UPSKILLING: DO EMPLOYERS DEMAND GREATER SKILL WHEN WORKERS ARE PLENTIFUL?

Alicia Sasser Modestino, Northeastern University*
Daniel Shoag, Harvard University
Joshua Ballance, Federal Reserve Bank of Boston

Abstract

Using a proprietary database of online job postings, we find that education and experience requirements rose during the Great Recession. These increases were larger in states and occupations that experienced greater increases in the supply of available workers. This finding is robust to controlling for local demand conditions and firm × job-title fixed effects as well as using a natural experiment arising from troop withdrawals as an exogenous shock to labor supply. Our results imply that the increase in unemployed workers during the Great Recession can account for 18 to 25 percent of the increase in skill requirements between 2007 and 2010.

JEL classifications: D22, E24, J23, J24, J63.

---

* Corresponding Author: School of Public Policy & Urban Affairs, Northeastern University, 360 Huntington Avenue, Boston, MA 02115, (617) 373-7998 a.modestino@neu.edu

The views expressed herein are those of the authors and do not indicate concurrence by the Federal Reserve Bank of Boston, or by the principals of the Board of Governors, or the Federal Reserve System. The authors thank David Autor, Bill Dickens, Chris Foote, Monica Galizzi, Brad Hershbein, Harry Holzer, Lisa Kahn, Yolanda Kodrzycki, Phil Moss, Paul Osterman, Jessica Wolpaw Reyes, Jonathan Rothwell, James Spletzer, Robert Triest, Jeff Zabel, and seminar participants from Amherst College, the Federal Reserve Bank of Boston, the Federal Reserve System Committee on Microeconomic Analysis, the Harvard Kennedy School, MIT Institute for Work and Employment Research, Northeastern University, the Society of Labor Economists, UMass Lowell, and Wellesley College for their valuable comments and insights. Special thanks to Dan Restuccia, Matthew Sigelman, and Bledi Taska of Burning Glass Technologies for supplying the data and providing insights regarding the collection methodology. We also thank the Russell Sage Foundation (award #85-14-05) and the NULab for Texts, Maps, and Networks for their generous support of this work. All remaining errors are our own.

I. Motivation: Shifting Employer Skill Requirements and Recruitment Intensity

During the slow recovery of the U.S. labor market from 2007-2012, there was a change in the relationship between the unemployment rate and the job vacancy rate known as the Beveridge curve. Following the Great Recession, the unemployment-to-vacancy-rate ratio was significantly higher than one would have projected (Diamond and Şahin 2015). Several explanations for this shift have been explored, including a mismatch between workers and vacancies across occupations and geographies, changes in the composition of job seekers, and changes in the motivation of job seekers.[1]

Another important factor that has been proposed to explain the shift in the Beveridge curve during the downturn and recovery period is a change in employer "recruitment intensity" per vacancy. In this context, recruitment intensity is described as a set of actions that employers can take to influence the rate of new hires, such as changes in "advertising expenditures, screening methods, hiring standards, and the attractiveness of compensation packages" (Davis, Faberman, and Haltiwanger 2012). For a given unemployment-to-vacancy ratio, actions that lower the recruiting intensity per vacancy also lower the fill rate, resulting in an upward shift in the Beveridge curve. This finding has sparked several theoretical models that endogenize this channel (Kaas and Kircher 2014; Mongey, Violante, and Gavazza 2015). Yet quantifying the magnitude of the change in recruitment intensity and why it occurs has been limited by the absence of direct measures of actions undertaken by employers (Diamond 2013; Rothstein 2012).

In this paper, we directly measure an important facet of recruitment intensity that shifted

[1] This extensive debate in the economics literature sparked numerous papers seeking to explain the apparent shift in the Beveridge curve. See Şahin et al. (2014); Barnichon, and Figura (2010); Shimer (2012); Fujita and Moscarini (2013); Hall and Schulhofer-Wohl (2013); and Mukoyama, Patterson, and Şahin (2014).

during the Great Recession — namely, the skill requirements employers use to screen candidates when filling a new vacancy. Indeed, media reports and employer surveys indicate that employer requirements increased sharply between 2007 and 2012, such that a college degree was considered a requirement for many occupations that previously required only a high school degree. This trend has colloquially become known as "upskilling."[2]

Anecdotal accounts suggest that upskilling during the Great Recession was driven to some degree by a sense among employers that "[t]he recession is a wonderful opportunity to acquire top talent" when workers are more plentiful.[3] In contrast, as the labor market has recovered, employers report that "managers have to be much more flexible now than during the recession because there's less talent available."[4] These sentiments are consistent with what we find using a proprietary dataset of 36.2 million online job postings aggregated by Burning Glass Technologies (BGT). Figure 1 shows that the percentage of vacancies requiring a bachelor's degree or higher rose by more than 10 percentage points between 2007 and 2010 and then fell as the labor market recovered. A similar relationship is observed for the percentage of postings requiring four or more years of experience. Clearly, there is a strong time-series correlation between employer skill requirements and aggregate labor market slack as measured by the national unemployment rate. Yet it is still unclear the degree to which these aggregate trends reflect a *causal* shift in recruitment intensity in response to the increase in the supply of workers during the recession.

---

[2] Almost one-third of employers said their educational requirements for employment had increased over the past five years and specifically that they were hiring more college-educated workers for positions that were previously held by high school graduates. See CareerBuilder. 2014. "Education Requirements for Employment on the Rise, According to CareerBuilder Survey." March. See also Rampell, Catherine. 2012. "Degree Inflation? Jobs That Newly Require B.A.'s" The New York Times, December 4.

[3] Barry Deutsch, chief executive of IMPACT Hiring Solutions, as quoted in Green, Alison. 2009. "How the Recession Has Changed Hiring." U.S. News and World Report, June 15.

[4] Similar sentiments are expressed in interviews with four dozen employers in Modestino, Moss, and Shoag (2017).

To test the upskilling hypothesis, we start by exploiting the variation in unemployment rates across states during the Great Recession. We find that employer skill requirements increased more within occupations in states experiencing greater increases in their unemployment rate. The relationship is economically important: Within a six-digit detailed occupation, a 1 percentage point increase in the state unemployment rate is associated with a 0.6 percentage point increase in the fraction of employers requiring a bachelor's degree and a 0.8 percentage point increase in the fraction of employers requiring four or more years of experience. These OLS estimates are robust to using alternative measures of labor market slack, such as labor supply/demand ratios, and to including occupation, state, and year fixed effects and their interactions, as well as controls for initial skill requirements and their frequency among the population. To our knowledge, these findings provide some of the first empirical evidence of a shift in recruitment intensity whereby employer skill requirements are driven—in part—by the available supply of labor.[5]

Although this baseline relationship between rising employer requirements and the supply of jobless people seeking work is intriguing, the variation in the unemployment rate over the business cycle is potentially correlated with other factors. These include short-term factors, such as changes in the demand for certain goods or services and credit availability, as well as longer-term trends such as changes in technology or other production processes.

To establish a causal relationship between changing employer skill requirements and the supply of job seekers, we employ two central identification strategies. First, to account for

[5] A more recent paper by Hershbein and Kahn (2018) uses the same dataset of online job vacancies to study long-term structural shifts in employer skill requirements. That paper estimates the change in the requirements caused by Bartik demand shocks and finds considerable persistence. In contrast, this paper focuses on the causal identification of the portion of the increase in employer skill requirements that is related to the increased availability of workers (as opposed to demand shocks) during the business cycle.

changes in the composition of employers and/or vacancies over time, we show that upskilling occurs even *within firm × job-title pairs*—not just within occupations.[6] Our findings show greater increases in employer skill requirements in states and occupations experiencing larger increases in the unemployment rate, both within and across firm × job-title pairs.

Second, we introduce a plausibly exogenous instrument for the number of searchers in a given state and occupation. Specifically, we make use of a natural experiment that represents a clear shock to labor supply: the drawdown of troops from Iraq and Afghanistan between 2009 and 2012. We show that these troop withdrawals led to an additional 200,000 to 300,000 veterans entering the U.S. domestic labor force each year and were not correlated with underlying labor market trends. Consistent with the upskilling hypothesis, we find that state × occupation cells receiving larger numbers of returning veterans correspondingly experienced a greater increase in their skill requirements. We can further purge this instrument of potentially confounding correlation with contemporaneous shocks by instrumenting for a veteran's current state of residence using his or her state of birth. Finally, we combine the veterans-by-state-of-birth shock with the within firm × job-title specification to produce estimates using both approaches. These relationships imply effects on the same order of magnitude as the non-IV results, confirming that an exogenous increase in the supply of job searchers leads firms to change their job posting requirements.

