A Shock by Any Other Name? Reconsidering the Impacts of Local Demand Shocks*

Sean Bassler
Amazon

Kevin Rinz
Federal Reserve Bank
of Cleveland

David Wasser
US Census Bureau

Abigail Wozniak
Federal Reserve Bank
of Minneapolis and IZA

March 27, 2026

Abstract

Over the last decade, research on labor market adjustment following local demand shocks has expanded to explore a wide variety of measured shocks. However, the worker adjustments observed in response to these shocks are not always consistent across studies. We create a harmonized set of annual commuting-zone-level shocks following the major approaches in the literature to investigate these differences. As one might expect, shocks of different types exhibit different geographic and temporal patterns and are generally weakly correlated with each other. We find they also generate different employment and migration responses, with trade-related shocks showing little response on either margin, while more general Bartik-style shocks are associated with economically meaningful changes in both employment and migration.

Keywords: labor demand shocks, employment, migration
JEL classification: J23, R23, J61

*Any views, opinions and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the US Census Bureau, the Federal Reserve Bank of Minneapolis, the Federal Reserve Bank of Cleveland, or the Federal Reserve System. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7512631: CBDRB-FY23-CES023-001, CBDRB-FY23-CES010-025, CBDRB-FY24-CES010-017, CBDRB-FY25-CES023-012). Bassler's work on this project was completed prior to employment at Amazon. Ken Cowles and Alli Tucker provided outstanding research assistance on this draft. During earlier stages of the project, the authors benefited from helpful research assistance from Wenbo Li and Dylan Piatt as well as funding from the Washington Center for Equitable Growth. Comments from seminar participants at Macalester College, the University of Arizona, the U-MN Population Center, the University of Illinois Urbana-Champaign, the Federal Reserve Bank of Cleveland, the Upjohn Institute, and North Carolina State University as well as conference participants at the IRP Summer Research Workshop at UW-Madison and the Federal Reserve System Committee on Regional Analysis annual conference all improved the paper.

1 Introduction

A large literature examines employment and migration responses following shocks to local labor demand. These include China import shocks, the Great Recession, plant closings, and natural resource booms and busts (Autor et al., 2013; Greenland et al., 2019; Yagan, 2019; Foote et al., 2018; Black et al., 2005; Wilson, 2022), as well as the shift-share shock to local employment growth, known as the Bartik shock, which aggregates shifts in demand across the local industry mix due to changing industry fortunes elsewhere (Bartik, 1991; Blanchard and Katz, 1992; Dao et al., 2017). Almost without exception, analysis in this literature focuses on responses to local shocks using a single measure or source for these shocks.[1] Yet researchers and policymakers alike have tended to view the results in this literature collectively, making the implicit assumption that diverse shocks and the impacts that follow are comparable and their lessons generalizable, a potentially convenient but complicated assumption for policymakers who want research to inform deliberations on novel policy questions.[2] In this paper, we test this view in several ways and offer insight into the general features of worker responses following shocks to local US labor markets.

We compare multiple constructs of local demand shocks in terms of timing, incidence, magnitude, and outcomes. We make these comparisons in order to inform researchers and policymakers about how best to build on and use the existing body of evidence. We are motivated by two concerns that, in our view, should be addressed in order to make sense of the accumulating studies. The first is the possibility of limited generalizability: the lessons and impacts following one set of local shocks might not apply directly to another. This would be the case if local characteristics like population composition or geography interact with shocks to determine later outcomes. Generalizability might also be limited if the shocks themselves are not ultimately comparable to one another in terms of magnitudes, affected workers, or other features. To assess whether results should generalize, shocks first need to be comparable to one another. We take that step in this paper in order to better assess generalizability.

A second concern stems from the possibility that consistent findings across contexts

[1] Exceptions to this have appeared only recently: Foschi et al. (2025) and Faber et al. (2022).

[2] The Biden administration, for example, was heavily influenced by research on the effects of local labor demand shocks, including the "China shock," in formulating and discussing its climate policies. See Chapter 7 of the 2022 Economic Report of the President for an early discussion of research (Council of Economic Advisers, 2022) and remarks by Jared Bernstein, then chairman of the Council of Economic Advisers, for subsequent framing in public discussion (Bernstein, 2024). See also Chapter 8 of the 2025 Annual Report to Congress of the U.S.-China Economic and Security Review Commission (U.S.-China Economic and Security Review Commission, 2025) and Chapter 1 of the 2019 OECD Economic Survey of the United States (OECD, 2019) for examples of policy discussions airing concerns that impacts like those attributed to the lower trade barriers with China in 2000 will generalize to future settings.

could be spurious. Repeated findings across contexts is typically an ideal scientific outcome – similar stimuli leading to similar impacts – and might suggest a robust relationship between shocks from various sources and subsequent local outcomes. However, if the observed shocks are not independent of one another, then this robustness may be an illusion. Instead of repeatedly finding the same result using many different sources of demand shocks, the existing literature may instead be documenting the same response over and over if the measured shocks are in fact highly related to one another.

To assess whether these concerns are warranted, we answer three main questions. We begin by asking: how similar are these shocks to one another in terms of timing, geographic incidence, and the magnitude of first-order employment impacts? The literature assumes that these observable shocks approximate exogenous shifts in local labor demand curves. This leads us to examine two types of similarity: similarity with respect to temporal and spatial patterns, and similarity with respect to resulting changes in short-run employment. We use a set of shocks drawn from the literature but reconstructed to be comparable in terms of time period and geography. If the shocks reflect shifts in labor demand that stem from different exogenous sources, then we would expect them to have little correlation with one another across space and time. We would also expect the shocks to lead to similar changes in local employment, once expressed in comparable units.

Testing this intuition is straightforward, but the literature currently lacks a comprehensive assessment of these questions. To provide answers, we build a harmonized panel data set of shock measures used in previous studies. We reconstruct each shock measure independently and extend them to cover additional time periods, frequencies, or geographies. The result is a data set of annualized shocks at the commuting zone (CZ) level spanning 1994 to 2016 that includes several well-known measures of local labor demand shocks.[3] These include variants of the China trade shocks and the local employment growth shock, or Bartik shocks. We also study local shocks that are cross-sectional, namely, the Great Recession and the tariff gap measure of the China trade shock, and develop a simple method for comparing these against time-varying (panel) local shocks, since both types are often applied to answer similar questions. Having created a set of commonly used shocks defined for a consistent time period and set of geographies, we turn to County Business Patterns (CBP) data to assess their similarities in terms of timing, location, and severity as measured by employment changes.

We find that the major shocks in the literature are temporally and often spatially distinct. Correlations among the employment growth, import competition, and export competition shocks are typically modest over the period we consider. The distribution of CZ-year shock

[3]The choice of period is a result of choices made in constructing our harmonized series.

magnitudes differs substantially across shocks, especially in the tails. The magnitudes of import and export competition shocks are largely concentrated around zero, with a small number of more extreme values that tend to re-occur within a limited number of geographic locations. The employment growth shock is more geographically diffuse and not notably driven by outliers. There is some similarity between these panel shocks and the cross-sectional shocks we consider: the tariff gap shock, for example, is more strongly correlated with the import competition shock in the early 2000s, as is the employment growth shock with the Great Recession shock in the late 2000s. This may be due to a combination of the more general time-varying shocks partially reflecting some of the same information captured by the one-time shocks and pre-existing trends in places affected by one-time shocks.

We take an additional step toward comparability by expressing all shocks in terms of their effects on employment growth. This makes the constructs directly comparable throughout their distributions and reveals several additional patterns. We find that the trade-related shocks we consider have modest or null effects on employment, and when put in comparable employment growth terms, there are essentially no large trade-related shocks. However, many places do experience substantial employment growth shocks repeatedly, but we find that the relationship between employment growth shocks and employment growth is stable across the distribution of shocks and across places exposed to relatively more shocks. We then explore the dynamics of the impacts of our rescaled shocks across both panel and cross-sectional constructs using event studies and local projections analysis.

After assessing the relationships of our shocks to one another, we proceed to our second question: do similar shocks lead to similar outcomes? To answer this, we study their impacts on a local outcome that has been the subject of much analysis: migration. We turn to a second unique data set to study the impact of shocks on our example outcome of migration. We construct a longitudinal panel data set from the universe of IRS tax filings augmented with internal Census data providing individual characteristics for the period 1989 to 2018. We use these data to measure migration flows into and out of commuting zones on an annual basis. Our migration flows series is unique in the literature in offering both a high level of precision and consistent construction over our entire period. These data differ from the public microdata commonly used in the literature studying migration following local demand shocks by providing a much larger sample, which is useful because sample size is critical to credibly identifying population and migration changes following shocks (Dao et al., 2017). Data from tax returns may not capture the behavior of frequent non-filers well, but it should be representative of prime-age workers with high labor force attachment, who consistently engage with the tax system. We restrict our analysis sample to people with tax return data available in at least 70 percent of years. We find that characteristics of members of

our longitudinal panel match those of a comparison sample from the American Community Survey reasonably well.

With this second data source in hand, we examine migration responses following shocks using different constructs. We examine outflows and inflows separately, as the literature has generally found that these respond asymmetrically following shocks. A given set of shocks to local demand may lead to similar employment shifts, yet still deliver different migration responses if the barriers to successful relocation depend on the time, place, or nature of the local shock. We find that small magnitude shocks lead to small migration changes. In particular, we see a minimal short-term migration response to import and export competition shocks, either in terms of inflows or outflows. By contrast, employment growth shocks have economically significant effects on migration: increased labor demand reduces outflows by nearly 3 percent and increases inflows by about 5 percent, relative to the average migration rate in our sample. There is some evidence that the response of inflows to the employment growth shock can be moderated by place- and person-level factors. Places with more college-graduate residents or higher January temperatures, for example, see larger increases in inflows due to increased labor demand, while places in which a larger share of residents were born locally see smaller increases in inflows. More generally, moderators vary across the shock constructs we consider, suggesting that the dynamics of migration responses differ across contexts and that caution is warranted when applying results from one source of local shocks to other contexts.

Finally, we ask what scholars and policymakers should take from the literature on local shocks. We discuss whether local labor demand shocks as commonly measured in the data lead to any general responses, or whether researchers and policymakers are better advised to apply lessons from one shock to another with caution. Here we conclude that researchers should be aware of relationships between and underlying trends in any measured labor demand shocks they are studying, as well as data and aggregation issues surrounding their construction; should understand whether any of these are cause for concern in their settings; and should use appropriate tools from their respective econometric literatures to address concerns as appropriate. Our paper provides guidance on the first item along with some guidance on the remaining three. Policymakers looking to respond to local labor demand shocks in an evidence-based way would be well served by seeking out evidence based on a shock that aligns as closely as possible with the situations they actually confront rather than attempting to extrapolate from dissimilar shocks, which may have induced different labor market responses or been moderated by factors that are not relevant.

This paper contributes to a large literature on population and employment adjustments following local shocks. Prominent examples include Blanchard and Katz (1992), Autor et al.

(2013), Amior and Manning (2018), and Yagan (2019). We add to this literature by asking how best to view results from the growing collection of studies when those studies are informed by different time periods, geographies, and local shock constructs. Our paper also contributes to the considerable portion of this literature devoted to understanding migration responses to local labor demand shocks. Foschi et al. (2025) are motivated by similar questions. This has produced divergent answers to the question of whether affected workers migrate from weaker to stronger markets in the wake of these shocks. For example, Autor et al. (2013) find no detectable local population adjustments following localized China trade shocks, while Wilson (2022) documents a substantial migration response to the oil industry boom in fracking counties. In contrast to Autor et al. (2013), Greenland et al. (2019) use a different measure of localized China shocks and find that these lead to out-migration from affected areas at long time horizons. Other studies also come to conflicting conclusions about the migration response to changes in local economic opportunity. Dao et al. (2017) find substantial long-distance migration flows in response to surprise employment growth (or Bartik shocks). However, Yagan (2019) finds that outflows from heavily affected counties in the Great Recession often went to nearby counties.[4]

We inform this literature by testing the implicit assumption that local demand shocks generated from different sources and at different times lead to a consistent set of adjustments on both employment and migration margins. By examining the question directly in a consistent empirical framework, we can assess whether conflicting results in the literature stem from relatively minor differences in approach and time period or something more fundamental about different shock constructs.

This is related to, but distinct from, a well-known set of papers that examine the commonly used Bartik shock and ask under what conditions it meets identifying assumptions for place- and time-specific instruments (Borusyak et al., 2022b; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2025).[5] These papers focus on the local shift-share Bartik shock and adherence to the econometric assumptions behind its use. Our paper complements these papers by focusing on how one might aggregate results from studies of local shocks, even in cases where estimates are well-identified. In the language of Borusyak et al. (2025), common share weights or common shifts may both contribute to the relationships we identify across shock constructs. We return to the critiques from these authors in interpreting our results.

A second area to which our paper contributes is the growing literature on the role of place in providing economic opportunity distinct from the characteristics of individuals who

[4]See Jia et al. (2023) for a review.

[5]A closely related paper investigates issues of misspecification in studies of responses to local shocks (Borusyak et al., 2022a).

live and work there (Card et al., 2025; Gallagher et al., 2019). The divergent results mentioned above raise the question of why some changes in local opportunity appear to generate a substantial migration response while others do not. The divergence might simply be an artifact of the different data sources and time periods in each study. A more "apples to apples" comparison might produce more similar findings across contexts. However, the disparity might stem from more fundamental differences. For example, Molloy et al. (2019) find evidence for place itself in determining migration responses, namely distance to employment opportunity. Also, a key feature of measured shocks is their unequal geographic dispersion, which means that place factors may be correlated with shock severity (Goldsmith-Pinkham et al., 2020). Measured shocks likely affect different workers disproportionately, depending on the time period, industry context, and other construction details.

Finally, our paper provides important background for policymakers seeking to design adjustment policies for workers and places affected by declining local opportunity. The economic hardship that followed localized trade shocks in the 2000s has generated numerous calls for policies to assist affected families. However, important puzzles in the literature on the impacts of local economic shocks need to be resolved in order to design evidence-based policy to assist affected Americans.

The rest of this paper proceeds as follows. Section 2 describes the data. Section 3 compares different constructs of local labor demand shocks to each other in terms of distributional features, timing, location, and overall correlations. Section 4 compares shock constructs in terms of their employment impacts. Section 5 estimates effects of shocks on migration. Section 6 concludes.

2 Data Components

We assemble two main data components for our analysis. First, we extend the various measures of local labor demand shocks used in the previous literature. We harmonized these so that all shocks are defined for a consistent set of years and geographic areas. We also align base periods across shocks, where applicable. Second, we construct commuting-zone-level population counts and migration flows using annual address information from tax filers combined with demographic information available through the US Census Bureau. This section summarizes our construction of each component. Details are available in the Data Appendix.

2.1 Harmonized Set of Local Demand Shocks

To build our first data component, we replicate and extend measures of local labor demand shocks. We begin with measures of trade-related local shocks. These include shocks related to increased imports from China as well as the export-based trade shocks constructed by Feenstra et al. (2019). For the import-related shocks, we update the original China import trade shocks as constructed in Autor et al. (2013). In some analysis, we also use the modified measure of China import trade shocks in Acemoglu et al. (2016). As a final trade-related measure, we incorporate the tariff gap shock constructed by Pierce and Schott (2016b), which uses cross-sectional variation to identify import-related local labor demand shocks.[6]

We also incorporate local demand shocks measured using changes in employment or unemployment, including the standard Bartik-style employment growth shock and a cross-sectional shock based on local variation in the severity of the increase in the unemployment rate during the Great Recession, as defined by Yagan (2019). Later in the paper, we discuss limited extensions of our analysis to other local shock measures where data prevent us from fully adapting them into our harmonized framework. These include trade shocks originating in NAFTA and local robot adoption.

Table 1 provides an overview of the shocks we construct. This includes the assumed-exogenous source of variation, a narrative desription, and the reference that introduced each shock to the literature. Importantly, the final column of the table defines how each shock will be described in the text of this paper and labeled in tables and figures.

Most of these measures are constructed using a shift-share approach that weights a shift in imports, exports, employment, or tariff sizes by a geographic unit's exposure to the shift, measured by its industry shares of employment. The shifts are typically time-varying; the employment shares vary with geography. We illustrate this general approach using an industry-based shift-share shock below in Equation (1).

$$ADH_{it} = \sum_{j} \frac{L_{ijt}}{L_{it}} \frac{\Delta M_{jt}}{L_{jt}} \tag{1}$$

In the example above, i indicates geographic units and t time. The j subscript denotes

[6]The tariff gap shock measures exposure to Chinese import competition using changes in tariff policy in 2001. Before 2001, China was granted normal trade relation (NTR) tariff rates on an annual basis. This policy lowered the rates on Chinese imports, but the annual renewal of these rates was politically contentious and uncertain. This uncertainty protected US manufacturers from Chinese import competition, but it was eliminated when China joined the World Trade Organization (WTO) in 2001. At this time China was granted permanent NTR. The tariff gap shock measures commuting zones' exposure to changes in the NTR and non-NTR tariff rates.

industries and L employment at different levels of aggregation. Hence the first term in the sum is the local employment share by industry, which weights an industry-specific shift in the second term. In the case of the shock in Autor et al. (2013), for example, the shift is the per-worker change in dollars of imports by industry category over some period of time. This is weighted by a geographic unit's employment intensity in that category to generate a shock measure. Other shock measures are constructed by exchanging the underlying shift measure. A full description of the construction of all shock measures is given in the Data Appendix.[7]

We construct each shock using the data and methodologies of the original papers, with modifications to define originally decadal shocks at an annual frequency and harmonize industry and geographic classification across time. Employment information comes from the County Business Patterns data, which offers employment counts for year-by-industry-by-county cells. We follow Autor et al. (2013) and impute employment counts for observations that are missing them. Trade data come from Schott (2008). We use the NAICS 1997 vintage for industry classification and the 1990 commuting zone vintage for geographic information.

Several important shocks defined previously in the literature can be easily adapted to our harmonized framework. Details on construction of the shock measures is given in Appendix A and Appendix B. However, in the case of the China import shocks originally defined in ADH, we made a number of modifications to follow later authors and to extend the data over time. In Appendix B, we show how the import shocks as we construct them compare to those in the ADH replication files.

Distributional statistics for the shock measures as we have constructed them are summarized in Table 2. For comparability, they have been converted into z-scores. We have also harmonized the signs across shocks, so that a negative value always reflects a reduction in local labor demand. As described in Appendix A, the measures in 2 require extensive harmonization and imputation of underlying data series (with the exception of the Great Recession shocks). The table shows that, even after this cleaning, shocks defined by flows of imports and exports contain outliers. After trimming these, the normalized range of the shock constructs is broadly similar, though the upper values (99.5 percentile) for the trade flow shocks is about 50 percent larger than for other constructs. The final column shows that trade flow related panel shocks are more tightly clustered around zero, with a 90-to-10th percentile range that is half that of other standardized shocks.[8] Appendix Table C1 shows

[7]Table 1 includes two measures constructed with "incomplete shares": the construction- and manufacturing-focused employment growth shocks. If these were to be used for causal analysis, an additional control would be required, as described in Borusyak et al. (2025). We conduct our descriptive analyses without this adjustment.

[8]Subsequent estimates include these outliers, but results from re-estimation omitting them are available

descriptive statistics for the shocks without normalizing.

2.2 Population Flows

Our second data component consists of the population flows that are the dependent variable in our application. We use underlying population stocks from Census-IRS panel data available for 1989, 1994, 1995, and 1998-2019. These data provide detailed information on geographic location, demographics, and income with extensive coverage of the US population. Demographic information includes age, race, sex, and place of birth. It also has wage and salary information for every year a person filed a tax return. Hence, these data allow us to measure migration flows between commuting zones by demographic and earnings statistics of the migrants. Our flows reflect flows for highly attached prime-working-age people constructed from restricted-access Census Bureau microdata from the universe of tax returns combined with demographic information from sources internal to the Census Bureau. From these, we construct the migration outcomes analyzed in Section 5. Using flows constructed from a highly attached panel, as opposed to the full administrative data, reduces the underlying sample size but has important advantages. It focuses on the people most likely to respond to local labor demand shocks. It also reduces noise associated with transitions into and out of the labor market, as well as transitions into and out of the data by people migrating internationally.

We construct population stocks and ultimately migration flows from this administrative data set in three steps. First, we match observations at the person-year level to commuting zones of residence. Each record has information on an individual's residential location from zip codes on their 1040 forms, and a large majority can be linked with the county associated with the filer's exact address in the Census Bureau's Master Address File (MAF) using an anonymized internal address identifier (MAFID). We use both sources to determine an individual's location in a given year, but we prioritize information from the 1040 filings as these were completed by an individual in the year of interest and are available for essentially all 1040 filers. If an observation's 1040 zip code maps to a single commuting zone, then we use that commuting zone as the place of residence for that observation. If this is not the case, but the observation has been matched to a MAFID, then we use the county information associated with that MAFID to match the observation to a commuting zone. Using this method we match over two-thirds of the observations to commuting zones, as seen in Appendix Table C2. The table shows the number of people between the ages of 25 and 55

upon request. Re-estimation of correlations and employment regressions suggest that excluding outliers does not affect our results.

that we can match to commuting zones by year. In 2019, we directly matched 112 million out of 148 million observations to commuting zones using the 1040 zip codes and the MAF.

While the large majority of people who appear on a tax return can be matched to a commuting zone, a meaningful share of people we see in the tax data do not appear on returns each year. As a result, about 20 to 30 percent of theoretically matchable people cannot be matched to a commuting zone each year. In some cases, not filing a tax return is the correct thing to do (e.g., for people with earnings below the filing threshold, or for non-citizens who have emigrated from the United States), while in others it likely reflects noncompliance with tax laws.[9]

Given these gaps in observed residential location, we impose several restrictions on our analysis sample in our second step. We restrict our panel to people who are highly attached to the tax system, and typically also the labor market. We also use only observations between the ages of 25 and 55, as these are the years people are most attached to the labor market. We then drop people who can be matched to a commuting zone in less than 70 percent of the years that they are alive between the ages of 25 and 55. This panel allows us to track an individual's location throughout their prime working years. We assess the impact of these restrictions by comparing our sample to population data from the 1990 and 2000 Decennial Census, and the 2001-2019 American Community Survey, as seen in Appendix Tables C4 and C5. For people born in the 1940s and 1980s, the tables compare demographic characteristics of people in our panel, grouped by the share of their years that have commuting zone information available, to demographic characteristics of Decennial Census/ACS respondents. People with a high commuting zone match rate have demographics comparable to the population as a whole in the survey data.

For a small number of the remaining observations, we cannot match directly using the 1040 zip codes or the Master Address File. This occurs when an observation has a 1040 zip code that maps to multiple commuting zones while missing a Master Address File county. We could use a probabilistic method to impute commuting zone for these large zip codes, but as seen in the table, imputing commuting zones in this case does not net us many more matches. Hence, we drop observations that cannot be matched directly using a zip code or the Master Address File.[10]

[9] IRS estimates compiled by Guenther (2023) show that the net tax payer noncompliance rate averaged 14.1 percent between 2001 and 2021.

[10] We make two final remarks on our commuting zone matching method. First, the US Department of Housing and Urban Development has published multiple crosswalks that map zip codes to county mappings. These crosswalks are indexed by year. In theory, using more of these crosswalks has the advantage of addressing the time inconsistency in county and zip classification. However, we found that using more of these crosswalks did not yield a higher number of matches, and so we use only the 2021 crosswalk. Second, the data have zip codes from two sources: 1040 forms and 1099 forms. The zip codes from these two sources conflict frequently, and we believe that the 1040 forms are more accurate because they are more likely to

In our final step, we collapse our panel to get annual population counts and flows between commuting zones.[11]

3 Correlation and Incidence of Local Demand Shocks across Space and Time

We begin with a straightforward question: how related are commonly used local labor demand shocks? Some shocks should be related to one another mechanically by virtue of their construction. For example, the shocks to local employment growth derive from changes to industry-level employment. Therefore they seem likely to correlate with shocks that derive from changes to a subset of industries, such as shocks derived from changes in trade relationships, or with shocks that measure changes in local employment by other means, such as local unemployment spikes in the Great Recession shocks.

Although the potential mechanical relationships are obvious, to our knowledge they have not been systematically assessed in the literature. The extent of these relationships is ultimately an empirical question, and our aim in this section is to answer it. To do so, we first describe the correlation of shocks with one another and then with themselves (persistence). We conclude the section with an analysis of the spatial dimensions of the correlations we identify. We emphasize relationships across the panel shocks, but we include one-time shocks in some of the analysis that follows before turning to a deeper analysis of their relationships to panel shocks in the next section.

3.1 Correlation and Persistence of Local Shocks

We first examine correlations of each shock with the others. Table 3 reports simple correlations across our standardized shock constructs.[12] There is substantial variation in the correlations between pairs of shocks we consider, both within and across shock types. Turning first to the trade shocks, the two import-related shocks are fairly highly correlated with each other, with $\rho = 0.65$. Although this is a strong correlation, one might have expected

be supplied contemporaneously by the filer rather than the entity issuing the 1099. Also, we constructed migration rates using the 1099 zip codes to assign location, and they were implausibly high. Hence, we decided to discard the geographic information from the 1099 forms.

[11] Note that the administrative tax data are aligned to tax years, so the commuting zone of residence generally corresponds with the following calendar year. For example, the 2005 data file shows information from 2005 tax forms, even though the forms were generally filed in 2006. Hence, we shift the years in the data files up one to align an observation's commuting zone of residence to the timing of our shock measures.

[12] Results in the table are little changed if we use the raw constructs instead of z-scores.

an even higher coefficient given the relatively small difference between the construction of the two measures. Both variants are positively correlated with the tariff gap shock (during the 2000-2002 period), though not as strongly. Export shocks are negatively related to import-based shocks in panel settings but unrelated to the tariff gap shock of 2000. Alternative formulations of employment growth shocks based specifically on the construction and manufacturing sectors are also highly correlated with the economy-wide employment growth shock that we focus on, with $\rho \approx 0.75$.

