IMPROVING THE OCCUPATIONAL OUTLOOK SUPPLY AND DEMAND IN THE LONG-TERM EMPLOYMENT OUTLOOK

LONG-TERM EMPLOYMENT PROJECTIONS RESEARCH

Project Purpose

Attempts have been made over the years to incorporate occupational employment supply projections into the process that generates expected long-term occupational employment. These efforts have usually floundered because the gathering of data by occupation on people holding jobs— the so-called supply of labor— is difficult to pull off using current data collection mechanisms. Nevertheless, the original purpose of this project was to suggest the best way for states to do so. The end result of this research suggests that the effort to determine occupational supply is, in fact, futile— but for reasons unrelated to gathering data or projecting its future level.

This project, as originally conceived, was also an attempt to explore the best way to project occupational skills. The two projects were combined because neither is presently a part of the projections process. Upon reflection, however, it appears that this research into the role of supply and demand in long-term employment projections is not the best place for considering skills. How supply and demand are defined and treated is a fundamental quality of the method by which the future number of jobs is calculated. Skills, on the other hand, are merely a quality— a feature of a given occupation akin to training requirements— that are tagged onto the final results independently of the calculation itself. Therefore, the best way of estimating future skills can be handled within the project that considers occupational aggregation (in progress: see page 2 of Executive Summary).

The Need for Clarification

There has developed over the years a view that employment projections are undertaken in order to discover the future demand for various occupations. Under this notion, an attempt is made to quantify the number of openings to which an independent projection of supply ought also to be added in order to determine the likelihood that occupations will face a potential shortage or surplus of labor. Yet, as this research will

demonstrate, this is a faulty proposition from the start. It is faulty from the standpoint of economic theory in a strict sense but fails under more relaxed assumptions as well. It is also faulty in an empirical sense in that the data used to generate projections do not capture demand. If employment projections do not represent demand, then the whole premise of projecting supply also falls apart. Monitoring supply/demand and determining where shortages and surpluses will surface must be attempted in different ways, which belong outside the purview of the long-term projections process.

Where did the idea that the long-term employment projections process generates the number of jobs demanded by occupation originate? This is a surprising notion because economic theory suggests that new demand for any product or input typically occurs suddenly (and can be modeled by short-term production functions), while the supply of it changes or responds more slowly. The process for projections at the national level acknowledges this: it projects the size of the labor force as its first step, the results of which are then used as exogenous variables in the model of aggregate demand for the U.S. economy. In other words, future demand for labor is somewhat constrained by the expected supply of labor. The intermediate steps that determine GDP by product group and industry are termed demand. Since generating these values immediately precedes the calculation of expected employment counts, the demand label carries over and sticks, while the supply-side constraints remain hidden.

The idea that demand needs a supply counterpart at the state level probably also comes from the fact that data on jobs used in the process are collected from employers. Since they are the ones hiring (demanding) workers, the employment number must represent the amount of labor they demand. Under this logic, if these numbers were collected from workers, they would instead constitute supply. But the source of collection does not matter as much as how the question is asked. True demand for labor must include the number of people at work plus the unmet need, while true supply must include the number of people working plus those who want to work but are not. In fact, historical employment tallies used in making employment projections are the intersection of supply and demand: the quantity of labor at work. There is no need for a supply counterpart when employment as measured is not demand per se and already encompasses supply.

Problems with Projections: Economic Theory and Definition of Demand

Aside from collection issues, there are additional problems with treating employment projections as the demand for labor. Strictly speaking, occupational demand is not a set number of jobs. Demand is really a relationship between quantity and price that stipulates how much labor employers will choose to utilize at various wage levels. Its counterpart, supply, is the measure of the people who will offer their services to the market at various wage levels. The intersection of the two will determine both the quantity of labor that gets hired and the wages they are paid: the equilibrium point. This set of relationships can be visually represented by a set of intersecting lines (sometimes called curves), and it represents reality as well as any other model. Wages decreed at a particular level may keep the lines from intersecting, or constraints on the quantity of workers may cause the supply line to take on an abnormally steep slope and result in high wages, but these tend to be special cases in an economy like the U.S., and the model as a construct still holds. In any event, the model works better at the labor market level than at the national level, for both workers and employers can more easily adapt to constrained conditions by relocating to other labor markets. In the U.S, the main issue with supply/demand lies not in the model but in the time it takes the participants to adjust to changes and reach equilibrium.

