Do Recessions Accelerate Routine-Biased Technological Change?

Evidence from Vacancy Postings

Brad Hershbein
Upjohn Institute and IZA

Lisa B. Kahn
Yale University NBER and IZA

July 16, 2016

Abstract

Routine-biased technological change (RBTC), whereby routine-task jobs are replaced by machines and overseas labor, shifts demand toward high- and low-skill jobs, resulting in job polarization of the U.S. labor market. We test whether recessions accelerate this process. Exploiting a new database containing the near-universe of electronic job vacancy postings and cross-sectional variation in the MSA-level employment shock generated by the Great Recession, we estimates changes in skill requirements within occupations and firms. We find that postings in hard-hit metro areas have substantially larger increases in education, experience, cognitive, and computer skill requirements in 2010, relative to 2007, and that these increases persist through the end of our sample in 2015. We find important roles for both within-firm changes in skill demand and upskilling in firms that did not post in 2007. We also show that among publicly-traded firms in our data, those that upskill more also increase capital stock by more over the same time period. We argue that upskilling is driven primarily by firm restructuring of production toward more-skilled workers. Our result is unlikely to be driven by firms' opportunistically seeking to hire more-skilled workers in a slack labor market, and we rule out other cyclical explanations. We thus present the first direct evidence that the Great Recession precipitated new technological adoption.

JEL Classifications: D22, E32, J23, J24, M51, O33

---

*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, Joe Altonji, David Autor, Tim Bartik, David Berger, Jeff Clemens, David Deming, David Green, John Horton, Peter Kuhn, Fabian Lange, Steve Malliaris, Alicia Sasser Modestino, Daniel Shoag, Henry Siu, Basit Zafar, and seminar participants at the AEAs 2016, the Atlanta Fed, Brookings, Harvard, the LSE, MIT, Princeton, Rutgers, SOLE/EALE 2015 meetings, Trans Pacific Labor Seminar 2015, University of British Columbia, University of Chicago, University of Illinois, University of Notre Dame, University of South Florida, and the WEAI 2016 meetings. We are especially indebted to Dan Restuccia, Jake Sherman, and Bledi Taska for providing the Burning Glass data. Jing Cai provided excellent research assistance with CPS data. A previous version of this paper was titled "Is College the New High School? Evidence from Vacancy Postings."

1 Introduction

The employment shift from occupations in the middle of the skill distribution toward those in the tails is one of the most important trends in the U.S. labor market over the last 30 years. Previous research makes a compelling case that a primary driver of this job polarization is routine-biased technological change (RBTC), whereby new machine technologies and overseas labor substitute for middle-skill jobs in the U.S. and are in turn complementary to high-skill cognitive jobs.[1] Until recently, RBTC had been thought to be a gradual, secular phenomenon. However, a long theoretical literature beginning with Schumpeter's "creative destruction" (1939) suggests adjustments to technological change may be more episodic. In boom times, high opportunity costs, or frictions such as adjustment costs, may inhibit resources from being reallocated optimally in the face of technological change. Recessions lower the opportunity cost and can produce large enough shocks to overcome these frictions.[2]

Whether adjustments to new technology are smooth or lumpy is important for policy and for our understanding of recoveries. The recoveries from the last three U.S. recessions (1991, 2001, 2007–09) have been jobless: employment was slow to rebound following the recession despite recovery in aggregate output. The reasons for jobless recovery are not well understood, but a small theoretical literature points to adjustment costs as a potential mechanism, since they generate reallocation that is concentrated in downturns (Berger 2012, Koenders and Rogerson 2005, Jaimovich and Siu 2012). Such lumpy adjustment may leave a mass of displaced workers with the wrong skills for new production. Jaimovich and Siu (2015) provide suggestive evidence that countercyclical reallocation, in the form of RBTC, and jobless recovery are linked. They show that the vast majority of the declines in middle-skill employment have occurred during recessions and that, over the same time period, recovery was jobless only in these occupations. However, there is no direct evidence on how firms restructure in the face of technological change, and whether it is gradual or episodic. This paper aims to fill that gap.

In this paper we investigate how the demand for skills changes over the business cycle. We use a new data set collected by Burning Glass Technologies that contains the near-universe of electronically posted job vacancies in U.S. cities in 2007 and 2010–2015. Exploiting spatial variation in economic conditions, we establish a new fact: the skill requirements of job ads increase in metropolitan statistical areas (MSAs) that suffered larger employment shocks in the Great Recession, relative to the same areas before the shock and other areas that experienced smaller shocks. Our estimates imply that ads posted in a hard-hit metro area are about 5 percentage points (16%) more likely to contain education and experience require-

[1] See for example the seminal work of Autor, Levy, and Murnane (2003); Autor, Katz, and Kearney (2006, 2008); Goos and Manning (2007); and Autor and Dorn (2013).

[2] Many theoretical papers predict this phenomenon. See for example Hall (1991, 2005); Mortensen and Pissarides (1994); Caballero and Hammour (1994, 1996); Gomes, Greenwood, and Rebelo (2001); and Koenders and Rogerson (2005).

ments and about 3 percentage points (8-12%) more likely to state requirements for cognitive and computer skills. Drawing on the richness of our data, we show that the majority of this "upskilling" is driven by increases in skill requirements within firms and occupations, rather than a shift in the distribution of ads posted across firms or occupations. We also show a similar upskilling effect in employment, using American Community Survey (ACS) data, suggesting that firms successfully hire the more-skilled workers that they seek.

We then examine whether upskilling is driven by changes to production in a manner consistent with RBTC. In short, are firms changing what they do or simply changing whom they hire? For example, upskilling may instead occur because firms temporarily and opportunistically take advantage of a slack market to try to attract workers typically found higher up on the job ladder. As our data are best-suited to measure a shift in demand from middle- to high-skilled workers, rather than concomitant demand increases for low-skilled labor that also could be expected from RBTC, we concentrate our evidence on this margin. We present several analyses supporting the notion that accelerated RBTC brought on by the Great Recession drives upskilling.

First, we show that upskilling persists at nearly the same magnitude from early in the recovery through 2015. Even though most measures of local labor market conditions have converged back to their pre-recession levels by the end of our sample period, elevated skill requirements—especially higher in the occupational skill distribution where our data are most comprehensive—have not. We also document that this persistence is driven by the same firms that upskilled by 2010. Any purely cyclical explanation is thus insufficient to explain upskilling. Second, we show that upskilling occurs within firms and that among publicly traded firms in our data, those with larger increases in skill requirements also had larger increases in capital stock over the same time period, consistent with a substitution of routine-task workers with machines. Third, we find that MSAs where the Great Recession more severely affected college graduates experience an additional upskilling effect early in the recovery, suggesting that firms may attempt to hire "up" in a slack labor market. However, these effects converge back to their pre-recession levels, while the effect of the overall employment shock persists. Moreover, while there is some evidence that upskilling effects tend to fade for occupations at the bottom of the skill distribution, they are persistent in the top half of the occupational skill distribution, where our data are most comprehensive. Both findings are at odds with the notion that opportunistic upskilling on the part of firms drives our result. Finally, we rule out that changes in labor supply due to worker quits and formal schooling decisions drive upskilling.

Taken together, our results suggest that firms located in areas hit harder by the Great Recession were induced to restructure their production towards greater use of machines (or outsourced labor) and higher-skilled workers; that is, the Great Recession hastened the polarization of the U.S. labor market.

This paper is related to a number of important literatures. First, we provide direct evidence that recessions accelerate firm-level responses to technological change. This is consistent with the important, but suggestive, evidence provided by Jaimovich and Siu (2015) that 88% of the job loss in routine-task occupations since the mid-1980s has occurred around the time of an NBER-dated recession. The direct, demand-side evidence provided in our paper speaks to the many models in macroeconomics that assume adjustment costs and predict that recessions will be times of cleansing (Schumpeter 1939, Koenders and Rogerson 2005, and Berger 2012). These models are important for explaining business cycle dynamics, but have so far lacked strong empirical evidence.

Second, the Burning Glass job postings data provide a unique opportunity to measure changes in skill requirements both across and *within* occupations. In contrast, the extant literature on job polarization has focused on shifts across occupations and has therefore been unable to ascertain the importance of the intra-occupational margin. We find that, in response to worse local labor market conditions, demand for education, experience, and cognitive and computer skills increase within several different occupation groups, and throughout the skill distribution. We therefore present the first evidence, to our knowledge, that RBTC occurs within occupations in addition to the well-known adjustments across occupations.

This result helps to clarify work by Beaudry, Green, and Sand (2014 and 2016) and others documenting the "great reversal" in demand for cognitive skill. They show that since 2000, cognitive occupations have seen no gains in employment or wages, and that college graduates have become more likely to work in routine occupations than previously. They argue that a decrease in demand for cognitive occupations drove college graduates to take jobs lower in the occupational distribution, squeezing out the high school graduates who formerly held them. This is something of a puzzle, especially given the common belief that technological change continues and more-skilled workers earn a premium in the labor market (Card, Heining, and Kline 2013; Card, Cardoso, and Kline, forthcoming). Because we find that upskilling occurs within these occupations as well, we hypothesize part of the solution to this puzzle is that cognitive workers are being drawn into (formerly) routine-task occupations as the skill content of these occupations evolves. These changes make the occupations more-skilled and therefore likely more desirable than before, although probably still not as desirable as traditional high-skilled jobs.[3]

Third, we contribute to a growing literature exploiting data on vacancy postings. 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 information on the characteristics of a given vacancy or the firm that is posting it. Fewer studies have used vacancy data that contain information on the occupation or specific requirements

[3]Our analyses, however, do not explain why employment and wages have not grown in high-skill occupations. Deming (2015) proposes a compelling hypothesis that a rising importance of social skills, especially in conjunction with cognitive skills, can help account for this fact.

of the job posted, and these have generally used aggregate data (Sasser Modestino, Shoag, and Ballance 2016), narrow slices of the data (Rothwell 2014), or data that are limited to one vacancy source (Marinescu 2014, Kuhn and Shen 2013). To our knowledge, we are the first study to use data based on a near-universe of online job postings that covers every metropolitan area in the United States.

Fourth, we help answer an important question about the consequences of recessions for workers. It is well known that low-skilled workers suffer worse employment and wage consequences in recessions.[4] Evidence from past recessions shows that in downturns workers are more likely to take worse jobs, relative to their skills (Devereux 2002; Kahn 2010; von Wachter, Oreopoulos, and Heisz 2012). Some of this may result because workers apply to a broader set of jobs, even those for which they would normally be overqualified, when job-finding rates decline. However, it is also possible that firms actively seek a more-skilled worker than they could have attracted in a tighter market.[5] I.

Indeed, in a related recent paper, Sasser Modestino, Shoag, and Ballance (2016) use a small subset of vacancy data from 2007, 2010, and 2012 originating from the same sources as our micro-level data, but aggregated to the state level for only select occupations. They too find evidence of upskilling and argue that it is driven entirely by firms opportunistically seeking more skilled workers in a slack labor market. Our analysis differs in several key ways and we reach a different conclusion. First, we analyze data through the end of 2015 and show that since upskilling persists it is unlikely to be driven by purely cyclical phenomena, such as temporarily slack labor markets. Second, our data contain not just education and experience requirements (as in Sasser Modestino et al. 2016), but also specific stated skill preferences, which allow us to examine demand for cognitive and computer skills, arguably the most crucial skills required of a modern workforce. Third, we analyze the exhaustive set of occupations, and therefore provide a more thorough characterization of upskilling and important new facts on technological change within occupations. Fourth, firm identifiers in our data allow us to show that upskilling is driven both by within-firm changes in skill requirements and by changes in the composition of firms hiring, likely driven by firm deaths and births surrounding the Great Recession.[6] We also show that within-firm changes in skill requirements are accompanied by increases in capital stock, which is necessary for mechanization. Our analysis thus provides substantial contributions relative to the simultaneous work of Sasser Modestino et al. (2016), and allows us to conclude that technological change is indeed an important driver of upskilling.

In this paper, we demonstrate that during the Great Recession firms changed not only

[4]See for example Hoynes, Miller, and Schaller (2012), Forsythe(2016), and von Wachter and Handwerker (2009).

[5]There is mounting evidence that firms at the bottom of the job ladder benefit from increased retention of their incumbent workforce in a downturn (Moscarini and Postel-Vinay 2012, Kahn and McEntarfer 2015), but little is known about hiring dynamics.

[6]In our data, we do not observe firm deaths directly; rather, we observe which firms posted job ads to an online source before and after the recession.

whom they would hire in the recovery, but how they would produce. Instead of occurring gradually, with relatively few workers needing to be reallocated at any given time, technological change is episodic, resulting in a swath of displaced workers whose skills are suddenly rendered obsolete as firms ratchet up their requirements. The need to reallocate workers on such a large scale may help drive jobless recoveries. It also likely plays a role in the well-noted and marked decline in male employment-to-population ratios over the past 25 years, especially since these declines have been stair-step around recessions (Moffitt 2012).[7] The evidence provided in this paper is thus integral for understanding worker reallocation, and can help inform policymakers about the optimal mix during a downturn of worker retraining and subsidizing job search through unemployment insurance.

The remainder of this paper proceeds as follows. Section 2 introduces the data. Section 3 presents new facts on upskilling as a function of local labor market conditions. Section 4 offers additional evidence on persistence and heterogeneity that support episodic RBTC as the primary driver of upskilling. Section 5 rules out alternative explanations for our findings, and section 6 concludes.

2 Data

Our data come from a unique source: microdata from 87 million electronic job postings in the United States that span the Great Recession (between 2007 and 2015). These job postings were collected and assembled by Burning Glass Technologies, an employment analytics and labor market information firm. In this section, we describe the data and our particular sample construction. We provide a detailed examination of the sample's characteristics and representativeness in Appendix A.

2.1 Burning Glass Overview

Burning Glass Technologies (henceforth BG or Burning Glass) examines some 40,000 online job boards and company websites to aggregate the job postings, parse and deduplicate them into a systematic, machine-readable form, and create labor market analytic products. Thanks to the breadth of this coverage, BG believes the resulting database captures a near-universe of jobs that were posted online. Through a special agreement, we obtained an extract from BG, which covers every MSA in the United States in 2007 and from 2010 through 2015.[8]

[7] Supporting the notion that episodic restructuring drives stair-step declines in male employment, Foote and Ryan (2015) point out that middle-skill workers, the most vulnerable to RBTC, are most at risk of leaving the labor force when unemployed.

[8] The database unfortunately lacks postings from 2008 and 2009. Our extract was provided in February 2016. We also have data on jobs posted in Micropolitan Statistical Areas, which we do not use for lack of some of the labor market indicators in these areas, and substantial noise in the ones that are available. They represent 5.6% of all posted ads.

The two key advantages of our data are its breadth and detail. The broad coverage of the database presents a substantial strength over data sets based on a single vacancy source, such as CareerBuilder.com. While the JOLTS asks a nationally representative sample of employers about vacancies they wish to fill in the near term, it is typically available only at aggregated levels, and contains relatively little information about the characteristics of vacancies. In contrast, the BG data contain some 70 possible standardized fields for each vacancy. We exploit detailed information on occupation, geography, skill requirements, and firm identifiers. The codified skills include stated education and experience requirements, as well as thousands of specific skills standardized from open text in each job posting.[9] The data thus allow for analysis of a key, but largely unexplored, margin of firm demand: skill requirements within occupation.[10] Moreover, they allow for a firm-level analysis, which, as we show below, is key to disentangling between cyclical and structural explanations for upskilling.

However, the richness of the BG data comes with a few shortcomings. Notably, the database covers only vacancies posted on the Internet. Even though vacancies for available jobs have increasingly appeared online instead of in traditional sources, such as newspapers, one may worry that the types of jobs posted online are not representative of all openings. In Appendix A, we provide a detailed description of the industry-occupation mix of vacancies in BG relative to other sources (JOLTS, the Current Population Survey, and Occupational Employment Statistics), an analysis of how it has changed over our sample period, and various validity checks conducted on the data both by us and by other researchers. To briefly summarize, although BG postings are disproportionately concentrated in occupations and industries that typically require greater skill, the distributions are relatively stable across time, and the aggregate and industry trends in the quantity of vacancies track other sources reasonably closely.[11]

The data contain the occupation of the posting (at the 6-digit Standard Occupation Classification 2010 (SOC) level) and codes identifying the MSA where it is located. Burning Glass also collects the firm name, if available, for a given posting. Employer name is missing in 40% of postings, primarily from those listed on recruiting websites that typically do not list the employer.[12]

[9]For example, an ad might ask for a worker who is bilingual or who can organize and manage a team. BG cleans and codes these and other skills into a taxonomy of thousands of unique, but standardized requirements. Beginning with a set of pre-defined possible skills, BG searches text in an ad for an indication that the skill is required. For example, for team work, they search for the key words "team work" but also look for variations such as "ability to work as a team."

[10]Other private-sector firms, such as Wanted Analytics, used by the Conference Board's Help-Wanted Online Index, also offer disaggregated data, but not skill requirements.

[11]Appendix A also discusses how BG deduplicates multiple postings for the same vacancy. It is important to note that although BG's parsing algorithm has changed over time, each iteration is applied to all postings data, so our data set is consistent over our sample period.

