Assessing Individual Income Mobility Over Time

Rex Bodoia

2019 Senior Thesis

Department of Economics

Pomona College

Introduction

Countless research papers have been published over the past half century assessing income mobility in the United States. This is a popular topic rife with political implications that may shape policy debates, specifically ones concerned with taxation. The political Left and Right often quibble over the details, but mainstream politics on both sides tends to agree on the same fundamental tradeoff between competitive meritocracy and egalitarian redistribution. Due to the immense scope of this tradeoff, there is little consonance between the two sides when it comes to finding a specific line to draw.

Conservatives tend to support a smaller disparity between tax rates, specifically tax *breaks* for higher earners because it is argued these individuals earned their incomes fairly and taxing them at a higher rate disincentivizes their work and investment. Liberals tend to support a more progressive tax plan, specifically tax *increases* for the higher earners since it is presumed these individuals can afford to pay more in order to aid the lower earners through redistribution and other government spending. The underlying implications to these arguments are that the conservative side may be more justified if income mobility is high and every individual has a shot at success, while the liberal side is more justified given that income mobility is low and the rich get richer while the poor are condemned to perpetual poverty. These opposing attitudes toward governmental egalitarian redistribution at the expense of competitive meritocracy could therefore hopefully be conciliated in light of new information regarding income mobility.

It is worth noting that the above progression of policy arguments assumes that mobility is found to be *too high* or *too low*, but this requires some expectation of how income mobility should be. The previous context is laid out to emphasize the relevance of the issue, not to make a definitive judgement. The goal of this paper is not to make a statement about how income mobility should look because people would disagree over optimal levels of income mobility even if they had perfect information about the economy and policy outcomes. Instead of addressing any *proper* level of mobility, this paper simply aims to present the current state of the movement, structure, and variability of Americans' incomes over time. In doing so, several metrics are used in an attempt to quantify individual income mobility in a meaningful way. It is up to the reader to decide his or her own level of satisfaction with the results based on the dataset and methodology. The arguments above provide both a framework for how the policy debate might look once this decision is made and an explanation for the relevance of this topic.

Literature Review

Few recent papers have assessed individual income mobility. Instead the bulk of research has been focused on intergenerational mobility, specifically how a father's income correlates with his son's income. Until the 1990s, much of this research found reasonable levels of mobility, but has been criticized for using simplistic OLS regressions. Solon (1992) marked somewhat of a turning point after which previous methodology was repudiated and new methodology developed. These more recent studies such as Mitra and Ok (1998), Lee and Solon (2009), and Mutazashvili (2012) have found more alarmingly low levels of mobility. This intergenerational mobility research is interesting and may provide a more complete picture of the current state of income mobility in the US, but it is more directly pertinent to issues of education as well as death and estate taxes than it is to the issues discussed above, namely income taxation structure.

The purpose of this paper is instead to present a means of assessing how static or dynamic an *individual's* relative level of income is throughout his or her lifetime in the United States. In other words, is the top decile of earners a constant group of privileged individuals, or are the income strata more mercurial in nature? Several papers have evaluated topics along the same lines.

Veum (1992), Smith (1994), and Tsui (2009) found decreasing individual mobility over time, and, although these papers are relatively dated, they are worth exploring.

Smith in particular found *downward* mobility in income brackets among a significant portion of the population over time, arguing that more people are falling in terms of relative income than in the past. This issue of downward mobility however is nuanced, and it is important to identify exactly what is being discussed, specifically what constitutes mobility in each direction. Take for example a simplistic economy composed of ten people. Relatively speaking, they cannot all fall in income placement with no one to displace the downward movement. They could all earn lower wages in the subsequent year, but relatively speaking there would be zero mobility. Instead nine of them could descend in rank, which means one person must move from last to first. Thus, the economy can exhibit downward mobility in the sense that more people move down than up, but if we rank everyone in terms of income, the average rank shift is zero. In other words, the people who move up do so with an equal and opposite magnitude to the downward movements in terms of net effect.

It should also be noted that downward mobility is not inherently bad. For instance, is it better if nine people have constant income while the lowest earner shoots from last to first, or nine people retain constant income while the highest earner falls to last? In one case, there is much more *downward* mobility in the traditional sense of people falling relative to average income (or in terms of rank), but there is one person who is vastly better off and several people who are equally as well off. In the other case, there is plenty of *upward* mobility, but one person is much worse off while nine are equally as well off. Although this example is extreme, it illustrates how measuring directional relative mobility in isolation can be misleading. For this reason, the direction of mobility is not a primary concern, but instead a more general metric of movement is pursued.

Furthermore, it is important to consider changes in *absolute* income over time in addition to *relative* mobility for context in order to avoid misleading takeaways such as the one illustrated in the above example. Other research has explored the topic of absolute changes in earnings and these findings may provide valuable context for the relative mobility findings of papers such as this one.

Veum's research is more relevant to this topic than Smith's. His paper uses a ranking strategy in which every individual in the sample is ranked in terms of his or her income. For example, the highest earner is assigned a rank of 1, the next highest 2 and so on. Veum acknowledges that this net change in ranking will always be zero since one individual gaining rank means another losing rank and thus it would be meaningless. He proposes instead to assess the standard deviation of changes in relative position, which will be nonzero as long as ranks are not exactly the same and would measure how much overall movement there is from year to year. Furthermore, he divides the standard deviation of rank changes by their maximum possible values for the given sample size to standardize this metric accounting for sample size. He ultimately finds that mobility is decreasing over time and concludes that mobility in the US is on a downward trend.