That still does not make it easy to use, for collecting data on all occupations in all labor markets is a practical impossibility. In a pure sense, one would have to ask employers how many workers they would hire at each wage rate, and ask workers whether they would take employment at various rates; one would then sum the results. Clearly, the system for projecting employment does not do this.

Even if one takes liberties with strict supply/demand curves by ignoring the price variable for the moment and considers only the quantity of labor demanded and supplied, there is a disconnect between the ideal system and the way projections are constructed. One problem is definitional. As previously mentioned, demand would have to include both the number of people working and the number of unfilled openings in order to truly depict quantity. (The present way of measuring turnover does not come from the same source or even from the employer, as it should.) Any attempt to capture supply, no matter how comprehensive, would not match up.

The other problem is theoretical. Suppose analysts could have at their disposal the midpoint of both the supply and demand curves (both properly collected) and then measure the gap between the two as a surplus or a shortage. There is still an issue of how the two get resolved. For example, if there is a surplus of workers, then the only way one can predict it will be alleviated is for people to leave the occupation or that labor market. This is unrealistic because it is so slow; what often actually happens is that employers lower their wage rate and workers accept it rather than relocating. The reverse happens when a shortage typically prompts a rise in wage rates well ahead of any in-migration or surge in new training. In short, expecting, predicting or even encouraging gaps to be fully closed by quantity alone is not how the labor market works.

Given enough time, there will be adjustments by both prospective jobseekers and employers to clear supply and demand for just about all occupations. Wages are usually the first adjustment mechanism— followed by relocation, career changes and surges/retrenchments in training programs. Considering that employment projections cover a 10-year period, however, time is not the critical factor. In the long run, there ought not to be many shortages and surpluses.

Current, Short and Long-Term: The Distinctions are Important

Any forecast that uses industry and occupational data series from federal collection programs is measuring the future level at which supply and demand are expected to clear, and this is the best measure to have in any case. There is no need to collect and project supply data, at least in the long-term. The short-term is another story; here the case for collecting and forecasting supply and demand separately may make more sense. Markets cannot clear completely on price or quantity alone, and labor markets are notoriously slow because both training and relocation take time.

At any given moment, shortages and surpluses for demand for goods and services can easily exist—— and hence inputs to produce them can suddenly shift. These disparities, especially for shortages, could well persist for many years. Where does long become short? Ideally, the changeover point is the time it takes for individuals to absorb labor market signals, apply for training programs, complete them and attain gainful employment. This varies by occupation, of course. However, the maximum is probably

around six years, which, at any rate, is longer than either the worker or business relocation option— and well within the length of time associated with employment projections.

By themselves, employment projections cannot tell us which occupations will be in short supply; they can only tell us how many jobs there will probably be in an occupation. Analysts presume that wage signals will force movement into and out of occupations and areas, so training programs only have to keep up with employment growth (extra growth in new workers would only lower wages or force out-migration).

The assumption that the unemployment rate in the projected year returns to its base-year level or reaches its long-term equilibrium level is necessary, however, for theoretical consistency. States and areas need to adjust their projections if their business cycle is expected to behave differently than the nation's: changes in labor utilization would impact the quantity of workers (presumed to expand according to population and participation trends) and somewhat control aggregate demand. Likewise, states and areas that experience net migration rates that vary from the nation would be wise to work this information into their employment growth model. It may not be necessary to develop a system of top-down control, but it certainly is necessary to set parameters so projected employment can be recalculated if the sum of industry models adds up to an unrealistic result. By keeping in mind that employment projections represent the market clearing level and not demand for labor, the need for aggregate labor supply projections moves to the forefront. Since projections presume clearing markets, is there a way to answer the question of shortages and surpluses using different approaches and data sets?

