From Telegraph Messengers to Web Developers: Measuring Wages as Occupations Change*

Gabriel Ehrlich
University of Michigan

Kelsey O'Flaherty
University of Michigan

March 20, 2025

Preliminary and under frequent revision:
feedback most welcome!

Abstract

We evaluate the contributions of two features of occupational change—the entry and exit of occupations over time and shocks to relative supply and demand—to aggregate wage growth and the distribution of wages across occupations from 1940 to 2020. Our analysis builds on a two-sided CES model of the labor market for which we derive exact employer- and worker-side wage indices accounting for these forces, which increased workers' annualized real wage growth by 0.29 percentage points and reduced firms' real labor cost increases by 0.20 percentage points annually on average over this period. Inverting the system delivers relative supply and demand shifts for each occupation, allowing us to evaluate how each force contributed to wage polarization from 1970 to 2020. Demand shocks accounted for nearly the entire increase in the interquartile range of occupational wages over this period, with one-fifth of the increase coming from occupations appearing after 1970. Supply shocks, while large, had little impact on occupational wage polarization in the past fifty years.

*Ehrlich: gehrlich@umich.edu; O'Flaherty: oflahekj@umich.edu. We thank Ying Fan, Mel Stephens, Stephen Terry, and seminar participants at the Federal Reserve Bank of Cleveland and the University of Michigan for helpful comments.

1 Introduction

We study a model of occupational change that features three forces: entry and exit of occupations in the labor market, shocks to relative demand across occupations, and shocks to relative supply across occupations. The model rationalizes every shift in an occupation's relative wages and labor quantities in terms of a unique combination of supply and demand shocks. Doing so allows us to study each force's contribution to changes in the distribution of wages and calculate wage indices that account for changes in the set of available occupations in addition to relative demand and supply shocks. We study those questions in the context of the U.S. labor market from 1940 to 2020.

We document that the past eighty years have seen substantial entry of new occupations into the labor market. We estimate that approximately three-quarters of the occupations that existed in 2020 did not exist in 1940, and that these new occupations accounted for two-thirds of the wage bill in 2020. Our empirical approach allows us to identify exiting occupations in addition to entering occupations, but we find that occupational entry is more economically significant than exit, with entering occupations' share of the wage bill exceeding exiting occupations' by 80% on average.

We estimate that occupational entry increased workers' average real wage growth by 0.11 percentage points annually from 1940–2020 while reducing the corresponding firm-side measure by an average of 0.04 percentage points per year. Those effects are meaningful relative to the average annual real wage growth of 1.03% per year we measure using a traditional CES wage index of matched occupations (Sato, 1976; Vartia, 1976). Growth in occupational variety had its strongest positive effects from 1940–1970, when it added an average of 0.23 percentage points to workers' real wage growth annually while reducing firms' real wage costs by 0.09 percentage points per year. After 1970, occupational entry also contributed significantly to rising dispersion in the wage distribution: the entry of new occupations accounted for one-fifth of the increase in the interquartile range of occupational wages from 1970 to 2020.

Consistent with a large literature, we find that labor demand has polarized substantially over the past fifty years. We estimate that the interquartile range of relative labor demand across occupations more than tripled from 1970 to 2020, and relative labor demand shocks accounted for approximately four-fifths of the increase in the interquartile range of occupational wages in that time. We show that this increasing dispersion carries the the direct if surprising implication that firms' effective wage costs grew less quickly than a traditional wage index would suggest, by 0.16 percentage points annually on average. The effect of our relative shock adjustment is analogous to the well-known result that consumers benefit from more dispersed prices all else equal, as dispersion offers more opportunity for substitution.

We estimate that relative labor supply shocks across occupations are of roughly the same magnitude as labor demand shocks. From 1970 to 2020, workers became more willing to supply labor in occupations with a low share of routine tasks and a high share of social tasks, at the same time that the relative demand for such occupations was increasing (Deming, 2017). Accounting for these shifts in labor supply raises workers' effective real wages by an average of 0.18 percentage points annually from 1940–2020. Nonetheless, we find that labor supply shocks played essentially no role in the increase in the interquartile range of relative occupational wages over the past fifty years.

Our two-sided CES labor market framework delivers wage measures which answer more directly welfare-relevant questions than do simple average wages. The worker-side index answers the question, *"how much more do workers earn today for incurring a unit of disutility from supplying labor than they earned in the past?"* while the firm-side index measures *"how much more must firms pay today to procure a unit of effective labor input than they paid in the past?"* Our indices treat the entry of a new occupation as an economically consequential expansion of the choice sets for both firms and workers and the exit of an old occupation as a diminution of those choice sets. These indices also account for relative shifts in supply and demand across occupations analogously to shifts in relative wages.[1]

[1]Yet we necessarily normalize shifts in labor supply and labor demand to exclude changes that affect all occupations equally. Our indices would behave identically to a standard wage index in response to a sudden

To identify the entry of new occupations and the exit of old occupations, we construct an algorithm based on publicly available data. It suggests that occupational turnover has been immense over the past eighty years. We quantify the welfare implications of changes in the set of available occupations over time by adapting the approach of Feenstra (1994), which accounts for product turnover in CES price indices.

Our wage indices account for changing relative supplies and demands across occupations by adapting the approach of Redding and Weinstein (2020). An example helps to clarify the intuition for our relative shock adjustments. The industry occupational death rate in transportation, communications, and public utilities declined from over 20 per 100,000 workers in 1980 to roughly 10 in 1995, and the frequency of motor-vehicle related deaths dropped significantly (Centers for Disease Control and Prevention, 1999). These improving safety conditions reduced the disutility incurred from working as a bus driver or truck driver relative to other occupations. Our worker-side wage index accounts for such shifts (subject to our normalization that the average change across occupations is zero). Our firm-side wage index accounts for relative shifts in labor demand driven by analogous shifts in technology and market conditions.

We exploit the invertibility of our two-sided CES model of the labor market to recover relative shifts in occupational labor supply and demand and perform validation exercises that show these estimated shocks are not merely noise. Rather, they reflect fundamental economic changes that have important implications for the wage distribution and the measurement of aggregate wages.

Our paper relates to three primary literatures, the first of which develops exact CES price indices. As discussed above, we adapt the seminal contributions of Sato (1976), Vartia (1976), Feenstra (1994), and Redding and Weinstein (2020) to the context of the labor market. Our firm-side wage indices correspond to the CES consumer price indices already in the literature, while we develop analogous wage indices for the worker side of the market.

bout of indolence or industry on the part of all workers or improvement in productivity across all jobs.

Redding and Weinstein (2018) propose new estimators for the elasticity of substitution across products, which we apply to estimate the elasticity of substitution of labor demand across occupations. We propose an instrumental variables strategy that relies on long lags of our estimated demand shocks to estimate the worker-side elasticity of substitution.

The second related literature examines the changing types of work in the labor market. Lin (2011) examines geographical patterns in the growth of "new work" and provides a methodology for identifying new occupations using Census occupation classifications. Acemoglu and Restrepo (2018) use his classification to study the relationship of new work to automation and new task creation. Autor, Chin, Salomons, and Seegmiller (2022) link new work to labor augmenting and automating technological change by applying natural language processing techniques to patent text. Relative to these studies, we focus on estimating new work's effects on aggregate wages and the wage distribution. We also develop an algorithm to identify both entering and exiting occupations using publicly available data.

The third related literature studies how changes in the supply and demand for different worker types affects the distribution of wages. Key contributions include Goldin and Margo (1992), Katz and Murphy (1992), Bound and Johnson (1992), and DiNardo, Fortin, and Lemieux (1996). In the interest of tractability, this literature commonly makes two assumptions that simplify the supply side of the labor market. First, the stock of workers with varying skills or attributes (e.g., college education) is typically viewed as fixed within a period, i.e., perfectly inelastic with respect to wages. Second, workers of a given type reallocate their labor flexibly across sectors so that their supply responses to sectoral demand shocks are perfectly elastic. This modeling convention limits the ability to analyze how workers reallocate across sectors or occupations that are likely to be imperfect substitutes.

We introduce a model with upward-sloping labor supply across occupations, which allows us to distinguish adjustments along a given supply curve in response to changing wages from other shifts in relative supply. Our results imply that this distinction is economically consequential for the study of wage polarization. Imperfect substituability across occupations

has mitigated the increase in wage polarization that would have arisen from labor demand shocks if workers substituted perfectly elastically across occupations. We estimate that labor supply shocks are equally large as demand shocks across occupations and take place in occupations workers may find more appealing, in that they are less routine and involve more social interactions. These adjustments imply meaningful reductions in the disutility a representative worker incurs to earn a given wage over the period we study.

The remainder of this paper is organized as follows. Section 2 discusses currently published measures of aggregate wages as well as their conceptual and empirical limitations. Section 3 presents a two-sided CES model of the labor market and derives measures of relative supply and demand shocks as well as our model-implied wage indices. Section 4 describes the data and algorithm for identifying entering and exiting occupations. Section 5 describes our procedure for estimating elasticities of substitution across occupations, and Section 6 presents our estimated CES wage indices. Section 7 summarizes our estimated occupational supply and demand shocks, and Section 8 describes our decompositions of changes in the wage distribution. A final Section 9 concludes.

2 Currently Published Wage Indices

Wages play a central role in the coordination of economic activity and are themselves a key economic outcome, but wage measurement lags price measurement both empirically and conceptually. This section summarizes some prominent measures of aggregate wages for the U.S. economy, which suffer composition and substitution biases that official inflation statistics have worked to address. As we discuss, theoretically consistent wage estimates would also expand on published wage indices by accounting for the entry and exit of new and old types of work or for relative shifts in the supply and demand for different occupations.

Perhaps the best-known wage measure comes from the Bureau of Labor Statistics' (BLS') Current Employment Statistics survey, which reports the simple average hourly earnings

(AHEs) of production and nonsupervisory employees in private industries beginning in 1964 and of all employees in private industries beginning in 2006. Average earnings measures conflate changes in wages for a fixed labor force and changes in the mix of work being performed. This limitation was on acute display during the COVID-19 pandemic, as job losses in lower-paid sectors such as leisure and hospitality led to a sharp rise in average wages (e.g., Crust, Daly, and Hobijn, 2020). Less extreme versions of this problem also affected AHEs during previous recessions. For these reasons, simple averages are generally discouraged in the price index literature (e.g., International Monetary Fund, 2009).

Another prominent measure of employee compensation growth is the BLS' Employment Cost Index (ECI). The ECI is a Laspeyres-style index that holds the composition of work fixed over long horizons. By removing composition changes, it isolates average wage changes for the mix of workers in its base period. It is well-known from the price index literature (e.g., Boskin, Dulberger, Gordon, Griliches, and Jorgenson, 1996) that Laspeyres-style indices suffer from substitution bias, because they do not reflect consumers' ability to re-optimize their consumption bundles in response to price changes. The ECI is subject to similar limitations as it updates its employment weights at a very low frequency. For instance, the ECI reported average wage gains accruing to the 2012 workforce through through September 2022, despite the substantial changes in the mix of work over the intervening decade (Bureau of Labor Statistics, 2023). Moreover, as a traditional matched index, the ECI does not account for entering and exiting occupations or shocks to relative labor supply and demand.

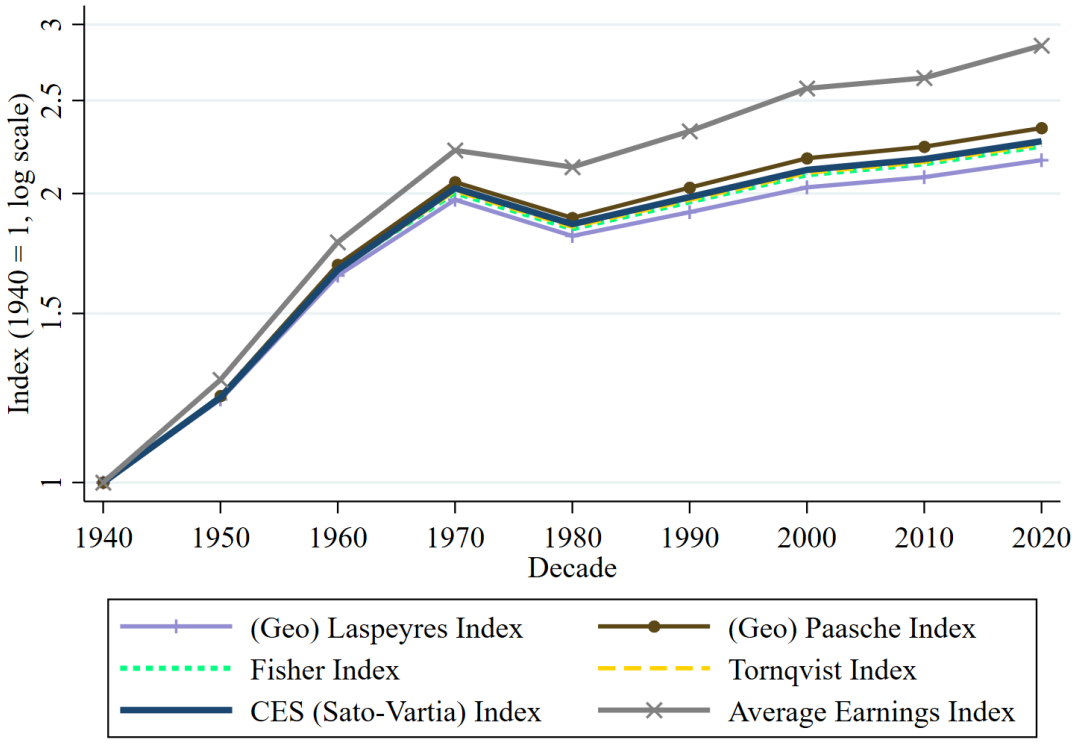
An alternative measure of wage growth is the Wage Growth Tracker (WGT) published by the Federal Reserve Bank of Atlanta (Federal Reserve Bank of Atlanta Center for Human Capital Studies, 2024).[2] The WGT measures the median twelve-month growth in the hourly wages of individuals in the Current Population Survey (CPS). Therefore, the WGT solves some of the problems faced by the AHE series and the ECI. By tracking individuals longitudinally, the WGT reduces the scope for the composition biases in the average earnings

[2]The WGT is based on a methodology developed in Daly, Hobijn, and Lucking (2012a) and Daly, Hobijn, and Wiles (2012b).

measures. Through its use of the CPS sample, the WGT implicitly updates its employment weights and accounts for the emergence of new occupations. Still, the WGT suffers from a mix of the limitations facing the other indices. For instance, Daly, Hobijn, and Wiles (2012b) show that composition effects are important drivers of median wage growth in the CPS.

Figure 1 presents several real wage indices calculated over the period 1940–2020. The indices are calculated from microdata from detailed Decennial Censuses from 1940 to 2000 and the American Community Survey for 2010 and 2020. All of the indices except the average earnings index treat three-digit occupation codes as their fundamental units of observation and are aggregated from occupational average wages and worker counts. Here and throughout the paper, we use the Consumer Price Index to deflate nominal wages. We describe our treatment of the data and our methodology for calculating price indices in detail in Section 4 below and the index formulas in Appendix A.

Figure 1: Traditional Real Wage Indices, 1940–2020



All indices deflated using Consumer Price Index (CPI). Employment and earnings are from decennial Census in 1940-2000 or 5-year ACS samples centered on 2009 and 2019.

The light purple line shows the chained geometric Laspeyres index, which cumulates average wage changes between decades using each occupation's share of the wage bill in the initial period as importance weights. It is conceptually similar to the Employment Cost Index in that it does not account for substitution behavior. The brown line is the chained geometric Paasche index, which uses wage bill shares in the end period and therefore overstates opportunities for substitution. The gold line shows the Tornqvist index, which is the geometric average of these Laspeyres and Paasche indices. The Fisher index, which averages the arithmetic Laspeyres and Paasche indices, is displayed in bright green and is nearly identical to the Tornqvist index. The similarity of the Tornqvist and Fisher indices is unsurprising because both are superlative indices, which means that they are exact for a flexible functional form (Diewert, 1978).

The dark blue line shows the Sato-Vartia index, which is exact for a CES production or utility function in the absence of occupational turnover and relative supply and demand shocks. Reinsdorf and Dorfman (1999) and Barnett and Choi (2008) show that the Sato-Vartia index is also superlative, making its similarity to the Fisher and Tornqvist indices unsurprising. The superlativeness of the Sato-Vartia index suggests that the CES form is likely to provide a reasonable approximation to the economy's underlying preference and technology structures for the purpose of constructing aggregate wage indices.[3]

The outlier in Figure 1 is the average earnings index shown in gray. It grew 26% (0.29% annually) more quickly from 1940 to 2020 than the superlative indices. The average earnings index's faster growth reflects many factors, prominently including the ongoing shift of the workforce into higher-paid occupations over the past 80 years. Many of those occupations are new, which highlights an important limitation of the traditional "matched-occupation" indices shown in Figure 1. These indices, due to their reliance on wage changes for each occupation, are unable to incorporate occupations when they first appear or when they exit.

[3]Furthermore, we note that the representative firm and household assumptions underlying the Sato-Vartia index need not be interpreted literally for the CES model to be relevant. For instance, Anderson, De Palma, and Thisse (1992, §3.7) show that CES preferences represent the expected utility of agents facing discrete choice problems with logit preferences.

Although these forces do affect the average earnings index, it does not capture them in an economically principled manner. The CES exact wage indices we propose below account for these forces directly in the context of well-posed firm and household optimization problems.

The faster growth of the Paasche index relative to the Laspeyres highlights the tension of using a single measure of wages as an indicator of both firm costs and worker earnings. For workers, the faster growth of the Paasche index reflects expected substitution patterns—workers' ability to substitute into occupations with higher wages increases their welfare, a possibility that the Laspeyres misses. Firms represent the opposite (demand) side of the labor market, so one might expect their substitution behavior to lead the Paasche index to be below the Laspeyres, in parallel with standard consumer theory. The faster growth of the Paasche index may therefore appear to pose a puzzle. This apparent puzzle is resolved if firms face shifts in labor demand alongside changes in wages that move them along given labor demand curves. Our proposed wage indices resolve these tensions by accounting for shifts in relative labor supply and demand curves across occupations.

3 Theoretical Setting and Wage Measurement

In this section, we propose a simple model of CES supply and demand over a variety of labor types, which we operationalize here as occupations but which in principle could correspond to many other categorizations. The set of occupations in the labor market shifts over time, as do their relative supply and demand curves. We propose empirically implementable exact wage indices for the firm and the household sides of the model. We then describe equilibrium in the model's labor market and show how to recover the relative supply and demand shifters for different occupations and calculate counterfactual relative wage distributions based on alternative constellations of those shifters. We defer a discussion of our empirical implementation to Section 4.

3.1 A Model of Labor Supply and Demand

We consider a perfectly competitive Walrasian general equilibrium model with a representative household and a representative firm. The household makes a labor supply decision based on a tradeoff between the utility it receives from consumption and the disutility it incurs from supplying labor. The firm produces output from labor and decides how much to employ in production. The labor market features a set of distinct occupations. The household's disutility from supplying labor and the firm's effective labor input are both CES aggregates over these occupations.

The representative firm produces output Y_t, which serves as the numeraire of the economy. The firm produces output according to the function

$$Y_t = F(A_t L_t^f), \tag{1}$$

where labor input L_t^f is a CES composite over the set of occupations $k \in \Omega_t$ available in period t, defined as

$$L_t^f \equiv \left[\sum_{k \in \Omega_t} (\alpha_{kt} \ell_{kt})^{\frac{\sigma_f - 1}{\sigma_f}} \right]^{\frac{\sigma_f}{\sigma_f - 1}}. \tag{2}$$

The ℓ_{kt} terms represent the quantity of labor input for each type of labor, A_t represents the level of labor-augmenting technological progress, and the α_{kt} parameters shift each labor type's relative efficiency.[4] σ_f represents the elasticity of substitution over labor types. The function $F(\cdot)$ obeys standard regularity conditions that we note in Appendix B.1.

The representative household in the economy experiences utility from consumption and disutility from labor according to the additively separable function

$$U\left(C_t, L_t^h\right) = G\left(C_t\right) - H\left(\Gamma_t L_t^h\right). \tag{3}$$

C_t is the household's consumption of the output good. Γ_t is an aggregate parameter that shifts the household's labor supply curve over time, such as growth in the labor force, and

[4] We will impose a normalization on the α_{kt} parameters to ensure that their interpretation is as relative efficiencies distinct from aggregate labor efficiency A_t.

L_t^h represents the household's CES composite quantity of labor supplied, defined as

$$L_t^h \equiv \left[\sum_{k\in\Omega_t} (\gamma_{kt}\ell_{kt})^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}. \tag{4}$$

The ℓ_{kt} terms represent the quantity of labor supplied for each occupation, and the γ_{kt} parameters shift the household's disutility from supplying each type of labor. σ_h represents the household's elasticity of substitution across occupations.[5]

A standard result (e.g. Pollak, 1989, p. 23) is that there exist CES composite wage rates W_t^f and W_t^h, given by

$$W_t^f = \left[\sum_{k\in\Omega_t} \left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}} \quad \text{and} \quad W_t^h = \left[\sum_{k\in\Omega_t} \left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}\right]^{\frac{1}{1+\sigma_h}} \tag{5}$$

such that the product of the firm's composite labor input L_t^f with the firm's composite wage rate W_t^f equals the firm's total wage bill, and the product of the household's composite labor supplied L_t^h with the household's composite wage rate W_t^h equals the household's total labor income.[6] The firm's wage bill will equal the household's wage income, which we denote Z_t:

$$Z_t \equiv \sum_{k\in\Omega_t} w_{kt}\ell_{kt} = W_t^f L_t^f = W_t^h L_t^h. \tag{6}$$

W_t^f represents the firm's effective cost of hiring one unit of effective labor as measured by the CES composite L_t^f. W_t^h represents the household's effective wage for supplying one unit of effective labor as measured by the CES composite L_t^h.

We assume that the representative household owns the representative firm, which remits its profits to the household as a lump-sum transfer Π_t, so that the household's budget constraint is given by

$$C_t = L_t^h W_t^h + \Pi_t. \tag{7}$$

A competitive general equilibrium of this model is a set of occupational wages w_{kt}, a set of occupational labor quantities ℓ_{kt}, and a consumption level C_t, such that in each period t:

[5] We again assume that the functions $G(\cdot)$ and $H(\cdot)$ obey standard regularity conditions discussed in Appendix B.2.

[6] We show how to derive these results in Appendices B.1 and B.2.

1. The household maximizes its utility taking occupation-specific wage rates as given;
2. The firm maximizes profits taking type-specific wage rates as given; and
3. The markets for goods and each type of occupational labor clear.

We show in Appendix B that the CES aggregate labor quantities L_t^f and L_t^h and the CES unit wage indices W_t^h and W_t^h will then be given by their respective definitions in this section.

3.2 Empirical CES Exact Wage Indices

The CES exact wage indices defined in the previous section cannot be constructed directly from labor market data because they require knowledge of the relative labor demand and supply shifters α_{kt} and γ_{kt} and the elasticities of substitution σ_f and σ_h. In this section, we propose empirically feasible exact wage indices for both the firm and the household. Each of the elements in the indices can be constructed directly from the data except for the elasticities of substitution σ_f and σ_h, the estimation of which we discuss in Section 5.

We begin by developing some additional notation. Define occupation k's share of the total wage bill s_{kt} as

$$s_{kt} \equiv \frac{w_{kt}\ell_{kt}}{\sum_{l\in\Omega_t} w_{lt}\ell_{lt}} = \frac{w_{kt}\ell_{kt}}{Z_t}. \tag{8}$$

Denote by $\Omega^{\star}_{t\text{-}1,t}$ the set of occupations that are present both period t-1 and in period t to distinguish from the set of all occupations present in period t, Ω_t. We call those "continuing" or "common" occupations. We can define the shares of the wage bill expended on these common occupations in periods t and t-1 as

$$\lambda_t \equiv \frac{\sum_{k\in\Omega^{\star}_{t\text{-}1,t}} w_{kt}\ell_{kt}}{\sum_{k\in\Omega_t} w_{kt}\ell_{kt}} = \sum_{k\in\Omega^{\star}_{t\text{-}1,t}} s_{kt} \quad \text{and} \quad \lambda_{t\text{-}1} \equiv \frac{\sum_{k\in\Omega^{\star}_{t\text{-}1,t}} w_{kt\text{-}1}\ell_{kt\text{-}1}}{\sum_{k\in\Omega_{t\text{-}1}} w_{kt\text{-}1}\ell_{kt\text{-}1}} = \sum_{k\in\Omega^{\star}_{t\text{-}1,t}} s_{kt\text{-}1}. \tag{9}$$

For any variable x_t, we will use $x_t^{\star}$ to denote the variable's value for common types. For wage bill shares across common labor types,

$$s^{\star}_{kt} \equiv \frac{w_{kt}\ell_{kt}}{\sum_{l\in\Omega^{\star}_{t\text{-}1,t}} w_{lt}\ell_{lt}} = \frac{s_{kt}}{\lambda_t} \quad \text{and} \quad s^{\star}_{kt\text{-}1} \equiv \frac{w_{kt\text{-}1}\ell_{kt\text{-}1}}{\sum_{l\in\Omega^{\star}_{t\text{-}1,t}} w_{lt\text{-}1}\ell_{lt\text{-}1}} = \frac{s_{kt\text{-}1}}{\lambda_{t\text{-}1}}. \tag{10}$$

We will also denote by $\tilde{x}_t$ for any variable x its unweighted geometric average across all types present in period t, and by $\tilde{x}^\star_t$ that average across common types. We will further denote by $\overline{\ln x}_t$ the simple average of a variable's log values in period t, i.e., $\overline{\ln x}_t = \ln \tilde{x}_t$, and by $\hat{x}_{kt}$ occupation k's log deviation from that log geomean in time t, i.e., $\hat{x}_{kt} = \ln x_{kt} - \overline{\ln x}_t$. $\overline{\ln x^\star_t}$ and $\hat{x}^\star_{kt}$ will refer to the corresponding concepts with respect to common occupations.

Our relative labor demand and supply shifters α_{kt} and γ_{kt} are not separately identifiable from the aggregate shifters A_t and Γ_t without additional assumptions.[7] We follow Redding and Weinstein (2020) in imposing a normalization that the geometric means of α_{kt} and γ_{kt} over the set of common labor types are constant between periods t-1 and t, i.e.,

$$\widetilde{\alpha}^\star_t = \widetilde{\alpha}^\star_{t\text{-}1} \Leftrightarrow \overline{\ln \alpha^\star_t} = \overline{\ln \alpha^\star_{t\text{-}1}} \Leftrightarrow \overline{\Delta \ln \alpha^\star_{kt}} = 0 \text{ and } \widetilde{\gamma}^\star_t = \widetilde{\gamma}^\star_{t\text{-}1} \Leftrightarrow \overline{\ln \gamma^\star_t} = \overline{\ln \gamma^\star_{t\text{-}1}} \Leftrightarrow \overline{\Delta \ln \gamma^\star_{kt}} = 0. \quad (11)$$

These normalizations underlie our interpretation of the α_{kt} and γ_{kt} parameters as *relative* shifters of labor demand and supply across occupations distinct from the *aggregate* shifters A_t and Γ_t.[8] A key insight of Redding and Weinstein (2020) is that shocks to these relative supply and demand parameters can affect aggregate wages.

We can now define empirically feasible CES exact wage indices for both the firm, $\Phi^{\text{CUWI},f}_{t\text{-}1,t}$, and the household, $\Phi^{\text{CUWI},h}_{t\text{-}1,t}$. Following Redding and Weinstein (2020), we call these indices CES unified wage indices, or "CUWIs," because they account for multiple aspects of change in the labor market in a unified fashion. In Appendix B, we derive that these indices take the forms

$$\Phi^{\text{CUWI},f}_{t\text{-}1,t} \equiv \frac{W^f_t}{W^f_{t\text{-}1}} = \left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{1}{\sigma_f - 1}} \frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t\text{-}1}} \left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t\text{-}1}}\right)^{\frac{1}{\sigma_f - 1}}, \quad (12)$$

$$\Phi^{\text{CUWI},h}_{t\text{-}1,t} \equiv \frac{W^h_t}{W^h_{t\text{-}1}} = \left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h + 1}} \frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t\text{-}1}} \left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h + 1}}. \quad (13)$$

[7] For example, suppose that every α_k doubled from period t-1 to period t. It would not be possible to distinguish such a movement from A_t doubling. Likewise, it would not be possible to distinguish every labor supply shifter γ_{kt} doubling from the aggregate labor supply shifter Γ_t doubling.

