The Distributional Impact of the Minimum Wage*

Erik Hurst
Chicago Booth
Erik.Hurst@chicagobooth.edu

Patrick Kehoe
Stanford and Minneapolis Fed
pkehoe@stanford.edu

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

Thomas Winberry
Wharton
twinb@wharton.upenn.edu

Abstract

We develop a general-equilibrium framework with rich worker heterogeneity and monopsony power in frictional labor markets to study the short-run and long-run implications of a large increase in the federal minimum wage of the magnitude currently debated in the United States. We ensure that our framework is simultaneously consistent with the evidence that changes in the relative price of labor have *small* employment effects in the short run when emanating from changes in the minimum wage, but *large* effects in the long run when emanating from changes in the price of capital. We find that an increase in the minimum wage to $15 per hour will lead to a drop in employment, income, and welfare for the bottom third of non-college-educated workers. Hence, a high minimum wage has perverse distributional impacts: although it increases labor income in the aggregate, it reduces the welfare of precisely the low-income workers that it is designed to help. Finally, we show that either a budget-equivalent change to the earned-income tax credit or to the progressivity of the current tax and transfer system does a much better job at redistributing resources to non-college-educated workers at the bottom of the wage distribution.

*The current draft is an updated version of the one presented at the BSE (Barcelona School of Economics) Summer Forum, June 2021. We thank Christian vom Lehn for helpful conversations and Brent Neiman and Kjetil Storesletten for sharing their data/codes. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Recent proposals have been advanced in the United States to increase the current federal minimum wage of $7.25 to at least $15 per hour. The goal of these proposals is to improve the welfare of workers currently earning less than this new minimum, especially those at the bottom of the wage distribution. In contrast to past changes in the minimum wage, such a proposed increase will have an impact on a large fraction of the U.S. workforce. Figure 1 shows the distribution of wages across individuals from the 2017-2019 American Community Survey (ACS) by whether or not an individual has obtained a bachelor's degree.[1] As seen from this figure, 40% of non-college-educated workers and 10% of college-educated workers currently earn a wage less than $15 per hour. Even within each education group, at the current wages, the minimum wage would have substantial effects—potentially very different for different workers. For example, for non-college-educated workers in the bottom 20% of their wage distribution, the proposed increase in the minimum wage would amount to an 80% increase in their wages, everything else equal. For those in the top 60%, however, the higher minimum wage would not bind at their current wages.[2]

FIGURE 1: Empirical Wage Distribution

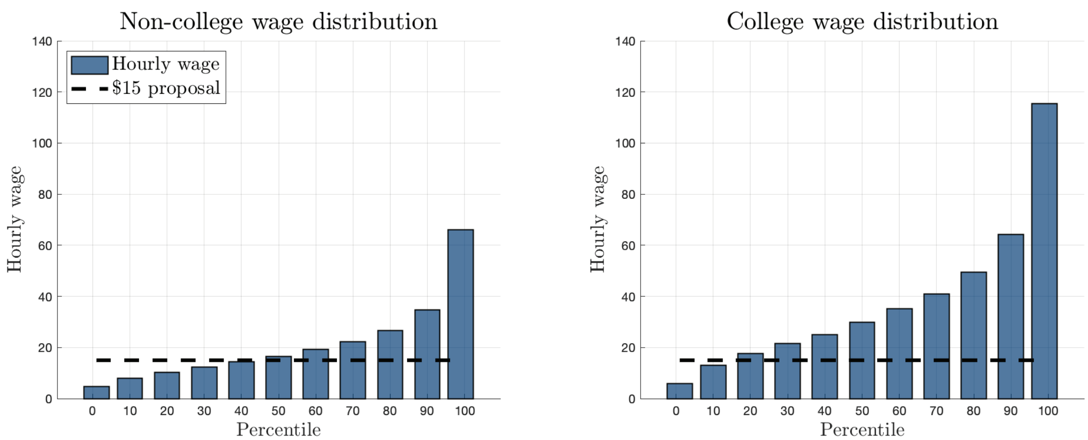

Notes: Wage distribution of full time employed workers over the age of 16 from the 2017-2019 ACS data.

The goal of this paper is to develop a framework that can capture this large within group-heterogeneity in the effects of the minimum wage so as to assess the redistributive impact of this policy. Our key finding is that although non-college-educated workers as a whole experience an increase in their labor income in response to a $15 minimum wage, nearly all of the workers that

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

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

this policy is designed to benefit are made worse off. In particular, all but the highest earning 5% of the 40% of non-college-educated workers currently earning less than the proposed minimum wage experience a reduction in welfare. We also find that, if the goal of the minimum wage is to help workers at the bottom of the wage distribution, other policies such as an expansion of the earned-income tax credit (EITC) or, more generally, an increase in the progressivity of the existing tax system dominate a minimum wage policy.

The framework we develop is consistent with three key features of the data. First, given that our goal is to explore the redistributive impact of the minimum wage, we ensure that our model matches the wage distribution from Figure 1. Second, at least since Robinson (1933), proponents of the minimum wage have argued that when firms have some degree of monopsony power in the labor market, increases in the minimum wage can help alleviate some of the distortions arising from it and lead to a desirable redistribution of firm profits to worker wages. Consistent with this idea, a large recent literature has documented that wages are marked down relative to the marginal product of labor.[3] To allow for this mechanism, we match the wage markdowns estimated by Berger, Herkenhoff and Mongey (2019). Third, a large empirical literature has found that increases in the minimum wage, which raise the relative price of labor, have at most a *small* effect on employment in the short run.[4] A separate literature, however, has found that declines in the relative price of capital, which also raise the relative price of labor, have *large* long-run effects on employment (see, for example, Krusell et al. (2000)). We embed in our model the feature that capital is of a putty-clay nature to capture the notion that the firms' adjustment to changes in the relative price of labor are small in the short run but potentially large in the long run. With this feature, our model simultaneously matches these short- and long-run empirical patterns.

We turn now to describing our framework in more detail. We allow for two dimensions of worker heterogeneity. First, workers are classified into two broad groups: non-college educated and college-educated. Second, within each group, workers differ in their productivity in the labor market. Such a productivity dispersion allows us to capture the rich heterogeneity in wages displayed in Figure 1 and, as we will argue, is critical for evaluating the distributional impact of the minimum wage. Importantly, this dimension of our model is consistent with the evidence that the vast 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)).

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

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

We incorporate monopsonistic competition among firms into a directed search environment that allows for endogenous worker non-participation and unemployment in the presence of matching frictions. Our monopsonistic directed search equilibrium generalizes the market utility approach that Montgomery (1991) used for competitive search.[5] Firms' monopsony power arises from the fact that the labor market consists of a large number of firms with imperfectly substitutable jobs, over which households have preferences. This specification is the natural analogue of firm monopolistic competition when consumers have preferences over differentiated goods.

This formulation is motivated by several reasons. First, the non-participation margin allows us to capture the idea that a higher minimum wage incentivizes non-participants to enter the labor market and search for a job. Second, directed search allows us to generalize the notion of a firm-specific labor supply curve in Robinson (1933) to incorporate dynamic wage contracts with multi-worker firms and endogenous job-finding rates. Finally, and most importantly, a search framework avoids the need to specify ad-hoc rationing rules in the allocation of workers to jobs when the minimum wage induces labor supply to exceed labor demand.

We capture the differing short-run and long-run substitutability of capital and labor through a putty-clay approach.[6] When capital is of this nature, ex-ante it is possible to build a machine with any desired ratio of capital to the labor of each worker type that lies on the frontier of the production function. Once a machine is built, however, it is clay-like in that the machine uses a fixed amount of each type of labor to operate it. That is, our production function is nested CES ex ante but Leontief ex post.

In the short run, this feature implies that given a stock of machines, the demand for each type of labor is inelastic because it is infeasible to substitute between existing capital and the various types of labor. Over time, though, new capital goods embodying new ratios of capital to labor can be installed. Hence, following a change in the relative price of labor, firms substitute away from any type of labor that a policy makes more expensive in the long run. With putty-clay capital, then, after the introduction of the minimum wage, firms do not immediately cut employment since they need the currently employed workers to operate the existing capital stock. But in the long run, firms

[5]For work on competitive search more generally, see also Moen (1997) and, for an extensive review of the literature, Wright et al. (2021).

[6]Our modeling of putty-clay capital builds on the work that dates back at least to Johansen (1959) and was extended by Calvo (1976). Sorkin (2015) and Aaronson et al. (2018) use a variant of the standard putty-clay setup to analyze the long-run effects of the minimum wage within a industry-equilibrium model. In their setup, the putty-clay aspect of technology applies to a firm as a whole, rather than at the level of each unit of capital as in our case. Hence, all inputs, including capital, once chosen are fixed forever within a firm, and productivity deterministically declines. Thus, after a firm is set up, the only decision for it is when to shut down. In contrast, we consider a general-equilibrium framework in which firms continually make decisions about which types of workers to hire or fire, and which mix of capital types to use. Aaronson et al. (2018) apply the setup in Sorkin (2015) to study the effect of the minimum wage on the restaurant industry.

shift their capital labor ratios to the preferred mixes governed by the underlying CES production function and reduce employment accordingly.

Empirically, we discipline both the short-run and long-run elasticities of substitution for capital and labor using changes in the relative price of labor to capital emanating from changes in the price of capital. We assume that the underlying production function is a nested CES function over non-college workers, college workers, and capital, as in the seminal work of Krusell et al. (2000). We estimate the long-run elasticities of input substitution for this production function using long-run changes in the relative price of labor to capital and in the distribution of income across education groups at the sectoral level. Consistent with Krusell et al. (2000), we find that college-educated workers are complementary to capital whereas non-college educated workers are substitutable to capital. We then provide evidence for our model's implications of small short-run input elasticities using temporary changes in the after-tax price of capital arising from the Bonus Depreciation Allowance, which is an investment tax incentive with differential impacts across sectors.[7] Consistent with our putty-clay model, a temporary decline in the relative price of capital has an insignificant effect on the distribution of labor income because firms cannot change their existing capital to labor ratios. The version of our model with standard capital is inconsistent with this finding.

Our first key finding is that, in the long run, large changes in the minimum wage disproportionately reduce the employment of low-income workers. In our model, small changes in the minimum wage may increase employment because firms set wages below their efficient level due to their monopsony power. However, if the minimum wage is set at a level that would offset the monopsony distortion for high-productivity workers, such a minimum wage would effectively price the low-productivity workers out of the labor market and lead to massive reductions in their employment.

In our baseline calibration, a $15 minimum wage per hour has effects like these. In particular, it reduces employment for 58% of non-college-educated workers, especially for the low-productivity ones. For example, the employment rates of workers who were originally making $7.25 per hour falls by 68%. Of course, households whose members keep their jobs will end up earning higher wages, but even so, the total labor income of such workers falls by 34%. In contrast, the substitution by firms toward higher-productivity workers implies that those who were initially earning over $14 per hour will be better off. We therefore conclude that a large minimum wage has perverse distributional impacts. Namely, it reduces employment of exactly the low-income workers that it is designed to help. Crucially, this occurs despite the fact that the minimum wage has a small effect on aggregate employment.

Our second key finding is that the short-run effects of the minimum wage are much smaller than

[7] House and Shapiro (2008) and Zwick and Mahon (2017), among others, use this sectoral variation to estimate the effects of the Bonus Depreciation Allowance on investment.

their long-run effects due to putty-clay frictions. Although a $15 minimum wage will ultimately reduce aggregate non-college employment by over 6% in the long run, in our model employment only falls by less than 1% four years after its introduction. In this sense, our model is consistent with the established evidence that observed changes in the minimum wage have a small impact on employment in the short run. But, conversely, this finding also implies that empirical estimates of the effects of the minimum wage in the first few years after its introduction will not detect its ultimate long-run consequences.

Our third key finding is that alternative existing policies are better-suited to raise the employment and welfare of low-income workers. The primary goal of a minimum wage policy is to increase the after-tax income of workers who earn relatively low wages. Moreover, a common rationale for why minimum wages may be beneficial is that they help offset the monopsony power of firms, which leads wages to be set at an inefficiently low level. A natural set of policies that can accomplish these goals are progressive taxes and transfers. We then start by examining the existing Earned-Income Tax Credit (EITC) policy, which is a refundable tax credit for low- to moderate-income working individuals and couples, and then turn to examining the effects of the overall progressivity built into the current U.S. tax code.

We evaluate how such policies perform in our environment both in terms of increasing the income of workers and offsetting the monopsony power of firms.[8] We find that a progressive tax and transfer system does a much better job of decreasing the labor market power of firms and redistributing resources to non-college-educated workers in the bottom half of the wage distribution. In particular, we show how the EITC leads to higher welfare for low-earning individuals relative to the minimum wage. We find a similar result from increasing the general progressivity of the U.S. tax system. Interestingly, even a simple uniform income tax cut, financed by a tax of the same amount on corporate profits, is preferable to the minimum wage if the goal is to help low-wage workers. Overall, our analysis shows that many other tax and transfer policies dominate a comparable change in the minimum wage if the goal is to simultaneously reduce monopsony power in the labor market and support the income of low earners.

Three papers are more closely related to our work. Flinn (2006) also examines the effects of the minimum wage within an equilibrium model of labor market search. In particular, Flinn (2006) studies an environment in which wages are determined by Nash-bargaining and the worker's bargaining power is lower than the level required for efficiency by the Hosios (1990) condition. Since workers' bargaining share is smaller than the elasticity of the matching function with respect to

[8]To make the magnitude of these policies comparable to those of the minimum wage policy, we first compute the drop in excess corporate profits that result from an increase in the minimum wage to $15, which amounts to an implicit tax on corporate profits. We then fund each of the alternative policies by an explicit tax on excess corporate profits that raises the same revenues.

the mass of the unemployed, the equilibrium is inefficient. In this case, an increase in the minimum wage is a crude way of effectively increasing the bargaining power of workers leading to increases in individual employment. Flinn (2006) highlights how his imposition of a inefficiently low bargaining share is somewhat akin to allowing firms to have some degree of monopsony power.

More closely related to our modeling of monopsonistic firm competition is a special case of the Cournot competition framework developed by Berger, Herkenhoff and Mongey (2019) in a non-search context. Building on their prior work, Berger, Herkenhoff and Mongey (2021) use their theoretical setup to pursue a normative analysis of the optimal level of the minimum wage. Beyond our focus on positive implications, we differ in three key respects from Berger, Herkenhoff and Mongey (2021). In particular, we study the distributional implications arising from worker heterogeneity whereas their study focuses on firm heterogeneity arising from differences in firm productivity. Given our focus on the distributional impact of the minimum wage across workers with different productivity, we comparatively evaluate alternative policies and illustrate that for lower-wage workers existing transfer policies alternative to the minimum wages—such as the EITC—dominate the minimum wage. Lastly, we highlight how a realistic putty-clay technology on the part of firms generates a prolonged labor market transition in response to the minimum wage, which implies that short-run regressions are uninformative about long-run labor market responses.

Most of the empirical literature evaluating the minimum wage exploits local changes in the minimum wage and examines short-run employment response for broad groups of workers. Cengiz et al. (2019), however, use local minimum wage changes to examine the extent to which the minimum wage affects the employment propensity of workers who were initially below the new minimum wage.[9] They find little employment declines for low-wage workers in the few years following a state-level minimum wage change. Our results are consistent with these findings for two reasons. First, most of the state-level changes in the minimum wage used by these authors to identify the employment responses to the minimum wage are quite small relative to the current proposed changes in the national minimum wage. Indeed, we find that small changes in the minimum wage have only small effects on the employment rates of low-productivity workers even in the long run. Second, and more importantly, like almost all papers in the literature exploiting cross-regional variation, Cengiz et al. (2019) only examine short-run employment effects. As we highlight, allowing for reasonable capital adjustment processes on the part of firms implies that employment effects are small in the first few years after a minimum wage increase even though long-run responses are large.

[9] 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 mechanisms we consider in our framework.

1 Model

We consider a dynamic general equilibrium model with labor markets subject to matching frictions characterized by several key features. First, consumers are heterogeneous in their *broad skill level* $b \in \{\ell, h\}$ where ℓ denotes low-skilled and h denotes high-skilled and within each skill type s, workers have an *ability level* z drawn from skill-type-specific discrete distributions with measures μ_ℓ and μ_h. In particular, we can index a consumer by i which denotes both broad skill group and ability level so that $i \in I \equiv I_\ell \cup I_h$ where the set abilities of low-skilled types is $I_\ell = \{i | z_i \in \{z_{\ell 1}, \ldots, z_{\ell K}\}\}$ and the set of abilities high skill types is $I_h = \{i | z_i \in \{z_{h1}, \ldots, z_{hK}\}\}$. For shorthand, we often let $i = (b, z_i)$ denote a skill-ability pair and let μ_i the measure of families of that type. In our quantitative analysis we identify the skills levels ℓ and h with non-college and college workers respectively.[10]

Second, firms are monopsonistically competitive in the market for labor. We extend the standard competitive search framework to capture this type of competition. Third, in production, capital is complementary with high-skilled labor and substitutable with low-skilled labor. Fourth, we focus on putty-clay capital, but contrast our model's predictions to a version with standard putty-putty capital. In our quantitative exercise we show how, after the introduction of the minimum wage, the putty-clay capital slows down the transition to the new steady state.

1.1 Production and Matching Technologies

Consider first the 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-skilled labor $\bar{n}_{\ell jt}$ and an aggregate of efficiency units of high-skilled labor $\bar{n}_{hjt}$. The capital accumulation law is $k_{jt+1} = (1-\delta)k_{jt} + x_{jt}$. Consumers view the labor supplied to these different firms as differentiated as we discuss below. In the equilibrium we construct all firms will make the same choices.

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

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

where

$$G(k_{jt}, \bar{n}_{hjt}) = \left[\lambda k_{jt}^{\frac{\alpha-1}{\alpha}} + (1-\lambda)(\bar{n}_{hjt})^{\frac{\alpha-1}{\alpha}}\right]^{\frac{\alpha}{\alpha-1}}. \tag{2}$$

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

The aggregates of low-skilled labor $\bar{n}_{\ell jt}$ and high-skilled labor $\bar{n}_{hjt}$ used by firm j are

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

Here for $i \in I_\ell$, n_{ijt} is the amount of low-skilled labor per family of type i supplied to firm j and $\mu_i n_{ijt}$ is the total amount of low-skilled labor from all families of type i supplied to firm j and similarly for $i \in I_h$. The outer nest in (1) is a CES production function over efficiency units of low-skilled labor $\bar{n}_{\ell jt}$ and capital-skilled labor aggregate $G(k_{jt}, \bar{n}_{hjt})$ and the inner nest is a CES production function over capital k_{jt} and efficiency units of high-skilled labor $\bar{n}_{hjt}$.[11]

Consider next the matching technology. We consider a directed search setting in which each firm j posts a measure of vacancies $\mu_i a_{ijt}$ directed at the measure of workers $\mu_i s_{ijt}$ of type i searching for it. Here a_{ijt} and s_{ijt} denote the vacancies and searchers per family i and $\mu_i a_{ijt}$ and $\mu_i s_{ijt}$ are the total vacancies and searchers for all families of type i. The cost of posting a total of $\mu_i a_{ijt}$ for type i workers is $\kappa_i \mu_i a_{ijt}$. The total matches created by $\mu_i a_{ijt}$ vacancies and $\mu_i s_{ijt}$ searchers is given by a constant returns to scale matching function

$$m(\mu_i a_{ijt}, \mu_i s_{ijt}) = B_i (\mu_i a_{ijt})^\eta (\mu_i s_{ijt})^{1-\eta}. \tag{3}$$

Now, if firm j posts $\mu_i a_{ijt}$ total vacancies for type i workers and families of type i send $\mu_i s_{ijt}$ total workers searching for that firm, then firm j is matched with a measure $m(\mu_i a_{ijt}, \mu_i s_{ijt}) = \lambda_f(\theta_{ijt}) \mu_i a_{ijt}$ of workers of type i where $\lambda_f(\theta_{ijt}) = \frac{m(\mu_i a_{ijt}, \mu_i s_{ijt})}{\mu_i a_{ijt}} = \frac{m(a_{ijt}, s_{ijt})}{a_{ijt}}$ so that $\lambda_f(\theta_{ijt})$ is the probability of a match per vacancy. We can also write these total matches as $m(\mu_i a_{ijt}, \mu_i s_{ijt}) = \lambda_w(\theta_{ijt}) \mu_i s_{ijt}$ where $\lambda_w(\theta_{ijt}) = \frac{m(\mu_i a_{ijt}, \mu_i s_{ijt})}{\mu_i s_{ijt}} = \frac{m(a_{ijt}, s_{ijt})}{s_{ijt}}$ is the probability of a match per family, referred to as the *job-finding rate*. Here we have used that the matching function is constant returns to scale and market tightness is $\theta_{ijt} = a_{jit}/s_{ijt}$.

1.2 Families

We set up a monopsonistic directed search equilibrium that generalizes the market utility approach that Montgomery (1991) used for competitive search. For competitive search see also Moen (1997)

[11] As Krusell et al. (2000) discuss, an alternative nesting of these three factors has the form $F(\bar{n}_{hjt}, \bar{G}(k_{jt}, \bar{n}_{\ell jt}))$ used in Stokey (1996). The parameters of either nesting can be chosen so as to imply any elasticity for low-skilled labor and capital and for high-skilled labor and capital. In particular, they can be chosen so that, as in the data, low-skilled labor and capital are substitutes whereas high-skilled labor and capital are complements. Given this setting, the nesting in (1) that the elasticity of low-skilled labor and high-skilled labor is the same as that between low-skilled labor and capital, which is consistent with the data, since low-skilled labor and high-skilled labor are estimated to be substitutes. The alternative nesting instead implies that the elasticity of low-skilled labor and high-skilled labor is the same as that between high-skilled labor and capital, which is not consistent with the data.

and, for an extensive review of the literature, Wright et al. (2021).

We represent the insurance arrangements among these consumers assuming that each consumer of type i belongs to one of a large number of identical families consisting entirely of the same skill and ability level, each of which has a large number of household members. Risk sharing within such families implies that each household member consumes the same amount of goods at date t, regardless of the idiosyncratic shocks that such a member experiences. (This type of risk-sharing arrangement in a search model is familiar from the work of Merz 1995 and Andolfatto 1996.)

1.2.1 A Family's Problem

The utility function of a family of type $i \in I$ 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 at t and s_{it} are the total searchers. In particular, $s_{it} = \sum_j s_{ijt}$ where s_{ijt} is the measure of family members of type i searching for firm j at t and n_{ijt} is the measure of family members in family i who work at firm j at t. The *index* of the disutility of work associated with the allocation (n_{ijt}) is

$$n_{it} = \left[\sum_j n_{ijt}^{\frac{1+\omega}{\omega}} dj \right]^{\frac{\omega}{1+\omega}} \quad \text{with } \omega > 0. \tag{4}$$

Notice that ω captures the degree of imperfect substitutability across jobs: the smaller is ω the less substitutable are jobs and when $\omega = \infty$ jobs are perfectly substitutable. Here we are adapting the standard way of modeling imperfect substitutability across differentiated consumer goods to modeling imperfect substitutability across differentiated jobs. (See Berger, Herkenhoff, and Mongey 2021 for a generalization of this framework along with discussion of various interpretations of what this differentiation is meant to capture.)

We consider directed search in which a firm of type j in period t chooses both the measures of vacancies (a_{ijt}) for all consumer types i to post at t and present value of wages (W^F_{ijt+1}) for all consumer types i who find the firm at t and begin worker at $t+1$. The family judges the attractiveness of the combined vacancy-wage pairs by their equivalent market tightnesses (θ_{ijt}) at t and present value wage offers (W^F_{ijt+1}).

The timing of events is that if a consumer of type i is employed at firm j in period $t-1$, then with probability σ that consumer separates from that job at the end of period $t-1$. At the beginning of period t, consumers without jobs decide whether to stay at home or search and, if they decide to search, which type j of firm to search for. The transition law for consumers of type i working at firm j is

$$n_{ijt+1} = (1-\sigma)n_{ijt} + \lambda_w(\theta_{ijt})s_{ijt} \text{ for all } j, \tag{5}$$

where $\lambda_w(\theta_{ijt})$ is the job-finding rate in firm j at t for a type i consumer, namely the probability that a consumer who searches for a job in firm j during period t, finds a match with firm j at the end of period t. Here a consumer who chooses to search for a job in period t and who finds a job starts working in period $t+1$, consumers who do not find jobs begin period $t+1$ with the option to search again or stay at home.