The Pitfalls and Promises of Job Vacancy Surveys

One approach may lie in continuing job vacancy surveys for a number of years. That would make them an important source of information about long-term occupational trends in addition to the value they have for identifying immediate job opportunities. Clearly, if unfilled vacancies exist year after year, then the labor market is not working properly on its own and either public policy to alter supply or stronger messages to job-seekers are warranted.

There are two problems with this approach. The obvious, albeit surmountable, problem is how to deal with turnover. Job openings result from new positions at expanding

organizations and from workers vacating positions— all of which is commonly known as turnover. Because there is a tendency for less desirable jobs to be vacated more quickly, any tally of openings that counts both new growth plus turnover without distinguishing between the two reasons is not a good guidepost for identifying shortages. Ranking occupations by total number of vacancies (or by share of employment currently vacant) is likely to send the wrong signal to both jobseekers and training providers by encouraging an increase of supply into an occupation that may in fact be so undesirable that it has many departures. To separate growth from replacement, some surveys have begun to adjust their vacancies by an assumed turnover factor. But these factors are a second-best approach, derived as they are from old sources, stemming from another area, and collected at another time of the year; they are also based on collection from a jobseeker survey, not a similar employer survey. All of this creates four potential sources of error. To the extent that replacement ratios are cyclical, seasonal and local (all probably true in most cases), turnover-adjusted vacancies as reported will be off the mark. Moreover, household surveys define turnover differently than employers do (the former as leaving an occupation, the latter as leaving a job). The adjustment process as it stands both understates replacement demand and ignores the cases where workers are enticed into switching employers while remaining in the same field.

This may not matter when the object of these surveys is decidedly immediate: to provide a sense of which occupations have the most and least current openings so job-seekers can be properly informed about their chances. But it has limited use in identifying even short-term shortages. The better approach would be to ask employers to disaggregate their occupational vacancies into new positions and refills. New positions as a share of total jobs in that occupation would be the gauge as to whether opportunities are growing and shortages might potentially be developing. A high replacement rate by itself tells us little.

Another, more insidious problem is that such surveys are perhaps self-reinforcing: businesses may report vacancies that get interpreted as shortages, but they only exist because employers are paying less than the market wage. Vacancies might diminish if wages were raised because fewer people in that area would depart the occupation or the location looking for greener pastures, and because new entrants would be enticed in. A

survey that does not account for wages in its analysis is flawed if the goal is truly to report shortage occupations. Also, there is no balancing employee survey that gauges underemployment. At the national level, the best sources of labor market tightness come from household surveys such as the unemployment rate survey, and from questions used to construct the consumer confidence index that ask whether jobs are plentiful or hard to get. Sources from employers like the help-wanted index and the survey of hiring intentions provide information on the strength of the job market, but without regard as to whether surplus, shortage or balanced conditions prevail.

The real promise of employer-based job vacancy surveys may be that they can isolate occupations for closer scrutiny— even allowing for flawed turnover adjustments. By tracking vacancies over time, this source may identify some occupations that show high numbers of openings beyond what turnover would suggest, and beyond what markets can quickly remedy. They are especially valuable if the survey captures wages offered. While lower-than-prevailing wages for vacancies may not tell us much (since younger, lower-paid incumbents tend to switch jobs more often), a high share of openings at or above prevailing wages could be an early signal that employers need to raise pay to attract workers— a sign of an impending shortage, at least in the short run. Should this condition persist, the warning sign of a chronic, lasting shortage could be even stronger.

But more than one-sided vacancy surveys are needed to fully identify shortages. An occupational household survey would help, but monitoring wage rates is cheaper and more reliable for providing the answer— it collects the clearing of supply and demand at the same time and from the same source. There is no messiness associated with trying to match the two.

