjustment is by definition borne by incumbents. The baseline model, incorporating the entry margin, provides a test of our theory.

For small shocks, the margin with higher elasticity have a larger effect. In the short-run (panel a), the extensive (entrant) margin accounts for half of the adjustment. This result confirms the first part of Corollary 1. Specifically, when slack is high for both margins, the high-elasticity entrants drive a large share of the response – the only "limit" being the raw number of entrants (N_t).

As the shock magnitude increases (moving to the right in each panel), the "entrant slack" defined as $E_t^s(1 - E_t^s)$ diminishes. The manufacturing sector becomes so unattractive that the entrant probability E_t^s approaches zero. The plots illustrate this mechanism clearly. The entrant share of adjustment falls monotonically with the shock size. Consequently, the marginal adjustment is increasingly borne by the less elastic, but larger, stock of incumbents, whose share of the adjustment rises accordingly.

The figure also reveals a nonlinear, dynamic interplay between the margins, resulting in *overshooting* over time. An adjustment share greater than 100% (or below 0%) indicates that the two margins are working in opposite directions. This is evident in the long-run (panel c) for small shocks, where the entrant share exceeds 100% while the incumbent share becomes negative. This dynamic demonstrates that high-elasticity entrants flee the sector rapidly that they stabilize the wage, which in turn temporarily draws a small number of low-elasticity

Figure 6: Decomposition of Manufacturing Adjustment by Shock, Horizon, and Margin



(a) Share: Short-Run (b) Share: Medium-Run (c) Share: Long-Run

(d) Magnitude: Short-Run (e) Magnitude: Medium-Run (f) Magnitude: Long-Run

Note: This figure plots the decomposition of cumulative employment adjustment in manufacturing following price shocks of varying magnitudes. The top row (a-c) plots the share of adjustment attributable to entrants (extensive) vs. incumbents (intensive). The bottom row (d-f) plots the total magnitude of adjustment in percentage points, decomposed into entrant and incumbent contributions. Short-run refers to margin in 1 year; medium-run refers to 5 years; long-run refers to 20 years after the shock.

incumbents into manufacturing. The net job loss is the sum of these opposing flows. This overshooting effect suggests a reason why incumbent reallocation can contribute negligibly to the net adjustment empirically (Section 2). With the entrants responding quickly, the intensive margin may not respond, or may even move marginally in the opposite direction, offsetting the extensive margin adjustment.

This experiment illustrates that incorporating entrant-incumbent heterogeneity changes the predicted adjustment path both qualitatively and quantitatively. This framework aligns well with stylized facts and highlights potential biases in canonical models regarding wage responses and welfare incidence.

6.3 Cross-State Evidence: Heterogeneity in Adjustment

Having isolated the aggregate mechanism, we now explore how regional heterogeneity shapes the response to trade shocks. We use a multi-region version of the model calibrated to the 50

U.S. states.

Model Setup and Calibration. We embed the framework from Section 6.1 into a 50-region model ($I = 50$).[11] The labor demand-side parameters (A^s, α^s, ρ^s) and supply-side costs (χ^{sn}, C^n) are calibrated to match regional sectoral wages, labor share of revenue, sector share of GDP, and sectoral employment shares for each state in the year 2000. We again compare the Baseline (entrant) model against the Canonical (no-entrant) model, which has the same calibrated parameters. We solve for the perfect-foresight transition path in all 50 states following the same 20% drop in manufacturing prices.

Findings. Figure 7 plots the cross-state results for the Baseline model (left column), highlighting the role of regional heterogeneity. The right column presents results from the Canonical model for comparison.

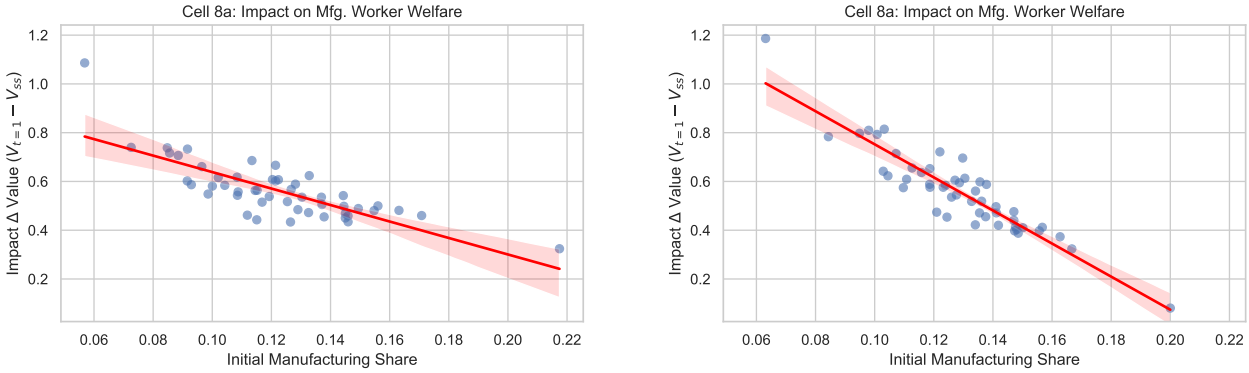
The entrant margin fundamentally alters the relationship between initial manufacturing dependence and the shock's cost. In the Baseline model, the negative correlation between a state's initial manufacturing share and the welfare loss on impact is significantly muted (Panel 7a, left). The high-elasticity entrant margin acts as a "buffer" that absorbs the initial shock, thereby protecting incumbent welfare. This welfare protection is achieved via a stronger quantity response. Panel 7b shows that high-manufacturing states experience a much larger total reallocation (a steeper negative slope), which aligns with Proposition 2.

The adjustment dynamics are also very different. Under the Baseline model, adjustment speed is relatively uniform across states (Panel 7c). This uniformity stems from the high-elasticity entrant and retirement margin. The composition of adjustment (Panel 7d) reflect a striking pattern. In states with a larger initial manufacturing share, the model predicts a positive net flow from the intensive margin. This occurs because the extensive margin (entrants and retirees) responds so strongly that it over-compensates for the shock. This allows existing workers with lower elasticities to fill the few remaining manufacturing jobs, even as the net employment adjustment remains negative.

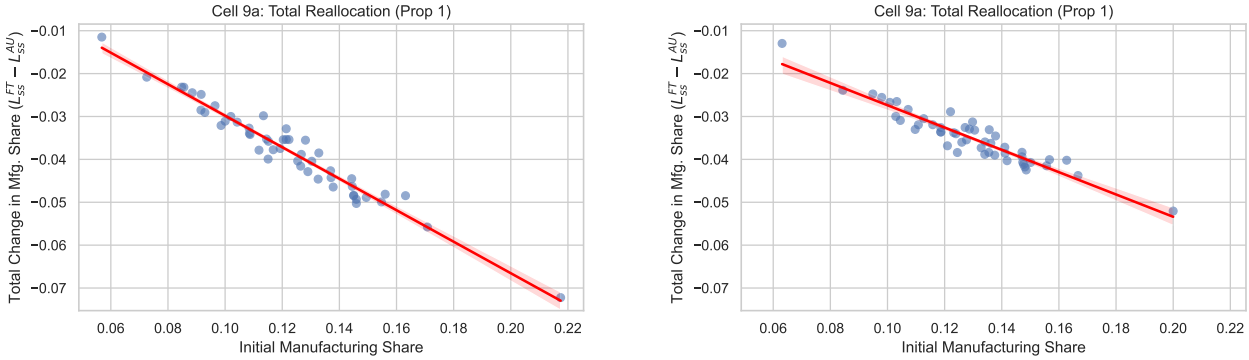
We next examine the characteristics of the service sector to which new workers adjust. Figure 8 plots the total manufacturing employment decline against the initial service sector wage in our simulation. We find a negative relationship between decline and service jobs. The largest manufacturing declines occur precisely in states where the initial service sector wage (a proxy for A^n) is lowest. This implies that the "rebound" in service jobs, which occurs mechanically as manufacturing declines, is concentrated in regions where existing service sector jobs are of lower quality. This result aligns with the empirical findings of Autor et al. (2025b), which find that regional recovery happened through increases in low-premium service jobs.

[11] We are in progress with incorporating a more full-fledged workhorse trade model with input-output linkages, imperfect substitution, and leaving the small open economy.

