

Top Earners: Cross-Country Facts

Alejandro Badel, Moira Daly, Mark Huggett, and Martin Nybom

We provide a common set of life cycle earnings statistics based on administrative data from the United States, Canada, Denmark, and Sweden. We find three qualitative patterns, which are common across countries. First, top-earnings inequality increases over the working lifetime. Second, the extreme right tail of the earnings distribution becomes thicker with age over the working lifetime. Third, top lifetime earners exhibit dramatic earnings growth over their working lifetime. Models of top earners should account for these three patterns and, importantly, for how they quantitatively differ across countries. (JEL D31, D91, H21, J31)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2018, 100(3), pp. 237-57.
<https://doi.org/10.20955/r.100.237-57>

1 INTRODUCTION

Over the past 100 years, the inequality of top incomes has followed a U-shaped pattern in the United States, the United Kingdom, and Canada. The recent increase in top-income inequality has become an important topic in academic, policy, and media discussions in these countries. In other countries, such as Denmark, France, and Sweden, income inequality also decreased strongly in the first half of the twentieth century but did not rebound strongly afterward. Figure 1 plots the top 1 percent income share for all these countries.¹

Wage and salary income play a very important role in shaping top-income inequality patterns. First, wage and salary income has been the largest component of top incomes in the United States and Canada in recent decades (see Piketty and Saez, 2003, and Saez and Veall, 2005). Second, income inequality patterns resemble earnings inequality patterns over time. For example, Figure 1 shows that both top income and earnings shares in the United States


Alejandro Badel is an adjunct professor at Georgetown University and was an economist at the Federal Reserve Bank of St. Louis; Moira Daly is a senior advisor at Copenhagen Business School; Mark Huggett is a professor at Georgetown University; and Martin Nybom is an associate professor at the Institute for Evaluation of Labour Market and Education Policy. The authors acknowledge financial support from Danish Social Science Research Council (FSE) grant No. 300279 and support from the Federal Reserve Bank of St. Louis. They also thank Bryan Noeth for research assistance.

© 2018, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.


Figure 1

Basic Top-End Inequality Facts


A. Top 1 Percent Income Share: United States, United Kingdom, and Canada



B. Top 1 Percent Income Share: France, Denmark, and Sweden



C. Earnings and Income Shares: United States



NOTE: The income measure excludes capital gains, and the earnings measure is based on wages and salaries. For the United Kingdom, the sampling unit was changed in 1990, and there is a jump in the series in that year.

SOURCE: Income comes from The World Wealth and Income Database. The earnings measure for the United States is from Piketty and Saez (2003 update).

have increased over time starting before 1980. For these reasons, discussions of the determinants of top-income inequality over time and across countries have focused on theories of top-earnings inequality.

The goal of this article is to document a common set of facts concerning the dynamics of the earnings distribution over the working lifetime. We focus on the United States, Canada, Denmark, and Sweden. For these four countries, administrative data on earnings are available to researchers under strict privacy protection arrangements. The datasets we employ have the following four common features: They are large, they do not truncate earnings, they cover

several decades, and, importantly, they track individuals over time. These features allow us to document the top of the earnings distribution by age or by birth cohort. They also allow us to observe the annual earnings of individuals for more than 30 years of their working lifetimes.

We find that the life cycle evolution of the earnings distribution for males follows three patterns, which are common across countries. First, top-earnings inequality increases over the working lifetime. Second, the extreme right tail of the earnings distribution becomes thicker with age over the working lifetime. Third, top lifetime earners exhibit dramatic earnings growth between their early and late working years.² There are important differences in the magnitudes of these facts across countries.

The patterns that we document provide empirical guidance for the specification and calibration of quantitative theoretical models aimed at understanding the distribution of earnings, income, and wealth within a given country. For many existing models of earnings distributions, these patterns also provide a challenge because these models lack forces generating extremely large earnings growth rates for top lifetime earners. The cross-country facts also provide a new challenge for quantitative theoretical work directed at understanding the underlying sources of cross-country differences in inequality. Ideally, a plausible quantitative theory should be able to account for cross-country differences in cross-sectional inequality and, simultaneously, account for the substantial cross-country differences in the three life cycle earnings facts that we document.

This article is closest to two literatures. First, there is a large literature that documents the life cycle evolution of the distribution of earnings, wages, and consumption.³ This literature documents how summary measures of dispersion, such as the variance of log earnings, wages, or consumption, vary with age based on survey data, controlling for time or cohort effects. Our work focuses on quantiles of the earnings distribution by age and properties of the top 1 percent by age. Focusing on quantiles is useful because these can fully describe a distribution. Much larger sample sizes and the lack of top coding allow us to address the behavior of the top 1 percent of the distribution by age. The very top of the distribution is critical for optimal tax theory (see Piketty and Saez, 2013, and Badel and Huggett, 2017) as specific statistics of the top of the distribution enter formulae that determine optimal top tax rates. Second, a recent literature uses administrative data to describe the top of the earnings distribution over time. See, for example, Guvenen, Kaplan, and Song (2014) and Guvenen et al. (2015). We differ because we focus on how three life cycle facts differ across countries.

This article is organized into four sections. Section 2 describes basic features of each dataset and provides inequality facts. Section 3 documents three facts that characterize the dynamics of earnings over the working lifetime. Section 4 discusses the ability of existing quantitative models of earnings and labor productivity to produce the three life cycle earnings facts that we document.

2 DATA

This section describes the earnings data, the samples, and some background facts.

2.1 Earnings Data

Our earnings data come from records kept by government agencies for administrative purposes. These datasets are not publicly available and are accessible only under special arrangements that protect personally identifiable information. Except for the United States, we directly access each country's microdata via the relevant statistical agency. For the United States, we lack access to the microdata, so we use the summary tables provided by Guvenen, Ozkan, and Song (2014) and Guvenen et al. (2015).

