

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

September 5, 2017

Abstract

We show that skill requirements in job vacancy postings differentially increased in MSAs that were hit hard by the Great Recession, relative to less hard-hit areas. These increases persist through at least the end of 2015 and are correlated with increases in capital investments, both at the MSA and firm-levels. We also find that effects are most pronounced in routine-cognitive occupations, which exhibit relative wage growth as well. We argue that this evidence is consistent with the restructuring of production toward routine-biased technologies and the more-skilled workers that complement them, and that the Great Recession accelerated this process.

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, Nick Bloom, Jeff Clemens, David Deming, David Green, John Horton, Hilary Hoynes, Peter Kuhn, Fabian Lange, Steve Malliaris, Adrien Matray, Alicia Sasser Modestino, Daniel Shoag, Henry Siu, Basit Zafar, four anonymous referees and seminar participants at the 2016 AEAs, the Atlanta Fed, Boston College, Boston University, Brookings, Georgetown University, Harvard, the LSE, MIT, NBER Summer Institute (Macro and Labor), Northwestern, Princeton, Rutgers, SOLE/EALE 2015 meetings, Trans Pacific Labor Seminar 2015, University of British Columbia, University of Chicago, University of Illinois, University of Kansas, University of Michigan, University of Notre Dame, University of Pennsylvania, University of Pittsburgh, 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 the 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 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 can 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 still relatively little direct evidence on how firms restructure employment in the face of technological change, and, in particular, whether this restructuring is gradual or episodic.³

In this paper we investigate how the demand for skills changed over the Great Recession. We use a new data set collected by Burning Glass Technologies that contains the near-universe of electronically posted job vacancies 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 experi-

¹See 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).

²Many 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).

³For example, Acemoglu and Restrepo (2017) find that the diffusion of industrial robots across U.S. commuting zones reduced aggregate employment and wages, and Harrigan, Reshef, and Toubal (2016) used matched employee-firm-level data from France to show that job polarization was more pronounced in firms with greater shares of technology-related occupations, but both focus on long run changes.

enced 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 requirements and about 2–3 percentage points (8-12%) more likely to state requirements for cognitive and computer skills. Moreover, the vast majority of this “upskilling” persists through the end of our sample in 2015. That is, even while most measures of local labor-market strength had converged back to pre-recession levels, differences in advertised skill demands remained. This is true holding constant a rich set of controls for the availability of skilled labor and the composition of ads across firms and occupations. In fact, we find that the very firms that upskilled early in the recovery drive the persistence later in our sample period.

These patterns collectively raise the possibility that a structural shift in the demand for skill occurred disproportionately in harder-hit MSAs. In particular, the skill requirements we explore (education, experience, cognitive, and computer) are known to complement routine-biased technologies (Autor, Levy, and Murnane 2003; Brynjolfsson and McAfee 2011). If a structural shift in line with RBTC is occurring, we would expect changes in these skill requirements to be accompanied by an accelerated adoption of such technologies, as well. Indeed, we find that increases in skill requirements are correlated with capital investments at both the MSA and firm levels. First, using the Ci Technology Database from Harte-Hanks, a market intelligence firm, we show that harder-hit MSAs exhibited a relative increase in IT investments, as measured by the adoption of personal computers, at the same time as they upskilled in job postings. These differences across MSAs emerge only after the Great Recession and, once again, persist through our sample period. Second, we link firms in our job postings database to those in the Harte-Hanks database, as well as to publicly traded firms in Compustat. We show that the firms increasing their capital investments, based on PC adoption and physical capital holdings, are also the ones that upskill the most. Thus, increased demand for labor skill appears closely linked to both general and IT capital investment.

If this increased investment is in fact related to routine-biased technologies, we would expect to see the strongest changes to labor characteristics for the jobs most susceptible to such technologies—routine ones. We thus additionally focus on different types of routine occupations, exploring joint changes in skill requirements, employment, and wages. Following Acemoglu and Autor (2011), we distinguish routine-cognitive occupations (e.g., clerical, administrative, and sales) from routine-manual ones (e.g., production and operatives), and we supplement the job ads data with Current Population Survey (CPS) and Occupation Employment Statistics (OES) data. For routine-manual occupations, we see evidence consistent with firms’ substitution of technology for labor—a sharp increase in layoff risk for harder-hit MSAs early on, and persistently depressed employment, with no particular impact on skill requirements. This is the traditional view exhibited in the polarization literature: employment losses concentrated in occupations we expect to be most readily replaceable by machines. Consistent with Jaimovich and Siu (2015), we show that these changes also

appear to be episodic around the Great Recession. However, in contrast to this conventional view of labor substitution, routine-cognitive occupations in harder-hit MSAs surprisingly exhibit only modest increases in layoff risk and no relative employment losses. Instead, we show that these occupations experience pronounced upskilling, as well as modest relative wage and employment growth after the recession. That is, rather than disappearing entirely, surviving routine-cognitive occupations appear to have become both relatively higher-skilled and more productive. These occupations thus became episodically less routine—and more cognitive—as a result of the Great Recession.

Taken together, our results suggest that firms located in areas more severely affected by the Great Recession were induced to restructure their production toward greater use of technology 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 evidence that the Great Recession spurred persistent changes in labor inputs, in a manner consistent with technological change. Several classes of models with adjustment costs can rationalize this result. For instance, firms may make productivity enhancing improvements in a recession because of a decline in the opportunity cost of restructuring (Hall 2005), a shift in managerial attention from growth to efficiency (Koenders and Rogerson 2005), or changes in the costs and benefits of making layoffs (Berger 2012, Jaimovich and Siu 2012). In addition, recessions may drive Schumpeterian cleansing, whereby, older, less-productive firms die, making way for newer, more-productive firms. Empirical support for adjustment-cost models has focused on the impacts of competition or trade shocks on productivity. For example, Bloom, Draca and Van Reenen (2016) show that increased Chinese import competition in Europe led to technological change within firms.⁴ Our paper adds to this literature by highlighting recession-induced changes in the firm-level demand for skill, which may have important consequences for labor market recoveries.

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 show that the bulk of upskilling occurs within occupations, suggesting this margin is quite important. Moreover, our finding that upskilling is concentrated within routine-cognitive occupations and is accompanied by relative wage growth implies that RBTC occurs both within and 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

⁴Additionally, Nickell (1996) provides evidence that increased competition is associated with faster total factor productivity growth; Syverson (2004a and 2004b) shows that productivity is higher in industries and geographies with greater substitutability of products across firms; and Bernard, Redding, and Schott (2011) show firms shift towards higher productivity products upon the liberalization of firm trade. Other examples are cited in each of these papers.

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 the fact that more-skilled workers still earn a sizable premium in the labor market (Card, Heining, and Kline 2013; Card, Cardoso, and Kline 2016). 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.⁵

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 (Davis, Faberman, and Haltiwanger 2012, 2013), 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 of the job posted, and these have generally used narrow slices of the data (Rothwell 2014), or data that are limited to one vacancy source (Kuhn and Shen 2013, Marinescu 2014). 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. Online job vacancies represent but one slice of the labor market, and, by their nature, will overrepresent growing firms (Davis, Faberman, and Haltiwanger 2013). Nonetheless, we show that linking vacancies to data on employment, wages, and capital investments—the last at the firm level—presents consistent evidence on how labor markets changed following the Great Recession.

We demonstrate that during the Great Recession firms changed not only 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, we find support that changes in demand for skill were episodic, resulting in a swath of displaced workers whose skills were suddenly rendered obsolete as firms ratcheted up their requirements. The need to reallocate workers on such a large scale may help drive jobless recoveries. It also has distributional consequences, given that low-skill workers are well known to suffer worse employment and wage consequences in recessions.⁶ Finally, this type of episodic reallocation 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 (Mof-

⁵Our analyses, however, do not explain why employment and wages have not grown in high-skill occupations. Deming (2017) proposes a compelling hypothesis that a rising importance of social skills, especially in conjunction with cognitive skills, can help account for this fact.

⁶See von Wachter and Handwerker (2008); Hoynes, Miller, and Schaller (2012); and Forsythe (2016).

fit 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, while section 3 summarizes our methodology. Section 4 presents new facts on changes in skill requirements as a function of local labor market conditions. Section 5 investigates how these changes are linked to capital investments. Section 6 examines cross-occupation heterogeneity in response to local labor market shocks on skill requirements, employment, and wages, with a particular focus on routine occupations. Section 7 concludes.

2 Data

Our data come from a unique source: microdata from nearly 100 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 these posting-level data for the years 2007 and from 2010 through 2015, covering every MSA in the United States.⁸

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

⁷Supporting 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.

⁸The database unfortunately lacks postings from 2008 and 2009. These years would be useful for completeness and for understanding the precise timing over which skill requirements changed; however, since hiring volume fell by one-third in 2008 and did not begin to recover until 2010 (per JOLTS), and our focus is on longer-term changes in hiring demand, additional data for the recession years is not integral for this paper. Our dataset 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.

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 mechanisms for upskilling.

However, the richness of the BG data comes with a few shortcomings. Notably, the database covers only vacancies posted on the Internet. First, Davis, Faberman, and Haltiwanger (2013) show that the distribution of vacancies in JOLTS overrepresents growing firms. Although roughly two-thirds of hiring is replacement hiring (Lazear and Spletzer 2012), vacancies in general will be somewhat skewed towards certain areas of the economy. Second, 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

Another downside of the BG data is that vacancies represent just one margin by which a firm may adjust labor inputs—through stated, but not necessarily realized, demand. For a complete picture, one would also like to see hires, separations, wages, and other measures (e.g., incumbent worker training, recruitment intensity (Davis, Faberman, and Haltiwanger 2013)). Thus, we also provide corroborating evidence on some of these margins using supplemental datasets, as described later.

We restrict our main BG sample to ads with non-missing employers that posted at least

⁹For 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.”

¹⁰Other 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. State vacancy surveys, conducted by a limited number of states, sometimes collect certain skill requirements, but cover only a few geographic areas and are generally not comparable across states (Carnevale, Jayasundera, and Repnikov 2014; Rothwell 2014).

¹¹BG 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.

10 ads over the sample period of 2007 and 2010-2015. Employer name is missing in 40 of postings, primarily from those listed on recruiting websites that typically do not reveal the employer.¹² 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, we have confirmed that the probability of satisfying this sample criterion (having a valid firm identifier) does not vary over the business cycle (see appendix A.8). 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. We choose these skill requirements for two reasons. First, they represent a broad swath of human capital measures in which both employers and economists have interest. Second, they reflect what the economics discipline has learned about technological change over the past 20 years (Autor, Levy, and Murnane 2003; Brynjolfsson and McAfee 2011). In particular, 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.¹³ 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 the ability to determine which tasks need to be performed at a given moment.

In accord with human capital theory, we believe more-educated workers or those with greater experience on the job will be better able to perform these functions.¹⁴ In Appendix

¹²When 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 a few problematic words, such as “Incorporated” or “Inc.” 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).

¹³This literature finds also that RBTC may indirectly affect low-skill, manual tasks (Autor and Dorn 2013). However, a downside of the BG sample is that low-skill jobs are underrepresented. We thus focus our analysis on the former margin, the degree to which employers shift demand from medium- toward high-skill tasks and workers.

¹⁴In 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

A.3, to cross-validate the data, we show that education requirements strongly correlate with average education levels of employed workers at the MSA and occupation levels.

We categorize cognitive and computer skill requirements based on the open text fields for skills. We designate an ad as requiring computer skills if it contains the key word “computer” or it is categorized as software by BG.¹⁵ We define cognitive skill requirements based on a set of key words chosen deliberately to match the non-routine analytical job tasks used in Autor, Levy, and Murnane (2003) and subsequently used by the majority of papers studying RBTC and polarization. We also ensure that the presence of these key words correlates with external measures of cognitive skill at the occupation level.¹⁶

This set of skills (education, experience, cognitive, and computer) aligns well with our priors on how jobs change with the availability of computers (Brynjolfsson and McAfee 2011). For example, a sales person who previously devoted most of his or her energy to client relations may now be required to use data analytics to better target packages to clients. This salesperson now needs computer and analytical skills, and some experience in the field may help in mapping data recommendations to practice. Similarly, thanks to machine vision technology, a quality control operator no longer need spend his or her time measuring and identifying the shapes of produced goods, but instead can be diverted to other tasks such as troubleshooting and making judgement calls in design optimization. This set of tasks requires higher cognitive function and intuition that can be gained by experience.¹⁷

Table 1 summarizes data for the primary regression sample.¹⁸ 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 associate degrees (not shown), master’s degrees, and professional degrees or PhDs. Converting the degrees to their modal equivalent years of schooling, the average education requirement, conditional on one being specified, is nearly 15 years.

The second column shows skill requirements averaged over 2010–2015. The third column

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.

¹⁵BG includes common software (e.g., Excel, PowerPoint, AutoCAD), as well as less common software and languages (e.g., Java, SQL, Python).

¹⁶Specifically, an ad is categorized as requesting a cognitive skill if any listed skills include at least one of the following phrases or fragments: “research,” “analy,” “decision,” “solving,” “math,” “statistic,” or “thinking.” The fraction of ads at the occupation level that contain each of these skills is strongly correlated with an O*NET measure developed by Deming (2017) meant to categorize cognitive occupations. We obtain this measure from Deming and Kahn (2017), who categorize a wide range of keywords found in the BG job ads into 10 general skills, including cognitive.

¹⁷It has been suggested by Deming (2017) and others that technology complements workers with interpersonal skills, since machines are still poor at reading and inferring human emotion. We have also analyzed changes in demand for a composite “social” skill requirement and obtained results very similar to those presented here on cognitive and computer skills.

¹⁸In 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.

Table 1: Summary Statistics

		mean (sd) in:		
		2007	2010-15	Change
Education Requirements:				
	Any	0.34 (0.06)	0.57 (0.05)	0.23***
	HS	0.09 (0.03)	0.20 (0.05)	0.10***
	BA	0.17 (0.05)	0.27 (0.08)	0.10***
	>BA	0.03 (0.01)	0.05 (0.01)	0.02***
	Years, Conditional on any	14.84 (0.40)	14.67 (0.44)	-0.18***
Experience Requirements:				
	Any	0.32 (0.06)	0.52 (0.07)	0.20***
	0-3	0.13 (0.03)	0.24 (0.03)	0.11***
	3-5	0.14 (0.03)	0.21 (0.04)	0.07***
	>5	0.05 (0.02)	0.08 (0.04)	0.03***
	Years, Conditional on any	3.52 (0.47)	3.34 (0.54)	-0.18***
Skill Requirements:				
	Any Stated Skills	0.73 (0.05)	0.91 (0.04)	0.18***
	Cognitive, conditional on any	0.22 (0.05)	0.34 (0.06)	0.11***
	Computer, conditional on any	0.28 (0.06)	0.39 (0.08)	0.11***
share of ads in 2010-2015:				
	Missing ACS match	0.08		
	Continuing firm	0.65		
	In Harte-Hanks, among continuing	0.78		
	In Compustat, among continuing	0.40		
		Mean	Min	Max
	# MSAs	381		
	Posts per MSA-year	21,779	132	1,231,417
	# Occupations (4-digit)	108		
	Posts per occupation-MSA-year	228	1	194,558
	# Firms	170,809		
	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, observations are weighted by the size of the MSA labor force in 2006. Missing ACS match is the share of weighted observations to MSAs that cannot be matched to the American Community Survey (weighted by the MSA labor force). Continuing firms are the fraction of 2010-2015 observations posted by a firm that also posted in 2007. In Harte Hanks (Compustat) among continuing firms are the share of weighted observations that post to a firm that can be matched to Harte Hanks (Compustat). All three statistics are calculated weighting by the 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.

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 (ppts), 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 ppts. These increases are again concentrated in the lower categories, so that the average, conditional on specifying any requirement, falls by about one-fifth of a year.

Finally, in 2007, 73% of weighted ads specify at least one specific, text-based skill requirement. Among these, 22% specify a cognitive skill requirement, and 28% have a computer requirement. In 2010–2015, 91% of ads have at least one text-based skill requirement, and the shares specifying cognitive skills or computer skills increase to roughly one-third and two-fifths, respectively. In regression analyses, we use the probability of posting a cognitive or computer skill 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.

These increases in stated skill demand could be driven by the national recession that took place between 2007 and the 2010–2015 period. 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.

The bottom panel of table 1 provides an idea of our sample coverage. We have a balanced panel of 381 MSAs, which contain an (unweighted) average of 21,779 posts per MSA-year. When we disaggregate to the four-digit occupation level, we have 108 occupations represented, with an average of 228 posts in each occupation-MSA-year.¹⁹ Finally, our data contain roughly 171,000 unique firms, which translate into an average of 14 posts in each firm-MSA-year.

¹⁹Though occupation is available in the BG data at the 6-digit Standard Occupation Classification (SOC) level, we restrict our attention to comparisons across and within 4-digit SOC codes, which provide more ads per occupation-MSA-year cell and ensure a balanced panel of occupation-MSAs across years in virtually all cases.

3 Methodology

Our goal is to understand how the Great Recession affected the demand for skill. Because we have only a short panel and need to worry about concurrent trends that may have affected online job ads (e.g., utilization, prices, pre-existing national trends in upskilling), we exploit cross-sectional geographic variation in the severity of the Great Recession. Our general approach is to examine temporal changes in skill requirements as a function of an MSA-level employment shock generated by the Great Recession.

Our initial regression specification is shown in equation (1). $outcome_{gmt}$ is any of several different measures associated with changes in labor skill demand (and eventually changes in production inputs more broadly, as discussed later) in MSA m , year t , and sometimes in subgroup g , (for example, occupation or firm). The left-hand side is the difference in the outcome variable between 2007 and year t . $shock_m$ is a measure of the local employment shock generated by the Great Recession, I^t are year dummies, $controls$ are additional control variables described in more detail below, and ε_{gmt} is an error term.

$$(1) \quad outcome_{gmt} - outcome_{gm2007} = \alpha_0 + [shock_m * I^t]\alpha_1 + I^t + controls + \varepsilon_{gmt}.$$

The variable $shock_m$ is fixed at the MSA-level for our entire sample period; we describe its construction in detail below. Through an exhaustive set of $shock_m$ -year interactions, the regression estimates the impact of the local employment shock on the *change* in skill requirements (or other outcomes) for a given MSA (and group) between 2007 and a subsequent year. The difference specification implicitly controls for time-invariant factors at the MSA (or group-MSA) level. We use 2007 as the base year in most analyses since this is the only pre-recession year available in BG. Such a specification allows us to empirically investigate the timing and persistence of upskilling in relation to local labor market shocks through the vector of coefficients, α_1 . The inclusion of year fixed effects (I^t) means we identify the key coefficients purely off of differences across metro areas in the employment shock, rather than relying on the national shock itself.

We cluster standard errors by MSA to address possible serial correlation within an area. For regressions at the MSA-year level, we weight observations by the size of the MSA’s 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 any (endogenous) changes in ads posted. When we further disaggregate to the MSA-year-group level, we weight cells by the product of the 2006 MSA labor force and the group’s observation share within MSA-year (the observation shares sum to unity), so that more aggregate regressions produce results identical to those using more disaggregated data when the underlying specification is the same.

The key explanatory variable, $shock_m$, is the MSA-specific change in projected annual employment growth between 2006 and 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 averaged over 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), who also used changes in projected employment growth to predict local shocks.²⁰

Specifically, we define projected employment growth, $\Delta\hat{E}_{mt}$ in equation 2, where for K 3-digit industries, ϕ is the employment share of industry k in MSA m at time τ (in practice, τ is the average of 2004 and 2005), $\ln E_{kt}$ is the log of national employment in industry k in year t , and $\ln E_{k,t-1}$ is the log of national employment in the industry one year prior.²¹

$$(2) \quad \begin{aligned} \Delta\hat{E}_{mt} &= \sum_{k=1}^K \phi_{m,k,\tau} (\ln E_{kt} - \ln E_{k,t-1}) \\ shock_m &= \Delta\hat{E}_{m2009} - \Delta\hat{E}_{m2006} \end{aligned}$$

We then define $shock_m$ as the change in projected employment growth from peak to trough (2006 to 2009). The calculated values of $shock_m$ range from about -0.12 to -0.04 across MSAs, but to make the coefficients easier to interpret, we renormalize this variable so that a one unit change is equal to the difference between the 10th and 90th percentile MSAs, -0.026 log points; a larger value corresponds to a worse economic shock.

We use this Bartik measure, instead of actual employment growth, 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.²² 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 toward more-skilled labor generates structural mismatch. Nonetheless, we examine the robustness of our results to other measures of labor demand.

The top left panel of figure 1 summarizes the relationship between the Bartik employment shock and actual annual (log) employment growth at the MSA level (obtained from the BLS State and Metro Area Employment program). We estimate equation 1, which nets out the

²⁰Some other papers utilizing a form of Bartik shock include Blanchard and Katz (1992), Notowidigdo (2013), and Yagan (2016).

²¹We obtain seasonally adjusted national employment for each 3-digit industry-month from the BLS Current Employment Statistics program, and take an unweighted average over months to obtain E_{kt} . 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>.

²²For 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.

actual employment growth rate in 2007 through the differences specification, for 2000–2015, controlling only for year fixed effects. The coefficients, α_1 , thus represent difference-in-difference estimates: the change in actual employment growth between a given year t and 2007, for a hard-hit MSA (90th percentile employment shock) relative to a less hard-hit MSA (10th percentile employment shock).

We plot the coefficients α_1 and 95% confidence-interval bars (results for figure 1 are also displayed in columns 1–4 of table C1). For example, the point estimate of -0.04 in 2009 (top left panel of figure 1) indicates that a one-unit change in the Bartik shock is associated with an additional 0.04 log-point drop in employment growth between 2007 and 2009. Employment growth actually peaked in 2006, so the peak-to-trough difference associated with a one unit increase in $shock_m$ is -0.053 ($= -0.040 - 0.013$). This is roughly double the Bartik 90-10 gap of -0.026 associated with a one unit change of $shock_m$ used in the regression. The actual BLS variables are likely substantially noisier than projected employment growth and also are influenced by other factors, such as supply shocks.²³

The figure also shows that the shock is episodic, such that employment growth (relative to that in 2007) looks similar across MSAs early in the decade, regardless of the size of the shock they will eventually face in the Great Recession. Hard-hit MSAs peak slightly higher than less hard-hit MSAs in 2005 and 2006, then experience a sharp dip in employment growth from 2007–2010, followed by a recovery.

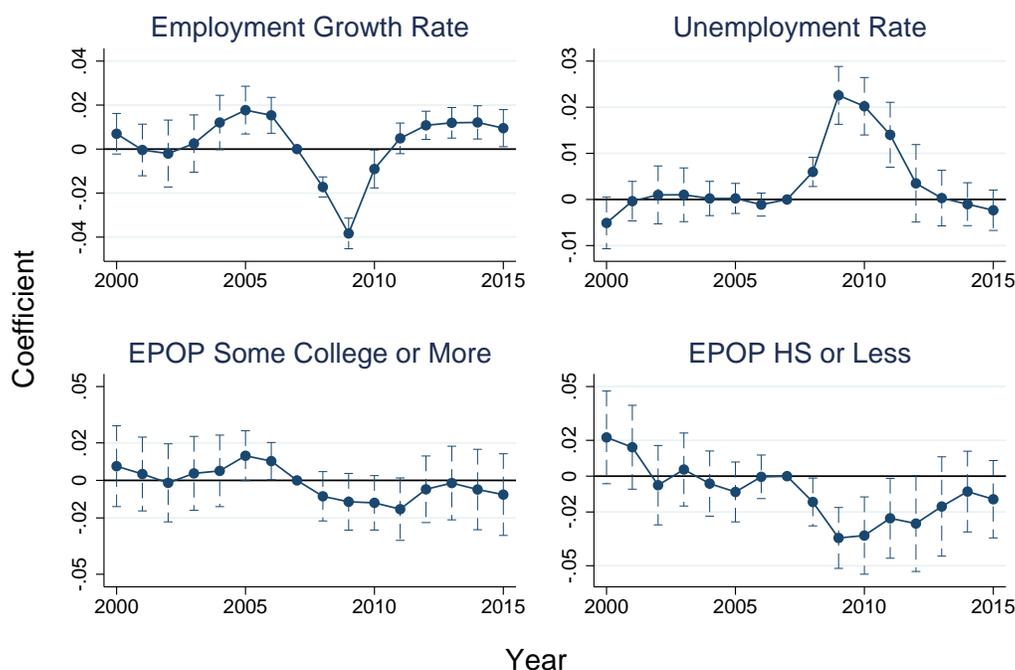
The Bartik shock measure is also highly correlated with movement in the unemployment rate (obtained from the BLS Local Area Unemployment Statistics program). The top right panel of figure 1 shows that a hard-hit MSA experiences an additional 2 percentage-point increase in the unemployment rate from 2007 to 2009, relative to a less hard-hit MSA. Again, areas look very similar in the period before the recession, and converge a few years after the recession ends.

