What is the Ideal Labour Code: A Measurement Perspective from the United States

Mike Horrigan
US Bureau of Labor Statistics

As with all public policy issues and debates, the quality of data that serves to inform the debate is of critical importance. However, as is often the case, there is usually a strong positive correlation between the questions and issues that are of greatest interest and the relative lack of data to shed light on those questions. The subject of this session, "What is the Ideal Labour Code" is a case in point. One of the recent trends in the United States and throughout the developed world is the emergence of alternative work arrangements that have, at least in part, replaced traditional long-term full-time employer-employee relationships and for which we have relatively little data. The use of on-demand workers, aided by the use of internet technology to bring together the providers of services or the producers of goods and consumers demanding those services or goods is one example. Two prominent examples of these practices in recent years is the creation of internet exchange markets for taxi cab services and for baby-sitting services in major cities across the globe. The data systems of the U.S. statistical agencies have not caught up in terms of having anything other than indirect measures of such economic activities as evidenced by trends in consumer spending or self-employment.

Other practices such as the use of third party contractors (also known as temporary help agencies) and self-employed contractors, while of longer standing use, are measured infrequently if at all in U.S. statistical systems. Household surveys that measure so-called 'contingent' work arrangements have not been conducted in the U.S. since 2005, effectively missing the potential changes in trends in these types of employment relationships since the Great Recession of 2007-2008. The use of these contingent worker arrangements as measured in establishment surveys is similarly limited. In the U.S., for example, we have data on the level of employment in the Temporary Help Services industry, a key component of these alternative work arrangements. However, what is unknown in these data series is the industry of placement of temporary help services workers. A careful examination of the changes in employment in the Temporary Help Services industry as compared to changes in total nonfarm employment shows that the former is a highly reliable leading indicator of changes in establishment employment levels. It appears obvious that firms are using temporary help workers as inventory buffers to respond quickly in terms of employment levels across business cycles. As the economy enters into a recessionary period, firms tend to lay off temporary help workers first, offering greater protection of their core workers. Similarly, as the economy begins to recover, firms are hiring the services of temporary help service workers first, prior to making a longer-term commitment to bringing on permanent workers. Understanding these trends would be greatly enhanced if we knew how many temporary help service workers were employed in different detailed industries.

More generally, beyond the business cycle aspects of using temporary help service workers as a buffer against risk, we have relatively little information on the variety of approaches that businesses are taking to producing their products including the mix of permanent, third party and self-employed contractors, the use of domestic outsourcing, and the use of global production supply chains. Firms are making profit maximization decisions on the choice of productive technologies and the mix of permanent and temporary workers used to produce their intermediate and final goods or provide services. As well, the decision to outsource (either using domestic or international sources) the production of some of their goods or the provision of some of their services also fundamentally changes the nature of employer-employee relationship. Again, there is relatively little information collected by U.S. statistical agencies on these changing business practices.

Globalization

Another singular measurement issue that has arisen in this context is the measurement of global production processes. While significant efforts are being expended to construct estimates of value-added trade, the statistical systems of developed countries are ill-equipped to capture, in real-time or in near real-time, the amount and value of global production supply chains. Two specific measurement issues have arisen in this arena. The first is the fact as domestic manufacturers shift their use of inputs from domestic to international sources, the price advantage associated with such outsourcing is not captured in U.S. import price indexes, leading to an overestimate of import price inflation and an overstatement of real GDP growth rates and real multi-factor productivity.

A second concern is the measurement of economic activity of so-called factoryless goods producers (FGPs). As defined by the U.S. statistical system, factoryless goods producers outsource all of the transformation steps traditionally considered manufacturing, but undertakes all of the entrepreneurial steps and arranges for all required capital, labor, and material inputs required to make a good. The characteristics of FGPs include: the establishment/firm owns or controls rights to intellectual property or design; can independently change product design; controls production, has the ability to control inputs, choose product lines, set output levels; owns the final product; sets price of final product; sells or arranges for sale of the final product; assumes entrepreneurial risk and is responsible for losses. Without a formal industry classification of factoryless goods producers, the potential for misclassifying their employment, sales, profits, etc. in wholesale trade rather than manufacturing is significant. In addition, the possibility exists that firms will be classified differently in different establishment survey frames and samples leading to inconsistent measurement of industry employment and wages (among other variables) across federal surveys.

