

Good Jobs, Bad Jobs: What's Trade Got To Do With It?

Jame Lake*

Southern Methodist University

Daniel L. Millimet[†]

Southern Methodist University & IZA

June 2, 2018

Abstract

We analyze whether local exposure to rising Chinese import penetration or routine-biased-technological-change (RBTC) can explain the extent of job polarization experienced by US local labor markets between 1990 and 2010. By defining jobs at a very disaggregate level, we have three main results. First, local RBTC exposure can explain job polarization (i.e., lower employment growth of ‘middle-quality’ jobs than ‘high-quality’ and ‘low-quality’ jobs) in US local labor markets; local exposure to rising Chinese import competition, if anything, actually engenders an anti-polarization effect. Second, local labor markets only experience job polarization when sufficiently exposed to RBTC. Third, local exposure rather than (national) occupation-specific exposure drives the degree of job polarization. This contrast has eluded prior studies conducted at the local labor market level; our analysis, utilizing location-job pairs as the unit of observations, is the first study to disentangle the impacts of local versus occupational exposure to shocks.

JEL: F13, J21, J31

Keywords: Trade Liberalization, China, Local Labor Markets, Job Polarization

*We are extremely grateful to the editor and two anonymous referees. We also thank Douglass Campbell, Maggie Chen, Martin Davies, David Drukker, Mina Kim, Bill Powers, Erick Sager, Andrey Stoyanov and Beyza Ural for useful comments and discussion as well as many seminar and conference participants at the DC Junior Trade Study Group, Fall 2015 Midwest Trade Meetings, Stata Texas Empirical Microeconomics conference 2015, George Washington University, University of Nevada Las Vegas, University of Richmond, USITC, and Washington and Lee University. Corresponding author: James Lake, Department of Economics, Box 0496, Southern Methodist University, Dallas, TX 75275-0496, USA. E-mail: jlake@smu.edu.

[†]Department of Economics, Southern Methodist University, Box 0496, Dallas, TX, 75205-0496. Email: millimet@smu.edu.

1 Introduction

Recent years have witnessed a renewed interest in two issues concerning the US labor market. The first issue is the impact of trade and technology on labor market outcomes. On one hand, the increased economic and political clout of China and the potential for trade deals of unprecedented size (e.g., the Trans-Pacific Partnership and the Transatlantic Trade and Investment Partnership) has raised concerns about the possible negative impacts of trade on labor market outcomes. On the other hand, automation and the growing usage of industrial robots has raised concerns about the possible negative impacts of technology on labor market outcomes. While an earlier literature found that technology rather than trade primarily explained various labor market phenomena, including rising wage inequality between skilled and unskilled workers, a recent literature documents substantial labor market impacts of trade stemming from China’s rapid accession in international markets since 1990.

The second issue is the disappearance of middle class jobs. Together with the relative rise in employment of low-skill and high-skill jobs, this has been labelled the ‘dumbbell’ or ‘hourglass’ economy in the popular press and job polarization in academia (Goos and Manning (2007); Samuel (2013)). Acemoglu and Autor (2011, p.1046) state that US and European Union labor markets have undergone “systematic, non-monotonic shifts in the composition of employment across occupations” resulting in “rapid simultaneous growth of both high education, high wage occupations and low education, low wage occupations.” In the language of Goos and Manning (2007), there has been simultaneous growth in “lousy” jobs and “lovely” jobs and a decline in “middling” jobs.

In this paper, we investigate the impact of exposure to trade and technology on the allocation of workers across jobs in the US and how these impacts vary with job quality as measured by job-specific wage and educational attainment.¹ We are particularly interested in whether exposure to trade or technology can explain job polarization defined as lower employment growth in middle-quality jobs relative to both low-quality and high-quality jobs. Specifically, we analyze the heterogeneous effects of *local* exposure to trade and technology on *local* employment growth between 1990 and 2010 across the job quality distribution. Here, local employment growth is the location-job specific change in the employment to working age population ratio between 1990 and 2010. By defining jobs at a very disaggregate level, we assess how trade and technology exposure

¹The notion of job quality is not intended to carry any normative connotations, but is rather a convenient way to describe a job’s position in the distribution of wages and/or education.

differentially affects local employment growth of ‘good’ versus ‘bad’ versus ‘middling’ jobs. To this end, we merge central insights from the two aforementioned literatures.

From the job polarization literature, we borrow the insight that employment growth can vary in interesting ways across the distribution of job quality. As in Goos and Manning (2007), we define a job as a detailed occupation within a broad industry. From our IPUMS Census data (Ruggles et al. (2017)), we use 381 detailed 3-digit IPUMS Census occupations and eight broad 1-digit NAICS industries, yielding 3048 possible jobs. Roughly speaking, we define a job’s quality as its percentile in the joint distribution of job-specific median wage and median education. A job’s median wage and median educational attainment have been used as independent measures of job quality (e.g., Autor et al. (2006)). Our approach of incorporating both into a single measure of job quality is motivated by the earlier quote in Acemoglu and Autor (2011, p.1046) describing high wage and high education jobs on one hand and low wage and low education jobs on the other hand.

From the literature focusing on the labor market impacts of trade and technology, we borrow the insight that local labor markets offer an appropriate setting to investigate the impacts of trade exposure. A central insight from Autor et al. (2013) and Autor and Dorn (2013) is that workers appear largely immobile across the 722 Commuting Zones (CZs) that span the continental US and, in turn, local labor market shocks generate labor market adjustment within CZs. Thus, our analysis focuses on the labor market impacts within CZs of local exposure to trade and technology.

As has become standard in the recent trade and labor literature, our measure of rising trade exposure revolves around the dramatic surge in imports from China. Further, to capture the supply side forces in China that have driven China’s dramatic rise in international markets, we instrument rising Chinese import penetration in the US with Chinese export growth to other high income countries. However, we borrow an important insight from Ebenstein et al. (2014, 2015) who find a worker’s occupational exposure to globalization is far more important than their industry exposure. Indeed, an occupational view of jobs, rather than an industry view, fits well with the empirical definition of a job used in the job polarization literature. Thus, we base our measure of local exposure to rising import competition from China on *occupational* exposure to rising Chinese import penetration and the occupational composition of CZs.

Similarly, our measure of local exposure to technological change revolves around occupational exposure to technological change and the occupational composition of CZs. We use data from Autor and Dorn (2013), in turn based on task data from the US Dictionary of Occupational Titles, that measures the extent of “routine-task intensity” for 3-digit occupations to compute the routine-task

intensive share of employment in each CZ. That is, our measure of local exposure to technological change captures the extent to which employment in a CZ is exposed to routine-biased technological change (RBTC). Indeed, according to Goos et al. (2014, p.2511), “... the literature seems to be settling on using the RTI [routine task intensity] measure as the best way to capture the impact of recent technological progress”. Further following Autor and Dorn (2013) and Autor et al. (2015), we instrument for CZ routine employment shares using the occupational composition of CZs in 1950. As shown by Autor and Dorn (2013), a CZ’s occupational composition in 1950, a period well before the dramatic rise of computerization, is a good predictor for the CZ’s adoption of workplace computers beginning in the 1980s and the associated displacement of routine-task intensive occupations. Important for our ability to distinguish between the impacts of local trade exposure and local RBTC, our data shows that the CZs strongly vulnerable to rising import competition tend not to be the CZs strongly vulnerable to RBTC.

Our first main result is that local exposure to RBTC, but not local exposure to rising import competition, drives job polarization. When comparing job-specific employment growth across CZs with different degrees of local exposure to rising Chinese import penetration, we find weaker employment growth in CZs facing stronger rises in Chinese import penetration. Moreover, we find that rising Chinese import penetration does not drive job polarization because the size of this depression is not larger for middle-quality jobs than low-quality and high-quality jobs. Indeed, if anything, this depression is actually *smaller* for middle-quality jobs than low- and high-quality jobs and thus engenders an “anti-polarization” effect. Ultimately, CZs facing stronger rises in Chinese import penetration have depressed employment growth across the job quality spectrum.

In contrast, the impact of local RBTC exposure varies noticeably across the job quality spectrum. When comparing job-specific employment growth across CZs with different degrees of local RBTC exposure, we find that CZs facing stronger exposure have stronger employment growth in both low-quality and high-quality jobs but weaker employment growth in middle-quality jobs. That is, differing degrees of RBTC across CZs explain differing degrees of job polarization across CZs.

Our second main result stems from our analysis at the CZ-job level. Unlike analyses at the national level or CZ level, this allows us to compare employment growth of different jobs within CZs for CZs facing weak RBTC exposure. Doing this reveals that such CZs actually display an “anti-polarization” pattern whereby employment growth is stronger in middle-quality jobs than both low- and high-quality jobs. In turn, our second main result is that CZs must be *sufficiently* exposed to RBTC for the emergence of job polarization at the CZ level.

Our third main result addresses whether our findings on job-specific employment growth within CZs are driven by local RBTC exposure and local trade exposure or, instead, by (national) occupation-specific RBTC exposure and (national) occupation-specific trade exposure. Fortunately for our ability to distinguish between local impacts and occupation-specific impacts of exposure to RBTC and rising import competition, our data reveal that local RBTC exposure and local trade exposure are essentially uncorrelated with occupation-specific RBTC exposure and occupation-specific trade exposure. Moreover, we find that local vulnerability drive our earlier results, with occupation-specific exposure having no statistically meaningful impact on the extent of polarization. Thus, local vulnerability and occupation-specific vulnerability have quite different implications for the reallocation of workers within CZs.

