

Trade and welfare (across local labor markets)*

Ryan Kim

Seoul National University

Jonathan Vogel

University of California, Los Angeles

Moises Yi

US Census Bureau

January 23, 2026

Abstract

What are the welfare implications of trade shocks? Theoretically, we provide a sufficient statistic that measures changes in welfare (to a first-order approximation) for the set of workers who start within a region, taking into account adjustment in frictional unemployment, labor force participation, the sectors to which workers apply for jobs, and the regions in which workers choose to live. Our theory is flexible; for instance, it allows for arbitrary heterogeneity in worker productivity and non-pecuniary returns (amenities) across unemployment, labor force non-participation, sectors, and regions. Empirically, we apply these insights to measure changes in welfare between 2000-2007 across workers who start in different commuting zones (CZs) in the U.S. in the year 2000. Finally, we identify the differential impact across CZs of a particular trade shock: granting China permanent normal trade relations.

*We thank David Atkin, David Baqaee, Kirill Borusyak, Ariel Burstein, Arnaud Costinot, Pablo Fajgelbaum, Katheryn Russ, and Pierre-Olivier Weill for helpful comments. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance (Project 6000266: CBDRB-FY25-CES024-001, CBDRB-FY26-CES011-003).

1 Introduction

Recent empirical work has identified substantial differential effects of international trade shocks (across, e.g., regions) on a broad set of labor-market adjustment margins, among them wages, labor force participation, unemployment, and labor allocations across sectors and regions. In response, a growing theoretical literature has emerged, with the objective of combining some of these empirical results and particular models to understand the welfare implications of international trade shocks.

To make progress, much of this theoretical literature has made four assumptions. First, it has imposed strong functional form restrictions, assuming—among other things—that worker preferences or productivities follow a particularly tractable distribution and that production functions are common across regions and sectors and take a simple parametric form. Second, it has typically incorporated a small subset of the adjustment margins that have featured prominently in the empirical analyses motivating it. Third, it has imposed particular assumptions on the economic environment (e.g., the strength and spatial nature of agglomeration, unemployment caused by wage rigidity or search and matching, etc.) that shape the response of each margin of adjustment to trade shocks. Fourth, it has measured changes in welfare using an approach that is not invariant to monotonic transformations of individuals' utility functions.

Our contribution is to revisit the results of the empirical literature, combining theory and data to identify the differential effect of a trade shock on welfare, rather than on various margins of adjustment. Unlike the existing theoretical literature, we measure welfare without imposing these strong restrictions.¹ We separate our approach into three steps.

First, we use theory to provide a sufficient statistic for changes in welfare for a set of agents, to a first-order approximation. Our approach does not impose strong functional forms, incorporates a broad range of adjustment margins the empirical literature has emphasized, allows for arbitrary changes in wages and unemployment probabilities, and measures welfare using a money metric that is invariant to monotonic transformations of individuals' utility functions (subject to risk neutrality). Second, we implement this theoretical approach using U.S. data, measuring welfare changes between the years 2000 and 2007 for the sets of workers initially living in each commuting zone (CZ) in the year 2000, regardless of where they live thereafter. To this point, our approach recovers changes in welfare for each group, but these changes are unconditional in the sense that they depend on all shocks to the economic environment. Third and finally, we identify the impact of a particular international trade shock—granting China permanent normal trade relations—on differential changes in

¹We do so at the cost of not providing an answer for the “missing intercept” problem, in the sense that whereas we measure changes in welfare in each region, we only estimate the impact of a trade shock on relative welfare changes across regions.

welfare across U.S. CZs using the same research design as in the reduced-form empirical literature, but replacing the change in a particular margin of adjustment with our sufficient statistic for the change in welfare.

We present our theoretical framework and results in Section 2. In our model, each agent chooses the region in which to live and the sector to apply for work—across all sectors in the labor force as well as home production (i.e., labor force non-participation)—to maximize expected utility. Agents who choose to participate in the labor force face a lottery over their employment status, the outcome of which is either obtaining employment in the sector to which they apply or becoming unemployed. Each agent is characterized by an arbitrary vector of productivities (one for each region-job pair, where the set of “jobs” include non-participation, unemployment, and each sector of employment) that shapes the agent’s pecuniary return and a vector of amenity values (one for each region-job pair) that shapes the agent’s non-pecuniary return.² Each region-job pair is characterized by a wage per efficiency unit that determines worker wages (together with worker productivity) and an employment probability that determines the likelihood that the agent will successfully obtain employment in the region-job to which he applies. We do not impose structure on the determination of goods prices, wages per efficiency unit, or employment probabilities; instead, we simply treat their changes as exogenous from the perspective of workers.

We derive a measure of the change in welfare, to a first-order approximation, for workers initially living in a particular region, r , in response to shocks to regional price indices in all regions and changes in wages per efficiency unit, productivity, and employment probabilities in every region-job pair. Our measure of welfare is the sum of equivalent variation across agents initially living in region r (before the realization of unemployment lotteries in the initial equilibrium) relative to the sum of their expected initial incomes.³

Our approach leverages envelope conditions where they apply: changes in labor allocations across regions, sectors, and labor force participation have no first-order effects on welfare. At fixed unemployment rates, welfare of workers starting in region r depends on changes in real productivity-adjusted wages per efficiency unit within each job in region r , weighted by the share of initial income of this region earned within each job. However, because agents do not choose to become unemployed, no envelope condition applies to changes in the probability of unemployment. The elasticity of welfare with respect to the probability of unemployment has two components: a pecuniary component arising from

²Heterogeneity in worker productivity across jobs is at the heart of a substantial literature in labor and international trade; see, e.g., [Roy \(1951\)](#) and Ricardo’s theory of comparative advantage. Heterogeneity in worker preferences across jobs is crucially important for explaining the wage distribution; see, e.g., [Heckman and Sedlacek \(1985\)](#) in labor and [Dix-Carneiro \(2014\)](#) in international trade.

³Throughout, we refer to equivalent variation. However, equivalent and compensating variation are identical to a first-order approximation, so either term applies to our analysis.

the agent-specific loss in income associated with unemployment and a non-pecuniary component arising from the agent-specific gap between the amenity value of employment in the sector to which the agent applies and of unemployment.⁴

In Section 3 we measure welfare changes over the years 2000 to 2007 for the set of workers initially living in each U.S. CZ in the year 2000. Measurement requires constructing a set of initial equilibrium shares and a set of shocks to the economic environment. We require the share of income workers in each CZ initially earn within each sector (including non-employment and each 6-digit NAICS industry) and within unemployment. For employment, we use the Longitudinal Business Database to measure labor income by CZ-sector. For unemployment, we construct state-specific unemployment replacement rates from the Department of Labor. For labor force non-participation, we combine data from the time use of non-market workers on non-market work from the American Heritage Time Use Study with state-level wage income in the private household sector (which proxies for non-market work).

We must also measure the shocks to the economic environment. We measure changes in regional price indices at the state level using data from Hazell et al. (2022). We use two approaches to measure the non-pecuniary (amenity) cost of being unemployed (the utility cost of being unemployed over-and-above the welfare effect of lost income). In our baseline, we set this to zero. In robustness, we instead match this parameter to evidence from a literature that calculates the amount of income necessary to compensate an individual for the change in self-reported, survey-based well-being associated with the loss of a job. Finally, we must measure the changes in productivity-adjusted wages per efficiency unit within each CZ-sector pair. Whereas changes in labor allocations across regions, sectors, and labor force participation have no first-order effects on welfare for given observed changes in wages per efficiency unit, changes in labor allocations do affect the measurement of changes in productivity-adjusted wages per efficiency unit. For example, changes in productivity-adjusted wages per efficiency unit cannot be measured using changes in average wages of workers employed in a region-job because average wages depend on both wages per efficiency units and the selection of workers into regions and jobs based on their efficiency units. To measure changes in productivity-adjusted wages per efficiency unit, we hold worker composition fixed, using a sample of workers who are employed in a common CZ-sector pair over time, leveraging worker-level panel data from the Longitudinal Employer-Household Dynamics program.

Finally, in Section 4 we conclude our analysis by revisiting the substantial empirical liter-

⁴If in practice an agent who is not participating in the labor force is not choosing non-participation—that is, if he is off his labor supply curve—then such an agent should be considered unemployed within our model. In empirical robustness, we consider this alternative mapping between model and data.

ature on local effects of trade shocks, replacing the literature's various dependent variables (changes in separate margins of adjustment) with our measured sufficient statistic for welfare changes. We focus on one particular change in U.S. trade policy, which eliminated potential tariff increases on Chinese imports: the granting of Permanent Normal Trade Relations (PNTR) to China, which passed Congress in 2000 and became effective upon China's accession to the World Trade Organization (WTO) in 2001, as studied in [Pierce and Schott \(2016\)](#) and [Handley and Limão \(2017\)](#). As shown in [Pierce and Schott \(2016\)](#), this policy led to a relative decline in U.S. employment in relatively exposed industries (those with a larger reduction in trade policy uncertainty). To identify the impact of CZ-specific exposure to granting China PNTR, we aggregate [Pierce and Schott](#)'s sector-level measure of exposure, as defined in 1990, up to the CZ level using CZ-sector weights, taking an approach that avoids what [Borusyak et al. \(2022\)](#) refer to as "incomplete shares."

Using a standard differences-in-differences empirical framework, we project CZ-level welfare changes on CZ-specific exposure to granting China PNTR. In our baseline specification, we find that a CZ at the 90th percentile of exposure experiences a 1.9% smaller welfare increase compared to a CZ at the 10th percentile of exposure. Most of this effect is driven by changes in productivity-adjusted wages per efficiency unit rather than by changes in the likelihood of unemployment. Incorporating non-pecuniary costs of unemployment (from which we abstract in our baseline measurement) alters this conclusion in the expected direction. In robustness, we show that doing so increases the magnitude of our estimate, going from a 1.9% smaller welfare increase to a 2.3% smaller welfare increase in the CZ at the 90th percentile of exposure than at the 10th. We also show that incorporating regional price indices is important for quantifying welfare changes. If we instead use a national price deflator, this increases our estimated differential impact of granting China PNTR. The intuition is straightforward: in practice, a region receiving a relatively negative trade shock experiences a relative decline in local prices, which partially offsets the direct negative effects of the trade shock on wages and unemployment. Finally, we document that alternative approaches to measuring welfare yield different implications of granting China PNTR.

We interpret our estimates as the causal effect of granting China PNTR. To strengthen this interpretation, we provide evidence that our baseline results do not reflect a continuation of differential pre-existing trends. This result is consistent with the evidence in [Pierce and Schott \(2016\)](#) that sectors that are more exposed to granting China PNTR do not experience employment losses before China's accession into the WTO. We also show that omitted contemporaneous shocks are likely not driving our results by documenting that our coefficient estimates are quite stable as we progressively add additional controls.

Literature. Our paper is motivated by and builds on a large and growing empirical literature studying the impact of international trade (and other) shocks on local labor market

inequality across a range of margins of adjustment; in the trade literature, see, e.g., [Topalova \(2010\)](#), [Kovak \(2013\)](#), [Autor et al. \(2013\)](#), [Autor et al. \(2015\)](#), [Dix-Carneiro and Kovak \(2017\)](#), [McCaig and Pavcnik \(2018\)](#), and [Dix-Carneiro and Kovak \(2019\)](#). Our objective is to replicate the research design developed in this literature, but use it to identify the relative *welfare* implications of a particular trade shock, by combining theory and data.

Our approach is most related to a trade literature inspired by [Porto \(2006\)](#)—see, e.g., [Atkin et al. \(2018\)](#) and [Borusyak and Jaravel \(2021\)](#)—that similarly calculates the group-specific equivalent variation of a trade shock under minimal functional form restrictions using a first-order approximation.⁵ This literature emphasizes envelope conditions and focuses on distributional changes in welfare induced by differential changes in consumer price indices. We similarly leverage envelope conditions, applied to both the consumption side and labor market adjustment, incorporating a broad range of adjustment margins and the fact that unemployed workers are not on their labor supply curves, so envelope conditions do not apply to changes in unemployment rates. Moreover, we show that even where envelope conditions apply and labor-market reallocation (across jobs and regions) is optimal, changes in labor allocations do affect the measurement of welfare.