Our primary regressions of interest involve changes in skill requirements within MSA using data that begins in 2007. Although the first-difference specification nets out differences across MSAs in the level of posted skill requirements, we cannot control for (or observe) any preexisting trends within MSAs in skill demand. Identification may be threatened if, for example, preexisting trends in upskilling were more prevalent in MSAs with industry mixes that would make them more or less susceptible to the demand shock.

The bottom panel of figure 1 helps speak to this concern by examining employment-to-population ratios (epop) by education group. We calculate these variables by MSA using CPS micro data, so they are naturally a bit noisier than the variables in the top panel (see appendix A.6 for details about this sample construction). The epops for workers with at least some college (bottom left) and for workers with a high school diploma or less (bottom right) are fairly similar across MSAs before the Great Recession. This is shown in the figure

²³Even though this is not the case for the Bartik shock, actual employment growth peaks slightly earlier for hard-hit MSAs, in 2005, but remains at almost the same magnitude in 2006.

Figure 1: Labor Market Variables and the MSA-Specific Employment Shock



Graph plots the coefficients on Bartik shock*year, as well as 95% CI bars. Regressions control for year fixed effects. Unemployment and employment growth rates are from the BLS LAUS. Employment-to-populations are author calculations based on the CPS.

by point estimates that are small in magnitude and generally statistically indistinguishable from zero prior to 2007. As with employment growth, the college epop does peak slightly higher in 2005 and 2006 for MSAs that will experience a worse shock; additionally, the epop for less-educated workers fared modestly better in these MSAs in 2000. However, differences are small and do not appear to be systematic trends. We explore the relationship between the shock measure and a range of additional labor market variables, below.

After 2007, both education groups experience relative drops in epops, though that for college workers is shallow and recovers quickly. The decline in epop for less-educated workers is both more severe and more sluggish to recover, and a gap of roughly 1 ppt still remains in 2015.²⁴ This lack of convergence may suggest that harder-hit areas had not fully recovered from the Great Recession by 2015. We revisit this lack of complete recovery below when we discuss our proposed mechanism.

To alleviate remaining concerns about differential pre-trends, we control, where specified, for a wide range of MSA characteristics—including demographics, educational attainment, and economic indicators—obtained from the American Community Survey (ACS), averaging years 2005 and 2006.²⁵ These controls help adjust for differences across MSAs in their

²⁴This finding is consistent with Yagan (2016) who uses IRS tax data to show that while unemployment rates had converged across U.S. commuting zones following the Great Recession, employment probabilities had not, holding constant a rich set of worker characteristics.

²⁵We chose years just prior to the Great Recession 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, migrated in the last year, is a high school drop out, has exactly a high school

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.

To summarize, we find that our constructed Bartik employment shock is episodic: although it is highly correlated with changes in employment growth rates and the unemployment rate during the Great Recession, the shock is not correlated with these labor market fundamentals before the Great Recession or several years into the recovery. As we cannot observe skill requirements before 2007, this is reassuring: the lack of pretrends in the labor market variables in figure 1, and for others described below, makes it less likely that areas were differentially trending in skill demand before 2007.

4 Skill Requirements and Local Employment Conditions

4.1 Main Results

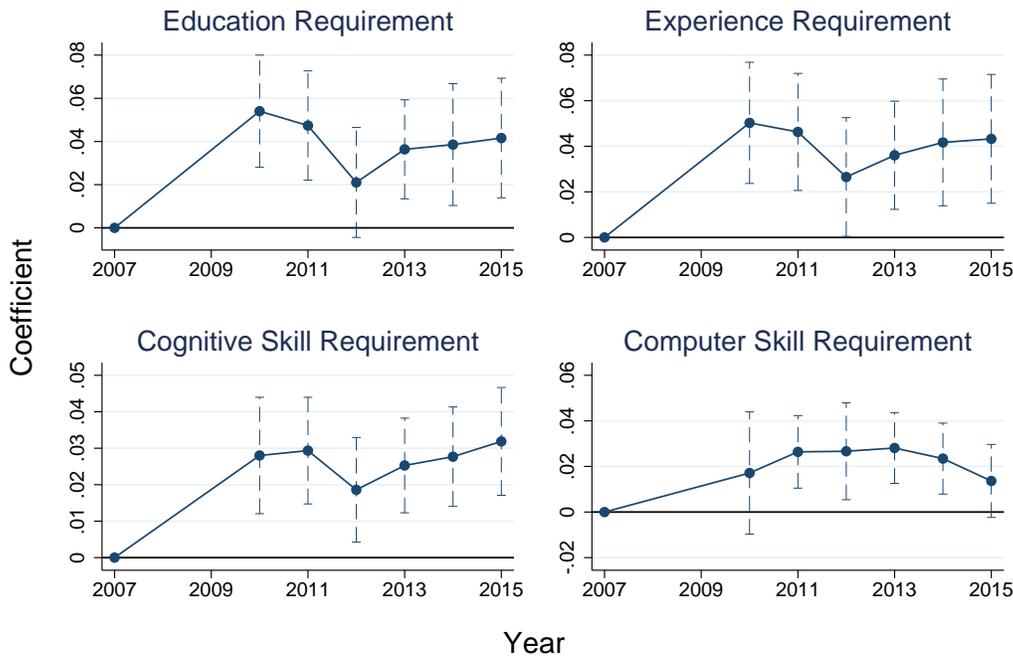
Figure 2 summarizes regression results from equation (1) for our four main dependent variables: the change in the share of ads posting any education requirement, any experience requirement, any cognitive requirement, and any computer requirement (results are also displayed in columns 5–8 of table C1). 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 95% confidence intervals. We use our preferred specification, which includes controls for MSA characteristics and year fixed effects.

Beginning with the top left panel, we find that, relative to 2007 levels, the probability of specifying any education requirement increases by 5.4 ppts in 2010 for an MSA experiencing a large employment shock (90th percentile) compared to an MSA experiencing a small shock (10th percentile). This difference-in-differences estimate implies an increase of 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 ppts larger than it was in 2007 for a hard-hit MSA, compared to a less hard-hit one. 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 2 exhibit remarkably similar patterns in both magnitudes and statistical significance. The probability of listing an experience requirement increases by 5.0 ppts (16%) between 2007 and 2010, and 85% of this increase remains in 2015. The

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 unmatched ads consisting of small MSAs not identifiable in the ACS. In such cases, we set the ACS controls to zero and include an indicator for not matching. In appendix B.2, we also include specifications that add controls for *changes* since 2000 in these variables.

Figure 2: Skill Requirements and the MSA-Specific Employment Shock



Graphs plot coefficients on MSA employment shock interacted with year, and 95% CI's (dashed lines). Regressions also control for year fixed effects and MSA characteristics.

probability of listing a cognitive requirement increases by 2.0 percentage points (12%), and this gap possibly widens slightly by 2015. Finally, the probability of listing a computer skill requirement also increases by roughly 2 ppts and remains elevated through 2015.

These patterns are in stark contrast to the labor market variables in figure 1. For employment growth, the unemployment rate, and the epop for college workers, hard-hit MSAs experience a more severe impact of the Great Recession that fully recovers within our sample time frame. For illustration, compare Detroit and Pittsburgh. The former, a hard-hit MSA, experienced a shock at about the 90th percentile, while the latter was at roughly the 10th percentile. Both MSAs had similar skill requirements in 2007; for example, in both areas about one-third of ads had an education requirement. By 2010, skill requirements increased in both MSAs, but Detroit's (actual, not predicted) increase was nearly 10 ppts larger for education and experience requirements and 2–4 ppts larger for cognitive and computer skill requirements. While unemployment rates had converged back to pre-recession levels in both MSAs, Detroit's elevated skill requirements persisted through 2015.

Figure 2 demonstrates that the case of Detroit and Pittsburgh is not isolated but systematic. In terms of their skill requirements, MSAs that looked similar before the Great Recession look quite different from each other in 2015, several years after the Great Recession ended.

To better understand the mechanisms underlying figure 2, table 2 provides regression results for within-occupation changes in skill requirements. In general, the distribution of

postings across high- and low-skilled jobs may vary for a variety of reasons. For example, differential job survival, use of word-of-mouth in recruiting, or time-to-fill across skill groups might generate patterns observed in figure 2, especially early in the recovery. However, we find that the primary driver of these patterns is increased skill requirements within similar types of jobs.

Each column of table 2 summarizes a separate regression of equation 1, at the occupation-MSA-year level, including MSA characteristics and year fixed effects. The results show significant upskilling effects of a magnitude comparable to the overall MSA-level effects.²⁶ For example, within occupation, the propensity to post an education requirement increases by 5.3 ppts in a hard-hit MSA, relative to a less hard-hit MSA, between 2007 and 2010. Although there is a temporary dip in 2012, at least three-quarters of this effect persists from 2013 through 2015. Similar patterns obtain for the remaining skill requirements. Indeed, we discuss below that these within-occupation increases in skill requirements completely account for the MSA-level upskilling effects found in 2; our upskilling results are not driven at all by changes in the occupation mix of postings. (This does not preclude variation in effects across occupations, and we examine such heterogeneity, with a particular focus on routine jobs, in section 6.)

In order to understand within-occupation skill demand changes along the intensive margin, we have also investigated the effect of the shock on specific levels of education and experience requirements. We present detailed results in Appendix B.1, but to summarize, we find effects throughout the distribution along expected channels: low-skilled jobs become more likely to ask for a high school diploma, higher-skilled jobs become more likely to ask for a college degree, and experience requirement increases are concentrated especially within the 1–5 year range.

One hypothesis for these results is that firms may become pickier when labor, and especially skilled labor, becomes more plentiful.²⁷ Then elevated skill requirements might reflect opportunistic behavior on the part of firms that cannot ordinarily attract (or afford) more-skilled workers in a tight market. This hypothesis would be compelling if the market for skilled workers remained more slack towards the end of our sample period, even while some labor market indicators had recovered. However, our results are similar when we include additional controls for local labor market variables by skill level, such as education-specific MSA-level unemployment rates, quit rates, and employment-to-population ratios. These controls account for changes in the supply of skilled labor due to, for example, differential quit behavior or changes in educational attainment brought on by the Great Recession or over the preceding decade (Charles, Hurst, and Notowidigdo 2015). We conclude that

²⁶We define occupations at the 4-digit SOC level. Virtually all ads posted in the 2010–2015 period are in occupation-MSAs that also posted in 2007, so we can generate within-occupation-MSA changes in skill requirements. The 0.36% of ads that cannot be matched back are dropped from these analyses.

²⁷Or, as in Menzio and Shi (2011), firms require a higher-quality match in a recession because of the negative productivity shock.

Table 2: Within-Occupation Changes in Skill Requirements

<i>Dependent Variable:</i>	<i>Education</i>	<i>Experience</i>	<i>Cognitive</i>	<i>Computer</i>
	(1)	(2)	(3)	(4)
<i>Shock*2010</i>	0.0528*** (0.0137)	0.0494*** (0.0135)	0.0276*** (0.00728)	0.0203** (0.00862)
<i>Shock*2011</i>	0.0478*** (0.0131)	0.0445*** (0.0134)	0.0284*** (0.00734)	0.0241*** (0.00718)
<i>Shock*2012</i>	0.0241* (0.0128)	0.0260* (0.0136)	0.0190*** (0.00695)	0.0207** (0.00851)
<i>Shock*2013</i>	0.0403*** (0.0121)	0.0364*** (0.0122)	0.0256*** (0.00642)	0.0250*** (0.00664)
<i>Shock*2014</i>	0.0430*** (0.0144)	0.0435*** (0.0140)	0.0265*** (0.00657)	0.0225*** (0.00675)
<i>Shock*2015</i>	0.0488*** (0.0143)	0.0465*** (0.0141)	0.0299*** (0.00727)	0.0131 (0.00800)
# Occ-MSA-Year Cells	192,842	192,842	177,971	177,971
R-Squared	0.045	0.069	0.040	0.034

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

Notes: Regressions are estimated at the msa-occupation (4-digit SOC) year level. The dependent variable is the MSA-occupation level annual change in skill requirements from 2007. All regressions control for year fixed effects and MSA characteristics from the ACS. Observations are weighted by the size of the MSA labor force in 2006 multiplied by the occupation's ad share in the MSA-year. 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.

opportunistic upskilling cannot be the primary driver of our results.²⁸

Appendix B.2 discusses these and other robustness checks in detail, which are summarized in appendix tables B1-B4. Our results broadly hold up to additional controls, different samples, variants of the Bartik shock, and different weights. For example, our estimates are robust to controls for occupation fixed effects and occupation-specific time trends, which allow occupations to systematically differ in their change in skill requirements from 2007, as well as in the slope of the change, across all MSAs. 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.

We have also explored heterogeneity within and across industry. We show in appendix table B5 that our results hold up to industry fixed effects and time trends, which is important in light of our identifying variation: industry composition in an MSA before the Great Recession. Our identification would be threatened by independent technology shocks concurrent to industries that experienced worse employment shocks or by systematic measurement error in industry shares (which could lead to spurious correlations in the shock across MSAs). The fact that our results obtain even within industry alleviates this simultaneity or measurement concern. We further show in appendix figure B3 that the upskilling effects tend to be concentrated in industries with locally consumed products, as would be expected given their greater sensitivity to local demand shocks.

Finally, one may be concerned about changes in the use of online job ads over our sample period. Rising familiarity with the internet, falling costs of posting jobs, increasing labor market tightness in the later period, and other factors may have brought more firms online to search for labor. The within-occupation and -industry results partially address the role that compositional changes in the use of online job ads may have on our results by restricting comparisons to similar types of jobs. However, they may not adequately control for heterogeneity across firms in, say, changes to their recruiting strategies or hiring needs. Such variations may be particularly pronounced during and after a recession.

In Appendix B.3 we conduct a formal decomposition exercise to apportion upskilling effects as a function of within- and between-firm responses. Indeed, we do find a large role for substitution between firms that stopped posting after 2007 and firms that began posting in 2010. As the latter post for higher skill requirements on average than the former, this substitution can account for nearly half of the full upskilling effect. However, we also find that nearly half of the effects shown in 2 can be attributed to changes in skill requirements within

²⁸In a pair of related papers and concurrent with our analysis, Sasser Modestino, Shoag, and Ballance (2016a and 2016b), using a version of our dataset, find evidence of upskilling in harder-hit US counties after the Great Recession and subsequent downskilling as markets improved, and argue that this pattern is driven entirely by firms opportunistically seeking more-skilled workers in a slack labor market. We disagree with this conclusion, which relies heavily on the small downward blip in 2012, seen also in our figure 2, rather than the more careful picture generated by using all available data. Examining heterogeneity within and across firms and occupations and other margins of adjustment—such as capital, employment, and wages—as we do in this paper, also implies more fundamental changes in production inputs and longer-lasting impacts.

firms, with a minimal role for compositional shifts across firms that post before and after the recession. This suggests that our results are not completely driven by the compositional changes mentioned above.

4.2 Discussion

We thus present strong evidence that employers in harder-hit MSAs were differentially induced to increase stated preferences for a range of skills. While most measures of local labor-market strength had converged back to pre-recession levels by 2015, differences in advertised skill demands remained. Furthermore, variation in the availability of skilled labor and compositional changes in the ads observed in our sample period are unlikely to explain the entire effects that we find.

This set of results raises the possibility that harder-hit MSAs differentially experienced a structural change in demand for skill. In particular, the skill requirements we investigate are complementary to routine-biased technologies. Did the Great Recession push an accelerated adoption of such technologies and accompanied hiring of cognitive workers to complement them? This could explain why skill requirements increase even within similar types of jobs. 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, in order to better reach and understand online readers, venerable journalistic organizations such as the *New York Times* now hire individuals with science training, not just journalism training, to be chief data officers (Greenfield 2014). It is also consistent with our finding that epops for workers with less education were slow to recover: rapid adoption of new technologies in hard-hit MSAs over this time period may have rendered certain worker skills obsolete, inducing labor force exit (also see Foote and Ryan 2015). Perhaps the epop of educated workers recovered rapidly precisely because of an increased demand for skill spurred by the Great Recession.

Several theoretical mechanisms may have induced firms to restructure. For example, in the classic Schumpeter (1939) cleansing model, this would occur because low-productivity firms shut down in the recession and resources are reallocated to firms with more-modern production technologies (see also Caballero and Hammour (1994, 1996) and Mortensen and Pissarides (1994)). Furthermore, this type of episodic restructuring could also occur because firms in harder-hit MSAs experience a greater negative product-demand shock that: (1) lowers the opportunity cost of adjusting production (Hall 2005), (2) shifts managerial attention from growth to efficiency (Koender and Rogerson 2005), (3) alters the costs of making layoffs (Mortensen and Pissarides 1994, Berger 2012)²⁹, and changes the incentives

²⁹Though not formalized, a sufficiently large negative product-demand shock could make layoffs worthwhile, offsetting any stigma or losses in terms of firm-specific human capital.

for a firm to invest in their workers’ human capital (Jaimovich and Siu 2012). We do not feel we have the ability to disentangle these mechanisms or provide strong support for any one model. Instead, we point out that these types of workhorse models in macroeconomics can rationalize the results that we see.

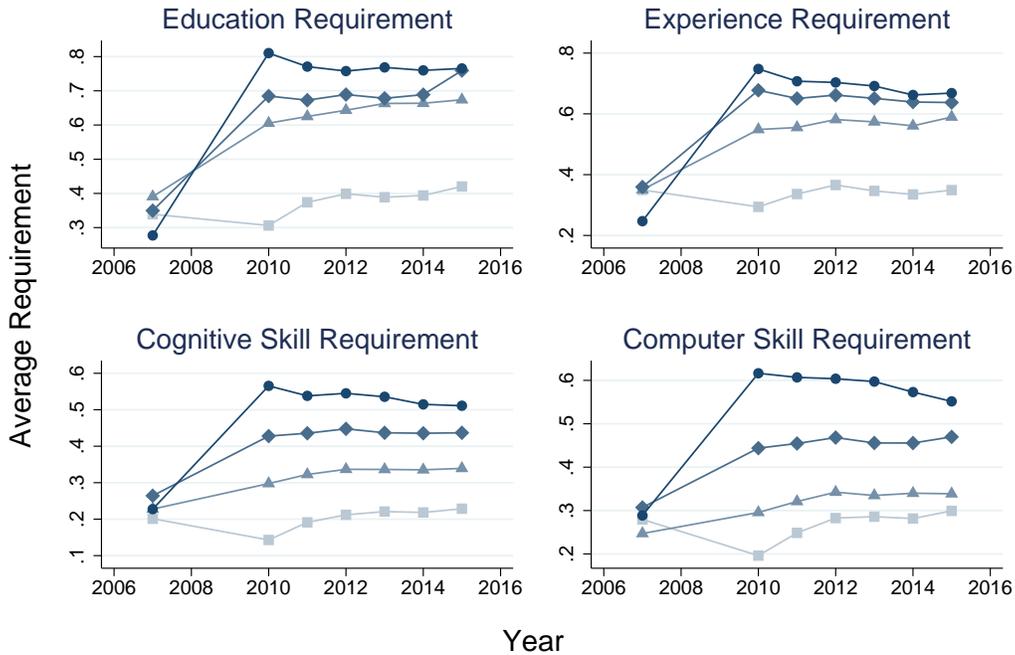
If firms are changing how they produce, and not simply whom they hire, changes to skill demand should persist among the same firms that initially upskilled. Note that the finding of within-firm upskilling, mentioned above, does not necessarily imply this point, as different firms may have upskilled at different times. The Burning Glass data provide an unprecedented glimpse at this margin at a detailed level, and we examine this prediction in figure 3. Here, we divide firms into (posting-weighted) quartiles based on changes in skill requirements between 2007 and 2010. We then plot the average skill requirements for each quartile over time.³⁰ Firms began at fairly similar average skill levels in 2007, although this similarity is not imposed by our exercise. 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, firms in the higher quartiles still had substantially higher skill requirements in their *new* ads than firms in the lower quartiles.

This within-firm persistence in upskilling holds up in regression analysis and is substantial. Though not shown, we find that, on average, 60–70% of a firm’s increase in skill requirements between 2007 and 2010 persists through 2015. Estimates are even larger when we instrument for the initial increase in upskilling with the Bartik shock. It could have been the case that the majority of firms increased skill requirements during the recession and reverted back later in the recovery (for example, in an attempt to opportunistically recruit while markets were slack), with higher skill demand in later years unrelated to the recession and driven by different firms. Instead, we find upskilling persists among the same firms both early and late in the recovery.

Futhermore, our finding that a substitution across old and new firms accounts for some of the upskilling effect could also be consistent with episodic restructuring. Substitution from failing (low-productivity) firms to new (high-productivity) firms is a hallmark prediction of “cleansing” models of recessions (Schumpeter 1939). In our data, we do not observe firm births and deaths, so the Schumpeter angle is difficult to fully assess. To provide some general intuition, we define a firm “closure” as when a firm posts in 2007 but not again in the post-recession period; we analogously define a firm “opening” as a firm that did not post in 2007 but did post in the later period. (We readily acknowledge we are abstracting away from hiring freezes and migration toward online job postings; we view this

³⁰We 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. Appendix A.8 shows that the probability of satisfying this restriction does not vary with the local labor market shock. Quartiles are defined separately for each skill measure, weighting by the firm average number of posts across 2007–2010.

Figure 3: Skill Requirements by Firm 2007-2010 Change



Graph plots average skill requirement by year and quartile of 2007-10 firm-level skill change. Circles, diamonds, triangles, and squares indicate skill change quartile from largest to smallest, respectively.

exercise as illustrative, not definitive.) We find that “opening” firms indeed have higher skill requirements than “closing” firms, even within occupation, and this is consistent with the Schumpeterian cleansing view.³¹

The episodic restructuring hypothesis has additional predictions that we can take to the data. First, if changes in skill requirements reflect changes in production inputs, we should see greater investments in capital, and in particular routine-labor replacing technologies, for firms located in harder-hit MSAs. Moreover, the two activities should be linked: the very firms upskilling in their labor demand should be the ones increasing their investments. Second, routine workers whose skills can be substituted with these technologies should experience an immediate contraction in labor demand, as evinced by a relative increase in layoffs of these workers in harder-hit MSAs during the recession, as well as relative employment declines in the recovery and beyond. In contrast, the occupations that are complementary to new technology should become more productive, from the increase in both physical and human capital, and thus should exhibit increases in relative wages. We explore these predictions in the next two sections.

³¹Moreira (2016), who shows that firms that begin in a recession are more productive than those that begin in an expansion, also provides support for the Schumpeterian view.

5 Capital

Under episodic restructuring, firms automate routine tasks with technology, which complements skilled labor. If this restructuring is occurring, then firms should also invest in physical capital around the time that they upskill. Information technologies (IT), in particular, have been linked to RBTC (Michaels, Natraj, and Van Reenan 2014). While investments in capital tend to be procyclical, and production of IT, in particular, has exhibited a secular decline (Byrne, Fernald, and Reinsdorf 2016), these trends could mask substantial heterogeneity.