The Best Alternative in the Long Run

Employment projections are of little value in suggesting shortage and surplus conditions, and no amount of supply-side tack-on will correct for this. The forgotten approach is to go back to theory and include price as well as quantity in any assessment of supply and demand. This is not to suggest that the employment projections program is not worthwhile— the expected market clearing level it generates provides valuable data in its own right. Nor is this a plea to project wages— for which an extrapolation of past trends is

risky. Instead, it's a call to track wages over time as a way of suggesting where and when labor market intervention may be warranted, to call attention to cases where neither employers, jobseekers nor training providers are reacting to their own to job market signals. Now that wage surveys are universally in place at the labor market level, the source to carry out this task exists.

Relative wage growth is what matters in this type of analysis. As long as some inflation exists in an economy and as long as labor productivity is positive, workers as a whole will tend to enjoy wage growth (the sum of the cost of living and productivity growth, plus or minus any bargaining power from a tight vs. loose market). But the relative supply and demand for types of workers will not be identical across all occupations. A change in the demand for workers or in the supply of them will first cause an abrupt shortage or surplus. Wages in that occupation will adjust, and these wage signals will prompt location and/or training decisions to be made. Over time, the market will again come into balance without the need for intervention, at least according to theory. One ought to see a divergence in wage growth for an occupation relative to overall wages, then a convergence to the mean.

In six to 10 years time, persistent relative wage growth above or below the norm would be a sign that markets are not adjusting fast enough on their own, that, for example, new programs are not being expanded at the right rate to meet the higher level of jobs in the economy— even after jobseekers flood into them. In these cases there are chronic shortages, and public intervention in education and training policy may be needed. This does not mean markets cannot clear given more time and no further changes, but it is probably a signal that demand is changing continuously and faster than supply.

Therefore, monitoring relative wage growth among occupations is the preferred method for identifying shortages and surpluses of labor in any given market and ought to be adopted. At the same time, state and local occupational time series of quantity could be developed. The potential benefits of this approach are more promising than any attempts to collect and project supply and to better collect and, potentially, project true demand. At any rate, the latter would only reveal intentions, whereas transactions are liable to yield surer results and generate more manageable conclusions.

Implications for State and Area Long-Term Employment Projections

The current research has demonstrated that long-term projections really represent the quantity of labor that will exist at the market clearing price; they do not represent the demand for labor warranting independently-projected supply values. We noted that the information collected from companies and used to generate projections is the number of people employed, not the demand for labor as such (which is jobs plus unfilled openings); supply side data will not match until demand-side data are collected and produced differently in the first place. We noted also that in the long-term, both prices and quantity should adjust to alleviate any surpluses or shortages, while we acknowledge that we do not know with certainty if this theory holds. Existing shortages seen in job vacancy surveys might indeed persist, but this could be more readily tested using relative wage movements than utilizing the current method of collecting vacancy data. **Hence, there is no need to change the way projections are now done or to add occupational supply projections**. However, an overall labor supply constraint of some sort should be incorporated into the process for those areas that do not mirror national labor force growth trends or do not expect a similar business cycle.

However, the same customers who rely on long-term projections for information about worthwhile careers and promising training programs may also want more information than merely the likely number of jobs in various occupations and their expected rate of growth. They may also want an assessment of whether the current training system is liable to meet this expected level easily, if employers are likely to fill those jobs with underqualified individuals (a shortage) or if some people who receive such training find themselves without local employment in their chosen field (a surplus). Projections analysts may decide to delegate this work to other labor market specialists or conduct this research themselves. In either case, the best methods are indirect ones and include: 1) tracking years of vacancies from well-developed surveys and 2) studying trends in relative wages from wage and staffing surveys. Fortunately, both are being collected using the same coding basis. From a process point of view, determining where shortages and surpluses are likely to surface in future years does not have a natural home in long-term employment projections work— but this additional, valuable information could be generated by the same knowledgeable analysts.