The Distributional Impact of the Minimum Wage in the Short and Long Run*

Erik Hurst
Chicago Booth
Erik.Hurst@chicagobooth.edu

Patrick Kehoe
Stanford and Minneapolis Fed
pkehoe@stanford.edu

Elena Pastorino
Stanford and Hoover Institution
epastori@stanford.edu

Thomas Winberry
Wharton
twinb@wharton.upenn.edu

July 2022
(original version: June 2021)

Abstract

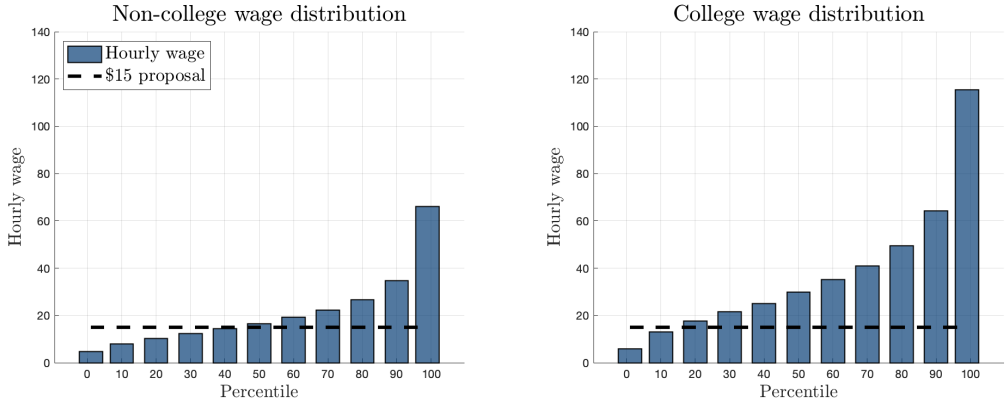
We develop a framework with rich worker heterogeneity, firm monopsony power, and putty-clay technology to study the distributional impact of the minimum wage in the short and long run. Our production technology is disciplined to be consistent with the small estimated employment effects of the minimum wage in the short run and the large estimated elasticities of substitution across inputs in the long run. We find that in the short run, a large increase in the minimum wage has a small effect on employment and therefore increases the labor income of the workers who were earning less than the new minimum wage. In the long run, however, the minimum wage has perverse distributional implications in that it reduces the employment, income, and welfare of precisely the low-income workers it is meant to help. Nonetheless, these long-run effects take time to fully materialize because firms slowly adjust their mix of inputs. Existing transfer programs, such as the earned income tax credit (EITC), are more effective at improving long-run outcomes for workers at the low end of the wage distribution. But combining existing programs with a modest increase in the minimum wage generates even larger welfare gains for low-earning workers.

*We especially thank Daron Acemoglu, David Autor, Mark Bils, Charlie Brown, John Campbell, Jeff Clemens, Ashley Craig, Amanda Kowalsky, and Richard Rogerson for many helpful comments as well as seminar participants at the Bank of Spain, the Barcelona Graduate School of Economics Summer Forum 2021, Berkeley, Cambridge, Chicago, the Federal Reserve Banks of Philadelphia, Richmond, and San Francisco, Harvard, Michigan, MIT, the NBER Summer Institute 2021, Princeton, Santa Cruz, the SED 2022 Meetings, Stanford, USC, Virginia, the Virtual East Asia Seminar Series, Wharton, and Wisconsin (Madison). The current draft is an updated version of the one presented at the Barcelona Graduate School of Economics Summer Forum in June 2021. 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 to increase the federal minimum wage from its current level of $7.25 per hour to at least $15 have been advanced in the United States. Their goal 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 magnitude would impact a large fraction of the U.S. workforce. As apparent from Figure 1, which shows the distribution of hourly wages from the 2017-2019 American Community Survey (ACS) by college attainment, about 40% of non-college-educated workers and 10% of college-educated workers currently earn a wage lower than $15.[1] The high degree of wage dispersion within each education group suggests that such a minimum wage would have very different effects on workers earning such different wages. For example, given the current wage distribution, a $15 minimum wage would dramatically increase the wages of workers in the bottom 20% of the non-college wage distribution but would not bind on workers in the top 60% of it. The variation in wages within each education group is an order of magnitude larger than the variation across groups.[2]

FIGURE 1: Empirical Wage Distributions



Note: Wage distribution of full-time employed workers over the age of 16 from the 2017-2019 ACS data. Non-college educated workers are defined as those workers with less than a bachelor's degree.

The goal of this paper is to develop a framework to assess the distributional impact of the minimum wage in both the short and long run. Given the substantial heterogeneity in wages documented in Figure 1, a key question is: to what extent will firms substitute away from workers for whom a large minimum wage binds? The empirical literature offers two sets of estimates of such elasticities of substitution. First, an influential literature has provided estimates of the degree to which different groups of workers are

[1] 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 can be found in the online appendix.

[2] Hence, in a model with only two types of workers, as represented by the median non-college-educated worker and the median college-educated worker, a $15 minimum wage would not bind on either education group, since their median wages are $16.40 and $29.80, respectively. This observation helps clarify why an analysis of the distributional impacts of the minimum wage must, at its heart, be about its differing impacts on workers within such groups.

substitutable over longer periods of time. Prominent studies include Katz and Murphy (1992), which estimates an elasticity of substitution *across* workers from different education groups of about 1.5, and Card and Lemieux (2001), which finds an elasticity of substitution among workers *within* an education group between 4 and 6. This literature highlights that workers with different levels of skills are fairly substitutable with each other in the long-run. Second, a separate and influential literature has documented only small employment effects one or two years following a minimum wage increase, implying a low elasticity of substitution across workers in the short run.[3] Taken together, these estimates imply that a study of the distributional impact of the minimum wage across workers must distinguish between short- and long-run effects.

Motivated by these observations, our framework incorporates rich worker heterogeneity, a production technology that is consistent with long-run estimates of input substitutability, and a putty-clay input adjustment process that is consistent with short-run estimates of input substitutability. We also allow for firms' monopsony power in labor markets subject to search frictions to capture the possibility that a higher minimum wage can potentially increase both employment and labor market participation.

Our main message is that a permanent increase in the minimum wage to $15 has beneficial effects for low-earning workers in the short run but detrimental effects for them in the long run. In the short run, even a sizable increase in the minimum wage induces only a small adjustment in the employment of workers who initially earn less than the new minimum wage. Hence, an increase in the minimum wage leads to an increase in labor income and welfare for such workers. Over time, though, as firms reorganize their production in response to the higher minimum, they start substituting away from these workers on whom the minimum wage binds and towards those on whom it does not. In light of the long-run costs of a large increase in the minimum wage, we also show that other policies such as an expansion of the earned income tax credit (EITC)—potentially combined with a small increase in the minimum wage—dominate a large change in the minimum wage in the long run for lower-income workers.

Four features of our model are key for our study. First, since our goal is to explore the distributional impact of the minimum wage, we incorporate rich worker heterogeneity in order to match the wage distributions in Figure 1. Specifically, we classify workers into two educational groups, namely, non-college educated and college-educated workers, and within each group, we let workers differ in their productive ability 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 a production technology that allows for rich substitution patterns among inputs in the long run. Specifically, we allow for firms to substitute (i) among workers of different productivities *within* an education group, (ii) among workers *across* different education groups,

[3] We discuss this literature in more detail later on in the paper.

and (iii) between labor and capital. As noted, we discipline the first two margins of substitutability by using estimates from Card and Lemieux (2001) and Katz and Murphy (1992), respectively.[4] Quantitatively, the most important elasticity for our results is the one governing the substitutability among workers within an education group. When this elasticity is high, firms have strong incentives to substitute away from low-productivity workers after the imposition of a $15 minimum wage.

The third key feature of our model is putty-clay technology, which leads to different degrees of short-run and long-run substitutability across inputs.[5] This type of technology captures the idea that when deciding on a production plan and so on new investments, firms can choose any desired ratio of capital to the labor of workers of each type that lies on the frontier of the underlying CES production function (putty phase). Once installed, however, each unit of capital is clay-like in that it requires a fixed amount of each type of labor to operate. That is, our production function is effectively CES in the long run but Leontief in the short run. As such, it implies that the short-run demand for each type of labor is inelastic because it is infeasible to substitute among the various inputs of production, once a production plan is established. In the short run, then, our model is consistent with the view that the minimum wage can be set quite high without adversely affecting employment, as long as it remains profitable for firms to continue to operate. Over time, though, as new capital goods embodying new ratios of capital to labor are installed, firms' capital-to-labor ratios shift to their preferred mixes governed by the underlying CES production function. Hence, this putty-clay specification provides a parsimonious way to capture the fact that firms face adjustment costs to altering their mix of *any* inputs.

The final key feature of our model is monopsonistic competition among firms, which we embed within a directed search framework that features endogenous labor market participation. Through this setup, we formalize the idea, dating back to at least Robinson (1933), that increases in the minimum wage can help alleviate some of the distortions arising from firms' wage-setting power and so lead to a desirable redistribution of firm profits to workers as 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. In this framework, firm monopsony power not only distorts the participation margin of workers deciding whether to look for a job, but also a hitherto neglected margin, namely, the job-creation margin of firms deciding whether to hire more workers. When we analyze tax policies, such as the EITC, we find their effects as well to subtly depend on how these policies affect both of these margins of labor market participation and job creation. The framework we propose 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. In our framework, indeed, 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

[4] We use the elasticities from Card and Lemieux (2001) and Katz and Murphy (1992) in our baseline estimation. We show our main results are sensitive to a range of other estimates in the literature.

[5] Our modeling of putty-clay capital builds on the work that dates back at least to Johansen (1959) and was extended by Calvo (1976). Our formulation follows that in Atkeson and Kehoe (1999).

vacancies until this excess supply is dissipated.

The degree of firms' monopsony power in the labor market is a crucial parameter in our quantitative exercise. We discipline it so that our model matches recent estimates on wage markdowns from Seegmiller (2021), Lamadon, Mogstad and Setzler (2022), Berger, Herkenhoff and Mongey (2022a), and Yeh, Macaluso and Hershbein (2022), which document that workers are paid, on average, between 15% and 35% less than their marginal products. Given the uncertainty surrounding the extent of firms' monopsony power on workers with different education and ability in this evolving literature, we perform an extensive analysis of the sensitivity of our results to the degree of firms' monoposony power.

We use our quantitative framework to study the distributional impact of a large increase in the minimum wage over time. In the short run, the minimum wage works as its proponents claim: it increases the wages of the workers at the lowest percentiles of the wage distribution with only small negative effects on their employment. Hence, it considerably boosts their labor income and consumption. The main force behind this result is our putty-clay technology, which leads to frictions to the adjustment of the mix of any inputs to production.

In the long run, however, a large minimum wage has perverse distributional consequences in that it lowers the employment, income, and welfare precisely of the low-wage workers it is designed to help, just as its skeptics fear. As a result, nearly one-fifth of all non-college workers—and 40% of non-college workers initially earning less than $15—experience a decline in labor income in the long-run. These earnings losses are concentrated among workers in the bottom 20% of the initial wage distribution, namely, workers initially earning less than about $10 per hour.

The economic mechanisms behind these long-run consequences are twofold. First, for low-wage workers, their "efficient" wage—the wage they would earn absent firm monopsony power—is much lower than $15 since their marginal productivity is fairly low. More generally, the heterogeneity in workers' productivity within education groups, as implied by the observed wage distributions in Figure 1, leads to efficient wages that are highly dispersed. Hence, a single minimum wage is a blunt tool in that it cannot eliminate monopsony distortions for all workers at once. In particular, choosing a high enough minimum to eliminate the monopsony distortion for the *average* worker requires setting it much above the efficient level of wages for *low-productivity* workers, thereby reducing their employment. Second, given the high long-run elasticity of substitution between non-college workers of differing abilities estimated by Card and Lemieux (2001), firms will eventually sharply substitute away from lower productivity non-college workers towards higher productivity ones in response to a $15 minimum wage. But firms with putty-clay technology find it optimal to adjust their workforce only slowly over time. Indeed, even four years after the higher minimum wage is introduced, only about one-quarter of the total long-run employment adjustment has taken place.

The overall impact of the minimum wage on worker welfare results from the combination of these short-run and long-run effects over the entire transition path to the new equilibrium with the higher minimum wage. In our baseline calibration, we find that all but the lowest 16% of current non-college workers

initially earning less than $15 benefit from the imposition of a $15 minimum wage. That is, for most low-wage workers, the benefits of higher wages and slow employment adjustments in the short run outweigh the larger employment losses in the long run. However, for new workers who enter the labor market ten years after the higher minimum wage is implemented, the benefits of a higher minimum have essentially dissipated, as firms have already started their adjustment away from workers of lower productivity. This comparison highlights how the benefits to low-wage workers from such a large minimum wage increase are concentrated over the first decade after its implementation and, as a result, may have more adverse impacts on future cohorts.

The difference between short- and long-run input substitutability implied by our putty-clay technology also allows our model to match many of the short-run estimates of the employment impact of the minimum wage (see, for example, the recent survey by Neumark and Shirley (2021)). We demonstrate this point by replicating in our model the effect of small, local changes in the minimum wage of the type usually studied in the empirical literature and showing that our model's response to such changes is in line with the employment elasticities typically estimated. Since only a small part of the eventual adjustment of employment occurs during the first few years after a minimum wage increase, this finding also implies that empirical estimates of the effects of the minimum wage over the first few years after an increase comes into effect do not help detect its ultimate long-run consequences.

Although most of our analysis focuses on a permanent increase in the real value of the minimum wage, we also examine the case in which the minimum wage steadily erodes in real terms over time. We interpret this case as capturing in a stylized way what happens between legislated changes of a nominal minimum wage that gradually becomes less binding due to a combination of inflation and growth in labor productivity. Interestingly, such a temporary minimum wage increase can improve the labor income of the lowest-productivity workers in every period in which it binds. This occurs because firms do not find it optimal to vary much their input mix in response to temporary changes in the minimum wage.

Given the long-run costs associated with the minimum wage, we conclude by showing that existing tools in the U.S. tax and transfer system, such as the EITC, are much more effective at raising the employment, income, and welfare of low-wage workers in the long run. Intuitively, these policies directly target low-income workers and are more successful at offsetting the monopsony distortions for these workers because they are easier to target and fine tune. Specifically, we first show how the EITC leads to higher welfare for low-wage workers relative to the minimum wage in a budget-neutral way. We then show that for low-wage workers, the existing EITC policy coupled with a modest increase in the minimum wage to about $9 is more beneficial than the EITC policy alone. The reason is that by offsetting the monopsony distortions for such workers, the minimum wage complements the redistributive effects of an EITC policy. We also find that combining increases in the overall progressivity of the U.S. tax system with small increases in the minimum wage leads to similar results.

Related Literature. Our paper is motivated in part by the empirical work on the short- and medium-run responses to minimum wage increases, which we discuss in greater detail in Appendix B. In terms of the short-run responses, Neumark and Shirley (2021) document that roughly 80% of the 109 published studies that they review find zero to small short-run negative employment effects in the two years after a minimum wage increase. In contrast to the vast amount of work on the short-run impact of the minimum wage, there has been almost none on its medium- and long-run impacts.[6] There are, however, a few recent exceptions. For example, Clemens and Strain (2021) document employment responses to the minimum wage over short- (up to 4 years after a minimum wage change) and medium-run (from 4 to 6 years after a minimum wage change) horizons. They find no significant effects in the short run but larger significant effects in the medium run for more sizable increases in the minimum wage. We discuss later in the paper how our model is consistent with their findings. Similarly, Meer and West (2016) estimate that an increase in the minimum wage reduces employment but such an effect takes several years to materialize. In response to a large and persistent minimum wage increase in Hungary, Lindner and Harasztosi (2019) estimate that employment elasticities are negative but small even four years after such a reform is introduced and that firms responded to it by substituting away from labor towards capital. Likewise, Clemens, Kahn and Meer (2021) exploit cross-regional variation to document that firms substitute away from low- productivity workers towards higher-productivity ones in response to minimum wage increases. These recent papers highlight that in response to minimum wage increases, firms adjust their input mix away from workers bound by the minimum wage and that such adjustments take time. These findings thus suggest that any study of the distributional effects of the minimum wage must allow for sluggish input substitutability on the part of firms. Our paper fills this gap in the literature.

Our framework builds on a long line of work that has studied the minimum wage through the lens of frictional models of the labor market. See, in particular, Eckstein and Wolpin (1990), Flinn (2006), Ahn, Arcidiacono and Wessels (2011), and, recently, Engbom and Moser (2021). We add to this literature by augmenting the search framework to allow for monopsonistic competition among firms, so as to capture an important margin along which the minimum wage can be beneficial in the long run, and by estimating the distributional effects of the minimum wage in both the short and long run. Most importantly, we provide a mechanism that reconciles the small short-run elasticities of substitution among inputs documented in the minimum wage literature with the large long-run ones typically estimated, which helps realistically discipline the distributional effects of the minimum wage at these very different time horizons.

In related work, Sorkin (2015) and Aaronson et al. (2018) use a variant of the standard putty-clay capital setup to argue that the effect of the minimum wage on employment is smaller in the short run than in the long run. However, these papers do not study the distributional impact of the minimum wage across workers, and therefore abstract from the rich worker heterogeneity that is the focus of our paper.

[6] The lack of long-run empirical evidence has been pointed out in Brown (1999)'s review article where he states that "[t]here is a simply a stunning absence of credible evidence—indeed, of credible attempts—[to identify] the long run effects [of the minimum wage]."

Our focus on this heterogeneity leads us to discipline our model using well-known estimates in the labor economics literature of the long-run elasticities of substitution among workers (see, for instance, Katz and Murphy (1992) and Card and Lemieux (2001)). The values of these elasticities turn out to be crucial for our quantitative results. In addition, since we are interested in large changes in the minimum wage that impact a significant fraction of the economy, we build a general equilibrium model parameterized so as to be representative of the entire U.S. economy. In contrast, Sorkin (2015) and Aaronson et al. (2018) consider small changes in the minimum wage, so they study an industry equilibrium model with an application to the restaurant industry. Within this industry, they assume that firms cannot adjust their inputs, so firms' decisions reduce to whether or not to shut down. In our model, firms continually decide which types of workers to hire and which mix of workers and capital to use, allowing us to shed light on how the minimum wage affects the intensive margins of firm's employment and capital investment decisions. Overall, we view our paper as building on the same basic insight as Sorkin (2015) and Aaronson et al. (2018), but scaling it up to study the distributional impact of large changes in the minimum wage allowing for input substitutability on the part of firms that affect the entire economy.

Our model of monopsonistic competition is the natural labor market analogue of the model of monopolistic competition in the goods market, adapted to a search setting. In our model, firms' monopsony power arises from an imperfect substitutability of jobs across firms in workers' preferences, as in Berger, Herkenhoff and Mongey (2022a). In contemporaneous work, Berger, Herkenhoff and Mongey (2022b) adapt this setup, which ports the Atkeson and Burstein (2008)'s model of Cournot competition in the goods market to a labor market setting, and pursue a normative analysis of the long-run optimal level of the minimum wage. Our paper differs from Berger, Herkenhoff and Mongey (2022b) in several ways. Most importantly, we focus on the distributional impact of the minimum wage across workers in an environment with rich worker heterogeneity, whereas Berger, Herkenhoff and Mongey (2022b) focus on the differential effect of the minimum wage across firms in an environment with rich firm heterogeneity. Given the different short-run and long-run elasticity of substitution across workers documented by a vast literature as discussed above, we heavily focus on the dynamics of the effect of the minimum wage over time. By the same token, our model abstracts from the rich oligopsonistic market structure studied in Berger, Herkenhoff and Mongey (2022b). Therefore, we view our two papers as complementary.[7]

2 Model

We begin by briefly highlighting the main features of our model. 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. When this is the case, large minimum wage increases can actually increase employment and labor market participation by potentially reducing the resulting monopsony distortions

[7] Given our focus on the redistributive effects of the minimum wage across workers with different observed and unobserved characteristics, we also study existing transfer programs like the EITC and how they interact with the minimum wage.

in the labor market. We embed this setup in an economy with labor markets subject to search frictions to be able to distinguish among the employment, unemployment, and non-participation effects of minimum wage policies. In particular, by incorporating a non-participation margin, we capture the idea that a higher minimum wage may incentivize non-participants to enter the labor market and search for jobs. Importantly, in such a framework, monopsony power not only distorts the participation margin of workers who look for jobs, but also an often neglected margin, namely, the vacancy-posting margin of firms deciding to hire more workers. We will show that the effect of tax policies such as the EITC subtly depend on how they affect these two margins.

Second, we include in the model rich worker heterogeneity, both within and across education groups, so as to explore the distributional impact of the minimum wage. We allow for this heterogeneity within a production technology consistent with the long-run evidence on the elasticity of substitution across multiple inputs. For example, we ensure that our model matches the long-run elasticity of substitution between college and non-college workers estimated by Katz and Murphy (1992). Likewise, we make our model consistent with the estimated elasticity of substitution across workers with different skills *within* an education group, documented in work such as Card and Lemieux (2001). We also let firms substitute between labor and capital.

Finally, to guarantee that our technology is consistent with the evidence of low input substitutability in the short run, we assume that firms operate technologies of the putty-clay type. Specifically, we capture in an intuitive way the notion that adjusting the ratio of any inputs is costly for firms in the short run by assuming that the production technology is embedded in the capital stock so its labor intensity is irreversible once capital is installed. The idea is that a new piece of capital can be built to be used in combination with low- and high-educated workers of any ability in any ratio. These ratios, however, are fixed after the new capital is installed. We view this putty-clay structure as a tractable way to allow for a complex set of adjustment costs for firms when they decide to alter their input mix over time.

In our quantitative exercises, 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 a higher minimum wage. Putty-clay capital thus helps the model 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, as consistent with a large literature on long-run input substitution patterns. We next discuss our model in greater detail.

2.1 Preferences, Production and Matching

We consider an infinite-horizon economy in discrete time populated by consumers and firms. We describe next consumers, firms, production, output and labor markets, and the timing of events in a period.

Consumer Heterogeneity and Preferences. Consumers are heterogeneous in two dimensions. First, they differ in their *education level* $g \in \{\ell, h\}$, where ℓ denotes the group of low-educated consumers (those

with less than a bachelor's degree) and h denotes the group of 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 education-specific discrete distributions. Ability differences among consumers allow us to match the observed wage distribution within each education group. We index a consumer by i, which denotes both 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.[8] Each type of family is composed of a large number of household members of the same education group and ability level. Risk sharing within such families implies that each member of a household of type i consumes the same amount of goods at date t, regardless of the idiosyncratic shocks that such a member experiences.[9]

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

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

where n_{ijt} is the number (measure) 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 at them and can be interpreted as arising from workers' idiosyncratic preferences over different firms, locations, or amenities. The smaller ω is, the less substitutable jobs at a same firm are. Here we adapt the standard way of modeling imperfect substitutability in consumers' preferences across differentiated goods to modeling imperfect substitutability in workers' preferences across differentiated jobs. This imperfect substitutability in preferences for jobs will generate an upward-sloping labor supply curve for each firm's jobs that is analogous to the downward-sloping demand curve for each firm's goods, which arises in standard models of monopolistic competition.[10] As we discuss below, ω is a key parameter that governs the extent of firms' monopsony power in the labor market. In our quantitative analysis, we discipline ω using estimates of the extent to which wages are marked down relative to workers' marginal products.

