Comments on "Labor Market Polarization Over the Business Cycle" by Foote and Ryan

Fatih Guvenen

University of Minnesota and NBER

August 6, 2014

1 Overview of Paper

One of the striking patterns in employment losses observed during the Great Recession (2007–2009) was its concentration among middle-skill workers and occupations. Interestingly, these were the very same types of workers whose employment share has been declining since the 1990s. This paper aims to understand if these two phenomena—the trend and the cycle in middle-skill employment—are related. Specifically, the paper investigates two alternative hypotheses. One possibility, proposed and analyzed concurrently to this paper by Jaimovich and Siu (2014), is that the two are related: large middle-skill employment losses in recessions result from the optimizing behavior of workers and firms in a search and matching framework in which the weak long-run trends in job prospects combines with a deep recession that makes it optimal for workers to quit and switch to other occupations with better long-run prospects. The second and simpler hypothesis is that middle-skill jobs happen to be predominantly located in cyclical industries (such as manufacturing and construction) and therefore are bound to shrink more during deep recessions. Overall, the evidence presented in the paper is more supportive of the second hypothesis with scant evidence in favor of the first. Along the way to testing these two hypotheses, Foote and Ryan also uncover a variety of simple but interesting empirical findings that are valuable contributions on their own.

I begin by reviewing what I view as the three main contributions of the paper and

then provide more detailed comments on each. The contributions of the paper can be grouped under three main headings.

I. New Historical Time Series on Occupational Employment

To explore the second hypothesis—that middle skills jobs could be concentrated in cyclical industries—the authors need to analyze historical time series of employment and unemployment by occupational groups. The official published occupation-level employment data is available only since January 1983. The authors extend these data all the way back to 1947 by constructing time series for quarterly employment by occupational categories by painstakingly going through published reports by the Census bureau. Similarly, official unemployment data by occupation is available only since January 2000, and the authors extend this back to 1957. This contribution is valuable independently of the rest of the paper and will provide useful input into a variety of academic and policy work. For example, using these time series, Foote and Ryan show that routine manual occupations have always been very volatile and cyclical (Figure 3) compared with others. Moreover, since 1950, routine manual employment has been shrinking almost monotonically, replaced with non-routine cognitive employment (Figure 4).[1] The long term perspective provided by these new data also allow the authors to conduct the dynamic factor analysis, discussed next.

II. Dynamic Factor Analysis

To investigate the second hypothesis, the authors conduct a simple dynamic factor analysis (Section 2). Basically, they answer the question: are the occupation-level employment fluctuations observed during the Great Recession consistent with past business cycles? The answer turns out to be largely yes. The substantial drop in middle-skill employment in recent recessions is almost entirely consistent with the cyclical sensitivity displayed by this group to historical GDP fluctuations, going back to the 1950s.

Specifically, the authors estimate a dynamic factor model with two equations ((1) and (2) in the paper): an autoregressive process for a latent factor in which GDP enters

---

[1]As I discuss below, I have some concerns about some aspects of the routine/nonroutine classification adopted in this literature and in this paper. But I still believe the construction of these new data series is a step forward.

as a forcing variable, and a second equation which links the employment level of a given skill group (routine/nonroutine, manual/cognitive) to the latent factor and includes controls for the manufacturing and construction share of the given skill group. These two equations are estimated using data for a given time period and are then used either to examine the in-sample fit or for making out-of-sample forecasts. The main finding is that the estimated model fits the behavior of manual occupations and especially routine-manual occupations well. The model predicts almost perfectly the magnitude of the drop in routine-manual employment during the last three recessions, even when it is estimated using only data before 1985.[2] The model fits less well the fluctuations in high skill employment, because this group displayed little business cycle variability before the 1980s and has become more cyclical since then (a fact previously documented by Castro and Coen-Pirani (2008)).

Overall, the conclusions of this analysis provide support to the idea that the Great Recession is not a fundamentally different type of recession when it comes to employment fluctuations. It is just a much deeper recession, but is otherwise similar in structure. This finding largely echoes the conclusion of Stock and Watson (2012) who conducted a large-scale VAR analysis of hundreds of macro variables and found that the Great Recession did not represent a structural break from previous US business cycles. This paper extends this analysis to employment fluctuations by skill categories to reach a similar conclusion.

