

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

February 25, 2016

Abstract

Routine-biased technological change (RBTC), whereby routine-task jobs are replaced by machines and overseas labor, shifts demand towards high- and low-skill jobs, resulting in job polarization of the U.S. labor market. We test whether recessions accelerate this process. In doing so we establish a new fact about the demand for skill over the business cycle. Using a new database containing the near-universe of electronic job vacancies that span the Great Recession, we find evidence of upskilling – firms demanding more-skilled workers when local employment growth is slower. We find that upskilling is sizable in magnitude and largely due to changes in skill requirements within firm-occupation cells. We argue that upskilling is driven primarily by firm restructuring of production towards more-skilled workers. We show that (1) skill demand remains elevated after local economies recover from the Great Recession, driven primarily by the same firms that upskilled early in the recovery; (2) among publicly traded firms in our data, those that upskill more also increase capital stock by more over the same time period, and (3) upskilling is concentrated within routine-task occupations – those most vulnerable to RBTC. Our result is unlikely to be driven by firms' opportunistically seeking to hire more-skilled workers in a slack labor market, and we rule out other cyclical explanations. We thus present the first direct evidence that the Great Recession precipitated new technological adoption.

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

*Correspondence: Brad Hershbein, W.E. Upjohn Institute for Employment Research, 300 S. Westnedge Ave., Kalamazoo, MI 49007. Email: hershbein@upjohn.org. Lisa Kahn, Yale School of Management, 165 Whitney Ave, PO Box 208200, New Haven, CT 06511. Email: lisa.kahn@yale.edu. We are grateful to Jason Abaluck, Joe Altonji, David Autor, Tim Bartik, David Berger, Jeff Clemens, David Deming, David Green, John Horton, Fabian Lange, Steve Malliaris, Alicia Sasser Modestino, Daniel Shoag, Henry Siu, Basit Zafar, and seminar participants at the AEAs, the Atlanta Fed, Brookings, Harvard, the LSE, MIT, Princeton, Rutgers, SOLE/EALE 2015 meetings, Trans Pacific Labor Seminar 2015, University of British Columbia, and University of South Florida. 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 towards those in the tails is one of the most important trends in the U.S. labor market over the last 30 years. Previous research makes a compelling case that a primary driver of this job polarization is routine-biased technological change (RBTC), whereby new machine technologies and overseas labor substitute for middle-skill jobs in the U.S. and are in turn complementary to high-skill cognitive jobs and low-skill manual 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). 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 occurred during recessions and that, over the same time period, recovery was jobless only in these occupations. However, there is no direct evidence on how firms restructure in the face of technological change, and whether it is gradual or episodic. This paper aims to fill that gap.

In this paper we investigate how the demand for skills changes over the business cycle. We use a new data set collected by Burning Glass Technologies that contains the near-universe of electronically posted job vacancies in U.S. cities in 2007 and 2010–2015. Exploiting spatial variation in economic conditions, we establish a new fact: the skill requirements of job ads increase in metropolitan statistical areas (MSAs) that suffered larger employment shocks in the Great Recession, relative to the same areas before the shock and other MSAs that experienced smaller shocks. Our estimates imply that experiencing a Great Recession-sized employment shock increases the probability that a job ad contains an education or experience

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); Mortensen and Pissarides (1994); Caballero and Hammour (1994, 1996); Gomes, Greenwood, and Rebelo (2001); and Koenders and Rogerson (2005).

requirement by nearly 12%. The average experience requirement increases by over half a year and the probability of requiring at least a bachelor’s degree increases by 37%. Drawing on the richness of our data, we show that the vast majority of this “upskilling” is driven by increases in skill requirements within firm-occupation cells, rather than a shift in the distribution of ads posted across firms or occupations. We also show a similar upskilling effect for realized jobs matches, using Current Population Survey (CPS) data, suggesting that firms successfully hire the more-skilled workers that they seek.

We then examine whether upskilling is driven by changes to production in a manner consistent with RBTC. In short, are firms changing what they do or simply changing whom they hire? For example, upskilling may instead occur because firms temporarily and opportunistically take advantage of a slack market to try to attract workers typically found higher up on the job ladder. As our data are best-suited to measure a shift in demand from middle- to high-skilled workers, rather than concomitant demand increases for low-skilled labor that also would 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 skill requirements remain elevated across MSAs even after the labor market has largely recovered. This persistence is driven by the same firms that upskilled early in the recovery. Any purely cyclical explanation is thus insufficient to explain upskilling. Second, we show that among publicly traded firms in our data, those with larger increases in skill requirements also had larger increases in capital stock over the same time period, consistent with a substitution of routine-task workers with machines. Third, we show that upskilling is concentrated in routine-task occupations – the types of jobs typically lost to polarization – and that the increases in education and experience requirements are accompanied by increases in demand for cognitive skills such as “problem solving” – the types of skills favored by RBTC. Fourth, we find that upskilling is concentrated in high-skilled firms, as defined by their skill requirements before the recession. This is at odds with the notion that upskilling is driven by firms attempting to hire “up” in a slack market. Finally, we rule out that changes in labor supply due to worker quits and formal schooling decisions drive upskilling.

Taken together our results suggest that firms located in areas hit harder by the Great Recession were induced to restructure their production towards greater use of machines or overseas 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 real-time 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 full picture of RBTC. We find that demand for education, experience, and cognitive skill increase in response to worse local labor market conditions, especially *within* routine-task (middle-skill) occupations. 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 hypothesize that a decrease in demand for cognitive occupations drove college graduates to take jobs lower in the occupational distribution, squeezing out the high school graduates who formerly held them. This is something of a puzzle, especially given the common belief that technological change continues and more-skilled workers earn a premium in the labor market (Card, Heining, and Kline 2013; Card, Cardoso, and Kline, forthcoming). We believe 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 help answer an important question about the consequences of recessions for workers. It is well known that low-skilled workers suffer worse employment and wage consequences in recessions.⁴ Evidence from past recessions shows that in downturns workers are more likely to take worse jobs, relative to their skills (Devereux 2002, Kahn 2010, von Wachter, Oreopoulos, and Heisz 2012). Some of this may result because workers apply to a broader set of jobs, even those for which they would normally be overqualified, when job-finding rates decline. However, it is equally possible that firms actively seek a more-skilled worker than they could have attracted in a tighter market.⁵ We find that the most impor-

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. He shows that occupations that require social skills, and especially occupations that require both social and cognitive skills, have grown in both employment and wages, even as occupations requiring cognitive but not social skills have not.

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

⁵There is mounting evidence that firms at the bottom of the job ladder benefit from increased retention of their incumbent workforce in a downturn (Moscarini and Postel-Vinay 2012, Kahn and McEntarfer 2015),

tant driver of upskilling is technological change, rather than temporary opportunistic hiring. Thus the downshift of workers into lower-quality jobs is more likely driven by worker search behavior.⁶

Finally, we contribute to a growing literature exploiting data on vacancy postings. Although several studies have used aggregate vacancy data, and even vacancy microdata, from the Bureau of Labor Statistics' Job Openings and Labor Market Turnover (JOLTS) survey (see, for example, Davis, Faberman, and Haltiwanger 2012), these data contain little information on the characteristics of a given vacancy or the firm that is posting it. Fewer studies have used vacancy data that contain information on the occupation or specific requirements of the job posted, and these have generally used aggregate data (Sasser Modestino, Shoag, and Ballance 2015), narrow slices of the data (Rothwell 2014), or data that are limited to one vacancy source (Marinescu 2014, Kuhn and Shen 2013). To our knowledge, we are the first study to use data based on a near-universe of online job postings that covers every metropolitan area in the United States.

We demonstrate that during the Great Recession firms changed not only whom they would hire in the recovery, but how they would produce. These changes may have occurred in the absence of the recession, but at a more gradual pace, with smaller numbers of workers needing to be reallocated at any given time. Instead, episodic technological change results in a swath of displaced workers whose skills are suddenly rendered obsolete. The need to reallocate workers on such a large scale likely drives jobless recovery. It also likely plays a role in the well-noted and marked decline in male employment-to-population ratios over the last 20 years, especially since these declines have been stair-step around recessions (Moffitt 2012).⁷ The evidence provided in this paper is thus integral for understanding worker reallocation, and can help inform policymakers about the optimal mix during a downturn of worker retraining and subsidizing job search through unemployment insurance.

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

⁶In a recent paper, Sasser Modestino, Shoag, and Ballance (2015) use a small subset of occupations from aggregated job postings data for 2007, 2010, and 2012 (originating from the same source as our micro-level data but aggregated to a single summary observation per state-year-occupation) to estimate the labor demand response to local labor supply shocks, exploiting identifying variation from the Great Recession and large-scale releases of military personnel. They too find evidence of upskilling and argue that it is driven by firms opportunistically seeking more skilled workers in a slack labor market. Our analysis shows that this conclusion is unlikely to be the most important driver of upskilling, which we show to be quite pervasive across time and space. Our ability to look further into the recovery (into 2015) and our ability to disaggregate to the job posting level, and in particular exploit information on firms, yields better and more broad-based understanding of the mechanisms driving upskilling.

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

and section 6 concludes.

2 Data

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

2.1 Burning Glass Overview

Burning Glass Technologies (henceforth BG or Burning Glass) examines some 40,000 online job boards and company websites to aggregate the job postings, parse and deduplicate them into a systematic, machine-readable form, and create labor market analytic products. Thanks to the breadth of this coverage, BG believes the resulting database captures a near-universe of jobs that were posted online. Through a special agreement, we obtained an extract from BG, which covers every MSA in the United States in 2007 and from 2010 through the second quarter of 2015.⁸

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

However, the richness of the BG data comes with a few shortcomings. Notably, the database covers only vacancies posted on the Internet. Even though vacancies for available

⁸The database unfortunately lacks postings from 2008 and 2009. Our extract was provided in September 2015. 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.

⁹Other 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.

jobs have increasingly appeared online instead of in traditional sources, such as newspapers, one may worry that the types of jobs posted online are somehow different. In Appendix A, we provide a detailed description of the industry-occupation mix of vacancies in BG.

To briefly summarize, we find that differences in aggregated industry composition, compared to JOLTS vacancies (which cover online and print ads), are small. However, differences are larger comparing the occupation composition in BG to that of employment stocks and flows. Perhaps not unexpectedly, BG has a much larger ad share in computer and mathematical occupations. It is also overrepresented among management, healthcare practitioners, and business and financial operations, although to lesser degrees.

A downside of the BG sample is that low-skill jobs are rare, relative to their employment share. For example BG data are underrepresented in transportation, food preparation and serving, production, and construction. We thus cannot obtain accurate estimates of the degree to which firms shift from middle-skill to low-skill, manual jobs. This is a dimension of RBTC and job polarization to which we cannot speak. However, we instead provide a compelling, rich analysis of the shift towards high-skilled workers, which is equally important, and perhaps more interesting given the recent secular decline in demand for cognitive occupations noted by Beaudry, Green, and Sand (2014, 2016).

Despite these differences in the distribution of occupations, we are comforted by the finding in Carnevale, Jayasundera, and Repnikov (2014) that BG fares well in terms of aggregate *fluctuations* of vacancies (relative to JOLTS). Furthermore, our estimate of the relationship between skill requirements of ad postings and local economic conditions will be internally valid to BG. The primary threat to internal validity is that the representativeness of jobs in BG may be changing over time in a way that is systematically correlated with local economic conditions. For example, if less-skilled jobs were gradually increasing their online presence, then improvements in economic conditions (which occur later in the sample period) would be associated with downskilling. However, we show in appendix A that the representativeness of occupations in the BG data hardly changes over our sample period.

2.2 Details and Construction of Analytic Data Set

In the raw data, there are two fields each for education and experience requirements: a minimum level (degree or years of experience) and a preferred level.¹⁰ 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.

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

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.¹¹

We restrict our sample to ads with non-missing employers that posted at least 10 ads over the sample period of 2007 and 2010 through June of 2015. After cleaning, our data contain 141,728 distinct employers.¹² Many of our analyses exploit firm-level information to distinguish among possible mechanisms for upskilling. We therefore choose to focus our entire analysis on the consistent sample of ads with non-missing firms, with a sufficient number of observations per firm to estimate firm-level characteristics. However, we have performed analyses not requiring firm-level information on the full data set and obtain very similar results. Moreover, we have confirmed that the probability of satisfying this sample criterion (having a valid firm identifier) does not vary over the business cycle. Appendix table A1 reports regression results that use our main specification (described in detail below) to estimate the probability of ads missing an employer (and the probability of not qualifying for other subsamples used throughout the analysis) as a function of local labor market conditions and controls. From column 1, the coefficient on the local conditions variable is tiny, both statistically and economically.¹³ Thus, our sample restriction should not confound the estimated relationship between local labor market conditions and the skill requirements of postings.

Table 1 summarizes data for the primary regression sample. In practice, we collapse our individual-level postings into cells at the MSA-month-year-occupation (4-digit SOC) level.¹⁴ This procedure produces nearly two million cells. The average cell is made up of about 25 individual postings, although there is considerable variation. For analyses that exploit firm-level variables, we further disaggregate to MSA-month-year-occupation-firm cells, of which there are 26 million, with an average of 2 individual postings per cell.

The remainder of the table produces weighted summary statistics. Here, and in our regression analyses, we weight by the product of the 4-digit SOC occupation’s ad share in a given MSA-year-month cell and the size of the MSA’s labor force in 2006. This enables us to preserve the time-varying occupation distribution within an MSA, and upweights both larger

¹¹When 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”).

¹²The 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).

¹³The coefficient of -0.0078 (with a standard error of 0.041) implies that when an MSA experiences an employment growth shock of comparable magnitude to the Great Recession, the probability of an ad missing a firm identifier falls by 0.78 ppts, off a base of 40%.

¹⁴We 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.