The U.S. summary tables are based on data from W-2 forms of wage and salary workers held by the Social Security Administration. Their earnings measure includes wages and salary, bonuses, and exercised stock options. The data consist of a 10 percent random sample of males with a Social Security number in the period 1978-2011. The summary tables include minimum, maximum, mean, and various percentiles of the earnings distribution for each year and include percentiles by age and year.

The earnings data for Canada come from the Longitudinal Administrative Databank (LAD) administered by Statistics Canada. LAD is a 20 percent random sample of the Canadian population covering the period 1982-2013. The earnings measure we employ is total earnings from T4 slips plus other employment income. T4 slips are issued by employers to the Canadian Revenue Agency and contain employment income and taxes deducted. T4 slips include wages, salaries and commissions, and exercised stock option benefits. Other employment income includes tips, gratuities, and director's fees not included in T4 slips.

The tax registers for Denmark are provided by Statistics Denmark. The sample period is 1980-2013. Over the sample period, the registers provide panel data on earnings for more than 99.9 percent of Danish residents between ages 15 and 70. We focus on individuals never classified as immigrants in the data. The earnings measure we employ is the sum of two variables in the registers. The first variable measures taxable wage payments and includes fringe benefits, jubilee and termination benefits, and the value of exercised stock options.⁴ It excludes contributions to pension plans and to ATP (the Danish Labour Market Supplementary Pension). The second variable is ATP contributions.

Earnings data for Sweden are provided by Statistics Sweden. We have access to earnings data for 1980, 1982, and 1985-2013. The data cover the entire Swedish population with taxable income in a given year. The earnings measure is based on taxable labor market earnings reported by individuals' employers to the national tax authority.⁵

2.2 Sample Selection

Cross-sectional samples are used to produce statistics by year or by age and year. Our cross-sectional samples for Canada, Denmark, and Sweden are designed to mimic the sample selection criteria employed in the U.S. sample. Thus, we employ harmonized samples that allow cross-country comparisons.

The U.S. cross-sectional sample includes an individual earnings observation in a given year t if (i) the individual is a male age 25 to 60, (ii) earnings are greater than a time-varying threshold, denoted \underline{e}_t^{US} , and (iii) self-employment income does not account for more than 10 percent of the earnings and does not exceed the \underline{e}_t^{US} threshold. The threshold \underline{e}_t^{US} employed by Guvenen et al. (2014, 2015) is defined as half the minimum hourly wage in year t times 520 hours.

Our cross-sectional samples for Canada, Denmark, and Sweden implement these three criteria: First, each sample includes only males age 25 to 60. Second, an earnings observation is included for a given country if it exceeds a threshold (\underline{e}_t^{CA} , \underline{e}_t^{DK} , \underline{e}_t^{SW}). Third, we implement the self-employment income criteria described above.⁶

We provide a method to obtain harmonized samples across countries. For each country $i \in \{CA, DK, SW\}$ and year t , we calculate the minimum earnings threshold as the product of a common factor at time t , denoted $factor_t$, and median earnings $median_t^i$:

$$\underline{e}_t^i = factor_t \times median_t^i.$$

The common $factor_t$ is based on the U.S. threshold and U.S. median earnings as follows:

$$factor_t = \underline{e}_t^{US} / median_t^{US}.$$

2.3 Background Facts

We document a number of earnings facts based on our cross-sectional samples. Figure 2 shows that, over the full sample period, the share of earnings obtained by the top 1 percent is substantially higher in the United States and Canada than in Denmark and Sweden. Furthermore, top-earnings shares trend upward in the United States and Canada over the sample period. Top-earnings shares in Denmark and Sweden also increased over the sample period but by much less than in the United States and Canada.⁷ The top income share patterns in Figure 1 resemble the earnings patterns we document.


Figure 2 shows that the earnings distribution above the median in both Denmark and Sweden is more compressed compared with that for the United States. Each of the 90-50 earnings ratios for Denmark and Sweden is about three quarters of the U.S. ratio, whereas each of the 99-50 earnings ratios for Denmark and Sweden is roughly half the U.S. ratio. Thus, compression is stronger above the 90th percentile in these countries. Dividing one-half by three quarters implies that the 99-90 ratios in Denmark and Sweden have been roughly two-thirds of the U.S. 99-90 ratio. Figure 2 also shows that earnings dispersion above the 50th percentile increases in all countries over time. Specifically, over the sample period, the 90-50 and the 99-50 earnings percentile ratios increase for all countries.

Figure 2 documents the evolution of the Pareto statistic of earnings at the 99th percentile over time. This statistic is defined as $\bar{e}_{99}/(\bar{e}_{99} - e_{99})$. That is, mean earnings beyond the 99th percentile, \bar{e}_{99} , divided by the difference between \bar{e}_{99} and the 99th percentile, e_{99} . The figure also shows that the Pareto statistic at the 99th percentile has trended downward in all countries over the sample period. A lower value for the Pareto statistic implies a thicker upper tail


Figure 2

Top-End Earnings Inequality Facts


A. Top 1 Percent Share




B. 90-50 Ratio



C. 99-50 Ratio



D. Pareto at 99th Percentile



NOTE: For the United States, the top 1 percent share and the Pareto statistic in each year are based on the assumption of a Pareto distribution within the top 1 percent and tabulated values for the 99th and 99.999th percentiles.

SOURCE: Authors' calculations based on the cross-sectional samples for each country.

in the sense that the mean, for observations above the threshold, is a higher multiple of the threshold. The Pareto statistic is particularly important in theories of taxation of top incomes or top earnings. It enters into formulas used to determine welfare- or revenue-maximizing top tax rates (see Piketty and Saez, 2013, and Badel and Huggett, 2017). Lower values of the Pareto statistic imply, other things equal, a higher revenue-maximizing top tax rate.