Our paper most closely relates to Autor et al. (2015) who also investigate the impacts of local exposure to import competition and RBTC on CZ employment growth.² Relative to their work, we make three contributions. These contributions stem from our focus on (i) the heterogeneous effects of local exposure to trade and RBTC according to job quality rather than occupation and (ii) employment growth for detailed jobs within CZs rather than CZ-level employment growth.

First, our analysis investigates job polarization while Autor et al. (2015) investigate occupational polarization, where occupational polarization is the phenomena of manual-intensive and abstract-intensive occupations having higher employment growth than routine-intensive occupations. Intuitively, job polarization and occupational polarization would be equivalent if routine-intensive occupations were equivalent to middle-quality jobs, while manual-intensive (abstract-intensive) occupations were equivalent to low (high) quality jobs. However, our data reflect a more nuanced mapping. To this end, define “low”, “mid” and “high” quality jobs as those in, respectively, the bottom 25%, middle 50% and top 25% of the job quality distribution. For manual- and routine-intensive occupations, respectively, low quality jobs are 52% and 21% of employment while middle quality jobs are 47% and 76% of employment. For abstract-intensive occupations, middle quality jobs are 30% of employment and high quality jobs are 70% of employment. Thus, while manual-, routine- and abstract-intensive occupations tend to map to, respectively, low-, middle- and high-quality jobs, the mapping is far from exact. As such, establishing that a phenomena generates (does not generate) occupational polarization does not indicate whether this phenomena generates (does

²The attempt to empirically separate the impacts of trade and technology is a common theme dating back to at least the 1990s debate over trade versus skill biased technological change as drivers of developed country wage inequality. But, obviously, at least at a broad aggregate level, trade and technology are intimately related. See, for example, Acemoglu (2002), Acemoglu et al. (2015), Burstein and Vogel (2017) and Impullitti and Licandro (2018).

not generate) job polarization.

Nevertheless, our analysis shows the two are closely related. Autor et al. (2015) find that (i) higher local RBTC exposure positively impacts the relative degree of *occupational polarization* across CZs but (ii) regardless of *occupational task-intensity*, higher local trade exposure negatively impacts relative employment growth across CZs. Our analysis shows that (i) higher local RBTC exposure positively impacts the relative degree of *job polarization* across CZs but (ii) regardless of *job quality*, higher local trade exposure negatively impacts relative employment growth across CZs.

Our second contribution is the ability to assess whether CZs facing high local RBTC exposure *actually* experience job polarization as opposed to merely saying whether such CZs experience *greater* job polarization *relative* to CZs facing weak local RBTC exposure. By comparing employment growth at the CZ-level, Autor et al. (2015) conclude that higher local RBTC exposure positively impacts the relative degree of *occupational polarization* across CZs. However, they cannot state whether CZs facing high local RBTC exposure actually experience occupational polarization. To be specific, it could be that (i) CZs facing weak local RBTC exposure experience “anti-polarization,” whereby manual- and abstract-intensive occupations have *lower* employment growth than routine-intensive occupations, and (ii) CZs facing strong local RBTC exposure experience a weaker degree of anti-polarization. As our analysis focuses on employment growth of detailed jobs within a CZ, we can assess whether a weakly exposed CZ experiences job polarization or anti-polarization. Indeed, we find that CZs facing weak local exposure to RBTC and Chinese import penetration experience anti-polarization. Nevertheless, the impact of higher local RBTC exposure is sufficiently strong such that CZs facing strong local RBTC exposure not only experience job polarization *relative* to CZs facing weak local RBTC exposure, but also experience *actual* job polarization. That is, local RBTC exposure alone can account for the job polarization phenomena within CZs.

Our third contribution is that we disentangle whether the impact of RBTC and trade exposure on job-specific employment growth within CZs stems from exposure of the CZ or exposure of the occupation. By examining employment growth at the CZ-level, analyses like Autor et al. (2015) cannot investigate this distinction. Indeed, we find that local exposure rather than (national) occupation-specific exposure drives the degree of job polarization within CZs. In turn, this suggests the job polarization phenomena is intimately connected with local exposure, rather than occupational exposure, and that geographic immobility rather than occupational immobility drives job polarization in US local labor markets.

Three broader strands of the literature also motivate our analysis. First, beginning with Autor

et al. (2013), the literature documents adverse local labor market effects of the rapid rise in import penetration from China since the early 1990s. The key finding in Autor et al. (2013) is that this surge in Chinese import penetration accounts for roughly 25% of the decline in the US manufacturing sector over the 1990-2007 period. Subsequent studies have documented similar effects in Norway (Balsvik et al. (2015)), Germany (Dauth et al. (2014)), and Spain (Donoso et al. (2015)).³ Our analysis documents a broader decline in US employment growth across the entire distribution of job quality.⁴

Second, beginning with Goos and Manning (2007), RBTC has been offered as a primary explanation for job polarization. Rather than the traditional concept of skill-biased technological change, the authors compellingly argue that the RBTC hypothesis of Autor et al. (2003) provides a valid explanation for job polarization in the UK. Specifically, technological change has reduced labor demand for routine-tasks and jobs that use these tasks intensively have disappeared, leading workers to move to either low or high quality jobs. Among others, Autor et al. (2006) and Autor and Dorn (2013) have made a similar argument for job polarization in the US and Goos et al. (2014) for Western Europe more generally. Offshoring has been posited as a possible secondary explanation in the literature, yet the typical conclusion here is that of Goos et al. (2014, p. 14): “RBTC is much more important than offshoring”. In contrast, Keller and Utar (2015) find evidence of job polarization created by Chinese import penetration when analyzing Danish worker-level data. Further, the authors argue that trade, rather than RBTC or offshoring, has the unique ability to explain job polarization in Denmark. Our local labor markets analysis suggests the opposite in the US, with RBTC rather than trade providing the unique explanation for US job polarization.

Third, the focus on workers’ occupations in Ebenstein et al. (2014) provides a natural link between the existing trade and local labor markets literature and the job polarization literature. Ebenstein et al. (2014, 2015) find that the exposure to globalization of a worker’s occupation, and not their industry, is the more salient determinant of their labor market outcomes. Intuitively, this follows from the fact that the occupational dimension of a job, and not the industry, will typically guide the job search processes of displaced workers. The focus on occupations link our analysis with prior studies of job polarization, as jobs are defined on the basis of either detailed occupations or the

³Interestingly, Shen and Silva (2018) find no adverse employment impacts in the US when looking at value added exports from China.

⁴Rather than look at impact of trade liberalization via imports on local labor market outcomes, a growing literature has looked at the impact via tariff liberalization (see, e.g., Hasan et al. (2007), Topalova (2007), McCaig (2011) and Hakobyan and McLaren (2016)).

cross between detailed occupations and aggregated industries. To compute an occupation’s trade exposure, Ebenstein et al. (2014, 2015) aggregate measures of industry-specific trade exposure using an occupation’s distribution of employment across industries. Similarly, our measure of local trade exposure aggregates (national) *occupation*-specific trade exposure using a location’s distribution of employment across *occupations*. This contrasts with prior studies on the impact of local trade exposure, e.g. Autor et al. (2013) among others, that aggregate *industry*-specific trade exposure using a location’s distribution of employment across *industries*.

The paper now proceeds as follows. Section 2 describes the empirical methodology and data. Section 3 presents the baseline results. Section 4 analyzes the relative impact of local shocks versus occupation shocks. Section 5 discusses numerous sensitivity analyses. Section 6 concludes.

2 Empirics

2.1 Empirical model

We assess the effects of local exposure to import competition and RBTC on employment growth across the job quality distribution in US local labor markets between 1990 and 2010. To do so, we build upon insights from the literatures on job polarization and the local labor market effects of trade exposure. To motivate our baseline specification, we first describe common empirical specifications from these literatures.

A typical specification in the job polarization literature (e.g., Goos and Manning (2007)) is

$$\Delta n_j = \beta_0 + \beta_1 q_j + \beta_2 q_j^2 + \varepsilon_j \tag{1}$$

where Δn_j is a measure of the national employment growth for job j between initial and terminal time periods, q_j is a measure of job quality for job j , and ε_j is a mean zero error term. Goos and Manning (2007) define jobs as either 3-digit occupations or the interaction of 3-digit occupations and 1-digit industries; Autor et al. (2006) define jobs as 3-digit occupations. Goos and Manning (2007) measure q_j using the median wage for job j in the initial time period, while Autor et al. (2006) measure q_j using the percentile position of job j in the national distribution of wages or education. In any case, interest has centered around the result that $\beta_1 < 0$ and $\beta_2 > 0$, producing a U-shaped relationship whereby employment growth for middle quality jobs is low relative to both low and high quality jobs.

While Section 2.2 discusses our data, we follow Goos and Manning (2007) and define jobs as the interaction of 3-digit occupations and 1-digit industries, yielding 2679 unique jobs. We follow the spirit of Autor et al. (2006) by defining a job’s quality as, essentially, its percentile position in the joint distribution of median education and median wage. Thus, quality varies from zero to one. Finally, we define the employment growth of a job as the change in its employment-to-working age population ratio. Column (1) in Table 1 and Figure 1 verify the stylized fact of job polarization at the national level in our data.

In contrast to the job polarization literature, a typical empirical specification in the literature analyzing the local labor market effects of trade exposure (e.g., Autor et al. (2013)) is

$$\Delta n_c = \beta_0 + \theta_1 \Delta T_c + x_c \delta + \varepsilon_c \quad (2)$$

where Δn_c is a measure of manufacturing employment growth in US local labor market c between initial and terminal time periods (often the change in the employment-to-working age population ratio), ΔT_c is a measure of the change in trade exposure faced by location c (often based on changes in import penetration from China), x_c is a vector of location-specific controls (e.g., location-specific demographics), and ε_c is a mean zero error term. Here, θ_1 is the coefficient of interest and captures the impact of local trade exposure on local manufacturing employment growth.