In this economy there is idiosyncratic uncertainty over finding jobs and losing jobs at the individual consumer level but since there are no aggregate shocks, and families consist a large number of members, there is no aggregate uncertainty at the family level. So our economy is a deterministic one in terms of aggregates. The budget constraint of family i is

$$\sum_{t=0}^{\infty} Q_{0,t} c_{it} \leq \sum_{t=0}^{\infty} Q_{0,t} \left[\sum_j W_{ijt}^F \lambda_w(\theta_{ijt-1}) s_{ijt-1} \right] + \psi_i \Pi_0, \tag{6a}$$

where $Q_{0,t}$ denotes the price of the homogenous good at date t in units of that good at date 0, Π_0 is the present value of all the firms' profits, and ψ_i is the share of those profits owned by the low-skilled family. To understand the term in brackets note that if s_{ijt-1} consumers of type i search for a job at firm j during period $t-1$, then $\lambda_w(\theta_{ijt-1})s_{ijt-1}$ find a job; these consumers start working in period t earn W_{ijt}^F in present value terms in units of period t goods. Since we have complete markets in financial markets, it is without loss of generality to adopt the convention that a firm fulfills its present value wage offer W_{ijt}^F by offering a constant wage during the associated match so that

$$W_{ijt}^F = w_{ijt} + (1-\sigma) Q_{t,t+1} w_{ijt} + (1-\sigma)^2 Q_{t,t+2} w_{ijt} + \ldots, \tag{7}$$

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

At time t, family i chooses consumption c_{it}, an employment plan (n_{ijt+1}) across firms, a search plan (s_{ijt}) across firms subject to the transition law for employment at each firm (5), the budget constraint (6a) and a nonnegativity constraint on searchers $s_{ijt} \geq 0$ where have substituted out $s_{it} = \sum_j s_{ijt}$ and n_{it} using (4). Dropping the i subscript designating consumer type for clarity and letting $\beta^{t+1}\mu_{jt+1}, \gamma,$and $\beta^t \chi_{jt}$ be the multipliers on the transition law, budget constraint, and nonegativity constraint, the first order conditions imply $\beta^t u_{ct} = \gamma Q_{0,t}$ so that

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

and

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

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

Here (8) is the standard Euler equation for consumption. To understand the next two equations, recall that when consumers from a family search for a job at firm j at t that begins in period $t+1$, that family gets paid a present value of wages W^F_{jt+1} for a commitment to work until an exogenous separation from that firm occurs which, in each period, happens with probability σ. In (9), μ_{jt+1} is the discounted marginal (dis)utility of having one more unit of consumers in a family work at firm j at $t+1$, $(1-\sigma)$ units at $t+2$, $(1-\sigma)^2$ more units at $t+3$ and so on. To express this disutility in consumption units, we define $W^N_{jt+1} \equiv \mu_{jt+1}/u_{ct+1}$ and substitute $Q_{t+1,t+2} = \beta u_{ct+2}/u_{ct+1}$ to write W^N_{jt+1} recursively as

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

Now, (10) is a rearranged version of the first order condition for the mass of consumers sent to search for firm, s_{jt}. Now substituting $\mu_{jt+1}/u_{ct+1} = W^N_{jt+1}$ and $Q_{t,t+1} = \beta u_{ct+1}/u_{ct}$ in (10) gives

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

Consider all jobs j such that the consumer is actively searching in that $s_{jt} > 0$ so that $\chi_{jt} = 0$. A consumer who searches at t for any job j immediately incurs some disutility, when expressed in consumption units is u_{st}/u_{ct}. The benefit of incurring that cost is that with probability $\lambda_w(\theta_{jt})$ this consumer finds a job at time $t+1$. In time t consumption units, the benefit of finding a job and starting to work at $t+1$ is getting paid W^F_{jt+1} in units of time $t+1$ consumption goods and incurring an expected present value of disutility of working of W^N_{jt+1} Expressed in time t consumption units, the expected benefit of finding a job is $Q_{t,t+1}\lambda_w(\theta_{jt})(W^F_{jt+1} + W^N_{jt+1})$. Clearly, then if a consumer actively searches for jobs j and k, so that s_{jt} and s_{kt} are both positive so $\chi_{jt} = \chi_{kt} = 0$, then from (12) we have that the value of searching at firms j and k must be equal in that the following condition must hold

$$\lambda_w(\theta_{jt})(W^F_{jt+1} + W^N_{jt+1}) = \lambda_w(\theta_{kt})\left(W^F_{kt+1} + W^N_{kt+1}\right). \tag{13}$$

1.2.2 The Participation Constraint

In our directed search setup we capture the monopsony power of each firm through a participation constraint that we add to the firm problem. Here we discuss how this constraint follows naturally from the family's first order condition for searching just discussed. In particular, in our construction of a symmetric equilibrium, each firm needs to think through the value of deviating from that symmetric allocation. Hence, we consider an *almost symmetric allocation* in which all firms but one, say firm j, are offering the same *common value* of $\lambda_w(\theta_{it})\left(W^F_{it+1}+W^N_{it+1}\right)$ to type i consumers and firm j is contemplating offering a potentially different value $\lambda_w(\theta_{ijt})(W^F_{ijt+1}+W^N_{ijt+1})$ to this type of consumers. Then for firm j to attract a consumer it must offer at least the common value. That is, firm j's offer of $(\theta_{ijt}, W^F_{ijt+1})$ must satisfy

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

or else firm j's offer will not attract any type i consumers. We refer to (14) as the *participation constraint* for firm j for type i consumers.

More precisely, each period t consists of two stages. In stage 1, each firm j post vacancies (a_{ijt}) for all consumer types i which, equivalently, determines the market tightness (θ_{ijt}) for all types, and commits to present value wage offers (W^F_{ijt+1}) for each consumer type i who find them and begin work at $t+1$ and each family decides on the total measure of searchers s_{it} to send to the market. In stage 2, after having observed all offers, the family decides the search plan (s_{ijt}) of which firms to direct these searchers to, so that $s_{it}=\sum_j s_{ijt}$. For a consumer of type i this means choosing which firm j to search for when confronted with the menu $(\theta_{ijt}, W^F_{ijt})$ for all j. These two stages should be thought of as occurring at the beginning of each period t.

We turn to characterizing the second stage of a symmetric equilibrium for workers which will be in the form of a participation constraint that firms face when they make offers at stage 1. We can solve forward the recursive expression (11) for the present value of the incremental disutility suffered from working one extra unit at firm j in the first period of the match and in subsequent period $(1-\sigma)$ units, and so on as

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

Now, the key way that the monoposonistic power of a firm shows up in a firm's problem is when it takes the derivatives of W^N_{ijt+1} with respect to its vacancies a_{ijt} and its market tightness θ_{ijt}. For example, the derivative of W^N_{ijt+1} with respect to a_{ijt} is given by

$$\frac{\partial W^N_{ijt+1}}{\partial a_{ijt}} = \frac{u_{nit+1}}{u_{cit+1}} \frac{\partial \left[(n_{ijt+1}/n_{t+1})^{\frac{1}{\omega}}\right]}{\partial a_{ijt}} + Q_{t+1,t+2}(1-\sigma)\frac{u_{nit+2}}{u_{cit+2}} \frac{\partial \left[(n_{ijt+2}/n_{t+2})^{\frac{1}{\omega}}\right]}{\partial a_{ijt}} + \dots \tag{16}$$

with

$$\frac{\partial \left[(n_{ijt+1}/n_{it+1})^{\frac{1}{\omega}}\right]}{\partial a_{ijt}} = \left(\frac{n_{ijt+1}}{n_{it+1}}\right)^{\frac{1}{\omega}-1} \frac{\lambda_f(\theta_{ijt})}{\omega n_{it+1}} \text{ and } \frac{\partial \left[(n_{ijt+2}/n_{it+2})^{\frac{1}{\omega}}\right]}{\partial a_{ijt}} = (1-\sigma)\left(\frac{n_{ijt+2}}{n_{t+2}}\right)^{\frac{1}{\omega}-1} \frac{\lambda_f(\theta_{ijt})}{\omega n_{t+2}},$$

where in the first expression we used that $n_{it+1} = (1-\sigma)n_{ijt} + \lambda_f(\theta_{ijt})a_{ijt}$ and in the second we used that $n_{ijt+2} = (1-\sigma)^2 n_{ijt} + (1-\sigma)\lambda_f(\theta_{ijt})a_{ijt} + \lambda_f(\theta_{ijt+1})a_{ijt+1}$. Proceeding in the same fashion we can show that

$$\frac{\partial W^N_{ijt+1}}{\partial a_{ijt}} = \lambda_f(\theta_{ijt})D_{ijt+1} \text{ and } \frac{\partial W^N_{ijt+1}}{\partial \theta_{ijt}} = \lambda'_f(\theta_{ijt})a_{ijt}D_{ijt+1}, \tag{17}$$

where D_{ijt+1} is recursively defined by

$$D_{ijt+1} \equiv \frac{1}{\omega}\frac{u_{nijt+1}}{u_{cijt+1}}\left(\frac{n_{ijt+1}}{n_{it+1}}\right)^{\frac{1}{\omega}-1} \frac{1}{n_{it+1}} + Q_{t+1,t+2}\,(1-\sigma)^2 D_{ijt+2}. \tag{18}$$

Notice that as ω is increased, labor becomes more substitutable across firms and, all else equal, the D_{ijt+1} terms fall and, hence, so do the derivatives in (17). As ω converges to infinity, labor becomes perfectly substitutable across firms, and the derivatives in (17) converge to zero. In this case, each firm's market power disappears and the equilibrium allocations and prices converge to those in the competitive equilibrium.

1.2.3 An Analogous Participation Constraint with Monopolistic Competition

Here we relate how we are using the participation constraint to an analogous way one could use it in a standard monopolistic competition framework. Recall that standard analyses of monopolistic competition derives 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. Instead, we could simply derive the first order conditions for a consumer in an almost symmetric equilibrium and impose the equality of marginal utility in buying another good from firm j at price p_j and the marginal utility of buying another good from the common market with price p for all goods $k \neq j$.

To that end consider a consumer's problem with a standard utility function over differentiated goods, namely to choose c_j to maximize $u(c)$ subject to $\sum_j p_j c_j \leq pc$, where $c = \left(\sum_j c_j^{\frac{\omega-1}{\omega}} dj\right)^{\frac{\omega}{\omega-1}}$

and p is the CES price index associated with this utility function. If we combine the first order conditions any two goods c_j and c_k that are consumed in positive amounts then

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j = u'(c)(\frac{c_k}{c})^{-\frac{1}{\omega}} - \lambda p_k, \tag{19}$$

where λ is the multiplier on the budget constraint. Now if we consider an almost symmetric equilibrium in which $c_k = c$ for all goods $k \neq j$ and p the associated price index then (19) implies that for good j to be consumed in positive amounts it must yield as least as much net marginal utility as all other goods which have price p so that

$$u'(c)\left(\frac{c_j}{c}\right)^{-\frac{1}{\omega}} - \lambda p_j \geq u'(c) - \lambda p, \tag{20}$$

which is a participation constraint for firms under monopolistic competition. Notice the similarity of this participation constraint for attracting a consumer to buy a good for firm j, and the participation constraint for attracting a searching consumer to the market (θ_j, w_j), namely (14).

1.3 Firms with Standard Capital

We begin by assuming that firms have the standard type of capital, which we sometimes refer to as *putty-putty* capital, in which each firm can instantly move along the nested CES production frontier by adjusting the ratios of aggregate low-skilled labor to capital and high-skilled labor to capital by hiring differing vectors of low-skilled and high-skilled labor of each ability level or investing in differing amounts of capital.

Given initial capital k_0 and an exogenous sequence of prices of investment goods in units of consumption goods, $\{q_t\}$, each firm chooses chooses sequences of market tightnesses $\{\theta_{ijt}\}$ for all i, total vacancies $\{\mu_i a_{ijt}\}$ for all i, to workers to employ $\{\mu_i n_{ijt+1}\}$, present value of wages $\{W^F_{ijt+1}\}$ for these workers, and new capital $\{k_{t+1}\}$ to maximize

$$\sum_{t=0}^{\infty} Q_{0,t}\left\{F\left(k_t, n_{\ell t}, n_{ht}\right) - q_t x_t - \sum_i W^F_{ijt}\lambda_f(\theta_{ijt-1})\mu_i a_{ijt-1}\right\} \tag{21}$$

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

$$\mu_i n_{ijt+1} \leq (1-\sigma)\mu_i n_{ijt} + \lambda_f(\theta_{ijt})\mu_i a_{ijt} \text{ all } i \tag{22}$$

and the participation constraints for attracting low-skilled and high-skilled consumers,

$$\lambda_w(\theta_{ijt})\left[W^F_{ijt+1}(w_{ijt}) + W^N_{ijt+1}\right] \geq \mathcal{W}_t(\theta_{ijt}, w_{ijt+1}) \text{ all } i, \tag{23}$$

where $\bar{n}_{st} = \left[\sum_{i\in I_s} z_i(\mu_i n_{ijt})^{\frac{1+\phi}{\phi}}\right]^{\frac{\phi}{1+\phi}}$ for $s = \ell, h$.

After some manipulation detailed in the Appendix, the first order conditions imply an Euler equation for capital accumulation

$$q_t = Q_{t,t+1}\left[F_{kt+1} + (1-\delta)q_{t+1}\right], \tag{24}$$

a vacancy-posting condition that generalizes the free-entry condition for vacancies familiar from competitive search models

$$\frac{\kappa_i}{\lambda_f(\theta_{ijt})} = (1-\eta)Q_{t,t+1}\left[y_{ijt+1} + W^N_{ijt+1} + \frac{\partial W^N_{ijt+1}}{\partial \theta_{ijt}}\right] + Q_{t,t+1}\frac{\partial W^N_{ijt+1}}{\partial a_{ijt}}a_{ijt} \text{ for all } i, \tag{25}$$

and a wage equation

$$W^F_{ijt+1} = y_{iht+1} + \frac{\partial W^N_{ijt+1}}{\partial a_{ijt}}a_{ijt} - \frac{\kappa_i}{Q_{t,t+1}\lambda_f(\theta_{ijt})} \text{ for all } i, \tag{26}$$

where y_{ijt+1}, the multiplier on the transition law (5), is the present value of the marginal product of labor of a mass one of workers of type i during a match in which the worker starts to work at $t+1$, and $(1-\sigma)^s$ of which remain employed at $t+s+1$, so y_{ijt+1} is given by

$$y_{ijt+1} \equiv F_{ijt+1} + Q_{t+1,t+2}(1-\sigma)y_{ijt+2} = \sum_{r=0}^{\infty} Q_{0,t+r+1}(1-\sigma)^r F_{ijt+r+1}, \tag{27}$$

where $F_{ijt+1} \equiv \frac{\partial F(k_{t+1},\bar{n}_{\ell t+1},n_{ht+1})}{\partial n_{ijt+1}}\frac{\partial \bar{n}_{it+1}}{\partial(\mu_i n_{ijt+1})}\frac{\partial(\mu_i n_{ijt+1})}{\partial n_{ijt+1}} = \mu_i\frac{\partial F_{ijt+1}}{\partial(\mu_i n_{ijt+1})}$ for all i.

Now, as we discussed earlier, the critical terms that capture monopsony power are $\partial W^N_{ijt+1}/\partial a_{ijt}$ and $\partial W^N_{ijt+1}/\partial \theta_{ijt}$ which capture that firm j realizes that when it increases its vacancies a_{ijt} and, thereby, the associated market tightness θ_{ijt}, on the margin, it will attract a higher fraction of the current searchers s_{it} from each family i to it and thus must compensate consumers accordingly. When we substitute the value for these derivatives from (17) into the vacancy posting condition (25) and the wage equation we get the form for them we use when solving the model, namely,

$$\frac{\kappa_i}{\lambda_f(\theta_{it})} = (1-\eta)Q_{t,t+1}\left[y_{it+1} + W^N_{it+1} + \lambda_f(\theta_{it})a_{it}D_{it+1}\right] \text{ for all } i \tag{28}$$

and

$$W^F_{it+1} = y_{it+1} + \lambda_f(\theta_{it})a_{it}D_{it+1} - \frac{\kappa_i}{Q_{t,t+1}\lambda_f(\theta_{it})} \text{ for } b \in \{\ell, h\} \text{ and all } i, \tag{29}$$

where D_{it+1} is given by (18). Since all firms are symmetric, the resulting equilibrium will be symmetric in that $n_{ijt} = n_{it}$ for all j, so that n_{it} will not only be a index of disutility of work but will also be the measure the total employment of consumers of type i. Since firms offer the same wages and market tightness so $\theta_{ijt} = \theta_{it}$ and families send an equal measure of searchers to all firms, so $s_{ijt} = s_{it}$ and the budget constraint and the transition law for employed consumers reduce to

$$\sum_{t=0}^{\infty} Q_{0,t}c_{it} \leq \sum_{t=0}^{\infty} Q_{0,t}W^F_{it}\lambda_w(\theta_{it-1})s_{it-1} + \psi_i\Pi_0$$

and $n_{it+1} = (1-\sigma)n_{it} + \lambda_w(\theta_{it})s_{it}$, where Π_0 denotes the present value of profits of firms, namely, the value of (21).

For an exogenous sequence of investment goods prices, $\{q_t\}$, a symmetric monopsonistic equilibrium with standard capital with $k_{j0} = k_0$ for all j is a collection of allocations of consumption, employment, searching, capital $\{c_{it}, n_{it}, s_{it}, \bar{n}_{it}, k_t\}$ vacancies and market tightnesses, $\{a_{it}, \theta_{it}\}$, and prices $\left\{W^F_{it+1}, Q_{0t}\right\}$ such that at these prices and allocations, i) consumer's decisions are optimal for each family i, ii) firm decisions are optimal, and iii) markets clear.[12]

Adding a Minimum Wage: Transition Experiment. So far, we have considered an economy without a minimum wage. Now we suppose that there is a minimum flow wage that a firm can offer of the form $w_{ijt} \geq \bar{w}$. Given a sequence of intertemporal prices $\{Q_{t,s}\}$ this flow constraint implies a constraint on the present value of wages of the form

$$W^F_{ijt} \geq \bar{W}^F_t \equiv \bar{w} + (1-\sigma)Q_{t,t+1}\bar{w} + (1-\sigma)^2Q_{t,t+2}\bar{w} + \ldots \tag{30}$$

Note that $\bar{W}^F_t$ is the smallest present value of wages consistent with meeting the flow constraint $w_{ijt} \geq \bar{w}$ in each period of the match. This present value depends on time because the intertemporal prices do. We add the constraint (30) to the firm's problem and note that the consumer's problem is identical to the earlier case. If for a consumer of type i the minimum wage constraint is slack, the first order conditions are the same as before. When the minimum wage constraint binds we set

[12] Here we have considered an exogenous sequence of investment good prices. Note, as is standard, we can interpret these prices as coming from an economy in which consumption goods are produced by $F(k_{ct}, n_{cut}, n_{cst})$ and investment goods are produced by $A_{xt}F(k_{xt}, n_{xut}, n_{xst})$ where, critically, these two production functions are the same up to multiplicative productivity shocks A_{xt}, which implies that the relative prices of investment to consumption is pinned down solely by the productivity shocks.

$W_{ijt}^F = \bar{W}_t^F$ and, as we show in the Appendix, the vacancy posting condition reduces to

$$\frac{\kappa_i}{\lambda_f(\theta_{it})} = (1-\eta)Q_{t,t+1}\left[\frac{\left(y_{it+1} - \bar{W}^F\right)W_{it+1}}{(1-\eta)W_{it+1} - \eta D_{it+1}\lambda_f(\theta_{it})a_{it}}\right].$$

When we consider the transition dynamics we start in a steady state and suppose there is a one-time unanticipated and permanent increase in the minimum wage relative to the wage in the steady state. Firms are allowed to fire existing workers if they wish. All workers retained by the firm and must be paid the larger of their existing wage and the minimum wage and all new hires must be paid at least the minimum wage.

1.4 Adding Putty-Clay Capital

Now consider an environment in which capital is putty-clay. The basic idea behind the model is most easily understood when all low-skilled consumers and all high-skilled 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 (or structure) with any ratio of low-skilled labor to capital and high-skilled labor to capital that lies on the frontier of the production function in (1) with just the aggregate inputs. However, once the machine is built it is clay-like, in that this machine uses a fixed amount of low-skilled labor and high-skilled labor to operate it, that is, it nested CES ex-ante but Leontief ex-post. In the short run, this implies that for a given stock of machines, demand for low-skilled labor and high-skilled labor is inelastic, because there is no way to substitute between existing capital and either type of labor. Over time, new capital goods embodying new labor to capital ratios can be installed so that, in the long run, firms can substitute away from the type of labor that policy makes more expensive, here low-skilled labor. The extension to the case when they are many ability levels z among the low-skilled and the high-skilled is immediate.

1.4.1 Two-Type Case

Suppose that all low-skilled workers and all high-skilled workers have the same level of ability, say $z = 1$, so that there are only two types of workers, low-skilled and high-skilled denote ℓ and h with measures μ_ℓ and μ_h respectively. For ease of notation only we set $\mu_\ell = \mu_h = 1$. We again suppress the firm index j. Here a type of capital is a $v = (v_\ell, v_h) \in V_\ell \text{x} V_h$ which specifies both the low-skilled labor to capital ratio, v_ℓ, and the high-skilled labor to capital ratio, v_h, to run *one* machine of type v at full utilization and, hence produce

$$f(v_\ell, v_h) \equiv F(1, v_\ell, v_h) = \left[\mu(\psi_\ell v_\ell)^\sigma + (1-\mu)(\lambda + (1-\lambda)(\psi_h v_h)^{\sigma/\rho}\right]^{1/\sigma}.$$

To run $k_t(v_\ell, v_h)$ of type $v = (v_\ell, v_h)$ at full utilization it takes $(v_\ell k_t(v_\ell, v_h), v_h k_t(v_\ell, v_h))$ units of low-skilled and high-skilled capital and hence produce $k_t(v_\ell, v_h) f(v_\ell, v_h)$ units of output.

More generally, if $\{k_t(v_\ell, v_h)\}$ units of capital of each type (v_ℓ, v_h) type is combined with $(n_{\ell t}(v_\ell, v_h), n_{ht}(v_\ell, v_h))$ units of low-skilled and high-skilled labor then total output is

$$y_t = \int_{v_\ell, v_h} \min\left[\frac{n_{\ell t}(v_\ell, v_h)}{v_\ell}, \frac{n_{ht}(v_\ell, v_h)}{v_h}, k_t(v_\ell, v_h)\right] f(v_\ell, v_h) dv_\ell dv_h \tag{31}$$

and the capital accumulation law for each type v of capital is

$$k_{t+1}(v) = (1-\delta)k_t(v) + x_t(v), \tag{32}$$

where $x_t(v)$ is investment in capital of type v. Given some initial vector of capital $\{k_0(v)\}$ that the firm owns, the firm chooses sequences of market tightnesses $\{\theta_{\ell t}, \theta_{ht}\}$, total vacancies $\{a_{\ell t}, a_{ht}\}$, total of workers $\{n_{\ell t+1}(v), n_{ht+1}(v)\}$ for each type of capital v, present value of wages $\left\{W^F_{\ell t+1}, W^F_{ht+1}\right\}$ and investment $\{x_t(v)\}$ for each type of capital to maximize

$$\sum_{t=0}^{\infty} Q_{0,t}\left\{\int_v [F(k_t(v), n_{ht}(v), n_{\ell t}(v)) - q_t x_t(v)]dv - W^F_{\ell t}\lambda_f(\theta_{\ell t-1})a_{\ell t-1} - \kappa_\ell a_{\ell t} - W^F_{ht}\lambda_f(\theta_{ht-1})a_{ht-1} - \kappa_h a_{ht}\right\} \tag{33}$$

subject to the transition laws for both types of workers (38), the participation constraints for workers (14), along with the transition law for each type of capital (32), the adding up constraints of uses of low-skilled and high-skilled labor, $n_{\ell t} \leq \int n_{\ell t}(v)dv$ and $n_{ht} \leq \int n_{ht}(v)dv$, the Leontief constraints on labor

$$n_{\ell t}(v) \leq v_\ell k_t(v) \text{ and } n_{ht}(v) \leq v_h k_t(v), \tag{34}$$

and the nonnegativity constraints on each type of investment $x_t(v) \geq 0$. To understand the constraints in (34), note that with a given capital stock $k_t(v_\ell, v_h)$ of type $v = (v_\ell, v_h)$ since the output of this stock is given by

$$\min\left[k_t(v_\ell, v_h), \frac{n_{\ell t}(v_\ell, v_h)}{v_\ell}, \frac{n_{ht}(v_\ell, v_h)}{v_h},\right]$$

if the firm uses more $n_{\ell t}(v) > v_\ell k_t(v)$ units of low-skilled labor with it, then the excess labor $n_{\ell t}(v) - v_\ell k_t(v)$ is wasted, so this is never optimal—likewise when $n_{ht}(v) > v_h k_t(v)$. Hence, we can impose the Leontief constraints in (34) and drop the min function. The non-negativity constraint $x_t(v) \geq 0$, implies that firms cannot disassemble their existing types of capital. This friction is key to generating frictional adjustment; otherwise, the model would collapse to the putty-putty model described in the previous section.

We now turn to characterizing some properties of the putty-clay economy. When we consider the transition dynamics we start in a steady state. Here we ask if we start in a steady state and there is a surprise increase in the minimum wage relative to the wage in the steady state, what are the conditions that guarantee that all existing capital is used and existing workers are used. We assume that capital left unused in period depreciates at the same rate δ as capital that is used.

In the model's steady state there is only one type of capital being used, but we prefer to state a slightly more general condition in which the state of the firm is a vector of capital stocks $\{\bar{k}(v)\}$, low-skilled workers $\bar{n}_\ell$, high-skilled workers $\bar{n}_h$, that satisfy

$$\int v_\ell \bar{k}(v)dv = \bar{n}_\ell \text{ and } \int v_h \bar{k}(v)dv = \bar{n}_h \tag{35}$$

as well as steady-state wages w_ℓ and w_h for low-skilled and high-skilled workers. Note that (35) simply means that the firm has just enough low-skilled and high-skilled workers to operate each unit of capital at full capacity.

In the next proposition we develop a sufficient condition for firms to choose to operate their existing capital at full capacity and retain their existing workers at the larger of their existing wages and the new minimum wage. In particular, consider a steady state with wages (w_ℓ, w_h) and surprise increase in the minimum wage to $\bar{w}$ and define $\bar{w}_b = \max\{w_b, \bar{w}\}$ to be the larger of the minimum wage and the existing wage for each type.

***Proposition* 1.** *Under (35), a sufficient condition for a unit of capital $\bar{k}(v)$ to be operated at full capacity is that the static profits from doing so are positive in that*

$$f(v) - \bar{w}_\ell v_\ell - \bar{w}_h v_h \geq 0. \tag{36}$$

Proof. Since we are only interested in developing a sufficient condition for capital to be operated at full capacity we proceed indirectly. To that end, imagine for a moment that workers can be *furloughed* in the sense that i. they can be laid off from the firm without pay, ii. they exogenously separate from the firm at the same rate σ as non-furloughed workers and, crucially, iii. as long as they have not separated they can recalled without the firm having to recruit them again. Such workers can also search while they are furloughed but this option will never be taken in our equilibrium. Under this assumption and that unused capital depreciates at the same rate as capital that is used, a firm's decision to lay off either capital or furlough workers is a static decision.