We first investigate whether harder-hit MSAs are more likely to invest in IT over the Great Recession. To measure IT investment, we use the Ci Technology Database from Harte-Hanks (now known as Aberdeen), a market intelligence firm. The Harte-Hanks database (hereafter, HH) is created from surveys and interviews with high-level IT staff at millions of businesses worldwide each year. They collect data primarily to sell to major IT firms like IBM, Dell, and Cisco.³²

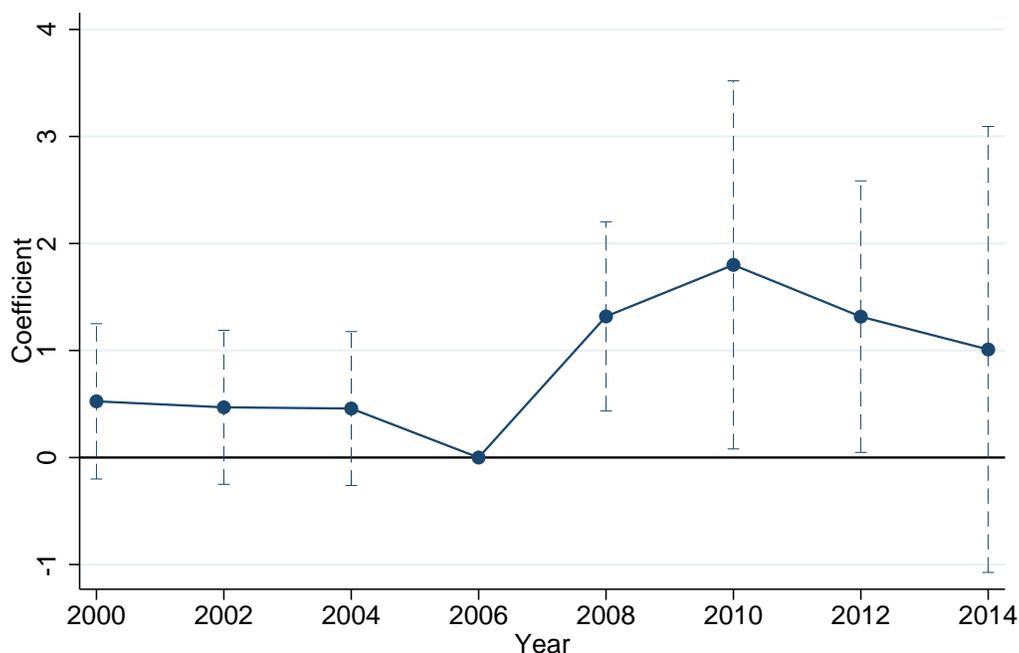
Following previous work using these data, our primary outcome measure is the number of personal computers (PCs) at a “site” (akin to business establishments). We have this measure consistently available in even years between 2000 and 2014, and we normalize by dividing by site employment in the pre-recession period.³³ We aggregate to the MSA-year level by taking an employment weighted average across sites.

Figure 4 (and column 9 of table C1) summarizes results from equation (1), with the MSA-level change in PCs per employee from the pre-recession period as the outcome. This graph provides evidence that firms located in harder-hit MSAs are more likely to intensify IT investment over the same time period. Our estimates imply that sites in a hard-hit MSA add an average of 1.5 PCs (per each pre-recession employee) between 2006 and 2012, relative to sites in less hard-hit MSAs. Though the confidence intervals are wide, this effect is statistically significant at the 5% level in 2008, 2010, and 2012. This differential increase experienced by hard-hit MSAs is substantial, roughly 60% more than the average increase

³²These data have been used in several studies. Bloom, Sadun, and Van Reenen (2012), for example, have used the HH data to show that US multinationals operating in Europe obtain higher productivity from IT investments than non-multinationals. Also using the data, Bloom, Draca, and Van Reenen (2016) show that Chinese import penetration increased technological change for exposed firms in Europe; Beaudry and Lewis (2014) show that variation in PC adoption across U.S. space can account for variation in declines in the gender pay gap; and Bresnahan, Brynjolfsson, and Hitt (2002) provide evidence that IT use, work organization that shifts more responsibility to workers, and worker skill are complements in production. We thank Nick Bloom for graciously sharing extracts of the HH data with us.

³³A measure of PCs per employee is desirable to better understand capital intensity (rather than simply growth in size), but as employment may be varying (endogenously) over this time period, we fix the normalization at a period before the Great Recession: the average of each available year among 2002, 2004, and 2006. This normalization means that variation in the outcome is strictly due to the numerator (total PCs), and ensures that greater employment losses in harder-hit MSAs will not mechanically induce a positive association between our PCs measure and the size of the shock. The fixed normalization requires that our sample be restricted to sites that are observed both prior and subsequent to the Great Recession, and this covers 65% of employment in HH across our sample years. Appendix A.8 shows that meeting this restriction is unrelated to our shock measure.

Figure 4: PC Adoption and the MSA-Employment Shock



Graph plots changes in IT investment (Harte Hanks survey) and 95% CI as a function of the MSA specific employment shock, controlling for year fixed effects and MSA characteristics. We aggregate site-level number of PCs per pre-recession employment to the MSA-year level taking an employment weighted average across sites. Dependent variable is the change in this measure from 2006.

across all MSAs (a 0.93 increase in PCs per employee off a base of 0.75).

By 2014 the point estimate has fallen somewhat and is no longer statistically significant, possibly reflecting the beginning of a more gradual catch-up of technology adoption in less hard-hit MSAs. However, the point estimate implies that harder-hit MSAs remain 1 PC per worker ahead of less hard-hit areas, relative to their pre-recession levels.³⁴

Furthermore, the estimated coefficients for 2000, 2002, and 2004 are all close to zero and statistically insignificant, implying that MSAs that would be severely affected had fairly similar IT investment trends before the Great Recession. If anything, there is a slight relative decline in per-worker PCs in these areas between 2000 and 2006, mostly in the last two years of this range, but these estimates are somewhat noisy. Thus, there is no evidence of a capital intensifying trend in harder-hit MSAs before the recession, as the modest relative movement goes in the opposite direction. As with the employment and unemployment rates in figure 1, it is comforting that the identifying assumptions of the Bartik shock appear to hold. Although we cannot observe skill requirements before 2007, that both employment-to-population ratios by education group *and* IT investments trend similarly across MSAs in the period before the Great Recession should reduce concern about preexisting trends.

³⁴These relatively large magnitudes are in part driven by long right tails in the distribution of PCs per worker across MSAs and years. To reduce the role of outliers, we have also estimated PC adoption on a sample trimmed of the top and bottom 1% of observations; we find qualitatively similar patterns of statistically significant increases in 2008–2012 that gradually decrease to insignificance by 2014, with point estimates that are somewhat smaller in magnitude.

In section 4, we showed not only that advertised job skill requirements increased and persisted in harder-hit MSAs, but that these increases occurred within firms. Since harder-hit MSAs also intensified their IT investments over the same time period, we next explore whether this investment and upskilling are linked at the firm level.

To do so, we link BG job ads at the firm level to two measures of investment from external data sources: PCs per worker from the HH database and capital holdings from Compustat North America by Standard & Poors (hereafter Compustat). Compustat is the most complete database of accounting and balance sheet data among publicly traded U.S. firms. Although PCs is a good proxy for overall IT investments, it may miss broader routine-labor replacing investments, such as new machinery, telecom infrastructure, or inventory management systems.³⁵ Thus, a firm’s overall holdings of property, plant, and equipment (hereafter, PPENT) from Compustat is a useful supplement.

We link both datasets to the BG data by firm name. See Appendix A.4 and A.5 for details on these mergers. In general, we can match more firms to HH than to Compustat, as the former is meant to cover all businesses while the latter is restricted to publicly traded companies. Among employers observed in both 2007 and the later period in BG (which cover 65% of postings), we are able to match about 80% of postings to firms in HH and 40% of postings to firms in Compustat. Appendix A.8 shows that the share of ads matching to these samples does not vary with the local employment shock.

We estimate upskilling regressions at the firm-MSA-year level, defined in 3. This equation allows for an additional interaction between the shock-by-year variables and the firm-level change in capital investments over the Great Recession ($Capital_f$). To reduce measurement error, we define these firm-level changes in capital investment as the difference between the average value in 2002, 2004, and 2006 and the average value in 2010, 2012, and 2014, for both PCs per worker (HH) and PPENT (Compustat).³⁶

(3)

$$outcome_{f_{mt}} - outcome_{f_{m07}} = \alpha_0 + [shock_m * I^t] \alpha_1 + [shock_m * I^t * Capital_f] \alpha_2 + I^t + X_m \beta + \varepsilon_{mt}$$

To make estimates easier to interpret, we normalize $Capital_f$ so that a one unit change equals the 90-10 percential differential in firm-level capital change. The 90th percentile

³⁵The Harte-Hanks database contains other measures of IT investment, including servers (for which we generally find results consistent with those from PCs), and specific types of software. Unfortunately, the latter are consistently available only from 2010 onward.

³⁶As before, for PCs, we normalize this difference by average employment in the pre-period. Note that even though observations are firm-MSA-year cells, the investment change is at the firm level, regardless of location. This is a necessary restriction of the Compustat data, which exists only for the firm and not individual establishments; for comparability, we aggregate sites in the HH data to the firm level, weighting by site employment. When we instead measure investment change at the firm-MSA level in the HH data, the results are qualitatively similar for most skill outcomes. However, both for comparability with the Compustat measure and to avoid additional noise from more demanding match criteria, we prefer defining investment change at the firm level. To limit the influence of extreme outliers in the PCs measure, we trim the top and bottom 2.5% of firms in the full HH database, which amounts to roughly 4% of weighted observations in our regressions.

firm in our sample added roughly two-thirds of a PC per worker (using fixed pre-recession employment levels) at each of its establishments and roughly tripled PPENT. In contrast, the 10th percentile firm lost nearly one-third of a PC per worker and dropped PPENT holdings by about 20%. Figure 5 plots the coefficients α_2 (see also table C2).

We find that firms with larger capital investments differentially and persistently increase their skill requirements. For example, the top left panel of figure 5 (sub-figure a) shows that between 2007 and 2010, a firm at the 90th percentile of increase in PCs per worker increased the likelihood of an education requirement in its job postings by 0.7 ppts more than a firm at only the 10th percentile of increase in PCs per worker, when both are exposed to a similar employment shock. This differential fluctuates somewhat, but persists and grows to about 1.0 ppt by 2015. This pattern and approximate relative magnitude hold for experience, cognitive skill, and computer skill requirements, with statistically significant differentials in most post-recession years, usually at the 1% level. Overall, we find that in harder-hit MSAs, skill requirements increase by roughly 30% to 50% more in high-investment firms than in low-investment firms, and these differentials hold over the post-recession period.

We find quantitatively similar effects when using the Compustat PPENT measure (sub-figure b). For example, in harder-hit MSAs high-investment firms increase the likelihood of specifying an education requirement by 0.6 to 0.8 ppts, relative to low-investment firms. This reflects roughly 35% greater responsiveness in upskilling. Results for the other skills variables are comparable, with most impacts significant 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. These patterns are consistent with both human and physical capital deepening at the firm level.

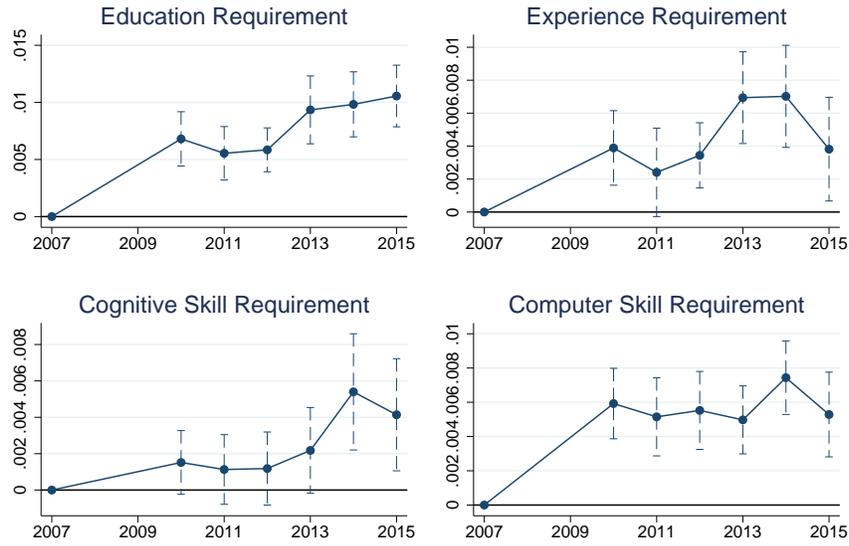
6 Routine Occupations

Thus far, we have provided evidence that MSAs more severely affected by the Great Recession experienced persistent increases in the skill demand of job postings as well as greater increases in capital. Moreover, the upskilling and capital investment occurred within the same firms. Both these findings are consistent with episodic restructuring. In this section we explore additional predictions of an RBTC-style restructuring: whether upskilling is more prevalent in more-routine occupations, and related trends in employment and wages for these occupations.

Indeed, the literature on job polarization has successfully linked employment and wage shifts across occupations to the tasks performed by workers in the occupations. Wages and employment have fallen for occupations in the middle of the skill distribution, which, being the most routine, are the kinds of occupations that can be replaced by machines or overseas

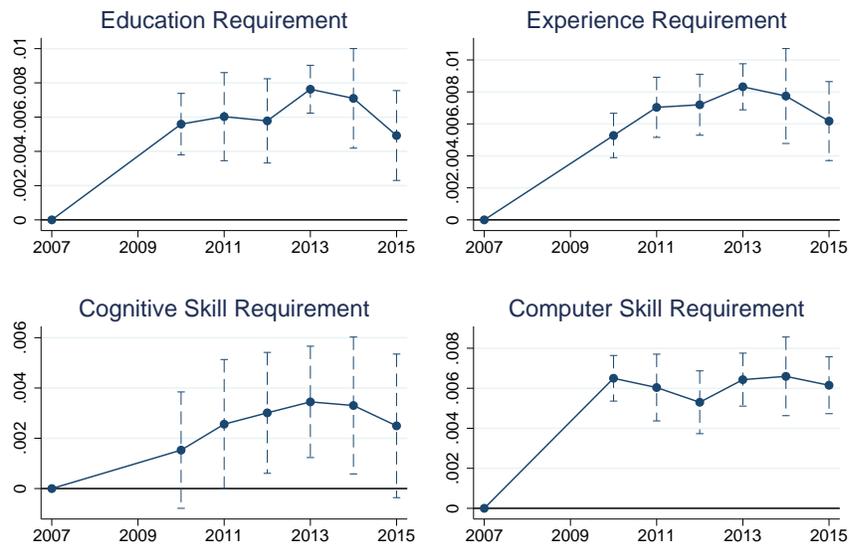
Figure 5: Differential Upskilling by 90-10 Change in Firm Capital Investments

(a) PCs (HH)



Graph plots the coefficients on the Bartik shock*year interacted with the firm-level capital change, fitted to the 90-10 differential and 95% CI bars. Regressions also control for year fixed effects and MSA characteristics.

(b) Capital Holdings (Compustat)



Graph plots the coefficients on the Bartik shock*year interacted with the firm-level capital change, fitted to the 90-10 differential and 95% CI bars. Regressions also control for year fixed effects and MSA characteristics.

labor.³⁷ Autor (2014) and Jaimovich and Siu (2015) have noted that employment continued to shift away from middle-skill occupations in the Great Recession. Our BG 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. Therefore, we next ask whether upskilling was relatively concentrated *within* routine occupations, the very jobs thought to be most affected by technological change in recent decades.

To determine an occupation’s routineness, we use Acemoglu and Autor’s (2011) routine-cognitive and routine-manual indices, derived from O*NET (see appendix A.7). These indices provide continuous scores based on the intensity of routine tasks performed, are simple to create, distinguish between tasks that use mental and physical capacities, and have been used in several other papers (e.g., Valletta 2015, Aaronson and Phelan 2017, Lewandowski 2017).³⁸ Routine cognitive tasks tend to be clustered in clerical, administrative and sales occupations, while routine manual tasks tend to be found in production and operative occupations. As a whole, employment in both types of occupations has been declining for at least the past two decades (Acemoglu and Autor 2011).

We begin by examining changes in skill requirements as a function of these routine index scores and the Bartik shock. To simplify our analysis, we focus on the top quartile of routine-cognitive and routine-manual occupations. The general pattern of results is similar when we allow for finer distinctions. We estimate regressions of the following form, where $Routine_o^i$ is an indicator equal to 1 if occupation, o , is in the top quartile of categorization, i , where $i \in \{cognitive, manual\}$. The parameter vector α_2 in equation (4) thus captures the additional effect of the shock, each year, for top-quartile routine occupations relative to the effect in occupations in the bottom three quartiles of the relevant routineness index. Henceforth, for exposition, we will refer to these top-quartile occupations as routine-cognitive

³⁷The 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).

³⁸The AA measures use O*Net to essentially update the original categorization of Autor, Levy, and Murnane (ALM; 2003), which used the predecessor, the Dictionary of Occupational Titles. Some papers in the literature (e.g., Autor and Dorn 2013; Autor, Dorn, and Hanson 2015b) use a simpler routine-manual-abstract categorization, still based on the original ALM categories, that does not allow for a distinction between routine-manual and routine-cognitive occupations; we find this distinction to be important. Jaimovich and Siu (2015) use broad occupation categories to generate their binary routine classification. For our purposes, the AA measures are preferable because they allow for finer (continuous) distinctions. The Spearman correlation between our adapted AA measures and the Jaimovich and Siu measures is 0.66 for routine manual but only 0.21 for routine cognitive, indicating that the AA measures likely avoid some miscategorization inherent in the binary definition.

or routine-manual, as appropriate.

(4)

$$outcome_{omt} - outcome_{om07} = \alpha_0 + [shock_m * I^t] \alpha_1 + [shock_m * I^t * Routine_o^i] \alpha_2 + I^t + X_m \beta + \varepsilon_{mt}$$

Figure 6 (and table C3) plots estimates of α_2 for each routineness index for the four skill requirements. The coefficients for routine-cognitive occupations are indicated with blue circles, while coefficients for routine-manual occupations (estimated with separate regressions) are shown as maroon squares.

The figure's primary pattern is a greater degree of upskilling in routine-cognitive occupations: these estimates are positive, statistically significant, and persistent in all cases (except for experience in 2015). For example, the blue circle in the top left panel in 2010 indicates that in hard-hit MSAs, job posts for routine-cognitive occupations saw about a 0.5 ppt larger increase in the probability of having an education requirement, relative to other occupations. For 2013 through 2015, the differential impact is roughly twice as large. As a whole, these differentials are approximately 25 (education/experience) to 50% (cognitive/computer) as large as the baseline within-occupation upskilling effects (table 2).

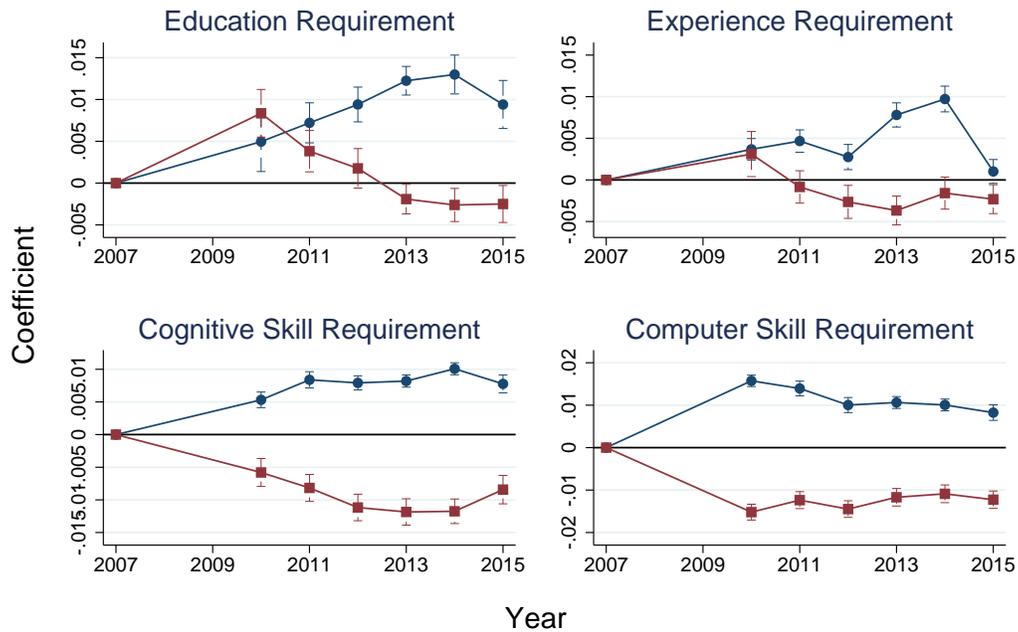
In contrast, routine-manual occupations do not exhibit a persistent differential in upskilling. In fact, in the case of cognitive and computer skills, these occupations experience relative downskilling compared to occupations that are not routine-manual. (That is, routine-manual occupations upskill less than other occupations.) For education and experience requirements, routine-manual occupations do exhibit temporary differential upskilling, which could reflect opportunistic behavior on the part of firms during a slack market. However, this effect fades quickly.

Upskilling thus appears to be relatively concentrated within routine-cognitive jobs. Our hypothesized explanation for this pattern is that the recession accelerated technological adoption, but that some types of jobs—routine-cognitive ones—could be made more complementary to the new technology with additional human capital, while labor for other types of jobs—routine-manual ones—was more subject to substitution by the new technology. For example, the use of data analytics may make a salesperson more productive by allowing her to better target customers' needs, but software alone will not close a sale—a salesperson capable of using the software is still needed to do the job. On the other hand, machine-vision technology may render obsolete the manual inspection of parts on an assembly line, essentially replacing that job.³⁹

To investigate this paradigm, we turn to the implications for employment and wages. If firms do not seek greater skills for routine-manual jobs because those jobs can be substituted with technology more readily than work with it, we would expect firms to disproportionately shed these types of jobs through layoffs, with little employment recovery over time. Since

³⁹Indeed, Hawkins, Michaels, and Oh (2015) present recent evidence of this type of capital-labor substitution in the Korean manufacturing sector.

Figure 6: Differential Upskilling for Routine Occupations



Blue circles = routine cognitive, maroon squares = routine manual

Graph plots the coefficients on Bartik shock*year interacted with an indicator for the top quartile of routine cognitive (blue circles) or routine manual (maroon squares), as well as 95% confidence interval bars. Regressions control for year fixed effects and MSA characteristics.

firms do seek greater skill for routine-cognitive jobs, these workers may complement new technology, and if so, their productivity and relative wages should rise. (Predictions for relative wages of routine-manual jobs and relative employment (or job loss) for routine-cognitive jobs are less clear cut, and depend on product demand.)

We explore all three margins—involuntary separations, relative employment, and wages—for routine-cognitive and routine-manual occupations. We examine the rate of involuntary job loss in the population (not just the unemployed), as measured in the CPS by the propensity to report being a job loser. For employment and wages, we use Occupational Employment Statistics (OES) data. Both data sets allow us to capture trends back to 2000, and the earlier years help to check the validity of our identifying assumption, that harder-hit MSAs would have been on a similar trend—in terms of skill demand—if not for the Great Recession. (See appendix A.6 for details on sample construction and an analysis of the overall impact of the Bartik shock on these outcomes.)

In Figure 7, we focus on the differential impacts of the Bartik shock on layoffs, employment, and wages for routine-manual (red squares) and routine-cognitive (blue circles) occupations.⁴⁰

⁴⁰OES data are based on a three-year moving average, so annual snapshots are not independent and trends are likely smoother than true annual snapshots would be. We have confirmed that this trait does not substantively affect our estimates, as our results do not change appreciably if we use data for every third year (which is independent) instead of annually. Additionally, we get similar results for employment and wages using the CPS or ACS, albeit with fewer included MSAs (CPS) or data years (Census/ACS). We prefer the use of OES for these outcomes for its fuller coverage across occupations, MSAs, and time.

Beginning with the top left panel, we find evidence of a large differential layoff effect for routine-manual occupations. At the peak in 2009, individuals whose current or most recent job was in a routine-manual occupation suffer an additional 1.5 ppt increase in involuntary separations, relative to those not in routine-manual occupations, due to a 90–10 percentile MSA shock. For comparison, the same-sized shock increased the probability of having been laid off for those not in routine-manual occupations by 0.8 percentage points in 2009 (see column 1 of table C4 for the full regression output); that is, individuals in routine-manual occupations experienced nearly triple the chance of involuntary separation as did individuals in other occupations. Although routine-cognitive occupations also experience a statistically significant differential increase, the magnitude is modest.⁴¹ Importantly, there appears to be little pre-trend for either routine occupation type, as, except for a tiny blip in 2003, the layoff rate differential is close to zero in all years from 2000 to 2007.

The top right panel (and columns 10 and 11 of table C1) shows how the share of MSA employment in each type of routine occupation varies with the Bartik shock. Following the Great Recession, there is a large and persistent drop in routine-manual employment and a steady and modest rise in routine-cognitive employment.⁴² At their trough, the employment share of routine-manual occupations fell by about 2 percentage points more in harder-hit MSAs, recovering only half of this gap by the end of the sample period. In contrast, routine-cognitive occupations differentially rise as a share of employment in harder-hit MSAs, (or, more aptly, experience smaller magnitude losses in employment share, relative to less hard-hit MSAs) though only modestly. Also, unlike for routine-manual occupations, it is harder to rule out a pre-trend, although it is small and generally statistically insignificant. This pair of results is generally consistent with Jaimovich and Siu (2015), who show that employment in routine occupations as a whole fell episodically, and more so in harder-hit U.S. states, in each of the past three recessions and did not recover fully.

