The Distributional Impact of the Minimum Wage*

Erik Hurst
Chicago Booth
Erik.Hurst@chicagobooth.edu

Patrick Kehoe
Stanford and Minneapolis Fed
pkehoe@stanford.edu

Elena Pastorino
Hoover, Stanford, SIEPR, and Minneapolis Fed
epastori@stanford.edu

Thomas Winberry
Wharton
twinb@wharton.upenn.edu

March 2022

Evolving

Abstract

We develop a general equilibrium framework with worker heterogeneity, monopsony power, and putty-clay frictions in order to study the distributional impact of large changes in the minimum wage. In the long run, we find that a high minimum wage has perverse distributional impacts in that it reduces the employment, income, and welfare of precisely the low-income workers it was meant to help. However, these long-run consequences take twenty years to fully materialize because firms slowly adjust the labor intensity of their capital stock. We also study the long-run impacts of alternative transfer programs, such as the earned-income tax credit, and find they are much more effective at improving outcomes for workers at the bottom of the wage distribution. In this context, a modest increase in the minimum wage is beneficial because it prevents firms from lowering the pre-transfer wages they pay to workers.

*We thank Christian vom Lehn for helpful conversations and Brent Neiman and Kjetil Storesletten for sharing their data/codes. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

Recent proposals have been advanced in the United States to increase the current federal minimum wage from its current level of \$7.25 to at least \$15 per hour. The goal of these proposals is to improve the welfare of workers currently earning less than this new minimum, especially those at the bottom of the wage distribution. In contrast to past changes in the minimum wage, an increase of this size would have an impact on a large fraction of the U.S. workforce. Figure 1 shows the distribution of wages across individuals from the 2017-2019 American Community Survey (ACS) by whether or not an individual has obtained a bachelor's degree.[1] As seen from this figure, about 40% of non-college-educated workers and 10% of college-educated workers currently earn a wage less than \$15 per hour. Furthermore, the dispersion of wages within education group suggests the minimum wage would have very different effects for different workers. For example, in the current wage distribution, a \$15 per hour minimum wage would nearly double the wages of workers in the bottom 20% of the non-college wage distribution but would not bind on workers in the top 40% of the distribution.[2]

FIGURE 1: Empirical Wage Distribution

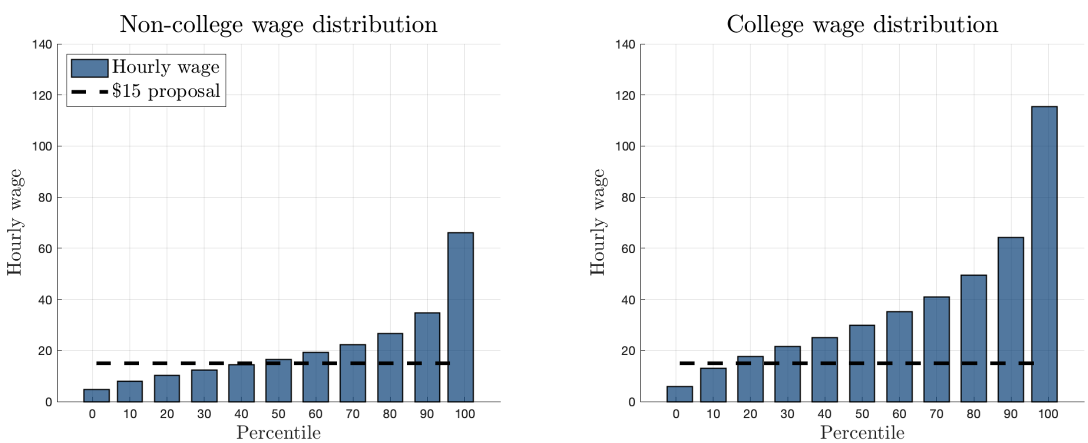

Notes: Wage distribution of full time employed workers over the age of 16 from the 2017-2019 ACS data.

In this paper we develop a framework that can capture this large within group-heterogeneity in the effects of the minimum wage and use it to assess the distributional impact of this policy in both the short run and the long run. Our main finding is that in the long run, although a large increase in the minimum wage can raise the labor income for non-college-educated workers as a whole, the

[1] Specifically, the sample includes all individuals over age 16 who report currently working more than 30 hours per week and worked at least 29 weeks during the prior year. A full discussion of the data used to construct this figure can be found in the online appendix.

[2] Note that if one used a model with only two types of workers, represented by the median non-college-educated worker and the median college-educated worker, to evaluate the proposed \$15 minimum-wage, then it would not bind on either group, since their median wages are \$16.40 and \$29.80 respectively. This helps clarify why an analysis of the minimum wage must, at its heart, be about the differing impacts on workers within such groups.

workers whom this policy is designed to benefit – namely those at the bottom of the wage distribution – are made much worse off. For example, nearly half of non-college workers experience a reduction in their employment in response to a $15 per hour minimum wage, and among those about third experience a reduction in welfare. In the short run this change has little effect because firms only slowly adjust the labor-intensity of their capital stock and these long-run consequences take about twenty years to fully materialize.

Given these perverse long-run distributional consequences, we also study other policies such as expanding the earned income tax credit (EITC). We find that the EITC dominates the minimum wage from the perspective of low-income households by substantially increasing both their employment and welfare. Given that an EITC system is in place, we find that a modest minimum wage, on the order of $9, further benefits these low-income households by preventing firms from lowering the pre-transfer wages they pay to their workers. In this sense, we find an important role for a modest increase in the minimum wage to support the EITC program. We find similar results for increasing the progressivity of the tax system more generally.

The framework we develop is consistent with three key features of the data. First, given that our goal is to explore the distributional impact of the minimum wage, we ensure that our model matches the within-education group wage distribution from Figure 1, especially for the low-income workers who are most affected by the minimum wage. Second, we also ensure that our model matches the evidence from a recent literature that documents that wages are marked down relative to the marginal product of labor.[3] Third, we are consistent with both the short-run and long-run evidence of the response of labor to a change in the relative price of labor to capital. The short-run evidence is provided by a large literature has found that increases in the minimum wage, which raise the relative price of labor, have at most a *small* effect on employment in the short run.[4] The long-run evidence is provided by a literature that finds that declines in the relative price of capital, which also raise the relative price of labor, have *large* long-run effects on employment (see, for example, Krusell et al. (2000)).

In terms of the evidence on the impact of the minimum wage, while there has been a fair bit of work on the short run impacts of the minimum wage, there has been almost none on the long run impacts. Indeed, as noted in Brown (1999)'s review article, "There is a simply a stunning absence of credible evidence—indeed, of credible attempts—[to identify] the long run effects [of the minimum wage]". The reason for this is that existing empirical evidence exploits small or local minimum

[3]See, for example, Manning (2021b) and the references therein.

[4]Extensively reviewing this literature is beyond the scope of our paper. However, prominent papers and surveys include, for instance, Brown (1988), Deere, Murphy and Welch (1995), Card and Krueger (1995), Kennan (1995), Neumark and Wascher (2008), Flinn (2010), Cahuc, Carcillo and Zylberberg (2014), Card and Krueger (2015), Neumark and Shirley (2021), and Manning (2021a).

wage changes using either time series or cross-state variation. The identifying assumptions of these studies allow research to examine the labor market effects of the minimum wage in the few years after minimum wage changes. Given that other drivers of forces that influence the labor market are changing over time, existing empirical strategies are designed to assess the short-run (1-5 year) effects of changes in the minimum wage.

In our model, the long-run elasticities of substitution between capital and various types of labor are important in determining the long-run effects of the minimum wage. Since there is essentially no direct evidence of the long run impact the relative price of labor to machines emanating from the changes in minimum wages, we discipline the the short- and long-run elasticities of substitution for capital and labor with estimates of how firms adjust capital to labor ratios in response to other well-identified changes in the relative of price of machines to workers. Once we do so, we can quantify these elasticities and then can use our framework the long-run effects of the minimum wage for different types of workers.

We turn now to describing our framework in more detail. Our model has three features which are key for our study. First, given that our goal is to explore the distributional impact of the minimum wage, we incorporate rich worker heterogeneity in order to match the wage distribution from Figure 1.[5] Workers are classified into two broad groups, non-college educated and college-educated workers, and within each group workers differ in their productivity in the labor market. Importantly, this feature of our model is consistent with the evidence that the majority of observed differences in wages across workers arise from heterogeneity in workers' characteristics, as opposed to heterogeneity in firms' attributes (see, for instance, Abowd, Kramarz and Margolis (1999) and Kline, Saggio and Sølvsten (2020)).

The second key feature of our model is monopsonistic competition among firms, which we incorporate into a directed search environment featuring endogenous labor market participation. Monopsony power captures the idea, going back to at least Robinson (1933), that increases in the minimum wage can help alleviate some of the distortions arising from monopsony power and lead to a desirable redistribution of firm profits to worker wages. Our monopsonistically competitive directed search equilibrium extends this idea to allow for multiple firms and frictional labor markets, which are central to modern treatments of labor market dynamics and is motivated by several reasons. Incorporating a non-participation margin allows us to capture the idea that a higher minimum wage incentivizes non-participants to enter the labor market and search for a job. Additionally, our directed search framework allows us to generalize the notion of a firm-specific labor supply curve in

[5]Note that if one used a model with only two types of workers, represented by the median non-college-educated worker and the median college-educated worker, to evaluate the proposed $15 minimum-wage, then it would not bind on either group, since their median wages are $16.40 and $29.80 respectively. This helps clarify why an analysis of the minimum wage must, at its heart, be about the differing impacts on workers within such groups.

Robinson (1933) to incorporate dynamic wage contracts with multi-worker firms and endogenous job-finding rates. Most importantly, this framework avoids the need to specify ad-hoc rationing rules in the allocation of workers to jobs when the minimum wage induces labor supply to exceed labor demand. Indeed, in our framework if at a given level of posted vacancies, the minimum wage induces labor supply to exceed labor demand, firms find it optimal to simply decrease vacancies until this excess supply is dissipated.

The last key feature of our model is the putty-clay technology, which captures differences in the short-run and long-run substitutability of capital and labor.[6] Specifically, when deciding on new investment, firms choose to build a unit of capital with any desired ratio of capital to the labor of each worker type that lies on the frontier of the production function. Once installed, however, the unit of capital is clay-like in that it requires that fixed amount of each type of labor to operate it. That is, our production function is nested CES ex-ante but Leontief ex-post.

The putty-clay technology implies that, given a stock of capital, the short-run demand of each type of labor is inelastic because it is infeasible to substitute between existing capital and the various types of labor. Hence, in the short run, our model is consistent with the view that the minimum wage can be set arbitrarily high without adversely affecting employment, as long as it remains profitable for firms to continue operation. Over time, though, new capital goods embodying new ratios of capital to labor can be installed. In particular, after the introduction of the minimum wage, firms do not immediately cut employment since they need the currently employed workers to operate the existing capital stock. But in the long run, firms shift their capital to labor ratios to the preferred mixes governed by the underlying CES production function and reduce employment accordingly. We choose parameters for this production function to ensure that the implied long-run substitutability is consistent with firms' observed response to the decline in the price of capital over the past forty years. Hence, our model captures the notion that firms' adjustment to changes in the relative price of labor are small in the short run but large in the long run.

We obtain four main results from our analysis. The first is that, in the long run, large changes in the minimum wage disproportionately reduce the employment, income, and welfare of low-income workers. In our model, the effect of the minimum wage on a particular worker depends on the gap between that workers' market wage and what their wage would be in the efficient allocation without monopsony power. The key tension in our model is that the distribution of efficient wages is highly dispersed, so that a single minimum wage cannot eliminate the monopsony distortion for all workers at once; setting the minimum wage high enough to eliminate the monopsony distortion for the average worker requires setting it above the efficient wage for low-productivity workers,

[6]Our modeling of putty-clay capital builds on the work that dates back atleast to Johansen (1959) and was extended by Calvo (1976).

reducing their employment. A critical component of our parameterization, therefore, is to discipline the distribution of efficient wages, especially for workers at the bottom of the wage distribution.

In our baseline model, a \$15 minimum wage has adverse long-run consequences for low-income workers. Employment falls for 44% of non-college-educated workers, especially the low-income ones. For example, the employment rates of workers who were originally making \$7.25 per hour falls by 68% in the long run; the members of that group who keep their jobs will end up earning higher wages, but even so, the total labor income of the group falls by 34%. In contrast, the substitution by firms toward higher-productivity workers implies that employment increases for those who were initially earning over \$13 per hour. In this sense, a large minimum wage has perverse distributional impacts; it reduces employment, income, and welfare of exactly the low-income workers that it is designed to help. Crucially, this result occurs despite the fact that the minimum wage raises non-college labor income in the aggregate.

Our second main result is that the short-run effects of the minimum wage are much smaller than their long-run effects due to the putty-clay frictions. Although a \$15 minimum wage will ultimately reduce aggregate non-college employment by over 8% in the long run, in our model employment only falls by 3% four years after its introduction, and it takes roughly twenty years for the capital stock to fully adjust to its new long run level. As we highlight throughout, these frictions allow our model to match many of the existing short-run estimates from the literature. We make this point concretely by replicating small, local changes in the minimum wage of the type usually studied in the empirical literature and showing that the model's response is in line with the small employment elasticities estimated in that literature. Since only a small part of the adjustment occurs during the first four years after the minimum wages, this finding also implies that empirical estimates of the effects of the minimum wage in the first few years after its introduction will not detect its ultimate long-run consequences. These results suggest that policy makers should not use the short run empirical estimates from the empirical minimum wage literature to forecast the long run impact of a large minimum wage change.

Our third main result is that existing policies within the tax and transfer system, like the EITC, do a much better job at raising the employment, income, and welfare of low-income workers. This occurs both because these policies directly target low-income workers and because they offset the monopsony distortion for these workers face in the labor market. We show how the EITC leads to higher welfare for low-earning individuals relative to the minimum wage. We find a similar result from increasing the general progressivity of the U.S. tax system. Interestingly, if the goal of policy is to help low-wage workers, even a simple uniform income tax cut is preferable to the minimum wage. Overall, our analysis shows that many other tax and transfer policies dominate a comparable change in the minimum wage if the goal is to simultaneously reduce monopsony power in the labor market and support the income of low earners.

Our paper contributes to a large literature which seeks to understand the quantitative effects of the minimum wage on US labor markets. Sorkin (2015) and Aaronson et al. (2018) use a variant of the standard putty-clay setup to analyze the short- and long-run effects of the minimum wage. Like us, these papers argue that putty-clay technology implies that the short-run effects of the minimum wage are smaller than the long-run effects. Our work differs from theirs in four key respects. First, we allow for rich worker heterogeneity — and extend the putty-clay model to incorporate this heterogeneity — in order to study the dynamics of the distributional impact of the minimum wage. Second, Sorkin (2015) and Aaronson et al. (2018) study an industry equilibrium model with an application to the restaurant industry. Instead we study a general equilibrium model to account for the fact that large changes in the minimum wage will bind for many sectors of the economy. Third, we discipline the long-run elasticities of substitution using the effects of the decline in the relative price of capital — which raises the relative price of labor — observed over the last forty years. Finally, our model of the putty-clay technology allows for firms to continually make decisions about which types of workers to hire and which mix of capital types to use instead of just focusing on the firm shut down condition.

Additionally, we build on existing work which studies the minimum wage through the lens of search models of the labor market. Important papers in this literature include Eckstein and Wolpin (1990), Flinn (2006), Ahn, Arcidiacono and Wessels (2011), and Lise, Meghir and Robin (2016). Our work extends this literature by (i) incorporating capital with the putty-clay feature which permits a meaningful distinction between the short- and long-run effects of the minimum wage, (ii) allowing for monopsony power in the labor market, (iii) closing the model in general equilibrium, and (iv) allowing for richer worker heterogeneity than is typically considered.

Our model of monopsonistic competition is the natural analog in the labor market of the monopolistic competition in the goods market adapted to a search setting. In recent work, Berger, Herkenhoff and Mongey (2022) explore a model with firm heterogeneity in monopsony power to pursue a normative analysis of the long-run optimal level of the minimum wage. We differ in three key respects from Berger, Herkenhoff and Mongey (2022). First, critical to our analysis is assessing the distributional impact of the minimum wage. As a result, we develop a model with rich worker heterogeneity. In contrast, Berger, Herkenhoff and Mongey (2022) abstract from this within-education group heterogeneity and instead mainly focuses on on firm heterogeneity arising from differences in firm productivity. Second, we highlight how putty-clay capital technology generates a meaningful distinction between the short- and long-run effects of the minimum wage, implying that one cannot infer its long-run effects from short-run estimates. Third, given our focus on the effects of the minimum wage across workers with different productivity, we show that existing transfer policies like the EITC dominate the minimum wage for low-income workers.

2 Model

Our model is characterized by four key features. First, we incorporate the notion of firm monopsony power in labor markets by allowing workers to view jobs at different firms as imperfectly substitutable with each other. This features allows for the possibility that large minimum wage increases can actually increase employment levels by potentially reducing the monopsony distortion in the labor market. Second, we embed this setup in an economy with labor markets subject to search frictions to be able to distinguish among the employment, unemployment, and nonparticipation effects of minimum wage policies. In doing so, our model results can guide researchers on the potential margins of adjustment to look for in response to large minimum wage changes. Third, we include in the model rich worker heterogeneity so as to explore the distributional impacts of the minimum wage. Finally, we develop a production technology consistent with the evidence on the substitutability between capital, skilled labor (college-educated workers) and non-skilled labor (non-college educated workers) both in the short and in the long run. Incorporating such features into our model allows for realistic margins of capital adjustment and worker substitution on the part of firms in response to changes in the minimum wage.

Formally, we consider a dynamic general equilibrium model in which labor markets are subject to matching frictions. In particular, we extend the standard competitive search framework to incorporate firm monopsony power. We allow for two groups of workers – low-educated and high-educated – as well as rich worker heterogeneity within each group. Firms operate technologies in which capital is complementary with high-educated labor and substitutable with low-educated labor, and of the putty-clay type. Specifically, we capture in an intuitive way the notion that adjusting capital is costly in the short run by assuming that technology is embedded in the capital stock so its labor intensity is irreversible. The idea is that a new machine can be built to be used in combination with low- and high-educated workers of any ability in any ratios, but these ratios are fixed once the machine is installed. In our quantitative exercise, we show how the presence of putty-clay capital slows down the transition of the economy to the new steady state after the introduction of the minimum wage and thus helps the model to reproduce the well-documented feature that employment responses to increases in the minimum wage tend to be muted in the short run but can be much larger in the long run. The following subsection discusses our model elements in greater detail.

2.1 Preferences, Production and Matching Technology

Consumer Heterogeneity and Preferences. Consumers are heterogeneous in two ways. First, workers differ in their group *education level* $g \in \{\ell, h\}$, where ℓ denotes low-educated consumers (those with less than a bachelor's degree) and h denotes high-educated consumers (those with a bachelor's degree or more). Second, within each group g, consumers are characterized by an *ability level* z drawn from group-specific discrete distributions. This second form of heterogeneity allows us to match the wage distribution within each education group. We index a consumer by i, which denotes a consumer's education group and ability level so that $i \in I \equiv I_\ell \cup I_h$, where $I_\ell = \{i | \; z_i \in \{z_{\ell 1}, \ldots, z_{\ell M}\}\}$ represents the set of abilities of low-educated consumers and $I_h = \{i | \; z_i \in \{z_{h1}, \ldots, z_{hM}\}\}$ represents the set of abilities of high-educated consumers. As shorthand, we let $i = (g, z_i)$ denote an education-ability pair.

The economy consists of a measure μ_i of families of each type i.[7] Each type of family consists of a large number of household members of the same skill group and ability level. Risk sharing within such families implies that each household member consumes the same amount of goods at date t, regardless of the idiosyncratic shocks that such a member experiences.[8]

The utility function of a family of type i is $\sum_{t=0}^{\infty} \beta^t u(c_{it}, n_{it}, s_{it})$, where c_{it} is the consumption of a representative family member, n_{it} is the index of the disutility of work of the family, and $s_{it} = \sum_j s_{ijt}$ are the total searchers where s_{ijt} denotes the number of family members of type i searching for jobs at firm j in period t. The index of the disutility of work is defined as

$$n_{it} = \left[\sum\nolimits_j n_{ijt}^{\frac{1+\omega}{\omega}} dj\right]^{\frac{\omega}{1+\omega}} \text{ with } \omega > 0, \tag{1}$$

where n_{ijt} is the number of family members who work at firm j in t. The parameter ω measures the imperfect substitutability of employment at different firms in terms of workers' disutility of work and can be interpreted as arising from workers' idiosyncratic preferences over firms, locations, or amenities. The smaller is ω, the less substitutable are jobs at a same firm. As we highlight below, ω is a key parameter in our quantitative analysis given that it governs the extent of monopsony power that firms have in the labor market.[9]

Firm Production. Consider first the production technology that will underlie the version of our model with standard capital and the version with putty-clay capital. In our economy, a large number of identical firms indexed by j produce the same homogeneous final good. Firm j uses capital k_{jt}, an aggregate of efficiency units of low-educated labor $\bar{n}_{\ell jt}$, and an aggregate of efficiency units of high-educated labor $\bar{n}_{hjt}$. The capital accumulation law is $k_{jt+1} = (1 - \delta)k_{jt} + x_{jt}$ where δ is the

[7] Letting J denote an integer number of firms and for each J we assume there is a total measure of consumers of type i, of $\mu_i J$, we are focusing on an economy with J large enough so that it is well-approximated by $J = \infty$.

[8] This type of risk-sharing arrangement in search models is familiar from the work of Merz 1995 and Andolfatto 1996.

[9] Note that here we adapt the standard way of modeling imperfect substitutability in preferences across differentiated goods to modeling imperfect substitutability in preferences across differentiated jobs. See Berger, Herkenhoff, and Mongey (2021) and Deb, Eeckhout, and Warren (2021) for related preferences along with a discussion of various interpretations of this differentiation.

depreciation rate and x_{jt} is the investment of new capital made by firm j in period t. As noted above, consumers view the labor supplied to these different firms as differentiated.

We follow Krusell et al. (2000) in using a nested CES production function over capital k_{jt}, an aggregate of low-educated labor $\bar{n}_{\ell jt}$ and an aggregate of high-educated labor $\bar{n}_{hjt}$ of the form

$$F(k_{jt}, \bar{n}_{\ell jt}, \bar{n}_{hjt}) = \left[\psi(\bar{n}_{\ell jt})^{\frac{\rho-1}{\rho}} + (1-\psi)G(k_{jt}, \bar{n}_{hjt})^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}}, \tag{2}$$

where

$$G(k_{jt}, \bar{n}_{hjt}) = \left[\lambda k_{jt}^{\frac{\alpha-1}{\alpha}} + (1-\lambda)(\bar{n}_{hjt})^{\frac{\alpha-1}{\alpha}}\right]^{\frac{\alpha}{\alpha-1}}. \tag{3}$$

The outer nest in (2) is a CES production function over efficiency units of low-educated labor $\bar{n}_{\ell jt}$ and capital-skilled labor aggregate $G(k_{jt}, \bar{n}_{hjt})$ and the inner nest is a CES production function over capital k_{jt} and efficiency units of high-educated labor $\bar{n}_{hjt}$.[10] Note that the parameters ρ and α capture the degree of substitutability among these inputs: the larger is ρ, the more substitutable low-educated labor is with the aggregate $G(k_{jt}, \bar{n}_{hjt})$ of capital and high-educated labor, whereas the smaller α is, the more complementary high-educated labor is to capital. ρ and α will be two additional key parameters in our quantitative work given that they govern the extent to which firms are will to substitute away from low educated labor towards high educated labor and capital when low educated labor becomes more expensive from the imposition of a large minimum wage.

The labor inputs $\bar{n}_{\ell jt}$ and $\bar{n}_{hjt}$ used by firm j are CES aggregates of the labor inputs of consumers of each skill group and are given by

$$\bar{n}_{\ell jt} = \left[\sum_{i \in I_\ell} z_i(\mu_i n_{ijt})^{\frac{1+\phi}{\phi}}\right]^{\frac{\phi}{1+\phi}} \text{ and } \bar{n}_{hjt} = \left[\sum_{i \in I_h} z_i(\mu_i n_{ijt})^{\frac{1+\phi}{\phi}}\right]^{\frac{\phi}{1+\phi}}. \tag{4}$$

Here, n_{ijt} is the amount of low-educated labor per family of type i supplied to firm j and $\mu_i n_{ijt}$ is the total amount of low-educated labor from all families of type i supplied to firm j. ϕ will be a key parameter in our assessment of the distributional effects of the minimum wage because it governs the extent to which firms are willing to substitute across z-types within an education group. If some low z workers become more expensive firms have the option to substitute towards higher z types in

[10] As Krusell et al. (2000) discuss, an alternative nesting of these three factors has the form $F(\bar{n}_{hjt}, \bar{G}(k_{jt}, \bar{n}_{\ell jt}))$ used in Stokey (1996). The parameters of either nesting can be chosen so as to imply any elasticity for low-educated labor and capital and for high-educated labor and capital. In particular, they can be chosen so that, as in the data, low-educated labor and capital are substitutes whereas high-educated labor and capital are complements. Given this setting, the nesting in (2) that the elasticity of low-educated labor and high-educated labor is the same as that between low-educated labor and capital, which is consistent with the data, since low-educated labor and high-educated labor are estimated to be substitutes. The alternative nesting instead implies that the elasticity of low-educated labor and high-educated labor is the same as that between high-educated labor and capital, which is not consistent with the data.

production; the strength of that substitution is governed by ϕ.

