Is College the New High School?

Evidence from Vacancy Postings*,†

Brad Hershbein
Upjohn Institute and IZA

Lisa B. Kahn
Yale University and NBER

December 15, 2014

*Correspondence: Brad Hershbein, W.E. Upjohn Institute for Employment Research, 300 S. Westnedge Ave., Kalamazoo, MI 49007. Email: hershbein@upjohn.org. Lisa Kahn, Yale School of Management, 165 Whitney Ave, PO Box 208200, New Haven, CT 06511. Email: lisa.kahn@yale.edu.

†We are grateful to Jason Abaluck, David Autor, David Deming, Alicia Sasser Modestino, and Daniel Shoag. We are especially indebted to Dan Restuccia for providing the Burning Glass data, and to Joe Altonji for providing us with O*NET measures.

Abstract

Do firms demand more skilled workers in a slack labor market? As the economy began to recover following the Great Recession, anecdotal evidence suggested that firms did indeed desire a more skilled worker than before. We use a new database of job vacancy postings, containing a near-universe of electronic posts in U.S. metro areas, collected by the firm Burning Glass over the period 2010–2013. The level of detail collected for each post allows us to analyze the education requirements of jobs as a function of local labor market conditions. We find evidence of upskilling—firms demanding more skilled workers—when the local unemployment rate is higher, i.e., in places where the recovery following the Great Recession was slower. Our estimates imply that firms demand three-quarters of a year more schooling, on average, in a large recession, compared to a boom. We find that at least two-thirds of this effect can be attributed to changes in skill requirements for the same firms and the same occupations, while the remainder is largely driven by shifts towards higher skilled occupations, not a shift in the distribution of firms hiring. Whether upskilling is a temporary response of firms to a slack labor market, or reflects a restructuring of production towards more skilled workers, is important for policy. Our findings that (1) the upskilling effect is concentrated among firms that always tended to have higher education requirements, (2) the distribution of vacancies shifts towards occupations that require more abstract tasks and upskilling is concentrated among these occupations, and (3) upskilling does not revert after the unemployment rate in a location has recovered, suggest that upskilling in the Great Recession reflects a more structural change due to growing polarization of the U.S. labor market.

1 Introduction

How does the hiring behavior of firms change in response to local labor market shocks? While it is well known that firms reduce their hiring in the aggregate when demand is weak, how firms and workers match during such a time, and the extent to which the distribution of their joint characteristics change, is less obvious. Evidence from past recessions shows that in downturns workers are more likely to take worse jobs, relative to their skills, and upgrade only as the economy improves.[1] Some of this may result because workers apply to a broader set of jobs, even those not targeted for them, when job-finding rates decline, but its is equally possible that firms actively seek a more-skilled worker than they could have attracted in a tighter market. Anecdotal evidence on the recovery following the Great Recession showed that firms did indeed demand a higher skilled worker for a given job.[2] Understanding the magnitude of this effect—and whether it is a temporary cyclical phenomenon or reflective of longer-term structural change in the U.S. labor market—has important implications for policies that annually allocate tens of billions of dollars toward subsidizing human capital attainment, worker retraining, and employment services.

In this paper, we use a new data set that is the near-universe of electronically posted job vacancies in the 100 largest metropolitan areas of the United States, containing detailed data on millions of postings at the vacancy level, to measure the education requirements of jobs as a function of local labor market conditions. Our data come from Burning Glass, a company which scrapes a wide range of online vacancy posts and systematizes information on firm, occupation, date, and location of the post, as well as unique information on the education and experience requirements of the job. We primarily exploit data from the 2010–2013 period, during which there was substantial cross-sectional variation in the speed at which locations were recovering from the Great Recession. We find that when local labor market conditions are worse, as measured by the MSA-level unemployment rate, firms demand higher-skilled workers. Our estimates imply that in a large recession, firms demand nearly three-quarters of a year more education, on average, compared to a boom.

We provide the first broad-based evidence of an upskilling effect in response to shocks in local labor market conditions.[3] Thanks to the level of detail in our data, we can decompose the upskilling effect into components driven by changes in the distributions of firms hiring,

[1] For example, Devereux (2002) shows that the education level of new hires within an occupation increases with the unemployment rate. Similarly, workers graduating from college in a downturn take jobs in worse occupations and firms compared to those graduating in booms and these poor initial matches can account for much of the negative earnings impacts for this group (Kahn 2010; von Wachter, Oreopoulos, and Heisz 2012; and Altonji, Kahn, and Speer 2014). Finally, the cyclical upgrading literature (e.g., Bils and McLaughlin 2001) emphasized that job matches in recessions were lower paying because high-paying industries are more cyclically sensitive.

[2] See Rampell (2012) and Burning Glass (2014).

[3] In a recent paper, Sasser Modestino and Shoag (2014) use a select set of occupations and U.S. states from aggregated Burning Glass data to estimate the labor demand response to local labor supply shocks, exploiting identifying variation from large scale releases of military personnel.

changes in the distribution of occupations demanded, and changes in skill requirements within firm-occupation cells. We find that at least two-thirds of the upskilling effect is driven by the last factor, upskilling at the firm-occupation level, while the remainder is largely driven by a shift towards occupations that tend to have higher skill requirements, with little-to-no impact from changes in the composition of firms.

We spend the remainder of the paper attempting to explain why upskilling occurs in slack labor markets and whether the phenomenon is likely cyclical or structural. We first show that downturns favor higher-skilled workers in that the returns to education widen in times of high unemployment, in terms of both earnings and the probability of being unemployed.[4] It is thus surprising that firms simultaneously increase the education requirements of new jobs. We consider two alternative explanations. First, lower-quality firms may attempt to hire workers they could not attract in a tighter labor market, resulting in a temporary upskilling of labor demand. Second, the recession may provide an opportunity for firms to permanently restructure their workforce away from those performing routine tasks—that could be instead performed by machines or by cheaper overseas labor—and towards more abstract tasks that require a higher skill level.

There is mounting evidence that firms at the bottom of the job ladder fare relatively better in a downturn than firms higher up on the job ladder. For example, Moscarini and Postel-Vinay (2013) show that small firms grow relative to large firms in slack markets, while Kahn and McEntarfer (2014) show the same for low-paying firms, compared to high-paying firms, and pinpoint that the effect is driven by a relatively larger reduction in voluntary quits at low-paying firms. Thus firms at the bottom of the ladder take advantage of a lack of outside opportunities for their incumbent workers, retaining workers who would have otherwise moved on to better opportunities. It is also reasonable, then, that these firms would employ a similar strategy for new hires as well.

By taking advantage of a pre-recessionary period in our data, from 2007, and the education distributions within specific firm-occupation cells, we can examine where in the skill distribution the increase in requirements is concentrated. If the job ladder explanation holds, we would expect to see an increase in the lower end of the distribution, as lower-quality firms and occupations become more competitive for higher-skilled workers. However, we show that the upskilling effect is concentrated in the upper two-thirds of the skill requirement distribution, slightly greater in the top third than the middle third, with no significant increase in the bottom third. It therefore seems unlikely that the increased demand for education in slacker labor markets is driven by lower-quality firms attempting to attract higher-skilled workers than they could ordinarily retain.

