The Race Between Education, Technology, and the Minimum Wage[*]

Jonathan Vogel[†]

January 2025

Abstract

What is the impact of the real minimum wage on wages and inequality? I develop a theory that nests "The Race Between Education and Technology." The theory predicts that the impact of changes in minimum wages is initially small but grows over time. I present empirical evidence of these dynamic effects: the elasticity of real wages grows substantially over a three-year period. I additionally show that minimum wages help rationalize a considerable decline in real wages of low-education workers in the 1980s and play a role in shaping the evolution of the U.S. college premium.

[*]I thank David Autor, Ariel Burstein, Arnaud Costinot, Nicole Gorton, Daniel Haanwinckel, Lawrence Katz, Thomas Lemieux, Christian Moser, Andres Santos, Liang Tan, and Pierre-Olivier Weill for helpful comments. This paper supersedes and expands the paper "The Race Between Education, Technology, and Institutions."

[†]University of California, Los Angeles (jvogel@econ.ucla.edu) and NBER

1 Introduction

What are the dynamic impacts of minimum wages on real wages and inequality? I present a theory in which the elasticity of real wages with respect to the real minimum wage is small on impact—with an increase in the minimum wage raising only the wages of workers whose wages are below the new minimum—and grows over time. I present empirical evidence consistent with these dynamic effects. And I use the model and empirical evidence to document that minimum wages—together with supply and demand—play a central role in shaping the evolution of the U.S. college wage premium and that declining real minimum wages in the 1980s in the U.S. caused a considerable decline in the real wages of workers without a high school degree.

To motivate the analysis, in Section 2 I show that the real minimum wage helps shape even a very aggregated measure of inequality—the college wage premium—even at the national level.[1] *The Race Between Education and Technology* or *the canonical model*—introduced by Tinbergen (1974) and Freeman (1976), operationalized by Katz and Murphy (1992), and named by Goldin and Katz (2008) and Acemoglu and Autor (2011)—provides the central organizing framework for studying the evolution of the college premium. It straightforwardly relates relative wages of more and less educated workers to their relative supply and the skill bias of demand. I estimate an extended version of the empirical canonical model, incorporating the real minimum wage. I find that the elasticity of the national college premium with respect to the real minimum wage is approximately -0.13, implying (for example) that the 26% decline in the real minimum wage between 1979 and 1989 increased the college premium by 3.4%, which is about a quarter of its observed increase.

Introducing the real minimum wage additionally alters a standard conclusion in the literature, that the rate of skill-biased technical change has declined dramatically in the 1990s and thereafter. The canonical model estimated on data spanning 1963 – 1987 (the set of years used in the seminal work of Katz and Murphy, 1992) predicts substantially more rapid increases in the college wage premium than actually occur in the data thereafter. Through the lens of the canonical model, this problem with its out-of-sample fit implies a substantial decline in the rate of skill-biased technical change. This issue is mitigated by incorporating the real minimum wage; its rise in the later period reduces the predicted growth rate of the college premium in the absence of changes in the rate of skill-biased technical change.

In Section 3, I present a generalized version of the canonical model that microfounds

[1] I focus in Section 2 on the college premium because its evolution plays a central role in shaping overall U.S. wage inequality. For instance, Autor, Goldin and Katz (2020) find that increased returns to post-secondary schooling generate 57% of the observed increase in wage variance between 1980 and 2017.

the national regression analysis of Section 2; rationalizes its findings of a substantial elasticity of the college premium with respect to the real minimum wage in spite of the small share of workers actually earning the minimum wage; introduces novel predictions on the dynamic implications of changes in the minimum wage on which the results in Section 2 are silent; and guides the subsequent empirical analysis in Sections 4 and 5, which uses disaggregated data and focuses directly on these dynamics. My objective is to maintain the simplicity and tractability of the original framework yet facilitate the study of the dynamic implications—both across steady states and in transitions—of real minimum wages on real wages and inequality. In addition to supply and demand, the model incorporates monopsony power, minimum wages, a job-ladder, dynamics, and unemployment. As in the canonical model, an aggregate constant returns to scale production function combines the output of different skills. Unlike the canonical model, the labor market is frictional. A worker who meets a potential new employer bargains over her wage with her current wage (or unemployment benefit) serving as her outside option.

In equilibrium, a worker exits unemployment at the minimum wage and slowly moves up the job ladder to higher wages as she matches with new employers over time.[2] In the steady state, the average real wage of a given skill group is a weighted average of the real minimum wage and that group's real value marginal product of labor. I characterize how changes in supply, demand, and the minimum wage affect real wages across steady states. I focus on a baseline case in which the minimum wage has no effect on unemployment, in which case changes in supply and demand shape value marginal products of labor (exactly as in a frictionless and competitive model with the same aggregate production function) whereas changes in the minimum wage affect wage markdowns; I additionally consider the case in which unemployment depends on the minimum wage. I show that, for any aggregate production function and any number of worker skills, the steady-state elasticity of a skill group's average real wage with respect to the real minimum wage can be expressed as the product of the share of that group's wage income earned at the minimum wage (*the minimum wage bite*) and a term that is strictly greater than one, which I refer to as the steady-state *magnification elasticity*. Under the additional assumptions of the traditional canonical model—two skills, a CES aggregate production function, constant factor-biased productivity growth rates—I micro-found the extended canonical model estimating equation of Section 2, to a first-order approximation. Theoretically, the elasticity of the college wage premium with respect to the real minimum wage estimated

[2]Here, I describe model implications focusing on workers in a skill group that exits unemployment at the minimum wage. In the model, other worker groups may exit unemployment at a wage higher than the minimum wage. The case described here is particularly relevant for lower-education and younger workers in the data.

in Section 2 equals the product of the long-run magnification elasticity and the difference between the minimum wage bite of college-educated and non-college-educated workers.

I also characterize the transition in response to a one-time increase in the real minimum wage. On impact, its increase raises real wages only for workers whose initial wage is bound by the new minimum wage. Hence, on impact the elasticity of a skill group's average real wage equals its minimum wage bite. Over time, however, workers initially at the new minimum wage slowly rise up the new job ladder, and the elasticity of the average real wage rises monotonically over time, with the time-varying magnification elasticity starting at one and converging to its steady-state level.

The model in Section 3 micro-founds and helps interpret the regression results not only in Section 2, but also in Sections 4 and 5, where the analysis is conducted at a more disaggregated level and focuses on dynamics. In Section 4 I show that the model of Section 3 implies that the change in the real wage of any group between period $t - T$ and period t in response to a one-time change in the real minimum wage can be expressed as a distributed lag model in which the independent variables are interactions between the time-varying initial minimum wage bites and lagged one-period changes in the real minimum wage. The lag weights to be estimated are exactly equal to the corresponding time-varying magnification elasticities for lagged shocks that occur between periods $t - T$ and t.

I estimate this regression defining a skill group in the data by the interaction of 100 labor groups (defined by age, gender, race, and education) and 50 regions (the U.S. states) and defining time periods in the data as half-year periods between 1979 and 2016. The dependent variable is the three-and-a-half-year change in the real wage of each of the region-group pairs for each time period. I instrument for the independent variables by replacing the time-varying minimum wage bite for a given region and group with its median value across time periods (leaving out the current period) to address issues of measurement error similar to those emphasized in a related context by Autor, Manning and Smith (2016). Finally, I use a "design-based" specification in the spirit of Borusyak and Hull (2023), controlling for the expectation of treatment under an assumption on the process of real minimum wage changes.

On impact, the magnification elasticity is 0.6. It rises slowly, peaking at 2.9 at a two- to three-year time horizon. Through the lens of the theory, this is consistent with a long-run magnification elasticity 2.9.[3] I then consider the implications of this magnification elastic-

[3] I consider sensitivity analyses that fit into three broad categories: I check for heterogeneity in magnification elasticities across worker groups by replicating the baseline empirical analysis separately for workers with different observable characteristics, I consider alternative sample years, and I include a broader set of controls.

ity for the college premium, revisiting the question in Section 2. Model results in Section 3, the average minimum wage bite across workers with and without a college education, and the long-run magnification elasticity of Section 5 together imply an elasticity of the college premium with respect to the real minimum wage of approximately -0.13, consistent with the estimates in Section 2. From both exercises, I conclude that the elasticity of the college premium with respect to the real minimum wage is sizable.

The theoretical and empirical results in Section 3 and 5 shed light on the impact of minimum wages not only on relative wages, but also on real wage levels. In the U.S., real wages of workers without a high school degree declined dramatically starting in the 1980s and into the early 1990s. In the extended canonical model a sizable decline in the real minimum wage (as occurred in the 1980s) substantially reduces the real wage of any labor group with a large minimum wage bite (as is the case for workers without a high school degree). According to the model and estimated long-run magnification elasticity, the 26% decline in the real minimum wage between 1979 and 1989 caused more than the entirety of the observed decline in the real average wage of workers without a high school degree between 1979 and 1992.

Additional Literature. In terms of static economic questions, my paper is perhaps most related to Autor, Katz and Kearney (2008), who estimate a regression very similar to my motivating empirics in Section 2 leveraging national-time series variation. They contend that the real minimum wage "does not much alter the central role for relative supply growth fluctuations and trend demand growth in explaining the evolution of the college wage premium" and that "institutional factors are insufficient to resolve the puzzle posed by slowing trend relative demand for college workers in the 1990s." They reach their conclusions because they find that the coefficient on the real minimum wage is negative, significant, and of a similar magnitude to my estimates in one of their two specifications (with a linear time trend) but smaller and insignificantly different from zero in their other specification (with a cubic time trend). I conclude differently. I show that my conclusions differ from theirs for three reasons: I consider a broader set of specifications, I have access to a longer sample, and I measure the real minimum wage using the average of minimum wages across states instead of the federal minimum wage (which introduces attenuation bias). More importantly, my empirical results additionally hold using state-and-worker-group-level variation. Based on these empirical findings, I conclude that relative supply growth fluctuations and trend demand growth remain crucial drivers of the college premium, consistent with Autor, Katz and Kearney (2008), but so too are changes in the real minimum wage. I also conclude that changes in the minimum wage go some ways towards resolving the puzzle posed by slowing trend relative demand for college workers in the 1990s. Finally, I also micro-found the regression they and I estimate.