Correlations between trade-based and employment-based shocks are generally weaker, suggesting that these measures may be capturing changes in different underlying economic conditions. However, comparing correlations across classes of shocks reveals some interesting patterns. Panel import-related trade shocks are essentially uncorrelated with local employment growth shocks, but the 2000 tariff gap shock (which is also import-related) and panel export-related shocks are positively related to local employment growth changes.

One might have expected the correlation between the employment growth and local unemployment shocks to be higher given that declining employment and increasing unemployment are often two sides of the same coin. One difficulty with making and interpreting this comparison is that employment growth shocks are time-varying, while Great Recession unemployment spikes are time-invariant. Table 3 correlates the unemployment rate spikes with shocks to local employment growth using values for 2008-2010. When they are instead correlated year-by-year, as in Figure 1, there are years in which the correlation is high–notably, years that fall during recessions.[13] The figure also shows that correlation of unemployment spikes was higher with employment growth shocks overall in recession periods than with a measure focused on either manufacturing or construction. Thus the greater correlation of these shocks in recessions is not fully explained by changes in those sectors. The generally low correlation between the two sources of shocks outside of recessions suggests that the Great Recession unemployment rate spikes were at least partially independent of local employment dynamics due to industry-level changes in both earlier and later periods.

Although much of the literature on local demand shocks in the US defines local markets using commuting zones, this choice entails tradeoffs. Some of these are explored in Foote et al. (2021), who examine the sensitivity of CZ definitions to alternative choices about how to aggregate the underlying commuting data. Other considerations may also be relevant. Data on employment and other local outcomes are often collected or reported for units other than the CZ. As a result, outcomes may be measured with less error using geographic groupings that correspond better to administrative units. At a more conceptual level, commuting may

[13]There was substantial overlap between the geographic footprints of the two recessions that occurred during the period covered by these data. See Hershbein and Stuart (2024) for more details.

be an imperfect indicator of integrated local economic activity. Although commuting may be driven by similar economic activity taking place across space, it may also arise as a means to diversify access to markets that function differently. For example, a worker who has strong family ties to her hometown may choose to drive an hour to a larger city for work, but this stems from differences in opportunity across the two markets, not economic similarity.

Because of these potential tradeoffs, we probe for differences in these relationships by level of geographic aggregation. We repeat the correlation analysis using shocks constructed with state of residence as the geographic dimension of the panel. Analyzing the behavior of shocks constructed at the two levels relative to others in the literature can inform assessments of these tradeoffs. Results are presented in the bottom panel of Table 3. Many correlations from the top panel are stronger when the underlying changes are aggregated at the state level as opposed to CZ. Export shocks are more negatively related to import shocks, and both are more strongly related to local employment growth shocks (though in opposing directions). Local employment growth changes are also more strongly related to unemployment rate spikes in the Great Recession. The one-time tariff gap shock is also more strongly related to the panel trade shocks.

What should we conclude from these patterns? One argument for the use of CZs is that they better approximate the geographic pattern of economic activity.[14] If this is true, we would expect the mechanical relationships between shock concepts to lead to higher correlations when shocks are constructed using CZ units. We see the opposite, which raises questions about how well CZs measure economic activity.

We explore these relationships using other methods in two appendix tables. First, to better understand the correlation of each shock with the others, we regress each panel shock on the complement of other panel shocks. The results are in Appendix Table C6 and align with expectations and with our findings in Table 3. Second, we use rank correlations in place of raw correlations in Appendix Table C7. The patterns in Table 3 are robust to these changes.

We also test for autocorrelation of each shock. We regress each shock on lags of one, two, and three years. We include four different specifications to illustrate how autocorrelations change when including year and commuting zone fixed effects:

$$shock_{it} = \beta_0 + \sum_{h=1}^{3} \beta_h shock_{i,t-h} + \Theta_i + \Theta_t + \epsilon_{it} \tag{2}$$

Here, β_0 is a constant, Θ_i is a commuting zone fixed effect, and Θ_t is a year fixed effect.

[14] A second argument is that they divide the entire geography of the US into exhaustive and exclusive units, but states do this as well.

Standard errors are clustered by commuting zone. The results are in Table 4. Panel A reports estimates from specifications including CZ and year fixed effects. In this panel, there is evidence of some positive serial correlation in the employment growth shocks but little in the trade shocks. Some level of positive serial correlation is to be expected for the local employment growth shocks, since the weight terms are correlated in adjacent years by construction, as they are made using five-year trailing averages of employment. The coefficient on the first lagged employment growth shock suggests that a one standard deviation larger shock in year $t-1$ is associated with a shock that is 0.15 standard deviations larger in year t. The coefficient on the two-year lag is less than half as large but remains significant. The trade shocks exhibit economically small reversion over two- or three-year periods.

Given the way our harmonized shocks are constructed and the stability of local industry shares over time, one might have expected more autocorrelation than our baseline analysis finds. We explore this in the remaining panels of the table, in which we exclude the fixed effects singly and then entirely. The degree of autocorrelation is somewhat sensitive to specification. Dropping commuting zone fixed effects leads to more autocorrelation in the employment growth shock, as shown in Table 4, but this is just another way of saying that local industrial composition is persistent and some industries have consistently below-average employment growth.[15] Figure C1 shows this second point to be true of manufacturing. Other differences in autocorrelation across specifications are generally not systematic, except that when shocks are measured over longer time horizons (four or more years) first-order autocorrelation is consistently negative and larger in magnitude than when shocks are measured over shorter horizons (including the one-year horizon used in our main analysis), as shown in Table C9 and Appendix Table C8. Mean reversion could provide a potential explanation for this pattern if modestly autocorrelated shocks build up over one several-year period before local economies re-equilibrate in the next period.

3.2 Spatial Incidence of Shocks

Examining these shocks across space also suggests that they capture somewhat different changes in local economic circumstances. We first explore this using visual evidence from heat maps of repeat large shocks, as in Figure 2. We define a shock as large if it is more than one standard deviation above or below its mean, where the mean and standard deviations are taken over all shocks of the same type that occurred over the entire sample period, 1994-2016.

[15] Indeed, re-estimating the regression in Table 4 using the manufacturing share of employment produces coefficients on the first lag that range from about 0.8 to 0.9, depending on which fixed effects are included.

The heat maps in the left column plot how frequently commuting zones experience large shocks that fall below the mean. Import-related large negative shocks are concentrated in the Southeast and Midwest regions. Employment growth shocks are more diffuse, with almost every commuting zone experiencing a few large shocks below the mean, but commuting zones that experience more repeat shocks of this type are also located in the Southeast. As for export-related shocks, few commuting zones experience large negative shocks of this type, and commuting zones that do are fairly dispersed across the Great Plains region.

The heat maps in the right column plot how frequently commuting zones experience large shocks above their mean. Very few commuting zones experience any positive import-related shocks. This is due to the steady increase in Chinese imports over the sample period. As a result, commuting zones tend to experience either negative shocks or shocks that are close to zero. The Western region has the highest concentration of large employment growth shocks above the mean, while the commuting zones that experience repeat large export-related shocks above the mean are more dispersed, but mostly located in the Great Plains region and the Rust Belt.

Overall, the geography of these shocks paints a fuller picture of the variation underlying the estimates in Table 4. Few places experience repeat large shocks of all kinds. This is consistent with the low to modest levels of correlation we see across the different shock types in Figure 3. Contrasting spatial variation in the export- and import-related shocks are also consistent with the literature that takes a more expansive view of trade-related shocks. That literature generally finds export-driven increases in employment that offset the declines in employment from import-related shocks (Feenstra et al., 2019). The figures and Table 3 simply show that these offsetting patterns materialized in different places.

Having documented where repeat large shocks occur, we also describe their role in the total variation delivered over the entire distribution of shocks. Do large shocks comprise most of the total variation, or only a portion? We assess this using scatterplots as shown in Figure 3. These figures compare the contribution of a CZ to the total variation in a given shock (indicated on the y-axis) to the CZ's percentile ranking in terms of the shock's CZ-level standard deviation (indicated on the x-axis).[16] The figures reveal a mix of intuitive and surprising findings. The figures show that CZs that experience more variability in a given type of shock account for larger shares of the total shocks, in absolute value. The trade-related shocks are also highly concentrated. CZs experiencing the most variation under these shocks account for shares of total variation that are an order of magnitude larger than the shares from the most variable CZs as measured by employment growth shocks.

The outsized role of variable places and the greater concentration of this role under trade-

[16]Specifically, the y-axis value for a CZ i' is calculated as follows: $\frac{\sum_t |shock_{i=i',t}|}{\sum_i \sum_t |shock_{it}|}$.

related shocks as opposed to employment growth shocks might have been expected from the descriptive analysis above. However, the magnitude of the differences, as shown in the scale differences on the y-axes, is quite large and difficult to intuit without this calculation. The most variable CZs under trade shocks contribute variation that is more than an order of magnitude larger than under employment growth shocks. By contrast, the scatter plots show much more consistency across places in contributed variation under employment growth shocks. The scatter plots for these show that variation accounted for by a given percentile of variability falls in a band that is fairly consistent across percentiles. This means that low-variability CZs often contribute as much total variation as high-variability CZs. This suggests that the composition of employment growth shocks differs across CZs, with some places experiencing many modest-sized shocks and others experiencing a few large shocks.

4 Defining and Comparing Shocks Based on Employment Impacts

We next compare shock constructs in terms of their effects on short-run employment growth. As all shocks are constructed to reflect shifts in local labor demand, we expect to observe employment changes following a shock. This would be consistent with findings in the literature that study the impacts of these shocks separately. However, estimating employment impacts of a range of shocks serves an important additional purpose in our study. Such estimates function as a rescaling of the diverse set of local demand shocks. This expresses all shocks in the same units and allows for direct comparisons of magnitudes across them.

4.1 The Effects of Shocks on Employment Growth

Expressing shocks in units of subsequent employment growth combines a measurement question - how to quantify local demand shocks - with a question about the impact of those shocks on an important outcome. This duality is central to thinking about local shocks but is not often discussed in analyses of single shocks. We believe it is important for the literature to grapple with the fact that we cannot observe shifts in the local demand curve for labor. Instead, we observe proxies or events that would be expected to shift the curve. We should then ask how to compare these events to one another. Our contribution here is to use employment growth impacts as a basis for comparison.

To compare employment impacts across shocks, we set up a simple regression. Let $g_{it} = \frac{n_{it}-n_{it-1}}{n_{it-1}}$ denote employment growth for commuting zone i in year t. We regress

annual employment growth on each annual shock, absorbing year and commuting zone fixed effects, and clustering by year and commuting zone, with commuting zones weighted by their employment level in 1990:

$$g_{it} = \beta \text{shock}_{it} + \Theta_t + \Theta_i + \epsilon_{it} \tag{3}$$

where Θ_t and Θ_i denote year and CZ fixed effects, respectively. β indicates percentage points of employment growth above or below that given by the CZ-level and national trends in employment growth absorbed by the fixed effects.

Results are reported in Table 5. As one might expect based on their construction, the Bartik shocks are strongly associated with employment growth. A one standard deviation EG shock is associated with a 2.7 percent increase in employment growth, year-over-year.

The trade shocks, which do not have the same mechanical relationship with employment growth, show smaller and statistically imprecise responses. The export-driven shock, which has the strongest relationship with employment growth among the trade shocks, is associated with roughly an order of magnitude less employment growth than the employment growth shocks. A standard deviation increase in this shock increases employment by 0.2 percent. The import-related shock of Autor et al. (2013) has an even smaller effect. A one standard deviation decrease in Chinese import intensity using this construct increases employment growth by 0.06 percent. The alternative import-related shock from Acemoglu et al. (2016), while constructed similarly, actually has a negative and statistically insignificant effect on employment growth. Because the two import-related shock constructs are highly correlated with each other but the measure in Autor et al. (2013) is more strongly correlated with employment growth, we focus on this as our measure of import-related shocks in the remainder of our analysis. Table 6 shows that these patterns are consistent across quantiles of the dependent variable distributions.

Although the shocks have been converted into z-scores, the standard deviations are defined in terms of each shock's original units and hence are not directly comparable to one another. To make shocks comparable to each other throughout the distribution, we use the regression coefficients β from Table 5 to express all shocks in the same units. These coefficients are semi-elasticities representing the percent change in employment given a one-unit (i.e., standard deviation) change in the shock. We multiply each shock by these semi-elasticities to rescale the shocks, allowing for a comparison of the distributions of shocks on the same scale: their relationships to changes in local employment.[17]

Figure 4 shows a histogram of the rescaled shocks. Comparing the rescaled distributions

[17]Key moments of the distributions of these rescaled shocks are summarized in Table C3.

across shocks reveals that the conclusions one might draw about their relative dispersions based on Table 2 are incomplete. The rescaled import- and export-related shocks remain tightly distributed around zero, but the rescaled employment growth shock appears much more diffuse than the version in Table 2, spread out by the relatively strong relationship between that shock and local employment growth. Similarly, the extreme values of the import- and export-related shocks are less evident in the rescaled versions, diminished by the weaker relationships between these shocks and employment growth and outweighed in this figure by the thick tails of the employment growth shock.

4.2 Spatial Incidence of Rescaled Shocks

We also examine how often large rescaled shocks are distributed across commuting zones. Figure 5 shows how often commuting zones are exposed to high-intensity rescaled shocks, that is, shocks expected to be associated with changes to employment growth above 0.5 percent or below -0.5 percent.[18]

The heat maps in the left column plot how frequently commuting zones experience large adverse shocks, specifically a rescaled shock below 0.5 percent. As they have little effect on employment, almost no commuting zone is hit by adverse high-intensity import- or export-related shocks using our rescaled measures. Conversely, many places are hit with repeated employment growth shocks below -0.5 percent. These are concentrated in, but not exclusive to, the Southeast region.

The heat maps in the right column plot how frequently commuting zones experience large positive shocks (above 0.5 percent). Again, very few commuting zones experience any high-intensity import- or export-related shocks. Meanwhile, many places experience repeated employment growth shocks above 0.5 percent, though the Southeast experienced these less often.

Given the outsized contribution of more highly variable places to the total variation in these shocks, it is natural to wonder whether these places also drive the employment response to these shocks. Table 7 shows estimates of the employment effects of the shocks we consider based on a specification that also includes an interaction between the shock value and an indicator for whether a place's within-CZ standard deviation of the given shock is in the top quartile. Despite the most volatile places being responsible for larger shares of variation in import and export competition shocks, this interaction is not statistically distinguishable from zero for any of the shock constructs, implying that the employment growth impacts of

[18]The thresholds were chosen to be symmetric around zero while also permitting some CZs to have detectable changes under the trade constructs.

the various shocks do not differ across high- and low-volatility places.

4.3 Comparing Time-Varying and One-Time Local Shocks

Another challenge to comparability across constructs is that some are defined as cross-sectional while others are time-varying, or panel. To address this challenge, we study the effects of all shocks on employment several years after they occur. With cross-sectional constructs, this approach is dictated by the data. With panel constructs, it is easy to define them as a series of one-time shocks. We can then compare the employment impacts generated by the series of one-time shocks directly to those generated by cross-sectional shocks.

Specifically, let k denote the year a shock occurred. For the Great Recession shock $k = 2007$, and for the tariff gap shock $k = 2000$. For the annual shocks, k can be any year as these shocks occur across all years. Let $y_{it} = \frac{n_{it} - n_{i,k-1}}{n_{i,k-1}}$ be commuting zone i's employment growth since year $k-1$. We measure the effect of a shock on y_{it} in the years before and after the shock occurs:

$$\tilde{y}_{it} = \sum_{h=-5}^{10} \delta_{kh} \mathbf{1}\{t = k + h\} shock_{ik} + \Theta_t + \Theta_i + \epsilon_{it} \tag{4}$$

where $shock_{ik}$ is the shock to commuting zone i in period k, $\mathbf{1}\{t = k + h\}$ is an indicator variable that equals 1 if year t occurs h years after the shock hit in year k, Θ_i is a commuting zone fixed effect (included for panel constructs only), and Θ_t is a year fixed effect. Standard errors are multiway clustered by year and commuting zone. Commuting zones are weighted by their employment level in 1990. We also measure the effect of shocks on employment growth before the shock occurs as a placebo test; the coefficients δ_{kh} for $h < 0$ should be close to zero, as a local shock should not affect employment outcomes before it happens.[19]

The key independent variable is $shock_{ik}$; its coefficient, δ_{kh}, measures the effect of the shock h years since or before the shock occurred in period k. More specifically, the coefficient δ_{kh} measures how much variation in the outcome in period $k+h$ can be explained by variation in the shock at time k. Hence, measurement of δ_{kh} stems from variation in the shock *across* commuting zones. We include year fixed effects to control for nationwide trends across time, and commuting zone fixed effects to control for differences across places. So, δ_{kh} is interpreted as the effect of a one standard deviation increase in the shock in year k on employment growth between years $k-1$ and h, above or below annual and local trends.

Figure 6 plots the corresponding values of δ_{kh} for the two cross-sectional shock constructs,

[19] Note that when including a CZ fixed effect, one of the δ_{kh} coefficients is omitted by construction. We omit $\delta_{k,-1}$.

the Great Recession unemployment spikes and the tariff gap, as well as employment growth shocks that occurred in the same year. Looking across y-axes of the panels, a first observation is that the size of subsequent employment effects differs across shocks. The heterogeneity in magnitudes of employment impacts that we identified in the panel constructs clearly extends to the cross-sectional constructs. Turning to the dynamics within each construct, Panel (a) shows the effect of the 2007 Great Recession shock. As expected, the effect of the shock before it occurs is not distinguishable from zero. After it occurs, there is a large, positive, persistent effect on employment. For instance, a one standard deviation increase in the shock above trend is associated with around a 4 percent increase in employment five years after the shock occurred.

Panel 6b shows the employment effect of the 2000 tariff gap shock. Unlike the Great Recession shock, there is a clear pre-trend under our specification in that the future shock has a significant negative effect on employment. The employment effect is remarkably similar before and after the shock occurs. As discussed in Greenland et al. (2019), the shock is correlated with pre-treatment local growth rates. We suspect the CZ fixed effect does not adequately control for these trends.

For comparison to the one-time shocks, Panels 6c and 6d show the 2007 and 2000 employment growth shocks. These shocks, like the tariff gap shock, also exhibit pre-trends. While the employment growth shocks and tariff gap shocks use different measures for industry-level shifts, they use the same employment shares to weight these shifts. Hence, we suspect that the underlying employment shares are correlated with CZ-level trends, causing both the tariff and employment growth shocks to fail the pre-treatment test.

These comparisons between time-invariant shocks and contemporaneous employment growth shocks raise the question of whether these specific instances of employment growth shocks are typical. Figure 7 suggests that similar analyses of most other instances of employment growth shocks look more like the 2000 case than 2007. Panel 7a shows that for many years, event study coefficients tend to drift up toward zero as event time approaches zero from below and continue along a similar trajectory in positive event time.

This drift is not universal in event study analysis of employment growth shocks. Panel 7b shows select years that do not exhibit this pattern of event study coefficients. Pre-trends in the years depicted here are fairly flat, and the trajectory of estimates does appear to change after at least one of these focal years. We discuss the potential significance of these being Economic Census years in the next subsection.

Figure 8 reports analysis of pre-trends in underlying shocks by focal year. These are

based on the following simple regression equation:

$$Shock_t = \beta_1 Time_t + \beta_2 Time_t \times PostShock_t + \varepsilon_t \tag{5}$$

Panel 8b reports the estimated trend in pre-focal year shocks (β_1 from equation 5), while panel 8a reports whether that trend changes after the focal year (β_2 from equation 5). Panel 8b confirms essentially flat pre-trends for the years depicted in panel 7b, though in only one year are these statistically insignificant. Consistent with the visual evidence in 7a, about half the centered event years have a positively sloped pre-trend that is both statistically significant and near or greater than 0.01, which exceeds many of the post-shock trend break coefficients reported in panel 8a. Figures 9 and 10 show that these patterns are present and similarly severe when using the Business Dynamics Statistics as a source of local employment by industry or when using state as the level of geography. Using state to define geography results in meaningfully fewer event years with significant pre-trends, but the problem is not eliminated.

Local projection estimates can also facilitate comparisons between time-varying shocks and one-time shocks. For one-time shocks, estimation is identical to the event study framework just discussed, for the post-shock period. For time-varying shocks, the local projection approach essentially averages effects of each annual shock at each horizon rather than focusing on individual annual instances of a shock. This approach can also provide insight into potential dynamic effects of time-varying shocks that may not be captured by specifications based on including lagged shock values.

Figure 11 shows local projection estimates of the employment effects of employment growth, import competition, and export competition shocks. Employment growth shocks have positive employment growth effects for two years concurrent with the shock. As in prior specifications, the employment effects of import and export competition shocks here are small and generally not statistically distinguishable from zero. The point estimate of the employment effect of the import competition shock does appear to grow more negative between seven and ten years after the shock, though it only becomes statistically significant in year ten. The effect of the export competition shock is small in magnitude and statistically insignificant at all horizons.

4.4 Measurement and Aggregation Issues

Three of the years with flat pre-trends in the period considered (2002, 2007, and 2012) happen to be years in which the Economic Census was conducted by the Census Bureau.

This is interesting because the Economic Census is one input into producing the CBP data used in this analysis. The CBP data are based on a combination of firm-level data from tax returns that capture how much economic activity is happening (e.g., employment and payroll totals) and data from surveys that capture where it is happening (e.g., by identifying which establishments firms own and where they are located). The Economic Census is the most comprehensive of these surveys, and the years in which it is conducted are the years in which information about where businesses exist and how they are connected is at its highest quality.

In non-Economic Census years there is less certainty about what economic activity belongs where and more potential for measurement issues to arise with less comprehensive contemporaneous surveys or Economic Census data that could become increasingly dated over time. Moreover, establishment births and deaths may be captured less quickly when surveys are less comprehensive. Lags in tax filings by firms that have just been created or are winding down operations could also result in CBP data missing changes in employment due to firm births or deaths.[20] If the allocation process leads employment in some firms to drift from the distribution measured in one Economic Census to the distribution measured in the next, that could contribute to non-zero pre-trends in the kind of event study analyses discussed above.

To investigate whether measurement issues related to the construction of the CBP data might influence our event study analyses, we compare the estimates presented above with similar estimates based on an alternative source of employment data and a higher level of geographic aggregation. The other data source we consider is the Business Dynamics Statistics (BDS) data, which can also be used to construct localized measures of year-over-year employment changes and provide a measure that is based only on employment changes within continuing establishments. This, of course, does not capture all changes in employment, since firm and establishment births and deaths are important margins for growth in overall employment. Still, this measure captures about 77 percent of total employment growth in the average CZ-year in our analysis period and circumvents measurement issues related to the difficulty of measuring the birth-death margin.

The alternative level of geographic aggregation we consider is the state level, since multi-unit firms are less likely to operate across states than across CZs, and even when they do operate across states, the existence of state-level data on firm employment (such as from the LEHD) likely leaves less room for error in allocation of economic activity across states

[20]Subsequent revisions to the data should largely address these possibilities, leaving them as a concern primarily for years of data since the most recent Economic Census. The CBP data we use entirely pre-date the 2017 Economic Census, making this issue less likely to complicate our analysis.

within the CBP data.

Figure 9 presents our alternative event study analyses of employment growth shocks by year using CZ-level BDS data, along with our estimates of the slope of the pre-trends by year. Neither visual inspection of the event study estimates nor the estimated slopes for the pre-shock periods suggest that changing data sources consistently reduces pre-trends in this type of analysis. In fact, pre-trend slopes using BDS data are further from zero for many shock years. This pattern suggests that the non-zero pre-trends we estimate are unlikely to be the result of any difficulty CBP data may have with timely, accurate measurement of establishment or firm births or deaths in this context.

Figure 10 presents estimates using CBP data aggregated to the state level. These estimates show much more balanced variation across years in the pre-shock coefficients (rather than consistently sloping up into the focal year as in the CZ case), and the pre-focal year trend in underlying shocks is consistently close to zero across focal years. Many years' post-shock event study coefficients continue to show employment responses to the shock, but the magnitudes are generally smaller than the corresponding CZ-based estimates.

One possible explanation for this pattern is that aggregating to the state level solves the allocation problem associated with assigning economic activity to CZs in the CBP data, which solves the pre-trend problem; the fact that the estimated employment effects shrink is incidental to correctly specifying the analysis. It is also possible, however, that in addition to sidestepping the need to allocate activity to CZs, aggregating to the state level sacrifices meaningful local variation in economic conditions, and as a result, all event study estimates are attenuated relative to the CZ case, which has the effect of making the pre-trends look better and the employment effects look smaller. Researchers should be mindful of the tradeoffs implicit in these competing explanations when determining the appropriate level of geographic aggregation to use in their analyses.

4.5 Comparisons with Other Shocks

This paper focuses on shocks that have played a noteworthy role in research on local labor market adjustment and that can be standardized in terms of time period and geography in order to facilitate the comparisons we make. A version of our approach could also be used to make partial comparisons between our focal shocks and other important shocks from recent decades that are not as consistently available. Mass layoffs, for example, are commonly studied in the labor literature, but official estimates of mass layoffs were discontinued in 2013 and were not produced for sub-state geographies.

The shift-share approach to constructing local labor demand shocks has also been used

to quantify the effects of the signing of the North American Free Trade Agreement (NAFTA) and to robotic automation. We constructed CZ-level shocks derived from NAFTA exposure following Benguria (forthcoming) and McLaren and Hakobyan (2016). These shocks are focused on 1995 because the majority of import value affected by NAFTA faced tariff schedule changes effective that year.[21] Because the data we use to estimate migration effects discussed in the next section are not consistently available until after 1995, we do not include a NAFTA construct in our above analysis. However, in results available upon request we find that NAFTA constructs are somewhat correlated with the import- and export-related shocks analyzed above, with ρ-values of 0.18 and -0.16, respectively. The NAFTA construct is also correlated with employment growth shocks in the second half of the 1990s. This relationship is about as strong as that between export-related and employment growth constructs.

Robot exposure has been studied by Acemoglu and Restrepo (2020) and Faber et al. (2022), with both studies finding that exposure reduces local employment and population. Both studies rely on data on exposure to robotic automation from European industries, matched to US employment shares by industry. The European data are available at seven-year intervals over our time period, making it difficult to define as either an annual panel or a cross-sectional shock, as we have done elsewhere in this paper.[22] For this reason, we have not developed the correlation and distributional analysis like that reported above for robot exposure.