Again, much of the data details are presented in Section 2.2. However, we follow Autor et al. (2013) by defining locations as commuting zones (CZs). Letting Δn_c be the change in manufacturing employment-to-working age population ratio between 1990 and 2010 and regressing Δn_c on our measure of local Chinese import penetration (henceforth Chinese IP, and instrumented using Chinese exports to other high income countries) along with local demographic controls and state fixed effects, a one standard deviation change in local Chinese IP yields a nearly one-half standard deviation change in Δn_c . Thus, our data conveys the well known and large impacts of local Chinese IP on local labor markets over the 1990-2010 period.

We combine the two specifications in (1) and (2) into the following baseline specification to examine the impact of local trade exposure on *job polarization* in local labor markets:

$$\Delta n_{jc} = \beta_0 + \beta_1 q_j + \beta_2 q_j^2 + \theta_1 \Delta T_c + \theta_2 \Delta T_c q_j + \theta_3 \Delta T_c q_j^2 + \gamma_1 R_c + \gamma_2 R_c q_j + \gamma_3 R_c q_j^2 + x_{jc} \delta + \varepsilon_{jc} \quad (3)$$

where Δn_{jc} is employment growth (i.e., the change in the employment-to-working age population ratio, as in Autor and Dorn (2013)) of job j in US local labor market c between 1990 and 2010. Henceforth, we slightly abuse terminology by using the term ‘employment growth’ to describe Δn_{jc} . As an explanation for the job polarization illustrated in Figure 1, we include the change in local Chinese IP between 1990 and 2010, ΔT_c , and its interactions with q_j and q_j^2 . Given RBTC is the standard explanation in the literature for job polarization, we also include local RBTC exposure, R_c , and its interactions with q_j and q_j^2 . The presence of q_j and q_j^2 uninteracted with ΔT_c or R_c captures residual explanations for local job polarization. Additionally, x_{jc} is a vector of controls including location-specific socioeconomic and demographic attributes of locations as well as state, industry, and occupation fixed effects.⁵ These allow general patterns of worker reallocation due to trends in location-specific socioeconomic factors as well as state-specific, industry-specific, or occupation-specific effects.

When estimating (3), we assume ε_{jc} is a mean zero error term. However, for inference, we do not assume independent and identically distributed errors. For a given job, we allow across-CZ correlation between unobserved shocks to employment growth. Specifically, we decompose ε_{jc} into a job-specific random effect, v_j , and an idiosyncratic error, u_{jc} . Thus, the errors ε_{jc} and $\varepsilon_{jc'}$ in two locations c and c' are correlated because of their common component v_j . In turn, estimation is performed using random effects (RE) generalized least squares.⁶ Finally, we allow for arbitrary heteroskedasticity via robust standard errors.

While our local labor markets approach follows the recent literature exploring the effects of trade exposure, perfectly integrated national labor markets effectively imply a single observation for each country-labor market pair (Goldberg and Pavcnik (2016)). With perfect mobility across locations for a given job (i.e., viewing jobs as distinct labor markets), one could thus consider the national-level alternative to (3), given by:

$$\Delta n_j = \beta_0 + \beta_1 q_j + \beta_2 q_j^2 + \theta_1 \Delta T_j + \theta_2 \Delta T_j q_j + \theta_3 \Delta T_j q_j^2 + \gamma_1 R_j + \gamma_2 R_j q_j + \gamma_3 R_j q_j^2 + \varepsilon_j, \quad (4)$$

⁵For time-varying variables in x_{jc} , we control for initial levels and changes over the sample period.

⁶Note, while inclusion of a job-specific random effect is feasible, allowing for arbitrary correlation across CZs within a job via clustered standard errors at the job level is not feasible due to perfect collinearity between regressors within clusters (e.g. ΔT_c and $\Delta T_c \cdot q_j$ are perfectly collinear within job clusters while q_j and $\Delta T_c \cdot q_j$ are perfectly collinear within CZ clusters). This is not an issue in RE estimation. However, RE estimation imposes exchangeability (i.e., a constant, intra-cluster correlation) on the covariance matrix within a given job. Note, RE estimation is equivalent to estimating a pooled ordinary least squares model with quasi-demeaned variables; for example, $\widetilde{\Delta n}_{jc} \equiv \Delta n_{jc} - \lambda \overline{\Delta n}_j$, where $\overline{\Delta n}_j$ is the sample mean of Δn_{jc} for job j , $\lambda \equiv 1 - \sqrt{1/[1 + N_c(\sigma_v^2/\sigma_u^2)]}$, and N_c is the number of CZs. We report the estimated value of λ in the results.

where $\Delta T_j \equiv \Delta T_k$ is the (national) change in Chinese IP for occupation k (defined below in (6)) and $R_j \equiv R_k$ is the (national) measure of RBTC for occupation k (defined below in Section 2.2).⁷ Table 1 presents the results from this estimation, with standard errors clustered by occupation. First, Column (1) shows that the estimates of β_1 and β_2 generating the pattern of job polarization illustrated in Figure 1 are individually statistically significant at conventional levels. Second, Columns (2)-(7) in Table 1 show that the impacts of Chinese IP and RBTC are imprecisely estimated, and never individually statistically significant at conventional levels, when relying only on between occupation variation at the national level.⁸

Our specification in (3) differs from the existing trade and local labor markets literature by assessing the impact of local trade exposure on the distribution of local employment across *narrowly* defined job types and permits *heterogeneous* impacts with respect to the initial quality of a job, q_j . Moreover, our focus on local employment growth within narrowly defined jobs allows us to augment (3) with measures of (national) *occupation-specific* trade exposure and RBTC, and their interactions with q_j and q_j^2 . In so doing, we are able to directly compare the relative importance of occupation-specific shocks (hence, invariant across locations within a job) to location-specific shocks (hence, invariant across jobs within a location). Disentangling these two channels is something that typical local labor market analyses cannot perform. However, this provides critical insights into the operation of labor markets by shedding light on the nature of spillovers along the location and job dimension as well as the implications of geographic versus occupational immobility.

Our primary interest lies in the θ and γ coefficients in (3). Because changes in local trade exposure and RBTC may be endogenous, our baseline approach takes numerous approaches to mitigate any such problems. First, we include a rich set of fixed effects. To control for industry- or occupation-specific shocks impacting employment growth, we include 1-digit industry fixed effects and 3-digit occupation fixed effects. Additionally, to control for region-specific shocks to employment growth we include state fixed effects.⁹

Second, to control for any other shocks that could affect employment growth and be correlated

⁷We utilize a slight abuse of notation to keep things simple. Jobs are indexed by j , which is the cross-product of 3-digit Census occupations (k) and 1-digit NAICS industries. Thus, each job j maps into an occupation k . National Chinese IP growth and RBTC are measured strictly at the occupation level.

⁸Note, two of our 381 occupations do not appear in our 1980 IPUMS Census data. Hence, per our discussion regarding construction of occupation-specific RBTC in Section 2.2, these two occupations have missing values for R_k . With eight industries, these two occupations account for 16 observations.

⁹In Section 5, we show our results are robust to replacing these fixed effects with job fixed effects and CZ fixed effects. Naturally, such a specification cannot identify β_1 , θ_1 and γ_1 . While the interaction coefficients govern the effect on polarization, the β_1 , θ_1 and γ_1 coefficients do aid with interpretation and especially the visual interpretation of our results so we maintain these fixed effects in our baseline specification.

with local exposure to rising Chinese IP, ΔT_c , or local exposure to RBTC, R_c , we instrument for both ΔT_c and R_c . Following Autor et al. (2013) and Autor et al. (2015), we instrument for Chinese IP using Chinese exports to high-income countries other than the US. Intuitively, the key idea is that the common component of Chinese exports across high income destinations is driven by productivity and other supply-side shocks in China rather than correlated import demand shocks across high-income countries. Additionally, we follow Autor and Dorn (2013) and Autor et al. (2015) and instrument for local RBTC exposure with the local share of employment in routine jobs based on a CZ’s 1950 occupational distribution. The key idea here is that this local routine employment share based on a CZ’s 1950 occupational distribution captures the long-run quasi-fixed component of a CZ’s routine employment share.

We also address several additional concerns via extensions or robustness checks. We begin by extending the baseline model to allowing for the effects of (national) occupation-specific Chinese IP growth and (national) occupation-specific RBTC on local employment growth to vary by job quality. Despite our inclusion of occupation fixed effects, one may be concerned about the omission of salient occupation-specific factors when modeling local job-specific employment growth. Specifically, while the occupation fixed effects capture any direct effects of occupation-specific attributes, they will not control for any differential effects by job quality. Moreover, by including both occupation- and CZ-specific factors in an extended model, each interacted with q_j and q_j^2 , we can assess the relative importance of (national) occupation-specific and (local) CZ-specific factors.

We then perform numerous robustness exercises. Despite our broad set of fixed effects, they will not account for (i) job-specific shocks, (ii) industry- or occupation-specific shocks that differentially affect locations, or (iii) state-level shocks that differentially affect jobs or locations within a state. As such, our results could reflect a spurious relationship between employment growth and trade exposure in the presence of secular job- or CZ-specific trends in employment growth. To partly address this issue, we replace our set of fixed effects with job and CZ fixed effects. As an alternative procedure, we augment (3) to include the lag of Δn_{jc} (specifically, employment growth between 1980 and 1990). Thus, we identify the model exploiting variation *conditional* on local job-specific employment growth in the prior decade.