A vast body of work studies the roles of labor-market institutions for shaping inequality. Much of the focus is on minimum wages (e.g., DiNardo, Fortin and Lemieux, 1996; Lee, 1999; Card and DiNardo, 2002; Teulings, 2003; Autor, Manning and Smith, 2016; Cengiz, Dube, Lindner and Zipperer, 2019; Dube, 2019; Chen and Teulings, 2022) and monopsony power (see Manning, 2013 for a summary). The analysis in my paper differs from this literature in a few ways. First, and most importantly, I focus on the time-varying impact of minimum wages on real wages and inequality whereas this literature does not focus on dynamics. Second, this literature often focuses on the impact of minimum wage changes on wage centiles (or the number of workers in each wage interval), whereas my empirical focus is on the impact of minimum wages on average real wages for each labor group. In this respect, my focus allows me to control explicitly for any potential changes in the observable composition of the labor force that might be correlated with changes in the minimum wage; but this comes at the cost of not studying within-group inequality empirically (because the data is too sparse to analyze changes in the distribution of wages within each labor group and state). Nevertheless, my short-run theoretical and empirical results are consistent with this literature's empirical results that the direct effect is the dominant mechanism through which the minimum wage affects inequality (see, e.g., Autor, Manning and Smith, 2016; Cengiz, Dube, Lindner and Zipperer, 2019). My theory, however, implies that the indirect effect strengthens over time. In my empirics, the indirect effect is stronger than the direct effect two to three years after the minimum wage change. Moreover, I show analytically that a job-ladder model can rationalize these empirical findings both across steady states and in the full transition.[4]

My paper is also related to a recent and fast growing macro-labor literature conducting counterfactual analyses to study the implications of minimum wage policies for inequality, efficiency, and welfare (e.g., Haanwinckel, 2023; Ahlfeldt, Roth and Seidel, 2022; Berger, Herkenhoff and Mongey, 2022; Hurst, Kehoe, Pastorino and Winberry, 2022; Engbom and Moser, 2022; Trottner, 2022). In contrast, I use my qualitative theory to guide and interpret my empirical analyses of the impact of minimum wages on real wages and wage inequality. In one respect, my analysis is closest to Haanwinckel (2023), who introduces a model of labor supply, labor demand, and the minimum wage to study the impact of changes in the minimum wage on inequality in Brazil. In another respect, my analysis is closest to Engbom and Moser (2022), who use a job-ladder model to study the

[4]While I show that a job ladder model can rationalize my empirical findings, I neither claim that only job ladder models are consistent with the data nor do I directly test that the indirect effects I find are generated by movements up the job ladder; such an exercise would require worker-level wage panel data.

implications of minimum wages in Brazil. But I am particularly interested in dynamic effects, from which these papers abstract. In this respect, my analysis is most related to Hurst et al. (2022). Their dynamics are driven primarily by putty-clay capital as opposed to job ladders.[5]

2 Motivation: national time series evidence

The Race Between Education and Technology (also known as the the canonical model) provides the central organizing framework for studying the evolution of the college premium. It relates relative wages of more (h) and less (ℓ) educated workers $w_{ht}/w_{\ell t}$ at time t, to their relative supply $Supply_{ht}/Supply_{\ell t}$ and demand $A_{ht}/A_{\ell t}$. Relative demand captures the extent to which technical change is skill biased. In the typical empirical implementation, wages and supplies are observed, whereas the skill bias of technology is unobserved. A linear time trend $\gamma \times t$ proxies for smooth changes in the rate of skill-biased technological change, leading to the following empirical specification

$$\log\left(\frac{w_{ht}}{w_{\ell t}}\right) = \alpha + \nu_L \log\left(\frac{Supply_{ht}}{Supply_{\ell t}}\right) + \gamma t + \iota_t \tag{1}$$

I consider an atheoretical (for now) extension of the canonical model that incorporates the value of the real minimum wage, m_t, as follows

$$\log\left(\frac{w_{ht}}{w_{\ell t}}\right) = \alpha + \nu_m \log m_t + \nu_L \log\left(\frac{Supply_{ht}}{Supply_{\ell t}}\right) + \gamma t + [...] + \iota_t \tag{2}$$

I refer to equation (2) as both *the race between education, technology, and the minimum wage* and, more compactly, as *the extended canonical model*.

Measurement. Here I briefly describe how I measure each variable in equations (1) and (2); Online Appendix A.1 contains details. I restrict my sample to the working-age population of those between 16 and 64 years old and define more educated workers as those with 16 or more years of education (college educated) and less educated workers as those with fewer than 16 years of education (non-college educated). I construct the composition-adjusted college premium by first measuring the log of the average hourly wage within each of 100 groups (defined by the interaction of 5 age bins, 2 genders, 2 races, and 5 education levels) in year t and then averaging across those groups with and those groups without college education using time-invariant weights. I similarly con-

[5] Clemens and Strain (2021) also focus on dynamics, but dynamic employment rather than wage effects. They document that medium-run employment effects exceed short-run effects.

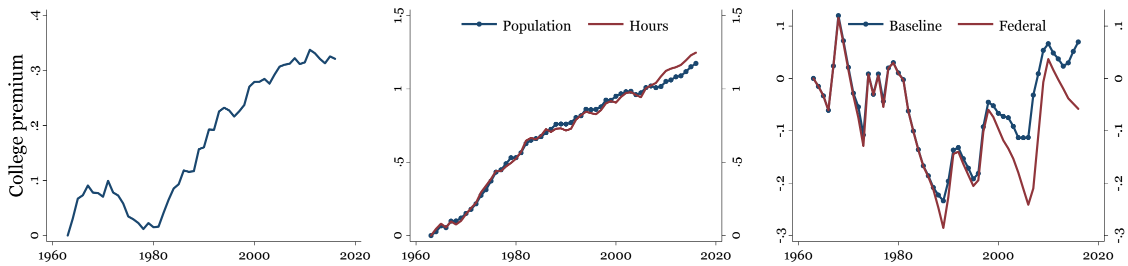

(a) The College Premium

(b) The Relative Supply of College Workers

(c) The Real Minimum Wage

Figure I: National Relative Wages, Relative Supplies, and Real Minimum Wages

Notes: Panels (a) and (b) display the national composition-adjusted college premium and relative supply of college labor. The relative supply of college labor is measured two ways: hours worked and populations. Panel (c) displays the baseline real minimum wage series, which averages minimum wages across states, and the real federal minimum wage; both series are deflated by the GDP deflator. All series are normalized to zero in 1963.

struct the supply of college and non-college workers as a composition-adjusted average of efficiency-unit hours worked across each of the groups with and without college education. Since hours worked in each year t for high- and low-education workers respond to contemporaneous deviations from the linear rate of skill-biased technical change, and since these deviations are included in the residual ι_t, I additionally construct an instrument for relative supply, using composition-adjusted populations across groups with and without college education. I measure these variables using the March Annual Demographic Files of the Current Population Survey from 1964 to 2017, which report earnings from 1963 to 2016.

I measure the real minimum wage in year t as the average minimum wage across states (the maximum of the legislated state and federal minimum wages), using time-invariant weights. In an alternative approach, I use the federal (Fair Labor Standards Act or FLSA) minimum wage. I deflate each series using the GDP deflator. I refer to my baseline measure as the real minimum wage and to my alternative measure as the real federal minimum wage.

Panels (a) and (b) of Figure I display the college premium and the two versions of the relative supply of college workers (measured using hours and populations), each normalized to zero in 1963. The particularly steep rise in the college premium starting in the early 1980s coincides with a decline in the growth rate of the relative supply of college workers, a fact first emphasized by Katz and Murphy (1992). Panel (b) additionally displays the strong (first-stage) relationship between the hours-worked and population-based measures of the relative supply of college workers. Panel (c) displays the two real minimum wage series (one measured using the average across states and the other using the fed-

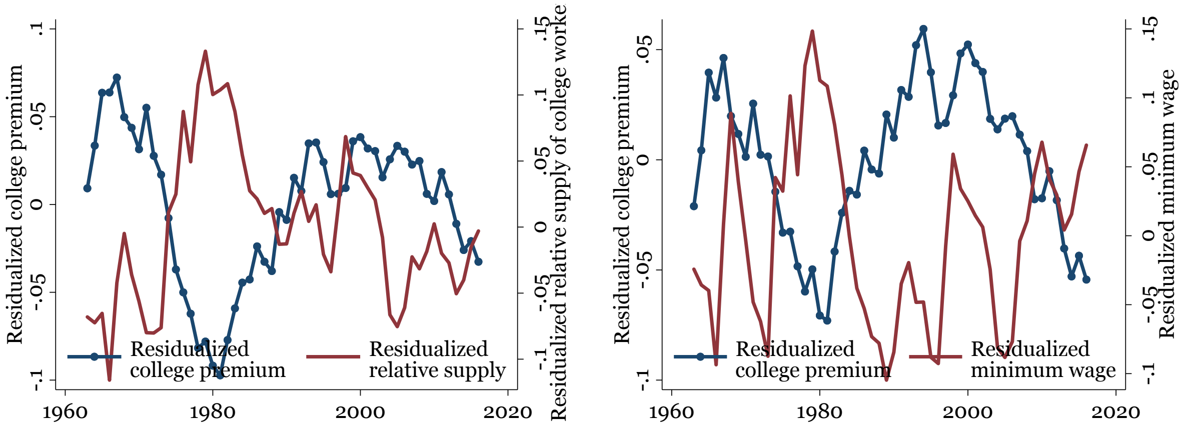

(a) Residualized Relative Supply, College Premium **(b)** Residualized Minimum Wage, College Premium

Figure II: Residualized Independent and Dependent Variables at the National Level

Notes: Panel (a) displays the national composition-adjusted college premium and relative supply, both residualized of the real minimum wage and a linear time trend. Panel (b) displays the national composition-adjusted college premium and real minimum wage, both residualized of composition-adjusted relative supply and a linear time trend. In each panel, the relative supply of college labor is measured using the population-based measure.

eral minimum wage), each normalized to zero in 1963. The two series move in lockstep until the late 1980s and diverge thereafter (especially in the 2000s) as more states set minimum wages above the federal level. Both real minimum wage series decline dramatically in the 1980s, as fixed nominal minimum wages are eroded by inflation; the real federal minimum wage experiences sharp swings in the 2000s as well, although these are less dramatic in my baseline series as states raise their nominal minimum wages to dampen the impact of inflation. Finally, there is substantial time variation in both real minimum wage series, which remain even after residualizing on high-dimensional polynomials of time.

Results. Before providing estimation results, I display the variation in the data that identifies ν_m and ν_L in the extended canonical model in equation (2). Panel (a) of Figure II displays the college premium and relative (population-based) supply of college workers, each residualized of the real minimum wage and a linear time trend. Panel (b) displays the college premium and real minimum wage, each residualized of the relative (population-based) supply of college workers and a linear time trend. The variation in Panel (a) identifies ν_L and the variation in Panel (b) identifies ν_m when estimating the reduced-form specification of equation (2), in which the relative supply of college workers is measured using relative populations. There is a clear negative relationship in each panel.

Column (1) of Table I displays results of estimating the canonical model of regression

	1963-1987	1963-2016	1963-1987	1963-2016			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Relative supply of college workers	-0.781	-0.411	-0.687	-0.632	-0.703	-0.608	-0.541
	(0.186)	(0.064)	(0.137)	(0.069)	(0.077)	(0.104)	(0.067)
Real minimum wage			-0.311	-0.220	-0.199	-0.133	
			(0.064)	(0.048)	(0.059)	(0.052)	
Real federal minimum wage							-0.161
							(0.044)
Time	0.027	0.016	0.021	0.021			0.019
	(0.006)	(0.001)	(0.005)	(0.002)			(0.002)
Constant	-0.004	0.007	0.017	0.009	-0.000	0.033	0.010
	(0.021)	(0.008)	(0.018)	(0.007)	(0.014)	(0.027)	(0.007)
Time Polynomial	1	1	1	1	2	3	1
Observations	25	54	25	54	54	54	54
R-squared	0.271	0.936	0.636	0.957	0.957	0.966	0.951

Table I: Regression Models for the National College Wage Premium

Notes: Results of estimating (1) in columns (1) and (2) and of estimating (2) in columns (3) – (7). All regressions are estimated using 2SLS, instrumenting for the "Relative supply of college workers" measured as the log relative supply of college-to-non-college hours worked using the population-based measure. The sample is 1963 – 1987 in columns (1) and (3) and 1963 – 2016 elsewhere. "Real minimum wage" and "Real federal minimum wage" are the logs of the real minimum wage (which averages across states) and real federal FLSA minimum wage. "Time Polynomial" refers to the degree of the polynomial of time; the coefficient on the linear trend on "Time" is omitted from the table whenever this polynomial is of degree 2 or greater.