Production Technology. 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

[8] Letting J denote the (integer) number of firms in the market and assuming that for each J, there is a total measure $\mu_i J$ of consumers of type i, here we focus on an economy with J large enough so that it is well-approximated by $J = \infty$.

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

[10] See Berger, Herkenhoff and Mongey (2022a) and Deb, Eeckhout and Warren (2021) for related preferences along with a discussion of various interpretations and alternative microfoundations of them. Intuitively, this specification can be primitively derived from idiosyncratic shocks to the value of working at any firm due to firm's different locations, amenities, and similar.

labor $\bar{n}_{\ell jt}$, and an aggregate of efficiency units of high-educated labor $\bar{n}_{hjt}$. The law of motion for capital accumulation is $k_{jt+1} = (1-\delta)k_{jt} + x_{jt}$, where δ is the 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 assume a nested CES production function over capital k_{jt}, the aggregate of low-educated labor $\bar{n}_{\ell jt}$ and the aggregate of high-educated labor $\bar{n}_{hjt}$ of the form

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

where

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

this production function will underlie both the version of our model with standard capital and that with putty-clay capital. The outer nest in $G(\cdot)$ is a CES production function over capital k_{jt} and an aggregate $G(\bar{n}_{\ell jt}, \bar{n}_{hjt})$ of the efficiency units of low- and high-educated labor, $\bar{n}_{\ell jt}$ and $\bar{n}_{hjt}$. The inner nest in $G(.)$ is a CES production function over $\bar{n}_{\ell jt}$ and $\bar{n}_{hjt}$. The parameters ρ and α are key as they capture the degree of substitutability among inputs: the larger ρ is, the more substitutable capital and the labor aggregate $G(\bar{n}_{\ell jt}, \bar{n}_{hjt})$ are, whereas the larger α is, the more substitutable low-educated and high-educated labor are. The parameter α will govern how much firms substitute away from low-educated labor towards high-educated labor when low-educated labor becomes more expensive from the imposition of a large minimum wage. In our quantitative exercise, we make this parameter consistent with the estimates of the elasticity of substitution between high- and low-educated workers from Katz and Murphy (1992). The parameter ρ, instead, will govern how much firms will substitute from the labor aggregate as a whole towards capital as the minimum wage increases the overall cost of labor. We will show the sensitivity of our results to using different estimates of capital-labor subsitutability from the literature.

The labor inputs $\bar{n}_{\ell jt}$ and $\bar{n}_{hjt}$ used by firm j are themselves CES aggregates of the labor inputs of workers of different abilities within each education group,

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

where n_{ijt} is the amount of low-educated labor of a 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. The parameter ϕ will play an important role in our assessment of the distributional effects of the minimum wage, because it governs the extent to which firms are willing to substitute *across* workers of differing ability, as indexed by z, within an education group. For instance, if some low-ability workers become more expensive due to a higher minimum wage, firms can substitute towards higher-ability workers in production. Our nested CES production structure for labor in (3) and (4) exactly follows the framework in Card and Lemieux (2001). Our setup differs from theirs only because we allow for capital in the overall production function in (2). To discipline ϕ, we rely on estimates from Card and Lemieux (2001), which finds a relatively high

elasticity of substitution between workers of differing ability within an education group.

Matching Technology. 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 searching for jobs at firm j, where a_{ijt} denotes the measure of vacancies posted by firm j aimed at each family of type i and s_{ijt} denotes the number of searchers from 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 of type i 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's jobs, 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 that a posted vacancy is filled or the *job-filling rate* and $\theta_{ijt} = a_{jit}/s_{ijt}$ 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 at firm j or the *job-finding rate*. This directed search framework ensures that labor market allocations are pinned down even when the minimum wage binds for some workers. It also provides a tractable way to distinguish among the responses of unemployment and labor force participation to changes in the minimum wage.

Timing. The timing of events within a period is as follows. Each period t consists of two stages. In stage 1, each firm j posts vacancies $\{a_{ijt}\}$ aimed at consumers of type i that determines the tightness $\{\theta_{ijt}\}$ of the markets for such consumers, and commits to a present-value of wages $\{W_{ijt}\}$ for each consumer of type i who is hired in t and begins to work in $t+1$. Each family chooses the total number of its members $\{s_{it}\}$ searching for jobs. In stage 2, after having observed all firms' offers, each family allocates its searching members among the j firms $\{s_{ijt}\}$, 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 $\{\theta_{ijt}, W_{ijt}\}$. At the end of period t, a proportion σ of matches exogenously terminate. These two stages should be thought of as occurring at the beginning of each period t.[11]

2.2 A Family's Problem

We describe here a family's problem, the implied optimality conditions, and how consumers' preferences over jobs give rise to firms' monopsony power.

[11] Note that we have separated across the two stages a household's decision about the total measure of its members searching for jobs in period t, s_{it}, and about the allocation of this measure across the existing firms, $\{s_{ijt}\}_j$. What matters is that when a firm j decides on the offer $\{\theta_{ijt}, W_{ijt}\}$ for any worker of type i in period t, it takes as given the measure of workers of type i in the market searching for jobs and so the tightness of the corresponding market, θ_{it}.

2.2.1 Setup and Optimality

Consumers of each 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. Thus, our economy is a deterministic one in terms of aggregates. Accordingly, the date-0 budget constraint of family i is

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

where $Q_{0,t}$ denotes the price of the homogeneous good in period t in units of that good in period 0; W_{ijt} is the present value of wages of newly employed workers; Π_0 is the present value of all firms' profits; and ψ_i is the share of profits of the firms owned by the family. To better understand the first term on the right side of (6), note that if s_{ijt-1} consumers of 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, start working in period t, and earn the present value of wages W_{ijt} in units of period-t goods. Since a consumer of type i employed at firm j in t separates from it in $t+1$ with probability σ, the transition law for consumers of type i employed at firm j in t 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 type-i consumers.

In period 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 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 $s_{ijt} \geq 0$, with $s_{it} = \sum_j s_{ijt}$ and n_{it} satisfying (1) for all t, in order to maximize the present value of its utility. Dropping the subscript i for simplicity and letting ζ, $\beta^{t+1}\nu_{jt+1}$, 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

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

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

$$-\frac{u_{st}}{u_{ct}} = \lambda_w(\theta_{jt}) \frac{\beta u_{ct+1}}{u_{ct}} \frac{\nu_{jt+1}}{u_{ct+1}} + \frac{\beta u_{ct+1}}{u_{ct}} \lambda_w(\theta_{jt}) W_{jt+1} + \frac{\chi_{jt}}{u_{ct}}, \tag{10}$$

where we have used $\beta^t u_{ct} = \zeta Q_{0,t}$ to derive (8), which is the standard Euler equation for consumption. To understand the next two equations, recall that the family of an employed consumer is paid the present value of wages W_{jt+1} for a match with firm j, which lasts until an exogenous separation occurs. In (9), ν_{jt+1} is the discounted marginal disutility from a marginal increase in the number of the family's members

who work at firm j in $t+1$, of whom $(1-\sigma)$, $(1-\sigma)^2$, and so on are still employed in $t+2$, $t+3$, and subsequent periods. To express this disutility in consumption units, we define $V_{jt+1} \equiv \nu_{jt+1}/u_{ct+1}$ and substitute $Q_{t+1,t+2} = \beta u_{ct+2}/u_{ct+1}$ into (9) to express V_{jt+1} recursively as

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

Further substituting $\nu_{jt+1}/u_{ct+1} = V_{jt+1}$ and $Q_{t,t+1} = \beta u_{ct+1}/u_{ct}$ into (10) gives

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

To understand this condition, note that a marginal increase in the number of consumers who search for jobs 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_{jt+1} in units of period- $t+1$ consumption goods, net of the present value of the disutility of work V_{jt+1}. Expressed in period-t consumption units, this expected net benefit is $Q_{t,t+1}\lambda_w(\theta_{jt})(W_{jt+1} + V_{jt+1})$, 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 at least as high as the value of searching for jobs at any other firm j' so that

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

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

2.2.2 The Participation Constraint and Firm Monopsony Power

We now examine how the constraint (13) simplifies in our symmetric equilibrium and how it encodes firms' monopsony power. In such an 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 $\lambda_w(\theta_{it})\,(W_{it+1} + V_{it+1})$ to type-i consumers and suppose that firm j contemplates offering a potentially different value, $\lambda_w(\theta_{ijt})(W_{ijt+1} + V_{ijt+1})$. Then, for firm j to attract a consumer, it must offer at least that common value. That is, firm j's offer must satisfy

$$\mathcal{W}_t(\theta_{ijt}, W_{ijt}) \equiv \lambda_w(\theta_{ijt})(W_{ijt+1} + V_{ijt+1}) \geq \mathcal{W}_t = \mathcal{W}_t(\theta_{it}, W_{it+1}) = \lambda_w(\theta_{it})\,(W_{it+1} + V_{it+1})\,. \tag{14}$$

We refer to this constraint as the *participation constraint* and note that is the symmetric version of (13). The second stage of an equilibrium is summarized by this constraint (14). Solving forward the recursive

expression in (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 $t+1$ yields

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

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

Although the supply curve of workers for a firm in period t is a dynamic object that depends on wages and market tightness in t as well as the expectations of these variables in all future periods, we can provide some intuition about it in steady state, assuming that preferences $u(c_i, n_i, s_i)$ are of the form $u(c_i - v(n_i) - h(s_i))$ as in Greenwood, Hercowitz and Huffman (1988), which we will use in our quantitative analysis. In this case, the participation constraint in steady state reduces to

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

Holding fixed θ_{ij} and differentiating this constraint with respect to w_{ij} and n_{ij}, we obtain 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) labor supply curve for firm j slopes upward in n_{ij}. But as ω becomes arbitrarily large, the slope of this curve converges to zero. This upward-sloping labor supply curve facing firms in our framework is akin to the simple static upward-sloping labor supply curve studied in Robinson (1933).

Firms' monopsony power, as captured by this labor supply curve, affects a firm's first-order conditions for offered wages and vacancies through the derivatives of workers' participation constraint, which capture how the present value of the disutility of work increases when any firm j changes its wages and vacancies, and so market tightness, to attract more workers, holding fixed the value of search in the common market.[12] Put differently, this upward sloping supply curve is the source of firms' monopsony power. Intuitively, firms realize that to attract additional workers of any type i, they need to compensate workers of the same type already employed for their increased disutility of work. As a result, they end up hiring fewer workers and paying lower wages in equilibrium. As we show below, an additional distortion arises in our framework, since firms' monopsony power also depresses job creation. In later sections, we discuss the extent to which the alternative policies we consider alleviate or compound all these monopsony distortions.

2.3 A Firm's Problem

We consider two versions of a firm's problem 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

[12] In Appendix A, we discuss how this upward-sloping labor supply curve for a firm's jobs that our model gives rise to is analogous to the downward-sloping demand curve for a firm's goods that arises in models of monopolistic competition.

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 a machine is built, it is possible to alter it to use it 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 at odds with the data. Motivated by this issue, we consider a second version of a firm's problem with *putty-clay* capital such that after a machine is built, its labor intensity is irreversible. This feature will imply that in this version, employment in the short run reacts much less to increases in the minimum wage than in the putty-putty version and, hence, its response is more in line with the data.

2.3.1 A Firm's Problem with Standard Capital

Consider a 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, 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 $\{W_{ijt}\}$ for these consumers, and new capital $\{k_{t+1}\}$ in order to maximize

$$\sum_{t=0}^{\infty} Q_{0,t} \left\{ F\left(k_{jt}, \bar{n}_{\ell jt}, \bar{n}_{hjt}\right) - q_t x_{jt} - \sum_i W_{ijt} \lambda_f(\theta_{ijt-1}) \mu_i a_{ijt-1} \right\}, \tag{17}$$

subject to the law of motion for capital $k_{jt+1} = (1-\delta)k_{jt} + x_{jt}$, 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

$$\lambda_w(\theta_{ijt})(W_{ijt+1} + V_{ijt+1}) \geq \mathcal{W}_t \equiv \max_{j'} \{\lambda_w(\theta_{ij't})(W_{ij't+1} + V_{ij't+1})\} \tag{19}$$

in each period. 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, searchers, and capital $\{c_{it}, n_{it}, s_{it}, \bar{n}_{it}, k_t\}$, vacancies and market tightnesses $\{a_{it}, \theta_{it}\}$, and prices $\{W_{it+1}, Q_{0t}\}$ 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.

Since families can perfectly insure the idiosyncratic risk of their members 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_{ijt} by offering a constant period wage w_{ijt} over the course of a match that begins at t so that

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

where $Q_{t,s}$ is the price of goods in $s > t$ in units of goods in t. Since $W_{ijt} = d_t w_{ijt}$ where $d_t \equiv$

$\left[1 + (1-\sigma)Q_{t,t+1} + (1-\sigma)^2 Q_{t,t+2} + \ldots\right]$, we can equivalently think of firms as choosing W_{ijt} or w_{ijt}.

2.3.2 A Firm's Problem with Putty-Clay Capital

Suppose now that 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. Once a machine is built, however, it is clay-like in that it requires a fixed amount of low-educated labor and high-educated labor to operate at full capacity, that is, the output technology is Leontief ex-post. Hence, given a stock of machines, demand for low-educated labor and high-educated labor is inelastic in the short run as long as total profits from operating the machines 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. Thus, in the long run, firms can substitute away from the type of labor that becomes more expensive, for instance, low-educated labor when the minimum wage increases, 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 of type i, a capital type $v = \{v_i\}$ denotes the education and ability intensity of capital, that is, how much labor of each education and ability capital needs in order to produce a certain amount of output. 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 are combined with $n_i = k(v)v_i$ units of labor for all i, then $f(v)$ units of output are produced according to the constant-returns-to-scale production function $F(\cdot)$ with $f(v)$ defined by

$$F(k, \{n_i\}) = kF(1, \{n_i/k\}) = kF(1, \{v_i\}) = kf(v).$$

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 $y(v) = \min\left[k(v), \{n_i(v)/v_i\}\right] f(v)$, were $\min\left[k(v), \{n_i(v)/v_i\}\right]/k(v)$ can be thought of as the utilization rate of the $k(v)$ units of capital of type v. The total output of a firm in period t is thus

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

Firms invest $x_t(v)$ units of output to accumulate type-v capital according to the 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, a firm chooses sequences of market tightnesses $\{\theta_{it}\}$, vacancies $\{a_{it}\}$, employed workers $\{n_{it+1}(v)\}$ for each type of capital v, present value of wages $\{W_{it+1}\}$, and investment $\{x_t(v)\}$ for each type of capital in order to maximize

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

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 $n_{it} \leq \int n_{it}(v) dv$ for the uses of labor of each type i, 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 $y_t(v) = \min[k_t(v), \{n_{it}(v)/v_i\}] f(v)$, if, say, the firm uses $n_{it}(v)$ units of type-i labor such that $n_{it}(v) > k_t(v) v_i$, then the excess labor $n_{it}(v) - v_i k_t(v)$ is wasted, so this is never optimal. Hence, we can impose the constraints in (23) and drop the Leontief function from the firm's problem. The non-negativity constraint $x_t(v) \geq 0$ implies that firms cannot disassemble their existing types of capital. Without this friction, the firm's problem would reduce to the putty-putty problem previously described.

2.4 Steady-State Properties

Here we discuss the steady state of the economy and discuss some key features of 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 the intertemporal prices of consumption that firms face become constant, firms invest in the unique type of capital that is ideally suited to their technologies at those prices and let all past capital depreciate. Eventually, all the old capital stock is replaced. In a steady state with putty-clay capital, firms invest in exactly the same type of capital as they do in a steady state with standard capital. In our baseline, we focus on Greenwood, Hercowitz and Huffman (1988) (GHH) preferences,

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

for each family i. Consider the steady state of this version of the model in which all variables, including the price of capital, are constants. 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

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

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 in (6). With GHH preferences, which imply no income effects, the steady state equations then split into two blocks. First, we can solve for the monopsonistically competitive search equilibrium wages $\{w_i\}$ and the associated allocations $\{\theta_i, s_i, n_i\}$ and k from (25)–(29). Then, given these allocations and prices, 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 $v'(n)/\omega$: for a given marginal product of labor and marginal disutility of work, a firm offers a smaller wage than under the competitive search equilibrium for our economy, in which case $v'(n)/\omega = 0$ since $\omega = \infty$. As apparent from the first-order condition in (27), this inefficiently low level of wages results in consumers searching too little for jobs. Firms' vacancy-posting condition (26) features both the indirect distortion from the inefficient level of wages and the direct distortion due to the term $v'(n)/\omega = 0$. Despite these two distortions have countervailing effects on the marginal benefits of posting vacancies, it turns out that in equilibrium firms post too few vacancies. Hence, firms create too few jobs, consumers search too little for them, and both wages and employment are lower than in the competitive search benchmark—the efficient case for our economy. Since a firm with monopsony power pays its workers only a fraction of their marginal products, a simple measure of firms' monopsony power is then the markdown of wages relative to workers' marginal products, $1 - w_i/F_{ni}$, or, equivalently, the percentage difference between a worker's marginal product and wage, $(F_{ni} - w_i)/F_{ni}$. In a slight abuse of language, we refer to w_i/F_{ni} as the *wage markdown*. Letting $\kappa(\theta_i) = (r+\sigma)\kappa_i/\lambda_f(\theta_i)$, we can combine the equilibrium vacancy-posting condition and the wage equation to show that the implied markdown for workers of a family of type i is

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

In (30), the *efficient component* of the markdown is defined as the markdown that would arise in the absence of firms' monopsony power. This efficient component corresponds to the level of the wage markdown needed for firms to recoup their vacancy-posting costs and hence earn zero expected profits per vacancy. More interesting is the *monopsony component*, which arises because firms with monopsony power set wages below their competitive search level and so their markdowns are larger than the competitive search ones. As this equation makes clear, for any given size of measured markdowns, the larger the efficient component is, the smaller the monopsony component is. In our quantitative exercises, we find that the overwhelming majority of the markdown in wages is due to the monopsony distortion.

A simple policy implication of this analysis is as follows. Since all distortions emanate from firms paying too low a wage to each consumer type, a policy that mandates that firms must pay a *type-specific* minimum wage $\underline{w}_i$ for each consumer type i, set equal to the wage for that worker type in the competitive search equilibrium, would fix all distortions and lead to efficient allocations. More precisely, define the competitive search equilibrium wages $\{w_i^*\}_i$, associated allocations $\{\theta_i^*, s_i^*, n_i^*, c_i^*\}$, and k^* that satisfy the competitive search equilibrium 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. As we consider a sequence of economies in which firms' monopsony power converges to zero in that ω diverges to infinity, the monopsonistically competitive search equilibrium wages converge to the competitive search equilibrium wages and allocations converge to those of the competitive search equilibrium. We summarize this discussion in the following result.

***Proposition* 1.** *If a minimum wage for each worker type is set equal to the competitive search equilibrium 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 of the competitive search equilibrium. Also, as ω diverges to infinity, wages and allocations in the monopsonistically competitive search economy converge to those of the competitive search equilibrium.*

For the second part of the proposition, note that as ω becomes arbitrarily large, the monopsonistically competitive search equilibrium wage in (28) converges to the competitive search equilibrium wage in (31). Hence, the distortions to wages vanish and, from an inspection of (26) and (27), so do the distortions to the vacancy-posting condition and the first-order condition for search.

The issue we will address in later sections 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. We will then analyze the other extreme, which corresponds to the minimum wage policies advocated in practice consisting of only *one* mandated minimum wage for all workers. In such a scenario, if large enough differences in skill and ability exist across workers, then a single minimum wage can have perverse distributional effects.

3 Quantification

We choose the parameters in our model to match key features of the U.S. labor market, which inform the mechanisms described above. We assume that a model period is one month in order to adequately capture worker flows in the labor market. We maintain that the utility function has the Greenwood, Hercowitz and Huffman (1988) form, namely,

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

where $\chi_{g,n}$ governs the disutility of work for education group $g \in \{\ell, h\}$, $\chi_{g,s}$ governs the disutility of search for each such group, and γ_n and γ_s help control the elasticity of labor supply and job search.[13]

3.1 Disciplining Key Features of the Model

We start by summarizing how we parameterize the key features of our model: the degree of firms' monopsony power, the distribution of worker heterogeneity, the elasticities of substitution across workers and with capital, and the putty-clay technology as summarized by the depreciation rate. Table 1 reports this set of parameters. The degree of monopsony power is crucial because it determines the potential for the minimum wage to increase employment and labor market participation. In our model, the degree of monopsony power is controlled by the substitutability across jobs in workers' preferences, ω. We discipline this parameter by targeting existing empirical estimates of the average wage markdown in the data. This target is informative because, as equation (30) shows, ω determines the size of the monopsony component of the markdown. As ω becomes large, firms' monopsony power falls to zero and the inefficient component of the wage markdowns vanishes. A growing literature has measured wage markdowns in the United States, providing estimates that, on average, workers are paid between 0.65 and 0.85 of their marginal product.[14] As a baseline, we target the value 0.75, which is midpoint of the estimated range. In Section 4.5, we show the sensitivity of our results to either higher or lower estimates of wage markdowns.

The distribution of workers' labor market productivity within each education group, captured by $\{\mu_i\}$, governs the degree of wage dispersion for each group and thus the distributional impact of the minimum wage across workers. For both education groups $g \in \{\ell, h\}$, we assume this distribution is lognormal with group-specific standard deviation σ_g, and choose the parameters σ_g to match the dispersion of the distributions of wages in Figure 1.[15] Specifically, we target the ratio of the 50th percentile to the 10th

[13] In principle, the preferences (32) may imply that households violate their time constraint in that $n_{it} + s_{it} > 1$. We ensure that the time constraint is always satisfied by augmenting (32) with a positive utility from leisure $\chi_\ell \log(1 - n_{it} - s_{it})$ and setting χ_ℓ to 0.01. Standard Inada conditions imply that households will not violate their time constraint. Otherwise, this term has minimal effects on our results. We view this specification as simply a technical device to ensure that the time constraint holds in each period without having to deal with an occasionally binding time constraint.

[14] See, for example, Berger, Herkenhoff and Mongey (2022a), Yeh, Macaluso and Hershbein (2022), Lamadon, Mogstad and Setzler (2022), and Seegmiller (2021). Manning (2021) provides a recent survey of the this literature.

