Should I Stay or Should I Go? The Response of Labor Migration to Economic Shocks

Foschi, House, Proebsting, and Tesar

Discussion by Brian K. Kovak
Carnegie Mellon University, Heinz College

Brookings Papers on Economic Activity, Spring 2025 Conference

Washington, DC

March 27, 2025

Contributions

- **Research questions**: What portion of local employment responses to labor demand shocks is accounted for by inter-regional migration, and how has that portion changed over time?
- **Answers**: Migration accounts for majority of employment adjustment w/in a few years of shock's arrival; reasonably constant effects from 1946 to 2016

Contributions

- 60 years of comparable evidence on population and employment responses to consistently-measured labor demand shocks
  - Complements large literature on secular changes in unconditional migration rates
  - Fills in and expands historical record on local responses to shocks (Blanchard Katz 1992, Dao Furceri Loungani 2017, Yagan 2019, and others)
- Four different shock measures, finding similar results for recent years
- Argues that migration is as important an adjustment mechanism as ever

Outline

- Parameter interpretation
- Addressing pre-trends
- Smaller effects in smaller regions

Parameter interpretation

Novel focus: population effect's share of employment effect

- Benefit: consistent units across various IVs
- Challenge: how should we interpret this empirical parameter? To what extent does it address concerns about "declining dynamism"?

- Interpret the parameter in a simplified version of Amior and Manning's (2018 AER) paper "The Persistence of Local Joblessness"
  - Modified Rosen-Roback model with frictional local population adjustment and endogenous labor supply conditional on location.
  - Simplified here to have constant wage elasticity of local population (rather than dynamic population adjustment leading to an error-correction model)

Parameter interpretation

Amior and Manning (2018) model

- Traded good with price P
- Housing price P_r^h, varies by region r
- Homothetic preferences: unique local price index $P_r = Q(P, P_r^h)$
- Local population: L_r, local employment N_r
- Workers earn W_r (taxed at rate τ) when employed, and B otherwise
- Let lower-case letters indicate logs

- Log-linearized housing demand
$$h_r^d = w_r - p + \kappa(n_r - \ell_r) + \epsilon^{hd}(p_r^h - p)$$
- Housing supply
$$h_r^s = \epsilon_r^{hs}(p_r^h - p)$$
- Labor demand
$$n_r^d = \epsilon^{nd}(w_r - p) + z_r^d$$
- Labor supply or "wage curve"
$$n_r^s = \ell_r + \epsilon^{ns}(w_r - p_r) + z_r^s$$
- Population supply curve [simplification]
$$\ell_r = \eta(w_r - p_r)$$

Parameter interpretation

Consider the effect of a shock to local labor demand, dz_r^d

Imposing market clearing equilibrium in housing and labor markets, and given the isoelastic population supply curve,

$$\frac{d\ell_r}{dz_r^d} = \frac{\eta}{\eta + \epsilon^{ns}} \frac{dn_r}{dz_r^d}$$

$$\gamma^{Pop} \equiv \frac{d\ell_r / dz_r^d}{dn_r / dz_r^d} = \frac{\eta}{\eta + \epsilon^{ns}}$$

Parameter interpretation

Consider variation over time in the migration elasticity

$$\gamma_t^{Pop} = \frac{\eta_t}{\eta_t + \epsilon_t^{ns}}$$

Implications

- Changes in γ_t^{Pop} reflect changes in both η_t and ϵ_t^{ns}
  - Extreme example: If both fall (declining dynamism, population aging, improved entertainment goods), can find constant γ_t^{Pop}
- Emphasizes the value of showing population and employment effects separately
- Rhetorically useful to separate out ϵ_t^{ns} from β_t^E $(= \eta_t + \epsilon_t^{ns})$ in Figure 5

Outline

- Parameter interpretation
- Addressing pre-trends
- Smaller effects in smaller regions

Addressing pre-trends

- Population pre-trends are pervasive in this literature, even in cases where shocks are unrelated to pre-trends in other economic outcomes. Examples:
  - China shock: Dauth Findeisen Suedekum (2014), Greenland Lopresti McHenry (2019), Autor Dorn Hanson (2021)
  - Housing net worth: Cadena Kovak (2016), Monras (2020)
  - Tariff Changes: Dix-Carneiro Kovak (2016)

