United States Earnings Dynamics: Inequality, Volatility, and Mobility

John M Abowd and Kevin L McKinney

Using data from the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) infrastructure files, we study changes over time and across sub-populations in the distribution of real labor earnings. This proposal builds on our recent work using a consistently defined population frame to examine changes in earnings inequality and mobility (Abowd, McKinney, and Zhao 2018), as well as earnings volatility (McKinney and Abowd 2019). Results from these two papers show that earnings inequality is increasing, worker mobility is declining, and for workers with weak labor force attachment, earnings volatility is increasing.

While much is known about earnings inequality, mobility, and volatility at the national level, little is known about changes at the sub-national level. We aim to create a single unifying framework to examine these types of earnings dynamics for various worker sub-populations stratified by age, gender, race/ethnicity, and geography, thus paralleling as much as possible the level of detail available in the LEHD's Quarterly Workforce Indicators (QWI) product (i.e. county/MSA level geography and eight separate age groups). While not directly related to job-to-job mobility, worker earnings mobility is often tied to a job change. Our work will allow us to see the effect on the earnings distribution in specific geographic areas and/or age groups as workers leave one area and move to another, changing the earnings distribution in both locations. For smaller sub-populations, disclosure avoidance is a concern, one we will address using both existing technology and new methods as appropriate.

Unlike other recent periods of changing inequality, changes over the last 10 years have been largely determined by changes in the patterns of worker entry/exit and earnings mobility not changes in the pay structure. Our methodology starts by dividing the earnings distribution into bins, for example the bottom (<20th percentile), the middle (20th to the 80th percentile), and the top (80th percentile or above) with an additional bin for eligible workers with no earnings. Each year we calculate every eligible worker's real earnings and assign them to one of the bins. We also assign each worker to a mutually exclusive bin pair or mobility edge connecting the worker's earnings during one period of time to the worker's earnings during another non-overlapping period of time. The minimum period of time would be one year, but longer periods of time would also be used to assess low frequency changes in earnings and inequality. Keeping the real earnings bin boundaries fixed at a consistent set of values over time, we measure variation in inequality by observing the changing share of earnings and workers in each bin. We track earnings mobility by observing the earnings and number of workers in each mobility edge. Earnings volatility is estimated by calculating statistics such as the mean and variance of the change in earnings for the set of workers in each mobility edge. Using standard formulas for combining the mean and variance across mutually exclusive data partitions (in this case the mobility flows), we relate changes at specific parts of the earnings distributions to the overall change. Taken as a whole, this approach will result in a much richer understanding of the workers most affected by changes in earnings inequality, mobility, and volatility.