10-17-2016

Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings

Brad J. Hershbein

W.E. Upjohn Institute for Employment Research and IZA, hershbein@upjohn.org

Lisa B. Kahn

Yale University, NBER, and IZA

Upjohn Author(s) ORCID Identifier:

<https://orcid.org/0000-0002-2534-8164>

Upjohn Institute working paper ; 16-254

****Published Version****

In *American Economic Review* 108(7): 1737-1772

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the [Labor Economics Commons](#)

Citation

Hershbein, Brad J. and Lisa B. Kahn. 2016. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." Upjohn Institute Working Paper 16-254. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp16-254>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings

Upjohn Institute Working Paper 16-254

Brad Hershbein

W.E. Upjohn Institute and IZA

Email: hershbein@upjohn.org

Lisa B. Kahn

Yale University, NBER, and IZA

Email: lisa.kahn@yale.edu

March 25, 2016

Revised August 2, 2016

Revised October 17, 2016

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, and that these differences across MSAs persist through the end of 2015. The increases are prevalent within occupations, more pronounced in the non-traded sector, driven by both within-firm upskilling and substitution from older to newer firms, accompanied by increases in capital stock, and are evident in realized employment. We argue that this evidence reflects the restructuring of production toward more-skilled workers and routine-labor saving technologies, and that the Great Recession accelerated this process.

JEL Classification Codes: D22, E32, J23, J24, M51, O33

Key Words: Job polarization, job postings, RBTC, recessions, routine-biased technological change, upskilling, vacancies

Acknowledgments:

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

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

October 17, 2016

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, and that these differences across MSAs persist through the end of 2015. The increases are prevalent within occupations, more pronounced in the non-traded sector, driven by both within-firm upskilling and substitution from older to newer firms, accompanied by increases in capital stock, and are evident in realized employment. We argue that this evidence reflects the restructuring of production toward more-skilled workers and routine-labor saving technologies, 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, Jeff Clemens, David Deming, David Green, John Horton, Peter Kuhn, Fabian Lange, Steve Malliaris, Adrien Matray, Alicia Sasser Modestino, Daniel Shoag, Henry Siu, Basit Zafar, and seminar participants at the AEAs 2016, the Atlanta Fed, Brookings, Harvard, the LSE, MIT, NBER Summer Institute (Macro and Labor), Princeton, Rutgers, SOLE/EALE 2015 meetings, Trans Pacific Labor Seminar 2015, University of British Columbia, University of Chicago, University of Illinois, University of Notre Dame, University of South Florida, and the WEAI 2016 meetings. We are especially indebted to Dan Restuccia, Jake Sherman, and Bledi Taska for providing the Burning Glass data. Jing Cai provided excellent research assistance with CPS data. A previous version of this paper was titled "Is College the New High School? Evidence from Vacancy Postings."

1 Introduction

The employment shift from occupations in the middle of the skill distribution toward those in the tails is one of the most important trends in the U.S. labor market over the last 30 years. Previous research makes 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 following the recession despite recovery in aggregate output. The reasons for jobless recovery are not well understood, but a small theoretical literature points to adjustment costs as a potential mechanism, since they generate reallocation that is concentrated in downturns (Berger 2012, Koenders and Rogerson 2005, Jaimovich and Siu 2012). Such lumpy adjustment may leave a mass of displaced workers with the wrong skills for new production. Jaimovich and Siu (2015) provide suggestive evidence that countercyclical reallocation, in the form of RBTC, and jobless recovery are linked. They show that the vast majority of the declines in middle-skill employment have occurred during recessions and that, over the same time period, recovery was jobless only in these occupations. However, there is no direct evidence on how firms restructure in the face of technological change, and whether it is gradual or episodic. This paper aims to fill that gap.

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

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).

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

We then examine whether upskilling is driven by changes to production in a manner consistent with RBTC. In short, are firms changing what they do or simply changing whom they hire? 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. In contrast, upskilling may instead occur because firms temporarily and opportunistically take advantage of a slack market to try to attract workers typically found higher up on the job ladder, with no accompanying changes to production. As our data are best-suited to measure a shift in demand from middle- to high-skilled workers, rather than concomitant demand increases for low-skilled labor that also could be expected from RBTC, we concentrate our evidence on this margin. We present several analyses supporting the notion that accelerated RBTC brought on by the Great Recession drives upskilling.

First, we show that upskilling persists at nearly the same magnitude from early in the recovery through 2015. Even though most measures of local labor market conditions have converged back to their pre-recession levels by the end of our sample period, elevated skill requirements—especially higher in the occupational skill distribution where our data are most comprehensive—have not. Any purely cyclical explanation is thus insufficient. Second, we show that upskilling is concentrated in non-traded sectors that should be most sensitive to fluctuations in local consumer demand. Third, we show that upskilling is driven both by continuing firms and by new firms that did not post in 2007, consistent with both Schumpeterian cleansing and changes in adjustment costs within firms. Fourth, we show that within continuing firms, upskilling effects are not only persistent, but, among publicly traded firms in our data, those with larger increases in skill requirements also had larger increases in capital stock over the same time period, consistent with a substitution of routine-task workers with machines. Finally, we rule out that our effects are driven by changes in labor supply due to unemployment, worker quits, or formal schooling decisions that would lead firms to opportunistically hire “up”.

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

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

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

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

Third, we contribute to a growing literature exploiting data on vacancy postings. Although several studies have used aggregate vacancy data, and even vacancy microdata, from the Bureau of Labor Statistics’ Job Openings and Labor Market Turnover (JOLTS) survey (see, for example, Davis, Faberman, and Haltiwanger 2012), these data contain little infor-

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

mation 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.

Fourth, we help answer an important question about the consequences of recessions for workers. It is well known that low-skilled workers suffer worse employment and wage consequences in recessions.⁴ Evidence from past recessions shows that in downturns workers are more likely to take worse jobs, relative to their skills.⁵ Some of this may result because workers apply to a broader set of jobs, even those for which they would normally be overqualified, when job-finding rates decline. However, it is also possible that firms actively seek a more-skilled worker than they could have attracted in a tighter market, the opportunistic upskilling mentioned above.⁶

We find some evidence consistent with this “opportunistic” upskilling: for some low-skilled occupations, upskilling effects early in the recovery fade by the end of our sample period. Moreover, additionally controlling for labor availability in the form of the local unemployment rate and the college-graduate-specific unemployment rate absorbs a fraction of our estimated upskilling effects early in the recovery, but not later, and MSAs where college graduates were particularly hard hit in the recession (and likely are relatively more plentiful in the job seeking pool) see additional upskilling effects early in the recovery that also fade. However, the majority of the upskilling effects persist through 2015, several years into the recovery, and we conclude that cyclical responses to labor market tightness cannot account for our main results and for upskilling effects across most of the occupation distribution.

Our conclusions contrast with those from a pair of related papers, concurrent with our analysis, that use a different version of our data to examine both relative increases in skill requirements early in the recovery (Sasser Modestino, Shoag, and Ballance 2016a) and any subsequent decreases in skill requirements as markets improved (Sasser Modestino, Shoag, and Ballance 2016b). These papers argue that upskilling in the immediate aftermath of the Great Recession was driven entirely by firms opportunistically seeking more-skilled workers in a slack labor market. While our results qualitatively accord with the empirical findings of the first paper, our more thorough analysis represents a substantial and new contribution. We examine a longer time horizon and provide a range of analyses that yield a richer characterization of upskilling and allow us to disentangle potential mechanisms.⁷

⁴See von Wachter and Handwerker (2009); Hoynes, Miller, and Schaller (2012); and Forsythe (2016).

⁵Devereux 2003; Kahn 2010; von Wachter, Oreopoulos, and Heisz 2012; Altonji, Kahn, and Speer 2016.

⁶There is mounting evidence that firms at the bottom of the job ladder benefit from increased retention of their incumbent workforce in a downturn (Moscarini and Postel-Vinay 2012 and 2013, Kahn and McEntarfer 2015), but little is known about hiring dynamics.

⁷First, in examining data through 2015, and finding persistent upskilling throughout, we rule out their purely cyclical explanation that is based upon data from only early in the recovery. Second, we charac-

Furthermore, we disagree with the analysis and conclusion of the second paper because the results are sensitive to a set of idiosyncratic data choices.⁸

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, technological change was 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 likely plays a role in the well-noted and marked decline in male employment-to-population ratios over the past 25 years, especially since these declines have been stair-step around recessions (Moffitt 2012).⁹ 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 upskilling as a function of local labor market conditions and explores heterogeneity in effects by occupation skill level and sector tradability. Section 5 examines heterogeneity in upskilling within and between firms. Section 6 provides a discussion of our results, and section 7 concludes.

2 Data

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

terize upskilling throughout the occupation distribution and for a wider set of skills, speaking to recent developments in the demand for cognitive skill. Third, we show that upskilling is concentrated in non-traded sectors, which is consistent with our argument of episodic restructuring, and difficult to rationalize with their opportunistic upskilling argument. Fourth, our firm-level analysis helps us identify Schumpeterian cleansing from other adjustment cost models, and confirms accompanied increases in capital stock.

⁸Notably, they overstate the degree of downskilling during the later recovery by not comparing it to the overall change in skill requirements from 2007. Additionally, they use only a fraction of available data, excluding roughly 35% of their weighted ads. Their results are particularly sensitive to this restriction (based on our own analysis) and their choice to assign equal weight to all counties, regardless of population or number of ads. Furthermore, the paper’s estimates are also sensitive to the use of one unusual year in the data to demarcate the beginning of the recovery. Section 4 has additional details.

⁹Supporting 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.