Table 1: Summary Statistics

	Mean	St Dev	Min	Max
Number of Postings (unweighted)	25.5	133.2	1	27353
Bartik employment growth	0.01	0.01	-0.08	0.05
Unemployment Rate	7.29	2.49	1.50	32.20
Year	2011	2	2007	2015
Month	6.27	3.44	1	12
Education Requirements:				
Any	0.53	0.22	0	1
Years, Conditional on any	14.57	1.67	12	21
AA	0.05	0.10	0	1
BA	0.26	0.22	0	1
MA	0.03	0.07	0	1
>MA	0.01	0.06	0	1
Experience Requirements:				
Any	0.50	0.22	0	1
Years, Unconditional on any	1.77	1.26	0	15
Years, Conditional on any	3.32	1.53	0.08	18
1-2 Years	0.20	0.14	0	1
3-5 Years	0.20	0.15	0	1
5-10 Years	0.07	0.09	0	1
Skill Requirements:				
None	0.12	0.19	0	1
Cognitive, conditional on any	0.32	0.20	0	1
Occupation Task Variables:				
DOT cognitive	0.55	0.50	0	1
DOT routine	0.36	0.48	0	1
DOT manual	0.09	0.29	0	1
Not in 2007-10 change sample	0.33	0.21	0	1
Not in Compustat and change sample	0.70	0.19	0	1
Not in 2007 match sample	0.29	0.21	0	1
# Occupation-MSA-Date Cells		1,966,837		

Notes: Burning Glass data 2007 and 2010-2015Q2. Sample is restricted to ads with non-missing firms that posted at least 10 ads over our sample period. Summary statistics are calculated weighting by the occupation's ad share in the MSA-date times the size of the MSA labor force in 2006. Bartik employment growth is the 12 month change in log employment in an MSA predicted by the 3-digit NAICS industry composition of the MSA in 2004-2005 and national industry-level employment changes.

occupations and larger MSAs, helping with precision and creating a more representative sample. We use the size of the labor force in a pre-period rather than labor market indicators throughout the sample to avoid simultaneity concerns.

In our data, 53% of the weighted ads list any education requirement. Among ads with an education requirement, roughly half (26% of all ads) specify minimum education of a bachelor’s degree, and another sizable fraction (about 18%) ask for a high school diploma. Relatively few ads specify an associate’s degree (5% of all ads), a master’s degree (3%), or a doctoral degree (1%). Converting the degrees to their modal equivalent years of schooling, the average education requirement, conditional on one being specified, is 14.6 years.

Similarly, half of the weighted ads ask for some amount of experience in the field. Among ads that specify a minimum amount of experience, the vast majority ask for between 1 and 5 years, with much of the remainder asking for between 5 and 10 years. Conditional on posting an experience requirement, the average ad asks for 3.3 years. If we treat ads without an experience requirement as zeroes, we find an unconditional average of 1.8 years. In regression analyses below, we will examine both intensive and extensive margins of changes in skill requirements.

The remainder of table 1 provides information on variables used in later analyses and will be discussed subsequently.

3 Skill Requirements and Local Employment Conditions

In this section we ask whether skill requirements of job vacancy postings respond to changes in local labor market conditions. We start with regression analysis, then perform a formal decomposition to determine the extent to which any change in skill requirements is driven by a changing composition of occupations and firms, or changes in skill requirements within observable cells. Finally, we perform a companion analysis with CPS data on new job starts to determine whether any changes in stated preferences for skills translate into changes among realized matches.

3.1 Methodology

We estimate regressions of the form specified in equation 1, where *skill* is any of several measures of the average skill requirements in occupation *o*, MSA, *m*, and date (month-year) *t*, discussed in more detail below.¹⁵ $I^{m*month}$ are MSA-by-calendar month fixed effects (capturing differences in levels across MSAs and MSA-specific seasonality), I^t are date (month-year)

¹⁵We also exploit firm identifiers in the data to estimate regressions at the MSA-date-firm level and the MSA-date-firm-occupation level, modifying the specification slightly as described below. The more disaggregated regressions weight by the firm’s ad share or the firm-occupation’s ad share in an MSA-date times the MSA labor force. Because of our weighting scheme, the more aggregate regressions produce results identical to those using more disaggregated data when the underlying specification is the same.

fixed effects, and ε_{omt} is an error term.

$$(1) \quad skill_{omt} = \alpha_0 + \alpha_1 inv_Emp_{mt} + I^{m*month} + I^t + \varepsilon_{omt}$$

The key explanatory variable, inv_Emp_{mt} , is based on the projected year-over-year employment growth rate for an MSA each month. 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).¹⁶

Specifically, 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) and $\ln E_{kt}$ 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 k in the same calendar month one year prior.¹⁷

The growth of employment is a natural measure of labor market conditions to use in our setting since both it and vacancy postings represent employment flow measures. Our choice of using the Bartik measure, instead of actual employment growth (as reported by the Bureau of Labor Statistics), is motivated by two reasons. First, actual employment growth at the MSA level is measured with substantial error, especially from month-to-month, 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. For example, MSAs with secular increases in population due to migration flows may experience employment changes that are higher than average but still have a weakening labor market. The Bartik shock addresses this issue. We note that other direct measures of local labor market tightness, such as the local unemployment rate, have similar shortcomings in terms of measurement error or reverse causality; for instance, an unemployment rate may be high precisely *because* a sudden demand shift towards more-skilled labor generates structural mismatch.¹⁸

In practice, we use the reduced-form Bartik shock but for ease of interpretation we invert the measure and scale it so that a one-unit change equals the drop in employment growth

¹⁶Other 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.

¹⁷We 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>. We choose a 12-month employment change because it is easy to interpret and it avoids the concern of seasonal adjustment (in our preferred specification we use seasonally unadjusted employment numbers). However, our results are robust to using shorter time horizons of employment growth (one month, three months, and six months).

¹⁸All that said, our estimates are qualitatively robust to using actual employment growth or the unemployment rate, although magnitudes are somewhat reduced, likely for the reasons indicated.

seen during the Great Recession, hereafter termed inverse employment growth.¹⁹

Since we control for month-year (I^t), the coefficient α_1 is identified purely off of cross-sectional variation in employment growth rates over time, rather than relying on the national shock generated by the Great Recession. This is a necessity given our short panel. In the raw data, we do see that skill requirements increase between 2007 and 2010. This coincides with the Great Recession, but might also be driven by preexisting national trends or changes in data quality. Date fixed effects absorb both confounding factors. We thus isolate changes in skill requirements driven only by differences across places in the employment shock generated by the Great Recession. There is substantial cross-sectional variation in the Bartik shock even within these controls across our 381 MSAs.

Our preferred specification does not control for occupation fixed effects, as the mix of occupations may be an important channel through which local labor market conditions affect the skill requirements in posted vacancies. Rather, our weighting strategy implicitly allows for both within and across occupation changes in demanded skills, but places more weight on relatively larger occupations within the MSA. Nonetheless, we explore the robustness of our results to a variety of different specifications, including controlling for occupation. Furthermore, the decompositions we perform below more fully address the specific mechanisms through which skill requirements may change.

We cluster standard errors by MSA to address possible serial correlation within an area.²⁰

3.2 Results

Table 2 reports regression results from equation (1) for several dependent variables related to the skill requirements of job postings. The dependent variable in column 1 is the share of ads with any education (panel A) or experience (panel B) requirement. The coefficient on inverse employment growth of 0.063 in panel A implies that when MSA employment growth falls by the amount seen in the Great Recession, the probability that an ad lists an education requirement increases by 6.3 ppts (12%), relative to that seen in the MSA before the recession, and relative to other MSAs that did not experience a recession. Panel B shows the probability of listing any experience requirement similarly increases by 5.7 ppts (11.4%). Both effects are significant at the 5%-level.

The remaining columns of table 2 explore other dependent variables, including the number

¹⁹From peak to trough over the Great Recession, national annual employment growth fell from about 2% to -6%. We thus divide Bartik employment by -0.08 and multiply by a factor to scale national growth numbers to local growth numbers. We obtain the scaling factor from a first-stage regression of local employment growth on the Bartik shock as well as our preferred set of controls. The coefficient on the Bartik shock is 2.18 with a standard error of 0.25. We use 2.18 as the scaling factor in all MSA-level regressions, regardless of controls, for consistency. We choose this method, rather than the full two-staged least squares, to save on computation time. Standard errors are about 10% larger when using the latter method and statistical significance is usually not affected.

²⁰We 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.

Table 2: Skill Requirements and Labor Market Conditions

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Education Requirements						
	Any	Yrs Condl	≥ AA	≥ BA	≥ MA	>MA
Inverse Emp Growth	0.0632** (0.0286)	0.396*** (0.0874)	0.121*** (0.0318)	0.109*** (0.0325)	0.0122*** (0.00438)	0.00954*** (0.00194)
# Occ-MSA-Date Cells	1,966,837	1,578,367	1,966,837	1,966,837	1,966,837	1,966,837
R ²	0.189	0.075	0.109	0.129	0.025	0.015
Panel B: Experience Requirements						
	Any	Years		≥ 1 Yr	≥ 3 Years	≥ 5 Yrs
		Condl	Uncond'l			
Inverse Emp Growth	0.0571** (0.0289)	0.373*** (0.120)	0.564*** (0.185)	0.0690** (0.0294)	0.0829*** (0.0274)	0.0430*** (0.0140)
# Occ-MSA-Date Cells	1,966,837	1,496,361	1,966,837	1,966,837	1,966,837	1,966,837
R ²	0.203	0.115	0.169	0.187	0.148	0.137

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

Notes: Each panel-column reports a different regression. All regressions include MSA-by-calendar month and date (month-year) fixed effects and cluster standard errors at the MSA level. Occupation is at the 4-digit SOC level. Observations are weighted by the occupation's ad share in the MSA-date times the size of the MSA labor force in 2006. Inverse employment growth is the 12 month change in log employment in an MSA predicted by the 3-digit NAICS industry composition of the MSA in 2004-2005 and national industry-level employment changes. This measure is normed to the size seen in a large recession by dividing by -0.08 (the size of the national employment change) and multiplying by 2.18 (a scaling factor to map national to MSA employment). Years education is calculated by imputing 12, 14, 16, 18, and 21 for high school, associate's, bachelor's, master's, and >MA, respectively. Unconditional years experience is calculated by imputing zeroes for ads with no requirement.

of years of education or experience listed and the probability that an ad asks for requirements at various thresholds. For example, column 2 of panel A shows that conditional on posting an education requirement, the average requirement increased by 0.4 years, and, in column 4, the probability of ads specifying at least a bachelor's degree (where ads with no education requirement are coded as zero) increases by 11 ppts (42%) in an MSA that experienced a severe recession. The other education thresholds also saw significant increases.

Similarly, column 3 of panel B shows the amount of experience required (adding in zeroes when none is specified) increases by 0.56 years after a severe recession, and this increase is seen throughout the distribution. More generally, across all the margins, effects are sizable and statistically significant at usually the 1%-level. These data provide strong evidence of a countercyclical increase in skill requirements of posted vacancies. We hereafter term this effect "upskilling."

In table 3 we examine the robustness of the upskilling effect to a number of different samples and controls. Column 1 replicates the results from column 1 of table 2. As a summary measure for upskilling, we use the probability of posting any education or experience requirement as an outcome, although results for the other measures are robust as well.

In column 2, we control for MSA-specific time trends. These trends help control for the possibility that different MSAs could be increasing their online representation of job postings at different paces. Also, the time trends help to control for any relationship between pre-recession industry composition and changes in the skill level of the economy. For example, areas that experienced a worse recession may have seen more rapid changing of skill demand because of RBTC even before the recession. That results are fully robust to MSA-specific time trends helps alleviate this concern.

In columns 3 through 8, we add a series of fixed effects: those for occupation, occupation-by-MSA, occupation-by-time, firm, firm-by-occupation, and firm-by-occupation-by-MSA, re-

spectively.²¹ Results are quite robust across these specifications and magnitudes drop off only slightly.

The purpose of these controls is to adjust for changes in the sample driven, say, by changes in the representativeness of the BG data. For example, some MSAs may use online job ads for a wider range of occupations than others. Some occupations may be increasing their representation on online labor markets faster than others. The fact that our estimates hold even within firm-by-occupation-by-MSA cells makes it quite unlikely that these data issues are driving our result. It is unlikely that firms would, on average, change their online strategy for a given occupation, in a given MSA, over a relatively short time period, for other reasons that happen to be correlated with local labor market conditions.

²¹To save on computational time, columns 2, 4, and 6–8 omit MSA-specific seasonality controls, though this makes very little difference.

Table 3: Skill Requirements and Local Labor Market Conditions – Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Panel A: Education Requirements					
Inverse Emp Growth	0.0632** (0.0286)	0.0683** (0.0343)	0.0400 (0.0270)	0.0409 (0.0269)	0.0648*** (0.00357)	0.0439* (0.0233)	0.0408* (0.0235)	0.0519* (0.0295)
# Cells	1,966,837	1,966,837	1,966,837	1,966,837	1,966,837	11,457,573	26,261,191	26,261,191
R-Squared	0.189	0.192	0.546	0.599	0.512	0.503	0.526	0.665
			Panel B: Experience Requirements					
Inverse Emp Growth	0.0571** (0.0289)	0.0603* (0.0343)	0.0517* (0.0282)	0.0569** (0.0283)	0.0967*** (0.00334)	0.0576** (0.0248)	0.0572** (0.0244)	0.0640** (0.0307)
# Cells	1,966,837	1,966,837	1,966,837	1,966,837	1,966,837	11,457,573	26,261,191	26,261,191
R-squared	0.203	0.208	0.535	0.584	0.492	0.445	0.485	0.636
MSA time trends		X						
Occupation FE's			X		X			
Occ x MSA FE's				X				
Occupation time trends					X			
Firm FE						X		
Firm x Occ FE's							X	
Firm x Occ x MSA FE's								X

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

Notes: See table 2. In column 6, regressions are estimated on firm-MSA-date cells. In columns 7-8, regressions are estimated on firm-occupation-MSA-date cells. Columns 2, 4, and 6-8 omit MSA-specific seasonality controls for computational ease, though still include MSA fixed effects.