3 EARNINGS FACTS

We document the evolution of the earnings distribution over the working lifetime with a focus on properties of the upper tail of the distribution.

3.1 Fact 1: Top-Earnings Inequality Increases with Age

We determine how the earnings distribution above the median evolves with age. For example, we calculate the 99-50 earnings percentile ratio $e_{99,j,t}/e_{50,j,t}$ for all ages j and all sample years t . We then estimate the time and age effects (α_t, β_j) or, alternatively, the cohort and age effects (γ_c, β_j) in the regressions below. An individual's birth year (i.e., cohort) is denoted c . Clearly, cohort c , current age j , and current year t are linearly related: $c = t - j$. The cohort-effects regression controls for cohort-specific effects that impact the 99-50 ratio for a cohort at any age, whereas the time-effects regression controls for time-specific effects that impact the 99-50 ratio for all age groups alive at that time. The variables D_j, D_t, D_c are dummy variables that take the value 1 when the observation occurs at age j , year t , or cohort c , respectively. We employ a full set of age, year, and cohort dummy variables:

$$\text{Time Effects: } e_{99,j,t}/e_{50,j,t} = \alpha_t D_t + \beta_j D_j + \varepsilon_{j,t};$$

$$\text{Cohort Effects: } e_{99,j,t}/e_{50,j,t} = \gamma_c D_c + \beta_j D_j + \varepsilon_{j,t}.$$

We use the estimated age effects $\hat{\beta}_j$ to describe how the 99-50 earnings percentile ratio evolves with age. We plot the estimated age coefficients adjusted by a constant $\hat{\beta}_j + k$. The constant k is chosen so that the height of the age profile at age 45 equals the empirical 99-50 ratio for those age 45 years in 2010 for each country.⁸


Figure 3 presents the results. The main finding is that the 90-50 and the 99-50 ratios tend to increase with age in all countries. In this sense there is fanning out in the top half of the distribution with respect to the median in all countries. The cohort-effects view produces a more dramatic pattern of fanning out compared with the time-effects view. The most striking pattern occurs for the 99-50 ratio. First, the 99-50 ratio is much larger at any age in the United States and Canada compared with Denmark and Sweden. Second, the 99-50 ratio roughly doubles from age 25 to age 55 in each country under the cohort-effects view. Thus, we conclude that there is growing earnings dispersion with age above the median and that this is driven by earnings beyond the 90th percentile.

Many studies have documented growth in summary measures of earnings or income dispersion with age for individuals or households based on dispersion measures such as the variance of log earnings or the Gini coefficient. The results in Figure 3 indicate that one reason summary measures display growing dispersion with age is because of the behavior of the very top of the distribution compared with the median.

To put these results into perspective, it is useful to characterize how real median earnings evolve with age.² Figure 4 provides the results of regressing real median earnings by age and time effects or age and cohort effects. Median earnings display a hump-shaped pattern with

Figure 3

Percentile Ratios: 90-50 and 99-50 Ratios by Age



NOTE: The figure plots the estimated age coefficients after adding a vertical shift term, so each figure is normalized to equal the data value of the 99-50 ratio or the 90-50 ratio at age 45 in the year 2010.


age in each country. Many previous studies have documented that male earnings or wage rates by age are hump-shaped over the working life.¹⁰

Figure 4 shows that median earnings in the United States and Canada approximately double with age from age 25 to 50. This holds regardless of whether one controls for time or for cohort effects. In contrast, for Denmark and Sweden the time-effects view implies that the median earnings profile is flatter, with less than a doubling of median earnings. Focusing on the time-effects view across countries reveals substantial differences in the timing of the peak of the earnings profile. For the United States and Canada, median earnings peak near age 50, whereas for Denmark and Sweden the peak occurs in the early 40s.


Figure 4

Median Earnings by Age


A. United States




B. Canada



C. Denmark



D. Sweden



NOTE: The figure plots the estimated age coefficients after adding a vertical shift term, so each figure is normalized to equal 100 at age 45.

3.2 Fact 2: The Upper Tail Becomes Thicker with Age


Next we analyze how the Pareto statistic at the 99th percentile evolves with age. This is a way to describe how the thickness of the upper tail of the earnings distribution evolves with age. To do so, we run the two basic regressions from the previous section after replacing ratios of earnings percentiles with the Pareto statistic for each age-year pair.

Figure 5 shows that the Pareto statistic declines with age in all countries. This holds in both the time- and cohort-effects regressions. Thus, the upper tail of the earnings distribution becomes thicker with age in each country in the sense that mean earnings beyond this thresh-


Figure 5

Pareto Statistic at the 99th Percentile by Age

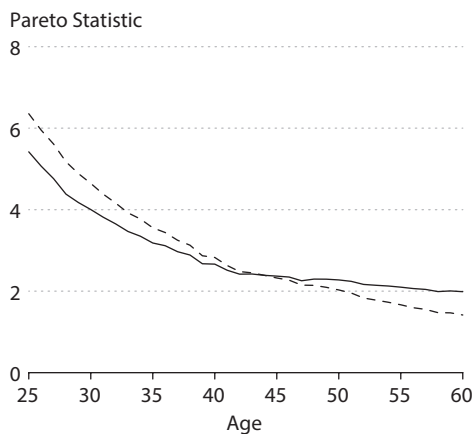
A. United States



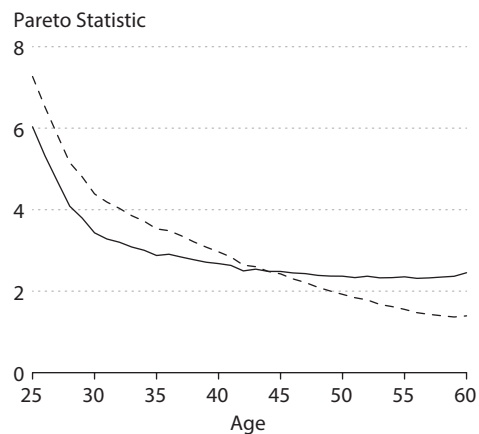
B. Canada



C. Denmark



D. Sweden



NOTE: The figure plots the estimated age coefficients after adding a vertical shift term, so each figure is normalized to equal the data value of the Pareto statistic at age 45 in the year 2010.


old is a growing multiple of the threshold with age. To the best of our knowledge, this fact has not been documented in the existing literature for a wide collection of countries.