2.1 Burning Glass Overview

Burning Glass Technologies (henceforth BG or Burning Glass) examines some 40,000 online job boards and company websites to aggregate the job postings, parse and deduplicate them into a systematic, machine-readable form, and create labor market analytic products. Thanks to the breadth of this coverage, BG believes the resulting database captures a near-universe of jobs that were posted online. Through a special agreement, we obtained an extract from BG, which covers every MSA in the United States in 2007 and from 2010 through 2015.¹⁰

The two key advantages of our data are its breadth and detail. The broad coverage of the database presents a substantial strength over data sets based on a single vacancy source, such as CareerBuilder.com. While the JOLTS asks a nationally representative sample of employers about vacancies they wish to fill in the near term, it is typically available only at aggregated levels, and contains relatively little information about the characteristics of vacancies. In contrast, the BG data contain some 70 possible standardized fields for each vacancy. We exploit detailed information on occupation, geography, skill requirements, and firm identifiers. The codified skills include stated education and experience requirements, as well as thousands of specific skills standardized from open text in each job posting.¹¹ The data thus allow for analysis of a key, but largely unexplored, margin of firm demand: skill requirements within occupation.¹² 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. 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

¹⁰The database unfortunately lacks postings from 2008 and 2009. Our extract was provided in February 2016. We also have data on jobs posted in Micropolitan Statistical Areas, which we do not use for lack of some of the labor market indicators in these areas, and substantial noise in the ones that are available. They represent 5.6% of all posted ads.

¹¹For 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.”

¹²Other 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.

time, and the aggregate and industry trends in the quantity of vacancies track other sources reasonably closely.¹³

The data contain the occupation of the posting (at the 6-digit Standard Occupation Classification 2010 (SOC) level) and codes identifying the MSA where it is located. Burning Glass also collects the firm name, if available, for a given posting. Employer name is missing in 40% of postings, primarily from those listed on recruiting websites that typically do not list the employer.¹⁴

We restrict our main sample to ads with non-missing employers that posted at least 10 ads over the sample period of 2007 and 2010-2015. After cleaning, our data contain 130,440 distinct employers.¹⁵ Many of our analyses exploit firm-level information to distinguish among possible mechanisms for upskilling. We therefore choose to focus our entire analysis on the consistent sample of ads with non-missing firms, with a sufficient number of observations per firm to estimate firm-level characteristics. However, we have performed analyses not requiring firm-level information on the full data set and obtain very similar results. Moreover, as we discuss later, we have confirmed that the probability of satisfying this sample criterion (having a valid firm identifier) does not vary over the business cycle. Thus, our sample restriction should not confound the estimated relationship between local labor market conditions and the skill requirements of postings.

2.2 Skill Requirements in Burning Glass

In our analyses, we exploit four categories of skill requirements: stated education and experience requirements, stated demand for skills that we classify as “cognitive,” and stated demand for computer skills. The RBTC literature emphasizes that new information technology or cheap overseas labor substitute for routine, algorithmic, middle-skill tasks. These new technologies are in turn complementary with high-skill cognitive, abstract tasks and may indirectly affect low-skill, manual tasks. 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. 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

¹³BG 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.

¹⁴When 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” (sometimes listed as “Inc”).

¹⁵The 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).

a given moment.

In accord with human capital theory, we believe more-educated workers or those with some amount of experience on the job will be better able to perform these functions.¹⁶ In Appendix 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.¹⁷ We define cognitive skill requirements based on a set of key words associated with non-routine analytical tasks, using the taxonomy developed by 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.¹⁸

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.¹⁹

Table 1 summarizes data for the primary regression sample.²⁰ In 2007, 34% of the

¹⁶In the raw data, there are two fields each for education and experience requirements: a minimum level (degree or years of experience) and a preferred level. Postings that do not list an education or experience requirement have these fields set to missing. We use the fields for the minimum levels to generate variables for the presence of an education or experience requirement as well as the number of years of education or experience required; the minimum is much more commonly specified than the preferred, and it is always available when a preferred level is listed.

¹⁷BG includes common software (e.g., Excel, PowerPoint, AutoCAD), as well as less common software and languages (e.g., Java, SQL, Python).

¹⁸Specifically, an ad is categorized as requesting a cognitive skill if any listed skills include at least one of the following phrases: “research,” “analysis,” “decision,” 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 Joe Altonji meant to categorize cognitive occupations. This check is important. For example, we find that presence of the phrase “problem solving,” which is often associated with non-routine tasks, is uncorrelated with the O*NET measure, and is instead present at high frequency in all kinds of posts. Deming and Kahn (2016) also use these measures to study firm heterogeneity in skill demands.

¹⁹It has been suggested by Deming (2015) and others that machines and overseas labor complement 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.

²⁰In 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.

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 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, more than one-sixth specify a cognitive skill requirement, and more than one-quarter 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-quarter 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, which would be consistent with our hypothesis. 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.

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)	14.66 (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	0.17 (0.05)	0.23 (0.06)	0.07***
	Computer, conditional on any	0.27 (0.06)	0.39 (0.08)	0.11***
share of ads in 2010-2015:				
Missing ACS match		0.08		
Continuing firm-MSA		0.54		
In Compustat, among continuing		0.47		
		Mean	Min	Max
# MSAs		381		
Posts per MSA-year		21,507	132	1,198,662
# Occupations (4-digit)		108		
Posts per occupation-MSA-year		225	1	189,642
# Firms		130,440		
Posts per Firm-MSA-year		14	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 observations to MSAs that cannot be matched to the American Community Survey (weighted by the MSA labor force). Continuing firm-MSAs are firm-MSAs that post in 2007 and 2010-2015 and In Compustat among continuing are the sharing of continuing firm-MSAs that can be matched to compustat. Both 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.

3 Methodology

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

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

The regression thus estimates the impact of the local employment shock generated by the Great Recession on the *change* in skill requirements for a given MSA between 2007 and each year of 2010 through 2015. Such a specification allows us to empirically investigate the timing and persistence of upskilling in relation to local labor market shocks. It also accords with several macroeconomic models that predict that productivity-enhancing improvements will be more likely to take place in recessions. In the classic Schumpeter (1939) cleansing model, this would occur because low-productivity firms shut down in the recession and resources are reallocated to new firms with more modern production.²¹ In addition, several classes of models can rationalize episodic restructuring within surviving firms as well. A negative product demand shock lowers the opportunity cost of adjusting production (Hall 2005), may shift managerial attention from growth to efficiency (Koender and Rogerson 2005), changes the costs and benefits of making layoffs (Mortensen and Pissarides 1994, Berger 2012),²² and alters incentives to invest in human capital (Jaimovich and Siu 2012).

The key explanatory variable, $shock_m$, is the MSA-specific change in projected year-over-year employment growth from 2006 to 2009, the national peak and trough years surrounding the Great Recession. We project employment growth in an MSA based on its employment shares in 3-digit NAICS industry codes in 2004 and 2005 and national employment changes at the 3-digit industry level. This type of shift-share method is sometimes referred to as a “Bartik shock,” following the strategy of Bartik (1991).²³ We use the Bartik measure,

²¹See also Caballero and Hammour (1994, 1996) and Mortensen and Pissarides (1994).

²²Though 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.

²³Specifically, we define Bartik employment growth as $\sum_{k=1}^K \phi_{m,k,\tau}(\ln E_{kt} - \ln E_{k,t-12})$, where for K 3-digit industries, ϕ is the employment share of industry k in MSA m at time τ (in practice, the average of 2004 and 2005), $\ln E_{kt}$ is the log of national employment in industry k in year-month t , and $\ln E_{k,t-12}$ is the log of national employment in industry in the same calendar month one year prior (both seasonally adjusted). We obtain national employment for each 3-digit industry from Current Employment Statistics. We construct ϕ using County Business Patterns data and the algorithm of Isserman and Westervelt (2006) to overcome data suppressions; the resulting county-level statistics are mapped to MSAs using the definitions provided by the Census Bureau and set by the Office of Management and Budget. See <http://www.census.gov/population/metro/data/def.html>. Other papers utilizing Bartik shocks include Blanchard and Katz (1992) and Notowidigdo (2013). The measure may be used directly as a regressor (reduced-form) or as an instrument for observed employment growth; in practice, this choice often does not matter much, and that is also true in our case.

instead of actual employment growth (as reported by the Bureau of Labor Statistics), for two reasons. First, actual employment growth at the MSA level is measured with substantial error, while the Bartik measure allows for more precision. Second, actual employment growth will reflect shocks to labor demand as well as other city-specific shocks, including those to labor supply, which may be problematic.²⁴ 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.²⁵

The variable $shock_m$ is fixed at the MSA-level for our entire sample period, and we examine its impact on the change in skill conditions for each year 2010–2015 through an exhaustive set of year interactions. Since we control for year fixed effects (I^t), we identify the key coefficients, α_1 , purely off of differences across metro areas in the employment shock generated by the Great Recession, rather than relying on the national shock itself. As noted above, this is a necessity given our short panel. We would have difficulty separately identifying coincident national trends, such as changes in the use of online job boards. 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.²⁶

The first-difference specification implicitly controls for differences across MSAs in posted skill requirements. Although our cross-sectional identification does not identify MSA fixed effects in the *change* in skill requirements, we do control 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 2006.²⁷ These controls help adjust for differences across MSAs in their preexisting tendency to upskill, to the extent that such a tendency is correlated with the skill distribution of the population or the health of its economy before the Great Recession.

In order to understand changes in skill requirements both within and across firms and jobs, we disaggregate our data to estimate changes in skill requirements within occupation-

²⁴For 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.

²⁵Nonetheless, our estimates are qualitatively robust to using actual employment growth or the unemployment rate, although magnitudes are somewhat reduced, likely for the reasons indicated.

²⁶Note, this normalization also implies that a larger value corresponds to a worse employment shock.

²⁷We 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 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. In such cases, we set the ACS controls to zero and include an indicator for not matching.

MSA, sector-MSA, firm-MSA, and firm-occupation-MSA cells. We categorize occupations based on 4-digit SOC codes.²⁸ These analyses are described in more detail below.

We cluster standard errors by MSA to address possible serial correlation within an area.²⁹ Finally, we weight observations in this regression by the size of the MSA labor force in 2006. This weighting scheme allows us to upweight areas with larger populations, helping with precision, while fixing the weight applied to each MSA-year. The latter ensures that we identify off of the same MSA weighting mix in each year, regardless of the overall changes in ads posted.³⁰

From the bottom panel of table 1, we estimate these regressions using all 381 MSAs, which contain an (unweighted) average of 21,507 posts per MSA-year. When we disaggregate to the four-digit occupation level, we have 108 occupations represented, with an average of 225 posts in each occupation-MSA-year. Finally, as noted above, our data contain 130,000 unique firms, which translate into an average of 14 posts in each firm-MSA-year.