Job polarization, whereby middle-skill routine jobs are lost to outsourcing or replaced

[4] That recessions tend to disproportionately impact lower skilled workers is well known. See for example Hoynes, Miller, and Schaller (2012), as well as von Wachter and Handwerker (2009); Berube (2010); Looney and Greenstone (Nov 2011), and the citations therein.

by machines, is perhaps the most important trend of the U.S. labor market over the last 30 years.[5] The shift towards higher-skill abstract tasks and lower-skill manual tasks was thought to have been a gradual, secular change, until recent work by Jaimovich and Sui (2014) pointed out that the vast majority of jobs lost to polarization were lost in recessions. Firm restructuring in response to technological change may then be more episodic. Indeed, a large literature in macroeconomics, starting with Schumpeter (1939), posits that recessions are a time of "creative destruction" because ordinarily frictions, such as adjustment costs, inhibit resources from being allocated optimally. Recessions can produce large enough shocks to overcome such frictions.[6] However, there is no direct evidence on firm restructuring that is due to the forces driving polarization, or whether this is gradual or episodic.

We use O*NET data on the task content of occupations and show that a shift in labor demand toward more-abstract tasks and away from more-routine tasks is exacerbated in places with worse labor market conditions. We further show that the upskilling effect is concentrated in the more-abstract occupations. Finally, we provide early evidence that the increase in education requirements does not revert when a local labor market recovers from the recession. In that sense the upskilling phenomenon appears to be structural, not cyclical. These facts together imply that polarization of the U.S. labor force, exacerbated by the Great Recession, is the most important driver of the upskilling of labor demand. We thus identify the first firm-level evidence that episodic restructuring, centered around recessions, drives polarization.

The remainder of this paper proceeds as follows. Section 2 describes the data. Section 3 presents baseline results on upskilling and provides a decomposition. Section 4 disentangles mechanisms for possible drivers of the upskilling, given the decomposition. Section 5 concludes.

2 Data

Our data come from a unique source: microdata from over 40 million electronic job postings in the United States that span the Great Recession (between 2007 and 2013). These job postings were collected and assembled by Burning Glass Technologies, an employment analytics and labor market information firm. Although several studies have used aggregate vacancy data, and even vacancy microdata, from the Bureau of Labor Statistics' Job Openings and Labor Market Turnover (JOLTS) survey (see for example Davis, Faberman, and Haltiwanger 2012), these data contain little more than the sheer number of vacancies in a give time period. Fewer studies have used vacancy data that contain information on the occupation or specific requirements of the job posted, and these have generally used aggregate

[5] See for example the seminal work of Autor, Katz, and Kearney (2006).

[6] Many theoretical papers predict this phenomenon. See for example Hall (1991), Mortensen and Pissarides (1994), Caballero and Hammour (1994, 1996) and Gomes, Greenwood and Rebelo (2001).

data (Sasser Modestino, Shoag, and Ballance 2014), narrow slices (Rothwell 2014) of the data, or data that are limited to one vacancy source (Marinescu 2014). To our knowledge, we are the first study to use data based on a near universe of online job postings that cover the 100 largest metropolitan areas in the United States. In this section, we first describe the data in some detail, before discussing our particular sample construction.

2.1 Burning Glass Overview

With the growth of the Internet, vacancies for available jobs have increasingly appeared online instead of in traditional sources, such as newspapers. The Conference Board, the publisher of the Help-Wanted Advertising Index since the 1950s, affirmed this trend when it discontinued the index in 2008, after having begun a Help-Wanted Online Index in 2005.[7] Several other private-sector firms also began to track online job postings in the 2000s by using web-crawling and data-scraping methods. One such firm is Burning Glass Technologies (henceforth Burning Glass), which examines some 15,000 online job boards and company web sites to aggregate the job postings, parse and deduplicate them into a systematic, machine-readable form, and create labor market analytic products. The resulting database covers approximately 40 million openings that were posted between calendar years 2010 through 2013, as well as 2007.[8] It is especially rich in the detail collected about each posting, containing some 40 possible standardized fields, although not all of these are populated for every vacancy. Through a special agreement, we obtained access to the complete job postings database from Burning Glass database, covering the 100 largest metropolitan statistical areas (MSA's) —those with the most total postings—in the United States.

The job postings database from Burning Glass has several advantages and disadvantages when compared with other sources of vacancies. A notable advantage over JOLTS is that Burning Glass records vacancies at the occupation, and not just industry, level. Because variation in human capital is much greater across occupations than industries, measures of employer demand at the occupation level offer more potential insight into firm hiring behavior over the business cycle. The near-universe of the Burning Glass data also offers sufficient sample size to exploit geographic variation in labor market conditions at the metro area, or labor market, level, whereas JOLTS has much coarser geography. While vacancy data from other private-sector firms, such as Wanted Technologies, used by the Conference Board's Help Wanted Online Index, also offer these advantages over JOLTS, the Burning Glass data present a unique strength: they also capture the educational and skill requirements for the vacancy. Thus they allow for analysis of a key but largely unexplored margin of firm response to changing labor demand: the extent to which employers change education requirements for a given occupation, relative to how they change the quantity of posted vacancies. Because

[7] See https://www.conference-board.org/data/helpwantedonline.cfm.

[8] The database unfortunately lacks postings from 2008 and 2009.

the data also contain industry and firm identifiers (albeit imperfectly, see below), it is also possible to see how this relationship varies for different types of firms.

However, the richness of the Burning Glass data does not come without a few shortcomings. First, the database covers only vacancies posted on the Internet, as opposed to JOLTS or state vacancy reports that directly survey a representative sample of employers. To the extent that vacancies from certain industries and occupations are less likely to be posted electronically, as might be the case for many less skilled jobs, they will be underrepresented in the data. Indeed, Rothwell (2014) compares the occupational distributions from an extract of the Burning Glass data to those from state vacancy surveys for select metropolitan areas for which the data are available. He finds that computer, management, and business occupations are overrepresented relative to the state vacancy surveys, while health care support, transportation, maintenance, sales, and food workers are underrepresented.[9] While this unrepresentativeness would be cause for concern if one were using the data to estimate economy-wide trends in occupational demand, it is less problematic for our purpose, estimating the educational requirements of jobs as a function of local labor market conditions. Because we control for date and MSA fixed effects in all specifications below, the lack of representation would be a problem for us only if the educational requirements for online versus offline postings changed over our sample period or within a location.

Second, the database may not be representative even of all vacancies posted online. As the construction of the database relies on Burning Glass's proprietary algorithms to weed out duplicates from cross-listed vacancies as well as count all true posted vacancies, weaknesses in the algorithm could change the representativeness of the data. This is a particularly salient issue for coding the specific fields within a posting, which depend heavily on being parsed correctly even when a unique posting is found. However, again, problems with deduplication will be a problem for us only to the extent that errors in identifying duplicates vary within our controls.

2.2 Details and Construction of Analytic Data set

In the raw data, there are two fields for education requirements: a minimum degree requirement and a preferred degree requirement.[10] Postings for which Burning Glass did not find an education requirement have these fields set to missing. In practice, we use the minimum degree requirement to generate variables on the years of school required and whether or not

[9] We also have compared the occupational representativeness of the Burning Glass job postings data to new hires as reported in the monthly Current Population Survey (CPS) over the same time period. For job postings and (weighted) CPS hires at the 5-digit SOC occupational level, we find reasonably high correlations ($r > 0.8$) for the majority of 2-digit SOC groups, with several above 0.9. Only two 2-digit SOC groups have $r < 0.5$. Consistent with Rothwell (2014), the correlations are much higher among more skilled occupation groups.

[10] Burning Glass recodes listed degrees into the modal years of schooling associated with those degrees; that is, a high school diploma is coded as 12, an associate's degree as 14, a bachelor's degree as 16, a master's degree as 18, and professional/doctoral degrees as 21.

there is a requirement, because it is much more commonly used than the preferred degree and is always available when a preferred degree is listed.