[8] Abe and Rao (2024) show that a normalization on the (potentially weighted) average log change is necessary and sufficient for the resulting index to be independent of units of measurement.

The CUWIs comprise three terms. The first term, $\left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{1}{\sigma_f - 1}}$ for the firm and $\left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h + 1}}$ for the household, represents an adjustment for turnover in labor types that we call the "Feenstra adjustment."[9] The Feenstra adjustment term collapses to one (i.e., no adjustment to wages) when there is no turnover in occupations or when the elasticities of substitution approach infinity, so that occupations are perfect substitutes. The second term in the CUWIs, $\frac{\widetilde{w}_t^\star}{\widetilde{w}_{t\text{-}1}^\star}$, is the Jevons (1865) index, which is the ratio of unweighted geometric mean wages across common occupations. The third term, $\left(\frac{\widetilde{s}_t^\star}{\widetilde{s}_{t\text{-}1}^\star}\right)^{\frac{1}{\sigma_f - 1}}$ for the firm and $\left(\frac{\widetilde{s}_t^\star}{\widetilde{s}_{t\text{-}1}^\star}\right)^{\frac{-1}{\sigma_h + 1}}$ for the household, adjusts the Jevons index for shifts in relative demand and supply for different occupations over time.[10]

Table A.1 summarizes our notation and formulas, including the various forms that the exact CES wage indices can take depending on whether there is turnover in occupations and whether there are shocks to relative supply and demand. Our firm-side wage indices have exact analogs in the price index literature. With fixed occupations and no relative supply or demand shifts, the firm and household indices coincide and take the form of the Sato-Vartia price index (Sato, 1976; Vartia, 1976). With changing occupations but no relative supply or demand shifts, the firm-side wage index takes the form of the Feenstra-adjusted Sato-Vartia (SVF) index (Feenstra, 1994). With changing occupations and shifts to relative labor supply and demand, the firm-side wage index takes the form of what Redding and Weinstein (2020) call the CES unified price index (CUPI).

Our household- and firm-side wage indices will differ when there are shocks to relative supply and demand across occupations or turnover in the set of occupations present in the labor market. Consider first the case of changes in the set of occupations present. If new occupations earn a larger share of the wage bill in period t than exiting occupations earned in period t-1, the Feenstra adjustment will be negative for firms and positive for workers.[11] To

[9] We choose this name because the firm-side term is equivalent to the adjustment to the Sato-Vartia price index to account for product variety proposed by Feenstra (1994).

[10] We show in Appendices B.1 and B.2 that the second and third terms together collapse to the traditional Sato-Vartia index in the absence of shocks to the relative demand and supply shifters α_{kt} and γ_{kt}.

[11] This statement relies on our assumption, maintained throughout the paper and estimated in the data,

build intuition for these adjustments, suppose that in period t-1 there were two occupations, artists and bakers. In period t, a third occupation, carpenters, enters, while no occupations exit. Then the Feenstra adjustment will *reduce* firms' measured labor costs and *increase* workers' measured wages, consistent with the idea that, in equilibrium, the entry of the new occupation must increase the welfare of both firms and workers.

The difference between the firm- and household-side adjustments for shifts in relative demand and supply is more subtle. Although the adjustments relative to the Jevons index must have opposite signs for the firm- and household-side indices, the same is not true for adjustments relative to the Sato-Vartia index. The log Sato-Vartia index is defined as

$$\ln \Phi^{SV}_{t\text{-}1,t} \equiv \sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \omega^{\star}_{kt} \ln \left(\frac{w^{\star}_{kt}}{w^{\star}_{kt\text{-}1}} \right), \quad \omega^{\star}_{kt} \equiv \frac{s^{\star}_{kt} - s^{\star}_{kt\text{-}1}}{\ln(s^{\star}_{kt}) - \ln(s^{\star}_{kt\text{-}1})} \Bigg/ \left(\sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \frac{s^{\star}_{kt} - s^{\star}_{kt\text{-}1}}{\ln(s^{\star}_{kt}) - \ln(s^{\star}_{kt\text{-}1})} \right). \tag{14}$$

Our CUWIs can be re-expressed in log terms as

$$\ln \Phi^{\text{CUWI},f}_{t\text{-}1,t} = \left(\frac{1}{\sigma_f - 1} \right) \ln \left(\frac{\lambda_t}{\lambda_{t\text{-}1}} \right) + \sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \omega^{\star}_{kt} \ln \left(\frac{w^{\star}_{kt}}{w^{\star}_{kt\text{-}1}} \right) - \sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \omega^{\star}_{kt} \ln \left(\frac{\alpha^{\star}_{kt}}{\alpha^{\star}_{kt\text{-}1}} \right) \tag{15}$$

$$\ln \Phi^{\text{CUWI},h}_{t\text{-}1,t} = \left(\frac{-1}{\sigma_h + 1} \right) \ln \left(\frac{\lambda_t}{\lambda_{t\text{-}1}} \right) + \sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \omega^{\star}_{kt} \ln \left(\frac{w^{\star}_{kt}}{w^{\star}_{kt\text{-}1}} \right) - \sum_{k \in \Omega^{\star}_{t\text{-}1,t}} \omega^{\star}_{kt} \ln \left(\frac{\gamma^{\star}_{kt}}{\gamma^{\star}_{kt\text{-}1}} \right). \tag{16}$$

The first terms in equations (15) and (16) are the Feenstra adjustments, the second terms are the Sato-Vartia index, and the third terms are the adjustments for shocks to relative demand and supply. These formulations clarify that the CUWIs' adjustments to the SVF index need not take opposite signs but will depend on the specific constellations of shocks to α_{kt} and γ_{kt}. If high-weight occupations tend to experience positive shocks to α_{kt} relative to low-weight occupations, the firm-side CUWI will be adjusted downward. Conversely, if high-weight occupations tend to experience negative shocks to α_{kt} relative to low-weight occupations, the firm-side CUWI will be adjusted upward. Analogous logic applies to the household-side index, which will be adjusted downward if high-weight occupations tend to experience positive shocks to γ_{kt} and will be adjusted upward if high-weight occupations tend to experience negative shocks to γ_{kt}.

that $\sigma_f > 1$ and $\sigma_h > 0$.

Furthermore, the shock adjustments to the firm-side index will tend to be downward if the dispersion of the relative demand parameters is increasing over time, while the shock adjustments to the household-side index will tend to be upward if the dispersion of the relative supply parameters is increasing. Because firms prefer lower wages all else equal, a downward shock adjustment to the firm-side CUWI corresponds to an increase in their welfare. Households prefer higher wages all else equal, so an upward shock adjustment to the household-side CUWI corresponds to an increase in their welfare. In other words, both sides of the market tend to benefit from increasing dispersion in the relative supply or demand parameters on their own side of the market. This relationship is analogous to the well-known result that consumers benefit from increasing dispersion in prices because it offers greater opportunities for substitution.

3.3 Equilibrium and Measuring Supply and Demand Shocks

We show in Appendix B that in the equilibrium of our two-sided CES labor market, the following relationships between labor quantities and wages hold for each occupation k on the demand and supply sides of its market, respectively:

$$\hat{\ell}_{kt} = (\sigma_f - 1)\,\hat{\alpha}_{kt} - \sigma_f \hat{w}_{kt} \quad \text{and} \quad \hat{\ell}_{kt} = -(1 + \sigma_h)\,\hat{\gamma}_{kt} + \sigma_h \hat{w}_{kt}. \tag{17}$$

We can therefore solve for each continuing occupation's equilibrium relative wages and labor quantities in terms of relative supply and demand parameters as

$$\hat{\ell}_{kt} = \frac{\sigma_h\,(\sigma_f - 1)}{\sigma_h + \sigma_f}\hat{\alpha}_{kt} - \frac{\sigma_f\,(\sigma_h + 1)}{\sigma_h + \sigma_f}\hat{\gamma}_{kt}, \quad \text{and} \tag{18}$$

$$\hat{w}_{kt} = \frac{\sigma_f - 1}{\sigma_h + \sigma_f}\hat{\alpha}_{kt} + \frac{\sigma_h + 1}{\sigma_h + \sigma_f}\hat{\gamma}_{kt}. \tag{19}$$

Appendix Figure B.1 displays equilibrium in our model of the labor market graphically.

Equations (18) and (19) show that our two equilibrium outcomes, wages and labor quantities, are determined jointly in a system of two linear equations. Given the labor supply

and labor demand elasticities σ_h and σ_f, we can recover estimates of the relative demand and supply shifters α_{kt} and γ_{kt}. Redding and Weinstein (2020) show how to invert their CES demand system to recover estimates of demand shifters for common goods. We follow Ehrlich et al. (2022) and Olivares (2018) in generalizing their approach to recover relative demand shifters for entrants and exiters in addition to common goods.[12]

Although our simple model of the labor market omits many potentially important deviations from the frictionless perfectly competitive supply and demand benchmark, we believe that this benchmark framework offers important advantages. It allows us to derive closed-form estimates of the changes in relative labor supply and demand for each occupation, which we can validate by comparison to well-documented trends in the labor market. Our system also permits us to compute counterfactual wages and labor quantities under alternative constellations of shocks enables model-consistent calculation of firm- and household-side wage indices. Finally, our model clarifies the assumptions necessary to estimate the elasticities of substitution in a general equilibrium setting.

Appendix B.4 describes how we can recover α_{kt} and γ_{kt} as

$$\alpha_{kt} = \widetilde{\alpha}_t^\star \left(\frac{w_{kt}}{\widetilde{w}_t^\star}\right) \left(\frac{s_{kt}}{\widetilde{s}_t^\star}\right)^{\frac{1}{\sigma_f - 1}} \lambda_t{}^{\frac{-1}{\sigma_f - 1}} \Rightarrow \Delta\hat{\alpha}_{kt}^\star = \Delta\hat{w}_{kt}^\star + \frac{1}{\sigma_f - 1}\Delta\hat{s}_{kt}^\star, \tag{20}$$

$$\gamma_{kt} = \widetilde{\gamma}_t^\star \left(\frac{w_{kt}}{\widetilde{w}_t^\star}\right) \left(\frac{s_{kt}}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}} \lambda_t{}^{\frac{1}{1+\sigma_h}} \Rightarrow \Delta\hat{\gamma}_{kt}^\star = \Delta\hat{w}_{kt}^\star - \frac{1}{1+\sigma_h}\Delta\hat{s}_{kt}^\star. \tag{21}$$

The left-hand-side versions of these formulas, for the levels of the supply and demand parameters, require knowledge of the geometric means $\widetilde{\alpha}_t^\star$ and $\widetilde{\gamma}_t^\star$. Without loss of generality, we normalize $\widetilde{\alpha}_t^\star$ and $\widetilde{\gamma}_t^\star$ to equal one in the first period of our analysis. Our normalization that $\widetilde{\alpha}_t^\star = \widetilde{\alpha}_{t\text{-}1}^\star$ and $\widetilde{\gamma}_t^\star = \widetilde{\gamma}_{t\text{-}1}^\star$ applies only to continuing occupations, so the geometric means $\widetilde{\alpha}_t^\star$ and $\widetilde{\gamma}_t^\star$ can (and empirically do) drift over time as new occupations enter and old occupations exit. Given our estimates of α_{kt} and γ_{kt}, we are able to calculate the geometric means $\widetilde{\alpha}_t^\star$ and $\widetilde{\gamma}_t^\star$ for the set of common types between periods t and $t+1$. We can then roll these

[12] That approach applies directly to recovering the relative labor demand shifters α_{kt}; it is straightforward to develop a parallel approach to recover the relative labor supply shifters γ_{kt}.

parameters forward to calculate α_{kt+1} and γ_{kt+1}. We discuss this procedure in more detail in Appendix B.4.

Note that equations (18) and (19) permit straightforward calculation of relative wages and labor quantities for any set of labor supply and demand parameters α_{kt} and γ_{kt}. We can therefore compute counterfactual relative labor quantities $\hat{\ell}_{kt}^{cf}$, relative wages $\hat{w}_{kt}^{cf}$, and wage bill shares s_{kt}^{cf} for counterfactual constellations of relative labor supply γ_{kt}^{cf} and demand α_{kt}^{cf}. Two counterfactual scenarios of particular interest are a scenario featuring the estimated relative demand shocks but no relative supply shocks, and a scenario featuring the estimated relative supply shocks but no relative demand shocks. We study how the distribution of relative wages would have evolved under these counterfactual scenarios in Section 8.

4 Data and Empirical Implementation

Our methodology requires data on employment levels and wages for a panel of occupations. Integrated Public Use Microdata Series (IPUMS, Ruggles et al., 2023) files furnish estimates of occupational employment and wages for each decade based on decadal Censuses and the American Community Survey. To link those occupations into a panel and explore the role of occupational entry and exit, we must bridge changing occupation codes across time while assessing when particular occupations appear or disappear. To do so, we utilize concordance tabulations published by the U.S. Census Bureau at the level of 3-digit occupation codes. This section describes first the occupational linking algorithm before discussing the construction of employment and wage data.

Several comprehensive crosswalks have been created to facilitate analysis of occupations across time.[13] An explicit goal of these projects has been to maximize longitudinal consistency, which requires minimizing entry and exit—sometimes entirely. The literature on new work has also developed measures of occupational turnover. Lin (2011) identifies new

[13]Notable examples include the IPUMS occ1950 and occ1990 variables described in Meyer and Osborne (2005) or the balanced panel of occupations constructed by Dorn (2009).

work using the Dictionary of Occupational Titles and the Classified Alphabetical Index of Occupations and Industries.[14] The new occupations identified in these documents are highly granular, having 5- to 9-digit occupation codes. Earnings and employment data are not available at this level of detail, so authors employing this technique impute the share of new work as the share of new titles appearing within each harmonized 3-digit occupation code.

These previous frameworks cannot be applied directly to the CES wage indices we developed in Section 3 for several reasons. First, they do not provide measures of occupational exit, which we show is a relevant phenomenon in the U.S. labor market. Second our wage indices require us to measure entering and exiting occupations' shares of the wage bill just after entry or just prior to exit. The occupational title-based approach does not provide a direct measure of new work's share of the wage bill. Furthermore, we would like to be able to capture the effective emergence and disappearance of entire lines of work (e.g., web developers on the one hand and telegraph messengers on the other) as represented by 3-digit occupation codes.

We therefore develop a new crosswalk of occupations that connects comparable occupations across time and identifies when obsolete occupations exit or new occupations enter. Our crosswalk applies a simple, flexible algorithm to comparability tables published by the Census Bureau, allowing other researchers to employ and update it in future work. The Census Bureau generates these comparability tables when it updates the occupational classification system each decade.[15] Those changes might involve simple re-labeling, splitting or consolidating similar occupations, or a more comprehensive re-organization. To help users understand the implications of each decade's changes, the Census Bureau classifies a sample of workers under both the current and previous occupational definitions. Our crosswalk re-

[14] Autor et al. (2022) have digitized earlier versions of this dictionary to extend this measure back to 1940.

[15] Tables can be found in U.S. Census Bureau (1968, Table 1) for 1950-1960, U.S. Census Bureau (1972, Table 1) for 1960-1970, U.S. Census Bureau (1989, Table 1) for 1970-1980, U.S. Census Bureau (2003, Table 2) for 1990-2000, U.S. Census Bureau (2011) for 2000-2010, and U.S. Census Bureau (2020, Tables E1-E3, Table F-1) for 2010-2020. The occupational schemes for 1950 and 1990 strongly resemble their predecessors and did not prompt such reports. A concise summary of 1980-1990 changes is summarized in U.S. Census Bureau (1994). We construct a mapping between 1940 and 1950 occupations using Tables 124 and 125 from U.S. Census Bureau (1953).

lies on a simple intuition: continuing occupations should exhibit a strong continuity across classification schemes, and a lack of continuity is the hallmark of new or disappearing work.

Table 1 shows several examples from the 1960–1970 comparability file (with condensed formatting). Each row grouping describes how many workers in a given 1970 category would have fallen into the corresponding classification in 1960. Many mappings in every decade are unambiguous: all actors and actresses in 1970 would also have been classified as actors in 1960. Dentists and mathematicians also exhibit complete continuity in classification from 1960 to 1970.[16]

Table 1: Example entries from 1960–1970 concordance files

1970 code	1970 occupation with 1960 components and codes	Employment
175	Actors	11,656
	Actors and actresses (10)	11,656
10	Chemical engineers	41,925
	Engineers, chemical (81)	41,315
	Engineers, mechanical (85)	610*
3	Computer programmers	8,283
	Prof, tech, and kindred workers nec (195)	8,283*
62	Dentists	83,198
	Dentists (71)	83,198
923	Health trainees	57,549
	Nurses, student professional (151)	57,549
30	Judges	11,140
	Lawyers and judges (105)	11,140*
31	Lawyers	207,118
	Lawyers and judges (105)	201,918*
	Prof, tech, and kindred workers nec (195)	5,200*
35	Mathematicians	7,685
	Mathematicians (135)	7,685
65	Physicians medical and osteopathic	232,826
	Osteopaths (153)	3,951
	Physicians and surgeons (162)	228,875
672	Spinners twisters and winders	177,416
	Spinners textile (710)	52,190
	Operatives and kindred workers nec (775)	125,226*

Notes: nec = "not elsewhere classified". *Also maps to other occupations.
Source: U.S. Census Bureau (1972).

[16]The example of "actors" mapping to "actors and actresses" is an example of a simple re-labeling. Our procedure also detects more complicated re-labelings that could pose problems for fuzzy text matching approaches. For instance, all 1970 "health trainees" would have been labeled "nurses, student and professional" in 1960.

To understand the linking algorithm in more complicated circumstances, consider that the concordance files define a collection of potential predecessor-successor pairs between every decade. In Table 1, the potential pairs include (10, 175), (81, 10), (85, 10), (195, 3), etc. For every potential predecessor-successor pair, we define two measures (A) and (B):

(A) For all employment this decade that would have previously been classed into a given predecessor code, what share is now classified in the given successor code?

(B) What share of employment in this decade's successor code would have been previously classed under this predecessor code?

Table 2 below computes these shares for select examples from above. These calculations show that while chemical engineers do not exhibit a perfect one-to-one mapping as do actors, 98.5% of individuals labeled "chemical engineers" in 1970 would have been labeled "engineers, chemical" in 1960 (measure A). At the same time, 100% of 1970 workers who would have been classified as "engineers, chemical" in 1960 are also classified as "chemical engineers" in 1970 (measure B). Our algorithm asserts that these occupations should be linked, the few matches of 1960 mechanical engineers to 1970 chemical engineers notwithstanding. Formally, we would like to classify the pair (81, 10), i.e., "engineers, chemical" and "chemical engineers," as a continuing occupation but not to link the pair (85, 10), i.e., "engineers, mechanical" and "chemical engineers."

This example points to our algorithm's core logic for classifying continuing occupations. We begin linking continuing occupations by applying a bi-directional threshold rule: shares (A) and (B) must both be at least N%. Such a rule would link the pairs (10, 175), (81,10), and (31, 105) for any sensible threshold N. At the same time, the bi-directional threshold rule eliminates the spurious predecessor-successor pairs in (85, 10) and (195, 31).[17] It would also classify "computer programmers" as a new occupation in 1970 because they represent a small minority of employment in the group from which they emerge, and it would classify

[17]The (195, 31) pair illustrates a common feature of spurious matches: that one member of the pair (here lawyers) already matches to another occupation under the bi-directional threshold rule.

"spinners textile" as an exiting occupation, because these workers are also a minority of the new "spinners twisters and winders" occupation into which they merge. We call this simple bi-directional threshold rule our "narrow continuer rule" for linking occupations.

Table 2: Bi-directional shares computed from 1960-1970 concordance files

Successor name	Succ. code	Predecessor name	Pred. code	(A)	(B)
Actors	175	Actors and actresses	10	100%	100%
Chem. engineers	10	Engineers, chemical	81	98.5%	100%
Chem. engineers	10	Engineers, mechanical	85	1.5%	0.4%
Comp. prog.	3	Prof, tech, & kindred work, nec	195	100%	2.4%
Judges	30	Lawyers and judges	105	5.2%	100%
Lawyers	31	Lawyers and judges	105	94.8%	97.5%
Lawyers	31	Prof, tech, & kindred work, nec	195	1.7%	2.5%
Spinners twist...	672	Spinners textile	710	16.0%	100%
Spinners twist...	672	Operatives and kindred workers	775	84.0%	3.5%

These examples all reflect appropriate classifications, but we expect this narrow continuer rule to overstate occupational turnover. The case of lawyers and judges highlights why. In 1970, the Census Bureau began to enumerate "lawyers" (31) and "judges" (30) separately from the previously grouped "lawyers and judges" (105). We would like our algorithm to avoid classifying judges as a new occupation in 1970 simply because of the additional detail in the Census Bureau's classification scheme. Similar issues arise when managers of occupations are first enumerated separately. A parallel problem occurs when occupations are combined, for instance when "osteopaths" (153) are grouped with "physicians and surgeons" (162) to form "physicians medical and osteopathic" (65).

To address this challenge, we also develop a "broad continuer rule" that aggregates occupations in its search for continuers. When looking for successor codes, it first sums measure (A) for all codes with $(B) > M\%$. Likewise, when looking for predecessor codes, it sums measure (B) for all codes with $(A) > M\%$. Because 100% of judges in 1970 would have been classified as "lawyers and judges" in 1960, our broad measure re-creates this aggregate category by combining the employment and earnings of lawyers and judges into one continuing occupation when comparing 1960 and 1970.[18]

[18]Our algorithm operates independently on each pair of successive decades, so when comparing 1970 to

To see why our broad continuer rule may understate occupational turnover, consider the case of programmers. Computer programmers, computer systems analysts, computer specialists, operation systems researchers/analysts, and many other occupations potentially representing new work in 1970 would have been classified as "professional, technical, and kindred workers nec" under the 1960 classification scheme. By pre-grouping such occupations, the broad continuer rule assumes that their enumeration in 1970 reflects increases in the level of detail the Census Bureau used for its classifications rather than the emergence of new work. We therefore treat our narrow and broad continuer rules as providing upper and lower bounds, respectively, on the degree of occupational turnover.

Table 3: Examples of Occupational Entry and Exit

Decades	Occupation	Classification
1950-1960	AC, heating, & refrigeration repair	Entrant
1950-1960	Service workers outside household nec	Exiter
1960-1970	Petroleum engineers	Entrant
1960-1970	Express & railway mail clerks	Exiter
1970-1980	Financial managers	Entrant
1970-1980	Telegraph messengers	Exiter
1980-1990	Nuclear engineer	Entrant
1980-1990	Apparel & fabric patternmakers	Exiter
1990-2000	Network & computer system admin	Entrant
1990-2000	Stevedores	Exiter
2000-2010	Web Developers	Entrant
2000-2010	Bridge & lock tenders	Exiter
2010-2020	Misc. media & communications workers	Entrant
2010-2020	Dredge, excavation, and loading ops	Exiter

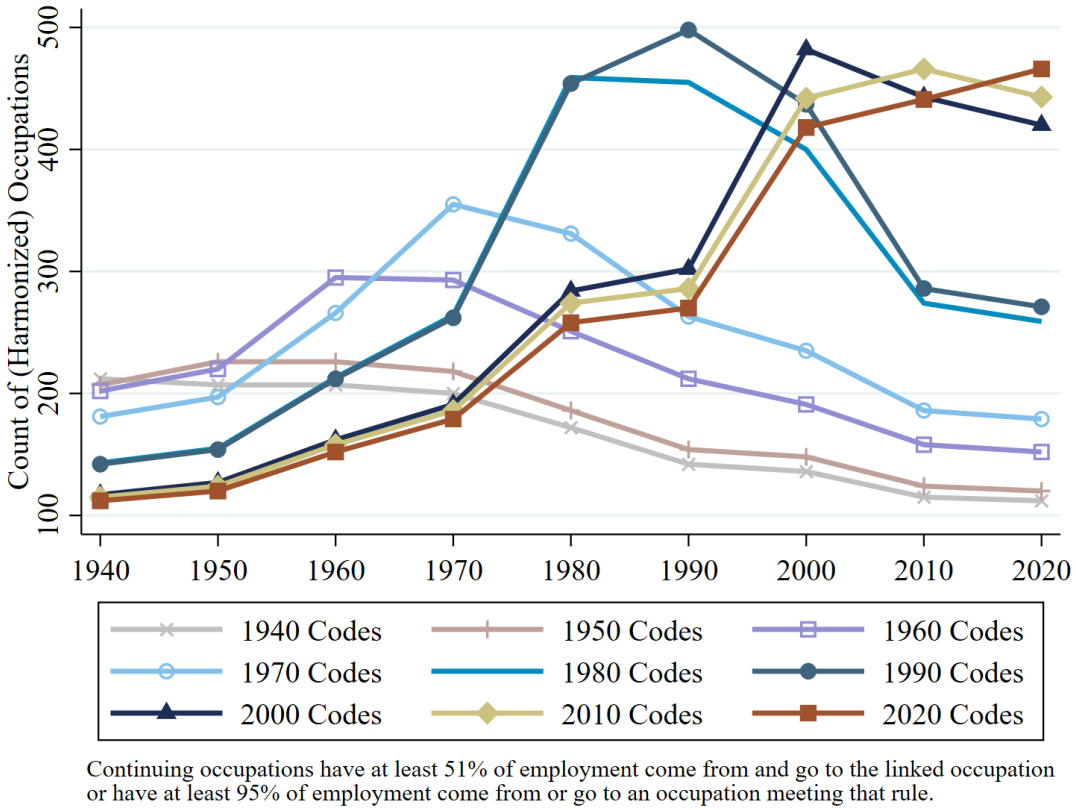
We adopt our broad continuer rule, setting $M = 95$ and $N = 51$, as our preferred algorithm for linking occupations over time. The algorithm identifies 112 occupations existing continuously from 1940–2020 for jobs such as bakers, carpenters, janitors, librarians, nurses, pharmacists, and truck drivers. Together, these occupations account for just over half of employment in 1940 and approximately one-third of employment in 2020.[19] Table 3 presents a list of occupations identified as entering or exiting in each decade under this measure. In

1980 we consider lawyers and judges to be distinct occupations.

[19]These statistics correspond fairly well to the estimated shares of new work reported by Autor et al. (2022, p. 12). Using an imputation methodology for shares of new types of work within continuing occupations, they estimate that new work accounts for approximately 60% of 2018 employment.

Appendix C.2, we describe the results of using our narrow continuer rule, and show that our conclusions are not sensitive to our choice of rule.

Figure 2: Counts of Distinct Occupations Identified by Our Algorithm, 1940–2020



Continuing occupations have at least 51% of employment come from and go to the linked occupation or have at least 95% of employment come from or go to an occupation meeting that rule.

Figure 2 illustrates the counts of occupations identified by our algorithm by decade. Each line in the figure displays the count of three-digit Census occupation codes that existed in the labor market as of each decade from 1940 to 2020. For instance, the lavender line labeled "1960 Codes" shows that there were nearly 300 3-digit occupations in the 1960 Census. Tracing the line to the left shows that our algorithm identifies roughly 200 of those occupations as also having existed in 1940 and roughly 150 of them as still existing in 2020. As expected, the number of existing occupations according to each decades codes always peaks contemporaneously. Additionally, there was significant growth in the number of 3-digit occupations identified by the Census Bureau from 1940 to 1990, but the number has receded a bit from then through 2020.

In Section 6, we show that occupational turnover had a much larger economic impact in the later part of our study period, when the number of occupation codes was roughly steady, than in earlier decades when the number of occupation codes was rising significantly. These results provide reassurance that our algorithm for identifying occupational turnover does not mechanically drive the estimated turnover adjustments to our wage indices because the Census Bureau categorized occupations at steadily increasing levels of refinement. The estimates seem instead to be tied to a genuine trend toward greater occupational specialization and the growth of new work over this period.