4 Skill Requirements and Local Employment Conditions

4.1 Main Results

Figure 1 summarizes regression results from equation (1) for our four main dependent variables: the change in the share of ads posting any education requirement, any experience requirement, any cognitive requirement, and any computer requirement. The figures plot the estimated impact of the Bartik shock on the change in skill requirements for each year, relative to 2007 (coefficients α_1), as well as 90% confidence intervals.

Beginning with the top left panel, we find that the probability of specifying any education requirement increases by 5.4 ppts, relative to the average requirement for ads posted in the same MSA in 2007, for an MSA experiencing a large employment shock (90th percentile), compared to an MSA experiencing a small shock (10th percentile). This increase is 16% of the average requirement in 2007 and is significant at the 1% level. The effect persists at fairly similar magnitudes and significance levels for subsequent years, with a small dip in 2012. In 2015, we estimate that the probability of posting an education requirement is still 4.1 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.³¹

²⁸We choose the 4-digit occupation aggregation primarily to save on computation time, but also to ensure larger cell sizes. Results are robust to disaggregating to the 6-digit level.

²⁹We have also estimated regressions clustering at the MSA-date level, which is the level of variation underlying α_1 , and obtain substantially smaller standard errors.

³⁰When we disaggregate across occupations, sectors, and/or firms within MSAs, we weight by the product of the MSA labor force and the ad share of the occupation, sector, firm, or firm-occupation in an MSA-date. Because of our weighting scheme, the more aggregate regressions produce results identical to those using more disaggregated data when the underlying specification is the same.

³¹As we reported in footnote 8, Sasser-Modestino et al. (2016b) find evidence of downskilling in education

Figure 1: Skill Requirements and the MSA-Specific Employment Shock

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

The remaining panels of figure 1 exhibit remarkably similar patterns in both magnitudes and statistical significance. The probability of listing an experience requirement increases by 5.0 ppts (16%), and 85% of this increase remains in 2015. The probability of listing a cognitive requirement increases by 2.0 percentage points (12%), and the gap across cities widens slightly by 2015. Finally, the probability of listing a computer skill requirement also increased by roughly 2 ppts and remained elevated through 2015.

Appendix figure D1 summarizes results for additional education and experience outcomes in order to understand changes in the intensive margin for these requirements. In figure D1a, there are similar-sized effects for increases in the probability of requiring a high school diploma and a bachelor's degree. These increases offset each other, resulting in no overall change in the years of education required, conditional on posting any requirement. Also, there is no change in the propensity to require a graduate degree. In some ways this is reassuring, since many professional jobs, such as lawyers and professors, always require higher degrees; these requirements would not change with improvements in technology. Figure

and experience requirements as local labor market conditions improved between 2010 and 2014. Their use of a first-difference specification examining changes from 2010 to 2012 and from 2012 to 2014 makes their results quite sensitive to the choice of the intermediate year, and their finding of downskilling comes entirely from the small downward blip in 2012. This blip is even larger in magnitude in the subset of data they use: they exclude all but the largest six-digit-occupation-by-county cells (omitting roughly 35% of weighted postings), and this exclusion is exacerbated by their choice to weight counties equally, regardless of population or number of postings. The more careful picture generated by using all available data, as captured in figure 1, provides strong evidence that the vast majority of the increase in skill requirements in locations hit harder by the Great Recession remains through 2015.

Figure 2: Labor Market Variables and the MSA-Specific Employment Shock

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

D1b exhibits a similar pattern for experience requirements. We see increases in both low experience requirements (2 years or less) and middle requirements (3–5 years). Beyond that we see little change. Again, this results in little change overall in the distribution of requirements, conditional on posting any.

For contrast, figure 2 summarizes estimates of equation (1) for a range of local labor market statistics as dependent variables.³² The top left panel shows the unemployment rate, and our estimate implies that a hard-hit MSA (90th-percentile Bartik shock) experiences an increase in the unemployment rate between 2007 and 2010 that is 2 ppts greater than a less hard-hit MSA (10th percentile). Over time, the impact of the shock declines in magnitude, and by 2015, unemployment rate differentials across MSAs have converged back to their pre-recession levels. We find similar convergence in employment levels (top right panel).³³

The bottom left panel shows employment-to-population ratio (epop) estimates, which run only through 2014 because of data availability. We find that a hard-hit MSA experiences a 1.25 ppt larger decline in its employment-to-population ratio in 2010, and only about 15% of

³²MSA unemployment rates and employment are obtained from the BLS Local Area Unemployment Statistics and Current Employment Statistics programs, respectively. Employment-to-population ratios are based on the authors' calculations from the ACS. At the time of this writing, ACS data for 2015 had not been released.

³³That is, we find that the portion of unemployment rates and employment predicted by Bartik employment growth converged back to their pre-recession levels by 2015, holding constant pre-recession MSA characteristics. Actual differences across cities were somewhat slower to converge.

this gap had narrowed by 2014.³⁴ There could be many reasons for this lack of convergence. For example, rapid adoption of new technologies in hard-hit MSAs over this time period could render a swath of worker skills obsolete, inducing labor force exit (also see Foote and Ryan 2015). This could be especially true among older workers for whom retraining may not be worthwhile. Alternatively, it could also reflect our age controls not fully capturing retirement effects from aging baby boomers. We thus also provide estimates using the *epop* among the prime-aged population (age 25–55) in the bottom right panel. Here we see similar initial effects in 2010 but find that two-thirds of the gap across cities had converged by 2015.

Clearly, MSA labor markets had not fully recovered from the Great Recession by the end of 2015. Yet, on most measures, there was substantial progress after 2010, with cities moving closer to their pre-recession differences. At the same time, in figure 1 we see almost no convergence in skill requirements.

Appendix tables D1a and D1b provide some robustness checks to the main skills results in figure 1, including accounting for the labor market variables just shown and applying different sample restrictions. Results underlying figure 1 are shown in column 1. In column 2, we additionally control for the change in the MSA unemployment rate in each year, relative to 2007, and the change in the unemployment rate for college graduates. These controls help address an important alternative hypothesis for the upskilling effects we find: firms may become pickier when labor, and especially skilled labor, becomes more plentiful.³⁵ If our estimates are driven by opportunistic search behavior on the part of firms, they should load on the unemployment rate variables that represent time-variation in the availability of skill. We find that including these controls does slightly reduce the impact of the Bartik shock early in the recovery (by about a fifth) for education and experience requirements, but that they have no effect on coefficients later in the sample period. This is intuitive and is driven by the fact that unemployment rates have converged back to their pre-recession levels by the end of the sample period, while skill requirements have not.

Column 3 of appendix tables D1a and D1b present the baseline results using the full sample of ads, including those with missing firm names. (As mentioned above, we focus on the 60% of ads that contain firm name because this sample allows us to later distinguish among mechanisms for upskilling.) One might be concerned, for example, that ads with firm names have differentially changing skill requirements relative to ads without firm identifiers, perhaps due to firm size or prestige. However, results on the expanded sample are quite similar to those in figure 1 for all dependent variables.³⁶

Thus, there is considerable evidence that MSAs with more severe labor demand shocks

³⁴These findings are 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.

³⁵Or, as in Menzio and Shi (2011), firms become temporarily pickier in a recession because of the negative productivity shock.

³⁶We 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 D2.

Figure 3: Relative Employment by Skill and the MSA-Specific Employment Shock

Based on Census and ACS data from 2000, 2005-2014. Graphs plot coefficients on MSA shock interacted with year, and 90% CI's (dashed lines). The top three panels show the change in relative employment rates (compared to dropouts) for high school graduates, some college, and at least a BA, respectively. The bottom right panel shows the change in the average experience level of employed workers. Regressions also control for a time trend.

during the Great Recession experienced larger and more persistent increases in the skill requirements of their job postings. Of course, stated preferences for skill are not necessarily reflected in realized employment. For example, given the relatively low cost of advertising electronically and the weak labor markets in 2010 and 2011, firms might post high skill requirements to gauge the available pool of labor, without necessarily expecting to immediately fill a position. Moreover, since the job postings data are available starting only in 2007, we cannot observe whether changes in stated skill requirements were simply a continuation of a previous trend. To address both these issues, we analyze changes in the skill level of employed workers using Census and ACS data that are available for a wider range of years (2000 and 2005–2014).³⁷

Figure 3 summarizes results for four dependent variables and a similar specification to equation (1).³⁸ The first three panels measure changes in relative employment rates across education groups. For example, the top left panel measures the change in the ratio of high-school-graduate to high-school-dropout employment rates. This relative employment rate averages 1.6 across our whole sample period. We find that in a hard-hit state (90th-percentile shock), the high-school-graduate–dropout relative employment rate increases by

³⁷In these years, 66% of MSAs (roughly 85% of the population) can be identified.

³⁸We naturally do not include the ACS controls in these regressions since they are highly collinear with the employment variables on the left hand side in contemporaneous years. Also, we include a time trend, rather than year fixed effects, because we have a smaller set of MSAs to separately identify the national trends. Results with year fixed effects are similar but noisier.

5–10 percentage points after 2007, relative to a less hard-hit state (10th-percentile shock). We find similar effects for the change in relative employment of workers with some college compared to dropouts (top right), and the change for workers with at least a bachelor’s degree relative to dropouts (bottom left).

In the bottom right panel, we examine the change in the average experience level of employed workers. Experience is simply years of school minus age minus six, which is different in spirit from the stated requirement in the BG data that an applicant needs experience in the field. However, more-experienced workers are still more skilled than less-experienced workers, and the measure is as close as the data allow. We find that in hard-hit states, average experience is about a fifth of a year higher than it was in 2007, relative to less hard-hit states.³⁹

Finally, we note that data from 2000, 2005, and 2006 show that the skill composition of employed workers was similar across MSAs and unrelated to how severely the Great Recession would affect areas. This pattern is reassuring that there are also no pretrends in stated skill requirements.

We thus present strong evidence that employers in harder-hit MSAs were induced to increase both stated preferences for a range of skills that are complementary to RBTC and employment of skilled, relative to unskilled workers. Furthermore, while most measures of local labor-market strength had converged back to pre-recession levels by 2015, differences in stated and realized skill demands remained. MSAs that experienced severe shocks in the Great Recession still look different from how they appeared in 2007, and different from other cities that experienced weaker shocks.