The data contain the occupation of the posting (at the 6-digit Standard Occupation Classification 2010 (SOC) level), and MSA codes.[11] Burning Glass also collects the firm name for a given posting. Employer name is missing in 41% of postings, and when it is available, it is highly non-standardized.[12] We use a fuzzy match algorithm based on a 90% match threshold to identify postings that belong to the same employer, despite having slightly different names.[13] The fuzzy match reduces the number of firms from roughly 600,000 to roughly 400,000.

In practice, we collapse our individual-level postings into cells at the MSA-month-year-occupation (6-digit SOC)-firm level, preserving the average share of postings that specify an education requirement, and the mean years of schooling among these. Although the full data include the year 2007 as well as 2010 through 2013, we restrict our main analysis to the latter contiguous time period. We do this primarily because it leaves us a pre-recession base period from which we can decompose changes in the observed education requirements into different compositional factors that evolved over the Great Recession.

We then merge our Burning Glass data with monthly MSA-level unemployment rates from the Bureau of Labor Statistics's Local Area Unemployment Statistics (http://www.bls.gov/lau/).

While we have only a short panel of data, the MSA-level unemployment rates vary substantially over this time period. For example, figure 1 shows the time series of unemployment rates for selected MSAs. Our time period begins at roughly the peak unemployment rates following the Great Recession. San Fransisco and Atlanta start out with the highest unemployment rates, closely followed by New York City. Boston, Dallas, and Washington, DC did not have as large unemployment rate increases during the Great Recession. Subsequently, San Fransisco recovered more rapidly than NYC and Atlanta. Boston also recovered more rapidly than Dallas and DC. These examples demonstrate the kind of variation we are exploiting. Holding constant MSA fixed effects and date fixed effects, we have very different movements in unemployment rates across MSAs.

[11] Occupation and MSA codes are almost never missing in our data. Burning Glass identifies occupation based on the job title for a given listing.

[12] Some of these missings are "true" in the sense that the posting did not actually list the employer name, while others stem from the algorithm not capturing the field.

[13] We use WinPure software to perform the fuzzy match. For example, this program links together "Bausch and Lomb", "Bauch Lomb", and "Bausch & Lomb".

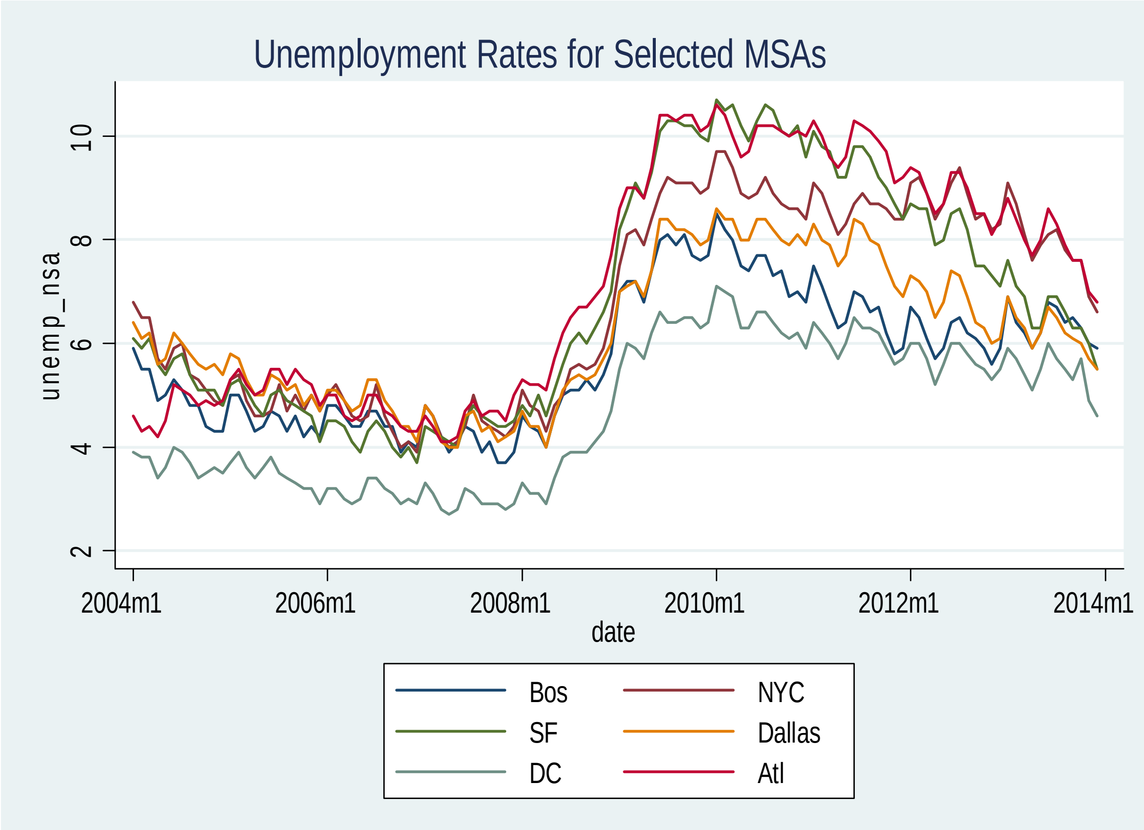

In our analysis, we weight observations by the size of the labor force in the MSA in 2007. This enables us to upweight larger MSAs, creating a more representative sample, and helping with precision. We use the size of the labor force in a pre-period rather than the number of postings throughout the sample to avoid simultaneity concerns.

Table 1 summarizes data for our primary regression sample (described in more detail below). Two-thirds of the ads in our preferred sample have an education requirement. The average education requirement is 9.65 years, when imputing a 0 for ads with no requirement, and 11.4 years conditional on having a requirement. The average MSA unemployment rate over this time period was 8.5% and ranged from 3.6% to 18.5%.

3 Degree Requirements and the Local Unemployment Rate

3.1 Overall Effect

We estimate regressions of the form specified in equation 1, where $shareED$ is the share of ads that have any education requirement in occupation, o; firm, f; MSA, m; and date (month-year), t.

(1) $$shareED_{ofmt} = \alpha_0 + \alpha_1 msa\_unemp_{mt} + I^m + I^t + \varepsilon_{ofmt}$$

The key explanatory variable is msa_unemp_{mt}, the MSA-level unemployment rate in month-year t. We control for date and MSA fixed effects, hence α_1 is identified purely off of cross-sectional variation in unemployment rates over time. This is a necessity given our short four-year panel, 2010-2013. We cluster standard errors by MSA-date, the level of variation underlying α_1. We weight regressions by the size of the labor force in the MSA in 2007 times the share of ads in an MSA-date belonging to a given firm-occupation pair.[14]

Table 2 reports regression results for a variety of specifications. Column 1 reports results for the whole sample; column 2 restricts the sample to identifiable firms; column 3 further restricts the sample to firm-occupation-MSA cells that can be linked to the data in 2007. The sample in column 3 (hereafter, our main sample) will be used for the remainder of the analysis so that we can decompose the effect into 2007 characteristics in a cell and changes. The sample size drops substantially with this restriction, so it is thus comforting that results are similar. Furthermore, we have analyzed the probability of being in the main sample as a function of the local unemployment rate and found no relationship.[15]

Column 3 shows that when the local unemployment rate increases by a percentage point (ppt), the share of ads with an education requirement increases by 0.0075, or by three-quarters of a ppt, significant at the 1%-level. In the Great Recession, the national monthly unemployment rate increased by roughly 6 ppts from trough to peak. Thus our estimates imply that in a large recession, $shareED$ increases by 0.042, a modest overall effect compared to its sample mean, 0.65. This effect can also be compared to the change in the total number of postings associated with a similarly large increase in the local unemployment rate, which we estimate to be a 13% reduction.