The particular trade shock on which we focus in our application builds on the industry-level shock introduced and studied in [Pierce and Schott \(2016\)](#) and [Handley and Limão \(2017\)](#), as briefly described above and expanded upon in Section 4. [Pierce et al. \(2024\)](#) study the effect of this shock on sectoral and regional wages whereas we focus on welfare.⁶

A growing theoretical and quantitative trade literature uses discrete choice models to study the impact of trade on wage inequality and welfare either at the national level—see, e.g., [Burstein et al. \(2019\)](#) and [Lee \(2019\)](#)—or across local labor markets—see, e.g., [Adão \(2015\)](#). Into this environment, [Caliendo et al. \(2018\)](#) and [Adão et al. \(Forthcoming\)](#) introduce a labor-leisure choice and [Kim and Vogel \(2021\)](#), [Galle et al. \(2023\)](#), and [Rodriguez-Clare et al. \(Forthcoming\)](#) additionally introduce frictional unemployment. Our approach dispenses with the functional form restrictions imposed in this literature. Moreover, we measure welfare using equivalent variation, which is cardinal and directly comparable across agents, avoiding the dependence on arbitrary cardinalizations of utility that arises in other approaches; see, e.g., the discussion in [Donald et al. \(2025\)](#).

⁵Our sufficient statistic approach also builds on the literature in public finance trying to isolate robust insights for welfare analysis across models—see, e.g., [Chetty \(2009\)](#)—which has been popularized in the trade literature by [Arkolakis et al. \(2012\)](#).

⁶An input into our welfare is changes in productivity-adjusted wages per efficiency unit. Our measurement of these at the CZ-sector level is related to, but distinct from the outcome variable in [Pierce et al. \(2024\)](#), which is the change in wage income of workers who start in a given CZ-sector pair. Nevertheless, we show that one of their empirical results—that changes in wage income are predicted by regional exposure to the shock rather than by direct sectoral exposure to the shock—applies to our measure of productivity-adjusted wages per efficiency unit as well.

Papers in this literature make different assumptions on the equilibrium determination of wages and employment, with [Adão et al. \(Forthcoming\)](#) incorporating agglomeration, [Kim and Vogel \(2021\)](#) and [Galle et al. \(2023\)](#) incorporating search-and-matching, and [Rodriguez-Clare et al. \(Forthcoming\)](#) incorporating downward wage rigidity. Our approach is consistent with each of these assumptions. Hence, our approach to measuring differential welfare effects of PNTR nests these mechanisms. Of course, our approach is not without cost. First, we cannot answer questions about counterfactual shocks; instead, our approach allows for measuring only observed changes in welfare. Second, whereas we measure changes in welfare for each region, we only estimate the differential impact of granting China PNTR across regions, rather than the level of these welfare effects.⁷

Our paper's objective is related to a macroeconomics literature providing better measures of economic welfare than GDP. Relative to this literature, we focus on changes in welfare (rather than levels), which allows us to generalize the underlying frameworks to incorporate arbitrary heterogeneity across agents, multiple sectors, multiple regions, and frictional unemployment (amongst others). At the same time, we abstract from changes in life expectancy, which play an important role in cross-country comparisons of welfare levels; see, e.g., [Jones and Klenow \(2016\)](#).

2 A theory of measurement

2.1 Model

There is a continuum of risk-neutral agents, denoted by $\omega \in \Omega$. Each agent chooses in which region $r \in \{1, \dots, R\}$ to live and in which sector $s \in \{0, \dots, S\}$ to apply for work, where $s = 0$ indicates the home production sector (labor force non-participation) and all other $s \in \{1, \dots, S\}$ involve participating in the labor force.⁸ An agent who chooses an (r, s) pair with $s \in \{1, \dots, S\}$ lives in region r and enters a lottery in which he is employed in s with probability $E_{rs} \in [0, 1]$ and is unemployed otherwise. An agent who does not participate in the labor force (choosing $s = 0$) does not become unemployed: $E_{r0} = 1$. We denote by $\mathcal{J} \equiv \{u, 0, 1, \dots, S\}$ the set of all possible "jobs," where this set includes unemployment.

An agent ω is characterized by an exogenous vector of productivities that shape his pecuniary reward to being in each region-job pair (r, j) and a vector of exogenous amenity values or compensating differentials that shape his non-pecuniary reward to being in each

⁷We assume away dynamic decisions for workers themselves, unlike [Caliendo et al. \(2018\)](#). However, unlike that paper and the related literature, we allow for the existence of dynamically optimizing firm and capital owners whose decisions shape worker welfare dynamically.

⁸In general, s could refer to firms, industries, occupations, ..., or their intersection. Given our empirical application, we refer to s as a sector throughout.

pair (r, j) . We denote by $\varepsilon_{rj}^\omega \geq 0$ the productivity or efficiency units of agent ω if in (r, j) and by η_{rj}^ω the amenity value experienced by agent ω if in (r, j) .

If agent ω is in (r, j) , his nominal income is given by $\varepsilon_{rj}^\omega \tilde{w}_{rj}$, where \tilde{w}_{rj} denotes the income per efficiency unit for all agents in (r, j) . We assume that this structure describes employment within sectors, non-participation (in which case agents operate a home-production technology), and unemployment (this assumption is consistent with agents in r all earning the same amount in unemployment, or not). Agents have common homothetic preferences across goods. Agent ω in (r, j) obtains

$$C_{rj}^\omega = \varepsilon_{rj}^\omega \tilde{w}_{rj} / P_r$$

units of real consumption, where P_r is the price index facing all agents living in r .⁹ Non-pecuniary benefits of being in (r, j) are additively separable from real consumption. The utility of agent ω in (r, j) who consumes C_{rj}^ω units of the final good is given by

$$U_{rj}^\omega = C_{rj}^\omega + \eta_{rj}^\omega$$

The previous expression contains a key assumption we impose throughout: agents are risk neutral over lotteries offering combinations of real consumption and amenities.¹⁰ Finally, expected utility for agent ω if choosing (r, s) is

$$\mathbb{E}[U^\omega(r, s)] = E_{rs} U_{rs}^\omega + (1 - E_{rs}) U_{ru}^\omega$$

Each agent ω knows ε_{rj}^ω , η_{rj}^ω , and \tilde{w}_{rj} in each (r, j) as well as E_{rs} in each (r, s) and chooses the (r, s) pair that maximizes his expected utility. Denote by

$$\mathcal{X}^\omega \equiv \{(r, s) | (r, s) \in \arg \max_{r', s'} \{\mathbb{E}[U^\omega(r', s')]\}\}$$

the set of agent ω 's expected utility maximizing choices and by

$$\mathbb{E}[U^\omega] \equiv \mathbb{E}[U^\omega(r, s) | (r, s) \in \mathcal{X}^\omega] = \max_{r', s'} \{\mathbb{E}[U^\omega(r', s')]\}$$

her expected utility if applying to *any* optimizing pair $(r, s) \in \mathcal{X}^\omega$. Finally, given agent ω 's

⁹Our theoretical results apply when productivity-adjusted wages per efficiency unit, probabilities of successful employment, and price indices are allowed to vary across workers with different observable characteristics (e.g., education, sex, etc.). We choose to omit this extra layer of notation given that we do not take this approach in our application.

¹⁰However, because we impose no structure on non-pecuniary benefits, it is without loss of generality to allow individuals to place different weights on real consumption relative to these non-pecuniary benefits.

particular optimizing choice $(r, s) \in \mathcal{X}^\omega$, his expected real income is given by

$$\mathbb{E}[RI^\omega] = E_{rs}C_{rs}^\omega + (1 - E_{rs})C_{ru}^\omega$$

where this real income may vary across choices in \mathcal{X}^ω , unlike expected utility.

2.2 Welfare changes for agents initially in a given region

Start from an initial equilibrium in which the set of agents living in region r is denoted by $\Omega_r \subseteq \Omega$ and the set of agents living in region r and applying to sector s is denoted by $\Omega_{rs} \subseteq \Omega_r$.¹¹ We consider the welfare implications, to a first-order approximation, for the set of agents who choose to live in region r in an initial equilibrium, Ω_r , in response to the following shocks: (i) changes in incomes per efficiency unit, w_{kj} , across all regions k and jobs j ; (ii) changes in probabilities of successfully finding employment, E_{ks} , across all regions k and sectors $s \geq 1$; (iii) changes in price indices, P_k , across all regions k ; and (iv) changes in productivity, ε_{ks} , across all regions k and sectors $s \geq 1$.¹² Agents treat these shocks as exogenous.

To measure welfare, we use equivalent variation. Equivalent variation for agent ω , Y^ω , is the real transfer received under the parameter values in the initial equilibrium (and before the realization of unemployment lotteries in the initial equilibrium) at which agent ω is indifferent, in expectation, between the initial and terminal parameter values.

Our measure of regional welfare is the sum of equivalent variation across agents in Ω_r relative to the sum of their real income in the initial equilibrium,¹³

$$EV_r \equiv \int_{\omega \in \Omega_r} \frac{Y^\omega}{RI_r} d\omega$$

where $RI_r \equiv \int_{\omega \in \Omega_r} \mathbb{E}[RI^\omega] d\omega$ is the sum of expected real income across all agents in Ω_r given their choices in the initial equilibrium. We refer to EV_r as the change in welfare for region r . Throughout, EV_r should be interpreted as the the percentage change in money-metric welfare between two years for workers who live in CZ r in the initial year, regardless of where they live or work in the terminal year. Because utility is quasi-linear in consumption and we measure welfare in money units, our EV-based welfare measure is cardinal and directly comparable across agents, avoiding the dependence on arbitrary cardinalizations of

¹¹The initial allocation of agents is an equilibrium only in the sense that each agent $\omega \in \Omega$ is choosing a region and sector to maximize his expected utility given initial parameter values.

¹²Each agent ω is associated with a value of ε_{rs}^ω for each rs . We consider a change over time in this productivity, ε_{rs} , that is common across agents within each rs pair.

¹³With a continuum of agents, expected real income summed across all $\omega \in \Omega_r$ equals realized real income.

utility that arises in other approaches.

In what follows, we consider the case in which the set of agents living in region r and applying to sector s is continuous in the shocks in a neighborhood around the initial equilibrium. Formally, we assume $|\mathcal{X}^\omega| = 1$ for almost all $\omega \in \Omega_r$. Welfare can then be expressed as the sum of four components

$$EV_r \equiv \sum_{s=0}^S \frac{RI_{rs}^e}{RI_r} d \log \left(\frac{w_{rs}}{P_r} \right) + \frac{RI_r^u}{RI_r} d \log \left(\frac{w_{ru}}{P_r} \right) \\ + \sum_{s=1}^S \frac{RI_{rs}^e - \frac{E_{rs}}{1-E_{rs}} RI_{rs}^u}{RI_r} d \log E_{rs} + \sum_{s=1}^S \frac{\eta_{rs}^e - \frac{E_{rs}}{1-E_{rs}} \eta_{rs}^u}{RI_r} d \log E_{rs} \quad (1)$$

where $d \log w_{rs} \equiv d \log \tilde{w}_{rs} + d \log \varepsilon_{rs}$ is the combination of changes in wages per efficiency unit and changes in productivity (productivity-adjusted wages per efficiency unit); $RI_{rs}^e \equiv E_{rs} \int_{\Omega_{rs}} C_{rs}^\omega d\omega$ is total real income earned in employment among those who applied to (r, s) in the initial equilibrium, $RI_{rs}^u \equiv (1 - E_{rs}) \int_{\Omega_{rs}} C_{ru}^\omega d\omega$ is total real income earned in unemployment among those who applied to (r, s) in the initial equilibrium, $RI_r^u \equiv \sum_{s=1}^S RI_{rs}^u$ is total real income earned in unemployment among those in region r in the initial equilibrium, $\eta_{rs}^e \equiv E_{rs} \int_{\Omega_{rs}} \eta_{rs}^\omega d\omega$ is total non-pecuniary benefits summed across employed agents who applied to (r, s) in the initial equilibrium, and $\eta_{rs}^u \equiv (1 - E_{rs}) \int_{\Omega_{rs}} \eta_{ru}^\omega d\omega$ is total non-pecuniary benefits summed across unemployed agents who applied to (r, s) in the initial equilibrium. The first term on the right-hand side of (1) corresponds to the change in welfare induced by changes in productivity-adjusted wages per efficiency unit within sectors, $s = \{0, \dots, S\}$, weighted by the initial share of real income earned within each sector. The second term corresponds to the change in welfare induced by changes in the real wage per efficiency unit within unemployment weighted by the initial share of real income earned within unemployment. The third and fourth terms correspond to the change in welfare induced by changes in employment probabilities, where the third term measures changes associated with the pecuniary, real consumption losses from unemployment and the fourth term measures changes associated with the non-pecuniary costs of unemployment.