(1)—which omits the real minimum wage—using 2SLS on the sample 1963 – 1987; I report robust standard errors in all regressions.[6] This is a similar specification and the same sample years included in the seminal work of Katz and Murphy (1992). Column (2) replicates this analysis, but using the full sample of 1963 – 2016. Consistent with past work, estimates are unstable across samples. The growth rate of skill-biased technical change γ falls from 2.7% per year to 1.6%; and the coefficient on relative supply β_L rises from -0.781 to -0.411 when changing the sample from the Katz and Murphy years to the full sample. Another way to state this result is that the canonical model estimated on the 1963 – 1987 sample systematically deviates from the data thereafter, predicting a sharper rise in the college premium than actually occurs. This result is well known—see, e.g., Acemoglu and Autor (2011)—and shown in Figure III. This deviation between the observed growth of the college premium and the model's prediction suggests a rapid slowdown in the rate of skill-biased technical change in the later period.

The extended canonical model goes some way towards fixing this well-known issue

[6] The first-stage is strong across all specifications in Table I, as suggested by panel (b) of Figure I. The first-stage KP F statistic is above than 70 across all columns.

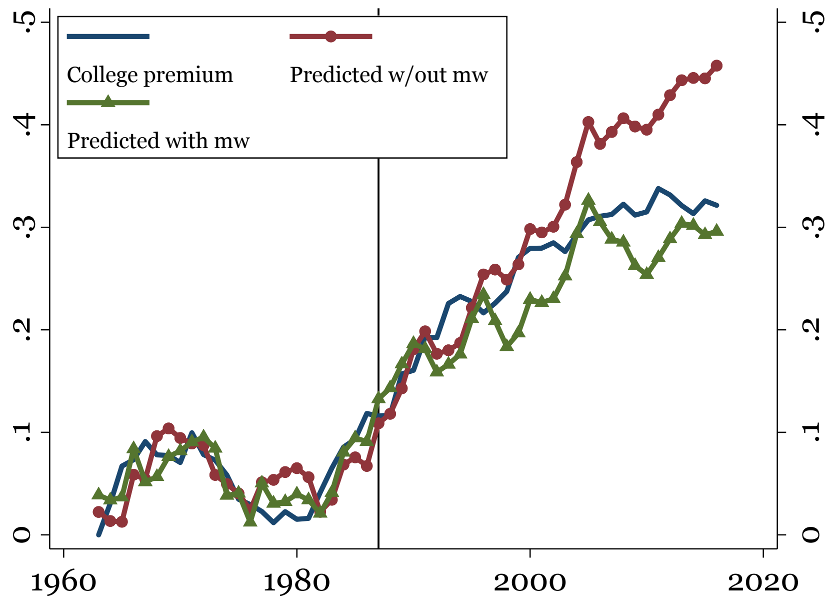

Figure III: Out-of-Sample Predictions of the Canonical Model and Extended Canonical Model at the National Level

Notes: The observed evolution of the national college premium, in solid blue, and its predicted evolution from estimating equation (1), in red circles, and equation (2), in green triangles, using 2SLS on the sample of years 1963 – 1987, instrumenting for the relative supply of college workers measured as the log relative supply of college-to-non-college hours worked using the population-based measure. The vertical line at 1987 indicates the final sample year in the estimation; all predictions thereafter are out of sample.

with the canonical model. Including the real minimum wage and estimating equation (2) leads to a substantial improvement in the model's out-of-sample fit, as shown in Figure III, reducing the implied slowdown in the rate of skill-biased technical change needed to match the observed college premium. This results from the fact that the real minimum wage is high in the 2000s—especially the baseline series that averages across states, as shown in panel (c) of Figure I—precisely when the growth rate of the skill premium is low. The extended canonical model also leads to more stable parameter estimates across samples than the canonical model. Columns (3) and (4) in Table I replicate columns (1) and (2), but estimating equation (2) instead of equation (1). The growth rate of skill-biased technical change γ is 2.1% when estimated on the 1963 – 1987 sample and on the full sample while the coefficient on relative supply β_L is -0.687 and -0.632 when estimated on the 1963 – 1987 sample and on the full sample.

The remaining columns (5), (6), and (7) are estimated on the full sample of years. Columns (5) and (6) include higher-dimensional polynomials of time. I include these columns to check if allowing for a more flexible evolution of the rate of skill-biased technical change dramatically affects estimated elasticities. Whereas the quantitative estimates do vary with the polynomial of time, the broad conclusions are robust.[7] Finally, column (7) replicates column (4), but replacing the real minimum wage (which averages minimum wages across states) with the real federal minimum wage. Since the federal minimum wage is a noisy measure of the minimum wage workers actually face, it is perhaps not surprising that the column (7) estimate is attenuated relative to the corresponding estimate in column (4).

The estimates in Table I highlight the importance of supply (education), demand (the skill bias of technology), and the real minimum wage in shaping the evolution of the U.S. college premium. The elasticity of the college premium with respect to the real minimum wage over the full sample ranges between -0.13 and -0.22, settling down to about -0.13 with cubic, quartic, and quintic polynomials of time, as shown in Table A.III in Online Appendix A.1. The -0.13 elasticity implies that the 26% decline in the real minimum wage between 1979 and 1989 caused a 3.4% increase in the college premium over this time period (see Figure A.I in Online Appendix A.1), which is one quarter of the observed 13.4% increase. The incorporation of the real minimum wage does not change the conclusion that changes in supply and skill-biased technical change remain dominant forces driving the evolution of the U.S. college premium.

Sensitivity. Here, I briefly describe results of three types of sensitivity exercises. Details and results are provided in Online Appendix A.1.[8]

First, I check if the results in Table I are robust to using alternative measures of the relative supply of college workers. The baseline in Table I instruments for the hours-worked measure of relative supply with the population-based measure. In Table A.I I estimate the reduced-form specification of equation (2) using OLS and measuring the relative supply of college workers with the population-based measure that serves as the instrument in my baseline. In Table A.II I estimate equation (2) using 2SLS, but replacing the hours-worked measure of the relative supply of college workers with the dual of wages, choosing supply of college and non-college workers such that the product of wages and supplies exactly matches wage income for college and non-college workers in each year. In both cases, results are consistent with the baseline: *(i)* the coefficient on the minimum wage is negative and significantly different from zero in all specifications and *(ii)* the estimates in the 1963 – 1987 sample are more similar to those in the full sample in the extended canonical model than in the canonical model.

[7] Reassuringly, my baseline empirical results in Section 5 are robust to allowing for arbitrary time-varying national technical change at the level of each of the 100 groups of workers across which I composition adjust in this section.

[8] Because I composition adjust wages and supplies, the baseline analysis controls for (observable) changes in the composition of the workforce.

Second, I aim to understand why my conclusions differ, as described in the Introduction, from Autor, Katz and Kearney (2008) (henceforth, AKK). AKK estimate a significant negative coefficient on their measure of the real minimum wage using a linear time trend and an insignificant negative coefficient using a third-degree polynomial of time. My results differ in part because I display results across a broader set of specifications, in part because I have access to a longer time horizon, and in part because I measure the real minimum wage using the average of minimum wages across states instead of the federal minimum wage.

Because AKK's conclusions on the impact of the real minimum wage on the college premium depend on the polynomial of time included in the regression, I begin by estimating the extended canonical model of equation (2) using my data, but incorporating higher-dimensional polynomials of time, up to and including a fifth-degree polynomial. Table A.III shows that the coefficient on the real minimum wage remains significantly different from zero across all specifications and stabilizes at -0.13 across a third-, fourth-, and fifth-degree polynomial of time. Next, in the top panel of Table A.IV I use data from AKK's replication package to show that outside of the specification with a cubic time trend, results that can be obtained using their data are similar to my results. I then show, in the bottom panel of Table A.IV, that the coefficient on the real minimum wage tends to be larger when I replace their federal minimum wage with my baseline measure of the minimum wage, similar to the comparison between columns (4) and (7) in my Table I. As there, using the federal minimum wage introduces bias, since it doesn't apply in many states, especially starting in the late 1990s.

To this point, I have studied the elasticity of the college premium—rather than the real wage of either education group—with respect to measures of supply, demand, and the real minimum wage, in keeping with the literature focusing on supply and demand. In Table A.V I decompose the estimated elasticities of the college premium into separate elasticities of college and non-college real wages. To do so, I estimate equation (2) replacing $\log(w_{ht}/w_{\ell t})$ with $\log w_{ht}$ and, separately, with $\log w_{\ell t}$. Each resulting elasticity of the college real wage minus the corresponding elasticity of the non-college real wage equals the elasticity of the college premium. Across each specification (linear, quadratic, and cubic), a higher real minimum wage raises the average real wage of workers without a college education, $w_{\ell t}$; and this effect is statistically significant in each specification. The impact of the real minimum wage on the average real wage of college-educated workers, w_{ht}, is not robust across specifications, varying from negative and statistically significant in the linear specification to positive and statistically insignificant in both the quadratic and cubic specifications. I conclude that increases in the real minimum wage robustly raise non-college-educated workers' real wages and have ambiguous effects on college-educated workers' real wages.[9]

Summary. Supply, demand, and the real minimum wage each appear to help shape the evolution of the U.S. college premium. Moreover, the extended canonical model (2) improves on the out-of-sample fit of the canonical model (1), implying a smaller slowdown in the rate of skill-biased technical change in the U.S. in the 2000s.

This motivating evidence on the relevance of real minimum wages for shaping the college premium leaves many questions unanswered. How can the minimum wage have such sizable effects when so few workers earn the minimum wage? If large spillover effects of changes in minimum wages—to workers not earning the minimum wage—are required to generate these elasticity estimates, how is this consistent with past empirical evidence showing stronger direct effects than wage-spillover effects? Finally, is the national time-series empirical strategy sufficiently compelling? I now turn to a model that—together with additional empirical exercises motivated and interpreted by it—helps address these and other questions.

3 Theory

In Section 3.1 I describe a simple job-ladder model with myopic workers and firms that nests the canonical model. In Section 3.2 I characterize the steady state of this model and provide comparative static results on changes in the minimum wage and labor supply and demand across steady states. In Section 3.3 I describe the transition to a one-time change in the economic environment, focusing on a change in the real minimum wage.

3.1 Setup

Time is discrete and indexed by t. The economy features S labor skills and types of firm, each indexed by s, with a type s firm hiring only skill s labor to produce type s output. The exogenous mass of skill s workers is L_{st}. There are many firms of each type, and output across type s firms is perfectly substitutable. Each employed skill s worker produces A_{st} units of type s output, where A_{st} is time-varying factor-specific productivity. Final output, Y_t, is produced by competitive producers. Final output is a constant returns to scale function of the output, Y_{st}, of each type s, $Y_t = Y\left(\{Y_{st}\}_{s=1}^{S}\right)$. The price of the

[9]For non-college-educated workers, these empirical results are consistent with my model's predictions in Section 3. For college-educated workers, the empirical results in the quadratic and cubic specifications are also consistent with my theory, although in the linear specification they are not.

final good is the numeraire and P_{st} is the endogenous (real) value marginal product of employed type s labor (the price of A_{st} units of type s output).