III. Flows Through Unemployment by Skill

Turning to the first hypothesis, the authors conduct several exercises. First, they examine the labor market flows by skill type. That is, where do middle skill workers transition either through job to job transitions or going through an unemployment spell first? To put this in context, it is useful to note that this is one of the main mechanisms at work in Jaimovich and Siu (2014), where middle-skill workers who realize during recessions that their future job prospects don't look to good choose to become unemployed, retrain, and search jobs in high-skill occupations where demand is strong. The present paper does not find any strong evidence to back up this mechanism. Instead, most of the flows from middle-skill jobs is to other middle-skill jobs or into non-participation.

---

[2] The level specification requires HP filtering, which can cause end point problems—coinciding with the Great Recession! The authors are aware of this, so they also estimate their model in first differences, which broadly confirms the findings from the levels specification described here. See Figure 8.

The non-participation margin is interesting because it is an often underemphasized dimension in the polarization debate. In Figure 14, which is one of the more interesting findings of the paper, the authors show that the large decline in the labor force participation rate of prime-age males since the 1980s is mainly concentrated among middle-skill workers, who increasingly left the labor market during this period.

I wish the authors had expanded a bit more on this point because it is an intriguing idea that can shed light on the broader debate on the participation trends in the last 30 years. In particular, the analysis in this section begs a true panel data structure (as opposed to synthetic panels used here), where one can track the same workers over long periods of time and see more clearly what kind of labor market history leads to their separation from the labor force. For these purposes, one could use the Longitudinal Employer-Household Dynamics (LEHD) project of the Census Bureau, which provides a very rich set of variables for a large number of US individuals. It has many appealing features that would be especially useful in this context: quarterly earnings and employment data, data on employers and locations, among many others. I hope the authors consider pursuing this research agenda in future work.

In the rest of my comments, I expand on three issues to provide context for the analysis in this paper. These are in turn:

- Wage- versus job-polarization can reveal different patterns

- How useful is the routine/nonroutine classification?

- A need for a theory of endogenous routinization.

2   Wage- vs. Job-Polarization

To understand the discussion about job polarization, which is the focus of this paper, it is useful to take a step back to see what came before it. In an influential paper, Juhn et al. (1993) documented a striking feature of the rising wage inequality between 1963 and 1988: inequality rose due to an almost perfect stretching out of the entire wage distribution over time. In other words, inequality grew in every part of the distribution because the higher a wage percentile was in 1963, the more wages grew in that percentile between 1963 and 1988. I shall refer to this phenomenon as "homogeneity." The left

Figure 1: Monotonic Rise in Inequality vs. Polarization

(a) Homogeneity in Inequality

(b) Wage Polarization
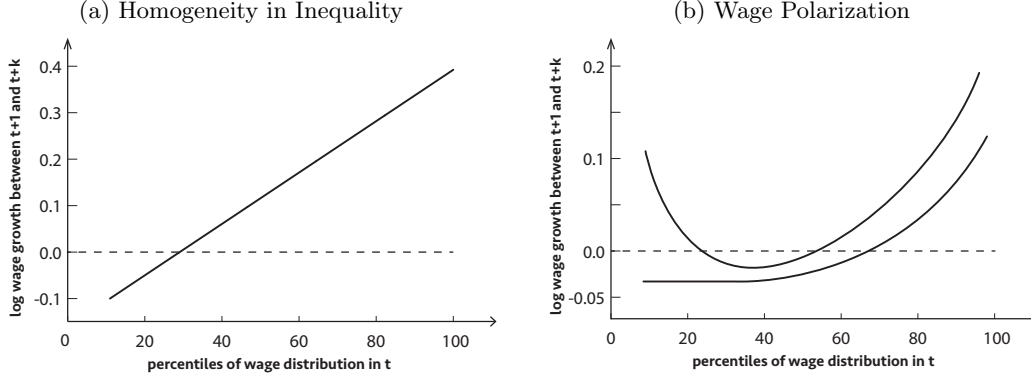
panel of Figure 1 plots the finding in Juhn et al diagrammatically, where this homogeneity can be seen.