Now in our model firms cannot furlough workers, but rather can only fire them. Note, however, that firing a worker is more costly to the firm than furloughing a worker, because furloughing allows for the possibility of recalling the worker at no cost whereas hiring workers in the future to replace currently fired workers involves costly vacancy posting. It is then immediate that if a firm prefers not to furlough any workers it also prefers not to fire any of them. So if we find conditions under which a firm prefers not to furlough workers then under these conditions a firm prefers not to to fire workers.

Let us now turn to the details of the proof. At an intuitive level, under the (35) condition we can think that for each type of capital v, the $\bar{k}(v)$ units of capital have exactly $n_\ell(v) = v_\ell \bar{k}(v)$ and $n_h(v) = v_h \bar{k}(v)$ workers ready to operate them. Hence, regardless of the size of the surprise increase in the minimum wage there is no other piece of capital for which these workers could be reassigned without leading to an excess of workers for that other piece of capital. So the only decision is to operate the given $\bar{k}(v)$ units of capital with the workers that are ready to operate them or let this capital be left unused and let the associated workers be furloughed. Now if the inequality in (36) is strict, then the firm makes strictly positive profits from operating it at full capacity, and if this inequality holds as an equality the firm is indifferent to operating it at any level, while if the reverse inequality to (36) holds, then operating this machine makes negative profits statically, and it is clearly optimal to shut it down completely. ■

In the monopsony problem with putty-clay capital, namely (33), the firm solves a problem with state consisting of a measure of capital goods by types, the possibly infinite state $\{k_t(v)\}$ together with the aggregate amounts of high-skilled and low-skilled labor $(n_{ht}, n_{\ell t})$. We now turn to showing that under a sufficient condition this problem can be drastically simplified. To that end, let $\hat{V}_t$ be the set of types of capital which are used at t in the monopsony problem, that is, $\hat{V}_t \equiv \{v = \{(v_\ell, v_h)|k_t(v) > 0\}$ where $\{k_t(v)\}$ is part of the solution to (33). We first define a restricted problem in which we replace the state $\{k_t(v)\}, n_{ht}, n_{\ell t}$ with a simple state consisting of two aggregates $y_t, n_{\ell t}$ where

$$y_t \equiv \int_v \min\left[k_t(v), \frac{n_{\ell t}(v)}{v_\ell}, \frac{n_{ht}(v)}{v_h}\right] f(v)dv \text{ and } n_{\ell t} \equiv \int n_{\ell t}(v)dv.$$

The *restricted* problem of the firm is to choose sequences of market tightnesses $\{\theta_{\ell t}, \theta_{ht}\}$, vacancies $\{a_{\ell t}, a_{ht}\}$, number of workers $\{n_{\ell t+1}, n_{ht+1}\}$, present value of wages $\left\{W^F_{t+1}, W^F_{ht+1}\right\}$ and investment $\{x_t(v)\}$ to maximize

$$\sum_{t=0}^{\infty} Q_{0,t}\left\{y_t - q_t\int_v x_t(v)dv - \sum_{i\in\{\ell,h\}} \lambda_f(\theta_{it-1})a_{it-1}W^F_{it} - \kappa_i a_{it}\right\} \tag{37}$$

subject to the participation constraints, the nonnegativity constraints on investment along with the following transition laws for labor and output

$$n_{it+1} = (1-\delta)n_{it} + \int v_i x_t(v)dv \tag{38}$$

$$y_{t+1} = (1-\delta)y_t + \int x_t(v)f(v)dv \tag{39}$$

***Proposition* 2.** *If for every $v \in \hat{V}_t$, the solution to the original firm problem (33) has full utilization of all capital in that $n_{it}(v) = v_i k_t(v)$ for $i = \ell, h$ then a solution to the restricted firm problem (37) is a solution to the original firm problem.*

Proof. Clearly the constraint set of the original problem is a relaxed version of the constraint set of the restricted problem. Hence, the maximized value of the original problem is at least as large as that of the restricted problem. To prove our result we only need show that, if a solution to the original problem has full utilization of capital in that

$$k_t(v) = \frac{n_{it}(v)}{v_i} \text{ for } i \in \{\ell, h\}, \tag{40}$$

then it satisfies the constraints of the restricted problem. Because when that is true then a solution to the restricted problem is feasible for original problem and, hence, the solution to the restricted problem coincides with that of the original problem.

Now a solution to the original problem that satisfied the capital accumulation law (32) and (40) must satisfy

$$\frac{n_{it+1}(v)}{v_i} = (1-\delta)\frac{n_{it}(v)}{v_i} + x_t(v). \tag{41}$$

Consider the transition law for output $y_t \equiv \int_v [n_{it}(v)/v_i] f(v)dv$ given by (39). Multiply (41) by $f(v)$ and then integrate over v

$$\int \frac{n_{it+1}(v)}{v_i} f(v)dv = (1-\delta)\int \frac{n_{it+1}(v)}{v_i} f(v)dv + \int x_t(v)f(v)dv,$$

which using the definition of y_t gives $y_{t+1} = (1-\delta)y_t + \int x_t(v)f(v)dv$ which is this transition law for output. Next, consider the transition law for aggregate labor, $n_{it} \equiv \int n_{it}(v_\ell, v_h)dv$, given by (38). To show that is satisfied, multiply (41) by v_i and then integrate over v_i to get

$$\int n_{it+1}(v)dv = (1-\delta)\int n_{it}(v)dv + \int v_i x_t(v)dv,$$

which using the definition of aggregate labor n_{it} gives $n_{it+1} = (1-\delta)n_{it} + \int v_i x_t(v)dv$ which is the transition law for aggregate labor. Hence, a solution to the original firm problem is feasible for the restricted firm problem and, hence, must solve it. ■

We next show that in the restricted monopsony problem a firm chooses to invest in at most one type of capital. To set up that result, note that the first order condition of the restricted problem with respect to $x_t(v)$ is given by

$$q_t = \eta_{yt} f(v) - \sum_{i\in\{\ell,h\}} \eta_{nit} v_i + \eta_{xt}(v) \text{ or } \eta_{xt}(v) = q_t - \eta_{yt} f(v) + \sum_{i\in\{\ell,h\}} \eta_{nit} v_i \geq 0 \tag{42}$$

and the complementary slackness condition for $x_t(v)$ is $\eta_{xt}(v)x_t(v) = 0$ where η_{nit} and η_{yt} are the multipliers on the transition laws for labor and output (38) and (39), and $\eta_{xt}(v)$ is the multiplier on the nonnegativity constraint on investment of type v.

***Proposition* 3.** *In the restricted monopsony problem at most one type of capital v has $x_t(v) > 0$.*

Proof. From the complementary slackness condition, investment $x_t(v) > 0$ only if $\eta_{xt}(v) = 0$. Furthermore, $\eta_{xt}(v) \geq 0$ so zero is its minimum value. Thus, to prove the proposition we need to show that $\eta_{xt}(v)$ has a unique minimum or equivalently that the term $h(v)$ defined to be $-\eta_{xt}(v)$ has a unique maximum. That is, $h(v) = \eta_{yt} f(v) - \sum_{i\in\{\ell,h\}} \eta_{nit} v_i - q_t$, has a unique maximum in v. Clearly, since $F(k, n_\ell, n_h)$ is strictly concave in (k, n_ℓ, n_h) then $F(1, n_\ell, n_h) = f(v_\ell, v_h)$ is strictly concave in (v_ℓ, v_h). Next, we show that if $f(v)$ is strictly concave in v then so is $h(v)$ But this is true because the second derivatives of h are proportional to the second derivatives of f with the positive proportionality constant η_{yt}. ■

The logic of Propositions 1, 2, and 3 generalizes that used in Atkeson and Kehoe (1999) in a model of energy use.

1.4.2 General Case

The extension to the general case is purely a matter of notation, with no conceptual differences. In particular, in the general case we have $\bar{n}_{st} = \left[\sum_{i\in I_s} z_i(\mu_i n_{ijt})^{\frac{1+\phi}{\phi}}\right]^{\frac{\phi}{1+\phi}}$ for $s = \ell, h$ are the aggregate labor inputs associated with each broad skill group. For notational simplicity, we again suppress j subscripts, with the understanding that each firm j internalizes the effect of its choices on W^N_{it+1} using the same formulas as in the neoclassical model.

Here with many types of labor $i \in I$, the continuum of capital types are indexed by $v = \{v_i\}_{i\in I}$ which is a vector with an entry for each type of labor i. A quantity $k(v)$ units of capital of type v provides $k(v)$ 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. If $n_i > k(v)v_i$, then the excess workers remain idle; if $n_i < k(v)v_i$, the excess capital remains idle. Therefore, n_i/v_i controls how much capital a worker of type i can support. Conditional on providing such capital services, the unit of capital produces $f(v)$ units of output, where $f(v) = F(k, \{n_i\})/k = F(1, \{n_i/k\}) = F(1, \{v_i\})$, where F is the production function from the neoclassical model.

The total output produced from type-v capital is $y(v) = \min\left[k(v), \{n_i/v_i\}\right] f(v)$, and the total output of the firm is

$$y_t = \int \min\left[k(v), \{n_i/v_i\}\right] f(v) dv.$$

Firms invest $x_t(v)$ units of output to accumulate type-v capital according to the capital accumulation equation (32) and are subject to the nonnegativity constraints $x_t(v) \geq 0$.

Propositions 1, 2, and 3 immediately extend to this general model. In our quantitative work below we use the extension of these three propositions to the general case to greatly simplify the computation.

1.5 Steady State

Here we discuss the steady state and some simple results that can be obtained in it.

First, note that it is immediate that the steady state of the model with standard capital and putty-clay capital are identical. The intuition is that once the factor prices and intertemporal prices facing the firms settle down, firms start investing in the unique type capital that is ideally suited to it at those prices and the firms let all past capital depreciate away. Eventually, all the old capital stock is replaced and in a deterministic steady state firms invest in exactly the same type of capital as they would have if they had standard capital rather than putty clay capital.

Second, in our baseline model we focus on GHH preferences with $u(c, s, n) = U\left[c - h(s) - v(n)\right]$. Consider the steady state of the model and let $\tilde{\sigma} = \frac{1-\beta(1-\sigma)}{1-\beta(1-\sigma)^2}\frac{\sigma}{\omega}$. As we show in the Appendix, in the steady state, the Euler equation for capital is

$$q\left[\frac{1}{\beta} - (1-\delta)\right] = F_k(k, n_\ell, n_h), \tag{43}$$

the vacancy posting condition is

$$\frac{\kappa_i}{\lambda_f(\theta_i)} = \frac{\beta(1-\eta)}{1-\beta(1-\sigma)}\left[F_{ni} - \left(1 + \frac{\tilde{\sigma}}{\omega}\right) v'(n_i)\right], \tag{44}$$

the search condition is

$$h'(s_i) = \beta\lambda_w(\theta_i)\frac{\eta\left[F_{ni} - (1+\frac{\tilde{\sigma}}{\omega})v'(n_i)\right]}{1-\beta(1-\sigma)}, \tag{45}$$

the steady-state wages satisfy

$$w_i = \eta \left[F_{ni} - \frac{\tilde{\sigma}}{\omega} v'(n_i) \right] + (1-\eta) v'(n_i), \tag{46}$$

the steady-state law of motion for labor is $\lambda_w(\theta_i)s_i = \sigma n_i$, and the consumption comes from the budget constraint. For later, it will be useful note that we can combine the vacancy posting condition and the wage equation so see that the implied markdowns for each type are given by

$$\frac{w_i}{F_{ni}} = \left(1 + \underbrace{\frac{\kappa(\theta_i)}{\frac{\eta}{1-\eta}\kappa(\theta_i) + v'(n_i)}}_{\text{efficient component}} + \underbrace{\frac{\frac{\tilde{\sigma}}{\omega} v'(n_i)}{\frac{\eta}{1-\eta}\kappa(\theta_i) + v'(n_i)}}_{\text{monopsony component}} \right)^{-1}, \tag{47}$$

where $\kappa(\theta_i) = (r+\sigma)\kappa_i/\lambda_f(\theta_i)$ and $r = 1/\beta - 1$. The efficient component of the markdown is the same that arises in a competitive version of the model, namely, the amount needed so that firms recoup their vacancy posting costs and, hence, earn zero expected profits on each vacancy. More interesting is the monopsony component, which captures that monopsonists push wages below their competitive levels and, hence, their markdowns are larger than the competitive markdowns. In our quantitative model, we will find that the overwhelming majority of the markdown in is due to the monopsony distortion.

With GHH preferences the system splits into blocks. The equilibrium allocations, namely, (θ_i, n_i, s_i) and k can be solved from the Euler equation for capital, vacancy posting condition, the search condition, the law of motion for labor. Wages are then given from (46) and, finally, consumption is then residually determined given from the budget constraint.

In our economy the only distortion is from monopsony power which directly distorts the vacancy posting condition, the search condition and wages via the term $\sigma v'(n_i)/\omega$. Notice as labor becomes more substitutable, in that ω increases, the distortion to both vacancy posting and searching fall. As labor becomes perfectly substitutable, in that ω increases to infinity, these terms converge to zero and the allocations and wages in this steady state converge to those of the competitive equilibrium.

2 A Policy Comparison: EITC and Other Tax Policies

A major goal of the minimum wage is to increase the (after-tax) income of those agents who earn relatively low wages. Moreover, a common rationale for why minimum wages may be beneficial is that they help offset the monopsony power of firms which leads wages to be set at an inefficiently low level. An alternative to the minimum wage to accomplish these goals is an Earned Income Tax

Credit (EITC) or, more generally, a progressive income tax. In our quantitative analysis, after we evaluate the minimum wage, we evaluate how such policies perform in our environment both in terms of increasing the income of agents and offsetting to some extent the monopsony power of firms. Here we describe how allowing such policies alters the equilibrium conditions.

All of the policies we consider can be represented by a potentially non-linear tax function of period wages w given by $T(w)$ which implies that after tax wage income is $A(w) = w - T(w)$. It is convenient to write $W_t^F = d_t w_t$ where $d_t \equiv 1 + (1-\sigma)Q_{t,t+1} + (1-\sigma)^2 Q_{t,t+2} + \ldots$ is the scaling factor that converts flow wages w_t over a match to their expected present values W_t^F and, for any tax schedule $T(w)$ on flow wages let $A_{ct}(W_t^F) \equiv d_t A(w_t)$ be the present value of after-tax wage income under the tax schedule $T(w)$ and the before-tax wage income w_t for the course of a match.

We incorporate a tax schedule that gives rise to after-tax income of the form $A_{ct}(W_t^F) \equiv d_t A(w_t)$ as follows. On the consumer side, the present value budget constraint of a family becomes

$$\sum_{t=0}^{\infty} Q_{0,t} c_{it} \leq \sum_{t=0}^{\infty} Q_{0,t} \sum_j \lambda_w(\theta_{ijt-1}) s_{ijt-1} A_{ct}(W_{ijt}^F) + \psi_i \Pi_0,$$

the first order condition for searching s_{ijt} becomes

$$-\frac{u_{sit}}{u_{cit}} = Q_{t,t+1} \lambda_w(\theta_{ijt}) \left(A_{ct}(W_{ijt}^F) + W_{ijt+1}^N \right) + \frac{\chi_{ijt}}{u_{cit}} \text{ for all } j,$$

where χ_{ijt} is the multiplier on the nonnegativity constraint on s_{ijt}. On the firm side, the participation constraint for firm j becomes

$$\lambda_w(\theta_{ijt}) \left(A_{ct}(W_{ijt+1}^F) + W_{ijt+1}^N \right) \geq \mathcal{W}_{it}(\theta_{it}, A_{ct}(W_{it+1}^F).$$

With standard capital, the firm's first order conditions give rise to the same Euler equation (24) as before. The wage and market tightness determined by

$$y_{it+1} = \frac{1-\eta}{\eta} \left[\frac{A_c(W_{it+1}^F) + W_{it+1}^N}{A_c'(W_{it+1}^F)} \right] - \lambda_f(\theta_{it}) \frac{a_{it} D_{it+1}}{A_c'(W_{it+1}^F)}$$

and

$$\eta \frac{\kappa_i}{Q_{t,t+1} \lambda_f(\theta_{it})} = (1-\eta) \frac{A_c(W_{it+1}^F) + W_{it+1}^N}{A_c'(W_{it+1}^F)}.$$

In a symmetric (across firms) steady state with GHH preferences these conditions simplify to an implicit wage equation

$$w_i = F_{ni} - \frac{\widetilde{\sigma}}{\omega} v'(n_i) - \frac{1-\eta}{\eta} \left(\frac{A(w_i) - v'(n_i)}{A'(w_i)} \right), \tag{48}$$

a vacancy posting condition

$$\frac{\kappa_i}{\lambda_f(\theta_i)} = \frac{\beta}{1-\beta(1-\sigma)} \left(F_{ni} - w_i - \frac{\widetilde{\sigma}}{\omega} \frac{v'(n_i)}{A'(w_i)} \right), \tag{49}$$

and a search condition

$$h'(s) = \beta \lambda_w(\theta) \frac{A(w) - v'(n)}{1-\beta(1-\sigma)}, \tag{50}$$

where $d = 1/\left[1-\beta(1-\sigma)\right]$.

3 Parameterization

We calibrate the model in order to match key empirical targets that inform the mechanisms described above. Whereas most of these parameters have straightforward empirical counterparts, there are no available estimates of the long-run elasticities of substitution ρ and α that are suitable for our model.[13] Therefore, a key component of our calibration is to estimate these long-run elasticities. We then validate our assumption that the short-run elasticities of substitution are smaller than these long-run elasticities.

3.1 Estimating Long-Run Elasticities of Substitution

We estimate the long-run elasticities of input substitution using sector-level time series on the relative price of capital and the distribution of income across households.

Data Sources. We measure the sector-level relative price of capital q_{st} using the Tornqvist index

$$\Delta \log q_{st} = \sum_{a=1}^{A} \omega_{ast} \Delta \log q_{at}, \tag{51}$$

where a indexes a type of capital good, ω_{ast} is the Tornqvist share of nominal investment expenditures on asset a, and $\Delta \log q_{at}$ is the growth rate of the price of asset a relative to the price index of

[13] The estimates in Krusell et al. (2000) and the following literature are derived from a model without search frictions, so the first-order condition for labor demand equates wages to marginal products. Inference is more complicated in our model because labor demand also depends on labor market tightness and thus the labor supply decisions of households. Nevertheless, we end up finding similar estimates to Krusell et al. (2000); our preferred estimates are $\rho = 2.03$ and $\alpha = 0.44$, while Krusell et al. (2000) estimate $\rho = 1.67$ and $\alpha = 0.67$.

consumption. Hence, our relative price series q_{st} varies both across sectors (due to differences in the mix of capital goods) and over time (due to changes in the relative prices of different goods). See Appendix C for details about how we measure these relative prices using the BEA Detailed Fixed Asset Tables. Using the same data, we also compute real investment expenditures i_{st} and the real value of the capital stock k_{st}. Our resulting series are available annually 1960-2019 for 61 sectors.

We combine these data with data on the distribution of households from the 1960-2000 U.S. decadal Censuses, the pooled 2010-2012 American Community Surveys (ACS), and the pooled 2017-2019 American Community Surveys (ACS). Our main outcome of interest is $\text{share}^{\text{col}}_{st}$, that is, the share of total income in a given sector s accruing to workers with at least a bachelor's degree in survey year t.[14] These data contain 195 sectors which are consistently defined over time. We then merge these data with the more aggregated capital price data using a manually-constructed crosswalk based on 2012 NAICS codes.

Results. Our goal is to isolate the effect of changes in the relative price of capital on various outcomes y_{st} using regressions of the form

$$\Delta y_{st} = b_0 + b_t + b_1 \Delta \log q_{st} + \varepsilon_{st}, \tag{52}$$

where b_0 is a constant, b_t is a time fixed effect, and b_1 is the coefficient of interest. We consider three different outcomes Δy_{st}: the change in sectoral log investment ($\Delta \log i_{st}$), the change in sectoral capital-to-wage bill ratio ($\Delta \log(k/wn)_{st}$), or the change in sectoral college income share ($\Delta \text{share}^{\text{col}}_{st}$). The response of the college income share is particularly informative about the degree of what Krusell et al. (2000) and others refer to as *capital-skill complementarity*. All else equal, if capital is more complementary to college-educated workers, then a decline in the price of capital would increase their labor income relatively more than non-college-educated workers' labor income. We control for aggregate trends in college attainment or labor-augmenting technical change through the time fixed effects b_t, so the coefficient b_1 is identified from cross-sectoral variation. Our key identifying assumption is that other determinants of y_{st} are orthogonal to changes in the relative price of capital q_{st} at the sectoral level.

For a sense of the underlying variation in the data, Figure 2 shows a scatterplot of the ten-year changes in the college income share $\Delta \text{share}^{\text{col}}_{st}$ against the ten-year changes in the relative price of investment Δq_{st} at the sector level. The scatterplot shows a clear negative relationship between these two variables. For example, a log-point decline in the relative price of capital is associated

[14] We restrict our sample to only include individuals with a moderate attachment to the labor force: over age 16 who report currently working 30 hours per week and reported working at least 29 weeks during the prior year. See Appendix A for details.

Figure 2: Scatterplot of Relative Price of Investment vs. College Income Share

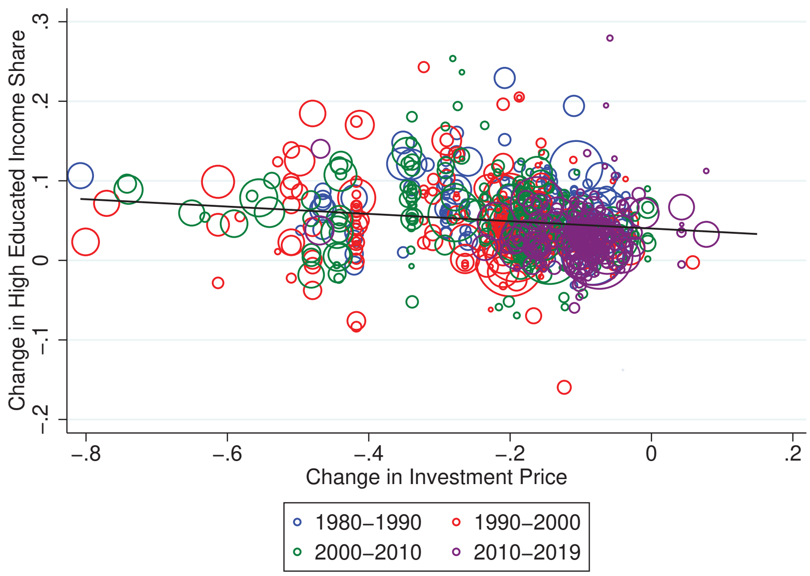

Notes: Scatterplot of ten-year change in college income share $\Delta \text{share}_{st}^{\text{col}}$ versus ten-year change in investment price $\Delta \log q_{st}$ using cross-industry variation. Size of circle reflects size of industry in year 2000 and different colors reflect different decades in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

with a 5 percentage point higher college income share. This negative relationship is robust across the four decades in our sample. Appendix C shows that a similar negative relationship exists between the relative price of investment and investment itself.

The left three panels of Table 1 quantify these relationships using our regression specification (52) over ten-year changes. Column (1) shows that a 1 log point decline in the relative price of capital would increase investment by about 1.37 log points; this user cost elasticity of -1.37 is in line with the consensus range estimated in the investment literature (see the discussion in Zwick and Mahon (2017)). Consistent with this finding, column (2) shows that a 1 log point decline in the relative price of capital would increase the capital-to-wage bill ratio by 0.47. Most importantly, column (3) shows that these changes are accompanied by a 8.1 percentage point increase in the share of income accruing to college workers, consistent with capital-skill complementarity. [15]

Appendix B contain three sets of robustness exercises on these empirical findings. First, we show that these results hold in each decade of our sample, and are therefore not driven by a particular subperiod of data. Second, we show that the results are robust to including the lagged relative

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

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

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

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

price of capital to capture path-dependence in response to past shocks. Third, we show that the investment results also holding when we use annual rather than decadal data—annual data are not available for the capital-to-wage bill ratio or college income share.

Calibration Targets. We will calibrate the model to match the twenty-year regression coefficients in columns (4) and (5) of Table 1.[16] Intuitively, these two targets are tightly related to the two long-run elasticities of substitution ρ and α. First, the response of the overall capital-to-wage bill ratio in column (4) is informative about the overall degree of substitutability between capital and labor. Second, the response of the college income share in column (5) is informative about the differential elasticity between college and non-college workers. See Appendix C for a more formal discussion of how these two coefficients map into values for ρ and α.

3.2 Calibration

We calibrate the model to match these long-run regression coefficients along with a number of other empirical targets. We calibrate the model in two steps. First, we exogenously fix a subset of parameter values based on external evidence. Second, we choose the remaining parameters—which are particularly important in determining the response to the minimum wage—in order to match

[16] We target twenty-year changes rather than ten-year changes because we assume the estimated responses reflect steady state comparisons of the model; see the discussion below for details.

Table 2: Fixed Parameters

Parameter	Description	Value
Households		
β	Discount factor	$(1.04)^{-1/12}$
γ	Labor disutility exponent	2.00
π_ℓ	Fraction of non-college households	0.69
ϕ	Elasticity of substitution across z	3.00
Firms		
δ	Capital depreciation (equipment + software, annualized)	0.15
Labor market frictions		
σ	Job destruction rate	2.8%
η	Elasticity of matching function w.r.t. vacancies	0.50

Notes: Parameters exogenously fixed in the calibration. A model period is one month.

moments in the data. We set a model period to be one month in order to adequately capture worker flows in the U.S. labor market.