- Population results are often sensitive to the way in which pre-trends are addressed, e.g. Autor Dorn Hanson (2021 *BPEA*)

Addressing pre-trends

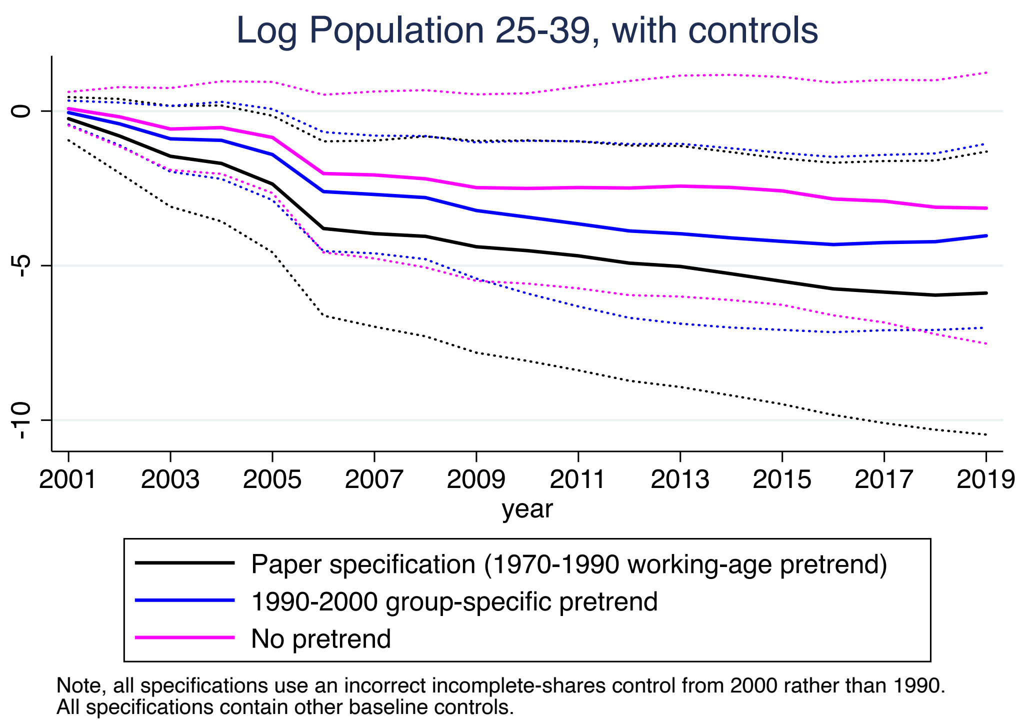

- Replicates ADH (2021) Figure 7, Panel C (the group with statistically significant effects)

- Original includes change in log working-age pop from 1970-1990
- Using group-specific pre-trend from decade preceding outcomes (1990-2000) cuts LR magnitude by 32%.
- Omitting pre-trend cuts LR magnitude by 47%

Discussant's analysis using replication files from Autor Dorn Hanson (2021)

Addressing pre-trends

How do Foschi et al. (2025) address this issue?

They estimate

$$\hat{Y}_{i,t+h} = \alpha_h^Y + \beta_h^Y \hat{Z}_{i,t} + \varepsilon_{i,t+h}^Y$$

where $\hat{x}_{i,t} \equiv x_{i,t} - \bar{x}_i - (\bar{x}_t - \bar{x})$ and $Y_{i,t+h} = \ln Pop_{i,t+h} - \ln Pop_{i,t-1}$

This double-demeaning is akin to including state and year fixed effects

$$\Delta_{t+h} \ln \mathrm{Pop_i} = \alpha_h + \beta_h Z_{i,t} + \eta_i + \varphi_t + \varepsilon_{i,t+h}$$

Addressing pre-trends

$$\Delta_{t+h} \ln \mathrm{Pop_i} = \alpha_h + \beta_h Z_{i,t} + \eta_i + \varphi_t + \varepsilon_{i,t+h}$$