Columns 4 adds occupation fixed effects and shows the point estimate on msa_unemp declines by roughly a third in magnitude. The same is true when we instead control for firm fixed effects in column 5. Thus it looks as though the majority of the effect is driven by changes in the skill requirements of jobs within firm or occupation. However, firm and occupation are measured with substantial noise. In order to better understand the degree to which the "upskilling" effect can be accounted for by a change in the composition of jobs or changes within jobs, we perform a formal decomposition, below.

Because our data also contain information on the level of schooling required, when there is a requirement, we can further describe how much upskilling occurs in a slack labor market. Columns 6 and 7 report results on the average years of schooling required in postings in a given firm-occupation-MSA-date cell. In column 6, we take the average years of schooling, imputing zeroes for ads that did not specify a degree requirement, while in column 7 we

[14] These weights allow us to preserve the time-varying firm-occupation distribution within an MSA-date, while still upweighting larger MSAs.

[15] In a regression like that specified in equation 1 with an indicator of being in the main sample as the dependent variable, we found a coefficient on msa_unemp of -0.0012 (with a standard error of 0.002). This number is neither economically nor statistically significant.

restrict the sample to ads with a degree requirement.[16]

We estimate that for a 1 ppt increase in the MSA unemployment rate, the average years of schooling required increases by 0.12 when imputing zeros for those with no requirement (column 6). In a large recession (6 ppt increase in the unemployment rate), average years of schooling would increase by nearly 0.75 years. In data from the American Community Survey (see below), we estimate the return to each additional year of school to be roughly 10%. Thus, we estimate the degree of upskilling in a large recession is equivalent to a roughly 7% pay increase. In column 7, restricting the sample to ads with an education requirement, the impact is a bit smaller, equivalent to a 4% pay increase, because we miss the extensive margin, which turns out to be a large margin of adjustment.

3.2 Firm-occupation decomposition

How does this upskilling pattern manifest? In this subsection, we decompose the upskilling effect into three components: a change in the distribution of firms hiring, a change in the occupation distribution of ads within a firm, and a change in the propensity to have an education requirement within a firm-occupation cell. In order to perform this decomposition, we take advantage of the 2007 data, excluded from the rest of the analysis, to estimate firm-occupation education requirements in a pre-period.

Equation 2 rewrites $shareED_{mt}$, the share of ads with an education requirement in MSA, m, and time, t, into 3 components of interest. $\frac{N_{ofmt}}{N_{fmt}}$ is the share of ads in a given firm, f, that are in a given occupation, o, in mt, and $\frac{N_{fmt}}{N_{mt}}$ is the share of ads in mt that are from a given firm, f. $shareED_{mt}$, the share of ads requiring education across all firm-occupation groups in mt, is simply a weighted average of $shareED_{ofmt}$, the share of ads requiring education in a given firm-occupation group.

$$(2) \qquad shareED_{mt} = \sum_{f}(\sum_{o} shareED_{ofmt} * \frac{N_{ofmt}}{N_{fmt}}) * \frac{N_{fmt}}{N_{mt}}$$

We measure these components in a pre-period as well (2007), so we can thus decompose the change in education requirements for a given MSA in time t as in equation 3. Here, the three components are, in order: (1) the change in the distribution of occupations with posted vacancies, fixing the within-occupation firm distribution and the 2007 share of ads in an occupation-firm-MSA with an education requirement; (2) the change in the distribution of firms posting vacancies between 2007 and time t, fixing the within-firm occupation distribution and the 2007 education share, and (3) the within-cell change in education requirements

[16]Note that these adjustments are performed on the individual-level posting data before collapsing into cells. Thus, the sample size is the same across these two specifications because every cell has at least one ad with an education requirement.

between 2007 and t, fixing the firm-occupation distribution in t.

$$\text{(3)} \qquad \begin{gathered} shareED_{mt} - shareED_{m07} = \\ \sum_f \sum_o shareED_{ofm07} * (\frac{N_{ofmt}}{N_{fmt}} \frac{N_{fm07}}{N_{m07}} - \frac{N_{ofm07}}{N_{fm07}} \frac{N_{fm07}}{N_{m07}}) + \\ + \sum_f \sum shareED_{ofm07} * (\frac{N_{ofmt}}{N_{fmt}} \frac{N_{fmt}}{N_{mt}} - \frac{N_{ofmt}}{N_{fmt}} \frac{N_{fm07}}{N_{m07}}) \\ + \sum_f \sum_o (shareED_{ofmt} - shareED_{ofm07}) * \frac{N_{ofmt}}{N_{fmt}} \frac{N_{fmt}}{N_{mt}} \end{gathered}$$

We then regress $shareED_{mt}$ and each of the three components on msa_unemp_{mt}, as well as date and MSA fixed effects and $shareED_{m07}$.[17] By definition the coefficients on msa_unemp_{mt} for each of the three components should sum to that for $shareED_{mt}$. This yields one example of a decomposition. However, the order in which we differenced the three components was completely arbitrary. There are in fact 6 possible combinations for decomposing $shareED_{mt}$-$shareED_{m07}$ into changes in the three components. We estimate all 6 possible models, and report the decompositions in table 3. Regression estimates underlying table 3 can be found in appendix table 1.

To illustrate how table 3 should be read, we discuss the decomposition in row 1. The first column reports the full effect of the impact of msa_unemp_{mt} on $shareED_{mt}$, 0.007. This is very similar to the figure reported above, and is the same for each decomposition.[18] Column (2) shows the impact of msa_unemp_{mt} on the component of $shareED_{mt}$ that is driven by within firm-occupation changes in skill requirements, 0.0064, or 91.1% of the total effect. Column (3) shows the impact of msa_unemp_{mt} on the component of $shareED_{mt}$ that is driven by changes in the distribution of ads across firms, an insignificant 0.00045, or 6.4% of the total effect. Column (4) shows the impact of msa_unemp_{mt} on the component of $shareED_{mt}$ that is driven by changes in the distribution of occupations within firms, a very small 0.00017, or 2.5%. Column (5) reports the order in which we decomposed the effect. Beginning with $shareED_{m07}$, we first differenced the distribution of occupations within firms, followed by the distribution of ads across firms, followed by the change in the within occupation-firm skill requirements. This is in fact the decomposition represented in equation 3.[19]

Across all 6 decompositions, the change in the within-firm-occupation skill requirements accounts for the majority of the overall change, though the size of the impact ranges from roughly 60% to 90%. The impact of a change in distribution of ads across firms is tiny,

[17] Empirically, $shareED_{m07}$ is not subsumed in the MSA fixed effect because it will vary slightly within MSA over time. This is because it reports the education requirement in an MSA in 2007, based on the distribution of firm-occupation cells represented in a given MSA-date.

[18] It differs slightly from the figure reported in table 2, column 3, because the decomposition regressions also include a control for $shareED_{m07}$, which, as mentioned above, varies slightly within MSA over time.

[19] As another example, here is the formula for the decomposition in row 2, firm-occ-within: $\sum_f \sum_o shareED_{ofm07} * (\frac{N_{ofm07}}{N_{fm07}} \frac{N_{fmt}}{N_{mt}} - \frac{N_{ofm07}}{N_{fm07}} \frac{N_{fm07}}{N_{m07}}) + \sum_f \sum_o shareED_{ofm07} * (\frac{N_{ofmt}}{N_{fmt}} \frac{N_{fmt}}{N_{mt}} - \frac{N_{ofm07}}{N_{fm07}} \frac{N_{fmt}}{N_{mt}}) + \sum_f \sum_o (shareED_{ofmt} - shareED_{ofm07}) * \frac{N_{ofmt}}{N_{fmt}} \frac{N_{fmt}}{N_{mt}}$. Since the within firm-occupation component was last in both row 1 and row 2, the coefficient is identical (column 2).

and only marginally significant when it is the last component to be differenced (rows 3 and 4). This suggests the distribution of ads across firms (as represented by their education requirements) is relatively constant over the business cycle.[20] The impact of the changing mix of occupations within firms is modest, but when differenced last (rows 5 and 6) can account for a third of the total impact.