[15] We restrict the sample to include all individuals aged 16 and above, exclude all individuals residing in group quarters and those who report themselves as being a student, which mirrors the sample restrictions used by the BLS to compute labor

TABLE 1: Parameters Governing Key Features of Model

Parameter	Description	Value	Discipline
Monopsony power			
ω	Substitutability across firms	2.74	Match literature's estimates of wage markdowns
Worker heterogeneity			
μ_i	Distribution of productivities z_i	lognormal	Match wage distribution from ACS
Long-run elasticities of substitution			
ρ	Elasticity of substitution between high-educated and low-educated workers	1.40	Fixed: Estimated value from Katz and Murphy (1992)
ϕ	Elasticity of substitution across workers within an education group	4.00	Fixed: Lower bound estimate from Card and Lemieux (2001)
α	Elasticity of substitution between capital and workers	1.00	Fixed: Cobb-Douglas
Putty-clay frictions			
δ	Depreciation rate	10% annual	Fixed: BEA data

Note: Baseline parameters governing key features of the model: monopsony power, worker heterogeneity, elasticities of substitution, and putty-clay frictions.

percentile of the wage distribution of each education group in order to precisely match the left tail of each distribution, which is most directly affected by a higher minimum wage.[16]

The production function parameters ρ, α, and ϕ govern the long-run substitutability between different types of workers and between workers and capital. The parameter ϕ in (4), which determines the elasticity of substitution across workers within an education group, controls the extent to which firms are willing to substitute away from low-wage workers *within* an education group in response to an increase in the minimum wage. Card and Lemieux (2001) estimate elasticities of substitution across workers of different ages, which we interpret as proxies of different levels of ability, in a range from 4 to 6. As a benchmark, we set $\phi = 4$ to match the lower end of this range and show in Section 4.5 that raising ϕ to 6 only amplifies the negative effects of the minimum wage in the long run for low-ability workers.[17]

We set the parameter ρ in equation (3), which determines the long-run elasticity of substitution between low- and high-education workers, to Katz and Murphy (1992)'s estimate of $\rho = 1.4$.[18] Hence, using the standard estimates from the literature implies that the elasticity of substitution among workers *between* education groups is smaller than the elasticity of substitution among workers *within* an education group, that is, $\rho < \phi$. We show in Section 4.5 that our key distributional results only change slightly when we use a higher value of this elasticity, such as Bils, Kaymak and Wu (2020)'s recent estimate of $\rho = 4$. These results confirm that the within-education-group elasticity of substitution among workers measured

market statistics. See Appendix B for more details.

[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.

[17] Card and Lemieux (2001) find similar estimates of the elasticity of substitution across workers of differing ability within an education group for both high- and low-education workers, so we set the same value of ϕ between education groups.

[18] Card and Lemieux (2001) also find an estimate of the cross-education group elasticity of substitution of about 1.4.

by ϕ is quantitatively more important than the between-education-group elasticity of substitution among workers measured by α. An active debate in the literature concerns the value of α, ranging from a value suggesting greater complementarity than Cobb-Douglas in Oberfield and Raval (2021) ($\alpha = 0.5$ to 0.7) to a value suggesting greater substitutability than Cobb-Douglas in Karabarbounis and Neiman (2014) ($\alpha = 1.25$). As a compromise between these estimates, we set $\alpha = 1$ (Cobb-Douglas) as our baseline, but show in Section 4.5 that our results are not very sensitive to values of α in this range.[19]

Finally, due to our putty-clay technology, the dynamic effects of the minimum wage depend on the speed with which firms adjust their capital stock. Since firms in our model cannot uninstall existing capital, the speed of this adjustment is largely determined by the depreciation rate of capital, δ. We set this rate to imply an annual depreciation rate of 10%, which roughly matches the aggregate depreciation rate for the U.S. economy.

3.2 Detailed Description of the Quantification Procedure

We now describe the details of our quantification procedure for the remaining parameters, which are less important for our main results. We proceed in two steps: we first exogenously fix a subset of additional parameters based on external evidence and then choose the remaining ones in order to match several informative statistics of the data. The parameters governing the search portion of the model—the cost of job posting, the parameters of the matching function, and those governing a household's disutility of search—mainly determine the degree to which a change in employment from an increase in the minimum wage manifests itself as a change in the labor force participation rather than in the unemployment.

TABLE 2: Other Fixed Parameters

Parameter	Description	Value
γ_n	Labor supply elasticity	1.00
π_ℓ	Fraction of non-college households	0.69
β	Discount factor	$(1.04)^{-1/12}$
σ	Job destruction rate	2.8%
κ_0	Vacancy posting costs	2 months of flow worker output
η	Elasticity of matching function w.r.t. vacancies	0.50
γ_s	Search supply elasticity	5.00
χ_s	Scale of search disutility	3.8×10^6

Note: Other parameters (in addition to ρ, α, ϕ, and δ discussed in Table 1) exogenously fixed in the calibration. A model period is one month.

Table 2 shows the parameters that we exogenously fix, in addition to the fixed parameters already discussed in Table 1. We set the parameter γ_n of the utility function, which governs the elasticity of labor supply, to 1, but we show that our results are robust to alternative values in Section 4.5. We fix the share of college-educated households in the population to $1-\pi_\ell = 31\%$ in order to match their proportion in the ACS data. We choose a value of $(1.04)^{-1/12}$ for households' discount factor β so that the annualized real

[19] In Appendix C, we study an alternative version of the model in which capital and non-college labor are substitutes but capital and college labor are complements, as in Krusell et al. (2000). The main results for this alternative specification are very similar to those discussed below.

interest rate r equals 4%. We select the value of the vacancy-posting cost so that it equals roughly two months of average flow worker output. We set the job destruction rate σ to 2.8% and the elasticity of the matching function with respect to the measure of searchers to $\eta = 0.5$. Finally, we note that there exists a locus of values for the parameters γ_s and χ_s governing the disutility of search that imply approximately identical steady-state moments but differ in the response of search effort to an increase in the minimum wage. We choose a pair of values on this locus that imply a relatively muted response of search effort to the minimum wage, as in the data; see the discussion in Section 4 below.

TABLE 3: Fitted Parameters

Parameter	Description	Value
Monopsony power		
ω	Monopsony power	2.74
Worker productivity distribution $\log\mathcal{N}(0,\sigma_b)$		
σ_ℓ	SD of non-college z	0.69
σ_h	SD of college z	0.78
Production function		
ψ	Coefficient on non-college labor n_ℓ	0.42
λ	Coefficient on capital k	0.30
Search frictions		
B	Matching function productivity	0.42
Labor Disutility		
$\chi_{\ell,n}$	Scale of non-college labor disutility	1.82
$\chi_{h,n}$	Scale of college labor disutility	2.43

Note: Parameters endogenously chosen to match the statistics in Table 4.

Table 3 contains the fitted parameters, which we choose to match the statistics in Table 4. We include the monopsony power ω and distribution of productivity z_i in this table, even though we already discussed them in Section 3.1, because the fitted parameters are all jointly determined. As already discussed, the degree of monopsony power ω is primarily determined by the average wage markdown in the data, and the dispersion of worker productivity σ_ℓ and σ_h is primarily determined by the ratio of the 50th percentile to the 10th percentile of the wage distribution of each education group. The scale parameters of the production function, ψ and λ, in (2) and (3) govern the aggregate labor share and the share of total labor income accruing to college-educated workers. The parameter B governs the efficiency of the matching function, which determines the steady state unemployment rate—we target a steady state unemployment rate of 5.9% to be consistent with the data before the Great Recession. Finally, the parameters $\chi_{g,n}$ of the disutility of work of each group $g \in \{\ell, h\}$ control the steady-state employment rates of each group.[20]

Table 4 shows that the model reproduces these targets extremely well. Importantly, the model matches the average wage markdown, and therefore the degree of firms' monopsony power, almost exactly. The unemployment rate, aggregate and college labor shares, the employment shares by education group, and

[20] As with many search models, there exists a set of values for the vacancy-posting costs κ_0 and the matching function productivity B that deliver similar steady-state unemployment rates. In order to resolve this indeterminacy, we exogenously fix κ_0 and endogenously choose B in order to match the average unemployment rate in the data. Our results are very similar if we instead exogenously fix $B = 1$ and endogenously choose κ_0 to match the same data.

TABLE 4: Targeted Statistics

Moment	Description	Data	Model
Average wage markdown			
$\mathbb{E}[w_{ni}]/\mathbb{E}[F_{ni}]$	Average wage markdown	0.75	0.75
Wage Distribution, ACS 2017-2019			
$w_{\ell 50}/w_{\ell 10}$	Non-college 50-10 ratio	2.04	2.00
w_{h50}/w_{h10}	College 50-10 ratio	2.30	2.06
Income shares			
$\mathbb{E}[w_i n_i]/Y$	Aggregate labor share	0.57	0.57
$\pi_h \mathbb{E}[w_{hz} n_{hz}]/\mathbb{E}[w_i n_i]$	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	0.47
$\mathbb{E}_h[n_i]$	College employment rate	0.62	0.61

Note: Statistics targeted using the parameters in Table 3. The average wage markdown is the midpoint of the range of estimated markdowns discussed in the main text. The average labor share is from Karabarbounis and Neiman (2014). The wage distribution targets, college-income share, and employment rates are calculated using the ACS 2017-2019 data described in Appendix B.

the non-college 50-10 wage ratio are almost identical in the model and in the data. The average job-finding rate is 0.44 in the model, which is similar to the value of 0.45 in the data (see Shimer (2005)).

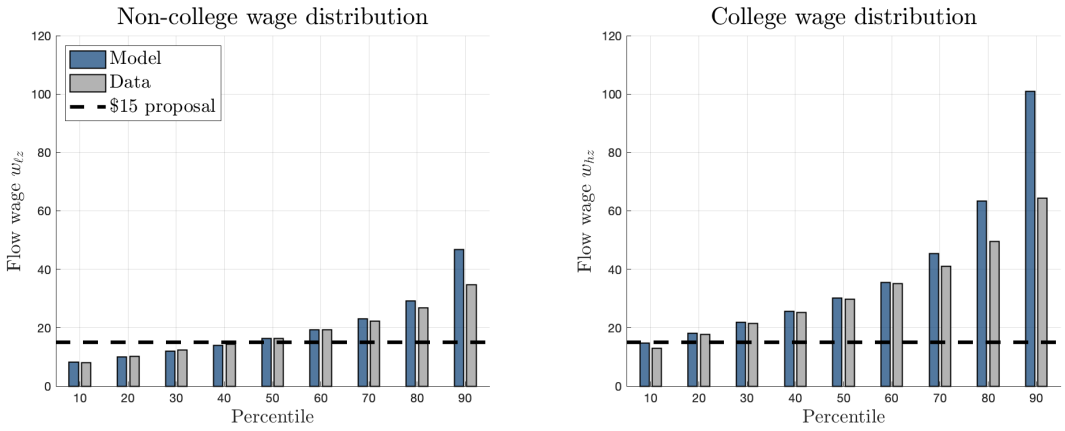
3.3 Implied Wage Distribution

Figure 2 compares the wage distributions in our calibrated model to those from the ACS data. The model matches the left tail of the wage distribution within education groups fairly well by construction, since we target their 50-10 ratio in our calibration. As a point of reference, a $15 minimum wage would bind for approximately 45% of non-college workers in both the model and the data. However, the required dispersion in worker productivity measured by σ_g and the shape of the lognormal distribution imply that the right tail of the two fitted wage distributions is more dispersed in our model than in the data. For example, the interquartile range of non-college wages is 2.1 in the data vs. 2.3 in our model. We are comfortable with this tradeoff between the fit of the model at the low and high end of the wage distributions of the two groups given that the effects of the minimum wage are primarily determined by the left tail of these distributions. Across education groups, the model predicts that the median college wage is 1.82 times larger the median non-college wage, which is in line with a ratio of 1.81 in the data.[21]

Figure 3 plots the model's implied wage markdowns as a function of the steady state wage $w_{\ell z}$ by worker education group and ability. Recall from Section 2 that the steady state markdown is given by the expression in (30). The efficient component of the markdown reflects the fact that firms must recoup the

[21] Curtis et al. (2021) estimates the firm-level response of capital and labor to the Bonus Depreciation Allowance, a temporary tax incentive for investment. They find that capital and labor increase roughly proportionally to this incentive. As a further validation of our model, we replicated the Bonus shock as a change in the after-tax relative price of capital, and found that capital and labor increase roughly proportionally due to our putty-clay technology. The Bonus only changes capital-labor ratios on new investment, which are small relative to the total size of the capital stock. In contrast, the model with standard capital cannot match this finding because firms adjust the capital-labor ratios on their entire stock of capital.

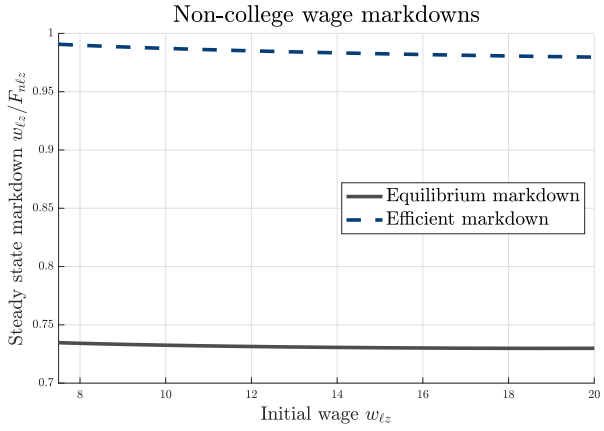
FIGURE 2: Implied Wage Distribution



Note: 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.

annuitized cost of recruiting a worker of any type i, $\kappa_i/\lambda_f(\theta_i)$, regardless of their monopsony power. In our calibration, the efficient component of the markdown is about 1-2% and therefore accounts at most for about four percent of the average markdown of 0.75. The remaining portion of the markdown is due to firms' monopsony power, which implies that firms earn substantial monopsony profits in equilibrium.

FIGURE 3: Implied Wage Markdowns



Note: Steady-state wage markdowns $w_{\ell z}/F_{\ell z}$ of selected z-types among non-college workers. "Equilibrium markdown" corresponds to our calibrated model. "Efficient markdown" corresponds to the equilibrium of the model without monopsony power ($\omega \to \infty$). The x-axis corresponds to the initial wage $w_{\ell z}$ earned by a particular z-type.

4 The Short-Run and Long-Run Effects of the Minimum Wage

We now use our quantitative model to study the distributional effects of large increases in the minimum wage. We assume that the minimum wage is unexpectedly introduced starting from the initial steady state

without a minimum wage.[22] We begin with the steady-states implications of such a policy change. We then investigate the dynamic path of the economy from the original steady state to the new one and highlight how our putty-clay technology slows down the transition between them. We next show that our putty-clay technology allows the model to match the small short-run effects of the minimum wage on employment estimated in the empirical literature. We then contrast the effects of a permanent increase in the minimum wage with those of a temporary one, which mimics the erosion by inflation and real productivity growth of a permanent increase in the *nominal* minimum wage. We end with some robustness exercises.

4.1 Long-Run Effects of the Minimum Wage

The minimum wage $\underline{w}$ modifies our steady state equilibrium conditions relative to those described in Section 2. In particular, for a household of type i for which the minimum wage binds, the steady-state wage equation in (28) no longer applies. We also replace w_i with $\underline{w}$ in the vacancy-posting condition in (26) and the first-order condition for search in (27). The remaining conditions are unchanged.

Aggregate Results. We begin our long-run analysis by studying the consequences of the minimum wage on aggregate non-college employment and labor income.[23] The left panel of Figure 4 shows that aggregate employment is a bell-shaped function of the minimum wage, which we refer to as an *employment Laffer curve*. As discussed in Section 2, this shape reflects the fact that a small increase in the minimum wage reduces the average monopsony distortion in the labor market, bringing wages and therefore employment closer to their efficient levels. A large minimum wage increase, however, pushes the affected workers' wages above their efficient level reducing employment.

The employment Laffer curve peaks at about \$9 per hour so that a \$15 minimum wage is well beyond this peak and reduces non-college employment in the long-run by about 12%. The left panel of Figure 4 also shows that the majority of these changes in employment are due to changes in the labor force participation rate rather than in the unemployment rate. To see how these rates are linked, note that

$$\underbrace{\Delta \log n_g}_{\text{employment}} \approx \underbrace{\Delta \log(s_g + n_g)}_{\text{labor force}} - \underbrace{\Delta \left(\frac{s_g}{s_g + n_g} \right)}_{\text{unemployment}}, \tag{33}$$

where n_g is the aggregate employment rate of education group $g \in \{\ell, h\}$, s_g is aggregate search effort of the group, $s_g + n_g$ is the labor force participation rate of the group, and $s_g/(s_g + n_g)$ is the unemployment rate of the group.[24] Across the range of minimum wages considered, the total change in non-college

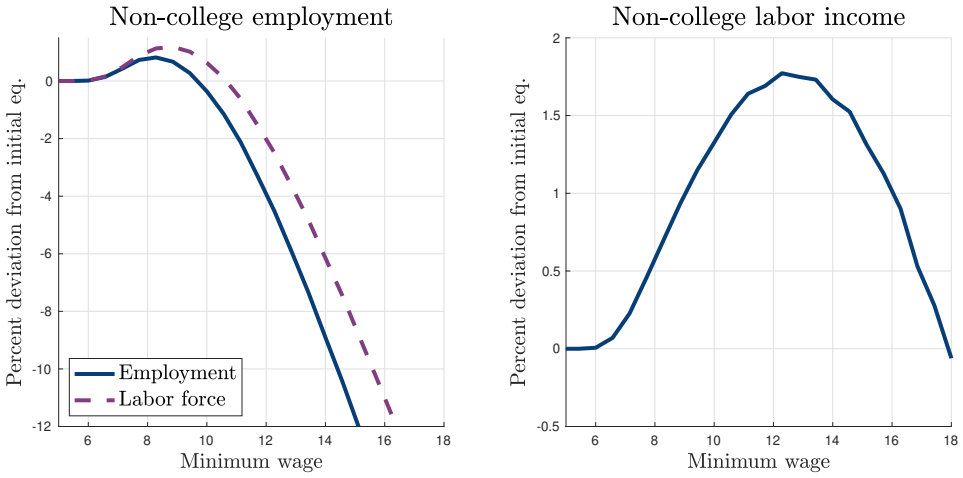
[22] With our calibrated wage distributions, the current national minimum wage of \$7.25 would barely bind so that the initial steady state is a useful approximation to the current national policy regime.

[23] We focus on non-college workers because the new minimum wage is barely binding for college workers. Appendix C shows how college employment and labor income vary with the minimum wage at both the aggregate and micro level.

[24] To derive this decomposition, observe that $n_g = \frac{n_g}{s_g + n_g} \times (s_g + n_g) = \left(1 - \frac{s_g}{s_g + n_g}\right)(s_g + n_g)$. Taking logs of both sides, it follows that $\log n_g = \log\left(1 - \frac{s_g}{s_g + n_g}\right) + \log(s_g + n_g)$. Using the fact that $\log x \approx x - 1$ for small x, this latter expression becomes $\log n_g \approx -\frac{s_g}{s_g + n_g} + \log(s_g + n_g)$. Finally, by taking differences, we obtain (33).

employment $\Delta \log n_\ell$ is primarily driven by changes in the participation rate $\Delta \log(s_\ell + n_\ell)$ rather than changes in the unemployment rate $\Delta(s_\ell/s_\ell + n_\ell)$. For example, a \$15 minimum wage reduces non-college employment by 11.7%, decreases the non-college labor force participate rate by 8.5%, and increases the non-college unemployment rate by 3.2 percentage points.

FIGURE 4: Long-Run Aggregate Minimum Wage Laffer Curves

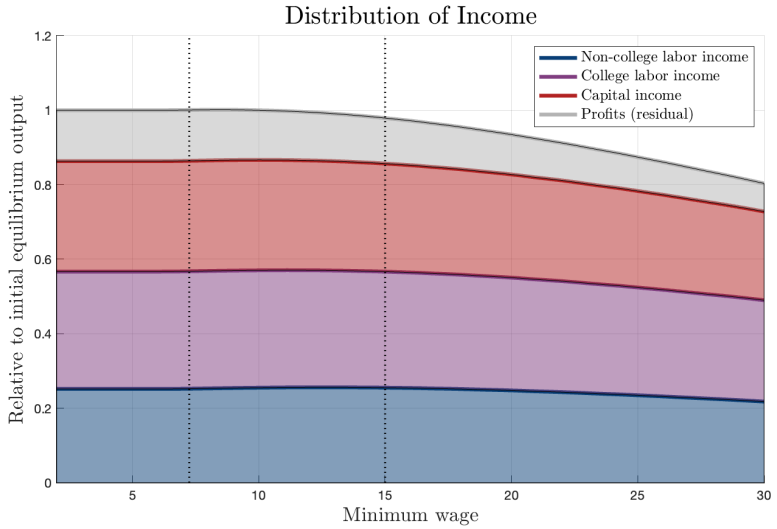


Note: Steady-state outcomes as a function of the minimum wage $\underline{w}$. The left panel plots the percentage change in aggregate non-college employment and the percentage change in of aggregate non-college labor force as a function of the minimum wage (see the decomposition (33)). The right panel plots the aggregate labor income of non-college workers as a function of the minimum wage. The x-axis is the level of the minimum wage $\underline{w}$ which binds on the same amount of initial workers as the associated level in the data.

The fact that the long-run disemployment effect of the minimum wage manifests itself through lower labor force participation rather than higher unemployment provides guidance for the outcomes that empirical work should focus on. Modeling endogenous search effort is not only necessary to distinguish among the implications of the minimum wage on employment and participation but also to account for the documented responses of employment and search activities to the minimum wage. Otherwise, the entire long-run change in employment would have to be driven by changes in the unemployment rate because all non-employed workers would immediately search for new jobs. Such an outcome, though, would contradict Adams, Meer and Sloan (2022)'s empirical finding that search effort does not significantly increase in the few months after an increase in the minimum wage.

The right panel of Figure 4 shows that aggregate labor income as a function of the minimum wage has also a standard Laffer curve shape. This *labor income Laffer curve* peaks at a higher level of the minimum wage than the employment Laffer curve, at about \$13 rather than about \$9. Intuitively, the minimum wage increases average wages even if it decreases employment. Specifically, a \$15 minimum wage increases non-college labor income by about 1.4% in the long run. So, although employment falls for non-college workers when a \$15 minimum wage is introduced, labor income for these workers, on average, increases.

FIGURE 5: Distribution of Aggregate Income

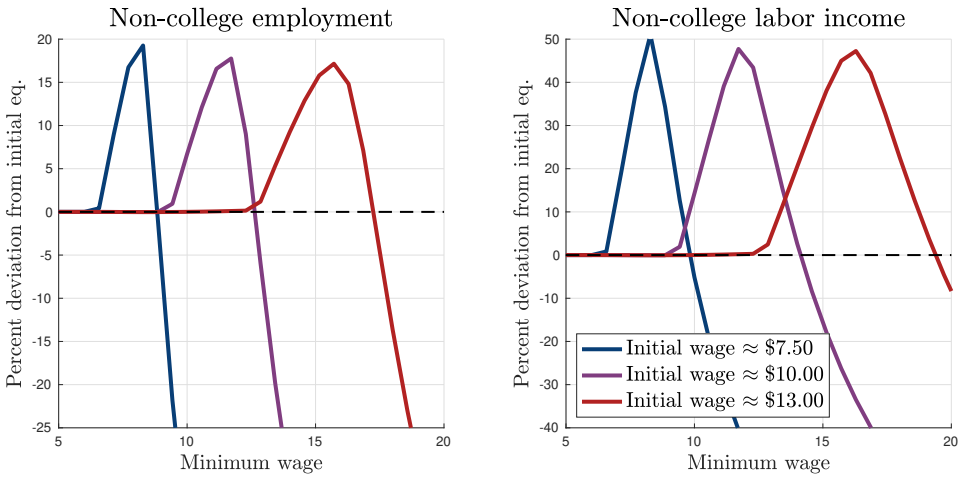


Note: Steady-state output and shares accounted for by non-college labor income, college labor income, capital income, and profits as a function of the minimum wage. The y-axis is normalized such that aggregate income equals 1 without the minimum wage. The x-axis is the level of the minimum wage that binds on the same amount of initial workers as in the data.