Finally, we explore alternative definitions of jobs, job quality and local trade exposure and investigate potentially important sources of heterogeneity along several dimensions: (i) age and cohort, (ii) local Chinese IP based on intermediate imports versus non-intermediate imports, and (iii) sample period.

2.2 Data

Estimating (3) requires definitions and data that can operationalize the concepts of local labor markets (c), jobs (j), local job-specific employment growth (Δn_{jc}), job quality (q_j), changes in local trade exposure (ΔT_c), local RBTC (R_c), the vector of controls (x_{jc}), and instruments for local trade exposure and local RBTC. The sample period spans 1990 to 2010 but, as part of the sensitivity analysis and our IV approaches, we also utilize data from 1980. Table A1 in the Appendix provides summary statistics.

Local labor markets (c) ConsPUMAs (Consistent Public Use Microdata Areas) are the most disaggregated geographic unit consistently defined over time and covering the entire US in IPUMS data (Ruggles et al. (2017)).¹⁰ Hakobyan and McLaren (2016) use ConsPUMAs when analyzing the local labor market impacts of NAFTA. However, we follow Autor et al. (2013), and the vast majority of the literature, by using commuting zones (CZs). Despite the need to concord IPUMS geographic units to CZs, CZs were explicitly designed to capture the boundaries of local labor markets in terms of commuting patterns. Our data include the 722 CZs covering the mainland US.

Job types (j) Prior job polarization studies define jobs as either detailed occupations or the cross-product of detailed occupations and industries. Using detailed 3-digit occupations and aggregate 1-digit industries, Goos and Manning (2007) define $370 \times 10 = 3700$ jobs and observe roughly 1600 in their data. We use 381 detailed 3-digit occupations (1990 IPUMS Census occupation codes) and eight 1-digit NAICS industry codes (we concord 1990 IPUMS Census 3-digit industry codes to SIC and NAICS industries), yielding 3048 possible jobs of which 2679 are observed in 1990.^{11,12} Thus, our sample has $722 \times 2679 = 1,934,238$ CZ-job observations.

Local job-specific employment growth (Δn_{jc}) The dependent variable captures changes in local job-specific employment shares between 1990 and 2010. Our employment data comes from IPUMS and, specifically, the 1980 and 1990 Decennial Census (5% sample) and the 2010 American Community Survey (1% sample). To begin, we compute the share of the working age population (aged 25 to 64 and not currently enrolled in school, institutionalized, or listing their occupation

¹⁰See <https://usa.ipums.org/usa/>.

¹¹See the Appendix for concordance issues.

¹²Note, we actually observe 2691 jobs in 1990. However, 12 jobs have missing data on job quality.

as military) employed in job j in location c in year t . Denoting this share by n_{jct} , we define $\Delta n_{jc} = n_{jc,2010} - n_{jc,1990}$.¹³

Job quality (q_j) To measure job quality and avoid confounding temporal labor reallocation across jobs with changes in the quality of jobs, we follow the existing job polarization literature. Specifically, we use a time invariant measure of job quality obtained from the initial period, 1990.¹⁴ Our primary measure of job quality is the Nam-Powers-Boyd (NPB) index of socioeconomic standing computed at the *national* level (i.e., the quality of a given job is constant across locations). We explore alternative measures in the sensitivity analysis.

The NPB index is a function of the median wage and median education level of a job, both of which have been used as independent measures of job quality (see, e.g., Autor et al. (2006)). Our simultaneous usage of both is in line with Acemoglu and Autor (2011, p.1046), who describe job polarization as the “simultaneous rapid growth of both high education, high wage occupations and low education, low wage occupations.” The NPB index, which varies from 0 to 1, is the approximate percentage of the labor force in jobs with a *lower* combination of median wage and median education (Nam and Boyd (2004)). Specifically, we begin by computing the national median wage and national median education level for each job in 1990 using the IPUMS data described above. We then convert these into empirical cumulative density functions (CDFs) using employment shares as weights. Finally, q_j is the average percentile of job j across the empirical CDFs for the median wage and median education level.

On the surface, sorting our 2679 observed jobs by the distribution of education and wages offers little transparency regarding the types of jobs that sit in various parts of the distribution. To give further insight, Table A2 in the Appendix describes the so-called good jobs and bad jobs across broad occupation groups and industries by splitting the sample into the bottom 25%, middle 50%, and top 25% of jobs according to the NPB index. Specifically, we show the distribution of low, middle, and high quality jobs across 1-digit NAICS industries and six occupation groups as defined

¹³Employment-to-working age population ratios account for the possibility that trade and technology exposure may contribute to nonemployment through unemployment, discouraged workers, retirement or disability. Additionally, they avoid the tricky issue that the distinction between out of the labor force and unemployment can be blurry and misreported in survey data. It also avoids econometric complications arising from the fact that job invariant, location-specific attributes (i.e., any x_{jc} that does not vary across j such as economic and demographic attributes of local labor markets) cannot affect all employment shares in the same direction if the shares must sum to one.

¹⁴Note, this means that only jobs observed in 1990 can be included in the analysis. The quality of any new jobs appearing in later years have missing quality. However, as stated above, 2679 jobs are observed in 1990. Only four jobs appear in 2010 that did not exist in 1990; 638 jobs observed in 1990 disappeared in 2010.

in Autor and Dorn (2013). Table A2 presents both the distribution of *jobs* and the distribution of *workers* across occupations or industries within each quality bin.

As expected, the data depict steady changes in the occupational and industrial composition as one moves up the NPB index. In terms of industries, low quality jobs are concentrated in (i) Wholesale/Retail Trade & Transportation/Warehousing, (ii) Educational/Health Care/Social Assistance Services, and (iii) Arts/Entertainment/Recreation, Accommodation/Food Services (45% of jobs and 65% of workers). This is unsurprising as one would expect these industries to be intensive in low-skilled labor. In fact, this is precisely the case as the two occupation groups of (i) Low Skill Services and (ii) Clerical, Retail Sales account for 42% of low quality jobs and 77% of workers in low quality jobs.

Middle quality jobs are skewed towards traditional ‘blue-collar’ industries and occupations. Here, (i) Manufacturing and (ii) Mining/Oil/Gas, Utilities/Construction account for 28% of jobs and 37% of workers (up from 15% and 9%, respectively, for low quality jobs). In terms of occupation groups, it is now (i) Managers, Professional, Technology, Finance, Public Safety, (ii) Clerical/Retail, and (iii) Transport, Construction, Mechanical, Mining, Farm that account for the bulk (70% of jobs and 80% of workers).

Finally, high quality jobs are dispersed primarily among (i) Manufacturing, which continues to be well represented (14% of jobs and 16% of workers), as well as (ii) Professional/Business services and (iii) Educational/Health Care/Social Assistance services (29% of jobs and 59% of workers). Unsurprisingly, high quality jobs across manufacturing and these service industries are dominated by the single occupation of Managers, Professional, Technology, Finance and Public Safety which accounts for 87% of jobs and 92% of workers (up from 22% and 20%, respectively, for middle quality jobs). Ultimately, the distribution of jobs according to the NPB index fits well with the notion of low quality jobs being dominated by low skill occupations/industries, middle quality jobs by blue-collar occupations/industries, and high quality jobs by professional occupations dispersed across several industries.

Local measure of Chinese import penetration growth (ΔT_c) Our measures of local trade exposure follow the approach popularized in Topalova (2007) and used recently elsewhere (e.g., Autor et al. (2013); Kovak (2013); Hakobyan and McLaren (2016)). However, while much of the prior trade literature dealing with local labor markets focuses primarily on industries, given that imports and trade policy are defined at the industry level, our analysis focuses on occupations as

discussed earlier. Thus, our starting point follows Ebenstein et al. (2014), who define occupational exposure to trade shocks as a function of an occupation’s employment composition across industries and the associated industry-specific trade shocks. We then define local exposure to trade shocks as a function of a CZ’s employment composition across occupations and the associated occupation-specific trade exposure.

To proceed, we first calculate the change in (national) Chinese IP at the industry level, following Acemoglu et al. (2016), as

$$\Delta T_s \equiv \frac{\Delta M_s}{Y_{s,1991} + M_{s,1991} - X_{s,1991}}, \quad (5)$$

where s indexes the 84 traded Census industries and the change in Chinese imports, $\Delta M_s \equiv M_{s,2010} - M_{s,1991}$, is normalized by domestic absorption in 1991 as proxied by domestic shipments, $Y_{s,1991}$, plus net imports, $M_{s,1991} - X_{s,1991}$. We obtain the necessary trade data from COMTRADE and the domestic shipments data from the NBER-CES Manufacturing Industry Database (Becker et al. (2013)).¹⁵ Table A3 in the Appendix lists the top 20 sectors in terms of growth in Chinese IP. As has been documented elsewhere, the rise in Chinese IP has been substantial. Across all 84 traded Census industries, the mean is an 11% point increase and 12 industries experienced at least a 25% point increase.

Next, we compute the change in (national) occupation-specific exposure to Chinese IP as

$$\Delta T_k \equiv \sum_s \omega_{sk,1990} \Delta T_s, \quad (6)$$

where $\omega_{sk,1990} \equiv L_{sk,1990}/L_{k,1990}$ is the 1990 employment share of Census occupation k in Census industry s computed using the 1990 Census data described above.^{16,17} ΔT_k is the trade exposure variable in our national-level regressions in Table 1.