Homogeneity can be generated by a simple model with a single skill index or a one-factor structure:

$$\log w_t^i = p_t \times \log s^i \Rightarrow \log w_t^i - \log w_t^j = p_t \times (\log s^i - \log s^j),$$

where i and j denote any two workers, s^i is a fixed skill index of worker i and p_t is the aggregate price of skill in period t. One can see that the log wage gap between any two workers will increase linearly as a function of the skill price, leading to a stretching out of the entire distribution with a rise in p_t.

Coming into the 2000s, however, subsequent papers that repeated Juhn et al's analysis with more recent data came upon a surprising realization: homogeneity was replaced by polarization during the 1990s. That is, while the wage distribution was stretched above the median of the distribution, it was in fact getting compressed below it. In other words, wage growth during the 1990s was faster at low wage percentiles compared with middle wage percentiles, shrinking the gap between the two. Therefore, what was once an upward sloping graph, became U-shaped during the 1990s. The right panel of Figure 1 shows two examples of how this may look like. Searching for an explanation, Autor et al. (2003) found that a similar U-shaped pattern emerged in employment growth when occupations were ranked by their average wages (at the beginning of the period) or by

the educational attainment of the workers in each occupation. This latter observation was referred to as "job polarization" to distinguish from "wage polarization."

The subsequent literature on polarization seems to have mostly focused attention on job polarization, as does the present paper. But studying employment without wages is likely to miss the big picture, for at least two reasons. First, all jobs are not created equal: there are "good jobs" and "bad jobs", for example, as Acemoglu (2001) called them. Thus, if one job were to be destroyed and two new jobs were created in a given year, it is hard to fully understand the implications without also knowing the associated wages. If the lost job paid $30 per hour whereas the two new jobs only paid $8 per hour each, this is a different picture than if the lost and created jobs were identical in the wages they pay. Second, the joint evolution of wages and employment (prices and quantities) provides much sharper insights into the driving forces of the observed phenomena. With employment data alone, it is not possible to tell if the underlying forces are on the demand side or on the supply side. One reason I am concerned about lacking data on wages is because after the 1990s there seem to be far less evidence of wage polarization than there is job polarization, as I document next.

No Wage Polarization in the 2000s

To complement the picture on employment, let us take a look at wage growth by wage percentiles.[3] Figure 2 plots a dynamic version of the figure in Juhn et al, by ranking workers at time t based on their average wage earnings in the previous 5 years ($t-5$ to $t-1$), and then computing the average wage *growth* for each percentile group going forward (from t to $t+k$ for some k to be determined).[4] The top panel performs this analysis when the future period t to $t+k$ covers one of the three full expansions experienced since the 1980s, namely the 1983–1990, 1992–2000, and 2002–2007 periods. The y-axis of these graphs have been adjusted by taking out age effects, so we should not read too much into the level (intercept) of this graph. Instead, we focus on the differences between groups ranked along the x-axis. Wage polarization in the 1990s is seen clearly in the top panel: workers in the bottom 10% of the past average earnings (PAE) distribution experienced

---

[3]Throughout this discussion I will refer to full year wage earnings as wages. It would be desirable to repeat the analysis here with hourly wages, but I suspect one would reach a similar conclusion.

[4]Plotting the same version as Juhn et al. makes no substantive difference to the conclusions drawn here.

6

Figure 2: Wage Earnings Growth, by Percentiles of Past Average Earnings Distribution. (Source: Guvenen et al. (2014))