The finding that employer skill requirements are driven—in part—by the available supply of labor has important implications for understanding the dynamics of the labor market. We document a novel feedback mechanism between labor supply and the selectivity of vacancies that operates within occupations and is consistent with macroeconomic models of

[6] It has been shown that firm × job-titles account for the vast majority of the variance in wages within occupations, making confounding compositional changes unlikely (Marinescu and Wolthoff 2016).

employer search decisions (Davis, Faberman, and Haltiwanger 2012) and heterogeneous workers (Shimer 2005; Albrecht and Vrooman 2002). Importantly, we find that upskilling occurs even within firm × job-title pairs, a notion that runs counter to some of the existing approaches to modeling changes in recruitment intensity as solely a compositional effect. Moreover, a related literature has explored worker entry and mobility during recessions, particularly for college graduates. These studies typically find that workers match at lower entry wages during recessions and have less steep wage trajectories over time (e.g., Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Moscarini 2001). We find that changes in employer requirements over the business cycle are consistent with—and even serve to reinforce—this effect.

II. Data: Using Job Postings to Measure Changes in Employer Skill Requirements

To study changes in employer hiring dynamics, we use a large, detailed dataset of online job postings. Over the past two decades, online vacancy data have been used by a number of researchers to study labor market dynamics (Kuhn and Skuterud 2004; Bagues and Labini 2009; Şahin et al. 2014; Marinescu 2017).

The advantage of using online vacancy data is that it allows analysis at a greater frequency and at more refined geographies than traditional employer surveys, such as the Job Opening and Labor Turnover Survey (JOLTS). This is because the data are constructed from measures collected by software that parses text contained in millions of job ads posted online daily. One potential drawback is that online vacancy data only capture vacancies posted on the Internet and may not be representative of the universe of job openings if vacancies from certain industries and occupations are less likely to be posted electronically. However, estimates show as of 2012 between 60 and 70 percent of job vacancies were posted online (Carnevale, Jayasundera, and Repnikov 2014). Other research shows that online job ads exhibit trends that are closely correlated with employer surveys over time as well as across industries and occupations (Templin and Hirsch 2013).

The main source of online job posting data used in this paper is collected and aggregated by Burning Glass Technologies. BGT aggregates detailed information daily on more than 7 million online job openings from over 40,000 sources including job boards, newspapers, government agencies, and employer sites.[7] These data are collected via a web-crawling technique that uses computer programs called "spiders" to browse online job boards and other websites, remove duplicates, and systematically text parse each job ad into usable data elements. BGT mines more than 70 job characteristics from free-text job postings, including job title, employer name, location, and the level of education and years of experience required. BGT then codes each posting to create education and experience categories as well as occupational groupings using the 2010 Standard Occupational Classification (SOC) hierarchy.[8]

BGT provides snapshots of the data in which vacancies are reported monthly and are pooled over the year without duplication. As such, these data are unique in allowing geographical analysis of labor demand by education and experience level over time. The data are available by state for detailed occupations—down to the six-digit SOC code level—in 2007, 2010, and 2012.[9] In total, our data represent roughly 36.2

[7] See http://www.burning-glass.com/realtime/ and the online appendix for more details on this data source.
[8] Carnevale, Jayasundera, and Repnikov (2014) audited a sample of job postings in the BGT database and found that the BGT coding for occupation, education, experience was accurate at least 80 percent of the time. This is likely an underestimate given algorithm improvements that have been retroactively and consistently applied since then.
[9] No data are available for 2008 and 2009. This is the main reason for using three-year differences (from 2007-2010), although we also show estimates over the whole five-year period.

million job postings across these three years. About half of the BGT postings also identify the employer name, which allows us to eliminate potential changes in the BGT sample composition or construction by also employing specifications that draw on a panel of firms and job titles over time, in addition to our OLS specifications that make use of the full cross-sectional sample.

Though we use the BGT data primarily as a dependent variable (meaning that random noise does not bias our regressions), it is important to understand the coverage patterns for interpreting the results. We explore this issue in detail in the online appendix, and other authors have also tested the robustness of these data (Carnevale, Jayasundera, and Repnikov 2014; Rothwell 2014; Hershbein and Kahn 2018). Despite differences in the sampling frame of the BGT data compared to state and national employer job vacancy surveys, the industry and occupation distributions are quite similar and are consistent over time.[10] As a robustness check, we also replicate our firm × job-title analysis using the Minnesota Job Vacancy Survey and find results similar to those generated by the BGT data (see Table B7 in the online appendix).

A. Changes in Employer Skill Requirements

Based on the education and experience fields parsed from the BGT online job postings, we construct several measures of employer skill requirements. Table 1 provides descriptive statistics for our dependent variables across the various samples we use in our identification strategies. Our first identification strategy controls for state and occupation fixed effects and makes use of the total sample of 36.2 million job postings, aggregated into state × occupation × year cells.[11] On average, there are roughly 500 to

---

[10] These comparisons show that the BGT data tend to slightly over-represent industries such as finance and slightly under-represent others such as food services. Similarly, occupations such as management are slightly over-represented while occupations such as food preparation are slightly under-represented.
[11] Job title is always populated for each posting so that all observations can be categorized by occupation.

600 postings for a given state × occupation cell in each year (2007, 2010, and 2012). It should be noted that these data exhibit a considerable amount of variation given the different employment levels of these occupations, even within state × occupation × year cells. The number of underlying observations available to construct some cells varies from as few as one posting to as many as 60,000 postings. To ensure that our dependent variables capture meaningful differences over time and accurately represent the state of the labor market, we drop observations with fewer than 15 total postings in a given cell which corresponds to eliminating the bottom 5 percent of the sample. In addition, since we are analyzing changes in the fraction of postings requiring a particular skill, we weight the observations by the occupation's share of total openings in the state in a given time period in all regressions. This ensures that our results are not driven by outlier occupations with few underlying postings, and that the regressions are not dominated solely by large states. [12]

We construct two primary dependent variables by state × occupation × year to measure the share of job postings requiring two dimensions of skill: educational attainment and years of experience. The BGT education categories range from the share of postings with no education requirement to the share requesting a graduate or professional degree — and all levels in between.[13] Required experience is measured continuously, although the most postings are concentrated at round numbers and half-year increments. Our primary dependent variables are the share of postings requiring a

[12] Our results are robust to more stringent restrictions, such as dropping observations for which there are fewer than 75 openings for a given occupation × state cell, as well as to alternative weighting schemes such as weighting observations by the minimum total openings in both periods.
[13] Some job postings list both a minimum ("required") and maximum ("preferred") requested educational qualification. For example, approximately 12 percent of job postings specify that a bachelor's degree is required but a graduate degree is preferred. The results in all of our specifications are qualitatively similar if instead we use the minimum qualifications required.

bachelor's degree or greater and the share of postings requiring at least four years of experience.[14] Prior to the Great Recession, roughly 13 percent of postings requested a bachelor's degree or higher in 2007, whereas 8 percent of postings requested four or more years of experience. Employer requirements along both dimensions of skill have increased over time, with most of the increase occurring between 2007 and 2010 during the height of the Great Recession (see Table 1).

We also make use of two other samples to implement our second identification strategy that controls for firm × job-title fixed effects. The first is a sample of BGT job postings that identify employer names, representing roughly half of the total sample (17.3 million postings). Occasionally, the data show an extremely large number of observations for the same firm × job title × state cell in a given month and year. We address this by eliminating extreme duplicates with more than 50 observations for a given firm × job title × state × year × month (less than one percent of the full sample with no missing data). The second is a panel of firm × job title × state observations over the three years constructed by collapsing the data to the firm × job title × state × year level and taking the mean of the education and experience requirements. This panel has the advantage of eliminating variation due to changes in firm or job-title composition while also taking a more conservative approach to high-frequency posters by weighting each firm × job title × state × year cell equally.[15] Summary statistics for both of these alternate samples show upskilling trends that are similar to when we aggregate the data

[14] As a robustness check, Table B3 in the online appendix also reports results from specifications that use the other skill levels and find results consistent with the upskilling hypothesis.

[15] Table A4 in the online appendix compares the industry and occupation distributions for the panel sample to that of the full sample with no missing data. Although certain industries and occupations account for a greater share of postings, there are no significant differences in the distributions across the two samples.

into state × occupation cells.

B. Changes in Labor Market Slack

Our basic empirical strategy is to explore the relationship between changes in employer skill requirements and changes in local labor market conditions over time.[16] To do this, we use two primary measures of labor market slack. The first measure is the state unemployment rate as reported by the Bureau of Labor Statistics. The second measure is a labor supply/demand ratio that varies within both state and broad occupation group. This measure is modeled on the Conference Board's Labor Supply/Demand Ratio for their Help Wanted On-Line (HWOL) dataset and is the ratio of the number of unemployed individuals relative to the number of job postings for a given state and occupation group. The numerator is estimated using data on unemployed individuals by state and six broad occupation groups from the American Community Survey (ACS).[17] The denominator is calculated using the number of BGT job postings by state and broad occupation group.