According to (6), *occupations* are highly exposed to Chinese IP if their employment is concentrated in *industries* highly exposed to Chinese IP. Table A4 in the Appendix lists the 20 occupations most exposed to growth in Chinese IP. Table A5 provides similar information for the six broad occu-

¹⁵Shipments data are only available for manufacturing industries and not all tradable industries. However, we do not set $\Delta IP_s = 0$ for non-manufacturing tradable industries. For these industries, we set ΔIP_s equal to the average ΔIP_s across all manufacturing industries.

¹⁶Using time-invariant employment shares mitigates endogeneity concerns due to employment composition responding to changes in Chinese IP over the sample period.

¹⁷We aggregate over all Census industries in (6), not just traded sectors, consistent with much of the literature (Topalova (2007); Topalova (2010); Hakobyan and McLaren (2016)).

pation groups defined in Autor and Dorn (2013). Table A5 also shows the distribution of occupations across 1-digit industries. All occupations except low skill services have a non-trivial share of their employment in manufacturing and/or agriculture. The two occupations of (i) Production, craft and (ii) Machine operators, assemblers have 64% and 76% of their employment in the manufacturing industry. In turn, they faced Chinese IP growth of 7% and 8.7%, respectively. Additionally, Transport, construction, mechanical, mining, farm has 23% of their employment in manufacturing and agriculture and faced Chinese IP growth of 2.3%. The two occupations of (i) Managers, etc. and (ii) Clerical, retail sales have 11% of their employment in manufacturing and faced Chinese IP growth of 1.7% and 1.4%, respectively. Given the differences across occupations in their exposure to Chinese IP growth, one might hypothesize that occupation-specific exposure to Chinese IP growth, as opposed to CZ-specific exposure, drives job-specific employment growth within CZs. Later, we explore this issue.

Finally, we compute the change in local exposure to Chinese IP as

$$\Delta T_c \equiv \sum_k \omega_{kc,1990} \Delta T_k, \quad (7)$$

where $\omega_{kc,1990} \equiv L_{kc,1990}/L_{c,1990}$ is the 1990 employment share of location c employment in Census occupation k computed using the 1990 Census data described above. Figure 2 illustrates the stark increase in local Chinese IP. Table A6 in the Appendix list the 20 CZs facing the largest increase in Chinese IP. Twelve of the top 20 CZs are located in Tennessee, Kentucky, and Virginia. Here, local Chinese IP growth ranges between 4% and 5% which is much higher than the overall mean of 2.6%. However, Panel A of Figure 3 illustrates that the top quartile of most exposed CZs cover a much broader swath of the US, stretching from the South into the Rust Belt and then westward through the Midwest. Conversely, the bottom quartile of most exposed CZs are concentrated in the Southwest and Southeast.

Local measure of routine biased technological change (R_c) A common concern in the trade and labor literature is adequately controlling for technological change which theoretically can have similar impacts as rising trade exposure. Given our 3-digit occupation fixed effects, we control for occupation-specific technological change. This could, for instance, take the form of skill-biased technological change, whereby low skill occupations suffer from technological change and high skill occupations benefit. It could also take the form of RBTC, whereby technological change hurts

occupations that rely heavily on routine tasks, as technology can automate these tasks relatively easily, and leaves the remaining occupations largely unaffected or even better off through positive complementarities created by automation. Our 1-digit NAICS industry fixed effects also control for broad industry-specific technological change.

That said, as argued in Autor and Dorn (2013), technological change may differentially affect locations depending on occupational composition. According to Goos et al. (2014, p.2511), “... the literature seems to be settling on using the RTI [routine-task intensity] measure as the best way to capture the impact of recent technological progress”. Indeed, Autor and Dorn (2013) show that locations endowed with a large share of routine-task intensive jobs in 1980 experienced larger RBTC in the form of sharper growth in the adoption of information and communications technology over the following 25 years.

Thus, to begin constructing our measure of location-specific RBTC (R_c), we start by taking the measure of occupation-specific Routine Task Intensity (RTI_k) in 1980 from Autor and Dorn (2013).¹⁸ Based on data from the US Dictionary of Occupational Titles, RTI_k measures the extent to which occupation k uses routine tasks relative to manual and abstract tasks in 1980. Following Autor and Dorn (2013) and Autor et al. (2015), we then use our Census employment data to define occupation k as routine-intensive if it lies in the top third of the 1980 employment-weighted distribution of RTI_k . Formally, we define $R_k = 1 [RTI_k > RTI^{P66}]$ as our occupation-specific RBTC variable (e.g. used above in the national level regressions). We then let \tilde{R}_c be the share of 1980 CZ c employment in routine-intensive occupations:

$$\tilde{R}_c \equiv \sum_k \omega_{kc,1980} \cdot R_k \quad (8)$$

where $\omega_{kc,1980} \equiv L_{kc,1980}/L_{c,1980}$ is the 1980 employment share of location c employment in Census occupation k . For ease of interpreting the coefficient estimates, the measure of R_c used in our analysis normalizes \tilde{R}_c in (8) to have a minimum value of zero:

$$R_c \equiv \tilde{R}_c - \min_c \tilde{R}_c. \quad (9)$$

Panel B in Table A6 lists the 20 locations most exposed to local RBTC. While the 20 locations most exposed to rising import competition with China are concentrated in the South, the 20

¹⁸See and <http://www.ddorn.net/data.htm>.

locations most exposed to RBTC stretch from the South to the Northeast to the Rust Belt and even as far west as Texas and California. Comparing Panels A and B of Figure 3 also depicts this stark contrast between the local exposure of trade and technology. While the Southwest, Northeast, Florida and parts of Texas are heavily exposed to RBTC, they face little trade exposure. On the other hand, the western part of the Midwest is relatively exposed to trade yet is barely exposed to RBTC. Indeed, the raw correlation between local trade exposure and local technology exposure is -0.147 and hence quite small and actually negative. Ultimately, the locations relatively vulnerable to trade are different than the locations relatively vulnerable to technology and this is helpful for distinguishing between the impacts of local exposure to trade and technology.

In part, this weak correlation between local trade exposure and local technology exposure across CZs stems from the fact that the burden of trade and technology falls on different occupations. This is illustrated by Table A4 which shows the 20 occupations most exposed to Chinese IP growth are very different than the 20 occupations most exposed to RBTC.

Covariates (x_{jc}) Using our Census data described above, we also control for the following time-varying and location-specific variables (both their 1990 levels and the change between 1990 and 2010): mean age and populations shares for US born, homeowners, at least a Bachelor’s degree, non-white, and speak two languages including English. The only time-invariant, location-specific variables are state fixed effects. The only time- and location-invariant attributes are the 3-digit Census occupation fixed effects (381 occupations) and the 1-digit NAICS industry fixed effects (8 industries).

Instruments We use instrumental variables (IV) estimation to address the potential endogeneity of local Chinese IP and RBTC. Construction of the first instrument follows Acemoglu et al. (2016), computed in four steps. First, we replace the numerator in (5) with the change in industry-level Chinese exports to eight non-US high income countries.¹⁹ Second, we use industry-level US domestic absorption from 1989 rather than 1991 in the denominator of (5).²⁰ Third, we use 1980 Census occupation-industry shares $\omega_{sk,1980} \equiv L_{sk,1980}/L_{k,1980}$, rather than $\omega_{sk,1990}$ from the 1990 Census, when aggregating to the occupation level in (6). Fourth, we use local 1980 Census employment weights $\omega_{kc,1980} \equiv L_{kc,1980}/L_{c,1980}$, rather than $\omega_{kc,1990}$ from the 1990 Census, when aggregating to the local level in (7). As discussed in Acemoglu et al. (2016), the instrument is relevant if Chinese

¹⁹The countries are Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland.

²⁰COMTRADE data for US imports is unavailable before 1991. So, we use USITC import data for 1989.

exports are correlated across high income countries and is valid if this correlation is driven by Chinese productivity and other supply-side shocks rather than correlated import demand shocks among high income countries.

Construction of the second instrument follows Autor and Dorn (2013). Specifically, we merely replace the 1980 Census employment weights $\omega_{kc,1980}$ in (8) with 1950 Census employment weights $\omega_{kc,1950} = L_{kc,1950}/L_{c,1950}$. As discussed in Autor and Dorn (2013), the instrument is relevant because a CZs occupational composition as far back as 1950, a period well before the dramatic rise of computerization, is still a good predictor for the CZ’s adoption of workplace computers beginning in the 1980s and the associated displacement of routine-task intensive occupations. Also as discussed by Autor and Dorn (2013), the validity of the instrument rests on the assumption that the routine share of a CZ’s employment in 1950 represents a long-run quasi-fixed component of a CZ’s vulnerability to RBTC that was determined prior to the post-1980 period of computerization.

Finally, note that we use interactions between the instruments and q_j and q_j^2 to instrument for the interaction terms involving the endogenous covariates.

3 Baseline results

Table 2 presents the baseline results. Column (1) regresses Δn_{jc} on q_j and q_j^2 only. Columns (2)-(5) display the RE estimates, while Columns (6)-(9) display the RE-IV estimates, after incorporating additional controls. Relative to Column (1), Columns (2) and (6) add local Chinese IP, ΔT_c , and its interactions with q_j and q_j^2 . Columns (3) and (7) add location-specific covariates. Columns (4) and (8) add local RBTC, R_c , and its interactions with q_j and q_j^2 . Finally, Columns (5) and (9) add state, 3-digit Census occupation, and 1-digit NAICS industry fixed effects.