Intuition. To understand equation (1) and its implications for measurement, it is useful to start from a simpler framework that abstracts from all margins of adjustment and then build back up to our baseline framework, step by step. First, consider a framework in which agents are identical, cannot choose where to live, there is a single sector, and there is no unemployment. In this case, we obtain the standard result that the change in welfare of

workers in region r equals the log change in real income for all agents in region r ,

$$EV_r = d \log \left(\frac{w_r}{P_r} \right)$$

In this calculation, the change in the regional real wage per efficiency unit w_r / P_r could be measured as the average change in real wage income across all workers in r , without concern for changes in the composition of these workers, since workers are not changing their region or sector. Second, additionally incorporate agent choice over region-sector pairs subject to worker-region-sector-specific productivities and amenities. Then

$$EV_r = \sum_{s=0}^S \frac{RI_{rs}^e}{RI_r} d \log \left(\frac{w_{rs}}{P_r} \right)$$

Since there is no unemployment, in this case RI_{rs}^e is both the earnings within employment of those who applied to (r, s) and their total earnings; for the same reason, RI_r is the sum across $s \geq 0$ of RI_{rs}^e , as there is no unemployment income. In spite of agents choosing (r, s) pairs, there is no additional term in welfare associated with changes in agent choices. This result follows from the combination of an envelope condition (which implies that agents who switch their choices are initially indifferent between choices) and the assumption that the measure of switchers is zero. This combination implies that the change in welfare associated with switching optimal choices is of second order. This envelope condition logic applies to welfare but not to earnings, since earnings are not necessarily equated across choices between which an agent is initially indifferent. Hence, changes in average wages among the employed within (r, s) over time do not measure changes in wages per efficiency unit, w_{rs} , unlike in the simpler model.

Third, additionally incorporate unemployment, but assume that unemployment probabilities are fixed. In this case,

$$EV_r = \sum_{s=0}^S \frac{RI_{rs}^e}{RI_r} d \log \left(\frac{w_{rs}}{P_r} \right) + \frac{RI_r^u}{RI_r} d \log \left(\frac{w_{ru}}{P_r} \right)$$

which is equivalent to the first two terms in equation (1). Relative to the simpler model above, total real income RI_r depends on incomes earned within sectors and unemployment and RI_{rs}^e now differs from the real income earned by those who apply to (r, s) . Moreover, the second term is incorporated because agents who apply to (r, s) for $s \geq 1$ may become unemployed, so changes in real income within unemployment must be included and must be weighted symmetrically to any other job.

Finally, additionally incorporate changes in the probability of employment. In this case,

we obtain equation (1), the third and fourth terms of which represent changes in expected total real income and compensating differentials, respectively, induced by changes in the probability of successfully finding employment, evaluated at the initial pecuniary and non-pecuniary costs of unemployment. The reason RI_{rs}^u is multiplied by $E_{rs}/(1 - E_{rs})$ in the third term is to construct the difference between real income for all agents who applied to (r, s) if they were all employed and if they were all unemployed (all multiplied by the probability of employment) in the initial equilibrium. The same applies to η_{rs}^u in the fourth term.

The third and fourth terms appear in welfare, but corresponding terms associated with changing probabilities of workers applying to different region-sector pairs do not, because an envelope condition applies to optimal choices of (r, s) . Hence, any change in the non-pecuniary component of welfare associated with switches across (r, s) pairs must exactly cancel out with changes in the corresponding pecuniary component. No equivalent envelope condition applies to unemployment because agents do not choose to become unemployed.

Discussion. Two points about our welfare measure, EV_r , are worth emphasizing. First, EV_r captures changes in welfare for a fixed set of agents over time. It is the percentage change in money-metric welfare between two years for workers who live in CZ r in the initial year, regardless of where they live or work in the terminal year. This is distinct from measuring changes in welfare across cohorts born into CZ r over time. For instance, if worker productivity rises with age, our measure of EV_r will capture this as an increase in worker welfare. But this increase in worker welfare does not imply that new cohorts born into CZ r are better off than previous cohorts. We view this as a desirable feature of EV_r , not a limitation.

Second, in a dynamic environment, well-being depends on wages, productivities, unemployment probabilities, and prices in the present and future periods. In this respect, EV_r captures a static measure of welfare. We impose this restriction throughout.

3 Measuring welfare changes

In Section 3.1 we impose additional restrictions on our framework that facilitate measurement. In Section 3.2 we describe how we measure EV_r given these restrictions. We define U.S. regions as mainland commuting zones (CZs) and we define jobs as the union of labor force non-participation, unemployment, and 6-digit NAICS codes. In Section 3.3 we provide results on the unconditional distribution of welfare changes across U.S. CZs over the period 2000-2007.

3.1 Restrictions and their implications

We impose two restrictions. First, we assume that region-sector-specific unemployment rates are common across sectors, varying only across regions, so that $E_{rs} = E_r$. We make this assumption because there is no dataset with an unemployment rate that is specific to the sector to which an agent applies.¹⁴

Second, we assume that the average compensating differential associated with successful employment in sector s among the applicants to (r, s) in the initial equilibrium, denoted by $\bar{\eta}_{rs} \equiv \frac{1}{|\Omega_{rs}|} \int_{\Omega_{rs}} \eta_{rs}^\omega d\omega$, minus the average compensating differential associated with unemployment among the applicants to (r, s) in the initial equilibrium, denoted by $\bar{\eta}_{ru|s} \equiv \frac{1}{|\Omega_{rs}|} \int_{\Omega_{rs}} \eta_{ru}^\omega d\omega$, is common across sectors s within region r : $\tilde{\eta}_r \equiv \bar{\eta}_{rs} - \bar{\eta}_{ru|s}$. We make this assumption because of a lack of data on variation in non-pecuniary costs of unemployment across workers who apply to different sectors.

These restrictions simplify the measurement of welfare changes for agents who are initially in region r . In particular, we obtain

$$EV_r = \underbrace{\sum_{j \in \mathcal{J}} \frac{I_{rj}}{I_r} d \log \left(\frac{w_{rj}}{P_r} \right)}_{\text{real wage}} + \underbrace{(1 + \eta_r) \frac{I_r^e}{I_r} \left[1 - \frac{\bar{I}_r^u}{\bar{I}_r^e} \right] d \log E_r}_{\text{unemployment}} \quad (2)$$

where the summation in the first term on the right-hand side of equation (2) is taken across all jobs (recall that jobs include each sector of employment, out of the labor force, and unemployment). In equation (2), I_{rj} denotes the total income of agents in region r who are in job j ; $I_r^e \equiv \sum_{s=1}^S I_{rj}$ is the total wage income earned by the employed; $I_r \equiv \sum_j I_{rj}$ is the total income across all jobs; and \bar{I}_r^e and \bar{I}_r^u are the average incomes of the employed and the unemployed, so that $\bar{I}_r^u / \bar{I}_r^e$ is a measure of the unemployment replacement rate. Finally $\eta_r \equiv \tilde{\eta}_r / (\bar{I}_r^e - \bar{I}_r^u)$ is the average non-pecuniary welfare loss of unemployment, $\tilde{\eta}_r$, relative to the average pecuniary welfare loss of unemployment, $\bar{I}_r^e - \bar{I}_r^u$; this represents a transformation of the utility cost of unemployment (via compensating differentials) into its corresponding nominal value. Each of these variables is measured in the initial period.

The first term on the right-hand of (2) represents the component of welfare associated with changes in productivity-adjusted real wages per efficiency unit; for compactness, we refer to this as the real wage component. The weight on a particular job in this component of equation (2) is the share of total region-level income that is earned within that job in the initial period. This term coincides with the first two terms in equation (1). The second term on the right-hand side of equation (2) represents the combination of the pecuniary and non-

¹⁴As described in the theory section, we can easily relax this assumption (and other related assumptions), e.g., allowing unemployment rates to vary not only across regions, but also across worker observables.

pecuniary components of changes in welfare associated with changes in the probability of successfully obtaining employment. In this term, $1 - \bar{I}_r^u / \bar{I}_r^e$ represents the average percent loss in income associated with unemployment in the initial period; and $1 + \eta_r$ magnifies this income loss to incorporate potential non-pecuniary losses associated with unemployment. This term coincides with the final two terms in equation (1).

Equation (2) can be used to measure changes in welfare between two arbitrary time periods for the set of agents choosing to live in region r in the initial period, wherever they live in the terminal period.

3.2 Data

In this section we briefly describe how we measure each term in equation (2). Additional details are provided in the Online Appendix. We define regions as mainland commuting zones (CZs) and we define sectors as the union of 6-digit 1997 NAICS codes and labor force non-participation. In our benchmark exercises we study the period 2000-2007. We also measure changes in welfare over a pre-period for placebo analyses, but focus on describing measurement over the benchmark period below.

Basic elements: We measure region-specific inflation, $d \log P_r$, using state-level CPI data from Hazell et al. (2022). We use the national Consumer Price Index for all CZs in states without state-level CPI data. In robustness, we apply the national CPI for all CZs. We measure the employment rate, E_r , as the ratio of the employed population relative to the labor force (one minus the unemployment rate) using the 2000 census and the combined 2006, 2007, and 2008 1% American Community Survey samples.

We use the Longitudinal Business Database (LBD) to measure the total income of agents in region r who are successfully employed within each sector s , I_{rs} for $s = \{1, \dots, S\}$; given these values, we solve for the total income of the employed, I_r^e . We construct unemployment income, I_r^u , by combining I_r^e with the state-specific unemployment replacement rate from the Department of Labor and the ratio of unemployed to employed workers from the 2000 census. We construct the income of agents not in the labor force, I_{r0} , by combining I_r^e with a state-specific home production replacement rate (constructed from time use data from the American Heritage Time Use Study and wage data from the 2000 census) and the ratio of non-participants to employed workers from the 2000 census. We construct total income I_r as the sum of I_{r0} , I_r^e , and I_r^u .

This leaves two remaining terms in equation (2): the non-pecuniary relative to the pecuniary cost of unemployment, η_r , and changes in productivity-adjusted wages per efficiency unit, $d \log w_{rj} = d \log \tilde{w}_{rj} + d \log \varepsilon_{rj}$, in each region-job pair.

The non-pecuniary cost of unemployment (η_r). Measuring η_r is difficult, as it reflects non-

pecuniary costs that cannot be recovered easily from agent choices given that agents do not choose to be unemployed. In our baseline, we set $\eta_r = 0$, assuming that workers are indifferent between unemployment and employment, conditional on income. This is a strong assumption that is at odds with introspection and a vast literature—cutting across economics, psychology, and marketing (among others)—that calculates the amount of income necessary to compensate an individual for the change in self-reported well-being associated with the loss of his job; see, e.g., [Winkelmann and Winkelmann \(1998\)](#) and [Blanchflower and Oswald \(2004\)](#).

In an alternative approach, we parametrize η_r to evidence in [Knabe and Ratzel \(2011\)](#). Their estimate provides a lower bound on the relative importance of the non-pecuniary component within this literature.¹⁵ They find that the transfer required to maintain well-being after entering unemployment is approximately three times the income loss associated with unemployment. In our framework, this translates into a value of $\eta_r = 2$, so that the sum of the pecuniary and non-pecuniary costs of unemployment are three-times larger than the pecuniary costs alone. We use this value for all r in robustness.