Big Data

These very specific data gaps about the changing nature of work can also be usefully viewed in a broader context. In recent years, the emergence of so-called ‘big’ data has presented the U.S. statistical system with a number of challenges and opportunities to closing significant data gaps, especially when these gaps are set against a backdrop of tighter and tighter budgets for data collection and dissemination. As well, the emergence of big data as a potential substitute for data collected by traditional survey methods figures importantly into these budget discussions.

But what are “Big Data”? It is a term that has an increasingly familiar ring, but at the same time defies easy description. One approach is to view big data as non-sampled data, characterized by the creation of data bases from electronic sources whose primary purpose is something other than statistical inference. For example, the Massachusetts Institute of Technology (MIT) Billion Prices Project digitizes posted Internet prices to construct estimates of daily price change. Hal Varian, chief economist at Google, has done highly innovative work using Google searches to create proxies for current economic activity. For example, to predict, at time (t), the level of initial claims for unemployment insurance (UI) at time (t+1), he constructs a model of distributed lag values of prior weeks’ initial claims data along with an index of searches made in the current week that are relevant to people looking for information about filing an initial claim. This is a clever combination of “official” government collected data with the construction of an indicator from a ‘big’ data source. Matthew Shapiro, along with other researchers at the University of Michigan, has used data from Twitter accounts in a model that also predicts the level of initial claims for unemployment insurance, where he isolates tweets that reference job loss. And yet another example of big data is scanner data such as the point of sale retail data bases and the household based purchase data from A.C. Nielson.

These innovative and exciting explorations of data would seem not to be the standard fare of an agency like the U.S. BLS. But are they? How does the BLS fit into this picture of the use of big data? From a non-sampled data point of view, the BLS makes extensive use of administrative data to draw stratified probability samples and to create weights for constructing estimates. The difference here is that this type of big data typically comprises the universe, and by definition, can represent (nearly), the entire population of establishments (the U.S. BLS Quarterly Census of Employment and Wages drawn from the universe of establishments reporting to the UI system) or households (the 2010 U.S. Decennial Census of household addresses). There are numerous other such administrative data bases such as those covering railroads, hospitals, medical claims, and auto sales that BLS also uses for its r surveys. For example, our item sample of used cars and trucks in the Consumer Price Index Program (CPI) is drawn from the universe data collected by JD Power and Associates. We use universe data on hospitals from the American Hospital Association to draw our samples of hospitals and data from the Agency for Healthcare Research and Quality to select the diagnosis codes used for pricing (DRGs – Diagnosis Related Groups) in the Producer Price Index Program (PPI).

In addition to using non-sampled universe files to draw samples and create sampling weights, increasingly we use this type of administrative data for the direct construction of population estimates. For example, the International Price Program (IPP) uses Energy Information Agency administrative data on crude petroleum for their import indexes; the PPI uses Department of Transportation administrative data on baggage fees in constructing airline price indexes; the PPI also uses a monthly census of all bid and ask prices and trading volume for all traded securities as of market close for 3 selected days of the month to construct price indexes for securities. Both the PPI and CPI use the universe file for Medicare Part B reimbursements to doctors by procedure code in the construction of health care indexes.

In other cases, administrative data are used to fill in missing data as an alternative method of imputation or in making statistical adjustments to improve the efficacy of estimates. For example, the Current Employment Statistics (CES) Survey uses administrative data from the Quarterly Census of Employment and Wages (QCEW) to impute for key non-respondents in the production of industry employment estimates by State. QCEW Data are also used in the development of the CES net birth-death model to account for the creation and death of firms between updates to the universe file that is used in constructing monthly employment estimates.