C. Changes in Veteran Labor Supply

As a source of exogenous variation in the number of job searchers, we make use of a natural experiment resulting from the large increase in the post-9/11 veteran labor force following troop withdrawals from Iraq and Afghanistan. The U.S. began withdrawing troops from these countries in 2009, and by 2012 approximately 1.6 million veterans had returned home and left active duty. Less than one quarter of

[16] We acknowledge that a large share of hires come from job-to-job transitions, and that this form of hiring varies more over the business cycle. While it would also be interesting to quantify the impact of job applications directly, we do not have access to such data.
[17] These groupings are similar to the major groups used in Current Population Survey and are the same as those used by HWOL to construct their labor supply/demand ratios. The BGT measure is highly correlated with the HWOL measure and provides similar results when used in our specifications. See the online appendix for further details.

veterans separating from the military during this period were disabled or retired, and more than half had applied for unemployment benefits.[18] As of 2010, the national unemployment rate for post-9/11 veterans who had recently served in Iraq or Afghanistan was 14.3 percent compared to 11.4 percent for veterans serving in other locations and only 9.4 percent for non-veterans.[19] Moreover, returning veterans are attractive job candidates with practical training, hands-on experience, well-developed teamwork and leadership skills, and even higher educational attainment than that of the total civilian population.[20]

To capture the change in veteran labor supply over this period, we use the ACS to estimate the change in the number of post-9/11 veterans in the labor force at the state level each year from 2007 through 2012. Table 2 shows that during this period, an additional 200,000 to 300,000 post-9/11 veterans joined the U.S. labor force each year. Moreover, veteran employment is concentrated among a select group of occupations that typically make use of the specialized skill set that comes from serving in the military. As a result, military-specific occupations such as protective services (e.g., police officers and sheriffs, security guards, and firefighters) and operations specialists (e.g., aircraft mechanics, logisticians, and computer support specialists) typically receive a disproportionate share of veteran employment (see Figure A5 in the online appendix).[21] This variation across states, occupations, and years provided a sizeable

[18] Department of Veterans Affairs. 2015. "Veterans Economic Opportunity Report 2015."
[19] Bureau of Labor Statistics. 2010. "Employment Situation of Veterans—2010."
[20] The educational attainment of post-9/11 veterans is higher than that of the non-veteran population, with a significantly lower share of high school dropouts and high school graduates with no college, a significantly higher share of individuals with some college or an associate degree, and similar shares of individuals with a bachelor's degree or higher. This is true even prior to separating from the military and taking advantage of veteran benefits to attend college. See the online appendix for further details on the veteran supply shock.
[21] These occupation shares are calculated using ACS 3-year 2007 PUMS to reflect pre-recession trends.

exogenous shock to the supply of skilled job searchers per posting for state × occupation cells receiving a disproportionate share of the veteran withdrawal (see Figure A6 in the online appendix). As such, the troop withdrawal from Iraq and Afghanistan creates a natural experiment from which we can measure the response of skill requirements to increases in labor supply.

To more formally capture this targeted impact of the increase in the supply of post-9/11 veterans on the labor market, we construct four measures of changes in the supply of veteran labor across state × occupation × year cells as reported in Table 2. Three of these measures are broadly similar. The first measure is simply the log difference in the number of post-9/11 veterans in the state labor force, as reported in the ACS Summary Files, multiplied by the occupation's share of veterans. One drawback to this measure is that the one-year ACS was not designed to measure high frequency changes in the number of post-9/11 veterans at the state level. As a result, the changes we measure are noisy and thus our estimates are subject to attenuation bias. An alternative approach is to take the log difference in the number of post-9/11 veterans at the national level and create state level variation by multiplying this change by each state's average share of the post-9/11 veteran labor force measured over time.[22]

Of course, the residence of veterans following the drawdown in Iraq and Afghanistan is potentially endogenous, as veterans can choose to return to places in the U.S. with better employment opportunities. This is particularly relevant during the period we are studying at the height of the Great Recession. To address this issue, we once again construct an allocation-based measure based on the national change in

[22] This approach is similar to the "initial immigrant share" methodology used by Card (2009) and others to study the impact of immigrants on native workers.

post-9/11 veterans. However, instead of allocating veterans based on each state's average share of veterans by their *current residence*, we allocate them based on each state's average share of veterans by *state of birth* (Charles, Hurst, and Notowidigdo 2018; Carneiro and Lee 2011). This measure of location is plausibly exogenous, as veteran state of birth from several decades ago is not likely to be correlated with changes in the current state of the labor market. Yet, places where many veterans were born do receive a larger labor supply shock, as many veterans return home to be near family.

Finally, to compare the results from the veteran specifications to the previous OLS estimates, we construct a veteran supply/demand ratio similar to the overall labor supply/demand ratio described above. This is defined as the ratio of the number of returning veterans relative to the number of job postings for a given state and occupation group. The numerator is constructed by multiplying the national change in the number of post-9/11 veterans by each occupation group's veterans' share and each state's share of veterans' birthplace. The denominator is the same as the overall labor supply/demand ratio.

III. Empirical Approach

We seek to explore the upskilling dynamic by measuring the degree to which the observed increase in employer skill requirements is related to the supply of job seekers. We use two primary sources of variation. For the specifications using the basic labor market measures, the identifying assumption is that the Great Recession affected some states more than others, allowing us to exploit the variation in local labor markets across states and time periods. For the specifications using the veteran supply shocks, the identifying assumption is that the timing of the troop withdrawal and the veteran's state of residence or birthplace was uncorrelated with underlying trends in skill requirements at the state × occupation level.

Although the BGT data provides detailed information on education and experience requirements posted by employers, there are two key limitations that shape our empirical approach. First, about half of the postings come from job boards where the employer name is not listed. Second, no BGT data exist for 2008 and 2009. As such, we use two different empirical specifications. Our first approach is to make use of all the underlying job postings, regardless of whether employer name is listed, by aggregating the BGT data into state × occupation cells and using a stacked difference specification due to missing data for 2008 and 2009 to show that the upskilling relationship holds in the cross-section. We then restrict the sample to observations with employer name to be able to control for changes in the composition of employers over time using a fixed effects model. Rather than pick one specification that eliminates half the observations — or the other, which does not control for employer × job-title fixed effects — we use both approaches to demonstrate that the results are the same regardless of the approach or the sample that is used.[23]

Given the considerable heterogeneity in both the level of required skill and labor availability across states and occupations, we first use a stacked difference specification, similar to that used by Autor, Dorn, and Hanson (2013). The sample is completely balanced with all state × occupation cells observed in both periods, and because the regression is specified in differences, this approach effectively allows for

[23] Table B8 in the online appendix provides analogous estimates for equation (2) using the cross-sectional sample.

differential nonlinear trends in skills across occupations. Specifically, we estimate the basic cross-sectional relationship between changes in employer skill requirements and changes in labor market slack using the following OLS specification:

$$\Delta S_{o,s,t} = \alpha + \beta \Delta U_{o,s,t} + \gamma X_{o,s,t} + \theta_o + \sigma_s + \tau_t + \varepsilon_{o,s,t} \quad (1)$$

Where for occupation o, in state s over time period t:

$\Delta S_{o,s,t}$ = percentage point change in skill requirements (e.g., education or experience)

$\Delta U_{o,s,t}$ = change in the labor availability measure (e.g., change in the state unemployment rate, the supply/demand index by state and broad occupation group, or the veteran shocks)

$X_{o,s,t}$ = vector of control variables (e.g., initial level of skill required in 2007, share of the population by education or experience level in 2000)

θ_o = occupation fixed effects to capture differences across occupations

σ_s = state fixed effects to capture differences across states

τ_t= time period dummy to capture changes in the general composition of vacancies

$\varepsilon_{o,s,t}$ = a stochastic error term

Using equation (1), we examine changes in employer requirements across occupations and states over time. Our vector of control variables ($X_{o,\sigma,t}$) includes the share of the state population with a bachelor's degree in 2000 and the average age of the state population in 2000 to account for both heterogeneity in the pre-existing pool of skilled labor available to employers, as well as the initial share of openings requiring a particular skill in 2007 to account for heterogeneity across state × occupation cells. The occupation, state, and time fixed effects are used to control for pre-existing trends by state and occupation. The coefficient of interest in equation (1) is β, which measures the increase in skill requirements related to changes in the availability of labor. If β is positive and significant, this suggests that employers are upskilling education and experience requirements in response to an increase in the supply of job searchers in a particular state and occupation. One potential limitation of this approach is that it ignores changes in the composition of firms and jobs that are potentially driven by factors other than the supply of job searchers. such as changes in technology and industry demand. It also cannot account for potential changes in the representativeness of the BGT data over time.