In order to illustrate how the minimum wage redistributes aggregate income across different groups, Figure 5 plots how the minimum wage affects the four components of aggregate income in the long run: non-college labor income, college labor income, capital income, and firms' profits, which, as discussed, primarily reflect firms' monopsony power. Consistent with the labor income Laffer curve in Figure 4, a \$15 minimum wage raises non-college labor income by 1.4%. Firm profits fall substantially, suggesting that the minimum wage successfully redistributes resources from firms to workers. This redistribution, however, has a cost: the minimum wage reduces total employment, the aggregate capital stock, and ultimately lowers total output by 2.2%. We turn next to investigate which workers pay the price for this redistribution.

Micro-Level Results. How does a higher minimum wage affect individual worker of different productivity within the non-college group? Figure 6 plots type-specific long-run employment and labor income Laffer curves for three non-college workers with different levels of ability z, namely, workers with initial wages of \$7.50, \$10, and \$13. Consider first the (red) curves for the worker initially earning \$13. As the minimum wage increases from \$13 to about \$16, employment for the most productive worker increases as the minimum wage locally offsets firms' monopsony power. But for the worker initially earning \$10 (violet lines), the same increase in the minimum wage drastically reduces this worker's employment. Indeed, the labor income of this worker of intermediate ability would have peaked a a minimum wage of about \$12. Such a minimum wage would have locally offset the monopsony distortions for this worker. Given the high long-run substitutability among non-college workers estimated by Card and Lemieux (2001), which our

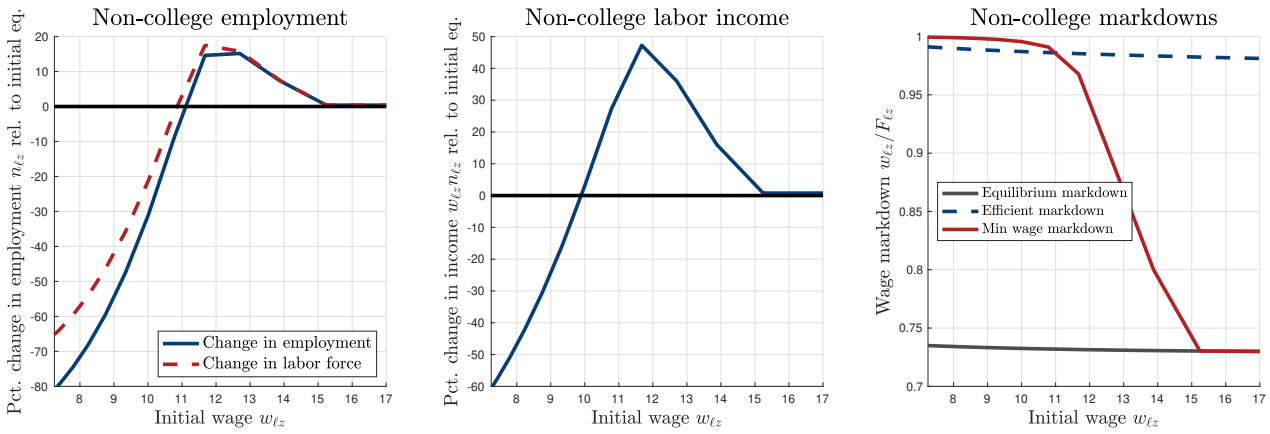
FIGURE 6: Disaggregated Long-Run Minimum Wage Laffer Curves



Note: 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 $\overline{w}$ which binds on the same amount of initial workers as the associated level in the data. Initial wages rounded to the nearest half dollar.

model replicates, any higher minimum wage simply leads firms to substitute away from such a worker. For the lowest-ability worker, who was initially earning \$7.50 (blue lines), employment would peak at a minimum wage of about \$9.40. Then, a minimum wage of \$15 is simply too high and would lead to an even more dramatic fall in employment for such a worker.

FIGURE 7: Distributional Effects of \$15 Minimum Wage Among Non-College Workers in Long Run



Note: Steady-state outcomes for selected z-types among non-college workers for a \$15 minimum wage. The left panel plots the percentage change in employment $n_{\ell z}$ relative to the initial steady state, the middle panel plots the percentage change in labor income $w_{\ell z} n_{\ell z}$ relative to the initial steady state, and the right panel plots the markdown $w_{\ell z}/F_{n\ell z}$ for three different parameterizations: (i) $\underline{w} = 0$ ("equilibrium markdown"), (ii) $\underline{w} = 0$ and $\omega \to \infty$ ("efficient markdown"), and (iii) $\underline{w} = \$15$ ("minimum wage markdown"). The x-axis corresponds to the initial wage $w_{\ell z}$ of a type-z worker in the initial steady state.

Figure 7 plots the long-run distributional consequences of a \$15 minimum wage for non-college workers,

as a function of the initial wage $w_{\ell z}$ earned by a type-z worker in the initial steady state. The left panel of the figure shows that employment falls for all non-college worker types who were initially earning less than $11 per hour, who account for 26% of all non-college workers and 62% of the workers on whom the minimum wage binds. The decline is largest among lowest-earning workers because their efficient wages are substantially below the $15 minimum. The middle panel shows that the same broad pattern holds in terms of labor income, although the set of workers whose labor income falls is smaller because wages rise even for workers whose employment falls. Indeed, we find that 17% of all non-college workers and 41% of non-college workers initially earning less than $15 experience declines in labor income. The first rows of Table 5 provides a summary of the aggregate and distributional effects of a $15 minimum wage in the long run for non-college workers (top panel) and for those non-college workers initially earning less than $15 (bottom panel). Overall, in the long run, the minimum wage disproportionately reduces the employment and income of precisely the group of workers that it is meant to benefit.

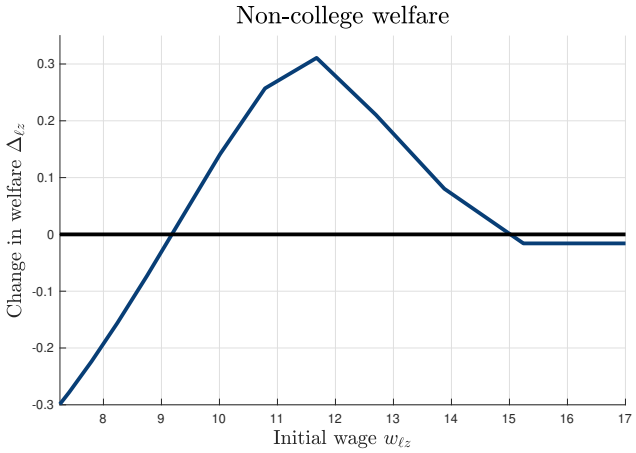
TABLE 5: Summary of Long-Run Findings for the Implementation of $15 Minimum Wage

Time Period	Δn_ℓ	$\Delta w_\ell n_\ell$	Share $(\Delta n_{\ell z} < 0)$	Share $(\Delta w_{\ell z} n_{\ell z} < 0)$
All non-college workers				
Long run (new steady state)	-11.7%	1.4%	0.26	0.17
Short run (first 5 years)	-3.5%	10.5%	0.26	0.00
Total effect, current cohorts	-6.6%	7.6%	0.26	0.06
Total effect, cohorts entering in 10 years	-9.1%	5.2%	0.26	0.17
Workers initially earning less than $15				
Long run (new steady state)	-28.0%	4.8%	0.62	0.41
Short run (first 5 years)	-8.6%	40.7%	0.62	0.00
Total effect, current cohorts	-16.9%	26.5%	0.62	0.16
Total effect, cohorts entering in 10 years	-22.1%	17.9%	0.62	0.41

Note: The top panel computes statistics for all non-college workers, whereas the bottom panel computes statistics for non-college workers earning less than $15 in the initial equilibrium. "Long run (new steady state)" entries compare the final steady state with $\underline{w} = \$15$ to the initial steady state, Δn_ℓ is the change in non-college employment, $\Delta w_\ell n_\ell$ is the change in aggregate non-college labor income, Share$(\Delta n_{\ell z} < 0)$ is the share of initial workers whose employment declines in the new steady state, and Share$(\Delta w_{\ell z} n_{\ell z} < 0)$ is the share of initial workers whose income declines in the new steady state. "Short run (first 5 years)" entries compute analogous statistics along the transition path up to 5 years after the introduction of the higher minimum wage. "Total effect, current cohorts" entries compute the present values of changes over the entire transition path. Δn_ℓ is the change in the present value of non-college employment, $\Delta w_\ell n_\ell$ is the change in the present value of aggregate non-college labor income, Share$(\Delta n_{\ell z} < 0)$ is the share of initial workers whose present value of employment declines relative to the initial steady state, and Share$(\Delta w_{\ell z} n_{\ell z} < 0)$ is the share of initial workers whose present value of labor income declines relative to the initial steady state. "Total effect, cohorts entering in 10 years" entries compute the same statistics but starting 10 years after the introduction of the higher minimum wage. Shares of workers along the transition path calculations are re-scaled to maintain comparability to the steady state calculations (which use a finer grid).

To further understand the mechanism behind these results, the right panel Figure 7 plots the wage markdown in the new steady state as a function of workers' ability, again expressed in terms of their initial wage $w_{\ell z}$. For low-earning workers, the minimum wage imposes lower markdowns than are efficient.

FIGURE 8: Welfare Effects of the $15 Minimum Wage



Note: Change in welfare for non-college workers $\Delta_{\ell z}$ from (34). The x-axis corresponds to the initial wage $w_{\ell z}$ earned by a worker of ability z in the initial steady state.

Such markdowns induce firms to decrease their employment to the point where their marginal product approximately equals the minimum wage. For middle-earning workers, the minimum wage lowers the wage markdown closer to its efficient level and so induces an increase in their employment—see the left panel of the figure. Since for high-earning workers the minimum wage does not bind, it does not affect their wage markdowns or employment.

To measure the effects of the minimum wage on welfare, we define the *long-run welfare change* 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 x^* denotes the value of variable x in the old steady state and $\widetilde{x}$ denotes its value in the new steady state with the higher minimum wage. Here Δ_i measures the percentage change in steady-state consumption that would make a household indifferent between the old and new equilibrium. Hence, Δ_i is positive if the policy makes the household better off and negative if the policy makes the household worse off.[25] Figure 8 shows that about one third of the households affected by the new minimum wage experience a significant decline in their welfare. Welfare increases for the remaining workers, despite their employment falling, because of their higher labor income.

It is worth noting that our model relies on a representative family construct and hence implies perfect risk-sharing against job loss within worker types. We conjecture that in an alternative environment with imperfect risk sharing, the welfare losses from the minimum wage would be even larger because workers who lose their jobs would experience a larger decline in consumption. Allowing for such imperfect risk

[25]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.

sharing is outside the scope of our framework because our fundamental decision-making unit is the family.[26]

4.2 Short Run vs. Long Run

Having explored the long-run consequences of the minimum wage, we turn to studying the transition path of the economy to the new steady state. Along this transition path, the minimum wage imposes a lower bound $w_{ijt} \geq \underline{w}$ 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

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

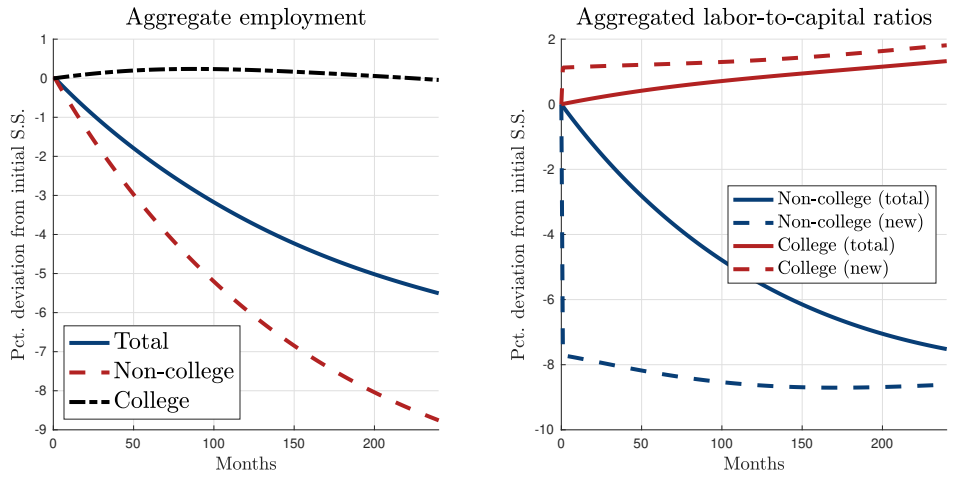
Note that $\underline{W}_t$ is the smallest present value of wages consistent with the minimum-wage constraint $w_{ijt} \geq \underline{w}$ in each period of a match, and depends on time because the intertemporal prices do. We add the constraint in (35) to a firm's problem. If for a consumer of type i the minimum wage constraint does not bind, first-order conditions are the same as before. When the minimum wage constraint binds, we set $W_{ijt} = \underline{W}_t$ so that a firm's wage offer satisfies (35). Upon the introduction of the new 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.[28] We turn to examine the dynamics of employment, income, and welfare.

Employment Dynamics. We first show how our putty-clay technology affects the response of employment to the introduction of a \$15 minimum wage. The left panel of Figure 9 plots the transition path of aggregate employment by education group, and shows one of our key results: it takes a long time for employment to adjust in response to a large increase in the minimum wage. For instance, four years after the introduction of a \$15 minimum wage, non-college employment has only decreased by one-quarter of its ultimate decline. Even twenty years out, the economy has not yet completely converged to the new steady state. The right panel of the figure shows the associated dynamics of the aggregate labor-to-capital ratios chosen by firms. After the imposition of the new minimum wage, firms immediately reduce the non-college labor intensity of newly installed capital (dashed blue line), namely, they invest in capital requiring much

[26] In particular, our model of monopsony power comes from the CES aggregate disutility of work in (1), which is only defined at the family level. Therefore, relaxing the assumption of perfect risk-sharing would require us to abandon the representative family assumption to explicitly model (and quantify) the nature of non-wage amenities and their distribution across individuals, which is outside the scope of this paper.[27] That said, perfect risk-sharing in our model only occurs within the fine level of a worker's education and ability. We do allow for imperfect insurance across education and ability types, which already helps isolate perverse distributional effects of large minimum wage policies.

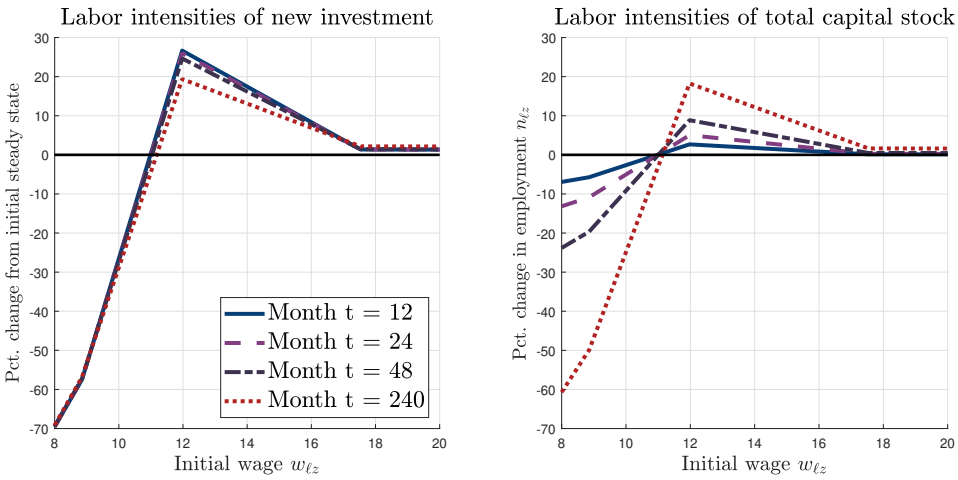
[28] If, instead, we allow a firm to lower the wage of an existing worker for whom the minimum wage does not bind, then a firm has an incentive to lower the wage until such a worker is just indifferent between quitting or remaining with the firm. We assume that the original contract the worker signed 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—workers with slightly higher productivity than those for whom the new minimum wage binds, who are the closest substitutes of the workers for whom the new minimum binds. We rule out this possibility because, without such a clause, the unexpected introduction of a minimum wage allows firms to renege on an existing wage contracts. 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.

FIGURE 9: Employment Dynamics Along the Transition Path in Response to $15 Minimum Wage



Note: Transition paths following an unexpected imposition of a $15 minimum wage. The left panel plots the employment of non-college workers, college workers, and total employment. The right panel plots the associated labor-to-capital ratios of new capital (dashed lines) and of the entire capital stock (solid lines).

FIGURE 10: Labor Intensities in Response to a $15 Minimum Wage at Various Time Horizons

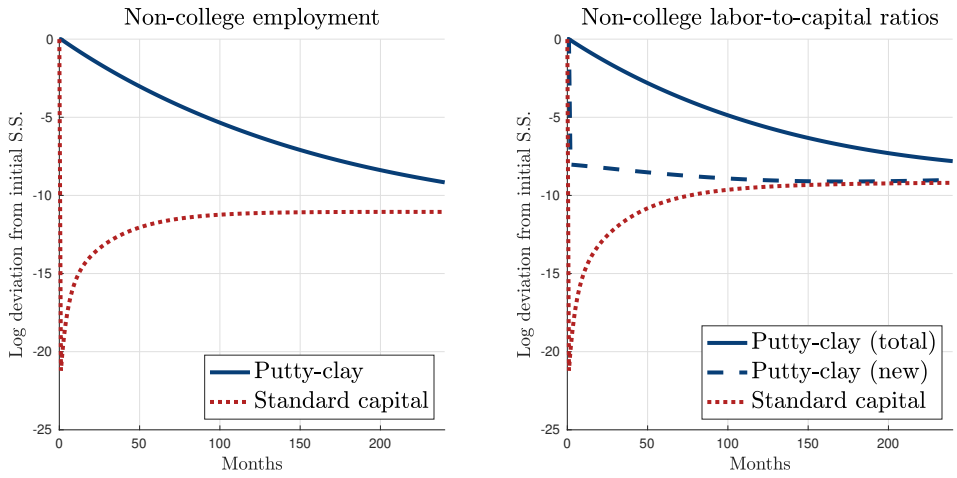


Note: Labor intensities along the transition path following an unexpected imposition of a $15 minimum wage starting from the initial equilibrium, as a function of the initial wage $w_{\ell z}$ earned by a non-college worker with ability z. The left panel plots non-college labor intensities on new capital goods as a function of worker ability, $v_{\ell z}$. The right panel plots the corresponding non-college labor intensities on the entire capital stock.

less non-college labor, but it takes much longer for the non-college labor intensity of the total capital stock to significantly decline (solid blue line), since it includes the previously installed capital.

In order to understand the role of the putty-clay technology in driving this dynamics, Figure 10 plots the labor intensities of new and existing capital over the transition path for individual non-college workers

FIGURE 11: Putty-Clay Capital vs. Standard Capital in Response to Imposition of $15 Minimum Wage



Note: Transition path following an unexpected imposition of a $15 minimum wage starting from the initial equilibrium without the minimum wage. The left panel plots the aggregated employment of non-college workers, college workers, and total employment. The right panel plots the associated labor-to-capital ratios. "Putty-clay" refers to the baseline model. "Standard capital" refers to the version of the model with standard capital described in Section 2.3.1.

in response to a $15 minimum wage.[29] The left panel plots the labor-to-capital ratios $v_{\ell z}$ of newly installed capital goods for each worker type (y-axis) against that worker type's initial wage $w_{\ell z}$ (x-axis). In response to a $15 minimum wage, firms immediately substitute away from low-ability workers towards higher-ability ones for the newly installed capital. The magnitude of such a substitution is fairly stable at various horizons along the transition. The right panel of the figure, though, shows that the implied dynamics for the labor intensity of the *total* stock of capital is much slower. The reason is that firms continue to operate their existing capital stock along the transition path, which requires the old steady-state mix of worker types chosen before the introduction of the $15 minimum wage. This old capital accounts for the majority of the aggregate capital stock early on in the transition, but as it depreciates away, it is replaced by new, less labor intensive capital. Hence, the depreciation rate δ is crucial in determining the speed of transition, as the distribution of capital types along the transition path in Appendix C highlights.

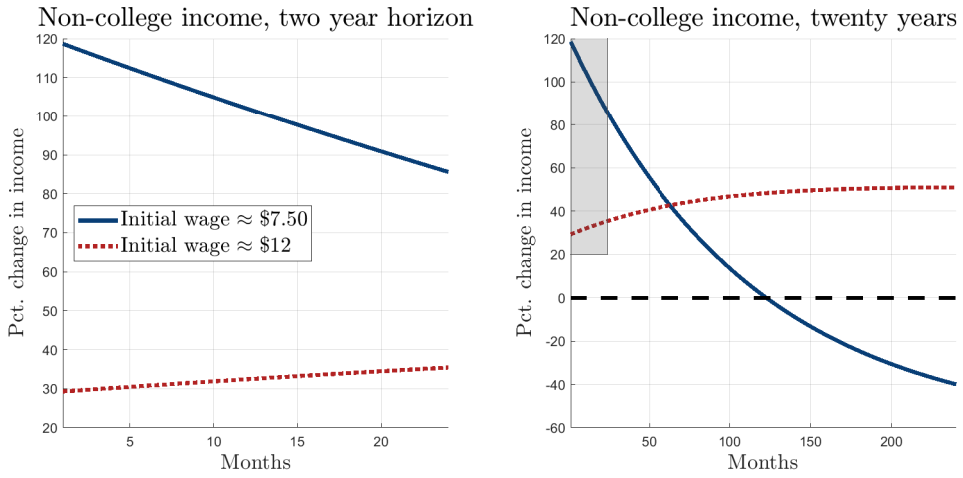
Figure 11 shows that the model with standard capital converges to the new steady state much more quickly than our model with putty-clay capital. In fact, in the standard model, total non-college employment immediately falls *below* its new steady-state level upon the introduction of the new minimum wage. This abrupt decline occurs because firms immediately fire the lowest-ability workers—they are not needed to operate the existing capital as in the putty-clay model. Over time, firms hire new workers of intermediate ability to replace the low-ability ones, but this hiring process takes time due to the search frictions in the labor market. This dynamics is nonetheless fairly rapid in the sense that non-college employment

[29] 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.

converges to its new steady-state level only after about four years.

