Graduate Labor Economics

Notes to Accompany Lecture 5: Tasks, Polarization, and the Future of Work

Brendan M. Price*
Federal Reserve Board

In this lecture we'll study the evolving task structure of the labor market—a literature launched by Autor et al. (2003)—and the implications of "routine-biased technical change" both for the polarization of the US occupational structure and, much more speculatively, for the future of work. Acemoglu and Autor (2011) and Autor et al. (2013) offer excellent introductions to the core ideas and fact patterns that have fueled research in this area.

Rather than offering color commentary on the material we'll cover in class, here I'll share some thoughts on an essential component of the data work that underlies Autor et al. (2003) and papers like it: the dark art and unglamourous science of crosswalking.

1 Classification schemes

- Lots of economic data are encoded in hierarchical classification schemes: examples include occupations, industries, product codes, geographic areas, medical diagnoses, and patent classes.
- Different data providers may use different classifications to represent the same underlying construct: for example, the Quarterly Census of Employment and Wages uses NAICS industry codes, whereas the Decennial Census uses a separate system of Census Industry Codes.
  - These differences sometimes represent different administrative priorities. The US Postal Service developed ZIP codes to facilitate mail delivery, but ZIP codes don't correspond neatly to particular land areas or political jurisdictions, so the Census Bureau uses its own system of ZIP Code Tabulation Areas for enumerating economic data.
  - Sometimes these differences reflect competing standards that were never harmonized (e.g., between countries), or lags in the adoption of harmonized coding regimes.
- These classification systems also change over time, both in fundamental ways (e.g., from 1987 SIC industry codes to 1997 NAICS codes) and in incremental ways (e.g., from the 1997 NAICS to the 2002 NAICS). Economic classifications tend to get more granular over time, reflecting the growing complexity of economic systems and the need for more detailed taxonomies.
- Occasionally, changes in classification systems are informative in their own right. Lin (2011) regards the emergence of new occupational job titles in Census data as a measure of "new work" and uses them to study the spatial and demographic characteristics of frontier jobs.

 Email: brendan.m.price@frb.gov. These notes are adapted from UC Davis course ECN 250A, which I taught in Winter 2018 and Spring 2019. Any views or opinions expressed here are my own and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. This version: 2020/03/21.

- But most of the time, changes in classification systems are just a nuisance to be finessed as expeditiously as possible through the use of "crosswalks", also called concordances or bridges.
  - Researchers often have to harmonize data series themselves—more on that below.
  - Data providers and research teams sometimes publish data that have already been harmonized. If done well, such pre-processing can save you a lot of time, but even expert crosswalking is often imperfect, and many "harmonized" data series exhibit telltale signs of imperfect stitching. In some cases, you may need to refine an existing crosswalk, or even to develop your own from the underlying data.
- Crosswalking can be tedious, painstaking work. Modern empirical projects routinely merge numerous data sources spanning long time periods, with many data seams that need to be stitched together. Before investing too much time, sweat, and tears, ask yourself:
  1. What's the right unit of analysis for my inquiry? A paper about the local labor market effects of, say, Lilliputian import competition might ideally be conducted at the level of commuting zones, each of which is an assemblage of counties—but maybe you can start with state-level data (which are easier to work with) to see if there's anything there.
  2. What's the right time period for my analysis? Since classifications change over time, it's common to face a tradeoff between a longer sample period and a more granular set of harmonized codes. Is it worth losing a lot of occupational or industry detail just to add a few extra years of data? Maybe yes, maybe no.
- For exploratory work, a quick-and-dirty crosswalking job will do, but sloppy stitching can lead to weird data artifacts and spurious conclusions. Do it right before getting in too deep.

2 Crosswalking

- Let's now consider the generic task of getting classification scheme A, with codes $a_1, \ldots, a_N$, to communicate with another scheme B, with codes $b_1, \ldots, b_M$. For concreteness, let's say these are industries.
- Crosswalking usually begins with one of the following:
  - Most commonly, you'll have access to some published concordance that provides a mapping between A and B. Don't reinvent the wheel: lots of crosswalks are available online, either buried in the recesses of government websites or listed on researchers' data pages. For example, David Dorn offers a wealth of industrial, occupational, and geographic crosswalks (`https://www.ddorn.net/data.htm`), and Peter Schott provides trade concordances (`https://sompks4.github.io/sub_data.html`).
  - Occasionally one comes across samples that are dual-coded under two classifications. Statistical agencies that are rolling out updated classifications will often use both old and new systems side-by-side in the same survey round, enabling them to construct concordances that account for, say, the share of firms in industry a_1 that are assigned the new industry code b_3. Autor et al. (2003) rely on the dual-coded "Treiman sample" to map 1970 task measures into 1980 occupations.
  - Relatedly, one might observe dual codings that are separated in time: for example, a plant-level survey dataset might use codes A in period 1 and codes B in period 2. If

the same plants are observed in both periods, one might construct a crosswalk based on the observed proportion of surviving a_1 plants that transition to code b_3. This approach is open to critique, however, both because of potential selection effects in the set of surviving plants and because industry affiliation responds to economic pressures like import competition (Bernard et al., 2006; Bloom et al., 2019).