The results in table 3 also help build intuition for whether upskilling is driven by changes in the composition of ads posted. Occupation controls do reduce the magnitude of the coefficient in the education regression, suggesting upskilling in education may be partially accounted for by the changing composition of occupations. However, the other results largely hold within occupation, firm, and firm-by-occupation-by-MSA, suggesting that compositional shifts across any of these dimensions are not driving upskilling. We can even restrict our sample to a balanced panel of ads (firm-by-occupation-by-MSA cells that appear in 2007 and later) and obtain similar results. However, since all of these variables are measured with substantial error, we perform a formal decomposition next.²²

3.3 Decomposition

In this subsection, we decompose the change in skill requirements of an MSA’s job postings over the Great Recession and its aftermath. The goal is to understand whether counter-cyclical upskilling is driven by compositional changes such as shifts in the types of firms hiring or shifts in the types of occupations being demanded, as compared to changes in skill requirements within firm-occupation cells.

The average skill requirement among ads in an MSA, m , date, t , is shown in equation 2, expressed as a function of six components.

$$(2) \quad skill_{mt} = C_{mt} \sum_f \left(\sum_o skill_{ofmt}^C * \frac{N_{ofmt}^C}{N_{fmt}^C} \right) * \frac{N_{fmt}^C}{N_{mt}^C} + (1 - C_{mt}) \sum_o skill_{omt}^{NC} \frac{N_{omt}^{NC}}{N_{mt}^{NC}}$$

C_{mt} is the fraction of ads in mt that are to continuing firms, i.e., firms that posted ads in the MSA in 2007 and in later years. Among continuing firms, skill requirements equal the average of skill requirements within occupation, o , and firm, f ($skill_{ofmt}^C$), weighted by the distribution of ads across occupations within a firm ($\frac{N_{ofmt}^C}{N_{fmt}^C}$), and the distribution of ads across firms ($\frac{N_{fmt}^C}{N_{mt}^C}$). Among non-continuing firms, fraction $(1 - C_{mt})$, skill requirements equal the average of skill requirements within occupations ($skill_{omt}^{NC}$) (averaged over firms), weighted by the distribution of ads across occupations ($\frac{N_{omt}^{NC}}{N_{mt}^{NC}}$).²³

Each of these six components can be measured in 2007 as well.²⁴ Regressing the full change, $skill_{mt} - skill_{m07}$ for all MSA-dates on inverse employment growth and controls

²²In appendix table A2, we present results on our main specification for different subsamples used below. Column 1 replicates the main education and experience results in our primary sample. Column 2 uses the sample of all job ads, including those with missing firms; reassuringly, results are very similar to those reported in column 1 of table 2. The other samples are described below.

²³Since, by definition, non-continuing firms cannot be matched to the same firm across 2007 and t , we do not distinguish changes in the distribution of ads across non-continuing firms. These firms include both those that post ads in 2007 but not afterward, as well as those that do not post in 2007 and only post afterward.

²⁴In 2007, the variables for non-continuing firms are defined among firms that post ads only in 2007 and not after.

yields virtually identical results to those reported in table 2.²⁵ But we can also consider counterfactual changes in skill requirements from 2007 to t by differencing any given component from its 2007 level, holding constant the other components. For example, $C_{mt} \sum_f (\sum_o (skill_{ofmt}^C - skill_{ofm07}^C) * \frac{N_{ofmt}^C}{N_{fmt}^C}) * \frac{N_{fmt}^C}{N_{mt}^C} + (1 - C_{mt}) \sum_o \bar{skill}_{omt}^{NC} \frac{N_{omt}^{NC}}{N_{mt}^{NC}}$ is the change in skill requirements between 2007 and t , attributed just to changes in the within occupation-firm skill requirement among continuing firms, holding constant all other components at their levels in t . We can regress this counterfactual skill change on inverse employment growth and the other controls in equation (1) to understand how much of the total cyclical responsiveness is attributed to a cyclical response in the within firm-occupation skill requirement.

A decomposition begins with $skill_{m07}$ and differences each of the six components, one at a time, between the two periods. After differencing one component, that component is fixed at its time- t value for the subsequent differences.²⁶ We can regress each counterfactual change on inverse employment growth (and controls) and the coefficients will sum to the coefficient on inverse employment growth in the actual ($skill_{mt} - skill_{m07}$) regression. Since the order of the decomposition matters, there are 720 possible combinations for decomposing the full effect into its six components.

Results on these 720 decompositions are summarized in table 4 for our primary dependent variables, the share of ads with an education or experience requirement. We report the mean fraction of the overall impact of inverse employment growth on changing skill requirements attributed to each component as well as the standard deviation of the fraction across the 720 decompositions. For example, the -0.26 in the first row and column means that -26% of the cyclical change in education requirements can be attributed to a change in the share of ads to continuing firms. This is because there are fewer continuing firms later in the time period (when economic conditions are worse than in 2007) and continuing firms tend to have

²⁵Since all regressions control for MSA fixed effects, they estimate the change in skill requirements in a given MSA as a function of employment growth. This will be almost identical to a regression of ($skill_{mt} - skill_{m07}$) on the same covariates except that $skill_{m07}$, in practice, depends on the specific continuing firms represented at time t , and therefore varies slightly within MSA over time. The coefficient for the share of ads with an education requirement for the change regression is 0.057 (compared to 0.063 in table 2). The coefficient on the share of ads with an experience requirement for the change regression is 0.046 (compared to 0.057 in table 2).

²⁶For example:

$$skill_{mt} - skill_{m07} =$$

$$\begin{aligned} & C_{m07} \sum_f (\sum_o skill_{ofmt}^C - skill_{ofm07}^C * \frac{N_{ofm07}^C}{N_{f07}^C}) * \frac{N_{f07}^C}{N_{m07}^C} + (1 - C_{m07}) \sum_o \bar{skill}_{om07}^{NC} \frac{N_{om07}^{NC}}{N_{m07}^{NC}} \\ & + C_{m07} \sum_f (\sum_o skill_{ofmt}^C * (\frac{N_{ofmt}^C}{N_{f07}^C} - \frac{N_{ofm07}^C}{N_{f07}^C})) * \frac{N_{f07}^C}{N_{m07}^C} + (1 - C_{m07}) \sum_o \bar{skill}_{om07}^{NC} \frac{N_{om07}^{NC}}{N_{m07}^{NC}} \\ & + C_{m07} \sum_f (\sum_o skill_{ofmt}^C * \frac{N_{ofmt}^C}{N_{f07}^C}) * (\frac{N_{f07}^C}{N_{mt}^C} - \frac{N_{f07}^C}{N_{m07}^C}) + (1 - C_{m07}) \sum_o \bar{skill}_{om07}^{NC} \frac{N_{om07}^{NC}}{N_{m07}^{NC}} \\ & + (C_{mt} - C_{m07}) \sum_f (\sum_o skill_{ofmt}^C * \frac{N_{ofmt}^C}{N_{f07}^C}) * \frac{N_{f07}^C}{N_{mt}^C} + (C_{m07} - C_{mt}) \sum_o \bar{skill}_{om07}^{NC} \frac{N_{om07}^{NC}}{N_{m07}^{NC}} \\ & + C_{mt} \sum_f (\sum_o skill_{ofmt}^C * \frac{N_{ofmt}^C}{N_{f07}^C}) * \frac{N_{f07}^C}{N_{mt}^C} + (1 - C_{mt}) \sum_o (\bar{skill}_{omt}^{NC} - \bar{skill}_{om07}^{NC}) \frac{N_{om07}^{NC}}{N_{m07}^{NC}} \\ & + C_{mt} \sum_f (\sum_o skill_{ofmt}^C * \frac{N_{ofmt}^C}{N_{f07}^C}) * \frac{N_{f07}^C}{N_{mt}^C} + (1 - C_{mt}) \sum_o \bar{skill}_{omt}^{NC} (\frac{N_{omt}^{NC}}{N_{mt}^{NC}} - \frac{N_{om07}^{NC}}{N_{m07}^{NC}}) \end{aligned}$$

Table 4: Decomposing the Impact of Local Economic Conditions on Skill Requirements

	Education Requirement	Experience Requirement
Fraction of impact attributable to a change from 2007 to t in:		
Share of ads to continuing firms	-0.26 (0.34)	-0.53 (0.38)
Among continuing firms:		
within firm-occ skill	1.02 (0.31)	1.49 (0.56)
occupation distribution	0.22 (0.32)	0.01 (0.54)
firm distribution	-0.21 (0.24)	-0.28 (0.34)
Non-continuing firms:		
within occ skill	0.13 (0.33)	0.15 (0.35)
occupation distribution	0.11 (0.13)	0.16 (0.14)

Notes: We decompose the impact of Bartik employment growth on skill requirements (reported in table 2, column 1) into 6 components. Continuing firms post at least 1 ad in 2007 and a later date. Non-continuing firms compare firms that only post ads in 2007 to firms that only post ads after 2007. Beginning with each component's value in 2007, we difference one component at a time by the t value. After differencing one component, that component is fixed at its t value for the subsequent differences. The sum of all counterfactual changes in skill requirements equals the actual change from 2007 to t. We regress each counterfactual change in skill requirements on Bartik employment growth and controls. The fraction of the full effect attributable to a change in a given component from 2007 to t equals the coefficient on Bartik employment growth in the counterfactual skill change regression divided by the coefficient in the actual skill change regression (the value reported in table 2, column 1 -- which contains MSA fixed effects and therefore estimates the change in skill requirements in a given MSA as a function of the Bartik shock). There are 720 possible orders for the decomposition. This table reports the mean fraction attributed to each component and the standard deviation in parentheses across all 720 combinations.

higher average skill requirements. Thus the effect goes in the opposite direction of explaining the main upskilling result.

By far, the largest and most robust contribution to the overall upskilling effect is a change in the within firm-occupation skill requirement of continuing firms (second row). Of the total countercyclical change, it accounts for 102% (149%) of the increase in education (experience) requirements. Regardless of order, the change within firm-occupation is the most important component, accounting for at least 60% or more of the total change.

Among the other components, we find small roles for the shifting distribution of occupations, both among continuing and non-continuing firms. For example, 11% (16%) of the increase in education (experience requirements) is accounted for by the fact that non-continuing firms post in more-skilled occupations when predicted employment growth falls. Similarly, 22% (1%) of the increase in education (experience requirements) is accounted for by the fact that continuing firms post in more-skilled occupations when predicted employment growth falls. These effects are less stable, though, with more variation across the different decomposition orders. These results closely accord with the reduced form analysis in table 3, where occupation fixed effects reduced the coefficient on inverse employment growth for the education requirement, but not for the experience requirement. We also find that non-continuing firms post higher skill requirements for a given occupation, compared to firms that posted only in 2007. This within-occupation skill increase among non-continuing firms accounts for 13% (15%) of the overall education (experience) skill increase.

Finally, we find that the distribution of postings across firms, among continuing firms, negatively explains the skill increase. That is, in a worse economy, among firms that continue to post vacancies, ads shift towards firms that tended to post lower skill requirements in 2007.

This is consistent with work by Kahn and McEntarfer (2015), who find that workers matching to jobs in downturns are more likely to match to low-paying firms than high-paying firms.

3.4 Does Upskilling Extend to CPS Job Matches?

Tables 2 and 3 provide robust, nationally comprehensive evidence of upskilling: firms facing worse local labor markets increase their stated skill requirements in job postings. However, the stated preferences of firms may be of less interest if they are not related to actual, realized hires. 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.

To address this possibility, we perform a companion analysis using data on realized job matches from longitudinally linked CPS data. We categorize new job matches in the CPS as individuals who from one month to the next changed employers, or transitioned from nonemployment to employment (as we do in appendix figures A2 and A3). We examine similar specifications to equation 1 except, given the smaller sample sizes in the CPS, we consider economic conditions across states rather than MSAs.²⁷

Table 5 reports results for education and experience variables that are analogous to those reported in table 2. Experience is simply years of school minus age minus six, which is different in spirit to 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 in the CPS, and it is as close as we can get. The first column for each dependent variable reports results spanning the same time period as the BG data, 2007–2015Q2, while the second ranges from 1999 through 2015, and thus spans the 2001 recession as well.²⁸

Overall we find evidence of upskilling among workers hired in downturns – workers hired when the state economy is worse are more educated and more experienced. Effects are especially pronounced for higher education groups and for experience. For example, we find that a drop in state employment growth equivalent to that seen in the Great Recession is associated with an increase in the share of workers hired with at least a master’s degree of about half a percentage point (off a base of 6%) and an increase in years of experience of about two-thirds. Both effects are significant at the 1%-level for the time period as a whole and at the 5%-level or better for the 2007–2015 period. Thus upskilling is seen in both stated firm preferences and in actual hires.

²⁷We control for state fixed effects, state-specific seasonality and an inverse employment growth measure at the state level (similarly divided by -0.08 but here scaled by 1.02 – a measure appropriate to map national employment changes to the state-level). We control for a quadratic time trend rather than date fixed effects because, with 50 states and the District of Columbia, we have fewer degrees of freedom than with 381 MSAs. We estimate regressions at the individual level and also control for gender-by-race/ethnicity fixed effects. Standard errors are clustered by state and observations are weighted with sample weights.

²⁸We use an average of 1998–99 state-level industry shares to define Bartik employment growth for 1999–2005 and an average of 2004–05 shares to define Bartik employment growth for 2006–2015. Results are similar using shares from 1998–99 for the whole time period.