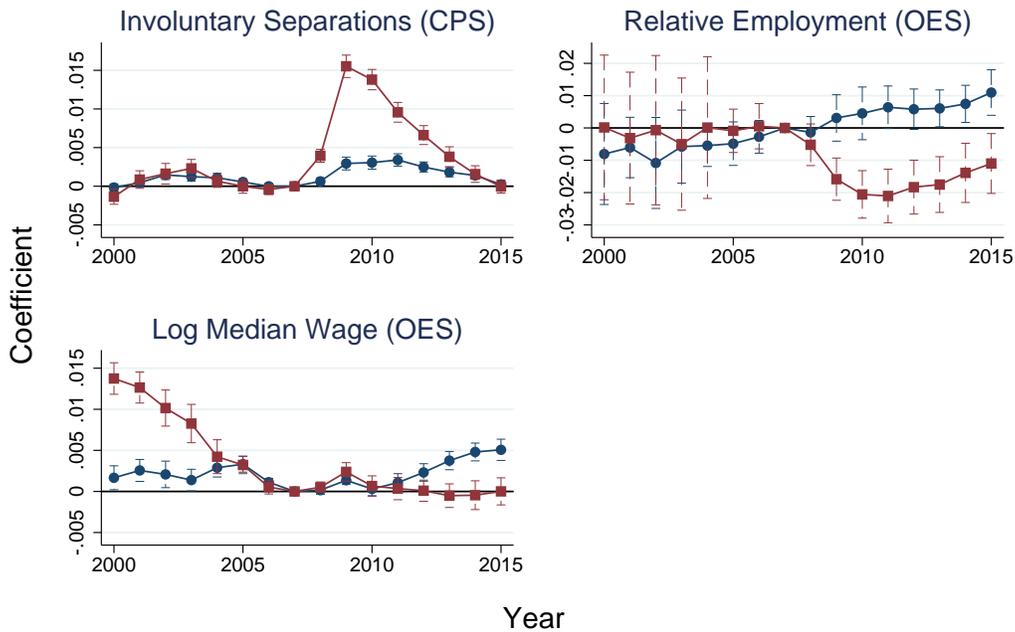
The bottom left panel (and column 2 of table C4) shows the differential impact on log median hourly wages. For routine-cognitive occupations, there is a slight but persistent rise in wages in harder-hit MSAs beginning after 2010. By 2015, the median routine-cognitive worker in a hard-hit MSA has experienced 0.5% faster wage growth than his peer worker in a less-hard-hit MSA. Conversely, routine manual occupations exhibit almost no post-recession change; wages evolve similarly in the subsequent period regardless of the MSA shock. In the single major exception to the absence of pre-trending, wages for routine-manual jobs were differentially trending downward before the Great Recession in areas that would experience a more severe shock, even though relative employment trended similarly, suggesting other factors (such as declining unionization) were possibly involved.

The sharpest predictions of our hypothesis—episodic increases in layoffs and persistent

⁴¹It is possible, but not central to our argument, that workers laid off from a routine-cognitive job have an easier time finding reemployment and thus do not report their current status as laid off.

⁴²For this result, we estimate versions of equation 1, using as dependent variables the share of employment in the MSA (relative to 2007) that is routine-cognitive or routine-manual.

Figure 7: Differential Employment and Wage Effects for Routine Occupations



Blue circles = routine cognitive, maroon squares = routine manual

Graph plots the coefficients on Bartik shock*year interacted with an indicator for the top quartile of routine cognitive (blue circles) or routine manual (maroon squares), as well as 95% confidence interval bars. Regressions control for year fixed effects and MSA characteristics.

decreases in employment share for routine-manual workers and increases in wages for routine-cognitive workers—are borne out by the data, and for these outcomes there is little evidence of differential pre-trending. To summarize, in harder-hit MSAs routine-manual occupations experience a sharp increase in layoff risk but there is no evidence of upskilling; rather, there appears to be downskilling accompanied by employment losses and flat wages compared with the same occupations in less-affected MSAs. For these occupations, the story is therefore consistent with firms’ substitution of technology for labor. This is the traditional view exhibited also in the polarization literature: employment losses concentrated in occupations we expect to be most readily replaceable by machines. Our contribution is to show that these changes appear episodic around the Great Recession, though we acknowledge that secular relative wage losses preceded the employment shock.

In contrast to this conventional view of labor substitution, routine-cognitive occupations in hard-hit MSAs surprisingly exhibit only a modest increase in layoffs and no loss in employment share, relative to less hard-hit MSAs. These changes were concomitant with a pronounced increase in upskilling. Even as the elevated differential risk of layoff declined, the differential in upskilling persisted, and was met with modest relative wage and employment growth after the recession. This pattern is quite consistent with an intensive margin restructuring due to technological adoption that effectively shifts out the labor demand curve.

These results are especially enlightening, given recent findings by Beaudry, Green, and Sand (2014, 2016) that more-educated workers have increased their presence 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.⁴³ 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 away from routine occupations, the tasks performed in the routine occupations that remain may be becoming less routine and more cognitive. Our work thus highlights a complementary hypothesis for why high-skilled workers are increasingly found in lower-skilled occupations: these latter occupations are becoming more skilled (and more highly paid), and it is possible that less-skilled workers are displaced because they are unable to perform the new duties required.

7 Conclusion

During the recovery following the Great Recession, anecdotal evidence suggested that the composition of new hires shifted toward higher-skilled workers, resulting in many workers being “overeducated” for their jobs (Burning Glass Technologies 2014). However, it was not clear how broad, deep, or enduring these effects were, or the extent to which they were driven by labor supply or labor demand responses. In particular, firms may have treated the recession as a time of “cleansing,” enabling them to restructure their production in a manner consistent with routine-biased technological change.

In this paper we draw upon detailed job postings data to provide comprehensive, broad-based evidence of upskilling—firms demanding higher-skilled workers—when the local economy suffers a recession. Using specific empirical skill measures that reflect what the discipline has learned about technological change and task-based production over the past 20 years (Autor, Levy, and Murnane 2003; Brynjolfsson and McAfee 2011), we show that job postings in harder-hit MSAs experienced larger increases in their education, experience, cognitive, and computer requirements following the Great Recession. These increases primarily reflected changes in demand *within*, and not across, occupations. Furthermore, skill requirements remained elevated through the end of our sample in 2015, even as most measures of labor market conditions had converged back to their pre-recession levels. Importantly, we find that the increases in skill requirements are accompanied by increases in capital investments, both at the MSA and firm-levels. We also show that upskilling is relatively concentrated in routine-cognitive occupations, which exhibit modest wage growth as well. In contrast, routine-manual occupations in harder-hit MSAs exhibit a sharp relative decline in employment shares following the Great Recession.

⁴³See also Castex and Dechter (2014) and Deming (2017).

We argue that the most likely explanation for this body of results is that the Great Recession did indeed provide firms a catalyst to restructure production according to a paradigm of routine-biased technological change. While firms may respond to changes in labor market conditions through posted skill requirements for a variety of other reasons, these cannot rationalize our full body of findings. For example, firms may worry that a flood of applicants early in the recovery will create a “bottleneck” in screening, and therefore raise requirements to signal that certain (unwanted) applicants need not apply. Alternatively, firms that typically cannot attract (or afford) more-skilled workers in a tight labor market may opportunistically seek them out in a slack one. However, while both of these cyclical behaviors may have been important early in the recovery, they cannot generate persistent, within-firm increases in skill requirements that stand up to controls for the availability of labor by skill group, occur concomitantly with greater investment in multiple measures of physical capital, and are concentrated in the types of occupations acknowledged to be most susceptible to routine-biased technological change. While it is possible that labor markets have been slower to recover than indicators such as employment growth, the unemployment rate, and education-specific employment-population ratios indicate, it is telling that we find little overall convergence in skill requirements as markets improve, even though the data *do* indicate such convergence in some (e.g., routine-manual) occupations.

Simply put, the evidence supports that shifts in skill requirements reflect technologically-driven changes in the means of production, not just changes in whom firms seek to hire. Our work is thus 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. For example, our result that routine-cognitive occupations are apparently becoming higher skilled and more productive can help to clarify studies by Beaudry, Green, and Sand (2014, 2016) and others documenting the “great reversal” in demand for cognitive skill. While it may be the case that employment in high-skill occupations did not grow, on average, over the past 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. This is indicated by the persistence in both the relative upskilling and wage growth in routine-cognitive occupations located in harder-hit MSAs. Our findings can thus help 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 past 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 et al. 2014; Autor, Dorn, and Hanson 2015a; Foote and Ryan 2015; Pierce and Schott 2015). Our results highlight that a worker’s 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. This is perhaps evident in the stair-step declines in male labor force participation that have tended to be concentrated around recessions (Moffitt 2012, Foote and Ryan 2015). If the changes to production instead occurred more gradually, workers would still need to be retrained, but over a longer time period, 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] Aaronson, Daniel, and Brian J. Phelan (Forthcoming), “Wage Shocks and the Technological Substitution of Low-Wage Jobs,” *Economic Journal*.
- [2] Acemoglu, Daron, and David Autor (2011), “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in David Card and Orley Ashenfelter (Eds.), *Handbook of Labor Economics: Volume 4B*. Amsterdam: Elsevier, pp. 1043–1171.
- [3] 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.
- [4] Autor, David (2014), “Polanyi’s Paradox and the Shape of Employment Growth,” NBER working paper #20485.
- [5] 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.
- [6] 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.

- [7] Autor, David, David Dorn, and Gordon Hanson (2015a), “The Labor Market and the Marriage Market: How Adverse Employment Shocks Affect Marriage, Fertility, and Children’s Living Circumstances,” MIT mimeo.
- [8] Autor, David, David Dorn, and Gordon Hanson (2015b), “Untangling Trade and Technology: Evidence from Local Labor Markets,” *Economic Journal*, 125(584): 621–646.
- [9] Autor, David, David Dorn, Gordon Hanson, and Jae Song (2014), “Trade Adjustment: Worker-Level Evidence,” *Quarterly Journal of Economics*, 129(4): 1799–1860.
- [10] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006), “The Polarization of the U.S. Labor Market,” NBER Working Paper #11986.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] Bartik, Timothy J. (1991), *Who Benefits from State and Local Economic Development Policies*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- [15] 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.
- [16] 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.
- [17] Beaudry, Paul, and Ethan Lewis (2014), “Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980–2000,” *American Economic Journal: Applied Economics*, 6(2): 178–194.
- [18] Berger, David (2012), “Countercyclical Restructuring and Jobless Recoveries,” mimeo.
- [19] Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott (2011), “Multiproduct firms and trade liberalization,” *Quarterly Journal of Economics* 126(3): 1271–1318.

- [20] Blanchard, Olivier J., and Lawrence F. Katz (1992), “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1–75.
- [21] Blinder, Alan S., and Alan B. Krueger (2013), “Alternative Measures of Offshorability, A Survey Approach,” *Journal of Labor Economics*, 31(no. 2 part 2): S97–S128.
- [22] Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012), “Americans Do IT Better: U.S. Multinationals and the Productivity Miracle,” *American Economic Review* 102(1): 167–201.
- [23] Bloom, Nicholas, Mirko Draca, and John Van Reenen (2016), “Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT, and Productivity,” *Review of Economic Studies*, 83(1): 87–117.
- [24] Brynjolfsson, Erik, and Andrew McAfee (2011), *Race Against the Machine*. Lexington: Digital Frontier Press.
- [25] Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt (2002), “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *Quarterly Journal of Economics* 117(1): 339–376.
- [26] Burning Glass Technologies (2014), “Moving the Goalposts: How Demand for a Bachelor’s Degree is Reshaping the Workforce,” mimeo.
- [27] Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf (2016), “Does the United States Have a Productivity Slowdown or a Measurement Problem?” *Brookings Papers on Economic Activity*, Spring.
- [28] Caballero, Ricardo, and Mohamad Hammour (1994), “The Cleansing Effect of Recessions,” *American Economic Review*, 84(5): 1350–1368.
- [29] Caballero, Ricardo, and Mohamad Hammour (1996), “On the Timing and Efficiency of Creative Destruction,” *Quarterly Journal of Economics*, 111(3): 805–852.
- [30] Card, David, Ana R. Cardoso, and Patrick Kline (2016), “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*, 131(5): 633–686.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] Charles, Kerwin K., Erik Hurst, and Matthew J. Notowidigdo (2015), “Housing Booms and Busts, Labor Market Opportunities, and College Attendance,” mimeo.
- [36] Clemens, Jeffrey (2015), “The Minimum Wage and the Great Recession: Evidence from the Current Population Survey,” NBER working paper #21830.
- [37] 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.
- [38] 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.
- [39] Davis, Stephen J., Jason Faberman, and John Haltiwanger (2013), “The Establishment-Level Behavior of Vacancies and Hiring,” *Quarterly Journal of Economics*, 128(2): 581–622.
- [40] Deming, David J. (2017), “The Growing Importance of Social Skills in the Labor Market,” NBER Working Paper #21473.
- [41] Deming, David J., and Lisa B. Kahn (2017), “Firm Heterogeneity in Skill Demands,” *Journal of Labor Economics*, forthcoming.
- [42] Foote, Christopher L., and Richard W. Ryan (2015), “Labor Market Polarization Over the Business Cycle,” *NBER Macroeconomics Annual*, 29: 371–413.
- [43] “Forsythe, Eliza (2016), “Why Don’t Firms Hire Young Workers in Recessions?” mimeo.
- [44] 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.
- [45] Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (2001), “Equilibrium Unemployment,” *Journal of Monetary Economics*, 48(1): 109–152.

- [46] Greenfield, Rebecca (2014), “Why the New York Times Hired a Biology Researcher as Its Chief Data Officer,” *Fast Company*, February 12.
- [47] Hall, Robert (1991), “Labor Demand, Labor Supply and Employment Volatility,” in NBER Macroeconomics Annual. Cambridge: MIT Press.
- [48] Hall, Robert (2005), “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1): 50–65.
- [49] Hawkins, William B., Ryan Michaels, and Jiyoung Oh (2015), “The Joint Dynamics of Capital and Employment at the Plant Level,” mimeo.
- [50] Hoynes, Hilary, Douglas L. Miller and Jessamyn Schaller (2012), “Who Suffers During Recessions?,” *Journal of Economic Perspectives*, 26(3): 27–48.
- [51] 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.
- [52] 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>.
- [53] 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>.
- [54] Jensen, J. Bradford, and Lori G. Kletzer (2005), “Tradable services: Understanding the Scope and Impact of Services Offshoring.” In J. Bradford Jensen, Lori G. Kletzer, Jared Bernstein, and Robert C. Feenstra (Eds.) *Brookings Trade Forum*. Washington, D.C.: Brookings Institution Press, pp. 75–133.
- [55] Kahn, Lisa B. (2010), “The Long-Term Consequences of Graduating from College in a Recession,” *Labour Economics*, 17(2): 303–316.
- [56] Kahn, Lisa B., and Erika McEntarfer (2015), “Employment Cyclicalities and Firm Quality,” NBER working paper #20698.
- [57] Keister, Roma, and Piotr Lewandowski, (2017), “A Routine Transition in the Digital Era? The Rise of Routine Work in Central and Eastern Europe,” *Transfer: European Review of Labour and Research*, 23(3): 263–279.
- [58] 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.

- [59] Kuhn, Peter, and Kailing Shen (2013), “Gender Discrimination in Job Ads: Evidence from China,” *Quarterly Journal of Economics*, 128(1): 287–336.
- [60] Lazear, Edward P., and James R. Spletzer (2012), “Hiring, Churn, and the Business Cycle,” *American Economic Review*, 102(3): 575–579.
- [61] Marinescu, Ioana (2014), “The General Equilibrium Effects of Unemployment Insurance: Evidence from a Large Online Job Board,” mimeo.
- [62] 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.
- [63] McCoy, Terrence (2016), “The Technology that Could Revolutionize the War on Hunger,” *Washington Post*, June 16.
- [64] Menzio, Guido, and Shouyong Shi (2011), “Efficient Search on the Job and the Business Cycle,” *Journal of Political Economy* 119(3): 468–510.
- [65] 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.
- [66] Moffitt, Robert (2012), “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations,” *Brooking Papers on Economic Activity*, fall.
- [67] Mortensen, Dale, and Christopher Pissarides (1994), “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61(3): 397–415.
- [68] 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
- [69] Nickell, Steohen J. (1996), “Competition and Corporate Performance,” *Journal of Political Economy*, 104(4): 724–746.
- [70] Notowidigdo, Matthew J. (2013), “The Incidence of Local Labor Demand Shocks,” University of Chicago, mimeo.
- [71] Pierce, Justin R., and Peter K. Schott (2015), “Trade Liberalization and Mortality: Evidence from U.S. Counties,” Yale mimeo.
- [72] Rothwell, Jonathan (2014), “Still Searching: Job Vacancies and STEM Skills,” Brookings mimeo.
- [73] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2016a), “Upskilling: Do Employers Demand Greater Skill when Skilled Workers are Plentiful?”, mimeo.

- [74] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2016b), “Downskilling: Changes in Employer Skill Requirements Over the Business Cycle”, mimeo.
- [75] Shimer, Robert (2004), “The Consequences of Rigid Wages in Search Models,” *Journal of the European Economic Association*, 2(2-3): 469–479.
- [76] Schumpeter, Joseph (1939), *Business Cycles*. New York: McGraw-Hill.
- [77] Syverson, Chad (2004a), “Market Structure and Productivity: A Concrete Example,” *Journal of Political Economy*, 112(6): 1181–1222.
- [78] Syverson, Chad (2004b), “Product Substitutability and Productivity Dispersion,” *Review of Economics and Statistics*, 86(2): 534–550.
- [79] U.S. Department of Labor (1977), *Dictionary of Occupational Titles, 4th edition*. Washington, D.C.: U.S. Department of Labor.
- [80] 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.
- [81] Yagan, Danny (2016), “The Enduring Employment Impact of Your Great Recession Location,” mimeo.

FOR ONLINE PUBLICATION

A Data Appendix

It is estimated that as of 2014 between 60% and 70% of all job postings could be found online (Carnevale, Jayasundera, and Repnikov 2014). Indeed, The Conference Board discontinued its-long-running, print-based Help-Wanted Advertising Index in 2008, after having begun a Help-Wanted Online Index in 2005 (HWOL).⁴⁴ 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 Industry-Occupation 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.⁴⁵ 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.⁴⁶

Figure A1 plots the distribution of BG ads across major industry groups, sorted from largest to smallest (solid bars), as well as the distribution of job vacancies in JOLTS (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.⁴⁷ Despite the sample differences, the industry distributions

⁴⁴See <https://www.conference-board.org/data/helpwantedonline.cfm>.

⁴⁵Rothwell (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.

⁴⁶Furthermore, 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.

⁴⁷Both data sets cover 2007 and 2010–2015. The BG distribution is from our primary estimation sample (notably excluding ads with missing firms), though we obtain similar results for the distribution across all ads. JOLTS data are based on a monthly, nationally representative sample of approximately 16,000 business

match each other reasonably well. BG is overrepresented in health care and social assistance, as well as in finance and insurance and education. It is underrepresented in accommodation and food services, public administration/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 United States. 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 since 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 (blue bars).⁴⁸ 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).⁴⁹

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

establishments drawn from unemployment insurance records; they count as a vacancy or job opening any position (including temporary and seasonal ones) that 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.

⁴⁸For clarity, we use 2-digit Standard Occupational Classification codes in the figure. The regression analyses use more granular codings.

⁴⁹All data sets cover 2007 and 2010–2015. The BG distribution is from our primary estimation sample, though, again, the distribution is similar for the full sample of ads. 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, reported a new employer, or reported changing occupations.

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 Skills Measures in BG

One of the most unique features of the BG data is the availability of skills measures. We argue that these stated preferences are informative about labor demand. Figures A4a and A4b crosscheck the average education requirements in BG with average education levels of employed workers at the MSA- and occupation-level, respectively. Using American Community Survey (ACS) data for overlapping sampling years, we rank both MSAs and occupations (four-digit SOC codes) by their average education of employed workers and plot the relationships between average education requirements and average education for 20 evenly sized employment bins, using smoothed local linear regression. As can be seen, at the levels of both MSA and occupation, the probability that an ad posts any education requirement is increasing with the average years schooling of employed workers (top left), as is the years of school conditional on any requirement (top right). Furthermore, the probability that an ad has a high school requirement is positively correlated with the share of workers that have exactly a high school diploma (bottom left), and the probability that an ad has a college requirement is positively correlated with the share of workers with exactly a bachelor’s degree (bottom right).

A.4 Harte-Hanks Sample

We are grateful to Nick Bloom for providing us with an extract of the Harte-Hanks (HH) database. To construct our merged BG-HH sample, we begin with @ BG firms that post

in the 2010–2015 period and can also be matched to at least one ad in 2007 (@% of firms and @% of ads). In this sample, we match a total of @% of BG ads (@% of firms) to HH firms; we did not attempt to match observations that did not meet the pre-post criterion.⁵⁰ We apply a multi-step approach. First, we match based on exact name, after regularizing firm names in both BG and HH (removing “inc.,” “LLC,” and other common suffixes, as well as punctuation and spaces). This accounts for @% of ultimately matched firms and @% of ultimately matched ads. Most of the remaining share of ultimate matches (@% of firms, @% of ads) are obtained by dropping one at a time common strings, such as “hotel,” “group,” or “insurance,” that might be part of the firm name in one dataset but not the other. We also obtain a small number of matches by removing any “s” at the end of the last word, either singularizing or removing possessives, and by replacing “univ” with “university”; these steps combined account for @% of eventual firm and @% of eventual ads matched. Finally, we match based on the first 10 characters of the firm name in each dataset, cleaning spurious matches by hand. This accounts for the remaining @% of total firm and @% of total ad matches.

A.5 Compustat Sample

We obtain Compustat data via Wharton Research Data Services. To construct our merged BG-Compustat sample, we again begin with @ BG firms that post in the 2010–2015 period and can also be matched to at least one ad in 2007 (@% of firms and @% of ads). In this sample, we match a total of @% of BG ads (@% of firms) to Compustat firms; we did not attempt to match observations that did not meet the pre-post criterion.⁵¹ We apply a multi-step approach. We first match based on exact name, after the same cleaning procedure described above (removing punctuation, spaces, and words that are sometimes abbreviated). This step accounts for @% of ultimately matched firms and 76% of matched ads. Second, we use a fuzzy-match program called Winpure and assign a match if the program determines at least a 95% probability of a match; this accounts for 5% of eventually matched firms and ads).⁵² Third, we add the sample of firms matched by Deming and Kahn (2017), which uses only BG firms posting in 2014 (@% of firms and @% of ads).

⁵⁰The HH database sampled roughly 500,000 U.S. sites each year prior to 2010, and roughly three million from 2010 onward. Although detailed sampling information is not available, and the HH database does not contain sampling weights, total employment among the sampled sites ranges from 39–57 million prior to the Great Recession and 101–114 million from 2010 onward. This represents about 30–40 percent of total payroll employment prior to 2010, and about 70–85 percent from 2010 onward.

⁵¹For 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.

⁵²We tried this procedure for the HH match, but with much less success, possibly owing to the greater variation in firm names in the larger HH database.

A.6 Current Population Survey and Occupation Employment Statistics Samples

Figure 7 in the main text explores the differential impact of the Bartik shock for routine-manual and routine-cognitive occupations, using Current Population Survey (CPS) micro-data to calculate layoffs and Occupation Employment Statistics (OES) data to calculate wages and employment.

We work with the basic monthly CPS and harmonize both MSA and occupation codes. To harmonize MSA codes, we begin with the 2013 Office of Management and Budget (OMB) CBSA delineations (which the CPS began using in May 2014) and work backward in time. The CBSA codes from May 2004 though April 2014 can be converted to the 2013 standard using Census delineation files; in most cases these are one-to-one code changes or absorptions of an MSA into a larger one, although some MSAs stop being identifiable in the CPS and others start being identifiable.⁵³ Prior to May 2004, the CPS used the older PMSA coding scheme. We apply a crosswalk initially developed by the Social Security administration and maintained by the NBER to convert PMSAs to 2013-vintage CBSAs⁵⁴. The more significant nature of the PMSA to CBSA switch results in several MSAs that do not cleanly map; the crosswalk maps 267 of 317 PMSAs into CBSAs at the exact county-level definition, although not all of these CBSAs are identifiable in the CPS from 2005 onward. Using these procedures, we consistently identify 149 MSAs between 2000 and 2015 in the CPS.