Fixed Parameters. Table 2 shows the parameters we exogenously fix. We set the discount factor $\beta = (1.04)^{-1/12}$ so that the annual real interest rate is $r = 4\%$. We assume that the utility function takes the Greenwood, Hercowitz and Huffman (1988) form[17]

$$\sum_{t=0} \beta^t \log \left(c_{it} - \chi \left[\frac{n_{it}^{1+1/\gamma}}{1+1/\gamma} + \frac{s_{it}^{1+1/\gamma}}{1+1/\gamma} \right] \right) \tag{53}$$

and set the exponent on the disutility of labor supply $\gamma = 2$.[18] We set the share of college-educated workers in the population to $1-\pi_\ell = 0.31$ in order to match the ACS data discussed in Appendix A. We set the elasticity of substitution across worker types z to $\phi = 3$; the literature provides relatively little empirical guidance on this parameter, so later we will perform sensitivity analysis with respect to different values. We set the capital depreciation rate $\delta = 0.014$ to imply an annual depreciation rate of 0.15 in order to match the average depreciation of equipment and software in the recent period (see Appendix A for details. We exclude structures from this calculation because they do not play a large role in capital-labor substitution (Krusell et al. (2000) make a similar assumption). We set the job destruction rate $\sigma = 0.028$ monthly and set the elasticity of the matching function with respect to vacancies to $\eta = 1/2$.

[17]The utility function imposes two restrictions beyond the Greenwood, Hercowitz and Huffman (1988) form: the exponent on the disutility of search effort s_{it} is the same as the exponent on labor supply n_{it} and the scale parameter χ_b is constant within broad group $b \in \{\ell, h\}$. We make these assumptions for the sake of parsimony but show below that the model nevertheless captures a number of salient features of the data.

[18]The exponent γ does not map directly into the Frisch elasticity of labor supply because, with search frictions, households' employment depends on labor market tightness and search effort in addition to the wage and marginal utility of consumption.

Table 3: Fitted Parameters

Parameter	Description	Value
Labor market frictions		
κ_0	Scale of vacancy posting cost	1.57
ω	Monopsony power	0.04
Worker productivity distribution $\log \mathcal{N}(\overline{\mu}_b, \sigma_b)$		
$\overline{\mu}_\ell$	Mean of non-college z (normalization)	0.00
σ_ℓ	SD of non-college z	0.60
$\overline{\mu}_h$	Mean of college z	0.83
σ_h	SD of college z	0.67
Production function		
α	Long-run elasticity of substitution b/t k and n_h	0.44
ρ	Long-run elasticity of substitution b/t n_l and $G(k, n_h)$	2.03
ψ	Coefficient on non-college labor n_ℓ	0.62
λ	Coefficient on capital k	0.58

Notes: Parameters endogenously chosen to match the moments in Table 4.

Fitted Parameters. Table 3 shows the parameters which we choose to match moments in the data. We group these parameters into three sets. The first set governs the degree of labor market frictions: the scale of vacancy posting costs κ_0 controls the strength of search frictions and the preference parameter ω controls the strength of monopsony power. The second set governs the distribution of idiosyncratic productivity z across households, which we assume is log-normally distributed within education group $z \sim \log \mathcal{N}(\overline{\mu}_b, \sigma_b)$. Finally, the last set of parameters govern the productivity function; the coefficients ρ and α control the long-run elasticities of substitution between capital and labor and the scale parameters ψ and λ control the distribution of income across factors.

Table 4 shows the moments that we target in our calibration.[19] The average unemployment rate is informative about the rate at which workers find jobs and, therefore, the scale of vacancy-posting costs κ_0. Conditional on a value for the vacancy posting costs, the average wage markdown primarily depends on the degree of monopsony power ω. There is a fast-growing literature measuring wage markdowns in the U.S., so as a benchmark we target the value from Berger, Herkenhoff and Mongey (2021) but will show how our results change given the range of values estimated in the literature.

We target moments of the wage distribution from the pool 2017-2019 wave of the ACS, which is described in detail in Appendix A. The college wage premium, defined as the median wage among college-educated workers relative to the median among non-college workers, is informative about the average productivity of college workers $\overline{\mu}_h$ relative to average productivity of non-college workers $\overline{\mu}_\ell$ (which we normalize). The 50-10 ratio within each group is informative about the dispersion of productivities σ_ℓ and σ_h. We target the 50-10 ratio rather than the standard deviation

[19]In each step of this moment-matching process, we choose the scale of the disutility of labor supply χ_ℓ and χ_h in order to match the aggregate employment rates of non-college and college workers.

Table 4: Targeted Moments

Moment	Description	Data	Model
Average unemployment rate			
$\mathbb{E}[s_i]/(\mathbb{E}[s_i]+\mathbb{E}[n_i])$	Average unemployment rate	5.9%	5.8%
Average wage markdown			
$\mathbb{E}[w_{ni}]/\mathbb{E}[F_{ni}]$	Average wage markdown	0.71	0.71
Wage Distribution, ACS 2017-2019			
$w_{h50}/w_{\ell 50}$	College wage premium	1.81	1.79
$w_{\ell 50}/w_{\ell 10}$	Non-college 50-10 ratio	2.04	2.04
w_{h50}/w_{h10}	College 50-10 ratio	2.30	2.30
Response to capital price decline (our data)			
$d\log\frac{k}{wn}/d\log q$	Response of capital-labor ratio	-0.47	-0.49
$d\text{share}^{\text{col}}/d\log q$	Response of college inc. share	-0.09	-0.07
Average income shares			
$\mathbb{E}[w_i n_i]/Y$	Aggregate labor share	0.57	0.64
$\pi_h\mathbb{E}[w_{hz}n_{hz}]/\mathbb{E}[w_i n_i]$	College income share	0.55	0.52

Notes: Moments targeted using parameters in Table 3. The *average wage markdown* is the payroll-weighted average markdown from Berger, Herkenhoff and Mongey (2021). The *wage distribution* targets are drawn from the ACS 2017-2019 data described in Appendix A. The *response to capital price decline* are the regression coefficients in columns (4) and (5) of Table 1. The *average labor share* is from Karabarbounis and Neiman (2014). Finally, the *college income share* is from the ACS 2017-2019 data described in Appendix A.

or interquartile range in order to precisely match the left tail of the wage distribution (which is most directly affected by the minimum wage).

As described in Section 3.1, the long-run regression coefficients from Table 1 columns (4) and (5) are informative about the long-run elasticities of substitution ρ and α (given values for these parameters, the average labor income shares are informative about the scale parameters μ and λ). We compute the regression coefficients in the model by comparing the values of the aggregate capital-to-wage bill ratio and the college income share from two different steady states corresponding to different levels of the relative price q^*.

This procedure depends on two key assumptions. First, we assume that our empirical results are driven by changes in the relative price of capital and not other confounding factors. Second, we assume that the effect of a price decline over a twenty-year period is well-approximated by comparing steady states in the model. We have validated this assumption in three ways. First, Table 1 shows that the regression coefficients are similar over the ten- and twenty-year horizons, indicating that the precise path of prices is relatively unimportant for the long-run elasticities. Second, Appendix B shows that our empirical results are robust to controlling for past price changes, indicating that the scope for path-dependence from past price changes is small. Third, Appendix C shows that after a permanent shock to the relative price of capital the model approximately returns to steady state within ten years, indicating that this steady state comparison is a good approximation to the entire

FIGURE 3: Calibrated Wage Distribution

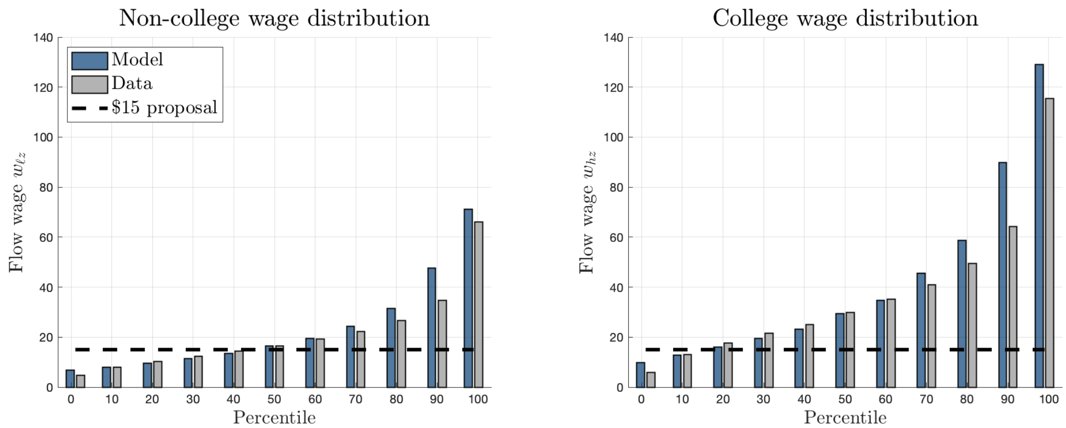

Notes: Wage distribution in calibrated model (blue bars) and the data (grey bars). The wage distribution in the data is drawn from the pooled 2017-2019 waves of the ACS, as described in Appendix A.

transition path.[20]

Table 4 shows that the model hits these targeted moments fairly well. Importantly, the model matches the average wage markdown and targeted moments of the wage distribution almost exactly. The model also matches the regression coefficients fairly well, although the college semi-elasticity is somewhat lower than in the data. The model's unemployment rate of 5.9% implies an average monthly job-finding rate of 0.42, close to the job-finding rate of 0.45 from Shimer (2005). Table 3 shows that our implied long-run elasticities $\rho = 2.03$ and $\alpha = 0.44$ are broadly in line with the estimates in Krusell et al. (2000), despite the fact that we use a different model and source of variation in the data. See Appendix C for a more formal discussion of identification in our model.

Figure 3 shows that the model matches the entire wage distribution fairly well, especially in the left tail of the distribution where the minimum wage will be most binding. For example, a $15 minimum wage would bind about 50% of non-college workers and 20% of college workers in both the model and the data. However, the dispersion of wages in the right tail of the distribution is somewhat higher in the model than in the data.

[20]In principle, we could estimate the long-run elasticities using an indirect inference exercise in which we incorporate sectoral heterogeneity into the model, feed in the realized path of relative prices for each series, solve for the transition path in response to these price changes, and choose parameters such that the resulting regression coefficients match the data. However, this procedure would require solving for the transition path at each step of our moment-matching exercise, which is computationally infeasible. In addition, the transition paths would still be sensitive to assumptions about the particular paths of sector-level prices in the future.

3.3 Validating Small Short-Run Elasticities of Substitution

While we have calibrated the model to match the long-run elasticities of substitution ρ and α by construction, a maintained assumption is that the short-run elasticities are Leontief (due to the putty-clay technology). Our goal in this section is to provide evidence in favor of this assumption using data on temporary changes in the relative price of capital. In the model with standard capital, temporary changes in the relative price will induce large changes in the capital-labor ratios due to intertemporal substitution. In the putty-clay model, however, capital-labor ratios can only adjust on newly purchased capital at the temporarily low price; since this new investment is a small fraction of the overall capital stock, the change aggregate capital-labor ratio will be relatively small.

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

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

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

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

$$\zeta_{st} = \theta_t + (1-\theta_t)\zeta_s. \tag{54}$$

Hence, the bonus generates time-series variation in the after-tax price of investment which differentially affects different sectors s due to the underlying heterogeneity in the baseline tax schedule ζ_s. The U.S. has used a $\theta_t = 30\%$, 50%, and 100% bonus at various points following the 2001 and 2008 recessions. Zwick and Mahon (2017) show that the Bonus Depreciation Allowance had a significant effect on firm-level investment.[21]

We will show that, even though this bonus has a sizeable effect on investment, it had an in-

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

significant effect on the capital-labor ratio. To examine the effect of the bonus on capital-labor substitution, we use the empirical specification

$$\text{share}^{\text{col}}_{st} = b_s + b(t) + b_1\zeta_{st} + \varepsilon_{st}, \tag{55}$$

where b_s is a sector fixed effect, $b(t)$ is a set of controls for aggregate conditions, and b_1 is the coefficient of interest.[22] As discussed, the putty-clay model predicts that $b_1 \approx 0$ because firms will cannot substantially adjust their total capital-to-labor ratios on college and non-college workers.

Table 5: Short-Run Response to Bonus Depreciation Allowance

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

Notes: Estimated coefficient b_1 from the regression (55) in the main text. Standard errors clustered by sector. "Sector-specific linear time trend" refers to controlling for a separate linear time trend in each sector.

Table 5 shows that the data support the putty-clay model's prediction of insignificant capital-labor substitution in the short run. Column (1) controls for $b(t)$ specified as a linear time trend and shows that the estimated coefficient b_1 is insignificantly different from zero with fairly tight standard errors. Column (2) shows that the estimate is nearly unchanged if we instead control for sector-specific linear time trends.[23]

Model Replication. We now show that the putty-clay model is quantitatively consistent with these empirical results while the model with standard capital is not. We replicate the Bonus Depreciation Allowance in our model as a transitory shock to the relative price of capital q_0 starting from an initial steady state with $q^* = 1$. We set q_0 so that the change in the relative price of capital equals the change in the tax-adjusted user cost $\frac{1-\tau_f\zeta_{st}}{1-\tau_f}$ from a 50% bonus. We assume the shock reverts back to steady state according to $\log q_{t+1} = \rho_q \log q_t$ and set the autocorrelation such that the half-life of the shock is one year. Finally, we compute the transition path in "partial equilibrium," that is, setting $Q_{t,t+1} = \beta$ for all t. This procedure assumes that the aggregate controls in

[22]The specification (55) is similar in spirit to Zwick and Mahon (2017), who put investment rather than the college income share on the left-hand side. Appendix B shows that the effect on investment in our more aggregated BEA data and finds comparable estimates to Zwick and Mahon (2017).

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

FIGURE 4: Response of College Income Share to Bonus Depreciation Shock

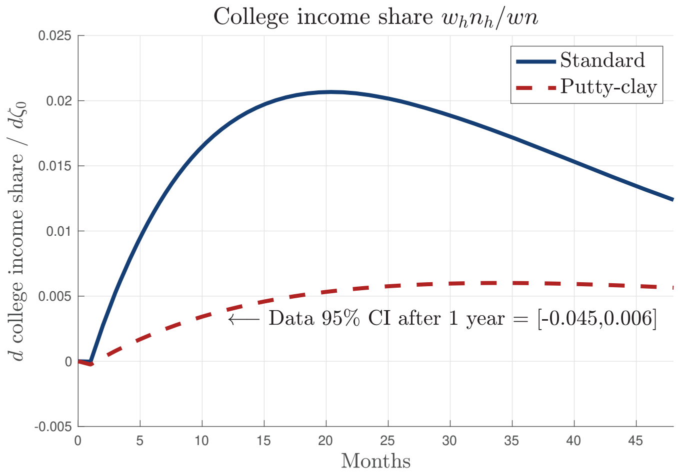

Notes: Partial equilibrium response of college income share to a transitory decline in q_t which mimics the bonus depreciation allowance. Specifically, the the economy starts in an initial steady state with $q^* = 1$ and then unexpectedly receives a shock q_0 such that the resulting decline in q_t equals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

the regression $b(t)$ absorb general equilibrium changes in the discount factor $Q_{t,t+1}$.

Figure 4 plots the response of the college income share relative to the implied change in ζ_0 in our model. One year after the shock, the college income share has increased by approximately 0.0045 relative to the increase in ζ_0, which is within the 95% confidence interval of the empirical coefficient b_1. In contrast, in a version of our model with standard capital, the college share increases by more than four times as much, putting it far outside the empirical confidence interval. Hence, the model with standard capital implies an unrealistically large degree of capital-labor substitution in the short run, while the putty-clay model does not. Appendix C shows that employment and capital-to-labor ratios are much more responsive in the standard model than in the putty-clay model.

4 Minimum Wage

We now use our calibrated model to study the effects of a minimum wage $\overline{w}$ at both the aggregate and microeconomic level. We assume the minimum wage is unexpectedly introduced starting from the initial steady state with $\overline{w} = 0$. Section 4.1 studies the new steady with $\overline{w} > 0$ in order to assess its long-run consequences. Section 4.2 then shows that the putty-clay technology slows down the transition from the initial equilibrium to this new steady state.

4.1 Long-Run Effects of the Minimum Wage

Aggregate Level. The top panel of Figure 5 shows that steady state aggregate non-college and aggregate college employment are both inverted-U functions of the minimum wage, which we refer to as a *Laffer curve*. As discussed earlier, this shape reflects the fact that a small minimum wage reduces the average monopsony distortion in the economy and increases employment but that a large minimum wage pushes wages sufficiently above their efficient level and decreases employment. The Laffer curve for non-college employment peaks at a lower level of the minimum wage than the curve for college employment because the average efficient wage among non-college workers is lower. A $15 minimum wage reduces non-college employment by 6% in our baseline calibration, but this number depends crucially on the degree of monopsony power $1/\omega$. For example, with a larger degree of monopsony power, consistent with the average markdown of 0.65 estimated by Hershbein, Macaluso and Yeh (2019), the $15 minimum wage does not significantly reduce non-college employment.[24]

The bottom panel of Figure 5 plots analogous Laffer curves for labor income. These Laffer curves peak at higher levels of the minimum wage than the employment Laffer curves because a higher minimum wage increases the average wage per worker even if it decreases employment. For example, a $15 minimum wage raises non-college labor income in our baseline calibration even though it lowers employment.

Figure 6 shows how these forces shape how the minimum wage shapes the distribution of aggregate income into its four components: non-college labor income, college labor income, capital income, and residual profits, which include both the vacancy posting costs and pure monopsony profits. In the initial equilibrium, profits are a substantial share of aggregate income due to the firms' monopsony power. The minimum wage reduces these profits by making workers more expensive and shrinking firms' markdowns. Relatively small levels of the minimum wage will nonetheless increase labor income, reflecting the labor income Laffer curves described above, but larger levels of the minimum wage will decrease labor income.

This discussion illustrates the key policy tradeoff from the aggregate perspective: an appropriately chosen minimum wage may reduce monopsony profits and increase labor income, but at the cost of decreasing employment and therefore total output in the economy. For example, a $15 minimum wage would increase aggregate non-college labor income by 4.1% but, because it reduces employment, would decrease total GDP by 1.6% (and the negative consequences for GDP worsen at higher levels off the minimum wage). One might think this tradeoff is worthwhile if the main

[24]For all values of monopsony power, the current national minimum wage $7.25 is barely binding given the observed distribution of wages which we use to calibrate the model.

FIGURE 5: Aggregate Minimum Wage Laffer Curves

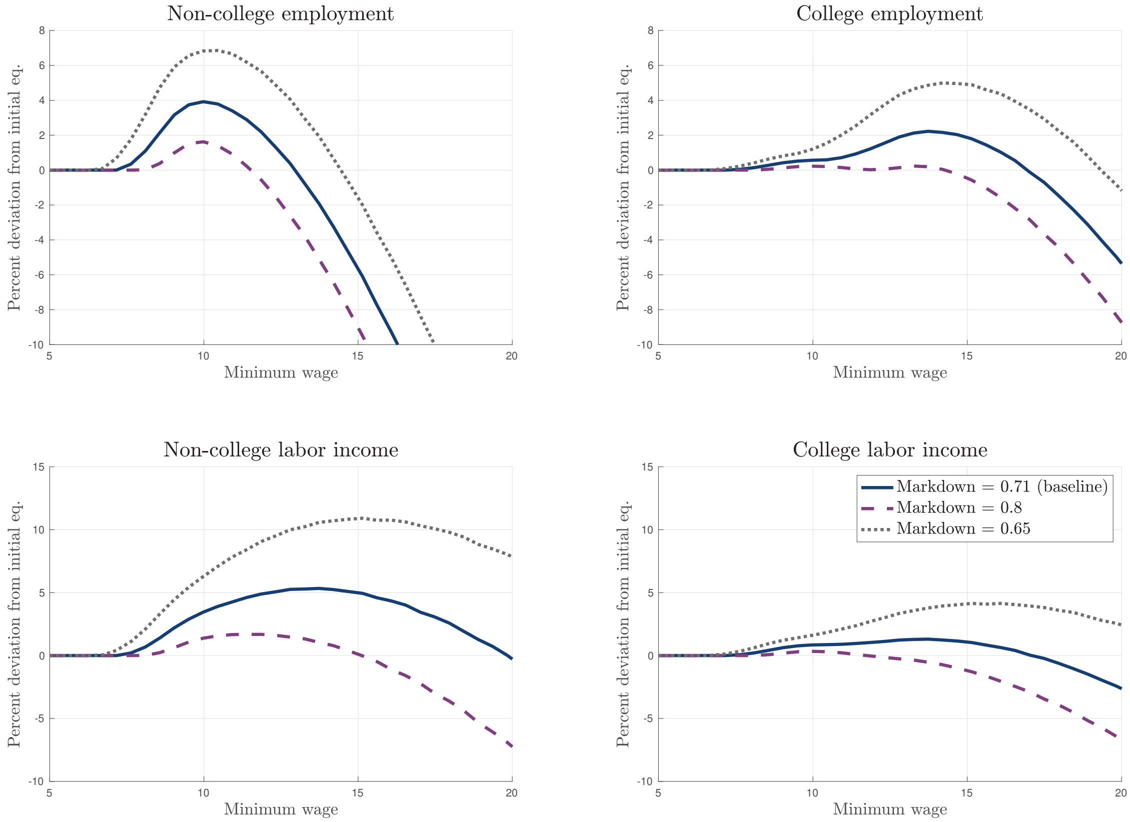

Notes: Steady-state outcomes as a function of the minimum wage $\overline{w}$. Top left panel plots aggregate employment of non-college workers and top right panel plots aggregate employment of college workers, both relative to their levels in the initial steady state with $\overline{w} = 0$. Similarly, the bottom left panel plots aggregate labor income of non-college workers and bottom right panel plots aggregate labor income of college workers, both relative to their initial steady states. The x-axis is the level of the minimum wage $\overline{w}$ normalized such that its level relative to the median non-college wage in the initial equilibrium is the same as in the data. Different markdowns correspond to different parameter values for ω.

goal is increase aggregate labor income of this particular group; however, we now show that such a policy would disproportionately harm the lowest-income workers within that group.

Micro Level. This simple example illustrates the tension inherent in setting a single minimum wage. If the desire it to help the non-college worker with the average level of skill, it has to be set fairly high, but doing so will severely hurt the low-skilled members of this group. If, instead, the desire is to help workers in this group with the lowest levels of skill, then it has to be set very low. But such a low level of the minimum wage will provide essentially no benefit to the vast majority of this group, who have higher skills.

Figure 7 plots analogous employment and labor income Laffer curves for non-college workers with certain levels of productivity z. The peaks of these Laffer curves are increasing in individual

Figure 6: Distribution of Aggregate Income

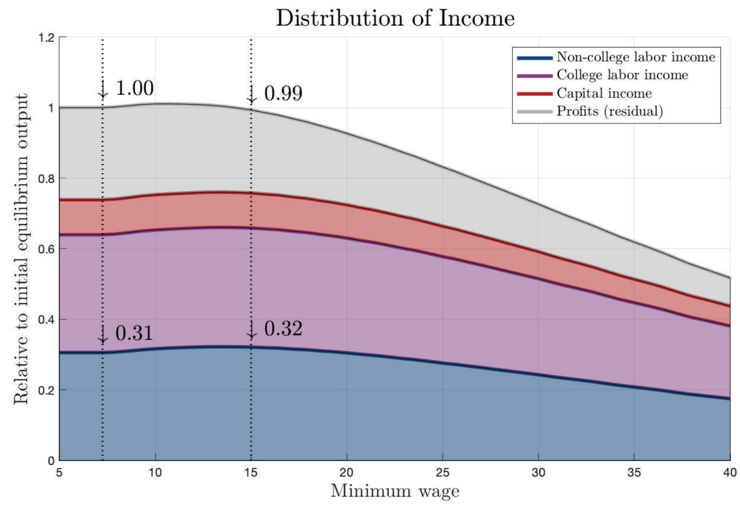

Notes: Plots steady-state GDP and the share accounted for by non-college labor income, college labor income, capital income, and residual profits as a function of the minimum wage $\overline{w}$. The y-axis is normalized such that aggregate income equals 1 without the minimum wage. The x-axis is the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

Figure 7: Disaggregated Minimum Wage Laffer Curves

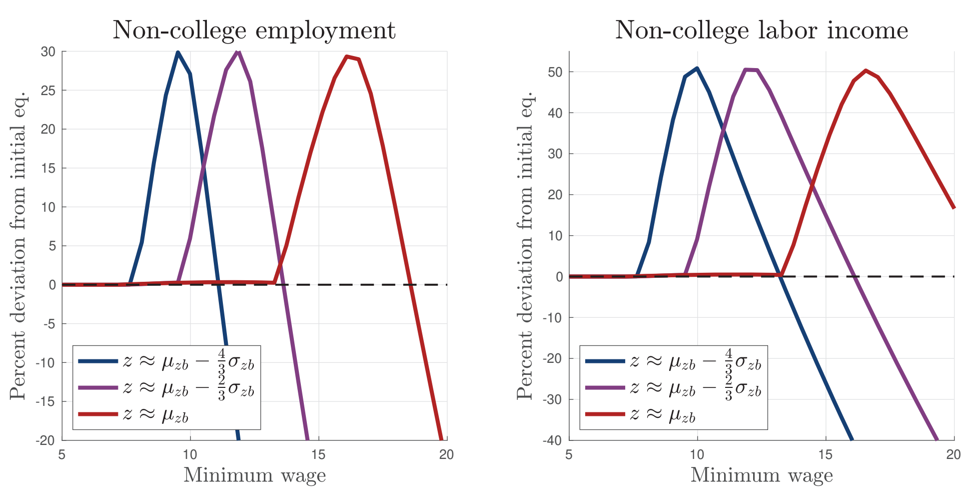

Notes: Steady-state employment (left panel) and labor income (right panel) of particular z-types among non-college workers as a function of the minimum wage. The x-axis is the level of the minimum wage $\overline{w}$ normalized such that its level relative to the median non-college wage in the initial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

Figure 8: Micro-level Effect of $15 Minimum Wage: Employment

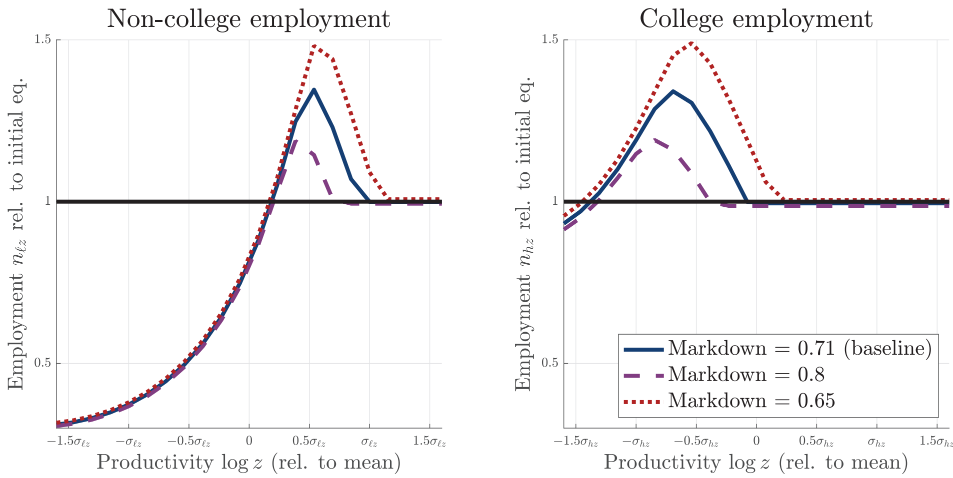

Notes: Steady-state employment of particular z types for a $15 minimum wage. The y-axis is normalized relative to employment level in the initial steady state without the minimum wage. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean. Different markdowns correspond to different parameter values for ω.

productivity z because the marginal products of those workers, and therefore their efficient level of wages, is increasing in z. This fact implies a distributional conflict across workers; setting the minimum wage high enough to alleviate the monopsony distortion for types with an average level of productivity z requires setting the wage above the efficient level for lower productivity levels. For example, the $15 minimum wage is around the peak of the employment Laffer curve of the mean-z type but implies that employment over lower-z types substantially falls. As with the aggregate curves, the labor income Laffer curves are shifted out and to the right because increasing the minimum wage mechanically increases the wage per unit of employment.