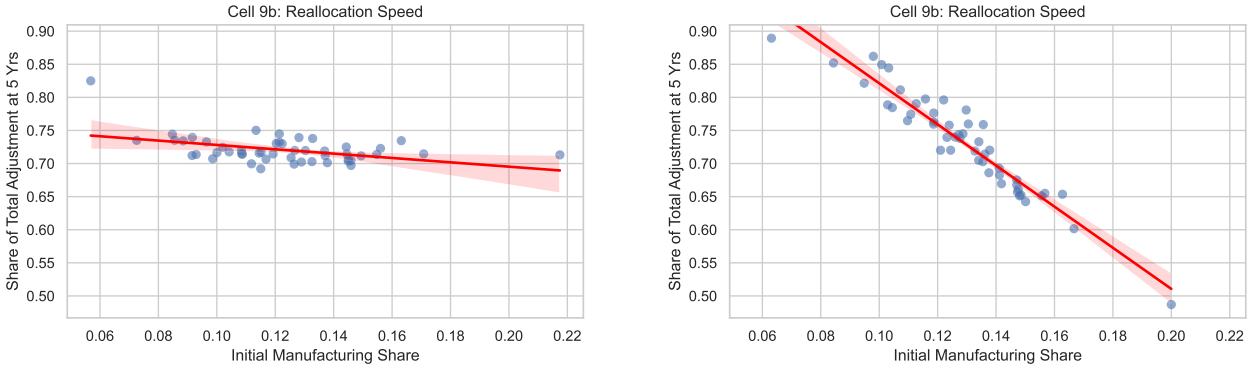
Figure 7: Cross-State Heterogeneity: Baseline vs. Canonical Model



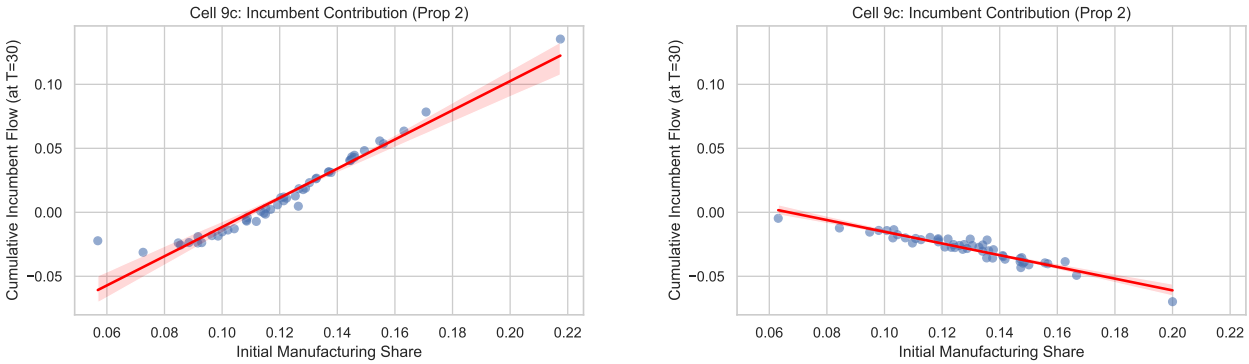
(a) Impact on Mfg. Worker Welfare



(b) Total reallocation (Prop 1)

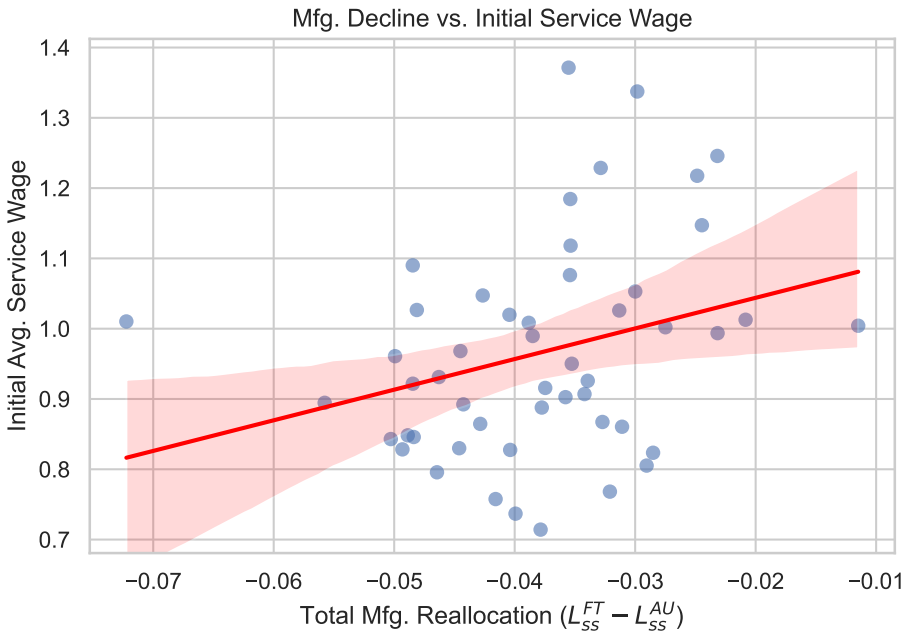


(c) Reallocation Speed (Share of Adj. at 5 Yrs)



(d) Intensive Margin Contribution in MFG decline (Prop 2)

Figure 8: Manufacturing Decline vs. Initial Service Sector Wage



Finally, we test the prediction in Section 4: local labor markets with a historically larger share of entrants in manufacturing (i.e., a higher baseline E_t^s) should exhibit more attenuated manufacturing wage responses and, simultaneously, larger employment responses to a negative shock. Figure 9 confirms this prediction. We plot the simulated wage and employment responses against the initial baseline entrant share for all 50 states. The results show that, in response to the same aggregate price shock, regions with higher initial entry rates experience smaller wage declines and larger employment adjustments.

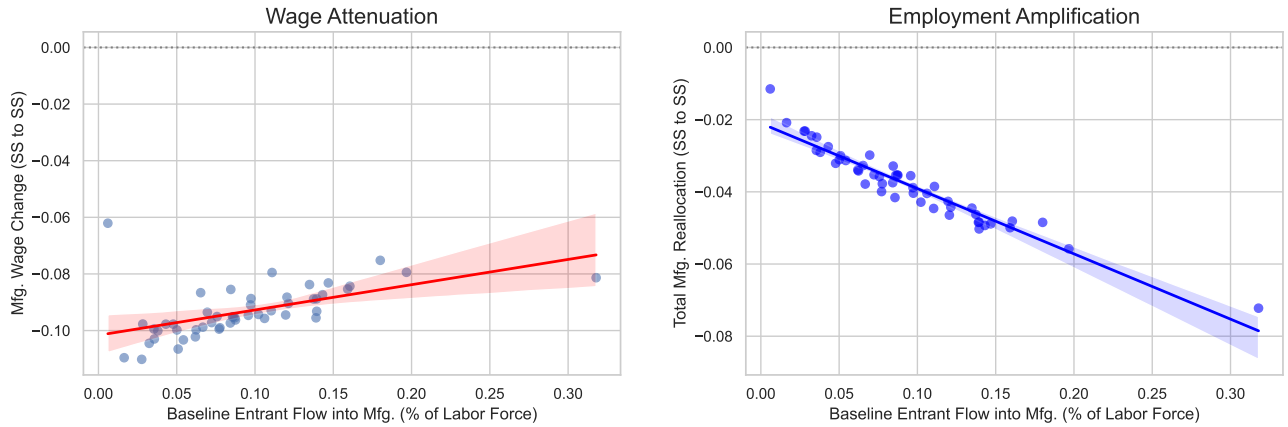
6.4 Calibration of the China shock

We now discipline the time path of the manufacturing price index so that the model replicates the rise in manufacturing import penetration associated with the China shock. Our strategy mirrors the inversion in Caliendo et al. (2019): rather than estimating high-dimensional fundamentals, we use the equilibrium conditions to find the sequence of exogenous shifters that matches observed trade outcomes.

Target and Procedure. Our target is the change in import penetration of Chinese manufacturing goods in the US (ΔIP_t) over the 2000-2012 window.[12] In our small open economy frame-

[12]Since we assume a small open economy, the model does not distinguish imports across origins. We calibrate the model so that the change in the aggregate manufacturing import penetration, $(C_t - Y_t)/C_t - (C_{2000} - Y_{2000})/C_{2000}$, matches the observed increase in import penetration specifically from China. The data target is $IM_t^{\text{mfg,china}}/(Y_t^{\text{mfg}} + IM_t^{\text{mfg}} - EX_t^{\text{mfg}}) - IM_{2000}^{\text{mfg,china}}/(Y_{2000}^{\text{mfg}} + IM_{2000}^{\text{mfg}} - EX_{2000}^{\text{mfg}})$, for $t =$ 2001-2012.