A worker can be employed or unemployed. Each period, a skill s worker-firm match faces an exogenous separation probability $\delta_s \in [0, 1]$, an unemployed skill s worker matches with a firm with probability $\lambda_{su} \in [0, 1]$, and an employed skill s worker meets a new firm with probability $\lambda_{se} \in [0, 1]$. An unemployed skill s worker receives real income $v_s < P_s$. If a worker meets a new firm, the worker and firm engage in generalized (asymmetric) Nash bargaining to determine if the worker moves from her current employer to the new firm and, if so, the worker's wage, which is fixed throughout the life of the contract (subject to the minimum wage, as described below). The worker's bargaining power is $\beta_s \in (0, 1]$. If an unemployed worker meets a firm, her outside option is her unemployment income v_s. If a worker employed at wage w meets a new firm, her outside option is to continue employment in her existing match at wage w.[10]

Whereas wages are fixed and not renegotiable, they are also subject to a minimum wage, m_t. If the minimum wage rises above an existing wage contract, the wage in that match rises to the minimum wage if the firm and worker find it optimal to maintain the match at this new wage; otherwise, the firm can endogenously fire the worker and the worker can quit to unemployment. I assume that the minimum wage satisfies $m_t < P_{st}$ for each s and t. This upper bound on m_t implies that each worker-firm match generates positive surplus, so firms always find it strictly profitable to hire workers at the minimum wage. This implies that changes in the minimum wage do not affect unemployment. I relax this assumption in Section 3.4.

Within period t, I assume exogenous separation shocks occur first and then new matches are realized for those workers who did not separate in t; bargaining occurs subject to the minimum wage that prevailed in the previous period. Finally, after matching and bargaining, the new minimum wage is set and employment and worker wages adjust accordingly.[11] I focus on the case in which the discount factor is zero, so that each worker maximizes her current wage and each firm maximizes its current profit.[12] This assump-

[10]I treat the matching probabilities as exogenous, as in, e.g., Postel-Vinay and Robin (2002) and Cahuc, Postel-Vinay and Robin (2006); this rules out the possibility that an increase in the minimum wage (which reduces employer profit) might reduce matching probabilities disproportionately more for worker skills that are more likely to earn the minimum wage. I rule out the possibility that a worker can exploit a new job offer to raise her wage with her current employer, consistent with counteroffers being uncommon empirically (see, e.g., Mortensen, 2003). Finally, unlike Shimer (2006), a worker's outside option is her current wage rather than unemployment; wages are, therefore, not renegotiation proof.

[11]I choose these timing assumptions so that the dynamic implications of the discrete-time model are similar to those in a continuous-time model.

[12]The model simplifies to the canonical model if $\beta_s = 1$ and $\delta_s = 0$ (in which case worker wages equal their value marginal product of labor and there is no unemployment), there are exactly two skill groups,

tion facilitates the analysis of the transition to aggregate shocks. I derive related steady-state results in a canonical (forward-looking) wage-posting model in Online Appendix B.3.

3.2 Steady state

In steady state, the exogenous time-varying parameters L_{st}, A_{st}, and m_t are all time invariant, as is the endogenous value marginal product of type s labor, P_{st}, and the density of real incomes across workers within each skill s, which I denote by $g_{st}(\cdot)$. Hence, in what follows, I omit time subscripts and re-introduce them in Section 3.3. I first solve for the steady-state distribution of wages across workers within skill s for a given P_s (and, therefore, for any aggregate production function and any number of skills). I then conduct comparative statics across steady states, again for given changes in P_s. Finally, I show how to solve for changes in P_s given any aggregate production function and derive the extended canonical model estimating equation (2) as an example. Online Appendix B.1 contains details and derivations.

Wages along the job ladder. Suppose that an unemployed skill s worker matches with a firm. If $m \geq (1-\beta_s)\, v_s + \beta_s P_s$, then the bargaining outcome is constrained by the minimum wage and the worker's wage is set to m; in this case, I refer to skill s as being "bound" by the minimum wage. Otherwise, the bargaining outcome is unconstrained and the wage is set to $(1-\beta_s)\, v_s + \beta_s P_s$. Each time an employed worker at wage w matches with a new firm, her wage rises to $(1-\beta_s)\, w + \beta_s P_s$.

This process implies the existence of a steady-state wage ladder $\{w_{j,s}\}_{j=1}^{\infty}$ for each skill s. The unemployment benefit is what the worker receives when she is below the job ladder, $w_{0,s} = v_s$. The wage on the first "rung" of the job ladder is $w_{1,s} = m$ if s is bound by the minimum wage and is $w_{1,s} = (1-\beta_s)\, v_s + \beta_s P_s$ otherwise. The wage on each successive rung $j+1$ on the job ladder is simply $w_{j+1,s} = (1-\beta_s)\, w_{j,s} + \beta_s P_s$ for $j \geq 1$. Consequently, wages rise as a worker moves up rungs of the job ladder, $w_{j+1,s} - w_{j,s} = \beta_s \left(P_s - w_{j,s}\right) > 0$, but proportionally less at higher rungs, $d\left(w_{j+1,s}/w_{j,s}\right) \big/ dw_{j,s} < 0$. In Online Appendix B.1 I solve explicitly for all wages on the job ladder, both in the case in which skill s is and is not bound by the minimum wage.

Distribution of workers across the job ladder and the average wage. In steady state, the rate at which workers exit a rung of the job ladder equals the rate at which they enter it. The probability that a skill s worker at any wage $w_{j+1,s}$ in period t does not work at this wage in period $t+1$ is $\delta_s + (1-\delta_s)\gamma_{se}$, where δ_s is the probability the worker exogenously

and the aggregate production function is CES.

separates from her firm and, if she does not, γ_{se} is the probability that she matches with a new firm. The probability that a skill s worker begins earning wage $w_{j+1,s}$ at date $t+1$ is $(1-\delta_s)\gamma_{se}g_s(w_{j,s})$ if $j>0$ and is $\gamma_{su}g_s(w_{j,s})$ if $j=0$, which is the probability that she was one rung below $j+1$ times the probability that she does not separate (if $j>0$) and the probability that she matches. Hence, the steady-state density across rungs can be defined recursively as

$$\begin{aligned}[\delta_s+(1-\delta_s)\gamma_{se}]\, g_s(w_{1,s}) &= \gamma_{su}u_s \\ [\delta_s+(1-\delta_s)\gamma_{se}]\, g_s(w_{j+1,s}) &= (1-\delta_s)\gamma_{se}g_s(w_{j,s}) \quad \text{for} \quad j\geq 1\end{aligned} \tag{3}$$

where $u_s = g_s(w_{0,s})$ is the unemployment rate. Using this recursive system I solve explicitly for the steady-state distribution of workers across the job ladder in Online Appendix B.1. As is apparent from the system of equations in (3), the steady-state distribution of workers across rungs j is invariant to whether or not the real minimum wage binds and its level.[13]

The average real wage among employed skill s workers is $w_s \equiv \frac{1}{1-u_s}\sum_{j\geq 1} w_{j,s}g_s(w_{j,s})$. Using the explicit solutions for the distribution of workers across the job ladder $g_s(\cdot)$—which is independent of whether or not the minimum wage binds—and wages along the job ladder $w_{j,s}$—which depends on whether or not the minimum wage binds—I solve for the average wage in Online Appendix B.1. If skill s is bound by the minimum wage, then the average wage is given by a weighted average of the minimum wage and the value marginal product of labor,

$$w_s = \frac{\delta_s}{\delta_s+\beta_s(1-\delta_s)\gamma_{se}}m + \left(1-\frac{\delta_s}{\delta_s+\beta_s(1-\delta_s)\gamma_{se}}\right)P_s \tag{4}$$

If, on the other hand, skill s is not bound by the minimum wage, then w_s is a weighted average of v_s and P_s.

Comparative statics across steady states for given changes in P_s. For any skill s that is not bound by the minimum wage, the partial elasticity (i.e., for a given change in P_s) of the average real wage with respect to the real minimum wage is zero. On the other hand, for a skill s that is bound by the minimum wage, the total derivative of w_s is given by

$$d\log w_s = M_s b_s \partial \log m + (1-M_s b_s)\,\partial \log P_s \tag{5}$$

[13] The result that the minimum wage does not affect the distribution of workers across rungs depends both on the assumption that $m_t < P_{st}$, which implies that the minimum wage does not affect the unemployment rate, and on the assumption that matching probabilities are exogenous.

The partial elasticity of the average wage with respect to the minimum wage, $M_s b_s$, is the product of two terms, the "magnification elasticity," denoted by

$$M_s \equiv \frac{\delta_s + (1-\delta_s)\gamma_{se}}{\delta_s + \beta_s(1-\delta_s)\gamma_{se}} > 1 \tag{6}$$

and the "bite" of the minimum wage, $b_s \equiv mg_s(m) \, / [(1-u_s)w_s]$, which is the share of all skill s wage income earned by workers earning the minimum wage.

Why do I refer to M_s as the magnification elasticity? Consider the impact of a change in m on the average wage of an arbitrary group of workers under the assumption that the minimum wage has no employment effects (an assumption I am imposing here) and under an additional assumption that an increase in the minimum wage only raises the wages of individual workers whose wages are bound by it (leaving all worker wages above the minimum wage unaffected). Under these assumptions, the elasticity of the average wage is exactly equal to its bite. This channel is what is often referred to as the "direct effect" of minimum wages in the literature. In this job-ladder model, the direct effect is active; however, at least in the long run (i.e., across steady states), there are wage spillovers up the distribution above the minimum wage.

To understand these wage spillovers, let $W_s(c)$ denote the steady-state wage at percentile c of employed skill s workers; and consider a skill s that is bound by the minimum wage. In response to an increase in the minimum wage, $W_s(c)$ rises disproportionately more for lower percentiles,

$$\frac{d\left[W_s(c)/W_s(c')\right]}{dm} > 0 \quad \text{for all} \quad W_s(c) < W_s(c')$$

The intuition for this result and its proof are straightforward. The steady-state share of workers on each rung of the job ladder is invariant to the value of the minimum wage. Moreover, in response to an increase in the minimum wage, wages at lower rungs (and, therefore, at lower centiles of the wage distribution) increase disproportionately more.

Finally, in equation (5) the elasticity of the average wage with respect to the value marginal product of labor, P_s, is less than one; this result applies whether or not skill s is bound by the minimum wage. The reason is that the average wage is a weighted average of m and P_s if s is bound by the minimum wage and is a weighted average of v_s and P_s otherwise. Hence, the average wage responds less than one-for-one to changes in P_s.

Determining changes in P_s and micro-founding the extended canonical model. For the steady-state analysis, it remains only to determine how changes across steady states in labor supply, labor demand, and the minimum wage affect value marginal products of labor, $d \log P_s$. Under the prevailing assumption that $m < P_s$ for all s, the share of workers who are unemployed is fixed across steady states. This implies that changes in value marginal products $d \log P_s$ depend only on changes in factor supply and demand. Hence, $M_s b_s$ in equation (5) is not only the partial elasticity of the average wage with respect to the real minimum wage, it is also the total elasticity.