To construct a panel of wages and labor inputs for each of these occupations, we rely on microdata from detailed decadal Censuses from 1940-2000 and five-year samples from the the American Community Survey thereafter.[20] IPUMS makes available samples covering between 1% and 5% of the total population, including frequency weights where necessary to make the sample representative of the whole population. A rich set of individual and household characteristics is available, of which we utilize several variables selected to maximize longitudinal consistency. For occupation, the linking algorithm above details our procedure for handling 3-digit Census codes as they appear in each decade. Labor supply in each occupation is defined as the count of full-time workers, identified as those employed for a minimum of 40 weeks in the preceding year and who report working at least 30 hours in the week preceding the survey. Wages are measured as the individual's pre-tax wage and salary income from the previous year.[21]

[20]Detailed questions about employment and income were no longer asked of any decadal Census respondent after the 2000 Census. The American Community Survey has replaced the decadal Census as a source of such information. It employs rolling 1% annual sample windows designed to capture 5% of the population over a five-year interval. To maintain continuity with the decadal Census approach, which asks respondents about income received in the previous year (so that the 2000 Census captures 1999 income) we use five-year samples centered on years preceding a decadal Census (2007-2011 and 2017-2021).

[21]The focus on full-time employment reflects our inability to distinguish, in this data, detailed changes in labor supply on the intensive margin, i.e. changes in hours worked per week or weeks worked per year. We therefore conceptualize occupational wages as average annual wage and salary earnings across individuals.

5 Estimating Elasticities of Substitution

Implementing our CES exact wage indices requires estimates of the elasticities of substitution σ_f and σ_h. One widely used approach to this challenge in the estimation of CES demand systems comes from Feenstra (1994). That approach entails an orthogonality assumption between supply and demand shocks within each period (and a heteroskedasticity assumption across occupations). We are skeptical that the orthogonality assumption can hold in our setting. For instance, Conzelmann, Hemelt, Hershbein, Martin, Simon, and Stange (2023) estimate that the higher education sector displays an elasticity of undergraduate credits to shifts in labor demand of around 1.3 at a five-year horizon, with responses increasing over time. We measure wages and labor quantities at a decadal frequency, which would appear to give ample time for relative supply and demand to respond to shocks to one another.

5.1 Firm Side

Redding and Weinstein (2018) propose what they call the "reverse-weighting" (RW) and "generalized reverse-weighting" (GRW) estimators, which rely on the structure of their CES demand system. We begin by discussing the reverse-weighting estimator, largely paralleling the presentation in Redding and Weinstein (2018). Their procedure estimates the elasticity of consumer demand across products, which corresponds to the firm or labor-demand side of our model. We therefore apply the RW and GRW estimators to estimating the firm-side elasticity of substitution σ_f. We can define CES exact wage indices in periods t-1 and t among continuing occupations as

$$W_t^{f\star} = \left[\sum_{k\in\Omega^\star_{t\text{-}1,t}} \left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}} \quad \text{and} \quad W_{t\text{-}1}^{f\star} = \left[\sum_{k\in\Omega^\star_{t\text{-}1,t}} \left(\frac{w_{kt\text{-}1}}{\alpha_{kt\text{-}1}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}. \tag{22}$$

Our firm-side CUWI can be written in terms of these indices as

$$\Phi^{\text{CUWI},f}_{t\text{-}1,t} = \left(\frac{\lambda_{t\text{-}1}}{\lambda_t}\right)^{\frac{1}{\sigma_f-1}} \left(\frac{W_t^{f\star}}{W_{t\text{-}1}^{f\star}}\right). \tag{23}$$

We can also express occupation k's shares of total labor income and of labor income from continuing occupations as

$$s_{kt} = \frac{\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}{\left(W_t^f\right)^{1-\sigma_f}} \quad \text{and} \quad s_{kt}^\star = \frac{\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}{\left(W_t^{f\star}\right)^{1-\sigma_f}}, \tag{24}$$

with parallel expressions holding in period t-1.

The Lloyd-Moulton indices (Lloyd, 1975; Moulton, 1996) are a pair of expressions that evaluate the cost of living from the perspective of tastes in periods t-1 and t, respectively. Both indices should theoretically be equivalent to equation (23) in the absence of relative supply shocks for the household side or demand shocks on the firm side. The "forward index" (Lloyd, 1975) evaluates the ratio of wages in period t-1 to wages in period t by substituting

$$s_{kt}^\star = \frac{\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}{\sum_{k\in\Omega_{t\text{-}1,t}^\star}\left(\frac{w_{kt\text{-}1}}{\alpha_{kt\text{-}1}}\right)^{1-\sigma_f}} \Rightarrow \sum_{k\in\Omega_{t\text{-}1,t}^\star}\left(\frac{w_{kt\text{-}1}}{\alpha_{kt\text{-}1}}\right)^{1-\sigma_f} = \left(\frac{1}{s_{kt\text{-}1}^\star}\right)\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f} \tag{25}$$

to generate the following expression equivalent to equation (23):

$$\Phi_{t\text{-}1,t}^F = \left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{1}{\sigma_f-1}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^\star} s_{kt\text{-}1}^\star\left(\frac{w_{kt}/\alpha_{kt}}{w_{kt\text{-}1}/\alpha_{kt\text{-}1}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}. \tag{26}$$

A parallel substitution for the *inverse* wage ratio $W_{t\text{-}1}^f/W_t^f$ yields the "backward index" (Moulton, 1996):

$$\Phi_{t\text{-}1,t}^B = \left(\frac{\lambda_{t\text{-}1}}{\lambda_t}\right)^{\frac{1}{\sigma_f-1}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^\star} s_{kt}^\star\left(\frac{w_{kt\text{-}1}/\alpha_{kt\text{-}1}}{w_{kt}/\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}. \tag{27}$$

Additional manipulation helpfully groups the terms as:

$$\Phi_{t\text{-}1,t}^F = \left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{1}{\sigma_f-1}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^*} s_{kt\text{-}1}^*\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^*} s_{kt\text{-}1}^*\left(\frac{\alpha_{kt}}{\alpha_{kt\text{-}1}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}, \tag{28}$$

$$\Phi_{t\text{-}1,t}^B = \left(\frac{\lambda_{t\text{-}1}}{\lambda_t}\right)^{\frac{1}{\sigma_f-1}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^*} s_{kt}^*\left(\frac{w_{kt\text{-}1}}{w_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}\left[\sum_{k\in\Omega_{t\text{-}1,t}^*} s_{kt}^*\left(\frac{\alpha_{kt\text{-}1}}{\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1+\sigma_f}}. \tag{29}$$

If there were no labor supply shocks, the forward index and the inverse of the backward

index would both be theoretically equivalent to $\Phi_{t\text{-}1,t}^{CUWI,f}$ in equation (23) and therefore to each other. Redding and Weinstein's (2018) RW estimator finds the value of σ_f that brings the forward and (inverse) backward indices as close to the CUWI as possible.[22] Appendix C provides the formal moment conditions for the reverse-weighting estimator of σ_f.

Redding and Weinstein (2018) provide two conditions under which the RW estimator will consistently estimate the true elasticity of substitution, but we are skeptical that either one is likely to apply in our context. The first condition for the consistency of the RW estimator is if relative labor demand shocks are small, i.e., $\alpha_{kt}/\alpha_{kt\text{-}1} \to 1$. We measure wages at the decadal frequency, over which we believe that relative labor supply and demand shocks can be economically large; we provide evidence to that effect in Section 7 below. More fundamentally, one of our primary interests in this paper is to study how CES wage indices that account for such shocks differ from traditional indices that do not. Indeed, in our simple model, relative supply and demand shocks are the only factors that drive shifts in relative wages for continuing occupations. Relying on the assumption that relative supply and demand shocks are small arguably therefore sits in some tension with our underlying motivation. The second condition for the consistency of the RW estimator of σ_f is if relative labor demand shocks are uncorrelated with wage changes (i.e., $\alpha_{kt}/\alpha_{kt\text{-}1} \perp w_{kt}/w_{kt\text{-}1}$). As noted in our discussion of the Feenstra (1994) estimator, there is evidence that labor supply and demand respond to each other over a decadal frequency, violating this assumption.

Redding and Weinstein (2018) develop their GRW estimator to account for cases in which neither condition for the consistency of the RW estimator holds. The GRW estimator assumes that the data generating process for changes in wages and relative demand is partitioned between a portion that loads on changes in wages and an uncorrelated error term. It relies on that assumed partition to construct counterfactual period-t wage bill shares that remove the changes in shares from period t-1 arising from the correlation between demand shifts and wage changes. These counterfactual "purged" shares are then substituted in place

[22] In doing so, it minimizes the contribution from the unobservable terms by setting the last bracketed expressions in equations (28) and (29) as close to 1 as possible.

of the observed shares s_{kt} in the moment conditions in equation (33) below.

Concretely, the GRW estimator assumes that relative demand shifts are correlated with wage changes according to the relationship

$$\ln\left(\frac{\alpha_{kt}}{\alpha_{kt\text{-}1}}\right) = \rho_{\alpha w}\left(\frac{\chi_\alpha}{\chi_w}\right)\ln\left(\frac{w_{kt}/\tilde{w}_t}{w_{kt\text{-}1}/\tilde{w}_{t\text{-}1}}\right) + \ln\left(\frac{\epsilon_{kt}}{\epsilon_{kt\text{-}1}}\right), \tag{30}$$

where χ_α and χ_w are the standard deviations of relative demand shifts and wage changes, respectively, and $\rho_{\alpha w}$ captures their correlation. We assume that $\ln(\epsilon_{kt}/\epsilon_{kt\text{-}1}) \perp \ln(w_{kt}/w_{kt\text{-}1})$. This assumption allows us to construct counterfactual expenditure shares $S^\star_{kt}(\sigma_f)$ that remove the change in shares arising from the correlation of demand shifts with wage changes. We can express $S^\star_{kt}(\sigma_f)$ as a function of the elasticity of substitution σ_f and observable data:

$$S^\star_{kt}(\sigma_f) = \frac{(w_{kt}/w_{kt\text{-}1})^{-\rho_{\alpha w}\left(\frac{\chi_\alpha}{\chi_w}\right)(\sigma_f-1)} s^\star_{kt}}{\sum_{\ell\in\Omega^\star_{t\text{-}1,t}}(w_{kt}/w_{kt\text{-}1})^{-\rho_{\alpha w}\left(\frac{\chi_\alpha}{\chi_w}\right)(\sigma_f-1)} s^\star_{\ell t}}, \text{ where} \tag{31}$$

$$\rho_{\alpha w}\left(\frac{\chi_\alpha}{\chi_w}\right) = \left(\frac{1}{\sigma_f - 1}\right)\left[\sigma_f - 1 + \frac{\chi_{ws}}{\chi^2_w}\right], \tag{32}$$

χ_{ws} is the covariance of wages and labor income shares, and χ^2_w is the variance of wages across labor types. We can substitute these counterfactual shares in place of the observed shares in equations (28) and (29) and set both equal to the firm-side CUWI to arrive at the following set of moment conditions with which to estimate σ_f:

$$M^{GRW}(\sigma_f, \text{data}) = \tag{33}$$

$$\begin{bmatrix} \frac{-1}{\sigma_f-1}\ln\left(\sum_{k\in\Omega^\star_{t\text{-}1,t}} s^\star_{kt\text{-}1}\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right)^{\sigma_f-1}\right) - \frac{1}{N^{\mathbb{C}}_{t\text{-}1,t}}\sum_{k\in\Omega^\star_{t\text{-}1,t}}\left[\ln\left(\frac{\tilde{w}_{kt}}{\tilde{w}_{kt\text{-}1}}\right) - \frac{1}{\sigma_f-1}\ln\left(\frac{\tilde{S}^\star_{kt}(\sigma_f)}{\tilde{s}^\star_{kt\text{-}1}}\right)\right] \\ \frac{1}{\sigma_f-1}\ln\left(\sum_{k\in\Omega^\star_{t\text{-}1,t}} S^\star_{kt}(\sigma_f)\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right)^{1-\sigma_f}\right) - \frac{1}{N^{\mathbb{C}}_{t\text{-}1,t}}\sum_{k\in\Omega^\star_{t\text{-}1,t}}\left[\ln\left(\frac{\tilde{w}_{kt}}{\tilde{w}_{kt\text{-}1}}\right) - \frac{1}{\sigma_f-1}\ln\left(\frac{\tilde{S}^\star_{kt}(\sigma_f)}{\tilde{s}^\star_{kt\text{-}1}}\right)\right] \end{bmatrix},$$

where $N^{\mathbb{C}}_{t\text{-}1,t} = |\Omega^\star_{t\text{-}1,t}|$ represents the number of common labor types across periods t-1 and t. The GRW estimator solves for σ_f using the unweighted generalized method of moments

(GMM) to minimize the squared values of the moment conditions in equation (33):

$$\hat{\sigma}_f = \arg\min_{\sigma_f} \left\{ M^{GRW'}(\sigma_f) \mathbb{I} M^{GRW}(\sigma_f) \right\}. \tag{34}$$

Redding and Weinstein (2018) show that the GRW estimator will be consistent for σ_f if the partition of labor demand shocks in equation (30) into term that is correlated with wage changes and an uncorrelated error holds.

5.2 Household Side

In this section, we describe our approach to estimating the household-side elasticity of substitution across occupations, σ_h. Taking first differences of the equation for labor supply curve in equation (17), we can write the change in occupation k's relative labor input as

$$\Delta \hat{\ell}^{\star}_{kt} = \sigma_h \Delta \hat{w}^{\star}_{kt} - (1 + \sigma_h) \Delta \hat{\gamma}^{\star}_{kt}. \tag{35}$$

Classic simultaneity issues will likely bias a simple ordinary least squares (OLS) regression of $\Delta \hat{\ell}^{\star}_{kt}$ on $\Delta \hat{w}^{\star}_{kt}$. To illustrate the difficulty, we can take first differences of equation (19) to write the change in occupation k's relative wage in terms of supply and demand shocks as

$$\Delta \hat{w}^{\star}_{kt} = \frac{\sigma_f - 1}{\sigma_h + \sigma_f} \Delta \hat{\alpha}^{\star}_{kt} + \frac{\sigma_h + 1}{\sigma_h + \sigma_f} \Delta \hat{\gamma}^{\star}_{kt}.$$

This expression shows that wages $\Delta \hat{w}^{\star}_{kt}$ are in part determined by and will generically be correlated with supply shocks, which appear as the error term $-(1 + \sigma_h) \Delta \hat{\gamma}^{\star}_{kt}$ in an OLS regression of equation 35.

We therefore turn to an instrumental variables strategy for estimating σ_h. This strategy requires locating a valid instrumental variable z_{kt} such that:

1. z_{kt} is correlated with $\Delta \hat{w}^{\star}_{kt}$ (instrument relevance); and

2. z_{kt} is uncorrelated with $\Delta \hat{\gamma}^{\star}_{kt}$ (exclusion restriction).

Our strategy will be to use long lags of labor demand shocks as an instrument for changes in relative wages. Consider the following data generating process for relative labor supply and demand shocks:

$$\Delta\hat{\alpha}^{\star}_{kt} = \sum_{\tau=1}^{T_{\alpha u}} \rho_{\alpha u,t-\tau} u_{\alpha t-\tau} + \sum_{\tau=1}^{T_{\alpha \nu}} \rho_{\alpha \nu,t-\tau} \nu_{\alpha t-\tau}, \tag{36}$$

$$\Delta\hat{\gamma}^{\star}_{kt} = \sum_{\tau=1}^{T_{\gamma u}} \rho_{\gamma u,t-\tau} u_{\gamma t-\tau} + \sum_{\tau=1}^{T_{\gamma \nu}} \rho_{\gamma \nu,t-\tau} \nu_{\gamma t-\tau}, \tag{37}$$

where the terms $u_{\alpha t-\tau}$ and $u_{\gamma t-\tau}$ are white noise errors, i.e., they are identically and independently distributed and uncorrelated with $\nu_{\alpha t-\tau}$, $\nu_{\gamma t-\tau}$, or each other, while $\nu_{\alpha t-\tau}$ and $\nu_{\gamma t-\tau}$ can be contemporaneously correlated.

Equations (36) and (37) describe a moving average process in which period t's relative supply and demand shocks are potentially correlated with past shocks and contemporaneous innovations on the other side of the market. Specifically, $\Delta\hat{\alpha}^{\star}_{kt}$ can show persistence for lags of up to $T_{\alpha u}$ periods and can be correlated with $\Delta\hat{\gamma}^{\star}_{kt}$ for up to $T_{\alpha \nu}$ periods; $\Delta\hat{\gamma}^{\star}_{kt}$ can show persistence for lags of up to $T_{\gamma u}$ periods and be correlated with $\Delta\hat{\alpha}^{\star}_{kt}$ for up to $T_{\gamma \nu}$ periods.

Our identifying assumptions are that

$$T_{\alpha u} > T_{\gamma \nu} \text{ and} \tag{38}$$

$$\exists\, T \text{ s.t. } T_{\gamma \nu} < T \leq T_{\alpha u} \text{ and } \rho_{\alpha u,t-T} \neq 0. \tag{39}$$

In words, these assumptions state that occupational labor demand shocks are more persistent than the shocks to labor supply that arise in response. Mathematically, equation (38) is our assumed exclusion restriction, and equation (39) is our instrument relevance assumption.

An example may clarify the economic interpretation of our identifying assumptions. If demand increases for computer programmers from 1970–1980, we might also expect there to have been positive demand shocks for computer programmers from 1980–1990 and 1990–2000. Equation (36) allows for this autocorrelation, but imposes that the autocorrelation

ends after $T_{\alpha u}$ periods. Imposing $T_{\alpha u} = 3$ requires the change in demand for programmers from 2010–2020 to be uncorrelated with the change from 1970–1980, i.e., $\Delta\hat{\alpha}^{\star}_{kt} \perp \Delta\hat{\alpha}^{\star}_{kt-4}$. Likewise, we might expect there to have been increases in the supply of computer programmers arising in response to the positive demand shock from 1970–1980. Our data generating process allows for such a response both contemporaneously and up to $T_{\gamma\nu}$ periods subsequently. Supposing that $T_{\gamma\nu} = 2$, our data generating process would allow for the supply curve for computer programmers to respond to the hypothetical 1970–1980 labor demand shock in the decades 1970–1980, 1980–1990, and 1990–2000, but not from 2000–2010, i.e., $\Delta\hat{\gamma}^{\star}_{kt} \perp \Delta\hat{\alpha}^{\star}_{kt-3}$.[23] Our instrument relevance assumption equation (39) in this example requires that labor demand shocks do persist for the third period, i.e., $\Delta\hat{\alpha}^{\star}_{kt} \not\perp \Delta\hat{\alpha}^{\star}_{kt-3}$.

Assuming $T_{\alpha u} = 3$ and $T_{\gamma\nu} = 2$ as in the example above, the long-lagged estimated labor demand shocks $\Delta\hat{\alpha}^{\star}_{kt-3}$ by construction serve as a valid instrumental variable for relative wage changes $\Delta\hat{w}^{\star}_{kt}$ in the equation

$$\Delta\hat{\ell}^{\star}_{kt} = \sigma_h \Delta\hat{w}^{\star}_{kt} + \varepsilon_{kt}. \tag{40}$$

We exploit this "long-lagged demand shock" strategy with $T_{\alpha u} = 3$ and $T_{\gamma\nu} = 2$ to estimate σ_h by two-stage least squares regression. We note that our framework to measure labor demand shocks is necessary for us to exploit this "long lags" strategy.

Table 4 presents the results of our instrumental variable estimation strategy for σ_h. Column 1 presents a baseline specification with no additional controls, which yields an estimated σ_h of approximately 1.9. Specifications which control explicitly for changes in female labor supply (column 2) and for more recent wage changes (column 3) imply broadly similar elasticities of 1.2 and 1.9. Replacing our instrument with the *fourth* lag of demand shocks (column 4) or weighting all occupations equally (column 5) yield elasticities of 2.9 and 4.0, respectively, but with weaker instruments than in the first three columns.

[23] Note that our assumed data generating process does allow supply shocks for computer programmers from 2000–2010 to be correlated with demand shocks for computer programmers that occur after 1970–1980, even though those demand shocks may be correlated with the demand shock from 1970–1980.

Table 4: Instrumental Variables Regressions for Relative Elasticity of Supply

	(1) Baseline	(2) Ctrl supply	(3) Ctrl lag wage chg	(4) $\Delta ln(\alpha_{t-4})$ as IV	(5) Unweighted
Second-Stage of $\Delta\hat{\ell}$ on $\Delta\hat{w}$					
Estimated σ_h	1.86***	1.17*	1.89***	2.94	4.03**
	[0.86,2.85]	[0.13,2.21]	[0.77,3.02]	[-1.43,7.31]	[1.02,7.04]
Change in female labor share		1.36***			
		[0.73,1.99]			
$\Delta\hat{w}_{t-1}$			-0.04		
			[-0.35,0.27]		
Constant	-0.02	-0.04*	-0.02	-0.03	-0.07**
	[-0.05,0.01]	[-0.06,-0.01]	[-0.05,0.01]	[-0.12,0.06]	[-0.11,-0.02]
First-Stage Instrumental Variables Regression of $\Delta\hat{W}$					
Lagged demand shock	0.16***	0.15***	0.15***	-0.04	0.09***
	[0.12,0.20]	[0.11,0.19]	[0.11,0.19]	[-0.08,0.00]	[0.05,0.14]
Change in female labor share		0.12			
		[-0.03,0.28]			
$\Delta\hat{w}_{t-1}$			0.09**		
			[0.02,0.16]		
Constant	0.01*	0.01	0.01*	0.02***	0.00
	[0.00,0.02]	[-0.00,0.01]	[0.00,0.01]	[0.01,0.03]	[-0.00,0.01]
N	574	574	574	448	574
adj. R^2	0.09	0.10	0.10	0.01	0.02
Fstat	60.9	48.9	48.8	3.4	14.6

95% confidence intervals in brackets. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Baseline approach uses the third lag of demand shocks $\Delta ln(\alpha_{t-3})$ as an instrument for current demand shocks. Estimation sample excludes decades for which the first-stage relationship between current and lagged demand shocks are negatively correlated (1980 and 2010). First-stage F-stat computed on excluded instruments. Except column (5) each observation weighted by its Sato-Vartia share of the wage bill.

Table 5: Estimated Elasticities of Relative Demand (σ_f) and Supply (σ_h)

	σ_f	σ_h
Estimate	8.15	1.86
95% CI	[7.2, 14.0]	[0.8, 3.8]

Note: estimates of σ_f derived using the Generalized Reverse Weighting procedure of Redding & Weinstein (2018). The third lag of the associated demand shocks are used as an instrument to estimate σ_h. Confidence intervals reflect bootstrapped standard errors over 100 replicate samples estimates. Shares are re-normalized and regressions re-estimated for each replicate.

5.3 Elasticity Estimates

Table 5 presents our baseline estimates for both elasticities of substitution σ_f and σ_h. The relative elasticity of labor demand implied by the Generalized Reverse Weighting estimator of 8.2, and as discussed above our instrumental variables approach implies an estimate of 1.9 for the relative elasticity of labor supply. These estimates suggest that it is easier for firms to shift production across occupations in response to changing relative wages than it is for workers to shift occupations, a result we consider intuitive. We discuss the robustness of our results to alternative estimation strategies in Section 6.2.

6 Estimated Wage Indices

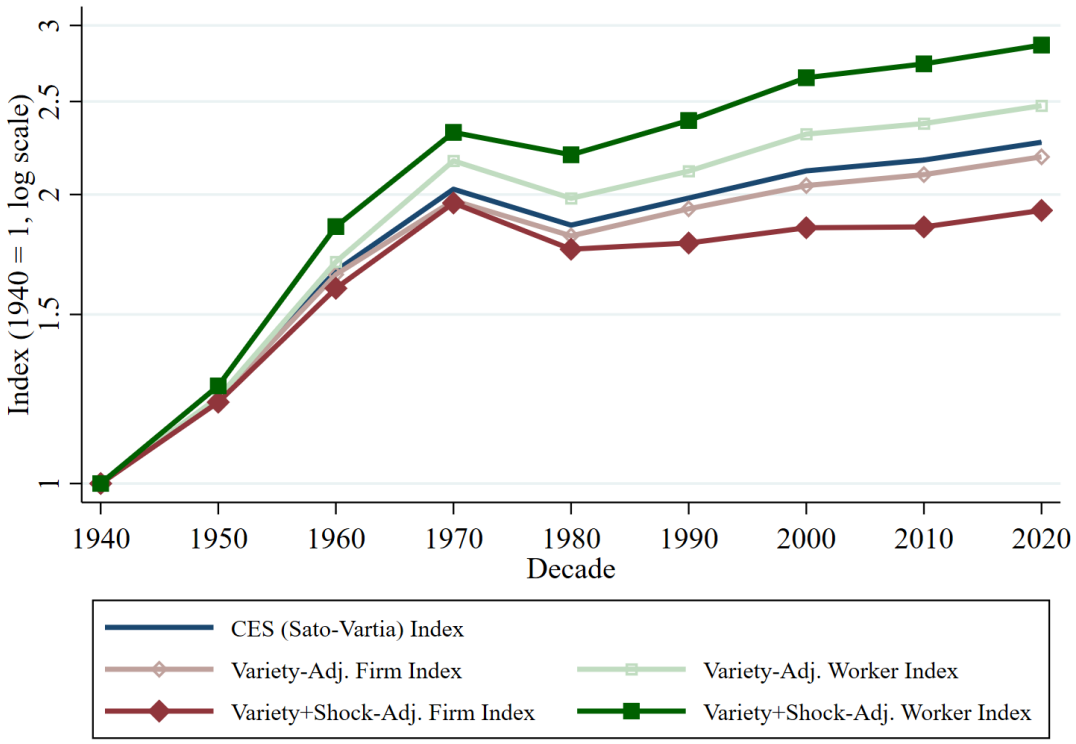
We begin this section by describing the results for our baseline wage indices in Section 6.1. In Section 6.2, we discuss the robustness of our results to alternative implementation choices.

6.1 Baseline Results

Figure 3 presents our estimated CES wage indices, all of which have been deflated by the Consumer Price Index and indexed to a value of one in 1940. The blue line in the figure shows the traditional Sato-Vartia index, which grew by a factor 2.3 from 1940 to 2020. That

cumulative growth masks an uneven path. The index approximately doubles from 1940–1970 (2.38 percent annually, on average) before declining 8.3 percent from 1970 to 1980. Growth slows thereafter, with the index rising 0.65 percent annually between 1980 and 2000 and just 0.34 percent annually between 2000 and 2020.

Figure 3: Estimated CES Wage Indices, 1940–2020



All indices deflated using Consumer Price Index (CPI). Employment and earnings are from decennial Census in 1940-2000 or 5-year ACS samples centered on 2009 and 2019.

As discussed in Section 2, the superlative Sato-Vartia index has important advantages over commonly used wage statistics but still faces limitations. It controls for composition effects by comparing wage changes within occupations, but this means it cannot evaluate the contributions of new or disappearing occupations. Moreover, while the Sato-Vartia index offers a principled account for substitution behaviors of workers and firms in response to changing relative wages, it does so assuming away changes in relative supply or demand. Yet those shifts in relative supply and demand must lie at the heart of any explanation of why relative wages change over time. Our CES unified wage indices account for these factors

in an interpretable manner derived from economic theory.

The light green line in Figure 3 displays our worker-side CES wage indices adjusted for the changing set of occupations in the labor market, while the light red line displays the corresponding firm-side index. They show that occupational turnover has benefited both households and firms over the past eighty years. We estimate that occupational turnover has increased worker-side effective real wages by 0.11 percentage points annually (9.2 percent cumulatively) relative to the Sato-Vartia index, while reducing them by 0.04 points annually (3.4 percent overall) on the employer side.[24]

The effect was strongest between 1940 and 1970: we estimate that it contributed 0.23 points annually (7 percent total) to worker wage growth over this period, during which time the number of occupations in our sample grew from 212 to 355. In a confirmation that this effect was not mechanically driven by Census enumerating a larger number of occupations, this same adjustment contributed to a *reduction* in worker wages from 1970-1980 of 0.04 percentage points annually even while the number of enumerated occupations grew again to 459. Occupational turnover contributed another 0.12 percentage points annually (2.7 percent overall) to worker wage growth between 1980 and 2000, but has had modestly negative effects thereafter. Our estimates, based on the relative wage bill shares of new and disappearing work, provide a unique quantification of the value of this process to workers and firms.