5 Migration Responses to Local Shocks

We have shown that local employment changes and their dynamics differ greatly depending on the demand shock construct under study. We have also shown that these constructs differ along other dimensions: the shapes of their distributions, geographic incidence, and the share of variance contributed by a small number of areas. The evidence also shows that these constructs each contain unique information. Although there is some correlation across construct measures, their correlations are well below one.

In this section, we consider whether the variation in employment impacts extends to other outcomes, in particular, to migration. We construct migration flows from individual-level data on the universe of tax filers merged to Census demographic information, as described in Section 2. These allow us to construct commuting-zone-level flows with a high degree of precision and enable comparisons across types of places and people. These in turn shed light on the important moderators of migration responses to shocks.

[21] Authors' calculations based on UN COMTRADE data and a concordance from Romalis (2007)
[22] Also, data measuring robot prevalence in the US by industry are not available until 2004.

5.1 A Statistical Model of Outflows and Inflows

We begin by outlining the statistical model that we use to measure the effect of a local shock on the probability that an individual leaves a commuting zone. Then, we aggregate the individual-level specification to the commuting zone level to measure the effect of a local shock on the share of the population migrating out of a commuting zone. The approach generalizes from outflows to inflows.

Let m_{it}^{out} be a dummy that equals 1 if person i migrated out of their commuting zone in year t,

$$m_{it}^{out} = \begin{cases} 1 & \text{if } j(i,t) \neq j(i,t-1) \\ 0 & \text{if } j(i,t) = j(i,t-1) \end{cases}$$

where $j(i,t)$ now denotes commuting zones, namely the commuting zone person i lives in at time t. A linear probability model can be used to estimate the effect (β_1) of a local shock $shock_{j(i,t),t}$ on whether a person migrates:

$$m_{it}^{out} = \beta_0 + \beta_1 shock_{j(i,t),t} + \Theta_{j(i,t)} + \Theta_t + \epsilon_{ijt} \tag{6}$$

Assuming $shock_{j(i,t),t}$ is exogenous to period t location choices, its coefficient measures the effect of a local shock on whether an individual migrates out of location j.[23] The CZ fixed effect $\Theta_{j(i,t)}$ controls for differences in migration propensity across places over the sample period. The year fixed effect Θ_t controls for differences in migration propensity across years.

To aggregate Equation (6), we take the expectation of both sides conditional on place and time to generate annual CZ-level outflows:

$$\mathbb{E}\left[m_{it}^{out}|j,t\right] = \mathbb{E}\left[\beta_0 + \beta_1 shock_{j(i,t),t} + \Theta_{j(i,t)} + \Theta_t + \epsilon_{ijt}|j,t\right]. \tag{7}$$

The left-hand side equals the out-migration rate, $\frac{m_{jt}^{out}}{n_{jt}}$, which is the share of people that migrate out from place j at time t. On the right-hand side, the conditional expectations of the constant β_0, $shock_{j(i,t)}$, and the fixed effects are all equal to themselves. Hence, Equation (7) simplifies to

$$\frac{m_{jt}^{out}}{n_{jt}} = \beta_0 + \beta_1 shock_{jt} + \beta_2 X_{jt} + \Theta_j + \Theta_t + \epsilon_{jt} \tag{8}$$

where the error term $\epsilon_{jt} \equiv \mathbb{E}[\epsilon_{ijt}|j,t]$ has an expectation equal to zero by assumption of shock

[23] We adjust our observation year to account for the fact that we observe a tax filer's location in the calendar year after their earnings year. Left- and right-side variables are constructed such that the time period $(t-1,t)$ refers to the same calendar years for all variables. For details, see the Data Appendix.

exogeneity and the law of iterated expectations.

The key independent variable is again the local shock $shock_{jt}$. After aggregating, β_1 measures the percentage point change in the outflow rate following a local shock. Controlling for place-level (Θ_j) and year-level (Θ_t) fixed effects allows migration rates to be higher or lower either across locations permanently over time or for all places from year to year. Allowing permanently higher or lower migration rates across locations via location fixed effects may seem surprising, since, if these exist, they imply that migration fails to reach equilibrium over a long horizon. The alternative of excluding location fixed effects implies that any deviations from the national migration rate in a year is due entirely to local shocks. Given evidence on very long-run adjustment patterns to local and cyclical shocks, we view including local fixed effects as the better approach (Amior and Manning, 2018; Hershbein and Stuart, 2024; Cajner et al., 2021).

We expand Equation (8) in two ways to reach our main estimating equation for migration flows. First, we control for shocks to other commuting zones following Borusyak et al. (2022a). This controls for changes in the value of other locations as outside options, since people living in commuting zone j are less likely to migrate if the potential places they could migrate to experience adverse shocks at the same time that their origin location experiences a negative shock. We add a weighted average of shocks to other commuting zones, where the weights are constructed with migration flows from before the shocks we consider were realized. Specifically, for commuting zone j, the weight on commuting zone j', denoted $\omega_{j,j'}$, is the percentage of out-migrants from j that moved to j' between 1989 and 1994. Second, we control for lags of both the shocks to each place and the weighted average of the shocks to other places. With these additional controls, Equation (8) becomes

$$\frac{m_{jt}^{out}}{n_{jt}} = \beta_0 + \beta_{1h}\ shock_{j,t-h} + \beta_{2h}\left(\sum_{j' \neq j} \omega_{j,j'} shock_{j',t-h}\right) + \Theta_t + \Theta_j + \epsilon_{jt} \tag{9}$$

To estimate a comparable version of the above equation for cross-sectional shock constructs, we modify (9) as follows:

$$\frac{m_{jt}^{out}}{n_{jt}} = \beta_0 + \beta_1\ post * shock_j + \beta_2\ post * \left(\sum_{j' \neq j} \omega_{j,j'} shock_{j'}\right) + \Theta_t + \Theta_j + \epsilon_{jt} \tag{10}$$

We estimate versions of Equation (9) and Equation (10) using both one-year and five-year changes ($h = 1, 5$) in outflows, i.e., migration rates, as the dependent variable.

5.2 Effects of Local Shocks on Outflows and Inflows

We report results from estimating (9) with one- and five-year out-migration rates in Table 8. The estimating equation includes the contemporaneous shock value and a one-year lag, as well as weighted averages of shocks to other CZs. The shock constructs are normalized to z-scores oriented such that an increase in the measured shock reflects greater labor demand. The first set of columns report the impacts of a one standard deviation improvement in a CZ's labor demand on one-year migration outflows using the panel shock constructs. The top row shows that only the employment growth construct has a significant impact on outflows, and at about 2.9 percent of the mean migration rate, this is economically modest. The previous year's employment growth shock has an even smaller effect, and we find no significant role for employment growth shocks in other CZs on outflows. Trade-related shocks to the origin CZ have no discernible effects on outflows at the one-year horizon. Impacts of these shocks on outflows are small, statistically insignificant, and precisely estimated. We can reject effects that are larger than 0.4 percent of the mean. For import-related shocks, trade shocks to other destination CZs are not significant at conventional levels, but export-related shocks in destination CZs (counterintuitively) reduce out-migration from the origin by about 2 percent.

The right three columns show impacts of the same one-year shocks over the medium-run—a five-year horizon. Results show that cumulative five-year out-migration is not affected by any of the panel shocks, either to the origin CZ or in the destination CZs. Appendix Table C10 reports estimates after including five years of lagged shocks. Including further lags does not make much difference for contemporaneous shock effects, and the lags after two years are never more than marginally statistically significant.

Results using inflows as the dependent variable of interest are reported in Table 9. The first three columns show that, as in the previous table, only employment growth shocks have a significant impact on flows. Impacts of trade-related shocks are again small, insignificant, and precisely estimated. However, employment growth shocks lead to increases in inflows that are about 5 percent of mean one-year migration rates, an effect that is twice as large as that on outflows. Unlike their effects on outflows, the effects of employment growth shocks on inflows persist into the medium-run. Five-year inflow rates remain about 3 percent above the mean following a one standard deviation increase in employment growth at the baseline year. Impacts of trade-related shocks on inflows are insignificant at the one- and five-year horizons. Shocks to other CZs have effects on in-migration that are not generally significant at conventional levels, though the sign on the point estimates is counterintuitive.[24] Appendix

[24]Direct comparisons with the effects of economic conditions outside the origin in Borusyak et al. (2022a) are not possible as that paper uses data from Brazil and does not report comparable specifications. The counterintuitive signs in our analysis may be related to overall procyclicality in migration, as demonstrated

Table C11 shows the in-migration results for the five-year lag specification.

Estimates from Equation (10) using cross-sectional shock constructs are reported in Table 10. Effects on outflows are shown in the top panel; effects on inflows are in the bottom panel. As was the case with the trade-related panel shocks, the tariff gap construct has no impact on outflows or inflows, at either the one- or five-year horizons. The Great Recession shocks have effects on inflows and outflows that are similar to those for the panel shocks to employment growth. A one standard deviation shock using the Great Recession construct increases inflows to a CZ by 0.20 percentage points over a one-year horizon; the estimate for panel shocks to employment growth is 0.19. At the five-year horizon, the effects of the Great Recession shock on inflows is still statistically significant but modestly smaller than the analogous effects of employment growth shocks. Strikingly, effects of the Great Recession shocks on migrant inflow and outflows parallel those of employment growth shocks closely even though the former are all negative overall. That is, an increase in local unemployment that is smaller by one standard deviation than increases elsewhere has the same effect on inflows as a one standard deviation increase in local employment growth.

5.3 Do Place and Person Characteristics Matter? Heterogeneity in the Migration Response to Shocks

We have shown above that shocks originating from different sources do not lead to equivalent changes in local employment, nor do they lead to similar changes in local population flows. One possible reason for this is that shocks may hit different types of places, and place characteristics may affect the response of migration flows. We explore this by allowing migration flow responses to differ with local characteristics in order to assess whether the impact of shocks differs with place factors.

We define a set of fixed local characteristics that may be linked with the spatial distribution of shocks or with responses of the local population flows to shocks. For succinctness, we call these moderators. This part of our analysis is descriptive in nature and only shows whether effects differ in the presence of the moderators we analyze.[25] We investigate eight moderators for which, based on the literature, we have a prior that the local characteristic seems likely to affect the local response following a shock to labor demand. These are: the local employment rate, the share of the population born locally (in-state), the share of college graduates, a measure of local college access, the share of low-income local children moving

in Molloy and Wozniak (2011) and Dao et al. (2017).

[25]In a setting with richer causal channels, some of these moderators may in fact be part of the causal process and therefore may more correctly be called "mediators." We do not attempt to determine which factors are moderators only and which are mediators.

into the top quintile of earnings as adults (an indicator of economic mobility), a January temperature indicator, local house prices, and market size. We describe these in more detail in Tables C12 - C15, which define each moderator, note its potential impact on migration flows, and document its origins in the literature. All are defined in a fixed based period, as noted in the tables.

We extend our migration flows regressions, equations (9) and (10), to include interactions of the shock constructs with a moderator and estimate these using our panel data set of migration flows. The full specifications are shown below. Equation (11) is used to estimate the effects of panel shocks and (12) is used to estimate the effects of cross-sectional shocks:

$$\frac{m_{jt}^{out}}{n_{jt}} = \beta_0 + \sum_{h=0}^{1} \beta_{1h}\ shock_{j,t-h} + \sum_{h=0}^{1} \beta_{2h} \left(\sum_{j' \neq j} \omega_{j,j'} shock_{j',t-h} \right) + \sum_{h=0}^{1} \beta_{3h} X_j shock_{j,t-h} + \Theta_t + \Theta_j + \epsilon_{jt} \tag{11}$$

$$\frac{m_{jt}^{out}}{n_{jt}} = \beta_1 \left(post_t \times shock_j \right) + \beta_2 \left(post_t \times X_j \right) + \beta_3 \left(post_t \times shockj \times X_j \right) + \Theta_t + \Theta_j + \epsilon_{jt} \tag{12}$$

Results from estimating the above two equations for outflows and inflows are reported in Appendix Tables C13 through C14. The tables report the main effects of all shocks along with coefficients on interactions of the local shocks with our moderators. Recall that our baseline estimates of migration flows in Tables 8 and 9 showed that inflows were much more responsive than outflows, and we found no evidence of migration responses to import- or export-related local labor demand shocks. Estimates of the expanded migration flow specifications in Appendix Tables C13 and C14 follow similar patterns. These estimates also show no statistically meaningful heterogeneity across places in these flows following trade-related shocks. Instead, detectable migration responses to local demand shocks—and place-level heterogeneity—are limited to the employment growth construct and the Great Recession severity construct. Inflows are similarly more responsive than outflows to local shocks, even after allowing responses to differ with values for the moderators.

We focus on the employment growth and Great Recession constructs, as these are the only two constructs that lead to migration responses. These two shocks generate employment impacts and inflow responses that are broadly similar. Does this similarity extend

to the dimensions of heterogeneity that we explore? For ease of interpretation, we visually summarize the key results for our moderator analysis to answer this question. These are reported in Figure 12.

Figure 12a reports estimates of β_{3h} in (11) using the employment growth construct, and Figure 12b reports β_{3h} in (12) using the Great Recession shock construct. Both report impacts on inflows. The figures show that higher January temperatures enhance inflow responsiveness under both shock constructs, while places with a local college and higher shares of local-born residents experienced less inflow responsiveness. Some determinants of inflow responsiveness differed across the constructs. Having a greater share of college graduates in the population enhanced inflow responsiveness to employment growth shocks but not to Great Recession shocks. Initial high employment growth areas saw lower inflows than other CZs following the Great Recession shocks but higher inflows following employment growth shocks.[26]

Results for analogous specifications using outflows are reported in Appendix Tables C13 and C15, and coefficients from the outflows regressions following employment growth shocks and the Great Recession construct in our baseline migration regressions were generally muted compared to inflow responses or statistically insignificant. Similarly, our analysis of place-level heterogeneity shows little meaningful heterogeneity in outflow responses. Places with more income mobility saw somewhat greater outflows in response to the Great Recession shock, but otherwise outflows are unresponsive to shocks across all types of places that we examine.

We also examine heterogeneity in migration flow responses by person-level characteristics. To do this, we create inflows and outflows for groups that divide the population in our migration sample along several dimensions. For example, to study heterogeneity in migration responses among foreign-born and US-born individuals, we divide the sample described in Section 2 into foreign-born and not-foreign-born and then construct annual CZ-level migration inflows and outflows for the two separate populations. We then report results from estimating our baseline migration equations on these separate populations in Appendix Tables C16 - C18.

We again summarize results by subgroups by plotting the group level $\beta_1 h$ coefficients from Equation (9) estimated on inflows in Figure 13. The figure plots responses to the employment growth and Great Recession constructs as these are again the only constructs where meaningful migration responses and heterogeneity follow. Figure 13a shows that inflows of younger (aged 25-35) and foreign-born individuals were more responsive to employment

[26]In unreported results, we estimated specifications that included interactions of local shocks and all moderators simultaneously. Our main conclusions are robust to this alternative approach.

growth shocks. Inflows for these groups are about twice as responsive as those for somewhat older (aged 45-55) and for US-born individuals, respectively, though the difference is not significant at conventional levels. Appendix Table C16 shows that outflows of foreign-born individuals are also more responsive following employment growth shocks. The coefficient is similar to that for non-white individuals, possibly because of a large degree of overlap in these groups. Outflows for other groups are not responsive to employment growth shocks—consistent with results in our baseline migration regressions—and outflow responsiveness among the foreign-born is about half that of inflows.

Person-level heterogeneity in inflows following the Great Recession shock is shown in Figure 13b. Similar to our findings above, inflows of US-born were much less responsive than those of the foreign-born. On other dimensions, however, responses to the Great Recession shock differ from those following employment growth shocks. Inflows to less affected CZs were greater among lower-earning (low AGI) individuals and those who were native to their origin CZ. The age gradient in inflow responsiveness also reversed following the Great Recession shock. These patterns may reflect the different nature or salience of the two shock constructs, but we also urge some caution in interpreting the results in this panel, as many of the underlying coefficients appear to be wrong-signed at the one-year horizon.

Our analysis of migration responses to the various shock constructs yields several findings. Using our large data set of regular tax filers, we affirm results from the recent literature that finds no medium-run migration responses to import-related local demand shocks. We find that this applies to export-related trade shocks as well. For shock constructs that do generate a migration response (employment growth and Great Recession severity), we find that inflows are much more responsive than outflows. However, both place- and person-level heterogeneity in migration responses differs across these constructs. Taken together, our analysis of migration flows implies that there is no single pattern of responses that follows all shock constructs, even after aligning definitions of geography, time interval, and base periods as in our framework.

6 Conclusion

A great deal of research has examined how local labor markets adjust to various labor demand shocks, but very little work has considered whether, how, or why the effects of different local labor demand shocks might be different. In this paper, we consider several constructs of local labor demand shocks that have figured prominently in the prior literature, standardize the construction of these shocks to make them directly comparable to each other, characterize the distribution of shocks across place and time, and compare shocks' employment and migration

effects estimated within our consistent framework.

Despite the tendency of some researchers and policymakers to view various local labor demand shocks as analytically interchangeable, our results suggest that there are important differences between the shocks we consider. Shocks differ meaningfully not only in their geographic concentration and intensity, but also in their relationships to employment growth. They are very much not interchangeable.

As a consequence, policymakers looking to research to inform their response to unfolding or anticipated local labor demand shocks should focus their attention on studies of prior shocks that resemble the situations they face as closely as possible rather than trying to draw lessons from aggregations of studies of shocks that differ substantially. A policymaker who faces a negative local labor demand shock driven by a national contraction in a locally important industry (which our analysis of employment growth shocks suggests would likely reduce local employment growth and in-migration while having more modest effects on out-migration) but formulates a response based on or heavily informed by research on import or export competition shocks (which our analysis suggests have little effect on employment growth or migration) would probably be making an error. On the other hand, a policymaker focused on analyses of employment growth shocks would be well prepared for the likely local consequences of such a shock.

Our results also reinforce lessons for researchers from the recent econometric literatures on causal analysis of employment growth shocks and difference-in-differences techniques. Our event study analysis reveals potentially concerning pre-trends for particular instances of employment growth shocks, as well as for the tariff gap shock. Researchers focused on making causal claims about the effects of these or similar local labor demand shocks should consider whether these concerns apply to their settings and use appropriate econometric tools to address them.

Finally, our event study analysis also reveals potentially meaningful differences in estimates based on whether shocks are constructed with commuting zones or states, with state-level construction seeming to mitigate some of the pre-trend concerns just described. Though state-level shocks may sacrifice potentially important local variation in labor demand shocks, they may also solve measurement problems that could tend to exaggerate differences in shock exposure across localities. Researchers should assess the degree to which local variation in the data used to construct local labor demand shocks is driven by measurement challenges as opposed to true differences in economic activity.

References

Acemoglu, Daron, David Autor, David Dorn, Gordon H Hanson, and Brendan Price (2016). "Import competition and the great US employment sag of the 2000s." *Journal of Labor Economics*, 34(S1), pp. S141–S198. doi:10.1086/682384.

Acemoglu, Daron and Pascual Restrepo (2020). "Robots and jobs: Evidence from U.S. labor markets." *Journal of Political Economy*, 128. doi:10.1086/705716.

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales (2019). "Shift-share designs: Theory and inference." *The Quarterly Journal of Economics*, 134(4), pp. 1949–2010. doi:10.1093/qje/qjz025.

Amior, Michael and Alan Manning (2018). "The persistence of local joblessness." *American Economic Review*, 108(7), pp. 1942–1970. doi:10.1257/aer.20160575.

Autor, David, David Dorn, and Gordon H Hanson (2013). "The China syndrome: Local labor market effects of import competition in the United States." *American Economic Review*, 103(6), pp. 2121–2168. doi:10.1257/aer.103.6.2121.

Bartik, Timothy J (1991). "Who benefits from state and local economic development policies?" doi:10.17848/9780585223940.

Benguria, Felipe (). "The impact of NAFTA on U.S. local labor market employment." *Journal of Human Resources, forthcoming*. doi:10.3368/jhr.0421-11621R3.

Bernstein, Jared (2024). "Remarks by CEA Chair Jared Bernstein at the Economic Club of New York." URL https://bidenwhitehouse.archives.gov/cea/written-materials/2024/04/30/chair-jared-bernsteins-remarks-at-the-economic-club-of-new-york/.

Black, Dan A., Terra McKinnish, and Seth Sanders (2005). "The economic impact of the coal boom and bust." *The Economic Journal*, 115, pp. 449–476. doi:10.1111/j.1468-0297.2005.00996.x.

Blanchard, Olivier Jean and Lawrence F. Katz (1992). "Regional evolutions." *Brookings Papers on Economic Activity*, 1992(1), pp. 1–75. doi:10.2307/2534556.

Borusyak, Kirill, Rafael Dix-Carneiro, and Brian Kovak (2022a). "Understanding migration responses to local shocks." *Available at SSRN 4086847*. doi:10.2139/ssrn.4086847.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022b). "Quasi-experimental shift-share research designs." *The Review of Economic Studies*, 89(1), pp. 181–213. doi:10.1093/restud/rdab030.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2025). "A practical guide to shift-share instruments." *Journal of Economic Perspectives*, 39, pp. 181–204. doi:10.1257/jep.20231370.

Cajner, Tomaz, John Coglianese, and Joshuah Montes (2021). "The long-lived cyclicality of the labor force participation rate." *Finance and Economics Discussion Series 2021-047*. doi:10.17016/FEDS.2021.047.

Card, David, Jesse Rothstein, and Moises Yi (2025). "Location, location, location." *American Economic Journal: Applied Economics*, 17(1), pp. 297–336. doi:10.1257/app.20220427.

CDC (2025). "North America Land Data Assimilation System (NLDAS) daily air temperatures and heat index (1979-2011) request: Version [dataset]." URL https://wonder.cdc.gov/nasa-nldas.html.

Chetty, Raj, Will Dobbie, Benjamin Goldman, Sonya R. Porter, and Crystal S. Yang (2024). "The opportunity atlas: Version [dataset]." URL https://www.opportunityatlas.org/.

Council of Economic Advisers (2022). *Economic Report of the President*. Executive Office of the President. URL https://bidenwhitehouse.archives.gov/cea/economic-report-of-the-president/.

Dao, Mai, Davide Furceri, and Prakash Loungani (2017). "Regional labor market adjustment in the United States: trend and cycle." *Review of Economics and Statistics*, 99(2), pp. 243–257. doi:10.1162/REST_a_00642.

Eckert, Fabian, Teresa C Fort, Peter K Schott, Natalie J Yang, et al. (2020). "County Business Patterns Database." doi:10.3886/E117464V1.

Faber, Marius, Andres Sarto, and Marco Tabellini (2022). "Local shocks and internal migration: The disparate effects of robots and Chinese imports in the U.S." *NBER Working Paper 30048*. doi:10.3386/w30048.

Feenstra, Robert C, Hong Ma, and Yuan Xu (2019). "US exports and employment." *Journal of International Economics*, 120, pp. 46–58. doi:10.1016/j.jinteco.2019.05.002.

Foote, Andrew, Michel Grosz, and Ann Huff Stevens (2018). "Locate your nearest exit: Mass layoffs and local labor market responses." *Industrial and Labor Relations Review*, 72. doi:10.1177/0019793917753095.

Foote, Andrew, Mark J Kutzbach, and Lars Vilhuber (2021). "Recalculating...: How uncertainty in local labour market definitions affects empirical findings." *Applied Economics*, 53(14), pp. 1598–1612. doi:10.1080/00036846.2020.1841083.

Fort, Teresa C, Shawn D Klimek, et al. (2016). "The effects of industry classification changes on US employment composition." *Tuck School at Dartmouth*.

Foschi, Andrea, Christopher L. House, Christian Proebsting, and Linda L. Tesar (2025). "Should I stay or should I go? The response of labor migration to economic shocks." *Brookings Papers on Economic Activity*, Spring. URL https://www.brookings.edu/articles/should-i-stay-or-should-i-go/.

Gallagher, Ryan, Robert Kaestner, and Joseph Persky (2019). "The geography of family differences and intergenerational mobility." *Journal of Economic Geography*, 19, pp. 589–618. doi:10.1093/jeg/lby026.

Gallin, Joshua Hojvat (2004). "Net migration and state labor market dynamics." *Journal of Labor Economics*, 22(1), pp. 1–21. doi:10.1086/380401.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020). "Bartik instruments: What, when, why, and how." *American Economic Review*, 110(8), pp. 2586–2624. doi:10.1257/aer.20181047.

Greenland, Andrew, John Lopresti, and Peter McHenry (2019). "Import competition and internal migration." *Review of Economics and Statistics*, 101(1), pp. 44–59. doi:10.1162/rest_a_00751.

Guenther, GaryPierce (2023). "Federal tax gap: Size, contributing factors, and the debate over reducing it." Technical report, Congressional Research Service. URL https://www.congress.gov/crs-product/IF11887.

Hershbein, Brad and Bryan A. Stuart (2024). "The evolution of local labor markets after recessions." *American Economic Journal: Applied Economics*, 16(3), pp. 399–435. doi:10.1257/app.20220132.

Jia, Ning, Raven Molloy, Christopher Smith, and Abigail Wozniak (2023). "The economics of internal migration: Advances and policy questions." *Journal of Economic Literature*, 61(1), pp. 144–80. doi:10.1257/jel.20211623.

Jorda, Oscar and Alan M. Taylor (2025). "Local projections." *Journal of Economic Literature*, 63, pp. 59–110.

McLaren, John and Shushanik Hakobyan (2016). "Looking for local labor market effects of NAFTA." *Review of Economics and Statistics*, 98, pp. 728–741. doi:10.1162/REST_a_00587.

Molloy, Raven, Spencer Perry, and Christopher L. Smith (2019). "Moving out to move up: Who leaves and where do they go?" *Pathways*.

Molloy, Raven and Abigail Wozniak (2011). "Labor reallocation over the business cycle: New evidence from internal migration." *Journal of Labor Economics*, 29, pp. 697–739. doi:10.1086/660772.