Income and Welfare Dynamics. Turning now to the implied dynamics of labor income and welfare, Figure 12 plots the dynamics of labor income for two types of non-college workers in response to a \$15 minimum wage: a low-ability worker, who was earning less than \$7.50 in the initial equilibrium and whose employment *decreases* in the long run due to the higher minimum wage, and an intermediate-ability worker, who was earning around \$12 in the initial equilibrium and whose employment *increases* in the long run due to the higher minimum wage. The left panel shows that right after the introduction of the higher minimum wage, labor income increases more for the low-ability worker than for the intermediate-ability worker relative to their initial levels of income. This result occurs because the imposition of a \$15 minimum wage immediately doubles low-ability workers' wages, but the putty-clay dynamics imply that their employment does not significantly decline yet. The right panel of the figure shows that their labor income gradually falls over time, as firms substitute away from these type of workers. After about ten years, low-ability workers' labor income falls below its initial level and remains below forever after.[30]

FIGURE 12: Labor Income for Two Types of Workers in Response to \$15 Minimum Wage



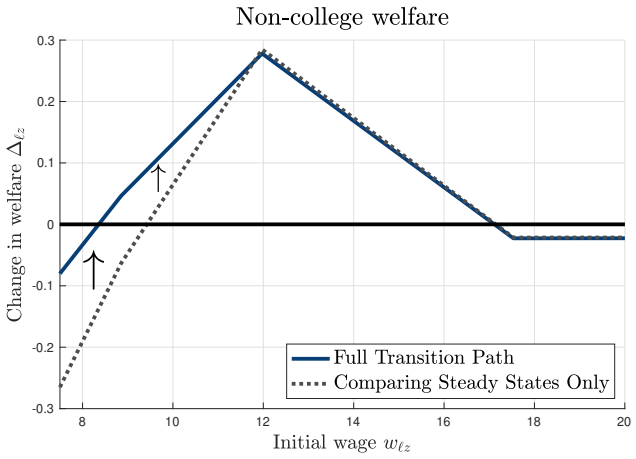
Note: Flow labor income $w_{\ell z,t} n_{\ell z,t}$ for two types of non-college workers along the transition path following the unexpected imposition of a \$15 minimum wage starting from the initial equilibrium without the minimum wage. The left panel plots these quantities over the first two years of the transition. The right panel plots these quantities over twenty years. The grey shaded area in the right panel is the area corresponding to the left panel.

The second rows of each panel of Table 5 summarize our short-run results, which focus on the employment and labor income changes for various types of workers five years after the imposition of a \$15 minimum wage. These results indicate that the effects of a high minimum wage substantially *change* over time. For low-wage workers, the minimum wage has sizable short-term benefits, but eventually large

[30] These results are consistent with the finding in Dube (2019) of a sharp increase in household income at the bottom of the distribution in the few years after the implementation of minimum wage increases.

long-term costs. Hence, cohorts of the lowest paid workers who start their careers after the introduction of the new minimum wage enjoy less of the short-term benefits and bear more of the long-term costs (the more so, the later these cohorts enter). Importantly, these dynamic effects arise from the slow response of employment to the increase in the minimum wage induced by the putty-clay technology.[31]

FIGURE 13: Welfare Effects of a $15 Minimum Wage



Note: Change in welfare due to a $15 minimum wage. The solid blue line plots welfare along the transition path computed as $\widetilde{\Delta}_i$ from (36). The dotted gray line plots welfare changes comparing only the new to the initial steady state computed as Δ_i from (34). The x-axis corresponds to the initial wage $w_{\ell z}$ of a type-z worker in the initial equilibrium.

Since the current cohort of lower-ability workers earn a higher income in the early phase of the transition to the new steady state, the steady-state-to-steady-state welfare comparisons in Section 4.1 overstate their decline in welfare along the transition path. To assess by how much, we calculate the consumption-equivalent *dynamic welfare change* taking into account the transition dynamics, $\widetilde{\Delta}_i$, from

$$\sum_{t=0}^{\infty} \beta^t u((1+\widetilde{\Delta}_i)c_i^*, n_i^*, s_i^*) = \sum_{t=0}^{\infty} \beta^t u(c_{it}, n_{it}, s_{it}), \tag{36}$$

where x^* denotes the value of variable x in the old steady state and the right hand side is evaluated along the entire transition path. Intuitively, $\widetilde{\Delta}_i$ is the change in period consumption that would make a household indifferent between the old steady state and the economy with the $15 minimum wage. Hence, when $\widetilde{\Delta}_i < 0$, a $15 minimum wage decreases welfare, whereas when $\widetilde{\Delta}_i > 0$, a $15 minimum wage increases welfare. The solid blue line in Figure 13 plots the welfare effects of a $15 minimum wage computed along the entire transition path and the dotted gray line plots the steady-state-to-steady-state welfare effects from Section 4.1. In both cases, the lowest-wage workers' welfare decreases with a $15 minimum wage. However, the set of workers whose welfare falls is smaller once we account for the transition dynamics of

[31] Notice that the share of initial workers who experience a decline in their employment (column 3) is the same over the various time periods we summarize because the set of worker types whose employment falls in response to a $15 minimum wage is the same over time. The only thing that changes is the magnitude of their decline. For example, all 26% of non-college workers who experience employment declines in the long run also experience a smaller employment decline in the short run.

the economy because of their short-term gains in terms of higher labor income described above, which are induced by the putty-clay technology. In fact, only 6% percent of the current cohort of non-college workers are made worse off after the introduction of a $15 minimum wage. In this sense, steady-state welfare comparisons overstate the true welfare change. The employment and income effects over the entire transition path for various types of workers are summarized in the third row of the two panels of Table 5.

We end this subsection by highlighting how welfare effects differ across different cohorts of workers, as summarized in the fourth rows of each panel of Table 5. For the current generation of workers, a $15 minimum wage is beneficial for about 94% of non-college workers and about 84% of workers initially earning less than $15. Although a large fraction of workers are eventually worse off, the benefits of the higher wages they receive in the short run essentially compensate their long-run losses. Future generations of lower-ability non-college workers, instead, will enjoy less and less of the short-run benefits.

4.3 Model Predictions vs. Existing Empirical Estimates

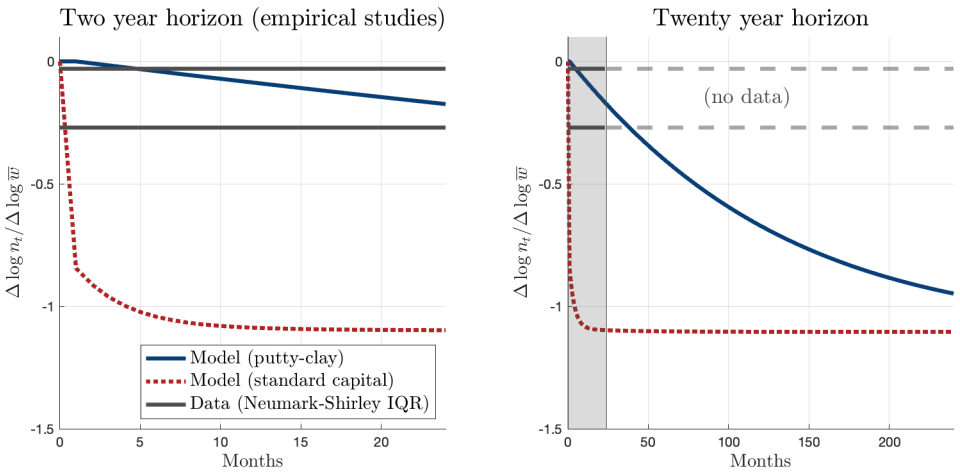
We now compare our model's predictions to the findings of a large empirical literature that has studied the impact of minimum wage increases on employment.[32] In a recent review article, Neumark and Shirley (2021) estimate the elasticity of employment in response to increases in the minimum wage through a meta-analysis of 109 published papers based on subnational variation in the minimum wage in the United States—the entire set of published studies in this literature.[33] As these authors highlight, essentially all of these papers examine the employment effects of relatively small minimum wage changes, that is, changes of $3 or less, and over relatively short time horizons. Specifically, the vast majority of papers examine employment changes up to two years following an increase in the minimum wage. Also, all these papers tend to focus on the employment effects of lower-earning workers, such as teenagers and young adults, who are most likely to be effected by changes in the minimum wage. Neumark and Shirley (2021) document that roughly 80% of the studies they review find zero to small short-run negative employment effects in the two years after a minimum wage increase. The interquartile range of the implied employment elasticities range from -0.03 to -0.27, that is, a one percent increase in the minimum wage tends to lead to between a 0.03% and a 0.27% decline in employment.

In order to investigate whether our model can match these findings, we perform a stylized model experiment to replicate the type of minimum wage increases considered by Neumark and Shirley (2021). Within the model, we consider a $3 increase in the minimum wage (from $7.50 to $10.50) to loosely capture a minimum wage increase of medium to large size typically studied in the data. We focus our results on non-college workers initially earning less than the new minimum wage to replicate the fact that

[32] Fully reviewing this literature is beyond the scope of this paper. However, some prominent papers in this extensive literature include Brown (1988), Card and Krueger (2015), Brown (1999), Neumark et al. (2008), and Cengiz et al. (2019). We discuss some recent papers in more detail Appendix B.

[33] Given the typical multiplicity of findings in this literature even within a given paper, the authors compute the elasticities of interest by identifying the core estimates that support the conclusions from each study, in most cases, as they explain, by relying on responses from the authors of the studies.

FIGURE 14: Stylized Replication of Empirical Work



Note: Transition path of low-wage, non-college employment following an unexpected \$3 increase in the minimum wage. Transition paths are computed assuming that $Q_{t,t+1} = \beta$ for all t. Solid grey lines correspond to the interquartile range of the estimates of the elasticity of employment to the minimum wage in Neumark and Shirley (2021). "Model (putty-clay)" refers to the baseline model. "Model (standard capital)" refers to the version of the model with standard capital.

most of the literature examines employment effects for the potentially affected sub-populations, such as teenagers or restaurant workers. Finally, we study a "small open economy" version of the model in that we hold consumption prices $Q_{t,t+1} = \beta$ fixed at each horizon along the transition path, in order to capture the idea that the minimum wage changes considered, being sub-national, are local in nature and so aggregate economic conditions should be treated as fixed—all the papers reviewed by Neumark and Shirley (2021) exploit local, as opposed to national, minimum wage increases.

The left panel of Figure 14 plots the time path of employment for initially low-earning workers in response to a \$3 minimum wage increase over the first two years of the transition. We normalize the change in log-employment $\log n_{\ell t}$ by the log-change in the minimum wage $\log \underline{w}$ in order to conform to the elasticities reported in Neumark and Shirley (2021). The short-run elasticity of employment with respect to the minimum wage—namely, the response of employment over the first two years of the transition—is around -0.2, which is well within Neumark and Shirley (2021)'s range. This is shown in the solid blue line of the figure, which essentially lies within the two gray horizontal lines that describe the interquartile range of the employment elasticities reported by Neumark and Shirley (2021). In contrast, the dashed red line in the figure shows that the model with standard capital implies an abrupt and large decline in employment after an increase in the minimum wage to \$10.5, because firms no longer need low-ability workers to operate the existing capital stock and, as a result, immediately substitute away from them. Hence, our putty-clay technology ensures that the short-run effects of small and local changes in the minimum wage predicted by our model are in accord with the data.

The right panel of the figure zooms out to longer horizons to illustrate the long-run elasticity of

employment with respect to the minimum wage generated by our hypothetical experiment. This long-run elasticity for low-wage workers is around -1.1, more than an order of magnitude larger than the two-year elasticity estimated in the data. In this sense, the short-run elasticities typically estimated are relatively uninformative about the long-run ones. In contrast, in the counterfactual model with standard capital, no meaningful difference exists between short-run and long-run employment elasticities. This result highlights the importance of the putty-clay technology and, more generally, of accounting for the adjustment of all inputs to correctly infer the long-run consequences of a minimum wage policy.

Almost all empirical papers in the minimum wage literature estimate short-run labor market responses to minimum wage increases. An exception is Clemens and Strain (2021), who estimate the employment effects of small (less than \$2.50) and large (more than \$2.50) state-level changes in the minimum wage in both the very short run (up to 4 years after the minimum wage change) and medium run (from 4 to 6 years after the minimum wage change). They find that (i) in both the short and medium run, small and large minimum wage changes have insignificant effects on employment and (ii) in the medium run, large minimum wage changes have negative and statistically significant effects on employment. Our model replicates these findings. In particular, as apparent from the right panel of of Figure 14, our five-year elasticity is much larger than our one-year one. Likewise, within our model, we find very small elasticities in both the short and medium run when we consider a minimum wage increases of only \$1, that is, an increase from \$7.50 to \$8.50.[34]

Overall, we view it as a strength that our model can match the recent empirical findings in Neumark and Shirley (2021) and the short- and medium-run findings in Clemens and Strain (2021). Importantly, because of putty-clay frictions, our model implies that the long-run responses of employment to increases in the minimum wage are likely to be much larger than the estimated short- and medium-run ones.

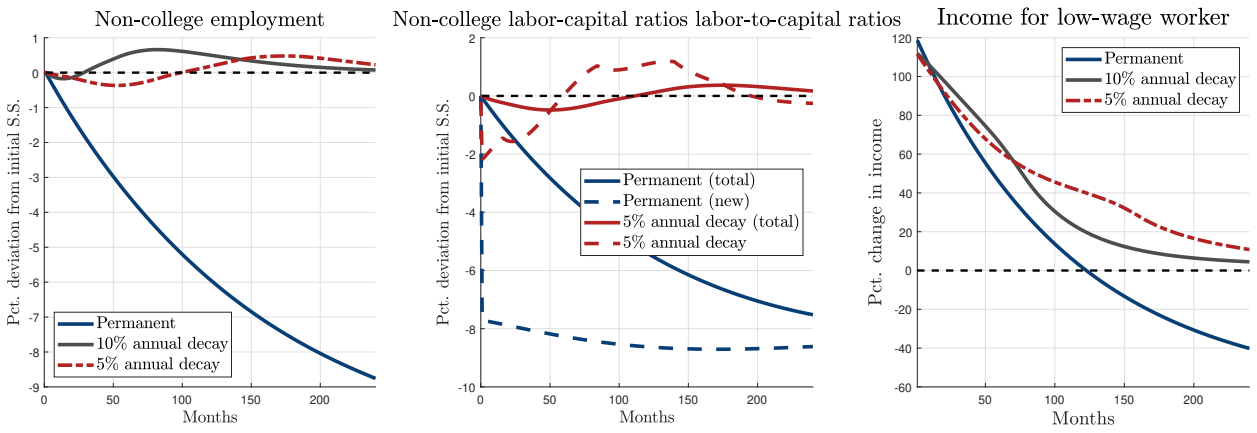
4.4 Temporary Changes in the Minimum Wage

Although our analysis so far has considered only permanent increases in the real minimum wage, actual legislation in the United States has changed the nominal minimum wage, so the real minimum has declined in between legislated changes due to general price inflation. In addition, productivity growth implies that the real value of the minimum wage may decrease relative to average wages in the economy over time. Historically, these periods in-between legislated changes in the minimum wage have lasted five to ten years, generating a saw-tooth pattern in the real value of the minimum wage relative to average wages (see, for instance, Sorkin (2015)). In this subsection, we consider a simple example in order to examine

[34] Clemens and Strain (2021) refer to small minimum wage changes as changes of about \$1 and to large minimum wage changes as minimum wage changes of about \$3 and focus on low-skill workers (aged between 16 and 25 years without a high-school diploma) and young workers (all those aged between 16 and 21). These authors document that in the short run, small minimum wage increases have effectively no impact on employment whereas large minimum wage increases have a small but insignificant effect on employment. In the medium run (4+ years ahead), small minimum wage increases have increased the employment rate in the medium run by 0.57 percentage points but this effect is statistically insignificant and not robust across specifications. Large minimum wage increases, instead, have decreased the employment rate by 2.65 percentages point, which is a robust and statistically significant effect.

the implications of this dynamics. Specifically, we model a temporary increase in the real minimum wage through a path for it that decays at constant rate g, $\underline{w}_t = \underline{w}_0(1+g)^t$. The decay rate g corresponds to the sum of inflation and real productivity growth over time. For illustration, we consider values of g equal to 5% and 10% annually, which we interpret as corresponding to 2% productivity grow and an inflation rate between 3% and 7%, respectively—these values of inflation can be thought of as the sum of 2% baseline inflation plus any additional inflation induced by the minimum wage. The process for $\underline{w}_t$ is a simple way to capture the periods in-between legislated changes in the nominal minimum wage.

FIGURE 15: Temporary vs. Permanent Minimum Wage Increases



Note: Transition paths following an unexpected imposition of a time-varying minimum wage $\underline{w}_t$ that starts at \$15 and then decays according to (**??**). The left panel plots the employment of non-college workers, college workers, and total employment. The right panel plots the associated capital-to-labor ratios. "Permanent" corresponds to the case of a constant real value of the minimum wage in Section 4.2. "10% annual decay" corresponds to the case of a decaying real value of the minimum wage in (**??**) with $g = 10\%$ annually. "5% annual decay" corresponds to a value of $g = 5\%$ annually.

The left panel of Figure 15 compares the implications of these temporary minimum wage paths to those of a permanent change in the minimum wage analyzed so far. The response of employment displays two important differences between the two cases. First, the temporary minimum wage paths feature non-monotonic dynamics: employment declines below its initial level immediately after the introduction of the minimum wage, but then increases above the initial level before reverting back to its steady-state level. This occurs because the real value of the minimum wage is initially relatively high—above the efficient level of wages for sufficiently many workers that it decreases employment—but later declines to a relatively low level—close to the efficient level of wages for a large enough set of workers that it actually increases employment. The decay rate g disciplines the length of these two phases. For either value of g, though, the real value of the minimum wage eventually falls enough that it becomes nonbinding, at which point the economy returns to the initial steady state.

The second difference between the permanent and transitory paths is that the overall response of employment to the transitory changes is an order of magnitude smaller than the response to the permanent

change. This difference is partly due to the fact that the present value of the transitory minimum wage increase is smaller than the present value of the permanent minimum wage increase. But even in the early phase of the transition, when these present values are similar, the irreversibility of capital implied by the putty-clay technology induces firms to reduce employment by much less in the transitory case than in the permanent case. This result occurs because firms anticipate that any new capital they install will be in use later on in the transition, when the decaying minimum wage implies that lower-productivity workers will once again be relatively inexpensive to hire. Consistent with this mechanism, the middle panel of Figure 15 shows that the labor intensities of newly installed capital change much less in the initial phase of the transitory paths compared to the permanent path and hence the desired mix of workers changes much less than in the permanent case.

The right panel of Figure 15 compares the paths of labor income for a low-wage worker initially earning around $7.50 in the permanent and temporary cases. Two opposing forces matter for this comparison: the real value of the minimum wage declines over time in the temporary case, dampening its impact on labor income relative to the permanent case, but employment declines more strongly in the permanent case, dampening its impact on labor income relative to the temporary case. These two forces roughly offset each other in the early stage of the transition, but over time income stays elevated in the temporary case whereas it declines in the permanent one. In this sense, a one-time unanticipated introduction of a decaying minimum wage allows a low-wage worker to reap the short-run benefits of the minimum wage without incurring the long-run costs. Note that we use this temporary case to illustrate the mechanisms of our model rather than suggest this is a policy that policymakers could exploit, especially repeatedly, because a full analysis would require a carefully modeling of the impact of this policy on ex-ante expectations.

4.5 Robustness and Additional Results

In this subsection, we study the sensitivity of our main results to alternative parameter choices and then discuss some back-of-the-envelope calculations for assessing the distributional effects of a $15 minimum wage across U.S. states and sectors.

Alternative Parameter Estimates. We first show how the long-run effects of the minimum wage depend on the key parameters highlighted in Section 3.1 such as the degree of monopsony power ω and the production function elasticities α, ρ, and ϕ. We also show that our results are robust to changes in other parameters, such as the labor supply parameter γ_n or the degree of search frictions.

Table 6 summarizes these results. The first three rows compare our baseline calibration to two alternative ones that correspond to three different levels of wage markdowns and so firms' monopsony power as measured by ω: our baseline value of 0.75, a larger markdown of 0.65, and a smaller markdown of 0.85. In each case, we recalibrate the model in order to match the new average wage markdown along with all of the other moments targeted in the baseline calibration. The larger markdown of 0.65, which

is consistent with the value estimated by Yeh, Macaluso and Hershbein (2022), generates a larger degree of firms' monopsony power and therefore increases the potential benefits of the minimum wage. At the aggregate level, a \$15 minimum wage decreases aggregate non-college employment by 1.5% compared to 12% in our baseline. Appendix C.4 also shows that, at the micro level, a higher degree of monopsony power expands the set of households who benefit from the minimum wage and reduces the employment losses for those who are hurt by it. In contrast, with a smaller degree of monopsony power, as consistent with the average markdown of about 0.85 estimated by Seegmiller (2021) or Lamadon, Mogstad and Setzler (2022), a \$15 minimum wage decreases aggregate non-college employment by nearly 20% rather than by 12% as in our baseline. With this lower level of monopsony power, roughly 50% of non-college workers will suffer employment losses or income declines in the long run. Collectively, these results highlight the importance of firms' monopsony power in motivating the use of the minimum wage as a policy tool.

The next set of entries in Table 6 explore the sensitivity of our results to different long-run elasticities of substitution estimated in the literature. First, we set a higher value of the across-education-group elasticity $\rho = 4$ from recent work by Bils, Kaymak and Wu (2020), as opposed to $\rho = 1.4$ in our baseline, and find that our main results are not very sensitive to the value of this parameter.[35] Second, we set a higher value of the within-education-group elasticity $\phi = 6$ from the upper bound of the range estimated by Card and Lemieux (2001). In this case, the minimum wage has a stronger negative effect on employment because it increases the ability of firms to substitute away from workers bound by the minimum wage. We also consider two different values of the elasticity of substitution between capital and labor α, corresponding to more complementary than Cobb-Douglas ($\alpha = 0.7$ from Oberfield and Raval (2021)) and more substitutable than Cobb-Douglas ($\alpha = 1.25$ from Karabarbounis and Neiman (2014)). In either case, the results are fairly similar to those from our baseline. Overall, these results underscore that the most important elasticity of adjustment on the part of firms in response to a \$15 minimum wage stems from within-education-group substitutability, as opposed to either across-education-group substitutability or the substitutability between capital and labor.

The next entry of Table 6 investigates how our results do not overly depend on the parameter γ_n for the elasticity of labor supply. We set this elasticity to a value of 0.5 as representative of the values in the "micro range" from Chetty et al. (2011), as opposed to our benchmark value of 1 from the "macro range", and again recalibrate our model given this lower elasticity. Our results are not sensitive to this change.

Finally, the bottom rows of Table 6 show that different values of the search frictions have minimal impact on our results. Since there is little external evidence on the magnitudes of these parameters, we simply perform a sensitivity analysis in which we increase their size by 50%.

[35]The findings in Bowlus et al. (2021) suggest an even greater substitutability between these workers than reported in previous studies. They document an elasticity of substitution between high-school- and college-educated workers that ranges between 3 and 8, depending on the years considered and the modelling of skill-biased technical change.