[12]When name is available, Burning Glass uses a proprietary algorithm to group name variants into a standard set: for example, "Bausch and Lomb", "Bausch Lomb", and "Bausch & Lomb" would be grouped together. We also perform some additional cleaning on firm name, removing any remaining punctuation and

We restrict our main sample to ads with non-missing employers that posted at least 10 ads over the sample period of 2007 and 2010-2015. After cleaning, our data contain 157,356 distinct employers.[13] Many of our analyses exploit firm-level information to distinguish among possible mechanisms for upskilling. We therefore choose to focus our entire analysis on the consistent sample of ads with non-missing firms, with a sufficient number of observations per firm to estimate firm-level characteristics. However, we have performed analyses not requiring firm-level information on the full data set and obtain very similar results. Moreover, as we discuss later, we have confirmed that the probability of satisfying this sample criterion (having a valid firm identifier) does not vary over the business cycle. Thus, our sample restriction should not confound the estimated relationship between local labor market conditions and the skill requirements of postings.

2.2 Skill Requirements in Burning Glass

In our analyses, we exploit four categories of skill requirements: stated education and experience requirements, stated demand for skills that we classify as "cognitive," and stated demand for computer skills. The RBTC literature emphasizes that new information technology or cheap overseas labor substitute for routine, algorithmic, middle-skill tasks. These new technologies are in turn complementary with high-skill cognitive, abstract tasks and may indirectly affect low-skill, manual tasks. As we note in the appendix, a downside of the BG sample is that low-skill jobs are underrepresented. We thus focus our analysis on the degree to which employers shift demand toward high-skill tasks and workers. High-skilled workers favored by RBTC may be required to work with computers and perform a more versatile set of functions. Indeed, the non-algorithmic tasks that complement routine-task performing machines or overseas labor, involve more complexity, problem solving and analytical skills, and determining which tasks need to be performed at a given moment.

In accord with human capital theory, we believe more-educated workers or those with some amount of experience on the job will be better able to perform these functions.[14] We take the skill measures for cognitive and computer from Deming and Kahn (2016), who categorize the BG open text fields into key groupings to measure firm heterogeneity in skill demands. The cognitive skill requirement is based on words associated with non-routine analytical tasks, using the taxonomy developed in Autor, Levy, and Murnane (2003) and

a few problematic words, such as "Incorporated" (sometimes listed as "Inc").

[13]The 10-ad restriction drops about 4% of job ads that list a firm name. However, employer names with very few ads are likely to be miscoded (for example, capturing a fragment of the city name).

[14]In the raw data, there are two fields each for education and experience requirements: a minimum level (degree or years of experience) and a preferred level. Postings that do not list an education or experience requirement have these fields set to missing. We use the fields for the minimum levels to generate variables for the presence of an education or experience requirement as well as the number of years of education or experience required; the minimum is much more commonly specified than the preferred, and it is always available when a preferred level is listed.

subsequently used by the majority of papers studying RBTC and polarization.[15] The computer skill requirement is based on the words "computer" and "spreadsheet" as well as specific software (e.g., Excel, PowerPoint, AutoCAD) and computer languages (e.g., Java, SQL, Python). We are particularly interested in whether a given ad contains any of these requirements, although we also examine the level of requirements for education and experience.

Table 1 summarizes data for the primary regression sample.[16] In 2007, 34% of the weighted ads list any education requirement (column 1, row 1). Among ads with an education requirement, half (17% of all ads) specify minimum education of a bachelor's degree, another quarter ask for a high school diploma, and the remainder are roughly evenly split between master's degrees and professional degrees or Ph.D.s. Converting the degrees to their modal equivalent years of schooling, the average education requirement, conditional on one being specified, is 14.76 years.

The second column shows skill requirements averaged over 2010–2015. The third column shows the within-MSA change in skill requirements across the two sample periods, and indicates statistical significance. The share of ads specifying an education requirement increased by 23 percentage points, on average. This is roughly evenly split across ads requiring high school and ads requiring college; because the proportional increase is slightly larger for high school, the overall (conditional) years of schooling falls slightly. All differences in means are statistically significant at the 1% level.

Experience requirements follow a very similar pattern to education requirements. In 2007, almost one-third of ads specify some amount of experience in the field. Among ads with a requirement, the vast majority ask for less than five years, with much of the remainder asking for between five and 10 years. Conditional on posting an experience requirement, the average ad asks for 3.5 years. In the later time period, the propensity to specify an experience requirement increases by 20 percentage points. These increases are again concentrated in the lower categories so that the average, conditional on specifying any requirement, falls by about one-quarter of a year.

Finally, in 2007, 73% of weighted ads specify at least one specific, text-based skill requirement. Among these, nearly one-quarter specify a cognitive skill requirement, and nearly one-third have a computer requirement. In 2010–2015, 91% of ads have at least one text-based skill requirement, with proportional increases between one-third and one-half in the share specifying a cognitive or computer requirement.[17]

These increases in stated skill demand could be driven by the national recession that took place between 2007 and the 2010–2015 period, which would be consistent with our hypothesis.

[15] An ad is categorized as requesting a cognitive skill if any skills requested include the following phrases: "problem solving," "research," "analytical," "critical thinking," "math," or "statistics."

[16] In the top two panels, observations are weighted as they are in our regression analyses: we give equal weight to ads within an MSA-year, but upweight larger MSAs, based on the size of the labor force in 2006.

[17] In regression analyses, we use the probability of posting a cognitive or computer requirement, conditional on posting a specific text-based skill, as dependent variables, rather than the unconditional probabilities, which might instead pick up a tendency for ads to become more verbose as postings costs decline.

Table 1: Summary Statistics

	mean (sd) in: 2007	2010-15	Change
Education Requirements:			
Any	0.34	0.57	0.23***
	(0.06)	*(0.05)*	
HS	0.09	0.20	0.10***
	(0.03)	*(0.05)*	
BA	0.17	0.27	0.10***
	(0.05)	*(0.08)*	
>BA	0.03	0.05	0.02***
	(0.01)	*(0.01)*	
Years, Conditional on any	14.76	14.48	-0.27***
	(0.39)	*(0.47)*	
Experience Requirements:			
Any	0.32	0.52	0.20***
	(0.06)	*(0.07)*	
0-3	0.13	0.24	0.11***
	(0.03)	*(0.03)*	
3-5	0.14	0.21	0.07***
	(0.03)	*(0.04)*	
>5	0.05	0.08	0.03***
	(0.02)	*(0.04)*	
Years, Conditional on any	3.51	3.24	-0.27***
	(0.45)	*(0.53)*	
Skill Requirements:			
Any Stated Skills	0.73	0.91	0.18***
	(0.05)	*(0.04)*	
Cognitive, conditional on any	0.22	0.33	0.11***
	(0.05)	*(0.06)*	
Computer, conditional on any	0.34	0.48	0.14***
	(0.08)	*(0.11)*	
share of ads:			
Missing ACS match	0.08		
Non-matching 2007 firm-MSA	0.46		
Not in Compustat and 2007 sample	0.29		
	Mean	Min	Max
# MSAs	381		
Posts per MSA-year	21,686	132	1,220,115
# Occupations (4-digit)	108		
Posts per occupation-MSA-year	227	1	192,645
# Firms	157,356		
Posts per Firm-MSA-year	13	1	16,413

Notes: Burning Glass data 2007 and 2010-2015. Sample is restricted to ads with non-missing firms that posted at least 10 ads over our sample period. In the top two panels, summary statistics are calculated weighting by the occupation or firm's ad share in the MSA-year times the size of the MSA labor force in 2006. *** indicates means are statistically significantly different from each other at the 1% level.

However, they could also be driven by a variety of other factors, such as changing composition of firms posting ads online or pre-existing national trends. Because of these issues and the relatively short panel we have to work with, our regression analyses always control for year dummies. We therefore fully absorb the overall change in skill requirements illustrated in table 1. Instead, we identify differences in the change in skill requirements across metro areas as a function of how they weathered the Great Recession.

3 Methodology

We estimate regressions of the form specified in equation (1), where *skill* is any of several measures of the average skill requirements (discussed in more detail below) in MSA, m, and year t ($2010 \leq t \leq 2015$), $shock_m$ is a measure of the local employment shock generated by the Great Recession, I^t are year dummies, X_m is a vector of MSA controls, and ε_{mt} is an error term.

$$skill_{mt} - skill_{m07} = \alpha_0 + [shock_m * I^t]\alpha_1 + I^t + X_m\beta + \varepsilon_{mt} \tag{1}$$

The regression thus estimates the impact of the local employment shock generated by the Great Recession on the *change* in skill requirements for a given MSA between 2007 and each year of 2010 through 2015. Such a specification allows us to empirically investigate the timing and persistence of upskilling in relation to local labor market shocks. It also accords with several macroeconomic models that posit episodic restructuring or production when the opportunity cost of doing so is lower. For instance, standard (s,S) models with convex adjustment costs predict lumpy adjustment to a changing environment, especially when conditions may alter worker-firm match surplus or the bargaining power of workers and firm management (Berger 2012). Relatedly, "pit stop" models of the business cycle express the notion that productivity-enhancing improvements are more likely to take place in downturns because of the relatively lower opportunity cost of doing so in times when profits are relatively low (Hall 2005, Koenders and Rogerson 2005). Finally, there could be asymmetries in the costs and benefits of laying off workers. The stigma or bad publicity to a firm making layoffs should decline in a downturn, and the losses in terms of firm-specific human capital that accompany layoffs might be outweighed by a general need to cut costs.

The key explanatory variable, $shock_m$, is the MSA-specific change in projected year-over-year employment growth from 2006 to 2009, the national peak and trough years surrounding the Great Recession. We project employment growth in an MSA based on its employment shares in 3-digit NAICS industry codes in 2004 and 2005 and national employment changes at the 3-digit industry level. This type of shift-share method is sometimes referred to as a "Bartik shock," following the strategy of Bartik (1991).[18] We use the Bartik measure,

[18] Specifically, we define Bartik employment growth as $\sum_{k=1}^{K} \phi_{m,k,\tau}(lnE_{kt} - lnE_{k,t-12})$, where for K 3-

instead of actual employment growth (as reported by the Bureau of Labor Statistics), for two reasons. First, actual employment growth at the MSA level is measured with substantial error, while the Bartik measure allows for more precision. Second, actual employment growth will reflect shocks to labor demand as well as other city-specific shocks, including those to labor supply, which may be problematic.[19] We note that other direct measures of local labor market tightness, such as the local unemployment rate, have similar shortcomings in terms of measurement error or reverse causality; for instance, an unemployment rate may be high precisely *because* a sudden demand shift towards more-skilled labor generates structural mismatch.[20]

Since we control for year fixed effects (I^t), we identify the key coefficients, α_1, purely off of differences across metro areas in the employment shock generated by the Great Recession, rather than relying on the national shock itself. As noted above, this is a necessity given our short panel. We would have difficulty separately identifying coincident national trends, such as changes in the use of online job boards. Interacting $shock_m$ with an exhaustive set of year dummies allows us to estimate the impact of the local employment shock on the change in skill requirements for any given year, relative to 2007. $shock_m$ ranges from about -0.12 to -0.04 across MSAs, but, in practice, to make the coefficients easier to interpret, we norm this variable so that a one unit change is equal to the difference across the 10th and 90th percentile MSAs, -0.026 log points.[21]

The first-difference specification implicitly controls for differences across MSAs in posted skill requirements. Also, in order to understand changes in skill requirements both within and across firms and jobs, we disaggregate our data to estimate changes in skill requirements within occupation-MSA, firm-MSA, and firm-occupation-MSA cells. We categorize occupations based on 4-digit SOC codes.[22] These analyses are described in more detail below.

Although our cross-sectional identification does not permit the use of MSA fixed effects, we do control for a wide range of MSA characteristics, including demographics, ed-

digit industries, ϕ is the employment share of industry, k, in MSA, m, at time τ (in practice, the average of 2004 and 2005) and lnE_{kt} is the log of national employment in industry k in year-month t and lnE_{kt-12} is the log of national employment in industry k in the same calendar month one year prior, both seasonally adjusted. We obtain national employment for each 3-digit industry from Current Employment Statistics. We construct ϕ using County Business Patterns data and the algorithm of Isserman and Westervelt (2006) to overcome data suppressions; the resulting county-level statistics are mapped to MSAs using the definitions provided by the Census Bureau and set by the Office of Management and Budget. See http://www.census.gov/population/metro/data/def.html. Other papers utilizing Bartik shocks include Blanchard and Katz (1992) and Notowidigdo (2013). The measure may be used directly as a regressor (reduced-form) or as an instrument for observed employment growth; in practice, this choice often does not matter much, and that is also true in our case.

[19] For example, MSAs with secular increases in population due to migration flows may experience employment changes that are higher than average but still have a weakening labor market. The Bartik shock addresses this issue.

[20] Nonetheless, our estimates are qualitatively robust to using actual employment growth or the unemployment rate, although magnitudes are somewhat reduced, likely for the reasons indicated.

[21] Note, this normalization also implies that a larger value corresponds to a worse employment shock.

[22] We choose the 4-digit occupation aggregation primarily to save on computation time, but also to ensure larger cell sizes. Results are robust to disaggregating to the 6-digit level.

ucational attainment, and economic indicators, obtained from the American Community Survey (ACS), averaging years 2005 and 2006.[23] These controls help adjust for differences across MSAs in their preexisting tendency to upskill, to the extent that such a tendency is correlated with the skill distribution of the population or the health of its economy before the Great Recession.

We cluster standard errors by MSA to address possible serial correlation within an area.[24] Finally, we weight observations in this regression by the size of the MSA labor force in 2006. This weighting scheme allows us to upweight areas with larger populations, helping with precision, while fixing the weight applied to each MSA-year. The latter ensures that we identify off of the same MSA weighting mix in each year, regardless of the overall changes in ads posted.[25]

From the bottom panel of table 1, we estimate these regressions using all 381 MSAs, which contain an (unweighted) average of 21,686 posts per MSA-year. When we disaggregate to the four-digit occupation level, we have 108 occupations represented, with an average of 227 posts in each MSA-occupation-year. Finally, as noted above, our data contain more than 150,000 unique firms, which translate into an average of 13 posts in each firm-MSA-year.

4 Skill Requirements and Local Employment Conditions

4.1 Main Results

Figure 1 summarizes regression results from equation (1) for our four main dependent variables: the change in the share of ads postings any education requirement, any experience requirement, any cognitive requirement, and any computer skills requirement. The figures plot the estimated impact of the Bartik shock on the change in skill requirements for each year, relative to 2007 (coefficients α_1), as well as 90% confidence intervals. Point estimates underlying the graphs are shown in column 1 of appendix tables A1a and A1b.

Beginning with the top left panel, we find that the probability of specifying any education requirement increases by 5.4 percentage points, relative to the average requirement for ads

[23]We chose years just prior to the Great Recession and that allow for MSA identification (prior to 2005, the ACS lacks substate identifiers). Specifically, we include the share of the population that is female, black, Hispanic, Asian, married, had migrated in the last year, is a high school drop out, has exactly a high school diploma, has some college, has exactly a bachelor's degree, is enrolled in school, is less than age 18, is age 19–29, is age 30–39, is age 40–49, and is age 50–64. We also control for the employment-to-population ratio and the average weekly wage of full-time workers. We can match all but 8% of weighted ads to the ACS (see the middle panel of table 1), with the unmatched consisting of small MSAs. In such cases, we set the ACS controls to zero and include an indicator for not matching.

[24]We have also estimated regressions clustering at the MSA-date level, which is the level of variation underlying α_1, and obtain substantially smaller standard errors.

[25]When we disaggregate across occupations and/or firms within MSAs, we weight by the occupation, firm or firm-occupation's ad share in an MSA-date times the MSA labor force. Because of our weighting scheme, the more aggregate regressions produce results identical to those using more disaggregated data when the underlying specification is the same.

Figure 1: Impact of MSA-Specific Employment Shock on Skill Requirements

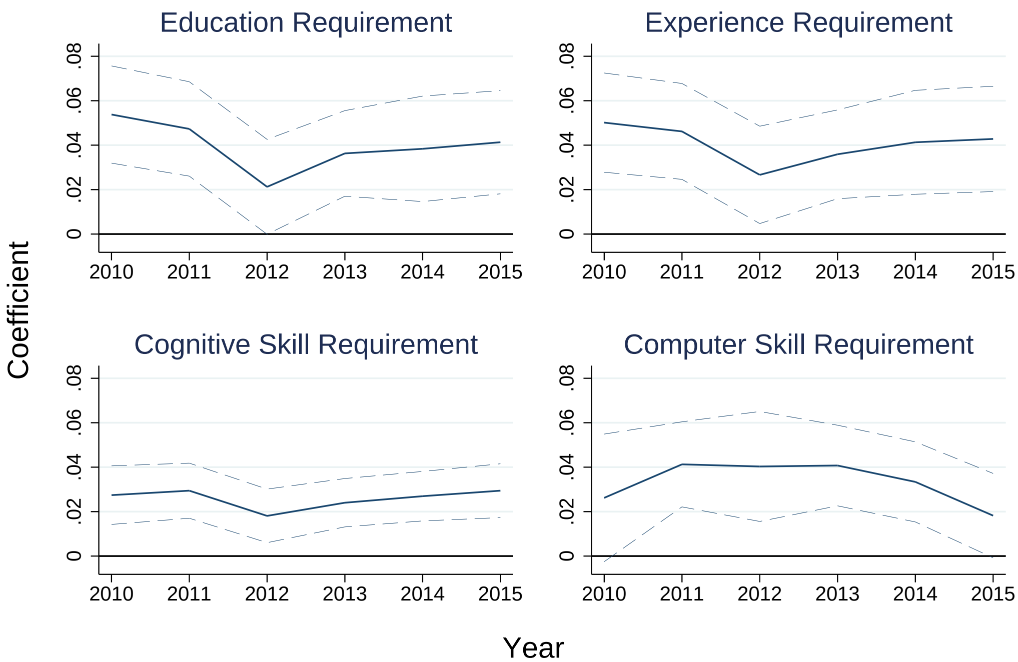

Graphs plot coefficients on MSA employment shock interacted with year, and 90% CI's (dashed lines).
Regressions also control for year fixed effects and MSA characteristics.

posted in the same MSA in 2007, for an MSA experiencing a large employment shock (90th percentile), compared to an MSA experiencing a small shock (10th percentile). This increase is 16% of the average requirement in 2007 and is significant at the 1% level. The effect persists at fairly similar magnitudes and significance levels for subsequent years, with a small dip in 2012. In 2015, we estimate that the probability of posting an education requirement is still 4.1 percentage points larger than it was in 2007 for MSAs at the 90th percentile of employment shocks, relative to the same difference for MSAs at the 10th percentile. That is, 76% of the initial upskilling effect in 2010 remains five years later. Estimates in each year except 2012 are significant at the 1% level.