OECD (2019). *OECD Economic Survey of the United States: Key Research Findings*. OECD Publishing. URL https://doi.org/10.1787/9789264310278-en.

Pierce, Justin R and Peter K Schott (2012). "A concordance between ten-digit US harmonized system codes and [sic/naics] product classes and industries." *Journal of Economic and Social Measurement*, 37(1-2), pp. 61–96. doi:10.3233/JEM-2012-0351.

Pierce, Justin R. and Peter K. Schott (2016a). "Replication data for: The surprisingly swift decline of US manufacturing employment." doi:10.3886/E112965V1.

Pierce, Justin R and Peter K Schott (2016b). "The surprisingly swift decline of US manufacturing employment." *American Economic Review*, 106(7), pp. 1632–1662. doi:10.1257/aer.20131578.

Romalis, John (2007). "NAFTA and CUSTFA's impact on international trade." *Review of Economics and Statistics*, 89, pp. 416–435. doi:10.1162/rest.89.3.416.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek (2020). "IPUMS USA: version 10.0 [dataset]." *Minneapolis, MN: IPUMS*, 10, p. D010. doi:10.18128/D010.V10.0.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek (2022). "IPUMS USA: Version 12.0 [dataset]." doi:10.18128/D010.V12.0.

Schott, Peter K (2008). "The relative sophistication of Chinese exports." *Economic Policy*, 23(53), pp. 6–49. doi:10.1111/j.1468-0327.2007.00195.x.

Tolbert, Charles M and Molly Sizer (1996). "US commuting zones and labor market areas: A 1990 update." Technical report. doi:10.22004/ag.econ.278812.

Triplett, Jack E, D Mark Kennet, Ron Jarmin, and Frank M Gollop (1995). "Do industrial classifications need re-inventing? An analysis of the relevance of the US SIC system for productivity research." *Advances in Classification Research Online*, pp. 145–172. doi:10.7152/acro.v6i1.12667.

U.S.-China Economic and Security Review Commission (2025). *2025 Annual Report to Congress*. URL https://www.uscc.gov/annual-report/2025-annual-report-congress.

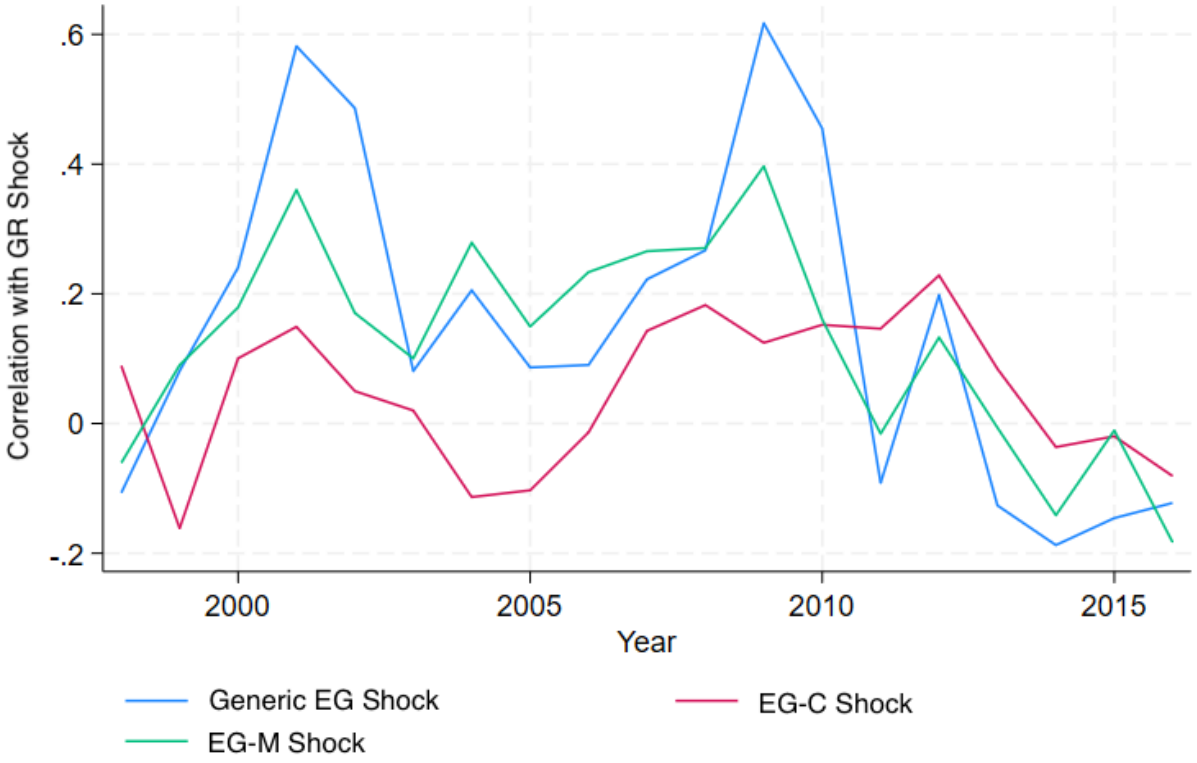
Wilson, Riley (2022). "Moving to economic opportunity: The migration response to the fracking boom." *Journal of Human Resources*, 57, pp. 918–955. doi:10.3368/jhr.57.3.0817-8989R2.

Winfield, Jake D. (2023). "Shifting sands: An analysis of college access deserts from 2001 to 2019." *SocArXiv*.

Wozniak, Abigail (2010). "Are college graduates more responsive to distant labor market opportunities?" *Journal of Human Resources*, 45(4), pp. 944–970. doi:10.3368/jhr.45.4.944.

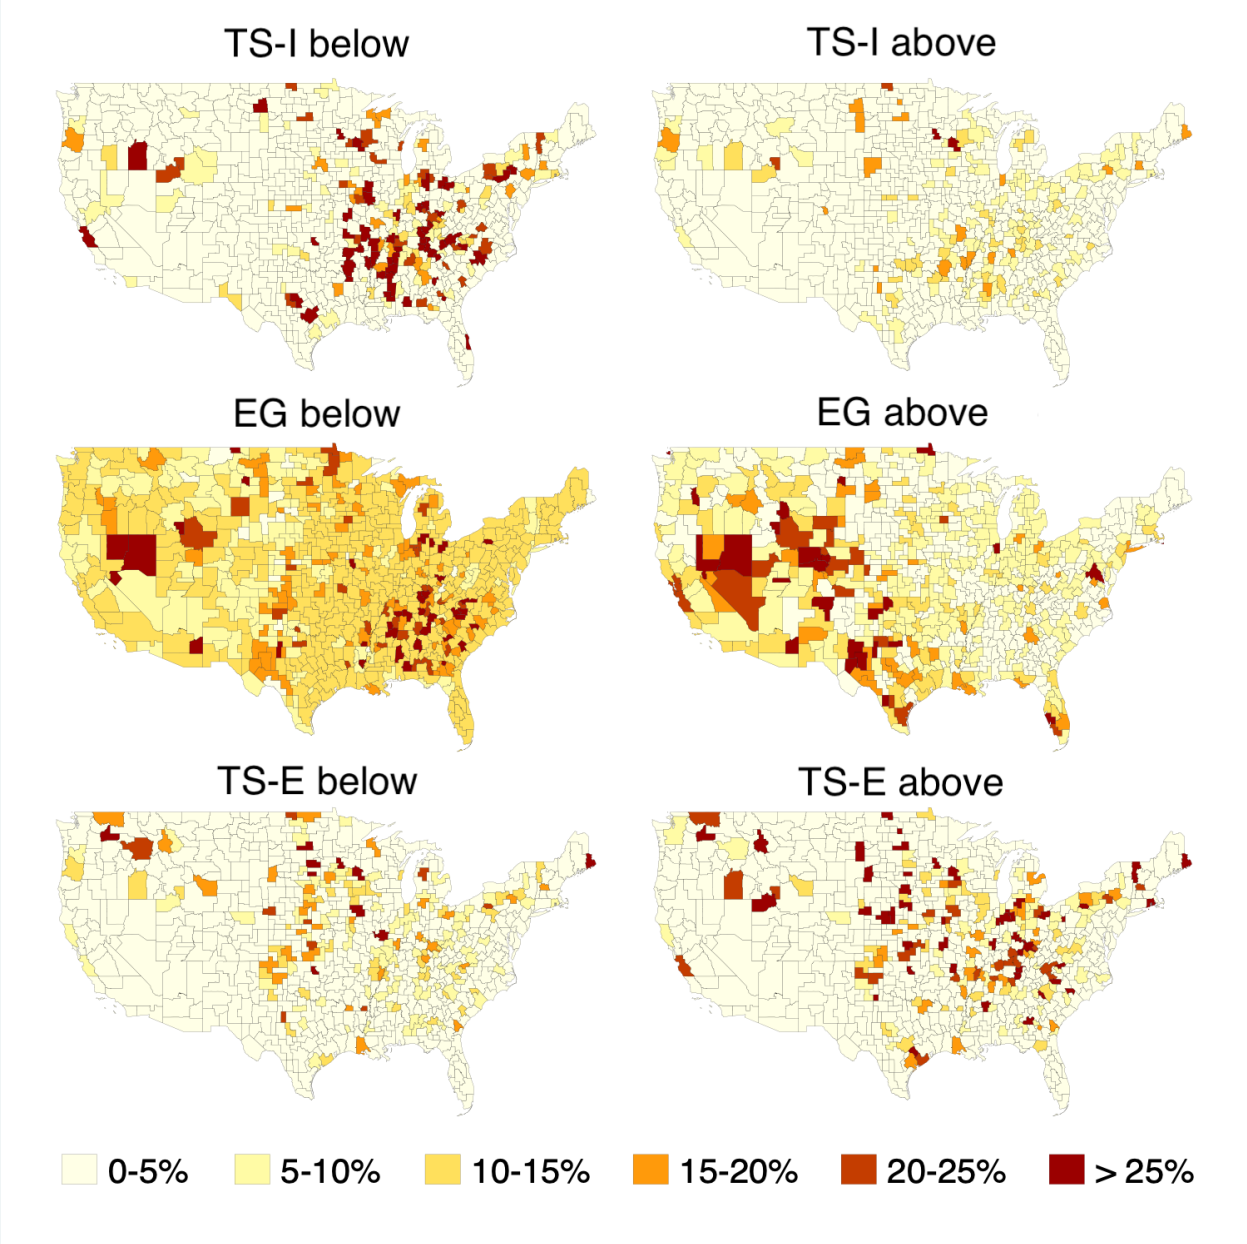
Yagan, Danny (2019). "Employment hysteresis from the Great Recession." *Journal of Political Economy*, 127(5), pp. 2505–2558. doi:10.1086/701809.

Figure 1: Employment growth and Great Recession Shock Correlations Across Time



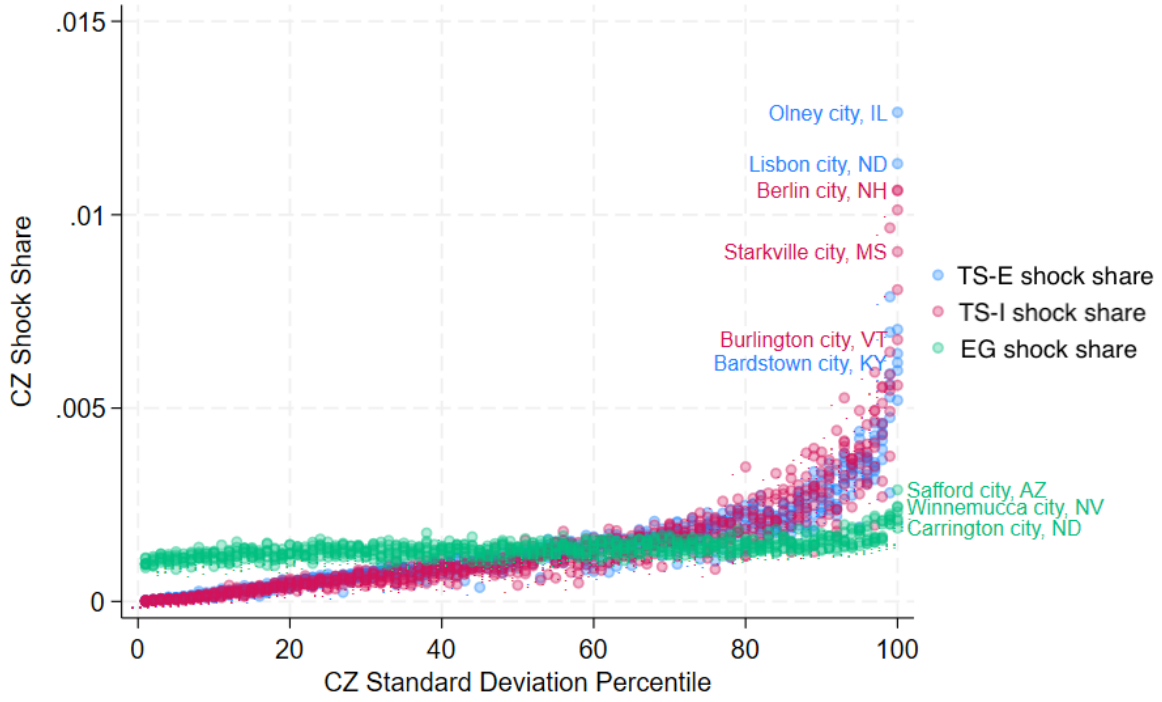
Notes: The figure shows the correlation between the standardized Great Recession shock and annual employment growth shocks between 1998 and 2016. Each annual employment growth shock is a one-year subset of the entire set of standardized employment growth shocks. As in Table 3, we also subset the employment growth shock to just construction and manufacturing industries. Underlying data are described in the Data Appendix.

Figure 2: Heat Map: Frequency of High-Intensity Shocks



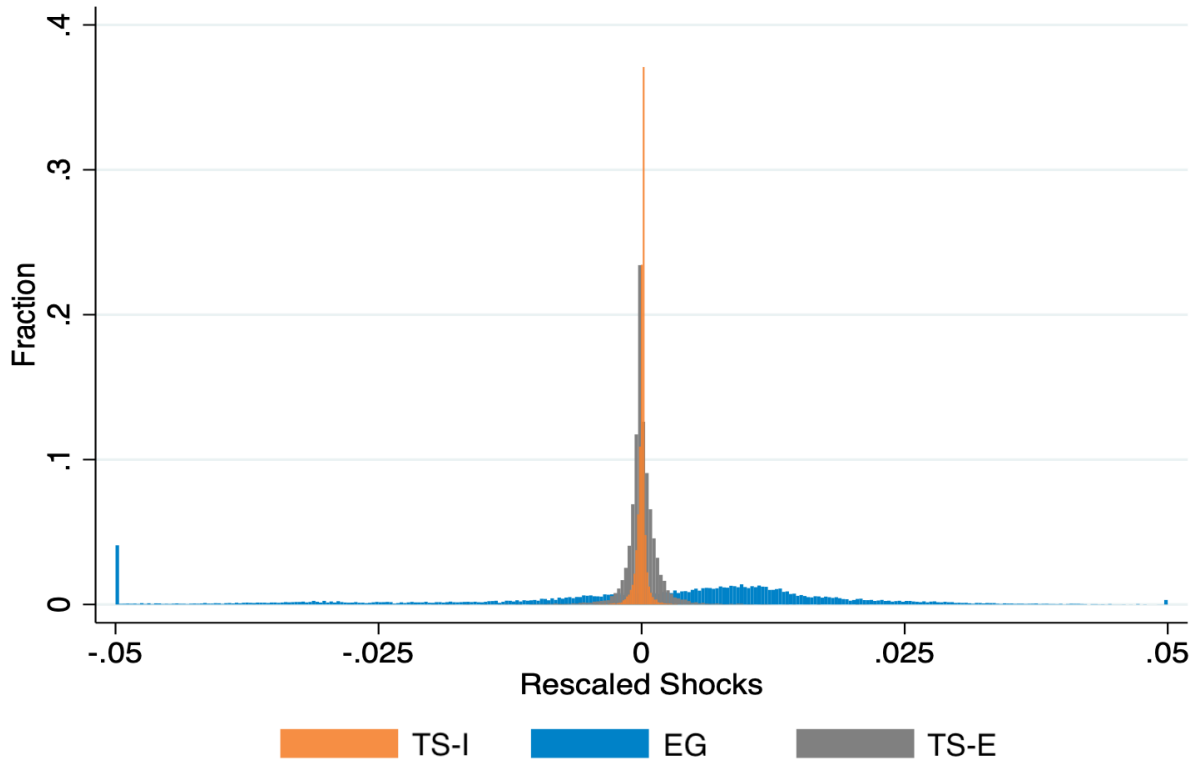
Notes: The figure shows the share of years between 1994-2016 that experienced shocks over one standard deviation below or above the mean value of shocks of the same type, which we call large shocks. The left column shows the share of years hit by large below-mean shocks, while the right column shows the share of years hit by large above-mean shocks. The three shocks pictured are the import-related trade shock (TS-I), employment growth shock (EG), and export-related trade shock (TS-E). Underlying data are described in the Data Appendix.

Figure 3: Shock Size vs Volatility



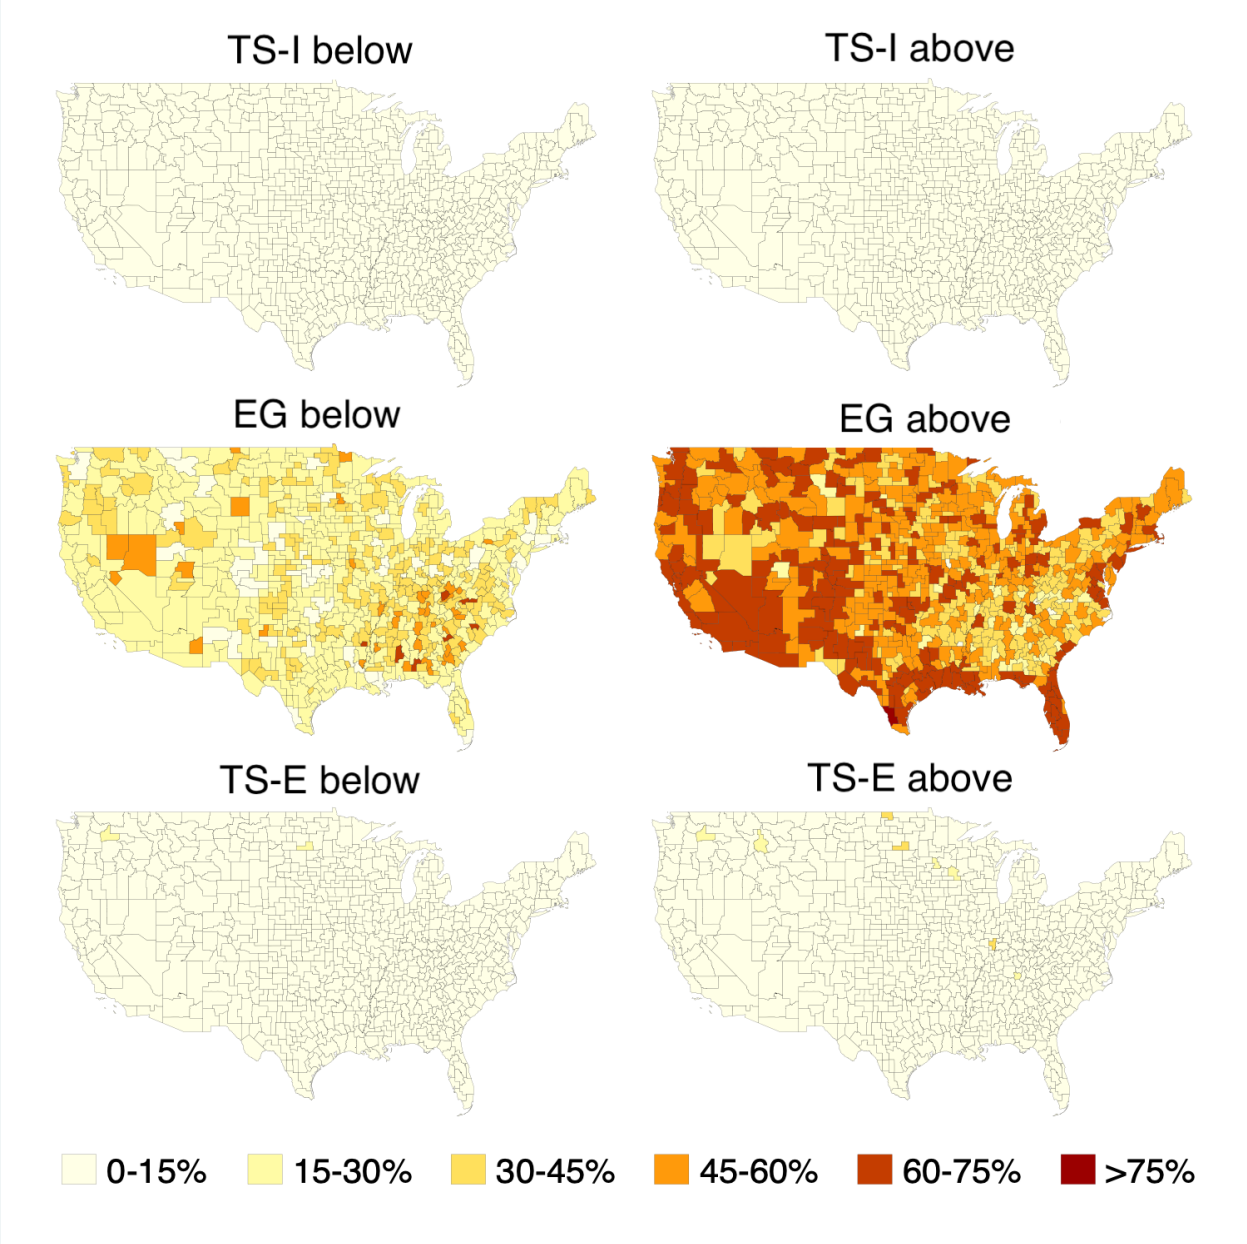
Notes: The Y-axis represents the sum of the absolute value of all shocks between 1994 and 2016 within a CZ over the sum of the absolute value of all shocks. The X-axis is the CZ's percentile ranking in terms of the shock's CZ-level standard deviation. For each shock, the CZs with the three largest standard deviations are labeled. While the TS-I.b shock is not represented on this figure, it closely matches the patterns of the TS-I shock share with the exception of the two largest shock-share observations, Berlin City, NH and Burlington City, VT. These CZs also have the largest standard deviation of the TS-I.b shock and are roughly three times larger than their TS-I shock-share observations.

Figure 4: Histogram of the Rescaled Shocks



Notes: The figure shows a histogram of the rescaled TS-I, EG, and TS-E shocks. We rescaled each shock using the coefficients from Table 5.

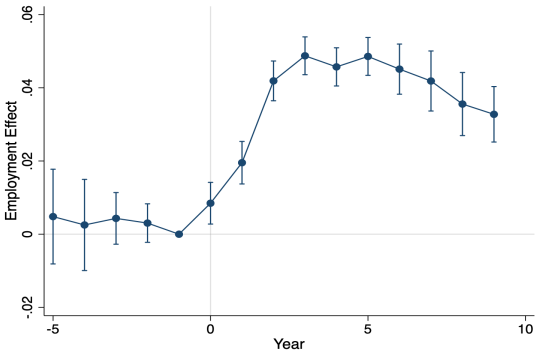
Figure 5: Heat Maps: Frequency of High-Intensity Rescaled Shocks



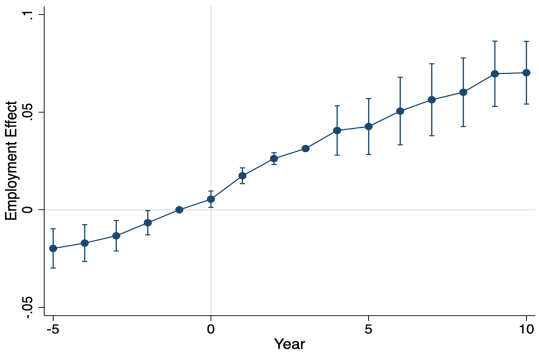
Notes: The figure shows the share of years between 1994-2016 that experienced high-intensity rescaled shocks by commuting zone. A high-intensity shock is a shock that has an effect above 0.5 percent or below -0.5 percent. The left column shows the share of years hit by shocks below -0.5 percent, while the right column shows the share of years hit by shocks above 0.5 percent. The three shocks pictured are the import-related trade shock (TS-I), employment growth shock (EG), and export-related trade shock (TS-E). Underlying data are described in the Data Appendix.