Table 5: Skills of New Job Matchers, CPS 1999-2015

Dependent Variable:	Years		Education						Experience Years	
	(1)	(2)	≥ AA		≥ BA		≥ MA		(9)	(10)
Inverse Emp Growth	0.00973 (0.0241)	0.0486** (0.0205)	0.00484 (0.00456)	0.0101*** (0.00370)	0.00924** (0.00416)	0.0117*** (0.00226)	0.00528** (0.00237)	0.00546*** (0.00175)	0.698*** (0.143)	0.608*** (0.133)
2007-2015	X		X		X		X		X	
Observations	215,246	455,128	215,246	455,128	215,246	455,128	215,246	455,128	215,246	455,128
R-squared	0.027	0.026	0.023	0.021	0.027	0.025	0.014	0.013	0.013	0.014

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

Notes: Dataset is CPS monthly data on new job matchers. Regressions control for a quadratic time trend, a full set of gender-race interactions, state fixed effects and state-specific calendar month dummies. Observations are weighted with sample weights and standard errors are clustered at state level. Inverse Emp Growth is predicted at the state-date level using 1998-99 industry shares for years 1999-2005 and 2004-05 industry shares for years 2006-2015. This measure is normed to the size seen in a large recession by dividing by -0.08 (the size of the national employment change) and multiplying by 1.02 (the scaling factor mapping national to state-level employment changes).

4 Is Upskilling Driven by RBTC?

The evidence presented above shows a substantial upskilling of labor demand in reaction to eroding local labor market conditions. The decomposition shows that the effect occurs primarily for the same firms and for the same types of jobs. This is true both for firms that post over the entire time period, and for firms that post only in 2007 or only afterward. Although there are several potential explanations for these phenomena, we argue that firms take the opportunity of the recession to restructure their production, responding to preexisting trends in technology or outsourcing.

Standard (s,S) models with convex adjustment costs predict lumpy adjustment to a changing environment (see for example Berger 2012). Relatedly, “pit stop” models of the business cycle express the notion that productivity-enhancing improvements are more likely to take place in downturns because of the relatively lower opportunity cost (see for example Hall 2005 or Koenders and Rogerson 2005). Finally, there could be asymmetries in the costs and benefits of laying off workers. For example, the stigma or bad publicity to a firm making layoffs should decline in a downturn, and the losses in terms of firm-specific human capital that accompany layoffs might be outweighed by a general need to cut on costs. Thus a number of intuitive models predict that adjustment to technological change will be relatively concentrated in downturns, yet no direct evidence exists.

While there are a number of cyclical reasons why firms might upskill their workforce in a slack labor market, discussed in more detail below, the primary implication of the episodic restructuring hypothesis is that upskilling in job ads will be sticky. Once a firm upgrades its workforce, presumably along with an upgrade in machines and technology or overseas infrastructure, those changes will remain. That is, if firms are changing what they do, not simply whom they hire, those changes will persist. In this section, we argue that the evidence is consistent with upskilling being driven primarily by RBTC. Specifically, we show that not only is upskilling persistent throughout the recession and recovery, the persistence is concentrated in the same firms that upskilled between 2007 and 2010. Additionally, we document that the firms that increased their skill requirements more in response to the recession also

increased their capital stock faster than other firms. Finally, we show that upskilling is concentrated in routine-task occupations – those that are most vulnerable to RBTC – and provide corroborating evidence that RBTC occurs *within* these occupations. In contrast, the extant literature on job polarization has examined shifts only across occupations, since the task content of jobs could be measured only at the occupation level in existing data sets.

4.1 Persistence of Upskilling

In the spirit of an impulse response function, we estimate how the impact of the Great Recession on skill requirements varies with time since the initial “shock”. For each MSA, we define the size of the shock ($shock_m$) as the change in Bartik employment growth from peak to trough of the recession, normed in the same way as our primary Bartik employment growth measure so that the coefficient reflects the impact of a Great-Recession-sized employment shock and can be compared to that in tables 2 and 3.²⁹ We then estimate regressions of the following form:

$$(3) \quad skill_{omt} = \alpha_0 + [shock_m * I_t^{year}] \beta + I^m -^{month} + I^t + \varepsilon_{ofmt}.$$

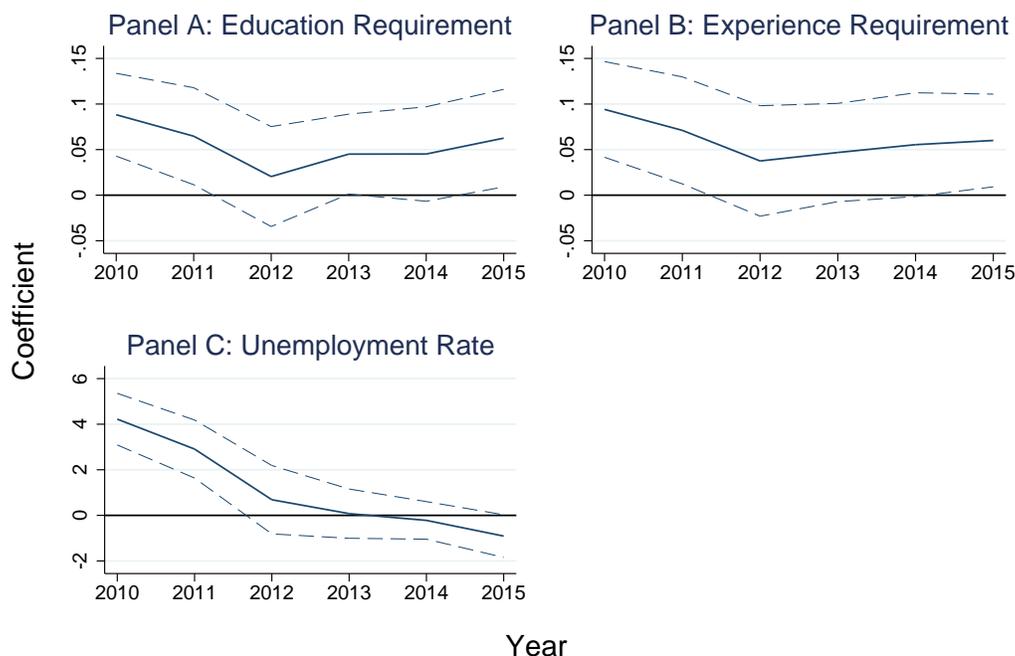
I_t^{year} is a vector of year dummies with 2007 as the omitted year, and the vector of coefficients, β , shows the effect of a unit shock for each year relative to 2007 (the main effect of $shock_m$ is subsumed in the MSA fixed effects). We plot these coefficients for three dependent variables in figure 1 along with their 90% confidence intervals. The regression otherwise includes our standard controls from before.

Panel A shows the impact of the shock on the probability that ads post an education requirement. The coefficient of 0.088 in 2010 means that the probability of posting an education requirement increased by nearly 9 ppts that year, relative to 2007, for an MSA that experienced a recession on par with that of the nation as a whole, compared with an MSA experiencing no shock. The magnitude of this coefficient is slightly larger than the regression results from table 2. This occurs because the peak-to-trough shock is highly correlated with Bartik employment growth in 2010; employment growth did not begin recovering in earnest until late in that year.

Of greater interest is how the impact of the shock persists once employment growth returns to normal. The graph shows some evidence of recovery between 2010 and 2012, with education requirements beginning to return to their 2007 levels (indicated by 0 on the graph). However, the coefficient then edges up before appearing to stabilize. Though estimates are

²⁹Specifically, we subtract the mean monthly employment growth in an MSA in 2006 from the mean monthly employment growth in 2009. (For the most recent business cycle, all MSAs experienced their peak and trough employment growth in 2006 and 2009, respectively). We then divide by -0.08 , the change in national employment growth from peak to trough, and multiply by 2.18, the scaling factor relating the Bartik measure to actual employment growth. Results are similar when we use the maximum growth rate in 2006 and the minimum in 2009.

Figure 1: Impact of MSA-Specific Employment Shock by Year



Graphs plot coefficients on MSA employment shock (change in emp growth from 2006-2009) interacted with year. Dashed Lines = 90% CI.

noisy, the pattern of persistence is clear. The coefficient estimate of 0.063 in the first two quarters of 2015 implies that almost three-quarters of the initial effect of the Great Recession on education requirements remains. Panel B shows a very similar effect for the share of ads with an experience requirement. There we find that nearly two-thirds of the initial impact on experience requirements remains.

For contrast, Panel C shows the impact of the *shock* on the MSA unemployment rate. The coefficient of 4.2 in 2010 implies that a city experiencing a severe employment growth shock from 2007–09 experienced a 4.2 ppt increase in its unemployment rate, relative to 2007, and relative to cities that did not have a shock. Over time, the impact of the shock declines in magnitude. And here, progress does not halt after 2012. By 2015 unemployment rates have converged back to their pre-recession levels.³⁰

Thus, while other labor market indicators, including employment and employment growth, recovered by the end of our sample period, the skill requirements of jobs did not. Of course, this could be because real fundamentals (such as the employment-to-population ratio) had not yet recovered. But the patterns exhibited in figure 1 are suggestive that the upskilling effect is sticky. As of 2015, MSAs that experienced severe shocks in the Great Recession still look different from how they appeared in 2007, and different from other cities that experienced weaker shocks.

³⁰That is, we find that the portion of unemployment rates predicted by Bartik employment growth converged back to their pre-recession levels by 2015. Actual unemployment rate differences across cities were somewhat slower to converge.

That upskilling persists into the recovery is suggestive that firms change what they do during the recession and do not revert back. If the main driver of upskilling is RBTC, we would expect that the *same* firms that upskilled in 2010 persisted in their skill requirements in 2015. It turns out this is, in fact, the case. We categorize firms in our data by their change in skill requirements between 2007 and 2010 and then examine their skill requirements years later. We find that the firms with the biggest increase in skill requirements during the recession maintain their higher skill requirements throughout our sample period.

Figure 2 provides an illustration of this firm-specific persistence in upskilling. We divide firms into quartiles based on changes in education (experience) requirements between 2007 and 2010; we then plot average education (experience) requirements by firm group over time.³¹ Firms began at fairly similar levels of skill requirements in 2007, when nearly 40% of ads had education (left panel) and experience (right) requirements.³² By construction there is a sharp contrast across firms in 2010, with the darker shaded lines representing firms with larger skill increases.³³ Interestingly, and not by construction, these firms remain spread apart throughout the sample period. By 2015 the different groups of firms still have very different skill requirements for the *new* ads that they post. Thus the persistence in upskilling is driven by the same firms that initially upskilled during the Great Recession.³⁴

We have examined these differences across firms in a number of alternative ways. We have split firms based on the MSAs where they post job vacancies and find that firms posting primarily in MSAs (in 2007) that experienced worse recessions have sharper increases in skill requirements that persist throughout the sample period. We have also divided firms based on their responsiveness to the Bartik employment shock from 2007 to 2010. The more-responsive firms increased skill requirements in 2010 (by construction) and *persist* with this increase throughout the sample (not by construction). Thus the finding that the persistence of upskilling is driven by the very same firms that initially upskilled is quite robust.

4.2 Upskilling and Investment

The RBTC story is that firms replace routine workers with either machines or cheaper overseas labor, both of which are complements with skilled labor. If mechanization is occurring,

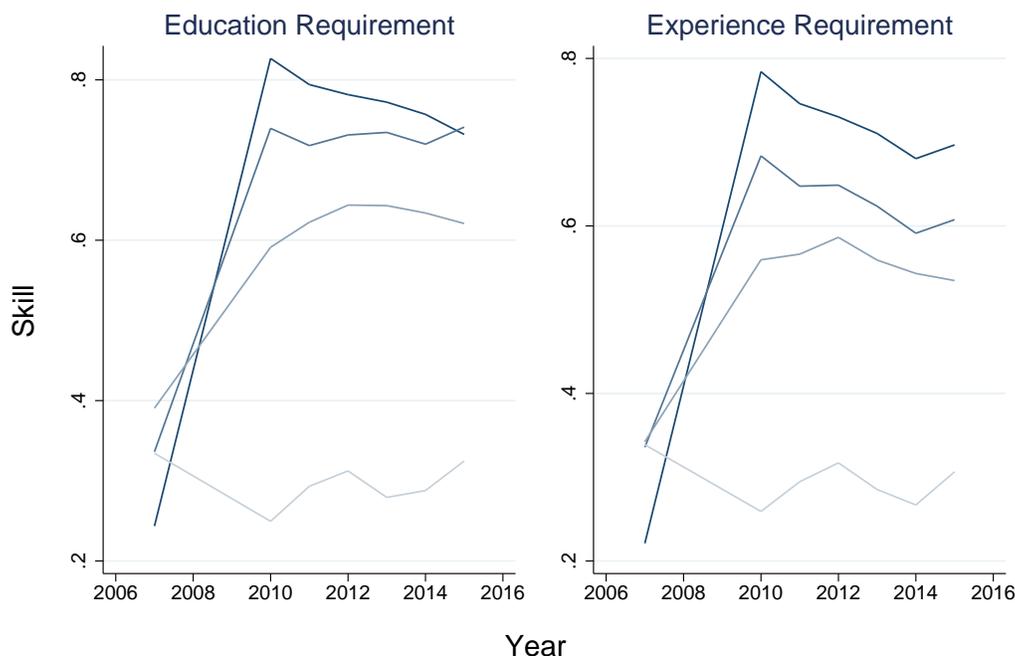
³¹This graph represents a subsample of our data, the 67% of weighted observations (7.4% of firms) with at least five observations in each of 2007 and 2010. We define quartiles using firms' average sample weights across 2007 and 2010. Thus, the quartiles represent roughly equal shares of postings (adjusted for MSA labor force size), grouped by the 2007–2010 change in firm skill demand. Entry into this subsample does not depend on our key explanatory variable, inverse employment growth (see column 2 in appendix table A1). Also, the main upskilling effect is similar in magnitude and significance in this sample as it is in table 2 (see column 3 in appendix table A2).

³²Note: This similarity across firm quartiles is not imposed by our exercise.