4.2 Heterogeneity within and across Occupations

Why do skill requirements change differentially across MSAs? It is possible that these differences result from selection: perhaps in harder-hit MSAs only higher-skilled jobs survive in the early recovery. An episodic RBTC-based explanation involving changes in production technology, however, would imply upskilling within occupations, even as demand also shifts across occupations, from routine to abstract. 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 scientists, not reporters, to be chief data officers (Greenfield 2014).

To investigate whether MSA-level changes in skill requirements are prevalent within and/or across occupations, we generate counterfactual skill distributions and estimate re-

³⁹This result is consistent with Forsythe (2016), who finds that the hire rate of young workers falls, relative to that of older workers, in recessions. She provides evidence based on wage data that these changes in hiring are driven by labor demand, rather than by labor supply.

Figure 4: Within and Across Occupation Changes in Skill Requirements

gressions like equation (1), using as the dependent variable the change in a counterfactual skill distribution from the actual 2007 MSA-level skill requirement. We plot the coefficients on $shock_m * I^t$ in figure 4. The solid blue line presents results where we hold constant the occupation distribution of ads within an MSA in 2007 and allow only the within occupation-MSA skill requirement to change. The dashed red line, in contrast, holds constant the within occupation-MSA skill requirement at its 2007 level and allows the occupation distribution of ads within the MSA to vary with time.⁴⁰

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

We also examine the robustness of the within-occupation changes in skill requirements to a range of specifications, based on equation (2), which has 4-digit-occupation-MSA-year cells as the unit of observation.

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

⁴⁰Virtually all ads posted in the 2010–2015 period are in occupation-MSAs that also posted in 2007, so we can generate change variables. The 0.36% of ads that cannot be matched back are dropped from these analyses. Our formal decomposition in section 5.1 addresses this concern. We also address the fact that order in the decomposition matters (for example, we could have instead benchmarked the non-changing variable at its time t level, rather than at 2007).

The results are reported in appendix tables D1a and D1b. Column 4 summarizes the baseline specification, while column 5 adds occupation fixed effects and occupation-specific time trends.⁴¹ These controls 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. The controls also help adjust for changes in the sample driven, say, by changes in the representativeness of occupations in the BG data. However, we obtain very similar results even when including the occupation-specific controls.

Finally, our data afford the unique opportunity to measure changes in skill requirements within occupations, while the bulk of work on polarization has measured shifts in employment and wages only across occupations. That literature has successfully pinpointed the kinds of occupations that can be replaced by machines or overseas labor using information on the tasks performed by workers at the occupation level.⁴² These occupations tend to be in the middle of the skill distribution, since those jobs tend to be the more routine. Autor (2014) and Jaimovich and Siu (2015) point out that employment shifted away from middle-skill occupations in the Great Recession. Though not shown, we also find this to be true of vacancy postings. We instead focus on our ability to measure task and skill content of jobs in the BG data to ask whether RBTC also occurs *within* middle-skill occupations.

To explore heterogeneity in within-occupation changes in skill requirements induced by the Great Recession, we estimate equation (2) separately for 20 occupation skill groups. As is common in the literature (Autor 2015), we define occupation skill by wage percentiles from a fixed point in time. In our case, we use the 2000 Census to divide occupations into ventiles, based on their average weekly wage of full-time workers.⁴³

⁴¹BG generates occupation codes using an imputation that is based on the job title and other characteristics of the ad. One might be worried about the accuracy of this imputation and, though the algorithm was constant over our sample period, the amount of measurement error induced by the fixed algorithm could vary over time. Reassuringly, the baseline results are nearly identical to the MSA-level results presented in figure 1.

⁴²The 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).

⁴³We obtain weekly wages (annual wage and salary income divided by weeks worked in the last year) of full-time, full-year workers (worked at least 48 hours in the last week and usually work at least 35 hours per week) from the 2000 Census and define ventiles based on Census sample weights. We do not follow the traditional approach of using the 1980 wage distribution both because we are interested in a more modern characterization of skills and because long-term institutional and structural changes in the labor market (e.g., the declining share of unions and the growth of information technology) have made the earnings distribution in 1980 a poorer proxy for skill 30 years later. However, we do obtain qualitatively similar results if we use the 1980 wage skill distribution, or if we classify occupations based on a measure of critical-thinking and problem-solving intensity derived from O*Net and developed by Joseph Altonji.

Figure 5: Upskilling by 2000 Skill Group

In figure 5, we plot the coefficients on $shock_m * I^{2010}$ (solid blue line) and $shock_m * I^{2015}$ (dashed red line), respectively, smoothed with local linear regression. These summary measures thus depict the size of the initial upskilling effect and the size of the remaining effect at the end of the sample period. In most cases the two sets of coefficients are quite close to each other, implying that the majority of the upskilling effects are persistent within occupation categories. Almost all coefficients are statistically different from zero (except for the lowest-wage occupations in the cognitive and computer specifications) and economically meaningful. We also see a fairly similar pattern across dependent variables: upskilling at relatively higher rates at low and high modes of the skill distribution, though in most cases (education, experience, computer), the low mode shows less persistence than the higher one. Together, these results suggest upskilling is broad-based across occupations, although the effects are particularly pronounced and persistent in the upper-middle of the skill distribution.⁴⁴

Beaudry, Green, and Sand (2014, 2016) document that more-educated workers have increased their employment in lower-skilled jobs since 2000. They term this shift, along with

⁴⁴Although minimum wage increases during the Great Recession have been shown to have sizable, negative employment effects for the low-skilled (Clemens and Wither 2014, Clemens 2015), our results are unlikely to be driven by a binding minimum wage that causes firms to seek more-skilled workers as a means of lowering the effective skill price in a downturn. Occupations higher in the skill distribution, where we show upskilling effects to be strongest, are unlikely to face binding minimum wages. More general downward wage rigidities (see Shimer 2004 and Hall 2005) may yield a similar effect; however, wages of new job matches have been shown to be much more cyclical, once selection is taken into account (Martins, Solon, and Thomas 2012).

stagnating employment in cognitive occupations, the “great reversal” in the demand for cognitive skill.⁴⁵ They hypothesize that lessened demand for cognitive occupations induces college graduates to take jobs lower in the skill distribution, squeezing out less-educated workers who formerly held these jobs. In light of the evidence above, we propose that any declining demand in cognitive occupations was accompanied by an increased demand for cognitive skill *within* routine-task occupations, and this shift accelerated in the Great Recession. Even as employment has shifted from routine to cognitive occupations, the remaining routine occupations themselves are becoming less routine and more cognitive. Our work thus highlights an alternative hypothesis for why high-skilled workers are increasingly found in lower-skilled occupations: these latter occupations are becoming more skilled, and it is possible that less-skilled workers are displaced because they are unable to perform the new duties required.

4.3 Upskilling and Traded Products

We exploit cross-sectional variation in how MSAs bore the Great Recession to understand the broader impact of recessions on skill demand. However, the models from the literature discussed above—which predict productivity-enhancing improvements will be concentrated in recessions—generate an additional prediction when using local shocks: upskilling effects should be 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 measure. To the extent that upskilling is driven by shifts from dying, low-productivity firms to new firms with modern production, or because surviving firms choose to adjust when the opportunity cost of temporarily stopping production is low or when layoffs are already necessary, the phenomenon should be more pronounced in our analytic strategy among firms selling products at the local level.

To classify sectors, we adopt two well-known measures of the degree to which production is “offshorable.” The first, from Blinder and Krueger (2013), is based primarily on workers’ survey responses to questions of location requirements to do their jobs. We create an index equal to the average of Blinder and Krueger’s three preferred measures, which are available at the two-digit NAICS level.⁴⁶ The second, from Jensen and Kleitzer (2005), 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

⁴⁵See also Castex and Dechter (2014) and Deming (2015).

⁴⁶Blinder and Krueger’s (2013) three preferred measures are based on (1) “self-classification”: whether the respondents believe their duties can be performed remotely or must be done on-site; (2) “inferred”: determined from questions about location requirements of the job; and (3) “externally coded”: how expert coders categorize jobs based on typical survey questions on industry and occupation.

output for that sector is more likely traded and need not be consumed locally. From Jensen and Kleitzer, we obtain the share of employment (in the two-digit NAICS sector) that can be categorized as *least* geographically concentrated.⁴⁷

In equation 3, we estimate the change in skill requirements within (two-digit NAICS) sector s , as a function of the Bartik shock, and we allow this effect to vary with the degree to which the sector produces locally consumed products ($local_s$). We augment our base controls with sector fixed effects, to control for the main effect of *local* and other unobservables that would generate variation in skill demand growth across sectors.

$$(3) \quad skill_{smt} - skill_{sm07} = \alpha_0 + [shock_m * I^t] \alpha_1 + [local_s * shock_m * I^t] \alpha_2 + I^t + I^{sector} + X_m \beta + \varepsilon_{omt}.$$

Figure 6 uses the Blinder-Krueger composite offshorability measure and fits impacts for non-traded (10th percentile) and traded (90th percentile) sectors with solid blue and dashed maroon lines, respectively. We find that increases in skill requirements are *larger* in the non-traded sectors. This is true for experience, cognitive skills and computer skills, where the differentials are statistically significant and effects for the traded sectors are essentially zero. In contrast, for education, both traded and non-traded sectors see significant increases in skill requirements, and the differences between them are small.

For comparison, appendix figure D2 shows results using the Jensen-Kleitzer geographic concentration measure. We find that the least concentrated sectors (and therefore likely the least offshorable) see substantially larger upskilling. This again holds for experience, cognitive skill, and computer skill, with magnitudes similar to the Blinder-Krueger measure, even though the measures are based on different methodologies.

These results help confirm the hypothesized mechanism for upskilling. Other explanations, such as temporary opportunistic upskilling, do not predict specific heterogeneous impacts across sectors, and are motivated by the availability of (homogeneous) skilled labor. The fact that we see little to no differential effect for education requirements across sectors suggests that upskilling may reflect multiple factors. However, the bulk of the evidence supports the hypothesis that upskilling effects are concentrated in the least-traded sectors—the ones most sensitive to local consumer demand.