Figure 6: Comparing One-Time and Time-Varying Shocks

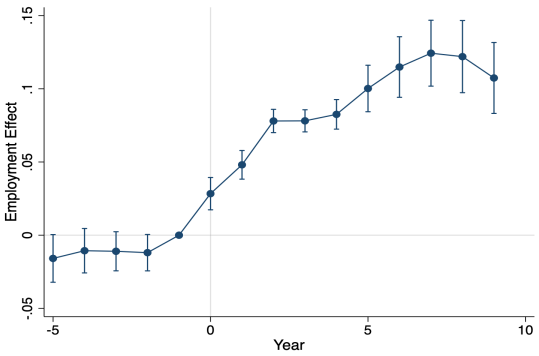
(a) 2007 GR Shock



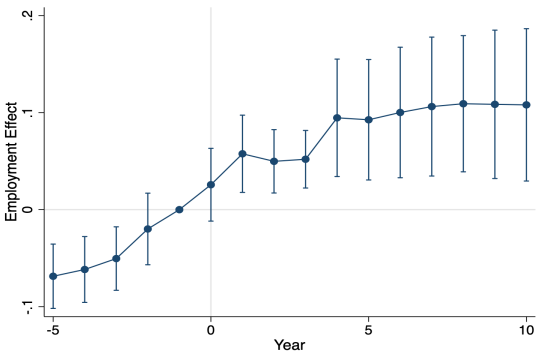
(b) 2000 TG Shock



(c) 2007 EG Shock

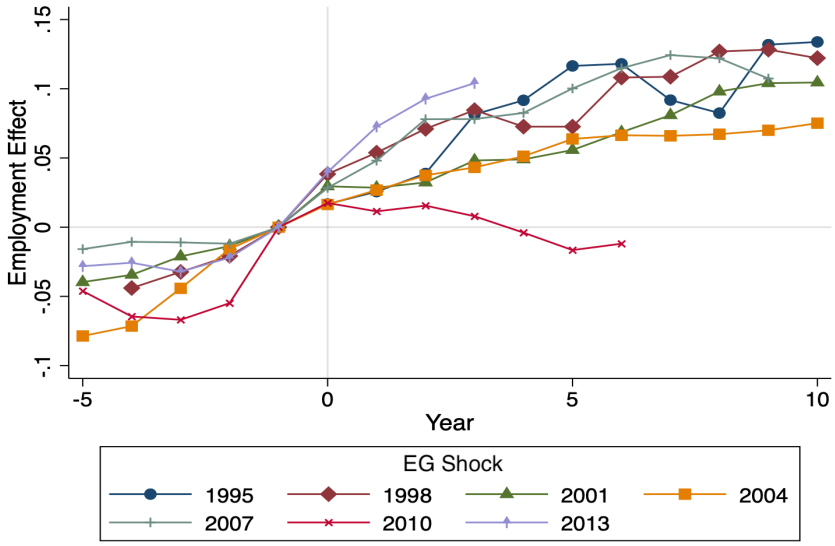


(d) 2000 EG Shock

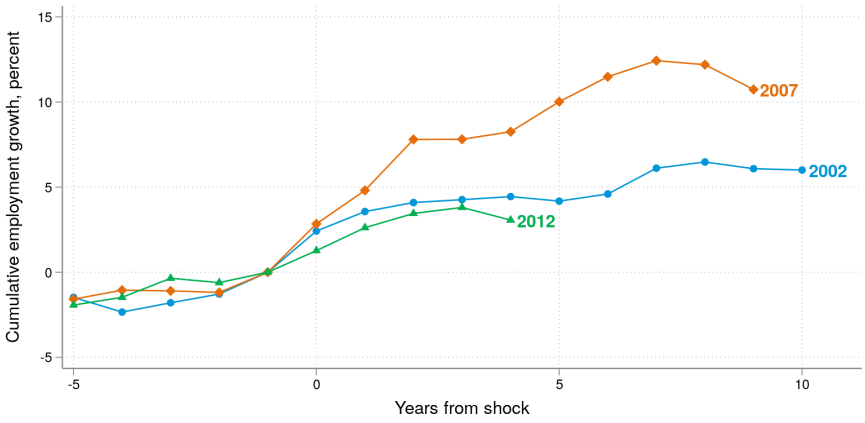


Notes: The figures plot the coefficients δ_{kh} from Equation (4), measuring the employment effect of a shock several years before and after it occurred. The x-axis is year before or since the shock occurred, h, with 0 being the year the shock occurred, k. The y-axis measures the coefficients δ_{kh}, and is interpreted as the effect of a one standard deviation increase in the shock in year k on employment growth between years $k-1$ and h. Shocks are converted to z-scores for comparability. The construction of the shocks and the employment data are discussed in the Data Appendix.

Figure 7: Employment growth Shock Event Studies



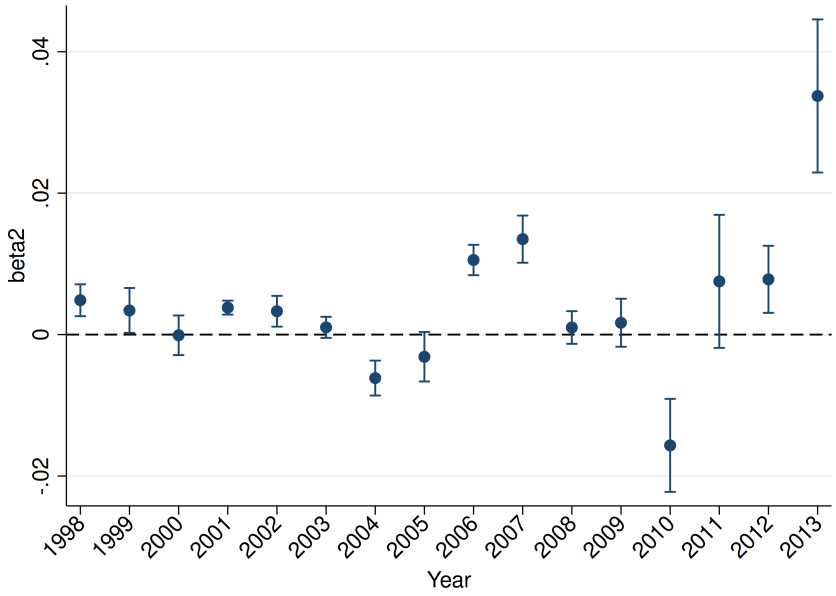
(a)



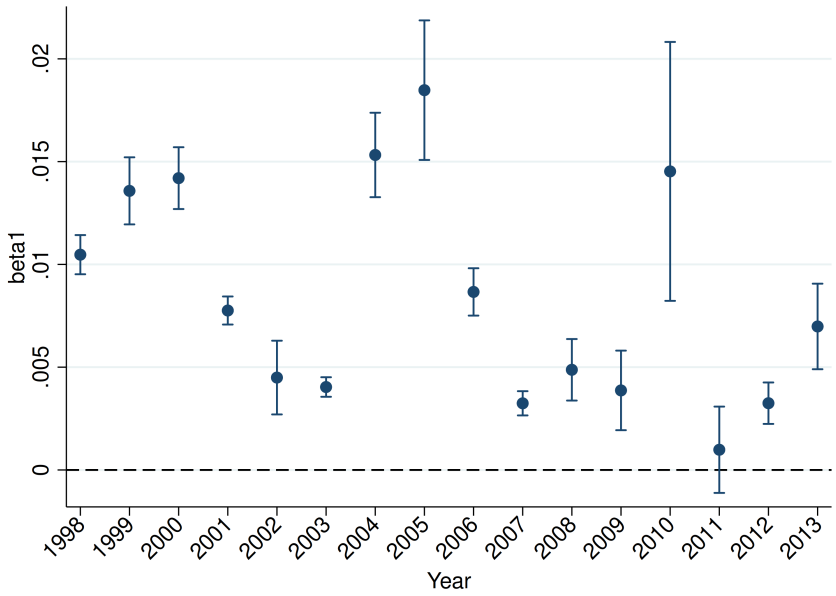
(b)

Notes: The figures show results from an event study regression for the employment growth shock. Panel A shows the results for every year in our data, while Panel B shows the result when we separate this analysis into Census years. Underlying employment data are the CZ-level data described in Section 2.

Figure 8: Pre-trend Estimates



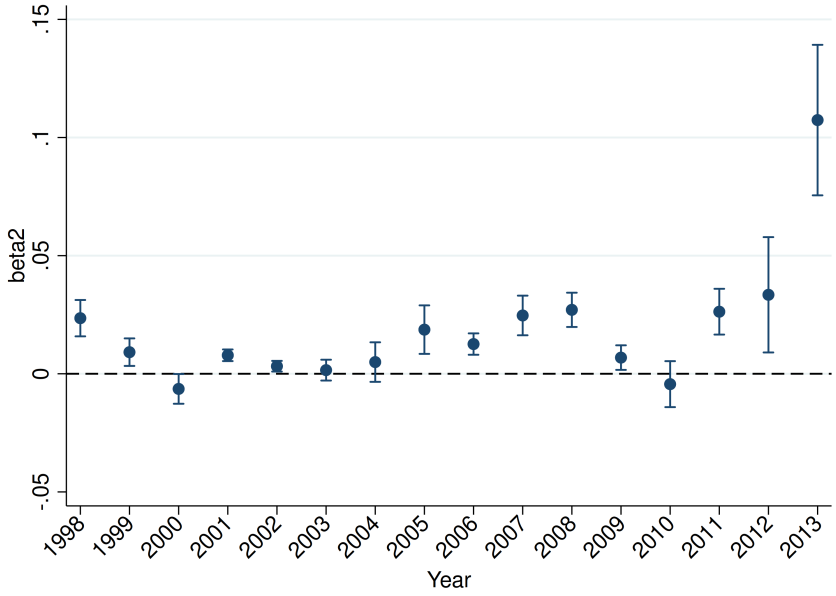
(a) β_2 Estimates: CZ-Level Data



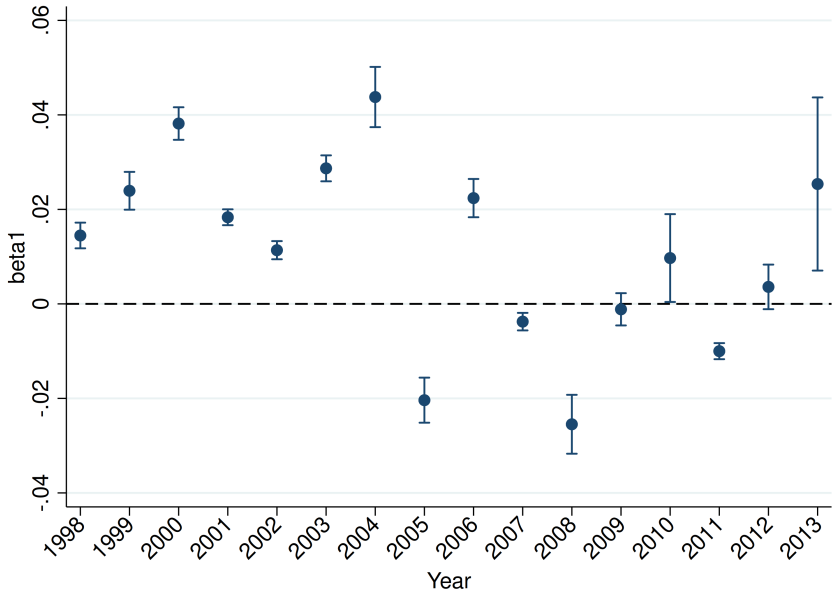
(b) β_1 Estimates: CZ-Level Data

Notes: We use employment shock event study estimates at the CZ level from the CBP data to estimate regessions capturing the significance of pre-trends in the data. We run a basic regression for each year for the employment shock that looks like: $EmploymentEffect \approx \beta_1(time) + \beta_2(time\,x\,post) + e$ where post takes a value of 1. Underlying data are described in Section 2.

Figure 9: Pre-trends Using BDS Data



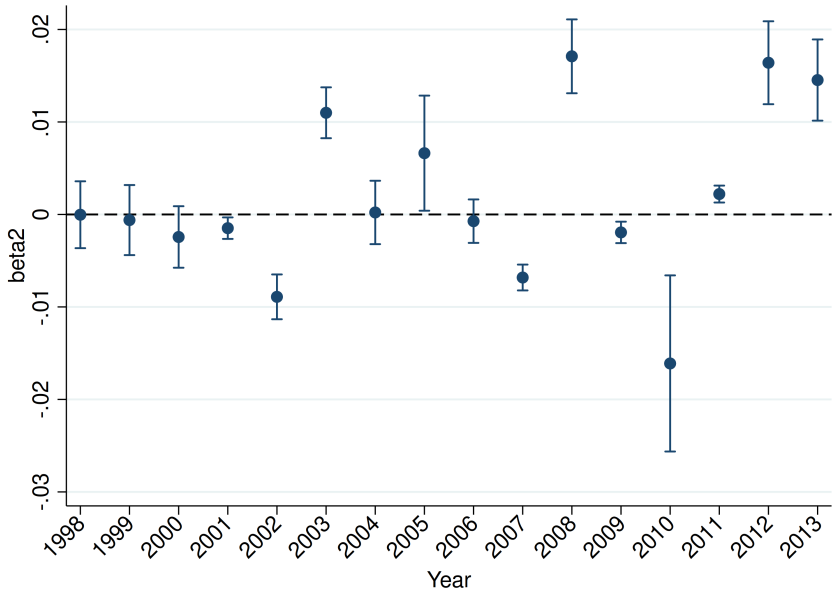
(a) β_2 Estimates using CZ-level BDS data



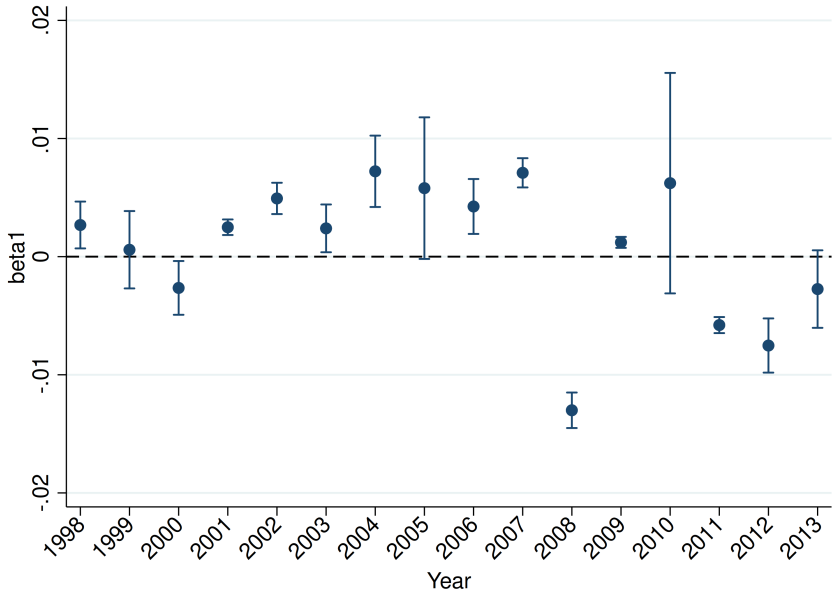
(b) β_1 Estimates using CZ-level BDS data

Notes: We use employment shock event study estimates at the CZ level from the BDS data to estimate regressions capturing the significance of pre-trends in the data. We run a basic regression for each year for the employment shock that looks like: $EmploymentEffect \approx \beta_1(time) + \beta_2(time\,x\,post) + e$ where post takes a value of 1. Underlying data are described in Section 2.

Figure 10: Pre-trends Using State-Level Data



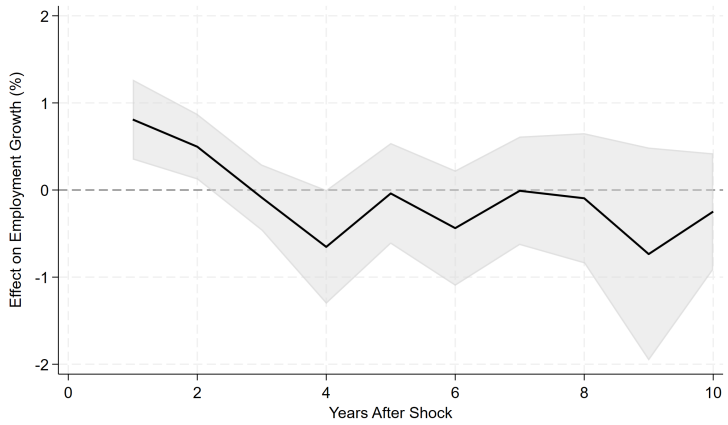
(a) β_2 Estimates Using State-Level CBP Data



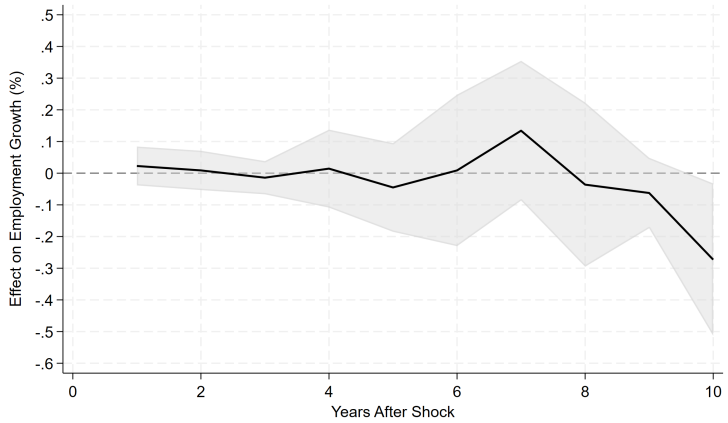
(b) β_1 Estimates Using State-Level CBP Data

Notes: We use employment shock event study estimates at the CZ level from the CBP data to estimate regressions capturing the significance of pre-trends in the data. We run a basic regression for each year for the employment shock that looks like: $EmploymentEffect \approx \beta_1(time) + \beta_2(time\,x\,post) + e$ where post takes on a value of one. Underlying data are described in Section 2.

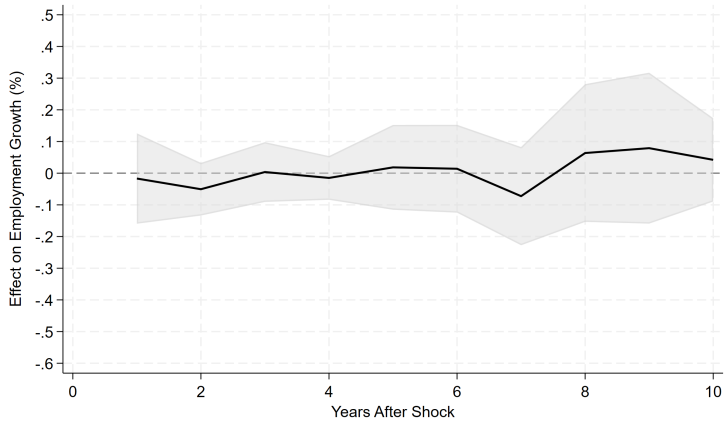
Figure 11: Local Projections Panel



(a) EG shock local projections



(b) TS-I shock local projections

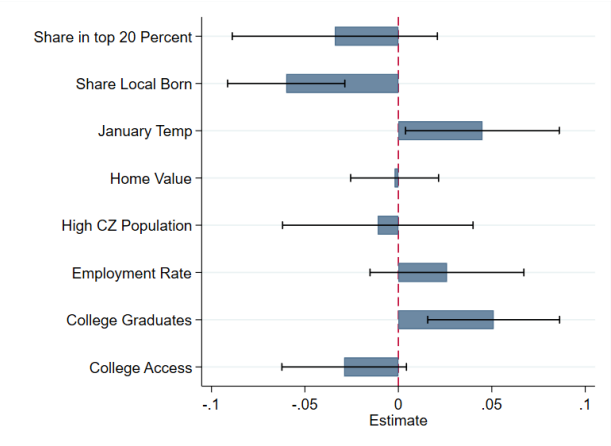


(c) TS-E shock local projections

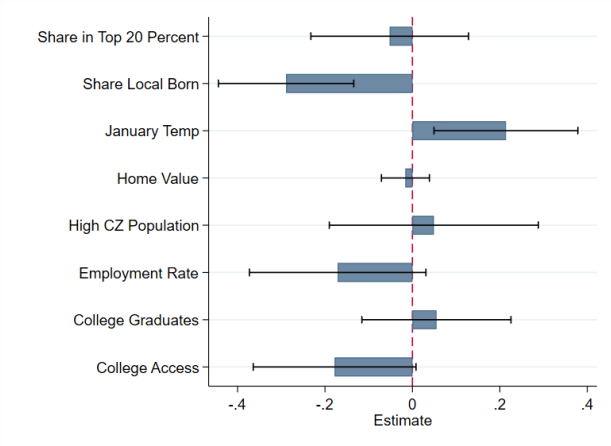
Notes: This figure shows the results for local projections for each time-variant shock (EG, TS-I, and TS-E). It uses a ten-year projection window and includes CZ and year fixed effects. We use the long difference approach specified by Jorda and Taylor (2025), although we include year and CZ fixed effects. Underlying data are described in Section 2.

Figure 12: Estimate of Shock Effects by Moderator

(a) EG Shock



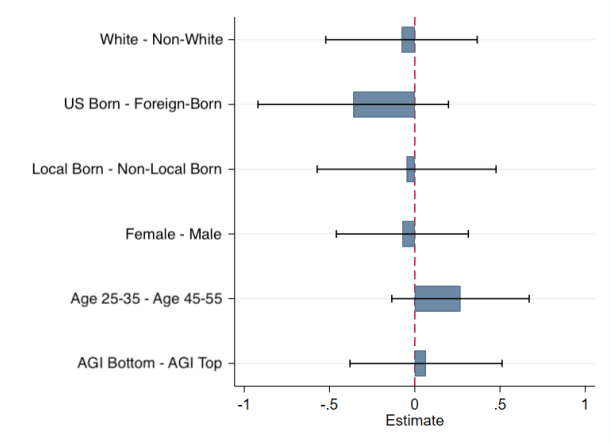
(b) GR Shock



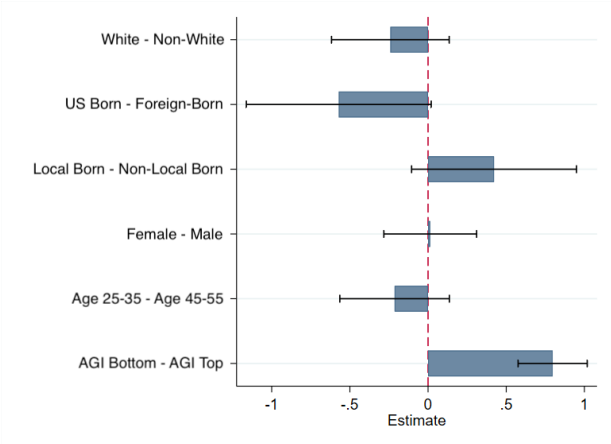
Notes: The figure summarizes CZ moderator interaction results of estimating specifications from Equation 10 using inflow rates as the dependent variable. Plot values and standard errors are described in panels B and E of Tables C13 and C15, respectively.

Figure 13: Estimate of Shock Effects by Subgroup

(a) EG Shock



(b) GR Shock



Notes: The figure summarizes the differences in results of estimating specifications from Equation 10 for population subgroups using inflow rates as the dependent variables. Plot values and standard errors are described in panels C and D of Table C16, respectively.

Table 1: Local Demand Shocks in the Harmonized Data Set

Panel Constructs

Source of variation	**Narrative definition**	**Introductory reference**	**Text (Exhibit) Label**
Flows of Chinese imports	Dollars of Chinese imports per worker	Autor et al. (2013)	Import-related trade shock (TS-I)
Flows of Chinese imports	Chinese import share of final demand	Acemoglu et al. (2016)	Import-related trade shock (TS-I.b)
Flow of US exports	Export share over total US shipments	Feenstra et al. (2019)	Export-related trade shock (TS-E)
US employment growth by industry	National industry employment growth weighted by local employment shares	Bartik (1991); Blanchard and Katz (1992)	Employment growth shock (EG)
US emp. growth in manufacturing	Manufacturing only variant of above	–	Employment growth shock (EG-M)
US emp. growth in construction	Construction only variant of above	–	Employment growth shock (EG-C)

Cross-sectional Constructs

Source of variation	**Narrative definition**	**Introductory reference**	**Text (Exhibit) Label**
Tariff rate changes upon China's 2000 accession to WTO	Industry tariff changes weighted by local employment	Pierce and Schott (2016b); Greenland et al. (2019)	Tariff gap shock (TG)
Great Recession	Excess local average unemployment rate in 2007-2009 over national average	Yagan (2019)	Great Recession shock (GR)

Notes: Trade absorption is defined as $\overline{Y_{jt}} + \overline{M_{jt}} - \overline{Ejt}$, the difference between the five-year trailing averages of trade shipments plus imports and exports.

Table 2: Descriptive Statistics of Normalized Shocks

Shock	N	min	p0.5	p10	p25	p50	p75	p90	p99.5	max	p90p10
TS-I	16721	−51.5	−3.9	−0.6	−0.2	0.1	0.3	0.4	3.0	22.6	1.1
TS-I.b	16721	−107.1	−1.6	−0.2	−0.1	0.1	0.1	0.2	0.8	6.5	0.4
TS-E	16721	−34.2	−4.0	−0.6	−0.2	−0.1	0.3	0.7	3.7	24.8	1.3
EG	16721	−7.2	−3.7	−1.4	−0.3	0.2	0.6	0.9	2.0	10.7	2.3
TG	727	−4.0	−3.5	−1.4	−0.6	0.1	0.7	1.1	1.9	2.0	2.5
GR	727	−3.9	−3.4	−1.4	−0.7	0.1	0.7	1.2	1.9	2.0	2.6

Notes: The table shows descriptive statistics for the annualized and time-invariant shocks. The annual shocks cover 1994 to 2016, while the tariff gap (TG) shock occurred in 2000 and the Great Recession (GR) shock in 2007. All commuting zones (1990 vintage) in the US are covered, except for Alaska, which we aggregate to the state level due to inconsistencies with its underlying borough classification, resulting in 727 commuting zones per year. All shocks are normalized to have means equal to zero and standard deviations equal to one over their entire sample periods. The shocks are also standardized so that negative shocks are associated with adverse shocks. The 'p90p10' column is the difference between the p90 and the p10. Underlying data are described in the Data Appendix.

Table 3: Correlation Matrix, Annual and Time-Invariant Shocks

Panel A: CZ-Level Correlations

Shock	TS-I	TS-I.b	TS-E	EG	EG-C	EG-M
TS-I	1					
TS-I.b	0.65	1				
TS-E	−0.33	−0.22	1			
EG	−0.08	−0.04	0.25	1		
EG-C	−0.10	−0.06	0.18	0.75	1	
EG-M	−0.06	−0.03	0.21	0.74	0.53	1
TG	0.52	0.66	0.35	0.75	0.16	0.46
GR	−0.03	−0.02	0.21	0.51	0.16	0.41

Panel B: State-Level Correlations

Shock	TS-I	TS-I.b	TS-E	EG	EG-C	EG-M
TS-I	1					
TS-I.b	0.88	1				
TS-E	−0.63	−0.55	1			
EG	−0.09	−0.03	0.36	1		
EG-C	−0.09	−0.04	0.27	0.84	1	
EG-M	−0.07	−0.02	0.33	0.88	0.70	1
TG	0.75	0.85	0.53	0.73	0.58	0.75
GR	0.12	0.19	0.30	0.61	0.31	0.55

Notes: The table displays the correlation matrix for the annual and time-invariant shocks. The correlations are calculated over the entire sample period, 1994-2016, for the annual shocks and over their respective years for the time-invariant shocks. We compare the TG shock to the 2000, 2001, and 2002 annual shocks. We compare the GR shock to 2008, 2009, and 2010 annual shocks. All shocks are converted into z-scores, and they are standardized so that negative shocks are associated with decreases in demand. Panel A displays correlations for shocks at the CZ level, while Panel B shows the correlations for state-level shocks.