This also implies that for any aggregate production function, steady-state value marginal products are determined exactly as in a competitive static model featuring the same aggregate production function, but with the level of supply in the present model, L_s, replaced with $(1 - u_s)L_s$. Endogenous value marginal products have been considered theoretically and empirically in a wide class of aggregate production functions, from the canonical model, which features a CES aggregate production function combining only two skill groups, to nested CES models featuring many skill groups as in, e.g., Card and Lemieux (2001), and beyond. Of course, it should be noted that while the determination of steady-state values of P_s are not affected by the job-ladder model, the relationship between the average wage and P_s is, as shown in equation (4).

Given my objective of studying the college premium and extending the estimating equation of the canonical model from equation (1) to (2), in the following proposition I show that imposing the same restrictions on the aggregate production function as in the canonical model yields equation (2); see Online Appendix B.1 for the formal derivation.

Proposition 1. *Suppose the aggregate production function is a CES combination of the output of two skills, h and ℓ, and parameters satisfy $\delta_h = \delta_\ell$, $\gamma_{he} = \gamma_{\ell e}$, and $\beta_h = \beta_\ell$. Then, to a first-order approximation around an initial steady state, the skill premium, $w_{ht}/w_{\ell t}$, is given by equation (2) under the assumption of linear growth in A_{ht} and $A_{\ell t}$ across time t, where $\nu_m = M(b_h - b_\ell)$ and b_h and b_ℓ are the minimum wage bites in the initial steady state.*

The assumption of common parameters in Proposition 1 is a sufficient condition for common magnification elasticities across skills, $M = M_s$. If the minimum wage only binds for lower-skill workers (for example), then this condition is not necessary to obtain the result, since in this case only the magnification elasticity for ℓ is relevant and $\nu_m = -M_\ell b_\ell$. Of course, under the additional assumption that parameters are common across h and ℓ, this is just a special case of the result in Proposition 1, as $M(b_h - b_\ell) = -Mb_\ell$.

3.3 Transition dynamics

Consider an economy in steady state entering date 0 that experiences a small, one-time increase in the real minimum wage from m to $m' > m$ during date 0. For a skill s that is not bound by the minimum wage at m', the change in real minimum wage has no effect.

The implications for a skill s that is bound by the minimum wage are summarized by the following proposition.[14]

Proposition 2. *Consider an economy in steady state entering into date* 0 *that faces a small one-time change during date* 0 *from* m *to* $m' > m$ *satisfying* $m' < (1-\beta_s)m + \beta_s P_s$ *for all* s*. For any skill* s *that is bound by* m*:*

i. *For all* $t \geq 0$*, two job ladders coexist, with first rungs* m *(old) and* m' *(new)*
ii. *At each rung* j*, the wage on the new job ladder is greater than the wage on the old ladder*
iii. *The share on each rung* j *summed across the two ladders is constant across* t
iv. *The share on rung* j *of the new ladder weakly increases in* t
v. *On impact, the elasticity of the average wage with respect to* m *equals* b_s
vi. *The elasticity of the average wage with respect to* m *rises in* t*, limiting to* $M_s b_s$

I prove Proposition 2 in Online Appendix B.2 and describe its intuition here. An employed worker's wage upon forming a new match depends only on her current wage, bargaining power, and value marginal product of labor; it does not depend on the new minimum wage unless her current wage equals the new minimum wage. Hence, any worker who was on the old job ladder and whose wage was not constrained by the new minimum wage remains on the old ladder until she enters unemployment; whereas any worker whose wage was constrained by the new minimum wage or who exited unemployment after its imposition is on the new ladder. This implies that each worker's wage in all future dates will be a wage that is either on the old job ladder or the new one, as stated in Result (i). Result (ii) follows from the fact that wages on steady-state job ladders are increasing functions of the corresponding minimum wage (for any skill s that is bound by the minimum wage).

On impact, the wage on the first rung on the old job ladder, m, becomes constrained by the new, higher minimum wage. Because surplus in each match remains positive, every worker earning m moves instantly to the first rung of the new ladder, m'. The increase in the minimum wage does not constrain the second rung of the old wage ladder or affect unemployment given the assumption that $m' < (1-\beta_s)m + \beta_s P_s$ for all s, so no other workers are affected on impact. This implies that the instantaneous, or impact elasticity of the average wage of skill s with respect to the minimum wage equals b_s, as stated in Result (v). It also implies that, at least on impact, the share of workers on each rung j summed across the two ladders equals the initial steady-state share. Result (iii) then follows from a proof by induction: if Result (iii) holds for a particular date $t \geq 0$, then

[14] If the real minimum wage experiences a small decrease (rather than increase), then Results (i), (iii), (iv), and (vi) of Proposition 2 are unchanged; Result (ii) is reversed, with the wage on rung j of the new job ladder below the wage on the old ladder; and, finally, the impact elasticity is zero in Result (v).

it must hold in the subsequent date $t+1$. Result (iv) is more complicated, but the basic intuition is simple: workers are exiting the old job ladder and entering the new one over time as they exogenously separate from their employers.[15]

Finally, Result (vi) follows from the previous results: wages on the new job ladder are higher than on the old one (Result ii) and workers are moving from the old to the new ladder over time (Result iv) at a fixed composition of workers across rungs (Result iii). Intuitively, the magnification elasticity grows over time as workers who start at the new minimum wage (either because their wage increased from m to m' on impact or because they exited unemployment after the minimum wage increase) slowly rise up to higher rungs of the new job ladder, each of which is higher than the corresponding rung on the old ladder. In the long run, the elasticity of the average wage with respect to the minimum wage limits to $M_s b_s$, as shown in equation (5), where $M_s > 1$ as shown in equation (6). These results apply for any number of skills and any aggregate production function.

Results (v) and (vi) of Proposition 2 imply that the elasticity of the average real wage of skill s between dates 0 and $T \geq 1$ (the impulse response) with respect to a one-time increase in the real minimum wage during date 1, can be expressed as

$$\log\left(\frac{w_{s,T}}{w_{s,0}}\right) \Big/ d\log m \equiv M_{s,T-1} \times b_{s,0} \tag{7}$$

where $b_{s,0}$ is the initial minimum wage bite (in period 0) before the minimum wage increase and where $M_{s,T-1}$ is the $T-1$-period magnification elasticity. The magnification elasticity is one on impact $M_{s,0} = 1$, is increasing over time $dM_{s,T}/dT > 0$, and limits to the steady-state value $\lim_{T\to\infty} M_{s,T} = M_s$ from equation (6). Equation (7) applies whether or not the minimum wage binds for s; but, of course, if it does not bind then the elasticity in equation (7) is zero, since $b_{s,0} = 0$.

These results also have direct implications for the time-varying elasticity of the skill premium. Impose the additional assumption in Proposition 1 that $\delta_h = \delta_\ell$, $\gamma_{he} = \gamma_{\ell e}$, and $\beta_h = \beta_\ell$, so that $M_T = M_{s,T}$ for each s. Then equation (7) directly implies

$$\log\left(\frac{w_{h,T}/w_{\ell,T}}{w_{h,0}/w_{\ell,0}}\right) \Big/ d\log m \equiv M_{T-1} \times (b_{h,0} - b_{\ell,0}) \tag{8}$$

for any $T \geq 1$ in response to a one-time increase in the real minimum wage during date 1.

[15] Result (iv) is more complicated because the growth of any rung j on the new ladder depends on the distribution of workers across rungs of the new ladder rather than the share of workers on the new ladder overall. The share of workers on rung j (of the new job ladder) in the previous period determines the number of workers leaving rung j in the present period and the share on rung $j-1$ in the previous period determines the number of workers entering rung j in the present period.

On impact, the elasticity of the skill premium with respect to the minimum wage equals the bite of the minimum wage for higher- minus for lower-skill workers. Over time, this elasticity grows. Equation (8) will help me interpret the estimated elasticity of the college premium with respect to the real minimum wage displayed in Table I.

3.4 The minimum wage and unemployment

A prevailing assumption throughout the theory presented above is that each worker-firm match generates positive surplus: $m_t < P_{st}$ for all t and s. This implies that changes in the minimum wage do not affect unemployment. Throughout this section, suppose that $m_t = P_{st}$ for s but not for any $s' \neq s$; and consider all but the knife-edge case in which unemployment is exactly equal to its level when $m_t < P_{st}$.

In this case, the wage ladder collapses for s (and only for s), with all employed s workers earning a real wage equal to m_t. This implies that the elasticity of w_{st} (and of P_{st}) with respect to m_t is then exactly one for skill s. Moreover, changes in the real minimum wage affect unemployment for skill s (and only for skill s). The elasticity of skill s employment with respect to m_t is determined by the elasticity of the residual labor demand curve for skill s induced by the aggregate production function at fixed employment for all other skills. That is, an increase in the real minimum wage raises the wage of s one-for-one and causes a decline in employment of skill s as the economy traces its residual demand curve for s. This is identical to the wage and employment impacts of a change in a binding minimum wage in a competitive labor market.

What is the impact of a change in the real minimum wage on the average wage of other skills? For all $s' \neq s$, equations (4), (5), and (6) continue to hold. However, since changes in m_t affect employment of skill s, they affect the value marginal product of skill s' and, therefore, their average wage $w_{s'}$. Specifically, whereas $M_{s'}b_{s'}$ remains the partial long-run elasticity of the real wage of s' with respect to m, it is no longer the total elasticity, as m also impacts $P_{s'}$. I return to these issues in Section 5.3.

4 Empirical approach

4.1 From theory to estimation

According to equation (7), if there is a small, one-time increase in the real minimum wage (starting from a steady state) that occurs during $t - T + 1$ or any period thereafter, then the change in the average real wage of skill s between periods $t - T$ and t is given by

$$\log \frac{w_{s,t}}{w_{s,t-T}} = \sum_{j=0}^{T-1} M_{s,j} b_{s,t-j-1} \log \frac{m_{t-j}}{m_{t-j-1}}$$

where, by assumption, the log change in the minimum wage is identically zero in all but one of the periods (indexed by j) in the above summation. The previous expression is a distributed lag model with lag weights that exactly equal the model's skill-specific and time-varying magnification elasticities, $M_{s,j}$. If the single minimum wage change occurs earlier than during period $t - T + 1$, then the above expression must be adjusted. If the small one-time increase in the real minimum wage occurs during period $t - T' + 1$ or any period thereafter, with $T' > T$, then the change in the average real wage of skill s between $t - T$ and t is given by

$$\log \frac{w_{s,t}}{w_{s,t-T}} = \sum_{j=0}^{T-1} M_{s,j} b_{s,t-j-1} \log \frac{m_{t-j}}{m_{t-j-1}} + \sum_{j=T}^{T'-1} (M_{s,j} - M_{s,j-T}) b_{s,t-j-1} \log \frac{m_{t-j}}{m_{t-j-1}} \tag{9}$$

This is, again, a distributed lag model. The one change from the previous expression is that the lag weights equal the magnification elasticities only for lagged shocks occurring during or after $t - T + 1$; for further lagged shocks, the lag weights are strictly less than the corresponding magnification elasticities. This is intuitive. Instead of studying the entire impulse response to a one-time change in the real minimum wage, for shocks that are lagged sufficiently the wage change between $t - T$ and t only captures the tail-end of the response, cutting off the initial periods after the shock.