Matching Technology. Consider next the technology by which matches between firms and consumers in the labor market are formed. We consider a directed search setting in which each firm j posts a measure of vacancies $\mu_i a_{ijt}$ directed at consumers of type i with measure $\mu_i s_{ijt}$ searching for jobs at firm j, where a_{ijt} denotes the measure of vacancies posted by firm j per family of type i and s_{ijt} denotes the number of searchers per family i. The cost of posting a measure $\mu_i a_{ijt}$ of vacancies for type i consumers is $\kappa_i \mu_i a_{ijt}$. The matches created by a measure $\mu_i a_{ijt}$ of vacancies and a measure $\mu_i s_{ijt}$ of searchers are determined by the constant-returns-to-scale Cobb-Douglas matching function

$$m(\mu_i a_{ijt}, \mu_i s_{ijt}) = B_i(\mu_i a_{ijt})^{\eta}(\mu_i s_{ijt})^{1-\eta}. \tag{5}$$

If firm j posts $\mu_i a_{ijt}$ vacancies for type i consumers and, in total, families of type i send $\mu_i s_{ijt}$ consumers searching for that firm, then firm j creates a measure $m(\mu_i a_{ijt}, \mu_i s_{ijt}) = \lambda_f(\theta_{ijt})\mu_i a_{ijt}$ of new matches with consumers of type i, where $\lambda_f(\theta_{ijt}) = m(\mu_i a_{ijt}, \mu_i s_{ijt})/\mu_i a_{ijt} = m(a_{ijt}, s_{ijt})/a_{ijt}$ is the probability of that a posted vacancy is filled or *job-filling rate* where $\theta_{ijt} = a_{jit}/s_{ijt}$ denotes market tightness. We can also express these new matches as $m(\mu_i a_{ijt}, \mu_i s_{ijt}) = \lambda_w(\theta_{ijt})\mu_i s_{ijt}$, where $\lambda_w(\theta_{ijt}) = m(\mu_i a_{ijt}, \mu_i s_{ijt})/\mu_i s_{ijt} = m(a_{ijt}, s_{ijt})/s_{ijt}$ is the probability that a consumer of type i finds a job with firm j, referred to as the *job-finding rate*. The directed search environment ensures that labor market allocations will still be pinned down even when the minimum wage binds for some workers.

Timing. The timing of events within a period is as follows. We assume that each period t consists of two stages. In stage 1, each firm j posts vacancies $\{a_{ijt}\}$ for consumers of type i that determines the tightness $\{\theta_{ijt}\}$ of the markets for such consumers, and commits to present-value wage offers $\{W^F_{ijt+1}\}$ for each consumer of type i who is hired in t and begins to work in $t+1$. Each family chooses the number of searchers $\{s_{it}\}$ looking for jobs. In stage 2, after having observed all firms' offers, families decide on the search plan $\{s_{ijt}\}$ of its members looking for jobs where $s_{it} = \sum_j s_{ijt}$. For a family of type i, such a plan specifies the number of consumers s_{ijt} who search for each firm j when confronted with the offers $\{\theta_{ijt}, W^F_{ijt}\}$. These two stages should be thought of as occurring at the beginning of each period t.

2.2 A Family's Problem

2.2.1 Preferences and Optimality

Consumers of any family i face the risk of not finding a job when looking for one and of losing a job when employed. But since there are no aggregate shocks and families consist a large number of members, there is no aggregate uncertainty at the family level. So our economy is a deterministic one in terms of aggregates. Accordingly, the date 0 budget constraint of family i is

$$\sum\nolimits_{t=0}^{\infty} Q_{0,t} c_{it} \leq \sum\nolimits_{t=0}^{\infty} Q_{0,t} \sum\nolimits_{j} W_{ijt}^{F} \lambda_w(\theta_{ijt-1}) s_{ijt-1} + \psi_i \Pi_0, \tag{6}$$

where $Q_{0,t}$ denotes the price of the homogenous good at date t in units of that good at date 0, W_{ijt}^{F} is the present value of wages of newly employed workers, Π_0 is the present value of all the firms' profits, and ψ_i is the share of those profits of the firms owned by a a family of type i. To better understand the first term on the right hand side of the equation, note that if s_{ijt-1} consumers in a family of type i search for jobs at firm j in period $t-1$, then $\lambda_w(\theta_{ijt-1})s_{ijt-1}$ of them find a job. Then, these consumers start working in period t and earn the present value of wages W_{ijt}^{F} in units of period-t goods. Since a consumer of type i employed at firm j in period t separates from it in period $t+1$ with probability σ, the transition law for consumers of type i working at firm j is

$$n_{ijt+1} = (1-\sigma) n_{ijt} + \lambda_w(\theta_{ijt}) s_{ijt} \text{ for all } j, \tag{7}$$

where $\lambda_w(\theta_{ijt})$ is the job-finding rate at firm j in t for a type i consumer.

At time 0, a family of type i chooses consumption c_{it}, the number $\{s_{ijt}\}$ of its members looking for jobs across firms, and the number $\{n_{ijt+1}\}$ of its members who are employed across firms subject to the budget constraint (6), the transition law for employment at each firm (7), and a nonnegativity constraint on the number of searchers $s_{ijt} \geq 0$, where we substitute out $s_{it} = \sum_j s_{ijt}$ and n_{it} using (1) for all t in order to maximize the present value of utility.

Dropping the i subscript designating a consumer's type for clarity and letting γ, $\beta^{t+1}\mu_{jt+1}$, and $\beta^t \chi_{jt}$ be, respectively, the multipliers on the budget constraint, the transition law for employed consumers, and the nonnegativity constraint on $\{s_{ijt}\}$, the first-order conditions for the problem of a family of type i with respect to consumption, the number of employed, and the number of searchers imply

$$\beta \frac{u_{ct+1}}{u_{ct}} = Q_{t,t+1}, \tag{8}$$

$$\frac{\mu_{jt+1}}{u_{ct+1}} = \frac{u_{nt+1}}{u_{ct+1}} \left(\frac{n_{jt+1}}{n_{t+1}}\right)^{\frac{1}{\omega}} + \beta(1-\sigma) \frac{\mu_{jt+2}}{u_{ct+2}} \frac{u_{ct+2}}{u_{ct+1}}, \tag{9}$$

$$-\frac{u_{st}}{u_{ct}} = \lambda_w(\theta_{jt})\frac{\beta u_{ct+1}}{u_{ct}}\frac{\mu_{jt+1}}{u_{ct+1}} + \frac{\beta u_{ct+1}}{u_{ct}}\lambda_w(\theta_{jt})W^F_{jt+1} + \frac{\chi_{jt}}{u_{ct}}, \tag{10}$$

where we have used $\beta^t u_{ct} = \gamma Q_{0,t}$ to derive (8), which is the standard Euler equation for consumption. To understand the next two equations, recall that when consumers search for jobs at firm j in t, the consumer's family is paid the present value of wages W^F_{jt+1} for a commitment to work until an exogenous separation from that firm occurs which, in each period, happens with probability σ. In (9), μ_{jt+1} is the discounted marginal disutility from a marginal increase in the number of the family's consumers who work at firm j in $t+1$, of whom $(1-\sigma)$, $(1-\sigma)^2$, ..., of whom are still employed in $t+2$, $t+3$, and so on. To express this disutility in consumption units, we define $W^N_{jt+1} \equiv \mu_{jt+1}/u_{ct+1}$ and substitute $Q_{t+1,t+2} = \beta u_{ct+2}/u_{ct+1}$ into (9) to write W^N_{jt+1} recursively as

$$W^N_{jt+1} = \frac{u_{nt+1}}{u_{ct+1}}\left(\frac{n_{jt+1}}{n_{t+1}}\right)^{\frac{1}{\omega}} + Q_{t+1,t+2}(1-\sigma)W^N_{jt+2}. \tag{11}$$

Now, (11) is a rearranged version of the first-order condition (9). Further substituting $\mu_{jt+1}/u_{ct+1} = W^N_{jt+1}$ and $Q_{t,t+1} = \beta u_{ct+1}/u_{ct}$ into (10) gives

$$-\frac{u_{st}}{u_{ct}} = Q_{t,t+1}\lambda_w(\theta_{jt})(W^F_{jt+1} + W^N_{jt+1}) + \frac{\chi_{jt}}{u_{ct}} \text{ for all } j. \tag{12}$$

To understand this condition, note that a marginal increase in the number of consumers who search for job at firm j in t leads to a corresponding increase in the disutility from searching u_{st}/u_{ct} when expressed in consumption units. This term is the left side of (12). The benefit of incurring this cost is that with probability $\lambda_w(\theta_{jt})$ such consumers find jobs in period $t+1$, and receive the present value of wages W^F_{jt+1} in units of time $t+1$-consumption goods, net of the present value of the disutility of working W^N_{jt+1}. Expressed in period-t consumption units, this expected net benefit is $Q_{t,t+1}\lambda_w(\theta_{jt})(W^F_{jt+1} + W^N_{jt+1})$, which corresponds to the first term on the right side of (12). For consumers who actively search in period t in that $s_{jt} > 0$, it follows that $\chi_{jt} = 0$, and so the last term on the right side of (12) is zero. Hence, for consumers who actively search for jobs at firm j in t, (12) implies that the value of doing so must be be at least as high as searching for any other firm j' so that,

$$\lambda_w(\theta_{ijt})[W^F_{ijt+1} + W^N_{ijt+1}] \geq \mathcal{W}_t \equiv \max_{j'}\{\lambda_w(\theta_{ij't})(W^F_{ij't+1} + W^N_{ij't+1})\}. \tag{13}$$

When we set up the firm's problem, when firm j makes employment offers to consumers, it understands that it will only attract them if this constraint is satisfied. Hence, this constraint, which arises from consumers' optimal search behavior will be the key constraint on firms when they make wage and vacancy decisions.

2.2.2 The Participation Constraint

We turn now to examine how this constraint specializes in our symmetric equilibrium and how it encodes the monopsony power of firms. In this equilibrium, each firm needs to anticipate what happens if it deviates from a symmetric allocation. Specifically, consider an allocation in which all firms but one, say firm j, offer the common value $\lambda_w(\theta_{it})\left(W^F_{it+1}+W^N_{it+1}\right)$ to type i consumers and firm j contemplates offering a potentially different value $\lambda_w(\theta_{ijt})(W^F_{ijt+1}+W^N_{ijt+1})$. Then, for firm j to attract a consumer, it must offer at least that common value. That is, firm j's offer must satisfy

$$\mathcal{W}_t(\theta_{ijt}, W^F_{ijt+1}) \equiv \lambda_w(\theta_{ijt})(W^F_{ijt+1}+W^N_{ijt+1}) \geq \mathcal{W}_t = \mathcal{W}_t(\theta_{it}, W^F_{it+1}) = \lambda_w(\theta_{it})\left(W^F_{it+1}+W^N_{it+1}\right). \tag{14}$$

We refer to this constraint as the *participation constraint* and note that is the symmetric version of (19).

The second stage of an equilibrium summarized by this constraint (14). Solving forward the recursive expression (11) for the discounted marginal disutility resulting from a marginal increase in the number of family i's members who work at firm j in $t+1$, of whom $(1-\sigma)$, $(1-\sigma)^2$, ..., are still employed in $t+2$, $t+3$, and so on, yields

$$W^N_{ijt+1} = \frac{u_{nit+1}}{u_{cit+1}}\left(\frac{n_{ijt+1}}{n_{it+1}}\right)^{\frac{1}{\omega}} + Q_{t+1,t+2}(1-\sigma)\frac{u_{nit+2}}{u_{cit+2}}\left(\frac{n_{ijt+2}}{n_{it+2}}\right)^{\frac{1}{\omega}} + \ldots \tag{15}$$

Monopsony power affects a firm's problem through the derivatives of W^N_{ijt+1} with respect to vacancies a_{ijt} and market tightness θ_{ijt}. (See Appendix for details.)

In this dynamic economy, there is no static supply curve as in the traditional textbook static model of labor supply, rather the supply curve in period t is a dynamic object that depends on wage and market tightness in period t as well as the expectations of these variables in all future periods. Nonetheless, we can gain some intuition about it by focusing on the steady state and assuming that the preferences $u(c_i, n_i, s_i)$ have the GHH form $U(c_i - v(n_i) - h(s_i))$ that we will use in our quantitative analysis. In this case, we can write the participation constraint in a steady state as

$$\frac{\lambda_w(\theta_{ij})}{r+\sigma}\left[w_{ij} - v'(n_i)\left(\frac{n_{ij}}{n_i}\right)^{\frac{1}{\omega}}\right] = \mathcal{W}_i.$$

Holding fixed market tightness θ_{ij}, we can differentiate this constraint with respect to w_{ij} and n_{ij} to find that

$$\frac{dw_{ij}}{dn_{ij}} = \frac{1}{\omega}\frac{v'(n_i)}{n_i}\left(\frac{n_{ij}}{n_i}\right)^{\frac{1}{\omega}-1} > 0. \tag{16}$$

In this sense, the (inverse) supply curve slopes upward in n_j. Moreover, as ω converges to infinity this slope converges to zero.

In sum, monopsony power affects a firm's first-order conditions for offered wages and vacancies through the derivatives of the participation constraint, which capture how the present value of the disutility of work changes when any firm j changes its wages and vacancies, and so market tightness, holding fixed the value of searching in the common market. In Appendix A, we discuss how the upward-sloping labor supply curve for each firm's jobs that our model gives rise to is analogous to the downward-sloping demand curve for each firm's goods that arises in models of monopolistic competition.

2.3 A Firm's Problem

We consider two versions of the model that differ by the type of capital that firms use. In the first version, we assume that firms use a standard type of capital, often referred to as *putty-putty* capital, which is homogeneous and whose substitutability with other inputs is the same before and after it is installed. Intuitively, if we think of each piece of capital as a machine, then this assumption implies that even after the machine is built, it is possible to alter it to be used with different amounts of labor. An issue with this version of the model is that it will have predictions for employment in response to an increase in the minimum wage that are inconsistent with the data. In particular, for the small increases in the minimum wage documented for the United States, changes in employment in the short run have been negligible. Counterfactually, the model with standard capital implies that similarly-sized increases in the minimum wage lead to non-neglible changes in employment in the short run.

Motivated in part by this issue, we consider a version with putty-putty capital, in which after a machine is built, it is not possible to alter its labor intensity. This feature will imply that in the short-run employment does not react much to modest increases in the minimum wage, hence, is more in line with the data.

In both versions, firms purchase capital and hire workers in frictional labor markets in order to produce. As discussed, we adopt a directed search approach to model how firms and consumers match to form employment relationships. Specifically, we assume that for each type of consumer i, firm j chooses both the measure of vacancies $\{a_{ijt}\}$ to post in t and the present value of wages W^F_{ijt+1} to offer. Consumer who match with this firm in t begin work in $t+1$. To capture the notion that a minimum wage can help offset the distortionary monopsony power of firms in labor markets, we develop a *monopsonistic directed search* setup that generalizes the market utility approach to competitive search introduced by Montgomery (1991). (For competitive directed search see also Moen (1997) and, for an extensive review of the literature, Wright et al. (2021).)

2.3.1 A Firm's Problem with Standard Capital

Consider the firm's problem with standard capital. Given an initial capital stock k_0 and an exogenous sequence of prices of investment goods $\{q_t\}$ expressed in units of consumption goods, each firm chooses sequences of tightnesses $\{\theta_{ijt}\}$ for markets for consumers of type i, the measures of vacancies $\{\mu_i a_{ijt}\}$ to post aimed at consumers of type i, the measures of consumers of type i to employ $\{\mu_i n_{ijt+1}\}$, the present values of wages $\{W^F_{ijt+1}\}$ for these consumers, and new capital $\{k_{t+1}\}$ in order to maximize

$$\sum_{t=0}^{\infty} Q_{0,t} \left\{ F\left(k_t, \bar{n}_{\ell t}, \bar{n}_{ht}\right) - q_t x_t - \sum_i W^F_{ijt} \lambda_f(\theta_{ijt-1}) \mu_i a_{ijt-1} \right\}, \tag{17}$$

subject to the law of motion for capital $k_{t+1} = (1-\delta)k_t + x_t$, the transition laws for employment for consumers of type i

$$\mu_i n_{ijt+1} \leq (1-\sigma)\mu_i n_{ijt} + \lambda_f(\theta_{ijt})\mu_i a_{ijt} \text{ all } i \tag{18}$$

and the *participation constraints* for attracting consumers of type i

$$\lambda_w(\theta_{ijt})(W^F_{ijt+1} + W^N_{ijt+1}) \geq \mathcal{W}_t \equiv \max_{j'}\{\lambda_w(\theta_{ij't})(W^F_{ij't+1} + W^N_{ij't+1})\}. \tag{19}$$

For this economy, given an exogenous sequence of investment goods prices, $\{q_t\}$, a *monopsonistically competitive search equilibrium* with standard capital and $k_{j0} = k_0$ for all j is a collection of allocations of consumption, employment, searching, capital $\{c_{it}, n_{it}, s_{it}, \bar{n}_{it}, k_t\}$, vacancies and market tightnesses $\{a_{it}, \theta_{it}\}$, and prices $\left\{W^F_{it+1}, Q_{0t}\right\}$ such that at these prices and allocations i) consumers' decisions are optimal for each family i, ii) firms' decisions are optimal, and iii) markets clear.

Note for later that since families can perfectly insure the idiosyncratic risk their members face and there are no aggregate shocks, it is without loss of generality to adopt the convention that a firm fulfills its present-value wage offer W^F_{ijt} by offering a constant period wage w_{ijt} over the course of a match that begins at t so that

$$W^F_{ijt} = w_{ijt} + (1-\sigma)Q_{t,t+1}w_{ijt} + (1-\sigma)^2 Q_{t,t+2}w_{ijt} + \ldots, \tag{20}$$

where $Q_{t,s}$ is the price of goods at $s > t$ in units of goods at t. Note that $W^F_{ijt} = d_t w_{ijt}$ where $d_t \equiv [1 + (1-\sigma)Q_{t,t+1} + (1-\sigma)^2 Q_{t,t+2} + \ldots]$ so that we can equivalently think of the firm as choosing W^F_{ijt} or w_{ijt}.

2.3.2 A Firm's Problem with Putty-Clay Capital

Now consider an environment in which capital is of the putty-clay type—we drop the subscript j denoting a firm for notational simplicity. The idea behind this version of the model is most easily understood when all low-educated consumers and all high-educated consumers have the same ability, so that there are only two types of consumers. Ex-ante capital is putty-like in that it is possible to build a machine with any ratio of low-educated and high-educated labor to capital that lies on the frontier of the production function in (2), that is, the output technology is CES ex ante. However, once the machine is built, it is clay-like in that it uses a fixed amount of low-educated labor and high-educated labor to operate, that is, the output technology is Leontief ex-post. Hence, in the short run given a stock of machines, demand for low-educated labor and high-educated labor is inelastic as long as total profits from operating the machine are positive, because a firm cannot substitute between existing capital and either type of labor. Over time though, new machines embodying new labor to capital ratios can be installed so that in the long run, firms can substitute away from the type of labor that becomes more expensive, in our case low-educated labor because of the minimum wage, towards both high-educated labor and capital.

More formally, consider the case of interest in which low-educated and high-educated workers differ in their ability level, z. With many types of labor i, a capital type $v = \{v_i\}$ denotes the skill and ability intensity of capital, that is, how much labor of each skill and ability capital needs in order to operate. Each v_i then specifies the type-i labor to capital ratio necessary to run a machine of type v at full utilization. As a result, $k(v)$ units of capital of type v provide $k(v)$ units of capital services only if, for all i, this capital is combined with at least $n_i = k(v)v_i$ units of labor for all i. If $n_i > k(v)v_i$, then the excess workers remain idle whereas if $n_i < k(v)v_i$, then the excess capital remains idle. If $k(v)$ units of capital is combined with $n_i = k(v)v_i$ units of labor for all i, then $f(v)$ units of output is produced, where

$$f(v) = F(k, \{n_i\})/k = F(1, \{n_i/k\}) = F(1, \{v_i\})$$

and F is our production function. More generally, if some arbitrary amount of labor $\{n_i(v)\}$ is combined with $k_i(v)$ units of capital of type v, then the total output produced with type-v capital is $y(v) = \min [k(v), \{n_i(v)/v_i\}] f(v)$. The total output of a firm in period t is thus

$$y_t = \int_v \min [k_t(v), \{n_{it}(v)/v_i\}] f(v) dv.$$

Firms invest $x_t(v)$ units of output to accumulate type-v capital according to the capital accumulation law

$$k_{t+1}(v) = (1-\delta)k_t(v) + x_t(v) \tag{21}$$

subject to the nonnegativity constraints $x_t(v) \geq 0$.

Given some initial vector of capital $\{k_0(v)\}$ that a firm owns and an exogenous sequence of prices of investment goods $\{q_t\}$ expressed in units of consumption goods, the firm chooses sequences of market tightnesses $\{\theta_{it}\}$, vacancies $\{a_{it}\}$, workers $\{n_{it+1}(v)\}$ for each type of capital v, present value of wages $\{W_{it+1}^F\}$, and investment $\{x_t(v)\}$ for each type of capital in order to maximize

$$\sum_{t=0}^{\infty} Q_{0,t} \left\{ \int_v [F(k_t(v), \{n_{it}(v)\}) - q_t x_t(v)] dv - \sum_i \left[W_{it}^F \lambda_f(\theta_{it-1}) a_{it-1} - \kappa_i a_{it} \right] \right\}, \tag{22}$$

subject to the transition laws for workers (18), the participation constraints for employed workers (19), along with the transition law for each type of capital (21), the adding up constraints of uses of labor of type i, $n_{it} \leq \int n_{it}(v)dv$, the Leontief constraints on labor

$$n_{it}(v) \leq v_i k_t(v), \tag{23}$$

and the nonnegativity constraints on each type of investment $x_t(v) \geq 0$. To understand the constraints in (23), note that with a capital stock $k_t(v)$ of type $v = (v_i)$ that can be used to produce the output

$$y_t(v) = \min \left[k_t(v), \left\{ \frac{n_{it}(v)}{v_i} \right\} \right] f(v),$$

if, say, the firm uses $n_{it}(v)$ units of type i labor with it such that $n_{it}(v) > v_i k_t(v)$, then the excess labor $n_{it}(v) - v_i k_t(v)$ is wasted, so this is never optimal. Hence, we can impose the Leontief constraints in (23) and drop the min function. The non-negativity constraint $x_t(v) \geq 0$, implies that firms cannot disassemble their existing types of capital. This friction is key to generating frictional adjustment–without it the model would collapse to the putty-putty model described in the previous section.

2.4 Steady State

Here we discuss the steady state and some implications from it. It is immediate that the steady state of the model with standard capital and the model with putty-clay capital are identical. Intuitively, once factor prices and intertemporal prices of consumption that firms face become constant, firms start investing in the unique type capital that is ideally suited to their technologies at those prices and let all past capital depreciate away. Eventually, all the old capital stock is replaced, and in a steady state with putty-clay capital firms invest in exactly the same type of capital as they would have in a steady state with standard capital.

In our baseline model, we focus on Greenwood, Hercowitz and Huffman (1988) (GHH) preferences with

$$u(c_i, s_i, n_i) = U\left[c_i - v(n_i) - h(s_i)\right] \tag{24}$$

for family i. Consider the steady state of this version of the model in which all variables, including the price of capital are constants. As we show in the Appendix, in the steady state, the firm's Euler equation for capital is

$$q\left[\frac{1}{\beta} - (1-\delta)\right] = F_k, \tag{25}$$

the firm's vacancy posting condition is

$$\frac{\kappa}{\lambda_f(\theta_i)} = \frac{F_i - w_i - v'(n_i)/\omega}{r+\sigma}, \tag{26}$$

a family's first-order condition for the number of searchers is

$$h'(s_i) = \frac{\lambda_w(\theta_i)\left[w_i - v'(n_i)\right]}{r+\sigma}, \tag{27}$$

and equilibrium wages satisfy

$$w_i = \eta\left[F_i - \frac{v'(n_i)}{\omega}\right] + (1-\eta)v'(n_i) \tag{28}$$

where $F_k = F_k(k, \bar{n}_\ell(n_i), \bar{n}_h(n_i))$,$F_i = F_i(k, \bar{n}_\ell(n_i), \bar{n}_h(n_i))$, and $r = 1/\beta - 1$. Finally, the steady-state law of motion for employment reduces to

$$\lambda_w(\theta_i)s_i = \sigma n_i \tag{29}$$

where $\bar{n}_\ell$ and $\bar{n}_h$ satisfy the symmetric steady-state version of (4). Consumption satisfies a steady state version of the budget constraint. With these preferences, which imply no income effects, the steady state equations split into blocks. First, we can solve for the monopsonistically competitive equilibrium wages $\{w_i\}$ and the associated allocations $\{\theta_i, s_i, n_i\}$ and k from (25)-(29). Then given these values we can solve for consumption $\{c_i\}$ from the budget constraint.