Our second approach addresses these concerns over changes in employer composition and data quality over time by focusing on within firm × job-title changes. Specifically, we control for changes in the composition of employers over time by using a fixed effects model for the more restricted sample of observations that identify employer name. Because we control for firm × job title × state fixed effects, this is similar to a first-difference model that is driven by changes over time.[24] We do this using the following specification for both the basic labor market measures as well as our veteran labor availability measures:

$$S_{fj,s,t} = \alpha + \beta U_{o,s,t} + \gamma X_{fj,s,t} + \eta_{fj,s} + \tau_t + \varepsilon_{fj,s,t} \qquad (2)$$

where for firm-job title pair *fj*, in state *s* over time period *t*:

$S_{fj,s,t}$ = a dummy variable for requiring a particular skill (either education or experience)

U_{ost} = the labor availability measure (either the state unemployment rate, the supply/demand index by state and broad occupation group, or the veteran shocks)

$X_{fj,s,t}$= vector of control variables (e.g., initial level of skill required in 2007)

$\eta_{fj,s}$= fixed effect for each firm × job title × state combination

[24] However, note that not all state-firm-job title cells are observed in all three years of the data (2007, 2010, 2012), so we cannot use the same stacked difference model as equation (1).

τ_t = time period dummy to capture changes in the general composition of vacancies

$\varepsilon_{fj,s,t}$ = stochastic error term

By focusing on within firm × job-title changes, these specifications remove the possibility that the observed upskilling relationship is due to changes in the composition of employers or jobs observed in the data. The downside to this specification is that roughly half of the BGT postings do not include employer names and therefore are dropped in this specification. Since the panel is unbalanced across the three years, the equation is specified in levels but includes firm × job-title fixed effects to be able to compare skill requirements for the same job at the same employer over time.

IV. Results

A. Baseline OLS Specifications

In this section, we establish the basic upskilling relationship in the cross-section as described in equation (1). Recall that a positive and significant β indicates that skill requirements rose more within occupations in states experiencing a greater increase in the unemployment rate. Of course, it would be naïve to infer causality solely from these cross-sectional relationships, given the potential for serious omitted variable bias.

Table 3 reports the results of these initial regressions using the variation in basic labor market measures across states and over time for our primary dependent variables: the share of postings requiring a BA or greater and the share of postings requiring at least four years of experience.[25] Each coefficient listed is from a separate regression, and standard errors for all regressions are clustered at the state level. Column (1) regresses the change in these skill requirements on the change in the state unemployment rate and the BGT supply/demand ratio

---

[25] Table B3 in the online appendix reports results for each education and experience category separately.

including only occupation × year fixed effects to allow for differential nonlinear trends in skills across occupations. Note that this effectively accounts for roughly 60 percent of the variance in both education and experience requirements.[26] The estimates in column (2) include our baseline set of controls, which have an extremely small impact on the coefficient of interest.[27] Column (3) allows for differential trends in skills across locations by including state fixed effects, and Column (4) also allows for differential trends for each state × occupation pair. The coefficients are larger with these controls but remain within the confidence intervals of the original specification. Across all four specifications, β is positive and statistically significant, indicating an increase in the share of postings requiring more education and experience in states and occupations experiencing a greater increase in the availability of workers.

To give one a sense of the magnitude of this relationship, Figure 2 plots the change in employer skill requirements versus the change in labor market slack for all state × occupation cells. Our OLS estimates indicate that within a six-digit detailed occupation, a 1 percentage point increase in the state unemployment rate is associated with a 0.64 percentage point increase in the share of job postings requiring a bachelor's degree and a 0.84 percentage point increase in the share of job postings requiring at least four years of experience. How large is the upskilling effect in terms of economic importance? In the context of the most recent downturn, our results imply that the nationwide increase in unemployment rates between 2007 and 2010 raised education requirements within occupations by 3.2 percentage points and raised experience requirements by 4.2 percentage points, respectively. Relative to the observed increases in skill requirements reported in Table 1 during this period, our estimates suggest that changes in employer skill requirements due to the increased availability of workers during the business

[26] See the R-squared statistics listed in columns (1) and (3) on Table B2 in the online appendix.
[27] See the individual coefficients on each control from column (2) as listed in Table B2 in the online appendix.

cycle can account for up to 30 percent of the total increase for education and nearly 50 percent of the total increase for experience.[28] Finally, we also perform two robustness checks by estimating the upskilling relationship across traded versus non-traded industries and substituting HWOL for the BGT supply/demand ratio—both of which produce results that are similar in magnitude or even stronger.[29]

B. Within Firm × Job-Title Specifications

As discussed earlier, another potential worry is that there were non-random changes in employer composition and/or data quality over time. To control for these potentially confounding factors, we explore changes in employer skill requirements within an individual firm and job-title pair over time using data at the job posting level. Again, we do this using both the change in the state unemployment rate and the change in the BGT supply/demand ratio as our measures of labor market slack.

Table 4 provides estimates of the upskilling relationship from equation (2) using within firm × job-title variation for the sample of BGT job postings that include employer name. As with the OLS specifications, each coefficient listed is from a separate regression, and standard errors for all regressions are clustered at the state level. For both education and experience requirements we find a positive and significant relationship between employer skill requirements and labor market slack—even when controlling for *the same job title at the same employer in the same state*. Columns (1) and (3) only control for year and firm × job title fixed effects, while columns (2) and (4) include fixed effects by year as well as by firm × job title × state cells. Both sets of regressions account for approximately 80 percent of the variation in employer skill

[28] See Table B15 in the online appendix, which compares estimates across all other specifications.
[29] See Tables B4 and B5 in the online appendix for more details on these results.

requirements.[30] Controlling for firm × job-title pairs within a state, we find that a one percentage point increase in the state unemployment rate raises the share of jobs postings requiring a bachelor's degree by 0.505 percentage points and the share of job postings requiring at least four years of experience by 0.483 percentage points.

The firm × job-title estimates in Table 4 are not dissimilar from the OLS estimates found in Table 3. However, as previously discussed, the two specifications weight occupations and states differently, making it difficult to make direct comparisons. We address this by providing additional estimates that re-weight each firm × job title × state cell in our panel sample using the state × occupation weights from our OLS regressions. This re-weighting procedure produces coefficients that are quite similar across the two sets of results, despite being estimated using different samples and different levels of aggregation. Finally, we perform several robustness checks by substituting the BGT labor supply/demand ratio with that reported by HWOL, estimating the five-year change over the entire period 2007-2012, and using data on actual vacancies reported by the Minnesota Job Vacancy Survey—all of which produce results that are similar in magnitude or even stronger.[31]

C. *Veteran Supply Shock Specifications*

As a source of exogenous variation in the availability of skilled workers, we make use of a natural experiment resulting from the large increase in the post-9/11 veteran labor force following troop withdrawals from Iraq and Afghanistan. With this identification technique, we investigate whether state × occupation pairs receiving a greater number of returning veterans in a given period saw greater increases in skill requirements. Yet one might be concerned that the return of these veterans was potentially correlated with underlying trends or other factors shifting

[30] See the R-squared statistics listed in the second two panels on Table B5 in the online appendix.
[31] See Tables B5-B7 in the online appendix for more details on these results.

skill requirements. We find no correlation between our constructed veteran supply shock measures and prior period trends in wages and employment.[32] Furthermore, the veteran supply shock is not correlated with the initial level of skill requirements measured in either the BGT job postings or in the population.[33] These robustness checks, combined with the lack of an easy-to-articulate omitted variable bias problem, further reinforces our confidence in this identification strategy.

While the timing of the drawdown from Iraq and Afghanistan was unlikely to be correlated with prior cross-sectional trends, it is possible for the residential location of veterans to be correlated with potential confounders if chosen endogenously. Therefore, we also allocate post-9/11 veterans by their state of birth as a proxy for their state of residence (Carneiro and Lee 2011; Charles, Hurst, and Notowidigdo 2018). We also test whether the correlation between the veteran supply shock and rising skill requirements can be accounted for by trends at the state, occupation, or state × occupation level. We can even test whether this correlation exists within individual firm × job-title pairings.

Using our different measures of the veteran supply shock, Panel A of Table 5 demonstrates that there is a strong, significant, and positive relationship between the sharp increase in the supply of returning veterans and the rise in employer skill requirements for both education and experience. Moreover, controlling for state fixed effects—which we do in columns (2) and (4)—has little impact on the results, suggesting that we are not picking up some hidden, underlying state-level trend.[34] As expected, we get more precise estimates when we use national estimates of the number of returning veterans and allocate them by state of residence or state of

[32] See Tables B9 and B12 in the online appendix.
[33] See Table B10 in the online appendix.
[34] See the online appendix for specifications without the baseline controls (Table B11) as well as specifications that include state × occupation fixed effects, allowing for a trend at the state × occupation level (Table B13).

birth. The magnitudes indicate that a one-standard-deviation increase in the supply of veterans increases the share of postings requiring a bachelor's degree and the share requiring at least four years of experience by roughly 1.6 percentage points.