Figure 9: Test of Mechanism: Baseline Entrant Share vs. Shock Response



(a) Wage Attenuation

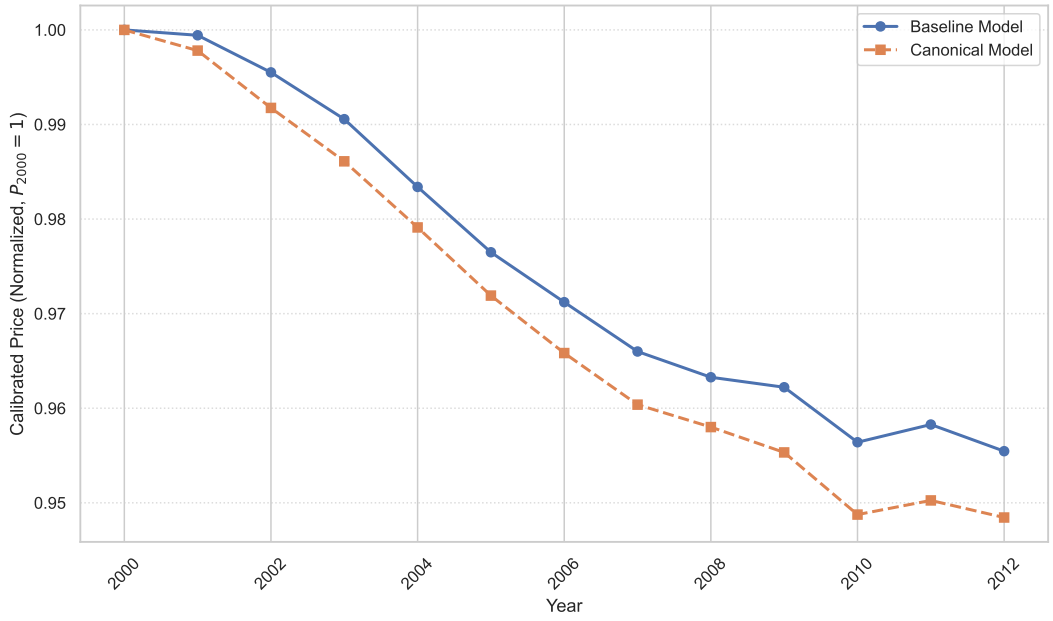
(b) Employment Amplification

work, the exogenous shifter is the sectoral goods price, P_t^s. We therefore invert a standard import-demand system to find the time series of equivalent price shifters $\{\xi_t^{Mfg}\}_{t=2000}^{2012}$ that would generate the observed ΔIP_t path. We then implement this sequence as the exogenous path for P_t^{Mfg}. We perform this inversion twice: first using our baseline model (with $\nu^E < \nu$) and second using the canonical model ($\theta = 1, N_t = 0$). This parallel calibration isolates the calibration wedge highlighted by Proposition 2: when the extensive margin is active and more elastic, the implied price changes needed to rationalize the data must be smaller.

Findings. As predicted, the model with the more elastic aggregate labor supply (the Baseline) should require a smaller price shock to rationalize the same data. Figure 10 plots the recovered price sequences. The results confirm our prediction: the price path inferred by the Baseline model is uniformly higher (less negative) than the path required by the Canonical model. By 2012, the price drop (the "China shock") implied by the Baseline model is 15.7 % smaller than the drop implied by the Canonical model. This finding is a direct consequence of the elastic entrant margin. Because the Baseline model's aggregate supply curve is more elastic, a milder price decline is sufficient to generate the same observed rise in import penetration.

Thus, explicitly incorporating worker entry and exit will lead to different welfare estimates for two primary reasons. First, conditional on the same shock ($\xi_t^s < 0$), our model predicts that wages move less and quantities adjust more, leading to smaller welfare losses. Second, conditional on matching the same data (ΔIP), the calibrated shock itself is smaller. Because both biases – the adjustment channel and the calibration wedge – move in the same direction, a model featuring the extensive margin will estimate smaller welfare costs (and larger benefits) from trade shocks compared to a canonical model that omits it.

Figure 10: Calibrated Manufacturing Price Paths (P_t^{Mfg})



7 Concluding remarks

We show the extensive margin driven by generational turnover – fewer young entrants replacing retirees – is the primary driver of the long-run decline in U.S. manufacturing, overshadowing the intensive margin of incumbent workers switching sectors. We formalize this mechanism in a quantitative general equilibrium model featuring heterogeneous mobility elasticities, where entrants are calibrated to be significantly more responsive to wage signals than incumbents.

Incorporating this entrant-incumbent distinction yields quantitative predictions that differ starkly from canonical models. First, the high-elasticity entrant margin provides a "buffer" against shocks, leading to wage attenuation (smaller wage declines) and quantity amplification (larger employment adjustments). This mechanism mutes the welfare losses for incumbents, particularly in regions with high initial manufacturing dependence. Second, the model replicates nonlinear adjustment paths: the entrant share of adjustment is largest for small shocks but diminishes as large shocks make the sector unattractive, forcing the less-elastic incumbent margin to bear a larger share of the adjustment. Third, when calibrating the model to match observed trade flows, a model with entry-exit requires smaller "China shock" price decline than a homogeneous-agent model. This suggests canonical models may systematically overestimate the magnitude of underlying shocks by omitting the elastic extensive margin.

These findings have direct policy implications. They suggest that policies targeting entrants, such as training incentives for new skills, may be more effective for accelerating struc-

tural change. Conversely, policies aiming to support incumbents might be better focused on mitigating hardship given their low mobility, rather than primarily on inducing reallocation. Our framework provides a tractable method for incorporating these generational dynamics into standard quantitative models of trade and structural change.

Several avenues for future research remain. Our analysis makes simplifying assumptions, such as abstracting from endogenous regional entry choice and other life-cycle dynamics like human capital accumulation. Furthermore, two open questions warrant a deeper investigation. First, the treatment of non-retirement movements into nonemployment requires refinement to distinguish between voluntary reallocation and involuntary exits, which have different welfare implications. Second, explicitly modeling unemployment and broader labor market frictions could provide a richer understanding of how these factors shape both incumbent reallocation and entrant decisions.

Addressing these issues is essential for a fuller picture of the incidence of trade shocks in labor market and structural transformation at large. The framework developed here, emphasizing generational entry and exit, offers a valuable lens for analyzing not only trade shocks but also other significant drivers of economic change, such as automation and shifts in sectoral demand. Understanding the distinct behavior of different cohorts within these transitions remains a critical task.

References

ACEMOGLU, D. AND P. RESTREPO (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244, publisher: The University of Chicago Press.

ADAO, R., A. COSTINOT, AND D. DONALDSON (2017): "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, 107, 633–689.

ADÃO, R., M. BERAJA, AND N. PANDALAI-NAYAR (2024): "Fast and Slow Technological Transitions," *Journal of Political Economy Macroeconomics*, 2, 183–227, publisher: The University of Chicago Press.

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 100, 1008–1045.

ARTUÇ, E. AND J. MCLAREN (2015): "Trade policy and wage inequality: A structural analysis with occupational and sectoral mobility," *Journal of International Economics*, 97, 278–294.

AUTOR, D., D. DORN, AND G. HANSON (2025a): "Trading Places: Mobility Responses of Native- and Foreign-Born Adults to the China Trade Shock," *ILR Review*, 78, 10–36, publisher: SAGE Publications Inc.

AUTOR, D., D. DORN, AND G. H. HANSON (2021): "On the Persistence of the China Shock," *NBER Working Paper*.

AUTOR, D., D. DORN, G. H. HANSON, M. R. JONES, AND B. SETZLER (2025b): "Places versus People: The Ins and Outs of Labor Market Adjustment to Globalization," .

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

BALBONI, C. (2025): "In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities," *American Economic Review*, 115, 77–116.

BLANCHARD, O. J. (1985): "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 93, 223–247, publisher: The University of Chicago Press.