(a) Expansions
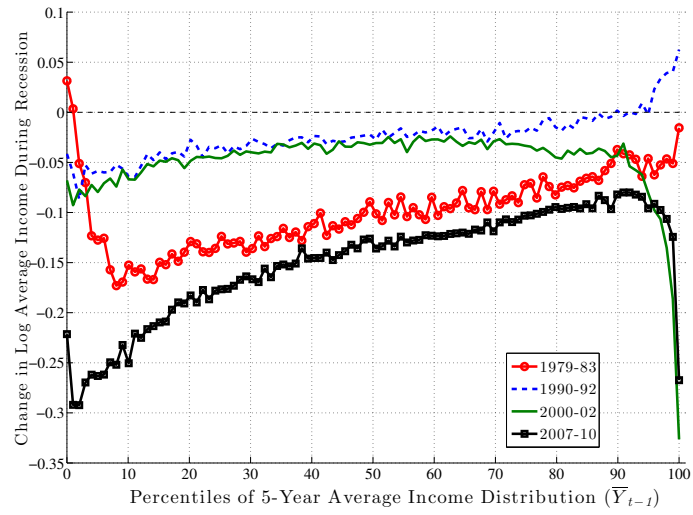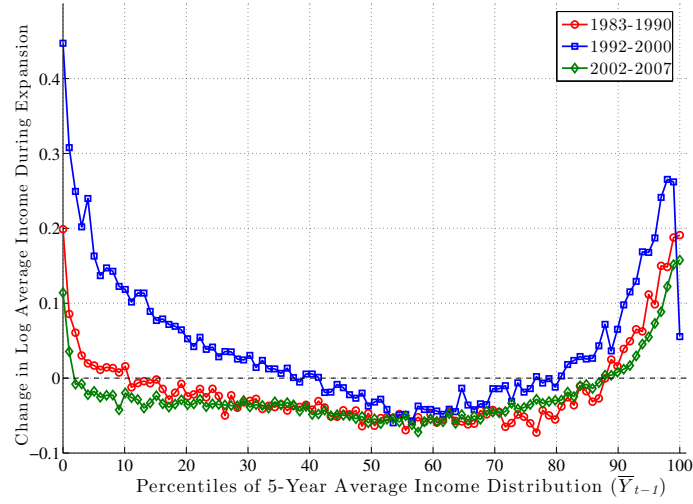
(b) Recessions

an average earnings growth that was more than 15% higher relative to workers at the median of the PAE distribution. We see much less evidence of this in the 1980s and 2000s expansions.

Turning to the bottom panel, we see *no* evidence of wage polarization during any one of the four recessions in the last 30 years. The deep recessions of 1980–1983 and the Great Recession of 2007–2010 display substantial homogeneity—a stretching out of the entire wage distribution. For example, during the Great Recession, workers in the bottom 10% of the PAE distribution saw their earnings fall by 12% more than those at the median. A similar picture emerges for the double-dip recession. Consequently, looking at the 2000s decade as a whole (combining the two recessions with the brief expansion) there is scant evidence of polarization in wage earnings during this time.

The picture that emerges from the previous discussion is that wage polarization is seen primarily during expansions and the long and strong expansion of the 1990s led to a period of rising wages at the bottom relative to the median. However, recessionary periods show no sign of such polarization, in contrast to the finding of job polarization by skill level documented in this paper.

To investigate if the issue is employment versus wages, we can plot a similar figure to those above, by ranking workers by their wage earnings in a given year and then computing the *employment rate* for workers in each group in the subsequent years. We consider a worker employed if his annual wage income in a given year exceeds a threshold, which is defined as the earnings level corresponding to one quarter of full-time work at half the legal minimum wage in that year.[5] Figure 3 plots the employment rate during the Great Recession for workers ranked by their income in 2006.[6] Employment rate is more or less flat above the median of the PAE distribution, but then declines precipitously as we move to the left: the employment rate is 97% at the median but only 80% for those at the 10th percentile of the PAE distribution.

To summarize, even though the emphasis in this paper and in the rest of the polarization literature is on middle-*skill* job losses, those who actually lose their jobs and and/or suffer wage income losses are not middle-*income* workers but rather low-income workers. In other words, there is no evidence that earnings losses in the Great Recession

---

[5]This definition is quite standard in the literature, going back to Juhn et al. (1993) and others.

[6]Ranking workers by PAE over 2001–2006 yields a very similar picture.