Equation (9) holds exactly in the model in the presence of a one-time increase in the real minimum wage. However, in moving from the model to the data, it is important to remember that reality differs from the model in many respects. Perhaps most importantly, the real minimum wage does not experience a one-time permanent change in any empirical context. Instead, it is always changing, driven by a combination of inflation and minimum wage policy.

4.2 Specification

In my baseline specification, I estimate the following version of equation (9),

$$\log \frac{w_{g,r,t}}{w_{g,r,t-T}} = \sum_{j=-L}^{T'-1} \mu_j b_{g,r,t-j-1} \left(\log \frac{m_{r,t-j}}{m_{r,t-j-1}} - \overline{\Delta \log m} \right) + \kappa_t b_{g,r} + \alpha_t + \varepsilon_{g,r,t} \tag{10}$$

I define time periods t as half years, January – June and July – December of each year between 1979 and 2016. Skill groups s in the model are mapped in the data to the combination of labor groups g, which correspond to the same 100 labor groups across which I composition-adjust wages and supplies in Section 2 (defined in Online Appendix A.2 by the interaction of 5 age bins, 2 genders, 2 races, and 5 education levels) and regions r, which correspond to the 50 U.S. states. Real minimum wages, $m_{r,t}$, vary across regions and time whereas minimum wage bites, $b_{g,r,t}$, vary across labor groups, regions, and time.

Relative to equation (9), there are four changes in my implementation of equation (10). First, I project changes in real wages not only on lagged one-period changes in real minimum wages, but also on future one-period changes (i.e., the summation begins with $j = -L$ rather than $j = 0$); I do this to check for pre-existing trends. Second, I estimate lag weights, μ_j, that are common across groups; in robustness, I estimate separate specifications for different groups.

Third, in estimating equation (10), I instrument for the interaction between the time-varying bite of the minimum wage and the change in the real minimum wage. I do this because there is measurement error in the minimum wage bite; moreover, measurement error in the time-varying bite will be correlated with measurement error in the average wage in either the numerator or the denominator of the dependent variable for particular values of j in the summation. This concern is very similar to that emphasized in a related context by Autor, Manning and Smith (2016). In the instrument, I replace $b_{g,r,t}$ with $b_{g,r}$, where $b_{g,r}$ is the median value of the time-varying bite for labor group g in region r across periods, excluding period t.[16]

Finally, I include two additional controls—a time effect, α_t, and the time-varying impact of the median value of the minimum wage bite, $b_{g,r}$—and I subtract from the one-period change in the real minimum wage in region r between $t - j - 1$ and $t - j$ the average one-period change in the real minimum wage across all states and all time periods between 1979 and 2016 (including 1994 and the first half of 1995), denoted by $\overline{\Delta \log m}$. Treatment in equation (10) is the interaction between shocks to the real minimum wage and the minimum wage bite, which fits into the "shock-exposure" framework of Borusyak and Hull (2023). As they argue, this specification may suffer from omitted variable bias.

In my context, the intuition is simple. In periods in which the real minimum wage tends to fall systematically over time (e.g., the 1980s) the interaction between changes in the real minimum wage and the minimum wage bite will be negatively correlated

[16] Whereas $b_{g,r}$ depends on t because it is the median across time omitting period t, I have omitted t for notational simplicity. In practice, given 73 time periods, using the leave-out version of $b_{g,r}$ or not makes little difference.

with the minimum wage bite. But the minimum wage bite could be correlated with other shocks in the residual. For instance, because the minimum wage bite tends to be higher for g, r pairs that are lower skill and the economy experiences skill-biased technical change (pushing up the wages of more skilled groups), expected treatment would be correlated with the impact of skill-biased technical change on real wage growth. Hence, there would be endogeneity.

Borusyak and Hull (2023) show that this type of omitted variable bias can be avoided by re-centering treatment (differencing realized treatment from expected treatment) or controlling for the expectation of treatment, either of which requires assumptions on the distribution of expected real minimum wage changes (in my context). Equation (10) re-centers treatment under the assumption that all one-period changes in real minimum wages are independent draws from a common distribution; it does so by subtracting from the realized one-period change in the real minimum wage in each region and time period its average value taken across all regions and time periods, $\overline{\Delta \log m}$. This addresses the endogeneity concern in my above example. The interaction between $b_{g,r}$ and a time effect in equation (10) additionally controls for any national time-varying shocks that are correlated with the median value of the minimum wage bite for labor group g in region r.

This approach provides additional benefits in my context. In the reduced-form specification, the interaction between (a non-leave-out version of) $b_{g,r}$ and a time effect absorbs changes in inflation, yielding identical results to an alternative specification replacing real minimum wages with nominal values. Moreover, this "design-based" approach additionally avoids negative *ex-ante* weights; see, e.g., Borusyak and Hull (2024).[17]

In robustness, I additionally control for the interaction between $b_{g,r}$ and state effects, which allows for expected changes in minimum wages to differ arbitrarily across states in a time-invariant way (since, in practice, different states are systematically more or less likely to set minimum wages above the federal level), as well as including additional controls.

I focus on three-and-a-half-year changes in real wages, setting $T = 7$ in equation (10). In addition to considering the impact of one-period lagged real minimum wage changes over this time horizon, I also control for lagged shocks going back farther in time, setting $T' = 13$. And I include three-and-a-half years of one-period lead real minimum wage changes, setting $L = 7$. To allow for correlation across time in the error terms, I clus-

[17]The result on negative weights in Borusyak and Hull (2024) is proven only for regressions with one treatment variable. I conjecture that the result extends to specifications featuring many treatment variables—as in regression (10)—if these treatments are uncorrelated with each other, after residualizing on controls.

ter standard errors by state. Finally, in all regressions, I weigh each observation by the product of (g, r) workers in t and $t - T$ divided by their sum.

4.3 Data and measurement

Here, I provide an overview of the data and measurement. Online Appendix A.2 contains details. To measure m_{rt}, I use the maximum value of the effective nominal minimum wage in region r in the corresponding half-year period, which I obtain from Cengiz et al. (2019). I deflate this by the maximum monthly GDP deflator in the same half-year period. To measure the bite of the minimum wage, I construct the share of wage income earned by workers at a wage no greater than 1.05 times the minimum wage; I vary this cutoff in robustness.

I measure the minimum wage bite, $b_{g,r,t}$, and average wages, $w_{g,r,t}$, using the CPS Merged Outgoing Rotation Groups (MORG). I switch from the March CPS, which I use in the national analysis in Section 2, both because the MORG CPS includes a larger sample, which is particularly useful when dividing the data across 50 states and 100 labor groups, and because individual wages can be measured with less error (see Lemieux, 2006), which is especially important for measuring the bite of the minimum wage. These benefits come at the cost of a shorter time frame, starting with 1979 instead of 1963. I additionally drop all of 1994 and the first eight months of 1995 given missing imputation flags; more generally, I clean the MORG CPS data following an approach similar to Lemieux (2006). After cleaning the data, there are approximately 12 (unweighted) worker observations per g, r, t (with 100 unique values of g, 50 of r, and 73 of t).[18] In my baseline, I use a balanced panel, only including g, r if there are observations in every t in the estimation sample. In robustness I include all g, r. Given outliers, I winsorize average wages at the 2nd percentile within each (g, r) pair.

The final sample year is 2016 because many municipalities begin setting minimum wages above the mandated state level at this point. For instance, New York City, Nassau, Suffolk, and Westchester counties begin setting minimum wages above the New York state level on the very last day of 2016 and Minneapolis begins setting minimum wages above the Minnesota level in 2018. However, some counties were already setting minimum wages above their states' levels beforehand. For instance, San Diego did so throughout 2015 and Los Angeles began doing so in the middle of 2016 (for large businesses). In robustness, I consider ending the analysis after 2010, to minimize the impact

[18]These small cells imply measurement error both in the dependent variable (where classical measurement error does not imply bias) and in the measure of the minimum wage bite in the independent variable (for which I instrument).

of endogeneity from mismeasured minimum wage changes.

5 Results

Section 5.1 displays baseline empirical results and discusses their implications for the college premium and the real wage of workers without a high school degree. Section 5.2 describes a range of sensitivity analyses for the baseline empirical results. Finally, Section 5.3 briefly discusses certain caveats.

5.1 Results and implications

In all that follows, I display results of estimating equation (10) as follows. For the impact magnification elasticity, I display the lag weight estimate μ_0 and its associated 90% confidence interval. This corresponds to the impact of current changes in the real minimum wage on real wages over a six month interval. For all other magnification elasticities, I display the averages of two consecutive lag weight estimates, thereby smoothing estimates across time differences. For $j > 0$, I display the average of the lag weight estimates μ_{j-1} and μ_j and its associated 90% confidence interval. Hence, the reported magnification elasticity at, e.g., two-and-a-half years corresponds to the average impact on real wages over the three-and-a-half-year period of changes in the real minimum wage between *(i)* two to two-and-a-half years ago and *(ii)* two-and-a-half to three years ago. I report magnification elasticities in the figures that correspond only to shocks that occurred within the three-and-a-half-year period of the change in the dependent variable (although I control for past changes). Hence, all figures show magnification elasticities up to three years, because the three-and-a-half year lag weight estimate would not correspond to the magnification elasticity, as it includes the impact of shocks that occurred between three-and-a-half and four years ago (before the time horizon in the dependent variable); see equation (9). Similarly, for any period $j < 0$, I display the average of $-\mu_j$ and $-\mu_{j-1}$; this both smooths estimates and also converts pre-trend estimates to their more traditional format, as discussed in Roth (2024).[19] In all that follows, I refer to these (linear combinations of) estimates as magnification elasticities.

I begin by estimating the reduced-form specification of equation (10), in which I replace $b_{g,r,t-1}$ in each right-hand-side covariate with $b_{g,r}$ and estimate the resulting regres-

[19] To understand this choice, suppose that there is a negative estimate for $j < 0$. This would imply that the real wage declines leading into a future increase in the real minimum wage. Displaying the negative value of the estimated lag weight (a positive point in the figure associated with this value of $j < 0$) then clarifies that the level of the real wage declines leading into the shock.

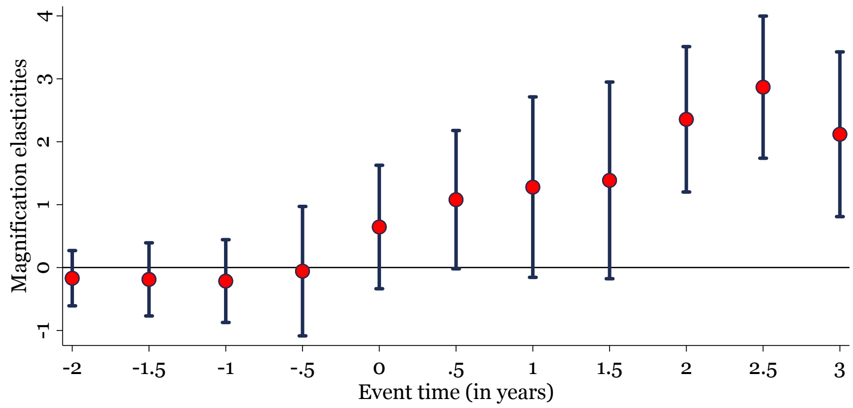

Figure IV: Time-Varying Magnification Elasticities

Notes: Results of estimating regression (10) using 2SLS. Results display μ_0 and the linear combinations of $(\mu_{j-1} + \mu_j)/2$ for $j > 0$ and of $-(\mu_j + \mu_{j-1})/2$ for $j < 0$, with corresponding 90% confidence intervals; robust standard errors are clustered by state.

sion using OLS. Figure A.III in Online Appendix A.2 displays the estimation results. On impact, the reduced-form coefficient is 0.5 (the estimated coefficient at $j = 0$) and rises to a maximum of 1.6. Figure A.III also shows no evidence of pre-existing differential trends.