The dark green and dark red lines in Figure 3 show our estimated worker- and firm-side CES wage indices, respectively. These indices further adjust the Feenstra-adjusted Sato-Vartia indices in the light lines to account for shocks to relative labor supply in the worker-side index and shocks to relative labor demand in the firm-side index. The relative demand adjustment reduced firm-side real effective wage growth from 1940–2020 by an average of 0.16 percentage points annually (12.1 percent overall), while the relative supply adjustment boosted worker-side real effective wage growth by 0.18 percentage points annually on average (15.7 percent total). Overall, our CUWIs imply that workers' real effective wages grew by a

[24]The effects are quantitatively larger for workers than for firms because their lower elasticity of substitution implies that new alternatives bring greater value relative to the existing occupation mix.

factor of 2.9 (1.32 percent annually) from 1940 to 2020, while firms' real effective wages rose by only a factor of 1.9 (0.82 percent annually).

We examine the supply and demand shocks underlying these adjustments thoroughly in Section 7, but here we provide some intuition for their effects on the estimated wage indices. Consider first the firm-side index, for which the shock adjustment reduces wage growth by an average of 0.24 percentage points annually (11 percent total) after 1980. That adjustment captures two important trends in the relative demand parameters α_{kt}. First, consistent with narratives of demand-driven wage polarization, we find that the relative demands for different occupations became much more dispersed: we estimate the cross-sectional standard deviation of α_{kt} increased by 74% from 1980 to 2020. Analogously to the well-known result that consumers are better off when they face more price dispersion because it affords more opportunities for substitution, the downward adjustment in our firm-side index reflects firms' gains from their ability to substitute toward occupations with lower effective wages.

Second, as reflected in the last term of equation (15), the scale of the relative shock adjustment depends on whether occupations experiencing large relative demand shocks contribute significantly to the total wage bill. Consider the contributions from the example occupations shown in Table 6. A downward adjustment to the index implies that positive shifts in relative demand from high-share occupations such as store managers and accountants/auditors have greater influence on the effective cost of the labor composite than do negative shifts in small-share occupations such as file clerks and fabricating machine operators.

Table 6: Example Relative Demand Shock Contributions: 1980-1990

Occupation	Sato-Vartia Share	Demand Shock: $\Delta ln(\alpha_{kt})$	Supply Shock: $\Delta ln(1/\gamma_{kt})$	Contribution to $ln(\Phi^{CUWI,f}_{1980,1990})$
Store proprietors & managers	2.70%	0.236	0.241	0.0064
Accountants & auditors	1.62%	0.129	0.092	0.0021
Nursing aides	0.71%	0.136	0.017	-0.0010
File clerks	0.12%	-0.144	-0.053	-0.0002
Fabricating machine operators	0.02%	-0.350	-0.066	-0.0001

Similar logic helps to explain why our shock adjustment on the household side is upward. Table 7 presents some illustrative examples. A net positive adjustment to the worker wage index implies that positive relative supply changes in large-share occupations such as truck driving outweigh the contributions from relative supply reductions in small-share occupations such as paving equipment operators.

Our finding of a simultaneous upward shock adjustment to the worker index and downward shock adjustment to the firm index suggests an alignment of firm and worker responses to occupational change. Occupations which firms have increasingly wanted to employ are also the occupations where workers have shown an increasing willingness to work. This alignment is not guaranteed at either the occupational or aggregate levels. Truck drivers and kindergarten/pre-kindergarten teachers, for example, both saw large increases in relative supply between 1980 and 1990, even as relative demand fell in those occupations. Conversely, workers were relatively less willing to supply work as paving equipment operators despite a slight increase in relative demand. In aggregate, although demand shocks $\Delta ln(\alpha_{kt})$ and supply shocks $\Delta ln(\gamma_{kt})$ exhibited a positive correlation of 0.15 between 1970 and 2020, this correlation was slightly negative from 1940 through 1960.

Table 7: Example Relative Supply Shock Contributions: 1980-1990

Occupation	Sato-Vartia Share	Demand Shock: $\Delta ln(\alpha_{kt})$	Supply Shock: $\Delta ln(1/\gamma_{kt})$	Contribution to $ln(\Phi^{CUWI,w}_{1980,1990})$
Truck drivers	2.95%	-0.109	0.188	-0.0056
Computer programmers	0.66%	0.209	0.169	-0.0011
Legal assistants	0.13%	0.286	0.331	-0.0004
Pre-K & Kindergarten Teachers	0.06%	-0.075	0.227	-0.0001
Paving equipment operators	0.01%	0.042	-0.405	0.0000

6.2 Robustness to Alternative Approaches

Table 8 summarizes the robustness of our elasticity and wage index estimates to alternative empirical implementations. Its first row displays our baseline results, while the second row

shows results using the narrow continuer rule for linking occupations across time discussed in Section 4. The supply-side elasticity of substitution using the narrow continuer rule is slightly lower than in the baseline estimates, producing a slightly larger worker-side index adjustment, but the results are quantitatively similar to the baseline estimates. Using the reverse-weighting estimator for the demand-side elasticity of substitution in the third row produces slightly higher estimated elasticity estimates on both sides of the market, but the results are again similar to our baseline estimates.[25] The fourth row presents results using the Feenstra (1994) approach to estimating the elasticities of substitution, which assumes that demand shocks are orthogonal to wage changes. This approach suggests a much higher household-side elasticity of substitution and a correspondingly lower worker-side wage index adjustment. In fact, the worker-side elasticity is estimated to be higher than the firm-side elasticity using this approach, a result we regard as intuitively implausible and suggestive of the likely violation of the identifying orthogonality assumption in our setting.

The fifth row highlights the importance of estimating the elasticity of supply explicitly for our purposes. It presents results using our baseline approach to estimating σ_f but assuming that labor supply is perfectly inelastic across occupations, i.e. assuming $\sigma_h \to 0$. This specification corresponds to the assumption in some of the literature on rising wage dispersion that the supply of workers of a given type (e.g., college educated males) is fixed within a period. Applying that assumption to occupations, which are a finer-grained and more elastic grouping, implies an adjustment to the worker-side wage index about four times larger than in our baseline estimates.[26]

[25] Note that using an alternative estimator for σ_f also changes our estimate for σ_h because it changes the estimates of $\Delta\hat{\alpha}^{\star}_{kt-3}$ used in our instrumental variables estimation approach.

[26] The natural complement to this exercise, an assumption of perfect elasticity, is not reported but by construction assumes no worker welfare shock adjustment. We show in a later exercise that such an assumption implies more worker re-allocation across occupations than is supported by the data.

Table 8: Robustness to Alternative Specifications

	σ_f	95% CI	σ_h	95% CI	Firm Index Adj. 1940-2020	Worker Index Adj. 1940-2020
Baseline (GRW + IV) estimation	8.15	[7.2, 14.0]	1.86	[0.8, 3.8]	–15.1%	26.2%
Narrow continuer rule	8.05	[7.3, 14.2]	1.51	[0.4, 3.9]	–16.4%	34.7%
RW + IV estimation	9.56	[9.0, 14.1]	2.21	[0.7, 5.1]	–13.5%	26.3%
Feenstra estimation	6.95	[3.4, 10.5]	8.15	[5.3, 11.0]	–17.0%	3.9%
Perfectly inelastic supply	8.15	[7.2, 14.0]	0.00	-	–15.1%	113.7%
Excl. sm. occs. from continuers	12.06	[11.0, 15.8]	1.86	[1.0, 4.1]	–13.9%	25.0%
Shock normalization weighted by $s^{\star}_{t-1}$	8.17	[11.1, 15.4]	1.96	[1.4, 5.1]	–13.5%	21.5%

Note: Baseline estimates of σ_f derived using the Generalized Reverse Weighting procedure of Redding & Weinstein (2018). The third lag of the associated demand shocks are used as an instrument to estimate σ_h. Confidence intervals reflect bootstrapped standard errors over 100 replicate samples for all estimates except Feenstra, which is parametric. Index adjustments refer to the difference between UWI and Sato-Vartia indices in 2020, each deflated by CPI. The line excluding small occupations from continuers (Excl. sm. occs. from continuers) reallocates occupations with Sato-Vartia wage bill shares falling below the first percentile from the set of continuing occupations to the set of entrants/exiters in estimation and index construction. Specification with shock normalization weighted by s_{t-1} imposes that s_{t-1} weighted log changes in $ln(\alpha)$ and $ln(\gamma)$ average to 0.

The sixth and seventh rows of Table 8 examine two issues that have been discussed in the context of CES price indices that adjust for demand shocks. Row six presents results in which we exclude occupations that account for small shares of the wage bill from the set of continuing occupations.[27] Redding and Weinstein (2020) and Ehrlich et al. (2023) employ alternative approaches to excluding economically minor products from the set of continuing goods in CES price indices and find that doing so has a substantial effect. In our context, the estimated elasticities and index adjustments are similar to the baseline results, suggesting that this "small shares issue" is less problematic in our context.

Finally, row seven presents results from imposing an alternative normalization on the supply and demand shocks. Our baseline normalization assumes that the simple average log supply and demand shocks are zero across continuing occupations. Kurtzon (2022) notes that many other normalizations are possible, and shows that these alternative normalizations can substantially affect the resulting CES price indices. Following Redding and Weinstein (2020), we re-estimate our indices imposing the normalization that the $s^{\star}_{t\text{-}1}$-weighted log average supply and demand shocks are zero. The results are very similar to the baseline estimates. We interpret the results in rows six and seven as suggesting that our CES wage indices are less sensitive to some of the issues that have been prominent in the discussion of the CES Unified Price Index, possibly because we work with a vastly smaller number of occupations than the number of products available in retail datasets.

[27] Specifically, we exclude occupations with Sato-Vartia shares below the first percentile in each decade pair from our set of continuing occupations. Note that these occupations will still affect the calculated wage indices through their inclusion in the Feenstra adjustment or "lambda ratio" terms in equations (12) and (13). They will be excluded from the second and third terms that apply to continuing occupations, however, as well as from the estimation procedure for elasticities of substitution, which use continuing occupations. Example occupations include athletes in the decades 1940–1950 or etchers and engravers in 2010–2020, which are arguably less informative about substitution in the labor market than occupations accounting for larger shares of the wage bill.

7 Estimated Relative Supply and Demand Shocks

This section presents our estimated relative supply and demand shocks. Section 7.1 describes patterns recovered across broad groups of occupations from 1980 to 2020, and Section 7.2 relates these shifts in relative supply and demand to the popular task framework of Deming (2017). Both illustrate that the shocks we recover track well-known trends over the last 40 years, and the latter provides additional insights into the nature of our index adjustments.[28]

7.1 Summarizing the Estimated Shocks

In this section we describe the relative supply and demand shifts across occupations that we estimate occurred from 1980 to 2020. We summarize the data by aggregating all occupations in the midpoint decade of 2000, as well as linked occupations in adjacent decades, to two-digit 2000 Standard Occupational Classification code groups.[29] In Appendix D, we estimate relative supply and demand shifts from 1940 to 2020 using 1980 SOC groups and discuss the challenges of summarizing these changes across long horizons.

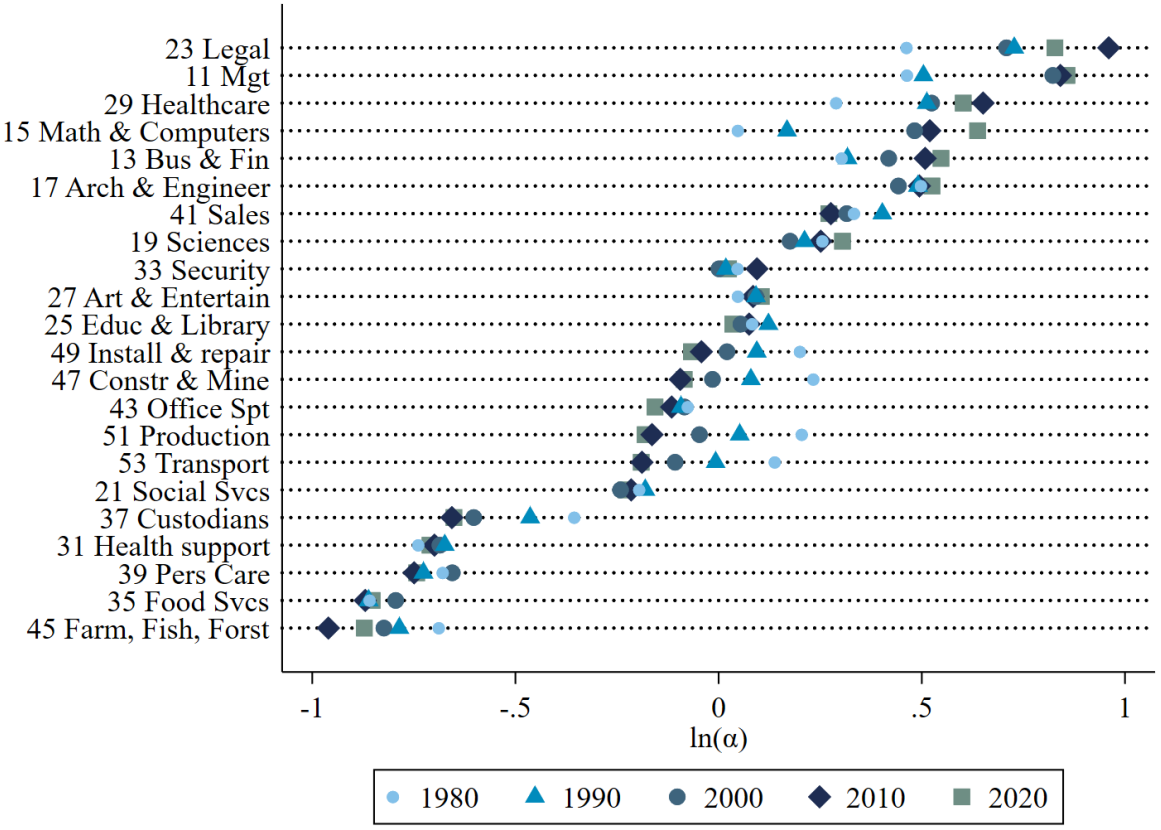
Figure 4 displays our estimated log relative demand shifters by decade.[30] Cross-sectionally, the occupational groups with the highest relative labor demand are the high-wage occupations of "Legal," "Management," "Healthcare," "Math and Computers," and "Business and Finance." Relative labor demand for occupations in these groups, which largely require non-routine analytic and/or interpersonal skills, has also risen significantly from 1980 to 2020. Sharp declines in relative demand can be seen in occupations more intensive in manual or routine tasks such as "Production," "Transportation," "Construction and Mining," "In-

[28] Appendix D.3 further validates these shocks by examining how the demand shocks relate to the trade exposure measures of Autor et al. (2013) and how supply shocks relate to changes in female labor supply.

[29] The U.S. Census Bureau provides a mapping of Census occupations to SOC codes in the crosswalk "Census 2000 Occupational Categories, With Standard Occupational Classification (SOC) Equivalents." The first two digits of a detailed SOC code identify its group, labels for which appear in Appendix B of Herman and Abraham (1999).

[30] As discussed in Section B.5, despite our normalizations that the average log changes in these parameters will be zero for continuing occupations, the averages across all occupations can drift over time with occupational entry and exit. We therefore normalize the estimated parameters by subtracting the simple log averages across all occupations present within a decade.

Figure 4: Estimated Relative Occupational Labor Demand, 1980–2020

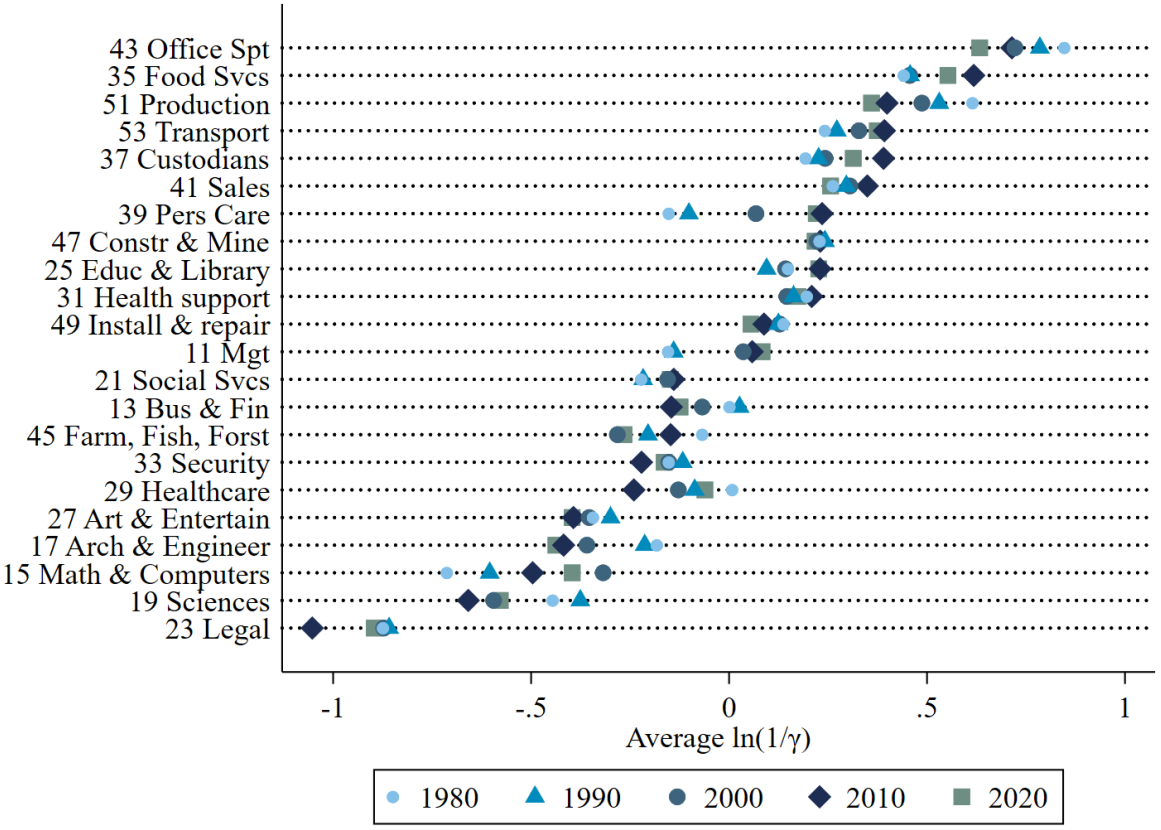


Marks represent (wage-bill weighted) averages across 3-digit Census occupations within given 2000 Standard Occupational Classification group. Our matching algorithm links continuing Census occupations in adjacent decades to 2000 groups.

stallation and Repair," and "Custodians." Several middle-demand service occupations saw comparatively stable relative demand over this period, including "Sales," "Security," "Arts and Entertainment," "Education and Library," and "Social Services." We view these patterns as generally consistent with the trends in routine-biased technological change documented, for instance, in Autor et al. (2003).

Figure 5 displays our estimated inverse log relative supply shifters by decade from 1980 to 2020; we display inverses so that higher numbers represent increases in relative labor supply. There are large cross-sectional differences across occupational groups in relative labor supply. The changes over time are also meaningful, although they are generally smaller than the shifts in relative demands. The "Food Services," "Office Support," and "Custodians" groups

Figure 5: Estimated Relative Occupational Labor Supply, 1980–2020



Marks represent (wage-bill weighted) averages across 3-digit Census occupations within given 2000 Standard Occupational Classification group. Our matching algorithm links continuing Census occupations in adjacent decades to 2000 groups.

have the highest relative supplies, while the "Math and Computers," "Sciences," and "Legal" groups have the lowest relative supplies. Relative supply has increased significantly in the "Custodians," "Personal Care," "Transportation," and "Math and Computers" groups. Occupations seeing large declines in relative supply include "Production," "Architecture and Engineering," "Office support," "Farming, Fishing, and Forestry," and "Sciences." Relative supply was roughly stable in most other groups, notably including those with large demand reductions like "Construction and Mining" or "Installation and Repair."

7.2 Relationship to O*NET Tasks

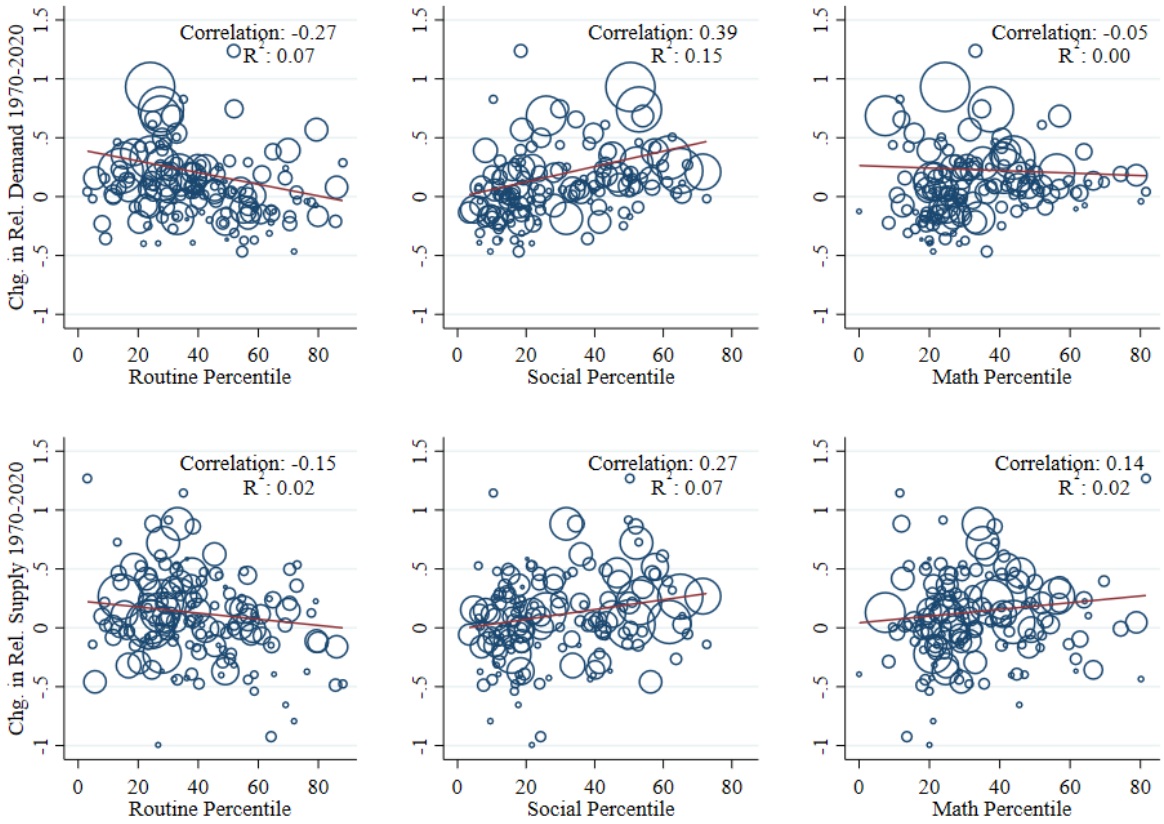
In this section, we validate and explore our estimated labor supply and demand shocks by relating them to occupational task content. Recent papers applying the task framework rely on O*NET, a survey conducted by the Bureau of Labor Statistics to study the work content of each occupation. Researchers use workers' answers to questions about their job to quantify how routine or social a job is, as well as how much analytic thought it requires. These quantifications are based on responses to questions about how often workers repeat the same activity, how often they need to be aware of and respond to others reactions, and whether the job requires knowledge and application of mathematics. An occupation's score in a given dimension is computed by averaging workers' reported intensities among relevant questions in the O*NET and scaling that average score by its decile across occupations.

Figure 6 summarizes the correlations between demand and supply shocks recovered by our model and the task content measures presented by Deming (2017).[31] The first row illustrates that our estimated demand shifts capture the key trends documented by Deming. Demand fell sharply among occupations with high concentrations of routine work. A key point Deming raises is that purely analytic occupations requiring little social interaction have seen slight decreases in demand, while occupations combining analytic and social skills have seen sharply increasing demand. The social and math panels in the top row of Figure 6 show that demand shocks in our model also reflect these patterns.

Workers in our model respond to higher wages in categories with rising demand by moving along an upward-sloping supply curve. At the same time, workers may shift their relative supply across occupations for other reasons, for example as features of the job make it relatively more or less appealing than other opportunities. The relative supply changes recovered by our model imply that, beyond the equilibrium responses to demand shocks, workers became relatively more willing to supply labor to occupations which were less routine and more social over this same period (i.e., the supply curves for these occupations shifted

[31] Since Deming provides decile scores computed to the first decimal, Figure 6 presents them as percentiles.

Figure 6: Correlations of Supply and Demand Shocks to Task Content

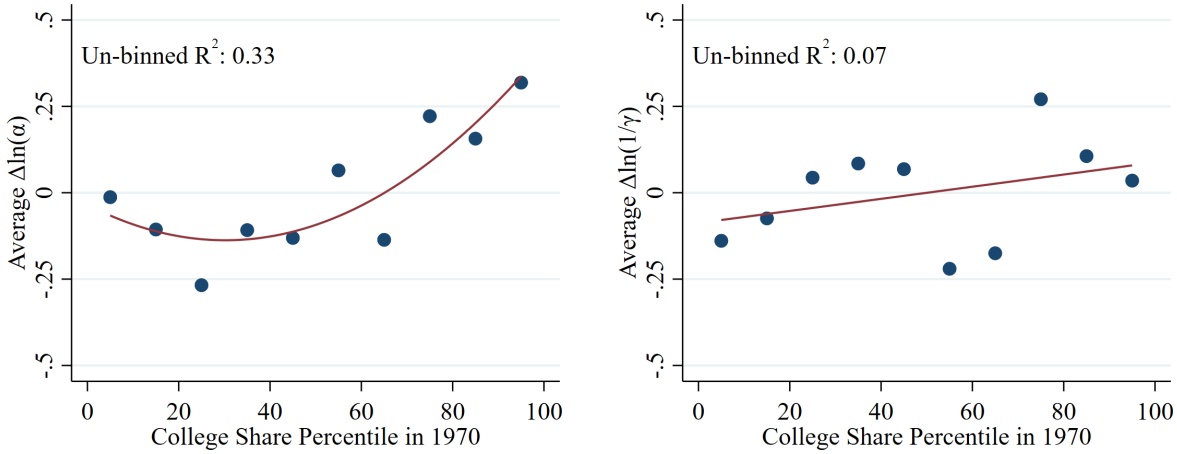


Note: Each observation represents a 1990 Census occupation weighted by its Sato-Vartia wage bill share. X-axes depict the occupation's percentile for a task type, from Deming (2017). These are correlated in the top row with $\ln(\alpha_{2020}) - \ln(\alpha_{1970})$ and in the bottom with $\ln(1/\gamma_{2020}) - \ln(1/\gamma_{1970})$.

outward). The positive shock adjustment in our household-side CUWI suggests workers incur less disutility supplying today's labor mix than they did in 1970, which may in turn reflect a growing preference for less routine, more social work.

Some portion of the supply adjustments seen in the bottom row of Figure 6 is also likely to reflect workers' increasing ability to perform the work for which there is growing demand. To evaluate the importance of this pattern, Figure 7 relates changes in demand and supply for a given occupation to the share of workers in that occupation with a college degree. The left panel shows that demand rose most quickly between 1970 and 2020 in occupations with high concentrations of college-educated workers. In fact, a quadratic function of 1970 college share can explain one-third of the variation in changing demand over this period.

Figure 7: Changes in Demand and Supply 1970-2020
by College Share Decile



Note: Each dot represents the average (weighted by wage bill share) among occupations within the given decile of 1970 college share. Adjusted-R^2 values from a quadratic regression of each occupation's demand or supply change from 1970-2020 on its 1970 college share.

If workers' supply shifts primarily reflected increasing educational investments in response to the rising demand for college-educated workers, then college shares should also explain a large portion of these supply changes. The right panel of Figure 7 demonstrates that this is not the case: college shares explain only 7% of changes in labor supply.[32]

Our results indicate that workers treat occupations as imperfect substitutes with distinct attributes. Their relative willingness to supply labor to these occupations varies over time, including for reasons that appear to be independent of demand and are not simply proxies for changing educational attainment. We interpret this evidence as indicating that it is important to model workers' supply preferences and decisions in a framework that can evaluate their welfare implications.

[32]Unreported results controlling for shares of female labor in each occupation do not change the relative importance of college in explaining demand and supply shocks. The patterns in Figure 7 are similar if we instead sort occupations by their 2020 college share percentiles.

