Partial Automation: Routine-Biased Technical Change, Deskilling, and the Minimum Wage

Mitch Downey
pmdowney@ucsd.edu
UCSD

October 13, 2015

Abstract

The job polarization literature emphasizes the pressure technological change exerts on middle-wage occupations that often involve routine, easily automated tasks. I argue that technology only *partially* automates these tasks. Often the tasks still require labor, but technology reduces their complexity enabling less skilled workers to do the same job. As a result, the costs of technology adoption are not only the costs of the technology itself but also of low-wage workers to use it. By raising the cost of low-wage labor, the minimum wage reduces the profitability of adopting such routine-biased technologies. To test this prediction, I exploit state variation in the minimum wage and industry variation in complementarity between low-wage workers and computers to identify heterogeneous effects of the minimum wage. For the average industry, I estimate that this complementarity leads a $1 decrease in the minimum wage to raise technology use by 30% and lower the routine share of the wage bill by 1 percentage point (3.3%), both relative to a counterfactual with no such complementarity. Many routine-intensive industries have high complementarity, making the minimum wage an important policy lever to influence the pace of routine-biased technical change.

Keywords: Minimum wage; Technology adoption; Wage inequality; Job polarization

JEL Classification Numbers: J24, J31, J38

1 Introduction

The literature on wage inequality has made great progress identifying forces driving inequality and estimating their effects on the wage structure. Yet there has been little consideration of how these drivers interact with or affect one another. Here, I emphasize a particular channel: the impact of the minimum wage on the profitability of adopting technologies that automate middle-wage jobs. The large and growing body of work on routine-biased technical change (RBTC) has shown that the expanding capabilities of information technology (IT) has primarily replaced routine tasks, which are relatively easy to codify and automate.[1] Yet many of these technologies still require some level of human oversight; the automation is only partial. Instead of completely replacing workers, these technologies often reduce the skills required of workers to perform these tasks. Machines may simplify a job but still need *someone* to operate them.

If IT needs an operator then the cost of technology adoption is not simply the cost of the equipment itself but also the cost of less-skilled labor to operate it. While the declining price of IT equipment has received considerable attention (Nordhaus, 2007), little research has considered how the price of low-wage labor affects technology adoption decisions. Institutional forces like the minimum wage and market forces like technology adoption have often been treated as competing explanations for rising inequality. Instead, I argue that they are connected. A low minimum wage makes it more profitable for a firm to adopt technologies that automate middle-wage workers' tasks.[2]

To empirically test this prediction, I use a variety of data from the Census, the Current Population Survey (CPS), and the Dictionary of Occupational Titles (DOT). The econometric strategy exploits cross-state heterogeneity in the minimum wage and cross-industry heterogeneity in complementarity between low-wage workers and technology. High complementarity industries are those where IT is beneficial for low-wage workers, enabling them to perform routine tasks (traditionally middle-wage tasks) and compete better with higher

[1] See Autor, Levy, and Murnane (2003). Goos and Manning (2007) first noted that routine tasks were concentrated in the middle of the wage distribution, which implies that technology should cause "job polarization": the decline in employment and wage growth in middle-wage occupations, relative to high- and low-wage ones. See Acemoglu and Autor (2011) and Autor (2013) for reviews, Autor, Katz, and Kearney (2008) and Autor and Dorn (2013) for evidence from the US, and Goos, Manning, and Salomons (2014) and Michaels, Natraj, and Van Reenen (2014) for international evidence.

[2] The focus of this paper is the ability of low-wage workers to use technologies to become more productive than they otherwise would be. In that way, it is a paper about the complementarity of low-wage workers and technology. In political debates, it is commonly argued that low-wage workers and technology are primarily substitutes, and that raising the minimum wage will lead firms to automate low-paying jobs. Despite the prominence of this narrative, only recently have economists formally considered such an effect (Phelan and Aaronson, 2015). This work complements that by demonstrating that complementarities between technology and low-wage workers also exist.

paid workers. To measure this complementarity, I use the CPS Computer Use Supplement to measure the share of low-skilled workers who use a computer and the DOT to measure the share performing tasks that would benefit from a computer.[3] To measure technology, I use the employment of IT workers (Tambe and Hitt, 2012).[4] Using state-year fixed effects to account for unobserved state-wide trends and shocks correlated with the minimum wage, and industry-year effects to account for industry-specific technology patterns, I show that high complementarity industries adopt more technology in low minimum wage states. As predicted by RBTC, the share of compensation paid to routine workers is also lower in these states.[5] The results imply that because of this complementarity, a \$1 decrease in the minimum wage increases technology use in the average industry by 30% and lowers the routine share of the wage bill by 1 percentage point (both relative to a counterfactual with no such complementarity).

Within the large literature on job polarization, wage inequality, and the labor market effects of technological change, this paper is most related to three strands.[6] The first is the literature on endogenous technical change. Beaudry, Doms, and Lewis (2010) summarize the key insight from endogenous technical change models: "When a major new technology becomes available, it is not ubiquitously or randomly adopted... Instead, it is adopted only in environments in which complementary factors are cheap and abundant." Here I emphasize that in many industries low-wage workers are a complementary factor and therefore the minimum wage directly affects the attractiveness of technology adoption.

This paper also relates to a small literature focused on how traditionally emphasized drivers of wage inequality affect one another. Bloom, Draca, and Van Reenen (2011) and Van Reenen (2011) show that import competition from China increases the adoption of labor-replacing technologies. Lewis (2011) shows that low-skilled immigration slowed manufacturing plants' technology adoption. Acemoglu, Aghion, and Violante (2001) develop a model in which skill-biased technical change accelerates the decline of unions by undermin-

[3] I define "low-skill" using both wages and education to help reduce measurement error.

[4] Like most of the literature, my focus is primarily on computers. Below I perform extensive validation exercises to show that IT labor is a good proxy for other measures of IT adoption.

[5] As shown elsewhere (e.g., Goos and Manning (2007); Autor and Dorn (2013)), routine occupations tend to be in the middle of the wage distribution (not the bottom). Thus, the direct mechanical effect of a low minimum wage should be, if anything, to *raise* the routine share of the wage bill by lowering the wages of lower paid workers (since evidence on spillover effects suggests that they're small (Autor, Manning, and Smith, 2015)).

[6] Throughout, I focus on the effects of technology on the wage distribution. However, technological change is not the only explanation for declining demand for routine occupations. Recent international trade research has shown that offshoring exerts particular pressure on routine-intensive jobs (Baumgarten, Geishecker, and Görg, 2013; Becker, Ekholm, and Muendler, 2013; Oldenski, 2012, 2014; Ottaviano, Peri, and Wright, 2013). See Autor, Dorn, and Hanson (2015), Ebenstein et al. (2014), and Gould (2015) for discussions of disentangling the effects of trade and technology.

ing the ability for high-skilled and low-skilled workers to form coalitions.[7] Those papers emphasize the effects of immigration, trade, and technology which are often difficult for policymakers to influence. This work complements that by showing how specific policy choices affect these broad market pressures.[8]

Finally, this work relates to a series of papers on who uses, or can use, technology. Supported by empirical evidence (Krueger, 1993), the inequality literature traditionally treats computers as used only by high-skilled workers.[9] Personal computers are commonly described as a skill-biased technological revolution and contrasted with a "deskilling" revolution like the assembly line. Yet technology developers have strong incentives to make their products accessible to as wide a range of potential users as possible. As computers have become increasingly commonplace, perhaps the costs of learning to use a computer for work have fallen and they have evolved from a skill-biased to a deskilling technology.[10] Beaudry, Green, and Sand (2013) show evidence that demand for high-skilled, college-educated workers has been declining since 2000. They present a model in which cognitive tasks are a stock, rather than a flow, so that demand for high-skilled workers increases during the adoption period when it first becomes available but falls when the technology becomes well-established. An alternative explanation is that over time computer technology has become as accessible to lower skilled workers as high-skilled ones.[11]

[7] Shim and Yang (2015) show that technology adoption and job polarization between 1980 and 2009 were concentrated in industries with high wage premia in 1980. If unions increased these wage premia, then there would also be an effect of unions on accelerating technology adoption.

[8] Other evidence could be assembled on how policy choices affect market pressures like technology. For instance, Fortin (2006) showed that states with more generous higher education funding policies slowed the growth of the college premium. The results of Beaudry et al. (2010), then, would imply that this would increase computer technology adoption.

[9] Caselli (1999) provides a useful framework for thinking about how the skill requirements of new technology determine its effects on inequality. In his model, the effect depends on the cost of learning to use a new (more productive) technology. If it is easy to learn, it is adopted by all workers, closing the productivity gap between the higher and lower ability workers and reducing inequality. If it is difficult to learn, only high ability (high wage) individuals adopt the technology and inequality rises.

[10] In the 1984 CPS Computer Use Supplement, only 4% of respondents with less than a high school education had a computer at home. In 2003, this figure was 44%, though still lower than among households with more education. If computer ownership affects workers' ability to learn to use a computer at work (Malamud and Pop-Eleches, 2011) then it is increasingly implausible to think that lower skilled workers lack the capacity to use a computer at work. In Figure B3 of the Results Appendix, I confirm that computers are more commonly used by higher paid workers, but show a four-fold in the computer use among low-wage workers during the 1980's and 1990's.

[11] To my knowledge, this is the first study to emphasize that low-skilled workers' technology use is important for RBTC. Nonetheless, many past studies have provided evidence consistent with this claim. Autor, Katz, and Krueger (1998) find that industries with larger increases in computer use from 1984 to 1993 also saw larger growth in the share of employment with less than a high school education (Table V). Using the PSID, Cortes (2014) shows that it is the lowest paid non-routine manual workers who are most likely to transition into routine occupations, and that this adverse selection has increased over time (Figure 4). Several specifications in Michaels, Natraj, and Van Reenen (2014) find that industries' information and

The structure of the paper is as follows. Section 2 presents a range of descriptive evidence to build intuition for the results presented later. In Section 3 I describe my econometric strategy before describing data and measurement in Section 4. Section 5 presents my results and summarizes several robustness checks and identification tests.[12] Section 6 concludes.

2 Low-wage workers and RBTC: Descriptive evidence

This paper argues that many transformative technologies only partially automate routine occupations. Often, instead of entirely replacing workers, technology allows firms to replace workers with lower-skilled substitutes. Here I present descriptive evidence to suggest that this framework helps explain important features of the data.

Consider three occupations that exemplify this process: material recording and production clerks, statistical clerks, and photographic process workers.[13] According to the Department of Labor's Occupational Outlook Handbook (OOH), material recording and production clerks work in establishments like warehouses to "keep track of information in order to keep businesses and supply chains on schedule. They ensure proper scheduling, recordkeeping, and inventory control." Thirty years ago, these tasks required a well-organized individual to manage complex filing systems and coordinate dozens of incoming and outgoing shipments daily. Today, software handles many of these complexities, and the OOH notes that materials recording clerks "increasingly use tablets and other handheld computers." Statistical clerks "compile data and tabulate statistics for use in statistical studies... [and] tabulate statistics using adding machine and calculator." Certainly Microsoft Excel has made simplified these tasks. Finally, photographic process workers "develop and process photographic images from film or from digital media." In the days of darkrooms, this required careful light restrictions, chemical blends, and complicated projection equipment. Compared to this, modern processing of digital photography is relatively simple.

Yet while technology has transformed these occupations, it has not replaced them. Figure 1 uses data from the CPS Outgoing Rotation Groups (ORG) to trace employment and wages in these occupations over time. Panel (a) shows that despite the effects of transformative

communications technology imports increase the low-skilled share of the wage bill by at least as much as the high-skilled (Tables 5 and 6). Acemoglu (2003) finds less evidence for SBTC-driven inequality in European countries where institutions raised the pay of low-skilled labor. Feng and Graetz (2015) show that occupations requiring the most on-the-job training in the 1970's saw the greatest declines in workers' education by 2008 (Figure 6).

[12] More robustness checks, identification tests, validation exercises, and descriptive analyses are included in the Results Appendix. These are summarized throughout the paper.

[13] Occupational codes are drawn from Autor and Dorn (2013). See the Data Appendix for additional detail. All three occupations are classified as routine by Autor and Dorn and are in the second and third quintiles of the 1990 wage distribution.

Figure 1: Partial automation in select occupations

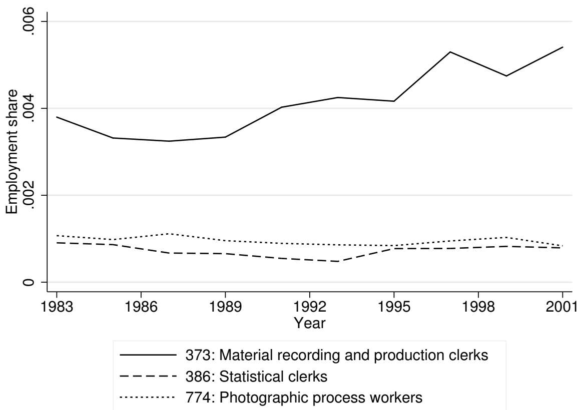

(a) Employment share

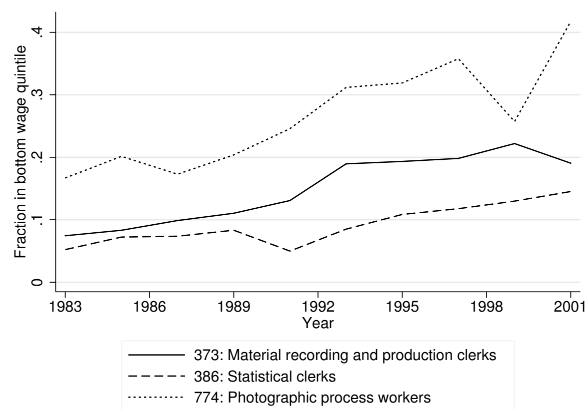

(b) Low-wage share

Source: Author's calculations based on CPS ORG. Occupation codes drawn from Autor and Dorn (2013). Among 330 occupations, the mean and median 1983 employment shares are .0030 and .0010, respectively. For these three occupations, they are .0037, .0009, and .0011, respectively. Thus, these occupations are fairly typical in terms of initial size. I focus on the 1983-2002 period because there are no major changes in occupational codes (see Figure B2). To reduce noise, the figure is based on two year bins. That is, 1983 is based on 1983 and 1984, 1985 is based on 1985 and 1986, etc.

technologies, these occupations' employment shares were constant or rising through the 1980's and 1990's, a period of rapid technological change. Rather than replacing workers, technologies seem to have simplified their skill requirements. Panel (b) shows the fraction of workers in the bottom wage quintile, a rough proxy for being near the minimum wage.[14] While employment shares were constant, the fraction of workers near the minimum wage nearly tripled for each occupation. This is consistent with the idea that technology has not automated these jobs, but enabled firms to replace workers with less skilled and lower paid alternatives.

Importantly, this effect will not always occur within occupations, particularly with coding schemes as fine-grained as Census codes. For occupations like cabinetmaker and furniture finisher or seamstress and sewing machine operator, differences in skill content and technology use lead to completely distinct codes. It is useful to consider job polarization and substitution between occupations within the same production process (proxied for by a narrowly defined

[14]In principle, one could define "near-minimum" wages by calculating the fraction of workers at or below it, but if the minimum wage raises the wages of other workers (spillover effects), this complicates the problem. One way to define "near" is to calculate the fraction earning less than the highest (nominal) minimum wage in place anywhere in the US. From 1979 to 2010, this fraction varies between 11% and 25%. Autor, Manning, and Smith (2015) estimate that the minimum wage affects wages up to around the 20^{th} percentile, although they cannot reject the hypothesis that this is pure measurement error. Compared to an absolute (wage-based) measure of nearness, a relative (percentile-based) measure avoids discrete jumps when the minimum wage changes and naturally accounts for wage trends over time and across the business cycle.

industry).[15] Autor and Dorn (2013) show that wage stagnation and employment contraction during the 1990's was concentrated in occupations in the 20^{th} to the 60^{th} percentiles of the 1990 wage distribution.[16] I refer to these two quintiles as "middle-wage" occupations. A natural question is: What happened to industries that intensively employed these workers?

Using the 1990 CPS ORG, I identify all industries in which 60% or more of workers are in these middle-wage occupations.[17] Across these 26 industries, these occupations make up over 70% of 1990 employment, more than double their employment share in other industries (34.5%).

Figure 2 plots four series. The solid line in Panel (a) shows middle-wage occupations' employment share in these industries. Consistent with the economy-wide trends, the share of industry employment in middle wage occupations fell by about 5-6% from 1990 to 2006.[18] The solid line in Panel (b) shows that these industries saw average log wages fall from 13 to 5 log points above those in other industries. These facts are consistent with the traditional understanding of RBTC and the automation of middle-wage occupations.

However, the dashed lines are less straightforward. The simplest RBTC model would imply that total industry employment would fall as these middle-wage occupations are automated. Instead, the dashed line in Panel (a) shows that these industries' share of national employment actually rose by 15% during this period.[19] Employment expanded, not contracted.[20] To better understand the wage trend, I estimate a simple Mincer regression in which log wages are modeled as a function of age, education, race, and other wage correlates, and compare the average *fitted* log wage between these industries and others, plotted as the dashed line in Panel (b).[21] Since the mid-1990's, these trends are remarkably similar. Since

[15]Autor and Dorn (2013) estimate that a third of job polarization can be accounted for by employment growth across industries (specifically, in the service sector). My focus is on the within-industry component of job polarization.

[16]Figure B1 uses data from Autor and Dorn (2013) to divide occupations into quintiles based on their 1990 wage, and traces the change in employment share and average wage (by quintile) from 1990 to 2005. The second and third quntiles are the only ones with declining employment shares and these also have the slowest wage growth.

[17]Since these occupations, by construction, comprise 40% of the national economy in 1990, they are *at least* 50% more common in these industries than nationally.

[18]I focus on this period for consistency with the Autor-Dorn data and to abstract from the 2007 recession effects.

[19]Certainly it is possible that expanding employment and increased automation are both driven by product demand shocks. My emphasis is that the marginal workers in these occupations are observationally different from the "original" workers.

[20]This result does not contradict the early industry level analyses of Autor, Levy, and Murnane (2003), who study changes in the task content, not overall employment, of routine-intensive industries. Since that work, most RBTC research has focused on occupations, abstracting from industries.

[21]Mincer wages are based on an OLS regression of log wages on years of education, sex, indicators for Black and White race, age, age squared, and marital status. The estimation sample is restricted 1990-2006 and workers between the 2^{nd} and the 98^{th} percentile. No coefficients are allowed to vary over time.