Table 4: Autocorrelation of Annual Shocks

Panel A: With CZ and Year Fixed Effects

	(1) TS-I	(2) EG	(3) TS-E
1 Year Lag	0.0374	0.156***	-0.0146
	(0.0251)	(0.0190)	(0.0212)
2 Year Lag	0.0000104	0.0607***	-0.0987***
	(0.0212)	(0.0178)	(0.0254)
3 Year Lag	-0.0504**	-0.00437	-0.0526**
	(0.0230)	(0.0139)	(0.0205)
Observations	14540	14540	14540
R^2	0.210	0.793	0.246

Panel B: With Only Year Fixed Effects

	(1) TS-I	(2) EG	(3) TS-E
1 Year Lag	0.124***	0.250***	0.0181
	(0.0306)	(0.0193)	(0.0213)
2 Years Lag	0.0844***	0.146***	-0.0670***
	(0.0234)	(0.0177)	(0.0235)
3 Years Lag	0.0400*	0.0899***	-0.0190
	(0.0229)	(0.0133)	(0.0190)
Observations	14540	14540	14540
R^2	0.147	0.776	0.222

Panel C: With Only CZ Fixed Effects

	(1) TS-I	(2) EG	(3) TS-E
1 Year Lag	0.0351	0.350***	-0.0259
	(0.0238)	(0.00858)	(0.0197)
2 Year Lag	-0.0158	-0.00930	-0.119***
	(0.0210)	(0.00876)	(0.0239)
3 Year Lag	-0.0578***	-0.107***	-0.0632***
	(0.0217)	(0.00803)	(0.0200)
Observations	14540	14540	14540
R^2	0.089	0.164	0.036

Panel D: With No Fixed Effects

	(1) TS-I	(2) EG	(3) TS-E
1 Year Lag	0.117***	0.377***	0.00260
	(0.0272)	(0.00843)	(0.0198)
2 Year Lag	0.0632***	0.00589	-0.0913***
	(0.0207)	(0.00872)	(0.0223)
3 Year Lag	0.0269	-0.0806***	-0.0334*
	(0.0230)	(0.00787)	(0.0187)
Observations	14540	14540	14540
R^2	0.020	0.141	0.010

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows how correlated shocks are with themselves, following the regression specified by Equation (2). We regress shocks on 1-, 2-, and 3-year lagged shocks and a constant. We alternate including CZ and year fixed effects in each specification. All shocks are converted into z-scores. We do not cluster standard errors by year or by commuting zone, and observations are not weighted. Data are from the shocks master data set, detailed in the Data Appendix. We run the regression on shocks to all commuting zones in the US (except for Alaska, which is aggregated to the state level) over the years 1994 to 2016.

Table 5: The Effect of Demand Shocks on Employment Growth

	(1) TS-I	(2) TS-I.b	(3) TS-E	(4) EG
shock	0.0643*	-0.00986	0.204***	2.719***
	(0.0340)	(0.0134)	(0.0360)	(0.356)
Observations	13813	13813	13813	13813
R^2	0.557	0.557	0.559	0.602

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows the effect of annual shocks on annual employment growth, corresponding to Equation (3). We regress employment growth at the CZ-year level on a local shock, commuting zone fixed effects, year fixed effects, and a constant. The shocks are converted into z-scores. Standard errors are clustered by year by commuting zone, and observations are weighted by CZ employment in 1990. Data cover all commuting zones in the US (with the exception of Alaska, which is aggregated to the state level) over the years 1994 to 2016. Employment data are taken from County Business Patterns, and construction of the shocks is detailed in the Data Appendix.

Table 6: The Effect of Demand Shocks on Employment Growth by Quantile

	TS-I	TS-I.b	TS-E	EG
Median	0.049*	–0.012	0.095**	2.251***
	(0.028)	(0.010)	(0.038)	(0.128)
10th percentile	0.023	–0.038**	0.134*	2.349***
	(0.053)	(0.019)	(0.071)	(0.234)
90th percentile	0.076	0.015	0.055	2.148***
	(0.054)	(0.019)	(0.073)	(0.248)
Mean	0.049	–0.012	0.095**	2.251***
	(0.040)	(0.035)	(0.039)	(0.085)

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows the effect of annual shocks on annual employment growth at different quantiles. The table shows the results from three different quantile regressions — median, 10th percentile, and 90th percentile — for each shock, and an unweighted mean regression for comparison. Each row corresponds to a distinct regression, and the entries are the shock's coefficient and standard error. We regress employment growth at the CZ-year level on a local shock, commuting zone fixed effects, year fixed effects, and a constant. The shocks are converted into z-scores. Observations are unweighted, and standard errors are not clustered. Data cover all commuting zones in the US (with the exception of Alaska, which is aggregated to the state level) over the years 1994 to 2016. Employment data are taken from County Business Patterns, and construction of the shocks is detailed in the Data Appendix.

Table 7: The Effect of Demand Shocks on Employment Growth in High Shock Volatility CZs

	(1) TS-I	(2) TS-I.b	(3) TS-E	(4) EG
Shock	0.427	0.02	0.067	2.642***
	(0.308)	(0.06)	(0.279)	(0.389)
High Volatility CZ	-0.360	-0.030	0.137	0.0836
	(0.301)	(0.554)	(0.269)	(0.136)
Observations	13813	13813	13813	13813
R^2	0.558	0.557	0.559	0.603

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

*Notes:*The table shows the effect of annual shocks on annual employment growth, corresponding to Equation (3). We regress employment growth at the CZ-year level on a local shock, commuting zone fixed effects, year fixed effects, and a constant. The shocks are converted into z-scores. Standard errors are clustered by year by commuting zone, and observations are weighted by CZ employment in 1990. Data cover all commuting zones in the US (with the exception of Alaska, which is aggregated to the state level) over the years 1994 to 2016. Employment data are taken from County Business Patterns, and construction of the shocks is detailed in the Data Appendix.

Table 8: The Effect of Local Shocks on Out-Migration Rates

	(1) TS-I	(2) EG	(3) TS-E	(4) TS-I	(5) EG	(6) TS-E
	1 Year Out-Migration Rate			5 Year Out-Migration Rate		
Shock to CZ	0.003	–0.108**	–0.008	0.007	-0.016	0.022
	(0.008)	(0.045)	(0.007)	(0.018)	((0.138)	(0.021)
Shock to CZ, 1 year lag	0.006	–0.007***	–0.006	0.021	0.047	0.031
	(0.010)	(0.048)	(0.009)	(0.019)	(0.096)	(0.018)
Average shock to other CZs	–0.046*	0.219	–0.074***	-0.194*	0.642	-0.092
	(0.026)	(0.205)	(0.023)	(0.107)	(0.679)	(0.111)
Average shock to other CZs, 1 year lag	–0.074	0.033	–0.088***	-0.152	0.903	-0.001
	(0.045)	(0.212)	(0.026)	(0.098)	(0.643)	(0.154)
Constant	3.753***	3.760***	3.717***	12.17***	12.42***	12.09***
	(0.009)	(0.034)	(0.006)	(0.034)	(0.216)	(0.008)
N	13086	13086	13086	12359	12359	12359
R^2	0.897	0.897	0.898	0.934	0.934	0.933

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows the effect of shocks on out-migration rates. As specified by Equation (9), we regress out-migration rates on a local shock, a one-year lag of that shock, a weighted average of the shock to other CZs, and a one-year lag of that weighted average. We include commuting zone and year fixed effects. Constants reflect the mean migration rate. Standard errors are clustered by commuting zone and by year. Migration data are described in Section 2.2, and data used to construct the shocks are described in the Data Appendix.

Table 9: The Effect of Local Shocks on In-Migration

	(1) TS-I	(2) EG	(3) TS-E	(4) TS-I	(5) EG	(6) TS-E
	1 Year In-Migration Rate			5 Year In-Migration Rate		
Shock to CZ	0.014	0.193***	0.002	-0.017	0.411***	0.009
	(0.008)	(0.042)	(0.014)	(0.022)	(0.137)	(0.028)
Shock to CZ, 1 year lag	0.006	0.146***	–0.008	-0.029	0.182*	0.016
	(0.008)	(0.048)	(0.014)	(0.024)	(0.098)	(0.026)
Average shock to other CZs	–0.010	0.202*	0.021	0.114	0.524	0.091*
	(0.019)	(0.113)	(0.016)	(0.066)	(0.466)	(0.049)
Average shock to other CZs, 1 year lag	0.010	0.262**	0.009	0.116*	-0.277	0.072**
	(0.030)	(0.114)	(0.017)	(0.056)	(0.428)	(0.031)
Constant	3.709***	3.825***	3.711***	11.93***	12.11***	12.01***
	(0.007)	(0.035)	(0.003)	(0.021)	(0.173)	(.003)
N	13086	13086	13086	12359	12359	12359
R^2	0.9332	0.9362	0.9327	0.9490	0.9496	0.9488

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports results of estimating specifications from Table 8 using inflow rates as the dependent variable. Standard errors are clustered by commuting zone and by year. Migration data are described in Section 2.2, and data used to construct the shocks are described in the Data Appendix.

Table 10: The Effect of Local Shocks on Out- and In-Migration

Panel A: Out-Migration Rates				
	(1)	(2)	(3)	(4)
	TG	GR	TG	GR
	1 Year Out-Migration Rate		5 Year Out-Migration Rate	
Shock to CZ	-0.001	0.012	–0.018	0.296*
	(0.041)	(0.046)	(0.141)	(0.164)
Constant	3.608***	3.779***	11.72***	12.21***
	(0.003)	(0.011)	(0.009)	(0.036)
N	8724	11632	8724	10178
Panel B: In-Migration Rates				
	(1)	(2)	(3)	(4)
	TG	GR	TG	GR
	1 Year In-Migration Rate		5 Year In-Migration Rate	
Shock to CZ	0.030	0.196***	-0.213	0.295**
	(0.033)	(0.052)	(0.142)	(0.147)
Constant	3.588***	3.792***	11.67***	12.1***
	(0.002)	(0.012)	(0.009)	(0.033)
N	8724	11632	8724	10178

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table reports results of estimating specifications from Equation 10 using outflow and inflow rates as the dependent variable. Standard errors are clustered by commuting zone and by year. Migration data are described in Section 2.2, and data used to construct the shocks are described in the Data Appendix.

Data Appendix

To analyze the effects of local shocks on migration, we construct multiple local demand shocks from the literature. While following their original methodologies as closely as possible, we construct annualized versions of each shock and extend them over two decades: 1994-2016. The resulting data set is a balanced panel of multiple shocks by year by commuting zone. Our main analysis uses this panel of shocks, combined with migration flows constructed from tax data, to measure and compare each shock's effect on employment and regional migration.

This document is organized as follows. Appendix A details the underlying data sources used to construct the shocks. Appendix B describes the methodology used to construct each shock. Appendix C contains additional figures and tables.

Appendix A Data Sources and Preparation

Appendix A.1 Delineating local labor markets and industries

To measure the effects of demand shocks on local labor markets, we follow the literature and use commuting zones to delineate local labor markets. A commuting zone is a cluster of adjacent counties between which many residents commute (Tolbert and Sizer, 1996). An advantage of using commuting zones to classify geographies is that every county in the US maps to a commuting zone, so the entire US is covered. We follow the literature and use the 1990 vintage instead of the 1980 or the 2000 vintage. We also aggregate Alaska into a single commuting zone because its boroughs changed borders many times over the past several decades, which makes using the 1990 commuting zones across all years difficult. After aggregating Alaska, there are a total of 727 commuting zones that cover the entire United States.

Constructing several of the demand shocks requires delineating employment and trade flows by industry. We use the North American Industry Classification System (NAICS) for industry classification. Several papers in the literature use Standard Industrial Classification (SIC) instead of NAICS.[27] We use the NAICS classification for two reasons. First, more years of underlying data used NAICS classification than SIC. Second, the NAICS classification groups establishments into industries by their product markets, which logically helps with mapping imports and exports, classified by products, to industries. In contrast, how SIC groups establishments into industries is inconsistent; Triplett et al. (1995) finds that roughly 25 percent of industries are defined by their products, 20 percent are defined by the markets they serve, and 20 percent are defined with no discernible methodology. Because of the availability of weighted crosswalks, as detailed in Appendix A.4, we specifically use the 1997 NAICS vintage.

[27] For example, Autor et al. (2013).

Appendix A.2 Measuring employment by industry, year, and commuting zone

We need employment measures by commuting zone, industry, and year to construct multiple local shocks. This section describes how we obtain this information from County Business Patterns (CBP) data. The resulting data series spans 1988 to 2016 and covers the entire United States.

CBP is an annualized data series with employment by industry and by county. The main advantages of using the CBP are its detailed breakout and extensive coverage. Industries are broken out into detailed 6-digit NAICS codes (or 4-digit SIC codes before 1998). The data series also covers every year since 1986, the entire US, and almost all private industries.[28]

While most observations, which are an industry-county-year, have exact employment counts, some do not for privacy-protection reasons. Observations missing an exact employment count instead have an upper and lower bound for employment. Each year before 2017, over 50 percent of observations are missing exact employment counts. (For the years after and including 2017, instead of creating an observation with a range for employment, an observation affected by privacy concerns is completely dropped from the data set.) The observations with missing employment counts tend to have relatively small employment levels.

We follow the literature and use the algorithm from Autor et al. (2013) and subsequently used by Acemoglu et al. (2016) and Feenstra et al. (2019) to impute the missing employment counts. The algorithm exploits the fact that the CBP provides firm counts by firm size bin for each observation. By year and at a national level, the algorithm estimates the average size of a firm in each firm size bin. Exact employment counts are then estimated using the estimated average firm sizes and the firm counts. We leave the specific details to the online data appendix of Autor et al. (2013).[29]

Another issue with the CBP data is that some employment is missing classification at the most disaggregate level. Before 1998, some county-level employment has 3-digit SIC codes but is missing a 4-digit classification. (SIC codes only have 4 digits, so the 4-digit SIC code is most similar to the 6-digit NAICS code.) The cause of this issue was that the original researchers who compiled the CBP data sometimes did not know an industry's most detailed classification but did know its classification at a lower, less-specific level (Eckert et al., 2020). We follow Acemoglu et al. (2016) and address this problem by proportionally allocating employment. For each year, we construct each 4-digit code's share of its 3-digit employment at the national level. Then we use these shares to allocate the 3-digit level employment to the 4-digit level.

[28] CBP excludes the following industries: Crop and Animal Production (NAICS 111,112); Rail Transportation (NAICS 482); Postal Service (NAICS 491); Pension, Health, Welfare, and Vacation Funds (NAICS 525110, 525120, 525190); Trusts, Estates, and Agency Accounts (NAICS 525920); Office of Notaries (NAICS 541120); Private Households (NAICS 814); and Public Administration (NAICS 92). Source: https://www.census.gov/programs-surveys/cbp/technical-documentation/methodology.html#par_textimage_36648475.

[29] An alternative option for imputing the missing employment counts is proposed by Eckert et al. (2020). The authors make a data set with imputed employment counts, along with a replication file, publicly available at http://fpeckert.me/cbp/. However, we attempted to run their replication files, and the resulting imputed data did not match those they published on their website. Hence, we found that using the data from Eckert et al. (2020) was not a viable option.

The CBP switches industry classifications multiple times. The most notable change was the switch from SIC to NAICS in 1997. Additionally, every five years, both classification schemes changed vintages and were updated to better delineate current economic activity. Following the literature, we address this problem by using concordances to map each classification scheme and its vintages to a single scheme: the 1997 NAICS vintage. The concordances we use are detailed in Appendix A.4.

Finally, counties are mapped to 1990 commuting zones (using the same adjustment for Alaska described above). The resulting data series has employment by commuting zone, by industry, and by year. It covers the entire continental US and almost all private industries, and spans 1986-2016. Across all years, industries are categorized using 6-digit NAICS codes.

Appendix A.3 Measuring manufacturing imports, exports, and shipments

We need measures of manufacturing imports, exports, and shipments to construct the trade shocks. The underlying data we use to construct these imports and exports are from Schott (2008).[30] These data are mainly drawn from US Customs Services. The data have imports to the US and exports from the US at 10-digit product level, spanning 1989-2021. The advantage of using trade data from Schott (2008) as opposed to other sources is its coverage and granularity. For comparison, UN Comtrade uses 6-digit HS codes and begins coverage on China imports to the US in 1992.[31]

After downloading the data, we map the 10-digit product-level codes to 6-digit NAICS codes using a concordance from Pierce and Schott (2012). This crosswalk maps 10-digit product codes to 6-digit NAICS codes and varies by year, using each scheme's active vintage. When using this crosswalk we had one minor issue. Of 414,000 year-HS code observations, around 100 observations (roughly 0.025 percent) do not have a matched NAICS code. To address this problem we assign these HS codes the NAICS codes of the adjacent HS codes. (This is part of the methodology Pierce and Schott (2012) use to create their crosswalk. They use a similar concordance provided by the Census as a starting point and fill in missing codes in various ways.) As NAICS vintages vary across time, we map each vintage to the 1997 vintage using the crosswalks described in Appendix A.4. We use the Personal Consumer Expenditure Chain Type Price Index to deflate nominal trade flows into real dollars.[32]

Shipments—which are similar to revenue—are directly available from the NBER-CES Manufacturing Productivity Database.[33] This data set has shipments by industry by year. Nominal shipments are deflated to real dollars using the Personal Consumer Expenditure Chain Type Price Index.

[30] Import and export data are available here: https://sompks4.github.io/sub_data.html.

[31] The official source of trade data from the US Census is available here: https://usatrade.census.gov/index.php. We use Schott's data as opposed to usatrade because Schott's data starts in 1989, while the USA Trade data start in 1992. Further, we found Schott's data easier to use.

[32] This index was downloaded from https://fred.stlouisfed.org/series/PCEPI in April 2022.

[33] Source: https://www.nber.org/research/data/nber-ces-manufacturing-industry-database.

Appendix A.4 Crosswalks

Appendix A.4.1 Mapping geographies to commuting zones

The tax data used to construct our migration flows classify geography using zip codes and counties. Additionally, the CBP data which we used to measure employment in Appendix A.2 use counties to delineate geographies. Hence, we need crosswalks to map zip codes and counties to commuting zones.

To map zip codes to counties, we use an annual crosswalk from the United States Department of Housing and Urban Development (HUD). In theory, when using multiple years of tax data, using the corresponding vintage of the crosswalk across each year would yield the most matches. In practice, we find that using multiple zip-to-county crosswalks does not yield more matches. Using the 2021 vintage across all years yields slightly more matches than using each vintage with its corresponding year. So, for simplicity, we only use the 2021 vintage of the zip-code-to-county crosswalk.

To map counties to commuting zones we use the crosswalk `cw_cty_czone.dta` from David Dorn's data page.[34] Due to county code changes over time, we slightly modify this crosswalk following Autor et al. (2013).[35] By construction, each county maps to a single CZ, except for several counties in Alaska. In Alaska, many counties have changed borders and names repeatedly over the past several decades. Due to this, we aggregate Alaska to the state level.[36][37]

Additionally, when constructing the education-level Bartik shock we use data from the American Community Survey and the Decennial Census. The most-detailed geographic level available in these data sets is the Public Use Micro Area (PUMA).[38] To map PUMAs to CZs, we use two weighted concordances: `cw_puma1990_czone.dta` and `cw_puma2000_czone.dta`. The former maps 1990 PUMAs to CZs, and the latter maps 2000 PUMAs to CZs. The source of these crosswalks is David Dorn's data page.[39]

Appendix A.4.2 Mapping the trade data to NAICS classification

The trade data from Schott (2008) is classified using product codes. Imports use 10-digit Harmonized Tariff Schedule (HTS) Code, while exports use 10-digit Schedule B Commodity Code. We map these 10-digit codes to 6-digit NAICS codes using two crosswalks from Pierce and Schott (2012): `hs_sic_naics_imports_89_118_20192105.dta` and `hs_sic_`

[34] Source: http://www.ddorn.net/data/.

[35] The changes to county codes across time are detailed here: https://www.ddorn.net/data/FIPS_County_Code_Changes.pdf.

[36] Alaskan county classification changes in the 1980s are detailed on a Census website https://www.census.gov/programs-surveys/geography/technical-documentation/county-changes/1980.html. Analogously named web pages exist for subsequent decades.

[37] For a discussion of how using commuting zone delineation for labor markets affects empirical results, see Foote et al. (2021).

[38] We download data from both these surveys from IPUMS USA (Ruggles et al., 2020). Note that these data do have a county variable, but many counties are not actually identifiable or available in these data due to privacy concerns.

[39] Source: http://www.ddorn.net/data/.

`naics_exports_89_118_20192105.dta`. Both crosswalks vary across time, using the active vintages each year.

Appendix A.4.3 Harmonizing industries to NAICS 1997 vintage

We harmonize industry classification across all years and data sets to the 1997 NAICS vintage. This requires several crosswalks, as the CBP uses SIC classification before 1998, and after 1998 it changes NAICS vintages every five years. Further, after mapping the product codes in the trade data to NAICS codes, the NAICS codes are classified using multiple vintages.

Our strategy for harmonizing industry classification is to use weighted crosswalks whenever they are publicly available. These crosswalks were created by other researchers using establishment-level microdata available through the Census Bureau. In years where industry classification changes, the Census Bureau publishes both the former industry code and current industry code of firms. Collapsing to sum employment over both former and current industry codes allows one to create a weighted crosswalk. The weights are the share of an industry's employment in one vintage that maps to an industry code in the other vintage.

However, for some industry code changes, weighted crosswalks are not publicly available. In these instances we use unweighted crosswalks from the Census Bureau. These crosswalks map industry codes from one vintage to another, but without weights. When a code maps to multiple codes from another vintage, we simply divide employment (or trade flows) evenly among all the codes to which it maps. As detailed below, the amount of employment and industries affected by this lack of weighting is small, so we do not expect it to impact our results.

To map each NAICS vintage to the 1997 vintage, we use three crosswalks. The first crosswalk maps the 2012 vintage to the 2007 vintage. The second maps the 2007 vintage to the 2002 vintage. The last crosswalk maps the 2002 vintage to the 1997 vintage. The latter two crosswalks are weighted and from the replication package of Acemoglu et al. (2016). The crosswalks are `naics02_naics97.dta` and `naics07_naics02.dta`.

We map NAICS 2012 codes to the 2007 vintage using the crosswalk `naics12_naics07.dta`. Since a weighted version of this crosswalk is not publicly available, we use an unweighted mapping; whenever a 2012 code maps to multiple 2007 codes, we divide employment or trade flows evenly among the 2007 codes. This should have a small economic impact as it affects a small share of employment. For instance, in 2012 only 2.7 percent of all private employees in the CBP data set worked in a 6-digit 2012 code that mapped to more than one 2007 code.

To map the 1987 vintage SIC to NAICS 1997 we use the crosswalk `sic87_naics97.dta`. To create this crosswalk, we start with a weighted crosswalk from Fort et al. (2016). However, this crosswalk is missing several industries. For such cases, we use an unweighted mapping; whenever a SIC 1987 code maps to multiple NAICS 1997 codes, we divide employment or trade flows evenly among the NAICS 1997 codes. This strategy should have a small economic impact as it affects a small share of employment. For instance, in 1995 only 6.6 percent of all private employees in the CBP data set worked in a 4-digit 1987 SIC code that mapped to more than one 6-digit 1997 NAICS code.

Appendix A.5 Construction of CZ moderators

The vector X_j of CZ moderators is made up of eight different variables. These variables are summarized in Table A1 and summary statistics in Table A2.

We construct measures of base period CZ-level characteristics to explore heterogeneity in migration responses following local shocks. These are enumerated, along with their sources, in Table A1. All moderators are constructed from publicly available data. They are aggregated to the CZ level from either counties or Census PUMAs. In either case, all CZ-level moderators are computed as the population-weighted mean of the corresponding county- or PUMA-level measure.

To ease interpretation of reported interaction coefficients, we converted each CZ moderator into a binary variable. College graduates, employment, the top 20 percent share, the share local born and January temperatures all take a value of 1 if they are greater than the median and 0 otherwise. College access, high CZ population, and median home prices are also converted to a binary variable based on cut-offs determined from their CDFs as described below.

Appendix A.5.1 Employment rate

The rate of employment is computed as total employed population (including military) between the ages of 16 and 64 (the sum of employed females and employed males) divided by the total population aged 16-64. These estimates were pulled directly from the Opportunity Atlas (Chetty et al., 2024). While Chetty et al. (2024) provide estimates for multiple years, we use the estimates for 2000 from the Decennial Census.

Appendix A.5.2 Local-born share

The local-born share is calculated as the share of a population that was born in the same state they reported as their residence in the 2000 Census. This was calculated using the 5 percent Decennial Census sample from Ruggles et al. (2022).

Appendix A.5.3 Top 20 percent share

The fraction of low-income children who grew up in a CZ who have average household income in 2014-15 (in their mid-30s) in the top 20 percent of the national income distribution for children born in the same year. This is the only income mobility measure we use from Chetty et al. (2024) and is calculated using income tax records.

Appendix A.5.4 College graduates

The number of people aged 25 or older who have a bachelor's degree, master's degree, professional school degree, or doctorate degree, divided by the total number of people aged 25 or older. These estimates were pulled directly from the Opportunity Atlas (Chetty et al., 2024). While Chetty et al. (2024) provide estimates for multiple years, we use the estimates for 2000 from the Decennial Census.