Column (1) confirms the pattern of coefficient estimates for job polarization at the local labor market level in the US: as we move through the distribution of job quality, employment growth initially decreases ($\beta_1 < 0$) but then increases ($\beta_2 > 0$). However, the estimated parameters are not individual statistically significant at conventional levels and visual inspection of Figure 4 reveals a nearly-monotonic relationship rather than a clear U-shaped relationship. Indeed, Figure 4 indicates negative employment growth for all low quality jobs. In terms of the magnitude of the employment effects, it is important to realize that, with 2691 jobs, the mean employment share across all location-jobs is 0.032%.²¹ As such, the predicted changes in employment growth are economically

²¹Note, 0.032% is less than $1/2691$ due to nonemployment; see footnote 13.

meaningful and with a decline (increase) of about 10% (22%) of this mean employment share when $q = 0$ ($q = 1$).

We now focus our discussion around the RE-IV results in Columns (6)-(9) as the RE and RE-IV estimates are qualitatively similar. Nonetheless, before turning to the parameter estimates, we note that the RE-IV specification tests perform very well. In particular, we easily reject the null of underidentification at the $p < 0.01$ level and the first-stage F -statistics on the excluded instruments are very large. Finally, despite the similarity in the OLS and RE-IV point estimates, we do reject the null of exogeneity in Columns (6)-(9) at the $p < 0.01$ level.

Turning to the parameter estimates, Column (6) adds local Chinese IP and the associated interactions with q_j and q_j^2 to the model. The point estimates for θ_1 , θ_2 , and θ_3 are neither individually nor jointly statistically significant at conventional levels. This suggests that trade with China is not a driver of job polarization within US local labor markets. Further, the sign pattern of the point estimates ($\hat{\theta}_1 < 0$, $\hat{\theta}_2 > 0$, $\hat{\theta}_3 < 0$) as well as the fact that $\hat{\theta}_1 + \hat{\theta}_2 + \hat{\theta}_3 < 0$ suggests an *anti-polarization* impact of local Chinese IP.

Formally, the marginal effect of local Chinese IP is

$$\frac{\partial \text{E}[\Delta n_{jc}|\cdot]}{\partial \Delta T_c} = \theta_1 + \theta_2 q_j + \theta_3 q_j^2, \quad (10)$$

where q_j varies from zero to one and $\text{E}[\Delta n_{jc}|\cdot]$ is the conditional expectation of local employment growth of job j given the covariates in the model. Based on the estimates in Column (6), we find that the marginal effect of local Chinese IP is negative for low and high quality jobs, yet positive for middle quality jobs.²² Moreover, controlling for the local Chinese IP variables increases the point estimates of β_1 and β_2 in absolute value. Based on these point estimates, job polarization in local labor markets would have been *more* pronounced had local Chinese IP stayed constant at 1990 levels. We illustrate this graphically below.

Column (7) adds CZ-specific socioeconomic and demographic controls. The point estimates are essentially unchanged. However, Column (8) shows that controlling for local RBTC is consequential. The point estimates for the impact of local Chinese IP remain neither individually nor jointly statistically significant at conventional levels; the sign pattern continues to indicate an anti-polarization effect. However, the point estimates for the impact of local RBTC are statistically significant at

²²Specifically, the marginal effect is positive for values of q_j between 0.60 and 0.92.

the $p < 0.01$ level. Formally, the marginal effect of local RBTC is

$$\frac{\partial E[\Delta n_{jc}|\cdot]}{\partial R_c} = \gamma_1 + \gamma_2 q_j + \gamma_3 q_j^2, \quad (11)$$

with the pattern of estimates ($\hat{\gamma}_1 > 0$, $\hat{\gamma}_2 < 0$, $\hat{\gamma}_3 > 0$) and the fact that $\hat{\gamma}_1 + \hat{\gamma}_2 + \hat{\gamma}_3 > 0$ indicating a *polarization* effect. That is, local RBTC increases employment growth of low quality and high quality jobs, but leads to lower employment growth of middle quality jobs.²³ Indeed, the impact of local RBTC is so pervasive that the point estimates on q_j and q_j^2 become very close to zero. In other words, as we illustrate visually below, our results imply that employment growth exhibits polarization *only* in locations sufficiently impacted by RBTC. Thus, consistent with the prior literature on job polarization at the national level, we find *local* RBTC to be the key determinant of job polarization in *local* labor markets.

Finally, Column (9) adds state, occupation, and industry fixed effects, thereby removing any unobservables along these dimensions that could influence local employment growth over the sample period. Yet, our estimates are nearly identical. Thus, our preferred specification, Column (9), confirms our findings above: local RBTC is the key explanation for job polarization within local labor markets and local Chinese IP has, if anything, an anti-polarization effect. In fact, not only does local Chinese IP have an anti-polarization effect based on the results in Columns (8) and (9), but the marginal effect is essentially non-positive across all values of q_j . That is, local Chinese IP reduces employment growth of all low and middle quality jobs, and has little impact on high quality jobs.

Figure 5 illustrates the polarization impact of local RBTC and the anti-polarization impact of Chinese IP. Based on the point estimates in Column (9), Panel A of Figure 5 shows the polarization impact of local RBTC by plotting $E[\Delta n_{jc}|\cdot]$ while fixing all covariates, including local Chinese IP, at their sample means except for local RBTC. Local RBTC is varied from zero (i.e., the minimum value of local RBTC), to its 10th percentile value (0.027), to its 90th percentile value (0.125). When local RBTC is held at its minimum value, and local Chinese IP is set at the sample mean, we do not see job polarization at the local level. As local RBTC grows we eventually see job polarization, with positive employment growth at the lower and upper tails and negative employment growth in the middle. This visually illustrates the polarization effect of local RBTC.

Conversely, Panel B of Figure 5 illustrates the anti-polarization impact of local Chinese IP. Based

²³Specifically, the marginal effect is negative for values of q_j between 0.37 and 0.87.

on the point estimates in Column (9), the figure plots $E[\Delta n_{jc}|\cdot]$ fixing all covariates, including local RBTC, at their sample means except for the change in local Chinese IP. The change in local Chinese IP is varied from zero (i.e., local Chinese IP is held constant at 1990 levels), to its 10th percentile value (0.019), to its 90th percentile value (0.033). When local Chinese IP is held constant at 1990 levels, and local RBTC is set at the sample mean, we see job polarization at the local level. As local Chinese IP growth is increased, we see employment growth declining across all jobs except those close to the top of the job quality distribution, with more extreme declines at the lower tail. This visually illustrates the anti-polarization effect of local Chinese IP. In fact, as the figure illustrates, sufficiently strong local Chinese IP growth (with local RBTC held fixed at the sample mean) essentially wipes out job polarization and generates a nearly-monotonic impact on employment growth across the job quality distribution with, essentially, non-positive employment growth for all jobs.

4 Local shocks versus occupation shocks

Our baseline analysis investigated whether and how local shocks, either local RBTC or local Chinese IP shocks, have impacted job polarization in local labor markets. In so doing, our preferred baseline specification includes occupation fixed effects which implicitly control for (national) occupation-specific Chinese IP growth and RBTC. However, these occupation fixed effects do not allow for heterogeneous impacts of (national) occupation-specific Chinese IP growth and RBTC by job quality. In other words, our baseline analysis does not allow for the possibility that occupation-specific shocks may impact job polarization in local labor markets. Given occupations differ in their distribution of employment across industries, as illustrated in Table A5, one may be concerned that these occupation-specific shocks have important impacts on job polarization in local labor markets.

Indeed, our approach of focusing on local job-specific employment growth rather than the typical approach of looking at local overall or just manufacturing employment growth helps address this issue. Specifically, our approach allows us to compare a given job in different locations and to compare different jobs in a given location. Moreover, an added benefit is that we can actually distinguish between the impact of (national) occupation-specific and (local) CZ-specific factors.

Formally, we now investigate the following expanded model

$$\begin{aligned} \Delta n_{jc} = & \beta_0 + \beta_1 q_j + \beta_2 q_j^2 + \theta_1 \Delta T_c + \theta_2 \Delta T_c q_j + \theta_3 \Delta T_c q_j^2 + \gamma_1 R_c + \gamma_2 R_c q_j + \gamma_3 R_c q_j^2 \quad (12) \\ & + \delta_2 \Delta T_j q_j + \delta_3 \Delta T_j q_j^2 + \phi_2 R_j q_j + \phi_3 R_j q_j^2 + x_{jc} \delta + \varepsilon_{jc}, \end{aligned}$$

where $\Delta T_j \equiv \Delta T_k$ is the (national) change in Chinese IP for occupation k (defined in (6)) and $R_j \equiv R_k$ is the (national) measure of RBTC for occupation k .^{24,25} Table 3 displays the results; Column (2) contains the new results and, for comparison, Column (1) repeats our baseline results from Column (9) of Table 2.²⁶

Two points stand out. First, adding the (national) occupation-specific variables interacted with q_j and q_j^2 leaves the coefficient estimates on the local variables and their interactions with q_j and q_j^2 unchanged, both qualitatively and in terms of statistical significance. Thus, the (national) occupation-specific variables and their local counterparts are essentially uncorrelated. This not only confirms our baseline analysis, but allows us to empirically distinguish the impacts of local versus occupational exposure to trade and technology on local job-specific employment growth across the distribution of job quality.

Second, the effects of (national) occupation-specific exposure to RBTC and Chinese IP are statistically insignificant at conventional levels. That is, occupation-specific exposure has no statistically meaningful impact on the extent of job polarization in local labor markets. Thus, our results suggest that *local* exposure to RBTC, not occupation-specific exposure to RBTC, drives job polarization within local labor markets.