Figure 3: Male Employment Rate, by PAE Distribution, Great Recession
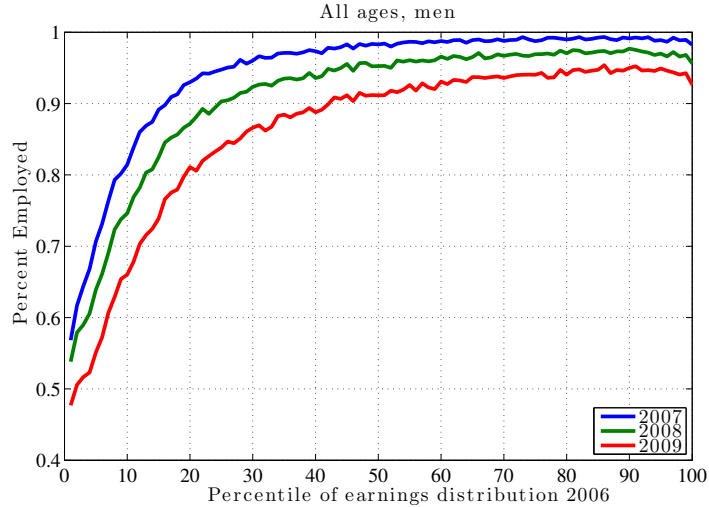
were concentrated in the middle; instead we witnessed remarkable homogeneity both in wage and employment losses.

So how do we reconcile the two sets of findings? One possibility is that those workers who suffer losses are simultaneously middle skill *and* low income. To see if this might be the case, it is useful to look at the wage distribution by skill type. Foote and Ryan have kindly produced the histograms of log hourly wages by skill type, shown in Figure 4. As seen here, the wage distribution for high skill workers (shown in dark) lies clearly to the right of all others, confirming that high-skill workers indeed are those with much higher wage earnings than the other groups. However, the densities for the other three groups—Routine cognitive, routine manual, and non-routine manual—overlap over the entire domain of wages, indicating only small differences in earnings across these groups. Thus, while routine manual workers are making slightly higher wages on average (by less than 10%) than non-routine manual workers, it is entirely possible that those who lose their jobs or suffer larger earnings losses are concentrated at the bottom end of their respective wage distributions, which generates the figures of earnings losses observed in the previous figures. In other words, a low-wage cashier job (classified as routine manual) and a low-wage waitress job (classified as non-routine manual) suffer similar fates to each other.

An important corollary to this explanation is pre-recession wage levels are a better

Figure 4: Histogram of Log Wages by Skill Type. (Source: Foote and Ryan (2014)).
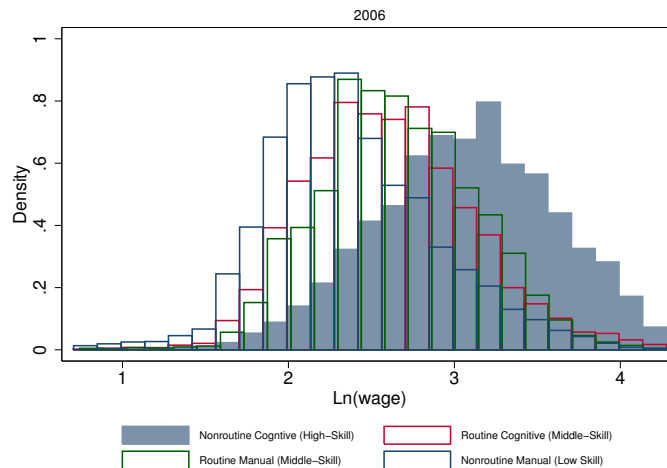
predictor of economic outcomes—such as employment and earnings losses—during the recession than a classification based on a somewhat arbitrary measure of skill. In other words, when job and wage polarization accompany each other, as they did in the 1990s, it might make sense to focus on one of them. But when they go in opposite directions, which seems to be the case in all four recessions shown above, extra caution should be exercised so as not to infer too much from changes by skills alone. Similarly, other results in the paper (for example, the dynamic factor analysis in Section 2) reveal that routine-manual workers behave very much like the industries that they work in, and non-routine manual workers behave like their own industries (services). So, again, one could argue that industry is a better measure of the economic prospects of a job than the skill level. These two pieces of evidence raise questions about the empirical content of the routine/nonroutine classification: once we group workers into industries and income groups, what additional insights are gained by the routine/nonroutine distinction? For the questions analyzed in this paper, the answer seems to be "not too much".