Figure IV displays the baseline empirical results for the time-varying magnification elasticities, corresponding to the estimates of the lag weights in equation (10) using 2SLS.[20] On impact, the magnification elasticity is approximately 0.6. It rises over time, peaking at 2.9 over the first two to three years (at $j = 2.5$). From this, I conclude that the long-run magnification elasticity is approximately 2.9.

A long-run magnification elasticity of 2.9 implies that the indirect effect of changes in the minimum wage (the effect on wages above the minimum wage) generates about 66% ($\approx 1 - 1/2.9$) of the overall wage effect (at their peak) of minimum wage changes. How do these estimates relate to results in the literature? To my knowledge, this is the first paper to explicitly study the time-varying effects of minimum wage changes on inequality. However, like this paper, Cengiz et al. (2019) quantify the direct and indirect effects of changes in the minimum wage on average wages, focusing within a \$4 (deflated to 2016) range around the minimum wage. To do so, they average wage changes over the five years following a minimum wage increase. They find that approximately 60% of the total wage effect of a change in the minimum wage is caused by the direct effect, with 40% caused by wage spillover effects above the minimum wage. My finding that the indirect effect generates 66% of the overall wage effect of minimum wage changes is a result about the maximum strength of the indirect effect (which occurs between two and three years after the shock itself), whereas their result is about its average effect over time. Taking

[20] The first-stage is very strong, with all SW F stats over 1,000.

a simple average of the magnification elasticities plotted in Figure IV for $j \geq 0$ yields a value of 1.7. Hence, a back-of-the-envelope measure of the share of the overall wage effect (averaged across time) accounted for by the indirect effect is about 40% ($\approx 1 - 1/1.7$), in line with the measure in Cengiz et al. (2019).

While Dustmann et al. (2021) does not focus on the dynamic implications of minimum wage changes on inequality, they do document that German regions more exposed to a national increase in the minimum wage experience wage growth for at least two years relative to less exposed regions. This is perhaps the most directly related empirical result to my own; and it is consistent with my empirical evidence at the state-and-group level.

Implications for the college premium and the relation to the national estimates. According to the theoretical result in equation (8), the long-run elasticity of the college premium with respect to the real minimum wage is determined by the product of the long-run magnification elasticity and the difference between the bite of the minimum wage between college and non-college workers. The average value of the minimum wage bite, $b_{g,r,t}$, taken across states r, time t, and all worker groups g with a college education (weighted by the number of workers in the g, r, t triplet and including only the balanced estimation sample of g, r pairs) is 0.6% and without a college education is 5.0%; see Table A.VI in Online Appendix A.2. Using a long-run magnification elasticity of 2.9, this implies a long-run elasticity of the college premium with respect to the real minimum wage of approximately $-.13$ ($\approx 2.9 \times (.050 - 0.006)$). This is squarely in the range of the national estimates displayed in Tables I and A.III. Hence, the national estimates in Section 2 and the estimates in Figure IV are consistent. Both imply sizable elasticities of the college premium with respect to the real minimum wage.

Implications for real wages of workers without a high school degree. The combination of the theoretical results in Section 3 and the empirical results in Figure IV shed light not only on relative wages, but also on real wage levels.

Real wages of workers without a high school degree (HSD) declined dramatically starting in the 1980s and into the early 1990s; see, e.g., Autor, Katz and Kearney (2008). This is a natural response to declining real minimum wages in the extended canonical model. According to equation (7), the long-run implication of the approximately 26% decline in the real minimum wage between 1979 and 1989 equals 26% times the long-run magnification elasticity (of 2.9) times the bite of the minimum wage among HSD groups. The average value of the minimum wage bite, $b_{g,r,t}$, taken across states r, time t, and all HSD worker groups g (weighted by the number of workers in the g, r, t triplet

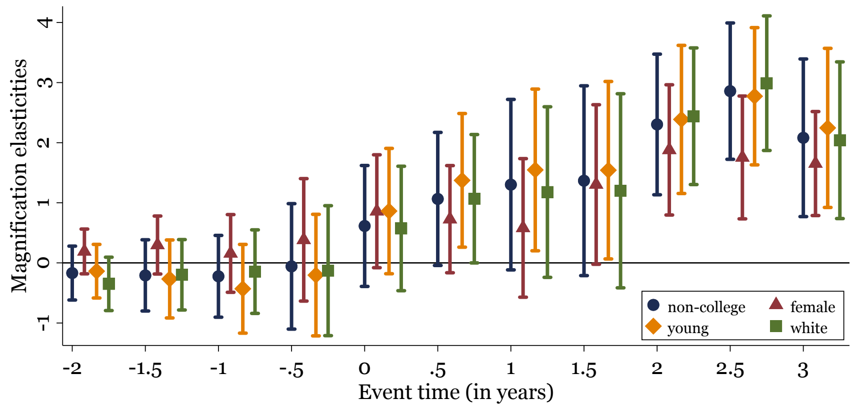

Figure V: Time-Varying Magnification Elasticities: Heterogeneity

Notes: Replicating Figure IV separately on the samples of non-college, female, young (those less than 36 years old), and white workers.

and including only the balanced estimation sample of g, r pairs) is approximately 13.4%; see Table A.VI in Online Appendix A.2. Hence, according to the model, the 26% decline in the real minimum wage between 1979 and 1989 caused an approximately 10.1% ($\approx 0.26 \times 0.134 \times 2.9$) decline in the real wage of workers without a high school degree. This is slightly more than the entirety of the observed decline in their real wages between 1979 and 1992.

5.2 Sensitivity analyses

To what extent are the results presented in Figure IV robust to a variety of alternative choices? I consider sensitivity analyses that fit into three broad categories: *(i)* labor groups included in the empirical analysis, *(ii)* the years included in the empirical analysis, and *(iii)* the particulars of the empirical specification in equation (10). The purpose of *(i)* is to investigate the extent to which magnification elasticities are heterogeneous across worker groups with distinct observable characteristics. The purpose of *(ii)* is to check if results are sensitive to the inclusion of sample years (towards the beginning of the sample) in which much of the variation in real minimum wages is common across states (see Dube et al. (2024) for a discussion of excluding these early years) and to the inclusion of sample years (towards the end of the sample) in which some municipalities apply a minimum wage above the statutory level set by their states. The purpose of *(iii)* is to check if the baseline results are broadly robust either to an alternative way of re-centering treatment—as described in Borusyak and Hull (2023)—or incorporating additional controls.

Estimation on subsamples of worker groups. The theory in Section 3 explicitly incorporates heterogeneity across worker groups in the elasticity of the real wage with respect to the real minimum wage in two respects: it predicts that the elasticity of the change in a group's real wage with respect to the real minimum wage equals the interaction between that group's minimum wage bite (which differs across groups) and the time-varying magnification elasticity (which depends on potentially group-specific parameters). The empirical specification directly incorporates the first dimension of heterogeneity by interacting the change in the real minimum wage with the observable group-specific minimum wage bite. However, in the specification displayed in Figure IV I impose a common lag weight across groups. To what extent do these estimates hide heterogeneous magnification elasticities? To answer this question, I replicate the empirical analysis separately by gender, by education (those without a college education and those with), by race (white and non-white workers), and by age (those younger than 36 and those 36 and older).

Figure V displays results on the subsamples of non-college educated workers, female workers, young workers, and white workers. Each replicates not only the qualitative patterns in the baseline Figure IV—a lack of pre-trends and an increase in estimated lag weights over time—but also broadly replicate the quantitative values of these estimates. Figure A.IV in Online Appendix A.2 replicates these results for the male and the non-white samples. The results for non-white workers are noisier, consistent with much smaller samples inducing more measurement error in the dependent variable whereas the results for male workers are very similar to the baseline except for a substantially larger point estimate at two years out.[21] Finally, results for college-educated workers and for older workers (aged 36 and above) are each insignificantly different from zero for almost all time horizons, consistent with the theory in which the impact of changes in the minimum wage are negligible for labor groups that do not work at the minimum wage; see Figure A.V in Online Appendix A.2.

Sample years. There are two potential concerns related to the years included in the baseline sample. The first concern relates to the inclusion of later years, by which point there are certain municipalities with minimum wages above their states' statutory levels. Because my definition of the real minimum wage is state-and-time specific, this measurement error may bias results. The second concern relates to the inclusion of earlier years, when most of the variation in the real minimum wage is common across states. Whereas this is not an issue according to my theory, Dube et al. (2024) argue for excluding years

[21] I exclude each of these results from Figure V because their inclusion dramatically widens the scale of the y-axis, which makes interpreting estimates more difficult.

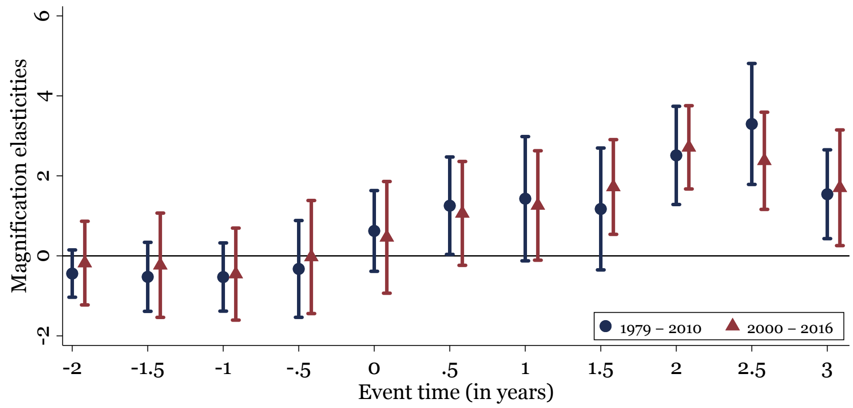

Figure VI: Time-Varying Magnification Elasticities: Sample Years

Notes: Replicating Figure IV separately on sample years 1979 – 2010 and 2000 – 2016.

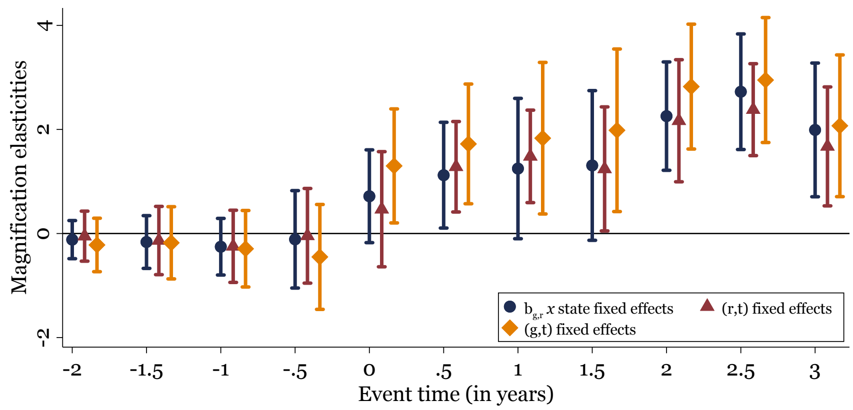

Figure VII: Time-Varying Magnification Elasticities: Sensitivity to Specification

Notes: Replicating Figure IV separately incorporating additional controls: the interaction between $b_{g,r}$ and state effects, (r, t) fixed effects, and (g, t) fixed effects. Also Replicating Figure IV re-centering treatment by demeaning all one-period changes in real minimum wages (in the independent variables and instruments) using the average one period change (averaging across states and time periods).

before 2000. Figure VI replicates Figure IV separately on the sample years of 1979 – 2010 and 2000 – 2016. Results are broadly similar.