These within-sector estimates also help address a concern about the Bartik shock, which uses for identification 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 that concomitant temporary employment declines during an adjustment period. In this case our results would not indicate that firms concentrated their *adoption of existing technologies*

⁴⁷Jensen and Kleitzer (2005) measure geographic dispersion across MSAs of employment in detailed industry categories, and designate the category “least concentrated” as those industries with a Gini coefficient less than 0.1.

Figure 6: Upskilling by Sector Tradability, based on Blinder-Krueger Offshorability

in recessions but rather that the occurrence of the innovations themselves was concentrated in the recession. The fact that our results obtain even within sector alleviates this concern. Holding constant the national change in skill requirements for a sector (and therefore also any coincident innovation in the sector), we still find evidence of upskilling.⁴⁸ That is, for a given sector, when a firm is located in a harder-hit MSA, it increases skill requirements, while a firm producing in the same sector but in a less hard-hit MSA does not. Our hypothesized mechanism is that the aggregate consumer demand shock generated by the Bartik instrument reduces adjustment costs and enables firms to adopt existing technologies and/or increases the exit of low-productivity firms.

5 Heterogeneity across Firms

The evidence presented thus far shows a persistent shift in job requirements toward workers with greater education, experience, and cognitive and computer skills in MSAs more severely affected by the Great Recession. This pattern is consistent with recessions accelerating adoption of routine-biased technologies. This accelerated technological adoption could be

⁴⁸In the estimates underlying figure 6 and appendix figure D2, we still find that the sector with average offshorability experiences statistically significant upskilling. Our results hold at the more disaggregated three-digit NAICS level, the same level underlying the Bartik shock. Furthermore, we have also explored using variation in housing price shocks to generate variation in local demand and obtained qualitatively similar results.

driven by substitution from failing low-productivity firms to new high-productivity firms, as a “cleansing” model of recessions would predict (Schumpeter 1939). In addition, surviving firms may concentrate the timing of new technological adoption around recessions because of lower adjustment costs. Many classes of models in macroeconomics generate these consequences, although sparse direct evidence exists in support of these models.

In this section, we decompose the upskilling effect into changes within and across firms and occupations to better understand which mechanisms are at play. We find that both between- and within-firm adjustments are important drivers of upskilling. We then focus specifically on within-firm changes in skill requirements, a unique capability of our data, in order to test additional predictions that support the hypothesis that upskilling is driven primarily by accelerated adoption of new technologies. In our argument, once a firm upgrades its workforce, presumably with a concomitant upgrade in machines and technology or overseas infrastructure, the changes should remain. That is, if firms are changing what they do, not simply whom they hire, those changes will endure. We thus investigate firm-level changes in skill requirements and other factors of production, notably use of capital, and find that the same firms that increased their advertised skill requirements also increased their capital stocks the most, another pattern consistent with restructuring of production.

5.1 Decomposing Upskilling Within and Across Firms

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.⁴⁹

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.⁵⁰ We hope to understand the extent to which substitution across dying and newly born firms affects overall changes in skill requirements, though it is important to note that we cannot distinguish whether a firm does not post in a given period because it has no vacancies (e.g., a hiring freeze) or because it does not exist.

In equation (4), 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;⁵¹ $\frac{N_{fmt}^C}{N_{mt}^C}$, the

⁴⁹Kahn and McEntarfer (2015) show that workers matching to jobs in downturns are more likely to match to low-paying firms than high-paying firms.

⁵⁰Here 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.

⁵¹By definition, $p_{mt}^C \equiv \frac{N_{mt}^C}{N_{mt}^C + N_{mt}^{NC}}$.

distribution of ads across continuing firms in mt ; $\frac{N_{fmt}^C}{N_{fmt}^C}$, 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;⁵² 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).

$$(4) \quad skill_{mt} = p_{mt}^C \sum_{fm \in C_t} \sum_o \frac{N_{fmt}^C}{N_{mt}^C} \frac{N_{ofmt}^C}{N_{fmt}^C} * skill_{ofmt}^C + (1 - p_{mt}^C) \sum_o \frac{N_{omt}^{NC}}{N_{mt}^{NC}} * skill_{omt}^{NC}.$$

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 these components. We provide the full details of this decomposition in appendix B, and here simply summarize key results.

In figure 7, we report the fraction of the overall impact of *shock* in the given year attributed to each component.⁵³ 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 $\frac{N_{fmt}^C}{N_{mt}^C}$.⁵⁴ 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.⁵⁵ Across dependent variables and years, each of these two components contributes

⁵²Since, by definition, non-continuing firms cannot be matched across time periods, we aggregate over all non-continuing firms.

⁵³Naturally, the order in which we perform the decomposition will matter for the relative importance of the components. Here each fraction represents the average contribution across all possible decomposition orders. More details are provided in the appendix.

⁵⁴Appendix table B1 separates these two effects, shown in columns 1 and 2. We find that the share of ads to continuing firms is the more important component, especially in the later years.

⁵⁵This 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

Figure 7: 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 components, averaged across all possible decomposition orders.

to roughly half of the upskilling effect. These results are thus consistent with both the firm birth-death substitution effect à la Schumpeter and the importance of within-firm episodic restructuring.

5.2 Within-Firm Changes

The decomposition shows that roughly half of the upskilling effect in each year can be accounted for by within-continuing-firm changes in skill requirements. This dimension merits further exploration since, while many of the models that generate episodic restructuring consider firm-level responses to the changing costs and benefits of adopting new technologies, there is essentially no direct evidence on the matter. In addition to the prediction that skill requirements should increase within firms (as shown above), these models generate two additional implications that we can test in our data: (1) increases in a firm's skill requirements should endure, and (2) upskilling should be accompanied by increases in capital stock.

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

firm-MSAs that post for different occupations. The summary measure plotted in figure 7 sums these two effects, but they are shown separately in appendix table B1, columns 3 and 5. 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 8: Skill Requirements by Firm 2007-2010 Change

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

time in figure 8.⁵⁶ 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. Instead we find that upskilling late in the recovery is driven by the same firms that initially upskilled during the Great Recession. These results thus provide further support for models of episodic restructuring.

An RBTC explanation would also suggest that firms automate routine tasks with ma-

⁵⁶We 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. Column 2 of appendix table D2 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.

chines (or cheaper, outsourced labor), which complement skilled labor. If mechanization is occurring, then firms should also invest in physical capital around the time that they up-skill.⁵⁷ We investigate directly whether the firms that upskilled simultaneously increased investment. We do so by linking the publicly traded firms in our data to Compustat North America by Standard & Poors (hereafter Compustat), the most complete database of U.S. firm accounting and balance sheet data.⁵⁸ While we cannot identify IT investment specifically in the Compustat data, we can measure a firm’s overall holdings of property, plant, and equipment (PPENT). This measure of capital stock includes IT investments and, under the RBTC hypothesis, should increase faster for firms mechanizing or using more sophisticated production technology.

We match firms to Compustat using firm name and provide details of the matching procedure in Appendix A.4.⁵⁹ Figure 9 provides a general sense of the relationship between upskilling and investment. We present binned scatter plots of the change in skill requirements between 2007 and 2010 on the rate of change in capital stock (PPENT) over the same time period.⁶⁰ There is a strong and highly significant positive relationship for all the skills measures. Firms that had larger increases in skill requirements between 2007 and 2010 also had larger increases in capital stock over the same time period.⁶¹

We also provide results from regression analysis, analogous to those presented in figure 1, comparing the impact of the Bartik shock across firms with different capital investment behavior. We estimate regressions similar to that in equation (1) but control for the firm’s change in capital stock between 2007 and 2010, and allow this variable to interact with the *shock*-year effects.⁶² Figure 10 fits the differential impact (and 90% confidence interval) of the Bartik shock for firms at the 90th and 10th percentiles of capital change. The former had a 60% increase in their capital stock while the latter experienced a 20% drop.

We find that firms with larger increases in capital stock had greater upskilling, and these differences persist. In harder-hit MSAs, skill requirements increase roughly 30% (education

⁵⁷This may seem at odds with the stylized facts on the cyclicity of investment. However, Jaimovich and Siu (2012), show that investment in information processing equipment and software bounced back immediately following the end of the NBER-dated recession (December 2007 through June 2009), while domestic private investment was much more sluggish. IT and computers are exactly the technologies that drive RBTC (Michaels, Natraj, and Van Reenan 2014).

⁵⁸We obtain these data via Wharton Research Data Services.

⁵⁹We can link 46% of weighted observations from the 2007–2010 change sample used in figure 8 to a business in Compustat. In column 3 of appendix table D2, we show that the change in the share of firms matching to Compustat does not vary with the Bartik shock.

⁶⁰We aggregate the firms into percentiles based on the magnitude of changes in skill requirements (weighted by the average number of ads posted in 2007 and 2010); we then plot for each percentile bin the ad-weighted average of this variable against the ad-weighted average of the change in capital stock.

⁶¹The figure would look similar if instead of actual firm upskilling we used predicted firm-specific upskilling measures based on variation in the Bartik shock.

⁶²Regressions are estimated at the firm-MSA-year level for the Compustat matched sample. Capital change varies at the firm level but we disaggregate to firm-MSA to take advantage of variation in local labor market conditions. These can be thought of as estimating the relationship between upskilling and capital change as a function of the Bartik shock in locations where the firm tends to post.

Figure 9: Firm-Level Changes in Skill Requirements and Capital Stock 2007–2010

Binned scatter of firm-level rate of change in capital stock and change in skill requirement from 2007 to 2010 for Compustat-matched sample. Scatter (but not lines) omits 2 outliers with > 2 capital stock change.

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

Graphs plots the coefficients on the Bartik shock*year interacted with the firm-level capital change, fitted to the 90-10 capital change differential, and 90% CI's (dashed lines). Regressions also control for year fixed effects and MSA characteristics.

and experience requirements) to about 50% (cognitive and computer requirements) more in high-investment firms than in low-investment firms. These differences are all significant, typically at the 1% level. Thus, throughout our sample period, firms with larger increases in capital stock around the time of the Great Recession also had larger increases in their posted skill requirements.

6 Discussion

Thus far, we have shown that job postings in harder-hit MSAs experienced larger increases in their education, experience, cognitive, and computer requirements following the Great Recession. Furthermore, these 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. These increased skill requirements are robust to controls for the local supply of skill availability and occur broadly within—and not just across—a wide range of occupations, and are concentrated in sectors that are more susceptible to local demand shocks and firms that also increased their capital stock. Furthermore, as best as we can measure, the changes in stated skill preferences translate to actual changes in realized employment.