Notice that firms' monopsony power distorts the wage equation and this distortion is captured by the term $v'(n)/\omega$. This term implies that for a given marginal product of labor and marginal disutility of work, a firm offers a smaller wage than under perfect competition, in which case $v'(n)/\omega = 0$ since $\omega = \infty$. As the search first order condition (27), the inefficiently low level of wages results in consumers searching too little for jobs. In the firm's vacancy posting condition, (26), there is both the indirect distortion from the inefficient level of wages and the direct distortion in terms of the presence of the term $v'(n)/\omega = 0$. So in equilibrium, these distortions imply that firms post too few vacancies and consumers search too little for jobs and both wages and employment are lower than in the competitive search case.

A firm with monopsony power pays its workers only a fraction of their marginal product. A simple measure of firms' monopsony power is then the markdown of wages relative to marginal product, $-(w_i/F_{ni} - 1)$, namely (the negative of) the percentage difference between a worker's wage and marginal product ($-(w_i - F_{ni})/F_{ni}$). In a slight abuse of language, we refer to w_i/F_{ni} as the *wage markdown*. We can combine the vacancy posting condition and the wage equation to show that the implied markdown for workers of a family of type i is given by

$$\frac{w_i}{F_{ni}} = \left[1 + \underbrace{\frac{\kappa(\theta_i)}{\frac{\eta}{1-\eta}\kappa(\theta_i) + v'(n_i)}}_{\text{efficient component}} + \underbrace{\frac{\frac{1}{\omega}v'(n_i)}{\frac{\eta}{1-\eta}\kappa(\theta_i) + v'(n_i)}}_{\text{monopsony component}}\right]^{-1}, \tag{30}$$

where $\kappa(\theta_i) = (r+\sigma)\kappa_i/\lambda_f(\theta_i)$. The *efficient component* of the markdown also arises under perfect competition and corresponds to the amount needed for firms to recoup their vacancy posting costs and, hence, earn zero expected profits per vacancy. More interesting is the *monopsony component*, which captures that firms with monopsony power set wages below their competitive level and, hence, their markdowns are larger than the competitive ones. In our quantitative model, we will find that the overwhelming majority of the markdown in is due to the monopsony distortion.

Consider for a moment a simple policy implication of this analysis. Since the distortions to the economy emanate from firms paying too low a wage for each consumer type, a policy that mandates that firms must pay a *type-specific* minimum wage $\underline{w}_i$ for each consumer type i, where $\underline{w}_i$ is set to equal the level in the competitive search equilibrium for that worker type, would fix the distortions and lead to efficient allocations. More precisely, define the competitive equilibrium wages $\{w_i^*\}_i$ and the associated allocations $\{\theta_i^*, s_i^*, n_i^*, c_i^*\}$ and k^* that satisfy the competitive wage equation

$$w_i^* = \eta F_i^* + (1-\eta)v'(n_i^*) \tag{31}$$

along with the conditions (25)-(27), (29) and a steady state version of the budget constraint.

Also note that since the only distortion is from the degree of monopsony power, as we consider a sequence of economies in which this monopsony power goes to zero, in that ω converges to infinity, then the monopsony wages converges to the competitive search wages and the allocations converge to those in the competitive search equilibrium.

We summarize this discussion with as follows.

***Proposition* 1.** *If a minimum wage for each type is set equal to the competitive wage for that type and that constraint binds, that is, $\underline{w}_i = w_i^*$, then wages and allocations in the minimum wage economy coincide with those in the competitive search equilibrium. Also, as ω converges to infinity, the wages and allocations in the monopsonistic competition economy converge to those in the competitive search equilibrium.*

To proof of the first part of the proposition is in the appendix. For the second part of the proposition note that as ω converges to infinity, the monopsony wage (28) converges to the competitive search wage (31). Hence, the distortions in wages go to zero and, from an inspection of (26) and (27) so do the distortions in the the vacancy posting condition and search first-order conditions induced by these distortions.

The issue we will address later is that although a very rich set of type-specific minimum wages could fix the distortions induced by firms' monopsony power, in practice setting such a complex system of minimum wages is infeasible. Instead, we shall analyze the other extreme, which corresponds to the minimum wage policies advocating in practice, namely, of only one mandated minimum wage for all workers. In such a scenario, if there are large enough differences in skill and ability across workers, then a single minimum wage can have perverse distributional effects.

3 Calibration

We calibrate the model in order to match key features of the U.S. labor market that inform the mechanisms described above. We then validate the calibrated model by showing that it reproduces non-targeted features of the data fairly well.

3.1 Utility Function

Before turning to our calibration, we complete the description of our model by assuming that the utility function takes the Greenwood, Hercowitz and Huffman (1988) form

$$u(c_{it}, n_{it}, s_{it}) = \sum_{t=0}^{\infty} \beta^t \log \left(c_{it} - \chi_{b,n} \frac{n_{it}^{1+1/\gamma}}{1+1/\gamma} - \chi_{b,s} \frac{s_{it}^{1+1/\gamma}}{1+1/\gamma} \right), \tag{32}$$

where $\chi_{b,n}$ controls the disutility of working for education group $b \in \{\ell, h\}$, $\chi_{b,s}$ controls the disutility of searching, and the exponent γ is related to the elasticity of labor supply.[11]

[11] In principle, the preferences (32) may imply that households violate their time constraint by choosing $n_{it} + s_{it} > 1$. We ensure that the time constraint is always satisfied by augmenting (32) with a positive utility from leisure

3.2 Key Model Parameters

Table 1 highlights the set of key parameters that are quantitatively important for determining the distributional effects of the minimum wage in both the short and long run. The potential for a binding minimum wage to increase output and employment depends on ω which, as discussed above, governs the degree of firms' monopsony power in the labor market. As ω tends to infinity, monopsony power falls to zero and the inefficient component of the wage markdowns vanishes. As a result, we discipline ω by targeting empirical estimates of the average wage markdown in the data. A growing literature has measured wage markdowns in the United States. As a baseline, we target the value 0.71 from Berger, Herkenhoff and Mongey (2019), but we show below how our results change given the range of values estimated in the literature.

According to our KORV production structure, the parameters ρ and α govern the substitutability between low-educated workers and capital and the substitutability between high-educated workers and capital, respectively. As the minimum wage binds on low-educated workers, firms may choose to substitute away from the suddenly more expensive low-educated workers towards both capital and higher-educated workers; ρ and α discipline the quantitative magnitudes of such substitution patterns. We set the elasticities of substitution between capital and the different types of labor, α and ρ, to the values estimated in Krusell et al. (2000). In principle, these estimates are not directly applicable to our environment because our model contains other features, such as search frictions and monopsony power, which Krusell et al. (2000) abstract from. However, we show below that their values $\alpha = 0.67$ and $\rho = 1.67$ produces realistic long-run substitution patterns within our model in response to changes in the relative price of capital.[12]

The model also allows firms to substitute among workers of differing ability within an education group and such a dimension of substitutability is governed by the parameter ϕ. The literature provides little guidance about the elasticity of substitution across workers within an education group, ϕ. As a baseline, we choose a value for *phi* of 3, which is similar to existing estimates of the elasticities of substitution *across* education groups (see, for example, Katz and Murphy (1992) or more recently Bils, Kaymak and Wu (2020)). Likewise, Card and Lemieux (2001) estimate an elasticity of substitution between workers of differing ages within an education group in the range of 4 to 6.

$\chi_\ell \log(1 - n_{it} - s_{it})$ and set χ_ℓ to 0.01. Standard Inada conditions imply that households will not violate their time constraint—otherwise this term has small effects on our results. We view this specification as a technical device to ensure that the time constraint holds in each period without having to deal with an occasionally binding constraint.

[12] We use cross-industry variation and detailed micro data from the U.S. Censuses and American Community Surveys to measure how changes in the price of capital affect both the industry's capital to labor ratio and the industry's share of high-educated employment out of total employment. We then target these empirical relationships as alternate approach to discipline α and ρ. Using this procedure that is model consistent, we obtain estimates of α and ρ that are very similar to the values in KORV. In light of this similarity, we relegate these results to a robustness exercise in Appendix C.

Table 1: Key Parameters Governing Distributional Effects of Minimum Wage

Parameter	Description	Value	Discipline
Preference Parameters			
ω	Extent of firm monopsony power	2.88	Targeted: Match empirical estimates of wage markdown from the literature
γ	Labor disutility exponent	1.00	Fixed: Literature
Production Parameters			
ρ	Elasticity of substitution between low-educated workers and capital	1.67	Fixed: Literature
α	Elasticity of substitution between high-educated workers and capital	0.67	Fixed: Literature
ϕ	Elasticity of substitution across workers within an education group	3.00	Fixed: Literature
δ	Annualized depreciation rate (equipment and software)	15%	Fixed

Notes: Key parameter that are important for determining the quantitative employment response by worker type to changes in the minimum wage.

We view our choice of $\phi = 3$ as a lower bound on the substitutability across workers of the same age *within* an education group, and show that our results about the negative effects of the minimum wage for low-z workers are even starker for higher values of ϕ.

As mentioned, the exponent γ in the utility function is related to the elasticity of labor supply. Therefore, γ determines how household labor supply responds for low-z workers after the imposition of a large minimum wage. In our baseline analysis, we set γ^{-1} equal to 1. We explore the robustness of our results to alternate values of γ.

The above set of parameters are the most relevant for the long-run distributional effects of a large change in the minimum wage. Differences between the short-run and long-run responses of the economy depend on the speed with with firms adjust their capital stock. In the putty-clay version of our model, the pace with which firms alter their capital stock in response to a change in the minimum wage is determined by the depreciation rate. We set the depreciation rate δ equal to 0.014 that implies an annual depreciation rate of 15%, which matches the average rate of depreciation of equipment and software in recent years (see Appendix B for details).[13,14]

[13]We exclude structures from this calculation because they do not play a large role in capital-labor substitution; Krusell et al. (2000) make a related assumption.

[14]As a baseline, we also set the disutility of search effort to $\chi_{b,s} = 100$. This parameter controls the sensitivity of

Table 2: Other Fixed Parameters

Parameter	Description	Value
π_ℓ	Fraction of non-college households	0.69
β	Discount factor	$(1.04)^{-1/12}$
σ	Job destruction rate	2.8%
η	Elasticity of matching function w.r.t. vacancies	0.50

Notes: Other parameters (in addition to ρ, α, ϕ, and γ) exogenously fixed in the calibration. A model period is one month.

For the most part, the remaining parameters of the model are less important in determining the magnitude of employment, participation, and labor income responses to large changes in the minimum wage. In the following two subsections, we discuss how we pin down these remaining parameters

3.3 Additional Fixed Parameters

We set a model period to be one month in order to adequately capture worker flows in the U.S. labor market. Throughout the calibration, we draw on data from the 2006-2019 American Community Surveys (ACS). We restrict the sample to include all individuals aged 16 and above, exclude all individuals residing in group quarters, and exclude all individuals who report themselves as being a student (mirroring the sample restrictions used by the BLS to compute labor market statistics). See Appendix B for more details about our data construction.

Table 2 contains the remaining parameters which we exogenously fix in our calibration. We set the share of college-educated households in the population to $1 - \pi_\ell = 31\%$ in order to match data from the ACS. We set households' discount factor β to $(1.04)^{-1/12}$ so that the annualized real interest rate r equals 4%. Finally, we set the job destruction rate σ to 0.028 monthly and set the elasticity of the matching function with respect to the measure of the unemployed to $\eta = 1/2$.

3.4 Additional Fitted Parameters

Table 3 contains the other parameters—in addition to ω discussed above—which we choose to match a number of statistics from the data. We assume that the distribution of idiosyncratic productivity z is log-normal within education group $b \in \{\ell, h\}$ with group-specific variance parameter σ_b, $z \sim \log \mathcal{N}(0, \sigma_b)$. Given the rest of the model, σ_l and σ_h determine the distribution of wages *within* each

search effort with respect to changes in the returns to search, and we find that our baseline value is consistent with small responses of the unemployment rate to changes in both the relative price of capital and the minimum wage. Appendix D shows that our main results are robust to different values of $\chi_{b,s}$.

Table 3: Other Fitted Parameters

Parameter	Description	Value
Worker productivity distribution $\log \mathcal{N}(0, \sigma_b)$		
σ_ℓ	SD of non-college z	0.69
σ_h	SD of college z	0.75
Production function		
ψ	Coefficient on non-college labor n_ℓ	0.59
λ	Coefficient on capital k	0.88
Search frictions		
B	Matching function productivity	0.16
Labor Disutility		
$\chi_{\ell,n}$	Scale of non-college labor disutility	2.29
$\chi_{h,n}$	Scale of college labor disutility	3.54

Notes: Other parameters (in addition to ω) endogenously chosen to match the statistics in Table 4.

education group. As in KORV, the scale parameters ψ and λ in the production function pin down the distribution of income *across* education group and the average capital to labor ratio. The parameter B governs the efficiency of the matching function. Finally, the parameters $\chi_{b,n}$ of the disutility of work of each education group control the steady-state employment rates of each group $b \in \{\ell, h\}$.[15]

Table 4 shows the statistics that we target in our calibration. As discussed above, the degree of monopsony power ω is primarily determined by the average wage markdown in the data. For the rest of the wage distribution, we exploit specific features of the distribution of wages by education group from the pooled 2017-2019 waves of the ACS. In particular, we target the ratio of the 50th percentile to the 10th percentile of the wage distribution of each education group in order to precisely match the left tail (which is most directly affected by the minimum wage).[16] We discipline the average wage differences across education group by targeting the aggregate labor share and the "college income share", namely, the share of total labor income earned by college workers. These moments discipline the scale parameters ψ and λ in the production function. The unemployment rate is informative about the rate at which searching members of households find jobs and, therefore, the productivity of the matching function B. We target a steady state unemployment rate of 5.9% to be consistent with the pre-Great Recession data. Finally, we target the average employment rates by education groups, which are informative about the disutilities of working for those groups, $\chi_{\ell,n}$ and $\chi_{h,n}$.

Table 4 shows that the model reproduces these targets fairly well. Importantly, the model matches

[15] As with many search models, there is a set of vacancy-posting costs κ_0 and matching function productivity B which deliver similar steady-state unemployment rates. In order to resolve this indeterminacy, we exogenously fix $\kappa_0 = 0.3$ and endogenously choose B in order to match the average unemployment rate in the data. Our value of $\kappa_0 = 0.3$ is roughly two times average monthly output per newly hired worker, in line with estimated vacancy-posting and training costs in the literature. Our results are extremely similar if we instead exogenously fix $B = 1$ and endogenously choose κ_0 to match the same data.

[16] We also place a low weight on the 90-50 wage ratios within each education group in our moment-matching procedure to ensure that we do not severely overpredict the dispersion of wages at the top of the distribution.

Table 4: Targeted Statistics

Moment	Description	Data	Model
Average wage markdown			
$\mathbb{E}[w_{ni}]/\mathbb{E}[F_{ni}]$	Average wage markdown	0.71	0.72
Wage Distribution, ACS 2017-2019			
$w_{\ell 50}/w_{\ell 10}$	Non-college 50-10 ratio	2.04	1.89
w_{h50}/w_{h10}	College 50-10 ratio	2.30	1.91
Income shares			
$\mathbb{E}[w_i n_i]/Y$	Aggregate labor share	0.57	0.56
$\pi_h \mathbb{E}[w_{hz} n_{hz}]/\mathbb{E}[w_i n_i]$	College income share	0.55	0.56
Unemployment rate			
$\mathbb{E}[s_i]/(\mathbb{E}[s_i]+\mathbb{E}[n_i])$	Average unemployment rate	5.9%	5.9%
Employment Rates			
$\mathbb{E}_\ell[n_i]$	Non-college employment rate	0.47	0.48
$\mathbb{E}_h[n_i]$	College employment rate	0.62	0.62

Notes: Statistics targeted using parameters in Table 3. The *average wage markdown* is the payroll-weighted average markdown from Berger, Herkenhoff and Mongey (2021). The *wage distribution* targets are drawn from the ACS 2017-2019 data described in Appendix B. The *average labor share* is from Karabarbounis and Neiman (2014). The *college income share* is from the ACS 2017-2019 data described in Appendix B. Finally, the *employment rate* targets are also drawn from the ACS 2017-2019 data described in Appendix B.

the average wage markdown, and therefore the degree of monopsony power, almost exactly. The unemployment rate, aggregate income shares, and employment shares almost identical in the model and in the data. The average job-finding rate is 0.46 in the model, which is similar to the value of 0.45 the data (Shimer (2005)). The model somewhat underpredicts the 50-10 wage ratios, but we show below that it fits the overall wage distribution well.

3.5 Validation

Before using the model to study the minimum wage and other policies, we first verify that it reproduces non-targeted features of the data well.

Wage Distribution. Figure 2 shows the percentiles of the wage distribution both in the ACS data and our calibrated model. Within education groups, the model matches the left tail of each distribution fairly well given that we targeted the 50-10 ratio in our calibration; as a point of reference, \$15 minimum wage would bind for about 50% of non-college workers in both the model and the data. However, the required dispersion of idiosyncratic productivity σ_b and the structure of the log-normal distributions imply that the right tail of the wage distribution is more disperse in our model than in the data (the interquartile range of non-college wages is 2.09 in the data vs. 2.35 in our model, for example). We are comfortable with this tradeoff given that the effects of the minimum wage are primarily determined by the left tail of the distribution. Across education groups, the model predicts

FIGURE 2: Calibrated Wage Distribution

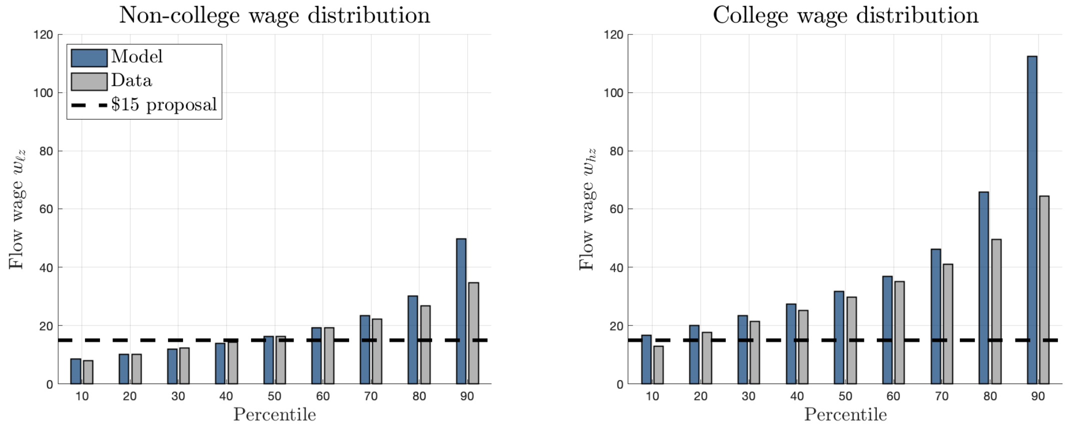

Notes: Wage distribution in calibrated model (blue bars) and the data (grey bars). The wage distribution in the data is drawn from the pooled 2017-2019 waves of the ACS, as described in Appendix B.

FIGURE 3: Calibrated Wage Markdowns

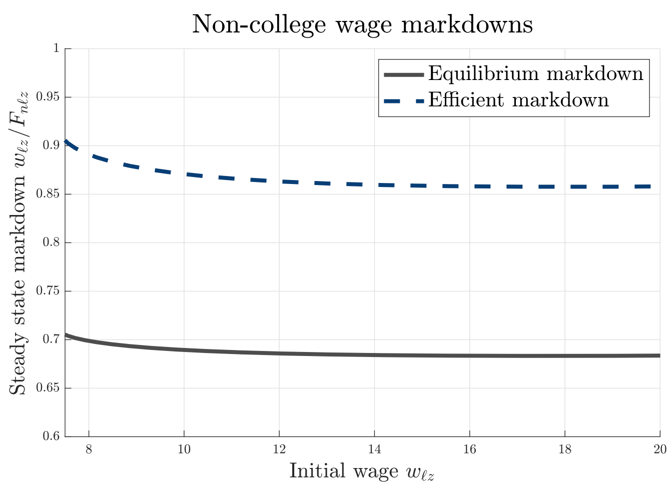

Notes: steady state wage markdowns w_i/F_{ni} of particular z types. "Equilibrium markdown" corresponds to our calibrated model. "Efficient markdown" corresponds to the model without monopsony power, i.e. $\omega \to \infty$. The x-axis is log individual productivity z, expressed in standard deviations relative to its mean.

the median college wage is 1.94 times than the median non-college wage, compared to 1.81 times higher in the data.

Figure 3 plots the model's implied wage markdowns as a function of worker productivity z. Recall from Section 2 that the steady state markdown is given by:

$$\frac{w_i}{F_{ni}} = \left(1 + \underbrace{\frac{\frac{(r+\sigma)\kappa_i}{\lambda_f(\theta_i)}}{\frac{\eta}{1-\eta}\frac{(r+\sigma)\kappa_i}{\lambda_f(\theta_i)} + v'(n_i)}}_{\text{efficient component}} + \underbrace{\frac{\frac{1}{\omega}v'(n_i)}{\frac{\eta}{1-\eta}\frac{(r+\sigma)\kappa_i}{\lambda_f(\theta_i)} + v'(n_i)}}_{\text{monopsony component}}\right)^{-1}.$$

In our calibration, the efficient component of the markdown is about 0.1 on average and therefore accounts for about one-third of the total average markdown of 0.71. This component reflects the fact that firms must recoup the annuitized cost of recruiting the worker, $\kappa_i/\lambda_f(\theta_i)$, and would exist even without monopsony power. The remaining two-thirds of the markdown is driven by the monopsony power, implying that firms earn substantial monopsony profits in equilibrium. Markdowns slightly decrease with individual productivity z because firms post more vacancies for those groups, lowering their job-filling rates $\lambda_f(\theta_i)$ and therefore increasing the vacancy-posting costs required to hire a worker $\kappa_i/\lambda_f(\theta_i)$, which decreases the monopsony component of markdowns.

Capital-Labor Substitution. Our model's long-run elasticities of substitution between capital and the different types of labor—governed by the values of ρ and α taken from Krusell et al. (2000)—imply a realistic degree of substitutability compared to long-run trends in the data. For example, over the last forty years, the relative price of equipment and software has declined by approximately one log point while the college income share has increased by more than 30 percentage points (according to our Census/ACS data). In our model, a similarly-sized decline in the relative price of investment implies the college income share increases by 13 percentage points, about a third of the aggregate decline.

These aggregate time trends are also driven by factors other than capital-skill complementarity, such as changes in educational attainment or skill-biased labor-augmenting technical change. In Appendix C, we control for these aggregate factors using sector-level variation in capital prices and college income shares to control for these aggregate factors. In particular, we estimate how the college income share has changed in sectors whose relative price of capital has declined by more compared to sectors in which it has declined by less, conditional on time fixed effects. Using this variation, we find that a one log point decline in the relative price of capital implies an increase in the college income share by 8.1 percentage points, closer to the 13 percentage points in our model.

In Appendix C, we also provide evidence in support of our model's small short-run elasticities of substitution coming from the putty-clay technology. We do so using temporary variation in the after-tax price of investment coming from the Bonus Depreciation Allowance, a countercyclical tax incentive that affects sectors in a heterogeneous fashion (Zwick and Mahon, 2017). In our model, temporary changes in the relative price of investment have a small effect on the college income share because they only induce capital-labor substitution on newly purchased capital, which is a small fraction of the aggregate capital stock. We show that the data are consistent with this prediction. In contrast, the model with standard capital implies a large response of the college income share as firms significantly change their capital-labor ratios due to intertemporal substitution, inconsistent, which is inconsistent with the data.

4 Minimum Wage

We now use our calibrated model to study the effects of a minimum wage $\overline{w}$ at both the aggregate and microeconomic level. We assume the minimum wage is unexpectedly introduced starting from the initial steady state without a minimum wage.[17] Section 4.1 studies the new steady with a positive minimum wage in order to assess its long-run consequences. Section 4.2 then shows that the putty-clay technology slows down the transition from the initial equilibrium to this new steady state.

4.1 Long-Run Effects of the Minimum Wage

The minimum wage $\overline{w}$ modifies our steady state equilibrium conditions relative to the model description in Section 2. The only difference is that, for household types i on which the minimum wage binds, we drop the equation for steady state wages, (28), and replace w_i with $\bar{w}$ in the vacancy posting condition (26) and the first-order condition for search (27). The Euler equation for capital, and the equations for households on which the minimum wage does not bind, are unchanged.

Aggregate-Level Results. We begin our long-run analysis by studying the consequences of the minimum wage on aggregate non-college employment and labor income.[18] The left panel of Figure 4 shows that aggregate employment is an inverted-U function of the minimum wage, which we denote a *Laffer curve*. As discussed in Section 2, this shape reflects the fact that a small minimum wage reduces the average monopsony distortion in the economy, bringing wages and employment closer to their efficient level. However, a large minimum wage pushes the affected workers' wages sufficiently above their efficient level, reducing employment. In our calibration, the Laffer curve peaks at about \$10 per hour and the \$15 minimum wage reduces employment by 9.2%.

[17] With our calibrated wage distributions, the current national minimum wage of \$7.25 is not binding, so that initial steady state is a useful approximation to the current policy regime.

[18] We focus on outcomes for non-college workers because, given our model's wage distribution, the minimum wage is barely binding for college workers. Appendix D shows that how college employment and labor income vary with the minimum wage, at both the aggregate and micro-level.