Changes in productivity-adjusted wages per efficiency unit ($d \log w_{rs}$). Measuring $d \log w_{rs}$ is also difficult. Productivity-adjusted wages per efficiency unit are not observed directly. The observed average wage within a region-sector-period depends both on the productivity-adjusted wage per efficiency unit and worker composition (the average value of ε_{rs}^ω across workers employed within rs in period t).¹⁶

To address these issues, we measure $d \log w_{rs}$ by leveraging panel data on individual earnings, which allows us to hold worker composition fixed. We take this approach using the Longitudinal Employer-Household Dynamics (LEHD) data, a restricted-use dataset constructed by the U.S. Census Bureau from quarterly earnings records submitted by employers to state unemployment insurance (UI) systems. The core LEHD data contain total quarterly earnings for each worker-employer pair, along with basic worker and establishment characteristics such as industry and geographic location. Coverage includes approximately 95% of private-sector employment and most state and local government workers, but excludes federal employees, military personnel, and the self-employed.

We construct a sample of workers from the LEHD who are employed in region r and sector s in both the initial and terminal periods. On this sample, which includes roughly 15 million unique workers in our baseline period, we measure $d \log w_{rs}$ as the average log change in nominal wage income. Finally, we choose changes in nominal returns to unemployment

¹⁵Their estimate of the non-pecuniary cost is lower than others because they incorporate into the pecuniary component not only current income losses associated with unemployment, but also future income losses.

¹⁶Adjusting wages for observable worker characteristics only partially addresses this issue of selection, since workers also select on unobservable margins; our ε_{rs}^ω captures both observables and unobservables.

and non-employment such that real wage changes across time are zero: $d \log(w_{rj}/P_r) = 0$ for $j \in \{u, 0\}$ in all r .

Our approach is distinct from measuring $d \log w_{rs}$ using changes in income for the set of agents *initially* employed in a given region-sector. In the presence of compensating differentials, a worker's income might increase or decrease discretely when reallocating across rs pairs. Hence, when a worker reallocates across jobs (whether into another sector, into non-employment, or into unemployment) or across regions, his change in income is generically not informative of changes in productivity-adjusted wages per efficiency unit within his initial sector of employment.

Given LEHD coverage over the sample period, our baseline includes approximately 600 mainland commuting zones and approximately 450 sectors.¹⁷ Given the large number of region-sector pairs, each value of $d \log w_{rs}$ is clearly measured with error. However, these estimates are then averaged within each region to construct approximately 600 values of the first component of equation (2), which should minimize concerns over measurement error at the rs level.¹⁸

3.3 Unconditional welfare changes across CZs

Figure 1 displays the distribution of welfare changes between 2000 and 2007 across CZs, with each CZ weighted by labor income for the employed, I_r^e .¹⁹ The median CZ experienced an increase in welfare equivalent to a 12.4% increase in real income in 2000. There is substantial dispersion in welfare changes, with an increase of 6% at the 10th and of 17.2% at the 90th percentiles.

Figure 1 also displays the distribution of changes in a more standard measure of welfare: the (log) change in average real income, measuring nominal wage income in each year using the average of the nominal wage taken across all employees (using the LEHD). This approach ignores changes in the composition of the workforce on observable and unobservable dimensions, changes in compensating differentials associated worker reallocation, and changes in unemployment.

The distribution of our measure of welfare changes is shifted to the right relative to the distribution of changes in real income for a simple reason: EV_r captures changes in welfare

¹⁷To comply with U.S. Census Bureau confidentiality and disclosure requirements, we do not report the exact counts of commuting zones and sectors. The number of observations reported here and elsewhere is rounded.

¹⁸For this reason, we do not focus on welfare at the more disaggregated level of agents initially applying to different (r, s) pairs.

¹⁹Because our welfare measures incorporate LEHD data, U.S. Census Bureau disclosure restrictions prohibit releasing statistics for individual commuting zones (and thus preclude presenting CZ-level maps). We therefore report only the distribution of welfare changes.

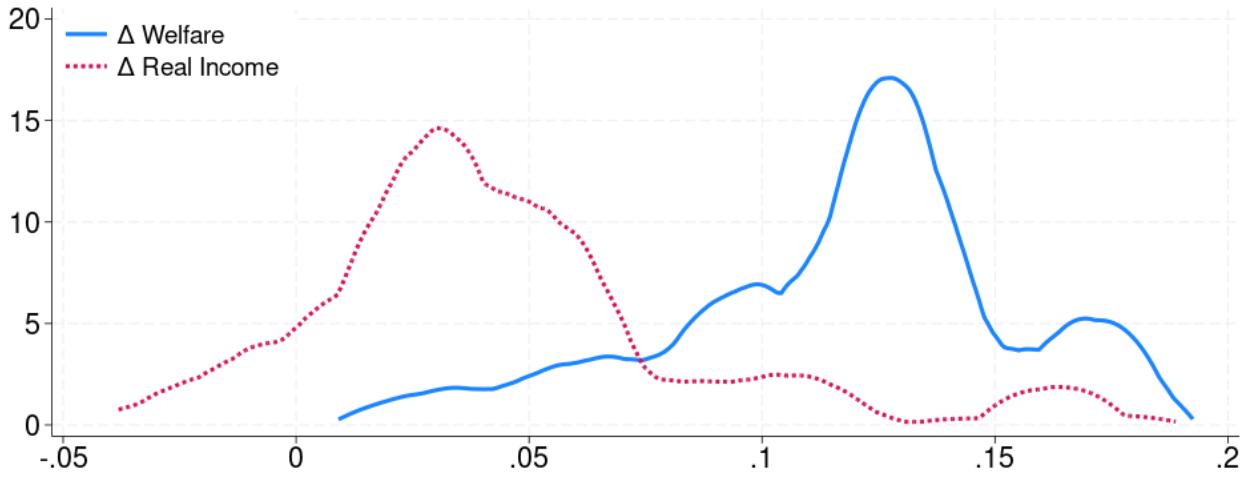


Figure 1: The distribution of changes in welfare across CZs, 2000-2007

for a fixed set of agents over time whereas the change in real income plotted in Figure 1 does not. As workers age, their wages tend to rise. In our framework, this is captured by an increase in productivity-adjusted wages per efficiency unit. This shifts the distribution of EV_r to the right. On the other hand, changes in real incomes plotted in Figure 1 do not capture increases in worker productivity as workers age, because it holds the age range (18-64) over which wages are averaged fixed in 2000 and 2007. If we instead construct an alternative change in real wage income that averages over workers aged 18-57 in 2000 and aged 25-64 in 2007 (which holds the age cohort fixed, as in the case of EV_r), the mean of the two welfare measures are very similar. This implies that the rightward shift between the two welfare measures in Figure 1 is driven by a focus on a fixed set of agents in EV_r versus a changing set of agents in real income, a point discussed in detail in Section 2.2.

To what extent is the variation in our welfare measure across U.S. commuting zones accounted for by variation in each of the CZ-specific components of welfare? To answer this question, we decompose the variance of welfare changes into the real wage component (recall that this contains changes in productivity and changes in real wages per efficiency unit), the unemployment component, and (two times) their covariance. Changes in real wages account for 93% of the variance, changes in unemployment for 2% of the variance, and the covariance between these components the remainder.

The unconditional variation in welfare across space is, of course, a function of all shocks. It remains to identify the impact of the trade shock in our particular application—granting China PNTR—on welfare. It also remains to identify the extent to which the welfare effects of granting China PNTR are transmitted through each of the welfare components. We next turn to these issues.

4 The welfare effects of granting China PNTR

In this section, we identify the causal impact of a trade shock—granting Permanent Normal Trade Relations to China—on the welfare of agents initially living across U.S. commuting zones. In Section 4.1 we describe the sector-level trade shock on which we focus and how we use it to construct the CZ-level trade shock. In Section 4.2 we detail our empirical design. In Section 4.3 we present our main results and sensitivity exercises. Finally, in Section 4.4 we compare our results to a set of alternative approaches.

4.1 The trade shock

Following [Pierce and Schott \(2016\)](#) (henceforth PS) and [Handley and Limão \(2017\)](#), we study the impact of granting Permanent Normal Trade Relations (PNTR) to China. As argued in these papers, China's accession into the WTO was not accompanied by marked reductions in U.S. tariffs. In practice, China's exports to the U.S. both before (since 1980) and after its accession to the WTO were subject to Normal Trade Relation (NTR) tariff rates reserved for WTO members.

However, whether or not these NTR rates would continue to apply to imports from China was subject to annual renewal. If these rates were not renewed, then tariffs on Chinese imports would have risen to non-NTR rates, from an average of about 6 percent to an average of about 31 percent in 1990. PS refer to the sector-specific gap between the NTR and the non-NTR rates as the 'NTR gap.' The U.S. Congress voted to grant China PNTR in October 2000 and this change in policy became effective at the end of 2001. The elimination of this tariff uncertainty then provides a heterogeneous reduction in trade policy uncertainty across sectors, measured by the initial gap as defined in 1990, which we denote by NTR gap_s .

We construct our CZ-level trade shock, which we denote by NTRG_r , as a region-specific weighted average of sector-specific NTR gaps,

$$\text{NTRG}_r \equiv \sum_{s=1}^S \omega_{rs} \text{NTR gap}_s$$

Weights, ω_{rs} , are employment shares measured using the 1980 LBD.²⁰ We normalize weights to sum to one across manufacturing sectors s in which the NTR gap_s is defined to avoid what [Borusyak et al. \(2022\)](#) refer to as "incomplete shares." Specifically, the weight ω_{rs} equals zero in any non-manufacturing sector s or any sector s in which the NTR gap_s is undefined. In any manufacturing sector in which the NTR gap_s is defined, the weight ω_{rs} equals employ-

²⁰We use the LBD instead of the LEHD because the LEHD did not exist in 1980 and had limited state coverage until the 1990s.

ment in sector s in region r in year 1980 relative to employment in region r in year 1980 summed across all manufacturing sectors in which NTR gap $_s$ is defined. Figure 2 displays the histogram of NTRG $_r$ across CZs.

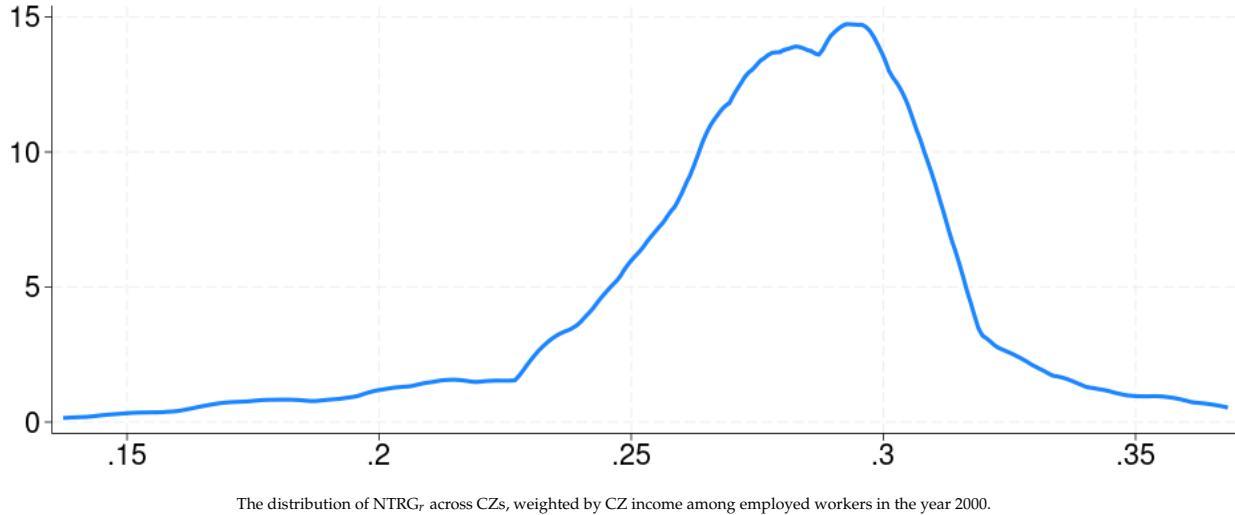


Figure 2: The distribution of NTRG $_r$ across CZs.