How do these results compare to the OLS results described in previous sections? To explore this, we use our veteran-specific supply/demand ratio. As seen in the first row of Panel B in Table 5, changes in this measure again strongly correlate with upskilling for both education and experience requirements, and the magnitude is similar to the results in Panel A. More importantly, our veteran supply/demand ratio can also be used as an instrument for the aggregate BGT labor supply/demand ratio used in Table 3. This instrument captures information about both the number of searchers and the demand by occupation and state. As such, it has sufficient power to surpass traditional weak instrument benchmarks as shown by the first stage F-statistics reported at the bottom of Panel B in Table 5. These IV estimates—using the change in the supply of veterans by *state of birth*—are, if anything, larger than those reported in Table 3. A one percentage point increase in the labor supply/demand ratio is associated with a 0.54 percentage point increase in the share of employers requiring a bachelor's degree and a 0.65 percentage point increase in the share of employers requiring four or more years of experience. In contrast, the previous OLS coefficients are likely to be attenuated due to both measurement error in the state labor supply/demand measures and endogeneity issues associated with hiring decisions made during the recession — problems that are addressed by using the IV estimates.[35] Moreover, we provide a robustness check accounting for the endogeneity of labor force participation among

[35] It is possible to benchmark how much of the attenuation is plausibly due to measurement error versus endogeneity, both of which tend to attenuate the point estimate of upskilling toward zero. Following Autor, Dorn, and Hanson (2013), we assess the degree of measurement error by substituting the HWOL labor supply/demand ratio for the BGT measure. Doing so increases the OLS coefficient on the education requirement from 0.172 to 0.262 (Table B5, column (4), accounting for 24.2 percent of the difference between the OLS and IV estimates, with the remaining difference (0.262 versus 0.544) associated with the correction for endogeneity.

returning veterans and reports results that are very similar, if not stronger, than those using our unadjusted veteran supply shock measures.[36]

Finally, in Panel C of Table 5, we combine both the within firm × job-title approach with the veteran supply shock analysis, which effectively compares upskilling within firm × job title × state groupings that received veteran shocks of different magnitudes. The coefficients in Panel C of Table 5 are similar in magnitude to those in Panel A, confirming that opportunistic upskilling occurs within individual firm × job-title cells even when using an exogenous shock to labor supply. Of course, the veteran supply shocks matter most for certain occupations and locations, which may not be generalizable. Still, the fact that this natural experiment provides results similar to our basic OLS regressions using the business cycle variation as well as the within firm × job-title fixed effects specifications strongly suggests a causal link between rising employer skill requirements and the increased supply of available workers.

There are several reasons firms may decide to raise skill requirements when confronted with an increase in the supply of job seekers. One intuitive possibility is that firms have a higher likelihood of finding a skilled worker for a mid- or low-skilled position when the unemployment rate is high. This intuition is consistent with the fact that the number of available high-skilled job seekers per vacancy is strongly correlated with the overall unemployment rate, as the composition of the pool of unemployed workers becomes more skilled during recessions (Mueller 2015). It is also consistent with the fact that the college wage premium for new hires is pro-cyclical, making it more likely that a high-skilled worker will accept a job offer for a mid- or

[36] Table B14 estimates upskilling using a measure of veterans in the labor force (excluding non-participants) as well as only non-disabled veterans. We are concerned, though, about the potential endogeneity of this margin (see Autor, Duggan, and Lyle 2011) and hence we use only the unscaled versions here.

low-skilled position during a recession.[37]

Conclusion

The persistent weakness of the U.S. labor market following the Great Recession continues to puzzle both researchers and policymakers alike. On the one hand, employers reported difficulty finding skilled workers to fill open positions, suggesting the potential for some degree of labor market mismatch across industry, occupation, or geography. Yet on the other, economists find that the lack of real wage growth observed even within industries and occupations with relatively strong demand suggests little role for labor market mismatch. More recently, the literature has explored the possibility that a decrease in "recruitment intensity" per vacancy during the recession may have led to an upward shift in the Beveridge curve such that a higher vacancy rate prevailed for a given unemployment rate during much of the recent recovery (Davis, Faberman, and Haltiwanger 2012). Yet to date, the application of this theory has been limited by the absence of direct measures of recruiting intensity across employers.

In this paper, we measure one channel along which recruitment intensity may have shifted during the Great Recession—in the skill requirements employers use to screen candidates when filling a new vacancy—and find evidence of opportunistic upskilling. Using data on millions of online job vacancy postings, we find that employer requirements rise for both education and experience when job seekers are more plentiful—even when controlling for year, occupation, and state fixed effects, among other covariates. This pattern is found using multiple measures of labor availability, and is robust to using both online job vacancy data as well as that from a state-level employer survey. Moreover, we find that unemployment-related upskilling occurs even within firm × job-title pairs, suggesting that changes in recruitment intensity do not

[37] See Figures C1-C3 in the online appendix, where we also present a simple partial-equilibrium model of the upskilling mechanism as well as some instructive correlations between upskilling and worker skills.

simply reflect a shift in the composition of employers or the positions that they seek to fill. We also use a natural experiment based on troop withdrawals from Iraq and Afghanistan as a source of exogenous variation in the availability of skilled workers and find a similar pattern of employer upskilling—again within firm × job-title pairs.

Our finding that weaker labor markets lead to rising employer skill requirements has important implications for models in labor and macroeconomics that aim to explain the dynamics of the labor market during recessions. Our estimates suggest meaningful effects such that employer upskilling during the Great Recession could potentially account for 18 percent of the total increase in education requirements and 25 percent of the increase in experience requirements observed between 2007 and 2010, with somewhat smaller impacts over the whole five-year period 2007-2012. [38] To our knowledge, these findings provide some of the first empirical evidence of a shift in recruitment intensity whereby employer skill requirements are driven—in part—by the available supply of labor. As such, search and matching models should incorporate upskilling over the business cycle, as it represents a novel feedback mechanism between labor supply and the selectivity of vacancies.

Given that our estimates do not account for most of the change in skill requirements, we recognize that opportunistic upskilling in response to the increased availability of job searchers is only one of the forces explaining the rise in skill requirements over the period we study. For example, prior research has demonstrated that changes over the business cycle are subtle and complicated because both cyclical and structural factors can interact (Jaimovich and Siu 2018; Charles, Hurst, and Notowidigdo 2016). Hershbein and Kahn (2018) find supporting evidence of

[38] This smaller magnitude over 2007-2012 is likely due to employer skill requirements reversing as the economy began to recover in some states as shown in Figure 1. For more evidence of downskilling during the recovery, see Modestino, Shoag, and Ballance (2016).

this, showing that the change in employer demand for skill responds persistently to local industry demand shocks that appear to be related to technological and capital changes that permanently affect the demand for education and experience. We believe their work complements our findings by measuring the degree to which structural forces drive employer upskilling compared to the estimates that we present in this paper that capture employer responses to the business cycle. Still, by demonstrating that employer upskilling is associated with opportunistic hiring when labor markets are slack, we provide important evidence that at least some of what is labeled as "structural mismatch" is at least partially cyclical and likely to revert.

REFERENCES

Albrecht, J. & Vroman, S. 2002. "A Matching Model with Endogenous Skill Requirements" *International Economic Review* 43.1 (2002), 283–305.

Autor, D., Dorn, D. & Hanson, G. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6), 2121–68.

Autor, D., Duggan, M. & Lyle, D. 2011. "Battle Scars? The Puzzling Decline in Employment and Rise in Disability Receipt among Vietnam Era Veterans." *American Economic Review* 101 (3), 339–44.

Bagues, M. & Labini, M. 2009. "Do On-line Labor Market Intermediaries Matter? The Impact of *AlmaLaurea* on the University-to-Work Transition." In *Studies of Labor Market Intermediation,* ed. David H. Autor, 127–54. Chicago: University of Chicago Press.

Barnichon, R. & Figura, A. 2010 "What Drives Movements in the Unemployment Rate? A Decomposition of the Beveridge Curve." FEDS Working Paper 2010–48.

Card, D. 2009. "Immigration and Inequality." Richard T. Ely Lecture, *American Economic Review: Papers & Proceedings* 99(2), 1–21.

Carneiro, P. & Lee, S. 2011. "Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000." *American Economic Review* 101(6), 2309–49.

Carnevale, A., Jayasundera, T. & Repnikov, D. 2012. *Understanding Online Job Ads Data.* Washington, D.C.: Georgetown University Center on Education and the Workforce.

Charles, K., Hurst, E. & Notowidigdo, M. 2018. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review* 108 (10), 2947–94.

Davis, S., Faberman, J. & Haltiwanger, J. 2013. "The Establishment-Level Behavior of Vacancies and Hiring," The Quarterly Journal of Economics, Oxford University Press, 128(2), 581–622.

Diamond, P. 2013. "Cyclical Unemployment, Structural Unemployment." *IMF Economic Review* 61.3 (2013): 410–55.