What Veum fails to account for is that he is using longitudinal panel data, meaning he is following the same sample of individuals over time. This is important because the age of individuals in his sample is necessarily increasing across time. It is likely that mobility does in fact decline with age as individuals become more entrenched in their career paths, and if this is the case, then it can be expected that mobility diminishes over time for the same group of people. This is because they are more financially mobile as young adults investing in education or small businesses, or simply searching for career paths in which they excel. As these individuals grow older, they complete their educations and find high paying jobs or realize the ultimate outcomes of their businesses, either earning them consistently high incomes or forcing them to choose other career paths. Either way they tend to become less financially mobile as they grow older and this decline in mobility cannot be extended to the entire US population or labor force. The overall US labor force does not have a systematically increasing average age in the way that a static group of individuals does. Instead the oldest retire and are replaced by young newcomers and so an extrapolation of Veum's results to Americans is tenuous.

This disconnect between a longitudinal set of panel data and the general US population is difficult to circumvent given that tracking the same people over time is necessary for assessing individual income mobility over time. Unfortunately, this is precisely why observations are connected with age and cannot be extrapolated to the general population. It may be impossible to separate age and time trends to discern what effect each has, but a statement can be made about an effect that is caused by one or both of these factors. Which one has a stronger effect can only be surmised through reason. For example, Veum's research implies downward mobility as a result of either increasing age, a current trend over time in the US economy, or both. There may be more reason to suspect that mobility decreases with age than that there is currently an aggregate downward trend in US income mobility over time. Unfortunately, the experimental design cannot pinpoint the magnitude of each effect due to data constraints, so any explanation is only speculation.

Data

This paper uses panel data from the National Longitudinal Survey of Youth (NLSY), a program of the US Bureau of Labor Statistics. The dataset uses survey data from a randomly selected, static group of 12,686 individuals from the US population who were born between 1957 and 1964. The survey began in 1979 when the individuals were between the ages of 14 and 22 years old. These individuals have reported their levels of income annually from 1979-1993 and biannually from 1994-2014 for a total of 26 interview years. This will allow for the tracking of individuals' incomes over time and a comparison to overall US income levels.

The initial sample is composed of three subsamples. The first is a cross-sectional sample of 6,111 respondents designed to represent the noninstitutionalized civilian segment of the US population. The second is a supplemental sample of 5,295 Hispanic/Latino, black, or economically disadvantaged non-black/non-Hispanic respondents. The third is a sample of 1,280 respondents designed to represent the population serving in one of the four branches of the US military. Together these three subsamples make up the entire NLSY79 cohort, but the first subsample is the only one of interest for this paper. Although there are many important avenues to explore regarding minority groups and the economically disadvantaged, the goal of this paper is to analyze trends in the general US population. With this in mind, the first subsample is the most relevant to use for this income mobility analysis.

Respondents are asked a variety of questions during the interviews about topics including income during the last calendar year, educational experiences, and family backgrounds. This paper uses some of this information in an attempt to determine how earnings change over time. The prediction is that individuals generally do not remain in the same income brackets for long periods of time, but instead tend to enter the workforce in lower brackets and move upward.
Furthermore, movement between brackets from year to year in both directions is predicted to be common.

There are several reasons for an individual's income to increase over time relative to overall levels. Even more contribute to the fact that the same income brackets are occupied by different individuals over time. First, individuals often gain valuable experience throughout their careers, which increases their market values and warrants higher wages as they age. Consequently, there is an upward trend in individual earnings over time. A second important cause of mobility is that individuals often invest in their own human capital through education. During the educational period, individuals tend to earn little or no income, however the more time they invest in education, the higher their delayed earnings tend to be afterward. For this reason, individuals with the highest earning potentials may spend years in lower income brackets before realizing the fruits of their labor and entering higher income brackets. The final and more nebulous reason for mobility is arbitrary luck. Firms and people alike have good years and bad years due to what this paper will refer to as pure luck. All these factors may lead to some level of income mobility and specifically in an upward direction for individuals throughout their lifetimes.

Methodology

For the purposes of the analysis, the original dataset requires filtering. Because respondents sometimes refuse to answer survey questions or do not know answers, the data are not complete for every year. Furthermore, 444 respondents pass away between 1979 and 2014 and retention rates for interviews generally tend to fall throughout the lifespan of the survey. In terms of the data, a respondent could be recorded with one of four possible values for a given year's income: their numeric estimate of income in US dollars, a "refusal" (denoted -1), "don't know" (-2), "valid-skip" (-4), or "non-interview" (-5). This complicates the analysis to some degree as issues arise in the absence of complete data, especially given that the omissions may be a result of systematic bias. These issues are addressed below.

The first round of analysis is performed on a subsample containing only the individuals who report their incomes in every year of the survey. Unfortunately, this reduces the original sample size from 6,111 to 1,090 (17.8%) individuals. It is important to note that there may be a self-selecting bias involved. If failure to respond is not random but is instead due to embarrassment about lower earnings, uncertainty from less well-defined compensation structures, or any other nonrandom reason(s), there may be a systematic bias involved. If embarrassment of income or unclear compensation tends to cause those who fail to respond in some years to have lower incomes, this would bias the results of the analysis upward. Table 1 in the appendix provides summary statistics for three periods of this subsample: the first year, a middle year, and the final year.