To harmonize occupation codes, we begin with the 2010 SOC coding scheme and work backward. The CPS uses 2010 Census occupation codes starting January 2011; these are converted to 2010 SOC codes using a Census crosswalk.⁵⁵ From January 2003 through December 2010, the CPS uses 2000 Census occupation codes. These occupation codes are first converted to the 2000 SOC system using a crosswalk maintained by IPUMS at the University of Minnesota.⁵⁶ The 2000 SOC codes are then converted to the 2010 SOC system using a BLS crosswalk (see https://www.bls.gov/soc/soc_2000_to_2010_crosswalk.xls) for occupations that match one-to-one and were simply recoded or combined; for occupations that split at the 6-digit level, we applied a stochastic crosswalk based on empirical shares observed in the IPUMS versions of the 2009 ACS (which contains the 2000 SOC) and the 2010–2012 ACS (which contains the 2010 SOC).⁵⁷ Finally, the CPS uses 1990 occupation codes prior to January 2003. To convert CPS years 2000 through 2002, we use the NBER version of CPS extracts files released by BLS.⁵⁸ These files contain 2000 Census occupation

⁵³See <https://www.census.gov/geographies/reference-files/time-series/demo/metro-micro/delineation-files.html>.

⁵⁴See <http://www.nber.org/data/cbsa-msa-fips-ssa-county-crosswalk.html>.

⁵⁵See https://www.census.gov/people/io/files/2010_OccCodeswithCrosswalkfrom2002-2011nov04.xls.

⁵⁶See https://usa.ipums.org/usa/volii/census_occtooccoc.shtml.

⁵⁷Occupations that split were randomly assigned to one of the splits based on the empirical distribution of the splits in 2010–2012. Since these splits almost never cross the 4-digit SOC level, which we use in all analyses, measurement error from this stochastic assignment is trivial.

⁵⁸See http://www.nber.org/data/cps_extract.html.

codes for CPS years 2000 through 2002, allowing use of the above procedure.

To determine layoff status, we first drop observations living in a non-metropolitan area or without a valid occupation code, as we cannot match an MSA shock or occupational routine-ness measure to these out-of-universe observations. We assign the status of an involuntary separation to individuals who answer the reason for unemployment question (pruntype) either “job loser/on layoff” or “other job loser.” Although this question is asked only of the currently unemployed, we include in the universe all members of the experienced labor force (those with a valid occupation code). Note that this is a “stock” variable; we do not restrict layoffs to individuals who had been employed the previous month.

We also use CPS microdata to calculate employment-to-population ratios (bottom panel of figure 1) and quit rates by education group (used as controls in appendix B.2). Note that these variables will be noisier than the unemployment rates and employment series calculated by the BLS (and used in the top panel of figure 1).⁵⁹ We calculate employment-population ratios according to the standard definition (the sample-weighted share of employed persons divided by the population). To measure quit rates, we longitudinally link observations and define a quit as an individual who was either (a) employed in months t and $t+1$ but reported having changing employers in month $t+1$ (using the variable `puiodp1`), or (b) employed in month t , unemployed in month $t+1$, and gave as the reason for unemployment being a “job leaver” (using the `pruntype` variable). The denominator is the (weighted) count of longitudinally matched individuals who were employed in month t .

We use OES data to calculate the wages and employment shares used in figure 7.⁶⁰ To make data comparable across years (from 2000 to 2015), we must use crosswalks for both MSA and occupation codes.

Data from 2000 to 2011 use SOC 2000 codes, which we map to SOC 2010 codes using the same procedure as with the CPS data. OES also uses a small number of temporary occupation codes in 2010 and 2011, which we must drop.

MSAs from 2005 onward use the OMB 2013 delineation, but years 2000–2004 use the 1999 delineation. We map old MSA codes to the new ones using data from the IPUMS versions of the 2000 Census and 2005–2011 ACS samples, where both measures are available, and keeping the modal new MSA match (based on sample weights)⁶¹.

To better understand how layoffs (CPS) and wages (OMB) vary with our key right-hand-side variables, we provide within-occupation-MSA estimates of equation 1. These are summarized in figure A5. The differential (rather than main) effects on these variables for routine-cognitive and routine-manual occupations (described in equation 4) are plotted in

⁵⁹BLS does not use microdata directly to calculate unemployment rates or payroll employment for local areas. Rather, these estimates are derived from time-series errors-in-variables models, using as primary inputs unemployment insurance claims and state-level CPS and payroll employment survey estimates. See <https://www.bls.gov/lau/gen4models.pdf>. Hence, demographic breakdowns are unavailable.

⁶⁰Annual data at the MSA-occupation level can be obtained from the BLS: <https://www.bls.gov/oes/ttables.htm>.

⁶¹See https://usa.ipums.org/usa-action/variables/MET2013#comparability_section.

figure 7.

We estimate that between 2007 and 2009, a hard-hit MSA experiences an additional 1.2 ppt increase in layoff probability (the share of the sample that reports being involuntarily unemployed), relative to a less hard-hit MSA. This is roughly two-thirds of the average involuntary separation rate observed across the sample period. Effects remain elevated in 2010, before dropping to zero by 2012. Wage effects, in contrast, move more slowly. Although issues of selection preclude us from making causal statements, we find that hard-hit MSAs experience a drop in wages about 0.5% to 1.5% greater than less hard-hit MSAs, with this differential the largest over 2012–2014. As with the other labor market measures in figures 1, 4, and 7, estimates preceding the Great Recession are close to zero. Thus, layoffs and wages were evolving similarly across MSAs, regardless of the size of the shock they would receive.

A.7 Autor-Acemoglu routineness measures

Acemoglu and Autor (AA; 2011) use O*Net job attributes to define six standardized measures capturing the task content of occupations: non-routine cognitive analytical, non-routine cognitive interpersonal, non-routine manual physical, non-routine manual interpersonal, routine manual, and routine cognitive. We focus on the last two. The routine manual index is created by summing studentized versions of three attributes (“Controlling Machines and Processes [4.A.3.a.3],” “Spend Time Making Repetitive Motions [4.C.2.d.1.i],” and “Pace Determined by Speed of Equipment [4.C.3.d.3]”) and re-studentizing the sum. The routine cognitive index is created by summing studentized versions of three attributes (“Importance of Being Exact or Accurate [4.C.3.b.4],” “Importance of Repeating Same Tasks [4.C.3.b.7],” and (reverse-scaled) “Structured v. Unstructured Work [4.C.3.b.8]”) and re-studentizing the sum. See AA for additional details. We depart slightly from AA in not using employment weights when creating the indices at the 6-digit SOC level, as we are interested in the distribution of routineness across occupations, not workers. In practice, this difference does not matter much. We aggregate to the four-digit SOC occupation category by taking an OES-employment weighted average across nested 6-digit occupations.

A.8 Sample Restrictions and the Bartik Shock

Table 1 examines how the probability of meeting certain sample criteria varies with the MSA-employment shock over time. Because we use a first-difference specification (see equation 1), the most relevant threat to internal validity would be if the *change* in the probability of meeting our sample criteria (from the base period) varies systematically with the Bartik shock. We generally find that this is not the case.

Column 1 explores the probability of ads missing firm name, and thus being excluded from our main sample. On average, 39% of ads are missing firm name, likely because they

are posted to a recruiter’s website. However, we find that the change in this probability from 2007 does not vary meaningfully with the Bartik shock. Most coefficients in column 1 are small in magnitude, change sign, and are statistically insignificant. The one exception is in 2012, where the estimate implies that hard-hit MSAs are 6.6 ppts less likely to have ads with a firm name, relative to 2007, than less hard-hit MSAs. This relative decline is significant at the 1% level. As pointed out in section 4, however, 2012 appears to be an unusual year in the BG data, and we express caution in overinterpreting any finding from this one year.

Column 2 explores the probability of being excluded from the sample of firms used for figure 3 (showing that upskilling in the later period is driven by the same firms that upskilled initially), by firms not posting at least five ads each in 2007 and 2010. About one-third of weighted observations in our main sample do not satisfy this criterion. We again find that the change in the share of ads within an MSA is unrelated to this sample restriction except for the anomalous 2012 (and weakly—at the 10% level—in 2011).

Column 3 looks at a related criterion: the probability of exclusion from the sample used for figure 5 (the differential upskilling effects for firms that have invested heavily in capital). For this sample we need to observe a firm in 2007 and again at least once in any later year. Among weighted observations in the firm sample, 31% do not meet this criterion. However, as can be seen, the change in satisfying the restriction does not vary with the Bartik shock. Aside from the one in 2012, coefficient estimates are statistically insignificant and small in magnitude.

In columns 4 and 5, we investigate whether the share of ads that can be matched at the firm level to Harte Hanks and Compustat, respectively, varies over time and with the Bartik shock. As noted in the text, @% and @% of ads can be matched to HH and Compustat, respectively, conditional on meeting the criteria in column 3. Among all ads with firm identifiers, match rates are 50% and 70%, respectively. Relative to these rates, the estimates in columns 4 and 5 are small in magnitude; none is statistically insignificant, and signs vary.

Finally, Column 6 explores whether a employment-weighted site (akin to an establishment) is present in the database at least once in 2002, 2004, or 2006, among the entire sample of firms in HH (not just those matching to BG). This restriction is necessary to generate PCs normalized by pre-recession employment at the site level. Overall, 35% of employment in HH do not meet this restriction. In general, we find that the change in this probability from 2006 does not vary with the Bartik shock. Point estimates are small and statistically insignificant, with one marginal exception in 2000.

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

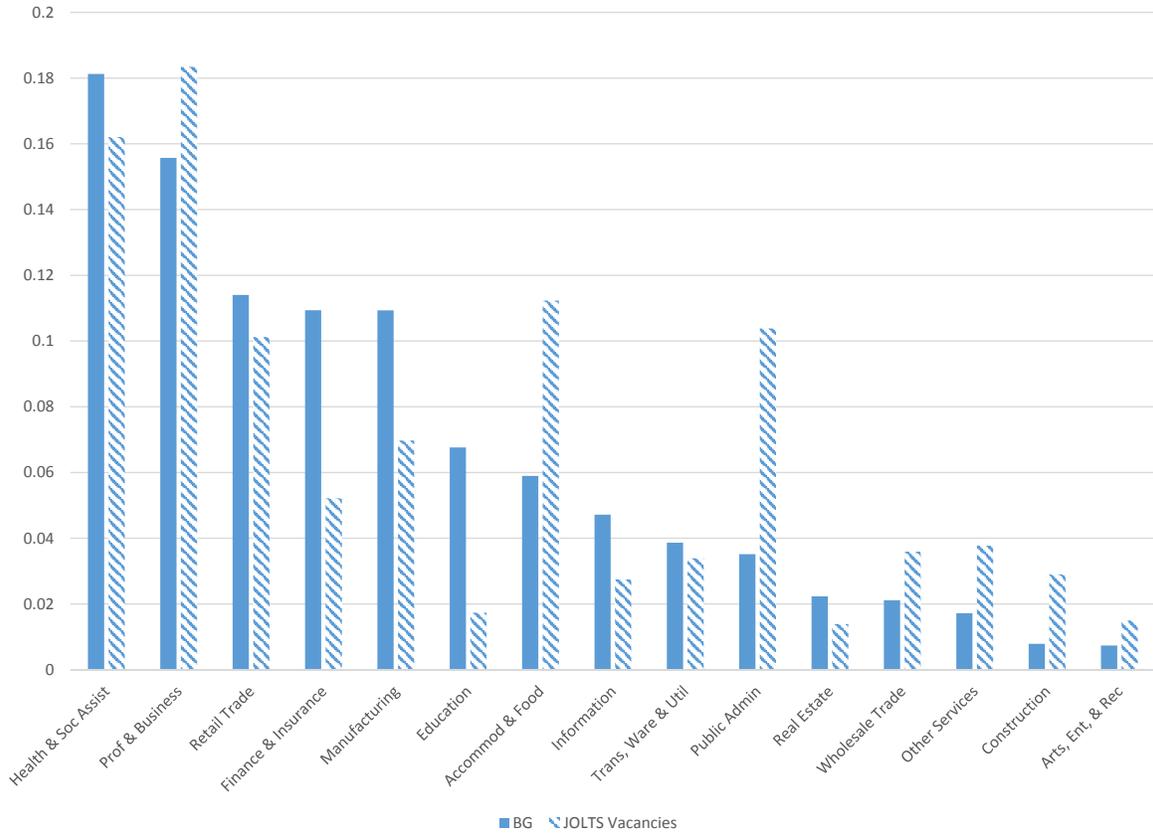


Figure A2: Occupation Distributions: BG, New Jobs (CPS) and Employment (OES)

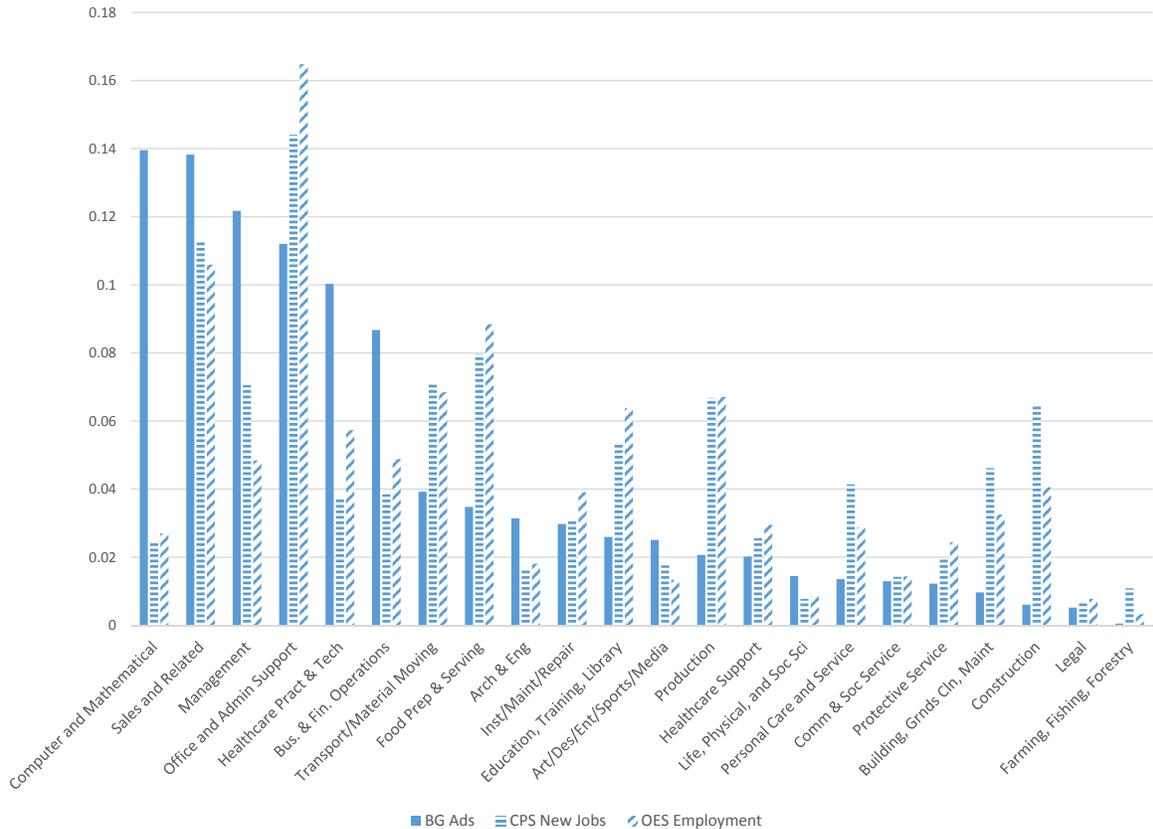
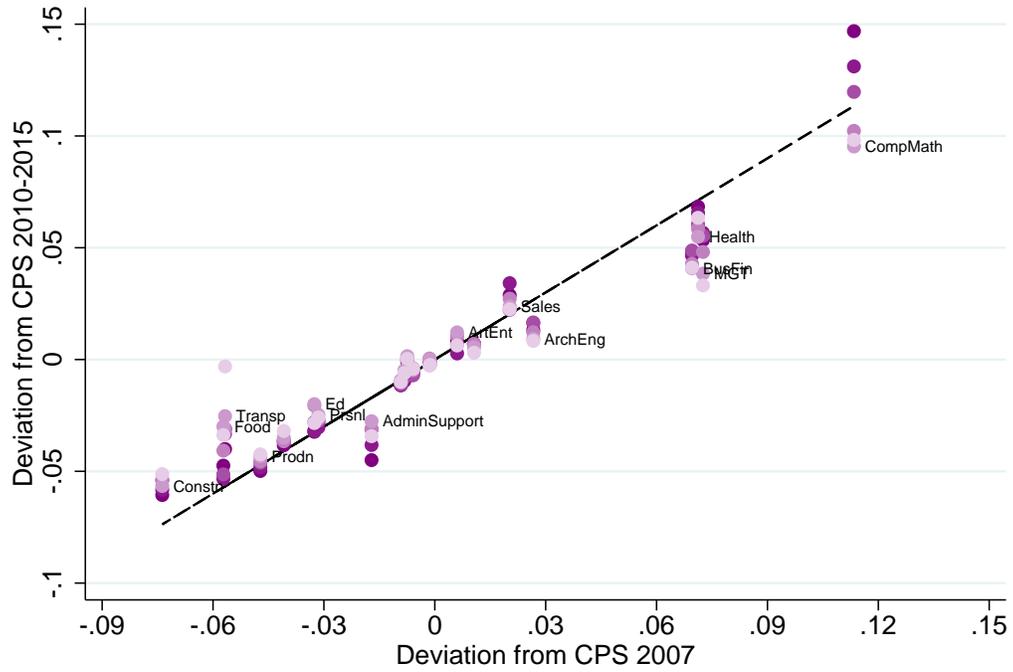


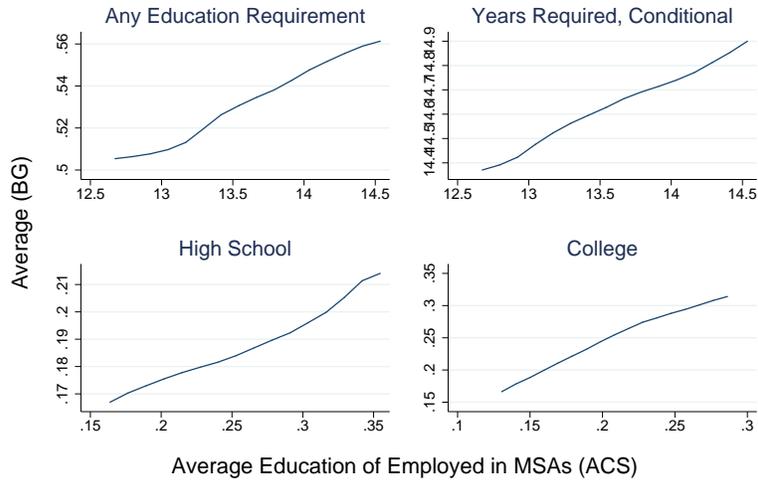
Figure A3: Representativeness of BG Occupations, Relative to New Jobs (CPS)



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.

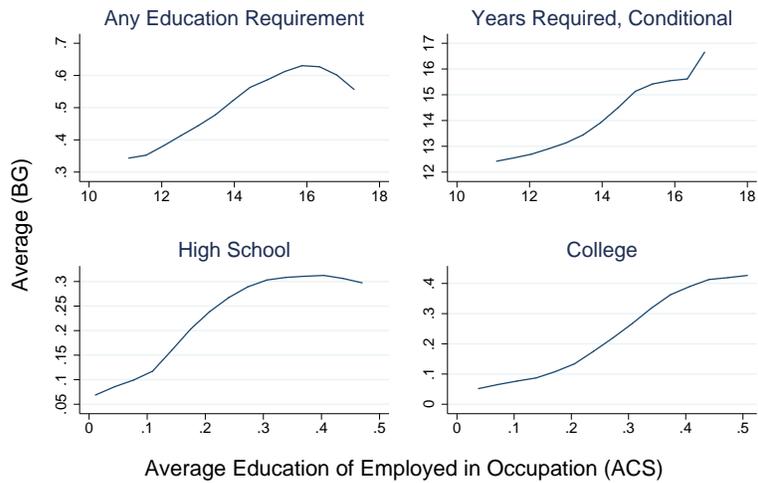
Figure A4: Comparison of BG Education Requirements and ACS Employment

(a) by MSA



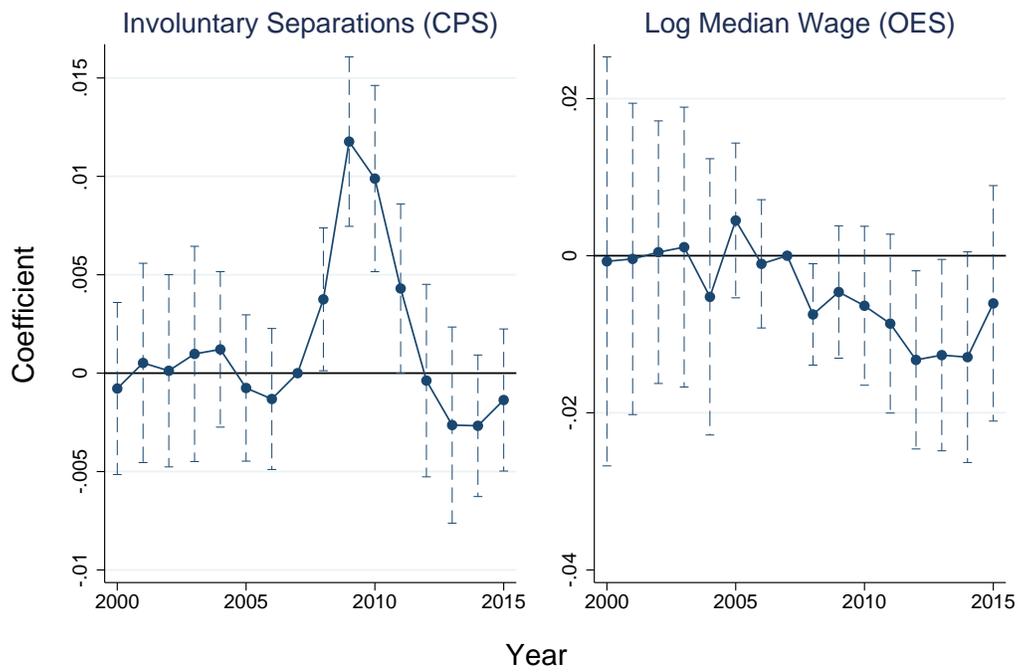
Smoothed local linear regression of MSA-level education requirement on ACS education percentile. Top two panels show average years school of employed workers in the MSA. Bottom shows share with high school (left) and college (right).

(b) by Occupation



Smoothed local linear regression of occupation-level education requirement on ACS education percentile. Top two panels show average years school of employed workers in the occupation. Bottom shows share with high school (left) and college (right).

Figure A5: Layoffs and Wages: Full Effects



Graph plots the coefficients on Bartik shock*year, as well as 95% confidence interval bars. Regressions control for year fixed effects and MSA characteristics.