To interpret this specification, difference across adjacent time periods

$$\Delta_{t+h} \ln Pop_i - \Delta_{t-1+h} \ln Pop_i = (\varphi_t - \varphi_{t-1}) + \beta_h(Z_{i,t} - Z_{i,t-1}) + (\varepsilon_{i,t+h} - \varepsilon_{i,t-1+h})$$

Counterfactual assumption: local Pop growth persists (plus common term across locations) when Bartik shock remains constant (e.g. constant natl. industry growth)

Equivalent to linear control for Pop pre-trend with coefficient α_h fixed to 1

$$\Delta_{t+h} \ln Pop_i = (\varphi_t - \varphi_{t-1}) + \beta_h\left(Z_{i,t} - Z_{i,t-1}\right) + \alpha_h \Delta_{t-1+h} \ln Pop_i + (\varepsilon_{i,t+h} - \varepsilon_{i,t-1+h})$$

Addressing pre-trends

This approach to pre-trends appears to be pivotal in driving the paper's findings

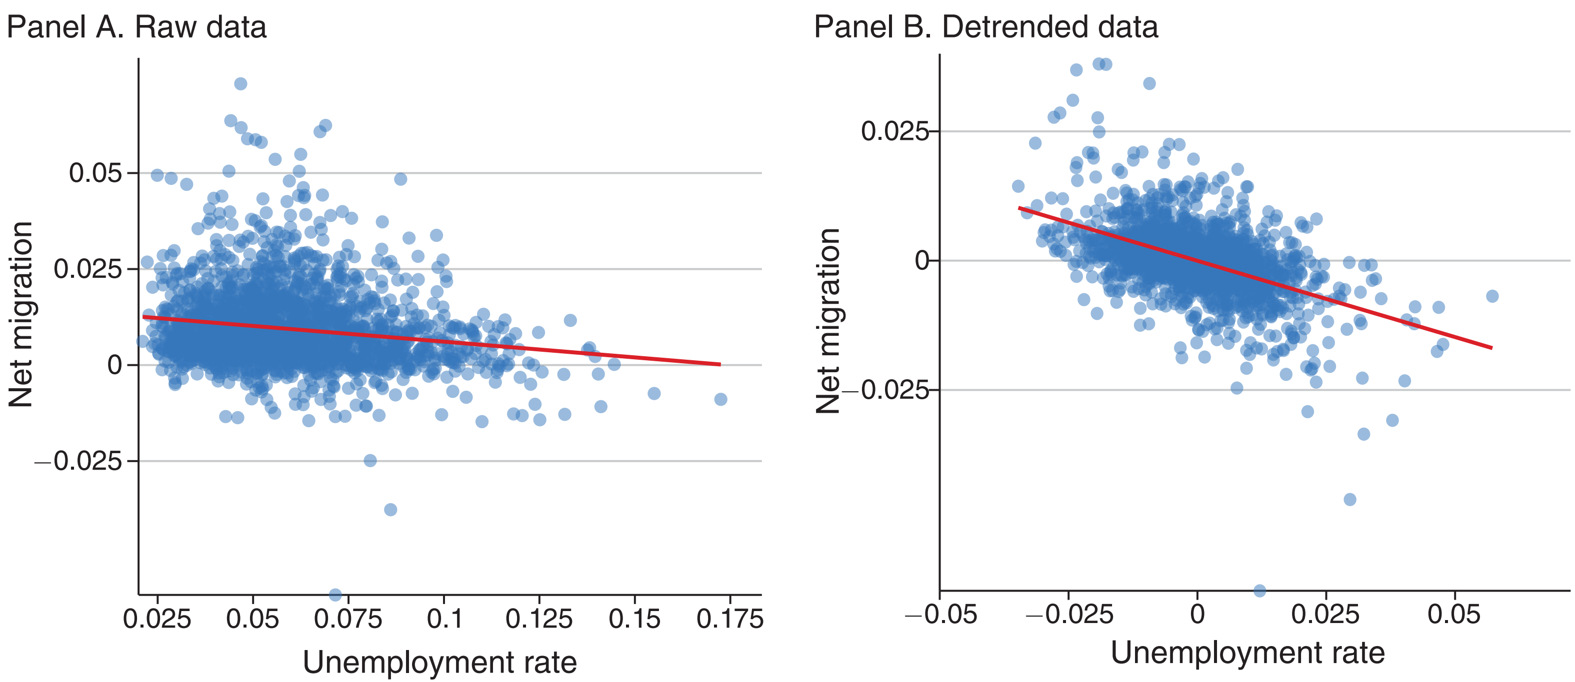

Foschi, House, Proebsting, Tesar (2023 AEA P&P), Figure 1

Addressing pre-trends

Suggestions

For this paper

- Examine results without demeaning (standard diff-in-diff)
- Try specification on slide 12 with Bartik level on RHS (as opposed to change)
- Estimate more general pre-trend model, allow α_h to differ from 1 (Ding and Li 2019 *Pol. Analysis*)

For the local shocks literature

- Examine sensitivity to pre-trend violations (Rambachan and Roth 2023 *REStud*)

For the econometrics literature

- Benefits and pitfalls of long timeseries of shocks and outcomes (e.g. Guren McKay Nakamura Steinsson 2021 *REStud*)

Outline

- Parameter interpretation
- Addressing pre-trends
- Smaller effects in smaller regions

Smaller effects in smaller regions

Paper p.20 "…the fact that the point estimates are even lower [with smaller region definitions] remains puzzling"

- Similar result in their related AEA P&P (2023)

- Explanations proposed in the text