One potential solution to this subsampling issue is to look at each year individually. The number of individuals who respond in every year is a relatively small subset of the sample, but the number of individuals who respond in any particular year is a much larger proportion of the sample. The problem with using such a subsample is that it is no longer the same group of individuals over time, and in fact the number of observations will be different from year to year. Instead, this self-selecting bias can be reduced by taking a subsample of individuals who respond every n years for some value $1 \geq n \geq 36$, where n is ideally a factor of 36 (the number of years the survey data spans). For example, a subsample of only the years 1979, 1984, 1989, …,2014 could be used. The idea behind this is that if the probability of an individual responding in any given year is roughly independent of those of other years, then simply reducing the number of years in which a valid response is required essentially reduces the number of criteria to be met and tends to increase the sample size. Using every n^{th} year allows the years to be evenly spaced. Interpolations can then be made about the years in between. This can be done with different granularities of time, keeping in mind a tradeoff between the number of individuals that remain in the sample and the amount of information lost between years.

Using this strategy, two more rounds of analysis are done in addition to the one that uses every year, with one sample that uses every sixth year – starting with 1979 and 1984 as the first two years to include 1979 despite a difference of only five years – and another sample that uses only the years 1979, 1996, and 2014 to retain early, mid-, and late career information. The differences between these three samples may be indicative of the size and direction of suspected bias. As mentioned earlier, the first sample contains 1,090 observations while the second and third retained 1,539 (25.2%) and 1,805 (29.5%) respondents, respectively. To be clear, the individuals who are in the second sample but not the first, fail to reveal their income in certain years that are used in the first sample. Tables 2 and 3 provide summary statistics for these second and third subsamples, respectively.

One final and more valuable subsampling strategy is employed and interpreted as the most insightful. This strategy uses 7-year time intervals instead of individual years. Specifically, individuals who respond in at least one of the seven years in each of the five periods are kept. Any respondent who fails to provide a valid response in every year of at least one of the five 7-year periods is dropped. The valid incomes of the remaining respondents are then averaged over each period so that each individual has an average income value for each period. For example, if an individual responds in three of the seven years of the first period, the average of those three years is computed and interpreted as his or her income in that period while the invalid years are ignored. The advantage of this strategy is straightforward: greater retention of observations. This new subsample contains 3,702 of the original 6,111 (60.6%) respondents and this higher retention rate should reduce loss of information and the effect of self-selection bias. Table 4 provides summary statistics for this subsample.

It should be noted that because the survey converts from annual to biannual collection of data after 1994, the first two and a half periods are inexorably different from the second two and half periods. Specifically, individuals may be more likely to have at least one valid response within the first two 7-year time intervals compared to the last three given that they have more chances for a valid response. This may not produce a bias with respect to certain individuals since there is no reason to suspect that this effects certain individuals more than others in any systematic way, nonetheless it is an inconsistency within the data to keep in mind. Furthermore, the lifespan of the survey does not partition evenly into 7-year periods. The periods used are ‘79-’84 (6-year period), ‘85-‘91, ‘92-’98, ’00-’06, ’08-’14. Although the process of choosing time interval length is to some degree arbitrary, this specific partition is chosen for a couple of reasons. First, 7 years is an appropriate amount of time in which to segment an individual’s working lifetime. It is granular enough to provide insight into different periods of someone’s life, yet it is a broad enough window to retain most of the original sample. Second, although the 36-year survey divides evenly into 6-year periods instead of 7-year periods, the second half of the survey only provides data for even years. Therefore, each period must necessarily be 3 years, 5 years, 7 years, etc. (since using 2002 and 2004 corresponds to the 3-year period of ’02-’04 and so on). Of these choices, 5- and 7-year intervals are both only offset from consistent period lengths by one year and 7 was chosen over 5 because, as a longer period, it allows for the retention of more observations.

With these subsamples, further analysis is done in order to identify what proportions of respondents retain their income placement over time. In particular, the number of individuals who remain in the top and bottom deciles between years is determined. The years of interest are the first year (or period) of the sample which corresponds to early or pre-career income, a middle period corresponding to mid-career earnings, and the final period which corresponds to late career income. This makes up the bulk of the analysis. Two methods, one theoretical and one simulation-based, are proposed for providing reasonable baselines to which the sample values of retention between deciles can be compared. The theoretical approach involves a chi-squared test comparing the sample values of total people who begin in the top decile in the first period and end up in each of the other deciles in the final period. For example, of the n/10 people who begin in the top decile, how many are in each of the ten deciles in the final period? This sampling distribution is compared to a chi-squared distribution where the expected value of individuals in each of the ten deciles in the later period is n/100. This is because the number of people beginning in the top decile is n/10, and 1/10 of this group (n/100) is expected to end up in each decile given that income distribution across periods is completely random and independent. This expected value of n/100 in each of the ten deciles thus constitutes the null hypothesis. The same test is applied for bottom decile earners. This process is then repeated for top and bottom decile earners in the mid-career period. There is thus a comparison of movement from initial period to middle period, initial period to final period, and middle period to final period for both the top and bottom deciles.

The simulation approach is somewhat analogous to the theoretical approach, but it uses computer simulation with pseudo-random number generation. As described above, it is expected that of the n/10 individuals who are in the top decile in a given period, n/100 is the expected number of individuals to remain in the top decile in a second period if yearly income is completely random and independent. This implies no correlation between income in one period and income in another, so it provides a useful baseline for what the expectation of complete sample mobility might look like. Although the expected value of this distribution is known to be n/100, the standard deviation and shape of the distribution are considered to be unknown, so a Monte Carlo simulation proves useful.