TABLE 6: Robustness of Long-Run Results to Alternate Parameter Values (Non-College Workers)

Parameterization	Δn_ℓ	Share($\Delta n_{\ell z} < 0$)	Share($\Delta w_{\ell z} n_{\ell z} < 0$)
Baseline	-11.7%	0.26	0.17
Monopsony power			
Markdown = 0.65 (Yeh, Macaluso and Hershbein, 2022)	-1.5%	0.15	0.07
Markdown = 0.85 (Lamadon, Mogstad and Setzler, 2022)	-19.3%	0.31	0.31
Production elasticities			
Cross-group $\rho = 4$ (Bils, Kaymak and Wu, 2020)	-11.0%	0.22	0.18
Within-group $\phi = 6$ (Card and Lemieux, 2001)	-14.1%	0.25	0.21
Capital-labor $\alpha = 0.7$ (Oberfield and Raval, 2021)	-10.0%	0.26	0.17
Capital-labor $\alpha = 1.25$ (Karabarbounis and Neiman, 2014)	-10.5%	0.22	0.18
Labor supply elasticity			
Labor supply elasticity $\gamma_n = 1/2$ (Chetty et al., 2011)	-13.0%	0.28	0.19
Search frictions (sensitivity to 50% above baseline)			
Vacancy posting costs κ_0	-12.4%	0.25	0.21
Matching function elasticity η	-11.9%	0.26	0.17
Job separation rate σ	-12.6%	0.26	0.21

Note: Change in steady state outcomes due to \$15 minimum wage; Δn_ℓ is the change in non-college employment, Share($\Delta n_{\ell z} < 0$) is the share of initial non-college workers who experience an employment decline, and Share($\Delta w_{\ell z} n_{\ell z} < 0$) is the share of initial non-college workers who experience an income decline. Rows correspond to different parameterizations of the model. "Baseline" refers to baseline calibration from Section 3. "Monopsony power" entries correspond to targeting different values of the markdown in the calibration (and simultaneously recalibrating all other fitted parameters as well). "Production elasticities" entries correspond to different values of the production elasticities (recalibrating all other fitted parameters as well). "Labor supply elasticity" corresponds to a different value of γ_n (recalibrating all other fitted parameters as well). "Search frictions" corresponds to setting the particular parameter to 50% above its baseline value, without recalibrating the other fitted parameters.

Additional Distributional Effects. Appendix C.5 uses our framework to explore the distributional impact of the minimum wage across different U.S. states or sectors of the U.S. economy, given that different states and sectors occupy different regions of the aggregate wage distributions in Figure 1. We show that states or sectors with lower average wages, such as the low-wage states of Mississippi or West Virginia or the low-wage sectors of Retail Trade and Personal Services, would be substantially worse off in the long run than high-wage states or sectors after an increase in the minimum wage to \$15. However, we emphasize that these results should only be viewed as illustrative because we do not formally model linkages or reallocation frictions across states or sectors.

5 Alternative Policies in the Tax and Transfer System

The previous section showed that although the minimum wage may benefit low-wage workers in the short run, it imposes significant costs on them in the long run. We now study how existing policies part of the U.S. tax and transfer system can better support these workers in the long run. Section 5.1 describes how we model the U.S. tax and transfer system and ensure that the alternative policies we consider are quantitatively comparable to the minimum wage. Section 5.2 examines the effects of the EITC, a particularly important component of the transfers to low-income households in the data, and Section 5.3 explores its interaction with the minimum wage. Section 5.4 evaluates the minimum wage in the context

of an approximation to the entire U.S. tax and transfer system.

5.1 Modeling Tax and Transfer Programs

Consider a general tax and transfer system $T(w_i)$, where $T(w_i)$ denotes labor income taxes so negative taxes $T(w_i) < 0$ indicate transfers. Let $A(w_i) = w_i - T(w_i)$ denote after-tax 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 for jobs in the labor market. Consider first how the system affects firms' labor demand, as summarized by the steady state vacancy posting condition

$$\frac{\kappa_i}{\lambda_f(\theta_i)} = \frac{1}{r+\sigma}\left[F_{ni} - w_i - \frac{1}{\omega}\frac{v'(n_i)}{A'(w_i)}\right], \tag{37}$$

where $A'(w_i)$ is marginal after-tax income. Equation (37) shows that a positive marginal tax rate, which implies $A'(w_i) < 1$, exacerbates monopsony distortions relative to our baseline model with $A'(w_i) = 1$, by increasing the magnitude of the last term in (37). To understand this mechanism, recall that a monopsony distortion arises because hiring a marginal worker increases the marginal disutility of work for all inframarginal workers, so a firm needs to compensate inframarginal workers with a higher wage thereby increasing its cost of hiring. A positive marginal tax rate reduces the after-tax wage that inframaginal workers receive and therefore increases the required before-tax wage that a firm must offer, further increasing firms' cost of hiring. Conversely, a negative marginal tax rate, namely, a tax credit which results in $A'(w_i) > 1$, alleviates monopsony distortions by reducing firms' costs of hiring and so bringing them closer to the planner's costs—that is, they reduce the last term on the right side of (37).

Whereas *marginal* tax rates affect monopsony distortions on labor demand, *average* tax rates influence households' labor supply by affecting their search decisions. The optimal search decision satisfies

$$h'(s_i) = \frac{\lambda_w(\theta_i)}{r+\sigma}\left[A(w_i) - v'(n_i)\right], \tag{38}$$

where $r = 1/\beta - 1$. A positive average tax rate, which implies $A(w_i) < w_i$, reduces the after-tax wage relative to our baseline model, thus depressing the incentive to search, as equation (38) shows. By contrast, negative average tax rates imply $A(w_i) > w_i$ and so increase the incentive to search. The marginal tax rate is irrelevant for households' search decisions because such decisions involve the extensive margin of whether or not to participate in the labor market. This discussion suggests that policies within the existing tax and transfer system, such as the EITC, that more directly target average and marginal wages may be more effective at alleviating monopsony distortions than the minimum wage.

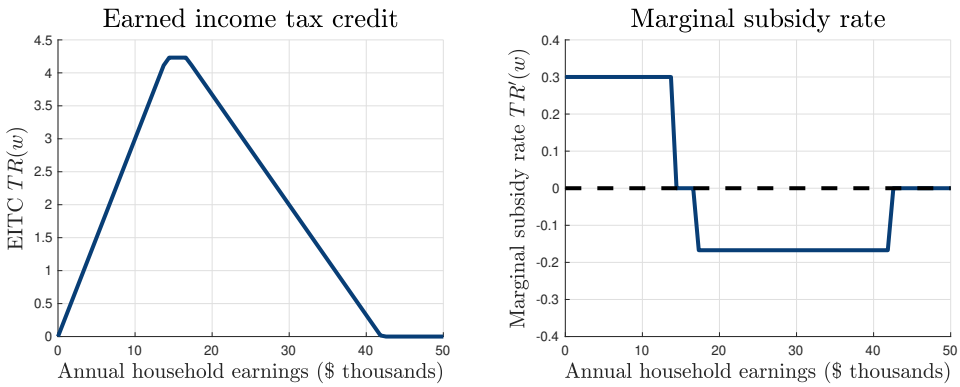
We ensure that the policies we consider are comparable to a \$15 minimum wage as follows. First, note that a \$15 minimum wage reduces firms' flow profits by an amount $\Delta\pi^*$, so we can think of the minimum wage as corresponding to an implicit tax on profits. For each of our alternative policies, we then assume that there is no minimum wage but that the government levies an explicit linear corporate income tax on

firms' profits of rate τ_f, which raises the same amount of revenues $\Delta\pi^*$. We assume that both investment and vacancy posting costs are fully deducted from corporate taxes, which implies that the tax τ_c does not distort any of the firms' marginal decisions, since it is a pure profit tax.[36] 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 denotes flow profits in the new equilibrium. Hence, each policy transfers on net $\Delta\pi^*$ from firms to households. Policies only differ in the schedule $T(w_i)$ that determines the distribution of these transfers across households.

5.2 Earned Income Tax Credit

The earned income tax credit (EITC) is one of the largest components of transfer payments in the United States. The empirical EITC schedule has several kinks due to the implied subsidy being phased in and out at different income levels. Specifically, the schedule entails three regions: (i) a *phase-in* region in which the subsidy (the tax credit) is proportional to household income, (ii) a *plateau* region in which the subsidy is capped at its maximum benefit, and (iii) a *phase out* region in which the subsidy is phased out.

FIGURE 16: Earned Income Tax Schedule



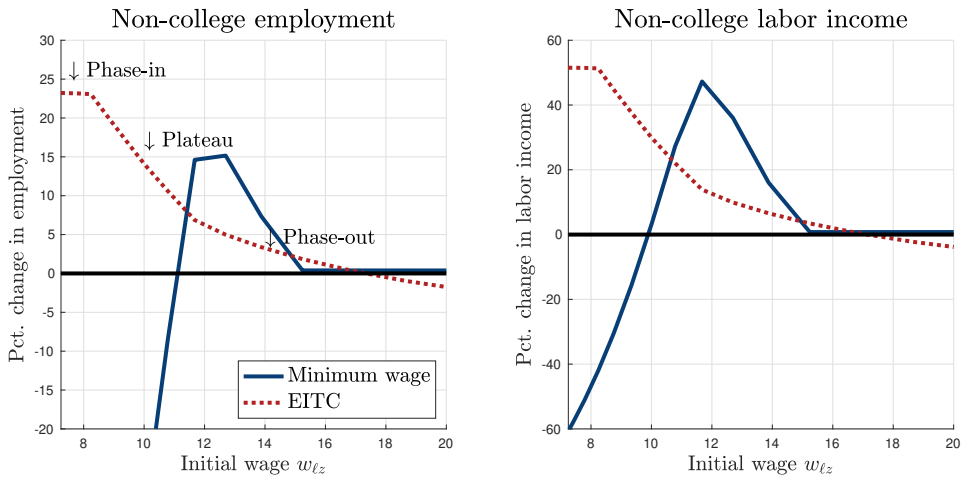
Note: The left panel plots an EITC schedule that is budget-equivalent to a $15 minimum wage. The right panel plots the implied marginal tax rate. In each panel, the x-axis rescales steady state labor income to annual earnings assuming each household works 1800 hours per year.

We choose our transfer system $T(w_i)$ to mimic this empirical schedule, but adjust magnitudes to ensure that it is budget-equivalent to a $15 minimum wage. Figure 16 plots our budget-equivalent schedule. We set the phase-in rate to 30%, as in the data. In the phase-in region, households face both a positive average subsidy rate (since the total tax credit is positive) and a positive marginal subsidy rate (since the credit is being phased in). In the plateau region, households face a zero marginal subsidy rate but still a positive average subsidy rate, since they are still receiving the benefit. Finally, we set the phase-out rate to approximately 18%. In the phase-out region, households faces a positive marginal tax rate, that is, a

[36] Our pure profit tax is nondistortionary because our model features a fixed number of active firms. If, instead, we allowed for free entry of firms, then reducing steady-state profits may reduce entry and therefore the total number of firms. Our profit tax ensures that the entry margin would be distorted by the same amount across all alternative policies we consider.

negative marginal subsidy rate, because each dollar of earnings reduces their transfer payments.

FIGURE 17: Effect of Earned Income Tax Credit on Employment

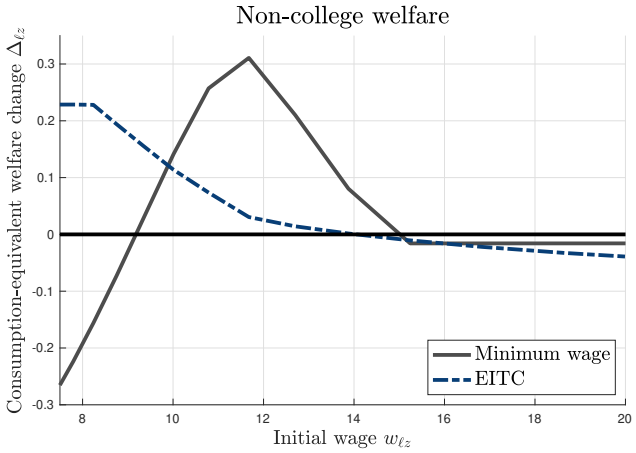


Note: Steady-state employment (left panel) and labor income (right panel) of selected z-types under the EITC. The y-axis is normalized relative to employment in the initial equilibrium (without any policies). The x-axis is the wage $w_{\ell z}$ of a z-type worker in the initial equilibrium without either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

Figure 17 shows how the EITC affects employment and income in the long run for non-college workers. The lowest-wage workers are in the phase-in region and so enjoy both the positive effect of the positive average subsidy on their marginal benefit of searching for jobs and of the positive marginal subsidy on firm's marginal benefit of posting vacancies. Workers with somewhat higher wages are in the plateau region, where the effect of the EITC is still positive but decreasing in initial wages because the credit becomes a smaller share of total wages. Importantly, in both of these regions, the EITC increases both employment and labor income, thereby helping exactly the low-wage workers whom the minimum wage hurts. Finally, workers in the phase-out region experience an even smaller increase, or perhaps even a decrease, in their employment, as the policy implies a positive marginal tax rate for them, as discussed.

Figure 18 compares the steady-state welfare effects Δ_i of a \$15 minimum wage and of the budget-equivalent EITC. The EITC dominates the minimum wage for low-wage workers, mirroring its effect on labor income described above. We show in Section 5.4 that a progressive tax and transfer system is even superior to the ETIC because it can finance larger transfers by raising taxes on high-income workers. In contrast, a higher minimum wage benefits middle-income households, who were earning wages close to the new minimum in the initial equilibrium. In this sense, a \$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 ones. By contrast, the EITC program and more generally the tax and transfer system both improve the welfare of low-income households at the expense of high-income ones.

FIGURE 18: Effects of Alternative Policies on Welfare



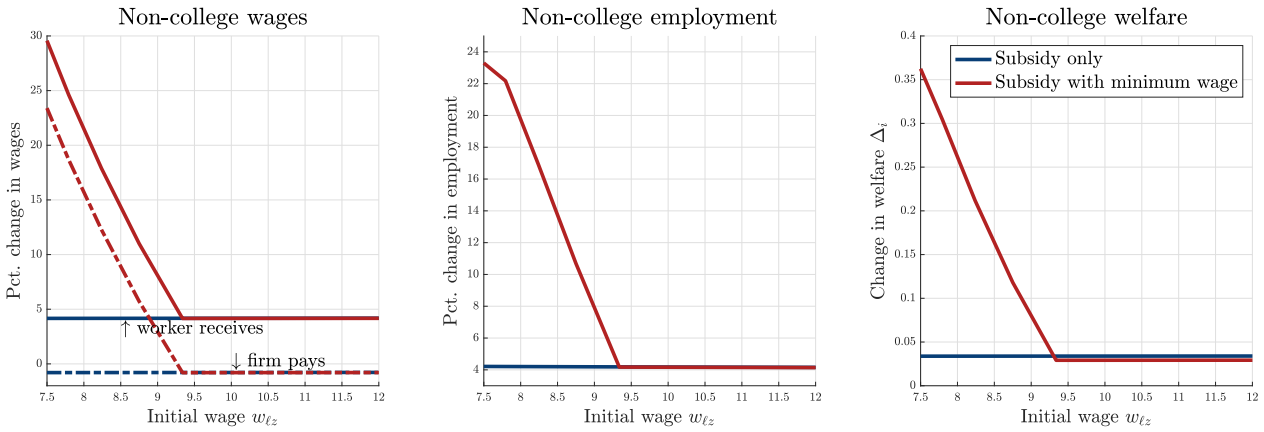
Note: Steady-state welfare change Δ_i from $u\left((1+\Delta_i)c_i^* - v(n_i^*) - h(s_i^*)\right) = u\left(\widetilde{c}_i - v(\widetilde{n}_i) - h(\widetilde{s}_i)\right)$, where the superscript $*$ denotes steady-state values in the initial equilibrium without the policy and the superscript $\widetilde{}$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

5.3 Interaction Between Transfer Policies and the Minimum Wage

A long-standing issue with the EITC and similar programs is that they partly benefit firms because firms are able to pay lower pre-transfer wages, given the increase in labor supply, and so appropriate part of the benefits. This issue has led some authors to suggest that a modest minimum wage may complement transfer programs because it prevents firms from reducing the wages they pay (see Neumark and Wascher (2011), Lee and Saez (2012), and Vergara (2022), for example.) Here, we reassess this argument in the context of our model. We confirm the intuition that a modest minimum wage is helpful in preventing firms from lowering pre-transfer wages, but also show its additional benefit in that it locally forces firms to pay workers more than the monopsony wage, thus stimulating employment. In order to illustrate this mechanism, we build intuition by first considering a constant transfer equal to 5% of households' labor income, that is, $T(w_i) = -0.05w_i$, which is identical to the phase-in region of an EITC with a phase-in rate of 5%. We then show that these results generalize to the full EITC schedule as well.

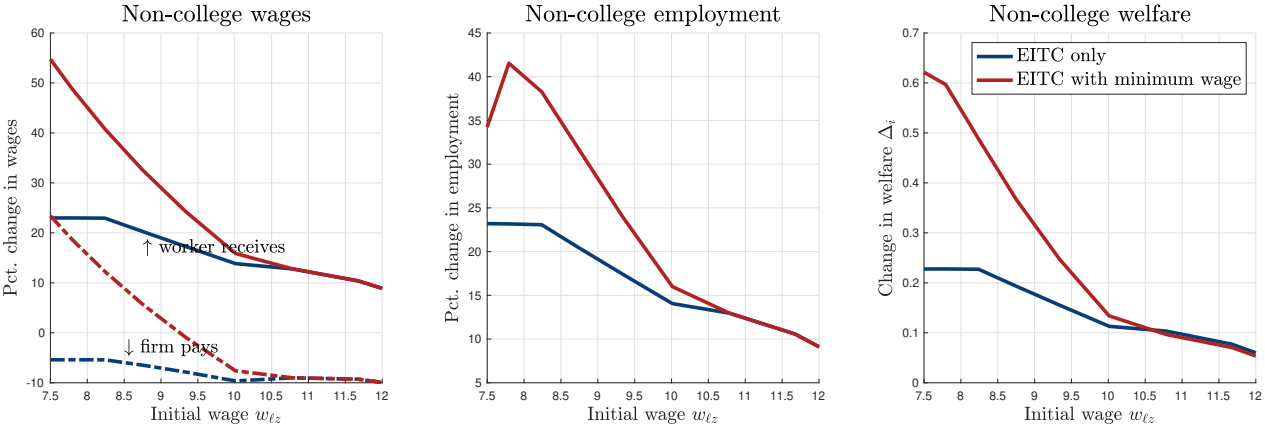
The blue lines in Figure 19 illustrate how this constant 5% transfer partly benefits firms. Specifically, the left panel shows that firms are able to pay their workers 1% less due to the transfer so that the post-transfer wages that households receive $(1+0.05)w_i$ only increase by 4% rather than 5%. As a result, employment only increases by 4%, reducing the overall increase in labor income that would have resulted if wages did not fall. The red lines in Figure 19 show that introducing a modest minimum wage of $\underline{w} = \$9.25$ leads to much larger increase in employment and labor income than the EITC program alone does. One reason is that this minimum wage prevents firms from lowering wages in response to the transfer to low-income workers. But the much more quantitatively important reason is that the minimum wage alleviates

FIGURE 19: Effect of Constant Transfer with and without $\underline{w} = \$9.25$ Minimum Wage



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

FIGURE 20: Effect of EITC with and without $\underline{w} = \$9.25$ Minimum Wage



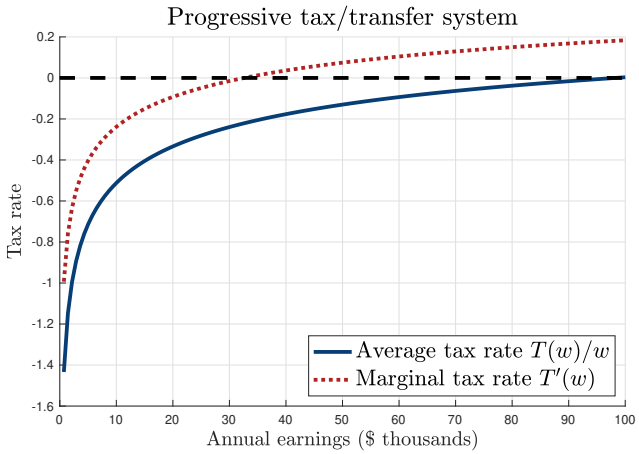
Note: Steady-state wages (left panel), employment (middle panel), and labor income (right panel) in response to an EITC system that is budget-equivalent to a \$15 minimum wage. The y-axis is the log-change relative to the initial equilibrium without any policies. The x-axis is the wage $w_{\ell z}$ of a z-type worker in the initial equilibrium without either policy.

the monopsony distortions these workers face.[37] Figure 20 shows that these insights extend to our budget-equivalent EITC schedule from Section 5. Although the effects of the EITC are more nonlinear due to the kinks in the EITC schedule, a \$9.25 minimum wage still substantially amplifies 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. The size of the minimum wage, however, must be chosen carefully. If the minimum

[37] Appendix D nets out the monopsony reduction in order to focus on the pure interaction between the transfer and the 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.

FIGURE 21: Progressive Tax and Transfer System



Note: 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 $\tau = 0.181$). In each panel, the x-axis rescales steady-state labor income to annual earnings assuming each household works 1800 hours per year.

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 D shows that a $12 minimum wage substantially *reduces* the benefit of the EITC for low-income non-college workers, since it significantly reduces their employment.

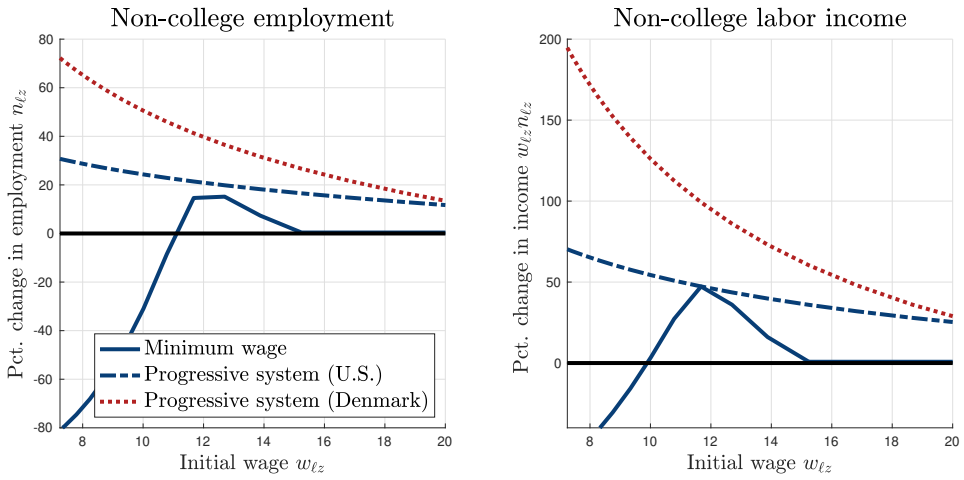
5.4 Progressive Tax and Transfer System and Minimum Wage

As discussed, we model the entire U.S. progressive tax and transfer system using the parametric tax function $T(w_i) = w_i - \lambda w_i^{1-\tau}$, where $T(w_i)$ is the labor income tax schedule, λ and τ are parameters that govern the level and progressivity of the system, and negative taxes $T(w_i) < 0$ indicate transfers. Heathcote, Storesletten and Violante (2017) show that $\tau = 0.181$ provides a good description of the overall progressivity of the U.S. system, except at the very bottom of the income distribution because transfers are phased in and out at various income levels, which introduces kinks in the implied tax rates.[38] We follow Heathcote, Storesletten and Violante (2017) in excluding the EITC and setting the progressivity parameter τ at 0.181. We then choose the scale parameter λ to ensure that the aggregate net transfer payment is budget-equivalent to a $15 minimum wage.