In summary, the increase in the share of ads with an education requirement associated with a higher local unemployment rate is overwhelmingly driven by a change in education requirements within firm-occupation cells. It is also partially driven by a shift in the distribution of occupations demanded by firms, towards those that typically have higher education requirements. A shifting composition of firms posting vacancies has only a modest effect, accounting for no more than a quarter of the total effect, and only in some specifications. That is, there might be a slight shift downturn towards firms that always have higher education requirements, but this effect is trivial.

4 Why Do Firms Increase Education Requirements in Slack Markets?

The evidence presented above shows a substantial upskilling of labor demand in reaction to eroding local labor market conditions.[21] The decomposition shows that the effect occurs primarily for the same firms and for the same types of jobs. As an illustrative example, in a better labor market, the pharmacy CVS might require their cashiers have a high school diploma, but in a worse one they might require a bachelor's degree.

Why would CVS decide they want a college-educated worker in a downturn? Has a college graduate become relatively cheaper? This is unlikely since it is well known that college graduates are somewhat sheltered from business cycle conditions.[22]

To investigate this, we estimate wage and unemployment regressions in the American Community Survey (ACS) as a function of education, local labor market conditions (the state unemployment rate), and an interaction.[23] These are reported in appendix table 2.

[20] This is consistent with work by Kahn and McEntarfer (2014) showing that the impact of the business cycle on hiring rates is similar across low- and high-paying firms (which likely also have fewer and more education requirements, respectively).

[21] More precisely, over our sample period, the coefficient on msa_unemp is identified primarily off of variation in the speed of improvement in labor market conditions.

[22] See for example Hoynes, Miller, and Schaller (2012) for an overview of the differential impact of business cycles on unemployment rates, across demographic groups, and how the Great Recession compared.

[23] For the years 2001-2012, we restrict the sample to those age 25-64 who are not enrolled in school and did not have an imputed value for education. For the wage regressions, we further restrict the sample to full-time (35+ hours), full-year (40+ weeks), employed workers and drop imputes on hours, earnings, or weeks worked. We adjust hourly earnings (total wage and salary earnings last year divided by hours times weeks worked) to 2006 dollars using the personal consumption expenditures (PCE) deflator, and restrict the wage sample to those earning at least $4 an hour (in 2006 dollars) and topcode at $250. Regressions control for gender; race/ethnicity (black, Hispanic, asian, other); interactions between gender, race/ethnicity, and a quadratic in age; as well as age, year, and state fixed effects. Regressions are clustered by state and weighted

We find that each additional year of school is associated with 0.10 higher log hourly earnings, and a 0.4 ppt reduced probability of being unemployed. We also find outcomes are worse when the local unemployment rate is higher (for those at the mean years of schooling, we estimate that wages fall by 0.65% for each point increase in the state unemployment rate, and the probability of being unemployed increases by 0.6 ppts). Most importantly, we find the effect of local labor market conditions on earnings and unemployment are offset for more educated workers. For example, comparing a college graduate to a high school graduate, we estimate an earnings gap of roughly 50%, that widens to 54% in a large recession, and an unemployment gap of 1 ppt that widens to 1.7 ppts in a large recession.

Thus, we estimate that more-educated workers are actually relatively more expensive in a slack labor market than in a tight labor market. However, in levels, the number of unemployed workers increases for all education groups in a downturn. Going back to our example, perhaps CVS always wants to hire college graduates, but in a tight labor market, it cannot attract them. As noted above, research suggests that firms at the bottom of the job ladder take advantage of the lack of outside opportunities afforded their workers in a downturn and increase worker retention.[24] If in a slack market, worse firms—those hiring disproportionately from the lower part of the skill distribution—similarly sought after new hires that they could not attract in a tighter market, we should see upskilling disproportionately occurring in firms at the bottom of the job ladder. We look at this effect below.

Alternatively the polarization of the U.S. labor market implies that middle-skill, routine jobs are being lost to outsourcing or being replaced by machines (see, for example, Autor, Katz, and Kearney (2006) and Autor and Dorn (2013)). This would certainly result in upskilling, i.e., a reduced demand for routine jobs. However, until recently, polarization had been thought to have occurred gradually over the last three decades. However, recent work by Jaimovich and Sui (2014) shows that jobs lost to polarization might be more episodic than previously thought. A literature in macroeconomics has emphasized that recessions can have a cleansing effect, purging the market of jobs that were no longer worthwhile. Adjustment costs imply that firms will not restructure their workforce until they have a big incentive to do so, like in a recession.[25]

Returning to our CVS example, a new technology—self-checkout machines—has become available and has the potential to reduce the labor bill for CVS, substituting a large number of

---

with sample weights.

[24] Moscarini and Postel-Vinay (2012) propose a dynamic search model whereby firms at the top of the job ladder can easily poach workers away from firms lower down on the job ladder. In response to a negative productivity shock (experienced by all firms), top firms stop poaching, allowing firms at the bottom of the job ladder to retain their workers. Consistent with this model, Moscarini and Postel-Vinay (2013) show that small firms grow relative to large firms in downturns. Similarly, Kahn and McEntarfer (2014) show that employment growth rates of low-paying firms are less cyclically sensitive than those at high-paying firms, precisely because low-paying firms experience a large decline in voluntary quits.

[25] Many theoretical papers assume a friction that generates lumpy adjustment. See for example Hall (1991), Mortensen and Pissarides (1994), Caballero and Hammour (1994, 1996) and Gomes, Greenwood and Rebelo (2001).

low-skilled clerks for one or two higher-skilled clerks who can operate the machines when they break down or customers get stuck. CVS waits until the local labor market is doing poorly because that is a time when it is more politically feasible to lay off workers. Polarization would predict a shift in job posts to less routine jobs, and it would also predict that the effects are not cyclical, but structural. When the economy recovers, we should not see a dropoff in upskilling. We test these below.

4.1 Job Ladder

The job ladder explanation is that upskilling is driven by a change in an educated worker's options, not by a change in a firm's preferences. A firm might always have preferred to hire a more skilled worker, but did not previously express a preference in a posting because they did not feel they could attract these workers. In a more slack labor market firms think they have a better chance of getting a higher skilled worker so they post a vacancy with an education requirement. This story implies the the upskilling effect should be concentrated among firms that historically tended not to post education requirements.

We divide firms into tertiles, based on their preexisting propensity to post an education requirement. We then estimate the differential impact of the unemployment rate on upskilling across tertiles. We estimate regressions as specified in equation 4.

(4)

$$shareED_{ofmt} = \alpha_0 + \alpha_1 msa\_unemp_{mt} + D\_ed_{f07}\alpha_2 + [msa\_unemp_{mt} * D\_ed_{f07}]\alpha_3 + I^m + I^t + \varepsilon_{ofmt}$$

Here, D_ed_{f07} is a vector of dummies indicating the tertile, which we we also interact with the MSA unemployment rate.

We define D_ed_{f07}in two ways. First, we simply take the firm-specific average share of ads with an education requirement in 2007, and divide firms into tertiles.[26] Second, we residualize the share of ads in a firm-occupation-MSA-date cell in 2007 on MSA, date, and occupation fixed effects, then divide firms into tertiles based on these residuals. We thus categorize firms based on their average skill requirement in 2007 or the average requirement holding constant the locations, dates, and occupations the firm tended to post in.