The simulation follows several steps. First, a random number between 0 and 1 is generated n times where n is the corresponding sample size (1090, 1539, 1805, 3702). These numbers represent the incomes of the n individuals. This is repeated once more so that each row (individual) is given two independent, random numbers representing incomes in two different periods. Finally, of those whose first number is in the top decile, a subset is taken of those whose second number is also in the top decile. In other words, the intersection of the two top deciles is observed. This can be used as a baseline with which to compare the intersection of the top two deciles that is actually observed in the NLSY sample. Since only doing this once provides a statistically weak baseline, the entire process is repeated 10,000 times to provide a sample distribution of these decile intersection values. This way the actual size of decile retention observed in the dataset can be compared to this experimental sample distribution to see how likely it would be to obtain such a result given that income is completely independent between years. It is then assumed that this distribution will not just be representative of the intersection of the two top deciles but can be extended to represent the intersection of any two deciles. This simulation process is useful because it provides a no-correlation baseline to which the sample value can be compared to how see how far it strays from "complete mobility".

These theoretical and experimental processes provide a reasonable upper bound for an expectation of mobility. Since they assume completely random, independent income rankings between years, they model a meaningful manifestation of complete mobility. Although it may be tempting to think of complete mobility as every rank changing as much as possible to produce the highest possible variance of ranks among the sample, this type of mobility can more accurately be thought of as very nonrandom and simply in the opposite direction. Instead of income in one year having no bearing on income in another and serving as a useless predictor, this incorrect high mobility situation would represent income in one period serving as an extremely strong predictor of a very different income in another period. In other words, the correlation would be highly negative rather than very small. This is why a retention rate of zero individuals instead of n/100 between deciles in two periods is not a reasonable baseline for high mobility.

One final avenue of analysis includes an exploration of the effect of education on earnings. A new variable, maximum education, is given to each individual in the sample representing the highest grade of education completed. The first version of this analysis is purely visual. A preliminary set of scatter plots is presented in Figures 3-6. Figures 3 and 4 plot average income in the initial period on the x-axis against average income in the final period on the y-axis. Figures 5 and 6 alternatively present average income in the middle period against average income in the final period. While these figures paint an interesting picture about the changing levels of mobility over an individual's lifetime, incorporating education into the equation provides more insight into what is occurring. Subsets of the 7-year period sample are taken according to different levels of max education. Figures 7-12 present new scatterplots analogous to the first set but with these education-based subsets instead.

The final and more rigorous assessment of mobility with regard to education is an OLS regression using the 7-year period sample. The natural log of average final period income is regressed on that of average initial period income as well as max education.

$$(1)\ \text{FinalInc} = \beta_0 + \beta_1\text{InitInc} + \beta_2\text{MaxEduc} + \varepsilon$$

The hope is to roughly assess what role both early career income and number of years of education play with regard to eventual earnings. Because initial period income likely changes with different levels of education, one more step is taken to include an interaction term of both these variables.

$$(2)\ \text{FinalInc} = \beta_0 + \beta_1\text{InitInc} + \beta_2\text{MaxEduc} + \beta_3\text{InitInc}*\text{MaxEduc} + \varepsilon$$

If early career income is strongly correlated with final income, then this is evidence of low mobility. If they are weakly correlated, then there is likely high mobility. If they are negatively correlated, then there may be a more complicated effect occurring such as an investment in education causing initial income to be low and final income to be high. Max education is predicted to be positively correlated with final period income.

Results

First, the average income of the sample for each year is compared to the average wage index in the US for the same years in Figure 1. Clearly the wages of the respondents do in fact increase consistently and steadily over time, as they age. They even surpass US averages in the early 90s when the participants are in their mid 20s and early 30s. This is important to note because it supports the claim that wages are rising over time in absolute terms and that they are strongly positively correlated with age.

The next step is the comparisons of sample decile intersectionality with a reasonable upper bound. The values of decile retention between time periods for both the top and bottom deciles in each of the four subsamples are shown in Table 5. Table 6 presents these values as a proportion of the overall sample size (value/n). The comparison to an upper bound can then be made using a series of chi-squared tests and Monte Carlo simulations. Table 8 shows the resulting test statistics of the chi-squared tests. Given 9 degrees of freedom, the chi-squared distribution gives a threshold of 21.67 at the 1% significance level, and all of the chi-squared tests yield much larger values. Thus, the null hypothesis that the distribution of individuals across the ten deciles is uniform with each containing 1/10 of the original top/bottom decile earners can be rejected at the 1% level.

Table 9 shows the means and standard deviations of the Monte Carlo simulation distributions for each subsample size, n. As predicted, the means are all very close to n/100. These mean values can then be used as a baseline with which to compare the actual values found in the sample from Table 5. Table 9 is the most important illustration of the simulation results. Each value is the proportion of the 10,000 iterations of the simulation that resulted in an intersection value at least as high as the one found in the actual sample for the corresponding decile, year, and subsample. This can be interpreted as a type of p-value, which indicates the probability that such a high retention rate would occur given completely random and independent yearly incomes. The histograms in Figure 2 provide visual representations of the simulation distributions. They look roughly normal as expected.

It is important to note that with the most valuable subsample, which uses 7-year periods, the highest p-value is 0.0001 and the others are all 0. This means that only 1 of the 10,000 iterations of the simulation produced such a large retention rate between deciles. It makes sense that the p-values decrease as the sample size increases because the distribution becomes tighter and the standard deviation decreases relative to the sample size, but it is striking nonetheless that these values are so unlikely. This serves as strong evidence that yearly incomes are not random and independent. Obviously, such a conclusion comes as no surprise, but it is important to note a significant and definitive distance from the upper bound of mobility.