8 Decomposing Changes in the Wage Structure

In this section, we revisit the literature on changes in the wage structure over the past fifty years through the lens of our CES model of the labor market. In Appendix E.3, we examine the "Great Compression" in wages from 1940 to 1950 (Goldin and Margo, 1992).

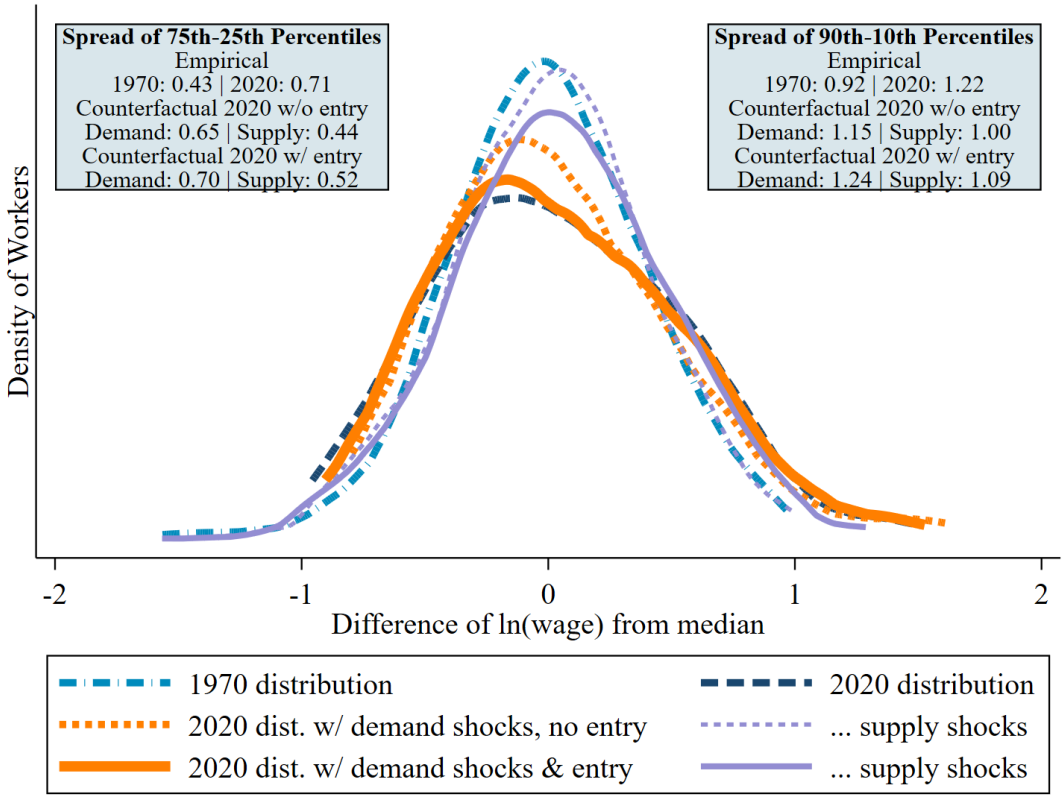
As we saw in in Section 3.3, given elasticities of substitution σ_f and σ_h we can compute relative supply and demand parameters γ_{kt} and α_{kt} and make straightforward calculations of relative wages $\hat{w}_{kt}$ and relative labor quantities $\hat{\ell}_{kt}$. These calculations in turn allow us to compute counterfactual wage distributions for any given constellation of relative supply and demand shifters. The counterfactuals we consider evaluate only demand or only supply changes in isolation. They compute the share of workers in each occupation and their relative wages in hypothetical scenarios where the demand shifts we estimate took place under stable supply, and conversely, in scenarios with shifting supply but stable demand.[33] We also examine the contributions of occupational entry by considering counterfactuals in which no occupations enter the labor market enter during the period under consideration.

Figure 8 presents a decomposition for 1970–2020. Each line summarizes an observed or counterfactual wage distribution as kernel-smoothed, employment-weighted densities of average occupational log wages relative to the median. The thick dashed light blue line shows the observed relative wage distribution in 1970, while the thick dashed dark blue line represents the 2020 wage distribution. It is well-known that individual wages have become more dispersed over the past 50 years (e.g., Autor et al., 2008); Figure 8 shows that this pattern also holds for occupational average wages. The interquartile range rose from 43 log points in 1970 to 71 log points in 2020, while the 90–10 spread increased from 92 to 122 log points.

The lavender lines in Figure 8 display counterfactual wage distributions which assume that the labor market faced the same relative supply shocks that we estimate for the actual

[33] Because we model only the average wage each occupation, these counterfactual wage distributions abstract from developments in wage dispersion within-occupation.

Figure 8: Decomposing Changes in the Wage Structure, 1970–2020



economy between 1970 and 2020, but the relative demand of each occupation remained at its initial level. For the thin dashed line, which excludes new occupations created since 1970, relative demands represent their 1970 values $\alpha_{k,70}$.[34] New occupations enter at their initial relative demands—which are subsequently held fixed—in the solid lavender line. Among the occupations that existed in 1970, supply shocks had effectively no influence on the interquartile range of wages. When including new work, changes in supply account for approximately one third of the increase in the interquartile range of wages across occupations. These results imply that supply shocks have done relatively little to push workers towards the extremes of the wage distribution, but that changes in supply among new occupations have contributed somewhat to growing wage inequality.

The orange lines in Figure 8 display the opposite counterfactual: the economy experiences

[34] By necessity, occupations exit in this counterfactual as they do in the data.

the relative demand shocks implied by our analysis, while relative supplies $\gamma_{k,t}$ remain fixed at their initial levels. Again, the dashed line represents an economy with only the occupations that existed in 1970. Among these occupations, shifts in relative demand explain about three quarters of the overall increase in the interquartile range of wages from 1970–2020. When new work is included, as shown in the solid orange line, nearly the entirety of the increase in the interquartile range can be explained by shifts in relative demand.

Occupational entry's effects in both scenarios are concentrated in the right-hand portion of the distribution. New work reduced the share of employment near the middle of the occupational wage distribution, raised it in high-wage occupations, and had little effect on the lower-wage occupations. These counterfactual results connect to our estimates in Figure 4 that relative labor demand declined from 1980 to 2020 in many traditionally blue-collar occupation groups while rising in many traditionally high-wage service occupation groups. In addition to shocks to the set of occupations that existed in 1970, new work also appears to have stretched the upper half of the wage distribution. We interpret these results as supporting the widely studied phenomenon of skill-biased technical change (e.g., Bound and Johnson, 1992; Katz and Murphy, 1992; Autor et al., 2008).

One of our important contributions to the literature on occupational wage dispersion is our treatment of the supply side of the labor market. Several seminal papers in this area (e.g., Bound and Johnson, 1992; Katz and Murphy, 1992; DiNardo et al., 1996) model workers of a given type (e.g., men with a high school diploma and 30-35 years of work experience) as being in fixed supply but allocating their labor elastically across individual sectors so that each worker of a given type earns a common wage across sectors. In Appendix E.2, we demonstrate that assumptions of fixed or perfectly elastic labor supplies generate starkly differing counterfactual changes in the occupational wage distribution, emphasizing the importance of disciplining such assumptions empirically.[35]

[35] Appendix Figures E.8 and E.9 display the results of our counterfactual decompositions under the limiting assumptions that $\sigma_h \to 0$ and $\sigma_h \to \infty$, respectively. The $\sigma_h \to 0$ exercise results in the opposite conclusion from our baseline: exogenous supply shocks account almost entirely for the rise in wage dispersion. The $\sigma_h \to \infty$ exercise requires a counterintuitive, strongly negative correlation of -0.85 between demand and

9 Conclusion

The nature of work has changed substantially over the past eighty years. The set of occupations has expanded dramatically, and there have been large shifts in supply and demand across occupations. We provide a novel evaluation of these developments' implications in the context of a general equilibrium model of the labor market. Building on recent developments in the price index literature, we propose theoretically exact household- and firm-side wage indices adjusted for these forces. To construct those indices, we also extend methods to identify exiting and entering occupations, estimate the elasticities of substitution across occupations, and measure shocks to relative occupational labor supply and demand. We show how to use our estimates to decompose changes in the wage distribution into changes rationalized by occupational turnover, relative supply shocks, and relative demand shocks.

Our CES Unified Wage Indices yield important new insights about how occupational change affects welfare. They imply that, on average, occupational turnover contributed 0.11 percentage points annually to household-side real wage growth from 1940 to 2020 and reduced firms' real wage increases by 0.04 points annually. Those effects were stronger from 1940–1970 than thereafter. Shifts in the relative supply and demand of occupations—arising from changing job characteristics, worker preferences, and firm production technologies—have also driven substantial welfare improvements. Relative supply shocks further increased household-side real wage growth by 0.18 percentage points annually, while relative demand shifts reduced firm-side wage growth by 0.16 percentage points annually on average from 1940–2020. Those adjustments are quantitatively significant relative to overall wage gains, which our traditional Sato-Vartia index estimates averaged 1.03 percent annually.

Finally, applying our methodology to assess the role of supply and demand shocks in wage polarization across occupations, we find that relative demand shocks to occupations that existed in 1970 and the entry of new occupations from 1970 to 2020 drove the increase in the dispersion of average occupational wages.

supply shocks to replicate the observed data.

References

N. Abe and D. P. Rao. On the normalization condition for cost of living comparisons under time-varying preferences. Technical Report DP24-1, January 2024. URL https://risk.ier.hit-u.ac.jp/English/2024/03/dp24-1_rcesr.pdf.

D. Acemoglu and P. Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6): 1488–1542, 2018.

D. Acemoglu, D. Autor, D. Dorn, G. H. Hanson, and B. Price. Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198, 2016.

S. P. Anderson, A. De Palma, and J.-F. Thisse. *Discrete choice theory of product differentiation*. MIT press, 1992.

D. Autor, F. Levy, and R. Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118:1279–1333, 2003.

D. Autor, D. Dorn, and G. Hanson. When work disappears: Manufacturing decline and the falling marriage market value of young men. *American Economic Review: Insights*, 1(2): 161–178, 2019.

D. Autor, C. Chin, A. M. Salomons, and B. Seegmiller. New frontiers: The origins and content of new work, 1940–2018. Working Paper 30389, National Bureau of Economic Research, August 2022.

D. H. Autor, L. F. Katz, and M. S. Kearney. Trends in us wage inequality: Revising the revisionists. *The Review of economics and statistics*, 90(2):300–323, 2008.

D. H. Autor, D. Dorn, and G. H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American economic review*, 103(6):2121–2168, 2013.

W. A. Barnett and K.-H. Choi. Operational identification of the complete class of superlative index numbers: An application of Galois theory. *Journal of Mathematical Economics*, 44 (7):603–612, 2008. doi: https://doi.org/10.1016/j.jmateco.2006.05.002. Special Issue in Economic Theory in honor of Charalambos D. Aliprantis.

A. B. Bernard, J. B. Jensen, and P. K. Schott. Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of us manufacturing plants. *Journal of international Economics*, 68(1):219–237, 2006.

M. Boskin, E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson. *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee*. U. S. Government Printing Office for the U. S. Senate Committee on Finance, Washington, DC, December 1996.

J. Bound and G. Johnson. Changes in the structure of wages in the 1980's: an evaluation of alternative explanations. *The American economic review*, 82(3):371–392, 1992.

Bureau of Labor Statistics. Introducing 2021 fixed employment weights and 2018 SOC codes for the Employment Cost Index, June 2023. URL https://www.bls.gov/eci/notices/2022/eci-2021-fixed-weights-and-2018-soc-update.htm. Accessed December 28, 2023.

Centers for Disease Control and Prevention. Achievements in public health, 1900-1999: Improvements in workplace safety – united states, 1900-1999. *Morbidity and Mortality Weekly Report*, 48(22):461–469, June 1999. URL https://www.cdc.gov/mmwr/preview/mmwrhtml/mm4822a1.htm.

J. G. Conzelmann, S. W. Hemelt, B. Hershbein, S. M. Martin, A. Simon, and K. M. Stange. Skills, majors, and jobs: Does higher education respond? Working Paper 31572, National Bureau of Economic Research, August 2023. URL http://www.nber.org/papers/w31572.

E. Crust, M. Daly, and B. Hobijn. The illusion of wage growth. *FRBSF Economic Letter*, 26:31, 2020.

M. Daly, B. Hobijn, and B. Lucking. Why has wage growth stayed strong? *FRBSF Economic Letter*, 10(2), 2012a.

M. C. Daly, B. Hobijn, and T. S. Wiles. Dissecting aggregate real wage fluctuations: Individual wage growth and the composition effect. Number 2011-23 in Federal Reserve Bank of San Francisco Working Paper Series. Federal Reserve Bank of San Francisco, 2012b.

D. J. Deming. The Growing Importance of Social Skills in the Labor Market*. *The Quarterly Journal of Economics*, 132(4):1593–1640, 06 2017. ISSN 0033-5533. doi: 10.1093/qje/qjx022. URL https://doi.org/10.1093/qje/qjx022.

W. E. Diewert. Superlative index numbers and consistency in aggregation. *Econometrica: Journal of the Econometric Society*, pages 883–900, 1978.

J. DiNardo, N. M. Fortin, and T. Lemieux. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2171954.

D. Dorn. *Data Appendix: Essays on Inequality, Spatial Interaction, and the Demand for Skills*. PhD thesis, University of St. Gallen, 2009.

A. Ebenstein, A. Harrison, M. McMillan, and S. Phillips. Estimating the impact of trade and offshoring on american workers using the current population surveys. *Review of Economics and Statistics*, 96(4):581–595, 2014.

G. Ehrlich, J. Haltiwanger, R. Jarmin, D. Johnson, and M. D. Shapiro. Re-engineering key national economic indicators. *Big Data for Twenty-First Century Economic Statistics*, 2022.

G. Ehrlich, J. C. Haltiwanger, R. S. Jarmin, D. Johnson, E. Olivares, L. W. Pardue, M. D. Shapiro, and L. Zhao. Quality adjustment at scale: Hedonic vs. exact demand-based price indices. Working Paper 31309, National Bureau of Economic Research, June 2023. URL http://www.nber.org/papers/w31309.

Federal Reserve Bank of Atlanta Center for Human Capital Studies. Wage growth tracker, 2024. URL https://www.atlantafed.org/chcs/wage-growth-tracker.

R. C. Feenstra. New Product Varieties and the Measurement of International Prices. *The American Economic Review*, 84(1):157–177, 1994.

C. Goldin and R. A. Margo. The great compression: The wage structure in the united states at mid-century. *The Quarterly Journal of Economics*, 107(1):1–34, 1992.

A. Herman and K. Abraham. Revising the standard occupational classification system. Technical Report Report 929, U.S. Department of Labor, Bureau of Labor Statistics, Washington, D.C., 1999.

K.-I. Inada. On a two-sector model of economic growth: Comments and a generalization. *The Review of Economic Studies*, 30(2):119–127, 1963.

International Monetary Fund. *Export and Import Price Index Manual: Theory and Practice*. Organisation for Economic Co-Operation and Develop, USA, 2009. ISBN 9781589067806. doi: 10.5089/9781589067806.069. URL https://www.elibrary.imf.org/view/book/9781589067806/9781589067806.xml.

W. S. Jevons. On the variation of prices and the value of the currency since 1782. *Journal of the Statistical Society of London*, 28(2):294–320, 1865.

L. F. Katz and K. M. Murphy. Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

G. Kurtzon. The Problem with Normalizing Preferences that Change in a Cost-of-Living Index. Technical report, June 2022. URL https://conference.nber.org/conf_papers/f168465.pdf.

J. Lin. Technological adaptation, cities, and new work. *The Review of Economics and Statistics*, 93(2):554–574, 2011. doi: 10.1162/REST_a_00079.

P. J. Lloyd. Substitution effects and biases in nontrue price indices. *The American Economic Review*, 65(3):301–313, 1975.

P. B. Meyer and A. M. Osborne. Proposed category system for 1960-2000 census occupations. Technical Report 383, U.S. Bureau of Labor Statistics, September 2005. URL https://www.bls.gov/osmr/research-papers/2005/ec050090.htm.

B. R. Moulton. Constant elasticity cost-of-living index in share relative form. *Bureau of Labor Statistics, Washington DC, December*, 1996.

E. Olivares. Derivation of product appeal in redding and weinstein (2018). Unpublished Note, 2018.

R. A. Pollak. *The Theory of the Cost-of-Living Index*. Oxford University Press, New York, 1989.

S. J. Redding and D. E. Weinstein. Measuring aggregate price indexes with demand shocks: Theory and evidence for CES preferences. Working Paper 22479 revision 2, National Bureau of Economic Research, August 2018. URL https://www.nber.org/system/files/working_papers/w22479/revisions/w22479.rev2.pdf.

S. J. Redding and D. E. Weinstein. Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences. *The Quarterly Journal of Economics*, 135(1): 503–560, 2020. doi: 10.1093/qje/qjz031.

S. J. Redding and D. E. Weinstein. Aggregation and estimation of constant elasticity of substitution (ces) preferences. *Princeton University, mimeograph*, 2023.

M. B. Reinsdorf and A. H. Dorfman. The sato–vartia index and the monotonicity axiom. *Journal of Econometrics*, 90(1):45–61, 1999.

S. Ruggles, S. Flood, M. Sobek, D. Backman, A. Chen, G. Cooper, S. Richards, R. Rogers, and M. Schouweiler. Ipums usa: Version 14.0 [dataset], 2023. URL https://doi.org/10.18128/D010.V14.0.

K. Sato. The Ideal Log-Change Index Number. *The Review of Economics and Statistics*, 1976. doi: 10.2307/1924029.

U.S. Census Bureau. *Detailed Occupation of Employed Persons, by Sex, for the United States: 1950 and 1940*, volume II, Characteristics of the Population, Part 1 of *U.S. Census of the Population: 1950*. 1953.

U.S. Census Bureau. *Changes Between the 1950 and 1960 Occupation and Industry Classifications–with Detailed Adjustments of 1950 Data to the 1960 Classifications*. John A. Priebe, Washington, D.C., July 1968. URL https://www.census.gov/content/dam/Census/library/publications/1968/demo/tech-18.pdf.

U.S. Census Bureau. *1970 Occupation and Industry Classification Systems in Terms of Their 1960 Occupation and Industry Elements*. John A. Priebe, Joan Heinkel, and Stanley Greene, Washington, D.C., 1972. URL https://www.census.gov/content/dam/Census/library/publications/1972/dec/1970-occupation-and-industry-classification-systems.pdf.

U.S. Census Bureau. *The Relationship Between the 1970 and 1980 Industry and Occupation Classification Systems*. Paula L. Vines and John A. Priebe, Washington, D.C., 1989. URL https://www2.census.gov/library/publications/1989/demographics/tech-paper-59.pdf.

U.S. Census Bureau. *Changes to the 1980 Industrial and Occupational Classification Systems for the 1990 Census of Populations*. Joanne Benetti and John Priebe, Washington, D.C., 1994. URL https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/changes-1980-ind-and-occ-classification.pdf.

U.S. Census Bureau. Census 2000 occupational categories, with standard occupational classification (soc) equivalents. https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/occ2000t.pdf, 2001. Accessed: 2022-03-30.

U.S. Census Bureau. *The Relationship Between the 1990 Census and Census 2000 Industry and Occupation Classification Systems*. Thomas S. Scopp, Washington, D.C., October 2003. URL https://www.census.gov/content/dam/Census/library/working-papers/2003/demo/techpaper2000.pdf.

U.S. Census Bureau. 2010 census occupation codes with crosswalk, November 2011. URL https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/2010-occ-codes-with-crosswalk-from-2002-2011.xls.

U.S. Census Bureau. *Recent Changes in the Census Industry and Occupation Classification Systems*. Julia B. Beckhusen, Washington, D.C., March 2020. URL https://www.census.gov/content/dam/Census/library/publications/2020/demo/acs-tp78.pdf.

Y. Vartia. Ideal Log-Change Index Numbers. *Scandinavian Journal of Statistics*, 3(3): 121–126, 1976.

Appendix

A Notation, Concepts, and Index Formulas

Table A.1: Summary of Notation, Concepts, and Formulas

Concept	All Occupations	Continuing Occupations
Occupational Wages	w_{kt}	$w^\star_{kt}$
Occupational Labor Quantities	ℓ_{kt}	$\ell^\star_{kt}$
Occupational Wage Bill Shares	s_{kt}	$s^\star_{kt}$
Geometric Mean of Var. x_{kt}	$\tilde{x}_t$	$\tilde{x}^\star_t$
Log Mean of Var. x_{kt}	$\overline{\ln x_t}$	$\overline{\ln x^\star_t}$
Log Dev. from Log Mean of Var. x_{kt}	$\hat{x}_{kt} = \ln x_{kt} - \overline{\ln x_t}$	$\hat{x}^\star_{kt} = \ln x^\star_{kt} - \overline{\ln x^\star_t}$

Concept	Firm Side	Household Side
Elasticities of Substitution	σ_f	σ_h
Supply and Demand Shifters	α_{kt}	γ_{kt}
Std. Dev. of Log Shocks	χ_α	χ_γ
CES Labor Composites (L_t)	$\left[\sum_{k\in\Omega}\left(\alpha_{kt}\ell_{kt}\right)^{\frac{\sigma_f-1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f-1}}$	$\left[\sum_{k\in\Omega}\left(\gamma_{kt}\ell_{kt}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}$
CES Unit Wage Indices (W_t)	$\left[\sum_{k\in\Omega}\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{1}{1-\sigma_f}}$	$\left[\sum_{k\in\Omega}\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}\right]^{\frac{1}{1+\sigma_h}}$

Empirical Wage Indices $\left(W_t/W_{t\text{-}1}\right)$	Firm Side	Household Side
Fixed Occupations, No Relative Shifts	$\prod_{k\in\Omega^\star}\left(\frac{w^\star_{kt}}{w^\star_{kt\text{-}1}}\right)^{\omega^\star_{kt}}$	
(Same for Firm and Household)	$\omega^\star_{kt} \equiv \frac{s^\star_{kt}-s^\star_{kt\text{-}1}}{\ln(s^\star_{kt})-\ln(s^\star_{kt\text{-}1})} \Big/ \left(\sum_{k\in\Omega^\star}\frac{s^\star_{kt}-s^\star_{kt\text{-}1}}{\ln(s^\star_{kt})-\ln(s^\star_{kt\text{-}1})}\right)$	
Changing Occ.s, No Relative Shifts	$\left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{1}{\sigma_f-1}}\prod_{k\in\Omega^\star}\left(\frac{w^\star_{kt}}{w^\star_{kt\text{-}1}}\right)^{\omega^\star_{kt}}$	$\left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h+1}}\prod_{k\in\Omega^\star}\left(\frac{w^\star_{kt}}{w^\star_{kt\text{-}1}}\right)^{\omega^\star_{kt}}$
	$\lambda_t = \sum_{k\in\Omega^\star} s_{kt}, \quad \lambda_{t\text{-}1} = \sum_{k\in\Omega^\star} s_{kt\text{-}1}$	
Fixed Occ.s, Relative Shifts	$\left(\frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t-1}}\right)\left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t-1}}\right)^{\frac{1}{\sigma_f-1}}$	$\left(\frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t-1}}\right)\left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t-1}}\right)^{\frac{-1}{\sigma_h+1}}$
Changing Occ.s, Relative Shifts	$\left(\frac{\lambda_t}{\lambda_{t-1}}\right)^{\frac{1}{\sigma_f-1}}\frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t-1}}\left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t-1}}\right)^{\frac{1}{\sigma_f-1}}$	$\left(\frac{\lambda_t}{\lambda_{t-1}}\right)^{\frac{-1}{\sigma_h+1}}\frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t-1}}\left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t-1}}\right)^{\frac{-1}{\sigma_h+1}}$

Note: the firm-side index with fixed occupations and no relative shifts corresponds to the Sato-Vartia price index (Sato, 1976; Vartia, 1976). The firm-side index with changing occupations and no relative shifts corresponds to the Feenstra (1994) index. The firm-side index with fixed occupations and relative shifts corresponds to Redding and Weinstein's (2020) "CES exact price index for common varieties," while the firm-side index with changing occupations and relative shifts corresponds to their "CES Unified Price Index" (CUPI).

B Model Details and Derivations

B.1 Firm-Side Wage Indices

In this section, we derive the theoretical firm-side CES exact wage index in equation (5) and the firm's relative labor demand condition in equation (17). Much of the discussion follows Redding and Weinstein (2020), with the difference that we solve the firm's problem as a two-step problem in which it first chooses its CES composite labor input L_t^f and then chooses occupational labor inputs ℓ_t^f.

As noted in equation (1), we assume that the representative firm produces output according to the function

$$Y_t = F(A_t L_t^f),$$

where as in equation (2), L_t^f is defined as

$$L_t^f \equiv \left[\sum_{k \in \Omega_t} (\alpha_{kt} \ell_{kt})^{\frac{\sigma_f - 1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f - 1}}.$$

First, the firm chooses composite labor input L_t^f to solve the problem

$$\max_{L_t^f} Y_t - W_t^f L_t^f, \tag{41}$$

which has first-order condition

$$\frac{\partial F(A_t L_t^f)}{\partial L_t^f} = W_t^f. \tag{42}$$

We assume that the function $F(A_t L_t^f)$ is strictly increasing and concave and follows the Inada (1963) conditions

$$\lim_{L_t^f \to 0} \frac{\partial F(A_t L_t^f)}{\partial L_t^f} = \infty \quad \text{and} \quad \lim_{L_t^f \to \infty} \frac{\partial F(A_t L_t^f)}{\partial L_t^f} = 0, \tag{43}$$

which guarantee an interior solution to equation (42). We denote the firm's optimal choice of L_t^f as L_t^{f*}. We can define the firm's wage bill as

$$Z_t \equiv \sum_{k \in \Omega_t} w_{kt} \ell_{kt} = W_t^f L_t^{f*}, \tag{44}$$

where we will verify below the second equality given our definition of W_t^f. The firm then solves a second maximization problem for its individual occupational demands ℓ_{kt} as

$$\max_{\ell_{kt}} L_t^f \text{ s.t. } \sum_{k \in \Omega_t} w_{kt} \ell_{kt} \leq Z_t. \tag{45}$$

We can write the firm's first-order condition for occupation k as

$$\ell_{kt} = {\alpha_{kt}}^{\sigma_f - 1} \left[\left(\frac{\sigma_f}{\sigma_f - 1} \right) \lambda^f w_{kt} \right]^{-\sigma_f}, \tag{46}$$

where λ^f is the Lagrange multiplier associated with the firm's second-step maximization problem. The first-order condition (46) must hold for every occupation k, so we can take the ratio of labor demands for any two occupations k and l as:

$$\frac{\ell_{kt}}{\ell_{lt}} = \left(\frac{\alpha_{kt}}{\alpha_{lt}} \right)^{\sigma_f - 1} \left(\frac{w_{kt}}{w_{lt}} \right)^{-\sigma_f} \Rightarrow \ell_{kt} = \ell_{lt} \left(\frac{\alpha_{kt}}{\alpha_{lt}} \right)^{\sigma_f - 1} \left(\frac{w_{kt}}{w_{lt}} \right)^{-\sigma_f}. \tag{47}$$

Note that, because this relationship holds for any two occupations k and l, it will also hold for occupation k and the geometric average across occupations:

$$\frac{\ell_{kt}}{\tilde{\ell}_t} = \left(\frac{\alpha_{kt}}{\tilde{\alpha}_t} \right)^{\sigma_f - 1} \left(\frac{w_{kt}}{\tilde{w}_t} \right)^{-\sigma_f}. \tag{48}$$

Taking logs and rearranging gives the condition

$$\hat{\ell}_{kt} = (\sigma_f - 1)\, \hat{\alpha}_{kt} - \sigma_f \hat{w}_{kt}, \tag{49}$$

which is the firm-side condition in the equilibrium relationships of equation (17).

Multiplying both sides of the right-hand condition in equation (47) by w_{kt} and summing over all occupations $k \in \Omega_t$ yields

$$\sum_{k \in \Omega_t} w_{kt} \ell_{kt} = \ell_{lt} \alpha_{lt}^{1-\sigma_f} w_{lt}^{\sigma_f} \sum_{k \in \Omega_t} \left(\frac{w_{kt}}{\alpha_{kt}} \right)^{1-\sigma_f}. \tag{50}$$

Now, we will "guess and check" that the firm's CES wage index W_t^f is given by

$$W_t^f = \left[\sum_{k \in \Omega_t} \left(\frac{w_{kt}}{\alpha_{kt}} \right)^{1-\sigma_f} \right]^{\frac{1}{1-\sigma_f}} \tag{51}$$

as asserted in equation (5). We will confirm that the firm's wage bill Z_t is equal to the hypothesized CES wage index W_t^f times its CES composite labor input L_t^f.