We believe the most likely explanation for these facts is that the recession accelerated adoption of existing technologies. Demand for skill changed because the means of production changed, and this process was episodic, catalyzed by the Great Recession. Theoretical models that can generate this kind of “lumpy” adjustment posit several possible mechanisms: substitution across old and new firms (Schumpeter 1939, Caballero and Hammour 1994, 1996), or greater efficiency to within-firm adjustments when the opportunity cost of adjusting is lower (Hall 2005), when managerial attention is more focused on efficiency (Koender and Rogerson 2005), or when layoffs are necessary, anyway (Mortensen and Pissarides 1994, Berger 2012). Using specific empirical skill measures that reflect what the discipline has learned about technological change over the past 20 years (Autor, Levy, and Murnane 2003; Brynjolfsson and McAfee 2011), we find evidence consistent with these mechanisms: about half of the upskilling effect is driven by substitution across old and new firms and half is due to within-firm changes. Within continuing firms, changes in skill requirements are persistent and accompanied by changes in capital stock, consistent with increased mechanization. Finally, upskilling is more pronounced in sectors more sensitive to local consumer demand, those for which restructuring costs are likely to be lower in downturns.

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 try to signal that certain (unwanted) applicants need not apply. However, this would not explain the accompanying changes in employment that we find, nor would it generate persistent increases in skill requirements independent of

the available supply of labor (as proxied by the unemployment rate).

It could instead be the case that firms that cannot attract (or afford) more-skilled workers in a tight labor market opportunistically seek them out in a slack one.⁶³ This mechanism would again predict short-lived upskilling effects that would be driven by the local availability of skilled labor. We do see some evidence consistent with opportunistic upskilling, though it cannot be the primary driver of our results. As we noted, controls for unemployment rates in our regressions do affect estimated coefficients on the Bartik shock early in the recovery (when unemployment is high) but have no effect on estimated coefficients in later years (when unemployment has recovered but skill requirements have not). Commensurately, holding constant the overall size of the Bartik shock on an MSA, a larger increase in the college graduate unemployment rate between 2007 and 2010 generates greater increases in skill requirements early in the recovery, but these effects dissipate.⁶⁴ Furthermore, there is no reason to think that opportunistic upskilling should be concentrated in the non-tradeable sector or accompanied by increases in capital stock.

Indeed, the concentration and persistence of the elevated skill demand provides substantial evidence in favor of our hypothesis, since most alternative explanations predict convergence when the economy recovers. While it is possible that labor markets are slower to recover than unemployment rates and employment levels suggest, it is telling that we find little overall convergence in skill requirements as markets improve, even though the data *do* indicate convergence in some occupations (primarily on the low end of the skill distribution; see figure 5).

In our view, the only remaining alternative hypothesis is that changes in labor supply later in the recovery induce changes in skill demand. For example, availability of skilled labor might increase later in the recovery due to differential quit behavior or changes in educational attainment brought on by the Great Recession. In either case changes in the availability of skilled labor would be correlated with the Bartik shock, but might be uncorrelated with unemployment rates.

First, we point out that this kind of explanation is hard to reconcile with the result that upskilling later in the recovery is driven by the same firms that upskilled early in the recovery (figure 8). If early upskilling is driven by bottlenecks and overall unemployment, while later upskilling is driven by quits and a more-educated workforce, then there is no particular reason why the two should be linked. Second, we have directly examined the timing of quits and educational attainment as a function of the Bartik shock. We conclude that these effects cannot be important drivers of our result. See appendix C for details.

Simply put, the evidence supports that changes in skill requirements reflect changes in

⁶³Even though it is well known that recessions disproportionately affect low-skilled workers (Hoynes, Miller, and Schaller 2012), making high-skilled workers *relatively* more costly and less available, it is still possible that high-skilled workers sufficiently reduce their reservation wage to become worthwhile hires to firms that generally employ the less skilled.

⁶⁴See appendix C.

production, not simply changes in whom firms ask for.

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. We present a range of evidence that consistently supports our hypothesis that this effect is primarily driven by an episodic restructuring on the part of firms and thus is likely to be long-lasting.

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

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

The U.S. economy has seen remarkable changes over the 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 2015; 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] Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer (2016), “Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success,” *Journal of Labor Economics*, 34(S1 Part 2): S361–S401.
- [2] Autor, David (2014), “Polanyi’s Paradox and the Shape of Employment Growth,” NBER working paper #20485.
- [3] Autor, David (2015), “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29(3):3–30.
- [4] Autor, David, and David Dorn (2013), “The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market,” *American Economic Review*, 103(5): 1553–1597.
- [5] Autor, David, David Dorn, and Gordon Hanson (2015), “The Labor Market and the Marriage Market: How Adverse Employment Shocks Affect Marriage, Fertility, and Children’s Living Circumstances,” MIT mimeo.
- [6] Autor, David, David Dorn, Gordon Hanson, and Jae Song (2014), “Trade Adjustment: Worker-Level Evidence,” *Quarterly Journal of Economics*, 129(4): 1799–1860.
- [7] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006), “The Polarization of the U.S. Labor Market,” NBER Working Paper #11986.
- [8] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2008), “Trends in U.S. Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics*, 90(2): 300–323.

- [9] Autor, David H., Frank Levy, and Richard J. Murnane (2003), “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 118(4): 1279–1333.
- [10] Barr, Andrew, and Sarah Turner (2015) “Out of Work and into School: Labor Market Policies and College Enrollment During the Great Recession,” *Journal of Public Economics*, 124: 63–73.
- [11] Bartik, Timothy J. (1991), *Who Benefits from State and Local Economic Development Policies*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- [12] Beaudry, Paul, David A. Green, and Benjamin M. Sand (2014), “The Declining Fortunes of the Young since 2000,” *American Economic Review*, 104(5): 381–386.
- [13] Beaudry, Paul, David A. Green, and Benjamin M. Sand (2016), “The Great Reversal in the Demand for Skill and Cognitive Tasks,” *Journal of Labor Economics*, 34(S1 Part 2): S199–S247.
- [14] Blanchard, Olivier J., and Lawrence F. Katz (1992), “Regional Evolutions,” *Brooking Papers on Economic Activity*, 1–75.
- [15] Berger, David (2012), “Countercyclical Restructuring and Jobless Recoveries,” mimeo.
- [16] 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.
- [17] Burning Glass Technologies (2014), “Moving the Goalposts: How Demand for a Bachelor’s Degree is Reshaping the Workforce,” mimeo.
- [18] Brynjolfsson, Erik, and Andrew McAfee (2011), *Race Against the Machine*. Lexington: Digital Frontier Press.
- [19] Caballero, Ricardo, and Mohamad Hammour (1994), “The Cleansing Effect of Recessions,” *American Economic Review*, 84(5): 1350–1368.
- [20] Caballero, Ricardo, and Mohamad Hammour (1996), “On the Timing and Efficiency of Creative Destruction,” *Quarterly Journal of Economics*, 111(3): 805–852.
- [21] Card, David, Ana R. Cardoso, and Patrick Kline (forthcoming), “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*.
- [22] 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.

- [23] 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.
- [24] 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.
- [25] 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.
- [26] Charles, Kerwin K., Erik Hurst, and Matthew J. Notowidigdo (2015), “Housing Booms and Busts, Labor Market Opportunities, and College Attendance,” mimeo.
- [27] Clemens, Jeffrey (2015), “The Minimum Wage and the Great Recession: Evidence from the Current Population Survey,” NBER working paper #21830.
- [28] 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.
- [29] 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.
- [30] Deming, David J. (2015), “The Growing Importance of Social Skills in the Labor Market,” NBER Working Paper #21473.
- [31] Deming, David J., and Lisa B. Kahn (2016), “Firm Heterogeneity in Skill Demands,” mimeo.
- [32] Devereux, Paul (2002), “Occupational Upgrading and the Business Cycle,” *Labour Economics*, 16(3): 423–452.
- [33] Foote, Christopher L., and Richard W. Ryan (2015), “Labor Market Polarization Over the Business Cycle,” *NBER Macroeconomics Annual*, 29: 371–413.
- [34] “Forsythe, Eliza (2016), “Why Don’t Firms Hire Young Workers in Recessions?” mimeo.
- [35] 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.
- [36] Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (2001), “Equilibrium Unemployment,” *Journal of Monetary Economics*, 48(1): 109–152.

- [37] Greenfield, Rebecca (2014), “Why the New York Times Hired a Biology Researcher as Its Chief Data Officer,” *Fast Company*, February 12.
- [38] Hall, Robert (1991), “Labor Demand, Labor Supply and Employment Volatility,” in NBER Macroeconomics Annual. Cambridge: MIT Press.
- [39] Hall, Robert (2005), “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1): 50–65.
- [40] Hoynes, Hilary, Douglas L. Miller and Jessamyn Schaller (2012), “Who Suffers During Recessions?,” *Journal of Economic Perspectives*, 26(3): 27–48.
- [41] 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.
- [42] 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>.
- [43] 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>.
- [44] 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.
- [45] Kahn, Lisa B. (2010), “The Long-Term Consequences of Graduating from College in a Recession,” *Labour Economics*, 17(2): 303–316.
- [46] Kahn, Lisa B., and Erika McEntarfer (2015), “Employment Cyclicalities and Firm Quality,” NBER working paper #20698.
- [47] 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.
- [48] Kuhn, Peter, and Kailing Shen. (2013), “Gender Discrimination in Job Ads: Evidence from China,” *Quarterly Journal of Economics*, 128(1): 287–336.
- [49] Marinescu, Ioana (2014), “The General Equilibrium Effects of Unemployment Insurance: Evidence from a Large Online Job Board,” mimeo.
- [50] 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.