The bottom row in Table 9 is also worth noting. This row displays the means of the columns above for each pair of time periods. This way the aggregate trends across the different subsamples with regard to different time period comparisons can be observed. It is clear that the strongest predictive relationships are between the middle and final time periods, whereas the first and last time periods have the weakest relationships. Once again, these findings come as no surprise given that the first and last time periods have a greater degree of separation and decreasing mobility over time is expected. A more surprising trend however is that the top decile has weaker retention than the bottom decile when looking at the first and middle or first and last time periods. Alternatively, the bottom decile has more retention than the top when the first time period is not involved. This may imply that there is more opportunity for those beginning with lower earnings than there is for those who are earning little money mid-career and that this effect is stronger than the inverse effect with high earnings. This would substantiate the hypothesis that education causes high potential earners to begin with low earnings in the same pool as those with low potential earnings.

Figures 3-6 provide an illuminating visual representation of this effect. Figures 3 and 4 show a scatterplot of the 7-year period subsample with average initial income on the x-axis and average final income on the y-axis. At first glance it looks like most people earn lower incomes in both periods while some people earn much higher incomes without much of a common trend. Drawing a trend line on this graph would be difficult given there is no clear correlation, but upon closer inspection it becomes apparent that two trend lines might be appropriate instead. Figure 4 shows that there may be two trends that represent alternate trajectories. There is a vertical trend which may represent people who pursue higher education. These individuals do not boast high earnings in the initial period, in fact they are almost entirely below the $15,000 threshold in the first period. Almost all of the highest ultimate earners find themselves on the left-hand side of this graph within this trajectory. Conversely, the majority of the high initial earners find themselves on the lower half of the graph. This implies that having high initial earnings may actually be a negative predictor of relative future success, while having low initial earnings provides little insight into future earnings since this is where the overlap of the two trajectories occurs. This is where other factors such as education likely play an essential role. Ultimately, this implies that income mobility is fairly high in young adulthood but tends to follow a pattern, particularly one hinged on highest level of educational attainment.

Figures 5 and 6 compare average mid-career earnings to final earnings and paint a markedly different picture. The trend is no longer a diversion of two trajectories, but is instead a single, more well-defined upward trend. This implies a positive correlation between mid-life income and final income and thus lower mobility in one's later years, which comes as no surprise; education is no longer part of the equation, and individuals have tended to settle in their eventual career paths. This downward trend in mobility is consistent not only with Veum's findings, but with the simulation comparison analysis of this paper.

Figures 7-12 show similar scatter plots, but with subsets of the sample based on education levels. Figures 7-9 show that attaining less than a college degree produces a flatter scatter plot. This means that initial income has a large spread with the potential to be high, but late career income is unlikely to be high relative to the entire sample. Figures 10-12 on the other hand show that when a college degree or higher level of education is attained, a taller and narrower scatterplot is produced. This can be interpreted as little to no potential for high initial earnings, but a much higher potential for future earnings. This is in line with the earlier interpretation of Figure 4 and provides some evidence that education stymies initial income in order to magnify ultimate income.

Tables 10 and 11 provide the results to the two regressions. All coefficients are significant at the 1% level which may in part be due to a large sample size. Moderately positive coefficients are observed for both the initial income and max education variables when the interaction is not included. When the interaction term is included, its coefficient is small and negative while the other two coefficients are larger. This may imply that both initial income and education have significant effects on final income and that the effect of initial income decreases with more education.

Conclusion

Although earlier success can predict later success, mobility does appear to exist, and few people are consistently winning or losing in terms of relative income levels. There appears to be an encouraging amount of opportunity in young adulthood, however by the time someone is at a mid-career point, it seems that he or she has little hope of completely turning things around. This decline in mobility as people age is an interesting phenomenon and one that people should be aware of. The early years of adolescence and young adulthood are crucial determinants of an individual's ultimate success, and consequently it is important for people to explore different career paths early, as well as pursue higher levels of education if it aligns with their interests. It may also be beneficial to follow talents and passions during this stage of life if one has hopes of becoming successful.

In terms of policy debate, there are some encouraging aspects of mobility in the US according to this sample of data. The same individuals are not always ahead or behind, however, it is clearly far from a random distribution from year to year. In fact, the evidence overwhelmingly supports a lack of complete mobility, a result which may favor redistributive policies. Unfortunately, the takeaway may be inconclusive since although retention rates between deciles prove to be far from random, they are even farther from 100%. It could be reasoned that a complete retention rate, or lack of mobility, is much worse than complete mobility is optimal, but it is difficult to pinpoint exactly how satisfied to be with these findings. Ultimately, it is a subjective question and one which the reader should decide for him or herself after some deliberation. The debate will need to continue in order to identify exactly where to draw that line, but it is encouraging to observe that the current state of income mobility is significantly different from either extreme. Hopefully this research in conjunction with other research regarding absolute income dynamics and educational attainment can be beneficial for such debate.