In order to more fully assess these distributional consequences, Figure 8 plots the effect of the $15 minimum wage on employment across all z-types. The left panel shows that, at this level, employment falls for all non-college workers with types below a threshold level of productivity, $\overline{z}$. This threshold implies that employment falls for 56% of all non-college workers. The decline in employment is the largest among the lowest-productivity workers, who earn the lowest wages in the initial distribution. Figure D.4 in the Appendix shows that a similar pattern holds for labor income, namely, it falls for about 40% of non-college workers and it falls the most for those who initially have the lowest wages. In this sense, the minimum wage disproportionately reduces the employment and income of precisely the group of workers that it is meant to benefit. In contrast, employment and labor income rise for a set of medium z-types of non-college workers because the minimum wage reduces the monopsony distortion to their wages. As Figures 8 and D.4 also show, increasing the

Figure 9: Micro-level Effect of $15 Minimum Wage: Markdowns

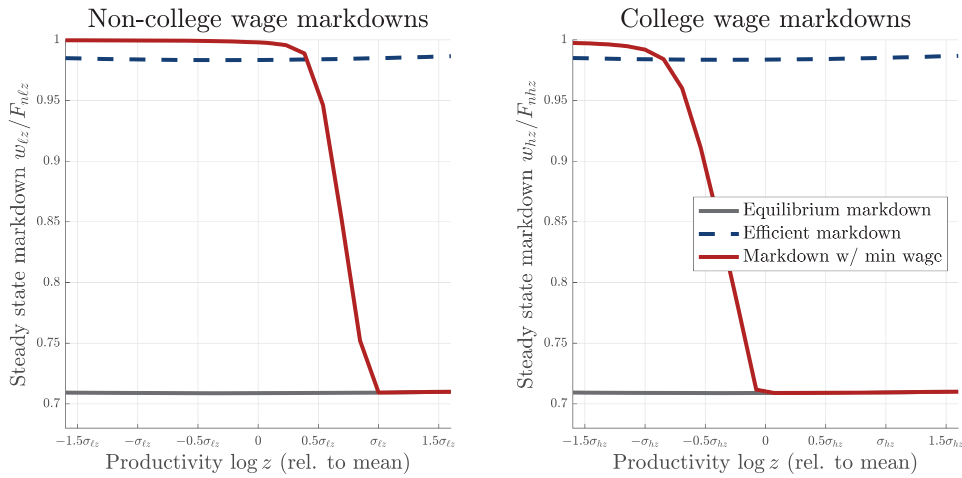

Notes: Steady-state wage markdowns w_i/F_{ni} of particular z types for three different parameterizations: (i) $\overline{w} = 0$ ("without minimum wage"), (ii) $\overline{w} = \$15$ ("with minimum wage"), and (iii) $\overline{w} = 0$ and $\omega \to \infty$ ("efficient markdowns"). The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

degree of monopsony power $1/\omega$ enlarges the set of workers who benefit, even without substantially affecting the cutoff levels of productivity below which employment and income fall.[25]

Figure 9 plots the associated wage markdowns w_i/F_{ni} as a function of individual productivity. Recall from (47) that markdowns are given by

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

While even the efficient markdown is below one in order to recoup the vacancy-posting costs $\kappa_i a_{it}$, the majority of the markdown in the initial equilibrium is due to the monopsony distortion. The effect of the minimum wage on markdowns falls into three cases which mirror its effect on employment. First, for low-z types, the minimum wage reduces the markdown to be lower than that in the efficient equilibrium, which implies that firms reduce employment, and thus vacancy-posting costs, to the point where wages approximately equal marginal products. Second, for middle-z types, the minimum wage shrinks the markdown closer to its efficient level, which is associated with an increase in employment. Finally, for high-z types, the minimum wage does not bind and therefore does not affect the markdown.[26]

[25] In contrast to its effects of non-college workers, the minimum wage increases employment among nearly all college workers simply because their productivity is on average higher than that of non-college workers.

[26] Employment of these high-z types nevertheless falls with the minimum wage due to the imperfect substitutability

Additional Results. Appendix D.1 contains three additional results about the long-run effects of the minimum wage, both at the aggregate and the micro level. First, we compute the aggregate Laffer curves for different values of the long-run elasticity of substitution ρ. Second, we compute the micro-level effect of the \$15 minimum wage on labor income rather than employment, as we did in Figure 8. Third, we show how these micro-level effects depend on the elasticity ρ.

4.2 Short Run vs. Long Run

We illustrate the role of putty-clay technology in slowing down the transition dynamics to the new steady state with a \$15 minimum wage. We solve for this path assuming, as we did in Proposition 3, that firms fully utilize all types of capital $k_t(v)$ and then verifying ex-post that it is optimal for them to do so.[27] Firms fully utilize capital because the variable profits from doing so are still positive, even with the introduction of the minimum wage. Since the initial capital stock is relatively labor-intensive and depreciates slowly over time, firms decrease their employment slowly over time as well. Hence, the presence of monopsony power—which ensures that variable profits are large enough to fully utilize existing capital—slows down the response to the minimum wage in our model.[28]

Figure 10 plots the paths of aggregated employment and the associated labor-to-capital ratios along the transition to the new steady state. The left panel shows that it takes non-college employment approximately twenty years to converge to its steady state value with a half life of approximately 100 months. The right panel shows that firms substitute away from these workers on the new capital that it installs, but in the initial periods of the transition that new capital is a small fraction of the overall capital stock (which remains relatively non-college labor intensive). In contrast, Appendix D.2 shows that non-college employment transitions much more quickly to the new steady state in the model with standard capital because firms can immediately substitute away from these workers. Hence, the putty-clay technology is the driving force behind these prolonged transitions in our model.

Figure 10 shows that firms substitute toward college employment more quickly than they substitute away from non-college employment. The right panel shows that the college labor-to-capital ratio jumps up upon the introduction of the minimum wage. This large jump implies that the total

across types ($\phi < \infty$). In this case, lower overall employment reduces the marginal product of all types, even if they are not directly affected by the minimum wage.

[27]Our model has three sources of variable profits (revenue net of wage payments): the sunk cost of capital, the sunk vacancy-posting cost, and the pure monopsony profits. Appendix D.2 shows that, upon the introduction of the minimum wage, these variable profits are positive for levels of the minimum wage below \$35 or so.

[28]Note that this rationale for full utilization is different than in Atkeson and Kehoe (1999). In that model, capital types are indexed by energy efficiency, and the relevant notion of variable profits is revenues net of energy costs. Although the only source of variable profits in that model is the sunk capital cost, energy costs are so small that even large changes in the relative price of energy does not push variable profits negative.

Figure 10: Transition Path to New Minimum Wage Steady State, Putty-Clay Model

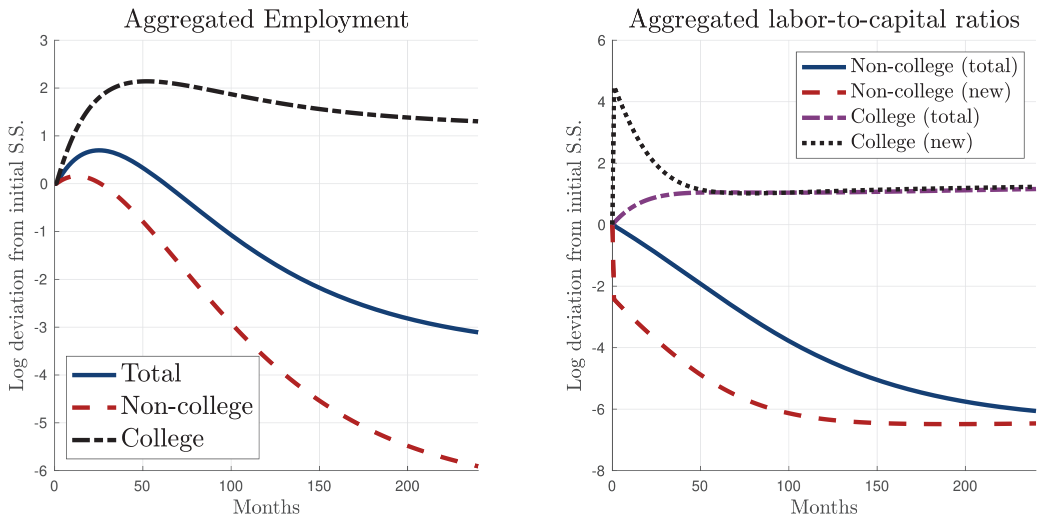

Notes: Transition path following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. Left panel plots the aggregated employment of non-college workers, college workers, and total employment. Right panel plots the associated labor-to-capital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

Figure 11: Distribution of Capital Types Along the Transition Path

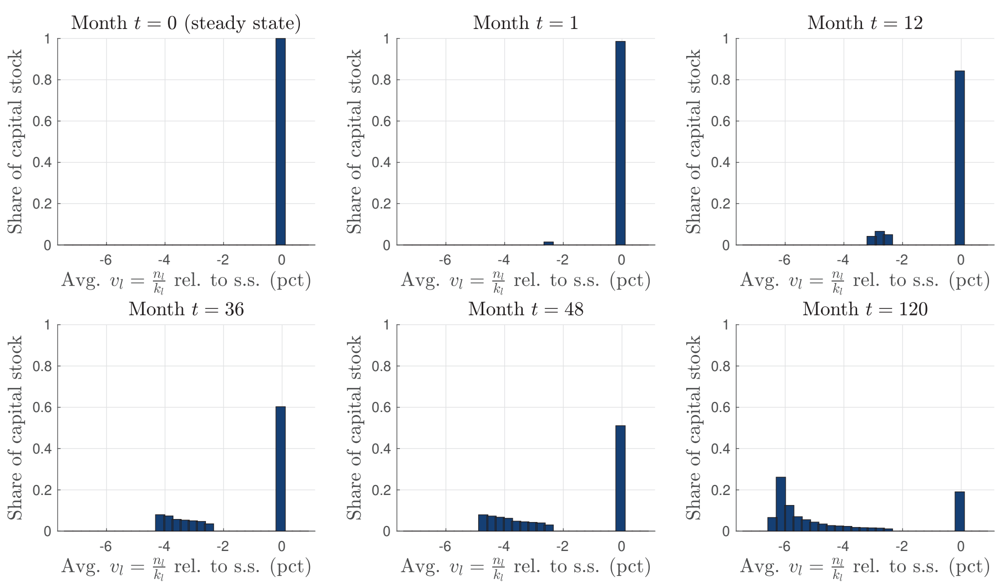

Notes: Distribution of capital types along the transition path following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. The x-axis indexes the type of capital by its average non-college labor to capital ratio.

college labor to capital ratio, and therefore college employment, converges relatively quickly to the new steady state.

Figure 11 plots the distribution of capital types in order to highlight the role of putty-clay technology. 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 such investment is a small share of the total capital stock. Over time, firms let the old capital stock depreciate away and continue to build up its investment in the new, less labor-intensive capital. Hence, the depreciation rate δ is crucial in determining the speed of transition.[29]

Together, we believe that these results highlight the importance of using our model to assess the long-run impact of the $15 minimum wage. For example, Figure 10 shows that, even four years after the introduction of the minimum wage, non-college employment has fallen by only one-fifth of its long-run amount. Hence, even the best-identified regressions using short-run data would not identify the long-run response. This result also suggests that putty-clay technology may be a reason why existing studies using short-run data do not find large effects of the minimum wage (although that could also be for the simple reason that historical changes in the minimum wage have been small).

5 Alternative Policies in the Tax and Transfer System

The previous section showed that, while the minimum wage may be successful in reducing the average monopsony distortion in the economy, it disproportionately reduces employment and labor income for low-income workers—a key policy goal of the minimum wage. We now study how alternative policies within the existing tax and transfer system may better achieve this goal. Section 5.1 describes how we model the tax and transfer system and ensure that the alternative policies are quantitatively comparable to a given change in the minimum wage. Sections 5.2 - 5.4 study the effects of three particular alternatives on employment and Section 5.5 compares their effects on welfare. Throughout, we study the steady state consequences of these policies to focus on their long-run effects.

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

5.1 Modeling Alternative Policies

We introduce the labor income tax system

$$T(w_i) = \underbrace{w_i - \lambda w_i^{1-\tau}}_{\text{progressive tax/transfer system}} - \underbrace{T_0(w_i)}_{\text{earned income tax credit}}, \tag{56}$$

where $T(w_i)$ is the labor income tax schedule, λ and τ are parameters, $T_0(w_i)$ is a parametric function that characterizes the earned income tax credit (described below), and negative taxes $T(w_i) < 0$ indicate transfers. We let $A(w_i) = w_i - T(w_i)$ denote *after-tax* labor income. In recent work Heathcote, Storesletten and Violante (2017) show that $\tau = 0.181$ provides a good description of the progressivity of the U.S. tax and transfer system, except at the very bottom because transfers are phased in and out at various income levels.[30] We explicitly model the earned income tax credit (EITC) schedule in order to capture these complications.

The tax and transfer system affects both the incentives for firms to hire workers and the supply of households searching in the labor market. To build intuition, note that the steady state version of the vacancy posting condition, which is effectively a firms' labor demand function, is now

$$\frac{\kappa_i}{\lambda_f(\theta_i)} = \frac{1}{r+\sigma}\left(F_{ni} - w_i - \underbrace{\frac{\widetilde{\sigma}}{\omega}\frac{v'(n_i)}{A'(w_i)}}_{\text{monopsony distortion}}\right), \tag{57}$$

where $A'(w_i)$ is marginal after-tax income.[31] Equation (57) shows that a positive marginal tax rate, $A'(w_i) < 1$, exacerbates the monopsony distortion relative to our baseline model with $A'(w_i) = 1$. Recall that the monopsony distortion arises because hiring a marginal worker increases the marginal disutility of labor supply for all the inframarginal hires, and the firm needs to compensates these inframarginal hires with a higher wage. A positive tax rate reduces the after-tax wage the inframaginal hires receive and therefore increases the required before-tax wage payment the firm must offer, further increasing the private marginal cost of hiring above the marginal cost of the planner. Conversely, a negative marginal tax rate, namely a tax credit which results in $A'(w_i) > 1$, alleviates the monopsony distortion by reducing firms' private marginal costs of hiring, bringing

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

[31] This expression involves the wage w_i because, unlike our baseline model, we do not have a closed-form expression w_i we can use to eliminate it from (57). Instead, the wage is determined by the equation $w_i = F_{ni} - \frac{\widetilde{\sigma}}{\omega}v'(n_i) - \frac{1-\eta}{\eta}\left(\frac{A(w_i)-v'(n_i)}{A'(w_i)}\right)$.

firms' marginal costs of hiring closer to those of the planner.

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

$$h'(s_i) = \frac{\lambda_w(\theta_i)}{r+\sigma}\left(A(w_i) - v'(n_i)\right), \tag{58}$$

where $r \equiv 1/\beta - 1$. Here a positive average tax rate, namely one that implies $A(w_i) < w_i$, reduces the after-tax wage payments of a job relative to our baseline model and therefore, as (58) indicates, reduces the incentive to search. For the search decision the marginal tax rate is irrelevant because the search decision is along the extensive margin; that is, to search or not to search. Conversely, negative average tax rates imply $A(w_i) > w_i$ and thus increase the incentive to search.

Hence, policies within the tax and transfer system are well-suited to alleviate the monopsony distortion in the economy. We ensure that these alternative policies are comparable to the \$15 minimum wage in the following way. First, note that the \$15 minimum wage reduces firms' flow profits by some amount $\Delta\pi^*$. Hence, we can think of the minimum wage as corresponding to an implicit tax on profits. For each of our alternative policies, we instead assume that there is no minimum wage but that the government levies an explicit corporate income tax τ_f on firms' profits which raises the same amount of revenues $\Delta\pi^*$. We assume that both investment and vacancy posting costs are fully deducted from corporate taxes, implying that the tax τ_c does not distort any of the firms' marginal decisions (because it is a pure profits tax). We then use these tax revenues to fund each of our alternative policies, that is, $\tau_f\pi = \Delta\pi^* = -\int T(w_i)di$ where π is flow profits in the new equilibrium. Hence, each of our alternative policies transfers, on net, $\Delta\pi^*$ resources from firms to households: they only differ in the schedule $T(w_i)$ that determines the distribution of these transfers across households.[32]

5.2 Uniform Tax Credit and Tax Cut

We start with a stylized policy that is useful to build intuition for our later experiments, namely, a uniform linear income tax credit of rate τ_c starting from an allocation with no taxes. This policy is a special cause of the general system (56) with the EITC component $T_0(w_i) = 0$, the progressivity parameter $\tau = 0$, and the scale parameter $\lambda = 1 + \tau_c > 1$. We then consider a uniform tax cut starting from an allocation with positive taxes.

Figure 12 shows that this tax credit increases employment by approximately 8% uniformly

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

Figure 12: Effect of Uniform Tax Credit and Tax Cut on Employment

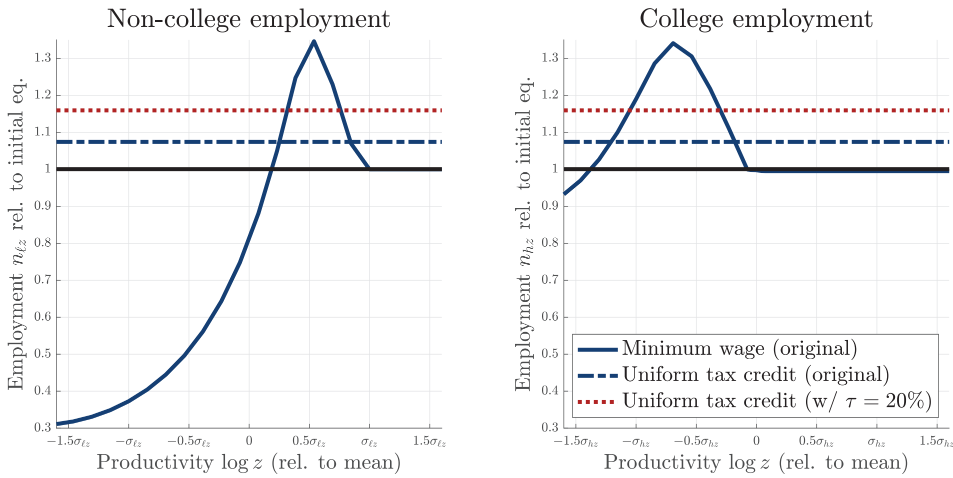

Notes: Steady-state employment of particular z types for three different policies: the \$15 minimum wage (blue line), the budget-equivalent tax credit (red line), and the budget-equivalent tax cut starting from a 20% tax rate (purple line). The y-axis is normalized relative to employment in the initial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

across workers. On the labor demand side, the tax credit, by increasing marginal after-tax income by $A'(w_i) = 1 + \tau_c > 1$, alleviates the monopsony distortion in (57). In this sense, it is analogous to the subsidy used to correct monopoly distortions in New Keynesian models with monopolistic competition. On the labor supply side, the tax credit increases average after-tax income by the same amount $A(w_i) - w_i = 1 + \tau_c$, increasing search effort in (58). These two forces increase employment equally across types z.

Of course, in reality labor income taxes are positive, so in Figure 12 we also consider the effect of a uniform tax cut starting from a positive tax rate $\tau > 0$. In particular, we compute the initial equilibrium of the model with a baseline tax rate of $\tau = 1 - \lambda = 20\%$, recompute the effects of a \$15 minimum wage in this alternative model, and then offer households a budget-equivalent tax cut. Figure 12 shows that the tax cut increases employment by even more than the baseline tax credit described above. This occurs because the positive tax rate is a further distortion reducing employment and the tax cut reduces this distortion, in addition to alleviating the monopsony distortions, as we described above).

Appendix E.1 contains two additional results. First, it shows that the uniform tax credit or tax cut also increases after-tax labor income uniformly across z-types. Second, it shows that the post-tax markdown $A(w_i)/F_{ni}$ uniformly shifts closer to the efficient markdown.

FIGURE 13: Progressive Tax and Transfer System

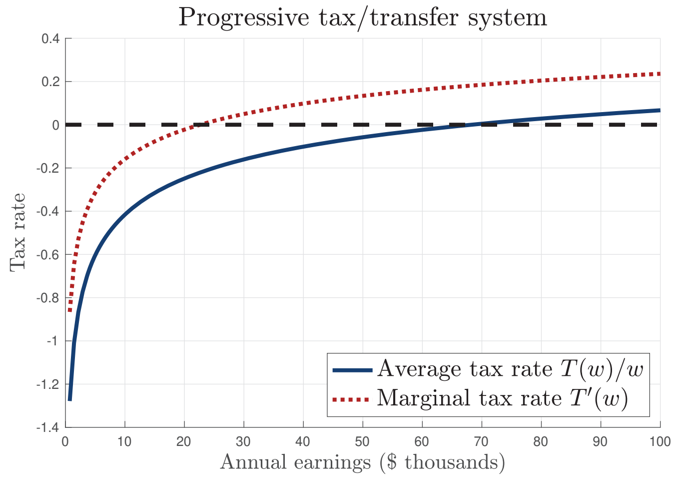

Notes: Average tax rates $T(w)/w$ and marginal tax rates $T'(w)$ from the budget-equivalent tax and transfer system described in the main text (with the U.S. level of progressivity $\tau = 0.181$). The x-axis rescales steady state labor income to annual earnings assuming each household works 1800 hours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

5.3 Progressive Tax and Transfer System

While the uniform tax credit/tax cut succeeds in alleviating the monopsony distortion, it does not target low-income workers in particular. Our second alternative policy targets these workers with larger transfers financed in part through raising taxes on high-income workers as in the progressive system used in the U.S. In particular, we continue to abstract from the EITC component by setting $T_0(w_i) = 0$ but we now set the progressivity parameter $\tau = 0.181$ in order to match the estimated progressivity from Heathcote, Storesletten and Violante (2017) for the United States and choose the scale parameter λ to ensure that the aggregate net transfer payment is budget-equivalent to the $15 minimum wage.