Figure 2: Industries intensively employing middle-wage occupations in 1990

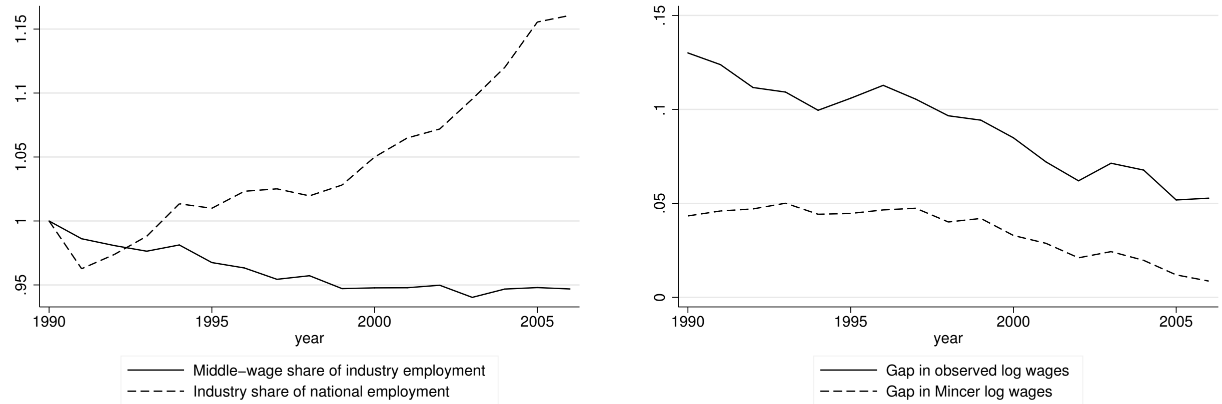

(a) Employment shares

(b) Relative log wages

Source: Author's calculations based on CPS ORG. Industry codes drawn from Autor, Dorn, and Hanson (2013). See the Data Appendix for more detail. Employment shares are normalized to 1 in 1990 because they are on different scales. In 1990, these industries (26 of 211) accounted for 14% of national employment (25% of middle-wage occupations' employment) and middle-wage occupations made up 70% of these industries' employment. Mincer wages are based on an OLS regression of log wages on years of education, sex, indicators for Black and White race, age, age squared, and marital status using data from 1990-2006. The estimation sample for this Mincer regression is restricted to workers between the 2^{nd} and the 98^{th} percentile. None of the coefficients are allowed to vary over time.

1993,[22] the gap in average wages fell by 6 log points, while the gap in Mincer wages fell by over 4 log points. Most of the decline in wages in middle-wage intensive industries is matched by a decline in observable worker characteristics. Figure 2 shows that the industries hardest hit by job polarization are employing more workers, and that these workers are of lower "quality" along observable dimensions.[23]

Figures 1 and 2 suggest that RBTC and job polarization are partly driven by compositional changes in the types of workers holding routine, middle-wage occupations. If technology allows less skilled workers to substitute for middle-skilled ones, then the minimum wage (and prevailing low-skilled wage generally) may have important implications for the pace of RBTC. Figure 3 shows this is the case. I combine the BEA capital stock data with the 1980 Census and the 2005 ACS. For each of the 56 industries available in the BEA

[22]In 1993, the CPS changed the measurement of education. Because years of education is such an important determinant of wages, this discrepancy may explain why the patterns diverge before 1993.

[23]It has long been understood that the wage changes documented job polarization literature could be driven by the changing composition of workers. This question is taken up directly by Böhm (2015) and Cortes (2014) who use panel data to study worker sorting. Both papers find composition changes driven by sorting responses to routine wages, yet that the composition-adjusted wage still fell (but see also Gottschalk, Green, and Sand (2015)). Those papers show that some of the change in routine wages is unexplained by composition changes; my emphasis is that some of the change *is* explained.

Figure 3: Time-variation in the routine-bias of technology and the minimum wage

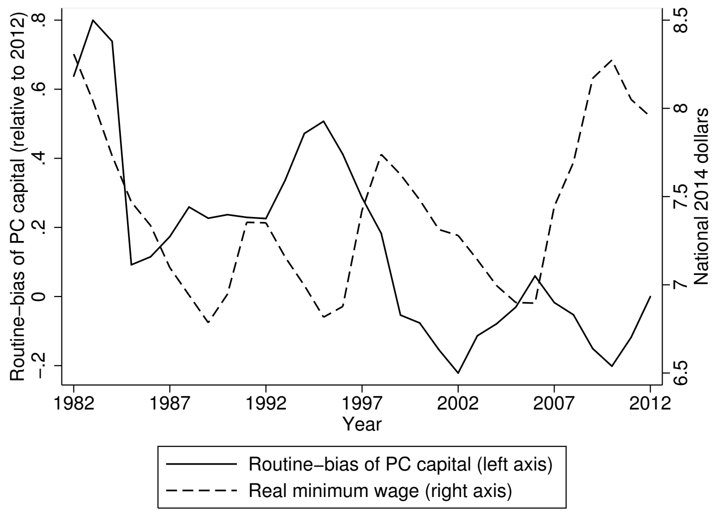

Source: The real minimum wage is the employment-weighted average minimum wage in national 2014 dollars. For each state, I calculate the effective minimum wage (the maximum of the state and federal minimum wage). I convert them to real dollars. I then take a weighted average across states, weighted by state employment shares. The routine-bias of PC capital is the estimated β_t series from equation (1) in the text. It captures time variation in the correlation between PC capital and declining routine employment. It is normalized to be zero in 2012.

data, I calculate real PC capital (PCK_{it}) and the change in log routine employment share from 1980 to 2005 ($\Delta \ln(RoutineEmpShare)_i$).[24]

The change in log routine employment share is a time-invariant measure of how severely an industry has been affected by RBTC. The goal of the exercise is to summarize the time varying degree to which differential PC adoption rates map to RBTC-affected industries. That is, when do the industries losing the most routine employment adopt the most technology? I estimate the following regression:

$$\ln PCK_{it} = \alpha_i + \delta_t + \beta_t \Delta \ln(RoutineEmpShare)_i + \varepsilon_{it} \tag{1}$$

The β_t series captures variation across time in the "routine-bias" of PC adoption. When β_t is larger, routine-declining industries adopted more technology than when β_t is small.[25] Figure 3 plots the resulting series of the routine-bias of technological adoption along with the real minimum wage.

[24] $\Delta \ln(RoutineEmpShare)_i \equiv \ln(RoutineEmpShare)_{i,2005} - \ln(RoutineEmpShare)_{i,1980}$

[25] Of course, with the full set of industry (α_i) and year (δ_t) fixed effects, the levels of the β_t series cannot be identified; it has to be normalized to some year. There must be some τ such that $\beta_\tau = 0$ and all other β_t are relative to year τ. I choose to set $\tau = 2012$ and identify the routine-bias of PC adoption in each year, relative to 2012.

Figure 3 illustrates a simple point. For industries that saw the greatest declines in routine employment, the over-time variation in technology adoption is closely and inversely related to the minimum wage. The periods when PC capital adoption became more heavily concentrated in these industries were almost all periods of falling minimum wages. Though this is only a simple time series correlation (with small T), the results in the remainder of the paper suggest this relationship may be causal.

3 Econometric strategy

My primary specification is as follows:

$$Tech_{ist} = \alpha_{st} + \delta_{it} + Z'_{ist}\gamma + \beta(Complementarity_{it} \times MinWage_{st}) + \varepsilon_{ist} \qquad (2)$$

In (2), $Tech_{ist}$ denotes technology use in industry i in state s at time t. Complementarity is the share of low-skilled workers who benefit from using technology. In the next subsection, I discuss measurement of technology and of low-skilled/technology complementarity in detail. Here, the important feature is that complementarity is measured at the industry level and the minimum wage ($MinWage$) is measured at the state level. Because α_{st} accounts for cross-industry state-specific factors correlated with the minimum wage[26] and δ_{it} accounts for arbitrary industry patterns of technology adoption over time, identification of β is driven solely by the interaction of complementarity and the minimum wage. Thus, β captures the heterogeneous effect of the minimum wage on technology use, differentially for high complementarity industries. While the advantage of this is that it makes weaker identification assumptions than a specification which assumes the minimum wage is exogenous, the disadvantage is that it does not identify the *direct* effect of the minimum wage, only heterogeneous effects. This is discussed in more depth below.[27]

[26] Allegretto, Dube, and Reich (2011) find that the minimum wage tends to be raised during periods of declining low-wage employment, though Sabia (2014) argues that this is at odds with research on the political determinants of minimum wage changes. While it may seem unlikely that minimum wage changes respond to shocks in technology use, because of the relationship between technology and inequality, states responding to trends in inequality may inadvertently respond to changes in technology adoption. Moreover, for political reasons, changes in the minimum wage may be correlated with changes in banking regulation or tax policy, both of which might affect investment in IT capital.

[27] This issue is common in wage inequality research. Much of the literature focuses on the effects of time-varying pressures (e.g., technology, Chinese imports, etc.) that are heterogeneous across different units (e.g., industries, commuting zones, etc.). The data is typically two dimensional (industry-year, state-year, etc.) and it is common to include a year fixed effect. As a result, one dimensional variation is perfectly absorbed and direct effects of time varying technology availability, for instance, can never be identified. Here, the data is three dimensional (industry-state-year), and the fixed effect scheme perfectly absorbs two dimensional variation.

For the intuition behind the identification strategy, consider two firms within the same industry but in different states. If industry codes are sufficiently precise, then both firms are engaged in the same economic activity and face the same production possibilities frontier. The chosen methods of production will depend on local factor prices. In equation (2), an estimate of $\beta < 0$ indicates that industries where low-wage workers benefit heavily from technology use less technology in states with a higher minimum wage (after α_{st} accounts for any sources of state heterogeneity that affect all industries equally). Note that this identification strategy is purely cross-sectional (identification comes from within st and it).[28]

Figure 4 communicates both the identification strategy and the results in two different ways. Panel (a) shows technology adoption across the minimum wage distribution for a set of industries with high complementarity and a set with low complementarity.[29] In my context, Panel (a) is comparable to a difference-in-difference plot, where "treatment" industries are those with high complementarity and "control" industries are those with low complementarity. The "parallel trends" assumption is that the groups are similar when the treatment is irrelevant ("pre-treatment"). This can be seen in the far right of the graph: when the minimum wage is sufficiently high, the extent to which low-wage workers benefit from technology is irrelevant because low-wage workers are so expensive. The difference between the treatment and control group is irrelevant, and technology use is similar for both sets of industries. As the minimum wage declines, however, the benefits from having high complementarity lead to a widening technology adoption gulf; high complementarity industries begin to use technology at a much higher rate than low complementarity ones.

Panel (a) flexibly captures state heterogeneity in the minimum wage, but imposes a lot of structure on industry heterogeneity by combining industries into only two groups. Panel (b) presents a complementary summary, which is restrictive in state heterogeneity, but flexible in industry heterogeneity. For each industry-year (it) in my main estimation sample, I regress technology in state s on the minimum wage in state s in the following specification:[30]

[28] In light of the Sorkin (2015) critique of the minimum wage literature (that technological adjustment takes time and, therefore, the long-run effects of the minimum wage are likely quite different from the short-run), a cross-sectional identification strategy has strong advantages. Most of the minimum wage variation that I use is persistent across states over time, allowing me to more credibly identify long-term effects.

[29] As discussed in the next section, my preferred measure of technology is IT workers per 100. Panel (a) shows this measure for a single year (2000) across states with different minimum wages (rounded to the nearest 50 cents) for two sets of industries: 45 industries which have below average low-wage/technology complementarity across all four of my metrics (described in detail below), and 31 industries that have above average complementarity across the four metrics.

[30] This estimating equation is similar to my main specification (equation (2)). Equation (2) includes state-year fixed effects that account for sources of heterogeneity that are correlated with the minimum wage and affect technology in all industries equally, and imposes that β_{it} varies linearly across complementarity.

Figure 4: Graphical intuition for the identification strategy

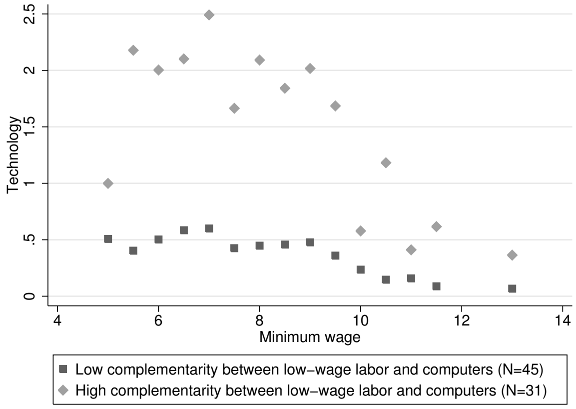

(a) Tech. across the minimum wage

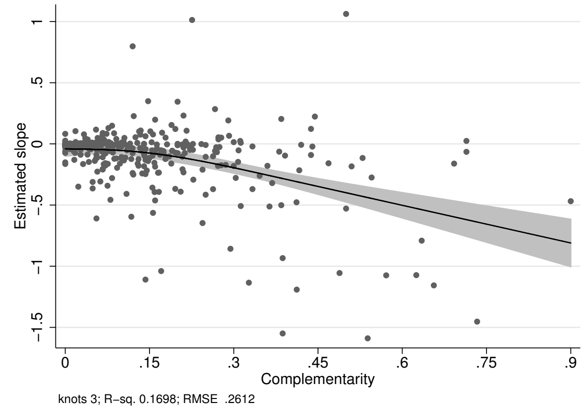

(b) Tech. gradients across complementarity

Panel (a) presents technology, measured as IT workers per 100, across the real minimum wage for the year 2000. As described below, nominal minimum wages are converted to real using both the Consumer Price Index and cross-state price deflators. Figures used in the plot correspond to 2014 national dollars, rounded to fifty cents for graphical purposes. Low complementarity industries have below average complementarity across all four measures described below; high complementarity industries have above average complementarity across all four measures. Panel (b) plots the slope coefficients (β's) from the following regression, separately estimated for each industry-year: $Tech_{ist} = \alpha_{it} + \beta_{it} MinWage_{st} + \varepsilon_{ist}$. The horizontal axis is the CUS-based complementarity among workers with less than a high school education (the primary measure used in the econometric analysis, described below). The curve is a restricted cubic spline with three knots, and the shaded area is the 95% confidence interval.

$$Tech_{ist} = \alpha_{it} + \beta_{it} MinWage_{st} + \varepsilon_{ist} \tag{3}$$

The estimated β_{it} is a linear summary of the technology adoption gradient in industry i and year t across states with heterogeneous minimum wages. Panel (b) then plots these estimated slopes across my main measure of complementarity, along with a non-parametric fit. The figure shows that very low complementarity industries exhibit little relationship between technology adoption and the minimum wage. As complementarity rises, however, industries begin exhibiting significantly negative slopes, indicating less technology adoption in higher minimum wage states. Panel (b) also provides support for the identification assumption: The only industries that systematically respond to cross-state minimum wage differences are those with significant complementarity. My econometric strategy essentially formalizes the insights from Figure 4.

My identification strategy is, of course, not perfect. The minimum wage is politically contentious and therefore highly endogenous. Causal inference has proved difficult for the literature. The key identification assumption is that the 205 private sector industries in the Census industry codes are precise enough that firms within the same industry-year are engaged in the same economic activity and face the same production possibilities frontier. The biggest threat to identification is that these firms are not valid counterfactuals for one another. In this case, rather than reflecting how the input mix responds to factor prices while holding the production function constant, observed differences in technology might result from firms (within industry-year) being engaged in more technology-intensive production in states where the minimum wage is lower.

Because of the state-year fixed effects, the direction of this bias is difficult to assess (it must be state heterogeneity that differentially affects high complementarity industries). One plausible expectation is that the structure of economic activity is more technology-intensive in high minimum wage states, and that this is particularly true in industries with more low-wage/technology complementarity.

This sort of unobserved heterogenetiy would impart a positive bias upon $\hat{\beta}$, biasing me against finding $\beta < 0$. In section 5.4, I look at cross-state variation in industrial composition and find suggestive evidence for this sort of heterogeneity. Though non-significant, the point estimates suggest that within-industry heterogeneity in economic activity, if anything, biases me against finding input mix responses (i.e., biases me against finding $\beta < 0$). I find no evidence that such heterogeneity "explains" or drives my results.

4 Data and Measurement

4.1 Measuring technology

Following Tambe and Hitt (2012), I proxy for technology using the share of employment in IT occupations.[31] Specifically, I use three occupations: computer scientists, computer software developers, and repairers of data processing equipment.

There are three concerns in using this variable. The first is the possibility that IT workers are directly affected by the minimum wage. However, Table B1 of Section B.2.1 in the Appendix shows that these workers are well-paid and well-educated, and therefore unlikely to be directly affected. The second is that the minimum wage affects total employment, the denominator of the IT labor share, and thereby affects my technology measure without affecting actual technology. This seems unlikely given the small employment effects found in the minimum wage literature. Also, as discussed below, I control for an industry-state-year employment and the fraction of industry employment near the minimum wage, further

[31] Section B.2.1 of the Data Appendix describes some advantages of a technology measure based on labor data, relative to other sources used in the literature.

mitigating these concerns.[32] Moreover, in robustness checks in the Results Appendix, I use different normalizations including one in which IT worker counts are divided by an exogenous measure of "potential" employment at the industry-state level. All of my results are robust to alternative normalizations (see Table B7).

The third concern with a technology measure based on IT labor is whether or not it correlates with other more standard measures of technology. Tambe and Hitt (2012) use firm-level data and show that IT labor is highly correlated with IT capital stock measured by Harte-Hanks. In section B.2.1 of the Appendix, I present several additional confirmatory analyses. I use the BEA's industry level data on IT capital stock and computers per worker (for 53 industries), and estimate elasticities of IT capital with respect to IT labor that are highly significant and range from 0.6-0.9, depending on the specification.[33] As noted by Tambe and Hitt, it is not obvious that IT capital is a *better* measure of technology. Results to a survey conducted by the technology magazine *InformationWeek* show that salaries and benefits comprise 32-35% of IT spending, the largest single component of IT budgets and more than double that of "new product/technology purchases." Unsurprisingly, then, when I combine *InformationWeek*'s estimates of total IT spending for 19 industries with my IT labor measure they are highly correlated.[34] Finally, Section B.2.1 also shows that the geographic variation in IT labor (after adjusting for cross-state differences in industrial composition) is highly correlated with the geographic variation in adjusted PC's per worker used by Beaudry et al. (2010) and Autor and Dorn (2013).

4.2 Measuring complementarity

I define industries to have high complementarity between low-skilled workers and technology if low-skilled workers' productivity is significantly enhanced by technology availability. In measuring this variable, there are two challenges: defining complementarity and defining low-skill. I discuss these briefly before turning to the benefits of leveraging multiple measures to address measurement error. Regardless of the definition of "low-skill" or source used to measure individual level complementarity, my industry level measure is the fraction of low-skilled workers who are computer complements.

[32] "Controlling" for total employment does not solve the problem because employment enters the IT labor share measure non-linearly (specifically as the denominator).

[33] One important result from the exercises in Section B.2.1 is that IT labor is a much stronger proxy in the cross-section than over time. While IT capital has increased dramatically over the last 30 years, IT employment has risen only slightly. Fortunately, my identification strategy exploits only cross-sectional variation.

[34] The *InformationWeek* sample is not representative. However, when the same survey asked respondents what steps they have taken to optimize business efficiency, around 80% reported "increased use of automation" (more than any of the other ten options) so it is certainly a sample of interest.

4.2.1 Defining complementarity

Consider a simple model in which capital costs are equal across all industries. Profit-maximizing firms will equalize the marginal benefit and marginal cost of capital investment. In industries where low-skilled workers benefit more from computers, firms will be more willing to undergo the investment cost. Thus, one can infer industry heterogeneity in the extent to which low-skilled workers *benefit from* computers by observing industry heterogeneity in how many low-skilled workers *actually use* a computer.