  - Measurement error in measuring local unemployment (applies to P&P)
  - Short-distance moves are mostly driven by non-market considerations (family proximity, housing needs, lifecycle considerations)

- Two additional potential explanations
  - Mismatch between geography of labor market equilibrium and shock measurement leads to mean-zero measurement error (maybe classical)
  - Correlated outside options bias may be more acute with finer geography

Smaller effects in smaller regions

Mismatch between geography of labor market equilibrium and shock

- Assume state labor markets are relevant economic equilibrium units. Correct shock is then state-level Bartik, but the county-level regression (incorrectly) uses county-level Bartik

$$z_i^s \equiv \sum_i \frac{Emp_{si}^{t_0}}{Emp_s^{t_0}} d \ln Emp_i \qquad z_i^c \equiv \sum_i \frac{Emp_{ci}^{t_0}}{Emp_c^{t_0}} d \ln Emp_i$$

- Results in mean-zero measurement error since county shocks average to state shock

$$z_i^s = \sum_{c \in s} \frac{Emp_c^{t_0}}{Emp_s^{t_0}} z_i^c$$

- Not due to measurement error per se (county data can be perfect), but still drives mean-zero error relative to the economically relevant shock, with possible attenuation bias
- To me, seems most relevant in comparing CZ (plausible labor market) vs. county (too small)

Smaller effects in smaller regions

Correlated outside options bias may be more acute with finer geography

- Apparent migration responses to local shocks are reduced when shocks are correlated across relevant sources and destinations (Borusyak Dix-Carneiro Kovak 2023)

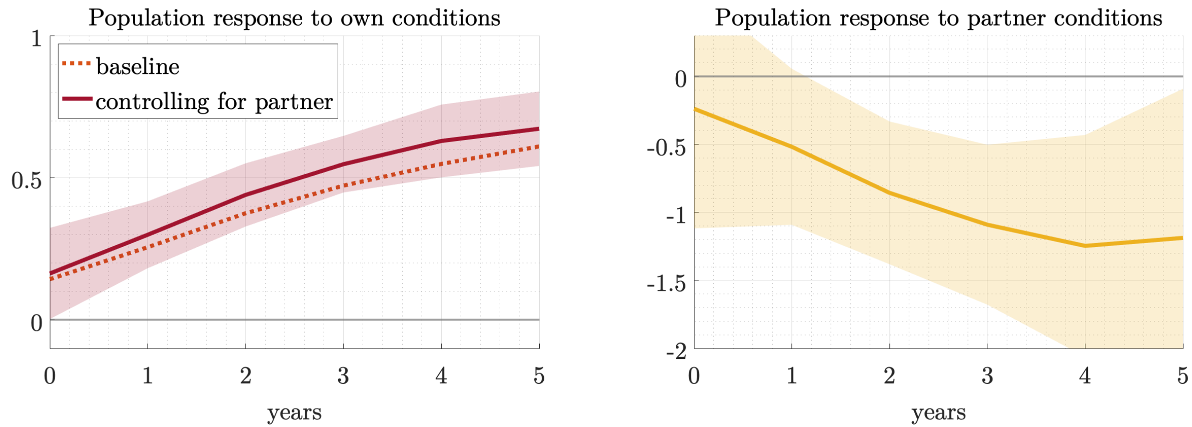

Figure 9: Controling for labor demand shock in effective partner

- The correlation may be stronger between finer geographic areas, which have more similar economic fundamentals (may explain state vs. CZ magnitudes)