Figure 13 plots the average tax rates $T(w_i)/w_i$ and marginal tax rates $T'(w_i)$ of our budget-equivalent system as a function of income. The fact that the system is progressive, in that $\tau > 0$, implies that marginal tax rates are higher than average tax rates throughout. The lowest-income households have both a negative average tax rate (indicating they are receiving transfers $A(w_i) > w_i$) and a negative marginal tax rate (indicating that the transfers are being phased in, in the sense that $A'(w_i) > 1$). For these households, the tax/transfer system both encourages search effort in (58) and reduces firms' monopsony distortion since $A'(w_i) > 1$. Middle-income households continue to receive transfers, which encourage search effort, but face positive marginal tax rates, which exacerbate the monopsony distortion. Finally, high-income households face both positive average and marginal tax rates, reducing their search effort and exacerbating monopsony distortions. In this sense, a progressive tax and transfer system differentially alleviates monopsony distortions for

FIGURE 14: Effect of Progressive Tax System on Employment

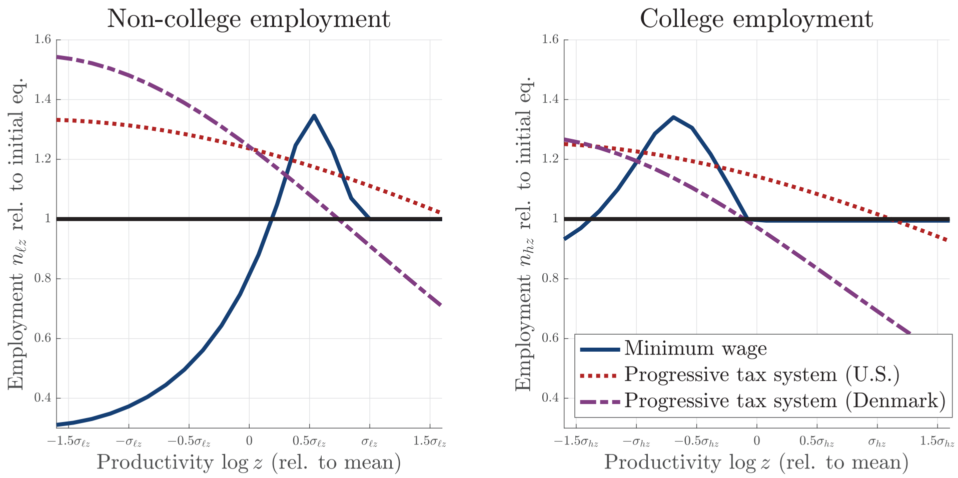

Notes: Steady-state employment of particular z types for three different policies: the \$15 minimum wage (blue line), the budget-equivalent progressive tax system with the U.S. level $\tau = 0.181$ (red line), and the budget-equivalent tax system with the Danish level $\tau = 0.463$ (purple line). The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

low-z and high-z types.

Figure 14 shows that this progressive system succeeds in substantially increasing employment of low-z types at the expense of decreasing employment for high-z types. Employment of the targeted low-z types increases between 20% to 40%, which is substantially larger than the 8% increase from the uniform tax credit. Importantly, the progressive system increases employment for all of the types of workers for whom the minimum wage would instead decrease their employment. (Figure E.3 in the Appendix shows similar pattern holds for labor income.)

For the sake of comparison, Figure 14 also shows the effect of an even more progressive tax and transfer system with $\tau = 0.463$ meant to capture degree of progressivity in Denmark.[33] This more progressive system further increases employment of the low-z types but decreases employment of the middle- and high-z types relative to the U.S. system.

Appendix E.2 contains two additional results about these progressive tax and transfer systems. First, it shows that the progressive tax/transfer system also increases after-tax labor income disproportionately for the low-z types. Second, it shows that the post-tax markdown $A(w_i)/F_{ni}$ shifts closer to the efficient level for the low-z types.

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

5.4 Earned Income Tax Credit

The results above show how a progressive tax and transfer system targeted at low-income households can alleviate the monopsony distortions they face. Empirically, the largest component of transfer payments is the Earned Income Tax Credit (EITC). The blue line in Figure 15 plots the 2014 EITC schedule for a household with two children drawn from Nichols and Rothstein (2016).[34] The schedule has three distinct regions which will play an important role in our analysis. In the first region, the tax credit is paid proportionally to the households' income at the *phase-in* rate of 40%. In this region, the household faces both a negative average tax rate (since the total tax credit is positive) and a negative marginal tax rate (since the credit is being phased in). Second, the EITC is eventually capped at its maximum benefit but is not yet being phased out; we refer to this region as the *plateau*. Finally, in the *phase out* region the EITC is phased-out at the proportional phase-out rate of 21.06% until benefits are exhausted. In this region, the household faces a positive marginal tax rate because each dollar of earnings reduces its transfer payment by 21.06%.

FIGURE 15: Earned Income Tax Schedule

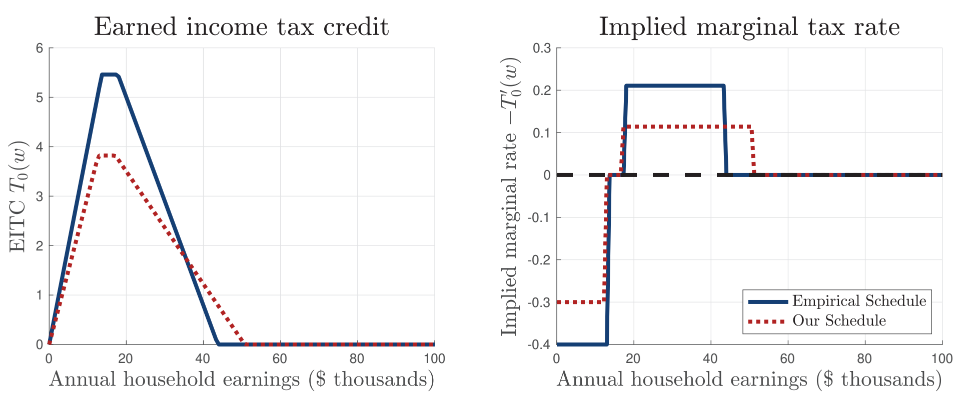

Notes: Blue line is the empirical EITC schedule from 2014 as a function of annual household income (from Nichols and Rothstein (2016)). Red line is the budget-equivalent schedule described in the main text; the x-axis rescales steady state labor income to annual earnings assuming each household works 1800 hours per year.

In order to understand how this kinked schedule affects employment, we implement a version of the policy designed to be budget equivalent to the minimum wage. In particular, we implement the policy plotted in red in Figure 15. We lower the phase-in rate to 30% and elongate the plateau in order to increase the mass of households in each region. We then solve for the value of the phase-out rate which makes the policy budget-equivalent to the $15 minimum wage (which implies a phase-out rate around 11%).

[34]The size of the EITC depends on some household characteristics, most importantly the number of children, from which we abstract in our analysis.

Figure 16: Effect of Earned Income Tax Credit on Employment

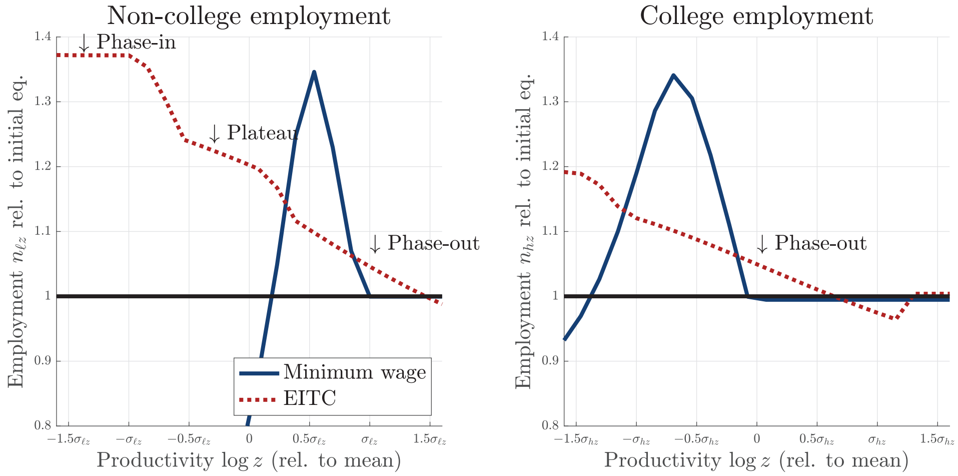

Notes: Steady-state employment of particular z types under the EITC. The y-axis is normalized relative to employment in the initial equilibrium (without any policies). The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.

Figure 16 shows how employment responds to each of these three regions of the EITC schedule. The lowest-productivity non-college workers are in the phase-in region and therefore receive both the positive effect of the negative average tax on labor supply and the negative marginal tax on labor demand. In fact, the phase-in region is isomorphic to the uniform tax credit studied above, but its effect is much larger because it is targeted at a small set of households. Employment is discontinuously lower in the plateau region because the marginal tax rate discontinuously jumps up from $-\tau_1$ to 0 (where τ_1 is the phase-in rate). However, employment is still higher than in the initial equilibrium without the policy because the average tax rate is still negative, increasing labor supply. Employment the jumps down again in the phase-out region because the marginal tax rate discontinuously jumps up from 0 to τ_2 (where τ_2 is the phase-out rate).

Hence, although the EITC is qualitatively similar to the smooth transfer payment from Section 5.3, its actual implementation creates discontinuous effects on employment. Appendix E.3 contains two additional results about the EITC: how it affects labor income and the after-tax markdown.

5.5 Welfare Effects of the Minimum Wage and Alternative Policies

We now compare the welfare consequences of the minimum wage and these alternative policies. Recall that we assumed that non-college workers of each skill level z share the idiosycratic risk to consumption from losing or finding a job among themselves—we make a similar assumption for college workers. These assumptions attenuate the welfare costs of losing a job and the welfare benefits from finding a job relative to a model without such risk sharing. Hence, we view the welfare

Figure 17: Effects of Alternative Policies on Welfare

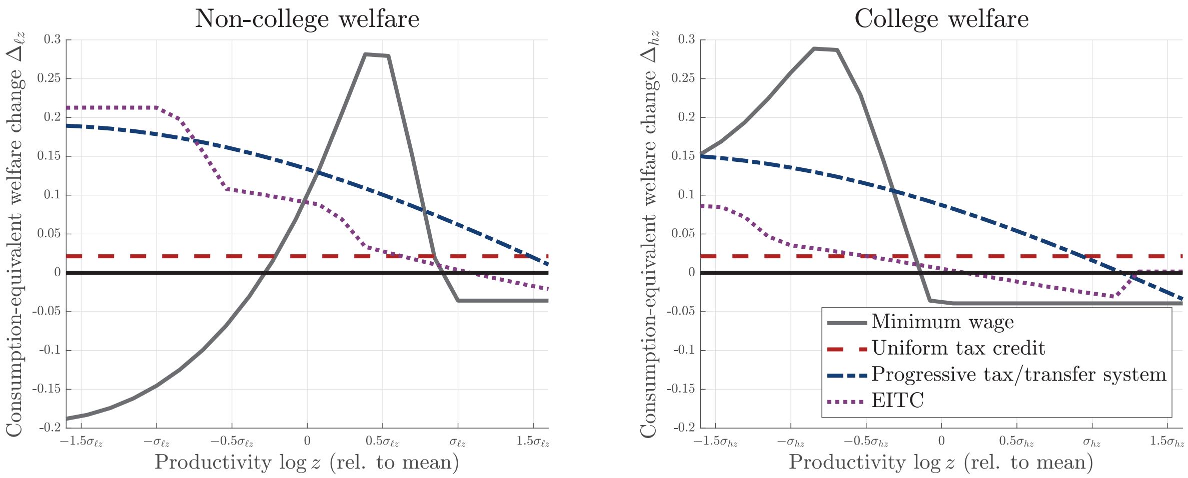

Notes: Steady-state welfare gains/losses Δ_i from $u\left((1+\Delta_i)c_i^* - v(n_i^*) - h(s_i^*)\right) = u\left(\widetilde{c}_i - v(\widetilde{n}_i) - h(\widetilde{s}_i)\right)$, where where the * superscript denotes the steady state values in the initial equilibrium without the policy and the $\widetilde{}$ superscript denotes values in the new equilibrium with the policy. The x-axis is log individual productivity $\log z$ relative to its mean value, expressed in standard deviations from the mean.

costs and benefits in our model with risk sharing as a lower bound on the actual ones.[35]

We define the welfare gain or loss for a particular household type i as the value of Δ_i which solves

$$u\left((1+\Delta_i)c_i^* - v(n_i^*) - h(s_i^*)\right) = u\left(\widetilde{c}_i - v(\widetilde{n}_i) - h(\widetilde{s}_i)\right), \tag{59}$$

where the * superscript denotes the steady state values in the initial equilibrium without the policy and the $\widetilde{}$ superscript denotes values in the new equilibrium with the policy. Intuitively, Δ_i is the percentage change in steady-state consumption that would make a household indifferent between living in the initial equilibrium and the new equilibrium. Hence, $\Delta_i > 0$ is positive if the policy makes the household better off and $\Delta_i < 0$ is negative if the household is worse off.

Figure 17 compares the welfare effects Δ_i of the \$15 minimum wage and the three budget-equivalent policies we have considered. The progressive tax and transfer system dominates the minimum wage for all non-college households in the bottom half of the productivity distribution (who make up around a third of all households in the economy). And among the lowest productivity types, the minimum wage has a large welfare cost while the progressive system has a large benefit. The EITC has an even larger benefit for those types because the phase-in region is particularly targeted toward them. However, the minimum wage dominates these other policies for households in the middle of the z distribution (which contain less than a third of non-college households)

[35]Note that 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.

because it reduces their monopsony distortion. The minimum wage also dominates the other policies for many of the relatively low-productivity college workers for the same reason (recall that college workers have higher productivity overall).

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

6 Minimum Wage Within a Tax and Transfer System

Given that the tax and transfer system influences the effective degree of monopsony power across workers, we now briefly study how the effect of the minimum wage depends on the existing tax and transfer system in place. We do so by re-calibrating the model to incorporate the empirical tax and transfer system in the U.S., as estimated by Heathcote, Storesletten and Violante (2017), and comparing the effects of the minimum wage in this environment to that in our baseline analysis from Section 4. In particular, we assume that labor income taxes are given by $T(w_i) = w_i - \lambda w_i^{1-\tau}$ (with no EITC) and set the progressivity parameter $\tau = 0.181$ and the scale parameter $\lambda = 0.894$ to match Heathcote, Storesletten and Violante (2017).[36] We then choose the same parameters as in Table F.1 to match the same moments as in Table 4 given this tax system. This strategy ensures that the average wage markdown, and therefore the average degree of monopsony power, is comparable across the two models. See Appendix F for details of the recalibration.

Figure 18 compares the effect of the $15 minimum wage in the recalibrated model to its effect in the baseline model without taxes. We focus here on the effect on labor income, but the effects on employment and welfare are similar and therefore relegated to Appendix F. As in the baseline model, the minimum wage substantially reduces labor income for the low-productivity types among non-college households. The decline in labor income is larger with the tax and transfer system, however, because these households receive negative marginal and average taxes, reducing their monopsony distortion before the introduction of the minimum wage. Hence, the $15 minimum wage pushes their wage further above its efficient level than in the baseline model. Conversely, the effect of the minimum wage on the middle-z types is better in the model with the tax and transfer system imposes positive taxes, and therefore exacerbates the monopsony distortion, on these types.

[36] Heathcote, Storesletten and Violante (2017) estimate the progressivity parameter τ from tax data alone, but their estimate of the scale paramter λ depends on their specific model of the income distribution, which differs from ours. However, we find that the distribution of tax rates is similar in our model as theirs; for example, our income-weighted average marginal tax rate is 0.41 vs. 0.34 in Heathcote, Storesletten and Violante (2017). Our unweighted average marginal tax rate is 0.34 and the average tax rate is 0.20.

Figure 18: Micro-level Effect of $15 Minimum Wage: Labor Income (with U.S. Tax System)

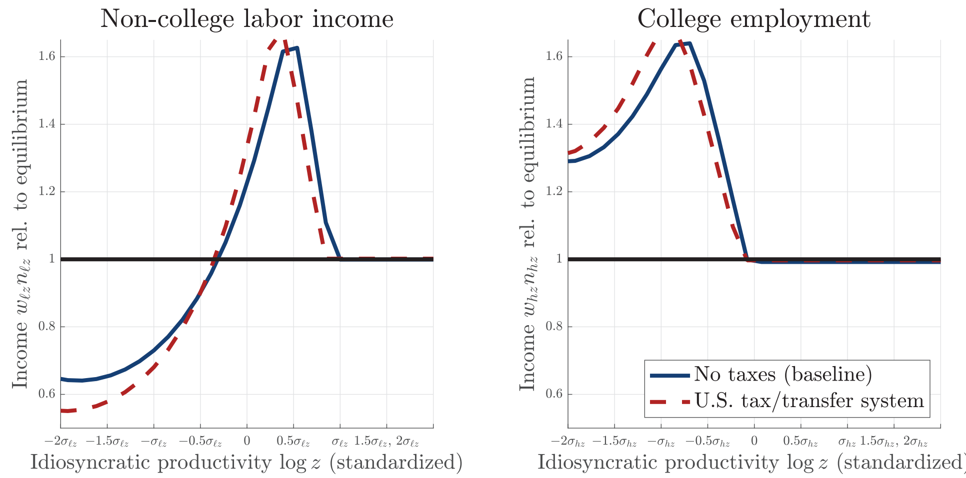

Notes: Steady-state labor income of particular z types given the $15 minimum wage. Blue line is the baseline model calibrated as in Section 3 (without the tax system) and red line is the recalibrated model as in Appendix F (with the tax system). The y-axis is normalized relative to employment in the 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.

While these differences are qualitatively in line with our theory, they are quantitatively fairly small. Therefore, we conclude that our baseline analysis is a useful guide for understanding the effects of the minimum wage on the economy.

7 Conclusion

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

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

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.

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

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

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

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

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

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

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.

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.

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

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

GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN (1988): "Investment, capacity utilization, and the real business cycle," *The American Economic Review*, pp. 402–417.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132(4), 1693–1754.

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

HERSHBEIN, B., C. MACALUSO, AND C. YEH (2019): "Concentration in US Local Labor Markets: Evidence from Vacancy and Employment Data," .

HOSIOS, A. J. (1990): "On the Efficiency of Matching and Related Models of Search and Unemployment," *The Review of Economic Studies*, 57(2), 279–298.

HUBMER, J. (2018): "The race between preferences and technology," *Unpublished Working Paper*.

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

KARABARBOUNIS, L., AND B. NEIMAN (2014): "The global decline of the labor share," *The Quarterly journal of economics*, 129(1), 61–103.

KENNAN, J. (1995): "The elusive effects of minimum wages," *Journal of Economic Literature*, 33(4), 1950–1965.

KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2020): "Leave-Out Estimation of Variance Components," *Econometrica*, 88(5), 1859–1898.

KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): "Capital-skill complementarity and inequality: A macroeconomic analysis," *Econometrica*, 68(5), 1029–1053.

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

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

Moen, E. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

Montgomery, J. D. (1991): "Equilibrium Wage Dispersion and Interindustry Wage Differentials*," *The Quarterly Journal of Economics*, 106(1), 163–179.

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

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

Nichols, A., and J. Rothstein (2016): *2. The Earned Income Tax Credit*. University of Chicago Press.

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

Shimer, R. (2005): "The cyclical behavior of equilibrium unemployment and vacancies," *American economic review*, 95(1), 25–49.

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

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

Zwick, E., and J. Mahon (2017): "Tax policy and heterogeneous investment behavior," *American Economic Review*, 107(1), 217–48.

A Data Appendix

This appendix contains details about our data sources and construction described in the main text.

A.1 BEA Data

We use the BEA Detailed Fixed Asset Tables to construct the relative price of investment, real investment expenditures, and the real value of the capital stock. These data contain sector-asset level observations for $S = 62$ sectors and $A = 93$ assets for the time period 1901-2019. We focus on the time period $1960 - 2019$ given the coverage of the Census/ACC data. Recall from the main text that that we define the growth rate in the price of investment relative to the price of consumption in sector s in year t using the Tornqvist index

$$\Delta \log q_{st} = \sum_{a=1}^{A} \omega_{ast} \Delta \log q_{at}, \tag{60}$$

where a indexes a type of capital good, ω_{ast} is a weight which sums to one within sector, and $\Delta \log q_{at}$ is the growth rate of the relative price of asset a. The Tornvqist weight ω_{ast} is the average of the share of nominal investment expenditures on asset a in periods t and $t-1$:

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

where P^I_{at} is the nominal price of asset a and $P^I_{at} i_{ast}$ are nominal investment expenditures in sector s on asset a. The growth rate in the relative of price $\Delta \log q_{at} = \Delta \log P^I_{at} - \Delta \log P^C_t$, where P^C_t is the price index of consumption goods.[37]

We measure the Tornqvist index (60) using data from the BEA Detailed Fixed Asset Tables, which contain. We compute the nominal price of asset a, P_{at}, as the ratio of nominal investment expenditures $P^I_{at} i_{ast}$ from Investment in the Private Nonresidential Fixed Assets table to real investment expenditures i_{ast} from the Fixed-Cost Investment in Private Nonresidential Fixed Assets.[38] We exclude artistic originals and R&D from our capital assets because they are not clearly related to capital-labor substitution.

Figure A.1 illustrates two key sources of variation in our relative price series $\Delta \log q_{st}$. First,

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

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

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

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

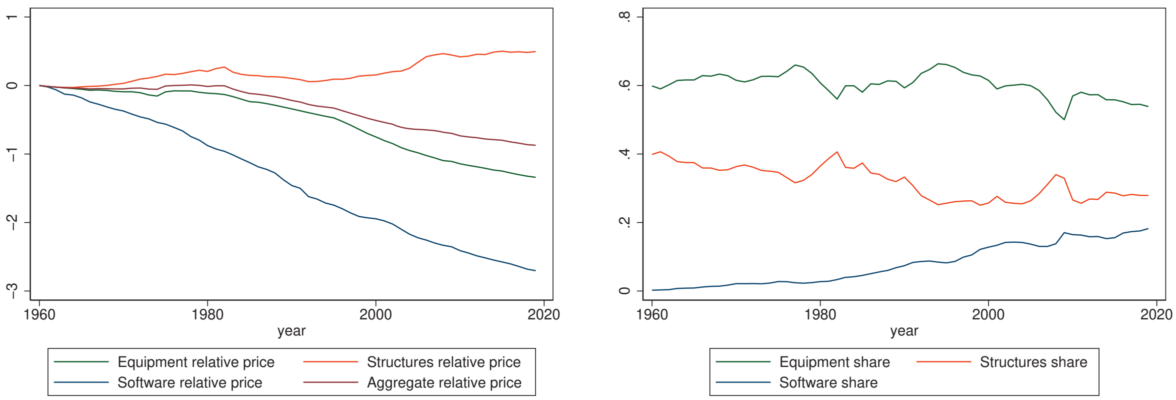

Notes: Left panel plots the level of the price index relative to 1960, $\log q_{st} \equiv \sum_{\tau=0}^{t} \Delta \log q_{s\tau}$, for three broad asset categories: structures, equipment, and software. The right panel plots the share of aggregate nominal investment expenditures on structures, equipment,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

the time-series variation in sector-level relative prices $\Delta \log q_{at}$ has tended to decline over time. We illustrate this variation in the left panel of the figure, which plots the level of the price index relative to 1960, $\log q_{st} \equiv \sum_{\tau=0}^{t} \Delta \log q_{s\tau}$, for three broad asset categories: structures, equipment, and software.[39] The relative price of equipment and especially software assets has substantially fallen over time, driving down the aggregate price index by more than 1 log point over this period (consistent with aggregate price series constructed in the literature, for instance, Hubmer (2018)). Second, there is variation in the investment expenditure shares ω_{ast} both across sectors and over time. The right panel of Figure A.1 illustrates the time series component of this variation by plotting the share of aggregate nominal investment expenditures in the three broad asset categories and shows that the share of expenditures on software has substantially risen over time. These two sources of variation interact to create substantial cross-sector heterogeneity in q_{st} over time.

Figure A.2 shows that these two sources of variation create significant heterogeneity in the long-run decline in the relative price across sectors. While most sectors' relative price falls between 0.5 and 1 log point between 1980 and 2019, many sectors relative price declines by 2 or more log points.[40]

We also construct other sector-level time series in a similar fashion. We compute the growth

[39] We define these three aggregated price series using a Tornqvist index analogous to (51) within that particular asset category.

[40] The only sector whose relative price increases over time is oil and gas extraction.

Figure A.2: Distribution of Long-Run Changes in the Relative Price q_{st}

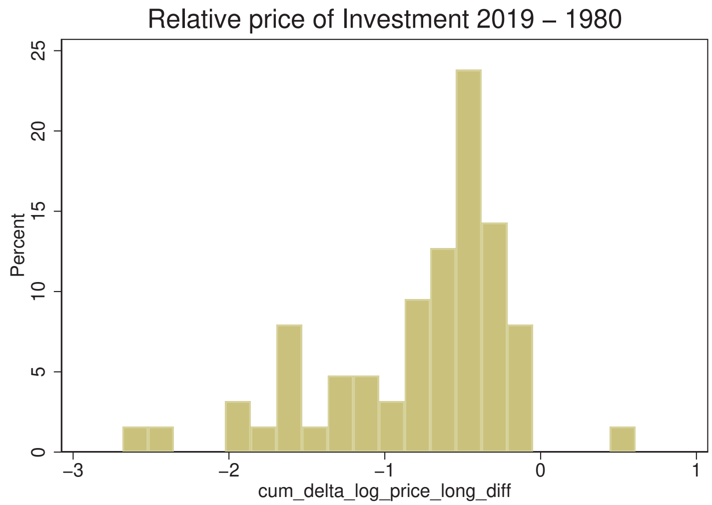

Notes: Histogram of long-run changes $\log q_{s,2019} - \log q_{s,1980}$ across sectors s.

rate of real investment at the sector level using a similar Tornqvist index to (51):

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

Second, we compute the growth rate of real capital stock in sector s as

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

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

Finally, as we discuss in Section 3.2, we calibrate the depreciation rate δ to match the depreciation rate of equipment and software in our BEA dataset. We compute an average depreciation rate δ_t by aggregating asset-level depreciation rates δ_{at} according to

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

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

A.2 Census American Community Survey (ACS) Data

To calibrate key model parameters and to test model predictions we used data from the 1970, 1980, 1990, and 2000 U.S. Censuses as well as the 2006-2019 American Community Surveys (ACS).[41] For all our analysis, we restrict the Census and ACS samples to (1) include all individuals aged 16 and above, (2) exclude all individuals residing in group quarters, and (3) exclude all individuals who report themselves as being a student. All data are weighted using the weights provided by the Census and ACS samples for the given time periods. Below we discuss how we use these data for various exercises in the paper.