It is interesting to compare the Pareto statistic in different age groups with the Pareto statistic in cross-sectional data previously documented in Figure 2. For the United States, the Pareto statistic at the 99th percentile in cross-sectional data is below 2 in the last two decades of the sample period. In Figure 5, it is below 2 in the United States for age groups above age 40, while it is above 2 for age groups below age 40. This suggests that the cross-sectional Pareto statistic for the United States is largely determined by the earnings distribution for males


Figure 6

Earnings Growth by Lifetime Earnings Group


A. United States




B. Canada



C. Denmark



D. Sweden



NOTE: The figure plots the ratio of mean group earnings at age 55 to mean group earnings at age 25, as well as the ratio of mean group earnings at age 55 to mean group earnings at age 30 for groups sorted by percentile of lifetime earnings.

SOURCE: U.S. data are from Guvenen et al. (2015). The results for the other countries are based on our calculations from country longitudinal data.

age 40 and older. The same patterns hold in Canadian data. Thus, the cross-sectional Pareto statistic seems to be driven by the tail properties holding for older earners in both countries.

3.3 Fact 3: Top Lifetime Earners Have Dramatic Earnings Growth

We now use the longitudinal feature of each dataset. For each male in the longitudinal sample, we compute lifetime earnings LE as follows: $LE^i = \sum_{t \in T} \frac{\max\{e_t^i, \underline{e}_t\}}{p_t}$, where e_t^i is individual i 's nominal earnings in year t , \underline{e}_t is the minimum earnings threshold used to con-

struct the cross-section sample, p_t is a country price index in year t , and T is the set of years for which earnings observations are available.¹¹ We then sort males in the longitudinal sample into 100 bins based on the percentiles of the lifetime earnings distribution. Bin 100 corresponds to males with lifetime earnings above the 99th percentile, whereas bin 1 corresponds to males with lifetime earnings below the 1st percentile. Appendix A.1 describes the construction of the longitudinal data samples.

Figure 6 contains two plots for each country. It plots the ratio of mean real earnings at age 55 to mean real earnings at age 25 for individuals sorted by lifetime earnings bins, as well as the ratio of mean real earnings at age 55 to mean real earnings at age 30. In both plots the grouping of individuals into lifetime earnings bins is unchanged. Thus, for a given country, the two plots differ only insofar as there is growth in real mean earnings for the group from age 25 to age 30.

Figure 6 documents that earnings growth is greater for groups with larger lifetime earnings. It also documents the remarkable fact that the highest lifetime earnings groups (i.e., groups in lifetime earnings bins 96-100) have a much larger earnings growth rate than those with lifetime earnings close to the median (i.e., those in bin 50). The top lifetime earnings bin in the United States and Canada have a 13- to 15-fold increase in earnings from age 25 to age 55. The top lifetime earnings bin in Denmark and Sweden have a seven- to ninefold increase in earnings from age 25 to age 55. Thus, there are large, systematic differences in group earnings growth rates over the working lifetime, particularly at the very top. The large differences at the top imply that in each country, top lifetime earners tend to become top earners late in the working lifetime. We anticipate that Fact 3 will be particularly useful in empirically disciplining quantitative theories of top earners. We conjecture that theories built on temporary sources of earnings variation will struggle to produce Fact 3.

4 DISCUSSION

We close the article by discussing the potential relevance of the three earnings facts that we document for economic models of the distribution of earnings and wage rates over the working lifetime. We do so by briefly discussing two prominent articles that offer a quantitative-theoretical account of the changes in U.S. cross-sectional inequality measures.

4.1 Models of Changes in Cross-Sectional Inequality

Heathcote, Storesletten, and Violante (2010) provide a quantitative-theoretical account for the changes in U.S. cross-sectional earnings, consumption, and hours inequality. The key exogenous driving force in their model is changes in transitory and persistent idiosyncratic productivity shocks. They measure the time-varying variances of these shocks from panel data on U.S. wage rates. They find that both transitory and persistent innovation variances increase over time. They then show that their model accounts for the rise in measures of U.S. household earnings and consumption dispersion, among other facts, based on the measured process for productivity shocks.

Kaymak and Poschke (2016) provide a quantitative-theoretical account for changes in U.S. top-end wealth inequality over the past half century. They consider these three exogenous sources for the increase in U.S. wealth inequality: changes in taxes, transfers, and productivity shocks. They measure changes in U.S. corporate, estate, and income taxes over time, and they measure changes in the level and progressivity in Social Security benefits. Finally, they calibrate an idiosyncratic productivity shock process to match evidence for the rise in U.S. earnings/wage dispersion over time. Their shock process captures persistent and transitory sources of variation. They find that the rise in wage dispersion, the change in taxes (i.e., decrease in some top tax rates), and the increase in transfers all contributed to the increase in top-end U.S. wealth inequality. They find a particularly important contribution from the increase in top-end wage dispersion.

4.2 Idiosyncratic Productivity Shocks

The articles by Heathcote, Storesletten, and Violante (2010) and Kaymak and Poschke (2016) stress the role of idiosyncratic productivity shocks. Productivity in their models corresponds to wage rates in the data. We now compare properties of the process used in these articles with the facts for earnings that we document.