Specification in equation (10). In my baseline analysis I control for a time effect and a time-varying effect of the median value of the minimum wage bite for r, g across all t, $b_{r,g}$. This controls for expected treatment under the assumption that there is an arbitrary time-varying national expectation of the change in the real minimum wage. Here I additionally control for the interaction between $b_{g,r}$ and state effects, which—in the shock exposure setting of Borusyak and Hull (2023)—allows for expected changes in minimum wages to differ arbitrarily across states in a time-invariant way.

I alternatively consider two distinct sets of controls in addition to the baseline controls. First, I incorporate r, t fixed effects. This absorbs the direct effect of changes in the real minimum wage. Second, I incorporate g, t fixed effects. This absorbs national, time-varying, and group-specific technical change; it also absorbs any changes in real value marginal products of labor across groups at the national level. Each of these results, displayed in Figure VII, is broadly similar to the baseline result in Figure IV.

Additional sensitivity. One concern related to the identification of the lag weights and, therefore, the magnification elasticities, is that they may not be robust to the definition of a minimum wage worker being a worker whose wage is no greater than an arbitrary cutoff (1.05 in the baseline) times the minimum wage.[22] Figure A.VI in Online Appendix A.2 shows that results change as expected when using a lower wage cutoff of 1.00 and a higher wage cutoff of 1.10. Finally, Figure A.VII in Online Appendix A.2 shows that results are very similar if I do not restrict to a balanced sample of (g, r) pairs.

5.3 Discussion

All models are abstractions.[23] I view the model in Section 3 as a useful abstraction because it provides insights into and simple intuition for the dynamic effects of the real minimum wage on real wages and inequality. Here, I first highlight where the empirical results in Figure IV deviate from the theory. I then discuss how changing my baseline assumption that minimum wages do not impact unemployment might affect the interpretation of the empirical results.

Perhaps the main deviation between regression results reported in Figure IV and the model's prediction is that the magnification elasticity between two-and-a-half and three-and-a-half years ($t = 3$ in Figure IV) is less than the magnification elasticity between two and three years ($t = 2.5$ in Figure IV), whereas effects are monotonic in the theory through $t = 3$. To understand this empirical non-monotonicity, I consider both a shorter time difference (three years) and a longer time difference (four years) in the definition of the dependent variable (whereas the baseline uses a three-and-a-half-year difference). These results—displayed in Figures A.VIII and A.IX in Online Appendix A.2—show that the decline in magnification elasticities is a robust result. Yet the timing of the decline appears to depend on the time difference in the dependent variable, with the decline occurring earlier when the time difference in the dependent variable is shorter and later when the time difference is longer.

[22]In a different empirical context, Derenoncourt and Montialoux (2021) define minimum wage workers as those earning no more than 1.15 times the minimum wage.

[23]The model in Section 3 abstracts from many mechanisms, including fairness or efficiency wage concerns (e.g., Grossman, 1983) and endogenous technological change (e.g, Hurst et al., 2022).

A final empirical concern that I consider is the possibility that the minimum wage affects unemployment. As discussed in Section 3.4, in this case changes in the minimum wage cause changes in value marginal products of labor. In my baseline empirical analysis, these changes in value marginal products of labor are in the residual. Hence, I identify the effect of changes in the real minimum wage, including their effects via marginal products of labor. If I controlled for changes in value marginal products (or changes in employment induced by the changes in the real minimum wage), then I would instead identify the partial effect. Since the total elasticity is the object of interest, in this respect my approach seems reasonable.

Nevertheless, it might be of independent interest to understand the extent to which the total and partial elasticities differ. Together with the theoretical results Section 3.4, the empirical literature focusing directly on the impact of minimum wages on unemployment can help answer this question. If the impact of minimum wages on employment is non-existent or very small—as in, e.g., Dube, Lester and Reich (2010) and Cengiz et al. (2019)—then any reasonable elasticity of labor demand will imply non-existent or tiny effects of minimum wages on value marginal products. In this case, estimated total elasticities will be very similar to partial elasticities. On the other hand, if changes in the minimum wage have large effects on unemployment, as might happen for large values of the minimum wage, then partial effects may deviate meaningfully from total effects.[24]

6 Conclusions

What is the impact of the real minimum wage on wages and inequality? I estimated a version of what Goldin and Katz (2008) refer to as the *The Race Between Education and Technology* and Acemoglu and Autor (2011) refer to as *the canonical model*, extended to incorporate the real minimum wage, which I refer to as *The Race Between Education, Technology, and the Minimum Wage* and *the extended canonical model*. Using national time-series data, I found that real minimum wages—together with supply and demand—play a role in shaping the evolution of the U.S. college wage premium.

I then presented a generalization of the theoretical model underlying the *The Race*

[24]Finally, note that the possibility that an increase in the minimum wage causes the least productive workers to exit employment is explicitly incorporated in the theory in Section 3.4, as the skills that are pushed into unemployment are precisely those with the lowest wages. If the empirical analysis is conducted after combining worker groups into certain aggregates (such as the college educated), then controlling for the change in the productive composition of these aggregate groups induced by changes in unemployment requires composition adjusting, as in, e.g., Katz and Murphy (1992), Card and Lemieux (2001), and my national analysis in Section 2. In my analysis in this Section, I do not aggregate up, and instead focus on 100 worker groups.

Between Education and Technology that generates my estimating equation as a first-order approximation around an initial steady state. The model also has implications for the dynamic effects of minimum wages on inequality. In the model, the elasticity of real wages with respect to the real minimum wage is small on impact and grows over time. Finally, I have provided empirical evidence consistent with these dynamic effects using data disaggregated across states, labor groups, and time periods. My estimates imply that declining real minimum wages in the 1980s in the U.S. caused a considerable decline in the real wages of workers without a high school degree and a substantial increase in the college premium. More generally, changes in minimum wages play a role in shaping the evolution of between-group inequality in the US.

While my empirical results are consistent with a job-ladder model, without worker-level panel data I have not provided direct evidence of the minimum wage affecting wages along the job ladder. I consider this to be an important avenue for future work.

References

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4 of *Handbook of Labor Economics*, Elsevier, 2011, chapter 12, pp. 1043–1171.

Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel, "Optimal Minimum Wages," *mimeo LSE*, 2022.

Autor, David, Claudia Goldin, and Lawrence F. Katz, "Extending the Race between Education and Technology," *AEA Papers and Proceedings*, May 2020, *110*, 347–51.

Autor, David H., Alan Manning, and Christopher L. Smith, "The Contribution of the Minimum Wage to US Wage Inequality over Three Decades: A Reassessment," *American Economic Journal: Applied Economics*, January 2016, *8* (1), 58–99.

_ , Lawrence F. Katz, and Melissa S. Kearney, "Trends in U.S. Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 2008, *90* (2), 300–323.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey, "Minimum Wages, Efficiency and Welfare," NBER Working Papers 29662, National Bureau of Economic Research, Inc January 2022.

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, November 2023, *91* (6), 2155–2185.

_ and _ , "Negative Weights are no Concern in Design-Based Specifications," NBER Working Papers 32017, National Bureau of Economic Research, Inc January 2024.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–273.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, March 2006, *74* (2), 323–364.

Card, David and John E. DiNardo, "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, October 2002, *20* (4), 733–783.

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

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chen, Yujiang River and Coen N Teulings, "What is the optimal minimum wage," CEPR Working Papers DP17026, The Centre for Economic Policy Research February 2022.

Clemens, Jeffrey and Michael R. Strain, "The Heterogeneous Effects of Large and Small Minimum Wage Changes: Evidence over the Short and Medium Run Using a Pre-Analysis Plan," NBER Working Papers 29264, National Bureau of Economic Research, Inc September 2021.

Derenoncourt, Ellora and Claire Montialoux, "Minimum Wages and Racial Inequality," *The Quarterly Journal of Economics*, 2021, *136* (1), 169–228.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 1996, *64* (5), 1001–1044.

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

_ , Michael Reich, Akash Bhatt, and Denis Sosinskiy, "Restaurant Employment, Minimum Wages, and Border Discontinuities," NBER Working Papers 32902, National Bureau of Economic Research, Inc September 2024.

_ , T. William Lester, and Michael Reich, "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," *The Review of Economics and Statistics*, November 2010, *92* (4), 945–964.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge, "Reallocation Effects of the Minimum Wage," *The Quarterly Journal of Economics*, 08 2021, *137* (1), 267–328.

Engbom, Niklas and Christian Moser, "Earnings Inequality and the Minimum Wage: Evidence from Brazil," *American Economic Review*, 2022, *112* (12), 3803–3847.

Freeman, Richard, *The Overeducated American*, Academic Press, 1976.

Goldin, Claudia and Lawrence F. Katz, *The Race between Education and Technology*, Harvard University Press, 2008.

Grossman, Jean Baldwin, "The Impact of the Minimum Wage on Other Wages," *The Journal of Human Resources*, 1983, *18* (3), 359–378.

Haanwinckel, Daniel, "Supply, Demand, Institutions, and Firms: A Theory of Labor Market Sorting and the Wage Distribution," *mimeo UCLA*, 2023.

Hurst, Erik, Patrick J. Kehoe, Elena Pastorino, and Thomas Winberry, "The Distributional Impact of the Minimum Wage in the Short and Long Run," NBER Working Papers 30294, National Bureau of Economic Research, Inc July 2022.

Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Lee, David S., "Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?," *The Quarterly Journal of Economics*, 1999, *114* (3), 977–1023.

Lemieux, Thomas, "Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?," *American Economic Review*, June 2006, *96* (3), 461–498.

Manning, Alan, *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press, 2013.

Mortensen, Dale T., *Wage Dispersion: Why Are Similar Workers Paid Differently?*, The MIT Press, 08 2003.

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, November 2002, *70* (6), 2295–2350.

Roth, Jonathan, "Interpreting Event-Studies from Recent Difference-in-Differences Methods," Technical Report 2024.

Shimer, Robert, "On-the-job search and strategic bargaining," *European Economic Review*, May 2006, *50* (4), 811–830.

Teulings, Coen N., "The Contribution of Minimum Wages to Increasing Wage Inequality," *The Economic Journal*, 2003, *113* (490), 801–833.

Tinbergen, Jan, "Substitution of Graduate by Other Labour*," *Kyklos*, 1974, *27* (2), 217–226.

Trottner, Fabian, "Misallocations in Monopsonistic Labor Markets," *mimeo UCSD*, 2022.

van den Berg, Gerard J. and Geert Ridder, "An Empirical Equilibrium Search Model of the Labor Market," *Econometrica*, 1998, *66* (5), 1183–1221.