Table A1: CZ Moderators

CZ Moderator	Description	Source
Employment Rate	The rate of employment for individuals between 16 and 64 years of age.	Chetty et al. (2024)
Share Local Born	The share of the population who reported being born in the same state as their residence in the 2000 Census.	Ruggles et al. (2022)
Top 20 Percent Share	The fraction of low-income children who grew up in a CZ who have average household income in the top 20 percent of the national income distribution at age 35.	Chetty et al. (2024)
College Graduates	The share of a CZ population with a four-year college degree or higher.	Ruggles et al. (2022)
College Access	A binary variable reflecting access to community or four-year colleges with high admission rates.	Winfield (2023)
High CZ Population	A binary variable that takes a value of 1 if the CZ is one of the top 40 largest CZs by population.	Ruggles et al. (2022)
January Temp	Average daily maximum January temperatures.	CDC (2025)
Home Prices	Median home value in the CZ.	Ruggles et al. (2022)

Table A2: CZ Moderators Summary Statistics

Moderator	mean	sd	min	p10	p25	p50	p75	p90	max
Employment Rate	0.69	0.072	0.39	0.59	0.64	0.69	0.74	0.77	0.83
Share Local Born	0.64	0.016	0.09	0.41	0.57	0.68	0.75	0.79	0.88
Top 20 Percent Share	0.10	0.03	0.03	0.05	0.07	0.10	0.13	0.18	0.42
College Graduate	0.18	0.06	0.06	0.11	0.14	0.16	0.21	0.27	0.43
January Temp	41.0	13.6	12.9	23.7	31.2	41.0	51.2	60.7	72.3
Home Prices (2000)	85376	45215	23500	46898	60804	77114	97600	125043	577500

Notes: These summary statistics are not weighted by CZ population. When we do so, the college graduation rate average is 24 percent. Home prices are in 2000 dollars. .

Appendix A.5.5 College access

College access is a binary variable to identify if an area has a high level of access to post-secondary educational institutions. The variable takes a value of 1 if a CZ has at least two community colleges or a four-year institution with at least a 75 percent rate of admission. This measure uses National Center for Education Statistics (IPEDS) data and was created by Winfield (2023).

Appendix A.5.6 High CZ population

The high CZ population measure is a dummy variable that takes a value of one if a CZ is in the top 41 of CZ populations; 41 was chosen because summing the populations of the top 41 CZs provides 50 percent of the US population in the 2000 Census. This measure was calculated using the 5 percent 2000 decennial census sample from Ruggles et al. (2022).

Appendix A.5.7 Average January temperatures

The January temperature data are from the CDC and represent the average daily max of January 2000 temperatures within a CZ. The non-binary CZ controls were converted to a dummy variable, taking a value of 1 if the value of a CZ was above the median.

Appendix A.5.8 Median home prices

We use the county-level median home price from the 2000 and 1990 Decennial Census. We aggregate to the CZ level. We convert this into a binary variable by assigning a value of one to each CZ observation that is in the top 50 percent of median home prices. All other CZ observations take a value of zero.

Appendix B Construction of Local Shocks

This section describes the methodologies used to construct each local shock, using the data from Appendix A. Appendix B.1 describes our shift-share methodology. Appendix B.2 details

background information on existing demand shocks in the literature. Appendix B.3 and Appendix B.4 detail the methodologies for how to construct the time-variant and the time-invariant shocks.

Appendix B.1 The shift-share approach

In our analysis every local shock (except the Yagan shock) uses a shift-share design. Since identification in this setting has been discussed extensively elsewhere (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022b; Adao et al., 2019), we only discuss the instrument briefly here.

The shift-share design imputes to each locality some change(s) in national conditions (e.g., industry-specific shocks) based on differential local exposure to those conditions (e.g., differing employment shares by industry at the CZ level). Generically, the shock to CZ i at time t is:

$$Shock_{it} = \sum_j \underbrace{\frac{L_{ijt}}{L_{it}}}_{\substack{\text{CZ-industry} \\ \text{level weight}}} \times \underbrace{\frac{\Delta Z_{jt}}{A_{jt}}}_{\substack{\text{Industry-level} \\ \text{Shock}}}$$

where $\frac{L_{ijt}}{L_{it}}$ is the share of employment in CZ i at time t that is in industry j, ΔZ_{jt} is the industry-level shock (for example, the change in employment or trade flows at the national industry level between t and $t-1$), and A_{jt} is a term that gives scale to the industry-level shock. The weight term is identical across each time-variant shock, while the industry-level term $\frac{\Delta Z_{jt}}{A_{jt}}$ varies. From here, we lag the employment weight terms and the normalization term,

$$Shock_{it} = \sum_j \underbrace{\frac{L_{i,j,t-1}}{L_{i,t-1}}}_{\substack{\text{CZ-industry} \\ \text{level weight}}} \times \underbrace{\frac{\Delta Z_{j,t-1,t}}{A_{j,t-1}}}_{\substack{\text{Industry-level} \\ \text{Shock}}}.$$

For annual versions of each shock, we use trailing five-year averages to smooth annual spikes, and the industry-level shock ΔZ_{jt} is taken as the change in Z_j between years $t-1$ and t,

$$Shock_{it} = \sum_j \underbrace{\frac{\overline{L_{i,j,t}}}{\overline{L_{i,t}}}}_{\substack{\text{CZ-industry} \\ \text{level weight}}} \times \underbrace{\frac{\Delta Z_{j,t-1,t}}{\overline{A}_{jt}}}_{\substack{\text{Industry-level} \\ \text{Shock}}}$$

where the overline denotes trailing five-year averages:

$$\overline{x_t} = \frac{1}{5}\sum_{k=1}^{5} x_{t-k}.$$

For decadal versions, the industry-level shock is taken as the change over the decade. The weights and normalization term A are taken from the first year of the decade. Due to

data constraints, we do not actually take the decadal shock over decades, rather we use three long time periods: 1990 to 1999, 2000 to 2007, and 2008 to 2016.

Appendix B.2 Overview of local demand shock measurement

An early and influential measure was proposed and analyzed by Bartik (1991), then implemented by Blanchard and Katz (1992). The "Bartik shock," or Bartik instrument, for labor demand movements is a shift-share measure of the percentage change in local employment predicted by growing local employment in various industries over time by the industry's national level growth rate, then normalized by a locality's base period employment level. Numerous studies have since used Bartik shocks to assess a range of outcomes following local demand shocks (e.g., Wozniak, 2010; Dao et al., 2017).

A second approach uses industry-specific booms and busts—usually in extractive industries—to measure local shocks (Wilson, 2022; Gallin, 2004). This is essentially a special case of the Bartik shock. An advantage is that it is often possible to directly support the assumption that the demand shift was truly unanticipated. For example, Black et al. (2005) compare counties with coal reserves to counties without such reserves within the US coal-producing region over a period in which global demand for coal rose and fell dramatically.

In the 2010s, researchers have turned their attention to the diverse local impacts of increased trade with China. The "China shocks" have been quantified in different ways. Autor et al. (2013), hereafter ADH, originally measured these as average local per worker dollars of Chinese imports. Subsequently, they modified their measure to normalize by local consumption, so the measure reflects an increase in local consumption intensity in Chinese goods (Acemoglu et al. 2016, hereafter AADHP). In both cases, inflows of Chinese imports in the numerator may reflect changes in local labor supply or local industry growth for reasons other than competition from Chinese imports. To isolate the unanticipated shifts in local demand for labor, both papers instrument for Chinese exports to US local markets using levels of Chinese exports to international markets.

Subsequent papers modify this approach to measuring local demand shocks arising from trade with China. Feenstra et al. (2019) examine industry- and commuting-zone-level employment changes after accounting for both rising import competition from China and expanding global exports. This takes a more expansive view of trade-related labor demand shocks than that in ADH. Pierce and Schott (2016b) quantify the trade shock using commuting-zone-level changes in tariff burden (weighted by local industrial composition) following China's accession to the WTO in 2000.

Appendix B.3 Time-variant shocks

We construct annual versions of each time-variant shock from 1992 to 2016. Each shock uses a shift-share design.

Appendix B.3.1 Autor, Dorn, and Hanson (2013) shock

The Autor et al. (2013) (ADH) shock is an extension of the decadal shock described in the authors' 2013 paper in the *American Economic Review*. The shock measures the effect of

a change in industry-level imports from China on a commuting zone by using its industry employment shares as weights. The annual ADH shock for commuting zone i in year t is

$$\Delta ADH_{it} = \sum_j \frac{\overline{L_{ijt}}}{\overline{L_{it}}} \frac{\Delta M_{jt}^{China}}{\overline{L_{jt}}}$$

where j denotes industries, ΔM_{jt}^{China} is the one-year change in the real US dollar value of imports from China in industry j between years t and $t-1$, L is employment, and the overline denotes the trailed moving five-year average starting with year $t-1$.

The data sources required to construct this shock are employment data from Appendix A.2 and the trade data from Appendix A.3.

Appendix B.3.2 Acemoglu, Autor, Dorn, Hanson, and Price (2016) shock

The Acemoglu et al. (2016) (AADHP) annual shock is identical to the ADH annual shock with one difference: AADHP normalizes changes in imports by absorption, while ADH uses industry-level employment. The annual AADHP shock to a CZ i at time t is

$$\Delta AADHP_{it} = \sum_j \frac{\overline{L_{ijt}}}{\overline{L_{it}}} \frac{\Delta M_{jt}^{China}}{\overline{Y_{jt}} + \overline{M_{jt}} - \overline{E_{jt}}}$$

where j denotes industry, L denotes employment, and ΔM_{jt}^{China} is the one-year change in imports into the US from China. The denominator on the rightmost term, $Y_{jt} + M_{jt} - E_{jt}$, is absorption: Y denotes industry shipments (sales), M denotes imports, and E denotes exports. The overline denotes the five-year trailing average.

The underlying data needed for construction of this shock are the employment data from Appendix A.2 and trade data from Appendix A.2.

Appendix B.3.3 Feenstra, Ma, and Xu (2019) export shocks

While Autor et al. (2013) analyze imports into the US from China, Feenstra et al. (2019) consider both Chinese imports and US exports. The Feenstra et al. (2019) (FMX) export shock is similar to the AADHP shock; the only differences are that US exports replace Chinese imports, and shipments replace absorption. The annual FMX export shock to a CZ i at time t is

$$\Delta FMX_{it} = \sum_j \frac{\overline{L_{ijt}}}{\overline{L_{it}}} \frac{\Delta X_{jt}}{\overline{Y_{jt}}}$$

where j denotes industry, L denotes employment, ΔX_{jt} is the one-year change in exports from the US to the rest of the world, and Y_{jt} is shipments. The overline denotes the trailing five-year average starting with year $t-1$.

The underlying data needed to construct this shock are the employment data from Appendix A.2 and trade data from Appendix A.2.

Appendix B.3.4 General Bartik shock

The Bartik shock uses industry employment shares at the CZ-level to measure the CZ's exposure to national-level changes in industry employment. The Bartik shock at time t in commuting zone i is

$$\Delta Bartik_{it} = \sum_j \frac{\overline{L_{ijt}}}{\overline{L_{it}}} \frac{L^i_{jt} - L^i_{jt-1}}{L^i_{jt-1}}$$

where j denotes industries and L denotes employment. The overline denotes the five-year trailing averages. L^i_{jt} is national industry-level employment without CZ i's contribution, $L^i_{jt} = L_{jt} - L_{ijt}$, which we use to avoid endogeneity concerns.

An issue with constructing the industry growth term is that because of imperfections with the industry crosswalk, sometimes industry growth rates are implausibly large. We find that dropping these years from our analysis does not significantly affect our results.

The general Bartik shock is constructed using only the employment data described in Appendix A.2.

Appendix B.3.5 Education-specific Bartik shocks

We construct two education-level Bartik shocks; one for workers with 4-year college degrees and another for those without. The annual Bartik shock for people of education level e in commuting zone i at time t is

$$\Delta Bartik_{eit} = \sum_j \frac{L_{eijt}}{L_{eit}} \frac{L^i_{jt} - L^i_{jt-1}}{L^i_{jt-1}}$$

where

$$\frac{L_{eijt}}{L_{eit}} = \begin{cases} \frac{L_{eij,1990}}{L_{ei,1990}} & \text{if } t \leq 2000 \\ \frac{L_{eij,2000}}{L_{ei,2000}} & \text{if } t \in (2000, 2007] \\ \frac{L_{eij,2007}}{L_{ei,2007}} & \text{if } t > 2007 \end{cases}$$

where j denotes industry and L denotes employment level. L^i_{jt} is the national industry-level employment without CZ i's contribution, $L^i_{jt} = L_{jt} - L_{ijt}$, which we use to avoid endogeneity concerns.

Compared to the general Bartik shock, the industry growth term is constructed with the same CBP data, while the weight term differs. Here, the weight term is constructed from the Decennial Census and the American Community Survey (ACS) from IPUMS, while in the general version it is constructed from the CBP data. This is due to the lack of education information in the CBP. While the Decennial Census and ACS have the necessary variables for the construction of the weights, these surveys are not available annually over the 1980s and 1990s. Hence, we use the 1990 Decennial Census, the 2000 ACS, and the 2006-2008 ACS, and construct the weights using the most recent year of data available.

The Decennial Census and ACS also use a different, less-detailed industry classification system than the CBP. We construct crosswalks to aggregate the CBP data, which use 6-digit NAICS codes, to the Census' and ACS's industry classification scheme. This process is

aided by a variable constructed by the IPUMS team, indnaics, which is essentially NAICS classification at varying digit levels.

The Decennial Census and the ACS also do not consistently provide county or commuting zone classification. Its finest geographic level is the public use micro area, so we use crosswalks detailed in Appendix A.4 to map public use micro areas to CZs.

Appendix B.4 Time-invariant shocks

Appendix B.4.1 Greenland, Lopresti, McHenry (2016) tariff gap shock

The tariff gap shock follows the same methodological structure as in Greenland et al. (2019), which builds upon Pierce and Schott (2016b). The tariff gap shock measures exposure to Chinese import competition using changes in tariff policy in 2001. Before 2001, China was granted normal trade relation (NTR) tariff rates on an annual basis. This policy lowered the rates on Chinese imports, but the annual renewal of these rates was politically contentious and uncertain. This uncertainty protected US manufacturers from Chinese import competition, but it was eliminated when China joined the World Trade Organization (WTO) in 2001. At this time China was granted permanent NTR. The tariff gap shock measures commuting zones' exposure to changes in the NTR and non-NTR tariff rates.

The tariff gap shock to a commuting zone is the sum of each industry's change in tariff rates before and after China's WTO entrance, weighted by the industry's commuting-zone-level employment share. The tariff gap shock for commuting zone i is

$$TG_c = \sum_j \frac{L_{cj,1992}}{L_{c,1992}} NTR\ GAP_j$$

where j denotes industry, L denotes employment, and $NTR\ GAP_j$ is the change in industry j's tariff rate after China joined the WTO,

$$NTR\ GAP_j = Non\ NTR\ Tariff_j - NTR\ Tariff_j.$$

The data and crosswalk files required to construct this shock are the employment data described in Appendix A.2 and data on the NTR gaps, provided by Pierce and Schott (2016a). While a data set with the tariff gap shock from Greenland et al. (2019) is publicly available, we reconstruct the shock so that the underlying industry classification is consistent with our other shocks.

Appendix B.4.2 Yagan (2019) Great Recession shock

The time-invariant Yagan (2019) Great Recession shock is simply the change in a CZ's unemployment rate between 2009 and 2007. The Yagan shock for CZ i is

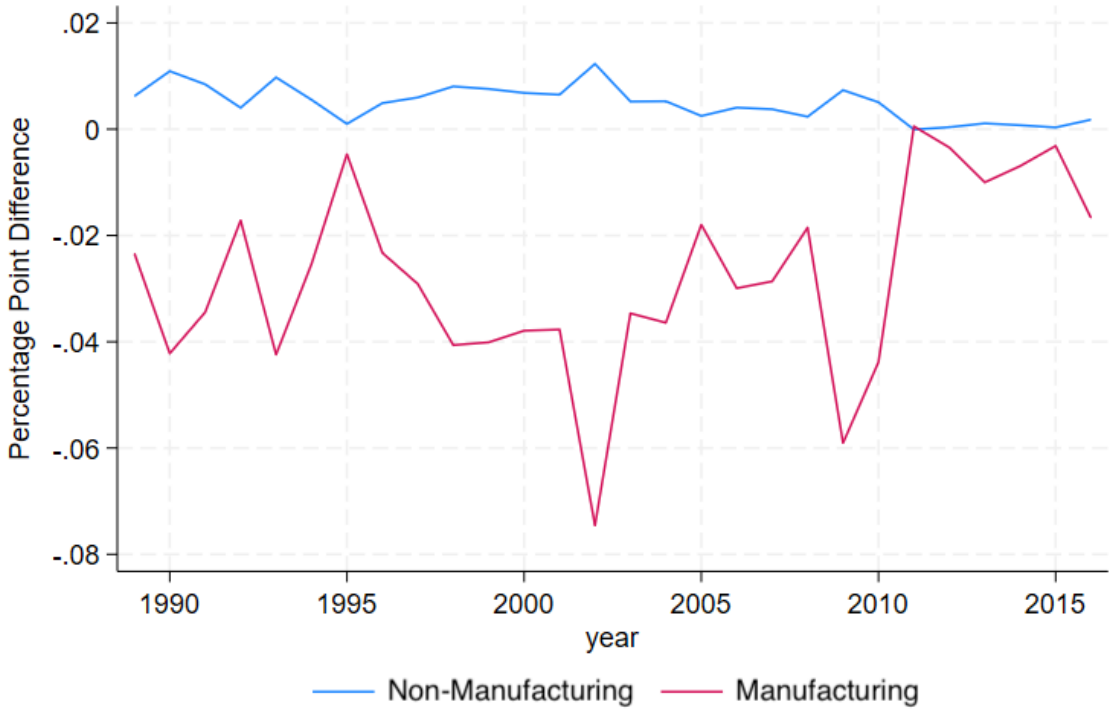
$$YS_c = UR_{c,2009} - UR_{c,2007}$$

where $UR_{c,t}$ denotes the unemployment rate in CZ i at year t. The Yagan (2019) replication material includes a data set with this shock. We directly take the shock from this data set

instead of reconstructing the shock because there is no underlying industry classification to harmonize with our other shock measures.

Appendix C Additional Figures and Tables

Figure C1: Employment Growth in Manufacturing Relative to Overall Growth



Notes: This figure shows percentage point difference between the employment growth in manufacturing and total employment growth between 1994 and 2016. The manufacturing sector is defined using the 1997 vintage NAICS code at the one-digit level.

Table C1: Descriptive Statistics of Annual Shocks

	count	min	p10	p25	p50	p75	p90	max
TS-I	16721	-11337.340	-85.447	6.198	70.187	218.670	469.485	26251.342
TS-I.b	16721	-0.026	-0.000	0.000	0.000	0.001	0.002	0.429
TS-E	16721	-0.191	-0.003	-0.001	0.000	0.002	0.005	0.140
EG	16721	-0.145	-0.021	0.002	0.013	0.020	0.028	0.236
EG, College	16721	-0.056	-0.008	0.010	0.017	0.023	0.029	0.107
EG, No College	16721	-0.100	-0.019	0.002	0.013	0.019	0.027	0.081
TG	727	-0.225	-0.130	-0.102	-0.078	-0.056	-0.041	-0.012
GR	727	-0.126	1.460	2.439	3.685	5.278	6.675	11.846

Notes: The table shows descriptive statistics for the annualized and time-invariant shocks. The annual shocks cover 1994 to 2016, while the tariff gap shock occurred in 2000 and the Great Recession shock in 2007. All commuting zones (1990 vintage) in the US are covered, except for Alaska, which we aggregate to the state level due to inconsistencies with its underlying borough classification, resulting in 727 commuting zones per year. Underlying data are described in the Data Appendix.

Table C2: Observations by Match Status

	1994	1999	2004	2009	2014	2019
a. All Observations						
Total	140	147	149	150	148	148
Matched directly with zip	88	96	98	99	98	104
Matched directly with zip or MAF	95	103	106	108	106	112
Matched directly or imputable	99	106	109	110	108	114
Cannot be matched	41	41	40	40	40	34
b. Observations Born in the US						
Total	112	116	116	115	114	112
Matched directly with zip	77	82	81	81	79	82
Matched directly with zip or MAF	84	89	89	89	86	90
Matched directly or imputable	87	91	91	91	88	91
Cannot be matched	25	24	25	25	26	21

Notes: The table shows the number of observations in millions by year and commuting zone (CZ) match status. An observation is a person-year. The table is constructed using restricted administrative tax data from the Census Bureau. Our sample is restricted to people aged 25-55. Matched is shorthand for whether an observation can be matched with a CZ. Observations can be matched directly if they have either i) a zip code that maps to a single CZ or ii) a county in the Master Address File (MAF). We can impute an observation's CZ if they cannot be matched directly but they have a zip code, which happens when their zip code maps to multiple CZs and their MAF county is missing.

Table C3: Descriptive Statistics of Annual Rescaled Shocks

	count	min	p10	p25	p50	p75	p90	max
TS-I	16721	-0.036	-0.000	-0.000	0.000	0.000	0.000	0.016
TS-I.b	16721	-0.001	-0.000	-0.000	-0.000	0.000	0.000	0.014
TS-E	16721	-0.074	-0.001	-0.001	-0.000	0.001	0.001	0.053
EG	16721	-0.154	-0.030	-0.006	0.005	0.012	0.020	0.230
EG, College	16721	-0.072	-0.023	-0.005	0.002	0.008	0.014	0.093
EG, No College	16721	-0.143	-0.036	-0.008	0.006	0.014	0.025	0.096

Notes: The table shows descriptive statistics for the annualized rescaled shocks. We rescale every shock so that they all have the same units, as described in Section 4. The annual shocks cover 1994 to 2016. All commuting zones (1990 vintage) in the US are covered, except for Alaska, which we aggregate to the state level due to inconsistencies with its underlying borough classification, resulting in 727 commuting zones per year. Underlying data are described in the Data Appendix.

Table C4: Demographics by Commuting Zone Frequency, Cohort Born in the 1940s

	Databank, Share of Years With Matched Commuting Zone										ACS
Demographics (Percents)	0-10	10-20	20-30	30-40	40-50	50-60	60-70	70-80	80-90	90-100	
Female	48	46	50	51	51	53	55	56	54	53	51
Male	52	54	50	49	49	47	45	44	46	47	49
Gender Missing	0	0	0	0	0	0	0	0	0	0	0
White, non-Hispanic	20	47	53	58	59	64	67	68	71	81	78
Black, non-Hispanic	5	15	13	13	13	11	11	10	11	7	10
Asian, non-Hispanic	3	3	4	4	4	4	5	5	4	3	3
Native American, non-Hispanic	0	0	0	0	0	0	0	0	0	0	1
Hispanic	9	11	11	12	12	11	11	11	10	7	7
Other	1	2	2	2	2	2	2	2	2	1	1
Race Missing	63	21	17	12	10	7	5	4	3	1	0
North	10	14	14	15	15	16	16	17	18	22	21
South	22	31	32	32	32	32	31	31	30	28	30
Midwest	10	16	17	19	19	20	20	20	22	27	25
West	7	11	11	11	11	11	12	12	13	11	12
US Territories	8	2	2	1	1	1	1	1	1	1	1
International	44	26	24	23	22	20	20	20	16	11	11
Birth Region Missing	0	0	0	0	0	0	0	0	0	0	0
Avg. Annual Wage	52910	28860	37680	38410	45480	47870	52230	53060	56640	59170	27000
Avg. Share of Years with CZ	0	14	24	34	42	51	65	75	85	100	
Observations (thousands)	7111	313	463	408	294	1012	826	815	1185	18680	636

Notes: The table displays demographic statistics binned by the share of years with matched commuting zones (CZs). The table is constructed from the 1989, 1994, 1995, and 1998-2019 Census Databank files, the 1990 and 2000 Decennial Census, and the 2001-2019 American Community Survey. Samples are restricted to people age 25 to 55, and the Databank sample is further restricted to people who lived past the age of 55. A person-year is considered matched if we can directly match it with a CZ through a 1040 form zip code or a Master Address File county. A person's share of years with a matched CZ is calculated as the percent of their years between ages 25 and 55 with a matched CZ. The *Native American, non-Hispanic* row includes non-Hispanic Native Americans, Native Alaskans, Pacific Islanders, and Hawaiians. Average Annual Wage is in nominal dollars.

Table C5: Demographics by Commuting Zone Frequency, Cohort Born in the 1980s

	Databank, Share of Years With Matched Commuting Zone										ACS
Demographics (Percents)	0-10	10-20	20-30	30-40	40-50	50-60	60-70	70-80	80-90	90-100	
Female	44	35	37	39	41	44	45	47	50	54	50
Male	56	65	63	61	59	56	55	53	50	46	50
Gender Missing	0	0	0	0	0	0	0	0	0	0	0
White, non-Hispanic	14	35	35	36	37	34	43	46	50	63	57
Black, non-Hispanic	7	20	18	17	16	13	15	15	15	9	13
Asian, non-Hispanic	2	2	3	3	3	3	4	5	5	5	6
Native American, non-Hispanic	0	1	0	0	0	0	0	0	0	0	1
Hispanic	10	16	14	13	13	12	15	16	16	14	20
Other	1	3	3	3	3	2	3	3	3	2	2
Race Missing	65	23	27	28	28	35	20	15	10	6	0
North	8	12	11	11	11	10	12	13	14	17	15
South	15	29	28	27	26	23	28	28	29	27	27
Midwest	8	15	14	14	14	13	16	16	18	22	20
West	9	18	16	16	15	14	17	18	19	19	18
US Territories	6	4	3	2	2	1	1	1	1	0	1
International	54	22	28	30	32	39	26	23	19	14	19
Birth Region Missing	0	0	0	0	0	0	0	0	0	0	0
Avg. Annual Wage	32760	22150	27200	30710	33950	37200	37350	40270	43410	59700	32200
Avg. Share of Years with CZ	1	14	25	34	43	53	65	74	85	99	
Observations (thousands)	5676	1444	1394	1464	1358	2527	2232	2480	4131	26360	3884

Notes: The table displays demographic statistics binned by the share of years with matched commuting zones (CZs). The table is constructed from the 1989, 1994, 1995, and 1998-2019 Census Databank files, the 1990 and 2000 Decennial Census, and the 2001-2019 American Community Survey. Samples are restricted to people age 25 to 55, and the Databank sample is further restricted to people who lived past the age of 55. A person-year is considered matched if we can directly match it with a CZ through a 1040 form zip code or a Master Address File county. A person's share of years with a matched CZ is calculated as the percent of their years between ages 25 and 55 with a matched CZ. The *Native American, non-Hispanic* row includes non-Hispanic Native Americans, Native Alaskans, Pacific Islanders, and Hawaiians. Average Annual Wage is in nominal dollars.