While earnings and wage rates are not strictly comparable, we think the comparison is still useful. Many age patterns in wage rate data also hold in earnings data. For example, Heathcote, Storesletten, and Violante (2005) show that the rise in both the variance of log earnings and the variance in log wage rates happens with age in U.S. data by similar amounts. In addition, cross-sectional inequality in log earnings and in log wage rates rises by a similar magnitude as documented by Heathcote, Storesletten, and Violante (2010). Finally, it is widely believed that productivity differences (i.e., earnings per work hour) are key in accounting for the earnings of top earners in U.S. data rather than work-hour differences.

The process used in each article is summarized below. Kaymak and Poschke (2016) model a worker's productivity as a finite Markov process, where the transition probabilities are given by the matrix Π . Productivity w takes on six values (z_1, \dots, z_6), where z_6 corresponds to an extraordinarily high level of productivity.¹² Heathcote, Storesletten, and Violante (2010) model log productivity as the sum of an age component μ_{j+1} , a persistent shock η_{j+1} , and a purely transitory shock v_{j+1} . The age component is common to all agents of age $j+1$, whereas the shock components are agent specific:

$$(KP) \text{Prob}(w_{j+1} = z' | w_j = z) = \Pi(z' | z)$$


$$(HSV) \log w_{j+1} = \mu_{j+1} + \eta_{j+1} + v_{j+1}, \text{ and } \eta_{j+1} = \rho\eta_j + \omega_{j+1}.$$

We now simulate 2 million wage histories from age 20 to age 60 using the Kaymak-Poschke process above. The inputs are an initial distribution, the workers matrix Π above, and the six productivity values.¹³ We highlight the implications of the Kaymak-Poschke process for Fact 3 from Section 3.3.

Figure 7 presents ratios of earnings across ages for different lifetime earnings groups in the Kaymak-Poschke model and in U.S. data from Figure 6. A measure of lifetime earnings is

Figure 7

Earnings Ratios: Kaymak-Poschke Model and U.S. Data



NOTE: Figure 7A plots the ratio of mean group earnings at age 55 to mean group earnings at age 25 in both U.S. data and the Kaymak-Poschke model. Figure 7B repeats this plot but uses data at age 55 and age 30. The horizontal axis sorts males and model agents by percentiles of lifetime earnings.

SOURCE: Kaymak and Poschke, 2016.

computed for each agent in the model based on earnings from age 25 to age 60. Agents are then placed into 100 bins according to their percentile of lifetime earnings. Thus, U.S. data and model data are treated symmetrically.

Figure 7 shows that the ratios in the model data are typically below the corresponding ratios found in U.S. earnings data. This holds most strikingly for the several highest lifetime earnings bins. For the highest bin in the Kaymak-Poschke model, the earnings ratio is below 2.5 for both the 55-25 age ratio and the 55-30 age ratio. In contrast, the corresponding ratios for the highest earnings bin in U.S. data are roughly 15 and 7.

One possible reason for the difference between the model and U.S. data is that there is mean reversion at the highest productivity state z_6 back to lower productivity states. Such mean reversion is one reason the model successfully concentrates a large fraction of wealth held by top wealth holders similar in magnitude to that found in U.S. data. Agents with shock z_6 save a large part of their labor income because this state is, to an important degree, transitory.¹⁴

We conjecture that models that rely only on purely temporary sources of earnings variation to account for the extreme right tail of the earnings distribution will also fail to produce the strong earnings growth for top lifetime earners documented in Figure 6. We suspect that simulations of productivity from the Heathcote-Storesletten-Violante model will also be below the patterns found in U.S. data for top lifetime earners. This is because their model relies on a persistent but mean-reverting component and a purely temporary component to account for the right tail of the productivity distribution.

We conjecture that models that allow for systematic differences in earnings growth over the working lifetime will be important to account for the earnings profiles of top lifetime

earners. Human capital models are promising in this regard. Specifically, some human capital models allow agents to permanently differ in learning ability. Those with high learning ability optimally choose steeper mean earnings profiles via an investment in skill formation. Badel, Huggett, and Luo (2018) provide a model with this feature that can produce the properties that we document in Facts 1-3. Learning ability differences also help produce Fact 2—the fall in the Pareto statistic with age. The mechanism is the same. High productivity agents make skill investments, even late in the working lifetime, and these investments are a force that cause the earnings of agents above the 99th percentile within an age group to grow faster than those at the 99th percentile. ■

APPENDIX A

A.1 Longitudinal Samples

For Canada, our raw data consist of all individuals in the LAD dataset. The LAD is a 20 percent random subsample from the Canadian population that either filed a T1 form or received Canadian child benefits in any year since 1982 and had a social insurance number.¹⁵ For Denmark we use tax registry data kept by Statistics Denmark. For Sweden we use tax registers kept in the Income and Taxation Register of Statistics Sweden. These data come from the Swedish Tax Agency, which collects information from virtually all persons who are Swedish citizens or who hold a residence permit.

We construct a longitudinal sample for Canada, Denmark, and Sweden. These three samples mimic the construction of the U.S. longitudinal sample described in Guvenen et al. (2015). The sample period is 1982-2013 for Canada, 1980-2013 for Denmark, and 1980, 1982, and 1985-2013 for Sweden. Thus, the sample period for each country spans a horizon of more than 30 years.

Our longitudinal sample for each of these three countries contains all individual histories that satisfy Conditions 1-4 below. The following notation is employed: e_t^i is individual i 's nominal earnings, \underline{e}_t is a minimum nominal earnings threshold, and se_t^i is individual i 's self-employment income. (1) The individual is male and age 24, 25, or 26 in the first year of the sample period. (2) The individual has a valid non-missing earnings observation in every year of the sample period. (3) There are more than 15 years for which $e_t^i > \underline{e}_t$. (4) There are less than 9 years for which $se_t^i > \max\{\underline{e}_t, 0.1e_t^i\}$.