Diamond, P. & Şahin, A. 2015. "Shifts in the Beveridge Curve." *Research in Economics* 69(1), 18–25.

Fujita, S. & Moscarini, G. 2013. "Recall and Unemployment." *American Economic Review* 107(12) 3875–916.

Hall, R. & Schulhofer-Wohl, S. 2015. "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers," *American Economic Journal: Macroeconomics* 10(1), 1–32.

Hershbein, B. & Kahn, L. 2018. “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings.” *American Economic Review* 108 (7), 1737–72.

Jaimovich, N. & Siu, H. 2018. “Job Polarization and Jobless Recoveries.” Review of Economics and Statistics (forthcoming).

Kaas, L. & Kircher, P. 2015. “Efficient Firm Dynamics in a Frictional Labor Market.” *American Economic Review* 105(10), 3030–60.

Kuhn, P. & Skuterud, M. 2004. “Internet Job Search and Unemployment Durations.” *American Economic Review* 94(1), 218–32.

Marinescu, I. 2017. “The general equilibrium impacts of unemployment insurance: Evidence from a large online job board.” Journal of Public Economics, 150, 14–29.

Marinescu, I. & Wolthoff, R. 2016. “Opening the Black Box of the Matching Function: The Power of Words.” NBER Working Paper No. 22508.

Modestino, A., Moss, P. & Shoag, D. 2017. “Upskilling During the Great Recession: Why Do Employers Demand Greater Skill When Workers Are Plentiful?” Report, Russell Sage Foundation.

Modestino, A., Shoag, D. & Ballance, J. 2016. “Downskilling: changes in employer skill requirements over the business cycle.” *Labour Economics* 41 C, 333–47.

Mongey, S., Violante, G. & Gavazza, A. 2015. “What Shifts the Beveridge Curve? Recruiting Intensity and Financial Shocks,” 2015 Meeting Papers 1079, Society for Economic Dynamics.

Mukoyama T., Patterson C. & Şahin A. 2018. Job Search Behavior over the Business Cycle. *American Economic Journal: Macroeconomics* 10(1), 190–215.

Oreopoulos, P., von Wachter, T. & Heisz, A. 2012. “The Short- and Long-Term Career Effects of Graduating in a Recession.” *American Economic Journal: Applied Economics* 4 (1): 1-29.

Rothstein, J. 2012. The Labor Market Four Years into the Crisis: Assessing Structural Explanations. *Industrial Labor Relations Review* 65(3), 467–500.

Şahin, A., Song, J., Topa, G. & Violante, G. 2014. “Mismatch Unemployment.” *American Economic Review* 104(11), 3529–64.

Shimer, R. 2012. “Reassessing the Ins and Outs of Unemployment,” *Review of Economic Dynamics, Elsevier for the Society for Economic Dynamics*, no. 15(2), 127–148.

Templin, T. & Hirsch, L. 2013. “Do Online Job Ads Predict Hiring?” New York: New York City Labor Market Information Services.

Figure 1. Aggregate Relationship between Changes in Labor Market Slack and Employer Skill Requirements

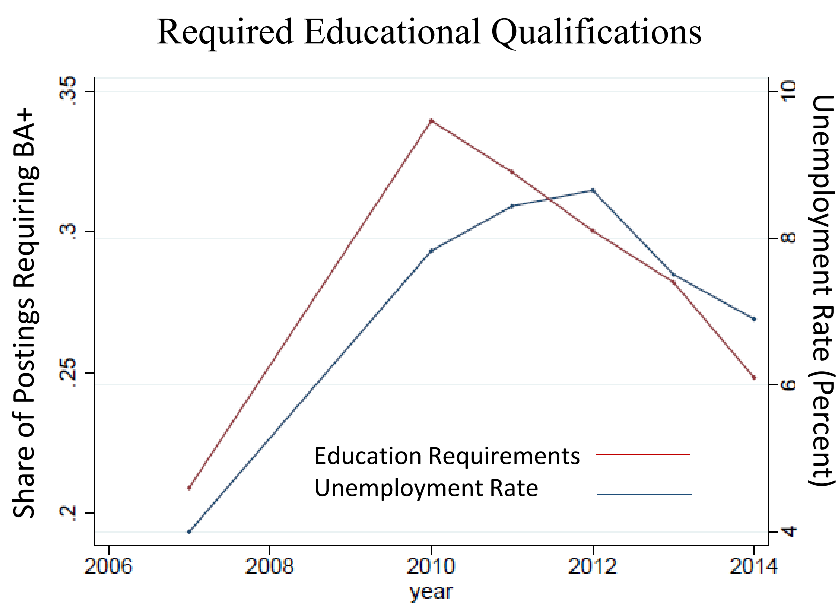

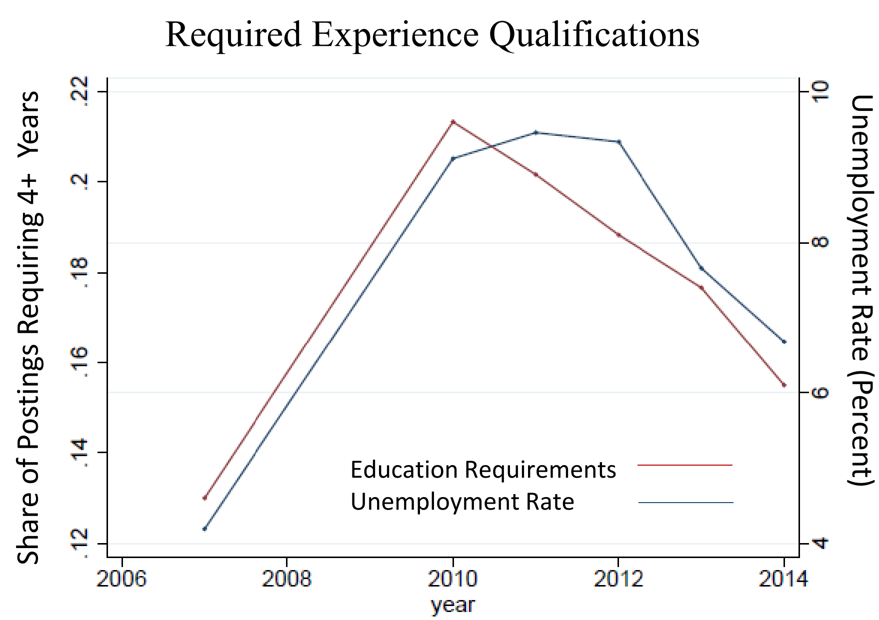

Note: Authors' calculations using data from Burning Glass Technologies and Bureau of Labor Statistics, 2007–2014.

Figure 2. Relationship between Changes in Labor Market Slack and Employer Requirements
By State and Occupation