- [51] McCoy, Terrence (2016), “The Technology that Could Revolutionize the War on Hunger,” *Washington Post*, June 16.
- [52] Menzio, Guido, and Shouyong Shi (2011), “Efficient Search on the Job and the Business Cycle,” *Journal of Political Economy* 119(3): 468–510.
- [53] 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.
- [54] Moffitt, Robert (2012), “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations,” *Brooking Papers on Economic Activity*, fall.
- [55] Mortensen, Dale, and Christopher Pissarides (1994), “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61(3): 397–415.
- [56] Moscarini, Giuseppe, and Fabien Postel-Vinay (2013), “Stochastic Search Equilibrium,” *Review of Economic Studies*, 80(4): 1545–1581.
- [57] Moscarini, Giuseppe, and Fabien Postel-Vinay (2012), “The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment,” *American Economic Review*, 102(6): 2509–2539.
- [58] 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
- [59] Notowidigdo, Matthew J. (2013), “The Incidence of Local Labor Demand Shocks,” University of Chicago, mimeo.
- [60] Oreopoulos, Phil, Till von Wachter and Andrew Heisz (2012), “The Short- and Long-Term Career Effects of Graduating from College in a Recession,” *American Economic Journal: Applied Economics*, 4(1): 1–29.
- [61] Pierce, Justin R., and Peter K. Schott (2015), “Trade Liberalization and Mortality: Evidence from U.S. Counties,” Yale mimeo.
- [62] Rothwell, Jonathan (2014), “Still Searching: Job Vacancies and STEM Skills,” Brookings mimeo.
- [63] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2016a), “Upskilling: Do Employers Demand Greater Skill when Skilled Workers are Plentiful?”, mimeo.
- [64] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2016b), “Downskilling: Changes in Employer Skill Requirements Over the Business Cycle”, mimeo.

- [65] Shimer, Robert. (2004), “The Consequences of Rigid Wages in Search Models,” *Journal of the European Economic Association*, 2(2-3): 469–479.
- [66] Schumpeter, Joseph. (1939), *Business Cycles*. New York: McGraw-Hill.
- [67] U.S. Department of Labor (1977), *Dictionary of Occupational Titles, 4th edition*. Washington, D.C.: U.S. Department of Labor.
- [68] 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.
- [69] von Wachter, Till, Philip Oreopoulos, and Andrew Heisz (2012), “The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates,” *American Economic Journal: Applied Economics*, 4(1): 1–29.
- [70] Yagan, Danny (2016), “The Enduring Employment Impact of Your Great Recession Location,” mimeo.

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).⁶⁵ Several other private-sector firms also began to track online job postings in the 2000s by using web-crawling and data-scraping methods. In this study, we employ data from one such firm, Burning Glass Technologies. This appendix discusses the representativeness of the data and investigates whether representativeness has changed over the time period of analysis.

A.1 Occupation-Industry Composition in BG

The BG database covers only vacancies posted on the Internet, as opposed to JOLTS or state vacancy reports that directly survey a representative sample of employers. To the extent that vacancies from certain industries and occupations are less likely to be posted electronically, as might be the case for many less-skilled jobs, they will be underrepresented in the data.⁶⁶ 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.⁶⁷

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.⁶⁸ Despite the sample differences, the industry distributions

⁶⁵See <https://www.conference-board.org/data/helpwantedonline.cfm>.

⁶⁶Rothwell (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.

⁶⁷Furthermore, 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.

⁶⁸Both data sets cover 2007 and 2010–2014 (the most recent full years available from JOLTS at the time of this writing). 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 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

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

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 (green bars).⁶⁹ 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, variety of means, of which posting a job ad (electronic or otherwise) is only one.

⁶⁹For clarity, we use 2-digit Standard Occupational Classification codes in the figure. The regression analyses use more granular codings.

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

diagonal lines).⁷⁰

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

⁷⁰All data sets cover 2007 and 2010–2014. (2014 is the most recent year for which OES data were available at the time of this writing.) 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.

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

period. On the x-axis we plot the deviation of the BG occupation share in 2007 from that 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. We have also investigated, to the best of our ability, the degree to which these stated preferences translate into realized hires.

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 Compustat Sample

Beginning with the 10,436 firms used to construct the 2007–2010 change sample (7.4% of firms and 66% of weighted observations, used in figure 8), we employ a sequential matching procedure. We first match based on exact name, after cleaning to remove punctuation and words that are sometimes abbreviated (e.g., “Inc.”). This step accounts for 81% of all Compustat-matched firms and 76% of matches based on our sample weights. 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 (5% of matched firms and matched weighted observations). Third, we add the sample of firms matched by Deming and Kahn (2016), which uses only BG firms posting in 2014. We thus match a total of 46% of weighted observations from the 2007–2010 change sample, or 30% of weighted observations in the sample as a whole—though we did not attempt to match observations that did not meet the 2007–2010 change criteria.⁷¹

⁷¹For context, the size of employment in Compustat is roughly half that of total employment in the U.S. For example, in 2014, the sum of employment listed in companies in Compustat was 70,505,000 and total payroll employment averaged 139,042,000. The Compustat employment figure includes both domestic and foreign workers, with no way to distinguish between the two. However, the employment comparison provides a useful benchmark.

Figure A4: Comparison of BG Education Requirements and ACS Employment

(a) by MSA

Average Education of Employed in MSAs (ACS)
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

Average Education of Employed in Occupation (ACS)
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).

B Firm-Occupation Decomposition

The decomposition explored in section 5.1 begins with equation (4), repeated here for convenience:

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

Using this equation, we decompose the change in skill requirements in MSA, m from 2007 to year t , into: 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 (5).

$$(5) \quad skill_{mt} = p_{mt}^C * \pi_1^c \sum_{fm \in C_t} \sum_{o \in CO^1} skill_{ofmt}^{CO^1} * \frac{N_{ofmt}^{CO^1}}{N_{fmt}^{CO^1}} \frac{N_{fmt}^{CO^1}}{N_{mt}^{CO^1}} + p_{mt}^C (\pi_2^c \sum_{o \in CO^2} skill_{omt}^{CO^2} * \frac{N_{omt}^{CO^2}}{N_{mt}^{CO^1}} + (1 - \pi_1^c - \pi_2^c) \bar{skill}_{mt}^{CO^3}) + (1 - p_{mt}^C) (\pi_1^{nc} \sum_{o \in NCO^1} \bar{skill}_{omt}^{NCO^1} * \frac{N_{omt}^{NCO^1}}{N_{mt}^{NCO^1}} + (1 - \pi_1^{nc}) \bar{skill}_{mt}^{NCO^2})$$

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).⁷² 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 defined by the within *ofm* average skill requirement ($skill_{ofmt}^{CO^1}$), the share of ads in this occupation, *o*, for the given *f**m* ($\frac{N_{ofmt}^{CO^1}}{N_{fmt}^{CO^1}}$), 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.⁷³ $skill_{omt}^{CO^2}$ 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 (5), 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.

As noted in the text, we 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 do this by generating counterfactual differences that allow one component to change from its level in

⁷²These weights vary by MSA-year, but subscripts are suppressed for clarity.

⁷³We 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*.

2007 to its level in t , holding all other components fixed at the level in either period.⁷⁴ 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.

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 (5) 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 B1. 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 $skill_{mt}^{CO3}$; and (6) the skill requirement among non-continuing firms, which combines $skill_{omt}^{NCO1}$ and $skill_{mt}^{NCO2}$. These differ from the six components listed in the main text in that (4) here combines all occupation-distribution elements, and we specifically isolate the change in skill requirements driven by continuing firms that move into and out of occupations (5).

⁷⁴For example, $p_{mt}^C(\pi_1^c \sum_{fm \in C} \sum_{to \in CO1} (skill_{ofmt}^{CO1} - skill_{ofm07}^{CO1}) * \frac{N_{ofmt}^{CO1}}{N_{fmt}^{CO1}} \frac{N_{fmt}^{CO1}}{N_{mt}^{CO1}})$ 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 .

Table B1: 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	0.074 (0.146)	0.055 (0.010)	0.419 (0.082)	0.042 (0.042)	0.077 (0.025)	0.333 (0.152)	-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	0.048 (0.149)	0.045 (0.008)	0.474 (0.118)	0.064 (0.045)	0.053 (0.038)	0.317 (0.131)	-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	-0.011 (0.319)	-0.021 (0.037)	0.599 (0.300)	0.053 (0.123)	-0.061 (0.067)	0.442 (0.255)	-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	0.044 (0.136)	-0.014 (0.027)	0.504 (0.087)	-0.010 (0.076)	0.088 (0.069)	0.388 (0.123)	-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)	-0.024 (0.033)	0.372 (0.059)	0.009 (0.062)	0.106 (0.049)	0.452 (0.187)	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 (0.171)	-0.069 (0.025)	0.318 (0.032)	-0.121 (0.042)	0.198 (0.044)	0.481 (0.178)	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 Alternative Explanations for Upskilling

We explore three possibilities for why firms may desire to hire more-skilled workers without otherwise adjusting their production technology: (1) firms temporarily and opportunistically seek higher-skilled workers in a slack labor market, (2) replacement hiring differentially requires more-skilled workers due to differential recoveries in quit rates (i.e., the job ladder needs to clear), and (3) population changes in educational attainment alter the availability of skill.

C.1 Opportunistic Upskilling

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

We augment equation (1) with controls for the change in the unemployment rate during the Great Recession among college graduates age 25 or older, obtained from the ACS. As with the Bartik employment shock, we interact the college unemployment rate shock with year dummies to allow it to affect changing skill requirements each year. We also norm this variable so that a one unit change equals the 90–10 percentile differential across MSAs. We plot both estimates in figure C1.

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

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

This set of results has two important implications. First, the impact of the shock to

Figure C1: Upskilling and the College Unemployment Rate

college graduates does not persist but fades to pre-recession levels, suggesting that some opportunistic upskilling may have occurred early in the recovery (Sasser Modestino, Shoag, and Ballance 2016a, 2016b). Second, and more important, the overall employment shock induces persistent upskilling, even holding constant the size of the shock to college graduates, implying that opportunistic upskilling does not drive our main results. The exception is for computer skills, where it may be difficult to separately identify the overall employment shock from the college graduate shock, since computer skill increases are likely heavily concentrated in the most-educated MSAs.⁷⁵

C.2 Quits and Replacement Hiring

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

⁷⁵We have also examined firm-level upskilling by the initial skill requirements of the firm (instead of the initial *change* in skill requirements, as above. If opportunistic upskilling is an important part of the story, we would expect to see increases in skill requirements concentrated among firms with low or middling initial levels; however, we find no evidence in support of this hypothesis.