We now provide a brief discussion of the specifics of imposing Conditions 1-4 in the longitudinal samples for each country. Condition 1 is straightforward to implement. All properties of mean earnings for groups by age are understood to be for the central age within the group. Condition 3 is straightforward to implement in each country. We simply employ the threshold used in the construction of each cross-sectional sample. We implement Condition 4 in Canada and Denmark by using the self-employment income measure described in Section 2 and employed in the construction of the cross-sectional sample. The longitudinal samples contain the following number of males after rounding to the nearest 100: (1) Canada 65,000, (2) Denmark 73,300, and (3) Sweden 143,400.

A.2 Pareto Statistic from SSA Data

Pareto statistics at the 99th percentile are not provided by Guvenen et al. (2014, 2015). Based on the statistics provided, we estimate the Pareto statistics for the United States in two different ways. First, for the Pareto statistics depicted in Figure 2D, we use the 99th and 99.999th percentiles of earnings, provided by Guvenen et al. (2014) for each sample year, to estimate the coefficient of a Type-I Pareto distribution for earnings above the 99th percentile. Such coefficient is the Pareto statistic. For the Pareto statistic at the 99th percentile by age group and year used to create the life cycle profiles in Figure 5, we employ the method described in Badel, Huggett, and Luo (2018), which uses the 95th and 99th percentiles that are provided by age group and year, to estimate a Pareto distribution for earnings above the 95th percentile.

A.3 Price Index Data Sources

Sources for price indexes are as follows:

Canada: Series number CPALCY01CAA661N, Consumer Price Index: Total, All Items for Canada[®], Index 2010 = 1, Annual, Not Seasonally Adjusted; <https://fred.stlouisfed.org>.

Denmark: Available from Statistics Denmark's StatBank Denmark; <http://www.statbank.dk/statbank5a/default.asp?w=1280>. Year 2000 was the base year used.

Sweden: Statistics Sweden's official consumer price index series; <http://www.scb.se/hitta-statistik/statistik-efter-amne/priser-och-konsumtion/konsumentprisindex/konsumentprisindex-kpi/pong/tabell-och-diagram/konsumentprisindex-kpi/kpi-faststallda-tal-1980100/>.

A.4 Descriptive Statistics

Tables A1-3 present descriptive statistics for Canada, Denmark, and Sweden, respectively, for the cross-sectional sample.

Table A1**Summary Statistics: Cross-Sectional Samples, Canada**

Year	<i>nobs</i>	<i>mean e</i>	<i>e50</i>	<i>e99</i>
1982	886,920	24,400	23,100	72,000
1983	886,310	25,400	24,200	74,800
1984	902,625	26,800	25,500	79,600
1985	914,700	28,100	26,700	84,300
1986	947,720	29,200	27,600	89,200
1987	956,945	30,700	28,800	95,300
1988	983,375	32,800	30,200	106,600
1989	1,012,465	34,700	31,600	114,500
1990	1,028,790	35,400	32,200	117,000
1991	1,022,650	36,000	32,900	119,300
1992	1,024,415	36,700	33,600	121,700
1993	1,028,755	37,300	33,800	124,900
1994	1,033,960	38,100	34,300	130,800
1995	1,044,510	39,100	34,900	139,000
1996	1,048,970	40,000	35,300	147,100
1997	1,058,555	41,900	36,100	160,500
1998	1,065,610	43,600	37,100	174,300
1999	1,084,320	45,100	38,100	182,600
2000	1,101,815	47,800	39,400	200,700
2001	1,140,225	49,000	40,200	210,500
2002	1,137,365	49,600	41,000	208,800
2003	1,149,010	50,800	42,000	214,300
2004	1,162,555	52,700	43,100	225,900
2005	1,177,270	55,200	44,400	243,500
2006	1,186,490	57,900	45,900	261,600
2007	1,199,525	60,000	47,400	273,200
2008	1,210,295	61,300	48,900	270,900
2009	1,201,615	59,500	48,000	257,800
2010	1,200,940	61,400	49,400	264,300
2011	1,221,400	63,700	51,100	275,600
2012	1,233,235	65,300	52,600	279,000
2013	1,241,750	67,000	53,900	284,200

NOTE: Earnings statistics are rounded to the nearest 100 for confidentiality. The notations *nobs*, *mean e*, *e50*, and *e99* denote the number of observations, mean earnings, and the 50th and 99th earnings percentiles, respectively.

SOURCE: Statistics Canada.

Table A2**Summary Statistics: Cross-Sectional Samples, Denmark**

Year	<i>nobs</i>	<i>mean e</i>	<i>e50</i>	<i>e99</i>
1980	871,620	118,228	113,579	294,873
1981	859,167	127,065	122,948	318,449
1982	866,315	141,548	137,413	352,795
1983	879,347	151,118	146,798	379,600
1984	890,302	160,415	154,354	408,569
1985	906,252	169,582	161,345	435,802
1986	917,972	181,792	172,378	464,136
1987	924,403	196,046	185,319	508,165
1988	926,431	206,465	195,574	535,390
1989	927,703	212,492	200,965	560,340
1990	936,043	217,956	206,008	579,697
1991	935,039	223,188	211,658	589,631
1992	943,109	228,499	217,469	605,339
1993	941,600	228,247	218,277	606,725
1994	951,024	239,598	227,380	647,593
1995	962,977	248,388	234,389	675,665
1996	972,286	255,076	240,615	695,749
1997	983,871	264,676	248,812	725,424
1998	998,120	273,832	255,626	765,677
1999	1,005,814	285,656	266,406	805,958
2000	1,011,325	296,760	275,335	850,898
2001	1,012,968	307,839	284,840	891,303
2002	1,009,869	315,493	292,833	917,934
2003	999,303	320,344	298,258	936,881
2004	993,586	328,709	306,149	960,951
2005	990,605	338,733	314,951	996,671
2006	989,524	351,789	325,779	1,040,338
2007	984,137	368,622	339,457	1,101,163
2008	969,799	386,238	353,782	1,160,732
2009	942,820	382,500	356,390	1,127,868
2010	923,739	402,285	365,001	1,304,489
2011	918,254	409,964	371,298	1,325,142
2012	913,586	417,780	376,711	1,355,580
2013	911,549	424,334	379,914	1,398,553

NOTE: The notations *nobs*, *mean e*, *e50*, and *e99* denote the number of observations, mean earnings, and the 50th and 99th earnings percentiles, respectively. All percentiles calculated from Danish data are six observation local averages, a confidentiality requirement.