4.2 Empirical design

Specification. Our empirical design compares changes in welfare between 2000 and 2007 across U.S. commuting zones, EV_r , that are more and less exposed to the trade shock, but similar in terms of other observable characteristics. We estimate

$$EV_r = \alpha + \beta NTRG_r + \mathbf{X}'_r \gamma + \epsilon_r \quad (3)$$

using OLS. The regression coefficient β identifies differential welfare changes across commuting zones following granting China PNTR. We interpret β as the causal effect of granting China PNTR. Alternatively, β could reflect differential pre-existing trends or omitted contemporaneous shocks. To address concerns of omitted variable bias, we include a vector of CZ controls, \mathbf{X}'_r , that we describe below. We test for pre-existing trends in Section 4.3.

We also estimate versions of equation (3) in which we replace the dependent variable with its two components: the real wage component and the unemployment component of equation (2). Given the properties of OLS, these specifications decompose welfare changes into these two components, with the sum of the coefficients in the decomposition equalling the corresponding value in the baseline specification.

In all regressions, we weight observations by each CZ's initial income among the employed, I_r^e . In our baseline we cluster standard errors by state; we show that [Borusyak et al.](#)

(2022) standard errors are very similar both in the baseline and in the decomposition.

Controls. A necessary condition to obtain a consistent estimate of the causal effect of interest, β , using OLS to estimate (3) is that $\mathbb{E}[\epsilon_r ntrg_r] = 0$, where $ntrg_r$ is the residualized value of $NTRG_r$ after projecting on a constant and \mathbf{X}_r .

The vector \mathbf{X}_r in (3) contains CZ-specific 1990 labor force and demographic composition that might be correlated with $NTRG_r$ and independently affect our dependent variables. Following Autor et al. (2013), this vector includes the manufacturing share; the share of the population that is college educated and the share that is foreign born; the share of employment among women; regional fixed effects for the nine Census divisions, which absorb trends that are common across CZs within broad geographic regions; the share of employment in routine occupations; and the average offshorability index of occupations.²¹ We take the values of these control variables directly from Autor et al. (2013). We also include two additional controls: the share of employment in the year 1980 in industries in which NTR gap_s is not defined (measured using the 1980 LBD) and the 1980 employment to population rate (measured using the 1980 census).

4.3 Granting China PNTR and local welfare

Baseline results. Panel A of Table 1 reports our baseline results on the impact of granting China PNTR on relative welfare across U.S. CZs between 2000 and 2007. The first column includes no controls. The second column controls for the initial employment composition of the commuting zone: the manufacturing share, the share of employment in sectors in which the NTR gap_s is undefined, and the labor force participation rate. The third column additionally controls for census division fixed effects. The fourth column incorporates demographic characteristics: the college share and foreign-born shares of the population and the female share of employment. The fifth and final column additionally includes controls for the routine occupation share and the measure of average offshorability of occupations. The impact of the $NTRG_r$ on welfare is broadly stable as we progressively add controls across columns. The column 5 point estimate implies that granting China permanent normal trade relations lowers the welfare of a CZ at the 90th percentile of exposure by 1.9 log points (or percent) relative to a CZ at the 10th percentile, as reported at the bottom of the table.

Panel B of Table 1 additionally decomposes the causal effect of granting China PNTR on relative welfare into its impact through changes in real wages per efficiency unit (combined with productivity) and its impact through changes in unemployment rates (across all

²¹ Routineness and offshorability are based on occupation-level data that are aggregated up to the CZ level. Routine occupations are intensive in tasks that follow a set of precise rules and procedures and are more easily substitutable with computers. Offshorable occupations are intensive in tasks that make little use of face-to-face interaction and are more easily executed from a distance.

five specifications).²² Approximately 90 percent ($\approx 0.224/0.248$) of the differential welfare effects of the trade shock across regions in our baseline specification (column 5) are transmitted through changes in real wages per efficiency unit. One reason why changes in real wages have a relatively larger differential effect on welfare across CZs is because the impact of changes in unemployment on welfare depend on the unemployment replacement rate, which averages approximately 43% across regions. Another reason is that in our baseline analysis we have set the non-pecuniary cost of unemployment to zero; we check the robustness of our results to this choice below. A final reason is that we have assumed in our baseline analysis that non-employed workers who are not actively looking for a job are on their labor supply curves; we also check the robustness of our results to this choice below.

	(1)	(2)	(3)	(4)	(5)
Panel A: EV					
NTRG _r	-0.214*	-0.190	-0.183**	-0.243***	-0.248***
Robust state clustered SE	(0.122)	(0.120)	(0.088)	(0.084)	(0.096)
$\mathbb{E}[EV_r : NTRG_{p90} - NTRG_{p10}]$	-.017	-.015	-.014	-.019	-.019
Panel B: EV Decomposition					
<i>Real wage:</i>					
NTRG _r	-0.175	-0.162	-0.162*	-0.211***	-0.224**
Robust state clustered SE	(0.118)	(0.118)	(0.085)	(0.082)	(0.093)
BHJ SE	(0.096)	(0.085)	(0.062)	(0.066)	(0.069)
<i>Unemployment:</i>					
NTRG _r	-0.039***	-0.027***	-0.021**	-0.032***	-0.024***
Robust state clustered SE	(0.013)	(0.009)	(0.009)	(0.008)	(0.007)
BHJ SE	(0.007)	(0.009)	(0.009)	(0.009)	(0.009)
Observations	600	600	600	600	600
manufacturing share ₋₁	✓	✓	✓	✓	✓
missing NTRG sector share ₋₁	✓	✓	✓	✓	✓
employment-to-population ₋₁	✓	✓	✓	✓	✓
college share ₋₁			✓	✓	✓
foreign born share ₋₁			✓	✓	✓
female share ₋₁			✓	✓	✓
routine occupation share ₋₁				✓	
average offshorability ₋₁					✓
regional FE			✓	✓	✓

Panel A reports the overall welfare effects across five specifications with different sets of controls. Panel B decomposes these effects into changes in the real wage and the unemployment components. Panel A reports robust standard errors clustered by state while Panel B reports robust standard errors clustered by state and [Borusyak et al. \(2022\)](#) standard errors. Regressions are weighted by each CZ's income of the employed, I^e in 2000. Manufacturing share refers to the share of employed income earned within manufacturing; college share and foreign born share refer to the relevant shares of the population; female share refers to the female share of employment; routine occupation share and average offshorability refer to the share of employment in the most routine and the most offshorable occupations. These controls are defined in 1990 and taken from [Autor et al. \(2013\)](#). Missing NTRG sector share measures the share of 1980 employment in sectors in which the NTR gap_r is not defined, measured using the 1980 LBD. Employment-to-population measures the employment to population rate, measured using the 1980 census. The number of CZ observations is approximate due to U.S. Census Bureau disclosure avoidance requirements.

Table 1: The welfare effect of granting China PNTR and its decomposition.

²²Panel A displays the overall welfare effects with robust standard errors clustered by state, while Panel B shows the decomposition into real wage and employment components with both robust standard errors clustered by state and [Borusyak et al. \(2022\)](#) standard errors, which are similar. BHJ standard errors for Panel A are available upon request. They are smaller than the robust SEs, meaning all results remain statistically significant, and follow similar patterns to those in Panel B. For example, in column 5, the BHJ SE for the overall welfare effect is 0.073 compared to the robust SE of 0.096.

	$d \log w_{rs}$		
	(1)	(2)	(3)
Sectoral NTR Gap _s	-0.056 (0.042)		-0.004 (0.033)
Regional NTRG _r		-0.573*** (0.146)	-0.569*** (0.137)
Baseline regional controls		✓	✓
$\mathbb{E}[d \log w_{rs} : \text{NTRG}_{p90(r)} - \text{NTRG}_{p10(r)}]$		-.055	-.054
$\mathbb{E}[d \log w_{rs} : \text{NTR Gap}_{p90(s)} - \text{NTR Gap}_{p10(s)}]$	-.015		-.0012
Observations	26500	26500	26500

Each column present results of regressing $d \log w_{rs}$, measured between 2000 and 2007, on either NTR Gap_s (in column 1), or on NTRG_r and the baseline set of controls included in column 5 of Table 1 (in column 2), or both (in column 3). Regressions are weighted by CZ-sector employment and standard errors are clustered by state. The number of CZ-sector observations is an approximation (because of reporting requirements).

Table 2: Changes in wages per efficiency unit, sector-level NTR Gap_s, and CZ-level NTRG_r.

What explains the strong correlation between the real wage component of regional welfare and the regional measure of exposure to granting China PNTR? One possibility is that the combination of changes in real wages per efficiency unit and productivity at the region-sector level, $d \log w_{rs}$, are similar across regions within a sector, but that some regions have higher shares of their incomes earned in more exposed sectors, where changes in $d \log w_{rs}$ are less positive. Another possibility is that changes in the combination of wages per efficiency unit and productivity at the region-sector level are less positive in more exposed regions within sectors. We investigate these possibilities by regressing $d \log w_{rs}$ separately and jointly on the sectoral and regional measures of the NTR Gap. Results are reported in Table 2. We find that the regional measure of the NTR Gap is strongly and negatively correlated with $d \log w_{rs}$ whereas the sectoral measure of the NTR Gap is not. From this, we conclude that the real wage component of regional welfare falls more in regions with a higher regional measure of exposure to granting China PNTR because the combination of changes in real wages per efficiency unit and productivity fall more within sectors in these regions. This conclusion is related to results in [Pierce et al. \(2024\)](#). Also using the LEHD, they find that granting China PNTR lowered incomes of workers initially employed in more exposed regions, but within a region, not more so in sectors that are more directly exposed. Our results in Table 2 differ from [Pierce et al. \(2024\)](#) because we measure changes in the combination of wages per efficiency unit and productivity (on the sample of workers who start and end employed in the same region-sector pair), which is what our theory requires for welfare analysis, whereas they measure changes in incomes for workers who start employed in a given region and sector, which is their focus.

Falsification. To verify that our results capture the effects of exposure to granting China PNTR rather than the continuation of pre-existing trends, we conduct a falsification exercise in which we regress changes in welfare over an earlier period on future exposure to granting

China PNTR. The earlier period we consider is 1992 through 2000. We use 1992 as the initial year because state-level coverage in the LEHD, which we use to measure changes in wages per efficiency unit in each region-sector pair, increases over time. Whereas in our baseline period 2000-2007 there are 34 states and approximately 600 CZs covered across all years, in our falsification period of 1992-2000 there are only 6 states and approximately 80 CZs.²³

	EV _r		
	2000-2007		1992-1999
	(1)	(2)	(3)
NTRG _r	-0.248*** (0.096)	-0.314** (0.141)	0.302 (0.180)
$\mathbb{E}[EV_r : NTRG_{p90}-NTRG_{p10}]$	-.019	-.024	.023
Baseline sample	✓		
Placebo sample		✓	✓
Observations	600	80	80

Column 1 replicates Column 5 in Table 1. Column 2 displays results of estimating this specification on the smaller sample of CZs observed in the LEHD over the 1992 - 2000 period. Finally, column 3 replaces the 2000-2007 change in welfare with the value between 1992 and 2000 and estimates on the smaller sample of column 2. The number of CZ observations is an approximation (because of reporting requirements).

Table 3: Falsification – Changes in welfare between 1992 and 2000.

We first document that the decrease in state and CZ coverage associated with measuring EV_r over the pre-period does not materially affect our baseline results over the 2000-2007 period. Column 1 of Table 3 replicates our baseline result from column 5 of Table 1, whereas column 2 of Table 3 displays results when estimating the baseline specification over the 2000-2007 period, but on the smaller sample of CZs. Point estimates are very similar. Although standard errors increase slightly, the result remains statistically significant at the 5% level. In column 3 we conduct our falsification exercise on the smaller CZ-level sample. This result provides little evidence consistent with a continuation of differential pre-trends across more and less exposed CZs. First, the estimated coefficient is insignificantly different from zero. Second, and more importantly, welfare was, if anything, rising in more exposed CZs before China was granted PNTR, suggesting a larger trend break in welfare changes than implied by our baseline estimation alone.