FIGURE 4: Aggregate Minimum Wage Laffer Curves

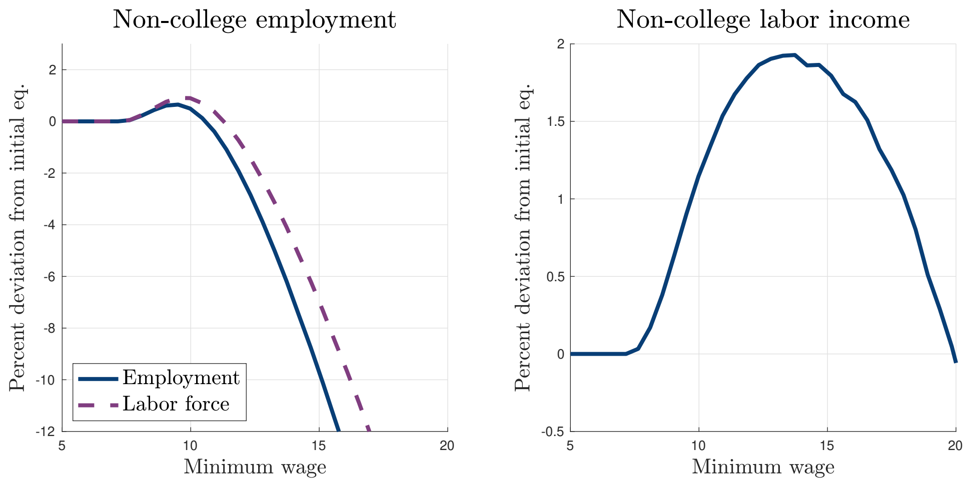

Notes: steady state outcomes as a function of the minimum wage $\overline{w}$. Left panel plots the log-change of aggregate non-college employment, the log-change of aggregate non-college labor force, and the change in the unemployment rate relative to their levels in the initial steady state with $\overline{w} = 0$ (see the decomposition (33)). Right panel plots aggregate labor income of non-college workers. The x-axis is the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

The left panel of Figure 4 also shows that the majority of this changes in aggregate employment are due to changes in the labor force participation rate rather than the unemployment rate. These outcomes are linked through the decomposition

$$\underbrace{\Delta \log n_b}_{\text{employment}} \approx \underbrace{\Delta \log(n_b + s_b)}_{\text{labor force}} - \underbrace{\Delta \frac{s_b}{s_b + n_b}}_{\text{unemployment}}, \tag{33}$$

where n_b is the aggregate employment rate of education group $b \in \{\ell, h\}$, s_b is aggregate search effort of the group, $n_b + s_b$ is the labor force participation rate of the group, and $\frac{s_b}{s_b+n_b}$ is the unemployment rate of the group. Across the range of minimum wages $\overline{w}$ considered, the total change in non-college employment $\Delta \log n_\ell$ is primarily driven by changes in the participation rate $\Delta \log(n_\ell + s_\ell)$ rather than changes in the unemployment rate $\Delta \frac{s_b}{s_b+n_b}$. For example, the \$15 minimum wage reduces non-college employment by 9.2%, reduces the non-college labor force by 6.7%, and increases the non-college unemployment rate by 2.5 percentage points.

The fact that the disemployment effects of the minimum wage manifest through lower labor force participation rather than higher unemployment provides guidance for the outcomes which empirical researchers should study. Modeling endogenous search effort is necessary to obtain this result; otherwise, the entire long-run change in employment would have to be driven by changes in the unemployment rate, which we view as implausible. Empirically, Adams, Meer and Sloan (2022)

Figure 5: Disaggregated Minimum Wage Laffer Curves

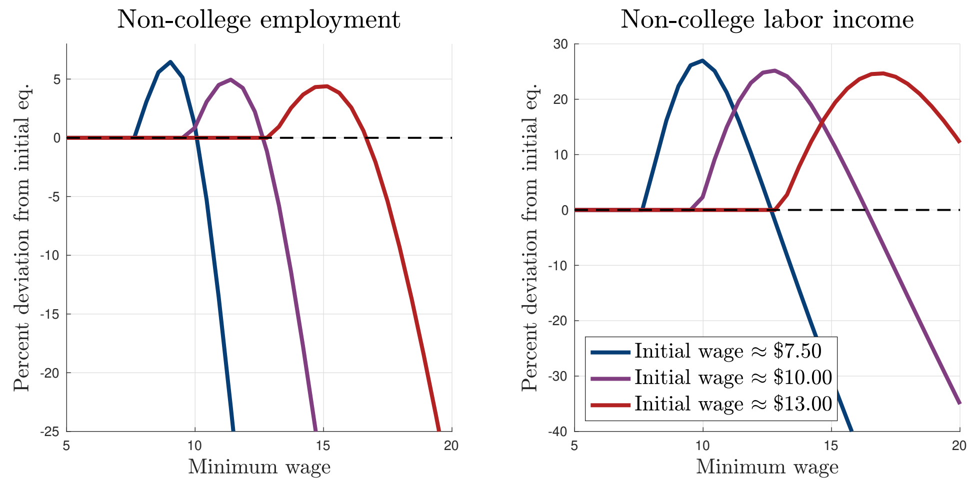

Notes: Steady-state employment (left panel) and labor income (right panel) of particular z-types among non-college workers as a function of the minimum wage. The x-axis is the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data. Initial wages rounded to the nearest half dollar.

find that search effort does not significantly increase in the few months after the introduction of the minimum wage.

The right panel of Figure 4 also shows that aggregate labor income is a Laffer curve as a function of the minimum wage. The labor income Laffer curve peaks at a higher level of the minimum wage than the employment Laffer curve because higher levels of the minimum wage increases the average wage per worker even if it decreases employment. In our calibration, the $15 minimum wage increases non-college labor income by about 1.7%.

Based on this aggregate outcome, it may be tempting to conclude that a $15 minimum wage will benefit non-college workers; however, we show in the next subsection that this conclusion is naive in the sense that it ignores the distribution of outcomes within non-college workers. Before moving on to that analysis, we note that Appendix D studies how the minimum wage shapes the aggregate distribution of income (between non-college labor income, college labor income, capital income, and residual profits). It shows that increasing the minimum wage within a certain range increases non-college labor income but, because it reduces employment, reduces firms profits and total GDP. In this sense, the minimum wage may redistribute profits from firms to workers, but at the expense of reducing aggregate economic activity.

Micro-Level Results. We now turn to studying how the minimum wage affects individual workers types z_i within the non-college group. Figure 5 plots employment and labor income Laffer curves for non-college workers with certain levels of productivity z. While the shapes of these Laffer curves

Figure 6: Distributional Effects of \$15 Minimum Wage Among Non-College Workers

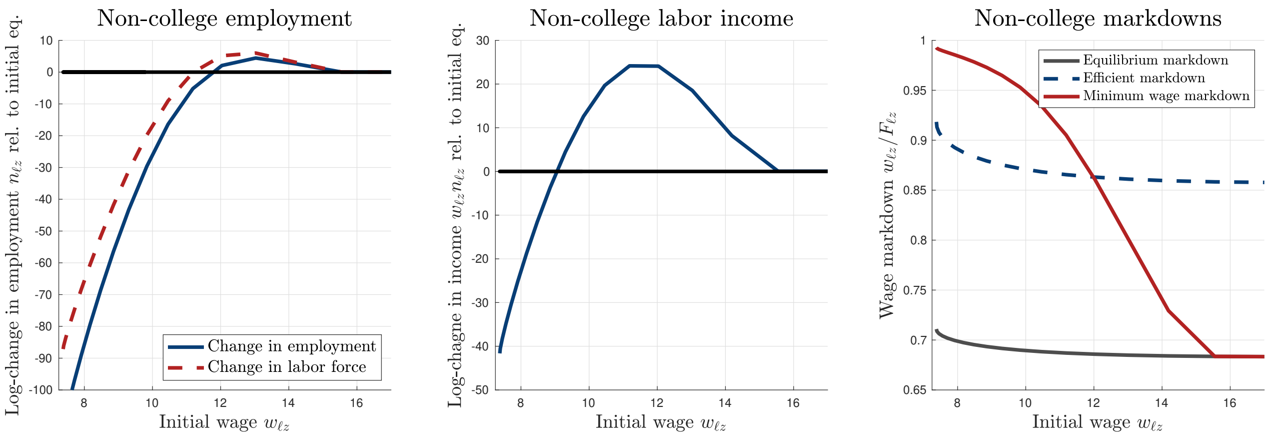

Notes: steady state outcomes for particular z-types among non-college workers for a \$15 minimum wage. Left panel plots the log-change in employment n_i, the middle panel plots the log-change in labor income $w_i n_i$ relative to the initial steady state without the minimum wage, and the right panel plots the the levels of the markdowns w_i/F_{ni} for three different parameterizations: (i) $\overline{w} = 0$ ("equilibrium markdown"), (ii) $\overline{w} = 0$ and $\omega \to \infty$ ("efficient markdown"), and (iii) $\overline{w} = \$15$ ("minimum wage markdown"). The x-axis corresponds to the initial wage $w_{\ell z}$ earned by a particular type z in the initial equilibrium.

mirror the aggregate ones plotted above, the peaks of the individual-level Laffer curves are increasing in the workers' productivity z because their efficient wage is increasing in z. This fact immediately implies a distributional conflict across workers: setting the minimum wage high enough to alleviate the monopsony distortion for the average z-type will decrease employment of the lower-z types. For example, setting the minimum wage to fix the distortion for the type making the equivalent of \$13 per hour in the initial equilibrium reduces employment of the types that were \$7.50 or \$10 per hour by more than 25%. Hence, there is a stark tension in setting one single minimum wage.

Figure 6 plots the distributional consequences of a \$15 minimum wage within the non-college group, as a function of the initial wage $w_{\ell z}$ earned by a particular type z in the initial equilibrium. The left panel shows that employment falls for all non-college workers who were initially earning less than \$12 per hour (who account for 44% of non-college households). The decline is largest among the lowest-wage workers because those workers have the lowest productivity z. The middle panel shows that the same broad pattern holds for labor income, although the set of workers whose labor income falls is smaller than for employment because wages rise for these workers even if their employment falls. Overall, though, the minimum wage disproportionately reduces the employment and income of precisely the group of workers that it is meant to benefit.

To further understand the mechanisms driving these results, the right panel Figure 6 plots the wage markdowns associated with particular z-types (again, expressed as a function of their initial wage $w_{\ell z}$). For low-z types, the minimum wage shrinks markdown to be lower than that in the efficient

Figure 7: Welfare Effects of the \$15 Minimum Wage

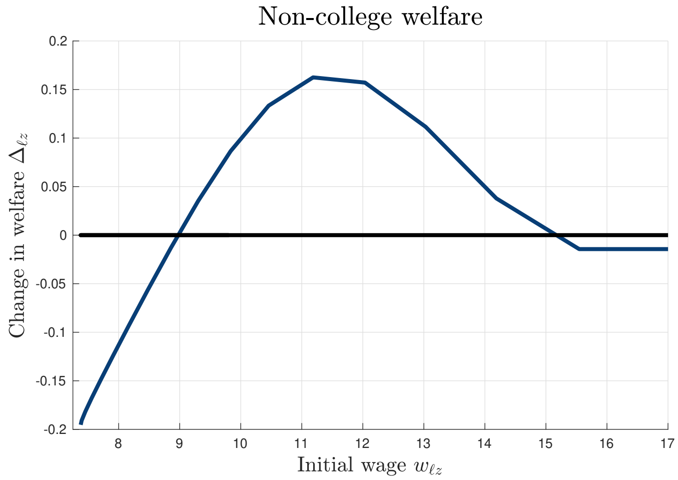

Notes: change in welfare for non-college workers $\Delta_{\ell z}$ as described in the main text. The x-axis corresponds to the initial wage $w_{\ell z}$ earned by a particular type z in the initial equilibrium.

equilibrium, which induces firms to reduce their employment to the point where wages approximately equal marginal products. For middle-z types, the minimum wage shrinks the markdown closer to its efficient level, which is associated with an increase in employment in the left panel. And for high-z types, the minimum wage does not bind and therefore does not affect the markdown or employment.

We define the change in welfare from the \$15 minimum wage as the value of Δ_i which solves

$$u\left((1+\Delta_i)c_i^* - v(n_i^*) - h(s_i^*)\right) = u\left(\widetilde{c}_i - v(\widetilde{n}_i) - h(\widetilde{s}_i)\right), \tag{34}$$

where the * superscript denotes the steady state values in the initial equilibrium without the minimum wage and the $\tilde{}$ superscript denotes values in the new equilibrium with the minimum wage. Δ_i is the percentage change in steady-state consumption that would make a household indifferent between living in the initial equilibrium and the new equilibrium. Hence, $\Delta_i > 0$ is positive if the policy makes the household better off and $\Delta_i < 0$ is negative if the household is worse off.[19] We view these welfare gains or losses as lower bounds on the actual ones because our model assumes perfect risk-sharing within each z-type family, so family members who lose their job are insured by those who do not. In a model without perfect risk-sharing, job loss would result in a larger drop in consumption and therefore in welfare.

With that caveat in mind, Figure 7 shows that about 1/3 of the households affected by the minimum wage experience a significant decline in their welfare. Welfare increases for the remaining workers — despite the fact that their employment falls — because wages and therefore labor income

[19] Our welfare analysis depends on the precise rule for distributing profits across households. We assume that profits are distributed in proportion to each households' share of total labor income.

rises for this group. Again, the fact that workers share risk within types imply that the workers who lose employment are compensated by the workers who retain their employment and are now earning the higher minimum wage; these welfare gains would shrink or even reverse if we relaxed the assumption of perfect risk-sharing within types.

Additional Results. Appendix D contains three additional sets of results. First, it shows how the effect of the minimum wage depends on the degree of monopsony power in the economy. Importantly, in the version of our model without monopsony power, any minimum wage unambiguously reduces employment and income, highlighting the importance of monopsony power in motivating the use of the minimum wage in the first place. Second, the appendix shows the effect of the minimum wage on college workers. Finally, it performs comparative statics with respect to key parameters governing labor supply, elasticities of substitution, and search frictions.

4.2 Short Run vs. Long Run

Now that we understand the long-run consequences of the minimum wage, we now study the transition path of the economy to that new study state. Along this transition path, the minimum wage imposes a lower bound $w_{ijt} \geq \bar{w}$ on the period wage that a firm can offer. Given a sequence of intertemporal prices $\{Q_{t,s}\}$, this constraint on period wages implies a constraint on the present value of wages of the form

$$W_{ijt}^F \geq \bar{W}_t^F \equiv \bar{w} + (1-\sigma)Q_{t,t+1}\bar{w} + (1-\sigma)^2 Q_{t,t+2}\bar{w} + \dots \tag{35}$$

Note that $\bar{W}_t^F$ is the smallest present value of wages consistent with satisfying the minimum-wage constraint $w_{ijt} \geq \bar{w}$ in each period of a match. This present value depends on time because the intertemporal prices do. We add the constraint (35) to a firm's problem—the consumer's problem does not change. If for a consumer of type i the minimum wage constraint does not bind, first-order conditions are the same as before. When the minimum wage constraint binds, we set $W_{ijt}^F = \bar{W}_t^F$ so that a firm's wage offer satisfies (35). Upon the introduction of the minimum wage, firms can fire existing workers or increase an existing worker's wage if desirable. All new hires must be paid at least the minimum wage and all workers retained by a firm must be paid the larger of their existing wage and the minimum wage.[20]

[20] If instead we allow a firm to lower the wage of an existing worker for whom the minimum wage does not bind, then a firm has an incentive to lower the wage until such a worker is just indifferent between quitting or staying with the firm. We assume that the original contract the worker signed contains a clause that specifies that the firm cannot lower the wage in the event that a minimum wage is introduced. We note that absent such a clause, in our baseline there is a small group of workers for whom the firm would like to lower wages. We rule out this possibility because, without such a clause, the unexpected introduction of a minimum wage allows firms to renege on the existing wage

FIGURE 8: Transition Path to New Minimum Wage Steady State

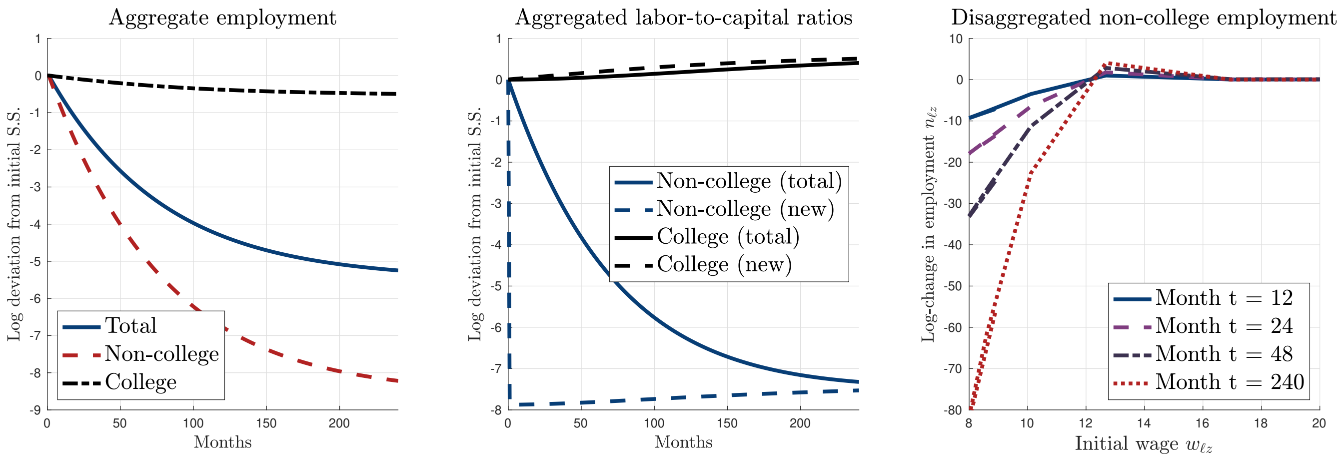

Notes: Transition path following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. Left panel plots the aggregated employment of non-college workers, college workers, and total employment. Middle panel plots the associated labor-to-capital ratios. Right panel plots the micro-level response of employment at different points along the transition path.

The $15 Minimum Wage. Figure 8 plots key features of the transition path following the introduction of the $15 minimum wage. Our main result, in the left panel, is that it takes nearly twenty years for employment to converge to its new steady state value. This result occurs because, as shown in the right panel, it takes a long time for firms to adjust their aggregate labor-to-capital ratios due to the putty-clay technology (although firms do immediately adjust the labor-to-capital ratios on new investment). In contrast, Appendix D shows that firms immediately adjust their aggregate labor-to-capital ratios in the model with standard capital, implying much faster transition dynamics. Hence, the putty-clay capital technology significantly dampens the effect of the minimum wage in the short run.

The right panel of Figure 8 plots the distributional effects of the minimum wage along the transition path (analogous to the steady state outcomes from Figure 6).[21] Even three years after the introduction of the minimum wage, employment of the low-z types — whose employment declines the most in the long run — has only declined by about 40% of its ultimate amount.

In order to highlight the role of the putty-clay technology in slowing down the transition, Figure 9 plots the distribution of capital types at different points of time. Before the introduction of the minimum wage, firms hold only one type of capital, namely the type that is optimal at the original

contracts that applied in its absence. More technically, we imagine that all agents believe that with probability ε a minimum wage will be introduced in the next period and the economy we consider is the limit of such an economy when ε converges to zero.

[21] The curves are less smooth along the transition than in steady state because we use a coarser grid of z-types when computing the transition paths (due to computational constraints).

FIGURE 9: Distribution of Capital Types Along the Transition Path

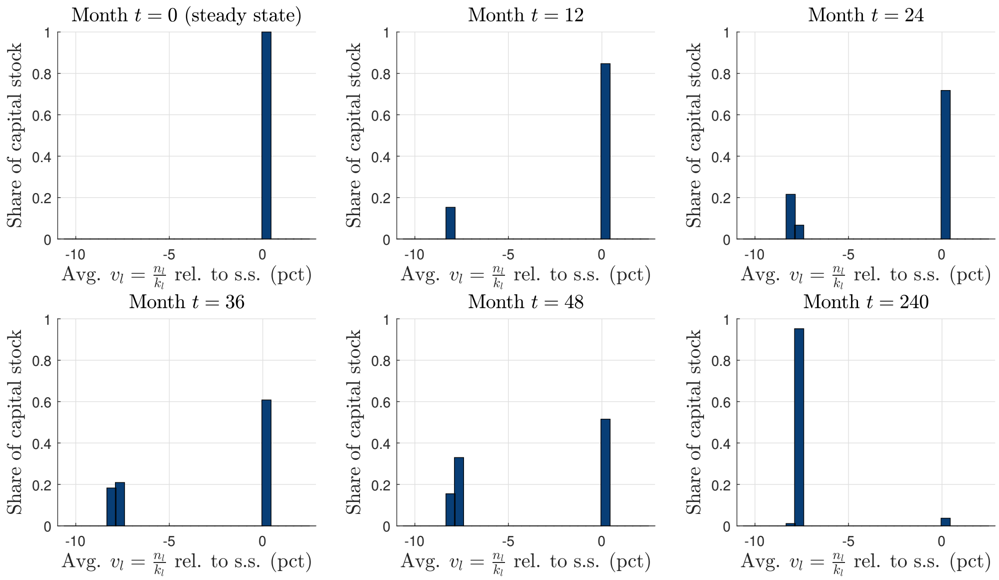

Notes: Distribution of capital types along the transition path following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. The x-axis indexes the type of capital by its average non-college labor to capital ratio.

steady state prices. The minimum wage induces firms to invest in less labor-intensive types of capital, but in the early stages of the transition that new capital accounts for a small share of the total capital stock. Over time, firms let the old capital stock depreciate away and continue to build up its stock of the new, less labor-intensive capital. Hence, the depreciation rate δ is crucial in determining the speed of transition.[22]

Implications for Empirical Work. Our results above show that the long-run effects of a \$15 minimum wage are substantially larger than the short-run effects, indicating that even the best-identified regressions using short-run data would not be a useful guide for estimating the long-run response. These results also suggest that the putty-clay technology may explain why existing empirical studies tend to find small effects of observed minimum wage changes on employment.[23]

[22] Card and Krueger (2015) argue against the idea that putty-clay technology can substantially slow down adjustment to the minimum wage based on the fact that establishments in the restaurant industry turn over every two years (indicating a high depreciation rate δ in that industry). While that argument may be correct in the context of small changes in the minimum wage, the large \$15 proposal would affect a much broader set of sectors in the economy and therefore requires the depreciation rate to be consistent with aggregate evidence from the BEA as in our calibration. That said, we excluded structures capital from this calculation in order to focus on the substitution between labor and equipment and/or software; if we had included structures, the implied depreciation rate would be even lower and therefore slow down the transition paths even more.

[23] Of course, there is disagreement in the literature about the sign of the response of employment to the minimum wage, depending on the size of the minimum wage considered, the source of variation, and the empirical methodology. Neumark and Shirley (2021) survey the literature and identify a range of estimated elasticities of employment with respect to the minimum wage. They place the estimated elasticities far below than the long-run effects of the \$15

FIGURE 10: Stylized Replication of Empirical Wor;

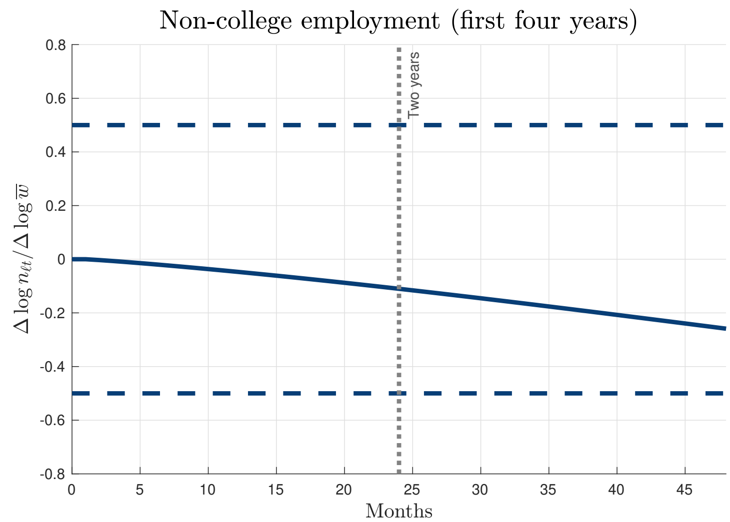

Notes: Transition path following an unexpected \$3 increase in the minimum wage. Transition paths computed in "partial equilibrium," i.e. with $Q_{t,t+1} = \beta$ for all t. Plots aggregated non-college employment over the full transition path. Dashed lines correspond to the set of minimum wage elasticities surveyed by Neumark and Shirley (2021).

However, these studies differ from our experiment above in two key respects: they typically consider "small" changes in the minimum wage, substantially below the \$15 experiment above, and they typically use local variation that will not generate economy-wide general equilibrium effects.