Figure C2: Quit Rates by Date and Education Group

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

lower-skilled workers recover.

However, we see no evidence of differential recovery in quit rates across skill groups. For instance, in figure C2 we plot smoothed time series of quit rates by education group from 1998–2015 using longitudinally-linked CPS data.⁷⁶ Besides the large and well-known secular decline in quits (Molloy, Smith, and Wozniak 2013), the figure shows the quit rate has a clear cyclical component, declining during NBER-dated recessions (indicated with vertical dashed lines) and recovering somewhat shortly thereafter. There is certainly no evidence in this figure that quit rates for higher-skilled workers (lighter lines) recover more quickly or become more pronounced later in the recovery than those for lower-skilled workers (darker lines); in fact, less-educated workers show faster recoveries. It thus seems unlikely that replacement hiring is driving the persistence in upskilling.

C.3 Educational Attainment

Firms may decide to upskill if they observe that skilled workers have become more plentiful in their local economy. Indeed, a long line of research has explored whether educational attainment responds to local labor market conditions. Past evidence generally finds mod-

⁷⁶We define a quit as a worker who reported switching employers between month t and month $t + 1$ (but was employed in both months), or one who reported being employed in month t and unemployed in month $t + 1$, and gave voluntary job leaving as a reason for the unemployment. To obtain quit rates we divide the count of quits by the total number of employed workers in month t .

erate effects on enrollment, but limited effects on attainment.⁷⁷ One exception is Charles, Hurst, and Notowidigdo (2015), who find that during the 2000s, the housing boom reduced educational attainment, primarily on the two-year college margin. This effect (and any symmetric rise in attainment when the housing bubble crashed) is less relevant for our sample of primarily higher-skilled jobs. In particular, the housing boom was a shock principally affecting low-skilled workers, while the Great Recession was broad-based.⁷⁸

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

D Supplemental Figures and Tables

⁷⁷Card and Lemieux (2001) find that local unemployment rates have small, positive impacts on high school attendance and completion, marginal impacts on college attendance, and no impact on college completion. Barr and Turner (2015), using more recent data, find that college enrollment has grown more responsive to the business cycle over time. Kahn (2010) finds that among cohorts graduating from college between 1979 and 1989, those graduating in a worse economy obtained an additional year of graduate school, on average. She also finds that economic conditions at time of high school graduation did not affect college completion. Altonji, Kahn, and Speer (2016) find similar modest effects over a broader time period.

⁷⁸For example, Altonji, Kahn and Speer (2016) show that the Great Recession affected recent college graduates more severely than past recessions had, and that higher-earning college majors lost much of their previous advantage in weathering worse labor market entry conditions.

Figure D1: Impact of MSA-Specific Employment Shock on Education Requirements

(a) Education Requirements

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

(b) Experience Requirements

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

Figure D2: Upskilling by Sector Tradability, based on Jensen-Kleitzer Offshorability

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.

Table D1a: Education and Experience Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
	Education Requirements					Experience Requirements				
Shock*2010	0.0543*** (0.0134)	0.0462*** (0.0150)	0.0454*** (0.0133)	0.0528*** (0.0137)	0.0518*** (0.0129)	0.0505*** (0.0136)	0.0430*** (0.0157)	0.0457*** (0.0147)	0.0494*** (0.0135)	0.0489*** (0.0130)
Shock*2011	0.0477*** (0.0129)	0.0416*** (0.0137)	0.0382*** (0.0124)	0.0478*** (0.0131)	0.0466*** (0.0126)	0.0461*** (0.0131)	0.0403*** (0.0139)	0.0361*** (0.0134)	0.0445*** (0.0134)	0.0437*** (0.0133)
Shock*2012	0.0218* (0.0130)	0.0194 (0.0130)	0.0202 (0.0128)	0.0241* (0.0128)	0.0234* (0.0124)	0.0268** (0.0133)	0.0242* (0.0134)	0.0204 (0.0138)	0.0260* (0.0136)	0.0250* (0.0134)
Shock*2013	0.0365*** (0.0117)	0.0376*** (0.0118)	0.0332*** (0.0127)	0.0403*** (0.0121)	0.0408*** (0.0114)	0.0359*** (0.0122)	0.0369*** (0.0123)	0.0311** (0.0131)	0.0364*** (0.0122)	0.0358*** (0.0121)
Bartik*2014	0.0385*** (0.0145)	0.0392*** (0.0147)	0.0304** (0.0153)	0.0430*** (0.0144)	0.0441*** (0.0137)	0.0411*** (0.0143)	0.0417*** (0.0144)	0.0301** (0.0149)	0.0435*** (0.0140)	0.0434*** (0.0137)
Shock*2015	0.0413*** (0.0141)	0.0436*** (0.0140)	0.0298** (0.0150)	0.0488*** (0.0143)	0.0532*** (0.0140)	0.0426*** (0.0144)	0.0448*** (0.0138)	0.0313** (0.0158)	0.0465*** (0.0141)	0.0478*** (0.0139)
# Cells	2,286	2,286	2,286	192,842	192,842	2,286	2,286	2,286	192,842	192,842
R-Squared	0.209	0.216	0.247	0.045	0.375	0.292	0.299	0.306	0.069	0.354
MSA-Level	X	X	X			X	X	X		
Unemployment Controls		X					X			
Includes missing firms			X					X		
Occupation-MSA Level				X	X				X	X
Occ FE and time trend					X					X

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

Notes: Regressions are estimated at the msa-year level (columns 1-3) or msa-occupation (4-digit SOC) year level (columns 4-5). The dependent variable is the MSA (1-3) or MSA-occupation (4-5) level annual change in skill requirements from 2007. All regressions control for year fixed effects and MSA characteristics from the ACS. Unemployment controls (column 2) include the annual change (from 2007) in the MSA overall and college-specific unemployment rates. Column 3 includes ads with no firm listed. Column 5 additionally includes occupation fixed effects and occupation-specific time trends. Observations are weighted by the size of the MSA labor force in 2006 (multiplied by the occupation's ad share in the MSA-year in columns 4-5). Standard errors are clustered at the MSA level. Shock is the change in projected year-over-year employment growth in the MSA from 2006 to 2009, divided by the 90-10 differential in the variable across all MSAs.

Table D1b: Cognitive and Computer Skill Robustness Checks

	Cognitive Requirements					Computer Requirements				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
Shock*2010	0.0201*** (0.00688)	0.0202*** (0.00704)	0.0143** (0.00718)	0.0207*** (0.00631)	0.0207*** (0.00567)	0.0168 (0.0137)	0.0122 (0.0133)	0.0161 (0.0112)	0.0203** (0.00862)	0.0212*** (0.00744)
Shock*2011	0.0210*** (0.00640)	0.0208*** (0.00665)	0.0165** (0.00682)	0.0207*** (0.00631)	0.0206*** (0.00578)	0.0261*** (0.00816)	0.0227*** (0.00865)	0.0263*** (0.00907)	0.0241*** (0.00718)	0.0239*** (0.00648)
Shock*2012	0.0128** (0.00638)	0.0120* (0.00636)	0.00866 (0.00602)	0.0136** (0.00628)	0.0140** (0.00576)	0.0265** (0.0109)	0.0253** (0.0112)	0.0261** (0.0102)	0.0207** (0.00851)	0.0205*** (0.00767)
Shock*2013	0.0128** (0.00562)	0.0128** (0.00564)	0.00921 (0.00633)	0.0137** (0.00542)	0.0143*** (0.00493)	0.0277*** (0.00789)	0.0283*** (0.00814)	0.0271*** (0.00815)	0.0250*** (0.00664)	0.0248*** (0.00596)
Bartik*2014	0.0142** (0.00576)	0.0139** (0.00575)	0.0106* (0.00616)	0.0150*** (0.00542)	0.0160*** (0.00498)	0.0229*** (0.00786)	0.0234*** (0.00790)	0.0199** (0.00878)	0.0225*** (0.00675)	0.0232*** (0.00608)
Shock*2015	0.0213*** (0.00675)	0.0213*** (0.00656)	0.0204** (0.00857)	0.0230*** (0.00697)	0.0252*** (0.00664)	0.0130 (0.00817)	0.0142* (0.00852)	0.0131 (0.00831)	0.0131 (0.00800)	0.0152** (0.00769)
# Cells	2,286	2,286	2,286	177,971	177,971	2,286	2,286	2,286	177,971	177,971
R-Squared	0.351	0.354	0.427	0.044	0.314	0.318	0.323	0.333	0.034	0.291
MSA-Level	X	X	X			X	X	X		
Unemployment Controls		X					X			
Includes missing firms			X					X		
Occupation-MSA Level				X	X				X	X
Occ FE and time trend					X					X

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

Notes: See panel A. Observations additionally restricted to ads with any specific skill requirements.

Table D2: Probability of Being Missing from a Subsample

	(1)	(2)	(3)
	Change in Pr(Missing from):		
	Firm Sample	2007-10 Change Sample	Compustat Sample
Shock*2010	-0.0273 (0.0178)	-0.0217* (0.0125)	0.0128 (0.0144)
Shock*2011	-0.0299 (0.0184)	-0.0119 (0.0118)	0.0239* (0.0139)
Shock*2012	-0.0656*** (0.0195)	-0.0151 (0.0118)	0.0333** (0.0159)
Shock*2013	-0.0155 (0.0202)	0.00369 (0.0115)	0.00888 (0.0107)
Bartiik*2014	0.00305 (0.0269)	0.0146 (0.0123)	0.0105 (0.0124)
Shock*2015	0.0269 (0.0224)	0.000999 (0.0124)	-0.00101 (0.0128)
# MSA-year Cells	2,286	2,286	2,286
R ²	0.265	0.241	0.594

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

Notes: Observations are MSA-year cells weighted by the MSA labor force in 2006. Dependent variables are the change from 2007 in the share of ads missing from the sample of non-missing firms (column 1), the sample of firms that post at least 5 ads in 2007 and 2010 conditional on being in the non-missing firm sample (column 2), and the Compustat sample, conditional on being in the 2007-10 change sample (column 3). Regressions also include year fixed effects and MSA characteristics obtained from the ACS. Standard errors are clustered at the MSA level.