BRINATTI, A. AND X. GUO (2024): "Third-Country Effects of U.S. Immigration Policy," *Working Paper*.

BUDI-ORS, T. (2025): "Rural-Urban Migration and Structural Change: A Reinterpretation," *Working Paper*.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13758.

CALIENDO, L., L. D. OPROMOLLA, F. PARRO, AND A. SFORZA (2023): "Labor Supply Shocks

and Capital Accumulation: The Short- and Long-Run Effects of the Refugee Crisis in Europe," *AEA Papers and Proceedings*, 113, 577–584.

CALIENDO, L. AND F. PARRO (2022): "Trade policy," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 5 of *Handbook of International Economics: International Trade, Volume 5*, 219–295.

CHARLES, K. K., E. HURST, AND M. SCHWARTZ (2019): "The Transformation of Manufacturing and the Decline in US Employment," *NBER Macroeconomics Annual*, 33, 307–372, publisher: The University of Chicago Press.

COSAR, A. K. (2013): "Adjusting to Trade Liberalization: Reallocation and Labor Market Policies," *Working Paper*.

DEKLE, R., J. EATON, AND S. KORTUM (2007): "Unbalanced Trade," *American Economic Review*, 97, 351–355.

DIX-CARNEIRO, R., J. P. PESSOA, R. REYES-HEROLES, AND S. TRAIBERMAN (2023): "Globalization, Trade Imbalances, and Labor Market Adjustment," *The Quarterly Journal of Economics*, 138, 1109–1171.

DONALD, E., M. FUKUI, AND Y. MIYAUCHI (2025): "Optimal Dynamic Spatial Policy," *Working Paper*.

FLOOD, S., M. KING, R. RODGERS, S. RUGGLES, J. R. WARREN, D. WARREN, A. CHEN, G. COOPER, S. RICHARDS, M. SCHOUWEILER, AND M. WESTBERRY (2024): "IPUMS, Current Population Survey: Version 12.0," .

HOBIJN, B., T. SCHOELLMAN, AND A. VINDAS (2018): "Structural Transformation by Cohort," *Working Paper*.

KAMBOUROV, G. AND I. MANOVSKII (2013): "A Cautionary Note On Using (March) Current Population Survey And Panel Study Of Income Dynamics Data To Study Worker Mobility," *Macroeconomic Dynamics*, 17, 172–194.

KIM, B., M. DE LA BARRERA, AND M. FUKUI (2024): "Currency Pegs, Trade Deficits and Unemployment: A Reevaluation of the China Shock," *Working Paper*.

KLEINMAN, B., E. LIU, AND S. J. REDDING (2023): "Dynamic Spatial General Equilibrium," *Econometrica*, 91, 385–424.

LOEWENSTEIN, G., T. O'DONOGHUE, AND M. RABIN (2003): "Projection Bias in Predicting Future Utility," *The Quarterly Journal of Economics*, 118, 1209–1248.

MONGEY, S. AND M. E. WAUGH (2024): "Discrete Choice, Complete Markets, and Equilibrium," *NBER Working Paper*.

OLNEY, W. W. AND O. THOMPSON (2024): "The Determinants of Declining Internal Migration," *NBER Working Paper*.

O'CONNOR, D. G. (2025): "Optimal Place-based Transfers for Local Recessions," *Working Pa-*

per.

PIERCE, J. R. AND P. K. SCHOTT (2016): "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 106, 1632–1662.

PORZIO, T., F. ROSSI, AND G. SANTANGELO (2022): "The Human Side of Structural Transformation," *American Economic Review*, 112, 2774–2814.

RODRIGUEZ-CLARE, A., M. ULATE, AND J. VASQUEZ (2025): "Trade with Nominal Rigidities: Understanding the Unemployment and Welfare Effects of the China Shock," *Journal of Political Economy*, publisher: The University of Chicago Press.

Appendix

A Derivations

A.1 Derivation of Equations in Section 3

We provide the step-by-step derivation for the incumbent and entrant moment conditions (Equations 6 and 7) presented in Section 3.2.

Incumbent Estimating Equation

The derivation for the incumbent's Euler equation (Equation (6)) is as follows. We begin with the incumbent's logit probability of moving from sector s to sector n, as given by Equation (2):

$$\mu_t^{sn} = \frac{\exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^n - \chi^{sn}\right)\right)}{\sum_k \exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^k - \chi^{sk}\right)\right)}$$

The probability of staying in sector s (i.e., moving from s to s) is:

$$\mu_t^{ss} = \frac{\exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^s - \chi^{ss}\right)\right)}{\sum_k \exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^k - \chi^{sk}\right)\right)}$$

We take the log-ratio of these two probabilities, and the denominator cancels. Using the normalization that the cost of staying is zero ($\chi^{ss} = 0$), we arrive at the incumbent's log-odds ratio:

$$\log(\mu_t^{sn}/\mu_t^{ss}) = \frac{1}{\nu}\left[\beta\theta(V_{t+1}^n - V_{t+1}^s) - \chi^{sn}\right] \tag{A.1}$$

Next, we derive an alternative expression for the incumbent's ex-ante value function (Equation 3):

$$V_t^s = u(w_t^s) + \nu \log\left(\sum_k \exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^k - \chi^{sk}\right)\right)\right)$$

We can rearrange the expression for the stayer probability μ_t^{ss} to solve for the log-sum term:

$$\sum_k \exp\left(\frac{1}{\nu}\left(\beta\theta V_{t+1}^k - \chi^{sk}\right)\right) = \frac{\exp\left(\frac{1}{\nu}\beta\theta V_{t+1}^s\right)}{\mu_t^{ss}}$$

Substituting this back into the value function equation, we get a recursive expression for the value function:

$$V_t^s = u(w_t^s) + \beta\theta V_{t+1}^s - \nu \log(\mu_t^{ss}) \tag{A.2}$$

This identity, derived by Artuç et al. (2010), shows that the ex-ante value function can be written in terms of the current wage, the discounted future value, and an "option value" term $\nu \log(\mu_t^{ss})$.

We first derive an identity for the difference in continuation values by substituting the log-odds ratio (Equation A.1) into the value function difference derived from Equation (A.2):

$$\begin{aligned} V_t^n - V_t^s &= [u(w_t^n) - u(w_t^s)] + \beta\theta(V_{t+1}^n - V_{t+1}^s) - \nu(\log \mu_t^{nn} - \log \mu_t^{ss}) \\ &= [u(w_t^n) - u(w_t^s)] + [\nu \log(\mu_t^{sn}/\mu_t^{ss}) + \chi^{sn}] - \nu(\log \mu_t^{nn} - \log \mu_t^{ss}) \\ &= u(w_t^n) - u(w_t^s) + \chi^{sn} + \nu \log(\mu_t^{sn}/\mu_t^{nn}) \end{aligned}$$

Advancing this identity one period ($t \to t+1$) yields an expression for the future value difference:

$$V_{t+1}^n - V_{t+1}^s = u(w_{t+1}^n) - u(w_{t+1}^s) + \chi^{sn} + \nu \log(\mu_{t+1}^{sn}/\mu_{t+1}^{nn}) \tag{A.3}$$

We substitute this back into the log-odds ratio (Equation A.1):

$$\log(\mu_t^{sn}/\mu_t^{ss}) = \frac{1}{\nu}\left[\beta\theta\left(u(w_{t+1}^n) - u(w_{t+1}^s) + \chi^{sn} + \nu \log(\mu_{t+1}^{sn}/\mu_{t+1}^{nn})\right) - \chi^{sn}\right]$$

Simplifying and grouping the χ^{sn} terms gives the estimating equation:

$$\log(\mu_t^{sn}/\mu_t^{ss}) = \frac{\beta\theta}{\nu}[u(w_{t+1}^n) - u(w_{t+1}^s)] + \beta\theta \log(\mu_{t+1}^{sn}/\mu_{t+1}^{nn}) - \frac{1-\beta\theta}{\nu}\chi^{sn}$$

Adding the expectation error Ω_{t+1}^{sn} yields Equation (6).