In order to investigate how our putty-clay technology shapes the existing empirical estimates of the effects of the minimum wage, we perform a stylized model experiment to replicate these two features of existing empirical work. First, we consider a \$3 increase in the minimum wage to loosely capture the size of minimum wage changes typically studied in the data. Second, we study a "small open economy" version of the model, i.e. we hold the consumption prices $Q_{t,t+1} = \beta$ fixed at each point of the transition path. We interpret this version of our model as capturing the local area affected by a given minimum wage change, holding other aggregate economic conditions fixed.

Figure 10 plots the path of non-college employment in response to this stylized empirical experiment. We normalize the change in log-employment $\log n_{\ell t}$ by the log-change in the minimum wage $\log \overline{w}$ in order to correspond with the elasticities estimated in the literature. Neumark and Shirley (2021) survey the empirical literature and place the majority of estimated elasticities — all estimated using short-run data — between -0.5 and 0.5. The figure shows that our model's short-run elasticity — the response over the first four years of the transition — is around -0.2, well within Neumark and Shirley (2021)'s range. Hence, we conclude that our putty-clay technology indeed ensures that the short-run effects of small and local changes in the minimum wage are consistent with the data.

minimum wage we show above.

5 Alternative Policies in the Tax and Transfer System

The previous section showed that, while the minimum wage may be successful in reducing the average monopsony distortion in the economy, it disproportionately reduces employment and labor income for low-income workers—exactly the group which the minimum wage attempts to help. We now study how alternative policies within the existing tax and transfer system may better achieve this goal. Section 5.1 describes how we model the tax and transfer system and ensure that the alternative policies are quantitatively comparable to a given change in the minimum wage. Section 5.2 studies the effect of the EITC, a particularly important component of transfers to low-income households in the data. Appendix E also studies the effects of a simple uniform tax credit as well as an approximation to the progressive tax and transfer system from Heathcote, Storesletten and Violante (2017). Throughout, we study the steady state consequences of these policies in order to focus on their long-run effects.

5.1 Modeling Tax and Transfer Programs

Although we focus our analysis on the EITC transfer program, we allow for a general tax and transfer system $T(w_i)$, where $T(w_i)$ is labor income taxes (so negative taxes $T(w_i) < 0$ indicate transfers). We let $A(w_i) = w_i - T(w_i)$ denote after-tax/transfer labor income, which may be greater or smaller than pre-tax income. The tax and transfer system affects both the incentives for firms to hire workers and for households to search in the labor market.

Consider how the tax and transfer system affects firms' labor demand, as summarized by the steady state vacancy posting condition:

$$\frac{\kappa_i}{\lambda_f(\theta_i)} = \frac{1}{r+\sigma}\left(F_{ni} - w_i - \underbrace{\frac{1}{\omega}\frac{v'(n_i)}{A'(w_i)}}_{\text{monopsony distortion}}\right), \tag{36}$$

where $A'(w_i)$ is marginal after-tax income. Equation (36) shows that a positive marginal tax rate, which implies $A'(w_i) < 1$, exacerbates the monopsony distortion relative to our baseline model with $A'(w_i) = 1$. To understand this mechanism, recall that the monopsony distortion arises because hiring a marginal worker increases the marginal disutility of labor supply for all the inframarginal hires, so the firm needs to compensate these inframarginal hires with a higher wage (increasing its private marginal cost of hiring above the planner's). A positive tax rate reduces the after-tax wage the inframaginal hires receive and therefore increases the required before-tax wage payment the firm must offer, further increasing the firms' private marginal cost of hiring. Conversely, a negative marginal tax rate, namely a tax credit which results in $A'(w_i) > 1$, alleviates the monopsony distortion by reducing firms' private marginal costs of hiring and brings firms' marginal costs of hiring closer to those of the planner.

While marginal tax rates affect the monopsony distortion on the labor demand side, average tax rates affect households' search decisions on the labor supply side. The optimal search decision can be written as

$$h'(s_i) = \frac{\lambda_w(\theta_i)}{r + \sigma} \left[A(w_i) - v'(n_i)\right], \tag{37}$$

where $r \equiv 1/\beta - 1$. A positive average tax rate, which implies $A(w_i) < w_i$, reduces the after-tax wage payments of a job relative to our baseline model and, as equation (37) indicates, reduces the incentive to search. The marginal tax rate is irrelevant for the search decision because that decision occurs along the extensive margin of whether or not to search. Conversely, negative average tax rates imply $A(w_i) > w_i$ and thus increase the incentive to search.

This discussion suggests that policies within the existing tax and transfer system, such as the EITC, are well-suited to alleviate the monopsony distortion in the economy. We ensure that the alternative policies we consider are comparable to the \$15 minimum wage in the following way. First, note that the \$15 minimum wage reduces firms' flow profits by some amount $\Delta\pi^*$, so we can think of the minimum wage as corresponding to an implicit tax on profits. For each of our alternative policies, we instead assume that there is no minimum wage but that the government levies an explicit corporate income tax τ_f on firms' profits which raises the same amount of revenues $\Delta\pi^*$. We assume that both investment and vacancy posting costs are fully deducted from corporate taxes, implying that the tax τ_c does not distort any of the firms' marginal decisions (because it is a pure profits tax). We then use these tax revenues to fund each of our alternative policies, that is, set $\tau_f \pi = \Delta\pi^* = -\sum_i T(w_i)$ where π is flow profits in the new equilibrium. Hence, each of our alternative policies transfers, on net, $\Delta\pi^*$ resources from firms to households; they only differ in the schedule $T(w_i)$ that determines the distribution of these transfers across households.[24]

5.2 Earned Income Tax Credit (EITC)

One of the largest components of transfer payments in the data is the EITC, whose schedule has kinks due to the credit being phased in and out at various income levels. Specifically, the EITC schedule has three regions: (i) the *phase-in* region in which the tax credit is paid proportionally to the household's income, (ii) the *plateau* region at which the tax credit is capped at its maximum benefit but is not

[24]Our pure profits tax is nondistortionary because there is a fixed mass of firms in our model. If instead we allowed for free entry, then reducing steady-state profits may reduce entry and therefore the total mass of firms. Our procedure here ensures that the entry margin would be distorted by the same amount across all our alternative policies and the \$15 minimum wage.

Figure 11: Earned Income Tax Schedule

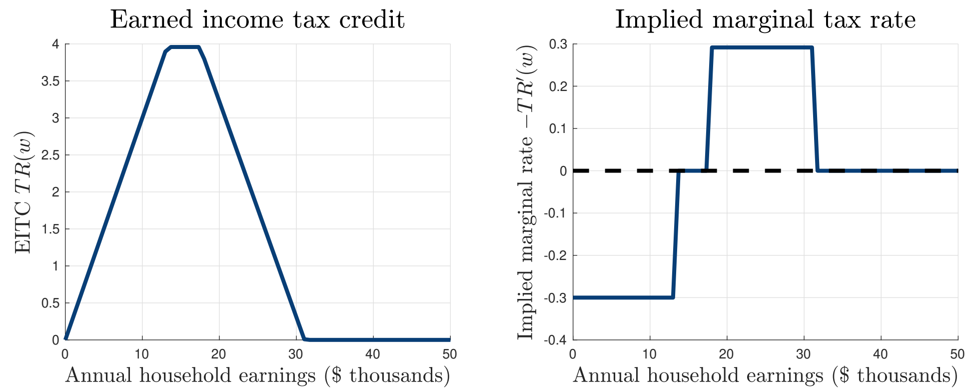

Notes: The left panel plots the EITC schedule which is budget-equivalent to the $15 minimum wage. The right panel plots the implied marginal tax rate. In each panel, the x-axis rescales steady state labor income to annual earnings assuming each household works 1800 hours per year.

yet being phased out, and (iii) the *phase out* region in which the credit is phased-out.

We set our transfer payments $T(w_i)$ to mimic this empirical ETIC schedule, but adjust the magnitudes to ensure that it is budget-equivalent to the $15 minimum wage. Throughout, we assume that there are no other components of the tax and transfer system, i.e. we set $\lambda = 0$. Figure 11 plots our budget-equivalent schedule. We set the phase-in rate to 30%; in the phase-in region, households face both a negative average tax rate (since the total tax credit is positive) and a negative marginal tax rate (since the credit is being phased in). In the plateau region, the household faces a zero marginal tax rate but still a negative average tax rate since they are still receiving the benefit. Finally, we set the phase-out rate to approximately 30%; in the phase-out region, the household faces a positive marginal tax rate because each dollar of earnings reduces its transfer payment.

The left panel of Figure 12 shows how employment responds to each of these three regions of the EITC schedule. The lowest-productivity non-college workers are in the phase-in region and therefore receive both the positive effect of the negative average tax on labor supply and the negative marginal tax on labor demand.[25] Employment falls in the plateau region because the marginal tax rate jumps up from $-\tau_1$ to 0 (where τ_1 is the phase-in rate). However, employment is still higher than in the initial equilibrium without the policy because the average tax rate is still negative, increasing labor supply. Employment falls again in the phase-out region because the marginal tax rate jumps up from 0 to τ_2 (where τ_2 is the phase-out rate). The right panel shows that similar patterns hold for labor income has well. Hence, although the EITC is qualitatively similar to the smooth transfer payment from Section E.1, its actual implementation creates kinks in the response of employment and income.

[25] In fact, the phase-in region is isomorphic to the uniform tax credit studied in Appendix E, but its effect is much larger because it is targeted at a small set of households.

Figure 12: Effect of Earned Income Tax Credit on Employment

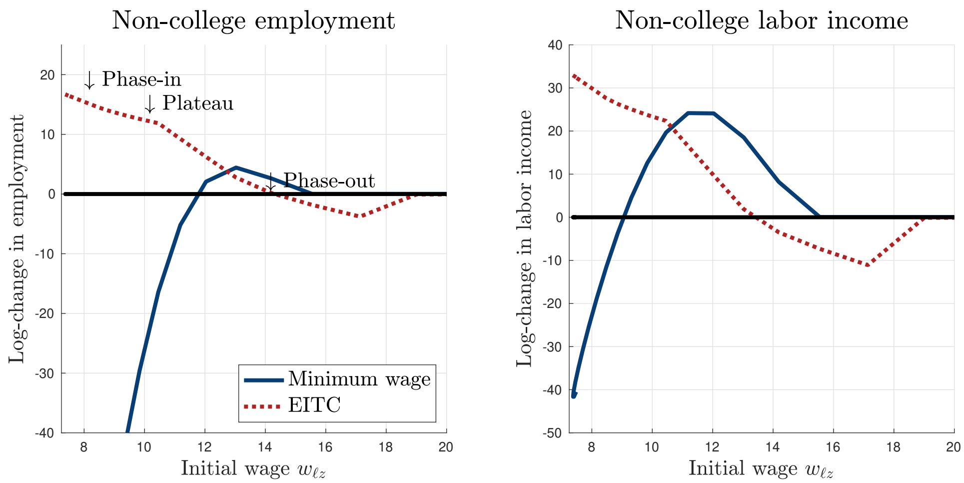

Notes: Steady-state employment (left panel) and labor income (right panel) of particular z types under the EITC. The y-axis is normalized relative to employment in the initial equilibrium (without any policies). The x-axis is the wage $w_{\ell z}$ earned by a particular z-type in the initial equilibrium without either policy.

Figure 13: Effects of Alternative Policies on Welfare

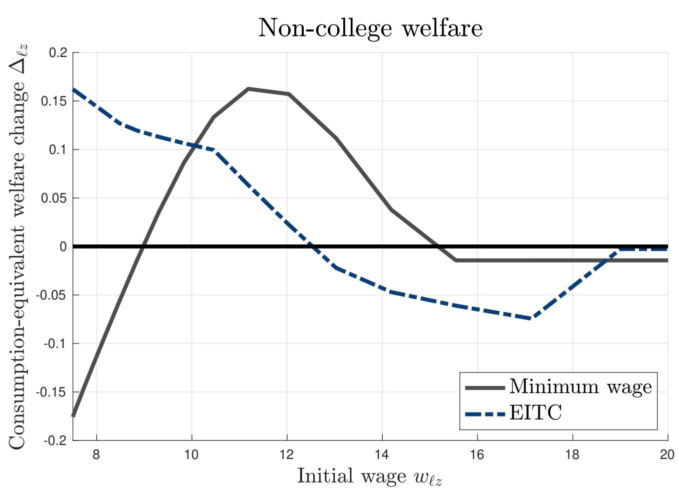

Notes: Steady-state welfare gains/losses Δ_i from $u\left((1+\Delta_i)c_i^* - v(n_i^*) - h(s_i^*)\right) = u\left(\widetilde{c}_i - v(\widetilde{n}_i) - h(\widetilde{s}_i)\right)$, where where the * superscript denotes the steady state values in the initial equilibrium without the policy and the $\widetilde{}$ superscript denotes values in the new equilibrium with the policy. The x-axis is the wage $w_{\ell z}$ earned by a particular z-type in the initial equilibrium without either policy.

FIGURE 14: Effect of Constant Transfer with and without $\overline{w} = \$9.25$ Minimum Wage

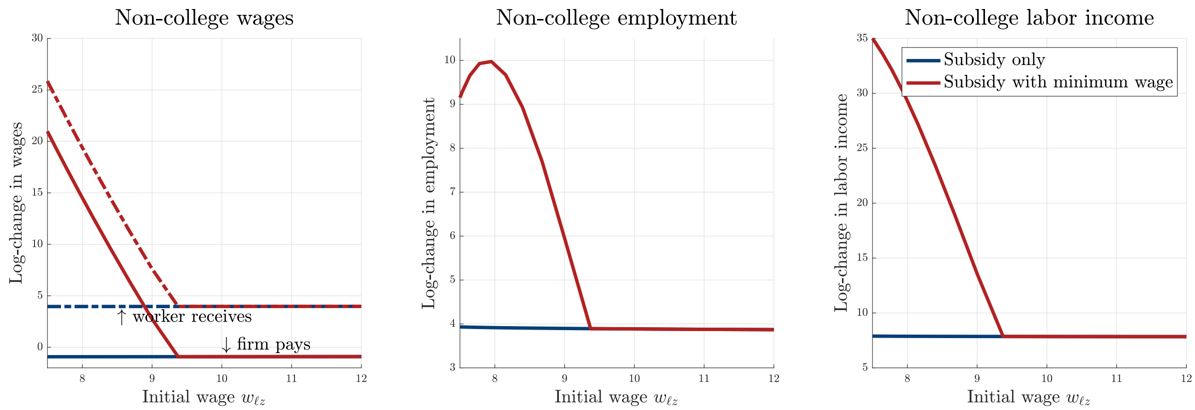

Notes: Steady-state wages (left panel), employment (middle panel), and labor income (right panel) in response to constant transfer $\tau_c = 5\%$ of labor income. The y-axis is the log-change relative to initial equilibrium (without any policies). The x-axis is the wage $w_{\ell z}$ earned by a particular z-type in the initial equilibrium without either policy.

Figure 13 compares the welfare effects Δ_i of the \$15 minimum wage and the budget-equivalent EITC. The EITC also dominates the minimum wage for low-income households, mirroring its effect on labor income described above. In contrast, the minimum wage benefits middle-income households, who were earning close to the minimum wage in the initial equilibrium. Appendix E shows that a progressive tax and transfer system is even better than the ETIC because that system can finance larger transfers by raising taxes on high-income workers.

In sum, the \$15 minimum wage is implicitly targeted at improving the welfare of middle-income households at the expense of dramatically reducing the welfare of low-income households. In contrast, the EITC program and the more general progressive tax and transfer system we consider both improve the welfare of low-income households at the expense of high-income households.

5.3 Interaction Between Transfer Policies and the Minimum Wage

A long-standing issue with the EITC and other types of transfer payments is that they partly benefit firms because they are able to pay lower pre-transfer wages (given the increase in labor supply). This issue has led some authors to suggest that a small minimum wage may complement transfer programs because it prevents firms from lowering the wages they pay their workers.[26] Our goal in this section is to explore this potential complementarity in our quantitative model. In order to clearly explain the economic mechanisms, we start by studying a constant transfer equal to 5% of households' labor income (i.e. $T(w_i) = -0.05w_i$, which is identical to the phase-in region of the EITC with a phase-in rate of 5%). We then show the results generalize for the full EITC schedule as well.

[26] See Neumark and Wascher (2011) or Lee and Saez (2012), for example.

Figure 15: Effect of EITC with and without $\overline{w} = \$9.25$ Minimum Wage

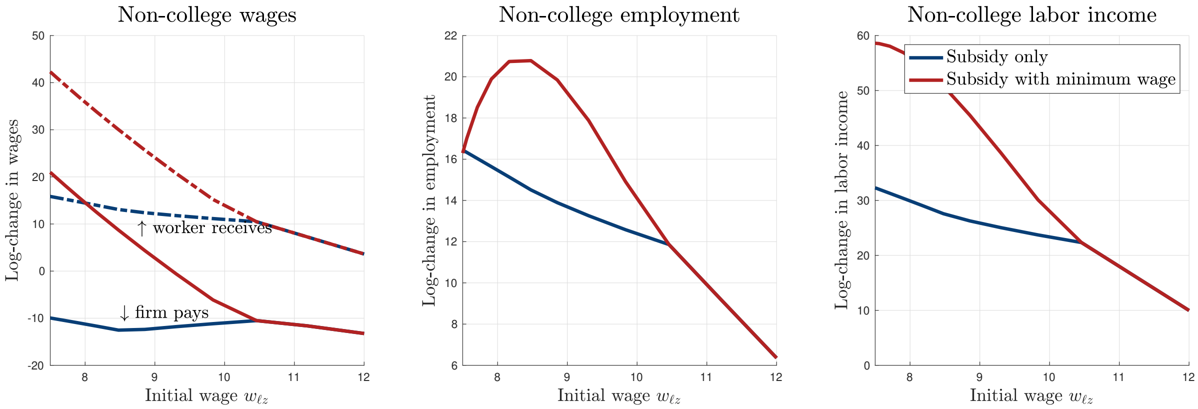

Notes: Steady-state wages (left panel), employment (middle panel), and labor income (right panel) in response to our budget-equivalent EITC system from Section 5. The y-axis is the log-change relative to initial equilibrium (without any policies). The x-axis is the wage $w_{\ell z}$ earned by a particular z-type in the initial equilibrium without either policy.

The blue lines in Figure 14 illustrate how our constant 5% transfer partly benefits firms. Specifically, the left panel shows that firms are able to pay their workers 1% less due to the transfer, implying that the post-transfer wages households receive $(1 + 0.05)w_i$ only increase by 4% rather than 5%. Employment therefore only increases by 4% as well, reducing the overall increase in labor income that would have resulted if wages did not fall.

The red lines in Figure 14 shows that a small minimum wage of $\overline{w} = \$9.25$ prevents firms from cutting wages in response to the transfer payment for low-income workers. The employment and labor income of these workers increases substantially more than under the transfer program alone. Of course, the minimum wage benefits these workers not only by preventing firms from reducing wages, but also because the minimum wage alleviates the monopsony distortions these workers face (as we have discussed throughout the paper). Appendix E.3 nets out the monopsony reduction in order to focus on the pure interaction between the transfer and minimum wage, and confirms that these interaction effect is positive for many of the workers on whom the minimum wage binds. In this sense, transfer programs and a small minimum wage are complementary policies because the minimum wage prevents firms from cutting wages in response to the transfer.

Figure 15 shows that these insights carry through in response to our budget-equivalent EITC schedule from Section 5. Although the effects of the EITC are more kinked due to the kinks in the EITC schedule, the $\overline{w} = \$9.25$ minimum wage substantially boosts the effect of the policy relative to the case without the minimum wage.

We conclude from this analysis that the minimum wage can have a valuable role in supporting transfer programs like the EITC. However, the size of the minimum wage must be chosen carefully;

if the minimum wage is set too high, it may end up hurting low-income workers for the reasons described throughout the paper. As an illustrative example, Appendix E.3 shows that a $\overline{w} = \$12$ minimum wage substantially reduces the benefit of the EITC for the lowest-income workers.

6 Conclusion

Many proposed changes in the national minimum wage policy currently contemplated in the United States entail increases in the minimum wage well outside of the range of past experience. Unfortunately, because of this feature, existing empirical work on changes in the minimum wage implemented in the United States are likely to be uninformative about the long-run labor market implications of such policies. In this paper, we develop a general-equilibrium framework with rich worker heterogeneity and firm market power in labor markets subject to search frictions to study the short-run and long-run effects of increases in the minimum wage of the magnitudes proposed.

We have argued that we can use the large changes in the relative price of labor, emanating from the observed changes in the price of capital, to discipline the long-run labor market responses to changes in the minimum wage. We have also ensured that our model is consistent with the evidence that changes in the relative price of labor have small employment effects in the short run, when emanating from changes in the minimum wage. Once we do so, we find that a large increase in the minimum wage has perverse distributional impacts in the long run: even though it increases aggregate labor income, it reduces the welfare of the low-income workers that it is designed to help. Our quantitative analysis shows that a better way to increase the welfare of low-productivity workers is through a progressive tax-transfer scheme that induces them to work more, as opposed to a large increase in the minimum wage that effectively prices them out of the market by making them unattractive for firms to hire. Within this context, a small minimum wage is a valuable tool which enhances the efficacy of these policies.

References

Aaronson, D., E. French, I. Sorkin, and T. To (2018): "Industry dynamics and the minimum wage: a putty-clay approach," *International Economic Review*, 59(1), 51–84.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251–333.

Adams, C., J. Meer, and C. Sloan (2022): "The minimum wage and search effort," *Economics Letters*, p. 110288.

Ahn, T., P. Arcidiacono, and W. Wessels (2011): "The Distributional Impacts of Minimum Wage Increases When Both Labor Supply and Labor Demand Are Endogenous," *Journal of Business and Economic Statistics*, 29(1), 12–23.

Berger, D. W., K. Herkenhoff, and S. Mongey (2021): "Labor Market Power," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-58).

Berger, D. W., K. F. Herkenhoff, and S. Mongey (2019): "Labor Market Power," Working Paper 25719, National Bureau of Economic Research.

——— (2022): "Minimum Wages, Efficiency and Welfare," Discussion paper, National Bureau of Economic Research.

Bils, M. J., B. Kaymak, and K.-J. Wu (2020): "Labor Substitutability among Schooling Groups," *Unpublished paper, University of Rochester*.

Brown, C. (1988): "Minimum Wage Laws: Are They Overrated?," *Journal of Economic Perspectives*, 2(3), 133–145.

——— (1999): "Minimum wages, employment, and the distribution of income," *Handbook of labor economics*, 3, 2101–2163.

Cahuc, P., S. Carcillo, and A. Zylberberg (2014): *Labor Economics*.

Calvo, G. A. (1976): "Optimal growth in a putty-clay model," *Econometrica: Journal of the Econometric Society*, pp. 867–878.

Card, D., and A. B. Krueger (1995): *Myth and Measurement: The New Economics of the Minimum Wage*.

Card, D., and A. B. Krueger (2015): *Myth and Measurement: The New Economics of the Minimum Wage-Twentieth-Anniversary Edition*. Princeton University Press.

Card, D., and T. Lemieux (2001): "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis*," *The Quarterly Journal of Economics*, 116(2), 705–746.

Deere, D., K. M. Murphy, and F. Welch (1995): "Employment and the 1990-1991 Minimum-Wage Hike," *The American Economic Review*, 85(2), 232–237.

Eckstein, Z., and K. I. Wolpin (1990): "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, 58(4), 783–808.

Flinn, C. J. (2006): "Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates," *Econometrica*, 74(4), 1013–1062.

——— (2010): *The Minimum Wage and Labor Market Outcomes*. MIT Press.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *The American Economic Review*, pp. 402–417.

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132(4), 1693–1754.

——— (2020): "Presidential Address 2019: How Should Tax Progressivity Respond to Rising Income Inequality?," *Journal of the European Economic Association*, 18(6), 2715–2754.

Hershbein, B., C. Macaluso, and C. Yeh (2020): "Monopsony in the US labor market," *Federal Reserve Bank of Richmond Working Paper*.

Hubmer, J. (2018): "The Race between Preferences and Technology," *Unpublished Working Paper*.

Johansen, L. (1959): "Substitution versus Fixed Production Coefficients in the Theory of Economic Growth: A Synthesis," *Econometrica*, 27(2), 157–176.

Karabarbounis, L., and B. Neiman (2014): "The global decline of the labor share," *The Quarterly Journal of Economics*, 129(1), 61–103.

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

Kennan, J. (1995): "The Elusive Effects of Minimum Wages," *Journal of Economic Literature*, 33(4), 1950–1965.

Kline, P., R. Saggio, and M. Sølvsten (2020): "Leave-Out Estimation of Variance Components," *Econometrica*, 88(5), 1859–1898.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000): "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68(5), 1029–1053.

Lee, D., and E. Saez (2012): "Optimal minimum wage policy in competitive labor markets," *Journal of Public Economics*, 96(9-10), 739–749.

Lise, J., C. Meghir, and J.-M. Robin (2016): "Matching, Sorting and Wages," *The Review of Economic Dynamics*, 19(S), 63–87.

Manning, A. (2021a): "The Elusive Employment Effect of the Minimum Wage," *Journal of Economic Perspectives*, 35(1), 3–26.

——— (2021b): "Monopsony in Labor Markets: A Review," *ILR Review*, 74(1), 3–26.

Neumark, D., and P. Shirley (2021): "Myth or Measurement: What Does the New Minimum Wage Research Say about Minimum Wages and Job Loss in the United States?," Working Paper 28388, National Bureau of Economic Research.