³³Firms in the top quartile (darkest line) had an average increase in education requirements of 0.63 from 2007 to 2010, with the remaining quartiles increasing by 0.38, 0.20, and -0.11 for the lowest quartile (lightest line). The analogous quartile changes for experience requirements are 0.60, 0.35, 0.18, and -0.11 .

³⁴Education requirements in firms with the largest increase in education requirements converge to those in the second highest quartile by 2015, but these groups remain spread apart from the other two groups.

Figure 2: Skill Requirements by Firm 2007-2010 Change



This graph plots average skill requirements by year and firm type. Firms are split into quartiles based on their change in skill requirements from 2007-2010. Darker shades indicate quartiles with larger changes.

then we should observe firms investing in physical capital around the time that they upskill. This, however, would be at odds with the stylized facts that (a) investment is in general low in a recession and (b) investment was particularly slow to recover following the Great Recession. 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). This type of investment actually surpassed its pre-recession level before June 2010, less than four quarters after the recession formally ended.³⁵ Interestingly, IT and computers are exactly the kinds of technology that drive RBTC (Michaels, Natraj, and Van Reenan 2014).

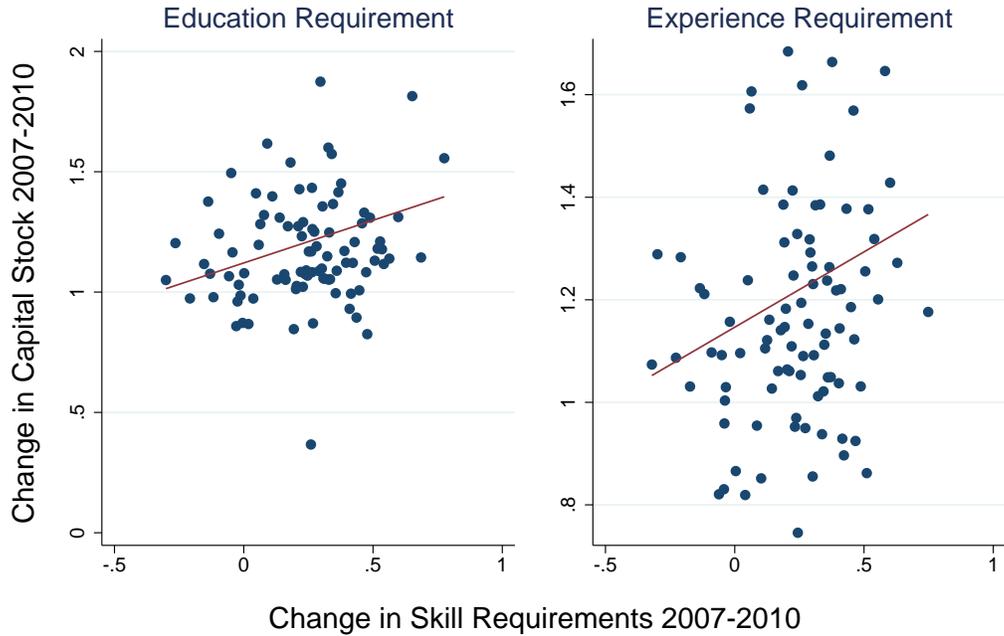
The evidence on general IT investment is again suggestive. In this subsection, we ask whether the firms that upskilled the most increased investment at the same time. We do this by linking the publicly-traded firms in our data to Compustat North America by Standard & Poors (hereafter Compustat), the most complete database of U.S. firm accounting and balance sheet data.³⁶ While we cannot distinguish IT investment in the Compustat data, we can measure a firm's overall holdings of property, plant, and equipment (PPENT). This measure of capital stock should include IT investments and should increase faster for firms that mechanize production in a manner consistent with RBTC. We can thus examine whether the firms that upskill most increase their capital stock over the same time period.

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

³⁵In contrast, domestic private investment as a whole took about four *years* to fully recover.

³⁶We obtain these data via Wharton Research Data Services.

Figure 3: Firm-Level Changes in Skill Requirements and Capital Stock 2007–2010



Graph plots average capital stock (PPENT) and skill requirement change within each ad-weighted percent of skill requirement change -- the firm-level change in average skill requirements from 2007 to 2010. Scatters omit 3 outlier points with capital stock changes > 2, but fitted lines do not.

procedure in Appendix A.3.³⁷ To provide a general sense of the data, figure 3 presents a binned scatter plot of the change in skill requirements between 2007 and 2010 on the rate of change in capital stock (PPENT) over the same time period.³⁸ As can be seen, there is a strong positive relationship for both the share of ads with an education requirement (left panel) and the share with an experience requirement (right panel). Firms that had larger increases in skill requirements between 2007 and 2010 also had larger increases in capital stock over the same time period.³⁹

We also provide results from regression analysis, analogous to those presented in figure 1, comparing the impulse response function across firms with different capital investment behavior. We estimate regressions similar to that in equation 3 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.⁴⁰ Figure 4 shows fitted effects for the impact of the shock on skill requirements, by year, for two types of firms: those at the 10th percentile (solid navy line) of

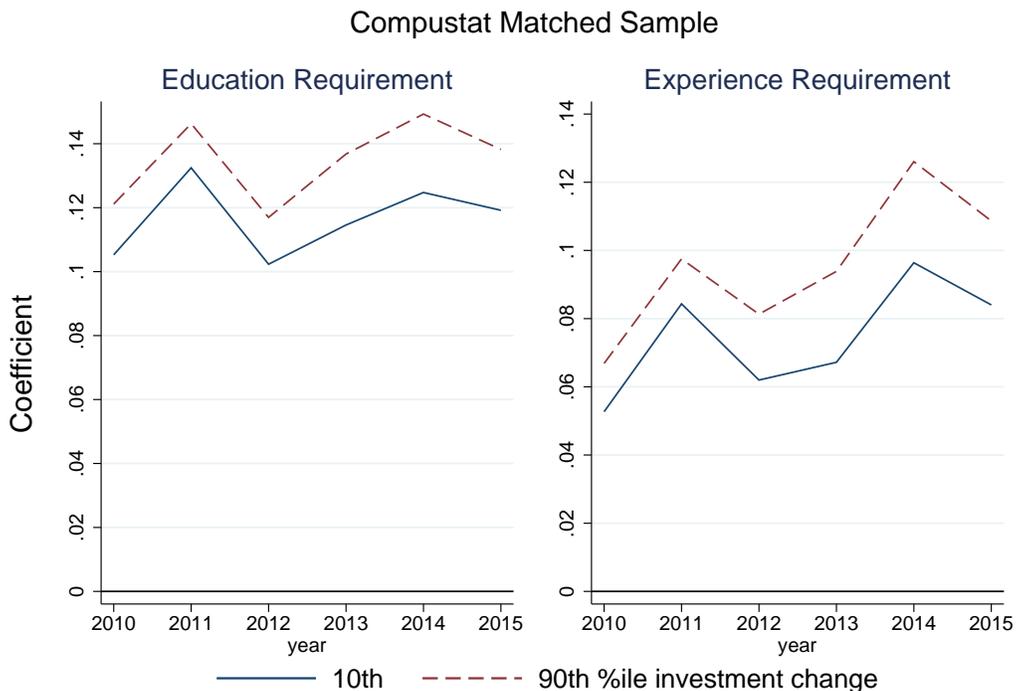
³⁷We can link 46% of weighted observations (18% of firms) from the 2007–2010 change sample used in figure 2 to a business in Compustat. In Appendix A.3, we discuss the fact that our ability to match does vary with local labor market conditions at the time an ad was posted, but fortunately all our results hold in the Compustat-matched sample.

³⁸To make the data easier to see, we aggregate the firms into percentiles based on the magnitude of changes in skill requirements; we then plot for each percentile bin the average of this variable against the average of the change in capital stock, weighting each bin by the number of ads.

³⁹The figure would look similar if instead of actual firm upskilling we plotted the predicted firm-specific upskilling measures defined above, which rely more on the impact of local labor market conditions on skill requirements.

⁴⁰Regressions are estimated at the firm-MSA-date (year-month) level for the Compustat matched sample.

Figure 4: Effect of MSA-Specific Employment Shock on Skill Requirements, by Firm Decile Change in Capital Stock



Graph plots fitted shock-year effects for firms at the 10th and 90th percentiles of investment change from a regression of skill requirements on the MSA employment shock interacted with year dummies and with firm-specific capital investment change from 2007 to 2010.

capital investment change between 2007 and 2010 and those at the 90th percentile (dashed maroon line). The former experienced a 20% drop in their capital stock while the latter had a 57% increase.

Both groups saw substantial increases in skill requirements in response to the MSA-specific employment shock, and these persisted well into the recovery. For example, firms at the 10th percentile of investment change increased their education (experience) requirements by 0.11 (0.05) ppts in 2010 in an MSA that experienced a large shock, relative to postings in that MSA before the shock and relative to other MSAs that did not experience a shock. However, firms with larger increases in capital stock had greater upskilling, and these differences persist. For education requirements, cyclicity is roughly 16% higher in high-investment firms than in low-investment firms, and experience requirements are around 30% more cyclical. These differences are all significant at the 1% level. Thus, throughout our sample period, firms with larger increases in capital stock around the time of the Great Recession also had larger increases in their posted skill requirements.

4.3 Upskilling and task content of jobs

The RBTC literature has been successful at pinpointing the kinds of occupations that can be replaced by machines or overseas labor using information on the tasks performed by workers. The original work by Autor, Levy, and Murnane (2003) and Autor, Katz, and Kearney

(2006) used the US Department of Labor’s *Dictionary of Occupational Titles* (DOT; US Department of Labor 1977) to categorize jobs 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 this grouping successfully predicted employment changes in the 1990s and has been used in a number of subsequent papers, including Autor and Dorn (2013).

Autor (2014) and Jaimovich and Siu (2015) point out that employment shifted away from routine-task occupations in the Great Recession, and, though not shown, we also find this to be true of vacancy postings. We instead focus on our unique ability to measure task and skill content of jobs in the BG data to ask whether RBTC also occurs *within* routine-task occupations. Panels A and B in table 6 present regression results where we augment our primary specification in equation 1 with controls for the DOT occupation categories and allow these characteristics to interact with inverse employment growth.

The coefficient on inverse employment growth in the first row is the effect for the omitted category, cognitive occupations. In column 1, estimates for these occupations are about 60% as large as those for all occupations on average (from column 1 of table 2) and are not statistically significant at conventional levels. However, the interaction term on routine is positive and statistically significant in both panels. Relative to the omitted category, cognitive, routine-task occupations have nearly double the cyclicalities of skill requirements. This relationship holds, although the magnitude is reduced somewhat, as we add occupation fixed effects (column 2), firm fixed effects (column 3), and firm-occupation-MSA fixed effects (column 4).

Upskilling is thus concentrated in routine-task occupations, and is only weakly present in cognitive jobs. The interaction of inverse employment growth and manual occupations is also positive and significant. However, as mentioned above, our data are best-suited to examine shifts from routine to cognitive tasks, rather than shifts from routine to manual, since the BG data are underrepresented in the latter. We thus refrain from making strong inferences about the extent of upskilling in manual occupations, and simply note that changing skill content in these jobs is a topic worthy of further study.

Beaudry, Green, and Sand (2014, 2016) document that more-educated workers have increased their employment in lower-skilled jobs and date this trend as beginning around 2000. They term this shift, along with stagnating employment in cognitive occupations, the “great reversal” in the demand for cognitive skill. They hypothesize that lessened demand for cognitive occupations induces college graduates to take jobs lower in the skill distribution, squeezing out the less-educated workers who formerly held these jobs. 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. That is, higher-skilled workers are more likely to be found in “routine” occupa-

Table 6: Skill Requirements and Occupation Task Content

	(1)	(2)	(3)	(4)
Panel A: Education Requirements				
Inverse Emp Growth (E)	0.0366 (0.0286)	0.0287 (0.0272)	0.0337 (0.0235)	0.0448 (0.0300)
E*Routine	0.0335*** (0.00336)	0.0194*** (0.00298)	0.0107*** (0.00211)	0.00950*** (0.00205)
E*Manual	0.0323*** (0.00510)	0.0269*** (0.00383)	0.0281*** (0.00459)	0.0259*** (0.00370)
Panel B: Experience Requirements				
Inverse Emp Growth (E)	0.0361 (0.0285)	0.0368 (0.0283)	0.0420* (0.0251)	0.0555* (0.0311)
E*Routine	0.0262*** (0.00254)	0.0302*** (0.00242)	0.0282*** (0.00284)	0.0169*** (0.00242)
E*Manual	0.0149*** (0.00479)	0.0167*** (0.00395)	0.00393 (0.00532)	0.0130*** (0.00356)
Panel C: Cognitive Skill Requirement				
Inverse Emp Growth (E)	0.0126 (0.0131)	0.0134 (0.0124)	0.0130 (0.0115)	0.0170 (0.0145)
E*Routine	0.0129*** (0.00207)	0.0120*** (0.00181)	0.0103*** (0.00258)	0.00875*** (0.00206)
E*Manual	0.0348*** (0.00313)	0.0427*** (0.00287)	0.0388*** (0.00459)	0.0290*** (0.00233)
Occupation FE's		X		
Firm FE			X	
Firm x Occ x MSA FE's				X

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

Notes: See tables 2 and 3. Dependent variables are the share of ads with an education (A), experience (B), or cognitive skill (C) requirement. Routine and Manual are *Dictionary of Occupational Titles* categorizations; the omitted category is cognitive. Main effects for Routine and Manual are included in all specifications that do not have occupation fixed effects. Panels A and B are estimated on 1,966,837 MSA-date-occupation cells in columns 1 and 2, on 11,457,573 firm-MSA-date cells in column 3, and on 26,261,191 firm-MSA-occupation-date cells in column 4. Observations in panel C are restricted to ads with any skills listed. Columns 1 and 2 are estimated on 1,860,387 MSA-date-occupation cells, column 3 on 10,447,091 firm-MSA-date cells, and column 4 on 23,824,039 firm-MSA-occupation-date cells.

tions, not only because of a labor supply shift, but because the jobs themselves are changing. Even as employment has shifted from routine to cognitive occupations, the remaining routine occupations themselves are becoming less routine and more cognitive.