The remaining panels of figure 1 exhibit remarkably similar patterns in both magnitudes and statistical significance. The probability of listing an experience requirement increases by 5.0 percentage points (16%), and 85% of this increase remains in 2015. The probability of listing a cognitive requirement increases by 2.7 percentage points (12%), and the gap across cities widens slightly by 2015. Finally, the probability of listing a computer skill requirement increases by 2.9 percentage points (7.7%) in more-affected, relative to less-affected, MSAs, and 69% of this increase remains in 2015.

Appendix figure A4 summarizes results for additional education and experience outcomes in order to understand changes in the intensive margin for these requirements. In figure A4a, there are similar sized effects for increases in the probability of requiring a high school diploma and a bachelor's degree. These increases offset each other, resulting in no overall change in the years of education required, conditional on posting any requirement (bottom right). Also, there is no change in the propensity to require a graduate degree. In some ways this is reassuring, since many professional jobs, such as lawyers and doctors, always require higher degrees; these requirements would not change with improvements in technology. Figure A4b exhibits a similar pattern for experience requirements. We see increases in both low experience requirements (2 years or less) and middle requirements (3–5 years). Beyond that we see little change. Again, this results in the distribution of requirements, conditional on posting any, changing little, with the average year requirement essentially unchanged (bottom right). Hereafter, we continue with the main four dependent variables to explore the heterogeneity and sensitivity of our estimates.

For contrast, figure 2 summarizes estimates of equation (1) for a range of local labor market statistics as dependent variables.[26] The top left panel shows the unemployment rate, and our estimate implies that a hard-hit MSA (90th percentile Bartik shock) experiences an increase in the unemployment rate between 2007 and 2010 that is 2 percentage points greater than a less hard-hit MSA (10th percentile). Over time, the impact of the shock declines in magnitude, and by 2015, unemployment rate differentials across MSAs have converged back to their pre-recession levels. We find similar convergence in employment levels (top right panel).[27]

The bottom left panel shows employment-to-population ratio estimates, which run only through 2014 because of data availability. We find that a hard-hit MSA experiences a 1.25 percentage point larger decline in its employment-to-population ratio in 2010, and only about 15% of this gap had narrowed by 2014.[28] There could be many reasons for this lack of convergence. For example, in our story, rapid adoption of new technologies in hard-hit MSAs over this time period could render a swath of worker skills obsolete, inducing labor force exit (also see Foote and Ryan 2015). This could be especially true among older workers for whom retraining may not be worthwhile. Alternatively, it could also reflect our age controls not fully capturing retirement effects from aging Baby Boomers. We thus also provide estimates using the employment-to-population ratio among the prime-aged population (age 25–55) in the bottom right panel. Here we see similar initial effects in 2010 but find that two-thirds of the gap across cities had converged by 2015.

Clearly, MSA labor markets had not fully recovered from the Great Recession by the end of 2015. Yet, on most measures, there was substantial progress after 2010, with cities

[26] MSA unemployment rates and employment are obtained from the BLS Local Area Unemployment Statistics and Current Employment Statistics programs, respectively. Employment-to-population ratios are based on the authors' calculations from the ACS. At the time of this writing, data for 2015 had not been released.

[27] That is, we find that the portion of unemployment rates and employment predicted by Bartik employment growth converged back to their pre-recession levels by 2015, holding constant pre-recession MSA characteristics. Actual differences across cities were somewhat slower to converge.

[28] These findings are consistent with Yagan (2016) who uses IRS tax data to show that while unemployment rates had converged aross U.S. Commuting Zones, employment probabilities had not, holding constant a rich set of worker characteristics.

Figure 2: Impact of MSA-Specific Employment Shock on Labor Market Variables

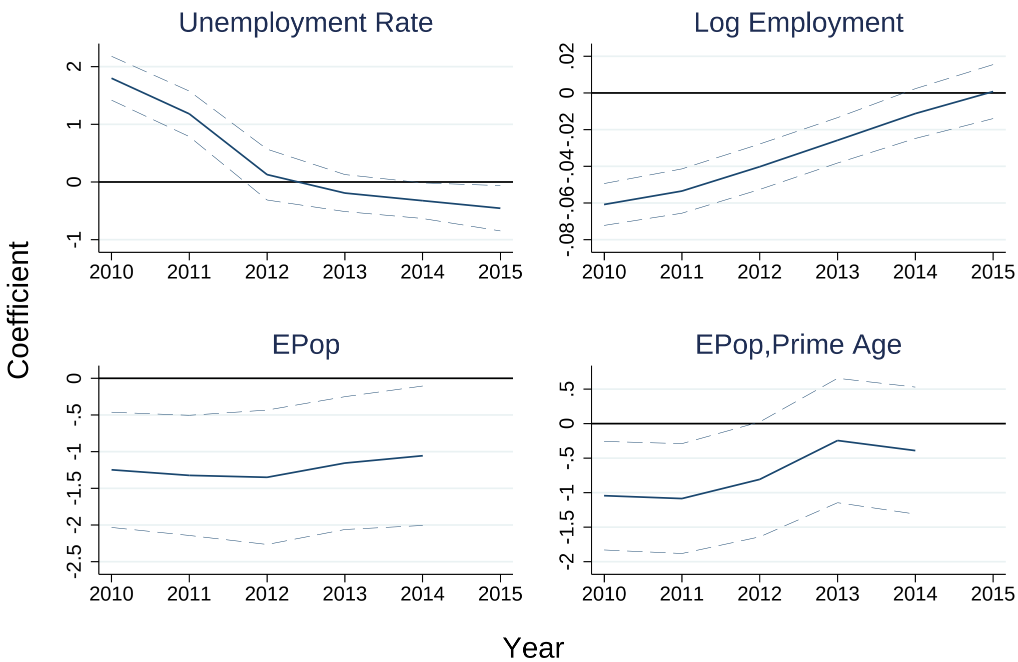

Graphs plot coefficients on MSA employment shock interacted with year, and 90% CI's (dashed lines).
Regressions also control for year fixed effects and MSA characteristics.

moving closer to their pre-recession differences. At the same time, in figure 1 we see almost no convergence in skill requirements. As of 2015, MSAs that experienced severe shocks in the Great Recession still look different from how they appeared in 2007, and different from other cities that experienced weaker shocks.

4.2 Heterogeneity within and across Occupations

Why do skill requirements change differentially across MSAs? It is possible that these differences result from selection: perhaps in harder-hit MSAs only higher-skilled jobs survive in the early recovery. An episodic RBTC-based explanation involving changes in production technology, however, would imply upskilling within, as well as across, occupations. For example, community and social service specialists at a food bank in Washington, D.C. might be required not only to interact with clients to assist with food security, but may have to understand and use database software and GIS, as well, to better serve them (McCoy 2016). Simultaneously, venerable journalistic organizations such as the *New York Times* now hire scientists, not reporters, to be chief data officers to reach and understand online readers (Greenfield 2014).

To investigate whether MSA-level changes in skill requirements are prevalent within and/or across occupations, we generate counterfactual skill distributions and estimate regressions like equation (1), using as the dependent variable the change in a counterfactual skill distribution from the actual 2007 MSA-level skill requirement. We plot the coefficients

Figure 3: Within and Across Occupation Changes in Skill Requirements

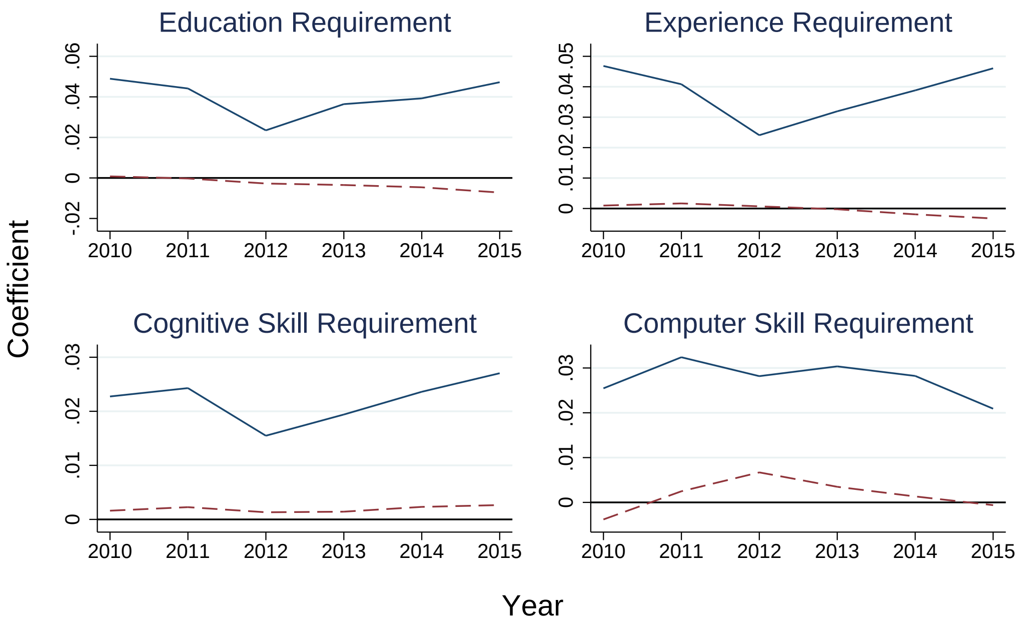

Blue solid line = within-occ-MSA change, red dashed line = across-occ-MSA change

Graphs plot impact of MSA shock on only change in each element, holding constant the other element at its 2007 level. Regressions also control for year fixed effects and MSA characteristics

on $shock_m I^t$ in figure 3. The blue solid line presents results where we hold constant the occupation distribution of ads within an MSA in 2007 and allow only the within occupation-MSA skill requirement to change. The red dashed line, in contrast, holds constant the within occupation-MSA skill requirement at its 2007 level and allows the occupation distribution of ads within the MSA to vary with time.[29]

As can be seen, the entirety of the effect is concentrated on the within-occupation margin, with no role for the shifting distribution of ads across occupations. Skill requirements increase in hard-hit MSAs because for the same 4-digit occupations, ads are more likely to list education, experience, cognitive skill, and computer skill requirements.[30]

Appendix tables A1a and A1b report regression specifications for a variety of within-occupation-MSA changes in skill requirements, based on equation (2).[31]

$$skill_{omt} - skill_{om07} = \alpha_0 + [shock_m * I^t]\alpha_1 + I^t + X_m\beta + \varepsilon_{omt} \tag{2}$$

[29] Virtually all ads posted in the 2010–2015 period are in occupation-MSAs that also posted in 2007, so we can generate change variables. The 0.36% of ads that cannot be matched back are dropped from these analyses.

[30] These skill distributions hold constant either the within–occupation-MSA or the across–occupation-MSA element at its 2007 level and allow the other element to change from 2007 to t. If we instead hold constant one element at its level in t, we obtain very similar results. These are shown in appendix figure A5.

[31] Occupations are categorized at the 4-digit SOC level.

Columns labeled 1 summarize the baseline specification.[32] Columns labeled 2 estimate equation (2) on the full BG sample, including ads that do not list the firm name. As mentioned above, we focus on the 60% of ads that contain firm name because this sample allows us to distinguish among cyclical and structural reasons for upskilling, as described in the next section. However, results for all dependent variables on the expanded sample are quite similar, if slightly smaller in magnitude.[33]

Column 3 augments the within occupation-MSA specification on our primary sample to include occupation fixed effects, while column 4 also ads occupation-specific time trends. Given the first-difference specification, these allow occupations to systematically differ in their change in skill requirements from 2007 and (in the latter case) the slope of that change. These could be important if some occupations are both more likely to upskill or accelerate upskilling because of preexisting trends and are disproportionately located in hard-hit MSAs. However, we obtain very similar results. These controls also help adjust for changes in the sample drive, say, by changes in the representativeness of the BG data.

Our data afford the unique opportunity to measure changes in skill requirements within occupations, while the bulk of work on polarization has measured shifts in employment and wages only across occupations. That literature has successfully pinpointed the kinds of occupations that can be replaced by machines or overseas labor using information on the tasks performed by workers at the occupation level.[34] These occupations tend to be in the middle of the skill distribution, since those jobs tend to be the more routine. Autor (2014) and Jaimovich and Siu (2015) point out that employment shifted away from middle-skill occupations in the Great Recession, and, though not shown, we also find this to be true of vacancy postings. We instead focus on our ability to measure task and skill content of jobs in the BG data to ask whether RBTC also occurs *within* middle-skill occupations.

To explore heterogeneity in within-occupation changes in skill requirements induced by the Great Recession, we estimate equation (2) separately regressions for 20 occupation skill groups. As is common in the literature (Autor 2015), we define occupation skill by wage percentiles from a fixed point in time. In our case, we use the 2000 Census to divide

[32] BG generates occupation codes using an imputation that is based on the job title and other characteristics of the ad. One might be worried about the accuracy of this imputation and, though the algorithm was constant over our sample period, the amount of measurement error induced by the fixed algorithm could vary over time. Reassuringly, the baseline results are nearly identical to the MSA-level results presented in figure 1.

[33] We also find no systematic relationship between the change in the share of ads with a missing firm and our key explanatory variables. See column 1, appendix table A2.

[34] The original work by Autor, Levy, and Murnane (2003) and Autor, Katz, and Kearney (2006) used the US Department of Labor's *Dictionary of Occupational Titles* (DOT; US Department of Labor 1977) to categorize occupations into manual, routine, and cognitive. They chose this categorization, arguing that new technologies can successfully replace American workers performing routine, algorithmic tasks, and are complementary to both manual and cognitive and analytical functions. Indeed, though coarse, this grouping successfully predicted employment changes in the 1990s and has been used in a number of subsequent papers, including Autor and Dorn (2013).

Figure 4: Upskilling by 2000 Skill Group

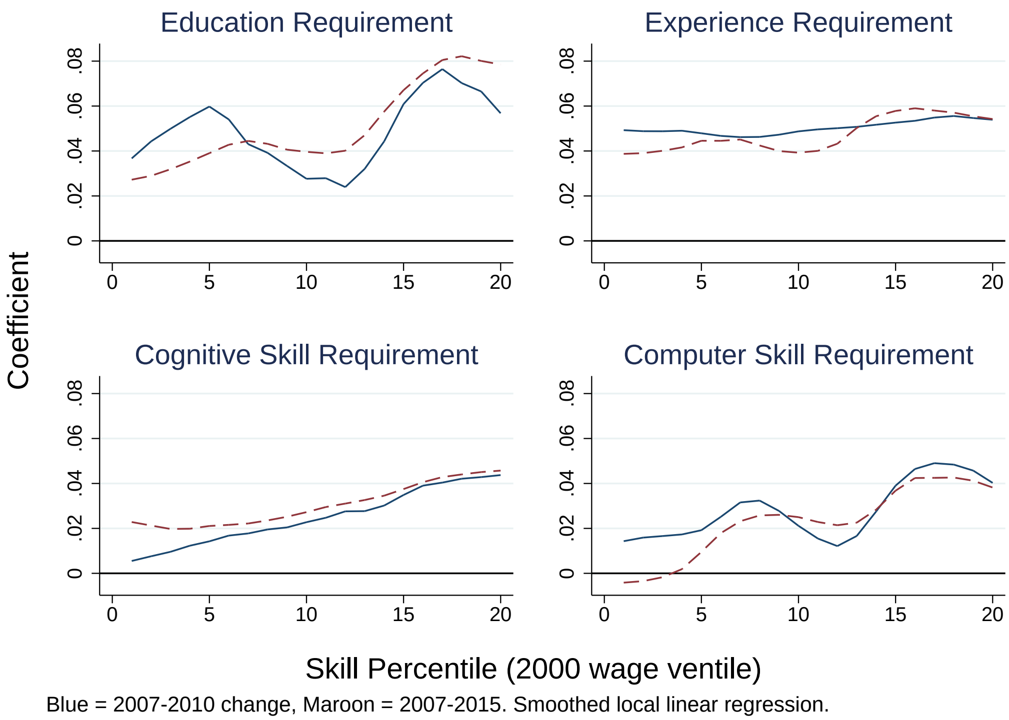

occupations into ventiles, based on their average weekly wage of full-time workers.[35]

In figure 4, we plot the coefficients on $shock_m * I^{2010}$ (solid blue line) and $shock_m * I^{2015}$ (dashed red line), respectively, smoothed with local linear regression. These summary measures thus depict the size of the initial upskilling effect and the size of the remaining effect at the end of the sample period. In most cases the two coefficients are right on top of each other, implying that the majority of the upskilling effects are persistent within occupation categories. Almost all coefficients are statistically different from zero (except for the lowest-wage occupations in the cognitive and computer specifications) and economically meaningful. We also see a fairly similar pattern across dependent variables: upskilling at relatively higher rates at low and high modes of the skill distribution, though in most cases (education, experience, computer), the low mode shows less persistence than the higher one. Together, these results suggest upskilling is broad-based across occupations, although the effects are particularly pronounced and persistent in the upper-middle of the skill distribution.[36]

[35]We obtain weekly wages (annual wage and salary income divided by weeks worked in the last year) of full-time, full-year workers (worked at least 48 hours in the last week and usually work at least 35 hours per week) from the 2000 Census and define ventiles based on Census sample weights. We do not follow the traditional approach of using the 1980 wage distribution both because we are interested in a more modern characterization of skills and because long-term institutional and structural changes in the labor market (e.g., the declining share of unions and the growth of information technology) have made the earnings distribution in 1980 a poorer proxy for skill 30 years later.