I use the Computer Use Supplement (CUS) to the CPS, which asks respondents if they "directly" use a computer at work.[35] I exploit industry heterogeneity in the share of low-skilled workers (defined below) who report using a computer directly at work.[36] The benefit of the CUS measure is that it is transparent and fairly straightforward to interpret as evidence that low-skilled workers benefit from using a computer. The disadvantage is that it is only asked during one month every five years, yielding small samples.

I complement this measure with a subjective measure of the potential for workers to benefit from a computer based on variables from the 4^{th} edition of the Dictionary of Occupational Titles (DOT). The DOT provides a wide range of information about occupations. One set of variables refers to work fields, defined as "organizations of specific methods either characteristic of machines, tools, equipment, or work aids, and directed at common technological objectives, or characteristics of the techniques designed to fulfill socioeconomic purposes." In other words, work fields describe specific tasks (methods and characteristics of techniques) involved in an occupation.

Each respondent selects up to four work fields from a list of 100. These 100 are combined into 28 groups, and I subjectively select three of these groups to represent complementarity:

1. Receiving, storing, issuing, shipping, requisitioning, and accounting for stores of materials
2. Preparing and maintaining verbal and/or numerical records
3. Providing, or effecting the transmission of, information to other persons, indirectly (by electrical or electronic media) or directly (by voice or written statement)

I believe one could perform these tasks far more effectively and efficiently using a computer. For each occupation in the DOT, I calculate the fraction of respondents that report

[35] The CUS is not available in every year. Table A1 shows how I map CUS years to Census years.

[36] In an influential paper, Krueger (1993) used this data to show that computer use increases as wages rise. Figure B3 in the Appendix shows computer use over the wage distribution for 1983, 1994, and 2003. In each year, the conclusions of Krueger hold: Higher wage individuals are more likely to use a computer. Yet while attention has focused primarily on the *slope* of the computer use-wage gradient, the intercept is non-trivial and rising appreciably over time. In 1984, only 5-10% of workers earning less than $10 per hour (about the 25^{th} percentile of wages) used a computer. By 2003, this had increased to 20-30%.

at least one of these work fields.[37] and merge these occupation level measures with the CPS ORG. I then calculate the fraction of low-skilled workers in an industry-year which have a "complementary" occupation. Note that the set of occupations does not change over time, but the fraction of low-skilled workers in these occupations might change. Thus, this measure of complementarity is still allowed to change over time.

Selecting appropriate work fields for a given purpose is inherently subjective, and the criticism of Autor (2013) is worth repeating: "While I have found that the task measures distilled from DOT and O*NET can serve as powerful proxies for occupational tasks, I am at best only moderately comfortable with these tools because their complexity and opacity places little discipline on how they are applied and interpreted." This criticism certainly applies to this work. For transparency, Table A3 lists all 100 available work fields. In addition to subjectivity, Autor (2013) points out that occupation level task measures ignore within-occupation heterogeneity and changes over time both in the set of tasks that make up occupations.[38] Nonetheless, the DOT complementarity measure is useful for understanding the degree to which an occupation's tasks can benefit from computer availability.

4.2.2 Defining low-skill

With these measures in hand, it is still necessary to define "low-skilled workers" to calculate the industry level complementarity measures described in Section 3. In the wage inequality literature, it is common to identify different skill levels using education, and I follow that approach here, defining "low-skilled" workers as those with less than a high school education.

However, in a study of the effects of the minimum wage, it also seems natural to use a definition of low-skill based on the wage. The challenge in doing so is that the minimum wage affects the wage distribution (both directly and through spillovers). Consider, for example, two identical workers, one who earns $9.75 in a state with a $6 minimum wage and one who earns $10.25 in a state with a $9.50 minimum wage. Their wages differ only due to the minimum wage in place (by assumption, they are identical workers). A wage-based definition of low-skill which draws a line at, say, $10 would classify only one of these workers as low-skill. This would induce measurement error in the complementarity variable that is correlated with the minimum wage in the states in which an industry locates.

[37] The 4^{th} edition of the DOT has around 60,000 respondents. Occupations are based on 1970 codes. See the Data Appendix for more detail.

[38] Autor and Handel (2013) provide evidence of within-occupation task heterogeneity. See Autor et al. (2003) and Autor and Dorn (2013) for detailed discussions of changes in occupations' task composition. Note that the 4^{th} edition of the DOT was revised in 1991 and replaced by O*NET in 1998, but that these datasets do not have work field variables comparable to the ones used here.

To address this concern, I develop latent wages based on a method proposed by Autor, Manning, and Smith (2010).[39] Section B.1.4 of the Appendix describes this method in detail and summarizes the results. In short, it involves assuming that the latent wage distribution (the wage distribution which would prevail without a minimum wage) is log-normal and using wages in the 50^{th}-75^{th} percentiles to estimate the parameters of the state-year wage distribution. Under the assumption that employment effects are small and that spillovers monotonically decline away from the minimum wage (which implies that there are no rank reversals), the percentiles of the observed state-year wage distribution can be used to create a one-to-one mapping from observed wages to latent wages. I then define low-skilled workers to be those with an estimated latent wage below $9.50, the point at which the effects of the minimum wage are less than 5% of the latent wage (see Figure B7). On average, this is just over 20% of the population.

4.2.3 Combining multiple measures

In total, then, I develop four industry level measures of low-skill computer complementarity: two based on the CUS and two based on the DOT; and two based on workers with less than a high school education and two based on workers with latent wages (estimated) below $9.50. These multiple measures help me address two important types of measurement error. First, there is "conceptual" measurement error. It is not straightforward to define complementarity and the measures described above have important limitations. The second is "population" measurement error. Worker skills are high dimensional and continuous, and identifying "low-skilled" workers is inherently problematic. My preferred specification uses DOT complementarity among low-wage workers as an instrument for CUS complementarity among low-education workers, helping eliminate both types of measurement error simultaneously.

[39] Autor, Manning, and Smith (2010) are interested in revising estimates from Lee (1999) of the wage inequality effects of the falling real value of the federal minimum wage during the 1980's. They propose two approaches to estimating the minimum wage effects at a given percentile of the wage distribution: one reduced form (similar to Lee, with additional corrections for endogeneity and unobserved heterogeneity) and one parametric (the approach used here). They find similar results for both methods. Because findings are so similar, the more recent version of the paper (Autor, Manning, and Smith, 2015) abandons the parametric approach I use here and includes only the reduce form approach (which requires fewer assumptions). For my purposes, unlike for the research question that Autor, Manning, and Smith focus on, the reduced form approach is not a viable substitute. Thus, I take comfort in the fact that the two approaches yield such similar results in Autor et al. (2010), lending credence to the validity of the parametric approximation.

4.3 Additional data details

The final dataset is based on IT labor measured in the 5% public-use samples of the 1980, 1990, and 2000 Census and the 2005 American Community Survey and the CPS ORG from 1980-2006.[40] The main estimation sample is that for which at least 10 respondents are available for each measure of complementarity is available (109 industries and 327 industry-years).[41] Note that this implicitly restricts my sample to industries with a non-trivial number of low-skilled workers.

My independent variables of interest are the interaction of industry level complementarity measures and the state level minimum wage. For simplicity later, let

$$X_{ist} \equiv Complementarity_{it} \times MinimumWage_{st}$$

I am, specifically, interested in how the technology effects of the minimum wage differ by the extent to which low-skilled workers benefit from computers. To ensure that my results are not driven by the differential response of industries with lots of low-wage workers or by industries where everyone benefits from technology, I control for the interaction of the minimum wage with the low-wage employment share and with complementarity measured among non-low-skilled workers. Because total employment appears in the denominator of the IT labor-based technology measure, I control for log employment.

Finally, I adjust all wages to be in real dollars not only using the CPI, as is standard, but also using state price deflators from the BEA (Aten and D'Souza, 2008). These price indices are available only at one point in time, so nominal (state level) wages are converted to real (national) wages by multiplying by a single state-specific factor (in addition to a year-specific price index for inflation). State-specific price indices have not been used in past minimum wage research, likely because they are only available at a single point in time. In

[40] Public-use census data and ACS data are from the Integrated Public Use Microdata Series (Ruggles et al., 2010). Decennial Census years are different than CUS years, and this is discussed in the Data Appendix. Industry codes are from Autor, Dorn, and Hanson (2013) and minimum wage data is from Meer and West (2013). I use the NBER version of the ORG (Feenberg and Roth, 2007). Several industries (e.g., IT consulting) were excluded because IT labor is an ineffective proxy for IT capital (as discussed in Section B.2.1). My sample ends in 2005 because the CUS has not collected data on computer use at work since 2003. For better or worse, this means that my estimation sample excludes the most recent round federal minimum wage increases (see Clemens and Wither (2014) for a discussion) and the large structural RBTC adjustments occuring during the Great Recession (see Jaimovich and Siu (2014) or Foote and Ryan (2014) for a discussion).

[41] In robustness checks, I relax this restriction, including any industry with at least one respondent available for each measure of complementarity. The results are similar. Consistent with significant measurement error, the coefficients are smaller and the standard errors are somewhat larger. Weighting by the respondent count, unsurprisingly, brings the results closer to those of my main specification. See Table B8 of the Results Appendix.

the standard state-year difference-in-difference models a state-specific price index provides no information that can be separated from the state fixed effect (indeed, if the logged minimum wage is used, the state effect is perfectly colinear with the log price index). With industries as an additional dimension of variation, this is not the case. Variation in price levels creates minimum wage variation that has differential effects on different industries and which is not entirely absorbed by the state fixed effect. Because my identification strategy exploits cross-sectional variation in the minimum wage to trace out differential technology adoption patterns (rather than time series variation to identify the effect of minimum wage changes), this additional variation is very useful. From 1980-2005, the federal minimum wage was binding during at least one month in anywhere from 37-49 states (including the District of Columbia). However, these states have very different price levels. The BEA state level price index estimates allow me to use this information.[42] Section B.1.3 of the Appendix provides additional information about how real wages are computed and the effect of state level price variation on observed minimum wages.

5 Results

5.1 Summary statistics and key correlations

Table 1 presents the summary statistics for the main estimation sample. As seen in Figure 5, there is substantial variation in complementarity across industries. For each of the four measures, 16-29% of low-skilled workers in the average industry are computer complements, and the interquartile range is at least as large (with the 90-10 range being nearly twice as large). Also note that many of these industries have high routine employment (on average, 34% of employees and 30% of the wage bill).

Finally, Table 1 highlights two important limitations of my empirical approach. First, the share of employment in IT occupations is fairly low (around .75% on average), which necessitates reliance on decennial Census data rather than, say, the CPS.[43] Second, the sample sizes available for the CUS are somewhat small, particularly for low-wage workers.[44] Those

[42] Handbury and Weinstein (2014) use detailed price scanner data to show that accounting for the availability of goods is important for comparing prices across cities, and they highlight that BEA estimates cannot account for this. While the BEA estimates are imperfect, they provide some useful information and are better than assuming constant prices across states. Price indices as detailed as those calculated by Handbury and Weinstein are not yet available at the state level (or for a sufficiently large set of cities). Figure B5 in Section B.1.3 of the Appendix shows that the BEA's price indices are highly correlated with state variation in the residual nominal wages of routine workers (correlation: .85). Thus, though imperfect, they do well in capturing variation in the prices of the inputs most relevant for this analysis.

[43] While median IT employment is zero, 48.9% of observations have positive IT labor.

[44] Sample sizes are about 25% higher for the less than high school population for two reasons. First,

Table 1: Summary statistics

Variable	Mean	SD	Percentiles 10th	25th	50th	75th	90th
Complementarity, CUS, $w < \bar{w}$	.204	.203	0	.053	.144	.289	.497
Complementarity, CUS, $< HS$	.158	.154	.016	.053	.114	.211	.368
Complementarity, CUS, any wage	.369	.233	.092	.173	.340	.525	.708
CUS sample size, $w < \bar{w}$	93	220	14	21	44	90	184
CUS sample size, $< HS$	117	222	18	28	55	114	223
Complementarity, DOT, $w < \bar{w}$	.288	.178	.086	.143	.250	.419	.538
Complementarity, DOT, $< HS$	.235	.174	.065	.105	.173	.337	.524
Complementarity, DOT, any wage	.260	.133	.103	.155	.235	.362	.445
DOT sample size, $w < \bar{w}$	1548	3322	211	387	732	1443	2840
DOT sample size, $< HS$	1077	1973	137	257	500	1042	1832
Routine emp. share, $w < \bar{w}$	.386	.234	.108	.227	.358	.513	.697
Routine emp. share, any wage	.340	.201	.103	.215	.299	.438	.594
Low-wage employment share	.286	.174	.077	.146	.247	.427	.539
Less than HS employment share	.195	.121	.052	.110	.175	.259	.372
Real minimum wage	7.96	1.36	6.35	7.03	7.81	8.63	10.01
IT workers per 100 employees	.753	1.91	0	0	0	.711	2.13
Routine share of wage bill	.296	.198	.092	.161	.246	.379	.575

Descriptive statistics are based on the main estimation sample of 14,152 industry-state-years. See Data Appendix for discussion of sample selection, and Sections 4.1-4.3 for discussion of variable measurement. While median IT employment is zero, 48.9% of observations have positive IT labor.

available for DOT complementarity are much higher since all that is required is occupational codes. Thus, in addition to the sources of measurement error discussed above, for the CUS variables I am additionally concerned about measurement error caused by sampling variability. As a result, my preferred specifications are based on using DOT complementarity to instrument for CUS complementarity, rather than the converse.

Table B5 of the Appendix contains the first stage regressions corresponding to this instrumentation scheme (as well as potential alternatives). The intuition, however, is better captured by Figure 5. Panel (a) shows that DOT complementarity is highly correlated with CUS complementarity, whether it is based on low-wage workers (dots) or low-education workers (crosses). At the same time, there is substantial deviation from the 45° line due to measurement error in defining complementarity. Panel (b) shows that complementarity among low-wage workers is highly correlated with complementarity among low-education

education is observed for any CUS respondent, whereas wages require that the respondent be successfully merged with the ORG. Second, education is almost never imputed, whereas wages are frequently imputed, particularly in more recent years (Hirsch and Schumacher, 2004). Following the literature, I do not use imputed wages. See the Data Appendix for more.

Figure 5: Correlations between multiple measures of low-skill complementarity

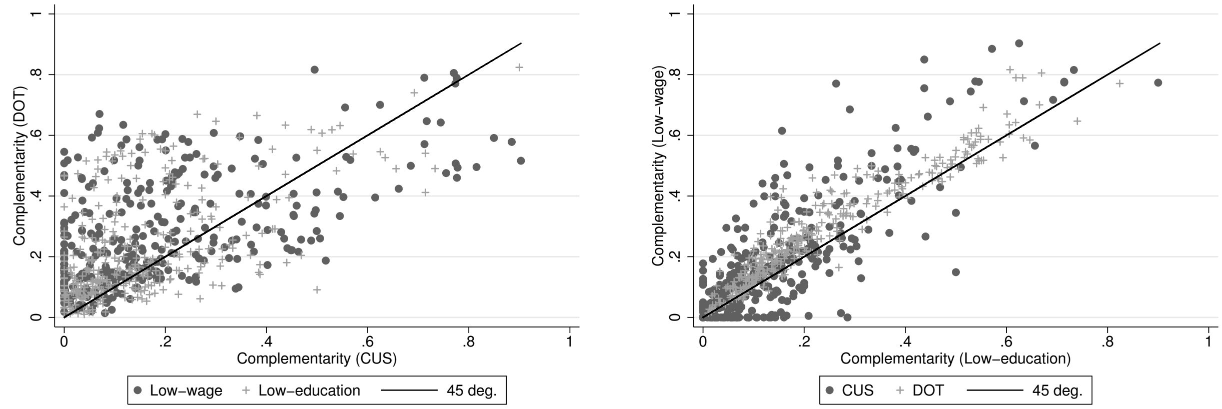

(a) Complementarity measures

(b) Skill measures

Panel (a) shows the relationship between DOT complementarity (y-axis) and CUS complementarity (x-axis) for industry-years in the main estimation sample. Panel (b) shows the relationship between computer complementarity among low-wage workers (y-axis) and complementarity among low-education workers (x-axis), also for the main sample.

workers, whether complementarity is based on the CUS (dots) or the DOT (crosses) measure. Again, much of the deviation from the 45° line is due to measurement error in defining low-skill.[45] My instrumentation strategy helps eliminate both types of measurement error at once.

A final important note is that my measures of complementarity are not unrelated to the Autor-Dorn routine indicator. In fact, they are positively correlated. Table 2 presents individual (columns 1 and 2) and occupation (columns 3-6) level regressions of computer use (columns 1-4) and DOT complementarity (columns 5 and 6) on the routine indicator from Autor and Dorn (2013). Individuals in a routine occupation are 1/3 more likely to use a computer. For low-wage workers, being in a routine occupation makes one 2/3 more likely. For the DOT-based measure of complementarity, the relationship is even more stark. Routine occupations have 3-4 times the DOT-based complementarity of non-routine occupations.

Initially, this seems to suggest that these are not valid measures of complementarity. The routine measure, after all, was developed to identify tasks and occupations for which technology was a substitute. However, the division between benefiting from technology and being replaced by technology is sometimes ambiguous. For instance, consider an accounting

[45] The CUS-based measure has significantly more variation from the 45° line than the DOT measure. Some of this is likely due to the fact that it is based on smaller samples than the DOT-based measure, and thus has more sampling error. My preferred specification is to instrument for CUS complementarity using DOT complementarity, rather than the other way around, to reduce the influence of this type of measurement error.

Table 2: Correlation between complementarity and routineness

DV :	Computer Use				DOT Compl.	
	Individual level		Occupation level			
	(1)	(2)	(3)	(4)	(5)	(6)
Routine	0.133*	0.090*	0.076**	0.117***	0.204***	0.243***
	(0.070)	(0.049)	(0.035)	(0.043)	(0.036)	(0.062)
Constant	0.398***	0.143***	0.396***	0.172***	0.085***	0.079***
	(0.043)	(0.020)	(0.021)	(0.021)	(0.012)	(0.020)
N	154390	29380	330	106	330	106
R^2	0.057	0.031	0.014	0.074	0.121	0.148
Sample	All	Low-wage	All	Low-wage	All	Low-wage
Year effects	Yes	Yes				

* $p < .10$, ** $p < .05$, *** $p < .01$. Occupation codes and routine indicator are taken from Autor and Dorn (2013). Low-wage is defined at the individual level as a latent wage below \$9.50 and at the occupation level as at least 16% of workers with latent wage below \$9.50 (approximately the mean). At the occupation level, computer use corresponds to the fraction of respondents who report using a computer. For individual level regressions, standard errors are clustered at the occupation level.

department 40 years ago which might have had five accountants using pencils and paper ledgers. Today, that department might have one accountant with Microsoft Excel. While the computer has automated four accountants, it has also made one accountant five times as productive. Additionally, I have argued here that technology often allows traditional medium-skilled tasks to be performed by low-skilled workers with technology. To the extent that this technology-enabled deskilling happens within occupation codes (as in Figure 1), we would expect to see workers in routine occupations benefiting from technology use, and expect this effect to be particularly strong among low-wage workers. This is exactly the pattern presented in Table 2. I do not believe that these positive correlations invalidate my measures of complementarity, but rather that they underscore the importance of *partial* automation of routine tasks that is the focus of this research.