Suggestion: try the outside-option correction at finer levels of geography

Conclusion

- Impressive data compilation and analysis over a huge span of time
- Places scattered datapoints from the literature in a consistent historical series
- Consistent results across multiple IVs helps overcome weaknesses in each
- Clear and consistent empirical work informing an important literature on internal migration and changes in economic dynamism

- Simple model of local labor markets helps clarify interpretation
- Understanding the role of pre-trends will link more clearly to prior literature and strengthen credibility of current results
- Auxiliary analyses may help understand smaller estimated effects with smaller region definitions

Smaller points

1. I was delighted to see the outside-option analysis in Section 3.1.7 and in your related *AEA P&P* paper. You explained the concern extremely clearly. In both papers, you describe the migration partners' shock as "a weighted average across likely destinations and origins of movers." It will be helpful to specify exactly how you calculate these weights. I also agree with the choice to implement the easy regional fix in this paper while mentioning the potential shortcomings in footnote 29.
2. I encourage more substantive interaction with Amior and Manning (2018 *AER*), who address a very similar set of questions using related methods. For example, I believe their model justifies the inclusion of the lagged Bartik shock in your control vector.
3. I didn't follow your response arguing that you don't need to account for the cross-sectional correlation driven by the fact that each location uses a common set of industry shocks when constructing the Bartik shocks. I do not find the claim that shares are exogenous plausible in this context – places specializing in vastly different industries are almost sure to exhibit differential outcome trends in the absence of shocks. The Bartik-robust inference in Borusyak, Hull, and Jaravel (2022 *REStud*) is so easy to implement, that I strongly recommend trying it.
4. I wasn't able to understand the timing of the industry changes you used to construct the Bartik shocks – please clarify in section 3.1.

Smaller points

5. In my mind, a big contribution of your paper is to place a wide range of other papers' results in context, with a consistent set of data and estimation strategy. You integrate closely with Dao et al. (2017), but I'd like to see some discussion of Yagan (2019 *JPE*) (see his Figure 2). I tried to see whether your results aligned with his, but they didn't seem to (though understanding the timing is challenging).
6. Footnote 6 now mentions studies finding that mortality relates to economic conditions, but there are also papers finding the fertility does as well. Black, Kolesnikova, Sanders, and Taylor (2013 *REStat*) find that energy price increases led to increased fertility in Appalachian coal-mining regions. A quick Google uncovered some other examples, such as the shale gas boom. None of this is likely to be quantitatively large enough to invalidate your approach, but it may be worth a quick mention for completeness.
7. As we discussed by email, it may be worth clarifying that the hats in (2) apply to the variables and the associated Deltas.
8. Figure 3 shows 90% confidence intervals, and Figure 4 shows one standard deviation error bands. In the absence of a strong reason otherwise, I suggest using 95% confidence intervals throughout, since readers usually expect them and may misinterpret the results as being more precise than they actually are.

Smaller points

9. On p.24, I didn't understand how the positive vs. negative shock measure was implemented. An additional sentence or footnote will be helpful here.
10. For the analysis in section 3.2.1, I couldn't tell if the Auerbach et al. (2020) sample omitted non-urban places. If so, could that explain the strange population results in this sample?
11. Is it worth flipping the sign of the y-axis in Figure 13 so they match the other shocks (given that the ADH shock is negative)? This may help make clearer that these results are largely in line with the others.
12. Typos
    1. P.13, 5th line from the bottom: "among many other Bound" -> "among many others Bound"
    2. P.30, 2 lines below equation: "relative lower shares" -> "relatively lower shares"