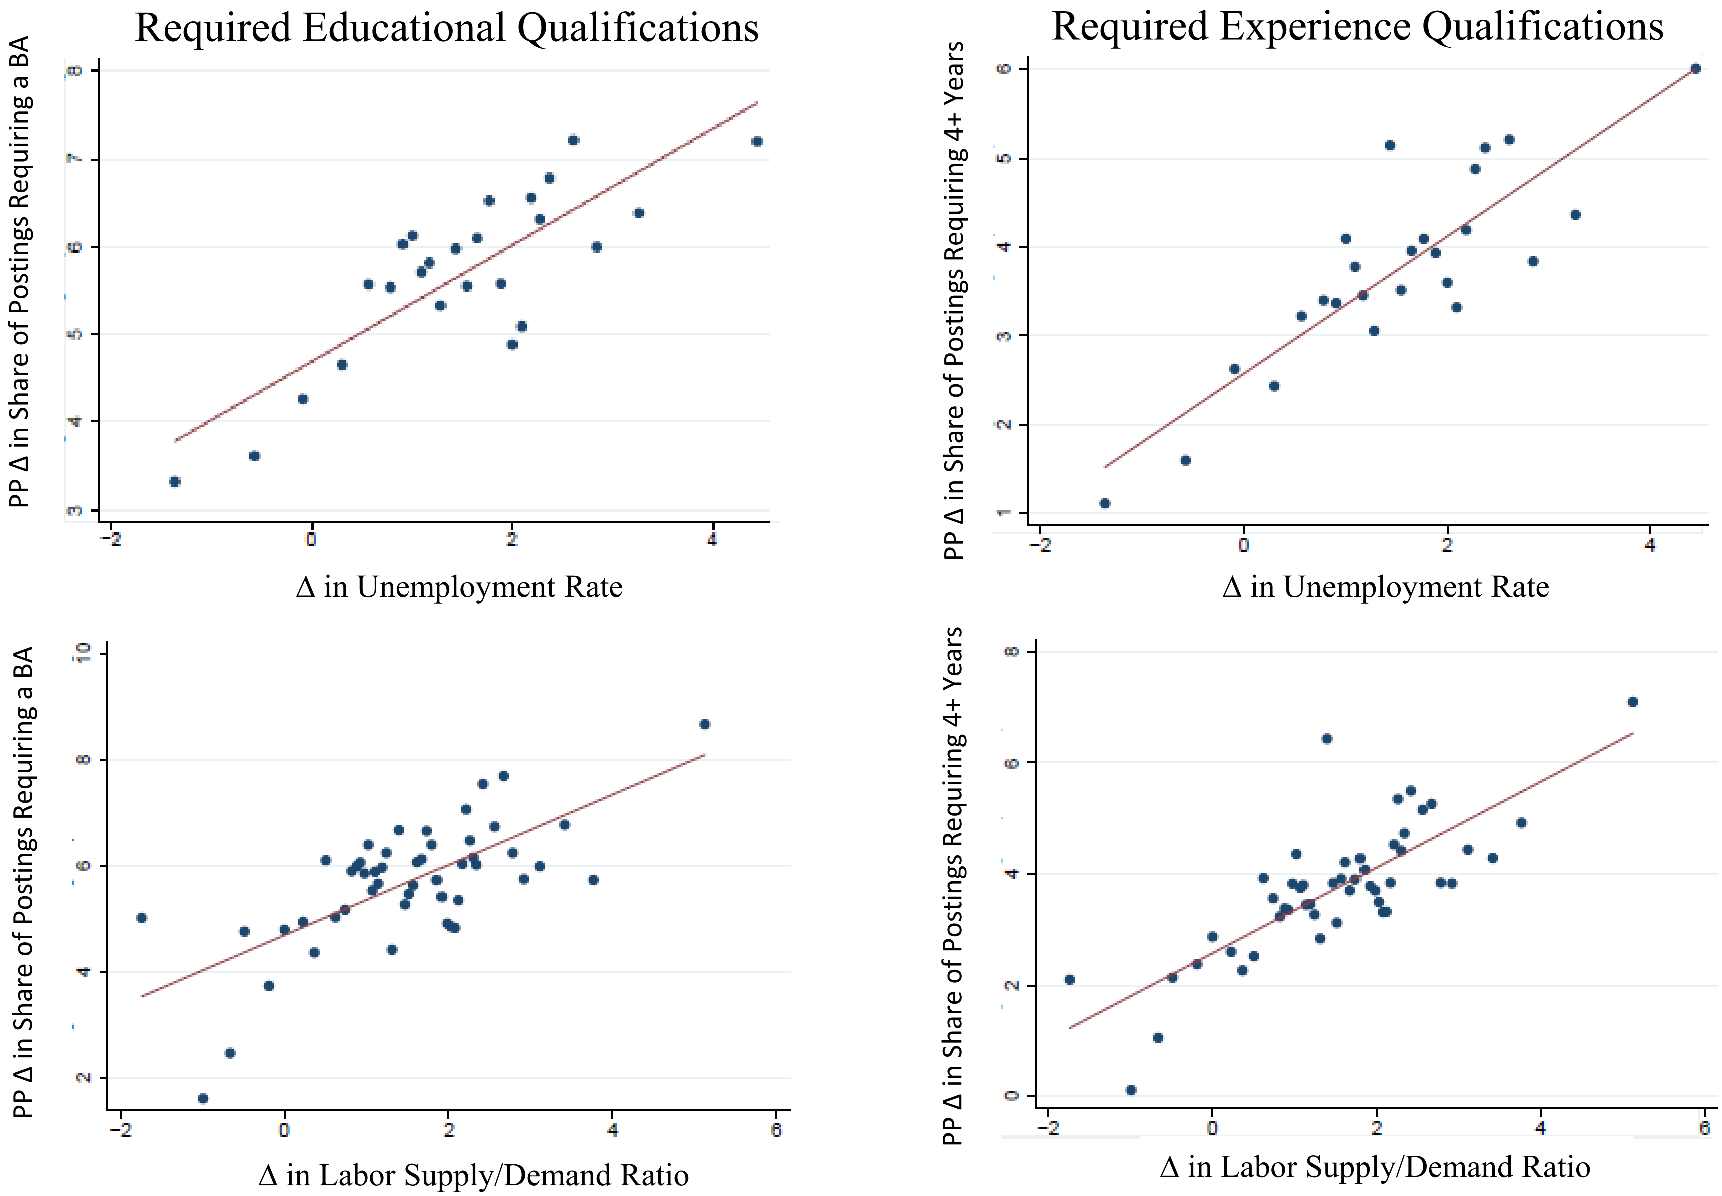

Note: Figures are binned scatterplots showing the baseline relationship between the percentage point change in employer requirements and the change in labor market slack (the BGT labor supply/demand ratio or percentage point change in the state unemployment rate).

Table 1. Summary Statistics for Employer Skill Requirements and Labor Market Slack

	2007	2010	2012	Δ 2007-2010	Δ 2010-2012
Panel A. Employer Skill Requirements					
Sample 1: Cross-sectional sample of all job postings aggregated by state × occupation cells					
Number of observations	18,694.00	19,470.00	18,970.00	18,694.00	18,970.00
Total number of job postings					
Mean	634.40	516.73	625.19	-97.46	95.49
Standard deviation	1,977.38	1,701.25	2,052.94	579.56	497.86
Mean percent of job postings requesting:					
Bachelor's degree or higher	12.93	22.43	24.95	9.83	2.34
Four or more years of experience	8.40	14.17	14.97	6.04	0.72
Sample 2: Cross-sectional sample of job postings with non-missing employer names					
Total number of job postings	5,684,530.00	4,441,522.00	6,973,321.00	-1,243,008.00	1,288,791.00
Percent of job postings requesting:					
Bachelor's degree or higher	23.23	33.38	34.42	10.15	1.04
Four or more years of experience	14.02	22.59	21.34	8.56	-1.25
Sample 3: Panel sample of repeated employer × job title × state observations					
Total number of job postings	506,914.00	295,783.00	474,884.00	-211,131.00	179,101.00
Percent of job postings requesting:					
Bachelor's degree or higher	20.98	31.07	31.85	10.09	0.78
Four or more years of experience	11.30	17.33	16.89	6.03	-0.44
Panel B. Labor Market Slack					
Unemployment rate by state					
Mean	4.43	9.07	7.59	4.64	-1.5
Standard deviation	0.87	1.95	1.61	1.51	0.76
BGT labor supply/demand ratio by state × occupation					
Mean	12.82	19.27	12.77	4.21	-4.33
Standard deviation	14.63	18.09	11.64	5.44	5.73

Note: Employer skill requirements are constructed using online job posting data provided by Burning Glass Technologies. The first sample uses data from all 36.2 million job postings aggregated into state × 6-digit Standard Occupation Code (SOC) cells containing at least 15 total postings. The second sample uses data from the subset of job postings that identify employer name and do not have more than 50 postings for the same employer × job title × state cell within a given year. The third sample includes only postings for which a given employer × job title × state cell is observed for more than one year. In all three of the above samples, we exclude postings that are missing a state FIPS code or are located in Washington D.C., Guam, and Puerto Rico. The state unemployment rate is the annual rate reported by the Bureau of Labor Statistics. The labor supply/demand ratio is an annual, state × occupation measure that is constructed following the methodology used by the Conference Board's Help Wanted OnLine Supply/Demand Index. Specifically, it is the ratio of the number of unemployed persons to the number of job postings by state for six broad occupation groups. See the data appendix for further details on sample and variable construction.

Table 2. Summary Statistics for Veteran Supply Shock Measures

Panel A: Annual Change in Post-9/11 Veteran Population, 2006-2012

Year	Number of Post-9/11 Veterans in the Labor Force	Change
2006	1,504,807	
2007	1,537,363	32,556
2008	1,559,495	22,132
2009	1,619,193	59,698
2010	1,927,541	308,348
2011	2,126,179	198,638
2012	2,330,987	204,808

Panel B: Constructed Veteran Supply Shocks

	Δ2007–10				Δ2010–12			
	Mean	**SD**	**Min**	**Max**	**Mean**	**SD**	**Min**	**Max**
Δ Log Post-9/11 Veterans by State × Occ Vet Share	0.3022	0.1379	-0.0682	0.6469	0.2440	0.1420	-0.1608	0.8052
Δ Log Post-9/11 Vets × State Vet Share × Occ Vet Share (Allocated By Residence)	0.7567	1.8045	0.0000	66.1589	0.6359	1.5128	0.0000	55.5906
Δ Log Post-9/11 Vets × State Vet Share × Occ Vet Share (Allocated By State of Birth)	0.7683	1.7782	0.0000	68.2561	0.6451	1.4889	0.0000	57.3528
Δ BGT Supply/Demand Ratio for Post-9/11 Veterans (Allocated by State of Birth)	0.1863	0.5284	-1.5770	1.6864	0.0405	0.5608	-1.5638	1.7363

Note: Panel A reports the change in the number of post-9/11 veterans in the labor force at the state level each year from 2007 through 2012 from the American Community Survey (ACS) Summary Files. Panel B reports four measures of changes in the supply of veteran labor across state × occupation × year cells constructed from the ACS. The first measure is the log difference in the number of post-9/11 veterans in the state labor force multiplied by the occupation's share of veteran employment estimated from the 2005-2007 ACS 3-year estimates. The second measure is the log difference in the number of post-9/11 veterans at the national level, multiplied by the state average share of the post-9/11 veteran labor force allocated by state of residence in each year as calculated from the ACS, multiplied by the occupation's share of veteran employment. The third measure recalculates the second measure using state of birth in place of state of residenceas calculated from the ACS. The fourth measure is a veteran supply/demand ratio which is constructed as the number of post-9/11 veterans in the labor force multiplied by an occupation's share of employees who are veterans divided by the broad occupation group's total job postings. See the online appendix for further details on variable construction.