Table 1: Probability of Being Missing from a Subsample

Dep Var:	Burning Glass: Change in Pr(Missing from):						Harte Hanks:	
	Firm Sample (1)	2007-10 match (2)	Firm pre-post match (3)	Harte Hanks (4)	Compustat (5)	Site pre- recession match (6)		
Shock*2000	-0.0273 (0.0178)	0.0136 (0.0146)	0.00921 (0.0125)	-0.0146 (0.0147)	-0.0139 (0.0113)	0.0148* (0.00781)		
Shock*2002	-0.0299 (0.0184)	0.0246* (0.0141)	0.0199 (0.0122)	-0.00598 (0.0154)	-0.000753 (0.00953)	0.00442 (0.00585)		
Shock*2004	-0.0656*** (0.0195)	0.0344** (0.0161)	0.0289** (0.0143)	0.00622 (0.0153)	0.00165 (0.00971)	0.00442 (0.00585)		
Shock*2008						0.00546 (0.00671)		
Shock*2010						0.0110 (0.00927)		
Shock*2011								
Shock*2012								
Shock*2013								
Shock*2014								
Shock*2015								
# MSA-Year Cells	2,286	2,286	2,286	2,286	2,286	2,667		
R-Squared	0.265	0.584	0.689	0.428	0.355	0.985		

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

Notes: Observations are MSA-year cells weighted by the MSA labor force in 2006; regressions also control for year effects and MSA characteristics; standard errors are clustered at the MSA level. We regress the change from 2007 (2006 in column 6) in the share of ads (site employment in column 6) meeting the indicated sample criterion. Column 1 is estimated on the full BG sample of ads and examines the change in the probability of having a missing firm name. Columns 2-5 are estimated on our primary non-missing firm BG sample and examine the change in the probability of being missing from: the sample of firms that post at least 5 ads in 2007 and 2010 (column 2), the sample of firms that post in 2007 and a later year (column 3), the Harte Hanks firm match sample (column 4), the Compustat firm match sample (column 5). BG observations are restricted to 2010-2015. Column 6 examines the change in shares of employment that are in sites (akin to establishments) that can be matched to at least one year in 2002, 2004, or 2006, in the Harte Hanks sample.

B Additional Results

B.1 Education and Experience Intensive Margins

Appendix figure B1 summarizes results for additional education and experience outcomes in order to understand changes in the intensive margin for these requirements. The top two panels of figure B1a show similar-sized increases in the probability of requiring a high school diploma and the probability of requiring 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 panel). Also, there is no change in the propensity to require a graduate degree (bottom left panel). This last finding is reassuring, since many professional jobs, such as lawyers and professors, have long required postgraduate degrees; these requirements would not be expected to change with improvements in technology. Figure B1b exhibits a similar pattern for experience requirements. We observe increases in experience requirements at the low (2 years or less) and middle (3–5 years) parts of the distribution; there is little change at the high end. As with education, this pattern results in little net change in total years of experience required, conditional on posting any requirement.

Are these increases in requirements plausible? For example if increases in college (high school) requirements were found in typically very low- (high-) skilled jobs, we might worry about the quality of the data. In figure B2 we explore heterogeneity in upskilling effects as a function of the average years of schooling of workers in the occupation, as measured in the ACS in 2005–06, before the Great Recession. We estimate separate within-occupation upskilling regressions (equation (1)), by ventile of this occupational education variable, for the change in the propensity to specify a high school diploma (left) or bachelor’s degree (right) requirement. We summarize these regression results by plotting the estimated coefficients on the Bartik shock for 2010 (blue, solid lines) and 2015 (maroon, dashed lines), smoothing the 20 estimates in each series with local linear regression.

Focusing first on the left panel, we find that increased demand for a high school credential is largest for occupations that tend to employ workers with less education before the recession. In fact, the strongest increases occur at the lowest ventiles, representing occupations whose workers have an average of about 11 years of schooling. Effects monotonically decline with the average education of occupations, and are essentially zero for occupations that typically employ college graduates. This pattern is highly consistent with the increased propensity to require a high school credential in job postings reflecting upskilling. It does not appear to be the case, for instance, that firms add the requirement to signal to workers that they do not need *even more* education.

The right panel shows that the increased propensity to require a bachelor’s degree is concentrated in occupations that tend to employ workers with at least some college, peaking for occupations that typically require a bachelor’s degree (16 years). Requirement

increases are non-existent for occupations that tend to employ high school graduates (12 years) or those with less education, and increases are smaller for occupations that tend to employ workers with post-graduate degrees (more than 16 years). This pattern of targeted increases in line with expectations is also reassuring that our estimates reflect upskilling. What’s more, the persistence in upskilling in our main estimates shows up where it is expected, as the blue and maroon lines essentially overlap; it is not the case that increased requirements are temporarily concentrated in overly high (or low) parts of the distribution.

B.2 Robustness Checks

Tables B1-B4 provide a range of robustness checks to the main within-occupation upskilling regression results from table 2. For each table, column 1 replicates the main result, and the remaining columns provide results from different specifications, discussed below.

Local labor market controls

In column 2, we control for additional MSA-level labor market variables. As noted in the text, controls proxying for the availability of labor across skill groups help to clarify the importance of opportunistic upskilling, as well as firm reactions to the availability of skilled labor, more broadly. We control for a wide range of labor market characteristics over our sample period, and changes in these characteristics from the previous decade. Specifically, drawing from the ACS, we include unemployment rates and employment-to-population ratios at the MSA level for five education groups (high school dropouts, high school graduates, those with some college, those with a BA, and those with more than a BA). To reduce measurement error from occasionally small samples, these rates and ratios are calculated as the average over 2005–07. Additionally, we include the change in the rates and ratios between 2005–07 and the current year-pair (2010–11, 2012–13, 2014–15), as well as the change between 2000 (using Census data) and 2005–07.

The results in column 2 include two additional sets of controls. The first are MSA-level quit rates and their changes, by education group, obtained from the CPS. Because the sample size of the CPS is much smaller than that of the ACS, we aggregate across broader education groups and pool more sample years. We distinguish between those with no more than a high school diploma and those with some college or more. Quit rates are averaged for each education group over 2005–07, and the changes are between 2005–07 and the current year-triple (2010–12, 2013–15), as well as between 2000–02 to 2005–07. The second set of controls complements the 2005–06 MSA characteristics already included in the main specification with the changes in these characteristics between 2000 and 2005–06: the share of the population that is female, black, Hispanic, Asian, married, 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, is age 50–64, the overall employment-to-population ratio, and the average weekly wage of full-time workers. These changes are calculated using the ACS and 2000 Census, and we include a dummy to capture cases in which MSAs are not identifiable in the ACS or CPS.

As shown in column 2, the inclusion of all of these local labor market controls does not have a substantive impact on the estimates, which are either reduced by between one-tenth and one-fifth (education and experience) or are essentially changed or even slightly larger (cognitive and computer). We have also explored including only subsets of these additional controls (these estimates are not shown in the tables). When we include just the education-specific unemployment rates, employment-population ratios, and their changes, the estimates are generally reduced by between one-tenth and one-fifth, and sometimes slightly more in the earlier years of the sample. However, they remain significant in both magnitude and a statistical sense. When we include only education-specific quit rates and their changes, the estimates increase modestly from baseline. A possible explanation for this pattern is that we find that quit rates of less-educated workers rebound more quickly than that of more-educated workers. Replacement hiring would then shift toward lower skilled workers, absent these controls. Finally, when we include only changes in MSA characteristics (in addition to the levels included in baseline), the estimates fall by about one-fifth for education and experience and about one-tenth for cognitive; they do not change for computer. In no case, therefore, does the general picture of our baseline estimates change.

Occupation controls

As noted in the text, results are robust to the inclusion of occupation fixed effects and occupation-specific time trends to allow for the possibility that some occupations may be both upskilling at a faster rate and disproportionately located in harder-hit MSAs. These controls also help adjust for possible changes in the sample due to shifting occupational mix in the BG data. Column 3 shows that we obtain very similar results even when including occupation-specific controls.

Weights

Our baseline within-occupation upskilling regressions in table 2 weight cells by the product of the occupation’s ad share in the MSA-year and the size of the MSA labor force in 2006. This procedure can reduce measurement error and attenuation bias by putting more weight on cells that have more ads while mitigating against the possible endogeneity of using the number of ads themselves as weights. However, if there are heterogeneous impacts of the shock on upskilling tied to the size of the MSA or occupation share, we may be identifying a local average treatment effect that may not hold across all cells. Column 4 of the robustness tables thus provides estimates when we instead allow equal weight across occupation-MSA-

years, limiting the sample to the 99% of ads that are in cells with at least 15 ads. We obtain qualitatively similar results that are generally still statistically significant, although, consistent with the possibility of attenuation bias, reduced in magnitude in some cases. This pattern also holds for the overall MSA-year level regressions, as in figure 2, and when we instead weight to match the occupation-MSA employment stock distribution in the OES or the national occupation flow (hiring) distribution in the CPS (not shown).

Ads with missing firm name

Our baseline results focus on the 60% of ads that contain firm name, as this sample allows us to distinguish among firm-level mechanisms for upskilling. One might be concerned, though, that ads with firm names have differentially changing skill requirements relative to ads without firm names, perhaps due to firm size, prestige, or use of a third-party recruiting firm. Column 5 shows the robustness of our results to using the full sample of ads, including those with missing firm names. Results on the expanded sample are quite similar for all dependent variables.⁶²

Alternative Bartik employment shocks

Our preferred measure of the local labor demand shock, used for our baseline estimates, is the change in projected annual employment growth between 2006 and 2009 (equation (2)). Columns 6–8 of the tables provide results using three variants of the Bartik employment shock. Rather than using fixed peak and trough years for all MSAs, Column 6 instead allows for MSA-specific peak and trough dates, using the calendar month with the largest 12-month employment growth between 2005 and 2007 as the peak date, and the calendar month with the smallest 12-month employment growth from 2008 onwards as the trough date. Column 7 uses the change in projected employment *levels* between 2006 and 2009, rather than the change in projected employment *growth*.⁶³ Column 8 uses the change between the average projected one-month employment growth in 2006 and the average in 2009, rather than the change in annual employment growth between the two years. The results are not especially sensitive to using MSA-specific cycle timing (column 6) or short-run projected changes in employment growth (column 8). Using projected changes in employment levels as the shock measure (column 7) reduces the magnitude of the point estimates by about one-half, but they are still statistically significant and of meaningful magnitude.

⁶²We 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 of appendix table 1.

⁶³We believe the latter approach better captures the suddenness of the shift in conditions between 2006 and 2009 as they pertain to flow employment, but the literature has sometimes used the levels approach (Bartik 1991). In practice, they are highly correlated, with an employment-weighted correlation of $r = 0.86$.

Industry Controls

In order to explore heterogeneity within and across industries, we disaggregate our data by industry (2-digit NAICS). Columns labeled 1 in table B5 estimate changes in skill requirements within industry-occupation cells as a function of the MSA employment shock.⁶⁴ We find that these estimates are very similar to the results presented in figure 2 and table 2. Firms in harder-hit MSAs differentially increase skill requirements in job postings, and this effect holds true within industry-occupation cells.

The second column of table B5 adds industry fixed effects and industry-specific linear time trends, and our estimates are essentially unchanged. These controls are particularly important given our identification: the interaction of national changes in employment growth by three-digit NAICS industry and MSA-level industry composition. Suppose that the industries driving negative employment shocks are precisely those that experience contemporaneous technology shocks—and concomitant temporary employment declines during an adjustment period. In this case our results would not indicate that firms concentrated their adoption of *existing* technologies during recessions but rather that the occurrence of the innovations themselves was concentrated during the recession. The fact that our results obtain even within sector alleviates this concern. Allowing for what essentially amounts to a quadratic secular change in skill requirements for each sector (given our first differences specification), we still find quantitatively similar evidence of upskilling.

We also investigate whether upskilling is more pronounced in sectors producing locally-consumed goods and services. Unlike the tradeable sector, the product demand for which is largely determined by markets farther away or diffused across many areas, firms producing locally-consumed goods and services are highly sensitive to local demand shocks, and thus should be more greatly affected by the variation we identify with the Bartik employment shock.

To classify sectors, we adopt Jensen and Kleitzer’s (2005) measure of the degree to which production is “offshorable,” which is based on geographic concentration of employment in the industry. Intuitively, if employment for a sector can be geographically concentrated (e.g., software developers in Silicon Valley), then output for that sector is more likely traded and need not be consumed locally. From Jensen and Kleitzer (2005), we obtain the share of employment in each two-digit NAICS sector that can be categorized as “*least* [geographically] concentrated,” which we denote $local_s$.⁶⁵

We then estimate versions of equation 1 at the occupation-sector-MSA-year level that

⁶⁴As before, we weight each cell by the product of the size of the MSA’s labor force in 2006 and the cell’s ad share in each MSA-year, where the ad shares here are over industry-occupation groups. We require a match at the MSA-industry-occupation level between 2007 and at least one later year; this restriction drops 6% of ads.

⁶⁵Jensen and Kleitzer (2005) measure geographic dispersion across MSAs of employment in detailed industry categories, and designate the category “least concentrated” to those industries with a Gini coefficient less than 0.1. We have also used the measure from Blinder and Krueger (2013), based primarily on workers’ survey responses on location requirements to do their jobs, to similar effect.

include sector fixed effects and allow for triple-interactions between the Bartik shock, year dummies, and $local_s$.

Figure B3 plots the estimates fitted at the 10th and 90th percentiles of $local_s$; these respectively capture the net upskilling effect for traded (dashed maroon line) and non-traded (solid blue line) sectors. With the exception of education, we find that increases in skill requirements are *larger* in the non-traded sectors, those in which local demand shocks should be the most salient for production.

B.3 Firm-Occupation Decomposition

We here explore the extent to which upskilling is driven by shifts in postings from old to new firms and changes within existing firms. By employing a formal decomposition, we investigate these margins simultaneously with shifts in ads across firms and those across occupations.⁶⁶

Define C_t as the set of firm-MSAs that post ads in both year t and in 2007. We hereafter refer to these as “continuing firms,” and the set of firm-MSAs that have posts only in 2007 or only in t as non-continuing firms. In our sample, 54% of weighted observations are to continuing firms.⁶⁷ We hope to understand the extent to which substitution across non-continuing firms versus changes within continuing firms affects overall changes in skill requirements.

In equation (5), we express the average skill requirement in MSA m and year t as a function of: p_{mt}^C , the share of ads in an MSA-year posted in continuing firms;⁶⁸ $\frac{N_{fmt}^C}{N_{mt}^C}$, the distribution of ads across continuing firms in mt ; $\frac{N_{ofmt}^C}{N_{fmt}^C}$, the distribution of ads across occupations for a given continuing firm; $skill_{ofmt}^C$, the average skill requirement for continuing firm f , posting in occupation o , MSA m , and year t ; $\frac{N_{omt}^{NC}}{N_{mt}^{NC}}$, the distribution of ads across occupations among non-continuing firms;⁶⁹ and \bar{skill}_{omt}^{NC} , the average skill requirement for occupation o , among all non-continuing firms in mt (that is, the average skill requirement in the occupation-MSA-year among firm-MSAs that posted either only in 2007 or only in period t).

$$(5) \quad skill_{mt} = p_{mt}^C \sum_{fm \in C_t} \sum_o \frac{N_{fmt}^C}{N_{mt}^C} \frac{N_{ofmt}^C}{N_{fmt}^C} * skill_{ofmt}^C + (1 - p_{mt}^C) \sum_o \frac{N_{omt}^{NC}}{N_{mt}^{NC}} \bar{skill}_{omt}^{NC}.$$

⁶⁶Kahn and McEntarfer (2015) show that workers matching to jobs in downturns are more likely to match to low-paying firms than high-paying firms.

⁶⁷Here we define “firm” as the group of ads with the same employer name in the same MSA, which allows us to take advantage of the cross-sectional variation in how MSAs bore the Great Recession. The set C_t is defined separately for each year from 2010–2015, though naturally there is substantial overlap in the set of continuing firms across years.

⁶⁸By definition, $p_{mt}^C \equiv \frac{N_{mt}^C}{N_{mt}^C + N_{mt}^{NC}}$.

⁶⁹Since, by definition, non-continuing firms cannot be matched across time periods, we aggregate over all non-continuing firms.

We then decompose the effect of the Bartik employment shock on the overall change in skill requirements at the MSA-year level ($skill_{mt} - skill_{m07}$), into effects attributable to changes in: p_{mt}^C , the share of ads in an MSA-year posted in continuing firms; $\frac{N_{fmt}^C}{N_{mt}^C}$, the distribution of ads across continuing firms in mt ; $\frac{N_{ofmt}^C}{N_{fmt}^C}$, the distribution of ads across occupations for a given continuing firm; $skill_{ofmt}^C$, the average skill requirement for continuing firm f , posting in occupation o , MSA m , and year t ; $\frac{N_{omt}^{NC}}{N_{mt}^{NC}}$, the distribution of ads across occupations among non-continuing firms; and \bar{skill}_{omt}^{NC} , the average skill requirement for occupation o , among all non-continuing firms in mt (that is, the average skill requirement in the occupation-MSA-year among firm-MSAs that posted either only in 2007 or only in period t).

In practice, this equation is not exact for two reasons. First, continuing firms do not necessarily post to the same set of occupations in each period (so $skill_{ofmt}^C$ would not be defined for some occupation-firm-MSA-year combinations but might be defined in, say, 2007). Second, the set of non-continuing firms does not post to the same set of occupations (so likewise \bar{skill}_{omt}^{NC} , which is the average skill requirement among all non-continuing firms posting in omt , would not be defined for some occupation-MSA-year combinations). To get around these issues, we simply aggregate up from the occupation-firm-MSA-year level to either the occupation-MSA-year level or the MSA-year level, the point where we get a match.

The exact definition is shown in equation (6).

$$\begin{aligned}
(6) \quad skill_{mt} = & p_{mt}^C * \pi_1^c \sum_{fm \in C_t} \sum_{o \in CO^1} skill_{ofmt}^{CO^1} * \frac{N_{ofmt}^{CO^1}}{N_{fmt}^{CO^1}} \frac{N_{fmt}^{CO^1}}{N_{mt}^{CO^1}} \\
& + p_{mt}^C (\pi_2^c \sum_{o \in CO^2} \bar{skill}_{omt}^{CO^2} * \frac{N_{omt}^{CO^2}}{N_{mt}^{CO^1}} + (1 - \pi_1^c - \pi_2^c) \bar{skill}_{mt}^{CO^3}) \\
& + (1 - p_{mt}^C) (\pi_1^{nc} \sum_{o \in NCO^1} \bar{skill}_{omt}^{NCO^1} * \frac{N_{omt}^{NCO^1}}{N_{mt}^{NCO^1}} + (1 - \pi_1^{nc}) \bar{skill}_{mt}^{NCO^2})
\end{aligned}$$

In the top two lines, we divide the set of ads to continuing firms into three groups: occupations that are posted in a given firm-MSA in both t and in 2007 (the set CO^1), occupations that are not posted in a given firm-MSA in both periods but are posted among other continuing firms in both periods (CO^2), and occupations that are posted in one period by continuing firms but not in the other period (CO^3). The ad shares for these three groups (π_1^c , π_2^c , and $1 - \pi_1^c - \pi_2^c$, respectively) sum to one within the set of ads to continuing firms (C_t).⁷⁰ Averaging across 2010–2015 in our data, 54% of weighted observations are to continuing firm-MSAs, of which 71% are to continuing occupations (CO^1), 28% are to non-continuing occupations that can still be matched to any other continuing firms (CO^2) and only 0.8% are to occupations that cannot be matched to any continuing firms (CO^3).

The first component, for continuing firm-occupations (CO^1) is straightforward and is

⁷⁰These weights vary by MSA-year, but subscripts are suppressed for clarity.

defined by the within *ofm* average skill requirement ($skill_{ofmt}^{CO^1}$), the share of ads in this occupation, *o*, for the given *fm* ($\frac{N_{ofmt}^{CO^1}}{N_{fmt}^{CO^1}}$), and the share of all ads in CO^1 that are by firm *f*, ($\frac{N_{fmt}^{CO^1}}{N_{mt}^{CO^1}}$). The second component yields the average skill requirement among occupations by continuing firms that do not post for the same occupation in both periods. The idea is that we would like to compare the firm-occupation-specific requirement across years. However, in some cases continuing firms post in a new occupation in *t*, so the comparison is not available. Instead, we simply fix occupations and aggregate over firms. We can then ask whether skill requirements are more strenuous for continuing firms that enter into the occupation, relative to those that exited the occupation.⁷¹ $\bar{skill}_{omt}^{CO^2}$ is the average skill requirement in occupation *o* among firm-MSAs that posted some ads in both periods, but posted in occupation *o* only in the given period, and $\frac{N_{omt}^{CO^2}}{N_{mt}^{CO^1}}$ is the ad share for the given occupation among all ads in the set (CO^2). Finally, the third component, $\bar{skill}_{mt}^{CO^3}$, is the average skill requirement in the MSA-year for ads posted by continuing firms in occupations that belong to neither CO^1 nor CO^2 (that is, an occupation where continuing firms either post only in 2007 or only in *t*).

In the third line of equation (6), we divide the set of ads to non-continuing firms into two groups: occupations that are posted in the MSA in both *t* and in 2007 (the set NCO^1) and those that are not (NCO^2), with weights π_1^{nc} and $1 - \pi_1^{nc}$, respectively. Skill requirements for the former are a function of the within-occupation average skill requirement among all ads posted by non-continuing firms in *mt* in occupations that can be matched ($\bar{skill}_{omt}^{NCO^1}$) and the share of ads that are posted to occupation *o*, ($\frac{N_{omt}^{NCO^1}}{N_{mt}^{NCO^1}}$). The latter component is the average skill requirement among ads posted to non-continuing firms in *mt* in occupations that cannot be matched ($\bar{skill}_{mt}^{NCO^2}$). Of the 46% of weighted observations that are to non-continuing firms in our data, 97.5% of observations belong to the former (matched) group.

To decompose the change in skill requirements for a given MSA, *m*, from 2007 to a given year *t* ($skill_{mt} - skill_{m07}$) into these components we generate counterfactual differences that allow one component to change from its level in 2007 to its level in *t*, holding all other components fixed at the level in either period.⁷² 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 responsiveness is attributed to a response in the within firm-MSA skill requirement.

⁷¹We could instead fix firms and aggregate over occupations. This allows us to ask whether the new occupations the firm enters into have different skill requirements than the occupations the firm left. However this yields far fewer matches, since it requires the firm-MSA to have at least one occupation it posted in 2007 but not *t* and at least one occupation it posted in *t* but not 2007. Many firm-MSAs either cease posting in some occupations after 2007 or begin posting to new occupations in *t*, but not both. Instead, aggregating over firms and fixing occupations requires that any continuing firm-MSA ceases posting in the occupation after 2007 and any continuing firm-MSA begins posting in the occupation in *t*.

⁷²For example, $p_{mt}^C(\pi_1^c \sum_{fm \in C_{to} \in CO^1} \sum (skill_{ofmt}^{CO^1} - skill_{ofm07}^{CO^1}) * \frac{N_{ofmt}^{CO^1}}{N_{fmt}^{CO^1}} \frac{N_{fmt}^{CO^1}}{N_{mt}^{CO^1}})$ is the change in skill requirements between 2007 and *t*, attributed to just changes in the within occupation-firm-MSA skill requirements among continuing firms, holding constant all other components at their levels in *t*.

A decomposition begins with $skill_{m07}$ and generates a counterfactual skill change distribution that allows only one of the components to vary. That component is then fixed at its value at time t and a second counterfactual skill change distribution is generated by allowing a second component to vary, keeping all components but the first two fixed at their 2007 level. This process continues until all components are at their time t values. 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 on the full change reported earlier.

Naturally the order of the decomposition affects the relative importance of each component. For p components in the decomposition, we have $p!$ possible orders. To reduce the state space, we combine many of the variables in equation (6) into a smaller set of components since they turn out to be empirically irrelevant.

We reduce the space to six components, resulting in 720 possible permutations. We estimate each decomposition and summarize results in Appendix table B6. Here we report the average fraction attributable to a given component, across all decompositions, as well as the standard deviation. The components are: (1) the share of ads among continuing firms, p_{mt}^C ; (2) the distribution of ads across continuing firms that post in the same occupation in both periods, $\frac{N_{ofmt}^{CO1}}{N_{fmt}^{CO1}}$; (3) the within-firm-occupation skill requirement for continuing firms, $skill_{ofmt}^{CO1}$; (4) the distribution of ads across occupations, which combines $\frac{N_{ofmt}^{CO1}}{N_{fmt}^{CO1}}$, $\frac{N_{omt}^{CO2}}{N_{mt}^{CO2}}$, $\frac{N_{omt}^{NCO1}}{N_{mt}^{NCO1}}$, π_1^c , π_2^c , and π_1^{nc} ; (5) the skill requirement among continuing firms posting to occupations they had not previously posted in, which combines $skill_{omt}^{CO2}$ and $skill_{omt}^{CO3}$; and (6) the skill requirement among non-continuing firms, which combines \bar{skill}_{omt}^{NCO1} and \bar{skill}_{omt}^{NCO2} .

To summarize results from the table, we report the fraction of the overall impact of $shock$ in the given year attributed to each component in figure B4. To make the graph easier to read, we combine some components together, focusing on the (empirically) most important.