Results are reported in table 4. The omitted firm category is the lowest tertile, i.e., firms least likely to post an education requirement. These are exactly the firms we would expect to upskill the most, since they should be the least able to attract skilled workers at their average wage in a tight labor market. Instead we find the opposite. The coefficient on msa_unemp can be interpreted as the impact for firms in the lowest tertile, while the interaction terms show the differential effect for higher tertiles.

Across all specifications in table 4, we find no significant effect for firms in the lowest

[26] To determine tertile cutpoints, we estimate the 33rd and 66th percentile of the distribution of firm effects in 2007, weighting by our regression weights, i.e., the size of the labor force in the MSA in 2007 times the share of MSA-date ads attributable to a given firm-occupation-MSA cell.

tertile. Point estimates are tiny and statistically insignificant. In contrast, the positive, sizable interaction terms imply that the largest upskilling effects come from firms that always had a higher tendency to post ads with an education requirement.

It seems then unlikely that firms taking advantage of a more plentiful base of higher skilled workers willing to do lower quality jobs is the mechanism driving upskilling. Of course, we cannot rule out that this effect is going on to some extent within the higher-tertile firms, perhaps as a downshifting of more skilled workers into lower skilled occupations.

4.2 Polarization

In our decomposition above, we showed that up to a third of the upskilling effect was driven by a changing distribution of occupations towards those that tend to have education requirements. In this section, we explore whether the shift in occupations is towards those that we expect would be most impacted by polarization. For this, we exploit O*NET task data. We explore two measures that have been used in the trade and outsourcing literature, first developed in Oldenski (2012): a measure of the intensity of tasks that are thought to be more routine, and intensity of tasks that are thought to be highly non-routine.[27] In practice the results reported below are robust to a variety of alternative metrics.

In table 5, we report regression results where we first estimate whether there has been a shift toward, or away from, vacancy posts in routine and non-routine occupations. We then ask whether upskilling has disproportionately occurred in routine or non-routine tasks. For this exercise, we collapse observations to the occupation-date-MSA cell level, aggregating over firms.

Column 1 of table 5 regresses the average non-routineness of the occupation on the MSA unemployment rate. The metric is normed to be mean zero and standard deviation one. We indeed find that vacancy postings shift towards non-routine occupations in times of higher unemployment. For a 1 ppt increase in the MSA unemployment rate, the intensity of non-routine tasks in the average posting increases by 1% of a standard deviation. Column 4 shows there is no impact on routine tasks.

Columns 2 and 3 divide occupations into tertiles based on our measure of non-routineness, where a higher tertile represents occupations where non-routine tasks are more important. We then estimate regressions similar to those specified in equation 4, asking whether the impact of the unemployment rate on upskilling is differential across the task content of the occupation requested in the ad. Here we find that upskilling is stronger in the more non-routine occupations, and weaker in the less non-routine occupations. In contrast, in columns

[27] We define routine as the first principal component of the following work activities, using both the level and intensity required: handling moving objects, controlling machine processes, operating vehicles and equipment, and general physical activity. For non-routine, we use: thinking creatively, making decisions and solving problems, consultation and advice, developing objectives and strategy, communicating inside the organization, and interacting with computers.

5 and 6, we divide occupations into tertiles based on our measure of routineness. We find upskilling is lowest in the most routine occupations and higher for less routine occupations. Thus the surviving routine jobs do not change their skill requirements as much as the less routine jobs, while as the distribution of job postings shifts towards non-routine jobs, these jobs also require more skill.

4.3 Is Upskilling Cyclical or Structural

In this subsection, we explore whether upskilling is likely to revert as locations recover from the great recession, or whether upskilling reflects a structural, and therefore more permanent, change. For this, we exploit variation in the degree to which MSAs have "recovered" from the Great Recession. We take the average unemployment rate in an MSA from 2004-2007 as the baseline and explore whether the upskilling effect varies for locations whose unemployment rate has fallen back down, close to the baseline level. In practice, we define "close to" as either within 125%, 110%, or 100% (equal to) of the baseline unemployment rate.[28]

In practice, even though we have only a very short panel, we do have some variation in the degree to which locations have recovered from the Great Recession. For example, in our weighted sample, one-quarter of observations have recovered to 125% of the pre-recession unemployment rate, one-fifth of observations have recovered to 110%, and one-sixth of observations have fallen to or below the unemployment level in the pre-reccession period.

We augment the regression specified in equation 1 with an indicator for whether an MSA unemployment rate has recovered and an interaction between recovery and the current unemployment rate. Regression results are reported in table 6. The main effect of "recovery" is positive, indicating that locations that have recovered have on average have higher-skilled postings. We find that once a location recovers, the upskilling effect gets stronger, not weaker. This is seen in the positive interaction terms between the unemployment rate and recovery in columns 2 and 3 (defining recovery as 110% and 100%, respectively). We estimate that the impact on $shareED$ per unit change in the unemployment rate is roughly 2.5 times larger (0.02 increase off a base of 0.0076) once a location has recovered. Importantly, since the regression controls for MSA fixed effects, this result is not driven by any correlation between $shareED$ and the speed of recovery in a location.

[28] We follow the approach of Smith (2014) who uses locations that have recovered from the Great Recession to understand whether the long-term unemployed will ever be employable in locations that have not recovered. Using Current Population Survey longitudinally-linked data, he shows that the long-term unemployed have higher transition rates between unemployment and employment in states that have recovered to their pre-recession unemployment rates. This suggests that the long-term unemployed in states that have yet to make a full recovery will eventually find employment.

5 Conclusion

In the Great Recession, hiring fell dramatically and hire rates have yet to fully recover. During the recovery, anecdotal evidence suggested the composition of new hires shifted towards higher skilled workers, resulting in more educated workers being "undermatched". However, it remained unclear how broad or deep these effects were, and the extent to which they were driven by labor supply or labor demand responses. In particular, workers may have applied to jobs lower in the skill distribution than they would have during a tighter labor market, but also firms may have tried to recruit workers from higher in the skill distribution than they would have before the recessionary period. While the former response is almost certainly cyclical and unlikely to persist as the economy recovers, the latter response could be both cyclical and more permanently structural. Firms that normally hire from lower in the skill distribution may try to exploit the limited job opportunities of somewhat more-highly-skilled workers, but this phenomenon should fade as the labor market strengthens; on the other hand, firms that previously hired from the middle of the skill distribution may use the recession as a time for "cleansing" to reorient the factors of production to more abstract, higher valued-added tasks, and this effect would endure even after recovery.

In this paper we show that "upskilling", an increase in the education requirements of a job in response to local labor market conditions, did indeed occur, both fairly broadly and deeply. Despite the widening of education gaps in wages and unemployment that occurs in a downturn, making more-skilled workers relatively more expensive, firms at the same time actively seek higher-skilled workers via a greater propensity to specify education requirements in vacancy postings. We show that this effect occurs primarily within firm-occupation cells and not across the distribution of firms posting vacancies. Finally, we provide evidence that the effect is primarily driven by the polarization of the U.S. labor market that accelerated during the Great Recession.

We can rule out a downshift in worker job ladders as a key driver of our results, though of course a partial collapse of the job ladder in downturns can have important impacts on workers through other channels. We have not yet discussed the possibility that downward nominal wage rigidities drive our results. Faced with a wage schedule that is too high, given a temporary negative productivity shock, a firm may decide to compensate by employing more-productive workers. This might generate a temporary shift in labor demand. However, wages of new matches have been shown to be much more cyclical (Martins, Solon, and Thomas 2010). Furthermore this would generate only a temporary upskilling effect, while we find the increased demand for higher-educated workers remains after a location has recovered.