Combining equations (50) and (51) and recognizing that the left-hand side of (50) equals Z_t gives

$$Z_t = \ell_{lt} \alpha_{lt}^{1-\sigma_f} w_{lt}^{\sigma_f} {W_t^f}^{1-\sigma_f} \Rightarrow \ell_{lt} = \alpha_{lt}^{\sigma_f - 1} w_{lt}^{-\sigma_f} \frac{Z_t}{{W_t^f}^{1-\sigma_f}}. \tag{52}$$

Multiplying both sides by w_{lt} and dividing by Z_t gives us the following expression for occu-

pation l's share of the wage bill:

$$s_{lt} \equiv \frac{w_{lt}\ell_{lt}}{\sum_{k\in\Omega_t} w_{kt}\ell_{kt}} = \frac{\left(\frac{w_{lt}}{\alpha_{lt}}\right)^{1-\sigma_f}}{\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}. \tag{53}$$

Re-arranging the right-hand portion of (52) gives the following expression for the firm's labor demand:

$$\ell_{kt} = \frac{Z_t}{w_{kt}} \frac{\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}{{W_t^f}^{1-\sigma_f}}, \tag{54}$$

which we can substitute into the definition of L_t^f to yield

$$L_t^f \equiv \left[\sum_{k\in\Omega_t} \left(\alpha_{kt}\ell_{kt}\right)^{\frac{\sigma_f-1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f-1}} = \left[\sum_{k\in\Omega_t}\left(\alpha_{kt}\frac{Z_t}{w_{kt}}\frac{\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}}{{W_t^f}^{1-\sigma_f}}\right)^{\frac{\sigma_f-1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f-1}}. \tag{55}$$

We can simplify equation (55) successively as follows:

$$L_t^f = Z_t\left[\sum_{k\in\Omega_t}\left(\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{-\sigma_f}{W_t^f}^{1-\sigma_f}\right)^{\frac{\sigma_f-1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f-1}} = Z_t\left(W_t^f\right)^{\sigma_f-1}\left[\sum_{k\in\Omega_t}\left(\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{-\sigma_f}\right)^{\frac{\sigma_f-1}{\sigma_f}}\right]^{\frac{\sigma_f}{\sigma_f-1}}$$

$$= Z_t\left(W_t^f\right)^{\sigma_f-1}\underbrace{\left[\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\alpha_{kt}}\right)^{1-\sigma_f}\right]^{\frac{\sigma_f}{\sigma_f-1}}}_{\left(W_t^f\right)^{-\sigma_f}} = Z_t{W_t^f}^{\sigma_f-1}\left(W_t^f\right)^{-\sigma_f} = \frac{Z_t}{W_t^f}$$

$$\Rightarrow Z_t = W_t^f L_t^f,$$

which was what we wanted to show. We have now verified that the firm's hypothesized labor cost index is correct and derived the firm's labor demand condition in equation (17). Finally, we note that the equation for each occupation's share of the wage bill in (53) corresponds to the consumer's share of expenditure on each product in equation (2) of Redding and Weinstein (2020).

These results suffice to show that the firm-side CES wage indices presented in Section 3.2 must follow the forms in Redding and Weinstein (2020) and, under the relevant conditions, the preceding literature (Sato, 1976; Vartia, 1976; Feenstra, 1994). In the interest of space, we do not discuss any further derivations for the firm-side indices here.

B.2 Household-Side Wage Indices

In this section, we derive the theoretical household-side CES exact wage index in equation (5) and the household's relative labor supply condition in equation (17). Much of the development again parallels Redding and Weinstein (2020) and the firm-side derivations above, but because the household's labor supply problem takes a different form than the firm's problem, the resulting conditions and wage indices also differ. We therefore also derive the household-side CUWI in equation (13) and show that it collapses to the Sato-Vartia index in the absence of turnover in the set of occupations and shocks to relative labor supply.

In Section 3, we specified the household's utility function as

$$U\left(C_t, L_t^h\right) = G\left(C_t\right) - H\left(\Gamma_t L_t^h\right),$$

where

$$L_t^h \equiv \left[\sum_{k \in \Omega_t} \left(\gamma_{kt} \ell_{kt}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}.$$

We also specified the household's CES exact wage index as

$$W_t^h = \left[\sum_{k \in \Omega_t} \left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}\right]^{\frac{1}{1+\sigma_h}}$$

and the household's budget constraint as

$$C_t = L_t^h W_t^h + \Pi_t.$$

We solve the household's maximization problem as a two-step problem with the first step given by

$$\max_{C_t, L_t^h} G\left(C_t\right) - H\left(\Gamma_t L_t^h\right) \quad \text{s.t.} \quad C_t = L_t^h W_t^h + \Pi_t. \tag{56}$$

Denoting the Lagrange multiplier associated with the household's budget constraint as λ^h, this problem has first-order conditions:

$$C_t : \frac{\partial G\left(C_t\right)}{\partial C_t} = \lambda^h \tag{57}$$

$$L_t^h : \frac{\partial H\left(L_t^h\right)}{\partial L_t^h} = \lambda^h W_t^h \tag{58}$$

$$\lambda^h : C_t = L_t^h W_t^h + \Pi_t. \tag{59}$$

Again in the spirit of Inada (1963), we assume that the function $G(\cdot)$ is monotonically increasing and concave with

$$\lim_{C_t \to 0} \frac{\partial G\left(C_t\right)}{\partial C_t} = \infty \quad \text{and} \quad \lim_{C_t \to \infty} \frac{\partial G\left(C_t\right)}{\partial C_t} = 0, \tag{60}$$

while the function $H(\cdot)$ is monotonically increasing and convex with

$$\lim_{L_t^h \to 0} \frac{\partial H\left(L_t^h\right)}{\partial L_t^h} = 0 \quad \text{and} \quad \lim_{L_t^h \to \infty} \frac{\partial H\left(L_t^h\right)}{\partial L_t^h} = \infty. \tag{61}$$

These conditions ensure interior solutions C_t^* and L_t^{h*}. We denote the household's labor income as

$$Z_t \equiv \sum_{k \in \Omega_t} w_{kt} \ell_{kt} = W_t^h L_h^{f*}, \tag{62}$$

The household's second-stage problem is then given by

$$\min_{\ell_{kt}} L_t^h \text{ s.t. } \sum_{k \in \Omega_t} w_{kt} \ell_{kt} \geq Z_t. \tag{63}$$

The household's first-order condition for occupation k can be written as

$$\frac{\sigma_h + 1}{\sigma_h} \gamma_{kt} \left(\gamma_{kt} \ell_{kt}\right)^{\frac{1}{\sigma_h}} = \lambda_k^h \, w_{kt}, \tag{64}$$

where λ_k^h is the Lagrange multiplier associated with the household's labor income constraint. The first-order condition (64) must hold for every occupation k, so we can take the ratio of labor supplies for any two occupations k and l as:

$$\frac{\ell_{kt}}{\ell_{lt}} = \left(\frac{\gamma_{kt}}{\gamma_{lt}}\right)^{-(1+\sigma_h)} \left(\frac{w_{kt}}{w_{lt}}\right)^{\sigma_h} \Rightarrow \ell_{kt} = \ell_{lt} \left(\frac{\gamma_{kt}}{\gamma_{lt}}\right)^{-(1+\sigma_h)} \left(\frac{w_{kt}}{w_{lt}}\right)^{\sigma_h}. \tag{65}$$

Again note that, because this relationship holds for any two occupations k and l, it will also hold for occupation k and the geometric average across occupations:

$$\frac{\ell_{kt}}{\tilde{\ell}_t} = \left(\frac{\gamma_{kt}}{\tilde{\gamma}_t}\right)^{-(\sigma_h+1)} \left(\frac{w_{kt}}{\tilde{w}_t}\right)^{\sigma_h}. \tag{66}$$

Taking logs and rearranging gives the condition

$$\hat{\ell}_{kt} = -\left(1 + \sigma_h\right) \hat{\gamma}_{kt} + \sigma_h \hat{w}_{kt}, \tag{67}$$

which is the household-side condition in the equilibrium relationships of equation (17).

Multiplying both sides of the right-hand condition in equation (65) by w_{kt} and summing over all occupations $k \in \Omega_t$ yields

$$\sum_{k \in \Omega_t} w_{kt} \ell_{kt} = \ell_{lt} \gamma_{lt}^{-(1+\sigma_h)} w_{lt}^{-\sigma_h} \sum_{k \in \Omega_t} \left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}. \tag{68}$$

Now, we will "guess and check" that the household's CES wage index W_t^h is given by

$$W_t^h = \left[\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}\right]^{\frac{1}{1+\sigma_h}} \tag{69}$$

as asserted in equation (5). We will confirm that the household's wage income Z_t is equal to the hypothesized CES wage index W_t^h times its CES composite labor supplied L_t^h.

Combining equations (68) and (69) and recognizing that the left-hand side of (68) equals Z_t gives

$$Z_t = \ell_{lt}\gamma_{lt}^{-(1+\sigma_h)}w_{lt}^{-\sigma_h}\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h} \Rightarrow \ell_{lt} = \gamma_{lt}^{1+\sigma_h}w_{lt}^{\sigma_h}\frac{Z_t}{(W_t^h)^{1+\sigma_h}}. \tag{70}$$

Multiplying both sides by w_{lt} and dividing by Z_t gives us the following expression for occupation l's share of wage income:

$$s_{lt} \equiv \frac{w_{lt}\ell_{lt}}{\sum_{k\in\Omega_t}w_{kt}\ell_{kt}} = \frac{\left(\frac{w_{lt}}{\gamma_{lt}}\right)^{1+\sigma_h}}{\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}. \tag{71}$$

Re-arranging the right-hand portion of (70) gives the following expression for the household's labor supply:

$$\ell_{kt} = \frac{Z_t}{w_{kt}}\frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{{W_t^h}^{1+\sigma_h}}, \tag{72}$$

which we can substitute into the definition of L_t^h to yield

$$L_t^h \equiv \left[\sum_{k\in\Omega_t}\left(\gamma_{kt}\ell_{kt}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}} = \left[\sum_{k\in\Omega_t}\left(\gamma_{kt}\frac{Z_t}{w_{kt}}\frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{{W_t^w}^{1+\sigma_h}}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}. \tag{73}$$

We can simplify equation (73) successively as follows:

$$L_t^h = Z_t\left[\sum_{k\in\Omega_t}\left(\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{\sigma_h}W_t^{h\,-(1+\sigma_h)}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}} = Z_tW_t^{h\,-(1+\sigma_h)}\left[\sum_{k\in\Omega_t}\left(\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{\sigma_h}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}$$

$$= Z_tW_t^{h\,-(1+\sigma_h)}\underbrace{\left[\sum_{k\in\Omega_t}\left(\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{\sigma_h}\right)^{\frac{\sigma_h+1}{\sigma_h}}\right]^{\frac{\sigma_h}{\sigma_h+1}}}_{(W_t^h)^{\sigma_h}} = Z_t(W_t^h)^{-(1+\sigma_h)}(W_t^h)^{\sigma_h} = \frac{Z_t}{W_t^h}$$

$$\Rightarrow Z_t = W_t^hL_t^h,$$

which was what we wanted to show. We have now verified that the household's hypothesized

theoretical CES wage index is correct and have derived the household's labor supply condition in equation (17).

We now derive the household-side CUWI in equation (13), where we again note that our discussion parallels Redding and Weinstein (2020), with details differing because we focus on the supply side rather than the demand side. We start by combining the definitions of λ_t and $\lambda_{t\text{-}1}$ in equation (9) with our equation for occupation k's share of wage income in equation (71) to write

$$\lambda_t = \sum_{k \in \Omega^\star_{t\text{-}1,t}} \frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{{W_t^h}^{1+\sigma_h}}, \qquad \lambda_{t\text{-}1} = \sum_{k \in \Omega^\star_{t\text{-}1,t}} \frac{\left(\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}\right)^{1+\sigma_h}}{{W_{t\text{-}1}^h}^{1+\sigma_h}}. \tag{74}$$

Combining our relationships from equation (10) that $s^\star_{kt} = \frac{s_{kt}}{\lambda_t}$ and $s^\star_{kt\text{-}1} = \frac{s_{kt\text{-}1}}{\lambda_{t\text{-}1}}$, with the relationships in equation (74) and re-arranging to solve for W_t^h gives

$$W_t^h = \frac{w_{kt}}{\gamma_{kt}} \left(\lambda_t s^\star_{kt}\right)^{\frac{-1}{1+\sigma_h}} \quad \text{and} \quad W_{t\text{-}1}^w = \frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}} \left(\lambda_{t\text{-}1} s^\star_{kt\text{-}1}\right)^{\frac{-1}{1+\sigma_h}}. \tag{75}$$

We now take the ratio of these theoretical CES worker-side wage indices to measure the change in wages from t-1 to t:

$$\frac{W_t^h}{W_{t\text{-}1}^h} = \left(\frac{\frac{w_{kt}}{\gamma_{kt}}}{\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}}\right) \left(\frac{\lambda_t}{\lambda_{t\text{-}1}} \frac{s^\star_{kt}}{s^\star_{kt\text{-}1}}\right)^{\frac{-1}{1+\sigma_h}} \tag{76}$$

This relationship will hold for every continuing occupation $k \in \Omega^\star_{t\text{-}1,t}$, so it will also hold for their geometric average:

$$\frac{W_t^h}{W_{t\text{-}1}^h} = \left(\frac{\frac{\widetilde{w}^\star_t}{\widetilde{\gamma}^\star_t}}{\frac{\widetilde{w}^\star_{t\text{-}1}}{\widetilde{\gamma}^\star_{t\text{-}1}}}\right) \left(\frac{\lambda_t}{\lambda_{t\text{-}1}} \frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t\text{-}1}}\right)^{\frac{-1}{1+\sigma_h}}. \tag{77}$$

Following Redding and Weinstein (2020), we now impose our normalization in equation (11) that the geomeans of the relative labor supply parameters across continuing occupations are equal between periods t-1 and t, i.e., $\widetilde{\gamma}^\star_t = \widetilde{\gamma}^\star_{t\text{-}1}$. Then equation (77) becomes

$$\frac{W_t^h}{W_{t\text{-}1}^h} = \left(\frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t\text{-}1}}\right) \left(\frac{\lambda_{t\text{-}1}}{\lambda_t} \frac{\widetilde{s}^\star_{t\text{-}1}}{\widetilde{s}^\star_t}\right)^{\frac{1}{1+\sigma_h}} = \left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h+1}} \frac{\widetilde{w}^\star_t}{\widetilde{w}^\star_{t\text{-}1}} \left(\frac{\widetilde{s}^\star_t}{\widetilde{s}^\star_{t\text{-}1}}\right)^{\frac{-1}{\sigma_h+1}}, \tag{78}$$

which is the form for $\Phi^{\text{CUWI},h}_{t\text{-}1,t}$ in equation (13).

Our final objective in this section is to show that the worker-side CUWI collapses to the Sato-Vartia index in the absence of turnover in the set of occupations Ω_t and shocks to relative supply. Our discussion will again parallel the development in Redding and Weinstein (2020), with minor differences because we focus on the supply side rather than the demand side of the market. We start by recalling our expression for an individual occupation's share of the wage bill in (71), expressing the denominator in terms of the theoretical CES wage

index W_t^h, and noting that it will also apply in period t-1:

$$s_{kt} = \frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{\sum_{k\in\Omega_t}\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}} = \frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{\left(W_t^h\right)^{1+\sigma_h}}, \quad s_{kt\text{-}1} = \frac{\left(\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}\right)^{1+\sigma_h}}{\sum_{k\in\Omega_{t\text{-}1}}\left(\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}\right)^{1+\sigma_h}} = \frac{\left(\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}\right)^{1+\sigma_h}}{\left(W_{t\text{-}1}^h\right)^{1+\sigma_h}} \tag{79}$$

Taking the ratios of the two wage income shares gives

$$\frac{s_{kt}}{s_{kt\text{-}1}} = \left(\frac{W_t^h}{W_{t\text{-}1}^h}\right)^{-(1+\sigma_h)}\left(\frac{\frac{w_{kt}}{\gamma_{kt}}}{\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}}\right)^{1+\sigma_h} \Rightarrow \frac{W_t^h}{W_{t\text{-}1}^h} = \frac{\left(w_{kt}/\gamma_{kt}\right)/\left(w_{kt\text{-}1}/\gamma_{kt\text{-}1}\right)}{\left(\frac{s_{kt}}{s_{kt\text{-}1}}\right)^{\frac{1}{1+\sigma_h}}}. \tag{80}$$

Taking logs and re-arranging gives

$$\frac{\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) - \ln\left(\frac{\frac{w_{kt}}{\gamma_{kt}}}{\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}}\right)}{\ln\left(\frac{s_{kt}}{s_{kt\text{-}1}}\right)} = \frac{-1}{1+\sigma_h}. \tag{81}$$

We now multiply both sides by $s_{kt} - s_{kt\text{-}1}$ and sum across all occupations to write

$$\sum_{k\in\Omega_t^\star}\left(s_{kt} - s_{kt\text{-}1}\right)\left[\frac{\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) - \ln\left(\frac{\frac{w_{kt}}{\gamma_{kt}}}{\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}}\right)}{\ln\left(\frac{s_{kt}}{s_{kt\text{-}1}}\right)}\right] = \underbrace{\sum_{k\in\Omega_t^\star}\left(s_{kt} - s_{kt\text{-}1}\right)\frac{-1}{1+\sigma_h}}_{=0}, \tag{82}$$

where we note that the right-hand side must equal zero because the wage shares in each period must sum to one. Then we can re-arrange the left-hand side to write

$$\sum_{k\in\Omega_t^\star}\frac{s_{kt} - s_{kt\text{-}1}}{\ln s_{kt} - \ln s_{kt\text{-}1}}\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) = \sum_{k\in\Omega_t^\star}\frac{s_{kt} - s_{kt\text{-}1}}{\ln s_{kt} - \ln s_{kt\text{-}1}}\ln\left(\frac{\frac{w_{kt}}{\gamma_{kt}}}{\frac{w_{kt\text{-}1}}{\gamma_{kt\text{-}1}}}\right), \tag{83}$$

which implies

$$\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) = \sum_{k\in\Omega_t^\star}\frac{\frac{s_{kt}-s_{kt\text{-}1}}{\ln s_{kt}-\ln s_{kt\text{-}1}}}{\sum_{l\in\Omega_t}\frac{s_{lt}-s_{lt\text{-}1}}{\ln s_{lt}-\ln s_{lt\text{-}1}}}\ln\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right) - \sum_{k\in\Omega_t^\star}\frac{\frac{s_{kt}-s_{kt\text{-}1}}{\ln s_{kt}-\ln s_{kt\text{-}1}}}{\sum_{l\in\Omega_t}\frac{s_{lt}-s_{lt\text{-}1}}{\ln s_{lt}-\ln s_{lt\text{-}1}}}\ln\left(\frac{\gamma_{kt}}{\gamma_{kt\text{-}1}}\right). \tag{84}$$

Recalling from equation (14) that the weights in the Sato-Vartia index are defined as

$$\omega_{kt}^\star \equiv \frac{s_{kt}^\star - s_{kt\text{-}1}^\star}{\ln(s_{kt}^\star) - \ln(s_{kt\text{-}1}^\star)} \Bigg/ \left(\sum_{k\in\Omega_{t\text{-}1,t}^\star}\frac{s_{kt}^\star - s_{kt\text{-}1}^\star}{\ln(s_{kt}^\star) - \ln(s_{kt\text{-}1}^\star)}\right),$$

we see that we can write the household's CES wage index as

$$\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) = \sum_{k\in\Omega_t^\star} \omega_{kt}^\star \ln\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right) - \sum_{k\in\Omega_t} \omega_{kt}^\star \ln\left(\frac{\gamma_{kt}}{\gamma_{kt\text{-}1}}\right). \tag{85}$$

It is straightforward to see that under the assumption of no time variation in the preference parameters, i.e. $\gamma_{kt} = \gamma_{kt-1} = \gamma_k \; \forall \; k$, the second term in equation (85) reduces to zero. Then the household's CES wage index can be written as

$$\ln\left(\frac{W_t^h}{W_{t\text{-}1}^h}\right) = \sum_{k\in\Omega_t^\star} \omega_{kt}^\star \ln\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right), \tag{86}$$

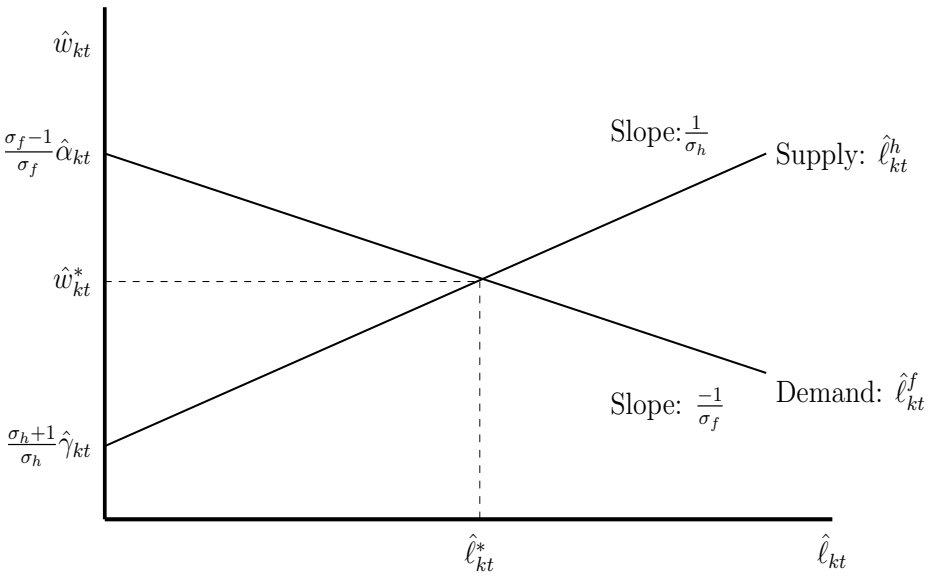
which is the Sato-Vartia index as seen in equation (14). From this result we see that, in the absence of shocks to relative labor supply and turnover in the set of occupations in the labor market, the household-side CES unified wage index defined in equation (13) collapses to the Sato-Vartia index and also equals the firm-side index under those conditions. From the above discussion, it is also straightforward to see that the household-side CES wage index in the presence of occupational turnover but in the absence of relative supply shocks can be expressed as

$$\ln \Phi_{t\text{-}1,t}^{\text{CUWI},h} = \left(\frac{-1}{\sigma_h + 1}\right) \ln\left(\frac{\lambda_t}{\lambda_{t\text{-}1}}\right) + \sum_{k\in\Omega_{t\text{-}1,t}^\star} \omega_{kt}^\star \ln\left(\frac{w_{kt}^\star}{w_{kt\text{-}1}^\star}\right). \tag{87}$$

We have therefore achieved our objectives for this section.

B.3 Graphical Representation of Relative Supply and Demand

Figure B.1: Graphical Equilibrium in the Labor Market



B.4 Recovering Relative Supply and Demand Shifters

In this section, we derive our expressions for the relative occupational demand and supply shifters α_{kt} and γ_{kt} in equations (20) and (21) of Section 3.3. We have already seen in Appendix B.1 that the functional forms of the firm (labor demand) side of our model correspond precisely to the forms for the consumer demand problem in Redding and Weinstein (2020). It thus follows that the expression for $\alpha^{\star}_{kt}$ must equal the expression derived in Redding and Weinstein (2020).

We therefore begin this section by deriving our expression for γ_{kt} in equation (21). In equation (71), we showed that any occupation k's share of the wage bill can be expressed as

$$s_{kt} = \frac{\left(\frac{w_{kt}}{\gamma_{kt}}\right)^{1+\sigma_h}}{\sum_{l\in\Omega_t}\left(\frac{w_{lt}}{\gamma_{lt}}\right)^{1+\sigma_h}},$$

where we have switched subscripts for notational convenience. Re-arranging this equation and expressing it in terms of continuing occupations gives

$$\gamma_{kt} = \frac{w_{kt}}{\sum_{l\in\Omega_t}\left(\frac{w_{lt}}{\gamma_{lt}}\right)^{1+\sigma_h}} s_{kt}^{\frac{-1}{1+\sigma_h}} \Rightarrow \gamma^{\star}_{kt} = \frac{w^{\star}_{kt}}{\sum_{l\in\Omega^{\star}_t}\left(\frac{w^{\star}_{lt}}{\gamma^{\star}_{lt}}\right)^{1+\sigma_h}} s^{\star}_{kt}{}^{\frac{-1}{1+\sigma_h}}. \tag{88}$$

Because the latter relationship holds for all continuing occupations, we can likewise write

$$\tilde{\gamma}^{\star}_{t} = \frac{w^{\star}_{t}}{\sum_{l\in\Omega^{\star}_{t\text{-}1,t}}\left(\frac{w^{\star}_{lt}}{\gamma^{\star}_{lt}}\right)^{1+\sigma_h}} s^{\star}_{kt}{}^{\frac{-1}{1+\sigma_h}}. \tag{89}$$

Taking ratios of our expressions for $\gamma^{\star}_{kt}$ and $\tilde{\gamma}^{\star}_{t}$ yields

$$\frac{\gamma^{\star}_{kt}}{\widetilde{\gamma}^{\star}_{t}} = \frac{w_{kt}}{\widetilde{w}^{\star}_{t}}\left(\frac{s^{\star}_{kt}}{\widetilde{s}^{\star}_{t}}\right)^{\frac{-1}{1+\sigma_h}}. \tag{90}$$

Equation (90) provides an expression for $\gamma^{\star}_{kt}$ normalized in terms of $\widetilde{\gamma}^{\star}_{t}$. We now follow parallel steps to the procedure in Ehrlich et al. (2022) and Olivares (2018) for relative appeal parameters to derive an expression for γ_{kt} that applies to entering and exiting occupations in addition to continuing occupations. To do so, we define:

$$\gamma^{c}_{kt} \equiv \frac{\gamma_{kt}}{\widetilde{\gamma}^{\star}_{t}} \quad \text{and} \quad \gamma^{a}_{kt} \equiv \frac{\gamma_{kt}}{\widetilde{\gamma}_{t}}, \tag{91}$$

i.e., γ^{c}_{kt} is relative labor supply for occupation k normalized by the geometric average for continuing occupations from time t-1 to time t, while γ^{a}_{kt} is relative labor supply for occupation k normalized by the geometric average for all occupations in the labor market at time

t. Then we can write

$$\frac{\gamma_{kt}^c}{\gamma_{kt}^a} = \frac{\frac{\gamma_{kt}}{\widetilde{\gamma}_t^\star}}{\frac{\gamma_{kt}}{\widetilde{\gamma}_t}} = \frac{\frac{w_{kt}}{\widetilde{w}_t^\star}\left(\frac{s_{kt}^\star}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}}{\frac{w_{kt}}{\widetilde{w}_t}\left(\frac{s_{kt}}{\widetilde{s}_t}\right)^{\frac{-1}{1+\sigma_h}}} = \frac{\widetilde{w}_t}{\widetilde{w}_t^\star}\left(\frac{\widetilde{s}_t}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}\left(\frac{s_{kt}^\star}{s_{kt}}\right)^{\frac{-1}{1+\sigma_h}}. \tag{92}$$

Recall from equation (10) that

$$\frac{s_{kt}}{s_{kt}^\star} = \lambda_t \rightarrow \frac{s_{kt}^\star}{s_{kt}} = \frac{1}{\lambda_t}.$$

Then we can write

$$\frac{\gamma_{kt}^c}{\gamma_{kt}^a} = \frac{\widetilde{w}_t}{\widetilde{w}_t^\star}\left(\frac{\widetilde{s}_t}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}\lambda_t{}^{\frac{1}{1+\sigma_h}} \rightarrow \gamma_{kt}^c = \gamma_{kt}^a\frac{\widetilde{w}_t}{\widetilde{w}_t^\star}\left(\frac{\widetilde{s}_t}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}\lambda_t{}^{\frac{1}{1+\sigma_h}}. \tag{93}$$

Substituting in our expression for γ_{kt}^a from equation (90) and recalling that $\gamma_{kt}^c \equiv \frac{\gamma_{kt}}{\widetilde{\gamma}_t^\star}$, we can write

$$\frac{\gamma_{kt}}{\widetilde{\gamma}_t^\star} = \frac{w_{kt}}{\widetilde{w}_t}\left(\frac{s_{kt}}{\widetilde{s}_t}\right)^{\frac{-1}{1+\sigma_h}}\frac{\widetilde{w}_t}{\widetilde{w}_t^\star}\left(\frac{\widetilde{s}_t}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}\lambda_t{}^{\frac{1}{1+\sigma_h}} \rightarrow \frac{\gamma_{kt}}{\widetilde{\gamma}_t^\star} = \frac{w_{kt}}{\widetilde{w}_t^\star}\left(\frac{s_{kt}}{\widetilde{s}_t^\star}\right)^{\frac{-1}{1+\sigma_h}}\lambda_t^{\frac{1}{1+\sigma_h}}, \tag{94}$$

which is an expression for the relative labor supply parameter γ_{kt} normalized relative to the geometric mean for all common labor types supplied in period t. It corresponds to equation (21) in Section 3.3, which was what we wanted to show.