Sensitivity. In this section we describe results of a range of robustness exercises, each summarized in Table 4. The first column of Table 4 replicates our baseline result from column 5 of Table 1.

In our baseline we assume that the non-pecuniary welfare cost of unemployment is zero, setting $\eta_r = 0$ in all CZs r . In our first sensitivity exercise, we instead set $\eta_r = 2$ in all CZs r , so that the non-pecuniary welfare cost of unemployment is twice the pecuniary component, as measured by Knabe and Ratzel (2011). The distribution of welfare changes between 2000 and 2007 remains similar to in our baseline case. The differential welfare impact of granting

²³See the Online Appendix for details on the empirical construction of EV_r in the pre-period.

	EV _r					
	(1)	(2)	(3)	(4)	(5)	(6)
NTRG _r	-0.248*** (0.096)	-0.296*** (0.104)	-0.254** (0.100)	-0.248** (0.100)	-0.290*** (0.087)	-0.314*** (0.120)
E[EV _r : NTRG _{p90} -NTRG _{p10}]	-.019	-.023	-.02	-.019	-.022	-.024
Baseline	✓					
Non-pecuniary cost of unemployment		✓				✓
NILF + unemployment, adjusted replacement			✓			✓
NILF + unemployment, baseline replacement				✓		
National price index					✓	
Observations	600	600	600	600	600	600

Column 1 replicates Column 5 in Table 1. The remaining columns estimate the same specification but with slight adjustments to the measurement of welfare. Column 2 incorporates a non-pecuniary cost of unemployment into the calculation of welfare, setting $\eta = 2$. Column 3 treats non-participants as unemployed and adjusts the unemployment replacement rate to include income both among the unemployed and those out of the labor force; NILF refers to not-in-the-labor-force. Column 4 treats non-participants as unemployed using the baseline measure of the unemployment replacement rate. Column 5 uses national rather than regional price indices.

Table 4: Sensitivity analysis of the welfare effect of granting China PNTR.

China PNTR rises slightly, as shown in column 2 of Table 4, driven by an increase in the unemployment component, which now accounts for a quarter of the overall welfare effect, as compared to a tenth in our baseline.

In our baseline, we assume that workers choose optimally between participating in the labor force and home production (labor force non-participation). Hence, we measure the probability of successfully finding a job using a regional version of the U-3 unemployment rate, which is the official unemployment rate reported by the U.S. Bureau of Labor Statistics. To check the sensitivity of our empirical results to the assumption that those not participating in the labor force are on their labor supply curves, in our second and third sensitivity exercises (reported in columns 3 and 4) we instead assume that agents who report being out of the labor force are off their labor supply curves, like the unemployed. In column 3 we re-measure the unemployment replacement rate, using an average of unemployment benefits and home production income, weighting by the shares of these populations. In column 4, we instead use our baseline unemployment replacement rate. Both results are very similar to our baseline results.

In column 5, we use a national rather than region-specific price index, imposing a single, national change in the price index, $d \log P_r = d \log P$. We measure the national price index using the national Consumer Price Index from the BLS. Welfare effects are slightly larger in this case than in our baseline. The intuition for this result is that regions receiving relatively bad shocks experience relative declines in their price indices that partially offset the negative effects in the labor market. Ignoring this differential price index effect therefore overstates the differential impact across CZs of granting China PNTR.

Finally, column 6 combines both adjustments from columns 2 and 3: we incorporate non-pecuniary costs of unemployment ($\eta = 2$) while also treating non-participants as unemployed (with adjusted unemployment replacement rates). The differential welfare impact

of granting China PNTR is slightly larger in this specification, with a coefficient of -0.314, reflecting the combined effects of both adjustments.

4.4 Alternative approaches to local welfare effects

Table 5 compares the results of estimating equation (3) measuring welfare changes using three alternative approaches (each of which implicitly uses a common national deflator). To facilitate these comparisons, we relate results to the version of our approach in which we assume a single, national change in the price index, $d \log P_r = d \log P$, displayed in column 5 of Table 4. We construct each alternative welfare measure using the LEHD.

	Alternative Measures of Welfare			
	(1)	(2)	(3)	(4)
NTRG _r	-0.266*** (0.085)	-0.332*** (0.093)	-0.169** (0.080)	-0.127** (0.061)
$\mathbb{E}[\Delta \text{Welfare}_r : \text{NTRG}_{p90}-\text{NTRG}_{p10}]$	-0.021	-0.026	-0.013	-0.0098

Each column replicates the specification of Column 5 in Table 1, but using a different measure of welfare (in place of our baseline measure ΔEV_r). In column 1, welfare is measured using the wage component of our baseline exercise in the case in which we use a national price index. In column 2, welfare is measured using the log of the change in the average nominal wage, with the average taken across all workers. In column 3, welfare is measured using the log of the change in the average wage on a sample of workers employed in the CZ in both periods. In column 4, welfare is measured using the log of the change in the average wage on a sample of workers initially employed in that CZ.

Table 5: Alternative approaches to measuring the welfare effect of granting China PNTR.

Because the first two alternative approaches focus exclusively on real wages, ignoring the impact of unemployment, column 1 of Table 5 displays the results of estimating equation (3) with the wage component of equation (2) as the dependent variable (replicating the wage component of column 5 of Table 1, Panel B, but using a national rather than a regional price index).

Column 2 presents results of estimating equation (3) measuring welfare as the change in (the logarithm of) average nominal wages relative to the price index (the change in the average real wage). This approach ignores changes in the composition of the workforce on observable and unobservable dimensions, changes in compensating differentials associated with worker reallocation, and changes in unemployment. Compared to column 1, the differential effect of granting China PNTR across CZs is larger.

Column 3 takes a step towards controlling for changing worker composition by fixing the sample of workers on which average wages are measured. Specifically, wages are constructed both in 2000 and 2007 on the sample of workers who were employed in the region in 2000 and remain employed in any location in 2007. This approach is similar to the wage component of our measure of equivalent variation in equation (2), although it ignores changes in compensating differentials associated with worker reallocation. Compared to column 1, the differential effect of granting China PNTR across CZs is smaller.

Our final alternative approach to measuring welfare takes a step towards incorporating unemployment. Specifically, wages are constructed both in 2000 and 2007 on the sample of workers who were employed in the region in 2000 whether or not they remain employed and regardless of where they live in 2007. Hence, this measure incorporates the impact on wage income of movements into unemployment. This approach is equivalent to a typical regression of changes in wage income for more- relative to less-exposed workers using worker-level panel data. Whereas this approach is reasonable for measuring wage income changes for initially employed workers, it does not measure welfare changes because it ignores changes in compensating differentials associated with worker reallocation (across sectors, regions, and between participation and non-participation) and it omits the impact of unemployment insurance on income. This result, displayed in column 4 of Table 5 should be compared with the result in column 5 of Table 4, since neither approach incorporates different changes in price indices across regions. We find a substantially smaller impact of granting China PNTR across CZs.

5 Concluding remarks

What are the welfare effects of trade shocks? This paper starts with the objective of answering this question without imposing strong functional-form restrictions and while using a research design identical to that used in the large reduced-form empirical literature interested in the effects of trade shocks on various margins of adjustment.

Motivated by this goal, we derive a sufficient statistic that measures changes in welfare consistent with labor-market adjustment along multiple margins while imposing minimal functional-form restrictions. In particular, we allow for an arbitrary distribution of non-pecuniary returns and productivity levels of each agent in each region-job pair. And we make no assumptions on region-sector production functions, firm competition, or labor-market matching technologies beyond assuming a single wage per efficiency unit and a single unemployment rate within each region-sector pair.

We apply these theoretical results to measure changes in welfare across U.S. commuting zones between 2000 and 2007 (and also between 1992 and 2000). Lastly, we identify the impact of a particular trade shock—granting China permanent normal trade relations—on changes in welfare across U.S. commuting zones. We find that CZs that are more exposed to granting China PNTR experience substantially smaller increases in welfare between 2000 and 2007. Moreover, the majority of this welfare effect is transmitted through changes in the combination of real wages per efficiency unit and changes in productivity.

We hope that our theoretical derivation of a welfare sufficient statistic and empirical approach to measuring its components has impacts beyond our particular empirical applica-

tion. Using our approach, a researcher can shed light on the impact of any well-identified shock on a measure of welfare that is consistent with a model that is substantially more general than those currently employed in quantitative work addressing similar issues.

References

ADÃO, R. (2015): "Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil," *Papers*, University of Chicago Booth School of Business.

ADÃO, R., C. ARKOLAKIS, AND F. ESPOSITO (Forthcoming): "General Equilibrium Effects in Space: Theory and Measurement," *American Economic Journal: Macroeconomics*.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.

ATKIN, D., B. FABER, AND M. GONZALEZ-NAVARRO (2018): "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy*, 126, 1–73.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–68.

— (2015): "Untangling Trade and Technology: Evidence from Local Labour Markets," *Economic Journal*, 0, 621–646.

BLANCHFLOWER, D. G. AND A. J. OSWALD (2004): "Well-being over time in Britain and the USA," *Journal of Public Economics*, 88, 1359–1386.

BORUSYAK, K., P. HULL, AND X. JARAVEL (2022): "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89, 181–213.

BORUSYAK, K. AND X. JARAVEL (2021): "The Distributional Effects of Trade: Theory and Evidence from the United States," NBER Working Papers 28957, National Bureau of Economic Research, Inc.

BURSTEIN, A., E. MORALES, AND J. VOGEL (2019): "Changes in Between-Group Inequality: Computers, Occupations, and International Trade," *American Economic Journal: Macroeconomics*, 11, 348–400.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2018): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Papers*, Yale School of Management.

CHETTY, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488.

DIX-CARNEIRO, R. (2014): "Trade Liberalization and Labor Market Dynamics," *Econometrica*, 82, 825–885.

DIX-CARNEIRO, R. AND B. KOVAK (2017): "Trade Liberalization and Regional Dynamics," *American Economic Review*, 107, 2908–2946.

DIX-CARNEIRO, R. AND B. K. KOVAK (2019): "Margins of labor market adjustment to trade," *Journal of International Economics*, 117, 125–142.

DONALD, E., M. FUKUI, AND Y. MIYAUCHI (2025): "Unpacking Aggregate Welfare in a Spatial Economy," NBER Working Papers 34075, National Bureau of Economic Research, Inc.

GALLE, S., A. RODRÍGUEZ-CLARE, AND M. YI (2023): "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *The Review of Economic Studies*, 90, 331–375.

HANDLEY, K. AND N. LIMÃO (2017): "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States," *American Economic Review*, 107, 2731–2783.

HAZELL, J., J. H. NO, E. NAKAMURA, AND J. STEINSSON (2022): "The Slope of the Phillips Curve: Evidence from U.S. States," *The Quarterly Journal of Economics*, 137, 1299–1344.

HECKMAN, J. J. AND G. SEDLACEK (1985): "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-selection in the Labor Market," *Journal of Political Economy*, 93, 1077–1125.

JONES, C. I. AND P. J. KLENOW (2016): "Beyond GDP? Welfare across Countries and Time," *American Economic Review*, 106, 2426–2457.

KIM, R. AND J. VOGEL (2021): "Trade Shocks and Labor Market Adjustment," *American Economic Review: Insights*, 3, 115–130.

KNABE, A. AND S. RATZEL (2011): "Quantifying the psychological costs of unemployment: the role of permanent income," *Applied Economics*, 43, 2751–2763.

KOVAK, B. (2013): "Regional Effects of Trade Reform: What is the Correct Measure of Liberalization?" *American Economic Review*, 103, 1960–1976.

LEE, E. (2019): "Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers," Tech. rep., Mimeo, University of Maryland.

MCCAIG, B. AND N. PAVCNIK (2018): "Export Markets and Labor Allocation in a Low-Income Country," *American Economic Review*, 108, 1899–1941.