5.2 Technology effects of the minimum wage

Table 3 presents my main results, estimating the differential effect of the minimum wage on technology adoption among high complementarity industries. In column (1), complementarity is measured as the fraction of workers with less than a high education who use a computer, and $X_{CUS}^{<HS}$ denotes its interaction with the minimum wage.[46] As predicted, the

[46]See Table B6 of the Appendix to see each of the four measures.

Table 3: Technology effects of the minimum wage

DV: IT Labor	(1)	(2)	(3)	(4)	(5)	(6)
$X_{CUS}^{<HS}$	-0.634***		-0.524***	-1.440**		
	(0.205)		(0.194)	(0.619)		
$X_{DOT}^{w<\bar{w}}$		-0.755**	-0.480			-0.253
		(0.324)	(0.308)			(0.360)
$RT_{it}^{w<\bar{w}} \times MW_{st}$					-0.187**	-0.015
					(0.085)	(0.087)
$RT_{it}^{w<\bar{w}} \times X_{DOT}^{w<\bar{w}}$						-0.551**
						(0.251)
N	14152	14152	14152	14152	14152	14152
R^2	0.732	0.731	0.732	0.730	0.731	0.732
Fixed effects	it, st	it, st	it, st	it, st	it, st	it, st
Controls	Yes	Yes	Yes	Yes	Yes	Yes
IV				$X_{DOT}^{w<\bar{w}}$		
F Stat.				20.829		

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 109$) level, are shown in parentheses. "IT labor" refers to IT workers per 100 employees, measured at the industry-state-year level. X variables are the interaction between the state level minimum wage and an industry level measure of complementarity between low-skilled workers and technology. The superscript refers to the definition of low-skill ($w < \bar{w}$ indicates low-wage, $< HS$ indicates less than high school education) and the subscript refers to the measure of complementarity (described in Section 4.2). Controls include total employment and the minimum wage interacted with the employment share of low-skilled workers and technology complementarity measures among non-low-skilled workers. The reported F statistic corresponds to the excluded instrument in the first stage regression, presented in the Results Appendix.

coefficient is negative, suggesting that industries in which low-skilled workers are more complemented by technologies tend to use less technology when the minimum wage is higher. Column (2) confirms this finding, now interacting the minimum wage with complementarity among low-wage workers according to the DOT-based measure. Coefficient magnitudes can be approximately compared, since the complementarity measures have similar standard deviations (see Table 1). I provide a more detailed interpretations of coefficient magnitudes below.

Column (3) includes both measures. While both coefficients shrink somewhat, the fact that they remain at least half as large in magnitude suggests that both measures contain significant independent variation. My interest, however, is in the common variation. Column (4) instruments for $X_{CUS}^{<HS}$ using $X_{DOT}^{w<\bar{w}}$. The first stage of this regression is presented in Table B5 of the Results Appendix, but the corresponding F-statistic is over 20, suggesting that this is a highly relevant instrument. As discussed above, this instrumentation strategy reduces measurement error both in the measurement of complementarity and in the definition of low-skilled labor. This more than doubles the coefficient on $X^{<HS}_{CUS}$. While the standard error also increases, the coefficient remains different from zero at the 5% level of significance.

Because the identification strategy includes state-year and industry-year fixed effects, the estimates can only be interpreted as heterogeneous effects of the minimum wage. It is helpful to return to the summary statistics in Table 1 and compare a high complementarity industry (at the 75^{th} percentile of computer use among workers with less than high school education) to a low complementarity industry (at the 25^{th} percentile). In the high complementarity industry, 21.1% of low-skilled workers are technology complements, while in the low complementarity it is only 5.3%. A \$1 decrease in the minimum wage (slightly less than one standard deviation), then, is expected to increase technology use in the high complementarity industry by .227 IT workers per 100 employees relative to the low complementarity industry.[47] This effect is about 30% of the mean of IT labor, large yet plausible. It suggests that the minimum wage is an important driver of technology decisions for high complementarity industries.[48]

As discussed above, my measures of complementarity are highly correlated with existing measures of routineness. A natural question is how routine-intensive industries respond to the minimum wage. Column (5) includes the share of low-wage workers in routine occupations, interacted with the minimum wage.[49] Because the standard deviation of the low-wage routine share is slightly larger than that of the complementarity measures, we might expect the coefficient to be somewhat smaller. Yet the coefficient is less than a third the size of those in columns (1) and (2). When including $X^{w<\bar{w}}_{DOT}$ and their interaction in column (6) it is clear that the significance of the routine effect is driven by the highly complementary industries. The mechanism driving the results seems to be the differential response of industries where low-skilled workers benefit from technology, although Column (5) suggests that routine-intensive industries often do have such workers. At the same time, the effects of complementarity are clearly stronger when routine workers are present.

A number of robustness checks are presented in the Results Appendix (see Section B.4). I briefly describe them here, leaving more detailed interpretations in the appendix. Table B6 presents the effect of each of the four complementarity measures. All are negative and

[47] $.227 = (.211 - .053) \times 1.440$

[48] Note that because the mean of the CUS-based complementarity among less than high school educated workers happens to be the same as its interquartile range, the magnitudes calculated by comparing industries at the 25^{th} and 75^{th} percentiles are the same as the magnitudes calculated by comparing the average industry with what would prevail without any complementarity.

[49] Other measures of routine intensity (e.g., among all employees, among low-education employees) produce very similar results.

three are statistically significant, with $X_{CUS}^{w<\bar{w}}$ having the smallest coefficient, consistent with having more measurement error because of the small samples. Table B6 also presents IV results for two alternate instrumentation strategies. Instrumenting for $X_{CUS}^{w<\bar{w}}$ with $X_{DOT}^{<HS}$ yields a slightly smaller coefficient than that in column (4) of Table 3; instrumenting for $X_{DOT}^{w<\bar{w}}$ with $X_{CUS}^{<HS}$ yields a larger (more negative) coefficient. I prefer the instrumentation scheme used in Table 3 to this alternative because I find it easier to interpret the CUS-based measure (computer users) and because the first stage is stronger.

Table B7 of the Results Appendix presents analogies to Table 3's columns (1) and (4), the main results. It uses two alternative normalization of the IT labor variable. First, it uses employment with a high school education or more in the denominator, as these workers are less likely to be directly affected by the minimum wage. The results are nearly identical. Second, it normalizes IT labor by "synthetic employment," which is the product of industry employment nationwide with the state's share of total employment. This is the share of employment that would result if every industry's employment were identically distributed across states. Thus, the denominator normalizing IT labor is even less likely to be affected by the minimum wage. The estimated coefficients (and their standard errors) are larger (more negative), suggesting that employment effects, if anything, dampen the estimates. Finally, my main estimation sample is restricted to that for which at least 10 respondents are available for each of the four complementarity measures. This is meant to reduce noise in measuring complementarity. Table B8 expands the sample to include any industry-year for which at least one respondent is available for each complementarity measure (196 industries). As expected, the coefficients shrink towards zero and some lose significance. Weights based on the number of respondents available increases the coefficients and shrinks their standard errors, making them similar to those in Table 3.

5.3 Routine wage bill effects of the minimum wage

Next, I turn to the second prediction of the partial automation framework. If the technology effects estimated in Table 3 are relevant for the automation of medium-skilled tasks, then by slowing technology adoption a higher minimum wage should raise the wage bill share accruing to routine workers in high complementarity industries.[50] Because the DOT measure of complementarity is so highly correlated with the Autor-Dorn measure of routine (see Table 2), I prefer to rely only on the CUS-based measure of complementarity. Columns (1)-(3) of Table 4 replicate columns (1), (2), and (4) of Table 3, using IT labor as a dependent variable

[50] Again, since routine occupations tend to be in the middle of the wage distribution and spillover effects are small, there is no reason to suspect that the direct mechanical effect of the minimum wage would be to raise the routine share of the wage bill.

Table 4: Impact on routine share of wage bill

	(1)	(2)	(3)	(4)	(5)	(6)
$X_{CUS}^{<HS}$	-0.634***			0.026***		
	(0.205)			(0.009)		
$X_{CUS}^{w<\bar{w}}$		-0.087	-1.004***		0.010	0.042***
		(0.178)	(0.337)		(0.009)	(0.015)
DV	IT Labor	IT Labor	IT Labor	Rt. Share	Rt. Share	Rt. Share
N	14152	14152	14152	14152	14152	14152
R^2	0.732	0.731	0.728	0.913	0.913	0.912
Fixed effects	it, st	it, st	it, st	it, st	it, st	it, st
Controls	Yes	Yes	Yes	Yes	Yes	Yes
IV			$X_{CUS}^{<HS}$			$X_{CUS}^{<HS}$
F Stat.			89.275			89.275

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 109$) level, are shown in parentheses. Dependent variable in columns (1)-(3), "IT Labor," is IT workers per 100 employees (measured at the industry-state-year level) and in columns (4)-(6), "Rt. Share," is routine share of wage bill (the share of all labor income in the industry-state-year paid to workers in routine occupations). The reported F statistic corresponds to the excluded instrument in the first stage regression, presented in the Results Appendix.

and instrumenting for $X_{CUS}^{w<\bar{w}}$ using $X_{CUS}^{<HS}$. The results are very similar to those of Table 3. A similar calculation based on the differential effect of a $1 minimum wage decrease on an industry at the 75^{th} percentile of complementarity, compared to one at the 25^{th} percentile, suggests an increase in technology of .237 (31% of the IT labor mean).

Columns (4)-(6) of Table 4 then present the same specifications using the routine share of the wage bill as a dependent variable. The coefficients on both complementarity measures are positive, though only that on $X_{CUS}^{<HS}$ is statistically significant. Again, measurement error likely biases both coefficients, expecially that of $X_{CUS}^{w<\bar{w}}$ toward zero. The IV results in Column (6) increase the magnitude considerably. Now, a $1 decrease in the minimum wage leads to a .99 percentage point decrease in the share of wages paid to routine worker, about 3.3% of the mean. This result suggests that the technology adoption decisions driven by the minimum wage have important labor market ramifications.

5.4 Identification tests

My goal is to hold fixed production possibilities (by looking within narrowly defined industries) and to see how the choice of inputs is affected by the minimum wage. The key identification assumption is that there is no cross-state unobserved heterogeneity in the economic activity performed by firms within the same narrowly defined industry such that firms in high complementarity industries naturally use more technology in low minimum wage states. If industry-year fixed effects are not sufficiently narrow, then it is possible that my results are driven by differences in economic activity and heterogeneous production possibilities, which I falsely attribute to differences in the chosen input mix. Here, I describe two tests that support my identification assumption.

First, a stronger version of this identification assumption is that there is *no* heterogeneity within the narrowly defined industries used here. This assumption obviously does not hold, but to understand the magnitude of violations I turn to the NBER-CES Manufacturing Industry Database (MID). The number of manufacturing industries in the Census codes (74) is similar to the number of manufacturing industries in 4-digit NAICS codes (86), but much smaller than the number in detailed 6-digit NAICS codes (473). In Appendix Table B4, I consider four technology-relevant variables from the MID and show that only 10-20% of variation across the 473 6-digit NAICS industries occurs *within* 4-digit NAICS. Thus, the level of details used in my analyses (based on Census industry codes) is sufficient to capture the vast majority of variation in the technology-intensity of economic activity.

However, the assumption that there is no within-industry heterogeneity is sufficient but not necessary to justify my econometric approach. Within-industry heterogeneity will only drive my results if firms in low minimum wage states naturally use more technology, and differentially so for high complementarity industries. This cannot be directly tested. However, if the composition of economic activity in these industry-state-years were naturally more technology-intensive within 3-digit industries, one might expect the same pattern within 2-digit industries. Because I can observe state heterogeneity in 3-digit industrial composition within 2-digit industries, this latter assumption can be tested.

To measure state heterogeneity in the 3-digit industrial composition, I proceed as follows. First, I calculate the average IT labor share for each 3-digit industry-year (averaging over all states). Formally, let IT_{ist} and E_{ist} be IT labor share and employment, respectively, in state s in year t in industry i (with E_{it} being total industry employment). My 3-digit measure of average IT labor share is simply

$$\bar{IT}_{it} = \frac{E_{ist}}{E_{it}} \sum_{s=1}^{N_s} IT_{ist}$$

Just as (lower case) i denotes 3-digit industry, let (upper case) I denote 2-digit industry. I then measure industrial composition within 2-digit industries (I) by taking a weighted

Table 5: Effects of the minimum wage on industrial composition

DV: $\tilde{IT}_{Ist}$	(1)	(2)	(3)	(4)
$X^{<HS}_{CUS}$	0.323		0.137	0.987
	(0.275)		(0.184)	(0.665)
$X^{w<\bar{w}}_{DOT}$		0.715	0.616	
		(0.523)	(0.466)	
N	10995	10995	10995	10995
R^2	0.966	0.966	0.966	0.965
Fixed effects	It, st	It, st	It, st	It, st
Controls	Yes	Yes	Yes	Yes
F Stat.				36.709

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 68$) level, are shown in parentheses. See equation (4) in the text for the definition of $\tilde{IT}_{Ist}$. X variables are the interaction between the state level minimum wage and an industry level measure of complementarity between low-skilled workers and technology. The superscript refers to the definition of low-skill ($w < \bar{w}$ indicates low-wage, $< HS$ indicates less than high school education) and the subscript refers to the measure of complementarity (described in Section 4.2). Controls include total employment and the minimum wage interacted with the employment share of low-skilled workers and technology complementarity measures among non-low-skilled workers.

average of $\bar{IT}_{it}$ within 2-digit industry, with weights corresponding to employment shares:

$$\tilde{IT}_{Ist} = \sum_{i \in I} \frac{E_{ist}}{E_{Ist}} \bar{IT}_{it} \tag{4}$$

The resulting $\tilde{IT}_{Ist}$ is a measure of the technological-intensity of industrial composition, where technology use is fixed to mechanically shut off changes in the choice of input mix within 3-digit industries (because $\bar{IT}_{it}$ is used). Variation in $\tilde{IT}_{Ist}$ comes only from variation in economic activity in the form of observable differences in industrial composition.

In Table 5, I use this industrial composition measure in an adaptation of my primary specification:

$$\tilde{IT}_{Ist} = \alpha_{st} + \delta_{It} + Z'_{Ist}\gamma + \tilde{\beta}(Complementarity_{It} \times MinWage_{st}) + \varepsilon_{Ist} \tag{5}$$

If higher minimum wage states naturally had a less IT-intensive composition of production within 3-digit industries (and especially for high complementarity industries) then this would bias β (from equation (2)) to be negative. This sort of *unobserved* heterogeneity within 3-digit industries should be matched by *observed* heterogeneity within 2-digit industries: one would expect high minimum wage states to have a less IT-intensive composition of 3-digit industries (particularly for high complementarity 2-digit industries). In this case, $\tilde{\beta}$ should be negative.

Table 5, replicating my main results in columns 1-4 of Table 3, shows this is not the case. None of the estimated coefficients are significantly different from zero. All point estimates are positive. Four of the five estimates are smaller in absolute value than their counterpart in Table 3. Overall, there is no evidence that high complementarity 2-digit industries have less IT-intensive economic activity (measured by 3-digit industrial composition) in high minimum wage states. This casts doubt on an interpretation of my estimated $\hat{\beta}$ that attributes it to bias induced by differences in economic activity, and supports my interpretation that it captures responses in the choice of input mix.

6 Conclusions

I have argued here that many of the transformative technologies implicated in the automation of middle-wage jobs do not fully eliminate the labor requirement. They entail only partial automation, reducing the skills required of workers to perform particular tasks. As a result, the cost of adopting labor-replacing technologies is not only that of the equipment itself, but also of low-wage workers to use it.

Exploiting state variation in the minimum wage and industry variation in low-skilled/technology complementarity, I have estimated one channel through which the minimum wage affects labor-replacing technology. Because of this complementarity, I estimate that a \$1 decrease in the minimum wage increases technology use in the average industry by 30% and decreases the routine share of the wage bill by 1 percentage point.

My identification strategy is, obviously, imperfect. I cannot fully rule out the possibility that there is some unobserved state heterogeneity that biases my results. However, for this bias to drive the results above, an omitted variable that increases technology use must satisfy the following four criteria:

1. It must vary within narrowly defined industry, although Table B4 suggests that only 10-20% of the observable variation in technology measures occurs within these industries.
2. It must disproportionately affect high complementarity industries (or it would be absorbed by the state-year fixed effect).

3. This unobserved technology driving force must be *negatively* correlated with the minimum wage.
4. Any difference in the content of economic activity must be fundamentally different at low levels of industry aggregation than at higher ones (because Table 5 finds no evidence that observable industrial composition favors technology in high complementarity industries in low minimum wage states).

My results suggest that the minimum wage should be *part* of the ongoing conversation about the labor market implications of technology change.[51] Card and Dinardo (2006) argue that a shortcoming of the technology-inequality literature is that it "has turned up surprisingly few insights into appropriate policy responses." Here, I identify one such policy lever. In popular press, it is commonly argued that raising the minimum wage will help the middle class (Holzer, 2015; Krugman, 2015), though the mechanism for this effect is often vague and the existing literature finds that spillover effects are small (Autor, Manning, and Smith, 2015). I identify a particular channel for that effect and suggest that the spillover effects likely are not instantaneous. My results show that the choice of technology responds to differences in the effective minimum wage which may have long-term implications for labor demand and the wage structure.[52]

The pressures from ongoing technological advances are unlikely to relent anytime soon, and there is no "silver bullet" in a policymaker's arsenal. Like most policies, the minimum wage has diverse and sometimes unforeseen effects. Nonetheless, I have argued here for a new set of considerations in deciding the role of the minimum wage in labor policy.

[51]It is tempting to conduct a full welfare analysis based on the results developed here. I believe this is premature. First, even without technology adoption effects, the minimum wage has a number of complicated effects (Hirsch, Kaufman, and Zelenska, 2015). A complete welfare analysis would be sensitive to assumptions regarding effects as diverse as the response of public programs (Boadway and Cuff, 2001), the incidence of price effects (MaCurdy, 2015), the effects on employment volatility (Brochu and Green, 2013; Gittings and Schmutte, 2014), the pace of structural transformation (Acemoglu, 2001), and the rate employer-provided training (Acemoglu and Pischke, 1999). Even valuing seemingly straightforward effects such as hiring rates for young workers (Meer and West, 2013) depends on whether markets for entry-level workers are efficient (Pallais, 2014), whether learning from employment spells is public (Kahn, 2013; Schönberg, 2007), and the non-market effects of youth employment (Gelber, Isen, and Kessler, 2014; Leos-Urbel, 2014). Even if one were to focus solely on the minimum wage effects emphasized here, the welfare implications of slowing technology adoption depend on the effects of technology on GDP and TFP growth (Graetz and Michaels, 2015), how declining labor demand affects public program dependency (Autor et al., 2013, 2014), and how inequality relates to geographic sorting and local amenities (Diamond, 2015). A full welfare analysis would be a complex ordeal that would depend on critical assumptions, with little transparency or discipline in its execution.