For completeness, we now show how to generalize the expression that Redding and Weinstein (2020) provide for recovering the relative demand shifters for continuing products, which corresponds to the expression for recovering relative demand shifters for continuing occupations in our context, to apply to entering and exiting occupations as well. We again follow the approach in Ehrlich et al. (2022) and Olivares (2018).

We start with the expressions

$$\frac{\alpha_{kt}^\star}{\widetilde{\alpha}_t^\star} = \frac{w_{kt}}{\widetilde{w}_t^\star}\left(\frac{s_{kt}^\star}{\widetilde{s}_t^\star}\right)^{\frac{1}{\sigma_f-1}} \quad \text{and} \quad \frac{\alpha_{kt}}{\widetilde{\alpha}_t} = \frac{w_{kt}}{\widetilde{w}_t}\left(\frac{s_{kt}}{\widetilde{s}_t}\right)^{\frac{1}{\sigma_f-1}}. \tag{95}$$

We also define

$$\alpha_{kt}^c \equiv \frac{\alpha_{kt}}{\widetilde{\alpha}_t^\star} \quad \text{and} \quad \alpha_{kt}^a \equiv \frac{\alpha_{kt}}{\widetilde{\alpha}_t} \tag{96}$$

and take their ratio to write

$$\frac{\alpha_{kt}^c}{\alpha_{kt}^a} = \frac{\frac{\alpha_{kt}}{\widetilde{\alpha}_t^\star}}{\frac{\alpha_{kt}}{\widetilde{\alpha}_t}} = \frac{\frac{w_{kt}}{\widetilde{w}_t^\star}\left(\frac{s_{kt}^\star}{\widetilde{s}_t^\star}\right)^{\frac{1}{\sigma_f-1}}}{\frac{w_{kt}}{\widetilde{w}_t}\left(\frac{s_{kt}}{\widetilde{s}_t}\right)^{\frac{1}{\sigma_f-1}}} = \frac{\widetilde{w}_t}{\widetilde{w}_t^\star}\left(\frac{\widetilde{s}_t}{\widetilde{s}_t^\star}\right)^{\frac{1}{\sigma_f-1}}\left(\frac{s_{kt}^\star}{s_{kt}}\right)^{\frac{1}{\sigma_f-1}}. \tag{97}$$

Using our results from equation (10), we can re-write the right-hand side as

$$\frac{\alpha_{kt}^{c}}{\alpha_{kt}^{a}} = \frac{\widetilde{w}_t}{\widetilde{w}_t^{\star}} \left(\frac{\widetilde{s}_t}{\widetilde{s}_t^{\star}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_t{}^{\frac{-1}{\sigma_f - 1}} \rightarrow \alpha_{kt}^{c} = \alpha_{kt}^{a} \frac{\widetilde{w}_t}{\widetilde{w}_t^{\star}} \left(\frac{\widetilde{s}_t}{\widetilde{s}_t^{\star}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_t{}^{\frac{-1}{\sigma_f - 1}}. \tag{98}$$

Substituting in our expression for α_{kt}^{a} and recalling our definition of α_{kt}^{c} yields

$$\frac{\alpha_{kt}}{\widetilde{\alpha}_t^{\star}} = \frac{w_{kt}}{\widetilde{w}_t} \left(\frac{s_{kt}}{\widetilde{s}_t} \right)^{\frac{1}{\sigma_f - 1}} \frac{\widetilde{w}_t}{\widetilde{w}_t^{\star}} \left(\frac{\widetilde{s}_t}{\widetilde{s}_t^{\star}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_t{}^{\frac{-1}{\sigma_h - 1}} \rightarrow \frac{\alpha_{kt}}{\widetilde{\alpha}_t^{\star}} = \frac{w_{kt}}{\widetilde{w}_t^{\star}} \left(\frac{s_{kt}}{\widetilde{s}_t^{\star}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_t^{\frac{-1}{\sigma_f - 1}}. \tag{99}$$

The right-hand equation corresponds to equation (20) in Section 3.3, which is what we wanted to show.

B.5 Tracking Relative Supply and Demand Distributions over Time

The discussion in Appendix B.4 focused on recovering relative supply and demand parameters in a two-period case. Here, we discuss how we track their distributions across multiple periods. In equation (11), we apply the normalization that the geometric means of the relative supply and demand shifters are equal across consecutive periods among the set of continuing occupations, i.e.,

$$\widetilde{\alpha}_t^{\star} = \widetilde{\alpha}_{t\text{-}1}^{\star} \Leftrightarrow \overline{\ln \alpha_t^{\star}} = \overline{\ln \alpha_{t\text{-}1}^{\star}} \Leftrightarrow \overline{\Delta \ln \alpha_{kt}^{\star}} = 0 \text{ and } \widetilde{\gamma}_t^{\star} = \widetilde{\gamma}_{t\text{-}1}^{\star} \Leftrightarrow \overline{\ln \gamma_t^{\star}} = \overline{\ln \gamma_{t\text{-}1}^{\star}} \Leftrightarrow \overline{\Delta \ln \gamma_{kt}^{\star}} = 0.$$

We will sometimes be interested in how the average levels of the relative labor supply and demand parameters change over time, for instance in the counterfactual wage distributions we discuss in Section 8. Our normalization does not apply to entering and exiting occupations, so the average relative supply and demand parameters can "drift" over time. For instance, modern office work may be both relatively more productive and less physically onerous than older production occupations. If so, the geometric average of the relative demand parameters in 2020, $\widetilde{\alpha}_{2020}$, would be higher than the average in 1940, $\widetilde{\alpha}_{1940}$. Conversely, the geometric average of the relative supply parameters in 2020, $\widetilde{\gamma}_{2020}$, would be lower than the average in 1940, $\widetilde{\gamma}_{1940}$, because the γ_{kt} parameters measure the relative disutility of labor from working in a particular occupation.

Without loss of generality, we can normalize the levels of $\widetilde{\alpha}_t^{\star}$ and $\widetilde{\gamma}_t^{\star}$ to equal one in any set of consecutive periods. We apply that normalization to the first two-decade pair of our study, 1940–1950, i.e., we normalize

$$\widetilde{\alpha}_{1950}^{\star} = \widetilde{\alpha}_{1940}^{\star} = 1 \Leftrightarrow \overline{\ln \alpha_{1950}^{\star}} = \overline{\ln \alpha_{1940}^{\star}} = 0 \text{ and } \widetilde{\gamma}_{1950}^{\star} = \widetilde{\gamma}_{1940}^{\star} \Leftrightarrow \overline{\ln \gamma_{1950}^{\star}} = \overline{\ln \gamma_{1940}^{\star}} = 0. \tag{100}$$

Consider the following three-period example spanning the decades 1940, 1950, and 1960. Let $\alpha_{k,t}^{\star,t-1,t}$ denote the relative demand of occupation k in period t for an occupation that was present in both periods t-1 and t. Likewise, let $\widetilde{\alpha}_t^{\star,t-1,t}$ denote the geometric mean of the relative demand parameters in period t over occupations that were present in both periods t-1 and t. Let $\gamma_{k,t}^{\star,t-1,t}$ and $\widetilde{\gamma}_t^{\star,t-1,t}$ denote the analogous concepts for relative supply. Also let $\lambda_{t,t\text{-}1}$ denote the share of labor income in period t going to occupations that were present in

both periods t-1 and t, and let $\lambda_{t\text{-}1,t}$ denote the share of labor income in period t-1 going to occupations that were present in both periods t-1 and t.

Following our results in Appendix B.4, we can write

$$\frac{\gamma_{kt-t}}{\widetilde{\gamma}^{\star,t-1,t}_{t-1}} = \frac{w_{kt-1}}{\widetilde{w}^{\star,t-1,t}_{t-1}} \left(\frac{s_{kt}-1}{\widetilde{s}^{\star,t-1,t}_{t-1}} \right)^{\frac{-1}{1+\sigma_h}} \lambda_{t-1,t}^{\frac{1}{1+\sigma_h}}.$$

Applying our normalization in equation (100), we can write

$$\gamma_{k,40} = \frac{w_{k,40}}{\widetilde{w}^{\star,40,50}_{40}} \left(\frac{s_{k,40}}{\widetilde{s}^{\star,40,50}_{40}} \right)^{\frac{-1}{1+\sigma_h}} \lambda_{40,50}^{\frac{1}{1+\sigma_h}} \quad \text{and} \quad \gamma_{k,50} = \frac{w_{k,50}}{\widetilde{w}^{\star,40,50}_{50}} \left(\frac{s_{k,50}}{\widetilde{s}^{\star,40,50}_{50}} \right)^{\frac{-1}{1+\sigma_h}} \lambda_{50,40}^{\frac{1}{1+\sigma_h}}. \tag{101}$$

Calculating $\gamma_{k,60}$ will not be as simple. In general, we expect $\widetilde{\gamma}^{\star,50,60}_{50} = \widetilde{\gamma}^{\star,50,60}_{60} \neq 1$, because the set of continuing occupations from 1950 to 1960 differs from the set of continuing occupations from 1940 to 1950. Note, though, that we can calculate $\gamma_{k,50}$ using equation (101). We can then calculate

$$\widetilde{\gamma}^{\star,50,60}_{60} = \left(\prod_{k \in \Omega^{\star}_{50,60}} \gamma_{k,50} \right)^{\frac{1}{\left|\Omega^{\star}_{50,60}\right|}}, \tag{102}$$

where $\left|\Omega^{\star}_{50,60}\right|$ is the number of continuing occupations from 1950 to 1960. Then we can calculate $\gamma_{k,60}$ as

$$\gamma_{k,60} = \widetilde{\gamma}^{\star,50,60}_{50} \frac{w_{k,60}}{\widetilde{w}^{\star,50,60}_{60}} \left(\frac{s_{k,60}}{\widetilde{s}^{\star,50,60}_{60}} \right)^{\frac{-1}{1+\sigma_h}} \lambda_{60,50}^{\frac{1}{1+\sigma_h}}. \tag{103}$$

We can repeat this "roll-forward" procedure for as many periods as we would like in order to calculate the relative labor supply parameters γ_{kt} normalized relative to the set of continuing occupations in the base decade pair 1940–1950.

A similar procedure will apply for the relative labor demand parameters α_{kt}. We can express $\alpha_{k,40}$ and $\alpha_{k,50}$ as

$$\alpha_{k,40} = \frac{w_{k,40}}{\widetilde{w}^{\star,40,50}_{40}} \left(\frac{s_{k,40}}{\widetilde{s}^{\star,40,50}_{40}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_{40,50}^{\frac{-1}{\sigma_f - 1}} \quad \text{and} \quad \alpha_{k,50} = \frac{w_{k,50}}{\widetilde{w}^{\star,40,50}_{50}} \left(\frac{s_{k,50}}{\widetilde{s}^{\star,40,50}_{50}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_{50,40}^{\frac{-1}{\sigma_f - 1}}. \tag{104}$$

Defining

$$\widetilde{\alpha}^{\star,50,60}_{60} = \left(\prod_{k \in \Omega^{\star}_{50,60}} \alpha_{k,50} \right)^{\frac{1}{\left|\Omega^{\star}_{50,60}\right|}}, \tag{105}$$

we can calculate $\alpha_{k,60}$ as

$$\alpha_{k,60} = \widetilde{\alpha}^{\star,50,60}_{50} \frac{w_{k,60}}{\widetilde{w}^{\star,50,60}_{60}} \left(\frac{s_{k,60}}{\widetilde{s}^{\star,50,60}_{60}} \right)^{\frac{1}{\sigma_f - 1}} \lambda_{60,50}^{\frac{-1}{\sigma_f - 1}}. \tag{106}$$

C Estimating Elasticities of Substitution

C.1 Demand Elasticities: Reverse Weighting Estimator

Formally, the reverse-weighting estimator for σ_f in Redding and Weinstein (2018) and Redding and Weinstein (2023) relies on the two moment conditions:

$$\frac{1}{1-\sigma_f}\ln\left(\sum_{k\in\Omega^{\star}_{t\text{-}1,t}} s^{\star}_{kt\text{-}1}\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right)^{1-\sigma_f}\right) - \frac{1}{N^{\mathbb{C}}_{t\text{-}1,t}}\sum_{k\in\Omega^{\star}_{t\text{-}1,t}}\left[\ln\left(\frac{\tilde{w}_{kt}}{\tilde{w}_{kt\text{-}1}}\right) - \frac{1}{\sigma_f-1}\ln\left(\frac{\tilde{s}^{\star}_{kt}}{\tilde{s}^{\star}_{kt\text{-}1}}\right)\right] = 0, \quad (107)$$

$$\frac{-1}{1-\sigma_f}\ln\left(\sum_{k\in\Omega^{\star}_{t\text{-}1,t}} s^{\star}_{kt}\left(\frac{w_{kt}}{w_{kt\text{-}1}}\right)^{-(1-\sigma_f)}\right) - \frac{1}{N^{\mathbb{C}}_{t\text{-}1,t}}\sum_{k\in\Omega^{\star}_{t\text{-}1,t}}\left[\ln\left(\frac{\tilde{w}_{kt}}{\tilde{w}_{kt\text{-}1}}\right) - \frac{1}{\sigma_f-1}\ln\left(\frac{\tilde{s}^{\star}_{kt}}{\tilde{s}^{\star}_{kt\text{-}1}}\right)\right] = 0. \quad (108)$$

The RW estimator minimizes the sum of the squared deviations of these conditions from zero.

C.2 Results of Alternative Occupational Linking Rule

Our results are not sensitive to using a narrower definition of continuing occupations. This "narrow continuer rule" applies only a bi-directional matching threshold without attempting to pre-group occupational splits. The set of occupations existing continuously between 1940–2020 drops slightly from 112 under our preferred measure to 101 under our narrow measure. Those "narrow" continuing occupations account for slightly less than half of employment in 1940 and approximately one quarter of employment in 2020, which we suspect underestimates the degree of continuity in occupations present in the labor market.

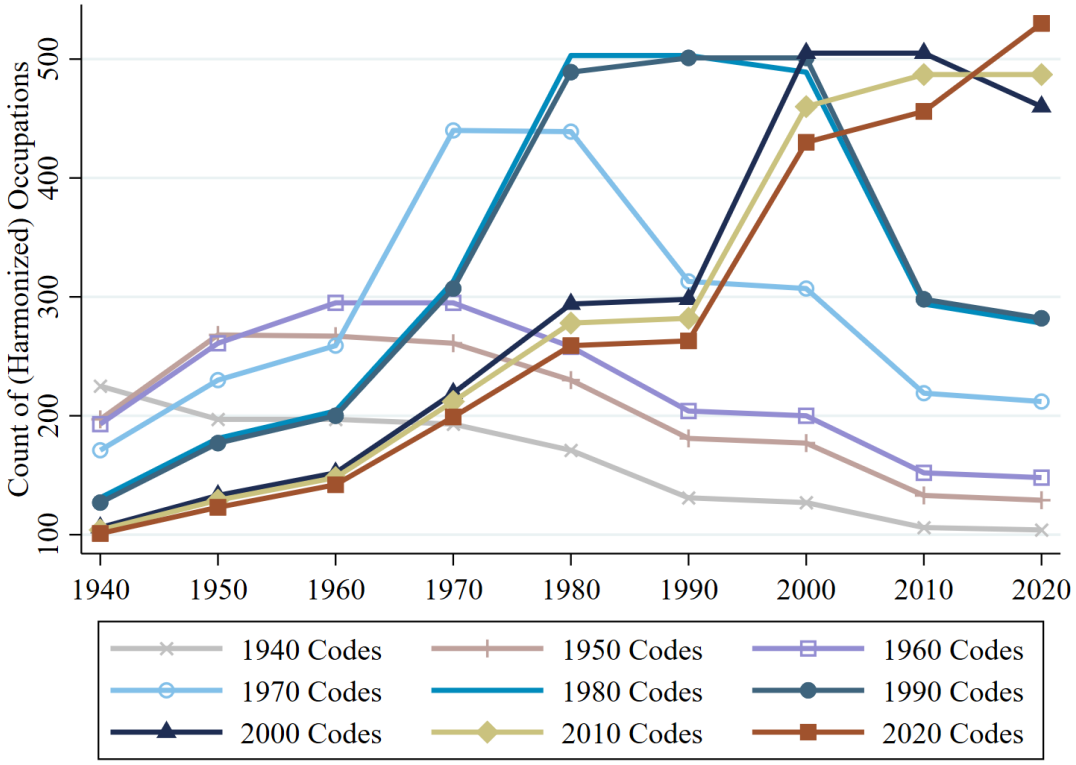
Figure C.2 illustrates the number of occupations identified in each period under the narrow rule. Relative to Figure 2, which employs our preferred measure, the narrow continuer rule tends to identify more occupations (as it does not attempt to pre-group) and generates a lumpier time series, which updates dramatically during major reorganizations of the classification scheme but little between these periods.

Despite these differences, the narrow continuer rule has little effect on the estimated elasticities of substitution, which are $\hat{\sigma}_f = 8.06$ and $\hat{\sigma}_h = 3.49$ under this classification. The wage index results are also little changed, with the resulting indices implying 49.5% growth in workers' real wages and 20.4% growth in firms' real wage index from 1940 to 2020, similar to our baseline results.

C.3 Estimating Occupational Exposures to International Trade

Much of the literature examining labor implications of trade shocks (e.g., Autor et al., 2019, 2022; Acemoglu et al., 2016; Ebenstein et al., 2014) uses a variation of the import penetration measure proposed by Bernard et al. (2006, p. 224). Acemoglu et al. (2016) compute an import penetration measure specific to China covering the period from 1991–2011. That measure is the product of two terms: an industry's share of imports coming from China, and the overall importance of imports in that industry's domestic consumption. Changes in import penetration, then, capture the effect of rising competition for domestically manu-

Figure C.2: Counts of Occupations Identified by an Alternative Algorithm, 1940–2020



Continuing occupations have at least 51% of employment come from and go to the linked occupation.

factured goods. These measures are computed for 392 harmonized 1987 Standard Industrial Classification codes (SIC87dd). To translate these measures from industry intensities to occupation intensities requires bridging several classification schemes.

Autor et al. (2019) provide a crosswalk from SIC87dd to a similarly harmonized set of 1990 Census industry codes (ind1990ddx).[36] A second mapping constructed by Autor et al. (2013) translates ind1990ddx codes to the 1990 Census industry codes available in our data. Finally, we follow Ebenstein et al. (2014, p. 585) in computing, for each occupation in 1990, its employment-weighted average of the change in import penetration across industries. Our matching algorithm connects this measure to continuing occupations in adjacent decades.

D Additional Results for Estimated Supply and Demand Shocks

As Section 3.3 describes, any pair of elasticities uniquely rationalizes each occupation's relative wage and employment through its relative demand and supply parameters α_{kt} and γ_{kt}.

[36]Some Census industry codes incorporate multiple SIC87dd codes. In such cases, we define the import penetration of a given Census industry code as the employment-weighted average of its constituent SIC87dd codes.

Table D.2: Properties of Estimated Demand and Supply Shocks

	Std. Dev. $\Delta ln(\alpha)$	Std. Dev. $\Delta ln(1/\gamma)$	Corr. $\Delta ln(\alpha), \Delta ln(1/\gamma)$	Corr. $\Delta ln(\alpha), \omega$	Corr. $\Delta ln(1/\gamma), \omega$
1940–2020	0.15	0.18	0.05	0.04	0.04
95% CI	[0.13,0.16]	[0.14,0.26]	[–0.34,0.27]	[0.02,0.06]	[0.02,0.05]
1970–2020	0.13	0.18	0.15	0.08	0.04
95% CI	[0.11,0.14]	[0.13,0.26]	[–0.22,0.33]	[0.06,0.10]	[0.02,0.06]

Subject to our normalization, which imposes that the average scale of α_{kt} and γ_{kt} among continuing occupations does not change between periods, we can use this framework to compare not only relative demand and supply across occupations at a moment in time, but also track how a given occupation's relative demand and supply evolve. The following discusses key properties of those supply and demand shocks: their size, historical incidence across occupation groups, and correlations with occupation size or external supply and demand measures.

D.1 Summary Statistics for Supply and Demand Shocks

Table D.2 presents, for the entire 1940-2020 window and separately for 1970-2020, the standard deviations of relative labor demand shocks $\Delta \ln \alpha_{kt}$ and inverse labor supply shocks $\Delta \ln \frac{1}{\gamma_{kt}} = -\Delta \ln \gamma_{kt}$ (for which a positive shock reflects a rightward shift in the supply curve). The table reveals several interesting patterns. First, it implies that relative labor supply shocks have a similar standard deviation as labor demand shocks. Second, supply and demand shocks appear positively correlated and became much more so in the period from 1970 to 2020, reinforcing the concerns that motivate our generalized reverse weighting and IV estimation approaches. The final two columns display the correlations between the Sato-Vartia weights and the estimated demand and supply shocks, respectively. As the last terms of equations 15 and 16 illustrate, those correlations determine the directions and magnitude of the shock adjustments to our CES unified wage indices. That positive demand shocks $\Delta ln(\alpha)$ and positive supply shocks $\Delta ln(\frac{1}{\gamma})$ both correlate positively with their occupations' share of the wage bill ω imply that these shocks tend to benefit workers and firms alike.

D.2 Long-Run Supply and Demand Shocks

Figures 4 and 5 displayed estimated relative supply and demand parameters for SOC2 occupation groups from 1980 to 2020. Because individual occupations enter and exit over time, the mix of occupations represented beneath an aggregate occupation group shifts over time. Those changes can affect the estimated group-level relative supply and demand parameters even if the continuing individual occupations within a group exhibit stable relative supply and demand.

To minimize the effects of this distortion when conveying trends from 1940–2020, Figures D.3 and D.4 map Census occupations into the midpoint 1980 Standard Occupational Classification groups, rather than the 2000 SOC groups applied in Figures 4 and 5. The 1980 classification enumerates 13 broad occupational groups, which can be found in Appendix B of U.S. Census Bureau (1989). For the summaries below, we further disaggregate the "Professional specialty occupations" group with constituent categories "Health diagnosing and treating," "Education" (Teachers, postsecondary; Teachers, except postsecondary; Librarians, archivists, and curators), "Science, engineering and math," (Engineers, architects, and surveyors; Mathematical and computer scientists; Natural scientists; Social scientists and urban planners), and "Other professionals" (all residual categories). Similarly, we separate "Food preparation and service" and "Personal service" from other "Service occupations, except protective and household" categories. Into this last group, we combine the small "Private household" and "Protective service" occupations into one category of "Other services."

Relative demand for occupational labor in the "Health diagnosing and treatment" group has risen from below average in 1940 to significantly above average in 2020. The "Legal," "Business and finance," and "Math and computers" groups have also seen large increases in relative demand. Conversely, non-service occupational groups such as "Precision craft and repair," "Construction and mining," "Transportation," "Assembly machine operators," and "Miscellaneous laborers" have all seen meaningful declines in relative labor demand. Some service occupation groups, notably "Administrative support and clerical" and "Food services," have also seen declines relative labor demand.

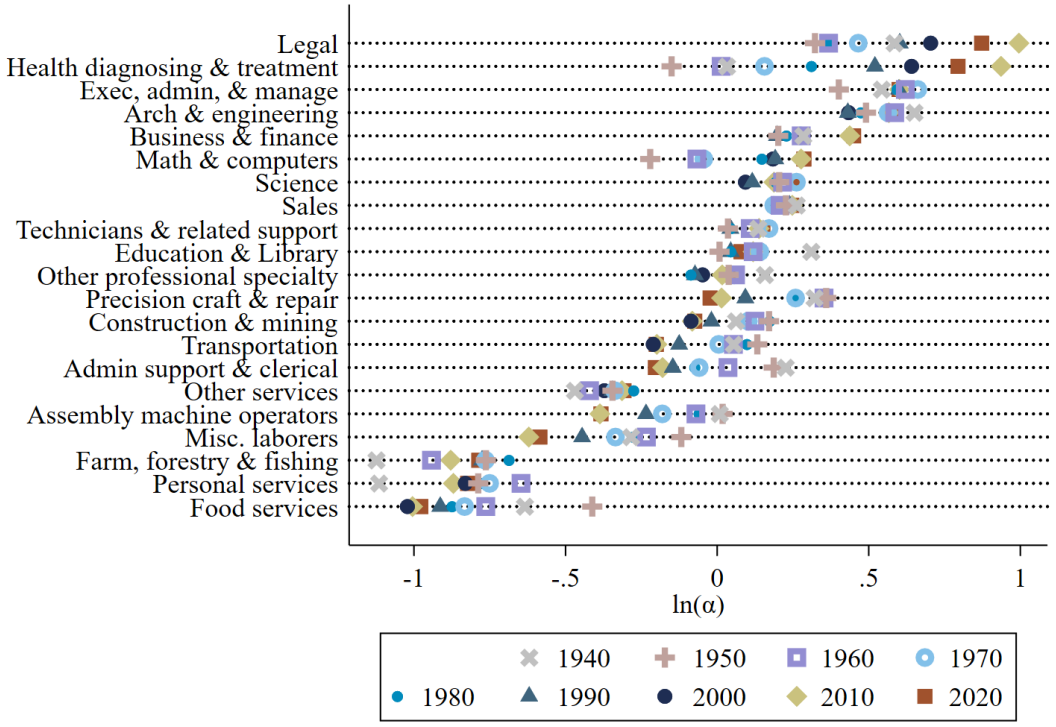
There are large cross-sectional differences across occupational groups in terms of relative labor supply, but the changes over time are generally smaller than the shifts in relative demands. The "Food services," "Administrative support and clerical," and "Other services" groups have the highest relative supplies. The "Math and computers," "Legal," and "Science" groups have the lowest relative supplies. The "Executive, administrative, and management" and "Health diagnosing and treatment" groups have both experienced noticeable declines in relative supply. The "Math and computers" group has experienced both upward and downward movements in relative supply across decades, with relative supply in 2020 settling near its 1960 value.

D.3 External Validation of Shocks

Figure D.5 illustrates the correlations between estimated supply shocks $\Delta ln(1/\gamma)$ for each occupation and that occupation's change in the share of workers who are female. Every decade exhibits a positive correlation, but the relationship weakens after the rise in female labor force participation tapers in 1990. The relationship between the estimated labor supply shocks and the change in the share of female workers was especially strong from 1970 to 1990, when the aggregate share of women in the workforce was rising rapidly and the occupations they pursued were becoming more diverse.

Figure D.6 displays results from a similar validation exercise for our estimated labor demand shocks. Each observation in the scatter plots represents our estimate of the change in an occupation's relative demand, $\Delta \ln \alpha_{kt}$, against an estimate of its exposure to trade with China. We estimate each occupation's exposure by mapping the industry-level estimates of Acemoglu et al. (2016) for the years 1991–2011 to our three-digit occupations, as in Ebenstein

Figure D.3: Estimated Relative Demand Shocks from 1940-2020



Marks represent (wage-bill weighted) averages across 3-digit Census occupations within given 1980 Standard Occupational Classification group. Continuing occupations in adjacent decades are linked to 1980 groups using our matching algorithm. See paper for details about slight grouping modifications and a comment about continuer selection.