3   In Search of A Theory of Routinization

Although the routine/nonroutine distinction provides a useful descriptive account of the job losses concentrated in manufacturing and some other middle-skill occupations during

the 1990s, it is not clear how to think about future disruptions in different occupations using such a classification. For example, should we expect most of the future routinization to happen in manufacturing? or, more generally, in middle-skill occupations? how about occupations at the low end, such as cooks and waitresses that are currently seen as non-routine—are they safe from routinization?

The main challenge in trying to answer these questions is that, in principle, (almost) every task is "routinizable" at a given cost. As I discuss in some examples below, many tasks that once might have been considered untouchable by routinization, such as diagnosing medical conditions, giving professional investment advice, designing complex web sites, and so on, are being increasingly routinized, perhaps not in a wholesale fashion, but in bits and pieces that makes them routinized over time. But what determines which tasks gets routinized and which ones do not? Autor et al. (2003) had the nice insight that the complexity of a task is an important ingredient into the routinization process and built a simple yet elegant model in which routinization was linked (exogenously) to rapid technological progress in information and communications technologies. Acemoglu and Autor (2011) went one step further and studied a richer task-based model. In one extension, they considered endogenous routinization that happens through directed technical change, in response to variation in supply of skills (building on Acemoglu (1998, 2007)). Although these analyses provide a useful first step, I believe there are other important forces that also operate in determining routinization. So in a way, I view this paper and others as a call for a deeper theory of endogenous routinization and want to discuss some thoughts on what I view as key trade-offs that such theories should capture.

3.1 Thinking Through Some Examples

A few years ago, Taiwan-based Foxconn, which is is one of the largest electronics contract manufacturing companies in the world, has faced a worker revolt stemming from complaints about low wages. After a series of suicides by workers and continued unrest, Foxconn was forced to substantially raise the wages it pays to its workers. Soon after, Foxconn announced that it was building new factories to replace about 1 million workers with various automated machines and robots.[7] This example is not unusual and repre-

---

[7]See for example Reuters (2010), Markoff (2012).

sents one of the key trade-offs faced by companies, especially in manufacturing where tasks are repetitive, require precision, and can be routinized by the use of machines and robots. It also represents one of many examples that provided the impetus to Autor et al. (2003)'s work as well as to the subsequent literature.

While examples of robots replacing workers in manufacturing is what gets the most attention, routinization happens across a very broad spectrum and can take many different forms. To get a better sense, it is useful to begin with some examples and discuss how the tasks that make each one up have been partly or fully routinized or are susceptible to routinization in the future.

Example 1: Waiter vs. Cashier. According to the classification adopted in this paper, a cashier is classified as a routine-cognitive occupation whereas a waiter is a nonroutine-manual one. The thinking seems to be that electronic cashier systems (increasingly used at public parking garages, movie theaters, etc.) are very effective substitutes for workers doing these jobs which replaces these workers. Let us take a closer look at what a waiter does. At a typical restaurant, a waiter visits a table 5 times to (i) welcome the guests and bring the menu and ask for drink orders, (ii) bring the drinks and take the food orders, (iii) bring the food, (iv) bring the check, and (iv) bring back the payment instrument (e.g., credit card). Except for (ii) and (iii), the other three tasks involve either taking an order or acting as a cashier and are very easy to routinize. In fact, many recent restaurants, especially at airports where service speed is essential, use iPads to replace these tasks and have managed to reduce waiters' trips to each table from five down to one or two. Clearly, this substitution significantly reduces the need for waiters.[8] Therefore, there is no presumption that waiters' current classification as non-routine is a good indicator that they will be immune to the forces of routinization.