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

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

Employment Rates. As part of our calibration, we match "full-time" employment rates by education group. By focusing on "full-time" employment, we measure workers with a strong attachment to the labor force. We define an individual as being "full-time" employed if (1) they are currently

[41] We downloaded these data directly from the IPUMS website: https://usa.ipums.org/usa/.

working at least 30 hours per week, (2) they reported working at least 29 weeks during the prior year, and (3) they reported positive labor earnings during the prior 12 month period. For our 2017-2019 sample, 46.8% of non-college individuals and 62.4% of college individuals worked full-time.

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

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

Long-Run Regressions. We exploit cross-industry variation to estimate how the share of income accruing to those with a college degree varies with changes in industry level price of investment, $\Delta log q_{st}$, defined above. In each decade, we collapse our Census/ACS samples to 230 three digit industries using 1990 industry codes which are harmonized across the Census/ACS years. The investment price data is at a higher level of industry aggregation than is the Census/ACS industry data. As a result, we create a cross-walk that maps the BEA sectoral investment price series to the Census/ACS industry codes. In our final sample, we have 62 sectors with a capital price series that we map to 195 Census/ACS industries; a given BEA capital price sector can be mapped to multiple Census/ACS industries. A complete description of our mapping procedure can be found in our replication package. Henceforth, we use industries and sectors interchangeably.

Figure B.1: Change in Investment Price vs. Change in Real Investment

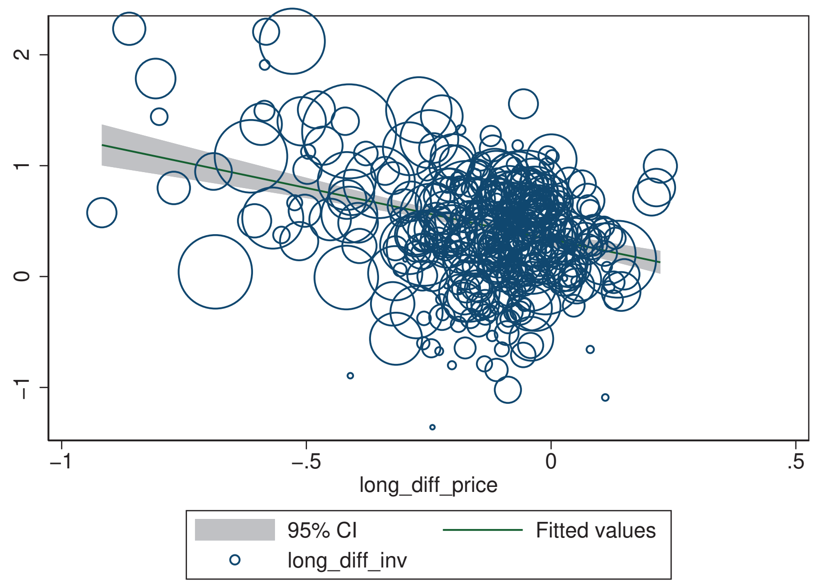

Notes: Scatterplot of ten-year changes in the relative price of capital q_{st} (x-axis) vs. the ten-year changes in log real investment $\log i_{st}$ (y-axis). Observations are weighted by their share of aggregate investment.

B Additional Empirical Results

This appendix contains additional empirical results mentioned in the main text.

B.1 Long-Run Investment Response to Investment Price Changes

In this subsection, we study the relationship between changes in the relative price of investment q_{st} and investment itself i_{st} in greater detail. We primarily focus on ten-year changes, as in the main text.[42] Figure B.1 shows a clear negative relationship between the decadal growth in the relative price and the decadal growth of investment. This negative relationship is consistent with our assumption below that changes in the relative price primarily reflect investment supply shocks, which allows us to trace out the slope of the investment demand schedule.

Table B.1 performs some additional analysis of our baseline regression specification

$$\Delta \log i_{st} = b_0 + b_t + b_1 \Delta \log q_{st} + \varepsilon_{st}, \tag{61}$$

where b_0 is a constant, b_t is a decade fixed effect, ε_{st} is a residual, and b_1 is the coefficient of interest.

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

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

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

Notes: Results from estimating (61) using ten-year changes in investment and the relative price of capital. Column (1) excludes the time fixed effects b_t. Column (2) includes the time fixed effects b_t and is equivalent to column (1) of Table 1 in the main text. Columns (3) includes time fixed effects and controls for the previous decades ten-year relative price growth. Columns (4) - (8) run (61) separately for each decade of data excluding time fixed effects.

Table B.2: Year-to-Year Investment Regressions

	(1)	(2)	(3)	(4)	(5)
$\Delta \log q_{st}$	-2.18***	-2.04***	-1.60***	-1.61***	-1.56***
	(0.136)	(0.172)	(0.166)	(0.134)	(0.120)
Observations	3758	3758	3758	3503	3250
R^2	0.825	0.871	0.946	0.942	0.939
Sector FEs	No	No	Yes	Yes	Yes
Time FEs	No	Yes	Yes	Yes	Yes
Horizon h	0	0	0	4	8
Lags?	No	No	No	No	No

Notes: Results from estimating $\Delta \log i_{st} = b_s + b_t + b_1 \Delta \log q_{st} + \varepsilon_{st}$ using annual data, where b_s is a sector fixed effect and b_t is a year fixed effect.

Column (1) shows that $b_1 = -0.93$ without the time fixed effects and column (2) shows that incorporating time fixed effects increases b_1 to -1.37 (the estimate in the main text). Column (3) shows that this result is robust to controlling for the last decade's investment price growth. Finally, columns (4)-(9) show that this negative relationship is fairly robust across decades, although it has weakened somewhat in the more recent period.

Table B.2 shows the results from running the investment regression (61) using year-on-year changes rather than decadal changes. Given the additional observations, column (2) adds sector fixed effects to the regression, but they don't substantially change the point estimate. Column (3) incorporates time fixed effects and shows that the estimated coefficient is comparable to our baseline results. Columns (4) and (5) replace the left-hand side with $\log_{st+h} - \log_{st}$ and show that the investment responses are persistent over time.

Table B.3: Response of Change in College Income Share to Change in Investment Prices, Cross Sector Variation, 10-Year Differences

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

Notes: Estimated α_1 from regression (52) for various time periods and various controls. Changes are 10-year (1980-1990 through 2010-2019). Data from Census/ACS. First three columns pool data across the 1980-2020 period allowing for four decadal changes per sector. The last four columns show results for each decade separately. Standard errors clustered at the sector level shown in parentheses.

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

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

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

B.2 Long-Run College Income Share Response to Investment Price Changes

Table B.3 shows alternative specifications of our long-run regression (52) referenced in the main text.

B.3 Short-Run Investment Response to Bonus Depreciation Allowance

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

$$\log i_{st} = b_s + b_t + b_1\zeta_{st} + b_2 \log q_{st} + \varepsilon_{st}, \tag{62}$$

where b_s is a sector fixed effect, b_t is a year fixed effect, and b_1 is the coefficient of interest. We also control for the relative price of investment $\log q_{st}$. Column (3) includes sector and time fixed effects and is therefore most directly comparable to Zwick and Mahon (2017)'s differences-in-differences specification. Our estimated $b_1 = 1.86$ is about half the size of Zwick and Mahon (2017)'s estimated 3.69, but still implies large effects of the bonus; for example, a 4.8 cent change in ζ_{st} (the average change in the early period according to Zwick and Mahon (2017)) would increase investment by about 9% for our estimate vs. 18% in Zwick and Mahon (2017)'s estimate.

C Calibration

This appendix contains additional results regarding the calibration of the model.

Identification of Long-Run Elasticities of Substitution. To understand identification in our framework, consider how a permanent decline in the relative price of capital q^* affects the two equations:

$$F_k\left(1, \{\frac{n_i}{k}\}\right) = q^*(r+\delta) \tag{63}$$

$$\frac{\kappa_i}{\lambda_w'(\theta_i)} = \frac{1}{r+\sigma}\left[F_{ni}\left(1, \{\frac{n_i}{k}\}\right) - v'(n_i) - \frac{\widetilde{\sigma}}{\omega}v'(n_i)\right]. \tag{64}$$

Equation (63) is the capital Euler equation and equation (64) is the firm's labor demand in steady state. First, the decline in q^* will increase the average capital-to-labor ratio in order to decrease the marginal product of capital in (63). The strength of this effect is captured empirically by our estimated capital-to-wage bill elasticity $d\log\frac{k}{wn}/d\log q$. Second, the higher capital-to-labor ratio will also increase the marginal products of labor and therefore equilibrium employment in (64). The magnitude of these changes for college vs. non-college workers depends on the degree of capital-skill complementarity. The response of the college share of income d college share $/d\log q$ is informative about the relative changes in marginal product using the wage equation $w_i = \eta\left(F_{ni} - \frac{\sigma}{\omega}v'(n_i)\right) + (1-\eta)v'(n_i)$.

Figure C.1 plots the values of these two key moments as a function of the elasticities of substitution ρ and α. The top right panel shows that the capital-to-wage bill elasticity primarily depends on the elasticity of substitution between capital and college labor α, and therefore primarily determines its value. Conditional on the value of α, the bottom left panel shows that the college income elasticity largely determines the elasticity of substitution ρ between non-college labor and the aggregate $G(k, n_h)$.

Figure C.1: Identification of Long-Run Elasticities of Substitution

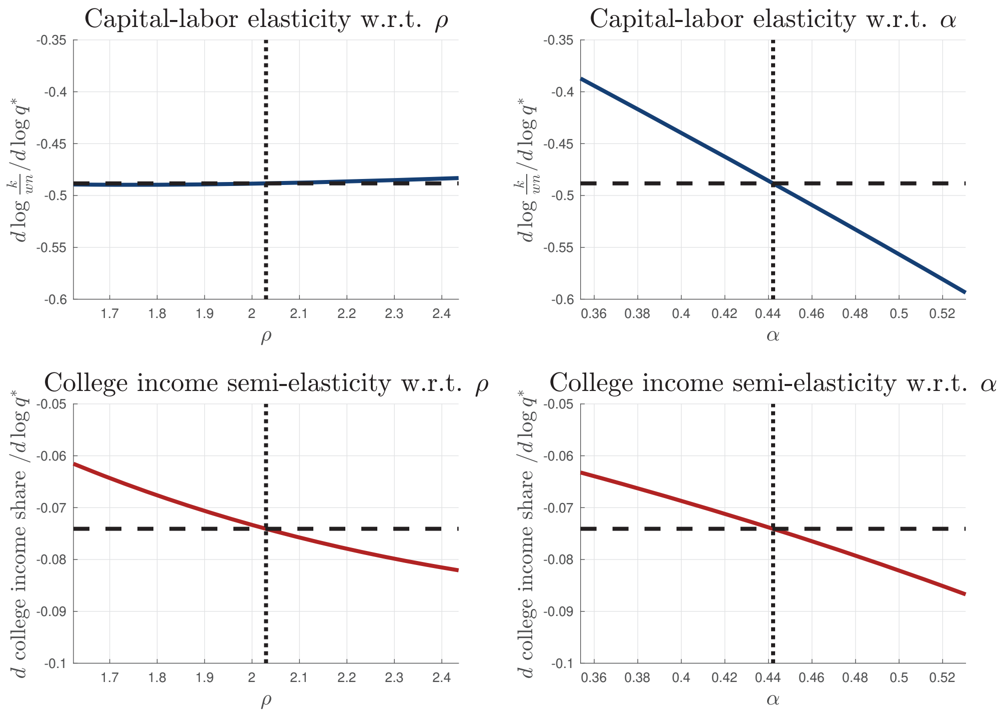

Notes: Calibration targets as function of parameters. "Capital-labor elasticity" refers to $d \log \frac{k}{wn} / d \log q$ and "college income elasticity" refers to d college share $/ d \log q$. The left column varies ρ holding $\alpha = 0.467$ fixed at its calibrated value. The right panel values α holding $\rho = 2.036$ fixed at its calibrated value.

Figure C.2: Response of Employment to Permanent q-Shock

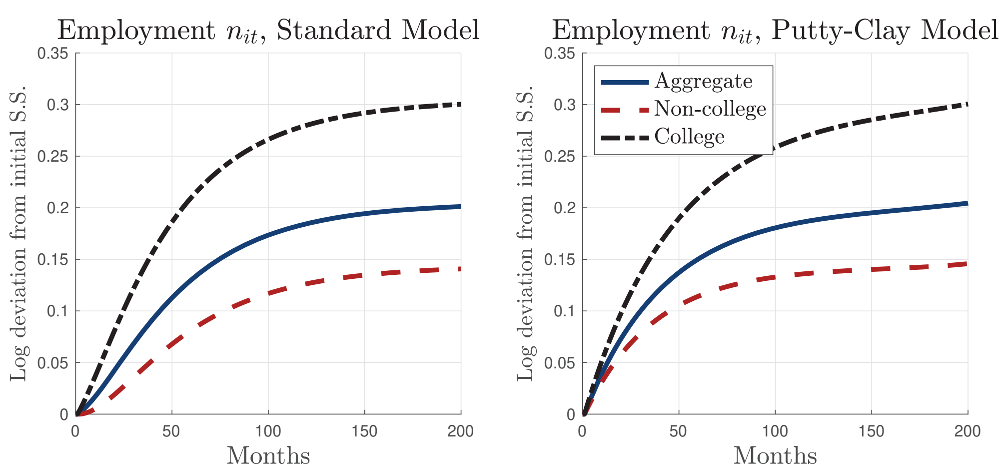

Notes: Transition paths following a permanent 1% decline in the relative price of capital q^*.

Transition Paths Following Permanent Shocks to q_t. Figures C.2 and C.3 plot the transition paths in response to a permanent decline in q^* in our calibrated model.

Figure C.3: Response of Capital-Labor Ratios to Permanent q-Shock

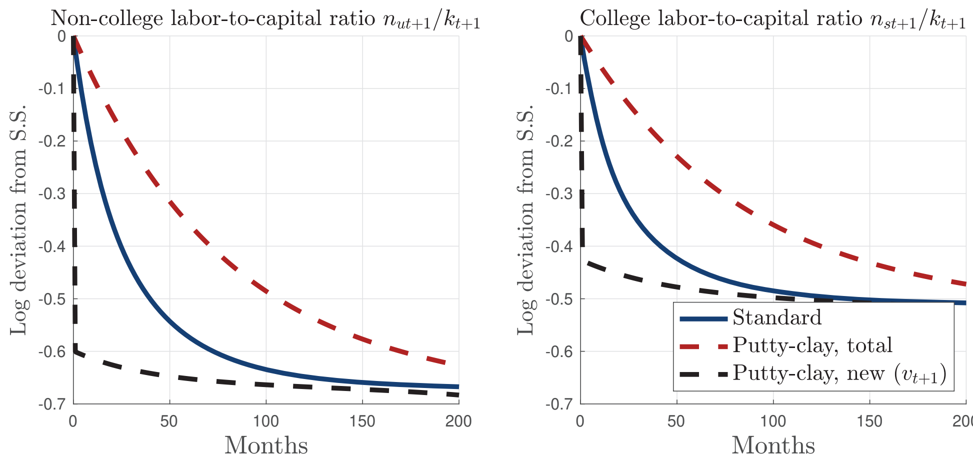

Notes: Transition paths following a permanent 1% decline in the relative price of capital q^*.

Figure C.4: Response of Employment to Bonus Depreciation Shock

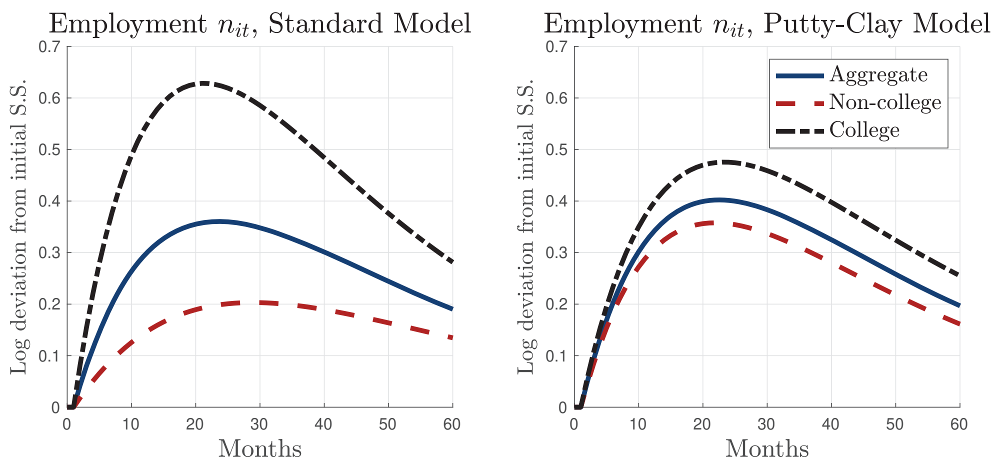

Notes: Transition paths following a transitory decline in q_t which mimics the bonus depreciation allowance.

Replicating Bonus Depreciation Shock in the Model. Figure C.4 plots the response of employment to the bonus depreciation allowance. Figure C.5 plots the response of labor-capital ratios to the bonus depreciation allowance.

D Minimum Wage in the Baseline Model

This appendix contains additional results regarding the effects of the minimum wage in the baseline model described in Section 4 of the main text.

Figure C.5: Response of Capital-Labor Ratios to Bonus Depreciation Shock

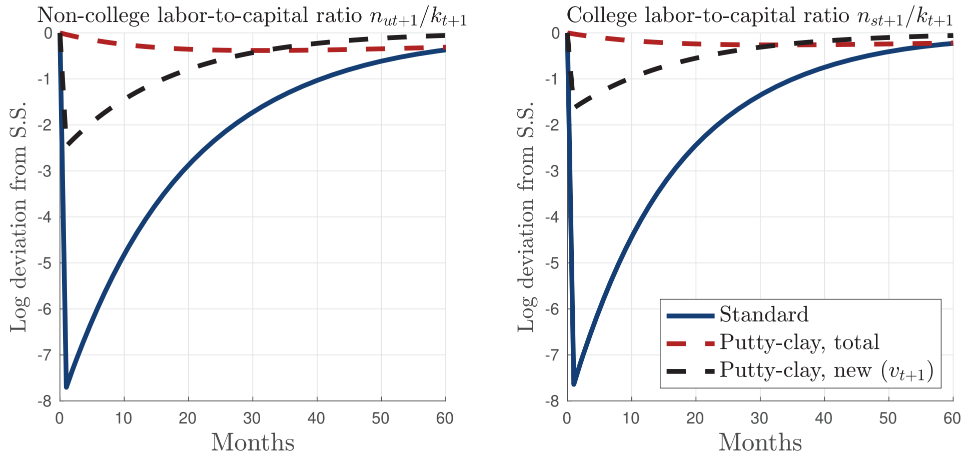

Notes: Transition paths following a transitory decline in q_t that mimics the bonus depreciation allowance.

Figure D.1: Aggregate Minimum Wage "Laffer Curves" (Comparative Statics w.r.t. ρ)

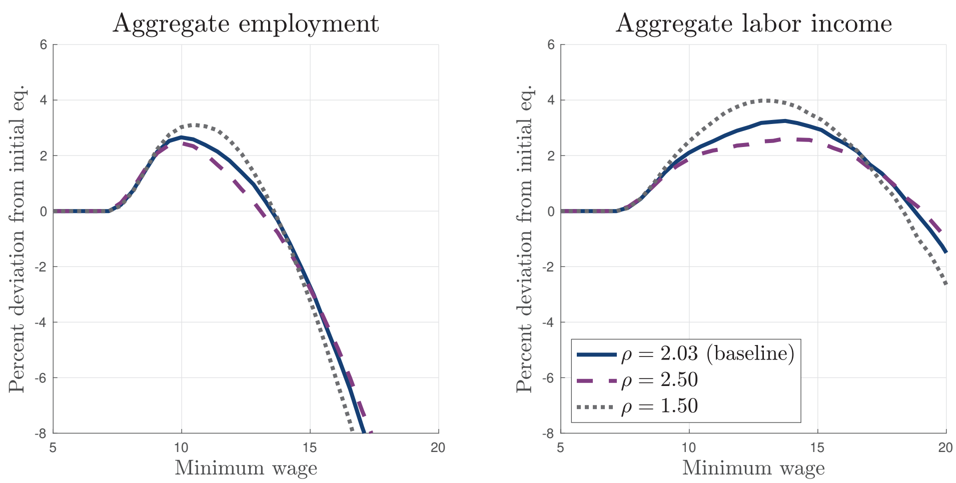

Notes: Steady-state aggregate employment (left panel) and aggregate labor income (right panel) as a function of the minimum wage.

D.1 Long-Run Effects of the Minimum Wage

Figure 5 shows the aggregate Laffer curves for different values of the elasticity of substitution ρ. Figure D.2 shows separate employment Laffer curves for non-college and college workers. Figure D.3 shows separate labor income Laffer curves for non-college and college workers. Figure D.4 plots the micro-level effect of the \$15 minimum wage on labor income. Figure D.5 shows how the micro-level effect of the minimum wage depends on the elasticity of substitution ρ.

Figure D.2: Employment“Laffer Curves,” College vs. Non-College

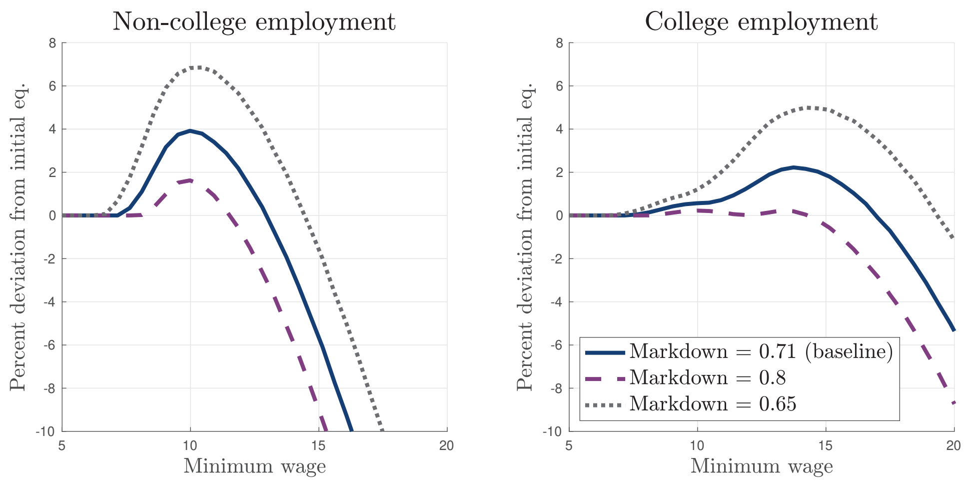

Notes: Steady-state aggregate college employment (left panel) and aggregate non-college employment (right panel) as a function of the minimum wage. Different markdowns correspond to different parameter values for ω.

Figure D.3: Labor Income“Laffer Curves,” College vs. Non-College

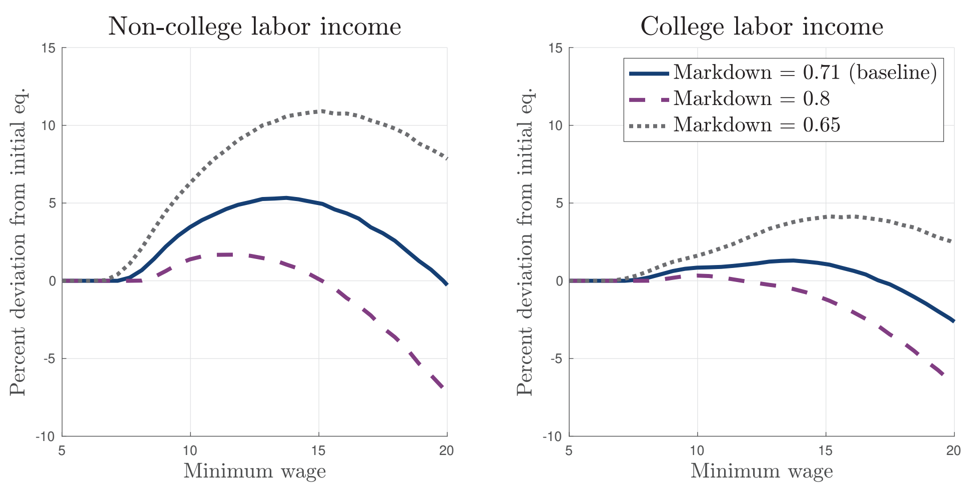

Notes: Steady-state aggregate college labor income (left panel) and aggregate non-college labor income (right panel) as a function of the minimum wage. Different markdowns correspond to different parameter values for ω.

FIGURE D.4: Micro-level Effect of \$15 Minimum Wage: Labor Income

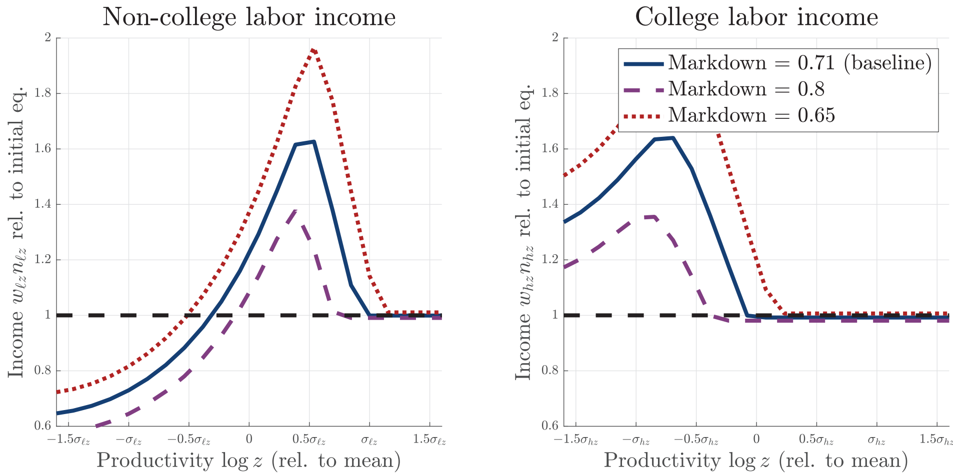

Notes: Steady-state labor income of particular z types for a \$15 minimum wage, relative to employment level without the minimum wage. Different markdowns correspond to different parameter values for ω.