Appendix

Tables

Variable	Obs.	Mean	Std. Dev.	Min	Max
’79 Income	1,090	3577.688	3593.091	0	27000
’96 Income	1,090	27339.45	25109.95	0	138775
’14 Income	1,090	53412.54	68540.77	0	370314

Table 1: Initial, middle, and final income summary statistics for subsample in which individuals have given valid responses in every year

Variable	Obs.	Mean	Std. Dev.	Min	Max
’79 Income	1,539	3513.413	3725.367	0	40000
’96 Income	1,539	26411.03	26746.73	0	138775
’14 Income	1,539	48921.48	66655.67	0	370314

Table 2: Initial, middle, and final income summary statistics for subsample in which individuals have given valid responses in every sixth year

Variable	Obs.	Mean	Std. Dev.	Min	Max
’79 Income	1,805	3501.058	3849.075	0	40000
’96 Income	1,805	26721.84	28092.23	0	138775
’14 Income	1,805	48605.70	67750.39	0	370314

Table 3: Initial, middle, and final income summary statistics for subsample in which individuals gave responses in each of those three years (’79, ’96, ’14)

Period	Obs.	Mean	Std. Dev.	Min	Max
’79-‘84	3,702	4875.25	4464.78	0	31493.75
’85-‘91	3,702	13288.27	10419.67	0	59659.86
’92-‘98	3,702	22227.87	19879.96	0	122336.5
’00-‘06	3,702	36526.66	39181.56	0	253983.0
’08-‘14	3,702	45511.69	56401.70	0	342156.0

Table 4: Summary statistics of average income over each 7-year period for subsample in which individuals have given valid responses in at least one year of every 7-year period

Number of Obs. (n)	Top in First and Last Period	Bottom in First and Last Period	Top in First and Middle	Bottom in First and Middle	Top in Middle and Last	Bottom in Middle and Last
1,090	12	17	14	20	62	22
1,539	17	26	28	31	82	26
1,805	21	28	32	40	95	42
3,702	58	91	87	113	224	131

Table 5: Numbers of individuals common to top and bottom deciles across different time periods for each subsample

Number of Obs. (n)	Top in First and Last Period	Bottom in First and Last Period	Top in First and Middle	Bottom in First and Middle	Top in Middle and Last	Bottom in Middle and Last
1,090	0.011	0.016	0.013	0.018	0.057	0.020
1,539	0.011	0.017	0.018	0.020	0.053	0.017
1,805	0.012	0.016	0.018	0.022	0.053	0.023
3,702	0.016	0.025	0.024	0.031	0.061	0.035

Table 6: Proportion of individuals in both deciles (expected value of random comparison baseline is 0.01)

Chi Squared Results	Initial vs Middle Period	Initial vs Final Period	Middle vs Final Period
Top Decile	226.243	118.027	1143.49
Bottom Decile	311.243	173.568	470.595

Table 7: Results of Chi-Squared tests for top and bottom deciles comparing initial and middle periods to the final period

Number of Obs. (n)	Mean	Standard Deviation
1,090	10.8791	2.94828
1,539	15.3422	3.46247
1,805	18.1470	3.84813
3,702	37.0056	5.50556

Table 8: Results of 10,000-iteration Monte Carlo simulations of number of common individuals between deciles for each sample size

Number of Obs. (n)	Top in First and Last Period	Bottom in First and Last Period	Top in First and Middle	Bottom in First and Middle	Top in Middle and Last	Bottom in Middle and Last
1,090	0.4063	0.0327	0.1850	0.0035	0.0000	0.0008
1,539	0.3578	0.0035	0.0007	0.0001	0.0000	0.0035
1,805	0.2652	0.0107	0.0009	0.0000	0.0000	0.0000
3,702	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
Mean:	0.2574	0.0117	0.04665	0.0009	0.0000	0.0011

Table 9: Proportion of the 10,000 simulations for each n that resulted in a number at least as high as the count in the sample

VARIABLES	(1) LogAvgFinalInc
LogAvgInitInc	0.385***
	(0.0282)
MaxEduc	0.348***
	(0.0203)
Constant	1.202***
	(0.344)
Observations	3,702
R-squared	0.126

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 10: Log avg final period inc regressed on log avg initial period inc and max educ

VARIABLES	(1) LogAvgFinalInc
LogAvgInitInc	1.056***
	(0.133)
MaxEduc	0.740***
	(0.0785)
Init∗Educ	-0.0522***
	(0.0101)
Constant	-3.807***
	(1.027)
Observations	3,702
R-squared	0.133

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 11: Now including interaction term between initial income and max education

Figures

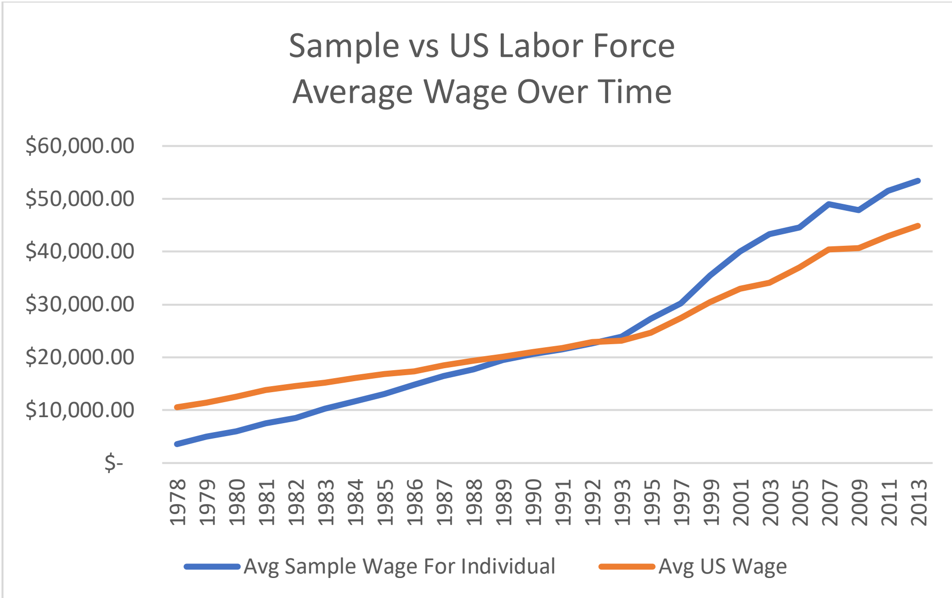

Figure 1: In blue is the average wage of the sample using every valid response from each year and in orange is the US average wage index