The specific skills encoded in the BG data allow us to test this hypothesis. For example, an ad might ask for a worker who is bilingual or who can organize and manage a team. BG cleans and codes these and other skills into a taxonomy of thousands of unique, but standardized requirements.⁴¹ Deming and Kahn (2016) categorize these skill measures into key groupings to measure firm heterogeneity in skill demands and relate that to RBTC and inequality across firms. Panel C of table 6 uses as its dependent variable one such grouping, the probability that an ad requests a cognitive skill.⁴² Conditional on posting any skill requirements, 32% of ads request a cognitive skill.⁴³

In our baseline specification (column 1), we find that in response to a Great-Recession-sized employment shock, ads in cognitive occupations are about 1.3 ppts more likely to request a cognitive skill although this effect is not statistically significant. More importantly, we see that the cyclical component is again stronger in routine occupations: adding the interaction term to the main estimate, the net effect is nearly twice as large and is highly statistically significant. The remaining columns show that these estimates are robust to including a range of fixed effects.

We do not believe that the results in table 6 are driven by a ceiling effect, that cognitive occupations have already reached practical upper limits in their skill requirements. Cognitive occupations are about 8–12 ppts more likely than routine occupations to specify an education, experience, or cognitive skill requirement. Given the mean of these variables ranges from 0.32 to 0.53, it seems there would still be room for firms to become more discerning even for relatively abstract occupations.

To summarize, we find that in a severe recession, job postings for routine-task occupations exhibit larger increases for education and experience requirements, with concomitantly larger increases in demand for explicit cognitive skills. We argue that this is driven by workers performing different and more-productive functions in these formerly routine occupations. If our hypothesis is correct, we would also expect their wages to increase. Our ability to measure wages in the BG data is limited since few job ads – online or not – post wages.⁴⁴ Instead, we look for evidence using new job matches in the CPS. Examining starting wages

⁴¹BG uses a top-down approach where they start with a set of pre-defined possible skills and search text in an ad for an indication that the skill is required. For example, for team work, they first search for the key words “team work” but also look for variations such as “ability to work as a team.”

⁴²An ad is categorized as requesting a cognitive skill if any skills requested include the following phrases: “problem solving”, “research”, “analytical”, or “critical thinking.” Special thanks to David Deming for sharing this measure.

⁴³Only 12% of ads do not list any specific skills. We use the conditional probability of listing a cognitive skill because we believe it is measured with less error. Nonetheless, results are similar when using the unconditional probability, and the likelihood of listing any skill requirement does not vary with local economic conditions (see column 4 of appendix table A1).

⁴⁴In our main sample, only roughly 10% of ads contain salary information.

Table 7: Log Wages of New Job Matchers, CPS 1999-2015

	(1)	(2)	(3)	(4)	(5)	(6)
Inverse Emp Growth (E)	0.00910 (0.0155)	0.00327 (0.0115)	0.0188* (0.0103)	0.00321 (0.0102)	-0.00372 (0.00843)	0.00971 (0.00766)
E*Routine	0.0449** (0.0177)	0.0389*** (0.0137)	0.0139 (0.0123)	0.0397*** (0.0131)	0.0345*** (0.0105)	0.0160* (0.00897)
E*Manual	0.0204 (0.0134)	0.0113 (0.0109)	-0.00218 (0.0109)	0.0374*** (0.00965)	0.0274*** (0.00909)	0.0136 (0.00837)
2007-2015	X	X	X			
Education and age controls		X	X		X	X
Occupation FE's			X			X
Observations	108,018	108,018	108,018	233,110	233,110	233,110
R-squared	0.248	0.366	0.437	0.243	0.368	0.437

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

Notes: See tables 5 and 7. Specification and sample are similar to those reported in table 5 except a worker must have a non-missing wage after starting the new job. We also control for month-in-sample fixed effects and, where specified, years of schooling dummies, a quadratic in age, and 4-digit SOC fixed effects.

over the business cycle is a notoriously difficult problem because of composition bias (Solon, Barsky, and Parker 1994). Nonetheless, in table 7, we provide suggestive evidence on the wage margin.

The table reports a log hourly wage analysis by DOT occupational category, analogous to that in table 6, except we use the CPS data set on new job matches.⁴⁵ Columns 1–3 report results over the same time period as our BG analysis, 2007–2015, while columns 4–6 use the longer sample, years 1999–2015.

We find that in cognitive occupations (the main effect of inverse employment growth) starting wages are roughly acyclical. This is likely because we cannot control for the differential selection of workers and firms matching over the business cycle. When we control for worker characteristics in column 2, the coefficient declines in magnitude. More importantly, we find a positive and statistically significant coefficient on the interaction of inverse employment growth with routine occupation. When a state experiences a Great-Recession-sized employment growth shock, workers matching to routine occupations earn roughly 4% more per hour than workers matching to a routine occupation in that state before the recession and than workers matching to routine occupations in states that experienced a weaker recession. Magnitudes are halved when occupation fixed effects are included, but effects are still sizable and marginally significant.

Beaudry, Green, and Sand (2014, 2016) highlight an important puzzle, that wages and

⁴⁵The monthly CPS collects wage data in Outgoing Rotation Groups (only one-quarter of the months an individual is surveyed). We define a starting wage as the wage collected in the first Outgoing Rotation Group survey after a worker started a new job, conditional on the worker being in the same new job in the month of wage collection. Unless otherwise specified, table 7 estimates similar specifications to those reported in table 5, except we also include month-in-sample fixed effects, to take into account the variation in the tenure length necessary to observe a wage. Education and age controls include years-of-schooling fixed effects and a quadratic in age.

employment have not grown in cognitive occupations since 2000. This demand shift likely drove some high-skilled workers into lower-skilled occupations. However, we document that, at the same time, demand for cognitive skill increased in these formerly routine-task occupations. In MSAs hit harder by the Great Recession, stated preferences for high-skilled workers, as well as their employment and wages increased in these occupations. Our work thus highlights an alternative hypothesis for why high-skilled workers are increasingly found in lower-skilled occupations: these latter occupations are becoming more skilled, and it is possible that less-skilled workers are displaced because they are unable to perform the new duties required.

5 Alternative Explanations for Upskilling

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

Early in the recovery firms may worry about a “bottleneck” effect, that the slack labor market will produce an unusual number of job applicants that can be costly to screen. Firms may try to communicate to certain applicants that a job is not for them by posting specific skill requirements, but this would not imply that firms change how they produce. We do not believe this screening hypothesis is a primary driver of our results for two reasons. First, it would seem surprising that firms end up hiring more-skilled workers (as shown in the CPS job matchers analysis) who are relatively more expensive. Second, this could not explain the persistence in upskilling once local economies have recovered. For example, we showed in figure 1 that unemployment rates eventually converged across MSAs experiencing different-sized shocks, while skill requirements did not.

If firms instead *intend* to hire more-skilled workers without otherwise adjusting their production technology, it presents something of a puzzle: Why would they choose to do so despite the fact that higher-skilled workers are, in general, more expensive, and become even relatively more costly in a downturn?⁴⁶ We explore three possibilities: that firms temporarily and opportunistically seek higher-skilled workers in a slack labor market, that replacement

⁴⁶That lower-skilled workers are more detrimentally affected by recessions is well known (Hoynes, Miller, and Schaller 2012). This implies that high-skilled workers become relatively more costly and relatively less available in downturns. Indeed, using American Community Survey data, we estimate that earnings and unemployment rate gaps widen across education and experience groups when local labor market conditions are worse.

hiring differentially requires more-skilled workers due to differential recovery in quit rates (i.e., the job ladder needs to clear), and population changes in educational attainment alter the availability of skill.

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. Also, we would expect this effect to be concentrated in lower-skilled firms, those that did not tend to post for skilled workers in a tighter labor market. However, we next show that the opposite is true.

In table 8, we augment our primary specification with controls for the average skill requirements of a firm in 2007 and allow these to interact with inverse employment growth. In column 1, we regress the share of ads in a firm-MSA-date with an education requirement on inverse employment growth, the firm's average education requirement in 2007 (normed across ads to be mean zero and standard deviation one), and an interaction. We include all other controls from equation 1.⁴⁷

In column 1, the main effect on inverse employment growth, 0.075, is the upskilling effect in education for a firm with the average skill requirement in 2007 and is quite similar to that found in other samples. The interaction effect of 0.038 means that a firm with a one-standard-deviation higher skill requirement in 2007 is 50% more cyclical than a firm with average skill requirements. The estimates in column 2 show that we obtain commensurate results when we split firms into two groups based on their 2007 requirements. Furthermore, the results in columns 3 and 4 indicate even greater cyclicity for experience requirements among the initially more-skilled firms.

This evidence does not comport with the primary mechanism underlying upskilling being opportunistic behavior among less-skilled firms. If in a recession firms take advantage of a more plentiful base of higher-skilled workers willing to do lower-quality jobs, we would expect upskilling to occur disproportionately at the lower end of the firm skill distribution, which is the opposite of what we find. It seems unlikely that the higher-skilled firms would be the ones acting opportunistically to attract more-skilled workers than they ordinarily could. We do know from past research that when workers match to jobs in recessions they tend to downgrade to a job typically held by a worker in a lower skill level (Devereux 2002). Based on the results from table 8, we infer that these downgrades are primarily driven by workers expanding their search (a supply-side effect), and not by opportunistic firms trying to hire

⁴⁷Naturally, for this regression, we must restrict the sample to firms with at least one ad in 2007. As shown in column 5 of appendix table A1, the probability of satisfying this criterion does not vary with economic conditions, and column 5 of appendix table A2 shows that our primary upskilling result is similar in magnitude to that from this subsample.

Table 8: Upskilling and Firm 2007 Skill Requirements

	(1)	(2)	(3)	(4)
	Education Requirement		Experience Requirement	
Inverse Emp Growth (E)	0.0748*** (0.0279)	0.0585** (0.0284)	0.0888*** (0.0288)	0.0500* (0.0298)
E*(2007 Firm Ave)	0.0377*** (0.00269)		0.0691*** (0.00353)	
2007 Firm Ave	0.176*** (0.00260)		0.127*** (0.00220)	
E*(Above Median Firm)		0.0213*** (0.00358)		0.0651*** (0.00444)
Above Median Firm		0.189*** (0.00302)		0.135*** (0.00281)
# Firm-MSA-Date Cells	5,719,492	5,719,492	5,719,492	5,719,492
R-Squared	0.180	0.152	0.125	0.115

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

Notes: See table 2. Sample is restricted to firms with at least 1 posting in 2007. "2007 Firm Ave" is the average education (columns 1-2) or experience (columns 3-4) requirement of the firm in 2007, normed to be mean 0 and standard deviation 1 in the weighted 2007 sample. Above Median Firm is an indicator that equals one if the firm had an above-median skill requirement in the 2007 weighted sample.

unusually highly-skilled labor (a demand-side effect).⁴⁸

Quits and Replacement Hiring

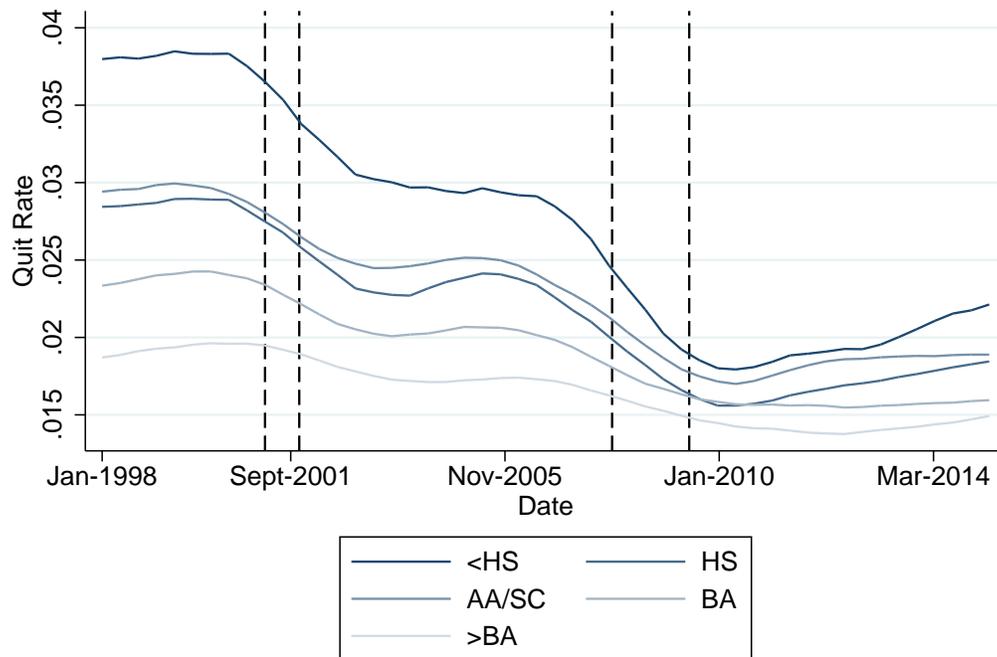
We have argued that the persistence in upskilling is key to disentangling cyclical explanations from structural ones. However, it could be that cyclical explanations (such as the bottleneck hypothesis) account for early upskilling and that labor supply factors drive upskilling later in the period. For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

However, we see no evidence of differential recovery in quit rates across skill groups. For instance, in figure 5 we plot smoothed time series of quit rates by education group from 1998–2015 using longitudinally-linked CPS data.⁴⁹ Besides the large and well-known secular

⁴⁸Although 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), the results presented here help rule out that a binding minimum wage causes firms to seek more-skilled workers as a means of lowering the effective skill price in a downturn. Firms on the top end of 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).