FIGURE D.5: Micro-level Effect of \$15 Minimum Wage: Employment (Comparative Statics w.r.t. ρ)

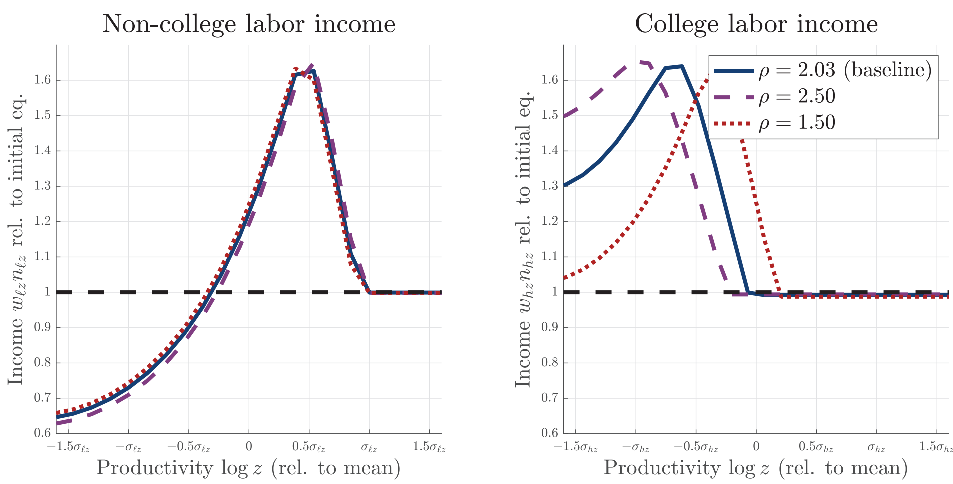

Notes: Steady-state labor income of particular z types for a \$15 minimum wage, relative to employment level without the minimum wage.

Figure D.6: Transition Paths of Employment to New Minimum Wage with Putty-Clay vs. Standard Capital

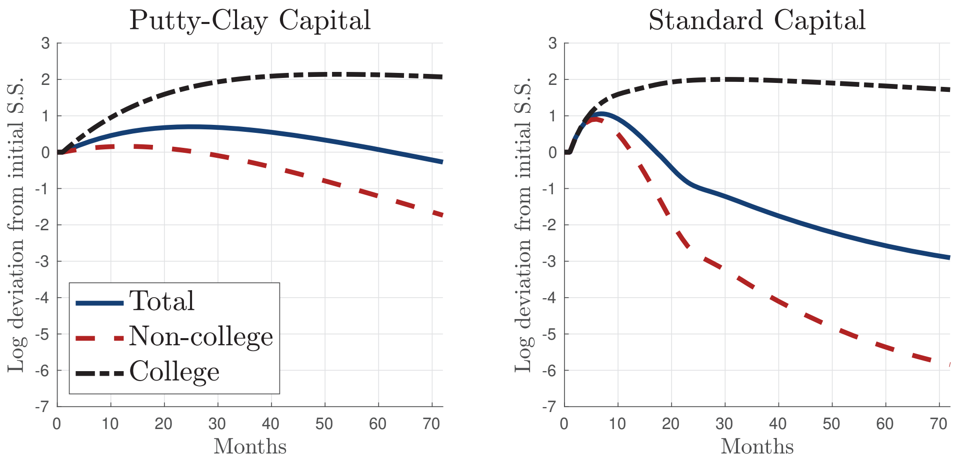

Notes: Transition paths of employment aggregated by group following an unexpected imposition of the minimum wage $\overline{w}$, starting from the initial equilibrium with $\overline{w} = 0$. The left panel is the model with putty-clay capital (as in the main text) capital and the right panel is the model with standard capital.

D.2 The Short Run vs. Long Run

Figure D.6 compares the transition paths of aggregated employment in the putty-clay model to the version of our model with standard capital. In the standard model, firms immediately reduce their employment of low-z types by doing no new hiring and instead letting those workers attrit at the job destruction rate σ.[43] These dynamics imply that non-college employment converges to the new steady state in approximately six years; by this point, the putty-clay model is less than 1/3 to the new steady state.

Figure D.7 shows the variable profits in the initial period from fully operating the initial capital stock as a function of the minimum wage.

E Alternative Policies in the Baseline Model

This appendix contains additional results regarding the effects of the alternative policies described in Section 5 of the main text.

[43]Firms do not hire the low-z types because the present value of their marginal products is strictly below the present value of minimum wage payments. Therefore, in principle these firms would like to furlough some of their initial stock of workers as well, and we are working on allowing for these furloughs in the standard model. The furloughs would strengthen our results because they imply that employment in the standard model converges to the new steady state even faster than is shown here. In this sense, the results in Figure D.6 represent a lower bound on the speed of transition in the standard model.

Figure D.7: Putty-Clay Shutdown Condition in Initial Period

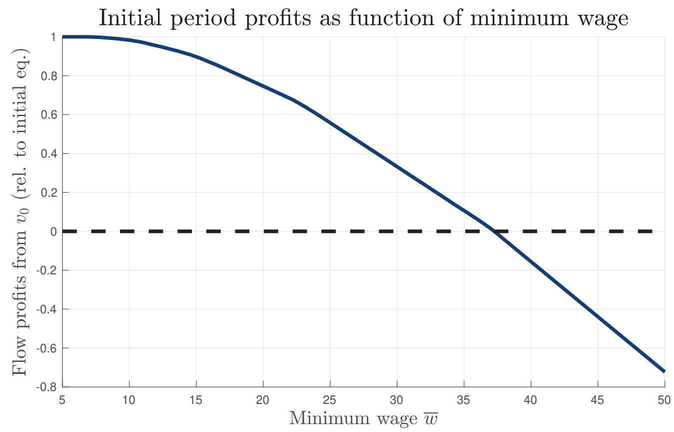

Notes: Initial period profits from fully operating steady state capital stock from the decentralized equilibrium as a function of the minimum wage. Normalized so that profits in initial steady state are 1.

Figure E.1: Effect of Uniform Tax Credit on Labor Income

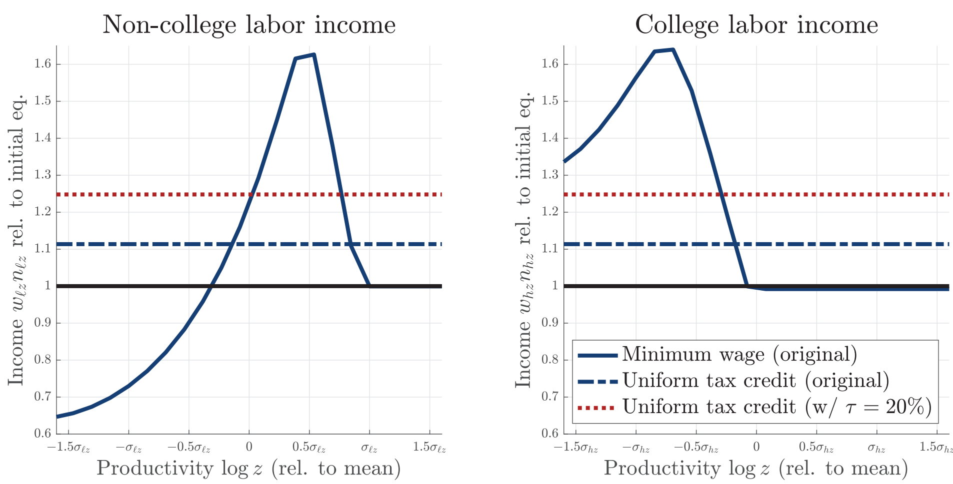

Notes: Steady-state labor income of particular z types relative to labor income level in the initial equilibrium (without minimum wage or alternative policy). The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

E.1 Uniform Tax Credit and Tax Cut

Figure E.1 shows the effect of the uniform tax credit and tax cut on labor income. Figure E.2 shows the effect on the post-tax markdown $A(w_i)/F_{ni}$.

E.2 Progressive Tax and Transfer System

Figure E.3 plots the effect of the progressive tax and transfer system on labor income. Figure E.4 plots the effect of the progressive tax and transfer system on the post-tax markdown $A(w_i)/F_{ni}$.

Figure E.2: Effect of Uniform Tax Credit on (Post-Tax) Markdown

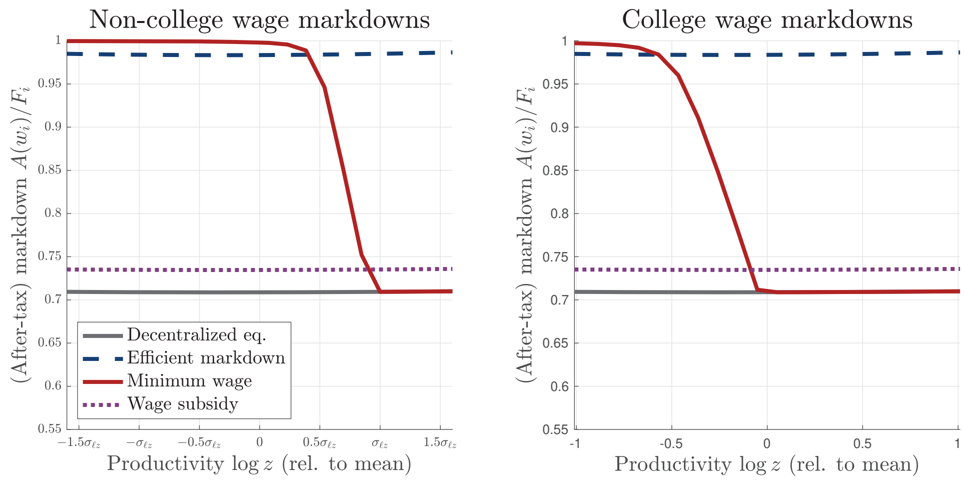

Notes: Steady-state post-tax wage markdowns $A(w_i)/F_{ni}$ of particular z types. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

Figure E.3: Effect of Progressive Tax and Transfer System on Labor Income

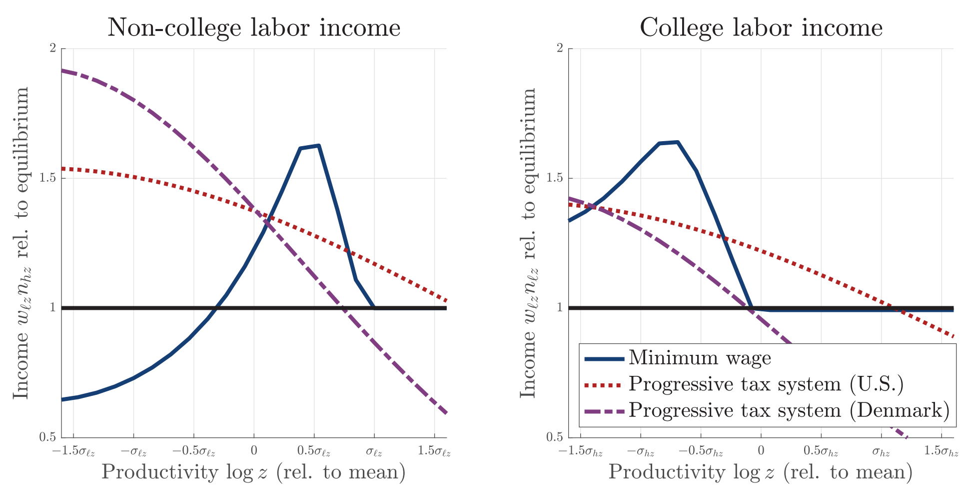

Notes: Steady-state labor income of particular z types relative to labor income level in the initial equilibrium (without minimum wage or alternative policy). Blue line is with the \$15 minimum wage, dashed red line is for the progressive tax system at the U.S. level of progressivity 0.181, and purple line is for progressive tax system at the Danish level of progressivity 0.463. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

Figure E.4: Effect of Progressive Tax and Transfer System on (Post-Tax) Markdown

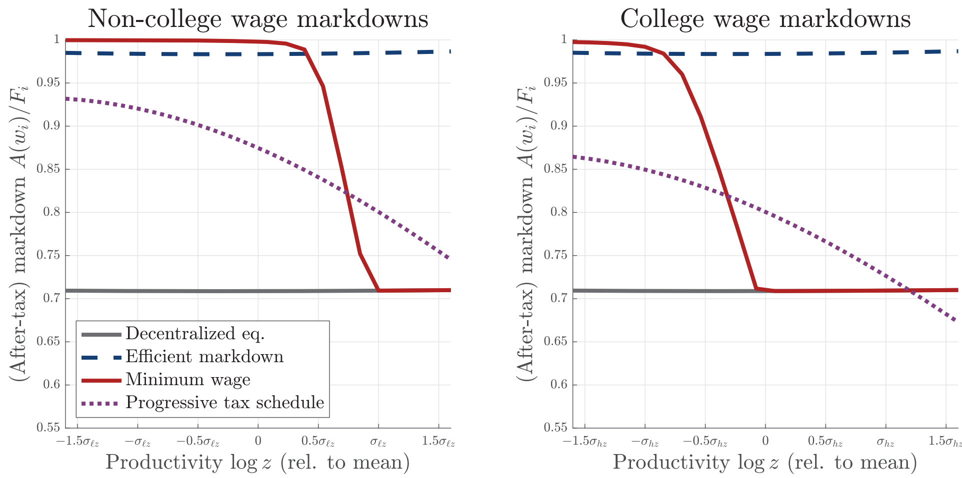

Notes: Steady-state post-tax wage markdowns $A(w_i)/F_{ni}$ of particular z types. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

Figure E.5: Effect of Earned Income Tax Credit on Labor Income

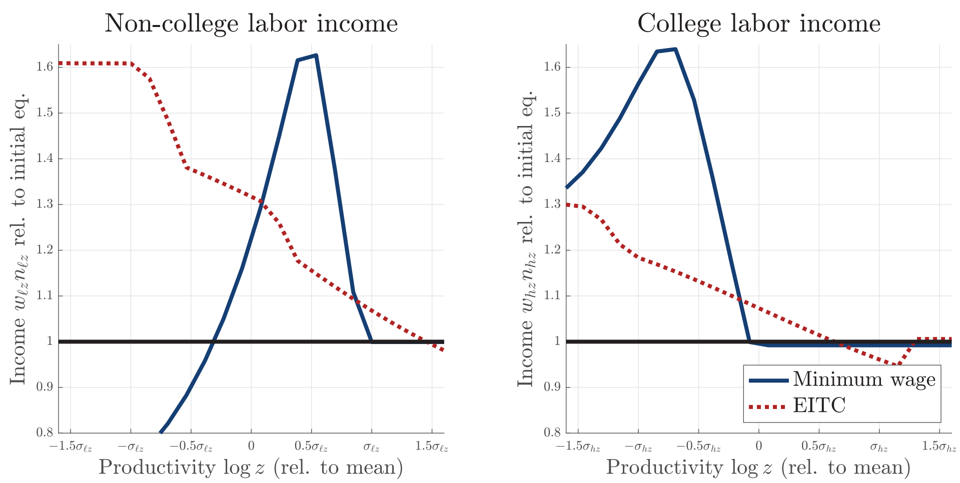

Notes: Steady-state labor income of particular z types relative to labor income level in the initial equilibrium (without minimum wage or alternative policy). Blue line is with the $15 minimum wage and dashed red line is for our EITC schedule. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

E.3 Earned Income Tax Credit

Figure E.5 shows the effect of the EITC on labor income. Figure E.6 shows the effect of the EITC on the past-tax markdown $A(w_i)/F_{ni}$.

Figure E.6: Effect of Earned Income Tax Credit on (Post-Tax) Markdown

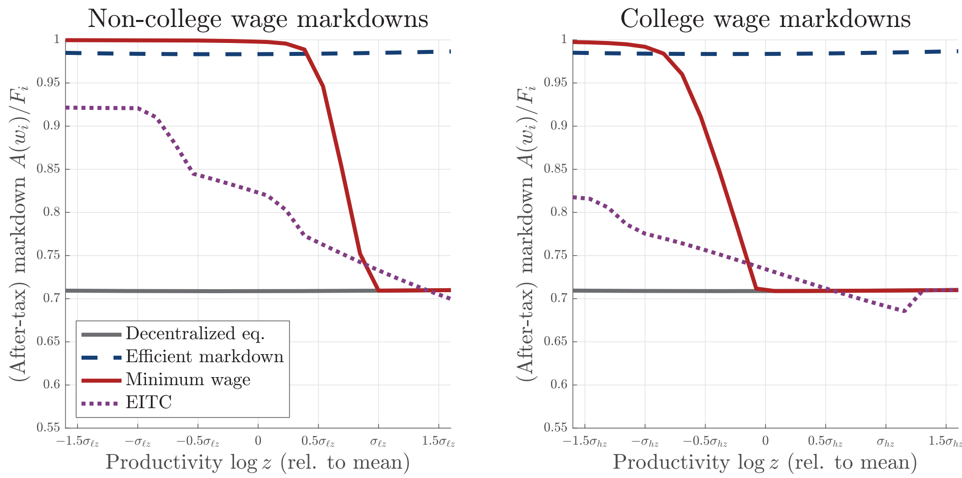

Notes: Steady-state post-tax markdown $A(w_i)/F_{ni}$. Steady-state post-tax wage markdowns $A(w_i)/F_{ni}$ of particular z types. The x-axis is log individual productivity z relative to its mean value, expressed in standard deviations from the mean.

F Minimum Wage Within a Tax and Transfer System

Table F.1 shows the fitted parameters in the calibration with the U.S. tax and transfer system. Table F.2 shows how well the model fits the empirical targets. Figure F.1 shows the effect of the \$15 minimum wage on employment. Figure F.2 shows the effect of the \$15 minimum wage on welfare. Finally, Figure F.3 plots our aggregate employment and labor income Laffer curves in our baseline model vs. full model.

Table F.1: Fitted Parameters (Model with U.S. Tax System)

Parameter	**Description**	**Value**
Labor market frictions		
κ_0	Scale of vacancy posting cost	1.58
ω	Monopsony power	0.04
Worker productivity distribution $\log \mathcal{N}(\mu_b, \sigma_b)$		
$\overline{\mu}_l$	Mean of non-college z (normalization)	0.00
σ_l	SD of non-college z	0.65
$\overline{\mu}_h$	Mean of college z	0.75
σ_h	SD of college z	0.71
Production function		
α	Long-run elasticity of substitution b/t k and n_h	0.47
ρ	Long-run elasticity of substitution b/t n_l and $G(k, n_h)$	2.04
ψ	Coefficient on non-college labor n_ℓ	0.61
λ	Coefficient on capital k	0.58

Notes: Parameters endogenously chosen to match the moments in Table F.2.

Table F.2: Targeted Moments (Model with U.S. Tax System)

Moment	Description	Data	Model
Average unemployment rate			
$\mathbb{E}[s_i]/(\mathbb{E}[s_i]+\mathbb{E}[n_i])$	Average unemployment rate	5.9%	5.6%
Average wage markdown			
$\mathbb{E}[w_{ni}]/\mathbb{E}[F_{ni}]$	Average wage markdown (Berger et al. 2021)	0.71	0.67
Wage Distribution, ACS 2017-2019			
$w_{h50}/w_{\ell 50}$	College wage premium	1.81	1.84
$w_{\ell 50}/w_{\ell 10}$	Non-college 50-10 ratio	2.04	2.06
w_{h50}/w_{h10}	College 50-10 ratio	2.30	2.29
Response to capital price decline (our data)			
$d\log\frac{k}{wn}/d\log q$	Response of capital-labor ratio	-0.47	-0.51
d college share $/d\log q$	Response of college inc. share	-0.09	-0.07
Average income shares			
$\mathbb{E}[w_i n_i]/Y$	Aggregate labor share	0.57	0.60
$\pi_h\mathbb{E}[w_{hz}n_{hz}]/\mathbb{E}[w_i n_i]$	College income share	0.55	0.53

Notes: Moments targeted using parameters in Table F.1.

Figure F.1: Micro-level Effect of \$15 Minimum Wage: Employment (Model with U.S. Tax System)

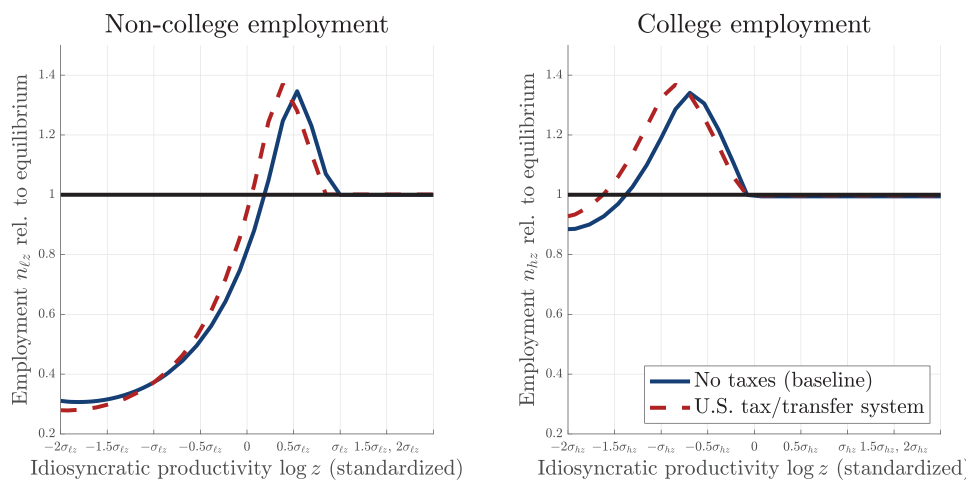

Notes: Steady-state employment of particular z types for a \$15 minimum wage in the model. The blue line corresponds to our baseline model without taxes from Section 3 and the red line corresponds to the model with taxes. The y-axis is normalized relative to the employment level in the initial steady state without the minimum wage. The x-axis is individual productivity relative to its mean value, expressed in standard deviations from the mean.

Figure F.2: Micro-level Effect of $15 Minimum Wage: Welfare (Model with U.S. Tax System)

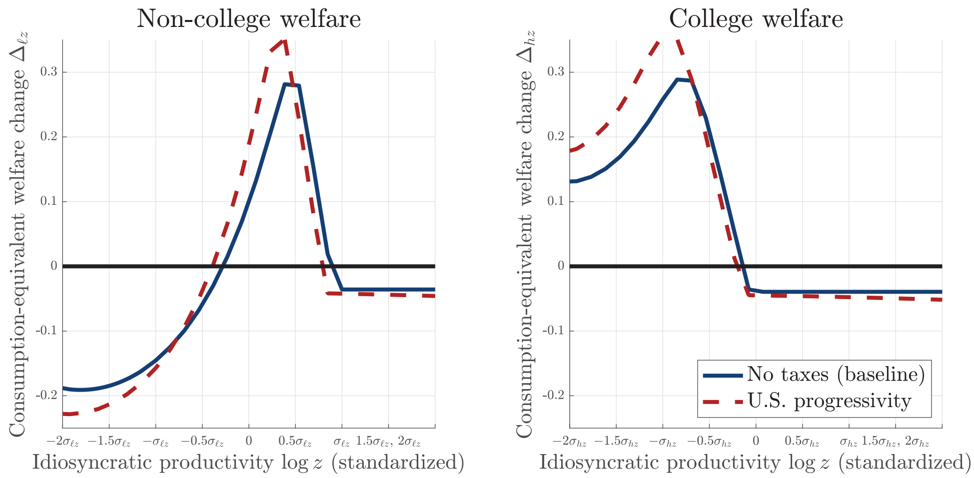

Notes: Steady-state welfare Δ_i of particular z types for a $15 minimum wage in the model. The blue line corresponds to our baseline model without taxes from Section 3 and the red line corresponds to the model with taxes. The x-axis is individual productivity relative to its mean value, expressed in standard deviations from the mean.

Figure F.3: Aggregate Minimum Wage Laffer Curves (Model with U.S. Tax System)

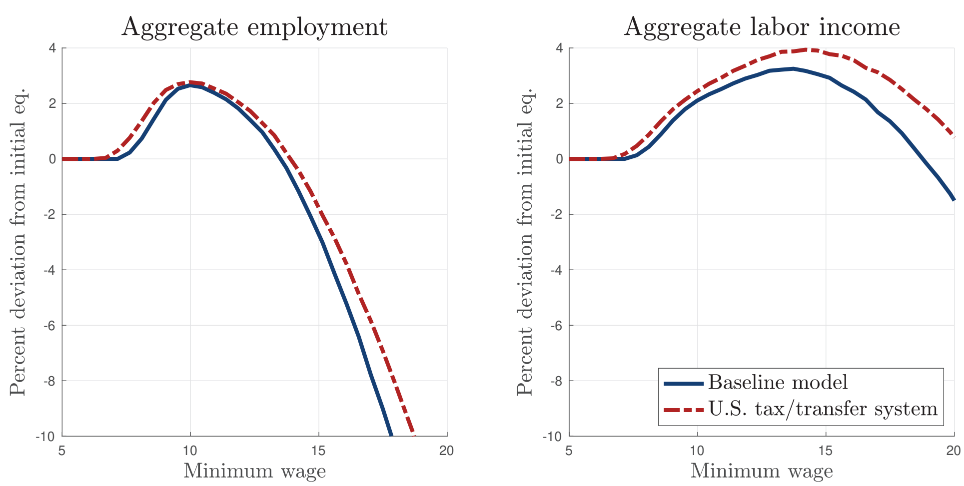

Notes: Steady-state aggregate outcomes as a function of the minimum wage $\overline{w}$. The blue line corresponds to our baseline model without taxes from Section 3 and the red line corresponds to the model with taxes.