Neumark, D., and W. Wascher (2008): *Minimum Wages.*

——— (2011): "Does a Higher Minimum Wage Enhance the Effectiveness of the Earned Income Tax Credit?," *Industrial and Labor Relations Review*, 64(4), 712–746.

Robinson, J. (1933): *The Economics of Imperfect Competition.*

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American economic review*, 95(1), 25–49.

Sorkin, I. (2015): "Are There Long-Run Effects of the Minimum Wage?," *Review of Economic Dynamics*, 18(2), 306–333.

Wright, R., P. Kircher, B. Julien, and V. Guerrieri (2021): "Directed Search and Competitive Search Equilibrium: A Guided Tour," *Journal of Economic Literature*, 59(1), 90–148.

Zwick, E., and J. Mahon (2017): "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107(1), 217–48.

A Model Appendix

We present here details omitted from the main text.

A.1 Analogy Between Monopolistic and Monopsonistic Competition

Here we discuss how the upward-sloping labor supply curve for each firm's jobs that our model gives rise to is analogous to the downward-sloping demand curve for each firm's goods that arises in models of monopolistic competition. In these latter models, consumers view each firm's good as imperfectly substitutable with any other. Then, the downward-sloping demand curve of consumers for a firm's good as a function of all firms' prices is the constraint that captures monopoly power in a firm's problem. Analogously, in our setup, workers view each firm's job as imperfectly substitutable with any other. Thus, the upward-sloping supply curve of searchers for a firm's jobs as a function of all firms' wages is the constraint that captures monopsony power in a firm's problem.

To elaborate, recall that standard analyses of monopolistic competition derive the static demand curve of consumers for goods of each type j and then impose this demand curve as a constraint on firm j's problem. Equivalently, one could derive the first-order conditions for a consumer's problem and impose the constraint that the marginal utility from buying good j is at least as high as that from buying any other good. More formally, the static part of a consumer's dynamic problem with a standard utility function over differentiated goods, given total expenditure pc, is to choose $\{c_j\}$ to maximize $u(c)$ subject to $\sum_j p_j c_j \leq pc$, where $c = \left(\sum_j c_j^{\frac{\omega-1}{\omega}} dj\right)^{\frac{\omega}{\omega-1}}$ is the consumption aggregate and p is the associated price index. The first-order condition for buying good j implies that

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j = \max_{j'}\left\{u'(c)\left(\frac{c_{j'}}{c}\right)^{-\frac{1}{\omega}} - \lambda p_{j'}\right\}, \tag{38}$$

where λ is the multiplier on the budget constraint. In a symmetric allocation with $c_{j'} = c$ and $p_{j'} = p$, (38) reduces to

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j \geq u'(c) - \lambda p, \tag{39}$$

which is the participation constraint under monopolistic competition. Notice the similarity of this participation constraint for attracting a consumer to buy a good from firm j and the participation constraint for attracting a searching consumer to the labor market (θ_j, w_j) created by firm j, namely (14). The main difference between the two constraints arises because choosing which good to buy given a level of expenditure is a static decision whereas searching for a firm offering a long-term labor contract is a dynamic decision.

B Data Appendix

This appendix contains details about our data sources and construction referenced throughout the main text.

B.1 Census American Community Survey (ACS) Data

In terms of household-level data, we use the 1970, 1980, 1990, and 2000 U.S. Censuses as well as the 2006-2019 American Community Surveys (ACS). All data are weighted using the weights provided by the Census and ACS samples for the given time periods.

Time Periods. We use data from the Census/ACS over three different time horizons. First, we use data from the 2017-2019 pooled ACS sample when calibrating moments key empirical moments. Second, we use decadal differences across the Census and ACS samples when estimating how sectoral income changes with decadal changes in the sectoral price of capital. We estimate these regressions over both 10 and 20 year horizons. These semi-elasticities allow us to validate the elasticities of substitution between high and low educated workers and capital. For the 1970, 1980, 1990, and 2000 time periods we use data from the U.S. Censuses. For the 2010 time period, we pool together data from the 2010-2012 American Community Surveys. For the 2020 data, we pool together data from the 2017-2019 American Community Surveys. We pool together data from the ACSs to increase the sample sizes when measuring sectoral changes over time. Finally, we use annual data from the 2000-2014 American Community Surveys when estimating the short run response of the sectoral share of income accruing to high educated workers in response to annual changes in the sectoral price of capital driven by changes in the U.S. tax code (described in Appendix C below).

Share of High Educated. We define two education groups within the paper: "college" and "non-college." We define college individuals as those individuals who report having a bachelor's degree or higher. During the 2017-2019 period, 31.3% of our sample had at least a bachelor's degree.

Employment Rates. As part of our calibration, we match "full-time" employment rates by education group. By focusing on "full-time" employment, we measure workers with a strong attachment to the labor force. We define an individual as being "full-time" employed if (1) they are currently working at least 30 hours per week, (2) they reported working at least 29 weeks during the prior year, and (3) they reported positive labor earnings during the prior 12 month period. For our 2017-2019 sample, 46.8% of non-college individuals and 62.4% of college individuals worked full-time.

Share of Income Earned by College Workers. For the 2017-2019 period, 37.8% of individuals working full-time were college educated. Conditional on being a full-time worker, mean annual earnings for college individuals totaled \$91,706 while mean annual earnings for non-college individuals total \$44,871. Given these numbers, we compute that 55.5% of all earnings of full-time workers accrued to workers with at least a bachelor's degree.

Wage Distribution. We also use the 2017-2019 pooled ACS data to compute the distribution of hourly wages within each education group. To do so, we compute hourly wages for our sample of full-time workers by dividing annual labor earning by annual hours worked. We compute annual hours worked as the multiple of weeks worked last year and reported usual hours worked. We make two other sample restrictions when computing the wage distribution. First, we restrict the sample to only those workers who report at least \$5,000 of labor earnings during the prior year. Second, we then truncate the distribution at the top and both 1% of the wage distribution. All wages are converted to 2019 dollars using the June CPI-U. From this data, we compute the median wage as well as the ratios of wages between the 10th percentile and the median and the ratio of wages between the 90th percentile and the median separately for each of the education groups. These moments are used as part of our calibration. We also show that even though only those three moments are targeted for each education group, our model matches the full distribution of wages for each education group quite closely.

B.2 Bureau of Economic Analysis (BEA) Data

Data Construction. In order to validate our elasticities of substitution between capital and labor, we combine the household-level income data described above with data on the relative price of investment, real investment expenditures, and the real value of the capital stock constructed from the BEA Detailed Fixed Asset Tables. We focus on the time period $1960-2019$ given the coverage of the Census/ACS data. We define the growth rate in the price of investment relative to the price of consumption in sector s in year t using the Tornqvist index

$$\Delta \log q_{st} = \sum_{a=1}^{A} \omega_{ast} \Delta \log q_{at}, \tag{40}$$

where a indexes a type of capital good, ω_{ast} is a weight which sums to one within sector, and $\Delta \log q_{at}$ is the growth rate of the relative price of asset a. The Tornvqist weight ω_{ast} is the average of the

FIGURE B.1: Sources of Variation in the Relative Price of Investment q_{st}

(a) Asset-level price series (b) Asset-level investment shares

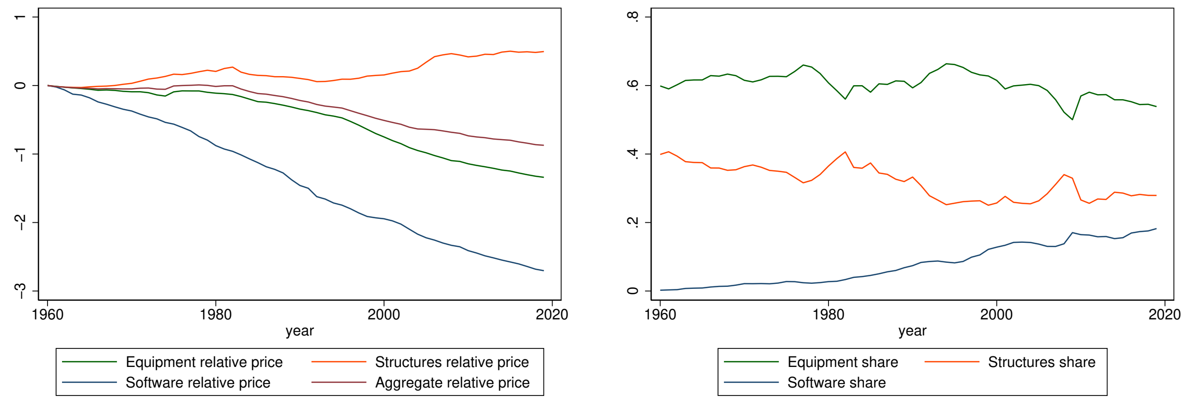

Notes: Left panel plots the level of the price index relative to 1960, $\log q_{st} \equiv \sum_{\tau=0}^{t} \Delta \log q_{s\tau}$, for three broad asset categories: structures, equipment, and software. The right panel plots the share of aggregate nominal investment expenditures on structures, equipment, and software assets.

share of nominal investment expenditures on asset a in periods t and $t-1$:

$$\omega_{ast} = \frac{P^I_{at} i_{ast}}{\sum_{a=1}^{A} P^I_{at} i_{ast}},$$

where P^I_{at} is the nominal price of asset a and $P^I_{at} i_{ast}$ are nominal investment expenditures in sector s on asset a. The growth rate in the relative price of investment is $\Delta \log q_{at} = \Delta \log P^I_{at} - \Delta \log P^C_t$, where P^C_t is the price index of consumption goods.[27] We compute the nominal price of asset a, P^I_{at}, as the ratio of nominal investment expenditures $P^I_{at} i_{ast}$ from Investment in the Private Nonresidential Fixed Assets table to real investment expenditures i_{ast} from the Fixed-Cost Investment in Private Nonresidential Fixed Assets.[28] We exclude artistic originals and R&D from our capital assets because they are not clearly related to capital-labor substitution.

Figure B.1 illustrates two key sources of variation in our relative price series $\Delta \log q_{st}$. First, the time-series variation in sector-level relative prices $\Delta \log q_{at}$ has tended to decline over time. We illustrate this variation in the left panel of the figure, which plots the level of the price index relative to 1960, $\log q_{st} \equiv \sum_{\tau=0}^{t} \Delta \log q_{s\tau}$, for three broad asset categories: structures, equipment, and software.[29] The relative price of equipment and especially software assets has substantially

[27] We measure the consumption price index using another Tornqvist index which aggregates over nondurable goods and services, constructed from the NIPA Tables 1.1.5 and 1.1.9.

[28] In practice, the data contain some variation in the implied price index P^I_{at} across sectors because the publicly available data does not disaggregate to the finest level of assets a available to the BEA. However, this variation is minor, and our results are robust to using the more aggregate asset-level price series from the NIPA Table 5.5.4. We prefer our baseline approach because it contains substantially more assets than the more aggregated NIPA series.

[29] We define these three aggregated price series using a Tornqvist index analogous to (40) within that particular asset category.

fallen over time, driving down the aggregate price index by more than one log point over this period (consistent with aggregate price series constructed in the literature, for instance, Hubmer (2018)). Second, there is variation in the investment expenditure shares ω_{ast} both across sectors and over time. The right panel of Figure B.1 illustrates this source of variation by plotting the share of aggregate nominal investment expenditures in the three broad asset categories and shows that the share of expenditures on software has substantially risen over time. These two sources of variation interact to create substantial cross-sector heterogeneity in q_{st} over time.

We construct other sector-level time series in a similar fashion. We compute the growth rate of real investment at the sector level using a similar Tornqvist index to (40):

$$\Delta \log i_{st} = \sum_{a=1}^{A} \omega_{ast} \Delta \log i_{ast}.$$

We compute the growth rate of real capital stock in sector s as

$$\Delta \log k_{st} = \sum_{a=1}^{A} \omega_{ast}^{k} \Delta \log k_{ast},$$

where k_{ast} is the real value of the stock of assets a in sector s and ω_{ast}^{k} is the share of the sector s nominal capital stock in sector s in asset a, both computed from the BEA detailed fixed assets tables.

We combine this BEA data with the household-level data from the Census/ACS at the sector level. Specifically, in each decade, we collapse our Census/ACS samples to 230 three digit sectors using 1990 sector codes which are harmonized across the Census/ACS years. The BEA data is at a higher level of sectoral aggregation than is the Census/ACS sector data. As a result, we create a cross-walk that maps the BEA sectoral investment price series to the Census/ACS sector codes. In our final sample, we have 62 sectors with a capital price series that we map to 195 Census/ACS sectors, i.e. a given BEA capital price sector can be mapped to multiple Census/ACS sectors.

Data Validation. In Appendix C, we study the relationship between long-run (decadal) changes in the relative price of investment q_{st} and in the college income share.[30] In order to validate our price series for this purpose, Figure B.2 shows a clear negative relationship between the decadal growth in the relative price and the decadal growth of investment. This negative relationship is consistent with our assumption below that changes in the relative price primarily reflect investment supply shocks,

[30] We remove the observations of the oil and gas extraction sector in 1980 and 2010 because they are influential outliers (the relative price of capital in these sectors substantially increased over time, likely due to extreme energy market fluctuations in those periods).

Figure B.2: Change in Investment Price vs. Change in Real Investment

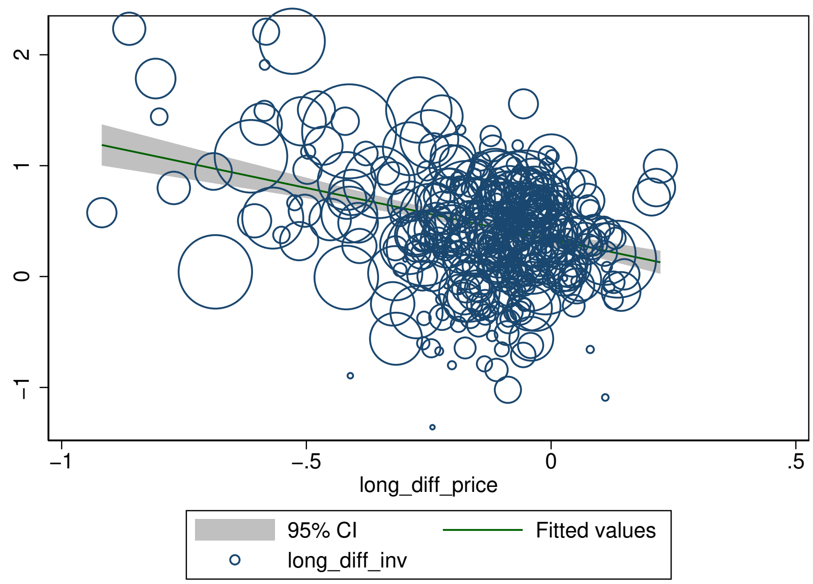

Notes: Scatterplot of ten-year changes in the relative price of capital q_{st} (x-axis) vs. the ten-year changes in log real investment $\log i_{st}$ (y-axis). Observations are weighted by their share of aggregate investment.

Table B.1: Robustness of Long-Run Investment Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
long_diff_price	-0.92***	-1.36***	-1.02**	-0.44	-1.64***	-1.72***	-2.22***	-0.63*	-0.68*
	(0.211)	(0.271)	(0.309)	(0.875)	(0.288)	(0.368)	(0.290)	(0.237)	(0.324)
Observations	370	370	306	62	61	62	62	61	62
R^2	0.133	0.388	0.432	0.004	0.356	0.266	0.496	0.106	0.068
Sample	Pooled	Pooled	Pooled	60-70	70-80	80-90	90-00	00-10	10-19
Time FEs?	No	Yes	Yes	No	No	No	No	No	No
Lagged Price?	No	No	Yes	No	No	No	No	No	No

Notes: Results from estimating (41) using ten-year changes in investment and the relative price of capital. Column (1) excludes the time fixed effects b_t. Column (2) includes the time fixed effects b_t. Columns (3) includes time fixed effects and controls for the previous decades ten-year relative price growth. Columns (4) - (8) run (41) separately for each decade of data excluding time fixed effects.

allowing us to trace out the slope of the investment demand schedule.

Table B.1 quantifies this scatterplot using the regression specification

$$\Delta \log i_{st} = b_0 + b_t + b_1 \Delta \log q_{st} + \varepsilon_{st}, \tag{41}$$

where b_0 is a constant, b_t is a decade fixed effect, ε_{st} is a residual, and b_1 is the coefficient of interest. Column (1) shows that $b_1 = -0.93$ without the time fixed effects and column (2) shows that incorporating time fixed effects increases b_1 to -1.37. Column (3) shows that this result is robust to controlling for the last decade's investment price growth. Finally, columns (4)-(9) show that this negative relationship is fairly robust across decades, although it has weakened somewhat in the more recent period.

Calibrated Depreciation Rate. In Section 3, we calibrate the depreciation rate δ to match the depreciation rate of equipment and software in our BEA dataset. We compute an average depreciation rate δ_t by aggregating asset-level depreciation rates δ_{at} according to

$$\delta_t = \sum_{a=1}^{\widetilde{A}} \frac{P^k_{at} k_{at}}{\sum_{a=1}^{\widetilde{A}} P^k_{at} k_{at}} \delta_{at},$$

where $\widetilde{A}$ denotes the set of equipment and software assets only and $P^k_{at} k_{at}$ is the nominal value the economy-wide capital stock in asset a. Each of these variables are recorded in the BEA detailed fixed asset tables. The average depreciation rate has increased over time, primarily due to the rising importance of software assets (which have higher depreciation rates than other assets). We target the value of the depreciation rate from 2019 to be consistent with the recent time period.

C Validating Long- and Short-Run Elasticities of Substitution

In our model, the long-run elasticities of substitution between capital and the different types of labor are governed by the parameters ρ and α from Krusell et al. (2000), while the short-run elasticities are Leontief due to the putty-clay technology. In this appendix, we present evidence in favor of both sets of assumptions.

C.1 Long-Run Elasticities of Substitution

We first describe our procedure for estimating long-run substitution patterns using sector-level time series on the relative price of capital and the distribution of income across households. Our goal is to isolate the effect of changes in the relative price of capital on various outcomes y_{st} using regressions of the form

$$\Delta y_{st} = b_0 + b_t + b_1 \Delta \log q_{st} + \varepsilon_{st}, \tag{42}$$

where b_0 is a constant, b_t is a time fixed effect, and b_1 is the coefficient of interest. We consider three different outcomes Δy_{st}: the change in sectoral log investment ($\Delta \log i_{st}$), the change in sectoral capital-to-wage bill ratio ($\Delta \log(k/wn)_{st}$), or the change in sectoral college income share ($\Delta \text{share}^{\text{col}}_{st}$). The response of the college income share is particularly informative about the degree of what Krusell

FIGURE C.1: Scatterplot of Relative Price of Investment vs. College Income Share

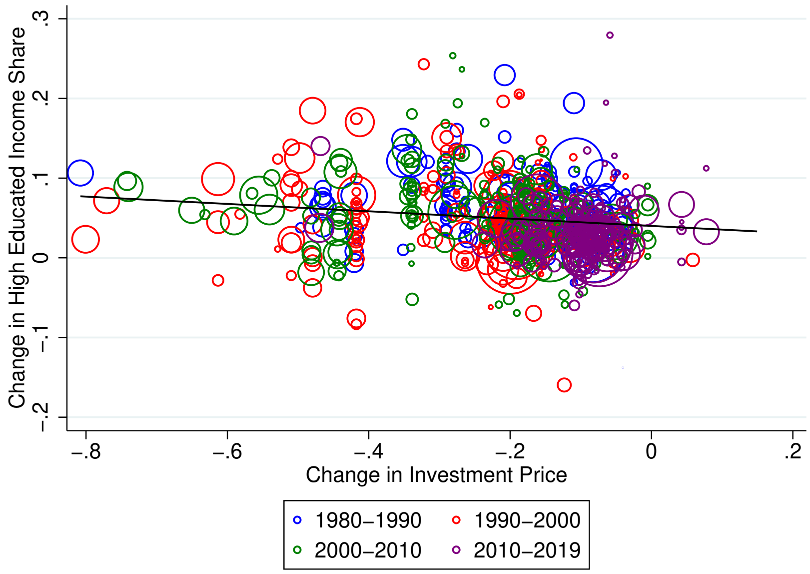

Notes: Scatterplot of ten-year change in college income share $\Delta\text{share}_{st}^{\text{col}}$ versus ten-year change in investment price $\Delta \log q_{st}$ using cross-industry variation. Size of circle reflects size of industry in year 2000 and different colors reflect different decades in our data. The weighted regression line through all the data has a slope coefficient of -0.046 with a standard error of 0.008.

et al. (2000) and others refer to as *capital-skill complementarity*. All else equal, if capital is more complementary to college-educated workers, then a decline in the price of capital would increase their labor income relatively more than non-college-educated workers' labor income. We control for aggregate trends in college attainment or labor-augmenting technical change through the time fixed effects b_t, so the coefficient b_1 is identified from cross-sectoral variation. Our key identifying assumption is that other determinants of y_{st} are orthogonal to changes in the relative price of capital q_{st} at the sector level.

For a sense of the underlying variation in the data, Figure C.1 shows a scatterplot of the ten-year changes in the college income share $\Delta\text{share}_{st}^{\text{col}}$ against the ten-year changes in the relative price of investment Δq_{st} at the sector level. The scatterplot shows a clear negative relationship between these two variables. For example, a log-point decline in the relative price of capital is associated with a 5 percentage point higher college income share. This negative relationship is robust across the four decades in our sample.

The left three panels of Table C.1 quantify these relationships using our regression specification (42) over ten-year changes. Column (1) shows that a 1 log point decline in the relative price of capital would increase investment by about 1.37 log points; this user cost elasticity of -1.37 is in line with the consensus range estimated in the investment literature (see the discussion in Zwick and Mahon

Table C.1: Long-Run Responses to Changes in the Relative Price of Investment

	10-Year Change			20-Year Change	
	(1)	(2)	(3)	(4)	(5)
	Δ Log Investment	Δ Log Capital to Wage Bill Ratio	Δ Log College Share	Δ Log Capital to Wage Bill Ratio	Δ Log College Share
Δ log q_{st}	-1.37	-0.467	-0.081	-0.467	-0.092
	(0.269)	(0.122)	(0.019)	(0.124)	(0.022)
R^2	0.390	0.685	0.169	0.477	0.165
Time FEs	yes	yes	yes	yes	yes

Notes: Response of change in sectoral investment (column 1), sectoral capital-to-wage bill ratio (columns 2 and 4), and the sectoral college income share (columns 4 and 5) to changes in the sectoral relative price of investment. Columns (1)-(3) measure ten-year changes between 1980 to 1990, 1990 to 2000, 2000 to 2010, and 2010 to 2020. Columns (4)-(5) measure twenty-year changes between 1980 to 2000 and 2000 to 2020. All regressions include period fixed effects. Standard errors clustered at the sectoral level and shown in parentheses. Regressions are weighted by the sectoral share of employment in 2000.

Table C.2: Response of Change in College Income Share to Change in Investment Prices, Cross Sector Variation, 10-Year Differences

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ log q_{st}	-0.047	-0.081	-0.063	-0.089	-0.096	-0.062	-0.072
	(0.019)	(0.019)	(0.023)	(0.039)	(0.029)	(0.026)	(0.037)
Adj. R-Squared	0.03	0.17	0.17	0.08	0.09	0.07	0.04
Time Fixed Effects	No	Yes	Yes	No	No	No	No
Control for Δ log $q_{s,t-1}$	No	No	Yes	No	No	No	No
Time Period	Pooled	Pooled	Pooled	1980s	1990s	2000s	2010s

Notes: Estimated α_1 from regression (42) for various time periods and various controls. Changes are 10-year (1980-1990 through 2010-2019). Data from Census/ACS. First three columns pool data across the 1980-2020 period allowing for four decadal changes per sector. The last four columns show results for each decade separately. Standard errors clustered at the sector level shown in parentheses.

(2017)). Consistent with this finding, column (2) shows that a 1 log point decline in the relative price of capital would increase the capital-to-wage bill ratio by 0.47. Most importantly, column (3) shows that these changes are accompanied by a 8.1 percentage point increase in the share of income accruing to college workers, consistent with capital-skill complementarity. [31]

Table C.2 contains three sets of robustness exercises on these empirical findings. First, compar-

[31] The regression estimate $b_1 = -0.081$ in Table C.1 is about twice as large as the slope of the regression line in the simple scatterplot (-0.046) due to our inclusion of time fixed effects. This finding highlights the importance of using sector-level variation in order to control for aggregate confounding factors.

Table C.3: Validating Long-Run Elasticities of Substitution

Long-Run Elasticity	**Data**	**Model**
Capital-to-wage bill elasticity	-0.467	-0.660
College income share semi elasticity	-0.081	-0.130

Notes: response of capital-to-wage bill ratio and college income share to long-run decline in the relative price of capital. "Data" column corresponds to the regression coefficients in columns (4) and (5) of Table C.1. "Model" column corresponds to comparing steady states with different values of q^*.

ing column (1) to column (2) shows that the coefficient on the college income share is somewhat smaller (but still significantly negative) if we exclude time fixed effects, illustrating the importance of controlling for aggregate factors. Second, column (3) shows that the results are robust to including the lagged relative price of capital to capture path-dependence in response to past shock. Finally, columns (4)-(7) show that these results hold in each decade of our sample, and are therefore not driven by a particular subperiod of data

Model Results. In order to compare our model to the data, we assume that the twenty-year changes in the right panel of Table C.1 identify the causal effect of a permanent change in the price of capital in our model. This procedure implicitly assumes that the heterogeneity in capital assets used to construct the capital price series can be approximated with one type of capital. While introducing capital heterogeneity would be an interesting extension of the model, it is outside the scope of this paper.