The shift in labor demand towards more-educated workers is likely to be permanent, not temporary. Thus, in addition to temporary income assistance, such as unemployment insurance benefits, programs that develop skills in in abstract tasks and critical thinking, as well as skilled manual tasks, are likely to be increasingly necessary to help displaced workers meet the changing labor demand of firms.

References

[1] Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer (2014) "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success," NBER working paper #20531.

[2] Autor, David, and David Dorn (2013), "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *American Economic Review*, 103(5): 1553-1597.

[3] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006), "The Polarization of the U.S. Labor Market," NBER Working Paper #11986.

[4] Bils, Mark and Kenneth J. McLaughlin (2001), "Inter-Industry Mobility and the Cyclical Upgrading of Labor," *Journal of Labor Economics*, 19(1), pp. 94-135.

[5] Burning Glass Technologies (2014), "Moving the Goalposts: How Demand for a bachelor's Degree is Reshaping the Workforce," mimeo.

[6] Caballero, R. and Hammour, M. (1994), "The Cleansing Effect of Recessions, *American Economic Review*, 84, pp. 1350-1368.

[7] Caballero, R. and Hammour, M. (1996), "On the Timing and Efficiency of Creative Destruction," *Quarterly Journal of Economics*, 111, pp. 805-852.

[8] Davis, Stephen J., Jason Faberman and John Haltiwanger (2010), "Labor Market Flows in the Cross Section and over Time," *Journal of Monetary Economics*, *Carnegie-Rochester Series on Public Policy* special issue, 59(1), pp. 1-18.

[9] Devereux, Paul (2002), "Occupational Upgrading and the Business Cycle," *Labour Economics*, 16(3), pp. 423-452.

[10] Gomes, J., J. Greenwood, and S. Rebelo (2001), "Equilibrium Unemployment," *Journal of Monetary Economics*, 48(1), pp. 109-152.

[11] Hall, Robert (1991), "Labor Demand, Labor Supply and Employment Volatility," in NBER Macroeconomics Annual (Cambridge: MIT Press)

[12] Hoynes, Hilary, Douglas L. Miller and Jessamyn Schaller (2012), "Who Suffers During Recessions?", *Journal of Economic Perspectives*, 26(3), pp. 27-48.

[13] Jaimovich, Nir, and Henry E. Siu (2012), "The Trend is the Cycle: Job Polarization and Jobless Recoveries," NBER Working Paper #18334.

[14] Kahn, Lisa B. (2010), "The Long-Term Consequences of Graduating from College in a Recession," *Labour Economics*, 17(2): pp. 303-316.

[15] Kahn, Lisa B., and Erika McEntarfer (2014), "Employment Cyclicality and Firm Quality," NBER working paper #20698.

[16] Marinescu, Ioana (2014), "The General Equilibrium Effects of Unemployment Insurance: Evidence from a Large Online Job Board," mimeo.

[17] Martins, P., G. Solon and J. Thomas, "Measuring What Employers Do About Entry Wages Over the Business Cycle: A New Approach," *American Economic Journal: Macroeconomics*, forthcoming.

[18] Mortensen, D. and C. Pissarides (1994), "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, pp. 397-415.

[19] Moscarini, Giuseppe and Fabien Postel-Vinay (2013), "Stochastic Search Equilibrium," *Review of Economic Studies*, 80(4), pp. 1545-1581.

[20] Moscarini, Giuseppe and Fabien Postel-Vinay (2012), "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment," *American Economic Review*, 102 (6), pp. 2509-2539.

[21] Oldenski, Lindsay (2012), "Export versus FDI and the Communication of Complex Information," *Journal of International Economics*, 87(2).

[22] Oreopoulos, Phil, Till von Wachter and Andrew Heisz (2012), "The Short- and Long-Term Career Effects of Graduating from College in a Recession," *American Economic Journal: Applied Economics*, 4(1) pp. 1-29.

[23] Rothwell, Jonathan (2014), "Still Searching: Job Vacancies and STEM Skills," Brookings mimeo.

[24] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2014), "Upskilling During the Great Recession: Do Employers Demand Greater Skill when Skilled Workers are Plentiful?", mimeo.

[25] Schumpeter, J. (1939), *Business Cycles* (New York: McGraw-Hill).

[26] Smith, Christopher L. (2014), "Using cross-state variation to assess the potential for additional improvement in measures of labor market conditions," *FEDS Notes*, June 2, 2014.

Table 1: Summary Statistics

	N	mean	stdev	min	max
Has an Education Requirement	4,298,109	0.65	0.42	0	1
MSA Unemployment Rate	4,298,109	8.52	1.95	3.6	18.5
Years of Schooling, Unconditional	4,298,109	9.65	6.35	0	21
Years of Schooling, Conditional	4,298,109	11.40	6.46	0	21
Log(Number of Postings)	4,298,109	1.55	1.36	0	7.74

Notes: This table reports summary statisics for columns 3-7 of table 2. See table 2 notes for details.

Table 2: Education Requirements of Vacancy Postings and Local Labor Market Conditions

Dependent Variable:	Share of ads with an education requirement					Average Years Required	
						0 if none	missing if none
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
MSA Unemployment Rate	0.00475***	0.00910***	0.00752***	0.00462***	0.00453***	0.118***	0.0650***
	(0.00138)	(0.00119)	(0.00136)	(0.00105)	(0.000954)	(0.0208)	(0.0198)
Constant	0.360***	0.478***	0.558***	0.618***	0.628***	8.219***	9.724***
	(0.0210)	(0.0181)	(0.0201)	(0.0145)	(0.0121)	(0.332)	(0.294)
Date and MSA FE's	X	X	X	X	X	X	X
Non-Missing Firm		X	X	X	X	X	X
2007 Firm-Occ-MSA Match			X	X	X	X	X
Occupation FE's				X			
Firm FE's					X		
Occ-Firm-MSA-Date cells	16,403,910	15,031,483	4,298,109	4,298,109	4,298,109	4,298,109	4,298,109
R-squared	0.017	0.008	0.010	0.216	0.343	0.014	0.021

*** p<0.01, ** p<0.05, * p<0.1

Notes: In columns 1-5 the dependent variable is the share of ads in a firm-occupation (6-digit SOC)-MSA-date (month-year) cell with an education requirement. In columns 6 and 7 the dependent variable is the average years of education required in the firm-occupation-MSA-date cell, with zeroes filled in for postings with no requirements (column 6) or left missing (column 7). Observations are weighted by the size of the labor force in the MSA in 2007 times the share of ads in a given MSA-date belonging to this cell. Standard errors are clustered at the date-MSA level.