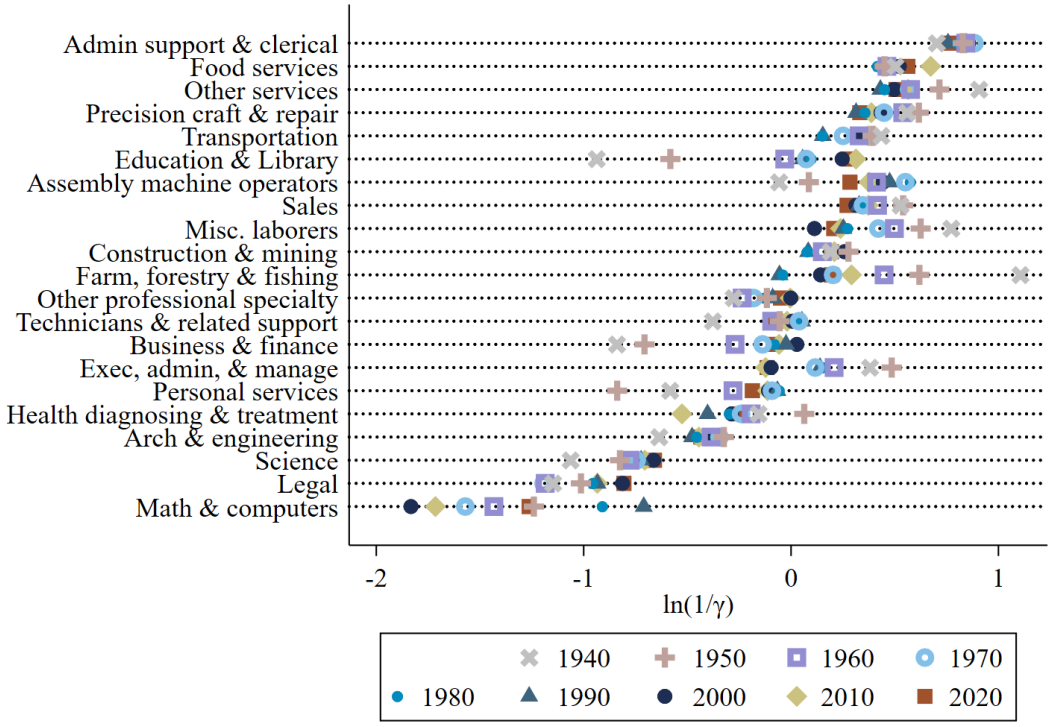
et al. (2014). More details on constructing occupational trade exposures are available in Appendix C.3. Figure D.6 plots these proxies of each occupation's exposure to import competition against our estimated relative demand shocks in each of the decade pairs from 1980–1990 to 2010–2020. There is an evident negative correlation between trade exposure and the estimated demand shocks in the decades 1980–1990 through 2000–2010, consistent with the results of Acemoglu et al. (2016). The correlation is much weaker from 2010–2020, after the period for which the trade exposures were estimated.

We interpret the evidence in Figures D.5 and D.6 as suggesting that our estimated relative labor supply and demand shocks reflect first-order trends in the labor market in a reasonable way. Our estimated shocks have a quantitative interpretation in terms of a simple but formal equilibrium model. They allow us to identify shocks to supply separately from shocks to demand and to examine counterfactual scenarios that decompose changes in the wage structure into their underlying sources.

E Additional Counterfactual Wage Distributions

To illustrate the importance of empirically disciplining labor supply in our framework, we repeat the counterfactual exercise of Section 8 using two alternate assumptions. In Figure E.8 the relative supply of labor in each occupation is assumed to be fixed or perfectly

Figure D.4: Estimated Relative Supply Shocks from 1940-2020



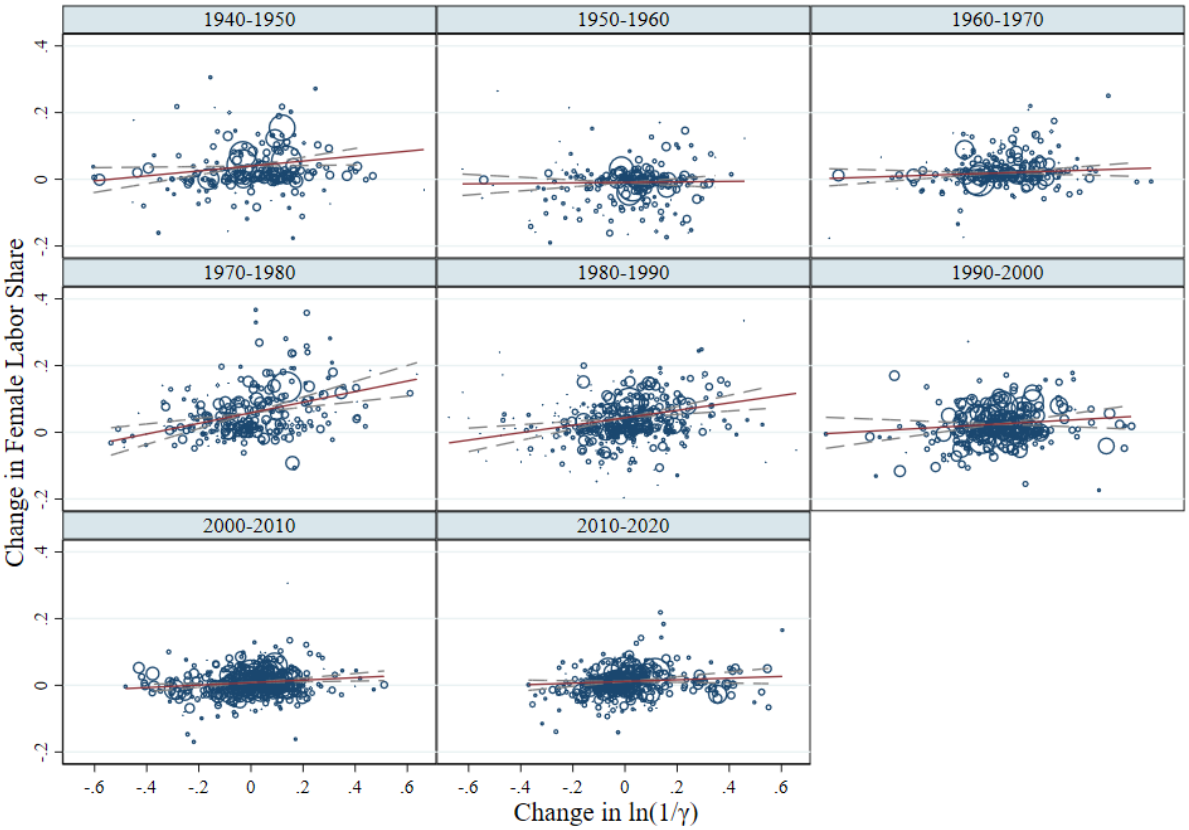
Marks represent (wage-bill weighted) averages across 3-digit Census occupations within given 1980 Standard Occupational Classification group. Continuing occupations in adjacent decades are linked to 1980 groups using our matching algorithm. See paper for details about slight grouping modifications and a comment about continuer selection.

inelastic (i.e. vertical). Figure E.9 instead depicts a perfectly elastic relative labor supply across occupations. As in Figure 8, the light and dark blue dashed lines depict the empirical wage distributions of 1970 and 2020, respectively. Lavender lines depict the counterfactual wage distributions we estimate would have occurred if only the supply shocks we estimate had taken place, while orange lines depict such counterfactuals under only demand shocks. Dotted lines exclude the force of occupational entry and depict distributions computed only among the occupations that existed in 1970. Solid lines permit entry, with occupations entering with their initial relative supply or demand (and, for the purpose of counterfactuals, maintaining those levels).

E.1 Wage Decompositions with Perfectly Inelastic Supply

With fixed relative labor supply, all changes in relative employment across occupations result from supply shifts. Changes in demand induce relative wage adjustments, but do nothing to re-allocate workers across occupations. Figure E.8 reveals that interpreting the data with such a paradigm generates results strongly counter to both our baseline and the literature's core conclusions about the employment implications of changes in technology and trade over this period. Such a framework sees almost the entire increase in wage dispersion over the past fifty years as resulting from *supply* shocks, with demand contributing little. This likely

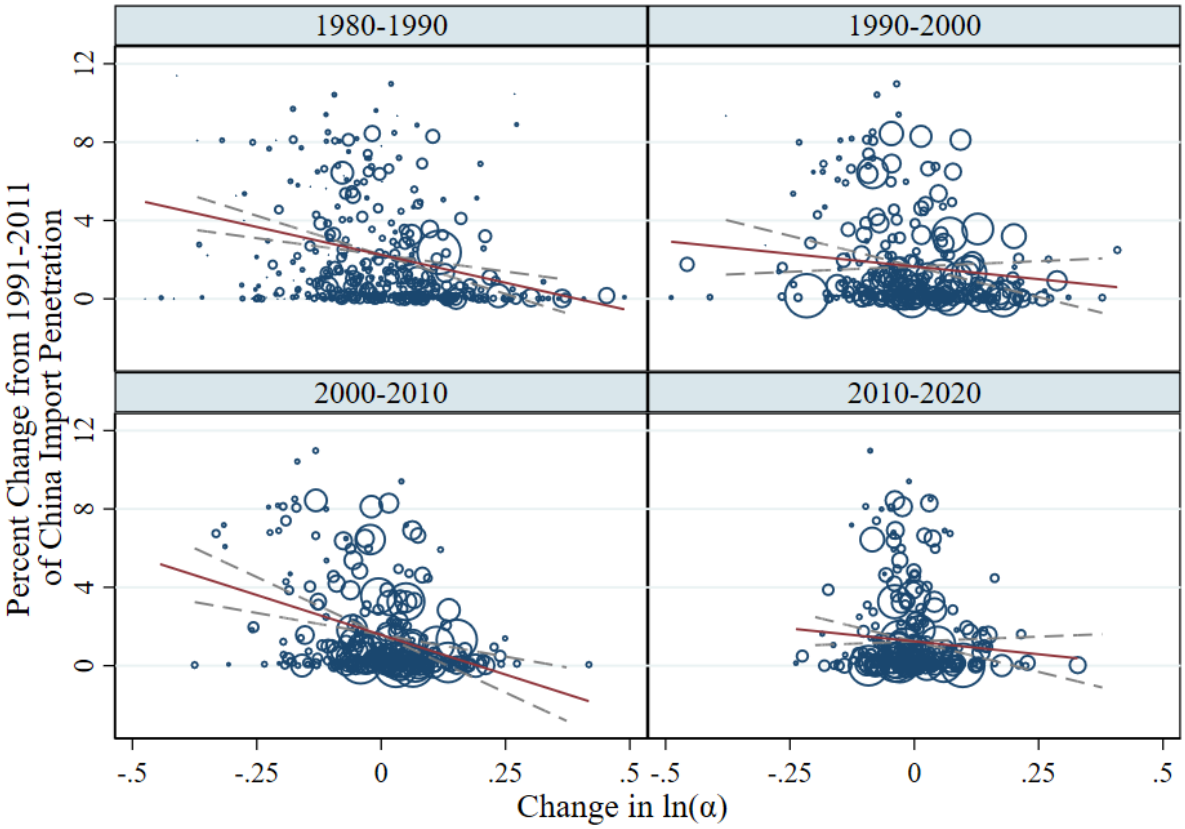
Figure D.5: External Validation of Estimated Relative Supply Shocks by Decade



Note: Dot size indicates weight, given by Tornqvist average of wage bill shares. Outliers removed beyond axis limits. When pooled across decades, the average slope is .06. Maximum slope is 0.26 in 1970-1980 and minimum is 0.04 from 2000-2010. Gray lines indicate 90% confidence intervals.

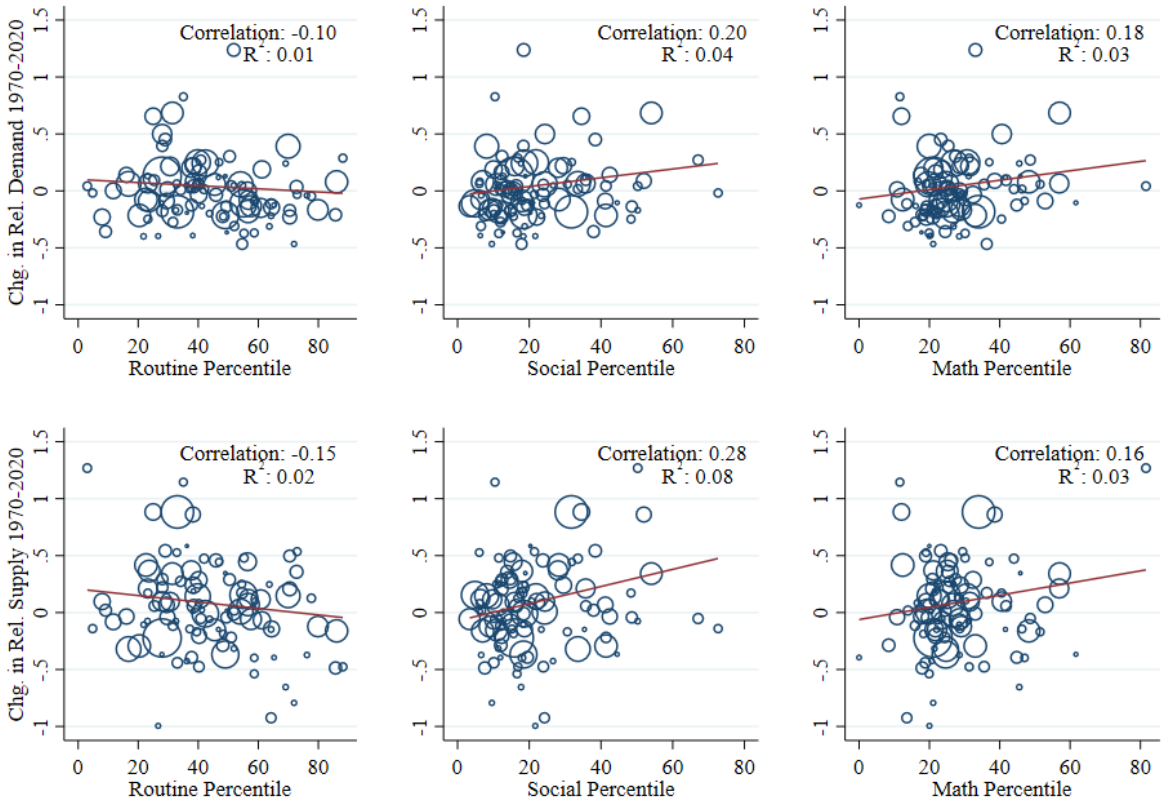
reflects that changes in relative employment, which the model with perfectly inelastic supply identifies as supply shocks, are significant contributors to rising occupational wage dispersion over this period.

Figure D.6: External Validation of Estimated Relative Demand Shocks



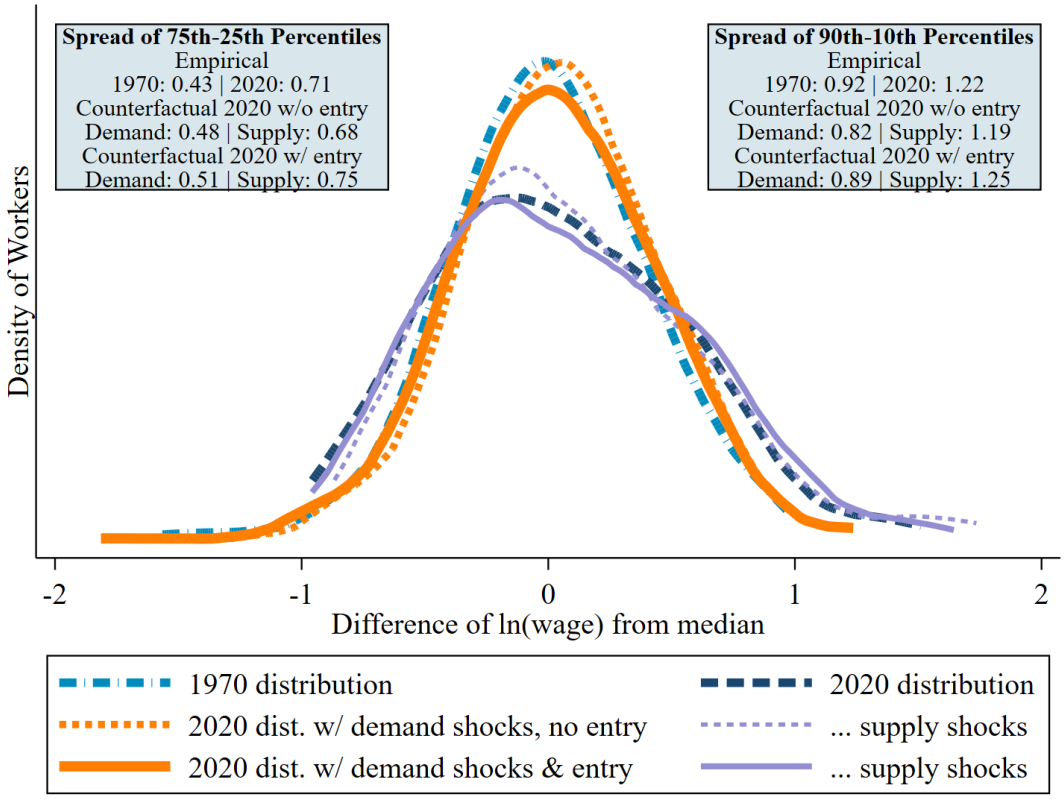
Note: Dot size indicates weight, given by Tornqvist average of wage bill shares. Outliers removed beyond axis limits. Slopes range from -8.2 in 2000-2010 to -2.6 in 2010-2020. Gray lines indicate 90% confidence intervals. Changes in import penetration estimated by Acemoglu et al. (2016) for SIC 1987 codes. Shown are occupation employment-weighted averages across linked 1990 Census industries.

Figure D.7: Correlations of Supply and Demand Shocks to Task Content
Among Occupations with College Shares Below 75th Percentile



Note: Each observation represents a 1990 Census occupation weighted by its Sato-Vartia wage bill share. X-axes depict the occupation's percentile for a task type, from Deming (2017). These are correlated in the top row with $\ln(\alpha_{2020}) - \ln(\alpha_{1970})$ and in the bottom with $\ln(1/\gamma_{2020}) - \ln(1/\gamma_{1970})$. These scatterplots and the regressions they summarize exclude all occupations with 2020 college shares above the 75th percentile of 62%.

Figure E.8: Counterfactual Relative Wage Distributions
Assuming Perfectly Inelastic Supply Across Occupations

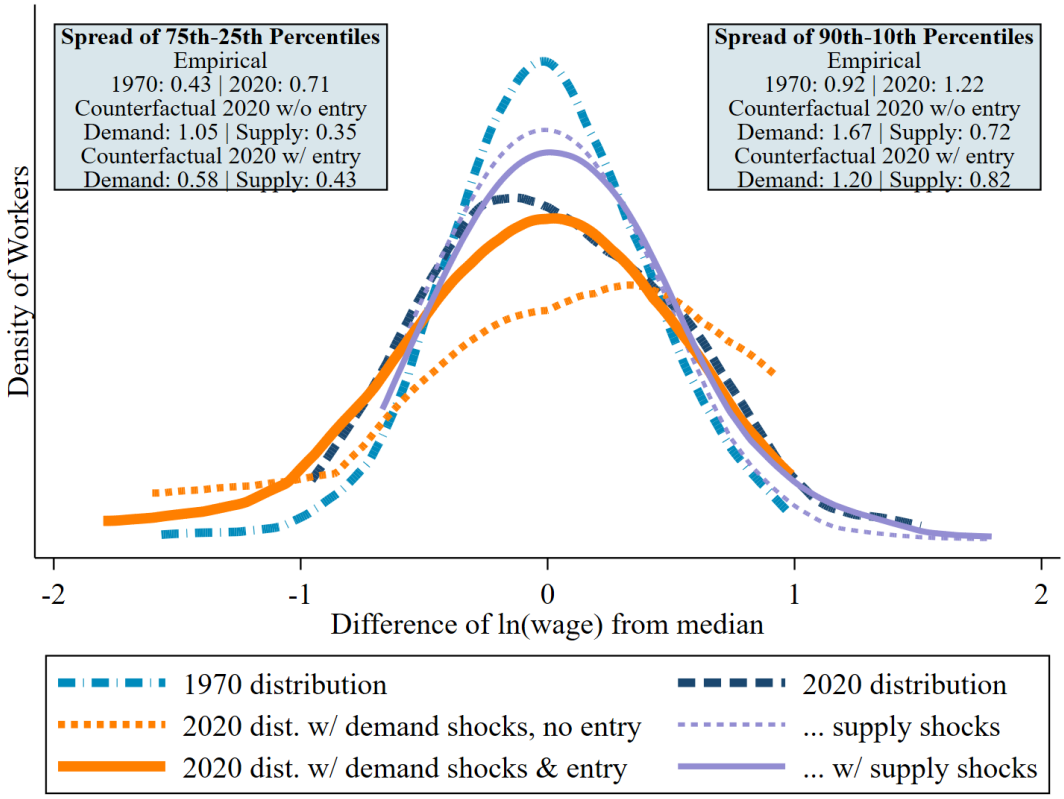


E.2 Wage Decompositions with Perfectly Elastic Supply

Figure E.9 repeats the exercise in Figure 8 of Section 8 under the assumption that relative supply curves are perfectly elastic. In that alternative model, each occupation has a wage for which the workforce will supply any (relative) amount of labor. Supply, therefore, entirely determines the prevailing wage of each occupation. The orange and lavender lines in Figure E.9 reveal patterns which, in combination, highlight the importance of accounting for upward-sloping supply. The orange lines, interpreted as demand-driven worker re-allocations, capture strong increases in demand in low-skill occupations over the period.[37] With perfectly elastic supply, those re-allocations take place without need for an equilibrating wage increase, creating a distribution that is more left-skewed than has been observed empirically (or than is implied by the similar counterfactual in our baseline estimates under upward-sloping supply). Conversely, the model interprets the wage changes that have taken place as arising from supply shocks. The distribution that would have resulted from those shocks in isolation, depicted in lavender, is right-skewed. At the wages that follow such shocks, there is no mass of employment in the left tail while a larger mass of workers move into top-wage jobs than has been empirically observed. To emphasize the difficulty this framework has in simultaneously explaining developments in demand and supply, we compute the implied correlation between demand and supply shocks recovered by this model. In our baseline model, we estimate a positive correlation of 0.18 between demand and supply shocks (see Table D.2), consistent with the intuition that demand and supply shocks respond to one another. In contrast, under the assumption of perfectly elastic supply, the empirical distribution of relative wages and shares would require such shocks to be strongly negatively correlated, with a coefficient of -0.85.

[37] The counterfactual considering only 1970 occupations results in a distribution which is more dramatically polarized, with mass at the middle shifting out in both directions.

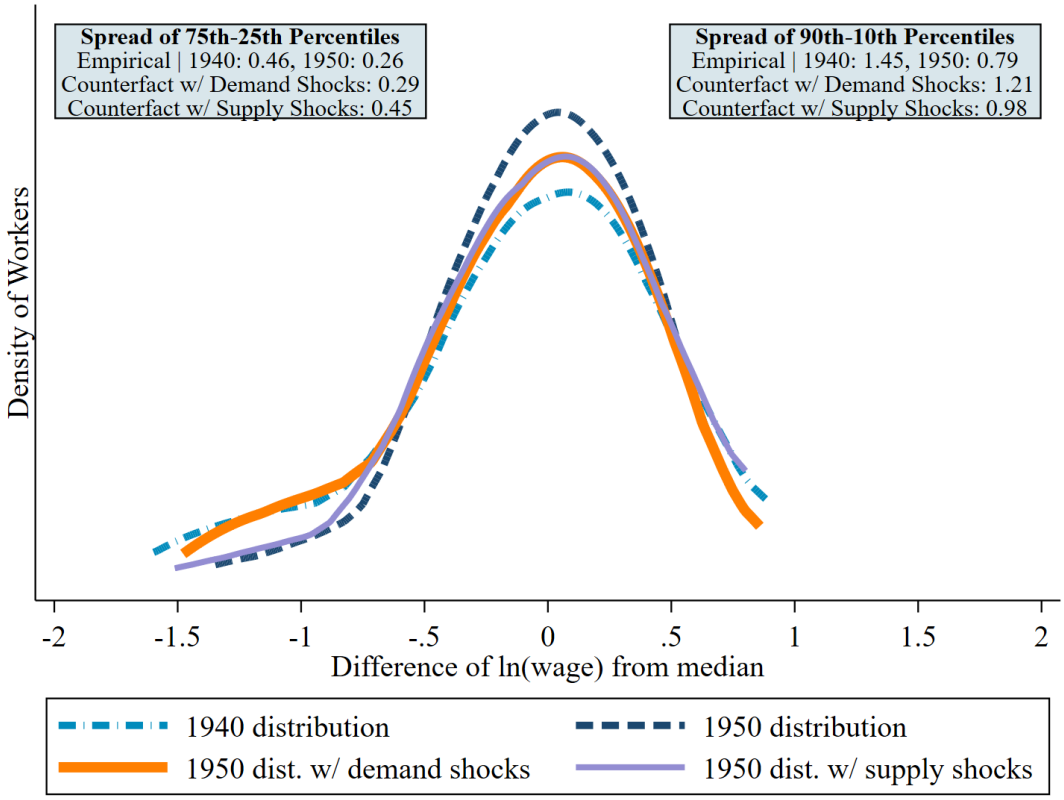
Figure E.9: Counterfactual Relative Wage Distributions
Assuming Perfectly Elastic Supply Across Occupations



E.3 Decomposing 1940–1950 Wage Compression

Figure E.10 presents the results from another counterfactual wage decomposition, examining the period from 1940–1950. By our measure, the 90th-10th percentile spread of log occupational average wages declined from 1.45 to 0.79 points. This compression is visually evident from comparing the dashed light blue line, which shows a kernel-smoothed version of the observed 1940 distribution of average wages across occupations, to the dashed dark blue line, which shows the observed 1950 distribution. The 1940 distribution has a lower mode and a much thicker left tail than the 1950 distribution. Goldin and Margo (1992) document that wage dispersion also fell along other dimensions such as education and job experience, in a phenomenon they dub the "Great Compression." Goldin and Margo (1992, abstract p. 1) find that "World War II and the National War Labor Board share some credit for the Great Compression, but much was due to an increased demand for unskilled labor when educated labor was greatly expanding."

Figure E.10: Decomposing Changes in the Wage Structure, 1940–1950



The solid orange line in Figure E.10 displays a counterfactual wage distribution in 1950 in which the relative demand parameters evolved according to our estimates from 1940 to 1950, but the relative supply parameters remained fixed at their estimated 1940 values. The solid lavender line shows the analogous counterfactual with shifts in relative supply but no

shifts in relative demand.[38] We estimate that the 90th-10th percentile spread would have compressed to 1.21 in the absence of supply shocks and to 0.86 in the absence of demand shocks. The counterfactual with supply shocks but no demand shocks also recreates the observed thinning of the left tail of the wage distribution from 1940 to 1950, while the counterfactual with demand shocks but not supply shocks does not. The interquartile range of log average occupational wages compressed from 0.46 in 1940 to 0.26 in 1950 in the observed data. We estimate that the interquartile range would have declined to 0.38 in the presence of the estimated demand shocks but no supply shocks, while it would actually have increased slightly, to 0.49, in the presence of the estimated supply shocks but no demand shocks. The interaction between the actual supply and demand shocks therefore acted to reduce the interquartile range by more than either set of shocks on its own.

Our analysis largely reinforces the key conclusion of Goldin and Margo (1992) that the Great Compression was driven by a combination of shifts in relative supply and demand in the labor market. Several important differences between our analyses make direct comparison difficult. They focus, for instance, primarily on white men, while we include all workers in our analysis; they consider educational and experience differences in the labor market in addition to the occupational differences we consider; and they are able to consider factors beyond shifts in relative supply and demand. Still, the basic consistency of our results with those of Goldin and Margo (1992) serves as an encouraging proof of concept for our methodology as we turn to more recent history, during which the drivers of changes in wage dispersion are arguably more contested.

[38] We assign occupations that we identify as having entered the labor market in 1950 their estimated relative supply and demand parameters at that time. We allow occupations to exit as identified in the data. Our algorithm indicates that occupational turnover was minimal from 1940 to 1950, so the treatment of entering and exiting occupations does not have a meaningful effect in this exercise.