[36]We obtain qualitatively similar results if we use the 1980 wage skill distribution instead of the one from 2000, or if we classify occupations based on a measure of critical-thinking and problem-solving intensity derived from O*Net and developed by Joseph Altonji.

Beaudry, Green, and Sand (2014, 2016) document that more-educated workers have increased their employment in lower-skilled jobs since 2000. They term this shift, along with stagnating employment in cognitive occupations, the "great reversal" in the demand for cognitive skill. They hypothesize that lessened demand for cognitive occupations induces college graduates to take jobs lower in the skill distribution, squeezing out less-educated workers who formerly held these jobs. In light of the evidence above, we propose that any declining demand in cognitive occupations was accompanied by an increased demand for cognitive skill *within* routine-task occupations, and this shift accelerated in the Great Recession. Even as employment has shifted from routine to cognitive occupations, the remaining routine occupations themselves are becoming less routine and more cognitive. Our work thus highlights an alternative hypothesis for why high-skilled workers are increasingly found in lower-skilled occupations: these latter occupations are becoming more skilled, and it is possible that less-skilled workers are displaced because they are unable to perform the new duties required.

5 Heterogeneity across Firms

While there are a number of cyclical reasons why firms might upskill their workforce in a slack labor market, discussed in more detail below, the primary implication of the RBTC-driven episodic restructuring hypothesis is that upskilling in job ads will be persistent. Once a firm upgrades its workforce, presumably with a concomitant upgrade in machines and technology or overseas infrastructure, the changes will remain. That is, if firms are changing what they do, not simply whom they hire, those changes will endure. In this section, we argue that the evidence is consistent with upskilling being driven primarily by RBTC. To do so, we investigate firm-level changes in skill requirements and other factors of production, notably use of capital.

We also note that the original Schumpeter (1939) cleansing model suggested that productivity improvements might be concentrated in recessions because low-productivity firms shut down and resources consequently can be reallocated to new, more modern firms. We thus explore this mechanism as well by decomposing upskilling into changes within and across firms. It is illustrative to perform a formal decomposition, which also allows us to address issues arising from our unbalanced panel of firms and possible measurement error in firm identification.[37]

[37] Also, though not shown, we find that our results are robust to include firm fixed effects, as well as firm-by-occupation fixed effects, suggested within-post changes in skill requirements is important.

5.1 Within-Firm Changes

Many of the models that generate episodic restructuring consider firm-level responses to the changing costs and benefits of adopting new technologies. These models generate three key implications that we can test in our data: (1) skill requirements should increase *within* firms; (2) demand for skills should remain elevated within firms; (3) upskilling should be accompanied by increases in capital stock.

We first divide firms into quartiles based on changes in skill requirements between 2007 and 2010; we then plot the average skill requirements for each quartile over time in figure 5.[38] Firms began at fairly similar levels of skill requirements in 2007, when roughly one-third of ads had education (top left panel), experience (top right), and computer (bottom right) requirements, and roughly one-quarter had cognitive skill requirements (bottom left).[39] By construction there is a sharp contrast across firm quartiles in 2010, with the darker shaded lines representing firms with larger skill increases. Interestingly, and not by construction, these quartiles remain spread apart throughout the remainder of the sample period, and by 2015, the higher quartiles still have substantially higher skill requirements in the *new* ads that they post compared with lower quartiles. Thus the persistence in upskilling is largely driven by the same firms that initially upskilled during the Great Recession.

Next, we provide a regression analysis that yields magnitudes and statistical significance for within-firm persistence in upskilling, and allows us to take advantage of our cross-sectional variation in how MSAs bore the Great Recession. In equation (3), we regress the within-firm-MSA change in skill requirements from 2007 to time $t \in (2011, 2015)$ on the firm-MSA change in skill requirements between 2007 and 2010.

$$(3) \qquad skill_{fmt} - skill_{fm07} = \alpha_0 + [(skill_{fm10} - skill_{fm07}) * I^t]\alpha_1 + I^t + X_m\beta + \epsilon_{fmt}$$

The coefficients, α_1, tell us how much of the initial within-firm-MSA upskilling effect persists to subsequent years in the recovery. However, our primary question is whether the recession induced firms to restructure. To get at this question, we can instrument for the 2007-2010 change ($skill_{fm10} - skill_{fm07}$) using the Bartik shock.[40] Results for both specifications are presented in table 2.[41]

In column 1 of table 2, the 0.82 coefficient in the first row implies that 82% of the average

[38] We exploit the subsample of firms in our data that post at least five observations in each of 2007 and 2010, comprising 66% of weighted observations—we explore upskilling within the rest of the sample in the next subsection.

[39] Note: This similarity across firm quartiles is not imposed by our exercise.

[40] Specifically, we instrument for the vector $(skill_{fm10} - skill_{fm07}) * I^t$ with the vector $(shock_m * I^t)$.

[41] These regressions necessarily restrict to the 51% of weighted firm-MSA cells in 2011–2015 that had matches in 2007 and 2010. For most years, the probability of making it into this sample does not vary with the Bartik shock, however, by 2015 harder-hit MSAs are slightly less likely to be missing from this sample, and this effect is significant at the 10% level (see appendix table A2, column 3). However, reassuringly, the baseline effects reported above are very similar within this sample. We explore upskilling for firms and occupations that did not post in 2007 in the next subsection.

Figure 5: Skill Requirements by Firm 2007-2010 Change

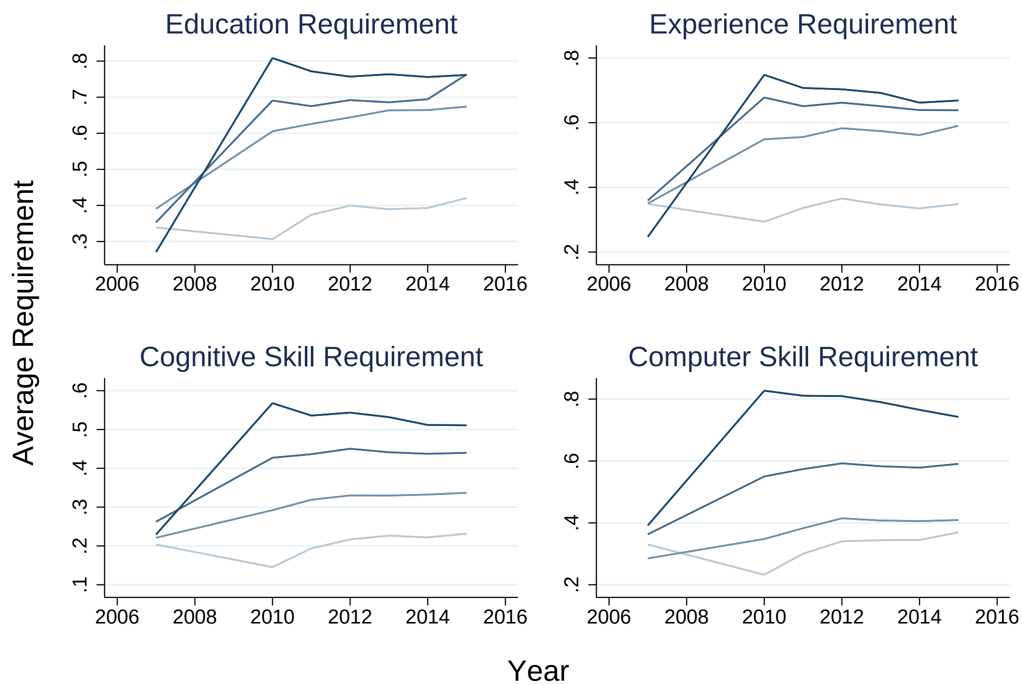

Graph plots average skill requirement by year and 2007-10 change quartile.

Table 2: Firm-Level Persistence in Early Upskilling

	Education		Experience		Cognitive		Computer	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
change*2011	0.820***	1.293***	0.836***	1.011***	0.786***	1.737	0.796***	1.130***
	(0.00375)	(0.393)	(0.00417)	(0.250)	(0.00524)	(1.589)	(0.00561)	(0.238)
change*2012	0.779***	0.482	0.779***	0.545*	0.736***	0.472	0.740***	0.943***
	(0.00491)	(0.397)	(0.00593)	(0.327)	(0.00626)	(0.464)	(0.00539)	(0.269)
change*2013	0.772***	0.948**	0.759***	0.586*	0.706***	0.210	0.711***	2.411
	(0.00459)	(0.384)	(0.00658)	(0.327)	(0.00700)	(1.005)	(0.00596)	(3.732)
change*2014	0.745***	0.631	0.734***	0.730**	0.650***	-0.923	0.658***	0.714
	(0.00586)	(0.453)	(0.00631)	(0.338)	(0.00783)	(2.037)	(0.00691)	(1.017)
change*2015	0.671***	0.936	0.671***	1.007	0.603***	0.290	0.627***	-0.721
	(0.00674)	(0.597)	(0.00826)	(0.750)	(0.00691)	(0.621)	(0.00863)	(4.397)
Observations	419,749	419,749	419,749	419,749	382,328	382,328	382,328	382,328
R-squared	0.566	0.481	0.558	0.516	0.495		0.514	

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

Notes: These regressions estimate firm-MSA changes in skill requirements from 2007 to years 2011-2015. "change" is the within firm-MSA change in skill requirement from 2007-2010. Columns labeled "IV" instrument for change*year with the Bartik shock*year. Regressions also include year fixed effects and MSA-level characteristics from the ACS. Regressions are weighted by the firm ad share within the MSA-year times the MSA labor force size in 2006. Standard errors are clustered by MSA.

initial firm-MSA change in skill requirements persists to new ads posted in 2011. This effect is significant at the 1% level and fades only slightly over the next 4 years, so that by 2015, 67% of the initial upskilling effect remains.

In column 2, labeled "IV", we estimate the relationship between early and later upskilling, using only variation in early upskilling driven by local labor market conditions. Using this approach, the magnitude of the coefficient increases substantially and the effects are even more persistent. Our estimates imply that the change from 2007 to 2010 in demand for education induced by the Bartik shock is fully kept for new ads in 2015, though the estimates are noisier.

OLS results for the other dependent variables are extremely similar – roughly two-thirds of the firms that upskilled between 2007 and 2010 maintained elevated skill requirements for *new* ads that they posted in subsequent years. IV results generally confirm these effects, meaning that the increase in skill requirements driven by firms located in worse economies is also persistent, though these effects are quite a bit noisier, especially in later years.

Thus, there is considerable evidence for increases in skill within firms and occupations, an important component of RBTC-driven episodic polarization. However, an RBTC explanation would also suggest that firms automate routine tasks with machines (or cheaper, outsourced labor), which complement skilled labor. If mechanization is occurring, then we should observe firms investing in physical capital around the time that they upskill. This, however, would be at odds with the stylized facts that (a) investment is in general low in a recession and (b) investment was particularly slow to recover following the Great Recession. Jaimovich and Siu (2012), however, show that investment in information processing equipment and software bounced back immediately following the end of the NBER-dated recession (December 2007 through June 2009). This type of investment actually surpassed its pre-recession level before June 2010, less than four quarters after the recession formally ended.[42] Interestingly, IT and computers are exactly the kinds of technology that drive RBTC (Michaels, Natraj, and Van Reenan 2014).

While this evidence on general IT investment is merely suggestive, we investigate whether the specific firms that upskilled the most also increased investment at the same time. We do so by linking the publicly-traded firms in our data to Compustat North America by Standard & Poors (hereafter Compustat), the most complete database of U.S. firm accounting and balance sheet data.[43] While we cannot identify IT investment specifically in the Compustat data, we can measure a firm's overall holdings of property, plant, and equipment (PPENT). This measure of capital stock includes IT investments and, under the RBTC hypothesis, should increase faster for firms mechanizing or using more sophisticated production technology.

We match firms to Compustat using firm name and provide details of the matching

[42] In contrast, domestic private investment as a whole took about four *years* to fully recover.

[43] We obtain these data via Wharton Research Data Services.

Figure 6: Firm-Level Changes in Skill Requirements and Capital Stock 2007–2010

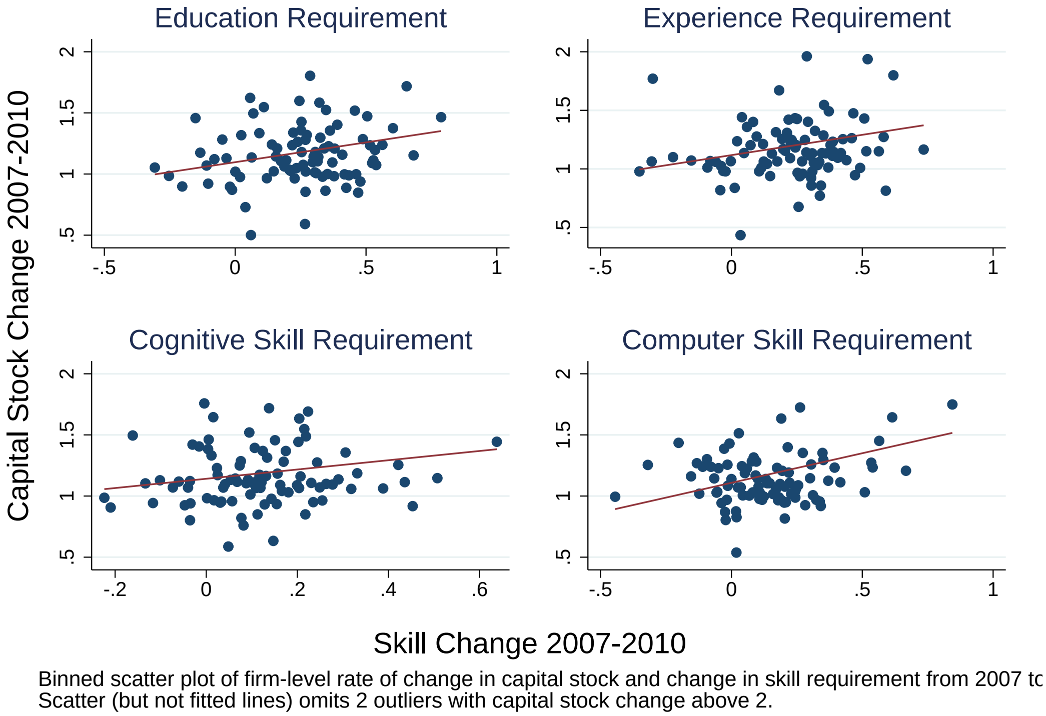

procedure in Appendix A.3.[44] To provide a general sense of the data, figure 6 presents binned scatter plots of the change in skill requirements between 2007 and 2010 on the rate of change in capital stock (PPENT) over the same time period.[45] As can be seen, there is a strong positive relationship for both the share of ads with an education requirement (left panel) and the share with an experience requirement (right panel). Firms that had larger increases in skill requirements between 2007 and 2010 also had larger increases in capital stock over the same time period.[46]

We also provide results from regression analysis, analogous to those presented in figure 1, comparing the impact of the Bartik shock across firms with different capital investment behavior. We estimate regressions similar to that in equation 1 but control for the firm's change in capital stock between 2007 and 2010, and allow this variable to interact with the *shock*-year effects.[47] We estimate within-firm-MSA skill changes in the Compustat sample

[44] We can link 46% of weighted observations (18% of firms) from the 2007–2010 change sample used in figure 5 to a business in Compustat. In column 2 of appendix table A2, we show that the change in the share of firms matching to Compustat does not vary with the size of the employment shock to the local labor market.

[45] To make the data easier to see, we aggregate the firms into percentiles based on the magnitude of changes in skill requirements (weighted by the average number of ads posted in 2007 and 2010); we then plot for each percentile bin the ad-weighted average of this variable against the ad-weighted average of the change in capital stock.

[46] The figure would look similar if instead of actual firm upskilling we plotted the predicted firm-specific upskilling measures defined above, which rely more on the impact of local labor market conditions on skill requirements.

[47] Regressions are estimated at the firm-MSA-year level for the Compustat matched sample.