  - Absent any of the above, researchers sometimes have to rely on their own subjective correspondences based on industry titles or descriptions. Don't do this unless you have to: it's subjective, error-prone, and time-consuming. But it's the option of last resort.

- Let's suppose we're lucky enough to have a published concordance. In the simplest cases, we observe a one-to-one, one-to-many, or many-to-one mapping between A and B. In these cases, the general principle is to map everything into the coarser classification. For example, suppose that $A = \{a_1, a_2\}$ and $B = \{b_1, b_2, b_3, b_4\}$. If code a_1 is associated with codes b_1, b_2, and b_3, while code a_2 corresponds exactly to code b_4, we can glue codes b_1, b_2, and b_3 together and express everything in terms of A. Since classifications tend to get more granular over time, this approach usually means mapping everything into older, historical classifications.

- Many crosswalks are fractional or probabilistic. For example, we might be told that 80% of employment in a_1 is associated with b_1 while 20% of employment is associated with b_2. Now suppose our data are encoded in system A in the early years of our sample but in system B in later years. We could convert everything into system B by apportioning all of the A industry employment counts into industries B according to the proportions given in the crosswalk.

  - What if we're interested in harmonizing industry-level data on value added, but the available crosswalk is expressed in terms of employment? As a practical matter, it's common for researchers to assume that employment weights are appropriate for outcomes like gross output or value added, though there is inevitably some measurement error here.

- Once you've applied a crosswalk to a given data series, do what you can to validate the crosswalk and correct any errors that emerge.

  - Classification hierarchies often impose lots of adding-up constraints: for instance, county-level employment should sum to state-level totals, employment in manufacturing subsectors should sum to the total manufacturing sector, and so forth. Make sure your data satisfy such constraints. If they don't, try to figure out why.
  - For time series data—for example, industry-level employment counts spanning a three-decade period—a useful check is to plot the harmonized data over time and look for discontinuities across seams. This often reveals obvious inadequacies in the crosswalk: for example, employment in the "tennis shoes" industry might fall sharply at a data seam, while employment in "shoes, n.e.c." rises by the same amount in the same year. Judicious aggregation can often fix these problems: in this case, we might combine "tennis shoes" with "shoes, n.e.c." throughout the full sample period.
  - If you observe the same data series in two separate datasets—for example, employment in textile manufacturing in both the County Business Patterns and the Current Population Survey—see whether the data line align well, both in levels and in trends.
  - These are necessary, not sufficient, conditions for a crosswalk to be working well. For example, looking for discontinuities in a harmonized time series will catch any differences in levels at the seam, but it's harder to spot trend breaks. Since published crosswalks

are often based on dual-coded samples, the accuracy of a crosswalk will tend to dwindle the further one gets from the seam between two classification regimes.

  - If you're worried about measurement error or crosswalking issues in a granular set of codes, consider aggregating to a higher level at which such issues tend to be less acute.

- Often we need to crosswalk from A to B and then from B to C. Crosswalks are typically transitive, but each step in a crosswalking chain will typically (as in the children's game of "telephone") entail additional measurement error, so you should be wary of crosswalks assembled from many imperfect links (and, correspondingly, subject them to extra validation exercises to persuade yourself of their accuracy).

- Many subtle issues arise in practice, and there are often ways to use auxiliary information or other data series to improve upon an off-the-shelf crosswalk or to assess its accuracy. Be careful, and be creative!

References

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4B, pages 1043–1171.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6):2121–2168.

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

Bernard, A. B., Jensen, J. B., and Schott, P. K. (2006). Survival of the Best Fit: Exposure to Low-Wage Countries and the (Uneven) Growth of U.S. Manufacturing Plants. *Journal of International Economics*, 68(1):219–237.

Bloom, N., Handley, K., Kurmann, A., and Luck, P. (2019). The Impact of Chinese Trade on U.S. Employment: The Good, The Bad, and The Apocryphal. Working paper.

Lin, J. (2011). Technological Adaptation, Cities, and New Work. *Review of Economics and Statistics*, 93(2):554–574.