[52]In this way, my emphasis is similar to Sorkin (2015), who shows that slow technology responses can lead to large differences between the short- and long-run effects of the minimum wage.

References

Acemoglu, D. (2001). Good jobs versus bad jobs. *Journal of Labor Economics 19*(1), 1–21.

Acemoglu, D. (2003). Cross-country inequality trends. *The Economic Journal 113*(485), F121–F149.

Acemoglu, D., P. Aghion, and G. L. Violante (2001). Deunionization, technical change and inequality. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 55, pp. 229–264. Elsevier.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics, Vol. 4, Part B*. Amsterdam: Elsevier.

Acemoglu, D. and J.-S. Pischke (1999). Minimum wages and on-the-job training. *NBER Working Paper*.

Allegretto, S. A., A. Dube, and M. Reich (2011). Do minimum wages really reduce teen employment? accounting for heterogeneity and selectivity in state panel data. *Industrial Relations: A Journal of Economy and Society 50*(2), 205–240.

Aten, B. H. and R. J. D'Souza (2008). Regional price parities comparing price level differences across geographic areas. *BEA Survey of Current Business Nov.*, 64–74.

Autor, D., D. Dorn, G. H. Hanson, J. Song, et al. (2014). Trade adjustment: Worker-level evidence. *The Quarterly Journal of Economics 129*(4), 1799–1860.

Autor, D. H. (2013). The "task approach" to labor markets: An overview. *Journal for Labour Market Research 46*(3), 185–199.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., D. Dorn, and G. Hanson (2015). Untangling trade and technology: Evidence from local labor markets. *The Economic Journal*.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *The American Economic Review 103*(6), 2121–2168.

Autor, D. H. and M. J. Handel (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics 31*(2 Part 2), S59–S96.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in US Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics 90*(2), 300–323.

Autor, D. H., L. F. Katz, and A. B. Krueger (1998). Computing Inequality: Have Computers Changed the Labor Market? *The Quarterly Journal of Economics 113*(4), 1169–1213.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Autor, D. H., A. Manning, and C. L. Smith (2010). The contribution of the minimum wage to u.s. wage inequality over three decades: A reassessment. *Centre for Economic Performance*.

Autor, D. H., A. Manning, and C. L. Smith (2015). The contribution of the minimum wage to u.s. wage inequality over three decades: A reassessment. *American Economic Journal: Applied Economics*, forthcoming.

Bartlesman, E. and W. B. Gray (1996). The nber manufacturing productivity database. National Bureau of Economic Research Technical Working Paper 205.

Baumgarten, D., I. Geishecker, and H. Görg (2013). Offshoring, tasks, and the skill-wage pattern. *European Economic Review 61*, 132–152.

Beaudry, P., M. Doms, and E. Lewis (2010). Should the personal computer be considered a technological revolution? evidence from us metropolitan areas. *Journal of Political Economy 118*(5), 988–1036.

Beaudry, P., D. A. Green, and B. M. Sand (2013). The great reversal in the demand for skill and cognitive tasks. National Bureau of Economic Research.

Becker, S. O., K. Ekholm, and M.-A. Muendler (2013). Offshoring and the onshore composition of tasks and skills. *Journal of International Economics 90*(1), 91–106.

Berman, E., J. Bound, and Z. Griliches (1994). Changes in the demand for skilled labor within us manufacturing: Evidence from the annual survey of manufactures. *The Quarterly Journal of Economics 109*(2), 367–397.

Bloom, N., M. Draca, and J. Van Reenen (2011). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *Review of Economic Studies*, forthcoming.

Boadway, R. and K. Cuff (2001). A minimum wage can be welfare-improving and employment-enhancing. *European Economic Review 45*(3), 553–576.

Böhm, M. J. (2015). The price of polarization: Estimating task prices under routine-biased technical change. *Working paper*.

Brochu, P. and D. A. Green (2013). The impact of minimum wages on labour market transitions. *The Economic Journal 123*(573), 1203–1235.

Card, D. and J. Dinardo (2006). The impact of technological change on low-wage workers: A review. In R. M. Blank, S. Danziger, and R. F. Schoeni (Eds.), *Working and poor: How economic and policy changes are affecting low-wage workers*, pp. 113–140. Russell Sage Foundation New York.

Caselli, F. (1999). Technological revolutions. *American economic review*, 78–102.

Caselli, F. and W. J. I. Coleman (2001). Cross-Country Technology Diffusion: The Case of Computers. *American Economic Review: Papers and Proceedings 91*(2).

Clemens, J. and M. Wither (2014). The minimum wage and the great recession: Evidence of effects on the employment and income trajectories of low-skilled workers. *NBER Working Paper*.

Cortes, G. M. (2014). Where have the middle-wage workers gone? a study of polarization using panel data. *Journal of Labor Economics*, forthcoming.

Diamond, R. (2015). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000.

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

Feenberg, D. and J. Roth (2007). CPS Labor Extracts. *National Bureau of Economic Research*.

Feng, A. and G. Graetz (2015). Rise of the machines: The effects of labor-saving innovations on jobs and wages. *Working Paper*.

Foote, C. L. and R. W. Ryan (2014). Labor-market polarization over the business cycle. In *NBER Macroeconomics Annual 2014, Volume 29*. University of Chicago Press.

Fortin, N. M. (2006). Higher-education policies and the college wage premium: Cross-state evidence from the 1990s. *The American Economic Review*, 959–987.

Gelber, A., A. Isen, and J. B. Kessler (2014). The effects of youth employment: Evidence from new york city summer youth employment program lotteries. *NBER Working Paper*.

Gittings, R. K. and I. M. Schmutte (2014). Getting handcuffs on an octopus: Minimum wages, employment, and turnover. *Working Paper*.

Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics 89*(1), 118–133.

Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *The American Economic Review 104*(8), 2509–2526.

Gottschalk, P., D. A. Green, and B. M. Sand (2015). Taking selection to task: Bounds on trends in occupational task prices for the u.s., 1984-2013. *Working Paper*.

Gould, E. D. (2015). Explaining the unexplained: Residual wage inequality, manufacturing decline, and low-skilled immigration. *Working Paper*.

Graetz, G. and G. Michaels (2015). Robots at work. *CEPR Discussion Paper No. DP10477*.

Handbury, J. and D. E. Weinstein (2014). Goods prices and availability in cities. *The Review of Economic Studies*.

Hirsch, B. T., B. E. Kaufman, and T. Zelenska (2015). Minimum wage channels of adjustment. *Industrial Relations: A Journal of Economy and Society 54*(2), 199–239.

Hirsch, B. T. and E. J. Schumacher (2004). Match bias in wage gap estimates due to earnings imputation. *Journal of Labor Economics 22*(3), 689–722.

Holzer, H. (2015). Interview. In R. Garver (Ed.), *Middle Class Jobs Aren't Gone; They've Just Moved*. The Fiscal Times (April 8, 2015).

InformationWeek (2006). Multiyear Comparisons Of InformationWeek 500 Survey Data Reveal IT Trends. *InformationWeek*.

Jaimovich, N. and H. E. Siu (2014). The trend is the cycle: Job polarization and jobless recoveries. *Working paper*.

Kahn, L. B. (2013). Asymmetric information between employers. *American Economic Journal: Applied Economics 5*(4), 165–205.

Krueger, A. B. (1993). How computers have changed the wage structure: Evidence from microdata, 1984-1989. *The Quarterly Journal of Economics 108*(1), 33–60.

Krugman, P. (2015). Knowledge isn't power. *New York Times Op-Ed February 23, 2015*.

Lee, D. S. (1999). Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage? *Quarterly Journal of Economics*, 977–1023.

Leos-Urbel, J. (2014). What is a summer job worth? the impact of summer youth employment on academic outcomes. *Journal of Policy Analysis and Management 33*(4), 891–911.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics 126*(2), 1029–1069.

MaCurdy, T. (2015). How effective is the minimum wage at supporting the poor? *Journal of Political Economy 123*(2), 497–545.

Malamud, O. and C. Pop-Eleches (2011). Home computer use and the development of human capital. *The Quarterly journal of economics 126*(2), 987–1027.

Meer, J. and J. West (2013). Effects of the minimum wage on employment dynamics. *Journal of Human Resources*, forthcoming.

Michaels, G., A. Natraj, and J. Van Reenen (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics 96*(1), 60–77.

Nordhaus, W. D. (2007). Two centuries of productivity growth in computing. *The Journal of Economic History 67*(01), 128–159.

Oldenski, L. (2012). The task composition of offshoring by us multinationals. *International Economics 131*, 5–21.

Oldenski, L. (2014). Offshoring and the polarization of the u.s. labor market. *Industrial & Labor Relations Review 67*, 734–761.

Ottaviano, G. I., G. Peri, and G. Wright (2013). Immigration, offshoring, and american jobs. *American Economic Review 103*(5), 1925–1959.

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review 104*(11), 3565–99.

Phelan, B. and D. Aaronson (2015). Wage shocks and technological substitution. Presentation at Allied Social Science Association 2015 Meeting.

Ruggles, S. J., T. Alexander, K. Genadek, R. Goeken, M. B. Schroeder, and M. Sobek (2010). Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]. Minneapolis: University of Minnesota.

Sabia, J. J. (2014). The effects of minimum wages over the business cycle. *Journal of Labor Research 35*(3), 227–245.

Schönberg, U. (2007). Testing for asymmetric employer learning. *Journal of Labor Economics 25*(4), 651–691.

Shim, M. and H.-S. Yang (2013). Business cycle properties of job polarization using consistent occupational data. *Working paper*.

Shim, M. and H.-S. Yang (2015). Interindustry wage differentials, technology adoption, and job polarization. *Working Paper*.

Smith, C. L. (2013). The dynamics of labor market polarization. *Working Paper, Board of Governors of the Federal Reserve System*.

Sorkin, I. (2015). Are there long-run effects of the minimum wage? *Review of economic dynamics 18*(2), 306–333.

Tambe, P. and L. M. Hitt (2012). The productivity of information technology investments: New evidence from IT labor data. *Information Systems Research 23*(3-part-1), 599–617.

U.S. Department of Labor (1972). Appendix I: Work Fields, DOT, Fourth Edition. *Handbook for Analyzing Jobs*, 85–87.

Van Reenen, J. (2011). Wage inequality, technology and trade: 21st century evidence. *Labour economics 18*(6), 730–741.

A Data Appendix

A.1 Basic data setup

Here I discuss the key data sources:

1. CPS ORG (Feenberg and Roth, 2007)
2. CPS CUS
3. Minimum wage data (Meer and West, 2013)
4. Census data (Ruggles et al., 2010)
5. DOT
6. BEA capital stock data

I also discuss issues in occupation and industry codes, and how various years were mapped to one another.

A.1.1 CPS ORG

The foundational data for the analyses above is the NBER extracts of the CPS ORG. I restrict the sample to those age 16-64 who are employed for pay in the private sector. For workers paid by the hour, I use the reported hourly wage. For other workers, I use usual weekly earnings divided by usual hours worked per week if available and hours worked last week otherwise. I never use allocated earnings variables. Wages are converted to real terms using the methods described in Section B.1.3. All calculations are based on the included earnings weights.

Because I am interested in low-wage workers, I do not drop or Winsorize low wages as is sometimes done with the CPS ORG. For my purposes, wages are primarily important only for determining whether workers are above or below the $9.50 latent wage cutoff.

To define workers with less than a high school education, prior to 1992, I use those whose highest grade attempted was less than 12 or who did not complete grade 12. For 1992 and afterward, "less than high school eductaion" also includes those who did highest grade completed was grade 12, but who do not have a diploma. It does not include those with a GED. This means that a GED holder who completed only 11^{th} grade will be "less than high school" before 1992, but not after. Conversely, someone who completed 12^{th} grade but without the credits necessary to receive a high school diploma will be "less than high school" after 1992, but not before. Since my identification strategy is entirely cross sectional, this is unlikely to be an issue.

A.1.2 CPS CUS

I use the Computer Use Supplements (CUS) for 1984, 1989, 1993, 1997, 2001, and 2003, obtained from ICPSR. The CUS has been conducted since 2003, but has not asked about computer use at work. I use the question "Do you directly use a computer at work?" See Krueger (1993) for a discussion of the CUS design. The specific question I use has not changed since.

To calculate computer use among workers with less than a high school education, I use full CUS sample that is employed in the private sector and the same education variables described above for the ORG. To calculate computer use among low-wage workers, I need to merge the CUS with the ORG, since wage data is not included in the CUS. Unfortunately, only one fourth of the CUS sample is expected to appear in the ORG, since the ORG is only conducted among respondents in their fourth or eighth wave of CPS interviews.[53]

To address this issue, I perform out-of-sample merging between the CUS and the ORG. I merge each CUS with the nearest ORG response if all individual identifiers *and* the industry of employment are unchanged. In other words, if an individual has their first interview in October 1989, I merge their October 1989 computer use response with their January 1990 wage response if they report the same industry in both surveys.[54] Because the merged wages and computer use responses are never more than three months apart, this is unlikely to induce much bias and it more than doubles the sample size.[55]

A.1.3 Other data

I use minimum wage data from Meer and West (2013), publicly available from Jonathan Meer's website. It is worth noting that Meer and West (2013) highlight several important limitations of previously used minimum wage datasets and take great care to provide a detailed monthly series. In aggregating to yearly data, I take a simple average of the minimum

[53]This issue is not addressed by Krueger (1993) because it is less important for the analyses therein. He is interested in whether wages are increasing in computer use after controlling for other observables. Thus, the full sample of wage earners is relevant and using only the ORG sample interviewed in October still leaves him with around 13,000 respondents per year. Sample size issues are more concerning for my analyses for two reasons. First, I focus specifically on usage by low-wage workers, which is obviously a minority of all wage earners. Second, I am interested in calculating industry-level averages for precise industries. This puts more strain on the data than documenting the relationship between wages and computer use, and so sample size constraints are more important.

[54]I only merge wave 4 wages with computer use from waves 1-4, and only merge wave 8 wages with computer use from waves 5-8. This is because there are eight months between wave 4 and wave 5. I never merge across this large gap.

[55]The CPS redesigns make it impossible to do some out-of-sample merging. For instance, respondents to the October, 1993 CUS were not administered the ORG until January, 1994, and the 1994 redesign changed the individual identifiers, making it impossible to merge October 1993 respondents with January 1994 respondents.

wage in effect during each month. Wages are converted to real terms using the methods described in Section B.1.3.

Census data, used in calculating IT labor and the change in routine employment share (for Figure 3 only), is from the Integrated Public Use Microdata Series (IPUMS) (Ruggles et al., 2010). I use the 5% PUMS for 1980, 1990, and 2000, and the 2005 American Community Survey (ACS). I define IT occupations as explained in the text and Section B.2.1 below, and calculate share share of respondents in the industry-state-year who are employed in an IT occupation, using person weights to calculate shares. Using the same weights, I also calculate the total share of wage and salary income reported by respondents holding a routine occupation.

I use the Dictionary of Occupational Titles obtained from Inter-University Consortium for Political and Social Research (ICPSR) as Study 7845 (originally published in 1977). For each occupation code (defined by 1970 Census codes), I calculate the fraction of respondents who report at least one of the three activities listed in the text (the DOT allows respondents to report up to four activities; I do not take into account the order in which activities are listed). I define this share of the DOT-based measure of complementarity. Note that this, as a share, is a cardinal and continuous measure. For Autor-Dorn occupation codes that have no matching 1970 Census code, I impute complementarity as an employment-weighted average of complementarity among other occupations within the same two-digit code (two-digit codes are available in the Appendix of Autor and Dorn (2013)).

To validate IT labor as a proxy, as well as for some calculations for Figure 3, I use the BEA's capital stock data, which includes PC capital (along with a chained price index to convert to real) for 56 industries from 1982 onward. Industry codes roughly correspond to two digit NAICS industries (with some more detail in, for instance, the manufacturing sector). To combine with labor data, I create a crosswalk between these codes and the Autor-Dorn-Hanson codes based on the ACS from 2003-2012 (also obtained from IPUMS), where many respondents are co-classified according to Census industry codes and NAICS. This works well because the industry codes in the BEA data are so coarse anyway, but would not work well for detailed NAICS codes.

A.1.4 Final years

Not all datasets are available for the same years. In particular, the CUS is not conducted in the same years as the census. Because IT labor is a small share of the total economy, and my identification strategy relies on relatively precise industry codes and separate technology measures for different states, the CPS is not a viable alternative to the Census (even the ACS is considerably noisier than the decennial Censuses). Table A1 shows how years were combined in the final analysis data.

Table A1: Years used for data construction

IT labor year	Minimum wage year	CUS years	DOT years
1980	1980	1984	1980-1984
1990	1990	1989, 1993	1989-1993
2000	2000	1997, 2001	1997-2001
2005	2005	2003	2002-2006

A.2 Sample selection

The Autor-Dorn-Hanson (ADH) industry codes include 205 private sector industries. With 51 states and four years, the potential sample size is 41,820. Here, I describe how I arrive at my main estimation sample. Table A2 presents the effects of each step of this selection process.

Not all industries are observed in all states. The set of industry-state-years with at least one Census respondent is 40,059. Based on the validation exercises above, I drop industries for which IT employment is a bad proxy for either IT spending or IT capital. These include (with ADH codes) telecommunications (440-442), computer and data processing services (732), and manufacturing of computers and related equipment (322) or machinery not specified (332). These are industries where minimum wage effects on IT labor likely do not reflect automation technology adoption, the underlying process in which I am interested.

Next, I drop industry-state-years with fewer than 30 Census respondents. One interpretation would be that these are very small industry-state-years and that any reasonable weighting scheme would *nearly* exclude them anyway. An alternative interpretation is that these are industry coding errors. As shown in Table A2, this is the most binding restriction.

Table A2: Effects of sample selection criteria

Criteria	Industry-state-years	Industries
Potential	41,820	205
Positive Census emp.	40,059	205
Drop ind. where IT labor is poor proxy	39,042	200
At least 30 Census respondents	26,894	200
Non-missing on all compl. measures	24,991	196
Too large industry-state-years	24,923	196
At least 10 complementarity respondents	14,152	109

It drops over 12,000 industry-state-years, though no industries are entirely dropped. This is not unreasonable. Industries are not uniformly distributed across the country. Many are only minimally present in certain states, particularly states with low population or population density. The fact that around two thirds of potential industry-state-year cells are filled with non-trivial employment is actually surprising.

Next, I restrict to the sample which is non-missing for each of the four complementarity measures described in the text. This implicitly drops industries with no employment of less than high school workers and no low-wage workers.

Next, I drop industry-state-years in which a single state accounts for more than 20% of national industry employment. The identification assumption is that industry level production processes do not respond to policy changes in a particular state (this is required for industry-state-years in the same industry with different minimum wages to be valid counterfactuals for one another). If a single state accounts for a large share of industry employment, this assumption may not be satisfied. This drops a trivial number of industry-state-years (68), which are mostly readily understandable examples (e.g., the motion picture industry in California).