Despite the statistical insignificance of the point estimates for occupation-specific exposure, Figure 6 illustrates these estimates by showing the effects of simultaneous changes in both local *and* (national) occupation-specific Chinese IP growth and RBTC based on the estimates in Column (2). In comparison to Figure 5 (derived from the baseline specification), we see that incorporating (national) occupation-specific exposure to RBTC in Panel A reduces employment growth across the job quality distribution. Incorporating (national) occupation-specific exposure to Chinese IP

²⁴As before (see footnote 7), we utilize a slight abuse of notation to keep things simple. Jobs are indexed by j , which is the cross-product of 3-digit Census occupations (k) and 1-digit NAICS industries. Thus, each job j maps into an occupation k . National Chinese IP growth and RBTC are measured strictly at the occupation level.

²⁵Note, occupation fixed effects absorb the uninteracted terms, ΔT_j and R_j .

²⁶The smaller sample size in Column 2 stems from two occupations having zero employment in 1980 and hence a missing value for R_k (see also footnote 8). Given the location-occupation-industry level of observation, Column 2 has $722 \times 2 \times 8 = 11,552$ fewer observations than Column 1.

growth in Panel B not only reduces employment growth across the job quality distribution, but also reduces the curvature of the graphs, indicative of a flattening out of the extent of polarization in local labor markets.

5 Sensitivity analyses

We perform numerous sensitivity analyses to assess the robustness of the baseline results.

Alternative specifications Table 4 presents results from a number of alternative specifications of our baseline model. For comparison, Column (1) replicates our preferred estimates from Column (9) in Table 2. Despite the broad set of fixed effects included in our baseline analysis, they will not absorb job-specific, location-specific or location-job specific trends. The latter two are a common concern within trade and local labor market analyses because economically declining locations may tend to specialize in import-competing goods. If this is the case, then these locations will experience the greatest changes in local trade exposure. In turn, this may generate a spurious relationship between local trade exposure and local labor market employment growth due to omitted location-specific secular trends.

To address these concerns, Column (2) replaces our baseline set of fixed effects with job and CZ fixed effects. While we can now only identify the interaction impacts of local exposure, our baseline results remain unaffected. If anything, the larger point estimates suggest a stronger polarization impact of local RBTC and a larger anti-polarization impact of local Chinese IP growth. Further, Column (3) addresses the possibility of secular CZ-job specific trends by controlling for the lag of Δn_{jc} (i.e. CZ job-specific employment growth from 1980 to 1990). The results reveal essentially no change in the estimates.²⁷ Thus, our baseline results hold even when *conditioning* on the CZ-job specific employment growth between 1980 and 1990.

Columns (4)-(8) are identical to our preferred baseline model, but explore alternative measures of job quality. In Column (4) we allow the quality of a given job to vary across regions; in contrast, recall our baseline measure of job quality is constant across the US. Specifically, we now compute job quality separately for each of the nine US Census regions.²⁸ In so doing, we allow for potentially

²⁷Note, the inclusion of the lagged dependent variable is likely to violate the assumption of strict exogeneity required in RE estimation. However, using (non-RE) IV produces estimates that are essentially unchanged (available upon request).

²⁸We do not attempt to measure job quality at a more disaggregate level than Census regions since many jobs are not observed. Even when computing regional measures of job quality, often a region does not contain a particular

important regional variation in real wages due to price differences or in educational attainment and nominal wages.²⁹

In Column (5) we revert back to a national measure of job quality, but instead use a time invariant measure based on the 2010 median wages and education levels observed in each job. This addresses the potential concern that the quality ranking of jobs may substantially change over the sample period and render our time-invariant notion of job quality based on data in 1990 misleading. Note, this change reduces the sample size as some jobs observed in 1990 no longer exist in 2010 and therefore have missing job quality (see footnote 14).

Finally, Columns (6)-(8) are identical to Columns (1), (4), and (5) except that the underlying quality measure (from either 1990 or 2010 and at either the regional or national level) is based solely on the median wage of a job, not the NPB index of median wages and median education levels.

Turning to the results in Columns (4)-(8), the statistical significance and qualitative results from our baseline specification are essentially unchanged. For interpretation, Figures 7A and 7B plot the coefficient estimates. While there is variation across the panels, the basic conclusion that local RBTC has a polarization impact and local Chinese IP has an anti-polarization impact is quite robust. In all cases, we find that local RBTC (local Chinese IP) increases (decreases) employment growth of low and high quality jobs and decreases (at best, has no impact on) employment growth of middle quality jobs.

Column (9) addresses the possible concern that our very disaggregated job definition leaves small sample sizes in many CZ-job cells. Thus, we follow Ebenstein et al. (2014) and simply define a job as an occupation. With our 381 3-digit Census occupations, we now have 381 jobs rather than 2679 jobs. Note that removing the eight 1-digit NAICS industries from the job definition reduces the number of jobs to roughly one-eighth. Indeed, the point estimates for the local RBTC variables in Column (9) are essentially scaled by a factor of eight and remain statistically significant at the 5% level. While the point estimates for the local Chinese IP variables are scaled by more than a factor of eight, they remain statistically insignificant at conventional levels and display the familiar anti-polarization pattern. Ultimately, the possibility of small sample sizes in our baseline CZ-job cells does not seem problematic.

job. In these cases, we use the national measure of its quality for the region.

²⁹See Marchand (2012) and Fajgelbaum and Khandelwal (2016) for recent empirical work emphasizing the importance of trade liberalization on prices.

Heterogeneous effects Table 5 displays the results from various extensions to our baseline model that explore possible heterogeneities in the determinants of local employment growth. Again, Column (1) repeats our preferred estimates from Column (9) of Table 2 for comparison.

In Columns (2) and (3) we assess the possibility of heterogeneous effects of local Chinese IP growth depending on the nature of the goods being imported. Specifically, using the UN’s Broad Economic Category Classification, we construct two measures of local Chinese IP growth: one based solely on imports of intermediate inputs and one based solely on imports of non-intermediate inputs. Formally, we create two alternative measures of ΔT_c that differ from our baseline measure through slightly varying (5). The first alternative measure of (5) only uses intermediate input imports. The second alternative measure of (5) only uses non-intermediate input imports. In terms of the results, we find the magnitude of the point estimates to be modestly larger in absolute value in Column (3), where local Chinese IP growth is measured using non-intermediate inputs. However, this is misleading since actual local Chinese IP growth is larger for intermediate input imports than non-intermediate input imports. Figure 8 reveals that the impacts of local Chinese IP growth are virtually indistinguishable in the two cases.

In Column (4) we assess the possibility of heterogeneous effects due to the Great Recession. Here, the changes in all variables are computed over the period 1990 to 2000. As shown in Figure 9, local Chinese IP growth has a more prominent anti-polarization effect while local RBTC has a slightly weaker polarization effect in the pre-Great Recession period. Thus, both in terms of statistical and economic significance, our baseline results are not driven by the Great Recession.

In Columns (5)-(7) we explore whether the effects of local Chinese IP growth and local RBTC differentially affect local employment growth across three different cohorts of workers: (i) ‘young’ individuals aged 25-44, (ii) ‘old’ individuals aged 45-64, and (iii) the ‘cohort’ of individuals aged 25-44 in 1990 and 45-64 in 2010. Formally, the dependent variable in all models, including our baseline specifications, can be written as $\Delta n_{jc} = n_{jc,2010} - n_{jc,1990}$, where $n_{jc,t}$ is the employment share in a CZ-job in year t . In the baseline model, employment shares in year t are computed using individuals aged 25-64. In Column (5), employment shares in year t are instead computed using only individuals aged 25-44. In Column (6), only individuals aged 45-64 are used. In Column (7), $n_{jc,1990}$ is computed using individuals aged 25-44, while $n_{jc,2010}$ is computed using individuals aged 45-64. Thus, Columns (5) and (6) allow us to assess whether the labor market impacts of local Chinese IP growth and local RBTC differentially impact younger and older workers. Column (7) allows us to assess the within cohort impacts of local Chinese IP growth and local RBTC.

While the general pattern of coefficient estimates is similar across the columns, confirming an anti-polarization impact of local Chinese IP growth and a polarizing impact of local RBTC, interesting differences arise when examining local Chinese IP growth in Figure 10. In particular, while local Chinese IP growth depresses employment in low quality jobs among young workers, old workers and the ‘cohort’ of workers, local Chinese IP growth only increases employment in high quality jobs for the ‘cohort’ of workers. Thus, over the life cycle, local Chinese IP growth allowed some workers to move into high quality jobs.

6 Conclusion

This paper investigates the heterogeneous effects of changes in local trade exposure and routine-biased technological change (RBTC) on the employment growth of good versus bad jobs across US local labor markets between 1990 and 2010. Our analysis yields three main, and robust, results.

First, we show that local exposure to RBTC, not local exposure to Chinese IP growth, drives job polarization in US local labor markets. This complements the findings in Autor et al. (2015) who find a similar result for occupational polarization in US local labor markets. Because of the inexact mapping between broad occupation groups and job quality, one should not take occupational polarization to imply job polarization (or vice versa).

Second, we show that job polarization only emerges in local labor markets that are sufficiently exposed to RBTC. Obtaining this result is only possible through the use of a CZ-job pair as the unit of our analysis; our analysis allows us to assess whether a CZ that is weakly exposed to RBTC experiences job polarization or not. In contrast, analyses at the CZ level can only determine relative degrees of polarization across CZs.