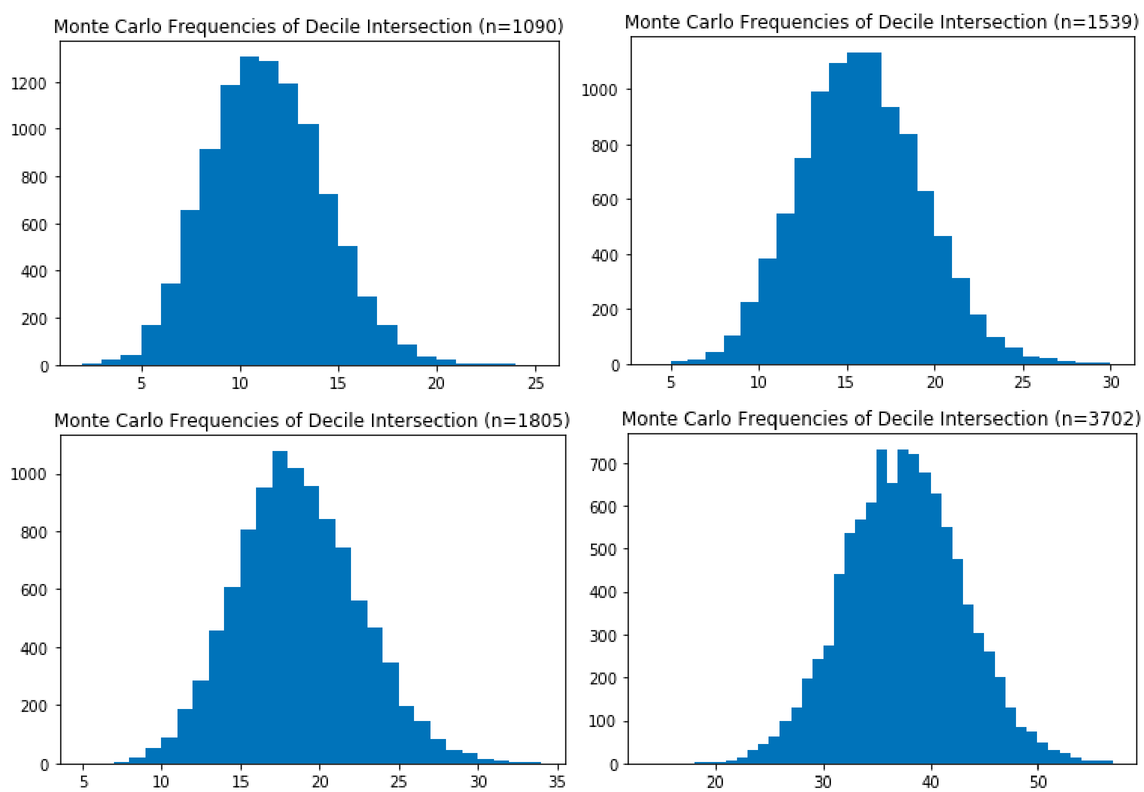

Figure 2: Each histogram shows the frequencies of each number of individuals that persists in the same decile over two periods for each of the four subsample sizes

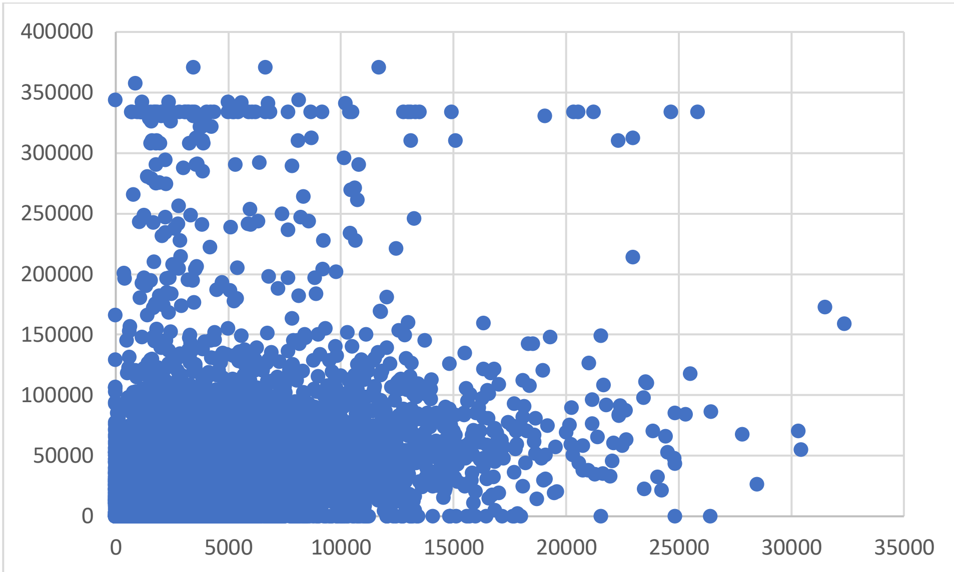

Figure 3: Scatterplot with Average Initial ('79-'84) Income on the x-axis and Average Final ('08-'14) Income on the y-axis