Table C.3 shows that our model is consistent with the long-run substitution patterns in the cross-sector data.[32] In both the model and the data, a decline in the relative price of capital induces firms to substitute toward capital and away from workers, although the substitution is somewhat stronger in the model than in the data. More importantly, however, the extent to which college workers disproportionately benefit from a decline in the relative price of capital—informative about the degree of capital-skill complementarity—is very similar in the model as in the data.

C.2 Validating Short-Run Elasticities of Substitution

Having shown that the model is consistent with rich patterns of capital-labor substitution in the long run, we now turn to validating the model's assumption that capital and labor are Leontief in the short run. We do so using temporary, rather than permanent, changes in the relative price of capital. In the model with standard capital, temporary changes in the relative price will induce

[32] We approximate the twenty-year changes with steady state comparisons in our model. We have verified that, following a permanent shock to the relative price of capital, the model approximately returns to steady state within ten years, indicating that this steady state comparison is a good approximation to the entire transition path.

large changes in the capital-labor ratios, due to intertemporal substitution. In the putty-clay model, however, capital-labor ratios can only adjust on newly purchased capital at the temporarily low price; since this new investment is a small fraction of the overall capital stock, the change aggregate capital-labor ratio — and therefore in the college income share — will be relatively small.

Empirical Results. Our source of temporary variation comes from the Bonus Depreciation Allowance, an investment stimulus policy that generates variation in the after-tax price of investment. To understand how the bonus affects investment incentives, we first define the present value of depreciation allowances per dollar of investment expenditures under the normal IRS tax schedule

$$\zeta_s = \sum_{t=0}^{T_s} \left(\frac{1}{1+r} \right)^t \widehat{\delta}_{st},$$

where r is the discount rate, $\widehat{\delta}_{st}$ is the fraction of investment expenditures of sector s that can be deducted from taxes t periods after purchase, and T_s is the tax life of assets in sector s. The depreciation schedule δ_s, and therefore the baseline present value ζ_s, varies across sectors due to heterogeneity in the mix of capital goods used in production.

The Bonus Depreciation Allowance increases the present value of tax deductions, and therefore decreases the after-tax price of investment, by allowing firms to immediately deduct some fraction θ_t of investment expenditures from their tax bill. In this case, the present value of depreciation allowances changes to

$$\zeta_{st} = \theta_t + (1 - \theta_t)\zeta_s. \tag{43}$$

Hence, the bonus generates time-series variation in the after-tax price of investment which differentially affects different sectors s due to the underlying heterogeneity in the baseline tax schedule ζ_s. The U.S. has used a $\theta_t = 30\%$, 50%, and 100% bonus at various points following the 2001 and 2008 recessions. Zwick and Mahon (2017) show that the Bonus Depreciation Allowance had a significant effect on firm-level investment.[33]

Table C.4 replicates Zwick and Mahon (2017)'s investment results using our sector-level BEA data. We focus on the 1998-2018 period, which is when the bonus was actually used, and run regressions of the form

$$\log i_{st} = b_s + b_t + b_1 \zeta_{st} + b_2 \log q_{st} + \varepsilon_{st}, \tag{44}$$

where b_s is a sector fixed effect, b_t is a year fixed effect, and b_1 is the coefficient of interest. We also control for the relative price of investment $\log q_{st}$. Column (3) includes sector and time fixed effects

[33]We downloaded the observations of ζ_{st} from Zwick and Mahon (2017)'s replication materials and merge them into the ACS data based on 4-digit NAICS codes.

Table C.4: Replicating Zwick and Mahon (2017) in Our BEA Data

	(1)	(2)	(3)
ζ_{st}	1.01	8.15***	1.86*
	(1.148)	(2.297)	(0.798)
$\log q_{st}$	-1.62***	-1.56***	-0.39*
	(0.225)	(0.211)	(0.152)
Observations	1178	1178	1178
R^2	0.767	0.789	0.979
Sector FEs	No	No	Yes
Time FEs	No	Yes	Yes

Notes: Estimated coefficients b_1 and b_2 from the regression $\log i_{st} = b_s + b_t + b_1\zeta_{st} + b_2 \log q_{st} + \varepsilon_{st}$, where b_s is a sector fixed effect and b_t is a time fixed effect. Regressions are weighted by the sector's share of aggregate investment. Standard errors are clustered at the sector level.

and is therefore most directly comparable to Zwick and Mahon (2017)'s differences-in-differences specification. Our estimated $b_1 = 1.86$ is about half the size of Zwick and Mahon (2017)'s estimated 3.69, but still implies large effects of the bonus; for example, a 4.8 cent change in ζ_{st} (the average change in the early period according to Zwick and Mahon (2017)) would increase investment by about 9% for our estimate vs. 18% in Zwick and Mahon (2017)'s estimate.

We now show that, even though the Bonus has a sizeable effect on investment, it had an insignificant effect on the college income share. To examine the effect of the bonus on capital-labor substitution, we use the empirical specification

$$\text{share}_{st}^{\text{col}} = b_s + b(t) + b_1\zeta_{st} + \varepsilon_{st}, \tag{45}$$

where b_s is a sector fixed effect, $b(t)$ is a set of controls for aggregate conditions, and b_1 is the coefficient of interest.[34] As discussed, the putty-clay model predicts that $b_1 \approx 0$ because firms will cannot substantially adjust their total capital-to-labor ratios on college and non-college workers.

Table C.5: Short-Run Response to Bonus Depreciation Allowance

	(1)	(2)
ζ_{st}	-0.020	-0.020
	(0.013)	(0.013)
Observations	3316	3316
R^2	0.993	0.995
Sector FEs	Yes	Yes
Aggregate Controls	Linear time trend	Sector-specific linear time trend

Notes: Estimated coefficient b_1 from the regression (45) in the text. Standard errors clustered by sector. "Sector-specific linear time trend" refers to controlling for a separate linear time trend in each sector.

[34]The specification (45) is similar in spirit to Zwick and Mahon (2017), who put investment rather than the college income share on the left-hand side.

FIGURE C.2: Response of College Income Share to Bonus Depreciation Shock

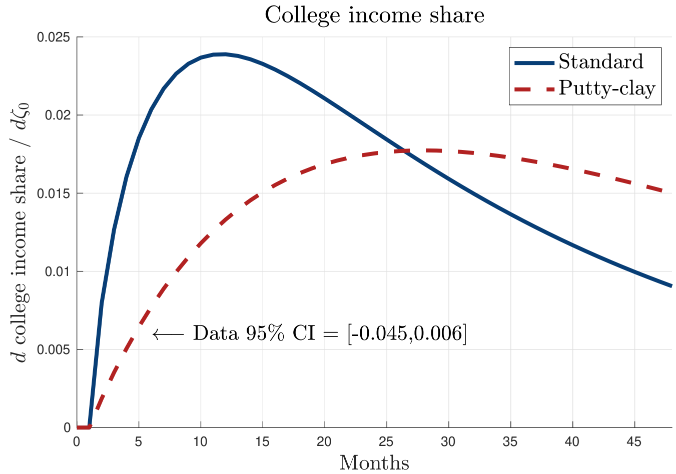

Notes: Partial equilibrium response of college income share to a transitory decline in q_t which mimics the bonus depreciation allowance. Specifically, the the economy starts in an initial steady state with $q^* = 1$ and then unexpectedly receives a shock q_0 such that the resulting decline in q_t equals $-\frac{\tau_f}{1-\tau_f}\Delta\zeta_{st}$, where $\tau_f = 35\%$ and $\Delta\zeta_{st}$ is the average change in ζ_{st} implied by a 50% bonus depreciation allowance in our data. The relative price of capital reverts back to steady state according to $\log q_{t+1} = \rho_q \log q_t$. We set ρ_q such that the half-life of the shock is one year. We assume $Q_{t,t+1} = \beta$ for all t.

Table C.5 shows that the data support the putty-clay model's prediction of insignificant capital-labor substitution in the short run. Column (1) controls for $b(t)$ specified as a linear time trend and shows that the estimated coefficient b_1 is insignificantly different from zero with fairly tight standard errors. Column (2) shows that the estimate is nearly unchanged if we instead control for sector-specific linear time trends.[35]

Model Replication. We replicate the Bonus Depreciation Allowance in our model as a transitory shock to the relative price of capital q_0 starting from an initial steady state with $q^* = 1$. We set q_0 so that the change in the relative price of capital equals the change in the tax-adjusted user cost $\frac{1-\tau_f\zeta_{st}}{1-\tau_f}$ from a 50% bonus. We assume the shock reverts back to steady state according to $\log q_{t+1} = \rho_q \log q_t$ and set the autocorrelation such that the half-life of the shock is one year. Finally, we compute the transition path in "partial equilibrium," that is, setting $Q_{t,t+1} = \beta$ for all t. This procedure assumes that the aggregate controls in the regression $b(t)$ absorb general equilibrium changes in the discount factor $Q_{t,t+1}$.

Figure C.2 plots the response of the college income share relative to the implied change in ζ_0 in our model. The college income share immediately and sharply rises in the model with standard capital,

[35]Controlling for $b(t)$ specified as a year fixed effect does not significantly change the point estimates, but increases the standard errors by an order of magnitude. This occurs because most of the variation in the policy variable ζ_{st} is at the aggregate level, so the fixed effects remove most of the useful variation.

FIGURE D.1: Distribution of Aggregate Income

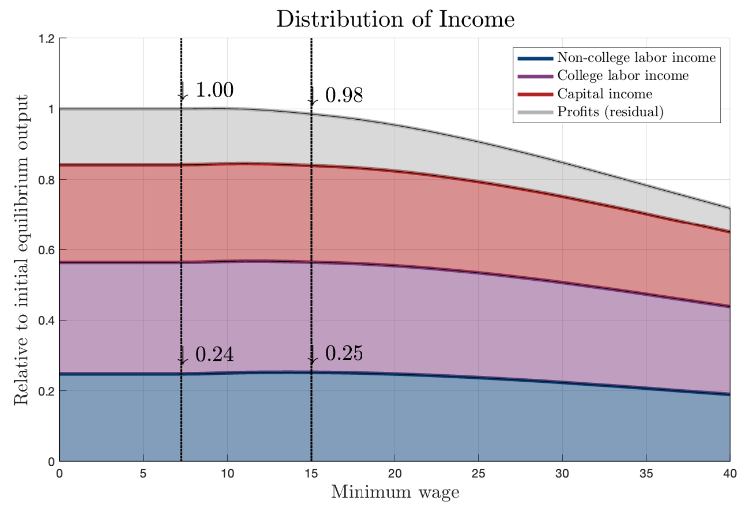

Notes: Plots steady-state GDP and the share accounted for by non-college labor income, college labor income, capital income, and residual profits as a function of the minimum wage $\overline{w}$. The y-axis is normalized such that aggregate income equals 1 without the minimum wage. The x-axis is the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

reflecting the immediate capital-labor substitution favoring college-educated workers. In contrast, the college income share rises much more slowly in the putty-clay model because firms can only substitute away from workers on newly purchased capital. To compare to the regression coefficient from (45), note that the data measures income cumulated over the course of the year, while Figure C.2 plots monthly income; we roughly account for this difference by comparing the model response six month after the shock. The putty-clay model implies that the college income share increases by 0.0077, just outside the 95% confidence interval of the data; the model with standard capital implies a three times larger response, far outside the model's confidence interval. Hence, the model with standard capital implies an unrealistically large degree of capital-labor substitution in the short run, while the putty-clay model does not.

D Additional Results About the Minimum Wage

This Appendix contains a number of additional results about the minimum wage referenced in the main text.

Figure D.1 plots how the minimum wage shapes the distribution of GDP into the four components mentioned in the main text: non-college labor income, college labor income, capital income, and residual profits, which include both the vacancy posting costs and pure monopsony profits. In the initial equilibrium, profits are a substantial share of aggregate income due to the firms' monopsony power. The minimum wage reduces these profits by making workers more expensive and shrinking firms' markdowns (see below for more details regarding the behavior of markdowns). Relatively small levels of the minimum wage will nonetheless increase labor income, effectively redistributing a portion of those lost profits to workers. However, larger levels of the minimum wage will decrease labor income, reflecting the Laffer curves described above.

This discussion illustrates the key policy tradeoff from the aggregate perspective: an appropriately chosen minimum wage may reduce monopsony profits and increase labor income, but at the cost of decreasing employment and therefore total output in the economy. For example, a \$15 minimum wage increases aggregate non-college labor income by 1.5% but, because it reduces employment, decreases total GDP by 1.6% (and the negative consequences for GDP worsen at higher levels off the minimum wage). One might think this tradeoff is worthwhile if the main goal is increase aggregate labor income of this particular group; however, as we show in the main text, such a policy would disproportionately harm the lowest-income workers within that group.

D.1 Role of Monopsony Power

Figure D.2 shows that effects of the minimum wage depend crucially on the degree of monopsony power ω. For example, with a larger degree of monopsony power—consistent with the average markdown of 0.65 estimated by Hershbein, Macaluso and Yeh (2020) as opposed to our baseline value of 0.71— the \$15 minimum wage decreases aggregate non-college employment by 5% (compared to 9.7% in our baseline) and increases aggregate non-college labor income by more than 4% (compared to 1.7% in our baseline). At the micro level, that degree of monopsony power enlarges the set of workers who benefit from the minimum wage and reduces the employment losses for the set of workers who are hurt by the minimum wage.

In contrast, if there is no monopsony power in the economy ($\omega \to \infty$), then the minimum wage unambiguously reduces employment and labor income for all workers. In this case, competitive search induces wages and employment to move to their efficient level, so any further increase in the minimum wage increases the private marginal cost of hiring workers above the social cost. This result highlights the importance of monopsony power in motivating the use of the minimum wage a policy tool.

D.2 Long-Run Effects of the Minimum Wage on College Workers

Figure D.3 plots our aggregate Laffer curves for college workers. The Laffer curve for employment is more complicated than for non-college workers because the minimum wage has both direct and indirect effects on college employment. The indirect effects occur for low values of the minimum wage (for which the minimum wage is not binding for any college workers). In this region, college

FIGURE D.2: Comparative Statics to Monopsony Power ω

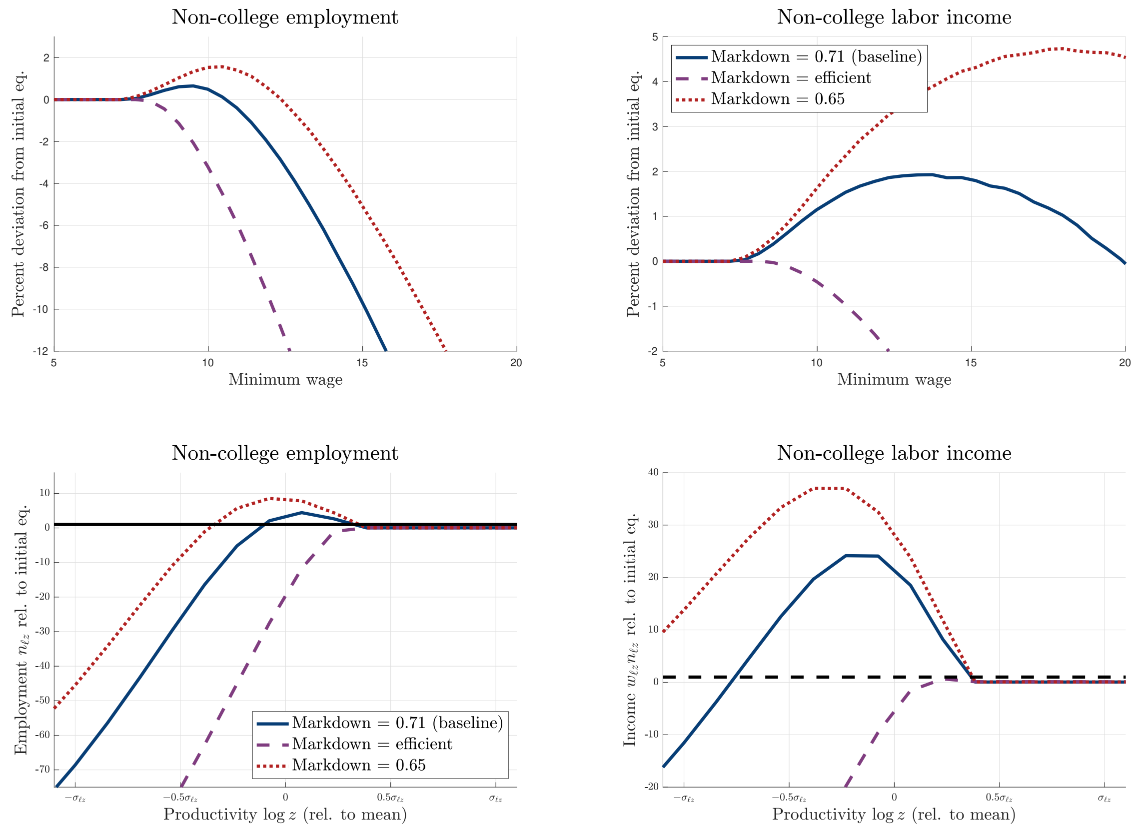

Notes: steady state outcomes for different markdowns. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure 6. Different markdowns correspond to different parameter values for ω. For each value of ω, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

employment falls because lower non-college employment reduces the marginal product of college workers, lowering their labor demand. However, eventually the minimum wage becomes binding for college workers as well, generating a hump shape as for the non-college workers in the main text.

D.3 Comparative Statics

This subsection contains a number of comparative statics to show how the effects of the minimum wage depend on key parameters. We focus on the long-run effects of the minimum wage unless otherwise noted. For the sake of parsimony, we show comparative statics about the distributional effects of the \$15 minimum wage (similar to Figure **??** in the main text), but we obtain similar implications for the aggregate Laffer curves.

Figure D.3: Long-Run Effects of the Minimum Wage on College Workers

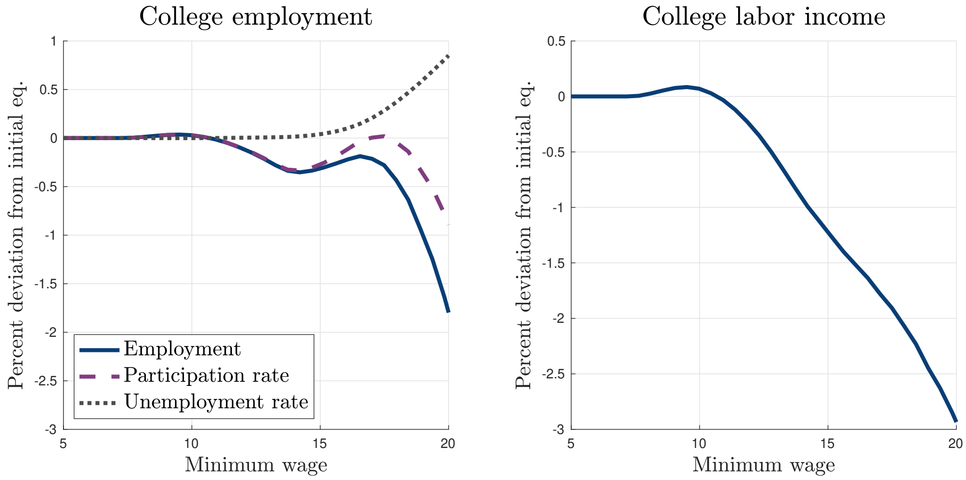

Notes: steady state outcomes as a function of the minimum wage $\overline{w}$. Left panel plots the log-change of aggregate college employment, the log-change of aggregate college labor force, and the change in the college unemployment rate relative to their levels in the initial steady state with $\overline{w} = 0$ (see the decomposition (33)). Right panel plots aggregate labor income of college workers. The x-axis is the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

Labor Supply. The first set of parameters we consider govern households' optimal labor supply. First, in Figure D.4 we show how the distributional effect of the minimum wage depends on the elasticity of labor supply, γ^{-1}. The response of employment, either negative (for the low-z types) or positive (for the high-z types) is smaller if the labor supply elasticity is smaller. However, the region over which the minimum wage has a positive effect is shifted left for less elastic labor supply because average wages are higher. The second labor supply parameter we consider is the disutility of search, $\chi_{b,s}$, which we fixed to $\chi_{b,s} = 100$ for both groups in the main text. Figure D.5 shows that our main results are largely insensitive to changes in this parameter value around our baseline.

Capital-Labor Substitution. As discussed in the main text, the long-run degree of "capital-skill complementarity" is given by $\rho - \alpha$. Figure D.6 shows that the effect of the minimum wage on non-college employment is generally more negative when capital-skill complementarity is higher (driven by a smaller value of α).

Within-Group Worker Substitution. Figure D.7 shows how our results depend on the elasticity of substitution between workers of different z types within the same education group, ϕ. The literature offers little guidance on the value of this parameter, so we consider values based on estimates of across-education group elasticities. Our baseline value of $\phi = 3$ is somewhat more substitutable than estimates from Katz and Murphy (1992). Figure D.7 shows that the minimum wage has less negative

FIGURE D.4: Long-Run Comparative Statics w.r.t. Labor Supply Elasticity γ^{-1}

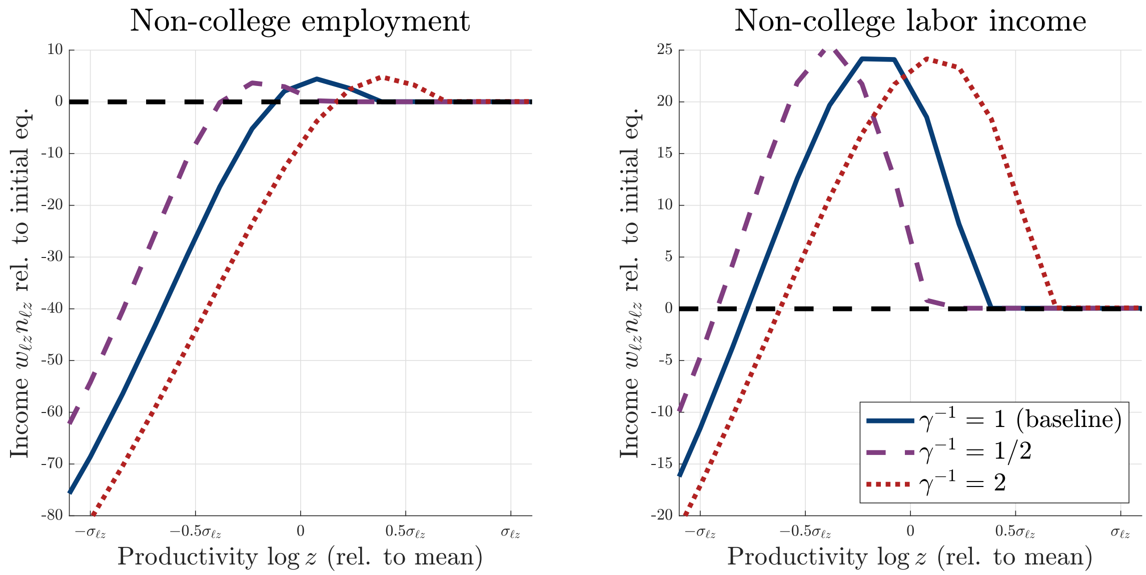

Notes: steady state outcomes for different labor supply elasticities γ^{-1}. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure **??**. For each value of γ^{-1}, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

FIGURE D.5: Long-Run Comparative Statics w.r.t. Search Cost $\chi_{b,s}$

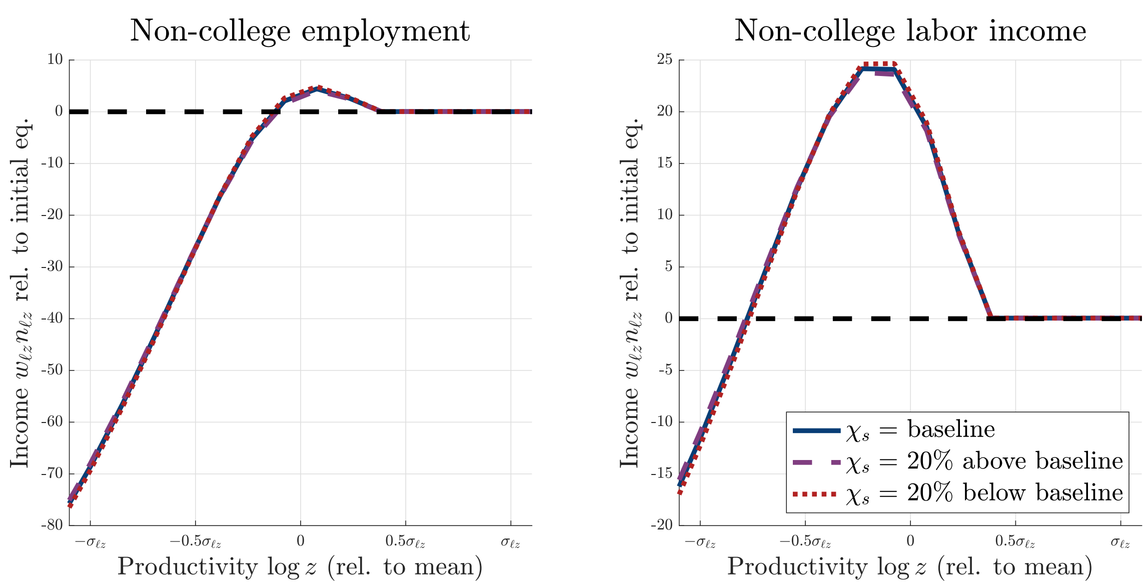

Notes: steady state outcomes for different disutilities of search $\chi_{b,s}$. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure **??**. For each value of $\chi_{b,s}$, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

Figure D.6: Long-Run Comparative Statics w.r.t. Elasticity of Substitution α

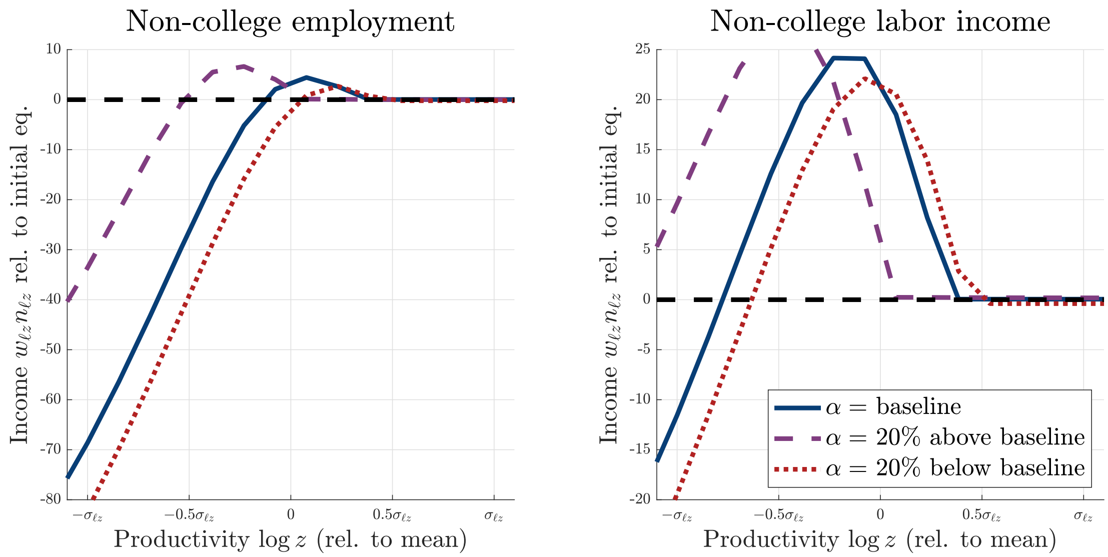

Notes: Notes: steady state outcomes for different elasticities of substitution α. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure ??. For each value of α, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

effects at $\phi = 2$ because it is harder to substitute away from the affected workers. More recently, Bils, Kaymak and Wu (2020) estimates an across-group elasticity of 4; using $\phi = 4$ implies that the minimum wage has even more negative effects because it becomes easier to substitute away from the affected workers.