⁴⁹We define a quit as a worker who reported switching employers between month t and month $t + 1$ (but

Figure 5: Quit Rates by Date and Education Group



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

decline in quits (Molloy, Smith, and Wozniak 2013), the figure shows the quit rate has a clear cyclical component, declining during NBER-dated recessions (indicated with vertical dashed lines) and recovering somewhat shortly thereafter. There is certainly no evidence in this figure that quit rates for higher-skilled workers (lighter lines) recover more quickly than those for lower-skilled workers (darker lines); in fact, the opposite appears more likely, with less-educated workers showing faster recoveries. It thus seems unlikely that replacement hiring is driving the persistence in upskilling.

Educational Attainment

Firms may decide to upskill if they observe that skilled workers have become more plentiful in their local economy. Indeed, a long line of research has explored whether educational attainment responds to local labor market conditions. Past evidence generally finds moderate effects on enrollment, but limited effects on attainment.⁵⁰ One exception is Charles,

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 .

⁵⁰Card 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.

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.⁵¹ Furthermore, to the extent that housing-bubble locations also experienced a larger employment shock in the Great Recession, our controls for MSA fixed effects and MSA-specific time trends will account for the preexisting changes in educational attainment.

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 estimate regressions similar to that shown in equation (3), regressing state-year educational attainment on states' peak-to-trough (inverse) Bartik employment shocks (analogous to figure 1), allowing this shock to interact with year. We focus on a young population (aged 18–32) whose educational attainment decisions should be most malleable. Our estimates (available upon request) imply that a state experiencing a Great-Recession-sized shock has no change in the probability that its young population attains at least some college and a (weakly) lower probability that it attains at least a bachelor's degree. Thus we conclude that changes in educational attainment induced by the Great Recession are unlikely to be driving our results.

6 Conclusion

During the recovery following the Great Recession, anecdotal evidence suggested that the composition of new hires shifted towards higher-skilled workers, resulting in many workers being “overeducated” for their jobs. However, it remained unclear how broad or deep these effects were, and the extent to which they were driven by labor supply or labor demand responses. In particular, workers may have applied to jobs lower in the skill distribution than they would have during a tighter labor market, but also firms may have tried to exploit the limited job opportunities of more-skilled workers to recruit from higher in the skill distribution than they would have before the recessionary period. Additionally, firms may have taken the opportunity of the recession as a time of “cleansing” to restructure their production in a manner consistent with routine-biased technological change. The magnitude of upskilling, and whether it is temporary or more permanent, has implications for our understanding of technological change and business cycles, as well as labor market and social insurance policies. In this paper we provide comprehensive, broad-based evidence of upskilling – firms demanding higher-skilled workers when local economic conditions are

⁵¹For 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.

worse. We argue that this effect is primarily driven by an episodic restructuring on the part of firms and thus is likely to be long-lasting.

Using a novel data set that captures detailed information on the near-universe of electronic job postings and that spans the Great Recession, we show that in metropolitan areas with larger employment shocks, job ads were 6 percentage points (nearly 12%) more likely to list education and experience requirements, and that the share of ads specifying at least a bachelor’s degree increased by 42%. These patterns are robust to a variety of alternative specifications, and even hold within firm-MSA-occupation cells. Moreover, the richness of our data allow us to examine how much of this demand-side upskilling is driven by changing composition of the firms advertising and the occupations they advertise for. Remarkably, we find that the vast majority is driven by upskilling within firm-occupation cells. Furthermore, we find a similar upskilling effect in Current Population Survey data on new job starts, suggesting that firms successfully hire the more-skilled workers that they seek.

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

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

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

cognitive workers still retain a substantial advantage over the low-skilled. They are drawn into formerly middle-skill jobs, which are becoming higher-skilled. We can thus explain why skilled workers still earn a premium in the labor market even though the returns to cognitive occupations appear to have diminished.

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

References

- [1] Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer (2016), "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success," *Journal of Labor Economics*, 34(S1 Part 2): S361–S401.
- [2] Autor, David (2014), "Polanyi's Paradox and the Shape of Employment Growth," NBER working paper #20485.
- [3] Autor, David, 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.
- [4] 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.
- [5] Autor, David, David Dorn, Gordon Hanson, and Jae Song (2014), "Trade Adjustment: Worker-Level Evidence," *Quarterly Journal of Economics*, 129(4): 1799–1860.
- [6] Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006), "The Polarization of the U.S. Labor Market," NBER Working Paper #11986.

- [7] 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.
- [8] 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.
- [9] 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.
- [10] Bartik, Timothy J. (1991), *Who Benefits from State and Local Economic Development Policies*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
- [11] 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.
- [12] 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.
- [13] Blanchard, Olivier Jean and Larry Katz (1992), “Regional Evolutions,” *Brooking Papers on Economic Activity*, 1–75.
- [14] Berger, David (2012), “Countercyclical Restructuring and Jobless Recoveries,” mimeo.
- [15] Burning Glass Technologies (2014), “Moving the Goalposts: How Demand for a bachelor’s Degree is Reshaping the Workforce,” mimeo.
- [16] Caballero, Ricardo, and Mohamad. Hammour (1994), “The Cleansing Effect of Recessions,” *American Economic Review*, 84(5): 1350–1368.
- [17] Caballero, Ricardo, and Mohamad Hammour (1996), “On the Timing and Efficiency of Creative Destruction,” *Quarterly Journal of Economics*, 111(3): 805–852.
- [18] 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*.
- [19] 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.

- [20] 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.
- [21] 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.
- [22] 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.
- [23] Charles, Kerwin K., Erik Hurst, and Matthew J. Notowidigdo (2015), “Housing Booms and Busts, Labor Market Opportunities, and College Attendance,” mimeo.
- [24] 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.
- [25] Clemens, Jeffrey (2015), “The Minimum Wage and the Great Recession: Evidence from the Current Population Survey,” NBER working paper #21830.
- [26] 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.
- [27] Deming, David J. (2015), “The Growing Importance of Social Skills in the Labor Market,” NBER Working Paper #21473.
- [28] Deming, David J., and Lisa B. Kahn (2016), “Firm Heterogeneity in Skill Demands,” mimeo.
- [29] Devereux, Paul (2002), “Occupational Upgrading and the Business Cycle,” *Labour Economics*, 16(3): 423–452.
- [30] Foote, Christopher L., and Richard W. Ryan (2015), “Labor Market Polarization Over the Business Cycle,” *NBER Macroeconomics Annual*, 29: 371–413.
- [31] 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.
- [32] Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (2001), “Equilibrium Unemployment,” *Journal of Monetary Economics*, 48(1): 109–152.

- [33] Hall, Robert (1991), “Labor Demand, Labor Supply and Employment Volatility,” in NBER Macroeconomics Annual. Cambridge: MIT Press.
- [34] Hall, Robert (2005), “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 95(1): 50–65.
- [35] Hoynes, Hilary, Douglas L. Miller and Jessamyn Schaller (2012), “Who Suffers During Recessions?”, *Journal of Economic Perspectives*, 26(3): 27–48.
- [36] 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.
- [37] 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>.
- [38] 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>.
- [39] Kahn, Lisa B. (2010), “The Long-Term Consequences of Graduating from College in a Recession,” *Labour Economics*, 17(2): 303–316.
- [40] Kahn, Lisa B., and Erika McEntarfer (2015), “Employment Cyclicalities and Firm Quality,” NBER working paper #20698.
- [41] 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.
- [42] Kuhn, Peter, and Kailing Shen. (2013), “Gender Discrimination in Job Ads: Evidence from China,” *Quarterly Journal of Economics*, 128(1): 287–336.
- [43] Marinescu, Ioana (2014), “The General Equilibrium Effects of Unemployment Insurance: Evidence from a Large Online Job Board,” mimeo.
- [44] 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.
- [45] 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.
- [46] Moffitt, Robert (2012), “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations,” *Brooking Papers on Economic Activity*, fall.

- [47] Mortensen, Dale. and Christopher Pissarides (1994), “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61(3): 397–415.
- [48] Moscarini, Giuseppe and Fabien Postel-Vinay (2013), “Stochastic Search Equilibrium,” *Review of Economic Studies*, 80(4): 1545–1581.
- [49] 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.
- [50] 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
- [51] Notowidigdo, Matthew J. (2013), “The Incidence of Local Labor Demand Shocks,” University of Chicago, mimeo.
- [52] Oldenski, Lindsay (2012), “Export versus FDI and the Communication of Complex Information,” *Journal of International Economics*, 87(2): 312–322..
- [53] 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.
- [54] Pierce, Justin R., and Peter K. Schott (2015), “Trade Liberalization and Mortality: Evidence from U.S. Counties,” Yale mimeo.
- [55] Rothwell, Jonathan (2014), “Still Searching: Job Vacancies and STEM Skills,” Brookings mimeo.
- [56] Sasser Modestino, Alicia, Daniel Shoag and Joshua Ballance (2015), “Upskilling During the Great Recession: Do Employers Demand Greater Skill when Skilled Workers are Plentiful?”, mimeo.
- [57] Shimer, Robert. (2004), “The Consequences of Rigid Wages in Search Models,” *Journal of the European Economic Association*, 2(2-3): 469–479.
- [58] Schumpeter, Joseph. (1939), *Business Cycles*. New York: McGraw-Hill.
- [59] Smith, Christopher L. (2014), “Using Cross-State Variation to Assess the Potential for Additional Improvement in Measures of Labor Market Conditions,” *FEDS Notes*, June 2.
- [60] Solon, Gary, Robert Barsky, and Jonathan A. Parker (1994), “Measuring the Cyclicity of Real Wages: How Important Is Composition Bias?” *Quarterly Journal of Economics*, 109(1): 1–25.

- [61] U.S. Department of Labor (1977), *Dictionary of Occupational Titles, 4th edition*. Washington, D.C.: U.S. Department of Labor.
- [62] 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.
- [63] 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.

A Appendix

It is estimated that as of 2014 between 60% and 70% of all job postings are found online (Carnevale, Jayasundera, and Repnikov 2014). Indeed, The Conference Board discontinued its long-running, print-based Help-Wanted Advertising Index in 2008, after having begun a Help-Wanted Online Index in 2005 (HWOL).⁵² 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.⁵³ 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.⁵⁴

Figure A1 plots the distribution of BG ads across major industry groups, sorted from largest to smallest (green bars), as well as the distribution of job vacancies in JOLTS (purple diagonal-lined bars). As mentioned, the BG database is meant to capture only electronically posted job ads; the universes of the data sources are thus not identical, but JOLTS is the best comparison available.⁵⁵ Despite the sample differences, the industry distributions match each other reasonably well. BG is overrepresented in health care and social assistance, as

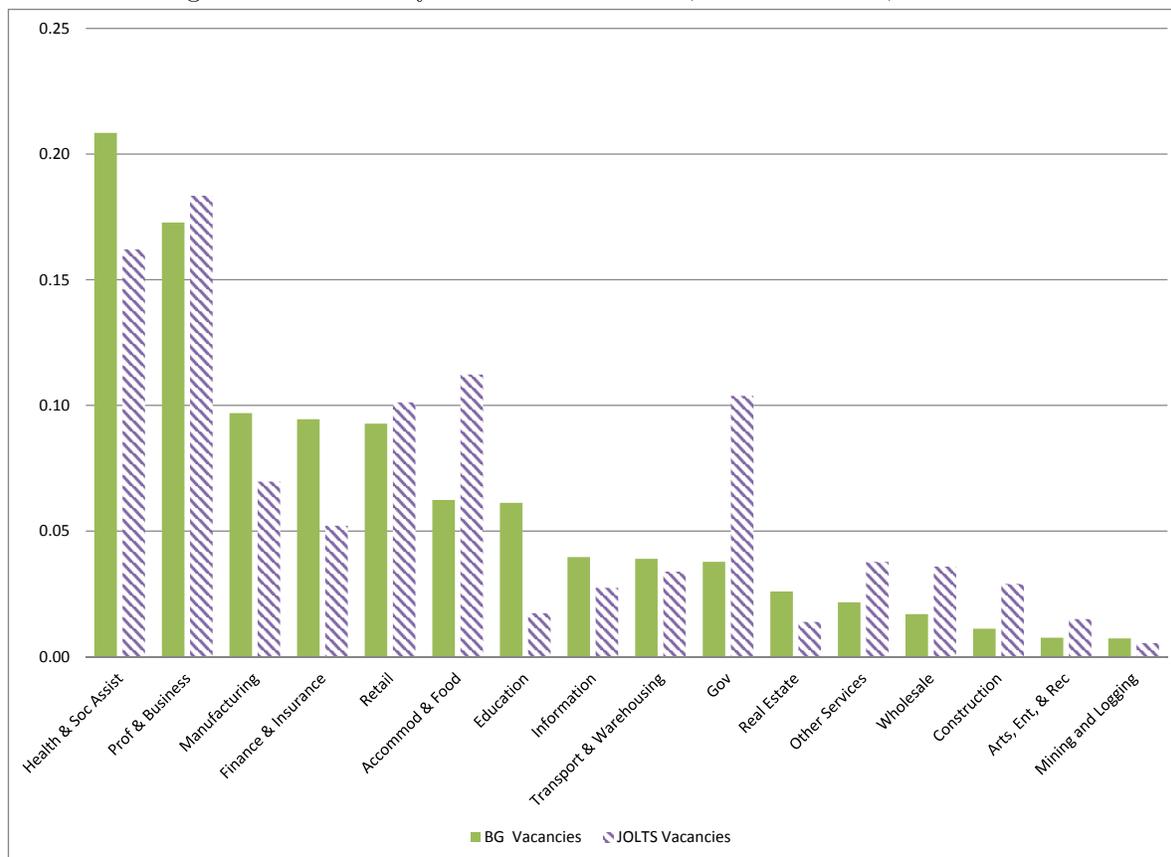
⁵²See <https://www.conference-board.org/data/helpwantedonline.cfm>.

⁵³Rothwell (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.