But what about the use of more "traditional" big data techniques? In fact, a survey of programs in BLS uncovered some intriguing forays into big data exploration. For example, the CPI uses web scraping techniques to collect input price information that is used to greatly increase the sample of observations we use to populate some of our quality adjustment models. So far, CPI has used this technique with quality adjustment models for televisions, camcorders, cameras and washing machines. We are also web scraping Current Procedural Terminology (CPT) codes, descriptions, and reimbursements for Medicare Part B quotes that are used in index calculation. CPI is also researching the use of web scraping for the collection of prices for cable TV services.

This latter example raises an obvious question: Why not just use web scraping to produce the CPI? The principal reason is the requirement that we select a bundle of goods and services that is a *statistically representative* sample of what consumers purchase and to *reprice* that *same* bundle month after month. Accomplishing this objective can be very challenging, especially having to account for changes in the quality characteristics of goods as well as goods that disappear off the shelves from one month to the next. The representative basket is updated on a regular basis to reflect changes in consumer tastes and preferences and the emergence of new products; however, the principle of constructing an inflation rate that is based on the rate of price increase for a known bundle of goods with statistically determined weights lies at the heart of what we do. While research may show the viability of using a web scraped source of data for a particular item, it needs to be done within the framework of this methodology. The Billion Prices Project, with all of its advantages in terms of the timeliness of a daily price index and large sample sizes, does not price the same representative bundle on a daily basis, nor does it have a source of sampling weights that are derived from the web sites for which it collects prices.

BLS, like many agencies, has been exploring the use of retail scanner data for many years. To date our most extensive use of scanner data has been in the realm of research, including comparative research between CPI data and scanner data. For example, we are currently conducting research that compares, for specific expenditure classes of items (e.g., fats and oils), the distributions of items selected in the CPI selection process with the distributions of those same items in the A.C. Nielson Homescan data base.

And one final example of big data, one that is unique even in terms of the examples given above, and one that has the potential to greatly affect our data collection systems, is the use of autocoding. For example, in the BLS Occupational Employment Statistics Survey Program, we currently require firms to identify, as relevant, the number of employees in each of more than 700 occupations and the associated number of employees falling into separate wage interval categories. We are planning to replace these burdensome survey forms with asking firms to report, for each employee, their wage level and job title. By employing big data autocoding techniques, we will translate these job titles to the correct occupation category. It is our hope that the burden on firms will be greatly reduced by allowing them to report information they already have in their data systems – a wage level and a job title for each employee.

And so, what is the future of the use of “big” data for the U.S. statistical system? I see one immediate potential, the use of big data to improve the quality of our estimates within our current methodological frameworks. This may include studies of comparability between official and big data derived estimates, the use of big data for modeling and imputation, and in some cases the use of big data for direct estimation. One important caveat, and one that is as relevant to the U.S. statistical system as it is to the practitioners of big data techniques such as Billion Prices and Google, is the need to create transparent methodological documentation (metadata) that describes the ways in which big data are used in the construction of any kind of estimate. Given rising costs of data collection and ever tighter resources, there is a need to consider the creative use of big data, including corporate data. However, the blending of estimates drawn from traditional statistical methods and the incorporation of larger universe data requires clear statements of how these estimates are developed and a perspective on potential sources of sampling and non-sampling errors that can produce biases in our estimates and threats to valid inference.

Conclusion

There are significant data gaps in our understanding of how the world of work is changing. How firms are changing the mix of productive technology and the use of various types of workers including third party contractors, self-employed contractors, and on-demand workers is not well understood. There are also significant gaps in our understanding of global production processes, including factoryless goods production. Finally, the emergence of big data as a possibility pathway to closing some of these data gaps, especially in an era of tight budgets for statistical agencies, is at a relatively early stage of exploration.