Search Frictions. Figure D.8 shows our results are fairly insensitive to the value fo the vacancy-posting costs, κ_0.

D.4 Comparison of Models with Putty-Clay vs. Standard Capital

Figure D.9 compares the transition dynamics in response to the \$15 minimum wage in our model with putty-clay capital technology to the version of our model with standard capital. The left panel shows that non-college employment converges to the new steady state much more quickly in the model with standard capital; in that model, employment falls by 5.4% in the first year after the introduction of the minimum wage, compared to only falling 1.1% over the same time period in our putty-clay model. The right panel shows that the differences in these dynamics are mirrored in differences in the dynamics of the labor-to-capital ratios across the two models. In the model with standard capital, the aggregate non-college labor-to-capital ratios falls immediately, while in our

Figure D.7: Long-Run Comparative Statics w.r.t. Elasticity of Substitution ϕ

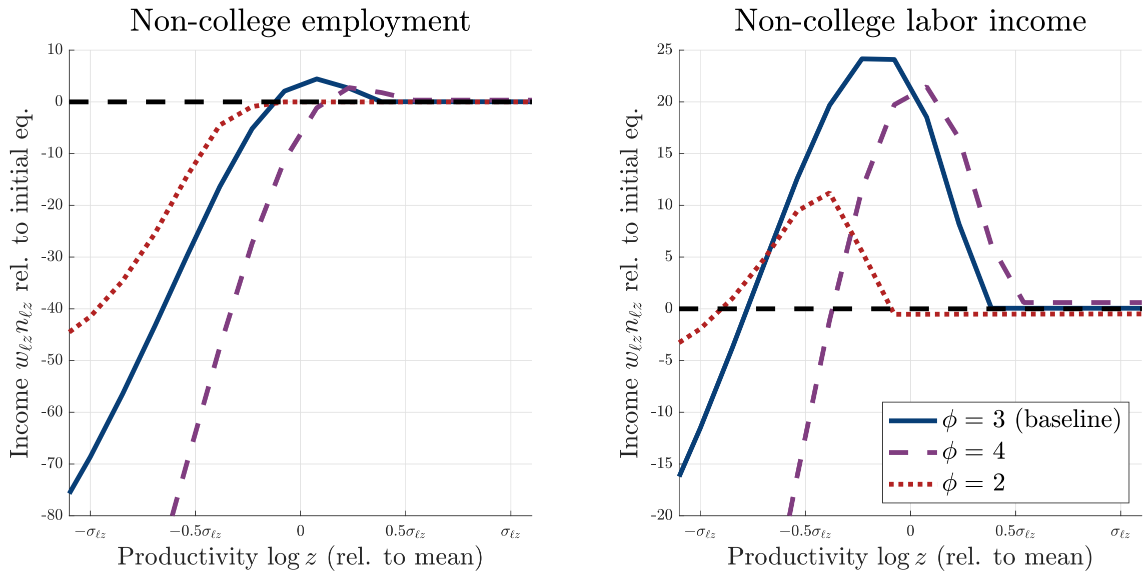

Notes: Notes: steady state outcomes for different elasticities of substitution ϕ. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure **??**. For each value of ϕ, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

Figure D.8: Long-Run Comparative Statics w.r.t. Vacancy-Posting Costs κ

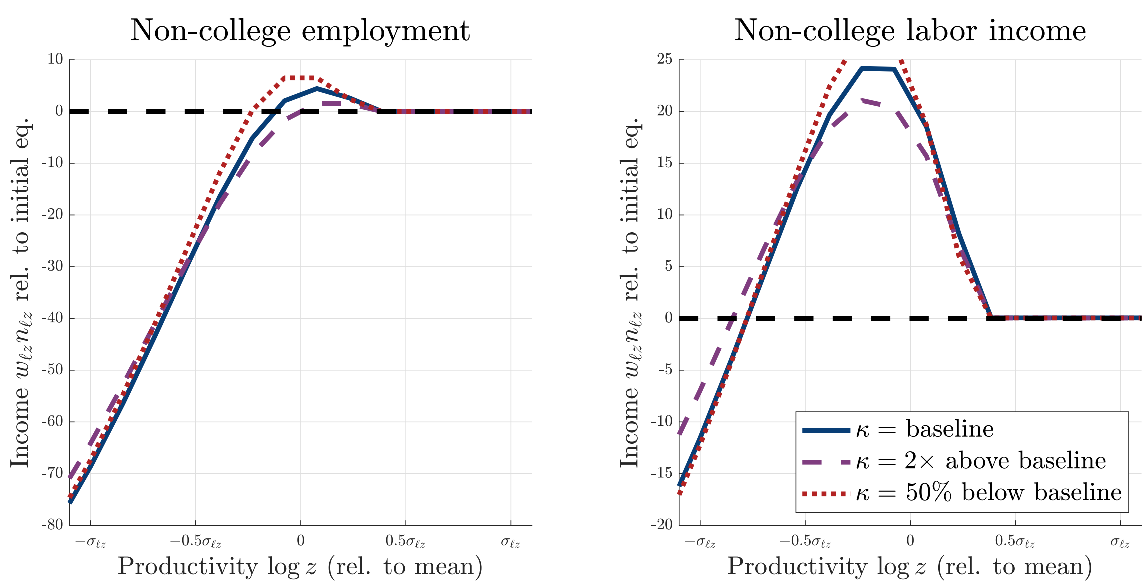

Notes: Notes: steady state outcomes for different vacancy-posting costs κ. Top panel plots aggregate non-college employment (left) and non-college labor income (right), as in Figure 4. Bottom panel plots the micro-level effect of the \$15 minimum wage as a function of individual productivity z, as in Figure **??**. For each value of κ, we compute the level of the minimum wage $\overline{w}$ whose level relative to the median non-college wage in the initial equilibrium is the same as in the data.

FIGURE D.9: Transition Path to \$15 Minimum Wage Steady State, Putty-Clay vs. Standard Capital

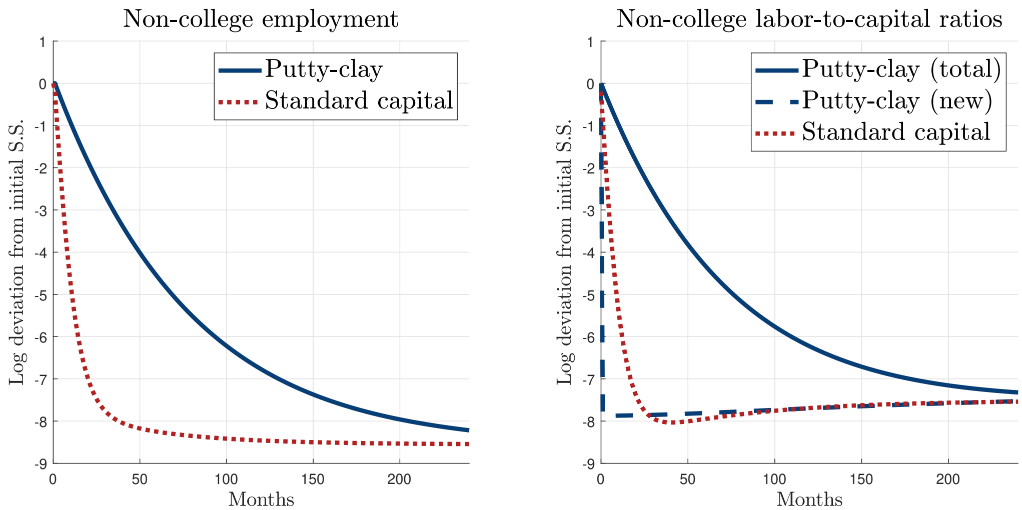

Notes: Transition path following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. Left panel plots the aggregated employment of non-college workers in our baseline model with putty-clay technology vs. the model with standard capital. Right panel plots the associated labor-to-capital ratios in the two models.

putty-clay model it takes much longer for the reasons described in the main text.

E Additional Results About Alternative Policies

E.1 Progressive Tax and Transfer System

We model the U.S. progressive tax and transfer system using the parametric form

$$T(w_i) = w_i - \lambda w_i^{1-\tau}, \tag{46}$$

where $T(w_i)$ is the labor income tax schedule, λ and τ are parameters, and negative taxes $T(w_i) < 0$ indicate transfers. In recent work Heathcote, Storesletten and Violante (2017) show that $\tau = 0.181$ provides a good description of the progressivity of the U.S. tax and transfer system, except at the very bottom because transfers are phased in and out at various income levels, inducing kinks in the transfer system.[36] For this analysis, we abstract from the EITC. We set the progressivity parameter $\tau = 0.181$ in order to match the estimated progressivity from Heathcote, Storesletten and Violante (2017) for the United States and choose the scale parameter λ to ensure that the aggregate net

[36] See pg. 1700: "... our tax/transfer scheme tends to underestimate marginal tax rates at low income levels ... marginal rates vary substantially across households, and some households simultaneously enrolled in multiple welfare programs face high marginal tax rates where benefits are phased out. Although our parametric functional form cannot capture this variation in tax rates at low income levels ..."

FIGURE E.1: Progressive Tax and Transfer System

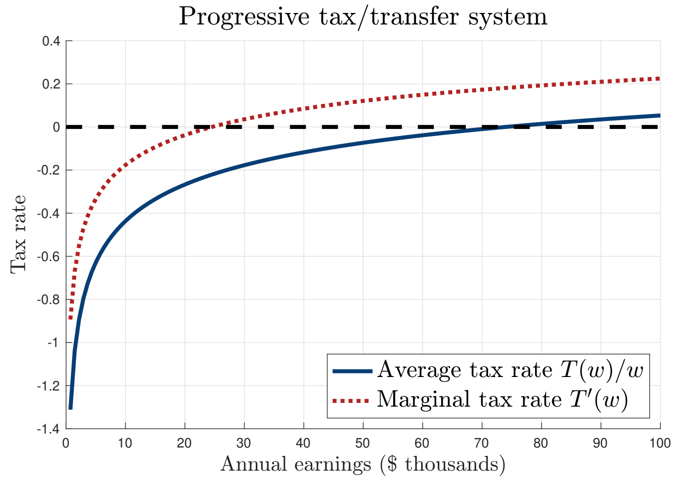

Notes: Average tax rates $T(w)/w$ and marginal tax rates $T'(w)$ from the budget-equivalent tax and transfer system described in the main text (with the U.S. level of progressivity $\tau = 0.181$). The x-axis rescales steady state labor income to annual earnings assuming each household works 1800 hours per year.

transfer payment is budget-equivalent to the \$15 minimum wage, as described in the main text.

Figure E.1 plots the average tax rates $T(w_i)/w_i$ and marginal tax rates $T'(w_i)$ of our budget-equivalent system as a function of income. The fact that the system is progressive, in that $\tau > 0$, implies that marginal tax rates are higher than average tax rates throughout. The lowest-income households have both a negative average tax rate (indicating they are receiving transfers $A(w_i) > w_i$) and a negative marginal tax rate (indicating that the transfers are being phased in, in the sense that $A'(w_i) > 1$). For these households, the tax and transfer system reduces the monopsony distortion in (36) and increases search effort in (37). Middle-income households continue to receive transfers, which encourage search effort, but face positive marginal tax rates, which exacerbate the monopsony distortion. Finally, high-income households face both positive average and marginal tax rates, reducing their search effort and exacerbating monopsony distortions. In this sense, a progressive tax and transfer system differentially alleviates monopsony distortions for low-z and high-z types.

The left panel of Figure E.2 shows that this progressive system succeeds in substantially increasing employment of non-college workers, especially for the low-z workers for whom the system is especially targeted. The right panel shows that similar results also hold for after-tax/transfer labor income. Labor income increases for all the z-types plotted in the figure, although it falls for some of the higher-z types not pictured, who earn higher wages (especially among college workers).[37]

[37] In Appendix E, we also study the effect of a budget-equivalent uniform subsidy (with $\tau = 0$) instead of the progressive subsidy (with $\tau = 0.18$) offered here. We find that the uniform subsidy increases the employment and labor income of equally for all types of workers. While that outcome is an improvement over the minimum wage for these workers, the progressive system is even more effective at increasing employment given its targeted nature.

Figure E.2: Effect of Progressive Tax System on Non-College workers

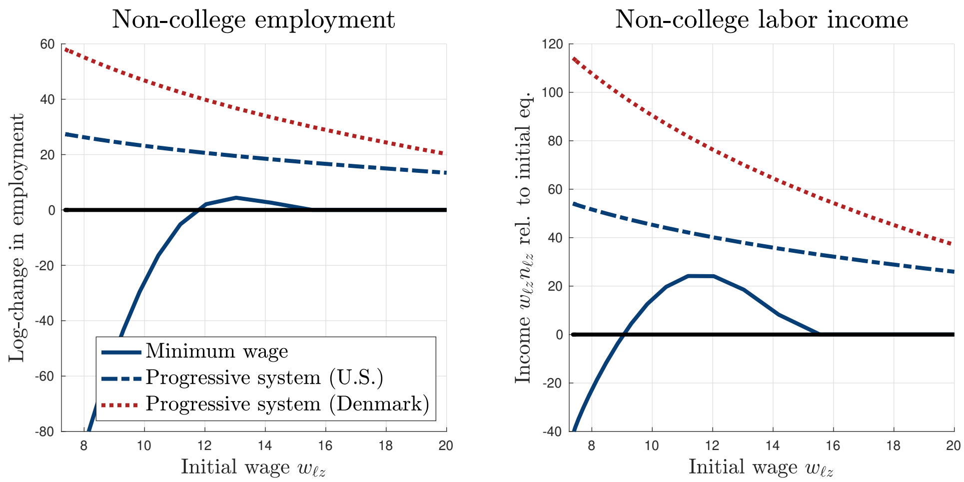

Notes: Steady-state employment (left panel) and labor income (right panel) of particular z types for three different policies: the \$15 minimum wage (blue line), the budget-equivalent progressive tax system with the U.S. level $\tau = 0.181$ (red line), and the budget-equivalent tax system with the Danish level $\tau = 0.463$ (purple line). The y-axis is normalized relative to employment in the initial equilibrium (without any policies). The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.

For the sake of comparison, Figure E.2 also shows the effect of an even more progressive tax and transfer system with $\tau = 0.463$ meant to capture degree of progressivity in Denmark.[38] This more progressive system further increases employment of the low-z types but decreases employment of the college workers more than the U.S. system (not pictured), highlighting the policy tradeoffs associated with different levels of progressivity.

E.2 Linear Income Tax Credit

As a simple example to build intuition, consider a uniform linear income tax credit of rate τ_c starting from an allocation with no taxes. This policy is a special cause of the general system (46) with the progressivity parameter $\tau = 0$, and the scale parameter $\lambda = 1 + \tau_c > 1$.

The left panel of Figure E.3 shows that this tax credit increases employment by approximately 1.5% uniformly across workers. On the labor demand side, the tax credit, by increasing marginal after-tax income by $A'(w_i) = 1 + \tau_c > 1$, alleviates the monopsony distortion in (36). In this sense, it is analogous to the subsidy used to correct monopoly distortions in New Keynesian models with monopolistic competition. On the labor supply side, the tax credit increases average after-tax income

[38] Heathcote, Storesletten and Violante (2020) estimate the progressivity of the tax system in Denmark and a number of other countries. Unfortunately, due to data limitations across countries, these estimates only include taxes but not transfers. We impute a value of the progressivity of the tax and transfer system by scaling our baseline progressivity parameter of the U.S. tax and transfer system by the ratio of the progressivity of the Danish to U.S. tax systems.

Figure E.3: Effect of Uniform Tax Credit and Tax Cut on Non-College Workers

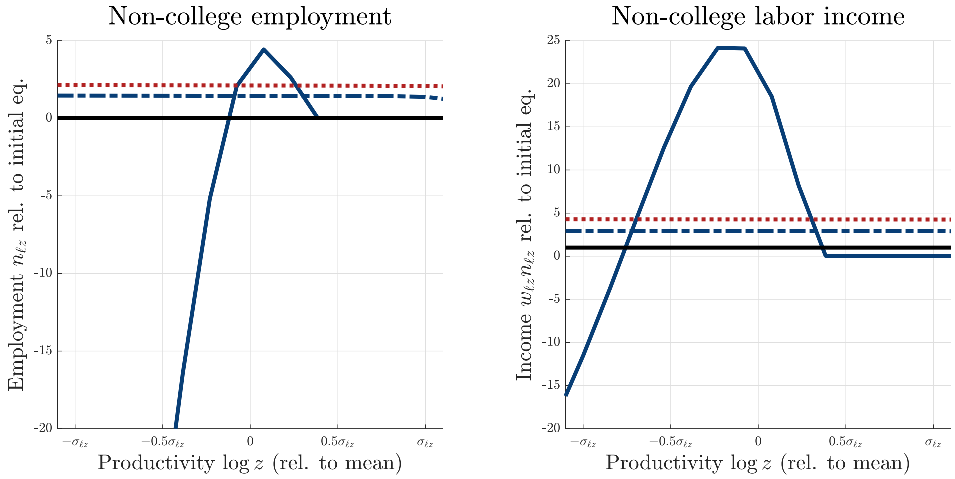

Notes: Steady-state employment or labor income of particular z types for three different policies: the \$15 minimum wage (blue line), the budget-equivalent tax credit (red line), and the budget-equivalent tax cut starting from a 20% tax rate (purple line). The y-axis is normalized relative to employment in the initial equilibrium (without any policies). The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.

by the same amount $A(w_i) - w_i = 1 + \tau_c$, increasing search effort in (37). These two forces increase employment equally across types z. The right panel of Figure E.3 shows that the subsidy unfiromly increases labor income as well.

Of course, in reality labor income taxes are positive, so we also consider the effect of a uniform tax cut starting from a positive tax rate $\tau > 0$. In particular, we compute the initial equilibrium of the model with a baseline tax rate of $\tau = 1 - \lambda = 20\%$, recompute the effects of a \$15 minimum wage in this alternative model, and then offer households a budget-equivalent tax cut. Figure E.3 shows that the tax cut increases employment by even more than the baseline tax credit described above. This occurs because the positive tax rate is a further distortion reducing employment and the tax cut reduces this distortion in addition to alleviating the monopsony distortions.

E.3 Interactions Between Transfer Programs and the Minimum Wage

Figure E.4 isolates the interaction between the constant transfer and the minimum wage by subtracting off the pure effect of the minimum wage in isolation. Figure E.5 shows how a $\overline{w} = \$12$ minimum wage interacts with our budget-equivalent EITC from Section 5.

Figure E.4: Effect of EITC with and without $\overline{w} = \$8$ Minimum Wage

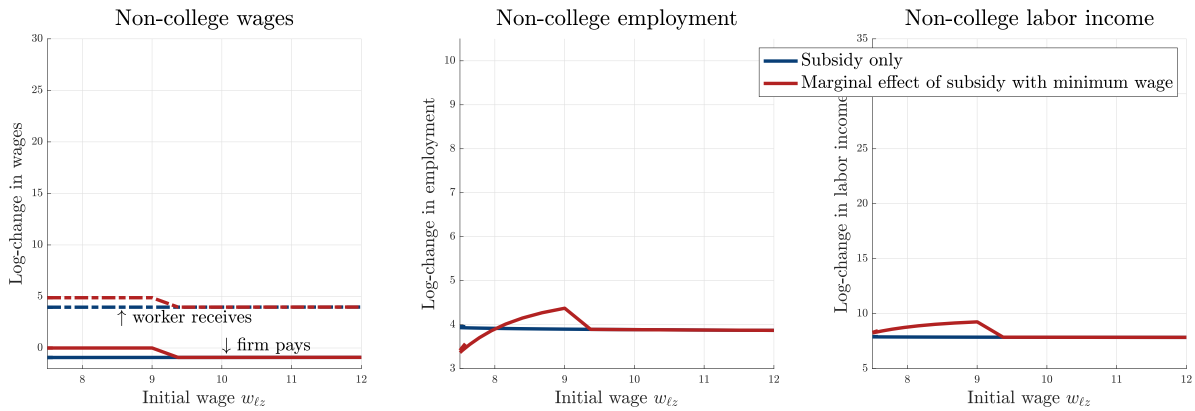

Notes: Steady-state wages (left panel), employment (middle panel), and labor income (right panel) in response to our budget-equivalent EITC system from Section 5. Lines with subsidy and minimum wage together have subtracted off the effect of the minimum wage in isolation. The y-axis is the log-change relative to initial equilibrium (without any policies). The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.

Figure E.5: Effect of EITC with and without $\overline{w} = \$12$ Minimum Wage

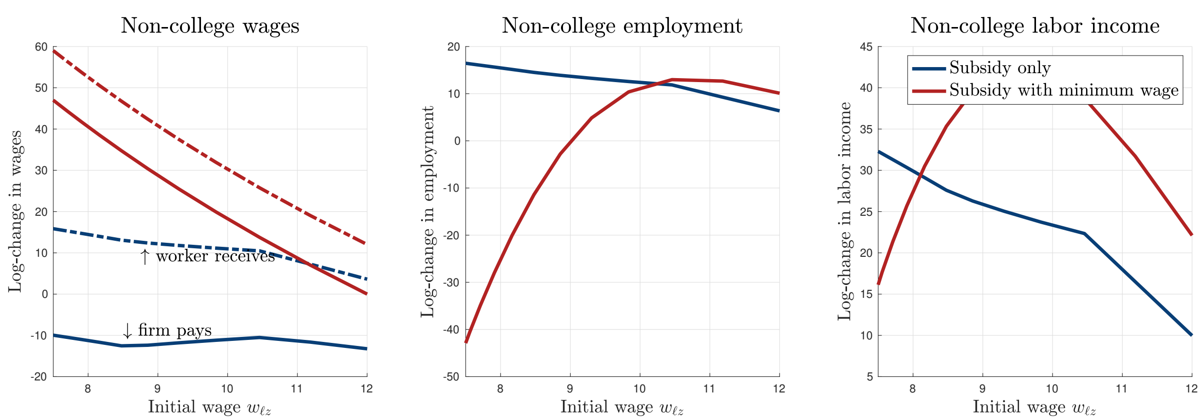

Notes: Steady-state wages (left panel), employment (middle panel), and labor income (right panel) in response to our budget-equivalent EITC system from Section 5. The y-axis is the log-change relative to initial equilibrium (without any policies). The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.