This comprises the sample used for robustness in Table B8. My preferred estimation sample, however, makes one additional restriction. I restrict to the sample with at least 10 respondents available on which to base the four complementarity measures. This reduces the sample by 10,000 industry-state-years (though again, the main results are robust to a broader sample), mostly from 2005 and 1980, since there is only one wave of the CUS available for those years (see Section A.1.4).

This restriction has two benefits. First, it reduces measurement error in the complementarity variables, which are already fraught with it. Second, it disproportionately drops industries with few low-skilled workers. My identification strategy includes a number of controls (including state-year effects), which capture the ways in which different types of industries might experience heterogeneous state level forces that affect technology adoption (e.g., state education systems) and may be correlated with the minimum wage (the state-year fixed effect captures forces that are constant for all industries). The industries driving the results are those with non-trivial shares of low-skilled workers. Thus, if I do not allow enough flexibility in how I control for other cross-state forces, I might be concerned that the coefficients on my "control" variables do not adequately "control" for the way high-complementarity industries respond. Restricting to industries with non-trivial low-skilled labor shares implictly allows for every coefficient (and the state-year fixed effects) to take different values for industries with and without significant low-skilled labor (since this flexibility would lead identification of the coefficients of interest to be entirely driven by the industries with low-skilled labor), giving the specification additional flexibility and providing extra assurance that the estimated effects are causal.

In other words, the state-year fixed effect accounts for any force that affects all industries equally. From an identification standpoint, reducing the estimation sample is equivalent to allowing the state-year fixed effects to take different values for industries with and without low-skilled labor. It reduces sample size and precision, but with the benefit of making weaker identifying assumptions. The results suggest that the remaining precision is adequate. Table B8 shows that the results are robust to including the observations dropped in the last step.

A.3 DOT Work Fields

Table A3 presents the full set of potential work fields available in the DOT. For more information, see U.S. Department of Labor (1972), or Appendix I of ICPSR Study 7845, which offers the following summary "The Work Fields, as listed below, have been organized into groups on the basis of similar technology, or overall socioeconomic objective." Codes used in this paper are denoted in bold.

Table A3: DOT Work Fields

Work Field	Description
Hunting-Fishing Logging Cropping Mining-Quarrying-Earth Boring Blasting Gardening	Securing, producing, or cultivating raw materials, products, or animals (livestock or game) on or below the surface of the earth; usually outdoor work
Loading-Moving Hoisting-Conveying Transporting Pumping	Moving materials (in solid, liquid, or gaseous form) or people, by hand and/or machine power
Stationary Engineering	Producing and/or distributing heat, power, or conditioned air
Cleaning Ironing Lubricating Butchering	Cleaning and maintainence work
Filling Packing Wrapping	Packing materials or products for distribution or storage
Abrading Chipping Boring Shearing-Shaving Milling-Turning-Planting Sawing Machining	Working with machines and/or other handtools to cut or shape wood, metal, plastics, or other materials, or objects made from these materials. Can also involve assembly of objects.

Table A3: DOT Work Fields

Work Field	Description
Fitting-Placing Folding-Fastening Gluing	Assembling materials, usually light
Bolting-Screwing Nailing Riveting	Assembling part or materials, usually of metal, wood, or plastics, by means of screws, nails, or rivets
Welding Flame Cutting-Arc Cutting Soldering	Bonding or cutting materials by means of a gas flame, electric arc, combination welding process, or soldering
Masoning Laying Troweling Calking	Working with brick, cement, mortar, stone, or other building materials (other than wood) to build or repair structures, or to assemble structural parts.
Upholstering Structural Fabricating-Installing-Repairing Electrical Fabricating-Installing-Repairing Electronic Fabricating-Installing-Repairing Mechanical Fabircating-Installing-Repairing Electro-Mechanical Fabricating-Installing-Repairing	All-round fabricating, installing, and/or repairing of interior fittings; structures; and electrical, electronic, and mechanical units. Involves combinations of other Work Fields.
Melting Casting Heat conditioning Pressing-Forging Die sizing Molding	Compounding, melting, heat conditioning, and shaping metal and plastics, by any method in which heat is a factor
Baking-Drying Crushing Mixing Distilling Filtering-Straining-Separating Cooking-Food preparing Processing-Compounding	Processing various materials, in sold, fluid, semi-fluid, or gaseous states, during production process or to prepare for distribution
Immersing-Coating Saturating Brushing-Spraying Electroplanting	Coating or impregnating materials and products to impart decorative or protective finish or other specific quality, as described under separate Work Field sections.
Combing-Napping Spinning Winding Weaving Knitting	Processing fibers from thread to fabric
Sewing-Tailoring	Joining, mending, or fastening materials with needle and thread and fitting and adjusting parts
Eroding Etching Engraving	Cutting designs or letters into materials or products by sandblasting (eroding), applying acids (etching), or action of sharp pointed tools (engraving)
Printing Imprinting	Transferring letters or designs onto paper or other materials, by use of ink or pressure, includes setting type and preparing plates.
Photographing Developing-Printing	Taking pictures and developing and processing film
Appraising Weighing	Evaluating or estimating the quality, quantity, or value of things or data; ascertaining the weight of materials or objects

Table A3: DOT Work Fields

Work Field	Description
Stock checking	**Receiving, storing, issuing, shipping, requisitioning, and accounting for stores of materials**
Recording Accounting-Recording	**Preparing and maintaining verbal and/or numerical records**
Laying out Drafting Surveying Engineering	Plotting, tracing, or drawing diagrams and other directive graphic information for use in design or production; designing and constructing machinery; structures, or systems
Research	Controlled exploration of fundamental areas of knowledge, by means of critical and exhaustive investigation and experimentation
Writing Painting Composing Styling	Creative work
Investigating Litigating	Dealing with people to gather information to carry out business or legal procedures
System communicating Information-Giving	**Provides, or efects the transmission of, information to other persons, indirectly (by electrical or electronic media) or directly (by voice or written statement)**
Accomodating Merchandising Protecting Healing-Caring Administering Teaching Entertaining Undertaking	Dealing with people to provide services of various types

B Results Appendix

B.1 Descriptive analyses

B.1.1 Industry and occupation codes: Consistency and characteristics

Much of this analysis is based on occupation codes from Autor and Dorn (2013) and industry codes from Autor, Dorn, and Hanson (2013) both obtained from David Dorn's website.[56] Figure B1 shows job polarization using data on these occupations. Occupations are classified into quintiles by their 1990 average wage, and the figure displays the average change in employment shares and real wages from 1990-2005. During this period, the average occupation from the lowest or highest two quintiles increased its employment share by 4-5%,

[56] I thank David Dorn for his transparency and for making so many valuable resources publicly available.

while occupations in the 20^{th}-60^{th} percentiles shrank. Further, these shrinking occupations have seen relatively stagnant wage growth, contributing to the "hollowing out" of the wage distribution.

Figure B1: Labor market changes by 1990 occupational wage quintile

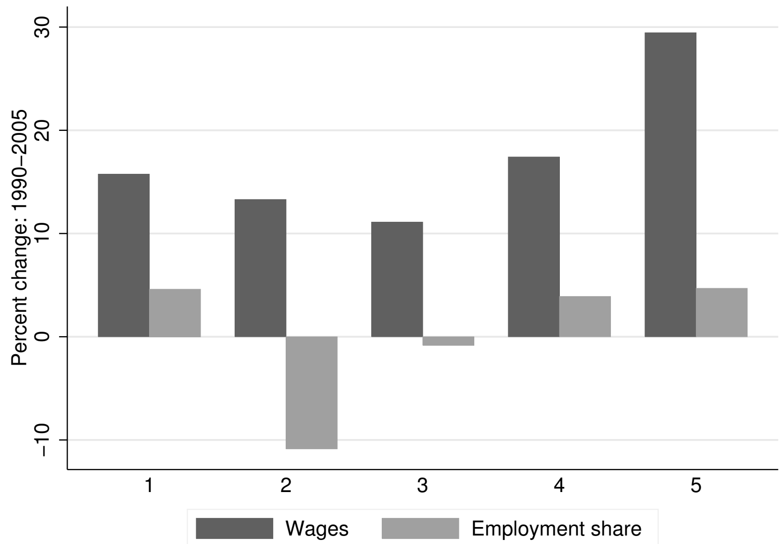

Source: Author's calculations based on data from Autor and Dorn (2013). Figure is based on 330 consistently defined occupations and wages and employment shares drawn from the 1990 Census and the 2005 American Community Survey. Quintiles are weighted by 1990 employment share so that each quintile represents approximately 20% of 1990 employment. Average wage and employment changes are also weighted by 1990 employment.

Figure B1 helps explain why recent research on wage inequality has decisively shifted attention towards market forces. Wage growth in the bottom quintile has far outpaced changes in the minimum wage, and while deunionization might slow wage growth in middle-wage occupations, it is not obvious how it could explain declining employment. The similar patterns in wage and employment changes in Figure B1 suggest the importance of labor demand and recent research has focused on the decline in demand for middle-wage workers.

Although the Autor-Dorn and Autor-Dorn-Hanson codes create a balanced panel, there are still inconsistencies when the census changes coding schemes.

One way to illustrate this problem is as follows. For each occupation or industry i, I calculate the change in estimated employment share from year $t-1$ to t ($EmpShare_{it} - EmpShare_{it-1}$). For each year, I then calculate the standard deviation of this one-year employment share change across all occupations or industries. This measure of the annual "dispersion" of employment growth is plotted in Figure B2.

The figure exhibits clear spikes, when the dispersion of employment growth increases by

Figure B2: Evidence of inconsistencies in occupation and industry codes

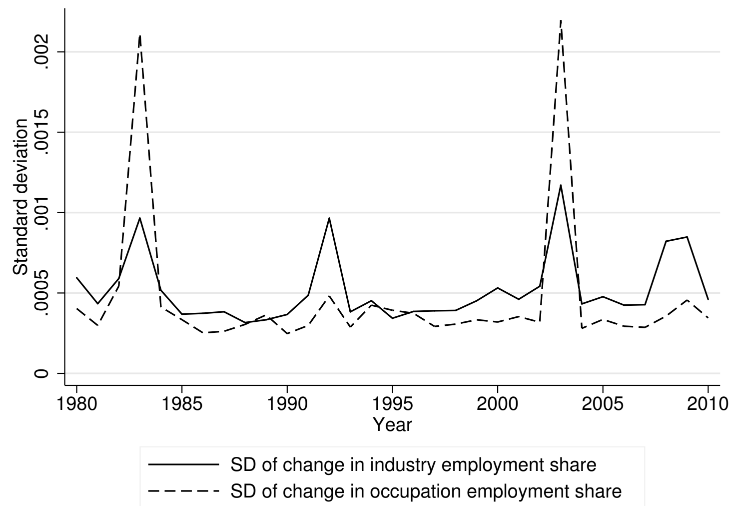

Source: Author's calculations based on CPS ORG. Figure plots the within year, cross-industry or cross-occupation standard deviation of $EmpShare_{it} - EmpShare_{it-1}$ based on Autor-Dorn occupation codes and Autor-Dorn-Hanson industry codes.

100-350% over its normal level, in 1983, 1992, and 2003 (all years during which the CPS changed its coding scheme). This implies that there are large breaks in the series; when the CPS codes are revised, many occupations and industries experience implausible shifts in their employment shares.

Shim and Yang (2013) and Smith (2013) have excellent discussions of the challenges in creating consistent occupational codes over time. Unfortunately, the solutions developed there (probabilistic crosswalks between different coding schemes) do not readily lend themselves to my application and the Autor-Dorn occupation and Autor-Dorn-Hanson industry codes remain the best available.

B.1.2 Computer use and wages

It is useful to consider how computer use relates to wages, and how this relationship has changed over time. The solid line of Figure B3 shows the use of computers across the wage distribution from 1984-2003 (along with the DOT complementarity measure as a dashed line). As emphasized by Krueger (1993), computer use clearly increases across the wage distribution. However, computer use among low-wage workers is not trivial, and it has increased appreciably over time. In 1984, only 5-10% of workers earning less than \$10 per hour (about the 25^{th} percentile of wages) used a computer. By 2003, this had increased to 20-30%.

Figure B3: Computer complementarity across the wage distribution

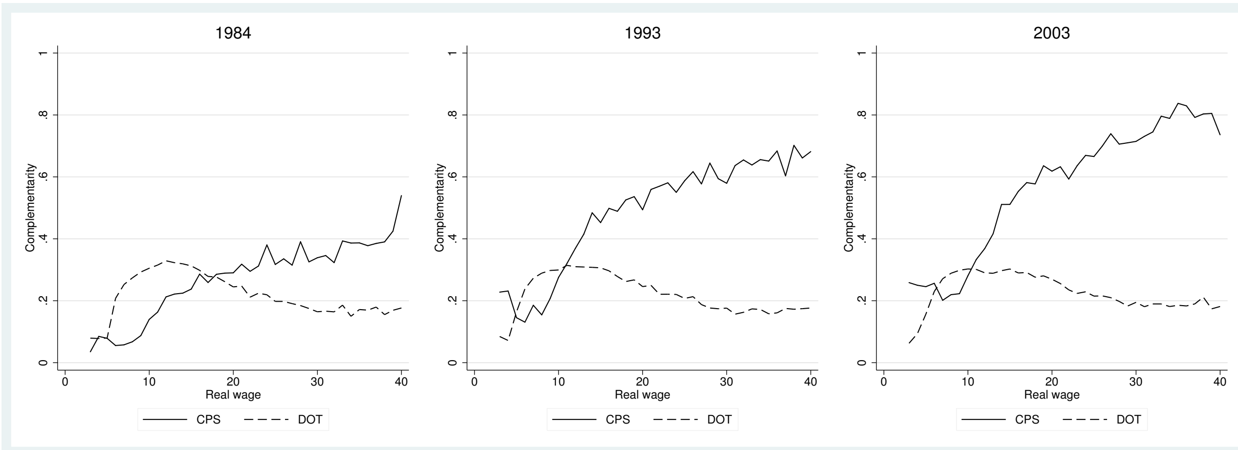

Source: Author's calculations based on CPS ORG. As described below, nominal minimum wages are converted to real using both the Current Population Index and cross-state price deflators. Figures used in the plot correspond to 2014 national dollars, rounded to integers for graphical purposes. See text for description of CPS and DOT measures of complementarity.

B.1.3 Real wages

The Meer and West (2013) minimum wage data and ORG wage data are both in nominal terms. I convert them to real terms in two steps. First, the BLS publishes regional consumer price indices (CPI's) for four regions. These indices capture differences in inflation across the regions, though they are not designed for comparison *across* regions. I use these CPI's to put all figures in 2014 dollars. In practice, the effects of using a regional CPI rather than the national are trivial. Regional variation in inflation is tiny.

The second step is far more consequential. I use BEA estimates of Regional Price Parities (RPPs), state level price indices that are explicitly desgined to compare prices across states (Aten and D'Souza, 2008). The RPPs are, essentially, only available for a single point in time (2005-2006). I use the 2005 state price index to convert all wages (the minimum wages and the wages paid to workers) to national dollars.

The effects of this approach are shown in Figure B4. Using both sets of price indices does slightly increase the standard deviation of the minimum wage, but much more important than its effect on the *level* of variation is its effect on the sources of variation. Two particular features of Figure B4 are useful to illustrate this point.

First, note the highlighted and labels state-years. The clearly separated line at the top of the figure is West Virginia. Like many states, West Virginia is often bound by the federal minimum wage. But prices are very low in West Virginia. The BEA estimates, for instance,

Figure B4: Real minimum wage with and without state level price adjustments

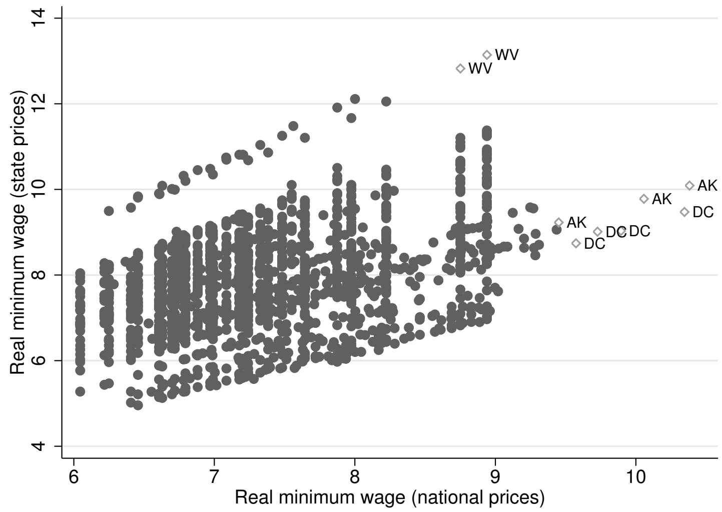

Figure plots the real minimum wage adjusted by state and time price indices (that used in the analysis) against the real minimum wage adjusted by the time price index only.

that the price level is more than a third lower than that of Virginia, another state often bound by the federal minimum wage. Using the national CPI alone makes it seem as though West Virginia has a quite typical minimum wage, but its low prices mean that the federal minimum wage actually exerts much more extreme pressure than elsewhere. Conversely, the District of Columbia and Alaska often have very high minimum wages. Yet these states are very expensive. Accounting for the different price levels adjusts the observed real minimum wage to be closely aligned with that of other states (they are outliers on the x-axis, not on the y-axis). Their high minimum wages are far less extreme. Thus, the advantage of the RPPs is not so much to increase overall variation in the real minimum wage, but to adjust that variation to more meaningfully reflect the underlying economic forces.

The second important advantage of the RPPs is shown by the series of perfectly vertical dots. These correspond to the federal minimum wage. Because my identification strategy is purely cross-sectional, using only the national CPI (or the minimally varying regional CPIs) to adjust for inflation would leave me with no usable identifying variation within the large set of states bound by the federal minimum wage. However, because these states have very different price levels, the RPP allows for meaningful variation for comparing across different states with the same nominal minimum wage at the same time. This is seen in the graph as the spread along the y-axis of points with the same value on the x-axis.

Because this paper focuses on firm decisions, state price indices are appropriate to use if and only if they represent the factor prices facing firms. To assess how reasonable this assumption is, I estimated the following regression using only workers in routine occupations in the 1980-2006 ORG:

$$\ln w_{it}^{n} = \alpha_{s(i)} + \gamma_{o(i)} + \delta_t + X_{it}'\beta + \varepsilon_{it} \tag{6}$$

where $\ln w_{it}^{n}$ is the log *nominal* wage of individual i in year t, X_{it} is a vector of individual characteristics,[57] the δ_t are year fixed effects accounting for inflation, the $\gamma_{o(i)}$ are occupation dummies to further reduce unobserved between individual heterogeneity, and the $\alpha_{s(i)}$ are state fixed effects.

Under the strong assumption that all cross-state heterogeneity is absorbed by $\gamma_{o(i)}$ and X_{it}, then the estimated $\alpha_{s(i)}$ series corresponds to the state level price indices specific to routine employment. This strong assumption is surely violated (e.g., education quality differs across states), yet it provides a useful $\alpha_{s(i)}$ series for a rough assessment of how well the BEA RPPs capture the relevant factor prices. Figure B5 plots these state fixed effects against the RPPs. An OLS regression yields an R^2 of .73; the series are clearly highly correlated. This suggests that the variation captured in the BEA RPPs is captures labor price variation that is meaningful for firm decisions.