Entrant Estimating Equation

The share of entrants choosing sector n is:

$$E_t^n = \frac{\exp\left(\frac{1}{\nu^E}\left(\beta V_{t+1}^n - C^n\right)\right)}{\sum_k \exp\left(\frac{1}{\nu^E}\left(\beta V_{t+1}^k - C^k\right)\right)}$$

Taking logs, we can express this as:

$$\log E_t^n = \frac{1}{\nu^E}\left(\beta V_{t+1}^n - C^n\right) - \log\left(\sum_k \exp\left(\frac{1}{\nu^E}\left(\beta V_{t+1}^k - C^k\right)\right)\right)$$

The log-sum term is common to all choices n and depends only on time t. We can therefore write the equation as:

$$\log E_t^n = \frac{\beta}{\nu^E} V_{t+1}^n - \frac{1}{\nu^E} C^n + \alpha_t' \tag{A.4}$$

where $\alpha'_t = -\log\left(\sum_k \exp(...)\right)$ will be absorbed by a time fixed effect in the estimation.

Substitute the incumbent's value function $V^n_{t+1} = u(w^n_{t+1}) + \beta\theta V^n_{t+2} - \nu \log(\mu^{nn}_{t+1})$ (Equation A.2 advanced one period) into (A.4):

$$\log E^n_t = \frac{\beta}{\nu^E}\left[u(w^n_{t+1}) + \beta\theta V^n_{t+2} - \nu \log(\mu^{nn}_{t+1})\right] - \frac{1}{\nu^E}C^n + \alpha'_t$$

This gives the expression:

$$\log E^n_t = \frac{\beta}{\nu^E}u(w^n_{t+1}) - \frac{\beta\nu}{\nu^E}\log(\mu^{nn}_{t+1}) + \frac{\beta^2\theta}{\nu^E}V^n_{t+2} - \frac{1}{\nu^E}C^n + \alpha'_t$$

To replace the V^n_{t+2} term, we advance Equation (A.4) one period and solve for $\frac{\beta}{\nu^E}V^n_{t+2}$:

$$\frac{\beta}{\nu^E}V^n_{t+2} = \log E^n_{t+1} + \frac{1}{\nu^E}C^n - \alpha'_{t+1}$$

We then rewrite the V^n_{t+2} term in the $\log E^n_t$ equation as $\beta\theta\left(\frac{\beta}{\nu^E}V^n_{t+2}\right)$ and substitute this expression in:

$$\log E^n_t = \frac{\beta}{\nu^E}u(w^n_{t+1}) - \frac{\beta\nu}{\nu^E}\log(\mu^{nn}_{t+1}) + \beta\theta\left[\log E^n_{t+1} + \frac{1}{\nu^E}C^n - \alpha'_{t+1}\right] - \frac{1}{\nu^E}C^n + \alpha'_t$$

Grouping the C^n terms and redefining the time fixed effect as $\alpha_t = \alpha'_t - \beta\theta\alpha'_{t+1}$, we get the final equation:

$$\log E^n_t = \frac{\beta}{\nu^E}u(w^n_{t+1}) + \beta\theta\log(E^n_{t+1}) - \frac{\beta\nu}{\nu^E}\log\mu^{nn}_{t+1} + \alpha_t - \frac{1-\beta\theta}{\nu^E}C^n \tag{A.5}$$

Adding the expectation error $\Omega^{n,E}_{t+1}$ gives the full equation.

Justification for Parameter Restriction. A direct estimation of Equation (A.5) faces three distinct challenges, leading us to adopt the behaviorally grounded restriction ($\nu^{\text{1-period}} = \nu^E$) as in the main text.

- **Structural Heterogeneity:** The parameter ν in the $\frac{\beta\nu}{\nu^E}$ term represents the entrant's expectation of their incumbent elasticity in the next period. The estimated ν (Table 2) captures the average elasticity across all incumbents, including those with significant sector-specific tenure. It is overly restrictive to assume a first-period incumbent's elasticity mirrors this long-run average. A more plausible assumption is that early-tenure elasticity remains highly comparable to entrant elasticity ($\nu^{\text{1-period}} \approx \nu^E$). Imposing this restriction simplifies the $\frac{\nu}{\nu^E}$ ratio to 1.

- **Behavioral Plausibility:** Even if early-tenure elasticity is strictly higher than entrant elasticity, Equation (A.5) requires entrants to perfectly foresee and internalize this shift. This introduces a behaviorally questionable artifact: if $\nu \gg \nu^E$, agents would be artificially drawn to sectors with low retention rates (low μ^{nn}) to capitalize on the magnified variance of their future taste shocks. A more realistic mechanism aligns with projection bias (Loewenstein et al., 2003), where entrants project their current high-elasticity state onto their near-term future selves, perceiving their future elasticity as ν^E.

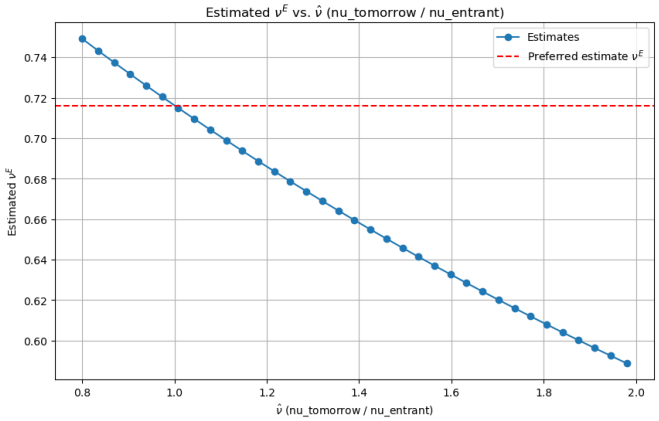
- **Empirical Identification:** Sectors with favorable unobservables typically exhibit both high wages (w_t^n) and high retention rates (μ_t^{nn}). Because these variables operate with opposite signs in the choice probability, estimating ν^E as a coefficient that responds simultaneously to both terms introduces collinearity and instability. Restricting the relative elasticity in option value stabilizes the moment conditions.

For these three reasons, we adopt the more stable and plausible specification by assuming $\frac{\nu^{\text{1-period}}}{\nu^E} = 1$ within the coefficient on the μ_{t+1}^{nn} term, yielding our final estimating equation:

$$\log E_t^n = \frac{\beta}{\nu^E} u(w_{t+1}^n) + \beta\theta \log(E_{t+1}^n) - \beta\theta \log \mu_{t+1}^{nn} + \alpha_t - \frac{1-\beta\theta}{\nu^E} C^n + \Omega_{t+1}^{n,E} \tag{A.6}$$

which give our preferred estimates of $\nu^E = 0.350$. As a robustness test, we plot estimates of ν^E for realistic levels of $\hat{\nu} = \nu^{\text{1-period}}/\nu^E$ as coefficients on μ_{t+1}^{nn}: the estimated ν^E are stable and consistent near our preferred estimate.

Figure A.1: Estimated ν^E for given $\hat{\nu}$



A.2 Proofs of Propositions in Section 4

This appendix provides the formal derivations for the lemmas and propositions presented in Section 4.

A.2.1 Proof of Lemma 1

Lemma 1. (Labor supply response to wages) *The marginal response of labor supply to future wages is given by*

$$x_t^s := \frac{\partial L_{t+1}^s}{\partial w_{t+1}^s} = \frac{\beta}{w_{t+1}^s}\left[\frac{\theta^2}{\nu}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns}) + \frac{1}{\nu^E}E_t^s(1-E_t^s)N_t\right] > 0.$$

Proof. The total labor supply in sector s at time $t+1$ is given by the law of motion (Equation 5), replacing the destination index n with s:

$$L_{t+1}^s = \sum_n \mu_t^{ns}\theta L_t^n + E_t^s N_t$$

We differentiate this with respect to V_{t+1}^s, using the standard logit derivative $\frac{\partial \mu_t^{ns}}{\partial V_{t+1}^s} = \frac{\beta\theta}{\nu}\mu_t^{ns}(1-\mu_t^{ns})$ and $\frac{\partial E_t^s}{\partial V_{t+1}^s} = \frac{\beta}{\nu^E}E_t^s(1-E_t^s)$:

$$\begin{aligned}\frac{\partial L_{t+1}^s}{\partial V_{t+1}^s} &= \sum_n \theta L_t^n\left(\frac{\beta\theta}{\nu}\mu_t^{ns}(1-\mu_t^{ns})\right) + N_t\left(\frac{\beta}{\nu^E}E_t^s(1-E_t^s)\right)\\ &= \beta\left[\frac{\theta^2}{\nu}\sum_n L_t^n\mu_t^{ns}(1-\mu_t^{ns}) + \frac{1}{\nu^E}N_t E_t^s(1-E_t^s)\right]\end{aligned}$$

From the value function identity (Equation A.2) and $u(w) = \log w$, the second term is $\frac{\partial V_{t+1}^s}{\partial w_{t+1}^s} = \frac{1}{w_{t+1}^s}$, holding future values fixed. Multiplying the terms gives the result:

$$x_t^s = \frac{\beta}{w_{t+1}^s}\left[\frac{\theta^2}{\nu}\sum_n L_t^n\mu_t^{ns}(1-\mu_t^{ns}) + \frac{1}{\nu^E}N_t E_t^s(1-E_t^s)\right]$$

□

A.2.2 When Does the Entry Margin Increase the Supply Elasticity?