Figure 7: Effect of MSA-Specific Employment Shock on Skill Requirements, Between Interdecile (90–10) Firm Change in Capital Stock

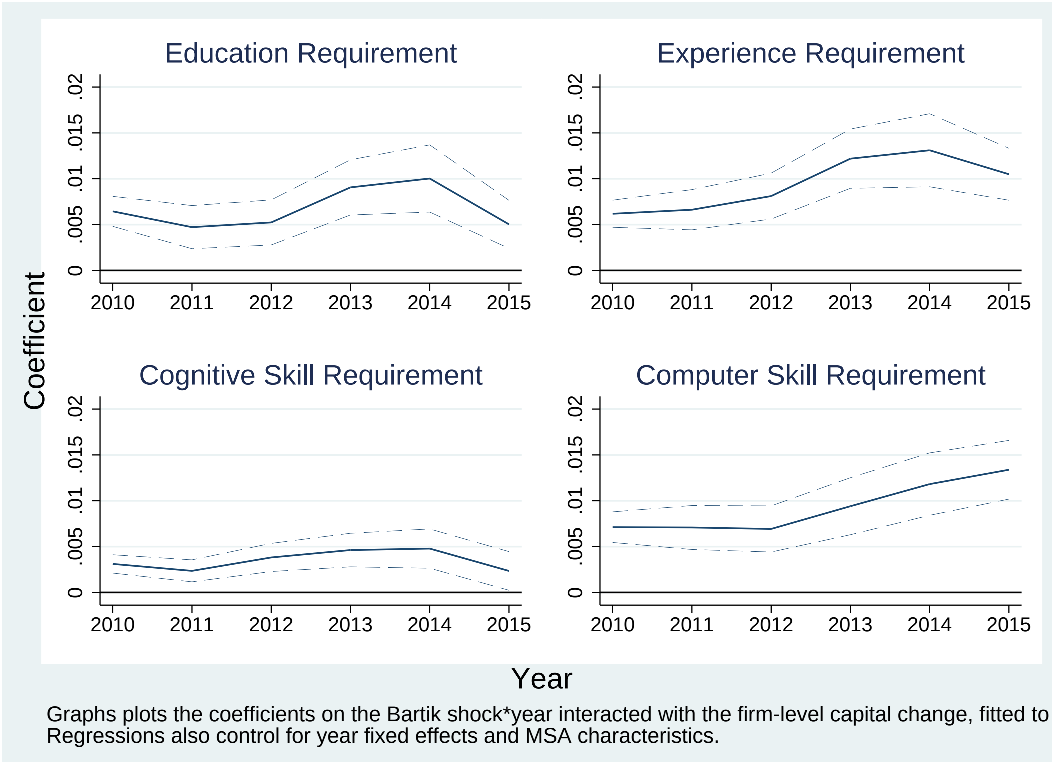

and thus drop firm-MSAs that cannot be matched back to 2007 (10% of weighted observations). Figure 7 fits the differential impact (and 90% confidence interval) of the Bartik shock for firms at the 90th and 10th percentiles of capital change. The former had a 60% increase in their capital stock while the latter experienced a 20% drop.

We find that firms with larger increases in capital stock had greater upskilling, and these differences persist. The increase in skill requirements is on average around 30% (education and experience requirements) to about 50% (cognitive and computer requirements) higher in high-investment firms than in low-investment firms. These differences are all significant, typically at the 1% level. Thus, throughout our sample period, firms with larger increases in capital stock around the time of the Great Recession also had larger increases in their posted skill requirements.

5.2 Within and Across Firm Changes

In principle, upskilling could occur within firm as adjustment costs decline and make it optimal for firms to adopt new technologies, and in the previous subsection we showed that this component is present. Another classic component of Schumpeterian cleansing is the idea that old, unproductive firms fail in a recession and resources are reallocated to new, more productive firms. Additionally, the distribution of jobs across firms may vary over the business cycle.[48]

To place relative importance on these explanations, we decompose skill changes into several components. Define C_t as the set of firm-MSAs that post ads in both year t and in 2007. We hereafter refer to this set as "continuing firms", and the set of firm-MSAs that have posts only in 2007 or only in t as non-continuing firms. With these categories, we can separate skill changes due to changes within existing firm from those due to changes in the mix of firm that post vacancies.[49]

In equation (4), we express the average skill requirement in MSA m, and year t as a function of four components: (1)p^C_{mt}, the share of ads in an MSA-year posted in continuing firms; (2) $skill^C_{fmt}$, the average skill requirement for firm f, posting in MSA m, and year t, where fm is a continuing firm; (3) $\frac{N^C_{fmt}}{N^C_{mt}}$, the share of ads posted by firm f, among all ads to continuing firms in mt; and (4) $s\bar{kill}^{NC}_{mt}$ is the average skill requirement among non-continuing firms in mt (that is, the average skill requirement in the MSA-year among firms that either did not post in the previous period, or posted only in 2007. The last component provides an impression of the substitution across firm death and birth margins in affecting skill requirements, though it is important to note that we cannot distinguish whether a firm does not post in a given period because it has no vacancies (e.g., a hiring freeze) or because it does not exist.

$$skill_{mt} = p^C_{mt} \sum_{fm \in C_t} skill^C_{fmt} * \frac{N^C_{fmt}}{N^C_{mt}} + (1 - p^C_{mt}) s\bar{kill}^{NC}_{mt} \tag{4}$$

Our regression analysis above uses $skill_{mt} - skill_{m07}$ as a dependent variable. But we can also consider counterfactual changes in skill requirements from 2007 to t by differencing any given component from its 2007 level, holding constant the other components. For example, $p^C_{mt} \sum_{fm \in C_t} (skill^C_{fmt} - skill^C_{fm07}) * \frac{N^C_{fmt}}{N^C_{mt}}$ is the change in skill requirements between 2007 and t, attributed just to changes in the within firm-MSA skill requirement among continuingfirms, holding constant all other components at their levels in t. We can regress this counterfactual skill change on the Bartik employment shock and the other controls in equation (1) to understand how much of the total cyclical responsiveness is attributed to a cyclical response in the within firm-MSA skill requirement.

A decomposition begins with $skill_{m07}$ and differences each of the four components, one at a time, between the two periods. After differencing one component, that component is

[48] Kahn and McEntarfer (2015) show that workers matching to jobs in downturns are more likely to match to low-paying firms than high-paying firms.

[49] Currently, we examine just the distribution of firms, but it is also possible to decompose this component into changes due to firms posting in only one period and changes due to firms that post in both periods but change the mix of occupations in which they post. We will examine these breakdowns in a later draft.

Figure 8: Decomposing Upskilling Into Within and Across Firm Components

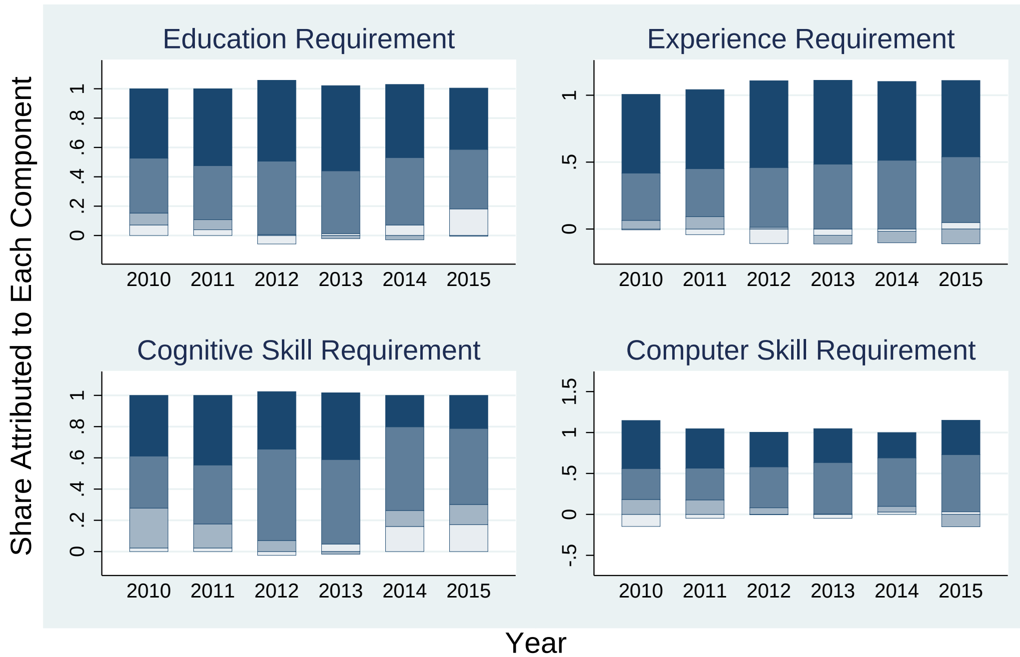

Darkest = within continuing firm, dark = across old-new firms, light = distribution across continuing firms, lightest = share continuing firms

We decompose the impact of the MSA-specific Bartik employment shock on the change in skill requirements from 2007 in each year. We then plot the share attributed to each of four components, averaged across 24 possible decomposition orders.

fixed at its time-t value for the subsequent differences.[50] We can regress each counterfactual change on the same variables in equation (1), and the coefficients on $shock * I^t$ will sum to the coefficient earlier reported. Since the order of the decomposition affects the relative importance of each component, there are 24 possible combinations for decomposing the full effect into its four components.

We summarize the results for our four primary dependent variables in figure 8. The bars report the mean fraction of the overall impact of *shock* in the given year attributed to each component. These fractions are averaged across the 24 decomposition orders. The darkest bar shows the fraction attributable to the within firm-MSA change in skill requirements from 2007 to t. The adjacent, slightly lighter bars show the fraction attributable to changes in the skill requirements of non-continuing firms. The two, more-lightly-colored bars show changes attributed to the within-MSA-year distribution of firms posting and the share of ads to continuing firms, respectively.

[50] For example: $skill_{omt} - skill_{om07} =$

$$\begin{gathered}
p^C_{omt} \sum_{fom \in C_t} skill^C_{fom07} * \frac{N^C_{fom07}}{N^C_{om07}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07} - p^C_{om07} \sum_{fom \in C_t} skill^C_{fom07} * \frac{N^C_{fom07}}{N^C_{om07}} + (1 - p^C_{om07}) s\bar{kill}^{NC}_{om07} \\
+ p^C_{omt} \sum_{fom \in C_t} skill^C_{fom07} * \frac{N^C_{fomt}}{N^C_{omt}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07} - p^C_{omt} \sum_{fom \in C_t} skill^C_{fom07} * \frac{N^C_{fom07}}{N^C_{om07}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07} \\
+ p^C_{omt} \sum_{fom \in C_t} skill^C_{fomt} * \frac{N^C_{fomt}}{N^C_{omt}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07} - p^C_{omt} \sum_{fom \in C_t} skill^C_{fom07} * \frac{N^C_{fomt}}{N^C_{omt}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07} \\
+ p^C_{omt} \sum_{fom \in C_t} skill^C_{fomt} * \frac{N^C_{fomt}}{N^C_{omt}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{omt} - p^C_{omt} \sum_{fom \in C_t} skill^C_{fomt} * \frac{N^C_{fomt}}{N^C_{omt}} + (1 - p^C_{omt}) s\bar{kill}^{NC}_{om07}
\end{gathered}$$

By far, the most important components are the within-continuing-firm changes and the changes driven by non-continuing firms. For education and experience, within-continuing-firmt changes contribute roughly half of the upskilling effect, with about 40% from changes in non-continuing firms, and little role for the other two components. For cognitive and computer skills, changes in the skill requirements of non-continuing firms tend to be the largest component contributing to upskilling, accounting for roughly 60% of the total effect, although there is some variation across years. The within-continuing-firm changes are the next largest component, at roughly one-third of the total effect, and the contribution due to the mix of firms and the share of continuing firms still playing small roles, although slightly larger than in the case of education and experience.

We therefore find that the Great Recession induced increased skill requirements within firm-MSA combinations that posted vacancies both before and after the recession, and that this effect accounts for roughly half of the overall increase in upskilling across MSAs over this time period. Given the persistence and breadth of this effect across skills, the results are quite suggestive that firms have changed the labor component of how they produce. However, we also find a sizable role for differences in skill requirements across non-continuing firms: those posting ads online only in the 2010–2015 period require more-skilled workers than those posting only in 2007. To the extent that human capital is associated with productivity, this is consistent with the notion that the Great Recession drove low-productive firms out of business, and generated replacement firms that produced with better inputs. Of course, we cannot pinpoint new and old firms or production technology directly.

6 Alternative Explanations for Upskilling

In this section, we discuss a number of alternative hypotheses that might explain upskilling. The primary goal is to understand whether upskilling reflects a larger change in what firms do and how they produce, or whether it reflects a more minor change in their mix of labor inputs (whom they hire) or their labor search strategy (whom they ask for). We thus discuss several possibilities for why firms might change their advertised demand for labor that are not driven by a restructuring during the Great Recession. We do not wish to suggest that these alternatives are completely unimportant. However, we conclude that they cannot explain the persistent upskilling effect we find.

If firms *intend* to hire more-skilled workers without otherwise adjusting their production technology, it presents something of a puzzle: Why would they choose to do so despite the fact that higher-skilled workers are, in general, more expensive, and become relatively even more costly in a downturn?[51] We explore three possibilities: (1) firms temporarily and

[51] That lower-skilled workers are more detrimentally affected by recessions is well known (Hoynes, Miller, and Schaller 2012). This implies that high-skilled workers become relatively more costly and relatively less available in downturns. Indeed, using American Community Survey data, we estimate that earnings and

opportunistically seek higher-skilled workers in a slack labor market, (2) replacement hiring differentially requires more-skilled workers due to differential recoveries in quit rates (i.e., the job ladder needs to clear), and (3) population changes in educational attainment alter the availability of skill.

It could also be the case that firms do not intend to hire the more-skilled workers they specify in their ads. Early in the recovery, firms may worry about a "bottleneck" effect, that the slack labor market will produce an unusual number of job applicants that can be costly to screen. They may try to signal to certain (unwanted) applicants that a job is not for them by posting specific skill requirements, but this would imply that firms do not change how they produce but rather that they make explicit requirements that had been only implicit previously. We do not believe this screening hypothesis is a primary driver of our results because the upskilling effect persists well into the recovery, even after labor markets have converged back to their pre-recession levels of labor market tightness. However, we also explore whether firms successfully hire the more-skilled workers they seek using ACS data on employment.

6.1 Opportunistic Upskilling

It could be that some firms always want to hire more-skilled workers, but in a tight labor market, they cannot attract (or afford) them. In slack economies, firms may opportunistically seek out higher-skilled workers. Again, this effect should be temporary, while we find that upskilling is persistent. However, as we have noted, labor markets have been slow to make a full recovery. To investigate this issue, we compare the upskilling effect across labor markets where employment shocks more or less severely affected high-skilled workers. If opportunistic upskilling take place, it should be more likely to occur in MSAs where high-skilled workers were more severely affected, and thus more readily available to take lower-skilled positions.

We augment equation (1) with controls for the change in the unemployment rate during the Great Recession among college graduates age 25 or older, obtained from the ACS. As with the Bartik employment shock, we interact the college unemployment rate shock with year dummies to allow it to affectchanging skill requirements each year. We also norm this variable so that a one unit change equals the 90–10 percentile differential across MSAs . We plot both estimates in figure 9.

First, we note that the overall employment shock and the college unemployment rate shock are only weakly correlated ($r = 0.14$). Highly educated MSAs such as New York City and San Jose/Santa Clara had larger increases in their college unemployment rates, while MSAs particularly affected by the housing bubble, such as Phoenix, did not.

Holding constant the overall size of the employment shock on an MSA, we find that

unemployment rate gaps widen across education and experience groups when local labor market conditions are worse.

Figure 9: Upskilling and the College Unemployment Rate

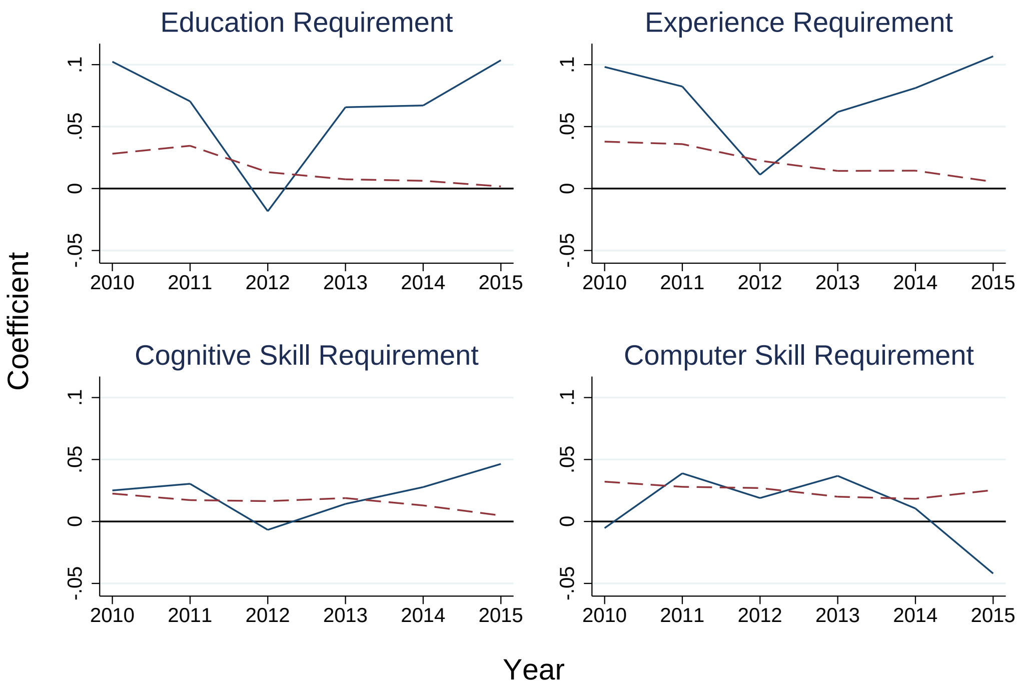

Blue solid line = bartik impact, Maroon dashed line = college unemployment rate impact

a larger increase in the college unemployment rate also generates upskilling. In the early years of the recovery, skill requirements see an additional modest increase that is statistically significant for all skill variables in 2010 and 2011. However, the impact of the college unemployment rate shock dissipates quickly for education, experience and cognitive requirements, and has converged back to pre-recession levels by 2015. The effect remains significant in 2015 only for computer skill requirements. Yet, the effect of the overall employment shock, holding constant the size of the college graduate unemployment rate shock, remains quantitatively large (and often larger than before), statistically significant, and persistent for each skill type except computer skills.