Figure 21 plots the average tax rate $T(w_i)/w_i$ and marginal tax rate $T'(w_i)$ of our budget-equivalent system as a function of labor income. The fact that the system is progressive ($\tau > 0$) implies that marginal tax rates are higher than average tax rates throughout. The lowest-income households face both a negative average tax rate, so they receive transfers ($A(w_i) > w_i$), and a negative marginal tax rate, so transfers are

[38] See page 1700 of their paper: "[O]ur 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 22: Effect of Progressive Tax System on Non-College Workers



Note: Steady-state employment (left panel) and labor income (right panel) of selected z types for three different policies: a \$15 minimum wage (blue line), the budget-equivalent progressive tax system with the U.S. level of $\tau = 0.181$ (red line), and the budget-equivalent tax system with the Danish level of $\tau = 0.463$ (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.

phased in ($A'(w_i) > 1$). For these households, the tax and transfer system reduces the monopsony distortion by (37) and increases search effort by (38). Middle-income households continue to receive transfers, which encourages their search effort, but face positive marginal tax rates, which exacerbates monopsony distortions. Finally, high-income households face both positive average and marginal tax rates, which reduces their search effort and exacerbates monopsony distortions. In this sense, a progressive tax and transfer system differentially affects the monopsony distortions faced by different workers.

The left panel of Figure 22 shows that this progressive system succeeds in substantially increasing the employment of non-college workers, especially for the low-wage workers whom the system especially targets. The right panel shows that similar results hold for after-tax labor income. Labor income increases for all the z-type in the figure, although it falls for some of the higher-z types not pictured who earn higher wages, especially among college workers. For sake of comparison, Figure 22 also shows the effect of an even more progressive tax and transfer system with $\tau = 0.463$ meant to capture the degree of progressivity in Denmark.[39] This more progressive system further increases the employment of low-z types but decreases that of college workers more than the U.S. system does (not pictured), which underscores the policy trade-offs associated with different levels of progressivity of a tax and transfer system.

[39] 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, excluding transfers. We impute a value of the progressivity of the tax and transfer system by scaling our baseline progressivity parameter for the U.S. tax and transfer system by the ratio of the progressivity of the Danish tax system to the U.S. one.

6 Conclusion

Existing empirical work on the minimum wage in the United States, which focuses on the one- to two-year impact on employment of increases in the minimum wage, finds very small elasticities of response for employment. Hence, this work suggests that, at least in the short run, the minimum wage successfully boosts the labor income of low-wage workers. By contrast, estimates of the long-run substitutability among workers within and across education groups, as in the classic studies of Katz and Murphy (1992) and Card and Lemieux (2001), imply large elasticities of substitution.

We have developed 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 currently proposed. The putty-clay framework we propose reproduces both these small short-run elasticities of substitution and the large long-run ones.

We find that in the short run, the impact of even a large increase in the minimum wage on employment is small so that it works as its proponents hope: it leads to a sizable increase in the labor income and welfare of workers at the low end of the wage distribution. Over time, however, firms substitute away from low-productivity workers with low education to high-productivity workers with low education as implied by the estimates of Card and Lemieux (2001). Hence, in the long run, a high minimum wage has perverse distributional implications in that it reduces the employment, income, and welfare of precisely the lowest-income workers it is design to support.

Our quantitative analysis shows that a better way to increase the welfare of low-productivity workers is through a progressive tax and transfer scheme, such as the EITC, that induces them to work more, as opposed to a large increase in the minimum wage that effectively prices them out of the labor market by making them too expensive for firms to hire. Nonetheless, we find that a modest minimum wage complements the EITC policy and performs much better than either a minimum wage policy on its own or an EITC policy on its own. In this sense, a modest minimum wage is a valuable tool that enhances the efficacy of many progressive tax and transfer policies.

In terms of fruitful extensions, note that our current framework holds the distribution of worker skills fixed in response to changes in the minimum wage. However, as we have shown, a large and permanent change in the minimum wage will induce permanent changes in wages and job-finding rates, which in turn affects the incentives for workers to acquire new skills. On the one hand, since the minimum wage equalizes wages for all workers at the low end of the skill distribution, the minimum wage may disincentivize these workers from investing in skills. On the other hand, since these workers also face lower job-finding rates, they may have an incentive to acquire more skills in order to increase their probability of finding a job. Hence, the role of endogenous skill acquisition is an interesting quantitative question that requires carefully modeling and pinning down how the skill acquisition process responds to these incentives. Such a model would have to confront the fact that workers who may in principle benefit the most from accumulating

more skills—those at the bottom of the skill distribution—may also face in practice the lowest returns to doing so. In fact, much work has found that low-wage workers tend to have lower levels of human capital as they face higher monetary or opportunity costs to acquiring skills (see, for instance, the papers reviewed by French and Taber (2011)). Such an analysis would be an ambitious and worthwhile endeavor not just for studying the minimum wage but also many other labor market policies. We think of these questions as an interesting avenue of future research.

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*, 212(110288), 1–8.

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.

Andolfatto, D. (1996): "Business Cycles and Labor-Market Search," *American Economic Review*, 86(1), 112–132.

Atkeson, A., and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.

Atkeson, A., and P. J. Kehoe (1999): "Models of Energy Use: Putty-putty versus Putty-clay," *American Economic Review*, 89(4), 1028–1043.

Berger, D., K. Herkenhoff, and S. Mongey (2022a): "Labor Market Power," *American Economic Review*, 112(4), 1147–93.

Berger, D. W., K. F. Herkenhoff, and S. Mongey (2022b): "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 Working Paper*.

Bowlus, A., L. Lochner, C. Robinson, and E. Suleymanoglu (2021): "Wages, Skills, and Skill-Biased Technical Change: The Canonical Model Revisited," *Unpublished Working Paper*.

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.

Calvo, G. A. (1976): "Optimal Growth in a Putty-Clay Model," *Econometrica: Journal of the Econometric Society*, 44(5), 867–878.

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.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 134(3), 1405–1454.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011): "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review*, 101(3), 471–75.

Clemens, J., L. B. Kahn, and J. Meer (2021): "Dropouts Need Not Apply? The Minimum Wage and Skill Upgrading," *Journal of Labor Economics*, 39(S1), S107–S149.

Clemens, J., and M. R. Strain (2021): "The Heterogeneous Effects of Large and Small Minimum Wage Changes: Evidence over the Short and Medium Run Using a Pre-Analysis Plan," Discussion paper, National Bureau of Economic Research.

Curtis, E. M., D. G. Garett, E. Ohrn, K. A. Roberts, and J. C. S. Serrato (2021): "Capital Investment and Labor Demand: Evidence from 21st Century Tax Policy," *Unpublished Working Paper*.

Deb, S., A. Eeckhout, Jan Patel, and L. Warren (2021): "The Contribution of Market Power To Wage Inequality," *Unpublished Working Paper*.

Dube, A. (2019): "Minimum Wages and the Distribution of Family Incomes," *American Economic Journal: Applied Economics*, 11(4), 268–304.

Dube, A., and A. Lindner (2021): "City Limits: What Do Local-Area Minimum Wages Do?," *The Journal of Economic Perspective*, 35(1), 27–50.

Dustmann, C., A. Lindner, U. Schönberg, M. Umkehrer, and P. vom Berge (2022): "Reallocation Effects of the Minimum Wage," *The Quarterly Journal of Economics*, 137(1), 267—328.

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

Engbom, N., and C. Moser (2021): "Earnings Inequality and the Minimum Wage: Evidence from Brazil," Discussion paper, National Bureau of Economic Research.

Flinn, C. J. (2006): "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates," *Econometrica*, 74(4), 1013–1062.

French, E., and C. Taber (2011): "Identification of Models of the Labor Market," *Handbook of Labor Economics*, 4, 537–617.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988): "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78(3), 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.

Jardim, E., M. C. Long, R. Plotnick, E. van Inwegen, J. Vigdor, and H. Wething (2022): "Minimum-Wage Increases and Low-Wage Employment: Evidence from Seattle," *American Economic Journal: Economic Policy*, 14(2), 263–314.

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.

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.

Lamadon, T., M. Mogstad, and B. Setzler (2022): "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market," *American Economic Review*, 112(1), 169–212.

Lee, D., and E. Saez (2012): "Optimal Minimum Wage Policy in Competitive Labor Markets," *Journal of Public Economics*, 96(9-10), 739–749.

Lindner, A., and P. Harasztosi (2019): "Who Pays for the Minimum Wage?," *American Economic Review*, 109(8), 2693–2727.

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

Meer, J., and J. West (2016): "Effects of the Minimum Wage on Employment Dynamics," *The Journal of Human Resources*, 51(2), 500–522.

Merz, M. (1995): "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, 36(2), 269–300.

Neumark, D., J. M. I. Salas, and W. Wascher (2014): "Revisiting the Minimum Wage–Employment Debate: Throwing Out the Baby With the Bathwater?," *Industrial and Labor Relations Review*, 67(S), 608–648.

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?," Discussion paper, National Bureau of Economic Research.

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

Neumark, D., W. L. Wascher, W. L. Wascher, et al. (2008): "Minimum Wages," MIT press.

Oberfield, E., and D. Raval (2021): "Micro Data and Macro Technology," *Econometrica*, 89(2), 703–732.

Robinson, J. (1933): "The Economics of Imperfect Competition," St. Martin's Press.

Seegmiller, B. (2021): "Valuing Labor Market Power: The Role of Productivity Advantages," Discussion paper, MIT Sloan Working Paper.

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.

Vergara, D. (2022): "Minimum Wages and Optimal Redistribution in the Presence of Taxes and Transfers," *arXiv preprint arXiv:2202.00839*.

Yeh, C., C. Macaluso, and B. Hershbein (2022): "Monopsony in the U.S. Labor Market," *American Economic Review*, 112(7), 2099–2138.

A Omitted Model Details

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 $\sum_j p_j c_j \leq pc$ with multiplier λ, 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

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j = \max_{j'}\left\{u'(c)\left(\frac{c_{j'}}{c}\right)^{-\frac{1}{\omega}} - \lambda p_{j'}\right\}. \tag{39}$$

In a symmetric allocation with $c_{j'} = c$ and $p_{j'} = p$, (39) reduces to

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j \geq u'(c) - \lambda p, \tag{40}$$

which is the participation constraint under monopolistic competition. Notice the similarity of this participation constraint for attracting a consumer to buy 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 is that 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 one.

B Data and Empirical Literature

This appendix contains details about our data sources and construction referenced throughout the main text as well as additional details on the empirical literature that has estimated the employment effects of the minimum wage. We use data from the pooled 2017-2019 American Community Survey (ACS). All observations are weighted using the weights provided by the ACS.

Share of College Workers. 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 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.

Estimates of the Employment Effects of the Minimum Wage. The recent survey by Neumark and Shirley (2021), which we reference in the main text, takes stock of the large body of work on the employment effects of the minimum wage in the United States and illustrates how, despite a lack of agreement in the literature on how to interpret the existing evidence, most studies report negative employment effects of minimum wage increases. Specifically, the authors argue that negative estimates are pervasive in the literature, that this evidence is stronger for teens, young adults, and less-educated workers, that the evidence from studies of directly-affected workers points even more strongly to negative effects, and that the evidence from studies of low-wage industries is less clear cut. These authors focus on the evidence that has relied on geographical (subnational) variation in the minimum wage within the United States assembling the entire set of published studies and identify the core estimates (more than a hundred) supporting the conclusions from each study. According to their preferred estimate from each study, a 1% increase in the minimum wage reduces employment by 0.151%, based on the mean such estimate across studies, or by 0.116%, based on the median such estimate across studies——estimates range from -1% to 1.7% (see their Figure 1). Instead, according to the median estimates from each study, the authors calculate that a 1% increase in minimum wage reduces employment by 0.133%, based on the mean such estimate across studies, or by 0.110%, based on the median such estimate across studies——estimates range from -1% to 1.7% (see their Figure 3).[40] Positive effects of the minimum wage on employment have been detected for

[40]Numerous papers have pointed out the challenges of measuring the effects of the minimum wage, which include isolating reliable identifying information, constructing accurate control groups, controlling for confounding factors, and accounting for potential issues of endogeneity in controls and selection in the sample examined, as well as the lack of robustness of results across alternative estimated specifications of the effects of interest. For instance, Neumark, Salas and Wascher (2014) revisit studies that claim that panel-data estimates commonly used are flawed, as they fail to account for spatial heterogeneity, and doing so supports the notion that minimum wages in the United States have not reduced employment. The authors' results, though, confirm the evidence of negative employment effects and so the trade-off at the heart of minimum wage policies between higher wages for some workers and job losses for others, which our distributional analysis highlights.

other countries, though. Dustmann et al. (2022), for example, study the impact of the introduction of a nationwide minimum wage in Germany that affected 15% of all employees and find that the minimum wage raised wages, did not not lower employment, and led to a reallocation of low-wage workers from smaller to larger, from lower- to higher-paying, and from less to more productive establishments.

We note that most of the empirical literature evaluating the minimum wage exploits local changes in it, which are quite small relative to the current proposed changes in the national minimum wage, and examines short-run employment responses among narrow groups of workers. Cengiz et al. (2019), for instance, which is one of the most recent papers reviewed by Neumark and Shirley (2021), rely on state-level minimum wage changes to examine the extent to which the minimum wage affects the employment of workers who were initially below the new minimum and find that the employment of low-wage workers declines only slightly in the few years following a minimum wage increases.[41] Similar impacts have been documented in other studies; when positive employment effects are detected, they are relatively small and concentrated among high-wage workers. See Dube and Lindner (2021) for analogous results.

Our results are consistent with this evidence for two reasons. First, we find that small changes in the minimum wage have only small effects on the employment rates of all workers, including low-productivity ones. Second, and more importantly, we find that even large changes have only small short-run employment effects. In line with these results, by exploiting state-level changes in the minimum wage between 1975 and 2012, Meer and West (2016) estimate that an increase in the minimum wage reduces employment but such an effect takes several years to materialize. When the authors regress annual changes in (log) employment of all non-agricultural employees on changes in (log) minimum wage, they find that an increase in the minimum wage has no contemporaneous effect on employment but it has a negative effect after one year. Specifically, a 1% increase in the minimum wage reduces employment by 0.03% after one year, by 0.06% after two years, and by 0.07% after three years (see their Table 4). The authors also regress long differences (from one to eight years) in (log) employment on long differences in (log) minimum wage and report that an increase in the minimum wage has no contemporaneous effect on employment but it has negative effects after one year. In particular, a 1% increase in the minimum wage reduces employment by 0.0387% after one year and by approximately 0.05% after two years—the effect is constant thereafter.[42]

As we highlight, allowing for reasonable frictions to input adjustment in the form of a putty-technology implies that employment effects are small in the first few years after a minimum wage increase even though long-run responses are large. These implications of our model are consistent with the evidence in Lindner and Harasztosi (2019). In response to a large and persistent minimum wage increase in Hungary, these authors find that employment elasticities are negative but small even four years after the reform and that firms responded to the minimum wage by substituting labor with capital.[43]

C Additional Results About the Minimum Wage

This appendix contains a number of additional results about the minimum wage.

[41] Clemens, Kahn and Meer (2021) also exploit cross-regional variation to document that firms switch away from low-productivity workers towards higher-productivity workers in response to minimum wage increases. This finding is consistent with the key adjustment mechanism in our model.

[42] Jardim et al. (2022) find similar non-linear effects of city-level changes in the minimum wage on hours worked. Namely, an increase in the minimum wage in Seattle from $9.5 to $11 did not impact hours worked but a further increase to $13 reduced hours worked and employment for low-wage workers, that is, workers earning less than $19 per hour across all industries. In particular, the authors study two successive increments in the minimum wage in the city between 2015 and 2016 and find that the first increase from $9.5 to $11 is associated with an increase in average wages by 1.7% and no statistically significant effect on hours worked (see their Table 6). They then document that the second increase from $11 to $13 is associated with an increase in average wages by 3.2% (see their Table 5) and a reduction in hours worked by 6.9% and in the number of jobs by 5.9% (see their Table 6).

[43] Thee authors propose a model of monopolistic competition in the output market to address these findings as well as the fact they document that negative employment effects are greater in industries where the pass-through of wage costs to output is likely lower, like in the tradable, manufacturing, and exporting sectors.

TABLE C.1: Targeted Statistics, Baseline vs. Capital-Skill Complementarity

Moment	Description	Data	Baseline Model	KORV Model
Average wage markdown				
$\mathbb{E}[w_{ni}]/\mathbb{E}[F_{ni}]$	Average wage markdown	0.75	0.75	0.75
Wage Distribution, ACS 2017-2019				
$w_{\ell 50}/w_{\ell 10}$	Non-college 50-10 ratio	2.04	2.00	1.90
w_{h50}/w_{h10}	College 50-10 ratio	2.30	2.06	2.06
Income shares				
$\mathbb{E}[w_i n_i]/Y$	Aggregate labor share	0.57	0.57	0.57
$\pi_h \mathbb{E}[w_{hz} n_{hz}]/\mathbb{E}[w_i n_i]$	College income share	0.55	0.56	0.55
Unemployment rate				
$\mathbb{E}[s_i]/(\mathbb{E}[s_i]+\mathbb{E}[n_i])$	Average unemployment rate	5.9%	5.9%	5.9%
Employment Rates				
$\mathbb{E}_\ell[n_i]$	Non-college employment rate	0.47	0.47	0.47
$\mathbb{E}_h[n_i]$	College employment rate	0.62	0.61	0.61

Note: Statistics targeted using parameters in Table 3. "Baseline model" corresponds to the model in the main text. "KORV model" refers to model with capital-skill complementarity described in the appendix text. The average wage markdown is the midpoint of the range of estimated markdowns discussed in the main text. The average labor share is from Karabarbounis and Neiman (2014).

C.1 Results with Capital-Skill Complementarity

In the main text, we assumed that capital is equally substitutable with college and non-college workers (see equation (2)). In very influential work, Krusell et al. (2000) argue that, instead, capital is more substitutable with non-college workers than it is with college workers, a configuration they refer to as "capital-skill complementarity." In this appendix, we show that our main results are robust to using this alternative specification. This robustness occurs because the distributional impact of the minimum wage is primarily determined by the substitutability of workers within an education group, not the substitutability of the two education groups with capital.

This alternative version of the model replaces the long-run production structure (2) and (3) with

$$F(k_{jt}, \bar{n}_{\ell jt}, \bar{n}_{hjt}) = \left[\psi(\bar{n}_{\ell jt})^{\frac{\rho-1}{\rho}} + (1-\psi)G(k_{jt}, \bar{n}_{hjt})^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{41}$$

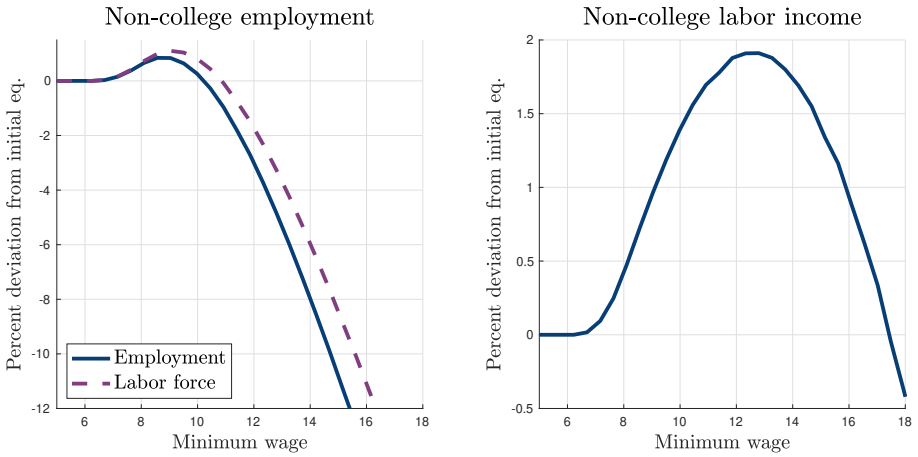
$$G(k_{jt}, \bar{n}_{hjt}) = \left[\lambda\,(k_{jt})^{\frac{\alpha-1}{\alpha}} + (1-\lambda)(\bar{n}_{hjt})^{\frac{\alpha-1}{\alpha}}\right]^{\frac{\alpha}{\alpha-1}}. \tag{42}$$

In this alternative formulation, the parameter α is the elasticity of substitution between capital and college labor and the parameter ρ is the elasticity of substitution between non-college labor and the capital-college labor bundle $G(k_{jt}, \bar{n}_{hjt})$. Krusell et al. (2000) estimate that $\rho > \alpha$, which implies that non-college labor is more substitutable with capital than is college labor. The aggregated labor $\bar{n}_{\ell jt}$ and $\bar{n}_{\ell jt}$ inputs depend on type-level productivity z_i with the within-group elasticity ϕ as in the main text.

We recalibrate the model with this alternative production structure to target the same statistics as the baseline model in the main text. Table C.1 shows that the alternative model matches the targets about as well as our baseline model. While the parameter values are also similar to the baseline model, we make one important change: we assume that capital depreciates at $\delta = 15\%$ annually rather than $\delta = 10\%$ annually as in the main text. We make this change in order to exclude structures from the specification of capital-skill complementarity, following Krusell et al. (2000).

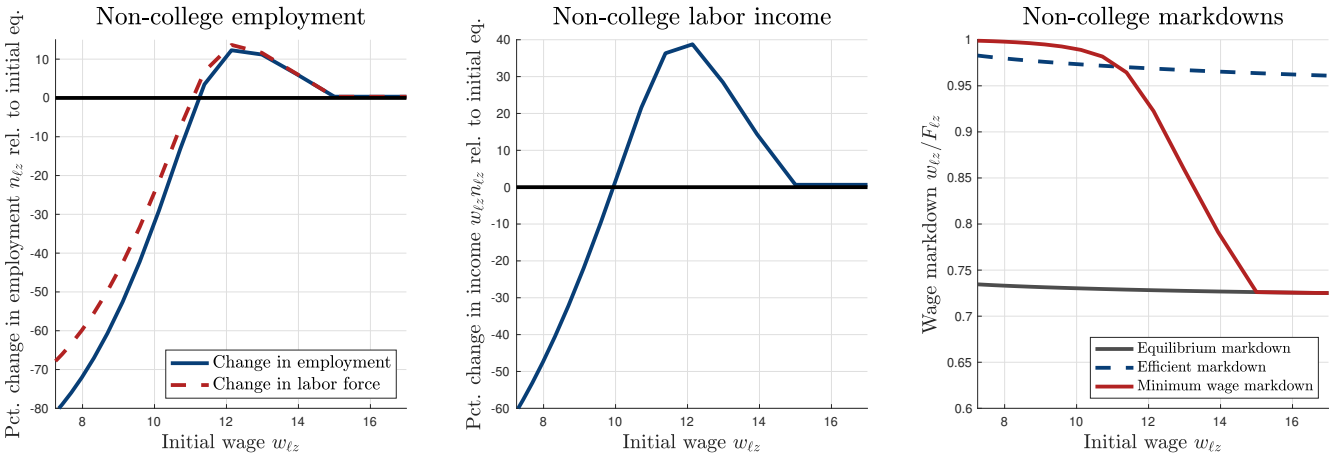
The long-run effects of the minimum wage in this alternative model are extremely similar to the baseline model in the main text. Figure C.1 plot the employment and labor income Laffer Curves for non-college workers; they are nearly identical to Figure 4 in the main text. For example, the \$15 minimum wage decreases non-college employment by about 12% in both cases. Figure C.2 shows that the distributional impact of the \$15 minimum wage across individual non-college types is also nearly identical

FIGURE C.1: Long-Run Aggregate Minimum Wage Laffer Curves w/ Capital-Skill Complementarity



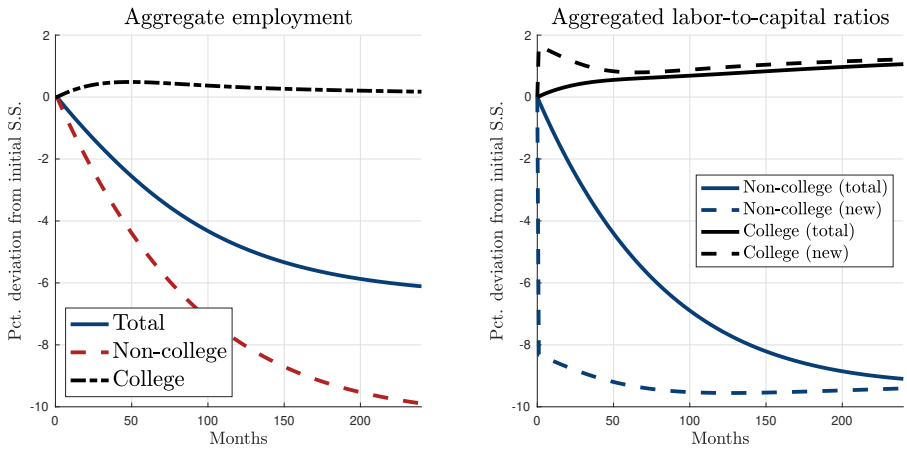
Note: Steady-state outcomes as a function of the minimum wage $\underline{w}$ in the model with capital-skill complementarity. The left panel plots the log-change of aggregate non-college employment, the log-change of aggregate non-college labor force, and the change in the unemployment rate relative to their levels in the initial steady state with $\underline{w} = 0$ (see the decomposition (33)). The right panel plots aggregate labor income of non-college workers. The x-axis is the level of the minimum wage $\underline{w}$ that relative to the median non-college wage in the initial equilibrium is the same as in the data.