PIERCE, J. R. AND P. K. SCHOTT (2016): "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 106, 1632–62.

PIERCE, J. R., P. K. SCHOTT, AND C. TELLO-TRILLO (2024): "To Find Relative Earnings Gains After the China Shock, Look Outside Manufacturing and Upstream," NBER Working Papers 32438, National Bureau of Economic Research, Inc.

PORTO, G. G. (2006): "Using Survey Data to Assess the Distributional Effects of Trade Policy," *Journal of International Economics*, 70, 140–160.

RODRIGUEZ-CLARE, A., M. UDATE, AND J. P. VASQUEZ (Forthcoming): "Trade with Nominal Rigidities: Understanding the Unemployment and Welfare Effects of the China Shock," *The Journal of Political Economy*.

ROY, A. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

TOPALOVA, P. (2010): "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India," *American Economic Journal: Applied Economics*, 2, 1–41.

WINKELMANN, L. AND R. WINKELMANN (1998): "Why are the Unemployed So Unhappy? Evidence from Panel Data," *Economica*, 65, 1–15.

Online Appendix

A Data Sources

In this section, we provide a detailed description of the raw data sources used in our analysis and the procedures undertaken for data cleaning and preparation. This includes information on how regions and industries in each dataset are converted to the commuting zones (CZs) and 6-digit 1997 NAICS codes (if applicable), any filtering criteria applied (e.g., age restrictions, sample selection), and methods for handling missing values or outliers.

Our analysis primarily relies on the following datasets:

- **IPUMS-USA (Census and American Community Survey):** We downloaded the following samples from the Integrated Public Use Microdata Series (IPUMS-USA):
 - **Census 1990 5% State sample**
 - **Census 2000 5% sample**
 - **2006–2008 American Community Survey (ACS) 3-year sample (3% density):** This sample combines the 1% ACS samples for 2006, 2007, and 2008 into a single dataset.

These datasets provide comprehensive demographic, labor force, and income information at the individual level. The cleaning and preparation of these data involved several steps, closely following the methodology outlined in Autor, Dorn, and Hanson (2013).

- **Variable Selection:** We retained key variables including year, state, county group, PUMA, person weight, metropolitan area codes, group quarters status, employment status, industry codes (NAICS and 1990 Census industry), weeks worked, usual hours worked, wage income, total income, and other income.
- **Weeks and Hours Worked Imputation:** For the ACS sample, weeks worked and hours worked are reported in intervals rather than as continuous values. We converted these to continuous variables by assigning the midpoint of each interval.
- **Missing Value Treatment:** We replaced sentinel values indicating missing or not applicable data with system missing values, following standard IPUMS coding conventions.

- **Income Topcoding:** Following [Autor et al. \(2013\)](#), wage income, total income, and other income were topcoded at 1.5 times their respective topcode thresholds for each survey year.
- **Hourly Wage Calculation and Truncation:** Hourly wages were calculated as wage income divided by the product of usual hours worked per week and weeks worked per year. Hourly wages were topcoded using year-specific maximum values derived from topcoded wage income assuming standard work hours (35 hours/week for 50 weeks/year). Hourly wages of zero were set to missing, and hourly wages were bottom-coded at the 1st percentile for each year to address outliers.
- **Sample Restrictions:** We applied several filtering criteria to define our analytical sample:
 - * We retained individuals aged 18 to 65 years.
 - * We excluded individuals in institutional or military group quarters.
 - * We excluded military personnel, identified using detailed employment status codes and 1990 Census industry codes for military-related industries.
 - * We dropped observations with missing education or missing/undefined industry codes, including those who last worked 5+ years ago or never worked.
- **Commuting Zone (CZ) Aggregation:** Individual-level data were aggregated to the commuting zone (CZ) level using IPUMS-provided geographic identifiers and the crosswalks available on David Dorn's website. Person weights were applied during aggregation to ensure representativeness.
- **IPUMS Time Use (American Heritage Time Use Study - AHTUS):** The American Heritage Time Use Study (AHTUS) data was downloaded from IPUMS Time Use. This dataset is used to measure time allocation for non-market work, specifically to derive the ratio of time use in non-market work by those not in the labor force relative to time use in market work by full-time employees.
 - **Sample Selection:** We selected individuals who were not employed and not unemployed, ensuring they were genuinely out of the labor force. We further restricted the sample to individuals aged 18 to 65 years.
 - **Non-Market Work Measurement:** We calculated the total time spent on non-market work by summing time allocated to undifferentiated domestic activities and civic/religious activities.
 - **Home Production Share Calculation:**

- * The numerator for the home production replacement rate is the average time use of non-participants on non-market work. Based on the 1993 AHTUS data, the average minutes spent on non-market work by those out of the labor force on a given day is 220.63 minutes.
- * The denominator is the average minutes spent on work by a full-time employee. This is approximated as 342.86 minutes (60 minutes \times 40 hours a week / 7 days a week).
- * The average share of hours worked by those in home production is approximately 0.64 ($= 220.63 / 342.86$).

– **Weighting and Aggregation:** Individual observations were weighted using record weights provided by AHTUS. These weighted averages were then aggregated to the national level by sample year to obtain the required time use ratios.

- **Confidential Longitudinal Employer-Household Dynamics (LEHD) Data:** We utilize confidential data from the Longitudinal Employer-Household Dynamics (LEHD) program. This dataset provides detailed information on worker earnings and employment spells, crucial for measuring changes in wages per efficiency unit by holding worker composition fixed.

The LEHD data are derived from quarterly earnings reports provided by employers to state unemployment insurance (UI) agencies and assembled by the U.S. Census Bureau into a national dataset. The core records contain total quarterly earnings paid by each employer to each worker, along with a small set of worker and establishment characteristics, including industry and location. Coverage extends to roughly 95% of private-sector employment and to state and local government workers, but excludes federal employees, members of the armed forces, and the self-employed.

For our main analysis, we extract all available LEHD earnings records for two reference quarters: 1999Q2 and 2007Q2. We restrict the sample in each quarter to individuals aged 18-57 in 1999Q2, ensuring workers are at most 65 years old in 2007Q2. We also restrict the sample to those employed in a single job (dropping those with multiple jobs within the quarter) to ensure a unique industry assignment. In all cases, industry and location are defined from the worker's unique employer in the quarter. For our main measures of wages per efficiency unit, we retain only workers who appear in both periods and who are employed in the same commuting zone and industry in both quarters. Observations with missing industry codes and/or commuting zone identifiers are dropped. The resulting sample represents full-quarter employment spells for a consistent set of workers over time, excluding movers across industries and/or

commuting zones, part-time workers with multiple concurrent jobs, and other atypical employment arrangements. All earnings measures correspond to total quarterly earnings for these workers. Due to state participation patterns in LEHD, our measures cover commuting zones in 34 states. Industry and geographic codes are harmonized to 6-digit 1997 NAICS and commuting zones using Census-provided crosswalks.

For the placebo exercises, we apply the same restrictions but instead focus on the period 1992Q2 to 1999Q2, which results in fewer observations due to declining state coverage in earlier years. When constructing alternative measures of wages per efficiency unit, we relax only the requirements that workers be employed in both periods and in the same commuting zone and industry in both quarters. For example, in one alternative we require only that workers be employed in 1992Q2 or 1999Q2; in another alternative, we require that workers be employed in both quarters but drop the requirement that they be in the same commuting zone and industry.

- **Confidential Longitudinal Business Database (LBD) Data:** The Longitudinal Business Database (LBD) provides comprehensive information on firm and establishment characteristics. This data is used for measuring total number of employed agents and the total wagebill within each sector (I_{rs}), as well as to construct commuting zone-level trade shocks.
 - **Data Sources and Years:** We utilize LBD data for the years 1980, 1990, and 2000. The raw LBD data includes county-level information for each year and records industries using 2002 6-digit NAICS codes. Establishments with zero employment are excluded from the sample.
 - **Geographic Concordance:** We convert the county-level information in the LBD to commuting zones (CZs) by applying crosswalks available on David Dorn's website.
 - **Industry Concordance:** To align with our defined sector classification, we crosswalk the 2002 6-digit NAICS codes to 1997 6-digit NAICS codes using a Census crosswalk file.²⁴
- **Commuting Zone Controls:** Commuting zone-level control variables used in our regression analysis—including the manufacturing share, college share, foreign-born share, female employment share, routine occupation share, and average offshorability—are taken directly from [Autor et al. \(2013\)](#).

²⁴This crosswalk involves handling instances where the mapping is not one-to-one. Specifically, among 1354 unique matches between 1997 NAICS and 2002 NAICS codes, approximately 300 are not one-to-one matches (i.e., a single 2002 NAICS code maps to multiple 1997 NAICS codes, or vice-versa). In such cases, we equally assign the industries.

- **Price Index Data (BLS CPI and State-Level CPI):** We use two primary sources for price index data.
 - **National CPI (CPI-U-RS):** The Consumer Price Index research series for all Urban Consumers using current methods (CPI-U-RS) was downloaded from the Bureau of Labor Statistics (BLS) website.
 - * **Cleaning and Processing:** The series was rebased to January 1990 to minimize approximation error when compared with state-level indices. Annual average CPI values were then calculated for the years 1990, 2000, and 2007. Log differences were computed for the periods 1990–2000 and 2000–2007. The 1990–2000 difference was scaled by a factor of 7/10 to make it comparable to a 7-year equivalent change, consistent with our benchmark period.
 - **State-Level Price Index:** We utilize state-level price indices available from [Hazell et al. \(2022\)](#), which provides regional inflation data. This data was sourced from their supplementary materials.
 - * **Cleaning and Processing:** The raw quarterly state-level CPI data was imported and quarterly inflation was calculated. Similar to the national CPI, the state-level series were rebased to 1990Q1. Annual average state-level CPI values were then computed for 1990, 2000, and 2007. Log differences were calculated for the 1990–2000 and 2000–2007 periods, with the 1990–2000 difference also scaled to a 7-year equivalent.
 - **Missing Data Handling and Merging:** State-level price indices were merged with commuting zone (CZ) data. To ensure comprehensive coverage across all CZs, for states where the state-level CPI data was unavailable (17 states, including the District of Columbia), the national CPI-U-RS difference was used as a substitute. This imputation strategy ensures that all CZs have a corresponding price index measure for both common and regional inflation analyses. Note that the simple average of state-level inflation across all available CZs is similar to the national inflation for both periods: 0.163 for state-level inflation versus 0.172 for national inflation in 1990–2000, and 0.182 for state-level inflation versus 0.186 for national inflation in 2000–2007.
- **Department of Labor (Unemployment Replacement Rate):** The state-level unemployment replacement rate reports were directly downloaded from the Department of Labor website: https://oui.dol.gov/unemploy/ui_replacement_rates.asp. This data is used to construct the unemployment income component.

- **Data Selection:** Specifically, we utilize "Replacement Ratio 2," which is defined as the Ratio of: Weighted Average Weekly Benefit Amount (WBA) / Weighted Average (Normal Hourly Wage x 40 Hrs.). The WBA represents the claimant's weekly benefit amount.
- **Coverage and Use:** The UI replacement rate covers a four-quarter moving average ending with a selected quarter-ending date. The data is available from 1997, and the 1997 unemployment replacement rate is used for both our main and placebo analyses.
- **Source Details:** This report uses the Unemployment Insurance Benefit Accuracy Measurement (BAM) data, stored in the UI database.
- **NTR Gap Data:** Sector-level NTR gaps are taken from [Pierce and Schott \(2016\)](#). The NTR gap is defined as the difference between the non-NTR tariff rate and the NTR tariff rate for each industry, measured as of 1990.

B Construction of Welfare Measures (EV_r)

This section provides detailed documentation on the construction of our welfare measure, EV_r , for both the baseline analysis period (2000–2007) and the pre-trend analysis period (1992–2000). We first describe the welfare equation and how each component is measured, and then discuss the specific years and scaling adjustments used for each period.