This set of results has two important implications. First, the impact of the shock to college graduates does not persist but fades to pre-recession levels, suggesting that some opportunistic upskilling may have occurred early in the recovery (Sasser Modestino, Shoag, and Ballance 2015). Second, and more important, the overall employment shock induces persistent upskilling, even holding constant the size of the shock to college graduates, implying that opportunistic upskilling does not drive our main results. The exception is for computer skills, where it may be difficult to separately identify the overall employment shock from the college graduate shock, since computer skill increases are likely heavily concentrated in the most-educated MSAs.[52]

[52]We have also examined firm-level upskilling by the initial skill requirements of the firm (instead of the initial *change* in skill requirements, as above. If opportunistic upskilling is an important part of the story, we would expect to see increases in skill requirements concentrated among firms with low or middling initial

6.2 Quits and Replacement Hiring

We have argued that the persistence in (firm-specific) upskilling is key to disentangling cyclical explanations from structural ones. However, it could be that cyclical explanations (such as the bottleneck hypothesis) account for early upskilling and that labor supply factors drive upskilling later in the period. For example, quits were slow to recover following the Great Recession. If quit rates among higher-skilled workers recovered more quickly, then some of the apparently persistent upskilling could be driven by replacement hiring higher up on the job ladder. Differences would eventually even out across groups as quit rates among lower-skilled workers recover.

However, we see no evidence of differential recovery in quit rates across skill groups. For instance, in figure 10 we plot smoothed time series of quit rates by education group from 1998–2015 using longitudinally-linked CPS data.[53] Besides the large and well-known secular decline in quits (Molloy, Smith, and Wozniak 2013), the figure shows the quit rate has a clear cyclical component, declining during NBER-dated recessions (indicated with vertical dashed lines) and recovering somewhat shortly thereafter. There is certainly no evidence in this figure that quit rates for higher-skilled workers (lighter lines) recover more quickly than those for lower-skilled workers (darker lines); in fact, the opposite appears more likely, with less-educated workers showing faster recoveries. It thus seems unlikely that replacement hiring is driving the persistence in upskilling.

6.3 Educational Attainment

Firms may decide to upskill if they observe that skilled workers have become more plentiful in their local economy. Indeed, a long line of research has explored whether educational attainment responds to local labor market conditions. Past evidence generally finds moderate effects on enrollment, but limited effects on attainment.[54] One exception is Charles, Hurst, and Notowidigdo (2015), who find that during the 2000s, the housing boom reduced educational attainment, primarily on the two-year college margin. This effect (and any symmetric rise in attainment when the housing bubble crashed) is less relevant for our sample of primarily higher-skilled jobs. In particular, the housing boom was a shock principally

levels; however, we find no evidence in support of this hypothesis.

[53] We define a quit as a worker who reported switching employers between month t and month$t+1$ (but was employed in both months), or one who reported being employed in month t and unemployed in month $t+1$, and gave voluntary job leaving as a reason for the unemployment. To obtain quit rates we divide the count of quits by the total number of employed workers in month$t$.

[54] Card and Lemieux (2001) find that local unemployment rates have small, positive impacts on high school attendance and completion, marginal impacts on college attendance, and no impact on college completion. Barr and Turner (2015), using more recent data, find that college enrollment has grown more responsive to the business cycle over time. Kahn (2010) finds that among cohorts graduating from college between 1979 and 1989, those graduating in a worse economy obtained an additional year of graduate school, on average. She also finds that economic conditions at time of high school graduation did not affect college completion. Altonji, Kahn, and Speer (2016) find similar modest effects over a broader time period.

Figure 10: Quit Rates by Date and Education Group

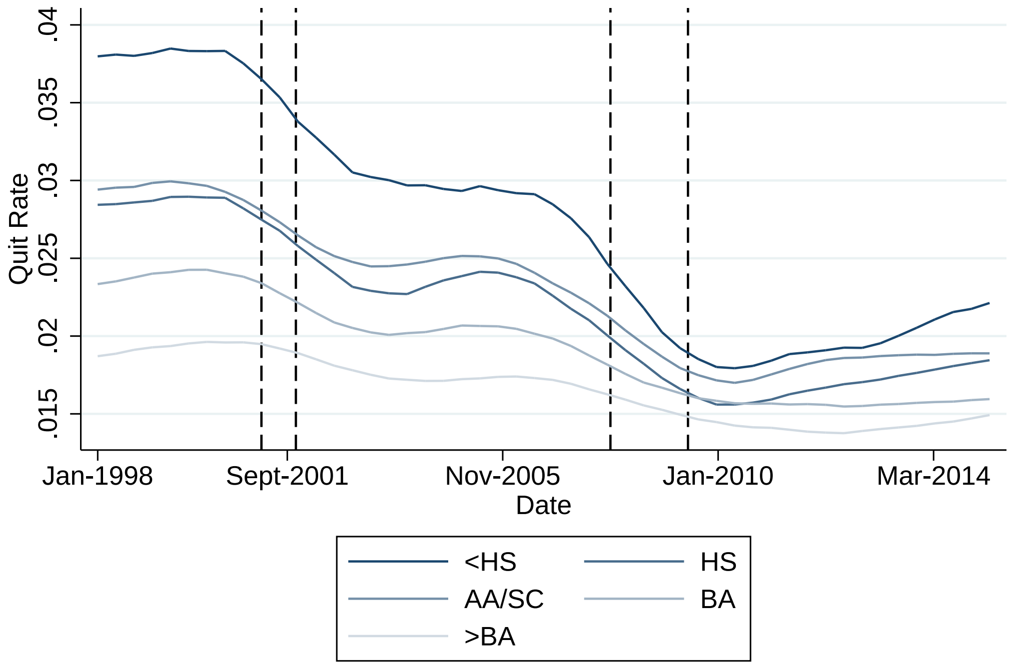

Graph plots local polynomial smoothed quit rates. Vertical dashed lines indicate NBER dated recessions.

affecting low-skilled workers, while the Great Recession was broad-based.[55]

Using American Community Survey data, we examine the relationship between labor market conditions and changes in educational attainment across U.S. states, and find only modest effects. We regress state-year educational attainment on state-level Bartik employment shocks, allowing this shock to interact with year, in a manner analogous to the MSA-level shocks used in equation (1). We focus on a young population (aged 18–32) whose educational attainment decisions should be most malleable. Our estimates (available upon request) imply that a state experiencing a Great-Recession-sized shock has no change in the probability that its young population attains at least some college and a (weakly) lower probability that it attains at least a bachelor's degree. Thus we conclude that changes in educational attainment induced by the Great Recession are unlikely to be driving our results.

6.4 Does Upskilling in Job Ads Extend to Employment?

We have provided robust evidence of upskilling: firms facing worse local labor markets increase their stated skill requirements in job postings. However, the stated preferences of firms may be of less interest if they are not related to realized employment. For example, given the relatively low cost of advertising electronically and the weak labor markets in 2010

[55] For example, Altonji, Kahn and Speer (2016) show that the Great Recession affected recent college graduates more severely than past recessions had, and that higher-earning college majors lost much of their previous advantage in weathering worse labor market entry conditions.

and 2011, firms might post high skill requirements to gauge the available pool of labor, without necessarily expecting to immediately fill a position.

To address this possibility, we perform a companion analysis using employment data from the American Community Survey (ACS). We ask whether the employment stock has become relatively more skilled since 2007 in locations that were harder hit by the Great Recession. We examine similar specifications to equation 1 except, given the smaller sample sizes in the ACS, we consider economic conditions across states rather than MSAs.[56]

Figure 11 reports results for four dependent variables. The first three panels measure changes in relative employment rates across education groups. For example, the top left panel measures the change in the ratio of high school graduate to high school dropout employment rates. This relative employment rate averages 1.6 across our whole sample period. We find that in a hard-hit state (90th percentile shock), the high school graduate-dropout relative employment rate increases by 10 percentage points after 2007, relative to a less hard-hit state (10th percentile shock). We find similar effects for the change in relative employment of workers with some college compared to dropouts (top right), and the change for workers with at least a BA relative to dropouts (bottom left).

Finally, in the bottom right panel, we examine the change in the average experience level of employed workers. Experience is simply years of school minus age minus six, which is different in spirit from the stated requirement in the BG data that an applicant needs experience in the field. However, more-experienced workers are still more skilled than less-experienced workers, and the measure is as close as the data allow. We find that in hard-hit states, average experience is about a fifth of a year higher than it was in 2007, relative to less hard-hit states.[57]

We therefore find that employed workers become increasingly more skilled in states that were harder hit by the Great Recession. Since these measures represent employment stocks, rather than flows, we would expect any impacts on changes in hiring behavior to accumulate gradually. Indeed our estimates continue to diverge as more time passes since 2007.[58]

[56] We use a state-level Bartik shock, defined in the same manner, and normalize to the 90–10 gap across states (0.014 log points). We control for a linear time trend, rather than year fixed effects because, with 50 states and the District of Columbia, we have fewer degrees of freedom than with 381 MSAs. For the same reason, we exclude the ACS characteristics. Standard errors increase when we include these controls, though point estimates are similar and impacts for most years are significant at at least the 10%-level. Standard errors are clustered by state and observations are weighted by the size of the state labor force in 2006. Estimation results using MSAs is qualitatively and quantitatively similar, albeit slightly less precise.

[57] This result is consistent with Forsythe (2016) who finds that the hire rate of young workers falls, relative to that of older workers, in recessions. She provides evidence based on wage data that these changes in hiring are labor demand, rather than labor supply, driven.

[58] We have also examined changes in employment flows using CPS data on new job starts. We find similar, though noisier, results, and prefer employment stocks data for two reasons. First, flows data have much smaller sample sizes, both because fewer workers change jobs than are employed and because the CPS is a smaller survey than the ACS. Second, the vast majority of hiring is replacement hiring and we might therefore expect changes in skill requirements to be masked, relative to aggregate changes in employment.

Figure 11: Change in Skills of Employed Workers, ACS 2005-2014

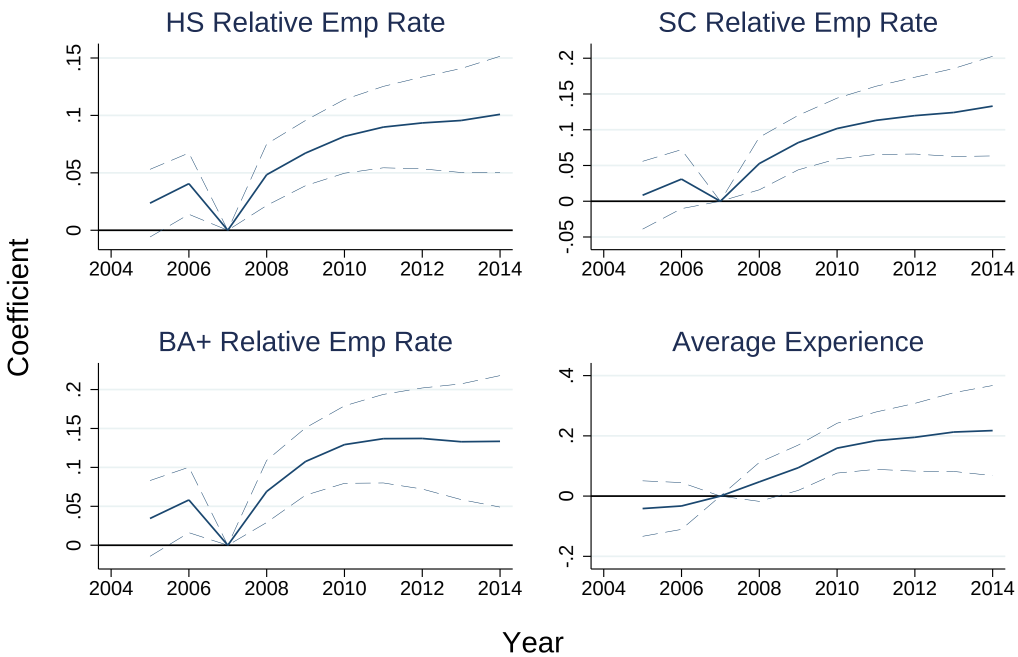

Graphs plot coefficients on employment shock interacted with year, and 90% CI's (dashed lines). The top three panels show the change in relative employment rates (compared to dropouts) for high school graduates, some college, and at least a BA, respectively.
The bottom right panel shows the change in the average experience level of employed workers. Regressions also control for a time trend.

7 Conclusion

During the recovery following the Great Recession, anecdotal evidence suggested that the composition of new hires shifted towards higher-skilled workers, resulting in many workers being "overeducated" for their jobs. However, it was 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, firms may have taken the opportunity of the recession as a time of "cleansing" to restructure their production in a manner consistent with routine-biased technological change. Additionally, firms may have tried to exploit the limited job opportunities of more-skilled workers to recruit from higher in the skill distribution than they would have before the recessionary period. In this paper we provide comprehensive, broad-based evidence of upskilling—firms demanding higher-skilled workers—when the local economy suffers a recession. We argue that this effect is primarily driven by an episodic restructuring on the part of firms and thus is likely to be long-lasting. Our findings that have important implications for our understanding of technological change and business cycles, as well as labor market and social insurance policies.

Using a novel data set that captures detailed information on the near-universe of electronic job postings and that spans the Great Recession, we show that in metropolitan areas with larger employment shocks, job ads were 5 percentage points (16%) more likely to list education and experience requirements and 3 percentage points (8–12%) more likely to list cognitive and computer skill requirements. These patterns are robust to a variety of alternative specifications, and even hold within firm-MSA-occupation cells. Moreover, the richness of our data allow us to examine how much of this demand-side upskilling is driven by changing composition of the firms advertising and the occupations they advertise for. We find that overall upskilling is driven by both changes in skill demand within firms, as well as by substitution across new firms posting ads in the recovery, relative to firms that stopped posting (and therefore likely closed) in the recession. Furthermore, we find a similar upskilling effect in American Community Survey data on employment, suggesting that firms successfully hire the more-skilled workers that they seek.

The upskilling effects do not fade as the labor market recovers. What's more, this persistence occurs among the same firms that we observe upskilling in the immediate aftermath of the recession. We merge our job postings data with firm-level investment data from Compustat and, in support of a retooling of production story, find that the firms exhibiting greater amounts of upskilling also had larger increases in their capital stock over the same time period. Finally, we show that upskilling occurred to a larger extent *within* more-routine occupations – those most susceptible to technological change.

This paper thus provides the first direct evidence that the Great Recession accelerated routine-biased technological change, and in so doing touches on a number of literatures and policy questions. It is consistent with the important, but suggestive, evidence provided by Jaimovich and Siu (2015) that the vast majority of employment lost in routine occupations was lost during recessions and never recovered. It also contributes to the many models in macroeconomics that assume adjustment costs and imply that recessions will be times of "cleansing" in terms of production (Schumpeter 1939, Koenders and Rogerson 2005, Berger 2012). As hypothesized by many, these kinds of episodic, productivity-enhancing changes can result in jobless recovery. Our findings are thus extremely relevant for policy makers, who allocate billions of taxpayer dollars to subsidize workers' job searches in a downturn.

We also demonstrate how electronic job postings data can provide a unique opportunity to understand real-time changes in skill demand both across and within occupations. This level of detail can provide new insight relative to earlier literature. We are able to show that the the middle-skill occupations *themselves* are changing. These occupations were likely already drawing in higher-skilled workers before the Great Recession, and as we show, this effect accelerated during the recession and subsequent recovery. This can help to clarify studies by Beaudry, Green, and Sand (2014 and 2016) and others documenting the "great reversal" in demand for cognitive skill. While it is certainly the case that employment in high-skill occupations has not grown, on average, over the last decade, our results show that cognitive workers still retain a substantial advantage over the low-skilled. They are drawn into formerly middle-skill jobs, which are becoming higher-skilled. We can thus explain why skilled workers still earn a premium in the labor market even though the returns to cognitive occupations appear to have diminished.

The U.S. economy has seen remarkable changes over the last 30 years, brought on by the computer revolution and globalization. These changes have led to great increases in productivity and wealth, but the benefits have not been shared across all workers. Indeed, mounting evidence suggests that a large population of workers, formerly employed in routine-task jobs, have suffered permanent labor market, health, and social consequences from the structural changes in the economy (Autor, Dorn, Hanson and Song 2014, Autor, Dorn, and Hanson 2015, Pierce and Schott 2015, Foote and Ryan 2015). Our results highlight that workers' ability to adjust to these changes may be especially difficult because the changes are episodic, concentrated in recessions. Thus large numbers of workers can find their skills depreciated at the same time. If the changes to production instead occurred more gradually, workers would still need to be retrained, but more gradually, and on a much smaller scale at any given time. Future policy work should be directed at understanding how to reallocate workers on a large scale following a recession.

References

[1] Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer (2016), "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success," *Journal of Labor Economics*, 34(S1 Part 2): S361–S401.

[2] Autor, David (2014), "Polanyi's Paradox and the Shape of Employment Growth," NBER working paper #20485.

[3] Autor, David (2015), "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29(3):3–30.

[4] 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.

[5] Autor, David, David Dorn, and Gordon Hanson (2015), "The Labor Market and the Marriage Market: How Adverse Employment Shocks Affect Marriage, Fertility, and Children's Living Circumstances," MIT mimeo.

[6] Autor, David, David Dorn, Gordon Hanson, and Jae Song (2014), "Trade Adjustment: Worker-Level Evidence," *Quarterly Journal of Economics,* 129(4): 1799–1860.