Third, we show that local exposure rather than (national) occupation-specific exposure to RBTC drives the extent of job polarization in local labor markets. Again, obtaining this result relies on our analysis being at the CZ-job level, rather than the CZ level, as it allows for comparisons of different occupations in the same CZ as well as comparisons of the same occupation across CZs. Ultimately, the degree of job polarization in US local labor markets is intimately tied to local exposure, rather than occupation-specific exposure, suggesting a prime role of geographic immobility rather than occupational immobility.

References

- Acemoglu, D., 2002. Directed technical change. *The Review of Economic Studies* 69 (4), 781–809.
- Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics* 4, 1043–1171.
- Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H., Price, B., 2016. Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics* 34 (S1), S141–S198.
- Acemoglu, D., Gancia, G., Zilibotti, F., 2015. Offshoring and directed technical change. *American Economic Journal: Macroeconomics* 7 (3), 84–122.
- Autor, D. H., Dorn, D., 2013. The growth of low-skill service jobs and the polarization of the US labor market. *The American Economic Review* 103 (5), 1553–1597.
- Autor, D. H., Dorn, D., Hanson, G. H., 2013. The China syndrome: Local labor market effects of import competition in the United States. *The American Economic Review* 103 (6), 2121–2168.
- Autor, D. H., Dorn, D., Hanson, G. H., 2015. Untangling trade and technology: Evidence from local labor markets. *The Economic Journal* 125, 621–646.
- Autor, D. H., Katz, L. F., Kearney, M. S., 2006. The polarization of the US labor market. *The American Economic Review*, 189–194.
- Autor, D. H., Levy, F., Murnane, R. J., 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118 (4), 1279–1333.
- Balsvik, R., Jensen, S., Salvanes, K. G., 2015. Made in china, sold in norway: Local labor market effects of an import shock. *Journal of Public Economics* 127, 137–144.
- Becker, R. A., Gray, W. B., Marvakov, J., 2013. NBER-CES manufacturing industry database: June 2013 revision. National Bureau of Economic Research.
- Burstein, A., Vogel, J., 2017. International trade, technology, and the skill premium. *Journal of Political Economy* 125 (5), 1356–1412.
- Dauth, W., Findeisen, S., Suedekum, J., 2014. The rise of the east and the far east: German labor markets and trade integration. *Journal of the European Economic Association* 12 (6), 1643–1675.
- Donoso, V., Martín, V., Minondo, A., 2015. Do differences in the exposure to Chinese imports lead to differences in local labour market outcomes? an analysis for Spanish provinces. *Regional Studies* 49 (10), 1746–1764.
- Ebenstein, A., Harrison, A., McMillan, M., 2015. Why are American workers getting poorer? China, trade and offshoring. NBER Working Paper No. 21027.
- Ebenstein, A., Harrison, A., McMillan, M., Phillips, S., 2014. Estimating the impact of trade and offshoring on American workers using the Current Population Surveys. *Review of Economics and Statistics* 96 (4), 581–595.
- Fajgelbaum, P. D., Khandelwal, A. K., 2016. Measuring the unequal gains from trade. *The Quarterly Journal of Economics* 131 (3), 1113–1180.

- Goldberg, P. K., Pavcnik, N., 2016. The effects of trade policy. NBER Working Paper No. 21957.
- Goos, M., Manning, A., 2007. Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics* 89 (1), 118–133.
- Goos, M., Manning, A., Salomons, A., 2014. Explaining job polarization: Routine-biased technological change and offshoring. *The American Economic Review* 104 (8), 2509–2526.
- Hakobyan, S., McLaren, J., 2016. Looking for local labor market effects of NAFTA. *Review of Economics and Statistics* 98 (4), 728–741.
- Hasan, R., Mitra, D., Ural, B. P., 2007. Trade liberalization, labor-market institutions and poverty reduction: Evidence from Indian states. In: *India Policy Forum*. Vol. 3. pp. 71–122.
- Impullitti, G., Licandro, O., 2018. Trade, firm selection and innovation: The competition channel. *The Economic Journal* 128 (608), 189–229.
- Keller, W., Utar, H., 2015. International trade and job polarization: Evidence at the worker level. Mimeo.
- Kovak, B. K., 2013. Regional effects of trade reform: What is the correct measure of liberalization? *The American Economic Review* 103 (5), 1960–1976.
- Marchand, B. U., 2012. Tariff pass-through and the distributional effects of trade liberalization. *Journal of Development Economics* 99 (2), 265–281.
- McCaig, B., 2011. Exporting out of poverty: Provincial poverty in Vietnam and US market access. *Journal of International Economics* 85 (1), 102–113.
- Nam, C. B., Boyd, M., 2004. Occupational status in 2000; over a century of Census-based measurement. *Population Research and Policy Review* 23 (4), 327–358.
- Ruggles, S., Genadek, K., Goeken, R., Grover, J., M., S., 2017. Integrated public use microdata series: Version 7.0 [dataset]. Minneapolis: University of Minnesota.
- Samuel, L. R., 2013. *The American Middle Class: A Cultural History*. Vol. New York. Routledge.
- Shen, L., Silva, P., 2018. Value added exports and US local labor markets: Does China really matter? *European Economic Review* 101, 479–50.
- Topalova, P., 2007. Trade liberalization, poverty and inequality: Evidence from Indian districts. In: *Globalization and Poverty*. University of Chicago Press, pp. 291–336.
- Topalova, P., 2010. Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics* 2 (4), 1–41.

Appendix

Here, we describe the various concordance issues that arise in the data.

Occupations and occupation groups The RTI variable from Autor and Dorn (2013) uses their self-compiled occupation variable *occ1990dd*. Further, the six occupation groups defined by Autor and Dorn (2013) collapse *occ1990dd*. However, we use the IPUMS Census variable *occ1990*. Thus, we concord from *occ1990dd* to *occ1990*. A further complication is that *occ1990dd* is based on the Census occupation variable *occ1990* which differs from the IPUMS Census variable *occ1990*. Nevertheless, we carry out the concordance using David Dorn’s concordance between the Census *occ1990* variable and *occ1990dd* (<http://www.ddorn.net/data.htm>) and the IPUMS concordance between its own Census *occ1990* variable and the Census *occ1990* variable (https://usa.ipums.org/usa/voliii/occ_ind.shtml).

Industries Industry level trade and shipments data is available at the 4-digit SIC level (from WITS COMTRADE and the NBER-CES Manufacturing database, respectively). However, we use the IPUMS Census industry variable *ind1990*. Thus, we concord from SIC to *ind1990*. To do so, we used a Census concordance that, in the past, was available at <http://www.cdc.gov/niosh/soic/pdfs/PT19Y99AppB.pdf>. While this url link is no longer active, interested readers can download this file, together with STATA and Excel versions, from the corresponding author’s website: http://people.smu.edu/jlake/data_code/SIC_IPUMS_industry_concordance.zip.

Locations While we use CZs as the definition of local labor markets, this variable is not in IPUMS Census data. Thus, we concord from the most disaggregated geographic unit in the IPUMS Census data to CZs.

In Census microdata, the most disaggregated level of geography needs to have at least 100,000 people (http://www.ddorn.net/data/Dorn_Thesis_Appendix.pdf, p.136). This gives rise to the notions of “county groups” (in the form of the CNTYGP97 and CNTYGP98 Census variables for 1970 and 1980) or “PUMAs” (in the form of the PUMA, PUMA1990 and PUMA2000 Census variables for 1990 onwards) whose definition changes over time to achieve the minimum population threshold for a geographical Census unit. CZs aggregate these most disaggregated geographical Census units in a way that carefully attempts to respect “local labor markets”.

Note that the PUMAYYYY variables and the PUMA variable convey somewhat different information. In 2010, the PUMA2000 variable is still the most disaggregated geographical unit (<https://usa.ipums.org/usa/vol11/2000pumas.shtml>) with the caveat that 3 PUMAs in LA were merged into one PUMA because of Hurricane Katrina affecting the Census population threshold described above. In the Census data, the PUMAYYYY variable is only non-missing in the year YYYY. But, the PUMA variable is non-missing for 1990, 2000 and 2010 and records the associated year-specific PUMAYYYY. However, the form of the numeric values differ across the PUMA and PUMAYYYY variables. Specifically, the PUMAYYYY variable contains state *and* PUMA information whereas PUMA numbers are not unique across states (i.e. a location according to the PUMA variable is really a (PUMA, STATEFIP) pair). Specifically,

$$PUMAYYYY = STATEFIP \times 10,000 + PUMA.$$

Before aggregating these disaggregated geographical units in the Census, CZs split these geographical units. The splitting process is explained in detail by David Dorn (http://www.ddorn.net/data/Dorn_Thesis_Appendix.pdf, pp.136-138). Moreover, David Dorn provides concordances that map from the various disaggregated geographic units in the Census to time-invariant 1990 CZs (<http://www.ddorn.net/data.htm>). Indeed, when using an n:n merge for the concordance from either (i) PUMA1990 to CZ1990 or (ii) PUMA2000 to CZ1990, there are not any PUMAs left unmatched nor are there any CZs left unmatched. The only subtlety arises because of the Hurricane Katrina issue above where PUMAs 221801, 221802 and 229105 are not in the Census microdata while PUMA 227777 (the newly aggregated PUMA) is not in Dorn's concordance. But, since the three old PUMAs all mapped to the same CZ then this is easily fixed manually. Having concorded PUMAs to CZs, one needs to adjust the person weights *perwt* that are in the Census microdata using the allocation factors *afactor* in David Dorn's concordance. Specifically, the new person weights are

$$perwt_cz = perwt \times afactor.$$