⁵⁴Furthermore, 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.

⁵⁵Both data sets cover 2007 and 2010–2014. JOLTS data are based on a monthly, nationally representative sample of approximately 16,000 business establishments drawn from unemployment insurance records; they count as a vacancy or job opening any position (including temporary and seasonal ones) for which work could start within 30 days and that the employer is actively trying to fill through a variety of means, of which posting a job ad (electronic or otherwise) is only one.

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



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

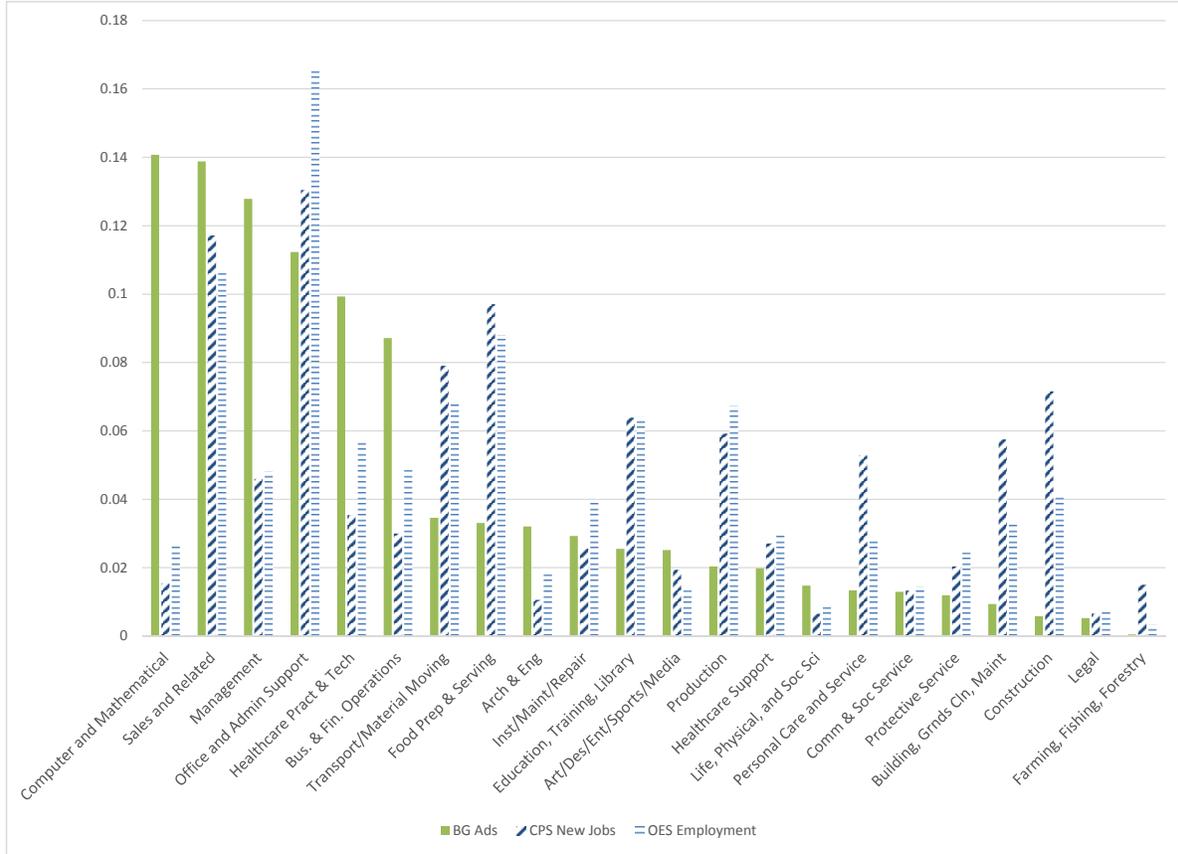
A great advantage of the BG data over the JOLTS is that they allow us to categorize jobs by occupation at a detailed level. We thus also compare the occupational distribution of BG job ads to both the stock and flow of employment in the U.S. We should not expect online job ads to precisely match either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

Figure A2 plots the distribution of BG ads across major occupation groups, sorted from largest to smallest (green bars).⁵⁶ We show the distribution of the stock of employment based on the Bureau of Labor Statistics' Occupational Employment Statistics (OES) data (light blue, horizontal lines). We also show the occupational distribution of new job starts (job flows) based on longitudinally linked Current Population Survey (CPS) data (dark blue, diagonal lines).⁵⁷

⁵⁶For clarity, we use 2-digit Standard Occupational Classification codes in the figure. The regression analyses use more granular codings.

⁵⁷All data sets cover 2007 and 2010–2014. (2014 is the most recent date for which OES data are available.)

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



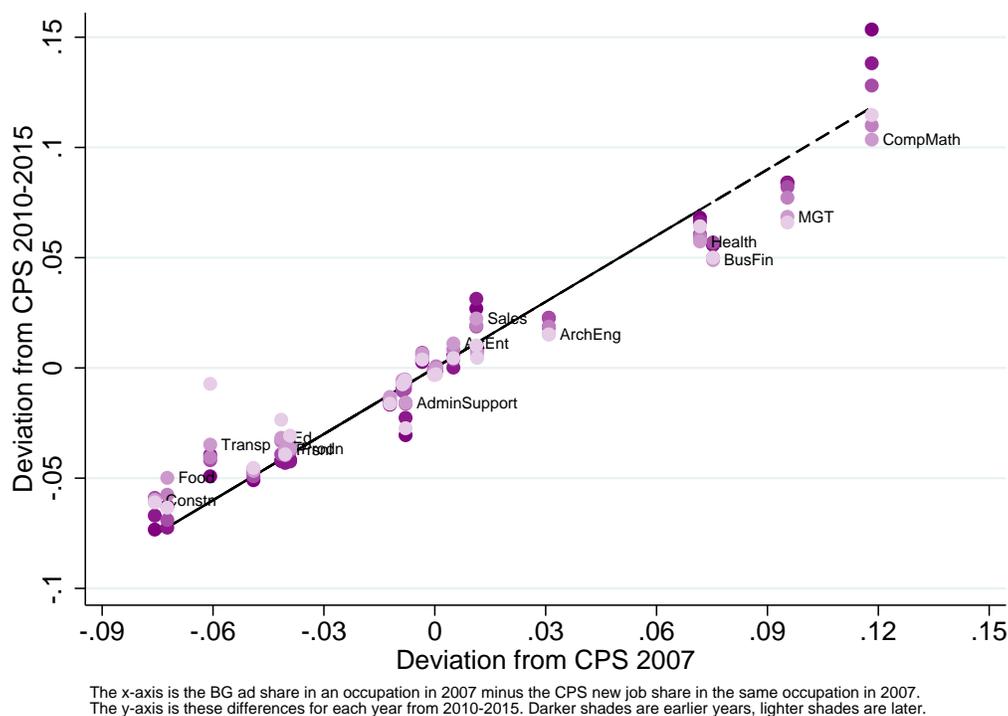
Perhaps not unexpectedly, BG has a much larger representation of computer and mathematical occupations, more than four times the OES and CPS shares. BG is also overrepresented among management, healthcare practitioners, and business and financial operations, although to lesser degrees. On the other hand,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

A.2 Representativeness of BG Data over Time

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

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

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



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

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

A.3 Compustat Sample

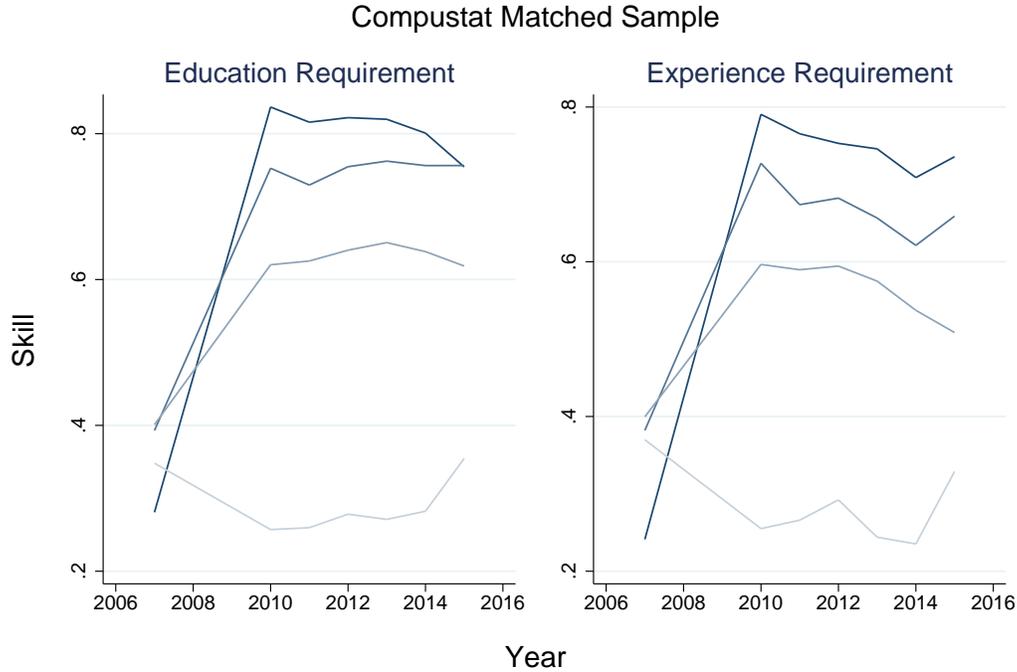
Beginning with the 10,436 firms used to construct the 2007–2010 change sample (7.4% of firms and 67% of weighted observations, used in figure 2), we employ a sequential matching procedure. We first match based on exact name, after cleaning to remove punctuation and words that are sometimes abbreviated (e.g., “Inc.”). This step accounts for 81% of all Compustat-matched firms and 76% of weighted observations. Second, we use a fuzzy-match program called Winpure and assign a match if the program determines at least a 93% probability of a match (5% of matched firms and matched weighted observations). Third, we add the sample of firms matched by Deming and Kahn (2016), which uses only BG firms posting in 2014. We thus match a total of 46% of weighted observations and 18% of firms from the 2007–2010 change sample, or 30% of weighted observations in the sample as a whole – though we did not attempt to match observations that did not meet the 2007–2010 change criteria.⁵⁸

The probability that an ad is posted by a firm that matches to Compustat is reported in column 3 of appendix table A1. We find that it does vary significantly with local labor market conditions. Our estimate implies that an ad posted in an MSA experiencing a Great-Recession-sized employment shock is 5 ppts more likely to be to a Compustat matched firm, relative to ads posted in the MSA before the recession and ads posted in other MSAs experiencing a lesser recession. This is likely because large, successful firms were better able to weather the recession.

Importantly, however, our key findings hold in this sample. Column 4 of appendix table A2 shows that the main upskilling effects are observed at similar magnitudes, and are not statistically different from our baseline sample. Furthermore, despite the fact that firms in Compustat are all relatively large and successful, we can still pick up important differences in their upskilling behavior. Figure A4 replicates figure 2 but for the Compustat Sample, plotting skill requirements by year across firms grouped into quartiles by their degrees of upskilling between 2007 and 2010. The pattern is very similar. Firms that increase skill requirements more between 2007 and 2010 (darker lines) maintain their elevated skill levels for the entire sample period. They look persistently different than the Compustat firms that did not increase skill requirements. These findings support our analysis of changes in capital stock among firms in the Compustat sample, suggesting that it will also be informative about the drivers of upskilling more generally.

⁵⁸For 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: Skill Requirements by Firm 2007-2010 Change – Compustat Sample



This graph plots average skill requirements by year and firm type. Firms are split into quartiles based on their change in skill requirements from 2007-2010. Darker shades indicate quartiles with larger changes.

Table A1: Probability of Being Missing from a Subsample

	(1)	(2)	(3)	(4)	(5)
	Firm	2007-2010 Change	Compustat	Skills	Firm posted in 2007
Inverse Emp Growth	-0.00780 (0.0411)	0.00809 (0.0268)	0.0523*** (0.0181)	0.0229 (0.0172)	0.0164 (0.0251)
# Occ-MSA-Date Cells	2,300,382	1,966,837	1,966,837	1,966,837	1,966,837
R ²	0.273	0.333	0.083	0.174	0.453

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

Notes: See table 2. Column 1 includes all ads; firm sample is ads with a non-missing firm that has posted at least 10 ads over 2007 and 2010-2015. Columns 2-4 restrict to the firm sample. "2007-2010 change" sample restricts to firms that post at least 5 ads in each of 2007 and 2010. Compustat sample further restricts to firms that can be matched to Compustat. "Skills" sample includes ads with at least one specific skill listed. "Firm posted in 2007" sample contains firms with at least 1 post in 2007.

Table A2: Skill Requirements and Local Labor Market Conditions – Subsamples

	(1)	(2)	(3)	(4)	(5)
Panel A: Education Requirements					
Inverse Emp Growth	0.0632** (0.0286)	0.0743** (0.0303)	0.0625** (0.0294)	0.0444 (0.0331)	0.0662** (0.0301)
# Occ-Msa-Date Cells	1,966,837	2,300,382	1,722,098	1,218,221	1,763,773
R-Squared	0.189	0.190	0.189	0.142	0.193
Panel B: Experience Requirements					
Inverse Emp Growth	0.0571** (0.0289)	0.0753** (0.0310)	0.0663** (0.0287)	0.0623* (0.0320)	0.0724** (0.0293)
# Occ-Msa-Date Cells	1,966,837	2,300,382	1,722,098	1,218,221	1,763,773
R-squared	0.203	0.217	0.185	0.140	0.188
Includes missing firms		X			
2007-2010 change			X	X	
Compustat				X	
Firm posted in 2007					X

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

Notes: See table 2. "2007-2010 change" sample restricts to firms that post at least 5 ads in each of 2007 and 2010. Compustat sample further restricts to firms that can be matched to Compustat. Column 5 restricts to firms with at least 1 post in 2007.