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

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

[9] Autor, David H., Frank Levy, and Richard J. Murnane (2003), "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics,* 118(4): 1279–1333.

[10] Barr, Andrew, and Sarah Turner (2015) "Out of Work and into School: Labor Market Policies and College Enrollment During the Great Recession," *Journal of Public Economics*, 124: 63–73.

[11] Bartik, Timothy J. (1991), *Who Benefits from State and Local Economic Development Policies*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

[12] Beaudry, Paul, David A. Green, and Benjamin M. Sand (2014), "The Declining Fortunes of the Young since 2000," *American Economic Review*, 104(5): 381–386.

[13] Beaudry, Paul, David A. Green, and Benjamin M. Sand (2016), "The Great Reversal in the Demand for Skill and Cognitive Tasks," *Journal of Labor Economics*, 34(S1 Part 2): S199–S247.

[14] Blanchard, Olivier Jean and Larry Katz (1992), "Regional Evolutions," *Brooking Papers on Economic Activity*, 1–75.

[15] Berger, David (2012), "Countercyclical Restructuring and Jobless Recoveries," mimeo.

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

[17] Caballero, Ricardo, and Mohamad. Hammour (1994), "The Cleansing Effect of Recessions," *American Economic Review*, 84(5): 1350–1368.

[18] Caballero, Ricardo, and Mohamad Hammour (1996), "On the Timing and Efficiency of Creative Destruction," *Quarterly Journal of Economics*, 111(3): 805–852.

[19] Card, David, Ana R. Cardoso, and Patrick Kline (forthcoming), "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics.*

[20] Card, David, Jörg Heining, and Patrick Kline (2013), "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128(3): 967–1015.

[21] Card, David, and Thomas Lemieux (2001) "Dropout and Enrollment Trends in the Post-War Period: What Went Wrong in the 1970s?" in J. Gruber (Ed.) *Risky Behavior Among Youths: An Economic Analysis*. Chicago: The University of Chicago Press, pp. 439–482.

[22] Carnevale, Anthony P., Tamara Jayasundera, and Dmitri Repnikov (2014), *Understanding Online Jobs Data: A Technical Report*, Washington D.C.: Georgetown University Center on Education and the Workforce.

[23] Castex, Gonzalo, and Evgenia K. Dechter (2014), "The Changing Role of Education and Ability in Wage Determination," *Journal of Labor Economics*, 32(4): 685–710.

[24] Charles, Kerwin K., Erik Hurst, and Matthew J. Notowidigdo (2015), "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," mimeo.

[25] Clemens, Jeffrey and Michael J. Wither (2014), "The Minimum Wage and the Great Recession: Evidence of Effects on the Employment and Income Trajectories of Low-Skilled Workers," NBER working paper #20724.

[26] Clemens, Jeffrey (2015), "The Minimum Wage and the Great Recession: Evidence from the Current Population Survey," NBER working paper #21830.

[27] Davis, Stephen J., Jason Faberman and John Haltiwanger (2012), "Labor Market Flows in the Cross Section and over Time," *Journal of Monetary Economics*, *Carnegie-Rochester Series on Public Policy* special issue, 59(1): 1–18.

[28] Deming, David J. (2015), "The Growing Importance of Social Skills in the Labor Market,"NBER Working Paper #21473.

[29] Deming, David J., and Lisa B. Kahn (2016), "Firm Heterogeneity in Skill Demands," mimeo.

[30] Devereux, Paul (2002), "Occupational Upgrading and the Business Cycle," *Labour Economics*, 16(3): 423–452.

[31] Foote, Christopher L., and Richard W. Ryan (2015), "Labor Market Polarization Over the Business Cycle," *NBER Macroeconomics Annual*, 29: 371–413.

[32] "Forsythe, Eliza (2016), "Why Don't Firms Hire Young Workers in Recessions?" mimeo.

[33] Goos, Maarten, and Alan Manning (2007), "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics,* 89(1): 118–133.

[34] Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (2001), "Equilibrium Unemployment," *Journal of Monetary Economics*, 48(1): 109–152.

[35] Greenfield, Rebecca (2014), "Why the New York Times Hired a Biology Researcher as Its Chief Data Officer," *Fast Company*, February 12.

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

[37] Hall, Robert (2005), "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1): 50–65.

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

[39] Isserman, Andrew M., and James Westervelt (2006), "1.5 Million Missing Numbers: Overcoming Employment Suppression in County Business Patterns Data," *International Regional Science Review*, 29(3): 311–335.

[40] Jaimovich, Nir, and Henry E. Siu (2015), "The Trend is the Cycle: Job Polarization and Jobless Recoveries," http://faculty.arts.ubc.ca/hsiu/work/polar20150503.pdf.

[41] Jaimovich, Nir, and Henry E. Siu (2012), "The Trend is the Cycle: Job Polarization and Jobless Recoveries," http://faculty.arts.ubc.ca/hsiu/research/polar20120331.pdf.

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

[43] Kahn, Lisa B., and Erika McEntarfer (2015), "Employment Cyclicality and Firm Quality," NBER working paper #20698.

[44] Koenders, Kathryn, and Richard Rogerson (2005), "Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries," *Federal Reserve Bank of St. Louis Review*, 87(4): 555–580.

[45] Kuhn, Peter, and Kailing Shen. (2013), "Gender Discrimination in Job Ads: Evidence from China," *Quarterly Journal of Economics*, 128(1): 287–336.

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

[47] Martins, Pedro S., Gary Solon and Jonathan P. Thomas (2012), "Measuring What Employers Do About Entry Wages Over the Business Cycle: A New Approach," *American Economic Journal: Macroeconomics*, 4(4): 36–55.

[48] McCoy, Terrence (2016), "The Technology that Could Revolutionize the War on Hunger," *Washington Post*, June 16.

[49] Michaels, Guy, Ashwini Natraj, and John Van Reenen (2014), "Has ICT Polarized Skill Demand? Evidence from Eleven Countries Over Twenty-Five Years."" *Review of Economics and Statistics*, 96(1): 60–77.

[50] Moffitt, Robert (2012), "The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations," *Brooking Papers on Economic Activity*, fall.

[51] Mortensen, Dale. and Christopher Pissarides (1994), "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3): 397–415.

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

[53] 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): 2509–2539.

[54] Molloy, Raven, Christopher L. Smith, and Abigail Wozniak (2013), "Declining Migration within the US: The Role of the Labor Market," FEDS Working Papers #2013-27

[55] Notowidigdo, Matthew J. (2013), "The Incidence of Local Labor Demand Shocks," University of Chicago, mimeo.

[56] Oldenski, Lindsay (2012), "Export versus FDI and the Communication of Complex Information," *Journal of International Economics*, 87(2): 312–322..

[57] 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): 1–29.

[58] Pierce, Justin R., and Peter K. Schott (2015), "Trade Liberalization and Mortality: Evidence from U.S. Counties," Yale mimeo.

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

[60] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2015), "Upskilling During the Great Recession: Do Employers Demand Greater Skill when Skilled Workers are Plentiful?", mimeo.

[61] Shimer, Robert. (2004), "The Consequences of Rigid Wages in Search Models," *Journal of the European Economic Association*, 2(2-3): 469–479.

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

[63] 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.

[64] Solon, Gary, Robert Barsky, and Jonathan A. Parker (1994), "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics*, 109(1): 1–25.

[65] U.S. Department of Labor (1977), *Dictionary of Occupational Titles, 4th edition*. Washington, D.C.: U.S. Department of Labor.

[66] von Wachter, Till, and Elizabeth Handwerker (2009), "Variation in the Cost of Job Loss by Worker Skill: Evidence Using Matched Data from California, 1991–2000," mimeo.

[67] von Wachter, Till, Philip Oreopoulos, and Andrw Heisz (2012), "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heteroegeneity in the Market for College Graduates," *American Economic Journal: Applied Economics*, 4(1): 1–29.

[68] Yagan, Danny (2016), "The Enduring Employment Impact of Your Great Recession Location," mimeo.

A Appendix

It is estimated that as of 2014 between 60% and 70% of all job postings are found online (Carnevale, Jayasundera, and Repnikov 2014). Indeed, The Conference Boarddiscontinued its-long-running, print-based Help-Wanted Advertising Index in 2008, after having begun a Help-Wanted Online Index in 2005 (HWOL).[59] Several other private-sector firms also began to track online job postings in the 2000s by using web-crawling and data-scraping methods. In this study, we employ data from one such firm, Burning Glass Technologies. This appendix discusses the representativeness of the data and investigates whether representativeness has changed over the time period of analysis.

A.1 Occupation-Industry Composition in BG

The BG 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.[60] It is also possible that the BG database is not representative even of online job postings, as comprehensiveness rests on the strength of the company's algorithms to code information in the ads and get rid of duplicates. Carnevale, Jayasundera, and Repnikov (2014) show that the occupation-industry composition of the BG data are similar to that of the Conference Board's HWOL. Moreover, the authors audited a sample of job postings in the BG database and compared them to the actual text of the postings, finding that the codings for occupation, education, experience were at least 80% accurate.[61]

Figure A1 plots the distribution of BG ads across major industry groups, sorted from largest to smallest (green bars), as well as the distribution of job vacancies in JOLTS (purple diagonal-lined bars). As mentioned, the BG database is meant to capture only electronically posted job ads; the universes of the data sources are thus not identical, but JOLTS is the best comparison available.[62] Despite the sample differences, the industry distributions match each other reasonably well. BG is overrepresented in health care and social assistance, as

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

[60]Rothwell (2014) compares the occupational distributions from an extract of BG to those from state vacancy surveys for select metropolitan areas for which 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 service workers are underrepresented.

[61]Furthermore, since BG regularly revises and attempts to improve its algorithms (applying them retroactively on the complete historical database of postings), and our extract is more recent than the one studied by Carnevale, Jayasundera, and Repnikov, it seems reasonable that their accuracy figure would be a lower bound for our sample.

[62]Both data sets cover 2007 and 2010–2014. JOLTS data are based on a monthly, nationally representative sample of approximately 16,000 business establishments drawn from unemployment insurance records; they count as a vacancy or job opening any position (including temporary and seasonal ones) for which work could start within 30 days and that the employer is actively trying to fill through a variety of means, of which posting a job ad (electronic or otherwise) is only one.

Figure A1: Industry Distributions: BG, JOLTS: 2007, 2010-2014

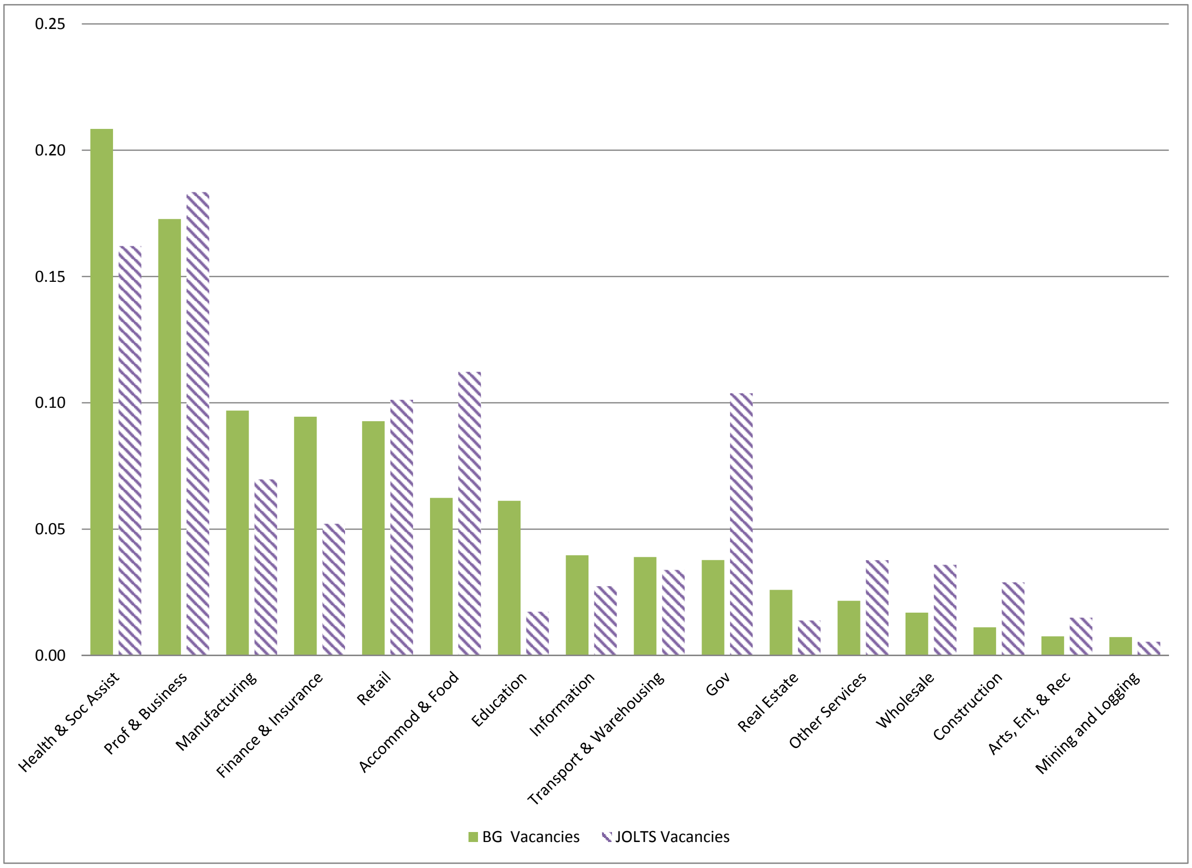

well as in finance and insurance and education. It is underrepresented in accommodation and food services, government, and construction. However, most differences are small in magnitude.

A great advantage of the BG data over the JOLTS is that they allow us to categorize jobs by occupation at a detailed level. We thus also compare the occupational distribution of BG job ads to both the stock and flow of employment in the U.S. We should not expect online job ads to precisely match either comparison group since occupations differ in turnover rates that would necessitate new hires (flows), and they also differ in the extent to which they use vacancy postings (rather than informal hiring channels) to fill a slot. However, these comparisons help build intuition for the BG data set.

Figure A2 plots the distribution of BG ads across major occupation groups, sorted from largest to smallest (green bars).[63] We show the distribution of the stock of employment based on the Bureau of Labor Statistics' Occupational Employment Statistics (OES) data (light blue, horizontal lines). We also show the occupational distribution of new job starts (job flows) based on longitudinally linked Current Population Survey (CPS) data (dark blue, diagonal lines).[64]

[63] For clarity, we use 2-digit Standard Occupational Classification codes in the figure. The regression analyses use more granular codings.

[64] All data sets cover 2007 and 2010–2014. (2014 is the most recent date for which OES data are available.)

Figure A2: Occupation Distributions: BG, New Jobs (CPS) and Employment (OES)

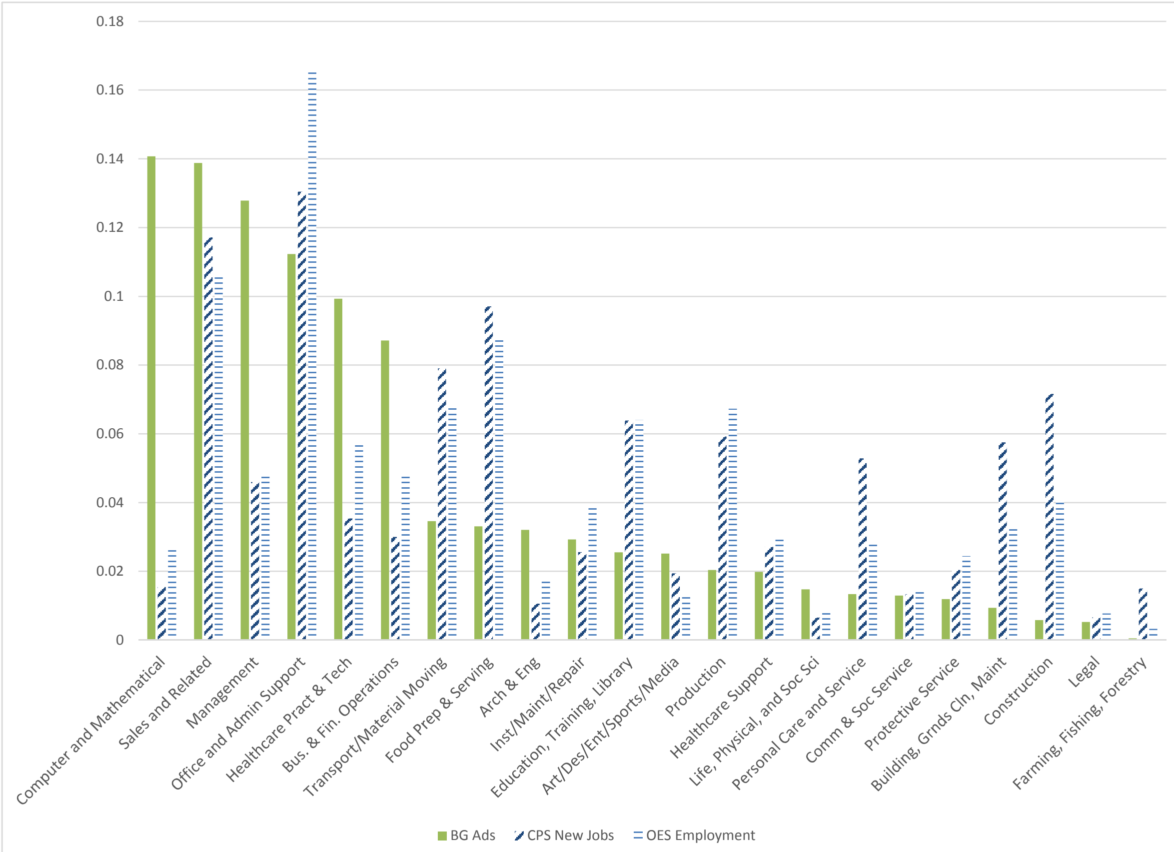

Perhaps not unexpectedly, BG has a much larger representation of computer and mathematical occupations, more than four times the OES and CPS shares. BG is also overrepresented among management, healthcare practitioners, and business and financial operations, although to lesser degrees. On the other hand, BG data are underrepresented in many of the remaining occupations, for example, in transportation, food preparation and serving, production, and construction. The OES and CPS distributions agree more closely, although there are notable gaps among occupations known to have very high (or very low) rates of turnover.