This section derives the condition under which the entry margin increases the aggregate labor supply elasticity x_t^s relative to a model with only incumbents. The comparison is between two models of the same economy: the canonical framework without cohort replacement ($\theta = 1$, $N_t = 0$) and our baseline with entry and exit ($\theta < 1$, $N_t > 0$).

Setup. From Lemma 1, the supply slope is

$$x_t^s = \frac{\beta}{w_{t+1}^s}\left[\frac{\theta^2}{\nu}\sum_n L_t^n\mu_t^{ns}(1-\mu_t^{ns}) + \frac{1}{\nu^E}E_t^s(1-E_t^s)N_t\right].$$

To isolate the role of cohort replacement, suppose the total labor force is held fixed via the demographic rule $N_t = (1-\theta)L_t$, where $L_t = \sum_s L_t^s$. A lower θ means faster turnover: more incumbents retire each period, replaced by an equal mass of entrants. The incumbents-only model corresponds to $\theta = 1$.

Derivation. Substituting $N_t = (1-\theta)L_t$ into x_t^s and differentiating with respect to θ:

$$\frac{dx_t^s}{d\theta} = \frac{\beta}{w_{t+1}^s}\left[\frac{2\theta}{\nu}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns}) - \frac{1}{\nu^E}E_t^s(1-E_t^s)L_t\right].$$

The first term captures the mechanical effect of θ on the incumbent margin (more survivors means more potential switchers), and the second captures the loss of entrants as θ rises. The entry margin increases the supply elasticity—that is, x_t^s is decreasing in θ—whenever $dx_t^s/d\theta \leq 0$, which requires

$$\frac{2\theta}{\nu}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns}) \leq \frac{1}{\nu^E}E_t^s(1-E_t^s)L_t. \tag{A.7}$$

Sufficient condition via Jensen's inequality. Define the employment-weighted average inflow share

$$q_t^s := \frac{1}{L_t}\sum_n L_t^n \mu_t^{ns},$$

which is the fraction of the incumbent pool that flows into sector s next period. Since $g(p) = p(1-p)$ is concave on $[0,1]$, Jensen's inequality implies

$$\frac{1}{L_t}\sum_n L_t^n \mu_t^{ns}(1-\mu_t^{ns}) \leq q_t^s(1-q_t^s).$$

Substituting into (A.7) and rearranging yields the sufficient condition

$$\theta \leq \frac{\nu}{2\nu^E}\,\frac{E_t^s(1-E_t^s)}{q_t^s(1-q_t^s)}. \tag{A.8}$$

Interpretation. Condition (A.8) is satisfied when three forces align: (i) entrants are more responsive than incumbents ($\nu/\nu^E > 1$), (ii) the turnover rate $1-\theta$ is not too small, and (iii) the entrant choice slack $E_t^s(1-E_t^s)$ is non-negligible relative to the incumbent inflow slack $q_t^s(1-q_t^s)$.

Back-of-the-envelope verification. Using our benchmark estimates $\nu \approx 2.44$, $\nu^E \approx 0.350$, and $\theta = 0.97$, the right-hand side of (A.8) becomes

$$\frac{\nu}{2\nu^E} \frac{E_t^s(1-E_t^s)}{q_t^s(1-q_t^s)} \approx 3.27 \times \frac{E_t^s(1-E_t^s)}{q_t^s(1-q_t^s)}.$$

Over the 1985–2019 sample period, observed entry-to-inflow ratios for manufacturing satisfy $E_t^s/q_t^s \in (0.5, 2)$. In such cases, the ratio of slacks $E_t^s(1-E_t^s)/q_t^s(1-q_t^s)$ remains bounded away from zero throughout this range. In all years, the right-hand side of (A.8) comfortably exceeds $\theta = 0.97$, confirming that the entry margin is an empirically relevant source of supply elasticity throughout the sample.

A.2.3 Proof of Proposition 2

Proof. This is a comparative statics exercise in the equilibrium of period $t+1$. The equilibrium is defined by two equations, (1) Labor Supply and (2) Labor Demand, which jointly determine L_{t+1}^s and w_{t+1}^s. Let $L = L_{t+1}^s$, $w = w_{t+1}^s$, and $\xi = \xi_{t+1}^s$.

$$L = L(w, \cdot) \tag{A.9}$$

$$w = f^s(L, \xi) \tag{A.10}$$

We analyze the system by total differentiation. Let $x_t^s = \frac{\partial L}{\partial w}$ be the labor supply elasticity from Lemma 1. Let $f_L^s = \frac{\partial f^s}{\partial L} < 0$ be the slope of the inverse labor demand, and $f_\xi^s = \frac{\partial f^s}{\partial \xi} > 0$ be the effect of the demand shifter.

Differentiating the system:

$$dL = x_t^s dw$$
$$dw = f_L^s dL + f_\xi^s d\xi$$

Solving this system yields the general equilibrium responses to the demand shock ξ:

$$\frac{dw}{d\xi} = \frac{f_\xi^s}{1 - f_L^s x_t^s} \quad \text{and} \quad \frac{dL}{d\xi} = \frac{x_t^s f_\xi^s}{1 - f_L^s x_t^s}$$

An active, elastic entry margin increases the aggregate labor supply elasticity x_t^s. We analyze the comparative statics with respect to x_t^s.

Proof of 1 (Wage Pass-Through):

$$\frac{dw}{d\xi} = \frac{f_\xi^s}{1 - f_L^s x_t^s}$$

The model with an active, elastic entry margin has a strictly larger x_t^s (as x_t^s is increasing in $1/\nu^E$ and N_t). The derivative of the wage response with respect to x_t^s:

$$\frac{\partial}{\partial x_t^s}\left(\frac{dw}{d\xi}\right) = \frac{\partial}{\partial x_t^s}\left(\frac{f_\xi^s}{1 - f_L^s x_t^s}\right) = f_\xi^s \cdot (-1) \cdot (1 - f_L^s x_t^s)^{-2} \cdot (-f_L^s) = \frac{f_L^s f_\xi^s}{(1 - f_L^s x_t^s)^2} < 0$$

(since $f_L^s < 0, f_\xi^s > 0$). A larger x_t^s reduces the response $\frac{dw}{d\xi}$, thus attenuating the wage decline $|\Delta w|$ from a negative shock $\Delta\xi < 0$.

Proof of 2 (Quantity Response): Substitute the expression for $\frac{dw}{d\xi}$ back into (1′) to solve for $\frac{dL}{d\xi}$:

$$\frac{dL}{d\xi} = \frac{x_t^s f_\xi^s}{1 - f_L^s x_t^s}$$

Taking the derivative, we get

$$\frac{\partial}{\partial x_t^s}\left(\frac{dL}{d\xi}\right) = \frac{f_\xi^s(1 - f_L^s x_t^s) - (x_t^s f_\xi^s)(-f_L^s)}{(1 - f_L^s x_t^s)^2} = \frac{f_\xi^s}{(1 - f_L^s x_t^s)^2} > 0$$

(since $f_\xi^s > 0$). A larger x_t^s increases the response $\frac{dL}{d\xi}$, thus amplifying the employment fall $|\Delta L|$ from a negative shock $\Delta\xi < 0$.