The lightest bar shows the fraction of the overall upskilling effect in each year attributable to changes in the distribution of firms, combining the share of ads to continuing firms p_{mt}^C with the firm distribution among continuing firms $\frac{N_{fmt}^C}{N_{mt}^C}$. This combines columns 1 and 2 of table B6. The next lightest bar shows the fraction attributable to changes in the occupation distribution, combining the occupation distribution among continuing ($\frac{N_{ofmt}^C}{N_{fmt}^C}$) and non-continuing ($\frac{N_{omt}^{NC}}{N_{mt}^{NC}}$) firms. Across all dependent variables, we find that changes in these firm and occupation distributions account for very little of the upskilling effects.

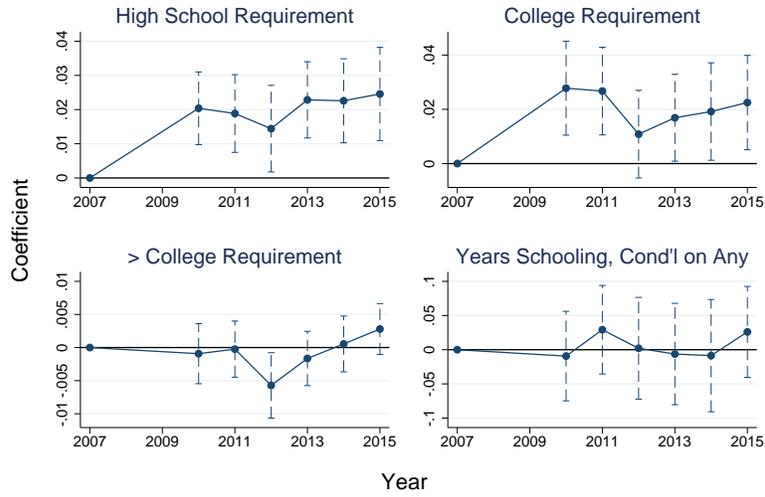
Instead, the vast majority of the upskilling effect is split across the two darker bars. The darkest bar shows the fraction attributable to changes in skill requirements of non-continuing firms for a given occupation (\bar{skill}_{omt}^{NC}). It compares, for each occupation, the skill requirements for firms that posted only in the later period (2010–2015) to the requirements for firms that stopped posting after 2007. The adjacent, next-darkest bars show the fraction attributable to the change in skill requirements between 2007 and t among continuing firm-

MSAs. This combines columns 3 and 5 of table B6.⁷³ Across dependent variables and years, each of these two components contributes to roughly half of the upskilling effect.

⁷³That is, this component includes both the change in skill requirements within firm-occupation-MSAs for firms that posted in a given *occupation* in both periods, and changes in skill requirements driven by continuing firm-MSAs that post for different occupations. Empirically, we find that both are important; the former is more important for education and experience requirements, while the two are roughly equally important for cognitive and computer requirements.

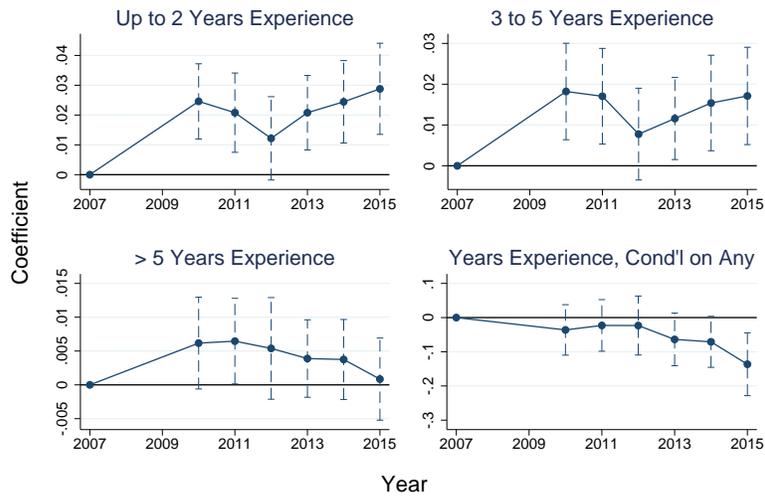
Figure B1: Impact of MSA-Specific Employment Shock on Education Requirements

(a) Education Requirements



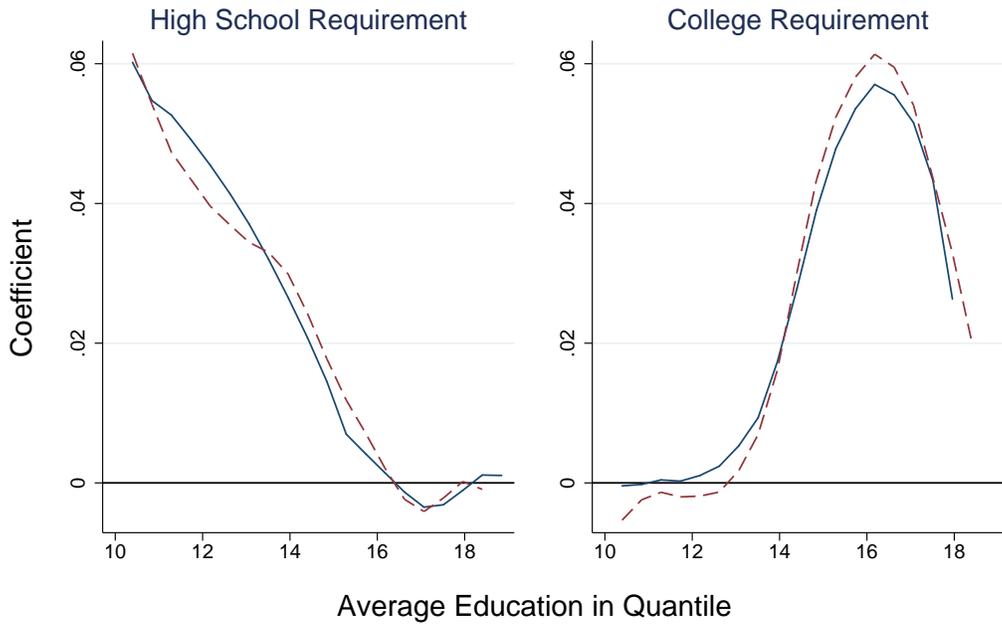
Graphs plot coefficients on MSA employment shock interacted with year, and 95% CI's (dashed lines). Regressions also control for year fixed effects and MSA characteristics.

(b) Experience Requirements



Graphs plot coefficients on MSA employment shock interacted with year, and 95% CI's (dashed lines). Regressions also control for year fixed effects and MSA characteristics.

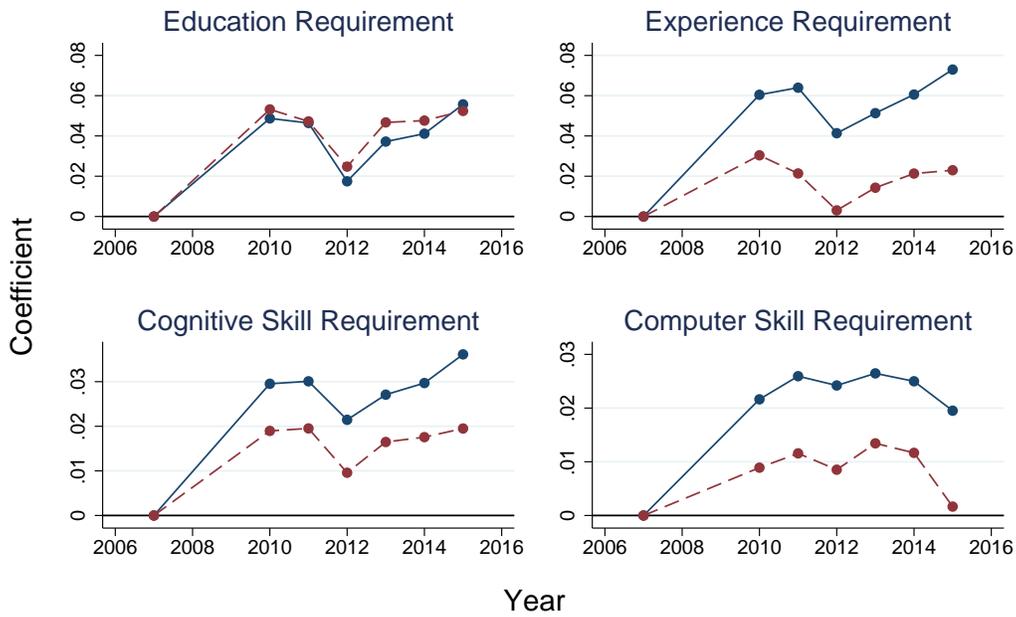
Figure B2: Change in Requirement by Occupation Education



Blue solid = 2007-2010 change, Maroon dash = 2007-2015.

We estimate separate within-occupation upskilling regressions for each ventile of the average years schooling in the occupation, (as measured by the ACS 2005-06 average). We plot the coefficients on the Bartik*2010 and Bartik*2015 year interactions for each ventile and smooth with local linear regression. Regressions also control for year fixed effects and MSA characteristics.

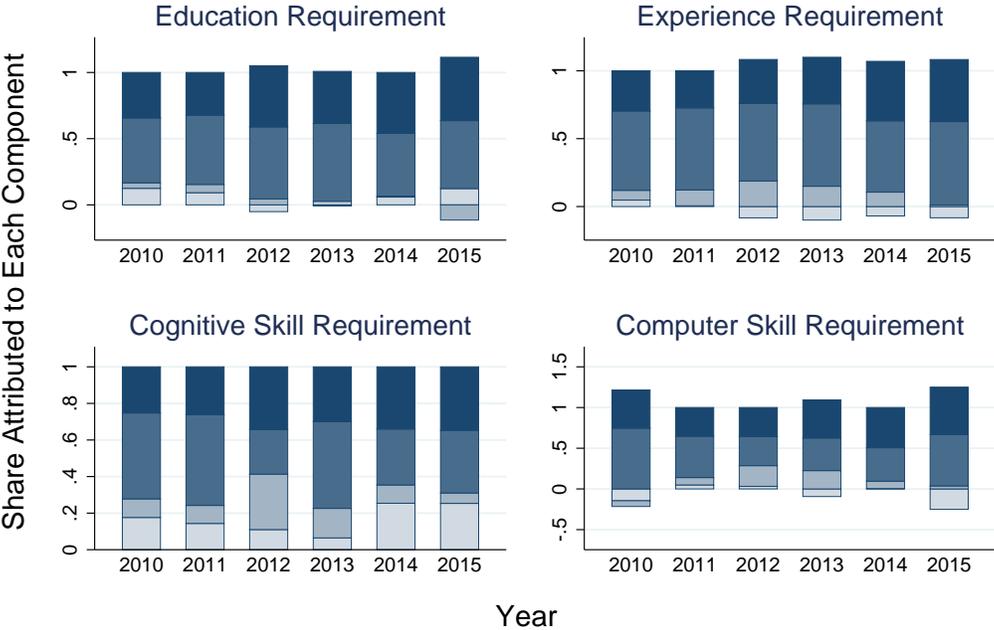
Figure B3: Upskilling by Sector Tradability



Blue line = local products, Maroon dashed line = traded products

Bartik shock*year effects interacted with offshorability, fitted for 90th and 10th percentiles. Regressions also control for sector and year fixed effects and MSA characteristics.

Figure B4: Decomposing Upskilling Into Within and Across Firm Components



Darkest = old-new firms within occ, dark = within continuing firms, light = occ dist, lightest = firm dist
 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 component, averaged across all possible decomposition orders.

Table B1: Robustness Checks: Education Requirement

<i>Dependent Variable:</i>	<i>Education Requirement</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Shock*2010</i>	0.0526*** (0.0135)	0.0528*** (0.0148)	0.0516*** (0.0128)	0.0333*** (0.0126)	0.0429*** (0.0126)	0.0521*** (0.0132)	0.0317*** (0.0118)	0.0465*** (0.0130)
<i>Shock*2011</i>	0.0475*** (0.0131)	0.0477*** (0.0144)	0.0462*** (0.0126)	0.0316** (0.0122)	0.0359*** (0.0119)	0.0483*** (0.0126)	0.0278** (0.0112)	0.0423*** (0.0124)
<i>Shock*2012</i>	0.0233* (0.0128)	0.0251* (0.0145)	0.0226* (0.0124)	0.0177 (0.0114)	0.0212* (0.0121)	0.0241** (0.0120)	0.00731 (0.0115)	0.0243** (0.0121)
<i>Shock*2013</i>	0.0400*** (0.0120)	0.0423*** (0.0131)	0.0406*** (0.0114)	0.0274*** (0.0101)	0.0343*** (0.0120)	0.0403*** (0.0116)	0.0215* (0.0110)	0.0359*** (0.0116)
<i>Shock*2014</i>	0.0429*** (0.0143)	0.0424*** (0.0140)	0.0440*** (0.0136)	0.0280** (0.0112)	0.0320** (0.0144)	0.0427*** (0.0136)	0.0218* (0.0131)	0.0374*** (0.0141)
<i>Shock*2015</i>	0.0488*** (0.0143)	0.0483*** (0.0149)	0.0532*** (0.0141)	0.0200** (0.0101)	0.0327** (0.0142)	0.0483*** (0.0140)	0.0291** (0.0122)	0.0413*** (0.0135)
# Occ-MSA-Year Cells	193,086	193,086	193,086	139,172	222,058	193,086	193,086	193,086
R-Squared	0.044	0.072	0.373	0.010	0.072	0.045	0.042	0.044
Labor Market Controls		X						
Occ FE and time trends			X					
Unweighted				X				
Includes Missing Firms					X			
Bartik w MSA-specific Peaks/Troughs						X		
Bartik in levels							X	
1 month employment change Bartik								X

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

Notes: Column 1 replicates estimates reported in table 2 of the text. All regressions control for year fixed effects and MSA characteristics from the ACS. Column 2 includes the change in the MSA characteristics from 2000 to 2005/6, as well as MSA-level education-specific employment, unemployment and quit rates in 2005-07, the change in these variables from 2005-07 to the current year, and the change in these variables from 2000-02 to 2005-07. Column 3 includes occupation fixed effects and occupation-specific linear time trends. Column 4 presents unweighted regressions (each occupation-MSA cell gets the same weight) and restricts to cells with at least 15 ads. Column 5 includes all ads, not restricting to those with non-missing firms. Column 6 defines the MSA employment shock using MSA-specific peak and trough years, rather than imposing 2006 and 2009, respectively, for all MSAs. Column 7 defines the MSA employment shock as the projected change in log employment from 2006 to 2009, rather than the projected change in employment growth). Column 8 defines the employment shock as the one-month change in employment growth, rather than year-over-year change. All shock variables are divided by the 90-10 differential in the variable across all MSAs.

Table B2: Robustness Checks: Experience Requirement

<i>Dependent Variable:</i>	<i>Experience Requirement</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Shock*2010</i>	0.0490*** (0.0134)	0.0453*** (0.0138)	0.0486*** (0.0130)	0.0295*** (0.0102)	0.0442*** (0.0138)	0.0495*** (0.0130)	0.0263** (0.0115)	0.0435*** (0.0131)
<i>Shock*2011</i>	0.0443*** (0.0134)	0.0406*** (0.0142)	0.0435*** (0.0133)	0.0245** (0.0106)	0.0340*** (0.0129)	0.0459*** (0.0129)	0.0237** (0.0114)	0.0406*** (0.0129)
<i>Shock*2012</i>	0.0253* (0.0136)	0.0257* (0.0148)	0.0244* (0.0134)	0.0139 (0.0105)	0.0199 (0.0136)	0.0266** (0.0129)	0.00502 (0.0119)	0.0275** (0.0129)
<i>Shock*2013</i>	0.0363*** (0.0122)	0.0370*** (0.0131)	0.0356*** (0.0121)	0.0172* (0.00881)	0.0312** (0.0124)	0.0366*** (0.0118)	0.0174 (0.0108)	0.0345*** (0.0119)
<i>Shock*2014</i>	0.0436*** (0.0140)	0.0418*** (0.0137)	0.0434*** (0.0137)	0.0206** (0.00963)	0.0303** (0.0139)	0.0431*** (0.0133)	0.0223* (0.0126)	0.0404*** (0.0140)
<i>Shock*2015</i>	0.0468*** (0.0142)	0.0450*** (0.0152)	0.0480*** (0.0139)	0.0150* (0.00866)	0.0328** (0.0147)	0.0463*** (0.0137)	0.0326*** (0.0124)	0.0431*** (0.0140)
# Cells	193,086	193,086	193,086	139,172	222,058	193,086	193,086	193,086
R-Squared	0.069	0.102	0.354	0.011	0.102	0.069	0.066	0.068
Labor Market Controls		X						
Occ FE and time trends			X					
Unweighted				X				
Includes Missing Firms					X			
Bartik w MSA-specific Peaks/Troughs						X		
Bartik in levels							X	
1 month employment change Bartik								X

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

Notes: See Table 2 and Appendix Table B1.

Table B3: Robustness Checks: Cognitive Skill Requirement

<i>Dependent Variable:</i>	<i>Cognitive Skill Requirement</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Shock*2010</i>	0.0275*** (0.00726)	0.0242*** (0.00785)	0.0270*** (0.00674)	0.0120* (0.00655)	0.0173** (0.00711)	0.0262*** (0.00702)	0.0161** (0.00641)	0.0236*** (0.00687)
<i>Shock*2011</i>	0.0281*** (0.00731)	0.0248*** (0.00754)	0.0272*** (0.00682)	0.0166*** (0.00600)	0.0200*** (0.00652)	0.0272*** (0.00707)	0.0151** (0.00635)	0.0258*** (0.00731)
<i>Shock*2012</i>	0.0186*** (0.00693)	0.0161** (0.00764)	0.0182*** (0.00648)	0.0117** (0.00574)	0.0141** (0.00596)	0.0175*** (0.00662)	0.0102 (0.00644)	0.0177*** (0.00660)
<i>Shock*2013</i>	0.0253*** (0.00642)	0.0229*** (0.00687)	0.0248*** (0.00625)	0.0163*** (0.00540)	0.0165*** (0.00564)	0.0242*** (0.00625)	0.0136** (0.00599)	0.0221*** (0.00622)
<i>Shock*2014</i>	0.0265*** (0.00657)	0.0242*** (0.00678)	0.0268*** (0.00635)	0.0154*** (0.00520)	0.0174*** (0.00630)	0.0251*** (0.00622)	0.0154** (0.00601)	0.0245*** (0.00655)
<i>Shock*2015</i>	0.0300*** (0.00730)	0.0278*** (0.00766)	0.0320*** (0.00711)	0.0171*** (0.00518)	0.0243*** (0.00777)	0.0284*** (0.00713)	0.0156** (0.00642)	0.0260*** (0.00733)
# Cells	178,176	178,176	178,176	135,878	207,552	178,176	178,176	178,176
R-Squared	0.040	0.050	0.287	0.006	0.075	0.040	0.039	0.040
Labor Market Controls		X						
Occ FE and time trends			X					
Unweighted				X				
Includes Missing Firms					X			
Bartik w MSA-specific Peaks/Troughs						X		
Bartik in levels							X	
1 month employment change Bartik								X

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

Notes: See Table 2 and Appendix Table B1.

Table B4: Robustness Checks: Computer Skill Requirement

<i>Dependent Variable:</i>	<i>Computer Skill Requirement</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Shock*2010</i>	0.0203** (0.00859)	0.0217*** (0.00777)	0.0212*** (0.00742)	0.0183*** (0.00605)	0.0181** (0.00784)	0.0206** (0.00816)	0.00715 (0.00753)	0.0163* (0.00871)
<i>Shock*2011</i>	0.0243*** (0.00716)	0.0257*** (0.00661)	0.0241*** (0.00647)	0.0206*** (0.00617)	0.0244*** (0.00694)	0.0248*** (0.00687)	0.0131** (0.00641)	0.0214*** (0.00722)
<i>Shock*2012</i>	0.0207** (0.00848)	0.0235*** (0.00817)	0.0205*** (0.00766)	0.0199*** (0.00622)	0.0224*** (0.00740)	0.0209*** (0.00803)	0.0112 (0.00764)	0.0196** (0.00844)
<i>Shock*2013</i>	0.0252*** (0.00664)	0.0283*** (0.00587)	0.0250*** (0.00596)	0.0250*** (0.00567)	0.0260*** (0.00649)	0.0252*** (0.00624)	0.0114* (0.00606)	0.0225*** (0.00702)
<i>Shock*2014</i>	0.0227*** (0.00679)	0.0246*** (0.00653)	0.0233*** (0.00612)	0.0234*** (0.00586)	0.0188*** (0.00709)	0.0224*** (0.00626)	0.0118* (0.00637)	0.0205*** (0.00713)
<i>Shock*2015</i>	0.0134* (0.00807)	0.0153** (0.00722)	0.0155** (0.00778)	0.0224*** (0.00497)	0.0134* (0.00703)	0.0136* (0.00759)	0.00582 (0.00780)	0.00936 (0.00808)
# Cells	178,176	178,176	178,176	135,878	207,552	178,176	178,176	178,176
R-Squared	0.034	0.049	0.292	0.008	0.055	0.035	0.033	0.034
Labor Market Controls		X						
Occ FE and time trends			X					
Unweighted				X				
Includes Missing Firms					X			
Bartik w MSA-specific Peaks/Troughs						X		
Bartik in levels							X	
1 month employment change Bartik								X

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

Notes: See Table 2 and Appendix Table B1.

Table B5: Within-Industry-Occupation Effects

<i>Dependent Variable:</i>	<i>Education</i>		<i>Experience</i>		<i>Cognitive Skill</i>		<i>Computer Skill</i>	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
<i>Shock*2010</i>	0.0514*** (0.0146)	0.0493*** (0.0144)	0.0489*** (0.0152)	0.0477*** (0.0148)	0.0248*** (0.00838)	0.0257*** (0.00811)	0.0194** (0.00928)	0.0164* (0.00859)
<i>Shock*2011</i>	0.0468*** (0.0142)	0.0454*** (0.0142)	0.0466*** (0.0150)	0.0457*** (0.0150)	0.0258*** (0.00810)	0.0263*** (0.00809)	0.0234*** (0.00797)	0.0200*** (0.00758)
<i>Shock*2012</i>	0.0214 (0.0141)	0.0199 (0.0140)	0.0274* (0.0152)	0.0259* (0.0151)	0.0163** (0.00777)	0.0170** (0.00775)	0.0213** (0.00992)	0.0180* (0.00939)
<i>Shock*2013</i>	0.0400*** (0.0135)	0.0411*** (0.0134)	0.0372*** (0.0138)	0.0377*** (0.0138)	0.0230*** (0.00702)	0.0231*** (0.00704)	0.0247*** (0.00791)	0.0218*** (0.00757)
<i>Shock*2014</i>	0.0434*** (0.0157)	0.0442*** (0.0154)	0.0466*** (0.0156)	0.0467*** (0.0154)	0.0249*** (0.00739)	0.0249*** (0.00739)	0.0231*** (0.00821)	0.0204*** (0.00784)
<i>Shock*2015</i>	0.0505*** (0.0157)	0.0559*** (0.0162)	0.0500*** (0.0156)	0.0552*** (0.0158)	0.0280*** (0.00816)	0.0293*** (0.00820)	0.0129 (0.00948)	0.0129 (0.00921)
Industry FE and trends		X		X		X		X

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

Notes: See table 2. Regressions are estimated at the msa-occupation (4-digit SOC)-industry (2-digit NAICS) year level. The dependent variable is the MSA-occupation-industry level annual change in skill requirements from 2007. All regressions control for year fixed effects and MSA characteristics from the ACS. Observations are weighted by the size of the MSA labor force in 2006 multiplied by the cell's ad share in the MSA-year. Standard errors are clustered at the MSA level.