Table 3: Decomposing the impact of the MSA-level unemployment rate on skill requirements

			Impact attributable to a change between 2007 and t in:			
		(1)	(2)	(3)	(4)	(5)
		Full Effect	within firm-occ skill requirement	distribution of ads across firms	distribution of occs w/in firm	Order of decomposition:
	coeff	0.00702***	0.0064***	0.00045	0.00017	occ-firm-within
1.	s.e.	(0.0011)	(0.0011)	(0.00055)	(0.00037)	
	% of total		91.1%	6.4%	2.5%	
	coeff	0.00702***	0.0064***	0.00019	0.00043	firm-occ-within
2.	s.e.	(0.0011)	(0.0011)	(0.00053)	(0.00037)	
	% of total		91.1%	2.8%	6.1%	
	coeff	0.00702***	0.0052***	0.00163*	0.00017	occ-within-firm
3.	s.e.	(0.0011)	(0.0012)	(0.00086)	(0.00037)	
	% of total		74.4%	23.2%	2.5%	
	coeff	0.00702***	0.0041***	0.00163*	0.0013***	within-occ-firm
4.	s.e.	(0.0011)	(0.0012)	(0.00086)	(0.00049)	
	% of total		58.0%	23.2%	18.9%	
	coeff	0.00702***	0.0043***	0.000194	0.0026***	firm-within-occ
5.	s.e.	(0.0011)	(0.0011)	(0.00053)	(0.00042)	
	% of total		60.8%	2.8%	36.4%	
	coeff	0.00702***	0.0041***	0.00040	0.0026***	within-firm-occ
6.	s.e.	(0.0011)	(0.0012)	(0.00086)	(0.00042)	
	% of total		58.0%	5.6%	36.4%	

*** p<0.01, ** p<0.05, * p<0.1

Notes: In columns 1-5 the dependent variable is the share of ads in a firm-occupation (6-digit SOC)-MSA-date (month-year)

Table 4: Upskilling and Firm Quality

Dependent Variable:	Share of ads with an education requirement			
	Firm Average		Residualized Firm Average	
	(1)	(2)	(3)	(4)
MSA Unemployment Rate (U)	0.00179	6.59e-06	0.00116	-0.000885
	(0.00160)	(0.00121)	(0.00161)	(0.00118)
U*2nd Skill Tertile	0.00653***	0.00445***	0.00820***	0.00649***
	(0.000981)	(0.000807)	(0.000953)	(0.000739)
U*3rd Skill Tertile	0.00830***	0.00821***	0.0104***	0.0101***
	(0.000993)	(0.000773)	(0.00102)	(0.000732)
Constant	0.490***	0.594***	0.494***	0.611***
	(0.0225)	(0.0165)	(0.0229)	(0.0164)
Date and MSA FE's	X	X	X	X
Occupation FE's		X		X
Occ-Firm-MSA-Date cells	4,298,109	4,298,109	4,298,109	4,298,109
R-squared	0.068	0.238	0.059	0.241

*** p<0.01, ** p<0.05, * p<0.1

Notes: See table 2. In columns 1 and 2, tertiles are calculated from the average share of ads with an education requirement in a firm in 2007. In columns 3 and 4, share of ads with education is first residualized on date, msa, and occupation fixed effects in 2007 and then tertiles are formed.

Table 5: The Task Content of Occupations and Polarization

Depndent Variable:	Non-Routineness Task	Non-Routineness Share of ads requiring education		Routineness Task	Routineness Share of ads requiring education	
	(1)	(2)	(3)	(4)	(5)	(6)
MSA Unemployment Rate (U)	0.00953***	0.00308***	0.00553***	0.00157	0.00833***	0.00765***
	(0.00134)	(0.000361)	(0.000270)	(0.00133)	(0.000362)	(0.000275)
U*2nd Tertile		0.00396***	-0.00138***		0.000885***	0.000830***
		(0.000230)	(0.000174)		(0.000235)	(0.000179)
U*3rd Tertile		0.00470***	0.00194***		-0.00494***	-0.00580***
		(0.000241)	(0.000182)		(0.000235)	(0.000180)
occupation-MSA-date cells	1,134,358	1,134,358	1,134,358	1,134,358	1,134,358	1,134,358
R-squared	0.015	0.216	0.562	0.031	0.238	0.562

*** p<0.01, ** p<0.05, * p<0.1

Notes: See table 2. An observation is an occupation-MSA-date cell. Routineness is defined as the first principle component of the following O*NET measures of both intensity and level requirements in the occupation: Handling moving objects, controlling machine processes, operating vehicles and equipment, general physical activity. Non-routineness uses the following tasks: Thinking creatively, making decisions and solving problems, consultation and advice, developing objectives and strategy, communicating inside the organization, and interacting with computers. Both metrics are normed to be mean zero and standard deviation one.

Table 6: Cyclical or Structural Upskilling

Dependent Variable:	Share of ads with an education requirement		
Explanatory Variable:	Local unemployment as recovered to X% of pre-recession unemployment rate		
	125%	110%	100%
	(1)	(2)	(3)
MSA Unemployment Rate (U)	0.00778***	0.00762***	0.00756***
	(0.00136)	(0.00136)	(0.00136)
Recovered	0.00581	0.0561**	0.0885***
	(0.00701)	(0.0230)	(0.0343)
Recovered*U	-0.00528	0.0173**	0.0223*
	(0.00337)	(0.00874)	(0.0121)
Date and MSA FE's	X	X	X
Occ-Firm-MSA-Date cells	4,298,109	4,298,109	4,298,109
R-squared	0.010	0.010	0.010

*** p<0.01, ** p<0.05, * p<0.1

Notes: See table 2. Recovered is an indicator equaling 1 if in a given MSA date the MSA unemployment rate is less that X% of the pre-recession (2004-2007) average, where X is indicated in the column heading.

Appendix Table 1: Decomposition Regression Results

	Overall	1. occ-firm-within			2. firm-occ-within			3. occ-within-firm		
		within	firm	occ	within	firm	occ	within	firm	occ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
MSA Unemployment Rate	0.00702***	0.00640***	0.000448	0.000173	0.00640***	0.000194	0.000427	0.00522***	0.00163*	0.000173
	(0.00110)	(0.00109)	(0.000553)	(0.000366)	(0.00109)	(0.000527)	(0.000366)	(0.00116)	(0.000861)	(0.000366)
Observations	4,700	4,700	4,700	4,700	4,700	4,700	4,700	4,700	4,700	4,700
R-squared	0.535	0.748	0.665	0.603	0.748	0.696	0.408	0.712	0.484	0.603

	4. within-occ-firm			5. firm-within-occ			6. within-firmocc		
	within	firm	occ	within	firm	occ	within	firm	occ
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)
MSA Unemployment Rate	0.00407***	0.00163*	0.00132***	0.00427***	0.000194	0.00255***	0.00407***	0.000395	0.00255***
	(0.00115)	(0.000861)	(0.000490)	(0.00108)	(0.000527)	(0.000415)	(0.00115)	(0.000862)	(0.000415)
Observations	4,700	4,700	4,700	4,700	4,700	4,700	4,700	4,700	4,700
R-squared	0.707	0.484	0.366	0.748	0.696	0.482	0.707	0.443	0.482

*** p<0.01, ** p<0.05, * p<0.1

Notes: See table 3.

Appendix Table 2: Education Gaps in Labor Market Outcomes and the Unemployment Rate, ACS Data

Dependent Variable:	Log (wages)	Unemployed
	(1)	(2)
Years of Schooling (Edu)	0.102***	-0.00400***
	(0.00107)	(0.000297)
State Unemployment Rate (U)	-0.00654***	0.00590***
	(0.00192)	(0.000415)
Edu*U	0.000915***	-0.000371***
	(0.000215)	(0.000116)
Constant	3.221***	0.0257***
	(0.0212)	(0.00658)
Observations	6,453,824	13,910,929
R-squared	0.327	0.013

*** p<0.01, ** p<0.05, * p<0.1

Notes: The data are from the American Community Survey 2001-2012, restricted to respondents age 25-64 with non-imputed education. Column 1 further restricts to full-time (35+ hours), full-year (40+ weeks), employed, workers and drop imputes on hours, earnings, or weeks worked. We adjust hourly earnings (total wage and salary earnings last year divided by hours times weeks worked) to 2006 dollars using the personal consumption expenditures (PCE) deflator, and restrict the wage sample to those earning at least $4 an hour (in 2006 dollars) and topcode at $250. Regressions control for gender; race/ethnicity (black, hispanic, asian, other); interactions between gender, race/ethnicity, and a quadratic in age; as well as age, year, and state fixed effects. Regressions are clustered by state and weighted with sample weights. Years of schooling and the state unemployment rate have been de-meaned within each sample.