Proof of 3 (Calibration): To match an observed $\Delta L^* < 0$, the required shock is $\Delta\xi \approx \Delta L^*/(dL/d\xi)$. From (2), a larger x_t^s increases the denominator $\frac{dL}{d\xi}$. Therefore, a smaller shock magnitude $|\Delta\xi|$ is required to generate the same ΔL^*. □

A.2.4 Proof of Proposition 3

Proof. From the law of motion (5), the total employment response in sector s to a change in continuation values is

$$\frac{\partial L_{t+1}^s}{\partial V_{t+1}^s} = \sum_n S_{I,n} \frac{\partial \mu_t^{ns}}{\partial V_{t+1}^s} + S_E \frac{\partial E_t^s}{\partial V_{t+1}^s},$$

where $S_{I,n} := \theta L_t^n$ is the surviving incumbent mass from sector n and $S_E := N_t$ is the entrant mass. From the logit choice probabilities, the derivatives are

$$\frac{\partial \mu_t^{ns}}{\partial V_{t+1}^s} = \frac{\beta\theta}{\nu}\mu_t^{ns}(1 - \mu_t^{ns}),$$
$$\frac{\partial E_t^s}{\partial V_{t+1}^s} = \frac{\beta}{\nu^E} E_t^s(1 - E_t^s).$$

Define the elasticity-slack product for each margin: $A_{I,n} := S_{I,n} \cdot (\beta\theta/\nu) \cdot \mu_t^{ns}(1-\mu_t^{ns})$ and $A_E := S_E \cdot (\beta/\nu^E) \cdot E_t^s(1-E_t^s)$. Since any demand shock $\Delta\xi_{t+1}^s$ affects employment only through the equilibrium change in V_{t+1}^s, the share of adjustment attributable to margin m is

$$\text{Share}_m = \frac{A_m}{\sum_n A_{I,n} + A_E}, \tag{A.11}$$

which is proportional to size $\times$ elasticity $\times$ slack.

Part (1): Small shocks. For small ΔV_{t+1}^s, the shares in (A.11) are evaluated at baseline values. Since $\beta/\nu^E \gg \beta\theta/\nu$ (from $\nu^E \ll \nu$), the entrant elasticity-slack product A_E is large relative to $A_{I,n}$ whenever E_t^s is bounded away from 0 and 1 (so that $E_t^s(1-E_t^s)$ is non-negligible) and the entrant mass S_E is positive. In this case, entrants account for a disproportionate share of the total employment adjustment.

Part (2): Large cumulative shocks. As $V_{t+1}^s \to -\infty$, both μ_t^{ns} and E_t^s converge to zero, so the slack terms satisfy $p_m^s(1-p_m^s) \approx p_m^s$ for each margin m. From the logit form, the choice probabilities decay exponentially:

$$\mu_t^{ns} \propto \exp\left(\frac{\beta\theta}{\nu} V_{t+1}^s\right), \qquad E_t^s \propto \exp\left(\frac{\beta}{\nu^E} V_{t+1}^s\right),$$

absorbing cost terms and other continuation values into the proportionality constants. The ratio of entrant to incumbent marginal adjustment is therefore

$$\frac{A_E}{A_{I,n}} \propto \frac{\exp(\beta V_{t+1}^s/\nu^E)}{\exp(\beta\theta V_{t+1}^s/\nu)} = \exp\left(\beta V_{t+1}^s\left[\frac{1}{\nu^E} - \frac{\theta}{\nu}\right]\right).$$

The bracketed term $B := 1/\nu^E - \theta/\nu$ is strictly positive since $\nu^E < \nu$ and $\theta \le 1$. As $V_{t+1}^s \to -\infty$, the exponent $\beta V_{t+1}^s \cdot B \to -\infty$, so $A_E/A_{I,n} \to 0$. The entrant margin's share of marginal adjustment vanishes, and further adjustment is borne increasingly by the less elastic incumbent stock. □

B Details for Empirics

B.1 Data Construction for the Monthly CPS

CPS microdata and linked panel. We use IPUMS CPS Basic Monthly microdata (Flood et al., 2024) from January 1982 through December 2019 and construct one-year links by matching respondents interviewed in the same calendar month one year apart. We restrict attention to civilian, non-institutionalized men ages 20–64 and compute all moments using CPS survey

weights. The CPS rotation design (months-in-sample 1–4 and 13–16) implies that many individuals can be observed at multiple 12-month intervals, so annual transition matrices are constructed at the person-year level using these within-person links.

Matching procedure and validation rules. We link individuals 12 months apart using the IPUMS CPS person identifiers and month-in-sample structure (same calendar month one year apart). We retain links only when sex, race, quarter of birth, and state match and age advances by one year (allowing standard minor reporting error). Because the CPS rotation pattern can generate multiple valid 12-month links for the same individual, we use a single canonical link, month-in-sample 4 to 16 (MIS 4→16), to construct annual transitions. We drop ambiguous one-to-many matches and any links failing these checks. Weekly earnings are taken from the outgoing rotation groups and are linked to subsequent one-year transitions using the month-in-sample structure.

State definitions and treatment of unemployment. In each month we assign individuals to one of four mutually exclusive states: manufacturing (M), agriculture (A), services (S), and non-employment (N). Sector states are defined using IPUMS's harmonized 1990 industry classification (`IND1990`), which provides a consistent mapping over time; given our coarse sector partition, remaining classification differences across CPS industry recodes are immaterial.

A key classification choice is the treatment of unemployment. We treat unemployed individuals who report a prior industry as *temporarily separated* and keep them attached to their reported sector, rather than coding them as having exited to N. This convention ensures that short unemployment spells during job search do not mechanically inflate extensive-margin flows; instead, sector switching is recorded when an individual moves between employed sectors across years. In the CPS, reported industry is near-universal ($> 95\%$) among the temporary unemployed but rare ($< 5\%$) among those out of the labor force, consistent with unemployment reflecting temporary separation and N capturing non-employment (Section 2).

Flow reconciliation (biproportional fitting / IPF). Linked CPS transitions produce a raw flow matrix that need not be exactly consistent with observed sectoral stocks in adjacent years: row sums may differ from L_t^i because of linkage gaps, and column sums may differ from L_{t+1}^j because the CPS is a repeated cross-section with changing population composition. To obtain an exact accounting decomposition, we reconcile the raw flows using biproportional fitting (iterative proportional fitting, IPF) on an augmented state space that includes an outside state R capturing entry to and exit from the CPS 20–64 universe.

Concretely, for each year t we begin with raw origin–destination counts $\tilde{F}_t^{ij}$ over $i, j \in \{M, A, S, N, R\}$. We then choose adjusted flows F_t^{ij} that (i) preserve the support of the raw

matrix (cells with $\tilde{F}_t^{ij} = 0$ remain zero), and (ii) match observed stocks in non-outside states:

$$\sum_j F_t^{ij} = L_t^i \quad \text{for } i \in \{M, A, S, N\}, \qquad \sum_i F_t^{ij} = L_{t+1}^j \quad \text{for } j \in \{M, A, S, N\}.$$

All margins involving R are left unconstrained and fit whatever is required for the adjusted matrix to satisfy these stock-flow identities. We implement IPF by alternating row and column rescaling of $\tilde{F}_t$ until the constraints are satisfied to numerical tolerance. The same procedure is applied within origin-date age bins when constructing age-specific decompositions.

B.2 Accounting identities

This appendix collects the accounting identities used to decompose changes in manufacturing employment into intensive and extensive margins, and to separate within-age flow changes from between-age compositional changes. All identities hold exactly for the reconciled flow matrix F_t^{ij} described in Appendix B.1.

Levels identity. Let L_t^i denote the number of individuals in state $i \in \{M, A, S, N\}$ at time t, and let F_t^{ij} denote the number transitioning from i to j between t and $t+1$. Augmenting the state space with an outside state R, the change in sectoral employment levels satisfies

$$L_{t+1}^i - L_t^i = \underbrace{\sum_{j \in \{M,A,S\}} \left(F_t^{ji} - F_t^{ij}\right)}_{\text{intensive}} + \underbrace{\left(F_t^{Ni} - F_t^{iN}\right)}_{\text{extensive}} + \underbrace{\left(F_t^{Ri} - F_t^{iR}\right)}_{\text{outside / residual}}. \tag{B.1}$$

Shares identity and the composition effect. Define total employment $L_t \equiv L_t^M + L_t^A + L_t^S$ and sectoral employment shares $s_t^i \equiv L_t^i / L_t$ for $i \in \{M, A, S\}$. Using (B.1) and allowing the employment denominator to change over time, the exact change in employment shares is

$$s_{t+1}^i - s_t^i = \underbrace{\frac{1}{L_{t+1}} \sum_{j \in \{M,A,S\}} \left(F_t^{ji} - F_t^{ij}\right)}_{\text{intensive}} + \underbrace{\frac{1}{L_{t+1}} \left(F_t^{Ni} - F_t^{iN}\right)}_{\text{extensive}} + \underbrace{\frac{1}{L_{t+1}} \left(F_t^{Ri} - F_t^{iR}\right)}_{\text{outside / residual}} + \underbrace{s_t^i \left(\frac{L_t}{L_{t+1}} - 1\right)}_{\text{composition effect}}. \tag{B.2}$$

The final term captures changes in shares arising from movements in the aggregate employment denominator L_t even when sectoral headcounts are unchanged.