Table B6: Decomposing Upskilling Within and Across Firms and Occupations

Share of coeff attributed to:	Education Requirements						Experience Requirements					
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
Shock*2010	0.070 (0.152)	0.055 (0.010)	0.415 (0.086)	0.042 (0.040)	0.075 (0.024)	0.343 (0.159)	-0.006 (0.126)	0.053 (0.098)	0.463 (0.149)	0.069 (0.051)	0.122 (0.045)	0.298 (0.129)
Shock*2011	0.046 (0.154)	0.046 (0.008)	0.471 (0.122)	0.062 (0.042)	0.052 (0.036)	0.323 (0.136)	-0.039 (0.134)	0.041 (0.046)	0.475 (0.133)	0.112 (0.062)	0.138 (0.048)	0.274 (0.119)
Shock*2012	-0.031 (0.342)	-0.020 (0.040)	0.613 (0.320)	0.045 (0.133)	-0.069 (0.067)	0.462 (0.273)	-0.076 (0.243)	0.009 (0.107)	0.550 (0.300)	0.184 (0.091)	0.016 (0.065)	0.318 (0.189)
Shock*2013	0.041 (0.142)	-0.014 (0.029)	0.503 (0.090)	-0.008 (0.074)	0.085 (0.066)	0.394 (0.129)	-0.018 (0.143)	-0.084 (0.066)	0.513 (0.137)	0.147 (0.108)	0.100 (0.075)	0.342 (0.119)
Shock*2014	0.085 (0.169)	-0.026 (0.031)	0.367 (0.056)	0.007 (0.061)	0.107 (0.051)	0.460 (0.186)	0.001 (0.178)	-0.075 (0.064)	0.448 (0.122)	0.104 (0.062)	0.091 (0.046)	0.431 (0.179)
Shock*2015	0.192 (0.174)	-0.069 (0.023)	0.319 (0.031)	-0.115 (0.040)	0.195 (0.044)	0.479 (0.179)	0.060 (0.167)	-0.153 (0.147)	0.426 (0.174)	0.007 (0.068)	0.203 (0.056)	0.457 (0.168)
Share of coeff attributed to:	Cognitive Skill Requirements						Computer Skill Requirements					
	(1)	(2)	(3)	(4)	(5)	(6)	(1)	(2)	(3)	(4)	(5)	(6)
Shock*2010	0.068 (0.107)	0.207 (0.037)	0.286 (0.062)	0.100 (0.079)	0.107 (0.066)	0.232 (0.099)	-0.178 (0.185)	0.033 (0.060)	0.575 (0.115)	-0.090 (0.166)	0.200 (0.159)	0.459 (0.167)
Shock*2011	0.088 (0.137)	0.147 (0.025)	0.269 (0.101)	0.131 (0.044)	0.097 (0.032)	0.268 (0.131)	-0.034 (0.167)	0.089 (0.070)	0.383 (0.109)	0.085 (0.110)	0.136 (0.089)	0.342 (0.163)
Shock*2012	0.116 (0.228)	0.145 (0.124)	-0.022 (0.174)	0.343 (0.097)	0.061 (0.019)	0.356 (0.206)	-0.004 (0.207)	0.040 (0.073)	0.289 (0.135)	0.250 (0.148)	0.072 (0.104)	0.353 (0.198)
Shock*2013	0.173 (0.169)	0.004 (0.181)	0.270 (0.216)	0.167 (0.151)	0.119 (0.096)	0.267 (0.094)	-0.035 (0.226)	-0.062 (0.040)	0.229 (0.058)	0.217 (0.216)	0.179 (0.141)	0.472 (0.195)
Shock*2014	0.301 (0.179)	0.174 (0.211)	0.008 (0.146)	0.034 (0.095)	0.104 (0.036)	0.379 (0.166)	0.062 (0.202)	-0.049 (0.048)	0.206 (0.049)	0.075 (0.181)	0.226 (0.155)	0.480 (0.197)
Shock*2015	0.244 (0.107)	0.138 (0.058)	0.150 (0.054)	0.065 (0.038)	0.128 (0.026)	0.276 (0.098)	0.091 (0.254)	-0.337 (0.087)	0.205 (0.142)	0.004 (0.262)	0.460 (0.196)	0.577 (0.120)

Notes: The table gives the share of the coefficient listed in the left column attributed to a change in the component labeled in the top row, averaged across 720 decomposition orders and the standard deviation across all orders in parentheses. Components are: (1) share of ads to continuing firm-MSAs, (2) firm distribution among continuing firms, (3) within firm-occupation-MSA skill requirement, (4) the distribution of ads across occupations, (5) the within-occupation skill requirement among continuing firms in non-continuing occupations, (6) within-occupation skill requirement among non-continuing firms.

C Tables Corresponding to Figures in the Text

Table C1: MSA-Level Outcomes and the Employment Shock

Dep Var:	Emp Growth (1)	Unemp Rate (2)	Epop College (3)	Epop High School (4)	Education (5)	Burning-Glass Skills: Experience (6)	Cognitive (7)	Computer (8)	PCs (HH) (9)	Routine-Cog (10)	Routine-Man (11)
Shock*2000	0.00698 (0.00470)	-0.00509* (0.00285)	0.0216 (0.0132)	0.00757 (0.0110)					0.525 (0.370)	-0.00806 (0.00798)	0.000161 (0.0114)
Shock*2001	-0.000397 (0.00599)	-0.000366 (0.00219)	0.0161 (0.0119)	0.00340 (0.0100)						-0.00605 (0.00477)	-0.00314 (0.0104)
Shock*2002	-0.00202 (0.00776)	0.000992 (0.00320)	-0.00504 (0.0113)	-0.00131 (0.0106)					0.469 (0.367)	-0.01108 (0.00718)	-0.000681 (0.0118)
Shock*2003	0.00252 (0.00664)	0.00101 (0.00297)	0.00371 (0.0104)	0.00375 (0.0101)					0.458 (0.367)	-0.00578 (0.00577)	-0.00501 (0.0104)
Shock*2004	0.0121* (0.00631)	0.000207 (0.00189)	-0.00414 (0.00928)	0.00507 (0.00972)						-0.00545* (0.00330)	0.000119 (0.0112)
Shock*2005	0.0177*** (0.00554)	0.000221 (0.00166)	-0.00882 (0.00853)	0.0131* (0.00678)						-0.00486 (0.00343)	-0.000883 (0.00342)
Shock*2006	0.0154*** (0.00416)	-0.00112 (0.00127)	-0.000365 (0.00620)	0.0103** (0.00507)						-0.00274 (0.00258)	0.000531 (0.00358)
Shock*2008	-0.0172*** (0.00233)	0.00597*** (0.00162)	-0.0144** (0.00681)	-0.00846 (0.00670)					1.318*** (0.451)	-0.00139 (0.00252)	-0.00521 (0.00329)
Shock*2009	-0.0383*** (0.00355)	0.0226*** (0.00320)	-0.0345*** (0.00867)	-0.0114 (0.00770)						0.00309 (0.00367)	-0.0158*** (0.00334)
Shock*2010	-0.00904** (0.00439)	0.0202*** (0.00318)	-0.0332*** (0.0110)	-0.0120 (0.00740)	0.0540** (0.0133)	0.0503*** (0.0136)	0.0280*** (0.00815)	0.0171 (0.0137)	1.800** (0.877)	0.00451 (0.00416)	-0.0206*** (0.00375)
Shock*2011	0.00488 (0.00355)	0.0140*** (0.00359)	-0.0235** (0.0114)	-0.0153* (0.00849)	0.0474*** (0.0129)	0.0463*** (0.0131)	0.0293*** (0.00746)	0.0264*** (0.00812)		0.00640* (0.00336)	-0.0210*** (0.00423)
Shock*2012	0.0108*** (0.00327)	0.00350 (0.00429)	-0.0265* (0.0137)	-0.00473 (0.00905)	0.0210 (0.0130)	0.0265** (0.0133)	0.0186** (0.00732)	0.0267** (0.0108)	1.316** (0.647)	0.00580* (0.00319)	-0.0183*** (0.00425)
Shock*2013	0.0120*** (0.00354)	0.000300 (0.00307)	-0.0169 (0.0141)	-0.00140 (0.0100)	0.0364*** (0.0117)	0.0360*** (0.0121)	0.0253*** (0.00662)	0.0281*** (0.00792)		0.00607** (0.00291)	-0.0175*** (0.00440)
Shock*2014	0.0121*** (0.00385)	-0.00102 (0.00237)	-0.00861 (0.0115)	-0.00485 (0.0109)	0.0386*** (0.0144)	0.0417*** (0.0142)	0.0277*** (0.00695)	0.0234*** (0.00794)	1.010 (1.063)	0.00745** (0.00293)	-0.0139*** (0.00467)
Shock*2015	0.00953** (0.00430)	-0.00234 (0.00223)	-0.0129 (0.0110)	-0.00757 (0.0111)	0.0416*** (0.0141)	0.0433*** (0.0144)	0.0319*** (0.00754)	0.0136* (0.00817)		0.0110*** (0.00359)	-0.0110** (0.00473)
# MSA-Year Cells	5,714	5,715	2,833	2,833	2,286	2,286	2,286	2,286	2,667	4,530	4,530
R-Squared	0.535	0.764	0.394	0.336	0.207	0.291	0.337	0.319	0.050	0.257	0.500

*** p<0.01, ** p<0.05, * p<0.1
 Notes: Columns 1-4 correspond to figure 1, columns 5-8 to figure 2, column 9 to figure 4, and columns 10-11 to the top right panel of figure 7. The dependent variable is the MSA-level annual change in the indicated variable from 2007 (2006 in column 9) - there are at most 381 distinct MSAs. All regressions control for year fixed effects (columns 5-11 also control for MSA characteristics). Observations are weighted by the size of the MSA labor force in 2006 and 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. Employment growth and unemployment rates (columns 1 and 2) obtained directly from BLS. Epop (columns 3 and 4) are author calculations based on the CPS - there are fewer MSAs available in the CPS. Skill requirements (columns 5-9) are from Burning Glass changes in skill requirements from 2007, where observations are restricted to 2010-2015. PCs per average pre-recession employment (column 9) is obtained from Harte Hanks; these regressions include only even years. Routine-cognitive and manual employment shares (columns 10 and 11, respectively) are obtained from CES. Routine-cognitive (manual) index of Acemoglu and Autor.

Table C2: Differential Upskilling by 90-10 Change in Firm Capital Investments

Dependent Variable:	PCs (HH)				Capital Holdings (Compustat)			
	Education (1)	Experience (2)	Cognitive (3)	Computer (4)	Education (1)	Experience (2)	Cognitive (3)	Computer (4)
Capital Shock*2010	0.00681*** (0.00121)	0.00389*** (0.00115)	0.00152* (0.000893)	0.00592*** (0.00105)	0.00559*** (0.000915)	0.00528*** (0.000708)	0.00153 (0.00118)	0.00650*** (0.000582)
Capital Shock*2011	0.00554*** (0.00119)	0.00241* (0.00137)	0.00113 (0.000977)	0.00515*** (0.00116)	0.00603*** (0.00131)	0.00704*** (0.000959)	0.00256* (0.00131)	0.00604*** (0.000850)
Capital Shock*2012	0.00585*** (0.000980)	0.00344*** (0.00101)	0.00118 (0.00102)	0.00552*** (0.00116)	0.00578*** (0.00126)	0.00720*** (0.000973)	0.00301** (0.00123)	0.00530*** (0.000801)
Capital Shock*2013	0.00934*** (0.00152)	0.00694*** (0.00142)	0.00218* (0.00120)	0.00497*** (0.00101)	0.00763*** (0.000714)	0.00832*** (0.000734)	0.00345*** (0.00113)	0.00643*** (0.000675)
Capital Shock*2014	0.00983*** (0.00145)	0.00702*** (0.00158)	0.00540*** (0.00163)	0.00744*** (0.00110)	0.00709*** (0.00148)	0.00774*** (0.00151)	0.00330** (0.00139)	0.00660*** (0.00100)
Capital Shock*2015	0.0106*** (0.00138)	0.00382** (0.00161)	0.00414*** (0.00157)	0.00529*** (0.00126)	0.00493*** (0.00134)	0.00617*** (0.00126)	0.00249* (0.00146)	0.00615*** (0.000725)
Shock*2010	0.0234 (0.0182)	0.0180 (0.0166)	0.00867 (0.00842)	0.0140 (0.0101)	0.0129 (0.0196)	-0.00132 (0.0169)	0.00467 (0.0103)	0.00356 (0.0110)
Shock*2011	0.0295* (0.0169)	0.0308* (0.0170)	0.00815 (0.00880)	0.0202** (0.00954)	0.0316* (0.0184)	0.0206 (0.0186)	0.00982 (0.0102)	0.0252*** (0.00898)
Shock*2012	0.0106 (0.0181)	0.00895 (0.0188)	0.00202 (0.00915)	0.0213* (0.0112)	0.0117 (0.0221)	0.00343 (0.0200)	0.00364 (0.00997)	0.0235** (0.0111)
Shock*2013	0.0191 (0.0173)	0.0129 (0.0184)	-0.00402 (0.00862)	0.0169 (0.0110)	0.0315 (0.0212)	0.0192 (0.0226)	-0.00422 (0.00887)	0.0213** (0.00961)
Shock*2014	0.0271 (0.0197)	0.0319 (0.0194)	-0.00345 (0.00977)	0.0182* (0.00997)	0.0388 (0.0238)	0.0410* (0.0235)	0.000217 (0.0114)	0.0142 (0.0107)
Shock*2015	0.0259 (0.0185)	0.0220 (0.0175)	0.0112 (0.0102)	0.0133 (0.0107)	0.0290 (0.0225)	0.0382* (0.0203)	0.00726 (0.0118)	0.0148 (0.0122)
# Firm-MSA-Year Cells	492,291	492,291	445,856	445,856	236,783	236,783	218,254	218,254
R-Squared	0.019	0.019	0.009	0.010	0.029	0.035	0.015	0.023

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

Notes: Columns 1-4 correspond to figure 5a and columns 5-8 correspond to figure 5b. The dependent variable is the firm-MSA-level annual change in the indicated variable from 2007. All regressions control for year fixed effects and MSA characteristics. Observations are weighted by the size of the MSA labor force in 2006 times the firm's ad share within the MSA-year and 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. Capital variable is the firm-level average change PCs per employee from 2002-2006 to 2010-2014 (columns 1-4) from Hartle Hanks or the firm-level average rate of change of capital holdings (property plant and equipment) over the same years (columns 5-8) from Compustat. Capital variables have been normed to the mean 0 so the main shock-year variables fit the effect for the average firm. Figures 5a and 5b fit capital-stock-year interactions to the 90-10 firm differential - therefore coefficients are multiplied by 0.97 (columns 1-4) or 2.27 (columns 5-8).

Table C3: Differential Upskilling for Routine Occupations

Dependent Variable:	Routine-Cognitive				Routine-Manual			
	Education (1)	Experience (2)	Cognitive (3)	Computer (4)	Education (1)	Experience (2)	Cognitive (3)	Computer (4)
<i>Routine *Shock*2010</i>	0.00496*** (0.00182)	0.00368*** (0.000657)	0.00532*** (0.000611)	0.0157*** (0.000683)	0.00835*** (0.00146)	0.00311** (0.00138)	-0.00580*** (0.00109)	-0.0152*** (0.000951)
<i>Routine *Shock*2011</i>	0.00721*** (0.00122)	0.00467*** (0.000685)	0.00838*** (0.000634)	0.0139*** (0.000883)	0.00382*** (0.00126)	-0.000845 (0.000993)	-0.00817*** (0.00106)	-0.0124*** (0.00102)
<i>Routine *Shock*2012</i>	0.00940*** (0.00106)	0.00275*** (0.000772)	0.00792*** (0.000551)	0.0100*** (0.000906)	0.00177 (0.00121)	-0.00264*** (0.00101)	-0.0112*** (0.00104)	-0.0145*** (0.000999)
<i>Routine *Shock*2013</i>	0.0122*** (0.000874)	0.00781*** (0.000741)	0.00820*** (0.000481)	0.0106*** (0.000715)	-0.00191** (0.000903)	-0.00367*** (0.000883)	-0.0118*** (0.00103)	-0.0117*** (0.00106)
<i>Routine *Shock*2014</i>	0.0130*** (0.00119)	0.00972*** (0.000788)	0.0101*** (0.000472)	0.0100*** (0.000707)	-0.00262** (0.00101)	-0.00159 (0.000983)	-0.0117*** (0.000960)	-0.0109*** (0.00106)
<i>Routine *Shock*2015</i>	0.00940*** (0.00147)	0.00102 (0.000731)	0.00777*** (0.000699)	0.00825*** (0.000933)	-0.00250** (0.00113)	-0.00231*** (0.000887)	-0.00843*** (0.00111)	-0.0123*** (0.00104)
<i>Shock*2010</i>	0.0514*** (0.0136)	0.0481*** (0.0134)	0.0263*** (0.00717)	0.0167** (0.00799)	0.0516*** (0.0135)	0.0486*** (0.0134)	0.0280*** (0.00728)	0.0218** (0.00858)
<i>Shock*2011</i>	0.0457*** (0.0130)	0.0431*** (0.0133)	0.0260*** (0.00716)	0.0207*** (0.00692)	0.0470*** (0.0131)	0.0444*** (0.0134)	0.0289*** (0.00735)	0.0256*** (0.00719)
<i>Shock*2012</i>	0.0204 (0.0127)	0.0245* (0.0135)	0.0162** (0.00681)	0.0176** (0.00825)	0.0230* (0.0128)	0.0257* (0.0135)	0.0202*** (0.00697)	0.0228*** (0.00851)
<i>Shock*2013</i>	0.0364*** (0.0120)	0.0340*** (0.0122)	0.0229*** (0.00637)	0.0220*** (0.00645)	0.0403*** (0.0120)	0.0368*** (0.0122)	0.0270*** (0.00642)	0.0268*** (0.00666)
<i>Shock*2014</i>	0.0393*** (0.0143)	0.0409*** (0.0140)	0.0237*** (0.00650)	0.0199*** (0.00662)	0.0433*** (0.0144)	0.0438*** (0.0140)	0.0281*** (0.00660)	0.0242*** (0.00680)
<i>Shock*2015</i>	0.0464*** (0.0143)	0.0465*** (0.0142)	0.0281*** (0.00722)	0.0113 (0.00798)	0.0492*** (0.0144)	0.0471*** (0.0142)	0.0312*** (0.00733)	0.0151* (0.00810)
# Firm-MSA-Year Cells	193,086	193,086	178,176	178,176	193,086	193,086	178,176	178,176
R-Squared	0.057	0.073	0.054	0.064	0.045	0.069	0.045	0.043

*** p<0.01, ** p<0.05, * p<0.1
Notes: Figure 6 plots the routine-shock-year interactions. The dependent variable is the occupation-MSA-level annual change in the indicated variable from 2007. All regressions control for year fixed effects and MSA characteristics. Observations are weighted by the size of the MSA labor force in 2006 times the occupation's ad share within the MSA-year and 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. "Routine" interactions in the left (right) panel are for routine-cognitive (routine-manual) occupations. These variables are an indicator for whether the occupation is in the top quartile of the routine-cognitive or routine-manual index of Acemoglu and Autor.

Table C4: Differential Employment and Wage Effects for Routine Occupations

Dependent Variable:	Routine-Cognitive		Routine-Manual	
	Layoffs	Log(Median Wages)	Layoffs	Log(Median Wages)
	(1)	(2)	(1)	(2)
<i>Routine*Shock*2000</i>	-0.000161 (0.000238)	0.00166** (0.000741)	-0.00135*** (0.000502)	0.0137*** (0.000980)
<i>Routine*Shock*2001</i>	0.000465* (0.000262)	0.00256*** (0.000682)	0.000866 (0.000579)	0.0126*** (0.000963)
<i>Routine*Shock*2002</i>	0.00147*** (0.000314)	0.00207** (0.000824)	0.00162** (0.000675)	0.0101*** (0.00112)
<i>Routine*Shock*2003</i>	0.00124*** (0.000291)	0.00138** (0.000671)	0.00232*** (0.000587)	0.00826*** (0.00119)
<i>Routine*Shock*2004</i>	0.00110*** (0.000293)	0.00289*** (0.000580)	0.000660 (0.000412)	0.00423*** (0.00105)
<i>Routine*Shock*2005</i>	0.000555*** (0.000205)	0.00332*** (0.000494)	-5.04e-06 (0.000458)	0.00322*** (0.000539)
<i>Routine*Shock*2006</i>	-4.54e-06 (0.000189)	0.00113*** (0.000237)	-0.000419 (0.000335)	0.000535 (0.000427)
<i>Routine*Shock*2008</i>	0.000638*** (0.000227)	0.000156 (0.000259)	0.00394*** (0.000422)	0.000512 (0.000328)
<i>Routine*Shock*2009</i>	0.00295*** (0.000423)	0.00136*** (0.000246)	0.0155*** (0.000739)	0.00239*** (0.000569)
<i>Routine*Shock*2010</i>	0.00306*** (0.000432)	0.000301 (0.000418)	0.0138*** (0.000672)	0.000659 (0.000631)
<i>Routine*Shock*2011</i>	0.00341*** (0.000409)	0.00107* (0.000558)	0.00957*** (0.000653)	0.000347 (0.000690)
<i>Routine*Shock*2012</i>	0.00249*** (0.000332)	0.00231*** (0.000542)	0.00662*** (0.000625)	9.11e-05 (0.000657)
<i>Routine*Shock*2013</i>	0.00182*** (0.000307)	0.00375*** (0.000566)	0.00382*** (0.000655)	-0.000514 (0.000733)
<i>Routine*Shock*2014</i>	0.00140*** (0.000248)	0.00480*** (0.000553)	0.00158*** (0.000548)	-0.000455 (0.000894)
<i>Routine*Shock*2015</i>	0.000242 (0.000244)	0.00507*** (0.000659)	-3.94e-05 (0.000429)	1.04e-05 (0.000841)
<i>Shock*2000</i>	-0.000694 (0.00223)	-0.00118 (0.0133)	-0.000130 (0.00225)	-0.00743 (0.0134)
<i>Shock*2001</i>	0.000453 (0.00258)	-0.00110 (0.0101)	0.000517 (0.00254)	-0.00632 (0.0101)
<i>Shock*2002</i>	-0.000182 (0.00250)	-8.73e-05 (0.00855)	-6.77e-05 (0.00244)	-0.00426 (0.00840)
<i>Shock*2003</i>	0.000735 (0.00277)	0.000687 (0.00912)	0.000634 (0.00274)	-0.00255 (0.00885)
<i>Shock*2004</i>	0.00101 (0.00201)	-0.00601 (0.00894)	0.00132 (0.00204)	-0.00703 (0.00882)
<i>Shock*2005</i>	-0.000844 (0.00189)	0.00363 (0.00498)	-0.000463 (0.00191)	0.00314 (0.00504)
<i>Shock*2006</i>	-0.00128 (0.00182)	-0.00137 (0.00416)	-0.000899 (0.00186)	-0.00105 (0.00423)
<i>Shock*2008</i>	0.00366** (0.00184)	-0.00757** (0.00329)	0.00302 (0.00184)	-0.00746** (0.00335)
<i>Shock*2009</i>	0.0112*** (0.00220)	-0.00505 (0.00428)	0.00829*** (0.00226)	-0.00536 (0.00438)
<i>Shock*2010</i>	0.00933*** (0.00241)	-0.00651 (0.00516)	0.00689*** (0.00244)	-0.00637 (0.00521)
<i>Shock*2011</i>	0.00364* (0.00219)	-0.00901 (0.00581)	0.00239 (0.00221)	-0.00852 (0.00582)
<i>Shock*2012</i>	-0.000811 (0.00252)	-0.0140** (0.00578)	-0.00152 (0.00251)	-0.0130** (0.00575)
<i>Shock*2013</i>	-0.00295 (0.00255)	-0.0138** (0.00620)	-0.00322 (0.00250)	-0.0122** (0.00618)
<i>Shock*2014</i>	-0.00293 (0.00183)	-0.0144** (0.00680)	-0.00276 (0.00182)	-0.0125* (0.00690)
<i>Shock*2015</i>	-0.00137 (0.00184)	-0.00763 (0.00756)	-0.00106 (0.00184)	-0.00581 (0.00777)
# Occ-MSA-Year Cells	226,191	376,897	226,191	376,897
R-Squared	0.029	0.607	0.049	0.610

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

Notes: Figure 7 (top left and bottom right panels) plots the routine-shock-year interactions. The dependent variable is the occupation-MSA-level annual change in the indicated variable from 2007. All regressions control for year fixed effects and MSA characteristics. Observations are weighted by the size of the MSA labor force in 2006 times the occupation's ad share within the MSA-year and 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. "Routine" interactions in the left (right) panel are for routine-cognitive (routine-manual) occupations. These variables are an indicator for whether the occupation is in the top quartile of the routine-cognitive or routine-manual index of Acemoglu and Autor. Layoffs are the rate of involuntary separations reported in the CPS. Log median wage is obtained from OES.