Example 2: Stock Broker. Now consider the stock-broker occupation, which in the 1980s was a high-skill non-routine cognitive occupation with a very high average

---

[8]A similar transformation happens for cooks: most chain restaurants have been moving food preparation to large facilities where food is prepped (vegetables are washed, skinned, chopped, etc.; breads and cakes are pre-baked, etc.) and is only finished off by cooks in the restaurant. Again, this economies-of-scale is likely to translate into lower demand for cooks, which are currently classified as non-routine manual workers.

income. The job consisted of three main tasks: (i) provide investors with detailed up-to-date information about the stocks, (ii) advise investors on investment strategies, and (iii) place the order on behalf of the customer. The substantial reduction in the cost of accessing information has reduced the value of (i), whereas new investors' preference for passive diversified investments (e.g., mutual funds) has reduced the need for (ii), and electronic trading made (iii) obsolete. As a result, the stock brokerage business has collapsed leading many brokers to lose their jobs permanently. Notice that in this example, software—and not hardware or machines—led to the demise of this profession, a trend that applies more broadly to many high-skill occupations going forward.

One can go through many more examples similar to the two discussed here to come to the conclusion that routinization is pervasive and there should be no presumption that it will be limited to the middle of the skill distribution where *currently routine* occupations reside. Now I will discuss two examples where routinization leads to a push for unbundling tasks that make up an occupation into distinct occupations (example 3) or the bundling of multiple tasks into one occupation (example 4). Also this push might come from the workers currently in that occupation (example 3) or from the buyers of the services of the occupation (example 4). Let us see how each one works.

Example 3: Unbundling a Dentist Occupation. At a typical dentist office 50 years ago, the dentist would see a patient for about 30–45 minutes, during which time she would (i) perform cleaning and other routine maintenance procedures, (ii) inspect the teeth (via x-rays etc), diagnose potential problems, and recommend treatments such as fillings, pulling a tooth, or performing root canal if needed. It is clear that the first set of tasks are quite routine and low risk. As such, they can easily be performed by a trained technician or assistant, whereas the second set of tasks require more experience and cognitive skills, and carry higher potential risks. Therefore, over time dentists have hired assistants who are trained to perform the first set of tasks which allowed them to focus their time entirely on the second set of tasks.[9] Because the second set of tasks are more critical, the hourly wage rate for them is much higher than the first, and the release of time from the routine tasks allowed dentists to see many more patients in the

---

[9]Part of this shift also happened due to a better appreciation of preventive care leading to increasing demand for these routine procedures.

same time frame, which allowed them to substantially increase their total earnings.[10] Furthermore, over time, a subset of the second type of tasks that involve mini surgeries that carry higher risks (such as pulling wisdom teeth, performing root canal, etc.) split into another occupation, which requires higher training and much higher wages, called oral surgeons. The bottom-line is that, in this example, the demand for splitting tasks into separate occupations (dentists assistant, dentist, and oral surgeon) is mainly driven by the workers who were originally part of the occupation.

Unbundling a Professor Occupation into Teacher + Researcher. Now consider a college professor and for simplicity let us focus on the two main tasks they perform: (i) teaching classes, and (ii) performing research. It could be argued that from a college's perspective the first task is what brings in most of the revenue, whereas the second one brings much less (especially true outside of professional schools). From a professor's perspective, research is the more enjoyable task (i.e., lower cost to produce) and she would like to spend most of her time doing that, but because the effective wage of that task is very low, she would like to keep the two tasks bundled so as to make a reasonable income. But when the total cost of education (tuition) goes up, the purchasers of this service, the students would start demanding that the two tasks are unbundled so that each teacher (who can be a lecturer) can devote most of her time to teaching (i. e., teaching more courses per unit time, once research time is driven to zero), reducing costs. This force seems to be one of the drivers behind initiatives such as online education and the increasing usage of lecturers and adjunct faculty in teaching.[11] Consequently, in this example, the demand for bundling comes from professors (the providers of the service), whereas the demand for unbundling is coming from students (the purchasers).

Taking Stock. As these examples illustrate, the incentives for bundling and unbundling tasks into occupations depend on a complex set of factors that involve the relative de-

---

[10]Since each dentist can now see more patients, this change reduces the demand for full-time dentists (assuming demand for dental services remain constant), and increases demand for dental assistants, which will feed back into the wages of each occupation, partially dampening the rise in the earnings of dentists.