FIGURE C.2: Distributional Effects of $15 Minimum Wage with Capital-Skill Complementarity



Note: Steady-state outcomes for selected z-types among non-college workers for a $15 minimum wage in the model with capital-skill complementarity. The left panel plots the percentage change in employment n_i relative to the initial steady state without the minimum wage, the middle panel plots the percentage change in labor income $w_i n_i$ relative to the initial steady state without the minimum wage, and the right panel plots the levels of the markdowns w_i/F_{ni} for three different parameterizations: (i) $\underline{w} = 0$ ("equilibrium markdown"), (ii) $\underline{w} = 0$ and $\omega \to \infty$ ("efficient markdown"), and (iii) $\underline{w} = \$15$ ("minimum wage markdown"). The x-axis corresponds to the initial wage $w_{\ell z}$ of a type-z worker in the initial equilibrium.

FIGURE C.3: Employment Dynamics Along the Transition Path with Capital-Skill Complementarity



Note: Transition paths following an unexpected imposition of the minimum wage $\underline{w}$, starting from the initial equilibrium with $\underline{w} = 0$, in the model with capital-skill complementarity. The left panel plots the aggregated employment of non-college workers, college workers, and total employment. The right panel plots the associated labor-to-capital ratios.

to its counterpart, Figure 7, in the main text.

Figure C.3 shows that the transition path in this alternative model is extremely similar to the baseline model in the main text as well. As in our baseline model, it takes employment more than twenty year to converge to its new steady state value due to the putty-clay technology. While the transition dynamics are somewhat faster in the alternative model than in the main text, that difference is due to the higher depreciation rate rather than the alternative production structure (recall that we are using $\delta = 15\%$ annually in this alternative model). As described in the main text, the depreciation rate controls the speed of transition because it determines how quickly the old, labor-intensive capital exits production.

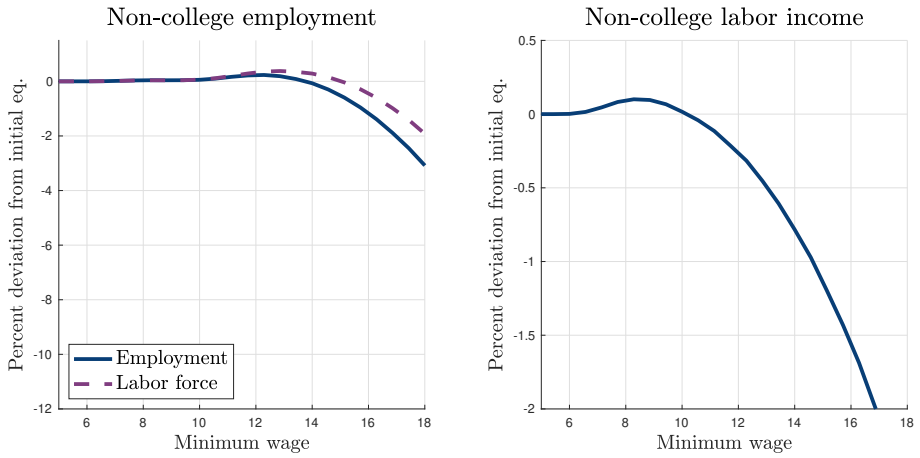
C.2 Long-Run Effects of the Minimum Wage on College Workers

The analysis in the main text focuses on non-college workers because the minimum wage primarily binds among those workers; relatively few college-educated workers earn below a $15 minimum wage for example. However, for the sake of completeness, Figure C.4, plots the aggregated employment and income Laffer curves for college workers. The left panel shows that college employment has the similar shape, but that it only starts to decline for higher levels of the minimum wage than for the non-college workers. In contrast, the right panel shows that labor income declines by more for college workers than non-college workers. This occurs because the lower non-college employment and lower capital stock reduce the marginal product, and therefore the wages, of the college workers.

C.3 Distribution of Capital Types Over Time

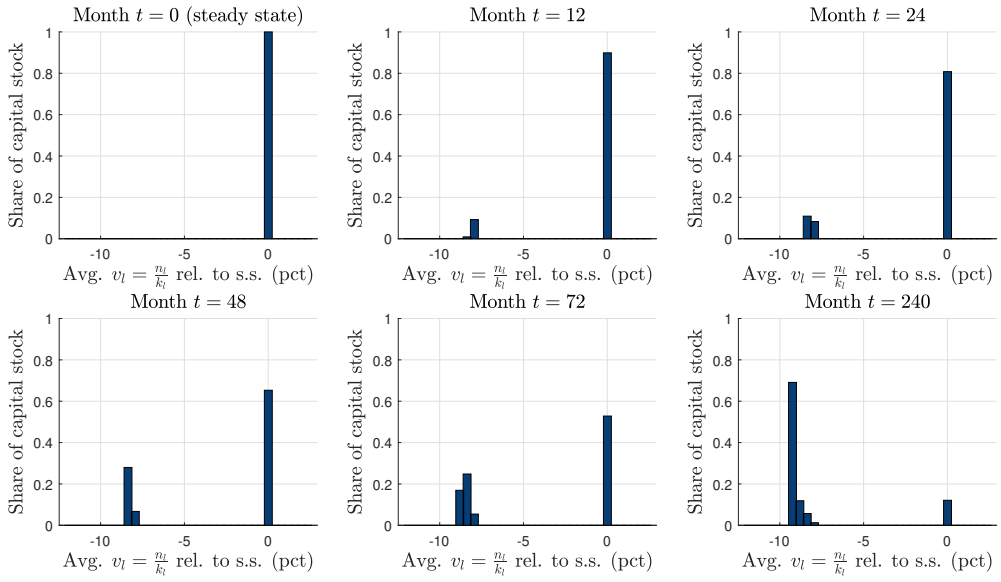
Figure C.5 plots the distribution of capital types along the transition path in response to the permanent $15 minimum wage. Before the introduction of the minimum wage, firms hold only one type of capital, namely the type that is optimal at the original steady state prices. The minimum wage induces firms 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. The employment dynamics are so prolonged because it takes time for the old capital stock to depreciate away; the bottom right panel shows that, even twenty years after the introduction of the minimum wage, more than 10% of the initial stock still exists.

FIGURE C.4: Long-Run Aggregate Minimum Wage Laffer Curves for College Workers



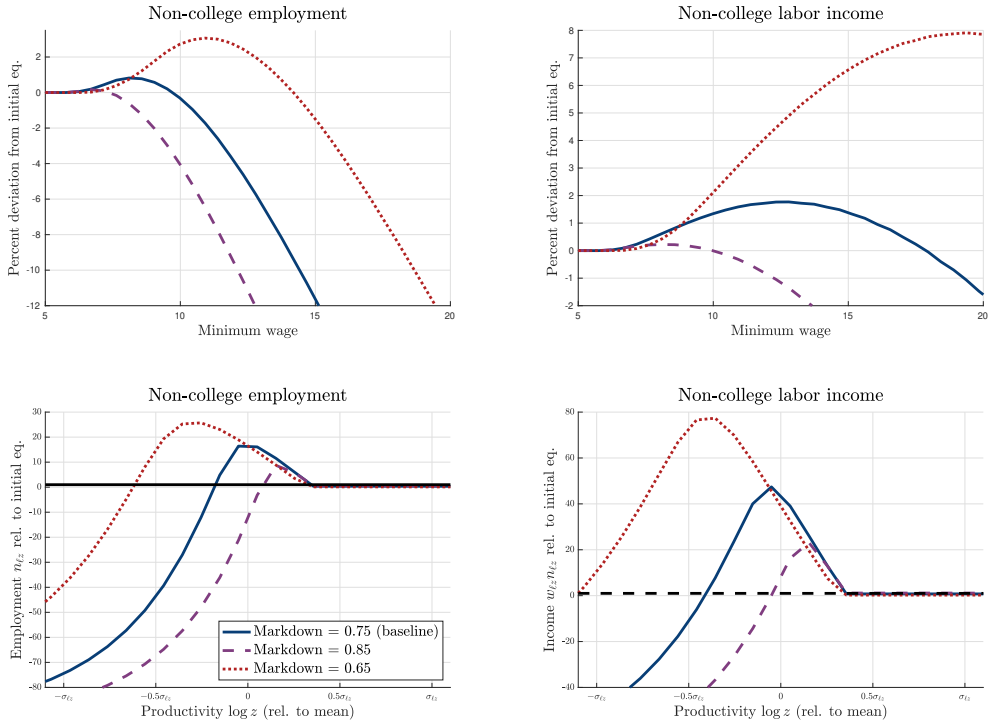
Note: Steady-state outcomes as a function of the minimum wage $\underline{w}$ in the model with capital-skill complementarity. The left panel plots the log-change of aggregate college employment, the log-change of aggregate college labor force, and the change in the unemployment rate relative to their levels in the initial steady state with $\underline{w} = 0$ (see the decomposition (33)). The right panel plots aggregate labor income of college workers. The x-axis is the level of the minimum wage $\underline{w}$ that relative to the median college wage in the initial equilibrium is the same as in the data.

FIGURE C.5: Distribution of Capital Types Along the Transition Path



Note: Distribution of capital types along the transition path. The x-axis indexes the type of capital by its average non-college labor-to-capital ratio.

FIGURE C.6: Comparative Statics to Monopsony Power ω in the Long-Run



Note: Steady-state outcomes for different values of the targeted wage markdown. The top panel plots aggregate non-college employment (left) and non-college labor income (right). The bottom panel plots the micro-level effect of a \$15 minimum wage as a function of ability z. Different markdowns correspond to different values of ω. For each such value, we compute the level of the minimum wage $\underline{w}$ that relative to the median non-college wage in the initial equilibrium is the same as in the data.

C.4 Comparative Statics With Respect to Monopsony Power

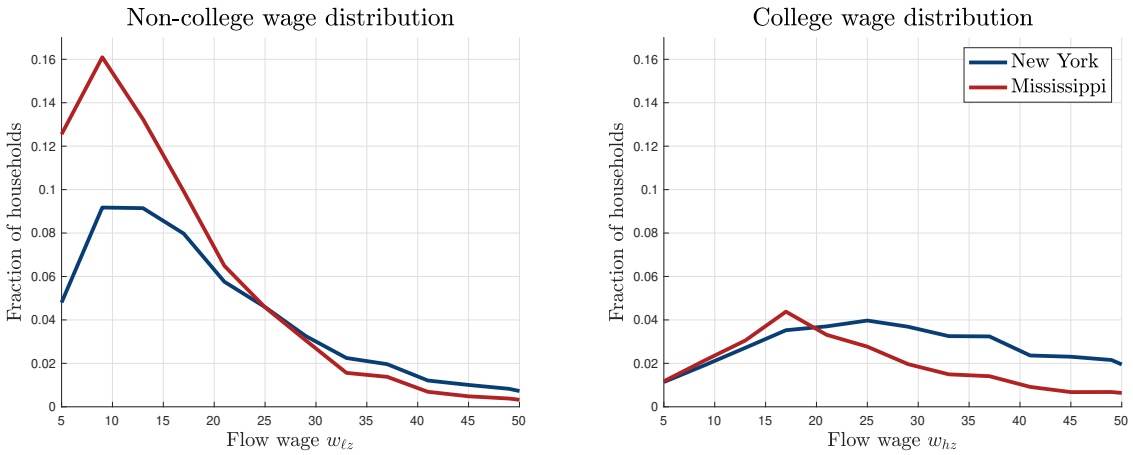
Figure C.6 plots how the long-run effects of the minimum wage depend on the degree of monopsony power. As described in the main text, we recalibrate the model to match three different markdown targets: our baseline value of 0.75, a larger markdown of 0.65, and a smaller markdown of 0.85. The top row shows that the aggregate Laffer curves are substantially shifted out and to the right for the larger degree of monopsony power, and are almost universally negative for the smaller degree of monopsony power. The bottom row shows that the set of workers who benefit from the \$15 minimum wage, either in terms of employment or labor income, is growing in the degree of monopsony power.

C.5 Additional Distributional Effects

In this appendix, we highlight the potential for a \$15 minimum wage to differential impact different states or sectors to the extent that these different states/sectors populate different portions of the aggregate wage distribution. Our model provides a mapping from an individual types' initial wage and their long-run change in employment or labor income in response to the minimum wage. The results in the main text aggregate up these changes using the aggregate wage distribution; now, we re-weight those individual-level changes by the distribution of wages in a given state or sector from the ACS. These results should only be viewed illustratively; in order to formally explore this heterogeneity, we would need to formally model linkages and reallocation frictions across states or sectors.

In order to illustrate the differences in wage distributions across states, Figure C.7 plots the distribution

FIGURE C.7: Distribution of Wages in Mississippi vs. New York



Note: Distribution of wages in Mississippi (red line) and New York (blue line). The left panel plots the distribution for workers without a college degree and the right panel plots the distribution for workers with college degree. The wage grid has been aggregated to $4 bins. Distributions across non-college and college workers sum to one within each state.

of wages in Mississippi and New York from our pooled 2017-2019 ACS data. The left panel plots the distributions for non-college workers and the right panel for college workers. Note that the distributions in these two panels have not been normalized, e.g. the entire wage distribution across both types of workers together sums to one. There are two important differences between these two states. First, New York has a much higher share of of college-educated workers (45%) compared to Mississippi (27%). Second, even within the two education groups, the distribution of wages in Mississippi is shifted left compared to New York. Both facts contribute to the stronger effects of the minimum wage in Mississippi in Table C.2.

The top panel of Table C.2 illustrates the distributional effects of the $15 minimum wage for selected states. The top three states correspond to those with the highest average hourly wage per worker in the ACS during 2017-2019 (Massachusetts, Connecticut, and New Jersey), while the bottom three states have the lowest average hourly wage per worker in the ACS during 2017-2019 (West Virginia, Arkansas, and Mississippi). According to this table, the $15 minimum wage will decrease employment by four times as much in Mississippi as in Massachusetts. This occurs because the distribution of wages in Mississippi is shifted left compared to Massachusetts; for example, the $15 minimum wage would bind on nearly 40% of workers in Mississippi but only 17% of workers in Massachusetts (Appendix C plots state-level wage distributions for comparison). Consistent with that, 27% of *all* workers (including college and non-college) in Mississippi would experience an employment decline vs. only 9% in Massachusetts. The other states line up in between these two extremes.

The bottom panel of Table C.2 shows the long-run results of a $15 minimum wage for the three sectors with the highest average hourly wage per worker (Finance/Insurance/Real Estate, Manufacturing, and Professional Services) and the three sectors with the lowest average hourly wage per worker (Entertainment, Retail Trade, and Personal Services). A $15 minimum wage will have a relatively small effect on the FIRE sectors because there are few workers (15%) in that sector who currently earn less than $15 per hour. Conversely, a $15 minimum wage will have much more dramatic effects on workers in the Retail Trade and Personal Services sectors—more than 50% of workers currently each less than $15.

Again, these results should be only considered illustrative; in order to assess fully the effects of a large change in the minimum wage across locations and sectors, one would need to have a model allowing for linkages across space and sectors. Additionally, one would potentially want to allow for the production technologies to differ across sectors. However, the back-of-the-envelope findings we highlight here suggest

TABLE C.2: Heterogeneous Long-Run Effects Across States and Sectors

State	Employment change	Share($w_i < \overline{w}$)	Share($\Delta n_i < 0$)	Share($\Delta w_i n_i < 0$)
Massachusetts	-3.9%	0.17	0.09	0.06
Connecticut	-4.4%	0.18	0.11	0.07
New Jersey	-6.0%	0.21	0.13	0.09
⋮				
West Virginia	-10.8%	0.37	0.23	0.15
Arkansas	-11.2%	0.19	0.39	0.16
Mississippi	-13.6%	0.40	0.27	0.19

Sector	Employment change	Share($w_i < \overline{w}$)	Share($\Delta n_i < 0$)	Share($\Delta w_i n_i < 0$)
FIRE	-3.4%	0.15	0.08	0.05
Manufacturing	-5.2%	0.25	0.16	0.13
Professional services	-6.2%	0.24	0.17	0.09
⋮				
Entertainment	-11.4%	0.38	0.50	0.39
Retail Trade	-17.2%	0.51	0.34	0.23
Personal Services	-18.2%	0.51	0.35	0.25

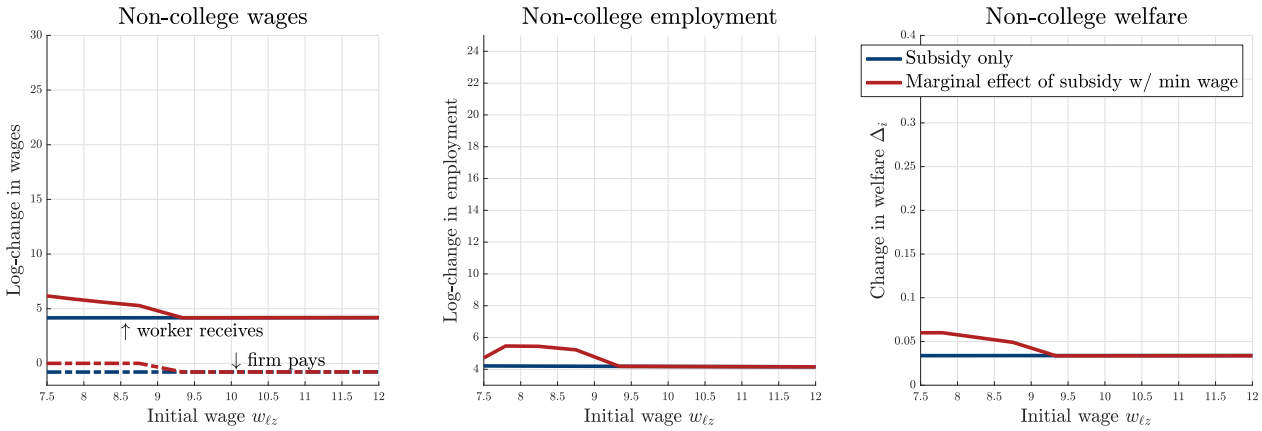
Note: Long-run effects of the $15 minimum wage across U.S. states and sectors. We compute the state-level outcome as the weighted value of that outcome across z_i types, using the distribution of z_i implied by that state/sector's wage distribution in the ACS. "Employment change" is change in total employment. "Share($w_i < \overline{w}$)" is the share of initial workers whose wage is below the $15 minimum wage. "Share($\Delta n_i < 0$)" is the share of initial workers whose employment declines due to the $15 minimum wage. "Share($\Delta w_i n_i < 0$)" is the share of initial workers whose labor income declines due to the $15 minimum wage.

that a large change in the minimum wage will disproportionately affect some states and sectors relative to others because of differences in the underlying wage distribution.

D Additional Results About Alternative Policies

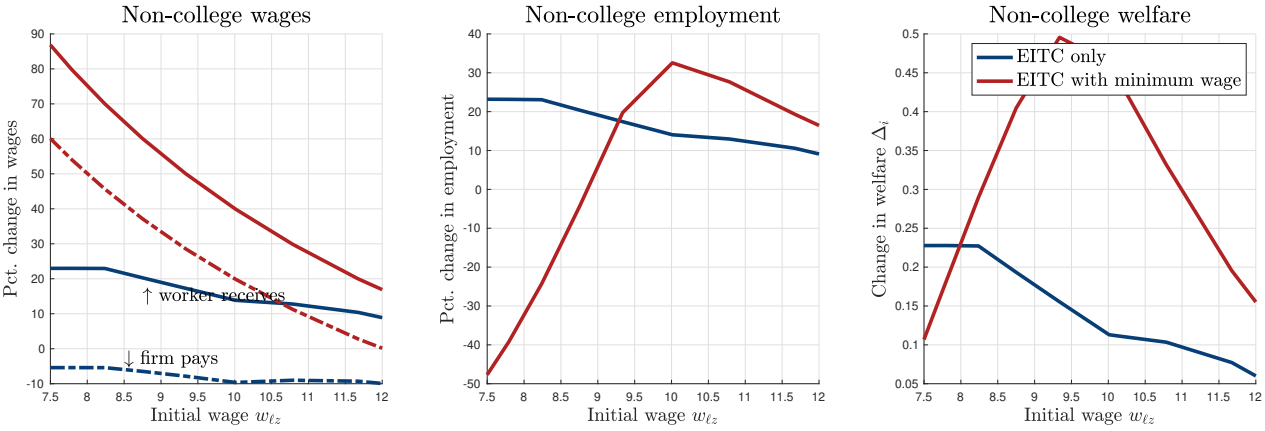
This appendix shows two additional results. First, Figure D.1 isolates the interaction between the constant transfer and the minimum wage by subtracting off the pure effect of the minimum wage in isolation. Second, Figure D.2 shows how a $12 minimum wage interacts with our budget-equivalent EITC from Section 5.

FIGURE D.1: Effect of EITC with and without $\underline{w} = \$9.25$ Minimum Wage



Note: 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 the initial equilibrium (without any policies). The x-axis corresponds to the initial wage $w_{\ell z}$ of a particular type-z worker in the initial equilibrium.

FIGURE D.2: Effect of EITC with and without $\underline{w} = \$12$ Minimum Wage



Note: 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 the initial equilibrium (without any policies). The x-axis corresponds to the initial wage $w_{\ell z}$ of a particular type-z worker in the initial equilibrium.