B.1 Welfare Equation

Changes in welfare for agents initially living in region r can be expressed as in equation [\(2\)](#), which we repeat here for completeness:

$$EV_r = \underbrace{\sum_{j \in \mathcal{J}} \frac{I_{rj}}{I_r} d \log \left(\frac{w_{rj}}{P_r} \right)}_{\text{real wage}} + \underbrace{(1 + \eta_r) \frac{I_r^e}{I_r} \left[1 - \frac{\bar{I}_r^u}{\bar{I}_r^e} \right] d \log E_r}_{\text{unemployment}}$$

where \mathcal{J} denotes the set of all jobs (each sector of employment, labor force non-participation, and unemployment). The first term represents the real wage component of welfare, and the second term represents the unemployment component. For the real wage component, we measure $CZ \times 6$ -digit NAICS sector income (I_{rs}) from the Longitudinal Business Database (LBD) and wage changes ($d \log w_{rs}$) from the Longitudinal Employer-Household Dynamics

(LEHD). For the unemployment component and for constructing CZ-level aggregates, we use IPUMS Census data.

B.2 Measurement of Welfare Components

We now describe how each component of equation (2) is measured in detail. The specific years and data sources used for each analysis period are described in Section B.3.

B.2.1 Income Shares

The real wage component of equation (2) requires measuring income shares I_{rj}/I_r for each job j . We describe the construction of each income component below.

Total Income (I_r). Total income in region r is the sum of employment, unemployment, and non-participation income:

$$I_r = I_r^e + I_r^u + I_{r0} \quad (\text{App.1})$$

Employment Income (I_r^e). Total employment income in region r is measured as the sum of total wagebill across all sectors:

$$I_r^e = \sum_{s=1}^S I_{rs} \quad (\text{App.2})$$

where I_{rs} is the total wage bill within CZ-sector pair (r, s) , measured from the LBD (for the real wage component) or IPUMS Census data (for the unemployment component).

Non-Participation Income (I_{r0}). Total income of non-participants in region r is constructed as:

$$I_{r0} = \bar{I}_{r0} \times N_{r0} \quad (\text{App.3})$$

where \bar{I}_{r0} is the average income of non-participants and N_{r0} is the number of non-participants. This can be rewritten as:

$$I_{r0} = \frac{\bar{I}_{r0}}{I_r^e} \times \frac{N_{r0}}{N_r^e} \times I_r^e \quad (\text{App.4})$$

where \bar{I}_{r0}/I_r^e is the home production replacement rate and N_{r0}/N_r^e is the ratio of non-participants to employed workers. We construct each component as follows. First, the home production replacement rate is constructed as the product of two terms:

$$\frac{\bar{I}_{r0}}{I_r^e} = \left(\frac{\text{time use, non-market work by NILF}}{\text{time use, work by employed}} \right) \times \left(\frac{\text{avg wage, private household sector}}{\text{avg wage, all sectors}} \right) \quad (\text{App.5})$$

where NILF refers to those not in the labor force. The first term captures the relative time devoted to home production, while the second term values this time using average wages in the private household sector (as a proxy for the market value of home production activities) relative to average wages across all sectors.

- **Time use ratio:** We measure this ratio using the American Heritage Time Use Study (AHTUS) data. The numerator is the average time spent on non-market work by those not in the labor force; based on the 1993 AHTUS data, this equals 220.63 minutes per day. The denominator is the average time spent on work by a full-time employee, approximately 342.86 minutes per day (40 hours per week divided by 7 days). The resulting ratio is approximately 0.64.
- **Wage ratio:** We measure this ratio using IPUMS Census data, where the private household sector is defined using the 1990 Census industry code for “Private households.” The ratio is computed as the average wage per worker in the private household sector divided by the average wage per worker across all sectors, measured at the state level. We measure this ratio at the state level rather than the CZ level because some CZ-level estimates yield implausible values due to small sample sizes; the state-level ratio is then applied uniformly to all CZs within each state.

Second, the ratio of non-participants to employed workers can be expressed as:

$$\frac{N_{r0}}{N_r^e} = \frac{L_r - M_r}{E_r M_r} \quad (\text{App.6})$$

where L_r is the total population aged 18–65, M_r is the labor force (employed plus unemployed), and $E_r = N_r^e / M_r$ is the employment rate. Note that $L_r - M_r$ equals the number of non-participants and $E_r M_r = N_r^e$.

- **Population ratio:** The ratio $(L_r - M_r) / (E_r M_r)$ is measured by commuting zone from IPUMS Census data.

Unemployment Income (I_r^u). Total income of unemployed workers in region r is constructed as:

$$I_r^u = \frac{\bar{I}_r^u}{\bar{I}_r^e} \times \frac{N_r^u}{N_r^e} \times I_r^e \quad (\text{App.7})$$

where $\bar{I}_r^u / \bar{I}_r^e$ is the unemployment replacement rate and $N_r^u / N_r^e = (1 - E_r) / E_r$ is the ratio of unemployed to employed workers.

- **Unemployment replacement rate ($\bar{I}_r^u / \bar{I}_r^e$):** State-specific unemployment replacement rate from the Department of Labor. We use “Replacement Ratio 2,” defined as the ratio

of weighted average weekly benefit amount to weighted average (normal hourly wage $\times 40$ hours).

- **Population ratio:** The ratio $(1 - E_r) / E_r$ is measured by commuting zone from IPUMS Census data.

B.2.2 Changes in Productivity-Adjusted Wages per Efficiency Unit

Sectoral Wages ($d \log w_{rs}$ for $s \geq 1$). For each sector $s \geq 1$, we measure $d \log w_{rs}$ using LEHD panel data. We construct a sample of workers who are employed in commuting zone r and sector s in both the initial and terminal quarters. On this sample, we compute:

$$d \log w_{rs} = \frac{1}{N_{rs}} \sum_{i=1}^{N_{rs}} (\log w_{i,t'} - \log w_{i,t}) \quad (\text{App.8})$$

where t and t' denote the initial and terminal quarters, the summation is taken over all N_{rs} workers who are employed in commuting zone r and sector s in both periods, and $w_{i,t}$ is total quarterly earnings for worker i in period t . This approach holds worker composition fixed, ensuring that measured wage changes reflect changes in productivity-adjusted wages per efficiency unit rather than changes in worker selection.

Unemployment and Non-Participation ($j \in \{u, 0\}$). As described in the main text, we choose changes in nominal returns to unemployment and non-participation such that real wage changes are zero:

$$d \log \left(\frac{w_{rj}}{P_r} \right) = 0 \quad \text{for } j \in \{u, 0\} \quad (\text{App.9})$$

which implies $d \log w_{rj} = d \log P_r$ for these jobs.

Handling Missing Wage Data. In constructing the real wage component of equation (2), there are some CZ-sector pairs for which changes in productivity-adjusted wages per efficiency unit are not defined (due to insufficient sample size, i.e., fewer than 20 workers employed in the same CZ-sector pair in both the initial and terminal quarters), yet the weight I_{rs} / I_r in the weighted average is not zero. For these CZ-sector pairs, we set the weights to zero and renormalize all other weights I_{rs} / I_r such that the sum of these weights across all sectors equals I_r^e / I_r .

B.2.3 Regional Price Index

We measure region-specific inflation, $d \log P_r$, using state-level CPI data from [Hazell et al. \(2022\)](#). For states without state-level CPI data (17 states including the District of Columbia),

we use the national CPI-U-RS as a substitute. As a robustness check, we also apply the national CPI-U-RS for all CZs.

B.2.4 Unemployment Component

The unemployment component of equation (2) requires the following elements:

Change in Employment Rate ($d \log E_r$). The employment rate E_r is measured as the ratio of employed population to labor force (one minus the unemployment rate) from IPUMS Census data.

Non-Pecuniary Cost of Unemployment (η_r). The parameter η_r represents the non-pecuniary welfare loss of unemployment relative to the pecuniary welfare loss. In our baseline analysis, we set $\eta_r = 0$. In robustness exercises, we set $\eta_r = 2$, which implies multiplying the unemployment component of welfare by a factor of 3 (i.e., $1 + \eta_r = 3$).

B.3 Analysis Periods

B.3.1 Baseline Period (2000–2007)

For our main analysis, we construct EV_r over the baseline period 2000–2007. Note that the specific years used for each component vary slightly depending on data availability and timing considerations, as detailed below.

Income Shares. We measure income shares using year 2000 data from the following sources:

- **Employment income (I_{rs}):** Longitudinal Business Database (LBD), year 2000 (for the real wage component); IPUMS Census 2000 (for the unemployment component).
- **Wage ratio for home production:** IPUMS Census 2000, measured at the state level.
- **Time use ratio for home production:** American Heritage Time Use Study (AHTUS) 1993. This ratio (approximately 0.64) is time-invariant in our analysis.
- **Population counts** (for non-participation and unemployment): IPUMS Census 2000, measured at the CZ level.
- **Unemployment replacement rate:** Department of Labor, 1997 (the earliest year available).

We use 2000 because the LBD provides the most reliable establishment-level data in decennial Census years due to enhanced coverage and accuracy during these benchmark periods. The IPUMS data similarly benefits from the larger sample sizes available in decennial Census years.

Changes in Productivity-Adjusted Wages per Efficiency Unit. We measure $d \log w_{rs}$ using LEHD data from 1999Q2 to 2007Q2. We use 1999Q2 as the initial quarter to ensure that any anticipation effects associated with the Congressional vote on PNTR in October 2000 and China's accession to the WTO in December 2001 are captured in the baseline period rather than the pre-trend period.

Because the LEHD wage changes span 8 years (from 1999Q2 to 2007Q2), we scale these changes by 7/8 to construct a 7-year equivalent measure that is comparable to the nominal 7-year baseline period (2000–2007).

Employment Rate. We measure $d \log E_r$ using the 2000 Decennial Census and the combined 2006–2008 ACS samples. The use of multiple ACS years increases sample sizes at the commuting zone level.

Price Index. We measure $d \log P_r$ using state-level CPI data from [Hazell et al. \(2022\)](#) for 2000–2007.

B.3.2 Pre-Trend Period (1992–2000)

For our falsification exercise, we construct EV_r over the pre-trend period 1992–2000 to verify that our baseline results capture the effects of granting China PNTR rather than the continuation of pre-existing trends. As with the baseline period, the specific years used for each component vary depending on data availability.

Income Shares. We measure income shares using year 1990 data from the following sources:

- **Employment income (I_{rs}):** Longitudinal Business Database (LBD), year 1990 (for the real wage component); IPUMS Census 1990 (for the unemployment component).
- **Wage ratio for home production:** IPUMS Census 1990, measured at the state level.
- **Time use ratio for home production:** American Heritage Time Use Study (AHTUS) 1993 (same as baseline).
- **Population counts** (for non-participation and unemployment): IPUMS Census 1990, measured at the CZ level.

- **Unemployment replacement rate:** Department of Labor, 1997 (same as baseline; the earliest year available).

We use 1990 because this is the only decennial Census year available for the pre-trend period, and the LBD and IPUMS data provide reliable coverage in this benchmark year comparable to that of the 2000 data used in our baseline analysis.

Changes in Productivity-Adjusted Wages per Efficiency Unit. We measure $d \log w_{rs}$ using LEHD data from 1992Q2 to 1999Q2. We use 1999Q2 as the terminal quarter (rather than 2000Q2) to ensure that any anticipation effects of China's WTO accession are captured by the baseline period and not the pre-trend period. The 7-year span from 1992Q2 to 1999Q2 requires no scaling adjustment since it already matches the 7-year baseline equivalent.

State coverage in the LEHD is substantially more limited in this earlier period due to the gradual expansion of state participation in the LEHD program over time. Whereas the baseline period covers 34 states and approximately 600 commuting zones, the pre-trend period covers only 6 states and approximately 80 commuting zones.

Employment Rate. We measure $d \log E_r$ using the 1990 and 2000 Decennial Censuses. Because this spans 10 years rather than 7, we scale the change in $\log E_r$ by 7/10 to construct a 7-year equivalent measure comparable to our baseline period.

Price Index. We measure $d \log P_r$ using state-level CPI data from [Hazell et al. \(2022\)](#) for 1990–2000. As with the employment rate, the 10-year price change is scaled by 7/10 to obtain a 7-year equivalent.