A.2 Representativeness of BG Data over Time

As noted in the text, our primary concern is that the representativeness of the sample changes over time. This would be a threat to internal validity in our analysis. Figure A3 gives a general sense of whether the representativeness of BG has changed over our sample period. On the x-axis we plot the deviation of the BG occupation share in 2007 from that

For these occupation descriptive analyses we apply the same sample selection criteria to the BG data as we do for the regression analyses. Most notably, we exclude ads with missing firm identifiers. We define a new hire in the CPS as an individual who, from month t to month $t+1$, transitioned from non-employment to employment or reported a new employer.

Figure A3: Representativeness of BG Occupations, Relative to New Jobs (CPS)

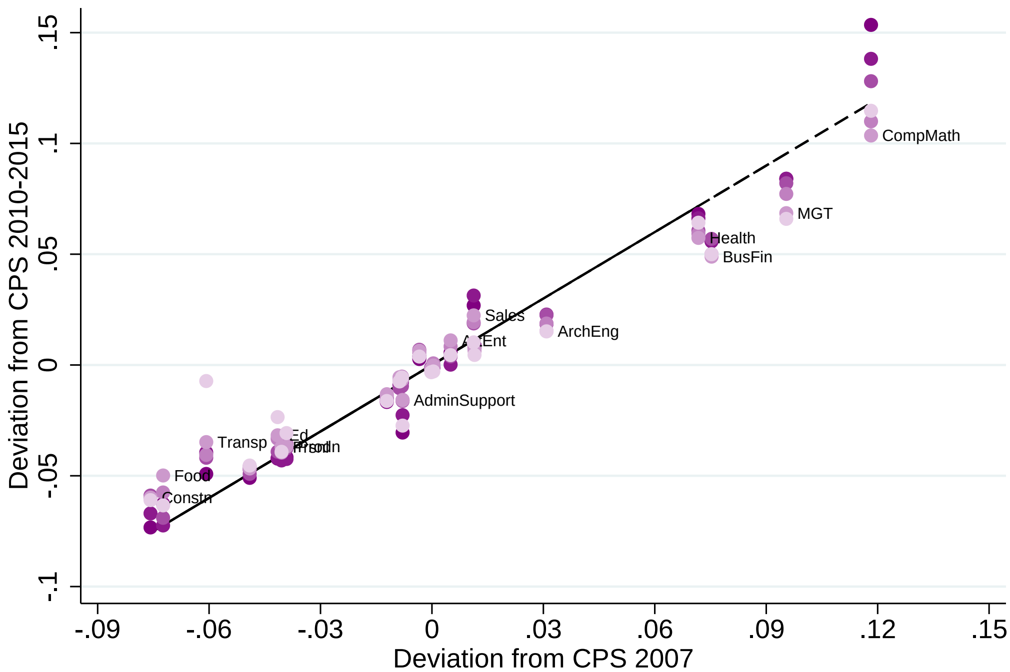

The x-axis is the BG ad share in an occupation in 2007 minus the CPS new job share in the same occupation in 2007.
The y-axis is these differences for each year from 2010-2015. Darker shades are earlier years, lighter shades are later.

occupation's share of CPS new job starts in the same year. For example, computer and mathematical occupations are shown on the far right at roughly 11 percentage points (ppts) overrepresentation in BG compared to CPS. Construction is on the far left at roughly 7 ppts underrepresented. On the y-axis we plot the deviation of the BG occupation share from its CPS share for each of the later years in the data. The markers are color-coded by year. The darkest markers plot the (2007,2010) representativeness pair for each occupation; the lightest markers plot the (2007,2015) representativeness pair. We also plot the 45-degree line as a benchmark: if representativeness of the BG data, relative to the CPS, remained constant over time, all markers should line up on the 45-degree line.

The figure shows that changes in representativeness over this time period are very small (most of the markers are close to the 45 degree line). To the extent that changes did occur, there is a tendency for them to have been in the direction of closer representativeness to the CPS. Computer and mathematical occupations, management occupations, and architecture and engineering occupations appear to have become less overrepresented while health care and business and finance look fairly unchanged; administrative support, food, transportation, and production occupations have become slightly less underrepresented. For most of these occupations, though, the differences are quite small.

A.3 Compustat Sample

Beginning with the 10,436 firms used to construct the 2007–2010 change sample (7.4% of firms and 67% of weighted observations, used in figure 5), we employ a sequential matching procedure. We first match based on exact name , after cleaning to remove punctuation and words that are sometimes abbreviated (e.g., “Inc.”). This step accounts for 81% of all Compustat-matched firms and 76% of weighted observations. Second, we use a fuzzy-match program called Winpure and assign a match if the program determines at least a 93% probability of a match (5% of matched firms and matched weighted observations). Third, we add the sample of firms matched by Deming and Kahn (2016), which uses only BG firms posting in 2014. We thus match a total of 46% of weighted observations and 18% of firms from the 2007–2010 change sample, or 30% of weighted observations in the sample as a whole – though we did not attempt to match observations that did not meet the 2007–2010 change criteria.[65]

[65] For context, the size of employment in Compustat is roughly half that of total employment in the U.S. For example, in 2014, the sum of employment listed in companies in Compustat was 70,505,000 and total payroll employment averaged 139,042,000. The Compustat employment figure includes both domestic and foreign workers, with no way to distinguish between the two. However, the employment comparison provides a useful benchmark.

Figure A4: Impact of MSA-Specific Employment Shock on Education Requirements

(a) Education Requirements

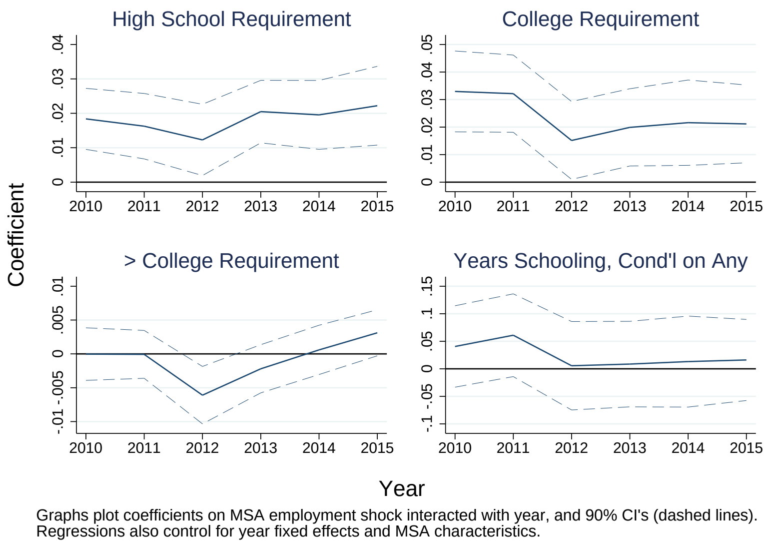

(b) Experience Requirements

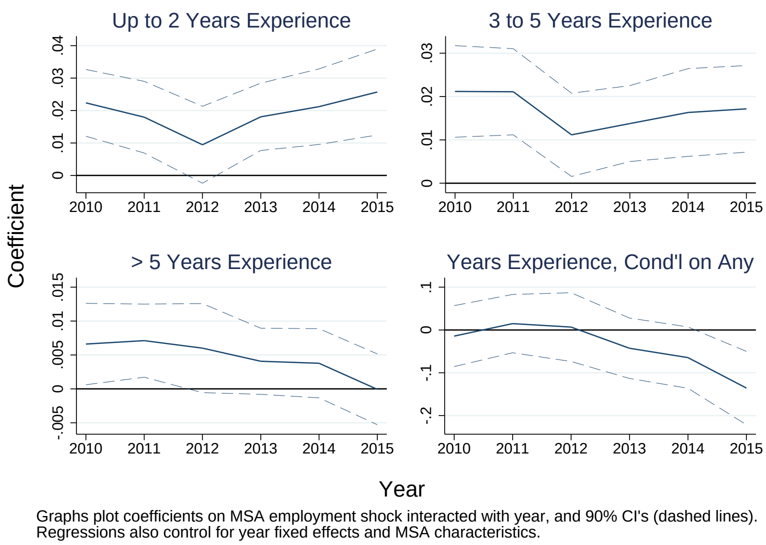

Figure A5: Within and Across Occupation Changes in Skill Requirements

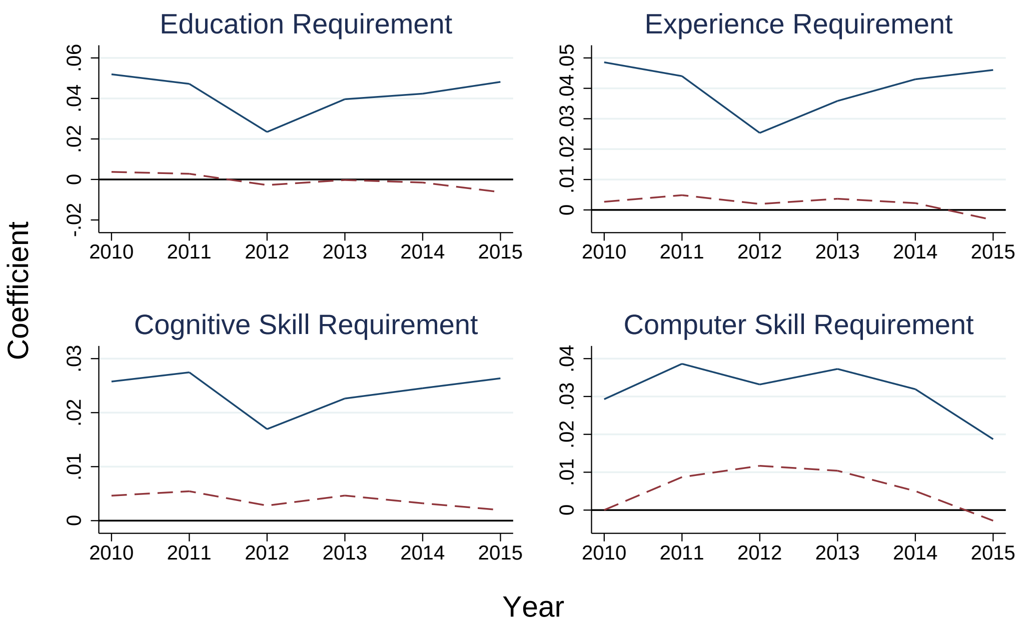

Blue solid line = within-occ-MSA change, red dashed line = across-occ-MSA change

Graphs plot impact of MSA shock on only change in each element, holding constant
the other element at its t level. Regressions also control for year fixed effects and MSA characteristics.

Table A1a: Education and Experience Robustness Checks

	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	Education Requirements				Experience Requirements			
Shock*2010	0.0523***	0.0429***	0.0519***	0.0514***	0.0490***	0.0442***	0.0489***	0.0485***
	(0.0136)	(0.0126)	(0.0129)	(0.0128)	(0.0134)	(0.0138)	(0.0130)	(0.0130)
Shock*2011	0.0474***	0.0359***	0.0466***	0.0461***	0.0443***	0.0340***	0.0439***	0.0435***
	(0.0131)	(0.0119)	(0.0125)	(0.0126)	(0.0134)	(0.0129)	(0.0132)	(0.0133)
Shock*2012	0.0235*	0.0212*	0.0230*	0.0228*	0.0256*	0.0199	0.0248*	0.0246*
	(0.0128)	(0.0121)	(0.0124)	(0.0124)	(0.0136)	(0.0136)	(0.0134)	(0.0134)
Shock*2013	0.0400***	0.0343***	0.0404***	0.0405***	0.0363***	0.0312**	0.0356***	0.0356***
	(0.0120)	(0.0120)	(0.0114)	(0.0114)	(0.0122)	(0.0124)	(0.0121)	(0.0121)
Bartiik*2014	0.0427***	0.0320**	0.0433***	0.0438***	0.0434***	0.0303**	0.0430***	0.0433***
	(0.0144)	(0.0144)	(0.0137)	(0.0136)	(0.0140)	(0.0139)	(0.0137)	(0.0136)
Shock*2015	0.0486***	0.0327**	0.0522***	0.0530***	0.0465***	0.0328**	0.0466***	0.0477***
	(0.0143)	(0.0142)	(0.0140)	(0.0141)	(0.0142)	(0.0147)	(0.0138)	(0.0139)
# Cells	193,013	222,058	193,013	193,013	193,013	222,058	193,013	193,013
R-Squared	0.044	0.072	0.362	0.374	0.068	0.102	0.342	0.354
Includes missing firms		X				X		
Occupation FE's			X	X			X	X
Occupation time trends				X				X

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

Notes: Regressions estimated at the msa-occupation (4-digit SOC) year level. Column 2 includes ads with no firm listed. Regressions also include year fixed effects and MSA characteristics obtained from the ACS. The dependent variable is the within occupation-MSA change in the probability that an ad lists an education (left panel) or experience (right panel) requirement from 2007. Observations are weighted by the occupation's ad share in the MSA-year times the size of the MSA labor force in 2006. Standard errors are clustered at the MSA level. Shock is the change in projected year-over-year employment growth in the MSA from 2006 to 2009, divided by the 90-10 differential in the variable across all MSAs.

Table A1b: Cognitive and Computer Skill Robustness Checks

	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	Cognitive Skill Requirements				Computer Skill Requirements			
Shock*2010	0.0261***	0.0168**	0.0254***	0.0255***	0.0294**	0.0234**	0.0293***	0.0304***
	(0.00721)	(0.00713)	(0.00663)	(0.00663)	(0.0120)	(0.0108)	(0.0105)	(0.0102)
Shock*2011	0.0276***	0.0196***	0.0267***	0.0266***	0.0388***	0.0336***	0.0377***	0.0379***
	(0.00724)	(0.00652)	(0.00673)	(0.00676)	(0.0106)	(0.0101)	(0.00945)	(0.00941)
Shock*2012	0.0172**	0.0134**	0.0168***	0.0167***	0.0333***	0.0322***	0.0322***	0.0320***
	(0.00679)	(0.00597)	(0.00633)	(0.00633)	(0.0123)	(0.0109)	(0.0109)	(0.0109)
Shock*2013	0.0230***	0.0159***	0.0224***	0.0223***	0.0374***	0.0344***	0.0367***	0.0364***
	(0.00628)	(0.00565)	(0.00604)	(0.00604)	(0.00980)	(0.00962)	(0.00863)	(0.00867)
Bartiik*2014	0.0250***	0.0170***	0.0251***	0.0252***	0.0321***	0.0258**	0.0327***	0.0324***
	(0.00632)	(0.00634)	(0.00609)	(0.00609)	(0.00997)	(0.0102)	(0.00890)	(0.00883)
Shock*2015	0.0270***	0.0234***	0.0284***	0.0288***	0.0187*	0.0190*	0.0213**	0.0210**
	(0.00716)	(0.00768)	(0.00691)	(0.00694)	(0.0110)	(0.00983)	(0.0101)	(0.0103)
# Cells	178,130	207,552	178,130	178,130	178,130	207,552	178,130	178,130
R-Squared	0.041	0.075	0.276	0.286	0.041	0.065	0.334	0.349
Includes missing firms		X				X		
Occupation FE's			X	X			X	X
Occupation time trends				X				X

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

Notes: See panel A.

Table A2: Probability of Being Missing from a Subsample

	(1)	(2)	(3)
	Change in Pr(Missing from):		Pr(missing from
	Firm Sample	Compustat Sample	2007-10 Change
Shock*2010	-0.0213	-0.0131	
	(0.0174)	(0.0116)	
Shock*2011	-0.0249	-0.00136	-0.00378
	(0.0176)	(0.0102)	(0.0145)
Shock*2012	-0.0639***	0.00151	0.00913
	(0.0194)	(0.0102)	(0.0155)
Shock*2013	-0.0100	0.00113	-0.0139
	(0.0203)	(0.0106)	(0.0128)
Bartiik*2014	0.00897	0.00808	-0.0166
	(0.0265)	(0.0104)	(0.0152)
Shock*2015	0.0319	0.00198	-0.0270*
	(0.0223)	(0.0110)	(0.0153)
# MSA-year Cells	222,058	2,286	1,905
R^2	0.093	0.364	0.682

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

Notes: In column 1, observations are occupation-MSA-year cells, while in columns 2 and 3 observations are MSA-year cells. Dependent variables in columns 1 and 2 are the change in the share of ads missing from the firm and Compustat sample, respectively, relative to 2007. The dependent variable in column 3 is the share of ads to firm-MSAs that posted in 2007 and 2010. Regressions also include year fixed effects and MSA characteristics obtained from the ACS. Observations are weighted by the size of the MSA labor force in 2006. Standard errors are clustered at the MSA level. Column 1 includes all ads while columns 2 and 3 restrict to ads with a non-missing firm.