Table C6: Regressing Shocks on Each Other

	TS-I		EG		TS-E	
	(1)	(2)	(3)	(4)	(5)	(6)
TS-I			0.000272	-0.00132	-0.316***	-0.241***
			(0.112)	(0.0117)	(0.0531)	(0.0375)
EG	0.000258	-0.00461			0.221**	0.166*
	(0.106)	(0.0407)			(0.0955)	(0.0821)
TS-E	-0.334***	-0.247**	0.247**	0.0489***		
	(0.0929)	(0.0985)	(0.113)	(0.0156)		
CZ & year FEs?	✗	✓	✗	✓	✗	✓
N	16721	16721	16721	16721	16721	16721
R^2	0.112	0.256	0.0609	0.787	0.160	0.276

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows how correlated shocks are with each other, reporting estimates from the equation $shock_{it} = \beta_0 + \sum_k \beta_k shock^k_{i,t} + \Theta_i + \Theta_t + \epsilon_{it}$, where *shock* is one of the annualized shocks in the master data set and $shock^k$ are the remaining other panel shocks. Θ_i is a commuting zone fixed effect, Θ_t is a year fixed effect, and β_0 is a constant. All shocks are converted into z-scores. Standard errors are clustered by year by commuting zone, and observations are not weighted. Data are from the shocks master data set, detailed in the Data Appendix. We run the regression on shocks to all commuting zones in the US (except for Alaska, which is aggregated to the state level) over the years 1994 to 2016.

Table C7: Ranked Correlations

Panel A: CZ-Level Ranked Correlations

Shock	TS-I	TS-I.b	TS-E	EG	EG-C	EG-M
TS-I	1					
TS-I.b	0.89	1				
TS-E	−0.45	−0.41	1			
EG	−0.02	−0.05	0.17	1		
EG-C	−0.08	−0.14	0.13	0.68	1	
EG-M	0.00	−0.02	0.15	0.69	0.44	1
TG	0.75	0.81	0.51	0.80	0.22	0.50
GR	−0.05	0.07	0.18	0.54	0.18	0.41

Panel B: State-Level Ranked Correlations

Shock	TS-I	TS-I.b	TS-E	EG	EG-C	EG-M
TS-I	1					
TS-I.b	0.86	1				
TS-E	−0.50	−0.46	1			
EG	−0.01	−0.05	0.20	1		
EG-C	−0.09	−0.19	0.16	0.77	1	
EG-M	0.01	−0.03	0.18	0.71	0.50	1
TG	0.83	0.93	0.64	0.74	0.56	0.81
GR	0.06	0.25	0.30	0.56	0.32	0.56

Notes: The table displays the correlation matrix for the annual and time-invariant shocks. The correlations are calculated over the entire sample period, 1994-2016, for the annual shocks and over their respective years for the time-invariant shocks. We compare the tariff gap shock to the 2000, 2001, and 2002 annual shocks. We compare the Great Recession shock to 2008, 2009, and 2010 annual shocks. All shocks are converted into z-scores, and they are standardized so that negative shocks are associated with decreases in demand. Panel A displays correlations for shocks at the CZ level while Panel B shows the correlations for state-level shocks.

Table C8: 5-Year Lagged Autocorrelations

	(1) TS-I	(2) EG	(3) TS-E
Shock Lagged 1 Year	0.0251	0.155**	-0.0406
	(0.0629)	(0.0732)	(0.0893)
Shock Lagged 2 Years	-0.00786	0.0461	-0.140*
	(0.0323)	(0.0490)	(0.0794)
Shock Lagged 3 Years	-0.0580	-0.0102	-0.0904
	(0.0442)	(0.0302)	(0.0873)
Shock Lagged 4 Years	-0.0331	-0.120**	-0.0833
	(0.0606)	(0.0479)	(0.0742)
Shock Lagged 5 Years	-0.112	0.0103	-0.239
	(0.0719)	(0.0618)	(0.251)
Observations	13086	13086	13086
R^2	0.214	0.792	0.287

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows how correlated shocks are with themselves, following the regression specified by Equation (2). We regress shocks on 1-, 2-, 3-, 4-, and 5-year lagged shocks and a constant. We include CZ and year fixed effects in each specification. All shocks are converted into z-scores. Data are from the shocks master data set, detailed in the Data Appendix. We run the regression on shocks to all commuting zones in the US (except for Alaska, which is aggregated to the state level) over the years 1994 to 2016.

Table C9: Employment growth Shock Autocorrelations Over Different Time Periods

	1 Year	2 Year	3 Year	4 Year	5 Year	6 Year	7 Year
1-Period Shock Lag	0.164***	0.171***	0.0677*	-0.322*	-0.236**	-0.229***	-0.413***
	(0.0225)	(0.0299)	(0.0356)	(0.169)	(0.0940)	(0.0588)	(0.0878)
Observations	14540	7997	5089	3635	2908	2181	2181
R^2	0.780	0.803	0.755	0.219	0.371	0.740	0.396

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table shows how correlated employment growth shocks are with themselves, varying the horizon over which the employment growth shock is constructed. We regress the employment growth shock on a one-period lag of the employment growth shock, with each period consisting of one to seven years, as indicated. We include CZ and year fixed effects in each specification. All shocks are converted into z-scores. We run the regression using shocks to all commuting zones in the US (except for Alaska, which is aggregated to the state level) over the years 1994 to 2016.

Table C10: The Effect of Local Shocks on Out-Migration

	(1) TS-I	(2) EG	(3) TS-E
	Out-Migration Rate		
Shock to CZ	–0.005	–0.129***	–0.005
	(0.008)	(0.043)	(0.007)
Shock to CZ, 1 year lag	0.002	–0.077	–0.003
	(0.008)	(0.057)	(0.007)
Shock to CZ, 2 year lag	0.005	–0.020	–0.001
	(0.011)	(0.045)	(0.011)
Shock to CZ, 3 year lag	0.004	–0.010	0.007
	(0.008)	(0.041)	(0.012)
Shock to CZ, 4 year lag	0.016*	0.033	0.014
	(0.009)	(0.035)	(0.015)
Shock to CZ, 5 year lag	0.016	0.077**	0.026
	(0.009)	(0.031)	(0.020)
Average shock to other CZs	–0.037	0.320	–0.083***
	(0.030)	(0.193)	(0.026)
Average shock to other CZs, 1 year lag	–0.052*	0.228	–0.078***
	(0.029)	(0.212)	(0.015)
Average shock to other CZs, 2 year lag	–0.074	–0.159	–0.099***
	(0.048)	(0.204)	(0.020)
Average shock to other CZs, 3 year lag	–0.048	0.174	–0.047
	(0.040)	(0.171)	(0.034)
Average shock to other CZs, 4 year lag	–0.058*	0.203	0.004
	(0.030)	(0.196)	(0.055)
Average shock to other CZs, 5 year lag	–0.048	0.461**	0.005
	(0.033)	(0.163)	(0.052)
Constant	3.771***	3.875***	3.689***
	(0.022)	(0.112)	(0.006)
N	12359	12359	12359
r2	0.8957	0.8981	0.897

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table displays the output from regressing out-migration rates on a local shock, one- to five-year lags of that shock, a weighted average of the shock to other CZs, and one- to five-year lags of that weighted average. We include a constant, commuting zone fixed effects, and year fixed effects. Standard errors are clustered by commuting zone and by year. Migration data are described in Section 2.2, and data used to construct the shocks are described in the Data Appendix. To put the coefficients in perspective, the average migration rate over our sample period is 3.818 percent.

Table C11: The Effect of Local Shocks on In-Migration

	(1) TS-I	(2) EG	(3) TS-E
	In-Migration Rate		
Shock to CZ	0.015	0.170***	–0.003
	(0.013)	(0.034)	(0.009)
Shock to CZ, 1 year lag	0.010	0.172***	0.005
	(0.007)	(0.051)	(0.016)
Shock to CZ, 2 year lag	0.006	0.124**	–0.005
	(0.008)	(0.054)	(0.016)
Shock to CZ, 3 year lag	–0.002	0.107*	0.002
	(0.008)	(0.055)	(0.012)
Shock to CZ, 4 year lag	–0.013	0.066	0.006
	(0.011)	(0.040)	(0.009)
Shock to CZ, 5 year lag	–0.012	0.037	0.008
	(0.013)	(0.040)	(0.009)
Average shock to other CZs	–0.020	0.084	0.017
	(0.017)	(0.107)	(0.019)
Average shock to other CZs, 1 year lag	0.008	0.077	0.037
	(0.020)	(0.107)	(0.024)
Average shock to other CZs, 2 year lag	0.009	0.135	0.021
	(0.026)	(0.096)	(0.024)
Average shock to other CZs, 3 year lag	0.036*	0.134	0.052
	(0.020)	(0.110)	(0.030)
Average shock to other CZs, 4 year lag	0.031	–0.033	0.026
	(0.028)	(0.156)	(0.030)
Average shock to other CZs, 5 year lag	0.047	–0.020	0.026
	(0.031)	(0.153)	(0.027)
Constant	3.661***	3.833***	3.684***
	(0.016)	(0.069)	(0.003)
N	12359	12359	12359
r2	0.9325	0.937	0.9324

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table displays the output from regressing in-migration rates on a local shock, one- to five-year lags of that shock, a weighted average of the shock to other CZs, and one- to five-year lags of that weighted average. We include a constant, commuting zone fixed effects, and year fixed effects. Standard errors are clustered by commuting zone and by year. Migration data are described in Section 2.2, and data used to construct the shocks are described in the Data Appendix. To put the coefficients in perspective, the average migration rate over our sample period is 3.818 percent.

Table C12: Migration Inflow Heterogeneity: Panel Shocks

Panel A: TS-I Shock Migration Inflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	0.011	0.014	0.004	0.000	0.012	0.008	0.017	0.008
	(0.015)	(.011)	(0.013)	(0.008)	(0.010)	(0.009)	(0.012)	(0.008)
Moderator	0.005	0.001	0.023	0.023	0.004	0.063	–0.005	0.013
Interaction	(0.018)	(0.020)	(0.025)	(0.014)	(0.013)	(0.048)	(0.016)	(0.010)
N	13086	13086	13086	13086	13086	13086	13086	13086

Panel B: EG Shock Migration Inflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	0.179***	0.206***	0.207***	0.158***	0.213***	0.199***	0.163***	0.196***
	(0.048)	(0.044)	(0.040)	(0.040)	(0.039)	(0.040)	(0.040)	(0.041)
Moderator	0.026	–0.060***	–0.034	0.051***	–0.029*	–0.011	0.045**	-0.002
Interaction	(0.021)	(0.016)	(0.028)	(0.018)	(0.017)	(0.026)	(0.021)	(0.012)
N	13086	13086	13086	13086	13086	13086	13086	13086

Panel C: TS-E Shock Migration Inflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	–0.002	0.016	–0.002	–0.001	0.015	0.003	–0.022	-0.002
	(0.023)	(0.017)	(0.014)	(0.011)	(0.025)	(0.012)	(0.014)	(0.012)
Moderator	0.005	–0.030	0.009	0.005	–0.020	–0.010	0.049**	0.013
Interaction	(0.023)	(0.022)	(0.024)	(0.017)	(0.021)	(0.043)	(0.020)	(0.011)
N	13086	13086	13086	13086	13086	13086	13086	13086

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above panels report coefficient values for β_1 and β_3 from equation (11), representing the effect of the shock and the shock-moderator interaction on one-year inflow rates. Moderators are defined in Table A1 and discussed in Appendix A.5.

Table C13: Migration Outflow Heterogeneity: Panel Shocks

Panel A: TS-I Shock Migration Outflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	–0.008	–0.006	0.024	0.014	0.014	0.011	0.006	0.013
	(0.013)	(0.009)	(0.016)	(0.016)	(0.016)	(0.011)	(0.009)	(0.011)
Moderator	0.017	0.030	–0.047*	–0.018	–0.014	–0.092	–0.009	-0.023
Interaction	(0.016)	(0.021)	(0.027)	(0.021)	(0.021)	(0.072)	(0.014)	(0.021)
N	13086	13086	13086	13086	13086	13086	13086	13086

Panel B: EG Shock Migration Outflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	–0.100**	–0.109**	–0.108**	–0.101*	–0.088**	–0.100**	–0.119***	–0.105**
	(0.045)	(0.049)	(0.047)	(0.045)	(0.046)	(0.046)	(0.045)	(0.045)
Moderator	–0.014	–0.003	–0.004	–0.014	–0.028	–0.025	0.016	-0.009
Interaction	(0.024)	(0.033)	(0.032)	(0.036)	(0.027)	(0.045)	(0.012)	(0.020)
N	13086	13086	13086	13086	13086	13086	13086	13086

Panel C: TS-E Shock Migration Outflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Shock	0.017	–0.018	0.004	0.006	0.002	0.002	–0.010	0.003
	(0.014)	(0.016)	(0.013)	(0.012)	(0.013)	(0.011)	(0.010)	(0.012)
Moderator	–0.036	0.022	–0.030	–0.025	–0.015	–0.069	0.005	-0.035
Interaction	(0.023)	(0.027)	(0.029)	(0.022)	(0.016)	(0.065)	(0.019)	(0.029)
N	13086	13086	13086	13086	13086	13086	13086	13086

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above panels report coefficient values for β_1 and β_3 from equation (11), representing the effect of the shock and the shock-moderator interaction on one-year outflow rates. Moderators are defined in Table A1 and discussed in Appendix A.5.

Table C14: Migration Inflow Heterogeneity: Cross-sectional Shocks

Panel D: TG Shock Migration Inflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Post x Shock	–0.027	–0.037	0.020	–0.104***	–0.051**	–0.079***	–0.093**	0.041
	(0.043)	(0.053)	(0.040)	(0.025)	(0.024)	(0.028)	(0.045)	(0.029)
Moderator	0.018	–0.054	–0.180***	0.056	0.026	0.012	0.100	0.009
Interaction	(0.059)	(0.066)	(0.069)	(0.061)	(0.050)	(0.119)	(0.062)	(0.035)
N	8724	8724	8724	8724	8724	8724	8724	8724

Panel E: GR Shock Migration Inflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Post x Shock	0.286***	0.306***	0.194***	0.159***	0.339***	0.181***	0.062	0.222***
	(0.087)	(0.072)	(0.073)	(0.050)	(0.073)	(0.038)	(0.052)	(0.048)
Moderator	–0.171*	–0.289***	–0.052	0.055	–0.178*	0.049	0.214**	-0.016
Interaction	(0.103)	(0.079)	(0.092)	(0.087)	(0.095)	(0.122)	(0.084)	(0.028)
N	11632	11632	11632	11632	11632	11632	11632	11632

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above panels report coefficient values for β_1 and β_3 from equation (12), representing the effect of the shock and the shock-moderator interaction on one-year inflow rates. Moderators are defined in Table A1 and discussed in Appendix A.5.

Table C15: Migration Outflow Heterogeneity: Cross-sectional Shocks

Panel D: TG Shock Migration Outflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Post x Shock	–0.013	–0.034	0.017	0.104***	0.115***	0.066**	–0.028	–0.071**
	(0.066)	(0.063)	(0.035)	(0.027)	(0.038)	(0.033)	(0.079)	(0.029)
Moderator	0.034	0.065	–0.030	–0.167	–0.150	–0.212*	0.049	0.084***
Interaction	(0.074)	(0.083)	(0.117)	(0.070)	(0.063)	(0.128)	(0.087)	(0.022)
N	8724	87241	8724	8724	8724	8724	8724	8724

Panel E: GR Shock Migration Outflows Moderators

	Employment Rate	Share Local Born	Share in Top 20 Percent	College Graduates	College Access	High CZ Population	January Temp	Home Value
Post x Shock	0.004	0.034	–0.030	0.014	–0.007	–0.035	–0.019	–0.045
	(0.103)	(0.051)	(0.072)	(0.032)	(0.048)	(0.049)	(0.037)	(0.039)
Moderator	0.015	–0.069	0.0260***	0.029	0.008	0.077	0.053	0.024
Interaction	(0.110)	(0.093)	(0.089)	(0.065)	(0.073)	(0.080)	(0.090)	(0.034)
N	11632	11632	11632	11632	11632	11632	11632	11632

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above panels report coefficient values for β_1 and β_3 from equation (12), representing the effect of the shock and the shock-moderator interaction on one-year outflow rates. Moderators are defined in Table A1 and discussed in Appendix A.5.

Table C16: Migration Inflow and Outflow Population Subgroups: Time-Invariant Shocks

Panel A: TG Shock Migration Outflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	–0.031	–0.060	0.035	0.297*	–0.098	–0.151	0.010	–0.046	–0.028	–0.031	0.046	–0.074
	(0.155)	(0.169)	(0.157)	(0.170)	(0.182)	(0.134)	(0.149)	(0.215)	(0.174)	(0.169)	(0.176)	(0.126)
Constant	11.39***	12.18***	11.94***	15.05***	17.54***	7.17***	12.47***	12.81***	8.309***	15.22***	14.11***	10.30***
	(0.059)	(0.066)	(0.069)	(0.061)	(0.050)	(0.119)	(0.062)	(0.035)	(0.035)	(0.035)	(0.035)	(0.035)
N	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724

Panel B: GR Shock Migration Outflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	0.403***	0.475***	0.410***	0.614***	0.281*	0.230*	0.394***	0.161	0.163	0.760***	0.606***	0.059
	(0.149)	(0.145)	(0.143)	(0.166)	(0.169)	(0.120)	(0.125)	(0.179)	(0.134)	(0.188)	(0.170)	(0.134)
Constant	11.85***	12.65***	12.47***	15.10***	17.95***	7.25***	13.02***	13.03***	8.71***	15.89***	14.72***	10.31***
	(0.031)	(0.030)	(0.030)	(0.035)	(0.035)	(0.025)	(0.026)	(0.037)	(0.028)	(0.039)	(0.035)	(0.028)
N	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451

Panel C: TG Shock Migration Inflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	–0.295*	–0.250	–0.220	0.362	–0.378*	–0.410***	–0.120**	–0.344	-0.500**	-0.202	0.033	–0.550***
	(0.161)	(0.174)	(0.151)	(0.287)	(0.206)	(0.154)	(0.154)	(0.228)	(0.226)	(0.187)	(0.186)	(0.191)
Constant	11.34***	12.14***	11.71***	15.64***	17.47***	7.11***	12.17***	13.02***	9.20***	15.46***	14.11***	10.77***
	(0.011)	(0.012)	(0.010)	(0.019)	(0.014)	(0.010)	(0.010)	(0.015)	(0.015)	(0.013)	(0.012)	(0.013)
N	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724	8724

Panel D: GR Shock Migration Inflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	0.023	0.009	–0.073	0.498*	–0.260	–0.046	–0.150	0.091	0.309**	–0.113	0.328***	–0.469***
	(0.096)	(0.117)	(0.019)	(0.301)	(0.164)	(0.071)	(0.109)	(0.158)	(0.152)	(0.222)	(0.025)	(0.117)
Constant	11.67***	12.44***	12.10***	15.22***	17.72***	7.11***	12.55***	13.06***	9.46***	15.93***	14.64***	10.52***
	(0.0204)	(0.025)	(0.019)	(0.064)	(0.035)	(0.015)	(0.023)	(0.034)	(0.032)	(0.047)	(0.025)	(0.025)
N	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451	9451

Notes: The above table shows the impacts of a time invariant shock on migration flows for a specific subgroup. Subgroups are defined using 1040 IRS tax data or data from the Decennial Census. Migration flows are measured for each subgroup prior to running the regression. Migration flows are measured across a five-year period. For all time-invariant shocks we create a dummy variable that takes a value of 1 if the measured migration year is within 10 years of the shock. This is then interacted with the standardized shock for each CZ to produce the post-shock variable used in the regression. Note local born is defined as an individual who resides in the same state they were born in. AGI bottom includes earners in the bottom third of adjusted gross income on their 1040 tax return; AGI top includes earners above the 80th percentile.

Table C17: Migration Inflow Population Subgroups: Time-Variant Shocks

Panel A: TS-I Shock Migration Inflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	–0.018	–0.015	–0.020	–0.002	–0.003	–0.020	–0.027	–0.016	0.009	–0.046	–0.011	–0.037
	(0.022)	(0.022)	(0.021)	(0.042)	(0.025)	(0.018)	(0.024)	(0.034)	(0.018)	(0.031)	(0.019)	(0.034)
Constant	11.56***	12.38***	11.97***	15.53***	17.64***	7.185***	12.43***	13.07***	9.38 ***	15.77***	14.37***	10.81***
	(0.021)	(0.022)	(0.023)	(0.030)	(0.028)	(0.019)	(0.028)	(0.0345)	(0.033)	(0.028)	(0.017)	(0.033)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359

Panel B: EG Shock Migration Inflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	0.376***	0.449***	0.336**	0.693**	0.578***	0.308**	0.415***	0.493**	0.455*	0.503***	0.424**	0.358**
	(0.125)	(0.153)	(0.133)	(0.253)	(0.173)	(0.111)	(0.135)	(0.182)	(0.227)	(0.142)	(0.160)	(0.162)
Constant	11.74***	12.55***	12.1***	14.89***	17.88***	7.429***	12.39***	13.55***	9.777***	16.06***	14.56***	10.85***
	(0.168)	(0.179)	(0.175)	(0.369)	(0.222)	(0.106)	(0.201)	(0.247)	(0.187)	(0.221)	(0.203)	(0.181)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359

Panel C: TS-E Shock Migration Inflows

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	0.007	0.012	0.010	0.030	0.006	0.006	0.013	0.012	0.023	0.010	–0.002	0.013
	(0.027)	(0.030)	(0.022)	(0.050)	(0.041)	(0.019)	(0.019)	(0.041)	(0.031)	(0.044)	(0.031)	(0.030)
Constant	11.61***	12.43***	12.03***	15.64***	17.71***	7.237***	12.52***	13.14***	9.401***	15.86***	14.42***	10.92***
	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)	(0.002)	(0.004)	(0.006)	(0.005)	(0.005)	(0.002)	(0.008)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above table shows the impacts of a time variant shock on migration flows for a specific subgroup. These specifications include in-migration rates regressed on a local shock, one-year lags of that shock, a weighted average of the shock to other CZs, and one-year lags of that weighted average. Subgroups are defined using 1040 IRS tax data or data from the Decennial Census. Migration flows are measured for each subgroup prior to running the regression. Migration flows are measured across a five-year period. For all time-variant shocks we create a dummy variable that takes a value of 1 if the measured migration year is within 10 years of the shock. This is then interacted with the standardized shock for each CZ to produce the post-shock variable used in the regression. Note local born is defined as an individual who resides in the same state they were born in. AGI bottom includes earners in the bottom third of adjusted gross income on their 1040 tax return; AGI top includes earners above the 80th percentile.

Table C18: Migration Outflow Population Subgroups: Time-Variant Shocks

	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Panel A: TS-I Shock Migration Outflows												
Post Shock	0.010	0.005	0.006	–0.017	0.013	0.001	0.010	–0.039*	0.001	–0.004	0.007	–0.006
	(0.018)	(0.020)	(0.018)	(0.022)	(0.024)	(0.012)	(0.024)	(0.021)	(0.013)	(0.033)	(0.022)	(0.014)
Constant	11.78***	12.6***	12.39***	15.19***	17.88***	7.378***	12.98***	12.99***	8.651***	15.73***	14.49***	10.55***
	(0.034)	(0.034)	(0.031)	(0.034)	(0.039)	(0.030)	(0.028)	(0.036)	(0.028)	(0.050)	(0.032)	(0.028)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359
Panel B: EG Shock Migration Outflows												
	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	–0.154	–0.156	–0.157	–0.325**	–0.145	–0.060	–0.068	–0.029**	–0.123	–0.195	–0.160	–0.107
	(0.129)	(0.148)	(0.121)	(0.126)	(0.148)	(0.113)	(0.132)	(0.133)	(0.107)	(0.175)	(0.130)	(0.113)
Constant	12.04***	12.84***	12.56***	15.01***	18.12***	7.715***	12.91***	13.56***	9.024***	16.02***	14.82***	10.39***
	(0.218)	(0.214)	(0.198)	(0.235)	(0.243)	(0.177)	(0.172)	(0.313)	(0.176)	(0.275)	(0.225)	(0.143)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359
Panel C: TS-E Shock Migration Outflows												
	Female	Male	US Born	Foreign Born	Age 25-35	Age 45-55	White	Non-White	Local Born	Non-Local Born	AGI Bottom	AGI Top
Post Shock	0.019	0.025	0.022	0.042	0.026	0.020	0.029	0.037*	0.020	0.031	0.012	0.019
	(0.021)	(0.021)	(0.024)	(0.027)	(0.018)	(0.020)	(0.030)	(0.018)	(0.022)	(0.030)	(0.014)	(0.023)
Constant	11.7***	12.52***	12.31***	15.17***	17.79***	7.324***	12.86***	12.96***	8.593***	15.64***	14.43***	10.46***
	(0.009)	(0.008)	(0.008)	(0.001)	(0.007)	(0.007)	(0.008)	(0.010)	(0.009)	(0.009)	(0.008)	(0.005)
N	12359	12359	12359	12500	12359	12359	12359	12359	12359	12359	12359	12359

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The above table shows the impacts of a time-variant shock on migration flows for a specific subgroup. These specifications include out-migration rates regressed on a local shock, one-year lags of that shock, a weighted average of the shock to other CZs, and one-year lags of that weighted average. Subgroups are defined using 1040 IRS tax data or data from the Decennial Census. Migration flows are measured for each subgroup prior to running the regression. Migration flows are measured across a five-year period. For all time-variant shocks we create a dummy variable that takes a value of 1 if the measured migration year is within 10 years of the shock. This is then interacted with the standardized shock for each CZ to produce the post-shock variable used in the regression. Note local born is defined as an individual who resides in the same state they were born in. AGI bottom includes earners in the bottom third of adjusted gross income on their 1040 tax return; AGI top includes earners above the 80th percentile.