SOURCE: Statistics Denmark.

Table A3**Summary Statistics: Cross-Sectional Samples, Sweden**

Year	<i>nobs</i>	<i>mean e</i>	<i>e50</i>	<i>e99</i>
1980	1,845,140	79,441	75,031	217,579
1981	—	—	—	—
1982	1,830,333	91,005	86,952	250,336
1983	—	—	—	—
1984	—	—	—	—
1985	1,615,820	118,882	111,795	311,643
1986	1,627,315	128,328	121,011	332,480
1987	1,644,682	138,457	130,327	363,534
1988	1,665,408	149,939	141,447	390,681
1989	1,691,587	164,961	156,272	423,839
1990	1,871,002	175,721	167,993	466,135
1991	1,898,011	187,592	178,511	524,915
1992	1,875,173	187,430	180,315	531,928
1993	1,840,234	189,948	183,420	554,667
1994	1,838,130	197,666	189,330	602,450
1995	1,856,135	204,850	197,545	594,136
1996	1,857,699	215,984	206,560	637,712
1997	1,860,797	225,708	215,075	672,172
1998	1,883,857	235,753	223,189	710,840
1999	1,914,785	244,286	229,501	745,901
2000	1,945,461	257,441	239,017	803,702
2001	1,962,558	270,076	248,633	853,088
2002	1,963,068	277,617	257,149	870,017
2003	1,945,148	282,512	263,059	880,435
2004	1,928,007	288,791	269,771	910,426
2005	1,914,243	299,104	278,085	946,117
2006	1,917,082	311,050	288,703	991,729
2007	1,913,805	325,196	300,023	1,043,752
2008	1,906,596	340,417	312,946	1,091,904
2009	1,875,741	342,493	316,552	1,097,314
2010	1,871,732	352,699	326,069	1,125,863
2011	1,885,636	365,852	336,988	1,166,034
2012	1,886,082	375,238	346,504	1,179,256
2013	1,886,746	381,534	353,416	1,193,581

NOTE: The notations *nobs*, *mean e*, *e50*, and *e99* denote the number of observations, mean earnings, and the 50th and 99th earnings percentiles, respectively.

SOURCE: Statistics Sweden.

NOTES

- ¹ Roine, Vlachos, and Waldenstrom (2009) and Alvaredo et al. (2013), among others, have documented inequality patterns over the past 100 years for many developed countries, including those in Figure 1.
- ² Lifetime earnings are defined as a present value (or weighted sum) taken over the full history of annual earnings of a worker's life. Top lifetime earners are those in the top 1 percent of the lifetime earnings distribution.
- ³ See Deaton and Paxson (1994), Storesletten, Telmer, and Yaron (2004), Heathcote, Storesletten, and Violante (2005) or Huggett, Ventura, and Yaron (2011) for the United States; Creedy and Hart (1979) and Blundell and Etheridge (2010) for the United Kingdom; Brzozowski et al. (2010) for Canada; and Domeij and Floden (2010) for Sweden.
- ⁴ This variable, labeled LOENMV in the registers, has changed coverage over time. For example, the value of exercised stock options were not included prior to 2000.
- ⁵ The earnings measure comes from Statistics Sweden variable ARBINK up to 1985 and from variable LONEINK thereafter. These measures include some labor-related benefits such as parental leave benefits and short-term sick leave benefits. Variable LONEINK includes income from closely held businesses starting in 1994. Part of the value of realized stock options are included in the earnings measure.
- ⁶ For Canada, self-employment income is measured with the LAD variable SEI, which measures the sum of net income from self-employment. For Denmark, self-employment income is measured with the Statistics Denmark variable NETOVSKUDGL. For Sweden it is measured with variable FINK, which measures net entrepreneurial income.
- ⁷ Domeij and Floden (2010) provide evidence that the 99-50 earnings percentile ratio, based on family earnings, rises from about 1.6 to 1.8 in Sweden between 1990 and 2000.
- ⁸ For the United States, the available summary tables contain data for $j \in \{25, 35, \dots, 55\}$, so estimating one age coefficient β_j for each $j = 25, 26, 27, \dots, 60$ is not possible. Therefore, we replace the age effects β_j in the regressions above with a third-order polynomial in age $P(j; \theta) = \theta_0 + \theta_1 j + \theta_2 j^2 + \theta_3 j^3$ and set the estimated age effects to $\hat{\beta}_j = P(j; \hat{\theta})$, where $\hat{\theta}$ are the estimated polynomial coefficients.
- ⁹ Appendix A.3 states sources for the price indexes that are used to deflate earnings.
- ¹⁰ For example, the *Review of Economic Dynamics* special issue on Cross-Sectional Facts for Macroeconomists in 2010 covers nine countries, and Lagakos et al. (2016) covers 18 countries.
- ¹¹ The set T^{US} is based on 1978-2011, T^{CA} is based on 1982-2013, T^{DK} is based on 1980-2013, and T^{SW} is based on 1980, 1982, and 1985-2013. Price indexes are in Appendix A.3.
- ¹² The process sets $(z_1, z_2, z_3, z_4, z_5, z_6) = (6.7, 19.2, 20.5, 58.4, 61.4, 1222)$.
- ¹³ The initial distribution mentioned above is the initial distribution of descendants productivity constructed from the relevant matrices in Section 4 of Kaymak and Poschke (2016).
- ¹⁴ The retirement period is also an important force in the Kaymak-Poschke model for wealth accumulation for those with high productivity realizations.
- ¹⁵ A person who is sampled in a particular reference year is also selected in all other available years.