[11]Obviously there are complementarities between teaching and research, but probably these are greater at the graduate teaching level, which may be better protected from routinization. Also, it is not clear how important research experience is for undergraduate teaching once one moves out of the top research institutions.

mand for each task (determining the wage rate for each task), the technology and skill requirements of each task (determining who can do it and at what cost) and the skill distribution in the population, among other factors. Thus, there are many more factors determining the direction of routinization than the supply of skills emphasized by Acemoglu and Autor (2011). Agents in an economy will demand the lowest price for the collection of tasks that they wish to purchase, while at the same time, the producers of each task will demand the highest wage for their efforts. If the demand or wage rate of a certain task increases significantly, this will put pressure on "routinizing" that particular task, as in the teaching example, so as to reduce its price. This is very much the idea of "directed technical change" applied to the routinization of tasks, whose wage has risen.

4   Conclusions

Foote and Ryan have produced a useful paper that is a nice blend of new data and a thoughtful analysis of two interesting hypotheses. The paper provides evidence that the middle-skill job losses observed during recent recessions have been consistent with the past cyclicality of the industries in which these occupations are located. There seems to be less evidence that the fall in middle-skill employment is due to workers leaving their jobs for other occupations with more robust demand as argued by Jaimovich and Siu (2014). Although this is not the last word on this debate, it provides a step forward in our understanding of the relevant issues. I believe that future editions of this debate would greatly benefit from a full-blown panel data analysis using recently available administrative datasets, such as the LEHD project of the Census Bureau.

One of the main conclusions I draw from the analysis in this paper is that perhaps classifying workers/occupations by broad categories of skills (low, middle, high) and a routine/nonroutine distinction is not as useful as it might first appear. When workers are classified by their recent wages instead (as seen from Figures 2 and 3), we see homogeneity—a single factor structure—during every recession both in wage growth and in employment, whereas grouping by skill categories reveals different patterns by decades. Similarly, as Foote and Ryan show, employment fluctuations by skill levels seem to track the behavior of the industry that the occupations reside in. Furthermore, even if the routine/non-routine classification provides a nice descriptive account of past trends, it is not clear how to link it to underlying forces to make predictions about future rou-

tinization of tasks and occupations. Acemoglu and Autor (2011) take one step in that
direction and I have outlined some thoughts on other factors that seem important for
determining the direction of future routinization. Finally, I find the results on declining
labor force participation of middle-skill workers to be especially intriguing and I hope
the authors pursue that analysis more fully, perhaps with better data, in future work.

References

Acemoglu, D. (1998, November). Why do new technologies complement skills? directed
technical change and wage inequality. *Quarterly Journal of Economics 113*(4), 1055–
1089.

Acemoglu, D. (2001, January). Good jobs versus bad jobs. *Journal of Labor Eco-
nomics 19*(1), 1–21.

Acemoglu, D. (2007, 09). Equilibrium bias of technology. *Econometrica 75*(5), 1371–
1409.

Acemoglu, D. and D. Autor (2011). Skills, Tasks and Technologies: Implications for
Employment and Earnings. In *Handbook of Labor Economics, Vol 4B*, Chapter 12, pp.
1043–1166. Elsevier.

Autor, D. H., F. Levy, and R. J. Murnane (2003, November). The Skill Content Of
Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of
Economics 118*(4), 1279–1333.

Castro, R. and D. Coen-Pirani (2008). Why have aggregate skilled hours become so
cyclical since the mid-1980s? *International Economic Review 49*(1), 135–184.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk.
Journal of Political Economy 122(3), 621–660.

Jaimovich, N. and H. E. Siu (2014). The trend is the cycle: Job polarization and jobless
recoveries. Working paper, Duke University.

Juhn, C., K. M. Murphy, and B. Pierce (1993, June). Wage inequality and the rise in
returns to skill. *Journal of Political Economy 101*(3), 410–42.

Markoff, J. (2012, August 18). Skilled work, without the worker (the ieconomy, part 6). New York Times, Retrieved from http://www.nytimes.com.

Reuters (2010, October 2). Foxconn to raise wages again at china plant. Reuters (Retrieved from www.reuters.com).

Stock, J. H. and M. W. Watson (2012). Disentangling the channels of the 2007-2009 recession. *Brookings Papers on Economic Activity*, 81–156.