Figure B5: Routine labor prices and state price indices

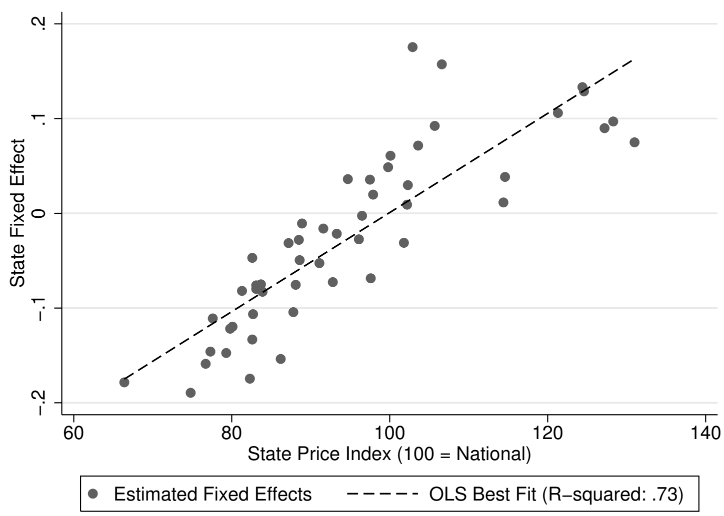

Figure plots the estimates state fixed effects (the $\alpha_{s(i)}$ series) from equation (6) against the BEA's state price indices (or RPPs) described above (Aten and D'Souza, 2008). In estimating equation (6), the sample is restricted to routine workers since, for the purposes of this paper, this is the most relevant factor price facing firms.

[57]These include years of education, sex, indicators for Black and White race, age, age squared, and marital status.

B.1.4 Latent wages

To estimate the complementarity between low-wage workers and technology, I first identify "low-wage" workers in an industry and then calculate either the DOT-based or the CUS-based complementarity measure among them. The most obvious way to do this is to define all workers with wage below some $\bar{w}$ as low-wage. As discussed in the text, however, since the minimum wage affects the wage distribution (potentially pushing some workers above $\bar{w}$) this approach would induce measurement error in the composition of low-wage workers in an industry that is a function of the minimum wage in the states in which the industry is located. To mitigate this issue, I use latent wages that (ideally) purge the wage distribution of the effects of the minimum wage.

I follow the approach developed in Autor, Manning, and Smith (2010) to estimate the latent wage. Five assumptions are important:

1. The latent wage distribution (that which would occur in the absence of a minimum wage) is log-normal, at least for wages below the 75^{th} percentile.[58]
2. The parameters (mean μ_{st} and standard deviation σ_{st}) of a particular state's wage distribution may vary over time, but only linearly.
3. The minimum wage has no effects on above-median wages.
4. The minimum wage has small, if any, employment effects.
5. The wage effects of the minimum wage decline as wages are further from the minimum.

The first two assumptions puts sufficient structure on the latent wage distribution that it could be estimated with a sample of unaffected workers. The third assumption (which finds empirical supported in the reduced form results in Autor, Manning, and Smith (2010, 2015)), combined with the first, states that workers from the 50^{th} to the 75^{th} can be used to estimate this distribution. And the last two assumptions imply that, given estimates of the parameters of the latent wage distribution, a worker's percentile in the observed wage distribution can be used to create a unique mapping to their latent wage.

Consider a set of workers indexed by j, observed in state s at time t. Let $p(j)$ be j's percentile in the state-year wage distribution and consider the sample such that $.5 \leq p(j) \leq .75$, which can be used to estimate:

$$\ln w_{jst} = \alpha_s + \delta_s t + \beta_s z_{p(j)st} + \gamma_s z_{p(j)st} t + \nu_{jst} \tag{7}$$

where $z_{p(j)st}$ is the z-score corresponding to $p(j)$ in the standard normal distribution.

The results of equation 7 can then be used to calculate the underlying parameters μ_{st}

[58] Autor et al. (2010) note that very high wages have much thicker tails than the normal distribution would imply.

and σ_{st} of the latent wage distribution because:

$$\mu_{st} = \alpha_s + \delta_s t \tag{8}$$

$$\sigma_{st} = \beta_s + \gamma_s t \tag{9}$$

Because they show that the minimum wage has no effect on above-median wages, they use workers who's wage is between the 50^{th} and the 75^{th} percentiles of the state-year wage distribution to estimate Equation 7. With the parameters of the latent wage distribution known, there is a one-to-one mapping of an individual's percentile in the state-year wage distribution (observed) to that individual's latent wage.

The intuition for the latent wage method is given by Figure B6 which plots the observed and latent wage distributions for Virginia and Washington. Because these states have similar observed wages between the 50^{th} and 75^{th} percentiles, they are estimated to have similar latent wages, shown in dashed lines. However, in 2006, the last year before the most recent wave of federal minimum wage increases, they had very different minimum wages, with Washington's set about 45% higher than Virginia's (the vertical lines). As such, Washington's observed wage distribution is much more compressed in the \$9-\$11 region, just above Washington's minimum wage (\$8.60). This compression, if ignored, would lead me to undercount low-wage workers in Washington relative to Virginia because more would be pushed just over the \$9.50 threshold.[59]

Further results are shown in Figure B7, which plots the "direct" and "total" effect of the minimum wage, on average, across different values of the latent wage.

The direct effect is the difference between the minimum wage in effect and the estimated latent wage (if this difference is positive), and the total effect is the difference between the observed wage and the latent wage. The figure shows the minimum wage spillover effects. For individuals with a latent wage as low as \$4, the minimum wage raises wages by an average of \$1.75, which is entirely due to the direct effects. However, as the latent wage increases to the \$7-\$9 range, spillovers (the total effect which is in excess of the direct effect) become more substantial. For an individual with an \$8 latent wage, the direct effects are only about 25 cents, but the total effect is three times this (and a 10% increase in the wage).

These effects are not huge, consistent with the rough calculations presented in the Theory Section, but they are not trivial either. Importantly, the slope of the wage percentile curve is steep in this region, suggesting that a significant mass of workers are near an \$8 hourly

[59] An alternative approach would be to define a certain percentile of the wage distribution to be low-skilled, say the bottom 20%. While less parametric, fixing the low-skilled employment share to be constant across states shuts off quality differences in state education, for instance, or any other factor that drives between-state skill differences.

Figure B6: Observed, latent, and minimum wages in two states

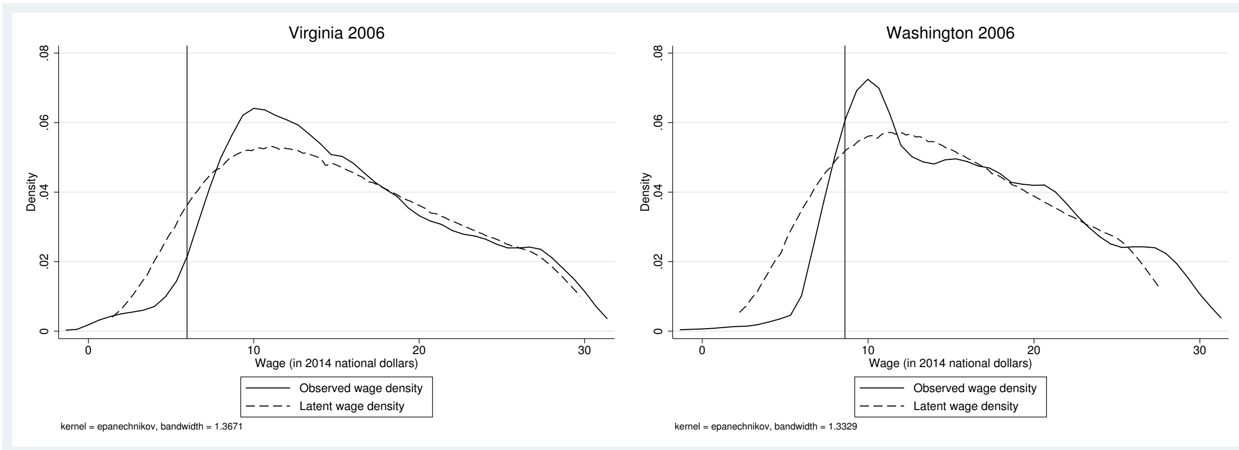

The figure shows observed real wages and estimated nominal wages for Virginia and Washington in 2006. The vertical bar shows the effective real minimum wage. In Virginia, this was $5.97 ($5.15 in nominal terms); in Washington it was $8.60 ($7.63 in nominal terms).

Figure B7: Minimum wage effects across latent wages

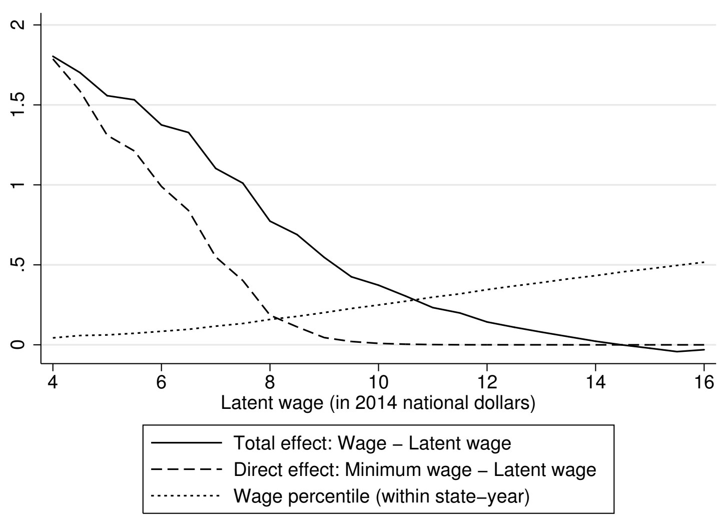

Figure displays the observed wage (and percentile) and minimum wage, both minus the estimated latent wage, across the latent wage. Latent wages are rounded to fifty cents for graphical purposes. All dollars correspond to national 2014 dollars (see Section B.1.3).

wage. The calculations underlying Figure B7 are the basis for my choice of \$9.50 as the cutoff latent wage to identify "low-wage" workers. Based on the average total effects, by latent wage with latent wages rounded to 10 cents (rather than 50, as used in the figure), the minimum wage effects at \$9.50 amount to about 5% of the latent wage, before continuous to decline monotonically.

B.2 Validation exercises and identification tests

B.2.1 IT labor as a proxy for technology

The identification strategy requires technology data at the industry-state level for precise industry codes. The precision of industry codes is important because the identification assumption that firms in the same industry but different states are only valid counterfactuals for each other will only be met if industries are defined narrowly enough. Much previous work has used Bureau of Economic Analysis (BEA) capital accounts data, but this data is not available at the state level. While the BEA publishes some industry-state level data, such as GDP and total compensation, none of the measures are ideal and the data is available only for coarse industry codes. Moreover, like the Harte-Hanks data use in some previous work examining geographic variation in technology adoption (Autor and Dorn, 2013; Beaudry et al., 2010), the BEA data is based on NAICS/SIC industry codes. [60] However, my measures of complementarity are based on labor data (either the CPS Computer Use Supplement or DOT variables for census occupation codes), where industries are coded using census industry codes. Thus, if technology is measured for NAICS/SIC industries, I would need a crosswalk to census industries, and these crosswalks require significant aggregation of industries. Using labor data to solves this problem.

I choose three occupations based on Autor-Dorn occupation codes to classify IT workers: computer scientists, computer software developers, and repairers of data processing equipment. Table B1 shows these occupations, their years of education, and their location in the state-year wage distribution. Even the least-educated, lowest-paid occupation (repairers) are well-education and well-paid. Median years of education is 14 and 95% are above the 28^{th} percentile of the wage distribution (recall that I estimate around 20% of workers to be significantly affected by the minimum wage). The other two occupations (which account for nearly 90% of IT workers) are even better educated and more highly paid. Thus, it is unlikely that employment in these occupations is "directly" affected by the minimum wage. Rather, changes in IT employment likely reflect changes in technology use.

[60] Caselli and Coleman (2001) use imports of computer technology to measure adoption. US import data is only available at the industry-state level since 2005, and is also based on NAICS/SIC codes.

Table B1: Wages and education of IT workers

	CS	Develop.	Repairer
Wage percentile			
5^{th}	.462	.410	.280
25^{th}	.753	.717	.559
50^{th}	.874	.850	.731
75^{th}	.939	.925	.852
Yrs. education			
5^{th}	12	12	12
25^{th}	14	14	12
50^{th}	16	16	14
75^{th}	18	18	15
Share	.505	.382	.113

As described in the text, survey results conducted by the technology magazine InformationWeek (2006) suggest that salaries and benefits comprise the largest share of IT budgets, more than twice the size of equipment. The article also reports, for 21 industries, the median respondents' IT budget as a share of revenue for 2002-2006. I use the ACS to calculate IT labor share for these broad industry aggregates. Figure B8 shows that reported IT budget shares are highly correlated with IT labor shares. Excluding outliers or accounting for year effects makes this correlation even stronger.

The biggest limitations of the IT budget data are that it is based on extremely coarse industry aggregates and that it is based on survey data without well-understood representativeness of businesses (though as mentioned in the text, it is certainly a sample of interest).

To augment this exercise, I combine industry level IT labor shares from the CPS with the BEA's data on PC capital stock (in real terms) for 56 industries from 1982 to 2006.[61] Initial inspection of the data showed that there were three industries in which the correlation was particularly and unacceptably bad. Both "Computer systems design and related services" and "Information and data processing services" were clear outliers in IT labor share, unsurprisingly. "Computer and electronics manufacturing" was a clear outlier in IT capital. These industries are excluded from both this validation exercise and from the main estimation sample (see Section A.2.

Table B2 presents the correlation with PC capital per worker and IT labor share. Column (1) shows that the two are highly correlated. IT labor share explains 21% of the variation in PC's per worker. Including year effects in column (2), the R^2 rises and the coefficient on IT labor share remains large and statistically significant ($t = 2.93$). Because the identification

[61]Results for broader definitions of IT capital are similar. In recent decades, the PC's share of total IT capital has grown to 70%.

Figure B8: IT budgets and IT labor

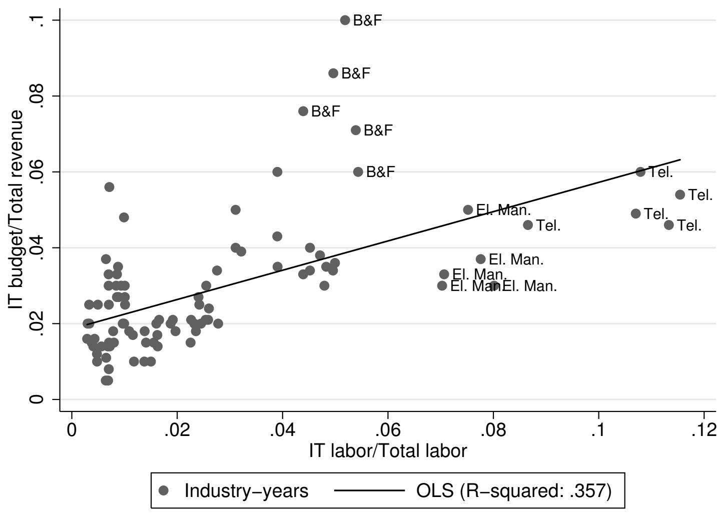

IT labor from American Community Survey (ACS). IT budgets from *InformationWeek*. Figure based on 19 broad industries from 2002-2006. I exclude two industries reported in the original *InformationWeek* data: "Consumer Goods" (which has no clear analog in either the Census or NAICS industry codes in the ACS) and "Information Technology" (a clear outlier with IT labor share of over 50%). Abbreviations "B&F," "El. Man.," and "Tel." stand for Banking and Financial Services, Electronics Manufacturing, and Telecommunications, respectively. Excluding Electronics manufacturing and Telecommunications raises the R^2 to .453. These outlier industries are excluded from the main estimation sample; see Section A.2 for a discussion.

strategy exploits cross-sectional variation only, this is the correlation of interest. Nonetheless, in the interests of completeness, column (3) includes industry effects and column (4) includes both industry and year effects. It is encouraging that the correlation in column (4) remains so strong ($t = 2.41$), suggesting that IT labor and computer capital increase together as industries become more technology intensive. Of course, as discussed above, there are many reasons to think that IT labor is a better measure than computer capital of technology upgrading, in its own right. Nonetheless, these validations are encouraging.

Panel A of Table B2, based on the levels of PC's per employment and IT labor share is the appropriate validation exercise because the actual analysis uses IT labor share in levels and because the PC capital has numerous zeros (IT labor also has a few), which are informative about the extensive margin of technology adoption. Nonetheless, the magnitudes are difficult to interpret. Thus, Panel B presents the same specifications based on the logs of IT labor and PC capital.

The coefficients in Panel B correspond to the elasticity of computer capital with respect to IT labor, evaluated at the intensive margin ($PC_{it}^{k} > 0$). Column (1) shows this elasticity is statistically indistinguishable from 1. This estimate is based on both cross-industry variation

Table B2: Correlation between IT labor share and PC's per worker

	(1)	(2)	(3)	(4)
DV :	PC's per 100,000 employment			
IT labor share	0.583***	0.407***	1.007***	0.537**
	(0.143)	(0.139)	(0.215)	(0.223)
N	1431	1431	1431	1431
R^2	0.213	0.468	0.409	0.608
DV :	ln(PC's per 100,000 employment)			
ln(IT labor share)	1.094***	0.631***	2.120***	-0.007
	(0.094)	(0.074)	(0.222)	(0.044)
N	1191	1191	1191	1191
R^2	0.238	0.880	0.384	0.965
Year FE		Yes		Yes
Ind. FE			Yes	Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Standard errors clustered at the industry level ($n = 53$) are shown in parentheses.

and general time trends. Column (2) includes year effects. While the elasticity falls (to .631) it remains large and highly significant ($t = 8.53$). This demonstrates IT labor is a useful proxy to understand the cross-sectional variation in technology adoption. Figure B9 presents this specification in graphical form, plotting residual log PC's per worker against residual log IT labor share (both net of year effects). The correlation is not driven by outliers.[62] My empirical strategy assumes that this high cross-industry correlation implies an equally high cross-state correlation. Columns (3) includes industry effects. The results are consistent with Panel A. Column (4) includes industry and year effects. The coefficient on log IT labor share drops to zero. The discrepancy between Panels A and B suggests that much of the information content in IT labor is about the extensive margin of computer capital purchases. I reiterate that the identification strategy is purely cross-sectional, since it is based on industry-year and state-year fixed effects, so column (4) should not cause concern.

The cross-industry correlation shown in Table B2 and Figure B9 is necessary but not sufficient to validate my use of IT labor as a proxy for technology because it does not establish whether IT workers are located where the PC capital is. Consider a multi-establishment firm that operates in multiple states. It may make technology use decisions for each establishment based on the factor prices in that state, and it may acquire PC capital for the establishments where it chooses more technology-intensive production processes, but it the IT workers to support those processes may be entirely concentrated in the firm's headquarters, rather than

[62] I choose this plot because the log transformation is useful in separating out the small numbers, and because the clustering of multiple industry-years at zero is graphically unhelpful.