REFERENCES

- Alvaredo, Facundo; Atkinson, Anthony B.; Piketty, Thomas and Saez, Emmanuel. "The Top 1 Percent in International and Historical Perspective." *Journal of Economic Perspectives*, Summer 2013, 27(3), pp. 3-20.
- Alvaredo, Facundo; Atkinson, Anthony B.; Piketty, Thomas and Saez, Emmanuel. The World Wealth and Income Database; <http://topincomes.g-mond.parisschoolofeconomics.eu/>.
- Badel, Alejandro and Huggett, Mark. "The Sufficient Statistic Approach: Predicting the Top of the Laffer Curve." *Journal of Monetary Economics*, May 2017, 87, pp. 1-12; <https://doi.org/10.1016/j.jmoneco.2017.02.001>.
- Badel, Alejandro; Huggett, Mark and Luo, Wenlan. "Taxing Top Earners: A Human Capital Perspective." Manuscript, 2018.

- Blundell, Richard and Etheridge, Ben. "Consumption, Income and Earnings Inequality in Britain." *Review of Economic Dynamics*, January 2010, 13(1), pp. 76-102; <https://doi.org/10.1016/j.red.2009.10.004>.
- Brzozowski, Matthew; Gervais, Martin; Klein, Paul and Suzuki, Michio. "Consumption, Income and Wealth Inequality in Canada." *Review of Economic Dynamics*, January 2010, 13(1), pp. 52-75; <https://doi.org/10.1016/j.red.2009.10.006>.
- Creedy, John and Hart, Peter E. "Age and the Distribution of Earnings." *Economic Journal*, June 1979, 89(354), pp. 280-93; <https://doi.org/10.2307/2231602>.
- Deaton, Angus and Paxson, Christina. "Intertemporal Choice and Inequality." *Journal of Political Economy*, June 1994, 102(3), pp. 437-67; <https://doi.org/10.1086/261941>.
- Domeij, David and Floden, Martin. "Inequality Trends in Sweden 1978-2004." *Review of Economic Dynamics*, January 2010, 13(1), pp. 179-208; <https://doi.org/10.1016/j.red.2009.10.005>.
- Guvenen, Fatih; Kaplan, Greg and Song, Jae. "How Risky Are Recessions for Top Earners?" *American Economic Review Papers and Proceedings*, May 2014, 104(5), pp. 1-6; <https://doi.org/10.1257/aer.104.5.148>.
- Guvenen, Fatih; Karahan, Fatih; Ozkan, Serdar and Song, Jae. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?" Federal Reserve Bank of New York *Staff Reports*, February 2015, Staff Report No. 710.
- Guvenen, Fatih; Ozkan, Serdar and Song, Jae. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, June 2014, 122(3), pp. 621-60; <https://doi.org/10.1086/675535>.
- Heathcote, Jonathan; Storesletten, Kjetil and Violante, Giovanni. "Two Views of Inequality over the Life-Cycle." *Journal of the European Economic Association*, April 2005, 3(2), pp. 765-75; <https://doi.org/10.1162/1542476054473306>.
- Heathcote, Jonathan; Storesletten, Kjetil and Violante, Giovanni. "The Macroeconomic Implications of Rising Wage Inequality in the United States." *Journal of Political Economy*, August 2010, 118(4), pp. 681-722; <https://doi.org/10.1086/656632>.
- Huggett, Mark; Ventura, Gustavo and Yaron, Amir. "Sources of Lifetime Inequality." *American Economic Review*, December 2011, 101(7), pp. 2923-54; <https://doi.org/10.1257/aer.101.7.2923>.
- Kaymak, Baris and Poschke, Markus. "The Evolution of Wealth Inequality Over Half a Century: The Role of Taxes, Transfers and Technology." *Journal of Monetary Economics*, February 2016, 77, pp. 1-25; <https://doi.org/10.1016/j.jmoneco.2015.10.004>.
- Lagakos, David; Moll, Benjamin; Porzio, Tommaso; Qian, Nancy and Schoellman, Todd. "Life-Cycle Wage Growth across Countries." *Journal of Political Economy*, April 2018, 126(2), pp. 797-849; <https://doi.org/10.1086/696225>.
- Piketty, Thomas and Saez, Emmanuel. "Income Inequality in the United States, 1913-1998." *Quarterly Journal of Economics*, February 2003, 118(1), pp. 1-39; <https://doi.org/10.1162/00335530360535135>.
- Piketty, Thomas and Saez, Emmanuel. "Optimal Labor Income Taxation," in A. Auerbach, R. Chetty, M. Feldstein, eds., *Handbook of Public Economics*. Volume 5. Amsterdam: Elsevier, 2013, pp. 391-474; <https://doi.org/10.1016/B978-0-444-53759-1.00007-8>.
- Roine, Jesper; Vlachos, Jonas and Waldenstrom, Daniel. "The Long-Run Determinants of Inequality: What Can We Learn From Top Income Data?" *Journal of Public Economics*, August 2009, 93(7), pp. 974-88; <https://doi.org/10.1016/j.jpubeco.2009.04.003>.
- Saez, Emmanuel and Veall, Michael. "The Evolution of High Incomes in Northern America: Lessons from Canadian Evidence." *American Economic Review*, June 2005, 95(3), pp. 831-49; <https://doi.org/10.1257/0002828054201404>.
- Storesletten, Kjetil; Telmer, Christopher and Yaron, Amir. "Consumption and Risk Sharing Over the Life Cycle." *Journal of Monetary Economics*, April 2004, 51(3), pp. 609-33; <https://doi.org/10.1016/j.jmoneco.2003.06.005>.