Table 3. Relationship Between Changes in Employer Skill Requirements and Changes in Labor Market Slack, 2007-2012

	(1)	(2)	(3)	(4)
	Percentage Point Change in the Share of Postings Requiring a Bachelor's Degree or Higher			
Δ State Unemployment Rate	0.698 ***	0.636 ***	0.907 ***	0.904 ***
	(0.225)	(0.203)	(0.244)	(0.243)
Δ BGT Labor Supply/Demand Ratio	0.125 ***	0.100 ***	0.105 ***	0.172 ***
	(0.0278)	(0.0310)	(0.0361)	(0.0472)
	Percentage Point Change in the Share of Postings Requiring 4 or More Years of Experience			
Δ State Unemployment Rate	0.837 ***	0.837 ***	0.960 ***	0.959 ***
	(0.143)	(0.135)	(0.172)	(0.171)
Δ BGT Labor Supply/Demand Ratio	0.174 ***	0.173 ***	0.151 ***	0.247 ***
	(0.0284)	(0.0263)	(0.0410)	(0.0525)
Occupation × Year Fixed Effects	Yes	Yes	Yes	Yes
Baseline Controls	No	Yes	Yes	Yes
State Fixed Effects	No	No	Yes	No
State × Occupation Fixed Effects	No	No	No	Yes
Number of observations	37,664	37,664	37,664	37,664

Note: Each coefficient listed is from a separate regression as specified by equation (1) in the text. Each regression uses the first sample listed on Table 1 that aggregated data from all 36.2 million job postings into state × 6-digit Standard Occupation Code (SOC) cells containing at least 15 total postings (for both years over which the change is measured), excluding those that are missing a state FIPS code or are located in Washington D.C., Guam, and Puerto Rico. Observations are weighted by the occupation's share of each state's total postings. Baseline controls include the initial (2007) share of postings requiring the skill measured as well as either the share of the state population with a Bachelor's Degree or greater in 2000 (for Panel A) or the average age of the population in 2000 (for Panel B). Standard errors (in parentheses) are clustered by state. * p<0.10, ** p<0.05, *** p<0.01.

Table 4. Relationship Between Changes in Employer Skill Requirements and Changes in Labor Market Slack, Using within Firm x Job Title Variation

	(1)	(2)	(3)	(4)	(5)	(6)
	Share of Postings Requiring a Bachelor's Degree or Greater			**Share of Postings Requiring 4 or More Years of Experience**		
	Panel A: Cross-sectional sample of job postings with non-missing employer names (N=17,099,373)					
State Unemployment Rate	0.156 **	0.318 *	-----	0.106 **	0.257 *	-----
	(0.074)	(0.176)		(0.045)	(0.133)	
BGT Labor Supply/Demand Ratio	0.056 **	0.079 ***	-----	0.044 **	0.075 ***	-----
	(0.026)	(0.021)		(0.019)	(0.023)	
	Panel B: Panel sample of repeated employer × job title × state observations (N=1,277,581)					
State Unemployment Rate	0.235 **	0.505 **	0.726 ***	0.204 **	0.483 ***	0.765 ***
	(0.120)	(0.195)	(0.240)	(0.175)	(0.153)	(0.166)
BGT Labor Supply/Demand Ratio	0.103 ***	0.113 ***	0.082 **	0.096 ***	0.106 ***	0.099 **
	(0.030)	(0.027)	(0.045)	(0.031)	(0.029)	(0.035)
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm × Job-Title Fixed Effects	Yes	No	No	Yes	No	No
Firm × Job-Title × State Fixed Effects	No	Yes	Yes	No	Yes	Yes
Using State × Occupation weights from OLS specifications	No	No	Yes	No	No	Yes

Note: Each coefficient listed is from a separate regression as specified by equation (2) in the text. Regressions using the BGT supply/demand ratio also include state × occupation fixed effects. Panel A uses the cross-sectional sample of job postings with non-missing employer names that do not have more than 50 postings for the same firm × job title × state cell within a given month and year. Panel B uses the panel sample of repeated employer × job title × state observations for which a given firm × job title × state cell is observed for more than one year. Both samples exclude postings that are missing state FIPS code or are located in Washington D.C., Guam, and Puerto Rico. Standard errors (in parentheses) are clustered by state. * p<0.10, ** p<0.05, *** p<0.01.

Table 5. Relationship Between Changes in Employer Skill Requirements and Veteran Supply Shock

	(1)	(2)	(3)	(4)
Panel A: OLS Estimates Using Veteran Supply Shock	**Change in the Share of Postings Requiring a Bachelor's Degree or Higher**		**Change in the Share of Postings Requiring 4 or More Years of Experience**	
Δ Log Post-9/11 Veterans by State × Occ Vet Share	1.382 *	1.590 *	1.674 *	1.580 *
	(0.771)	(0.857)	(0.844)	(0.905)
Δ Log Post-9/11 Veterans × State Vet Share x Occ Vet Share	0.421 **	0.364 ***	0.696 ***	0.469 ***
(Allocated by State of Residence)	(0.167)	(0.0964)	(0.160)	(0.106)
Δ Log Post-9/11 Veterans × State Vet Share x Occ Vet Share	0.477 **	0.312 ***	0.564 ***	0.392 ***
(Allocated by State of Birth)	(0.193)	(0.0914)	(0.132)	(0.0976)
Baseline Controls	Yes	Yes	Yes	Yes
Occupation, Year Fixed Effects	Yes	Yes	Yes	Yes
State Fixed Effects	No	Yes	No	Yes
Number of Observations	37,664	37,664	37,664	37,664
Panel B: IV Estimates Using Veteran Supply Shock	**Change in the Share of Postings Requiring a Bachelor's Degree or Higher**		**Change in the Share of Postings Requiring 4 or More Years of Experience**	
Δ BGT Supply/Demand Ratio for Post-9/11 Veterans	1.380 ***	1.317 **	1.840 ***	1.823 ***
(Allocated by State of Birth)	(0.514)	(0.621)	(0.506)	(0.646)
Δ BGT Supply/Demand Ratio	0.498 ***	0.544 **	0.582 ***	0.645 ***
IV with Δ BGT Supply/Demand Ratio for Post-9/11 Veterans	(0.153)	(0.204)	(0.142)	(0.208)
(Allocated by State of Birth)				
First Stage F-Statistic (for Δ BGT Supply/Demand Ratio)	16.4	21.7	16.4	21.7
Baseline Controls	Yes	Yes	Yes	Yes
Occupation, Year Fixed Effects	Yes	Yes	Yes	Yes
State Fixed Effects	No	Yes	No	Yes
Number of Observations	37,664	37,664	37,664	37,664
Panel C: Firm-Job Title Estimates for Veteran Supply Shock	**Change in the Share of Postings Requesting a Bachelor's Degree or Higher**		**Change in the Share of Postings Requesting 4 or More Years of Experience**	
Log Post-9/11 Veterans by State × Occ Vet Share	0.478 **		0.886 ***	
	(0.237)		(0.244)	
Log Post-9/11 Veterans × State Vet Share x Occ Vet Share	0.758 ***		1.201 ***	
(Allocated by State of Birth)	(0.204)		(0.319)	
Year Fixed Effects	Yes		Yes	
Firm × Job Title × State Fixed Effects	Yes		Yes	
Number of Observations	1,277,581		1,277,581	

Note: Each coefficient listed is from a separate regression. The first two panels use the same state × occupation sample as Table 3 while the third panel uses the same firm × job title sample as Table 4. Panel A uses the first three alternative measures of the veteran supply shock as described in Table 2. Panel B uses the fourth measure of the veteran supply shock--the veteran-specific supply/demand ratio from Table 2--and also uses this mesaure as an instrument for the previous BGT supply/demand ratio found in Table 1. Panel C uses the first and third measures of the veteran supply shock as described in Table 2. First stage F-statistics demonstrating the absence of weak instrument bias are reported for these IV regressions. See the online appendix for more detail on sample and variable construction. Standard errors (in parentheses) clustered by state. * p<0.10, ** p<0.05, *** p<0.01.