Age-specific decomposition and aggregation. For age groups $g \in \{\text{young, prime-age, old}\}$ defined at the origin date t, let $L_t^i(g)$ and $F_t^{ij}(g)$ denote the corresponding age-specific stocks and flows. Identities analogous to (B.1) and (B.2) hold within each g. We aggregate age-specific

shares into the overall share using the exact decomposition

$$s^i_{t+1} - s^i_t = \underbrace{\sum_g \bar{w}_t(g)\left[s^i_{t+1}(g) - s^i_t(g)\right]}_{\text{within-age (flows)}} + \underbrace{\sum_g \bar{s}^i_t(g)\left[w_{t+1}(g) - w_t(g)\right]}_{\text{between-age (composition)}}, \tag{B.3}$$

where $w_t(g) \equiv L_t(g)/L_t$ is the employment weight of age group g, $\bar{w}_t(g) \equiv \frac{1}{2}\left(w_t(g) + w_{t+1}(g)\right)$, and $\bar{s}^i_t(g) \equiv \frac{1}{2}\left(s^i_t(g) + s^i_{t+1}(g)\right)$.

Entry and exit labels. When summarizing the extensive margin by life-cycle stage, we define "exit" as the flow from manufacturing to non-employment among older workers, $F^{MN}_t(\text{old})$, and "entry" as the flow from non-employment to manufacturing among young workers, $F^{NM}_t(\text{young})$. All age groups are defined at the origin date t.

B.3 Robustness for Section 2

BLS "Fact 1." Figure B.1 reproduces Figure 1 using BLS employment totals for manufacturing headcount (and the corresponding aggregate employment denominator) in place of the CPS-linked headcount series. The same divergence is present and, if anything, more pronounced: manufacturing's employment share declines smoothly over the full sample, while manufacturing headcount is relatively stable through the late 1990s and then falls sharply after 2000. This confirms that the share–headcount break motivating Fact 1 is not an artifact of CPS linking.

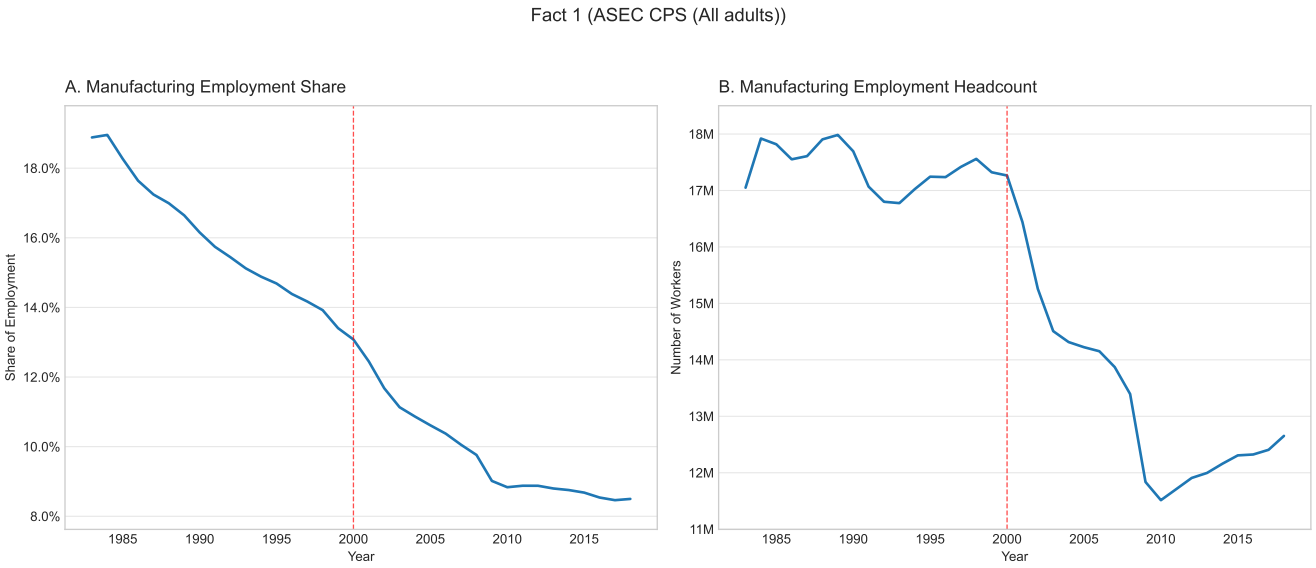


Figure B.1: Divergence of Manufacturing Employment Share and Headcount (BLS).

Using the LEHD dataset.

We conduct robustness exercises for the motivating findings using the Job-to-Job Flows (J2J) data from the Longitudinal Employer-Household Dynamics (LEHD) program.

Data Source and Sample Construction. The J2J dataset from the LEHD program, produced by the U.S. Census Bureau, tracks worker reallocation by linking state Unemployment Insurance (UI) wage records with firm-level data from the Quarterly Census of Employment and Wages (QCEW) and historical demographic records. Achieving near-national coverage by the year 2000, the dataset provides a quarterly panel that covers approximately 95% to 97% of all U.S. wage and salary jobs.

Defining Sectoral Entrants. To measure inflows into a sector sector from nonemployment, we utilize the LEHD J2J NEPersist variable (hires from persistent nonemployment). This variable identifies workers newly hired into a firm who had no reported UI earnings in the preceding quarter. By requiring a full quarter without earnings, NEPersist filters out short-term, frictional spells between jobs and isolates true structural entrants—capturing both school-to-work transitions and workers re-entering the labor force after extended displacement.

Sample Demographics and Sector Classification. Our primary analysis of "young workers" in this section maps to the LEHD age bin of 19–24 years old. To ensure our results are robust and strictly align with the broader CPS age definitions used in the main text, we also conduct a separate exercise incorporating the 25–34 age bin. Across all LEHD exercises, the manufacturing sector is consistently identified using the 2-digit NAICS classification 31–33.

Limitations on Sector-to-Sector Flows. While the J2J data precisely tracks total employment movements in and out of the manufacturing sector (separating hires from nonemployment versus hires from other jobs), a limitation of the data structure used in this exercise is the lack of detailed sector-to-sector transition matrices. Consequently, while we observe the aggregate intensive margin (workers switching into manufacturing from other employed positions), we do not trace the specific originating industries (e.g., Services or Construction) for these incumbent job-to-job switchers.

Results. [Available upon request.]

Using the March CPS.

Data and motivation. To validate that our findings are not an artifact of monthly CPS linking, we replicate the manufacturing decline decomposition using the CPS Annual Social and

Economic Supplement (ASEC, "March CPS"). The ASEC records respondents' industry at the March interview and, separately, the industry in which they spent the most time during the previous calendar year. This provides a retrospective one-year transition measure without relying on month-in-sample matching.

Sample and state construction. We follow the baseline sample restrictions as closely as possible (men ages 20–64; survey weights). We map both current and prior-year industries into the same four-state partition $\{M, A, S, N\}$ using harmonized industry codes. We treat nonemployment analogously by coding individuals as N when they report not working, and we construct annual transition matrices from prior-year state to current state.

Results. The ASEC-based decomposition yields the same qualitative conclusion as the linked monthly CPS: the manufacturing decline is dominated by the extensive margin, and age patterns indicate that exits are concentrated among older workers while entry by young workers declines. Appendix Figure A.X and Table A.X report the ASEC analogues of Figures 2–3.

[Results available upon request.]

C Life-cycle dynamics

We are extending the baseline entrant–incumbent framework to a finite-horizon life-cycle model in which mobility costs and retirement risk vary with age (or tenure). This extension nests the baseline model as a parsimonious approximation and clarifies how long-horizon entry decisions and short-horizon late-career decisions jointly shape sectoral adjustment.

[Details are work in progress and available upon request.]