Figure B9: Cross-industry variation in PC's and IT labor

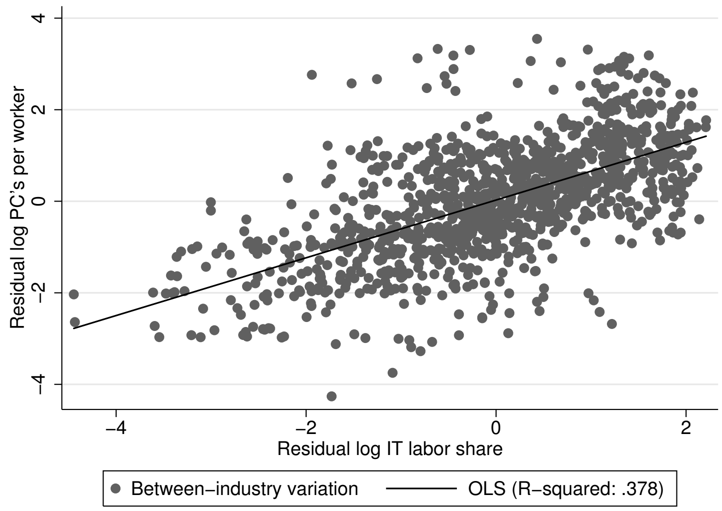

Figure displays residual log computer capital and IT labor, both net of year effects. These results correspond to column (2) of Panel B of Table B2.

the technology-intensive establishments' states. There may be perfect correlation between IT labor and PC capital at the industry level, but there may be no relationship in the geographic variation, making it inappropriate to use IT labor to study geographic heterogeneity in technology use.

To test this hypothesis, I use data from Autor and Dorn (2013). They use Harte-Hanks survey data which measures PC's per worker for many establishments at the Commuting Zone (CZ) level. Following Beaudry et al. (2010), who use the same data at the MSA level, they first regress PC's per worker on a full set of dummies for industry by year by establishment size. They then aggregate this adjusted measure of PC's per worker over establishments within a CZ. This is a purely geographic measure of PC intensity, adjusted for variation in industrial composition.

I first take a population weighted mean of CZs within each state to produce an adjusted state level measure of PC capital. I then create an adjusted state level measure of IT labor by regressing my industry-state-year measure of IT labor share on industry-year fixed effects, calculating residual IT labor share, and taking an employment weighted average over all industries within a state-year. Like the Autor-Dorn PC capital measure (aggregated to states), this is a purely geographic measure of IT labor share, adjusted for variation in industrial composition.

Figure B10 shows that these state level measures are highly correlated. It plots both series for the 48 states and 2 years (1990 and 2000) with non-missing PC measures in the

Table B3: Correlation between IT labor share and PC's per worker

DV: Residual IT labor	(1)	(2)	(3)	(4)	(5)	(6)
Residual PC's per worker	4.418***	4.081***	4.142***	3.818***	4.753***	4.348***
	(0.722)	(1.230)	(0.886)	(0.843)	(1.373)	(1.004)
N	48	48	96	48	48	96
R^2	0.360	0.206	0.237	0.363	0.305	0.303
Weights	None	None	None	Pop.	Pop.	Pop.
Years	1990	2000	1990, 2000	1990	2000	1990, 2000

* $p < .10$, ** $p < .05$, *** $p < .01$. Robust standard errors shown in parentheses.

Autor-Dorn data. Within both years, there is a strong positive correlation. This is formalized in the regressions displayed in Table B3.

Figure B10: Cross-state variation in PC's and IT labor

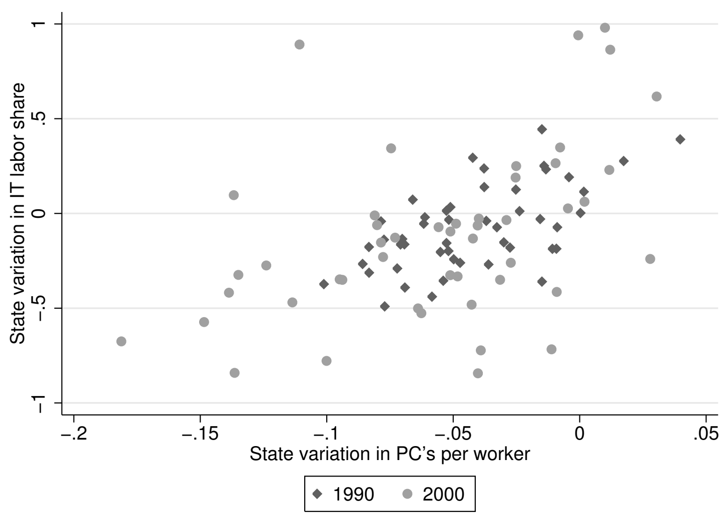

Figure displays residual PC's per worker and IT labor share, adjusted for industry-year fixed effects, across 48 states for 1990 and 2000. PC data from Autor and Dorn (2013).

Importantly, the units of the measures are not easy to interpret because both correspond to averages of residuals (after removing industry fixed effects). However, the coefficients are quite consistent across the two years and significantly and robustly positive, particularly after weighting by state population. Figure B10 and Table B3, along with the industry level exercises above, confirm that IT labor is an excellent proxy for technology use.

B.2.2 The degree of within-industry variation

Here I provide some support for my identification assumption. The identification assumption is that the Autor-Dorn-Hanson (ADH) industry codes based on Census industry codes are sufficiently precise that firms within the same industry are valid counterfactuals for one another (i.e., engaged in sufficiently similar economic activity and facing the same production possibilities frontier). This assumption is inherently difficult to test.

To gain some traction, I turn to the NBER-CES Manufacturing Industry Database (MID), based on the Annual Survey of Manufacturing (Bartlesman and Gray, 1996). I use the MID based on NAICS codes. The 6-digit codes in the MID are very precise, dividing manufacturing into 473 industries (462 of which are available before the NAICS system was adopted in 1997). These are much more precise than the ADH codes which have only 74 manufacturing industries. However, the 4-digit NAICS are very similar, with 86 manufacturing industries. One way to test how much variation exists *within* the ADH codes is to test how much of the 6-digit variation (across 473 industries) occurs within 4-digit NAICS codes (only 86 coarse industries).

To determine this, I estimate the following regressions:

$$\ln Y_{it} = \alpha_{It} + \varepsilon_{it} \tag{10}$$

where Y is an industry characteristics relevant for technology and wage inequality (discussed below), t denotes year (1980-2005, like my main sample), i denotes 6-digit NAICS code, and I denotes 4-digit NAICS code. Table B4 shows the R^2 from these regressions for four technology-relevant industry characteristics: Production workers share of employment (Berman et al., 1994), Average annual pay per worker,[63] Value added per worker, and Capital-labor ratio.

Panel A of Table B4 show that relatively little variation across 6-digit NAICS industries occurs within 4-digit industries. For production share of employment, for instance, only 17% of variation occurs within 4-digit NAICS ($.172 = 1 - .828$). This supports the claim that ADH Census-based industry codes (which are nearly as precise as 4-digit NAICS codes) are sufficiently narrow that firms within them serve as valid counterfactuals for one another. The variation in economic activity within these codes (at least that which is observable at the 6-digit NAICS level) is relatively small. Across the 4 technology/wage inequality relevant measures, between 80% and 90% of 6-digit variation is explained by the 4-digit codes.

[63]Hours worked is only available for production workers. Thus, it is not possible to calculate average hourly wage of workers. Results are similar when based on hour wages under the assumption that all non-production workers work 40 hours per week.

Table B4: Variation across 6-digit NAICS within 4-digit NIACS

DV (logs):	Production share	Average wage	Value added per worker	Capital-labor ratio
	Panel A: Annual			
R^2	0.828	0.897	0.810	0.830
N	12,111	12,111	12,111	12,111
N_{ind}	473	473	473	473
	Panel B: 25-year changes			
R^2	0.599	0.657	0.506	0.780
N	462	462	462	462

Panel A displays R^2 resulting from regressing 6-digit technology measures for a panel of 473 industries over 26 years (1980-2005) on 86 fixed effects for 4-digit industries by year (see equation (10) in the text). With the exception of 9 industries unavailable before 1997, the sample is a balanced panel. Panel B displays R^2 from regressing 25-year changes for 462 6-digit industries on 86 fixed effects for 4-digit industries (see equation (11) in the text). All regressions weight by employment.

For understanding technology adoption, it is also important to understand how well the codes explain *changes* in these measures during the sample period. I calculate 25-year changes from 1980 to 2005 (after smoothing over consecutive years to reduce measurement error) as follows:

$$\Delta \ln Y_i \equiv \frac{\ln Y_{i,2004} + \ln Y_{i,2005} + \ln Y_{i,2006}}{3} - \frac{\ln Y_{i,1979} + \ln Y_{i,1980} + \ln Y_{i,1981}}{3}$$

and then estimate:

$$\Delta \ln Y_i = \alpha_I + \varepsilon_i \tag{11}$$

where all notation follows that used in equation (10). Note that using differences, rather than levels, is a more difficult test.

While the R^2 values are lower, a sizeable share of the change is still explained by the 4-digit industry codes. Less than half of the variation in long-run changes (and as low as 22% in the case of the capital-labor ratio) occurs within these coarse industry codes. This supports that the identification assumption is not unreasonable.

B.3 First stages

Table B5 presents the first stages for all instrumental variables regressions contained in the paper. Odd columns exclude controls; even columns include them. All columns include fixed effects. Columns 1 and 2 instrument for $X_{CUS}^{<HS}$ using $X_{DOT}^{w<\bar{w}}$ (my preferred specification). Columns 3 and 4 instrument for $X_{CUS}^{w<\bar{w}}$ using $X_{DOT}^{<HS}$ (IV results shown in Column 5 of Table B6). Column 5 and 6 instrument for $X_{CUS}^{w<\bar{w}}$ using $X_{CUS}^{<HS}$ (the specification used in Table 4 of the main paper). All instruments are relevant ($p < .01$). While F-statistics are not shown in Table B5, they are all over 10 and are shown in every table presenting IV results.

Table B5: First stage regressions

	(1)	(2)	(3)	(4)	(5)	(6)
$X_{DOT}^{w<\bar{w}}$	0.436***	0.524***				
	(0.069)	(0.115)				
$X_{DOT}^{<HS}$			0.551***	0.474***		
			(0.108)	(0.106)		
$X_{CUS}^{<HS}$					1.084***	0.631***
					(0.053)	(0.067)
N	14152	14152	14152	14152	14152	14152
R^2	0.991	0.995	0.990	0.996	0.995	0.997
Fixed effects	it, st	it, st	it, st	it, st	it, st	it, st
Controls	No	Yes	No	Yes	No	Yes
DV	$X_{CUS}^{<HS}$		$X_{CUS}^{w<\bar{w}}$		$X_{CUS}^{w<\bar{w}}$	

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 109$) level, are shown in parentheses.

B.4 Robustness checks

Here, I discuss a series of robustness checks and supporting regressions that complement the results presented in Section 5 of the text. First, Table B6 presents my main specification, separately estimated with each of the four potential complementarity measures (columns 1-4). The results are similar for each, with the exception of $X_{CUS}^{w<\bar{w}}$. This is unsurprising, since this measure is based on, by far, the smallest samples and therefore likely has the most measurement error. Instrumenting with the other measures to account for this measurement error brings the coefficient magnitude in line with the others. This can be seen using $X_{CUS}^{<HS}$ as an instrument in column 3 of Table 4 in the text, or using $X_{DOT}^{<HS}$ as an instrument in column 5 of Table B6 (though here it is not statistically significant).

Table B6: Minimum wage effects for all complementarity measures

	(1)	(2)	(3)	(4)	(5)	(6)
$X_{CUS}^{w<\bar{w}}$	-0.087				-1.031	
	(0.178)				(0.626)	
$X_{CUS}^{<HS}$		-0.634***				
		(0.205)				
$X_{DOT}^{w<\bar{w}}$			-0.755**			-2.777***
			(0.324)			(0.982)
$X_{DOT}^{<HS}$				-0.488*		
				(0.283)		
N	14152	14152	14152	14152	14152	14152
R^2	0.731	0.732	0.731	0.731	0.728	0.727
Fixed effects	it, st	it, st	it, st	it, st	it, st	it, st
Controls	Yes	Yes	Yes	Yes	Yes	Yes
IV					$X_{DOT}^{<HS}$	$X_{CUS}^{<HS}$
F Stat.					20.124	15.856

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 109$) level, are shown in parentheses.

Column 6 of Table B6 uses $X_{CUS}^{<HS}$ to instrument for $X_{DOT}^{w<\bar{w}}$, the converse of my preferred instrumentation strategy. The coefficient is larger (more negative) than my preferred specification, but similar. I prefer my instrumentation strategy because I find the CUS complementarity measure easier to interpret and because the CUS is more subject to measurement error, making it a weaker instrument in the first stage.

Table B7 presents robustness checks which replicate the main effects, the heterogeneous effect of the minimum wage on technology adoption, given in columns 1 and 4 of Table 3 in the text. Columns 1 and 3 give OLS estimates; columns 2 and 4 give IV estimates (all with the same specification as in Table 3). The difference is in the normalization of IT labor. To reiterate the issue at hand, in the text IT labor is normalized by total employment in the industry-state-year. One might be concerned that the measured effects of the minimum wage on IT labor share are operating through employment effects, rather than through technology adoption (and the associated increase of IT employment).

Columns 1 and 2 normalize IT labor by employment with a high school education or more, since these workers are less likely to be directly affected by the minimum wage. The coefficients are nearly identical, suggesting that the estimated effects are driven by changes in the numerator (IT labor) rather than the denominator.

Columns 3 and 4 take an even stronger view. For each industry-state-year, I calculate

Table B7: Robustness

DV: IT labor	(1)	(2)	(3)	(4)
$X_{CUS}^{<HS}$	-0.637***	-1.407**	-1.116***	-2.140**
	(0.218)	(0.662)	(0.298)	(1.038)
N	14152	14152	14152	14152
R^2	0.700	0.699	0.539	0.537
Fixed effects	it, st	it, st	it, st	it, st
Controls	Yes	Yes	Yes	Yes
DV mean	0.855	0.855	0.829	0.829
IV		$X_{DOT}^{w<\bar{w}}$		$X_{DOT}^{w<\bar{w}}$
F Stat.		20.829		20.829
Normalization	Emp. with ≥HS ed.		Synthetic emp.	

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 109$) level, are shown in parentheses. The reported F statistic corresponds to the excluded instrument in the first stage regression, presented in Table B5 of the Results Appendix.

potential or "synthetic" employment as follows. Let Emp_{ist} be employment in industry i in state s in year t, and let $Emp_{st} = \sum_i Emp_{ist}$ be total state employment (across all industries) and $Emp_{it} = \sum_s Emp_{ist}$ be total industry employment (across all states). I define synthetic employment as:

$$\tilde{Emp}_{ist} \equiv \frac{Emp_{st}}{\sum_s Emp_{st}} Emp_{it}$$

In other words, synthetic employment allocates to each state the share of each industry's employment that corresponds to the state's share of national employment. It assumes that each industry is identically distributed across states. The advantage of this measure is that the denominator is even more likely to be exogenous to minimum wage effects, since it abstracts from the employment responses of particular industries in particular states. The disadvantage is that industries are *not* identically distributed across states. Constructing a denominator as if they were adds lots of noise to the dependent variable, and the estimates become much less precise.

Columns 3 and 4 present the results based on IT labor normalized by synthetic employment. The standard errors increase by 50-70% over their levels in Table 3 of the text. Nonetheless, the coefficients remain negative and statistically different from zero. In fact, the coefficients increase (become more negative), though they are not statistically distinguishable from the point estimates in Table 3.

As a last robustness check of my main specification, I show that my results are robust to the sample selection criteria Described in Section A.2. As noted there, I exclude all industry-years for which less than 10 respondents were available to calculate the complementarity measures. Table B8 shows the main results (columns 1, 2, and 4 of Table 3) estimated on the full sample for which at least one respondent was available to calculate the complementarity measure (the sample size effect of requiring at least one respondent is small; see Table A2 of the Data Appendix). The coefficients are negative, but only about two thirds as large as those in Table 3, consistent with expectations that measurement error biases coefficient estimates towards zero. The standard errors also shrink, but not by as much as the coefficients, so that most estimates become statistically insignificant.

Table B8: Full sample effects with and without weights

	(1)	(2)	(3)	(4)	(5)	(6)
$X_{CUS}^{<HS}$	-0.404***		-0.766	-0.636***		-0.787**
	(0.143)		(0.530)	(0.184)		(0.374)
$X_{DOT}^{w<\bar{w}}$		-0.323			-0.376*	
		(0.225)			(0.205)	
N	24923	24923	24923	24923	24923	24923
R^2	0.738	0.737	0.737	0.700	0.700	0.700
Fixed effects	it, st	it, st	it, st	it, st	it, st	it, st
Controls	Yes	Yes	Yes	Yes	Yes	Yes
DV mean	1.066	1.066	1.066	1.066	1.066	1.066
Weights	No	No	No	Yes	Yes	Yes
IV			$X_{DOT}^{w<\bar{w}}$			$X_{DOT}^{w<\bar{w}}$
F Stat.			18.294			36.646

* $p < .10$, ** $p < .05$, *** $p < .01$. Two-way clustered standard errors, at the state ($n = 51$) and industry ($n = 196$) level, are shown in parentheses. Weights are based on the number of respondents available for calculating the complementarity measures. See equation (12).

A less knife-edge way of dealing with small samples is to construct weights based on the number of respondents available for calculating the two primary complementarity measures. I use the following weights:

$$w = \left(\frac{1}{N_{CUS}^{<HS}} + \frac{1}{N_{DOT}^{w<\bar{w}}} \right)^{-1/2} \tag{12}$$

These would be the efficient GLS weights if $X_{CUS}^{<HS}$ and $X_{DOT}^{w<\bar{w}}$ were the dependent variables, although standard justifications for weighting do not include small samples for the *independent* variables (which, in particular, do not solve measurement error). This approach gives very little weight to industries with few low-skilled workers available for complementarity. The smallest value the weights can take is in this extended sample .71 ($.71 = 1/\sqrt{1+1}$)

when only one respondent is available for each measure, while the smallest value they can take in my main estimation sample is 2.23 ($2.23 = 1/\sqrt{1/10 + 1/10}$), more than three times as large.

Thus, the weighting approach is a less knife-edge way to address measurement error caused by small samples in measuring complementarity. This is not my preferred specification because it also weights the "intensive" margin of low-skilled employment. Consider, for instance, two industries that are the same size. This weighting scheme gives strictly more weight to the one with more low-skilled employment (since they have more complementarity respondents) which, by construction, gives less weight to industries with more non-low-skilled employment, which are precisely the industries that are useful for studying routine-biased technical change and job polarization.

The results in columns 4-6 are very similar to those in columns 1-3, but the standard errors are much smaller and all coefficients are statistically different from zero. The coefficient magnitudes themselves are smaller than those in the main results (Table 3). This is consistent with the claim that giving more weight to industries with more low-skilled workers (controlling for total employment) is less informative about routine-biased technical change.