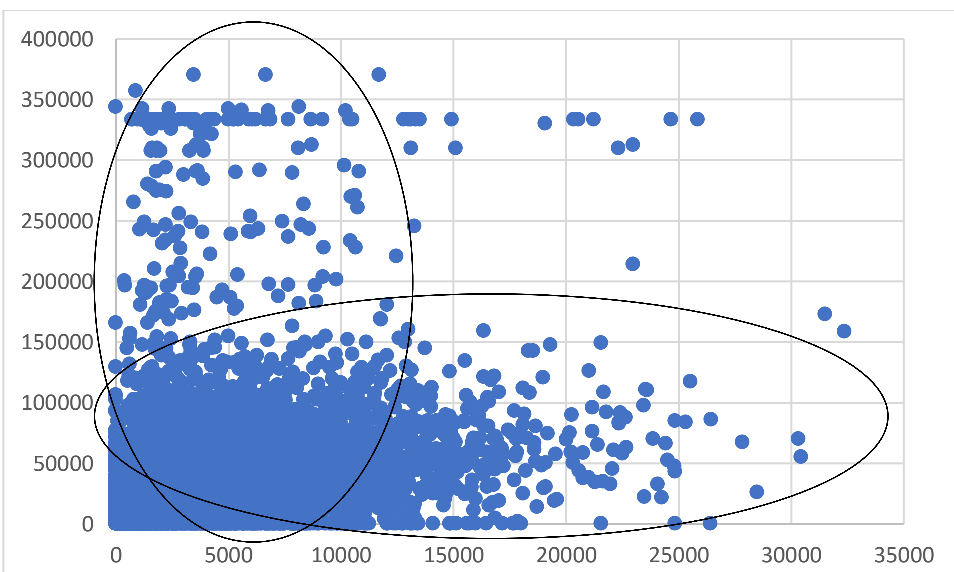

Figure 4: Same underlying scatterplot as in Figure 3; there are two opposing trajectories from starting point in first period

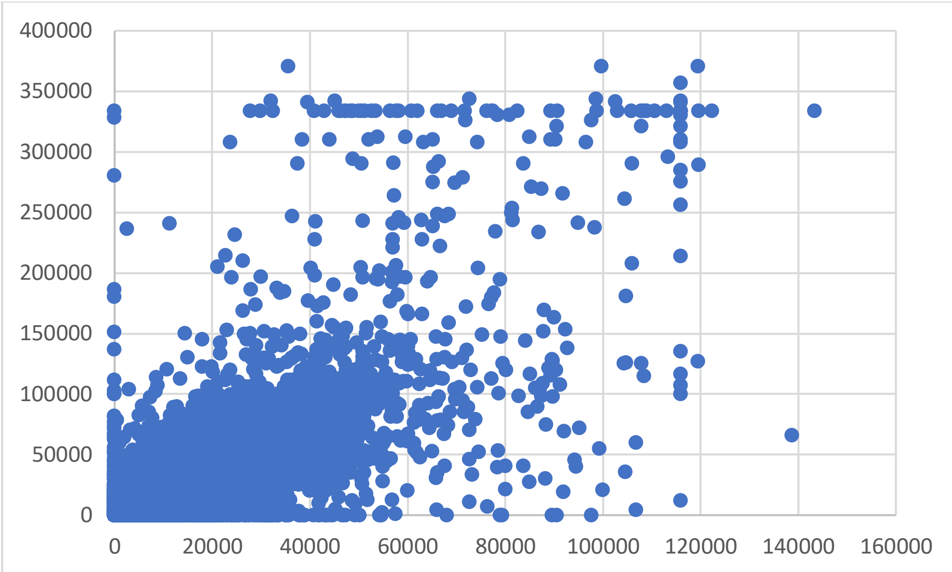

Figure 5: Scatterplot with Average Mid-Career ('92-'98) Income on the x-axis and Average Late Career ('08-'14) Income on the y-axis

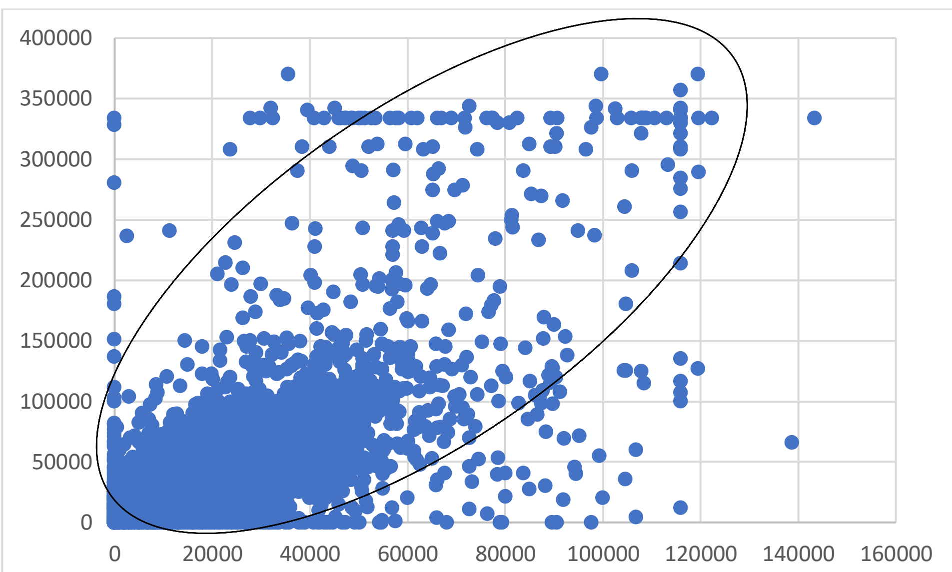

Figure 6: Same underlying scatterplot as in Figure 5; one somewhat well-defined correlation between mid and late career earnings

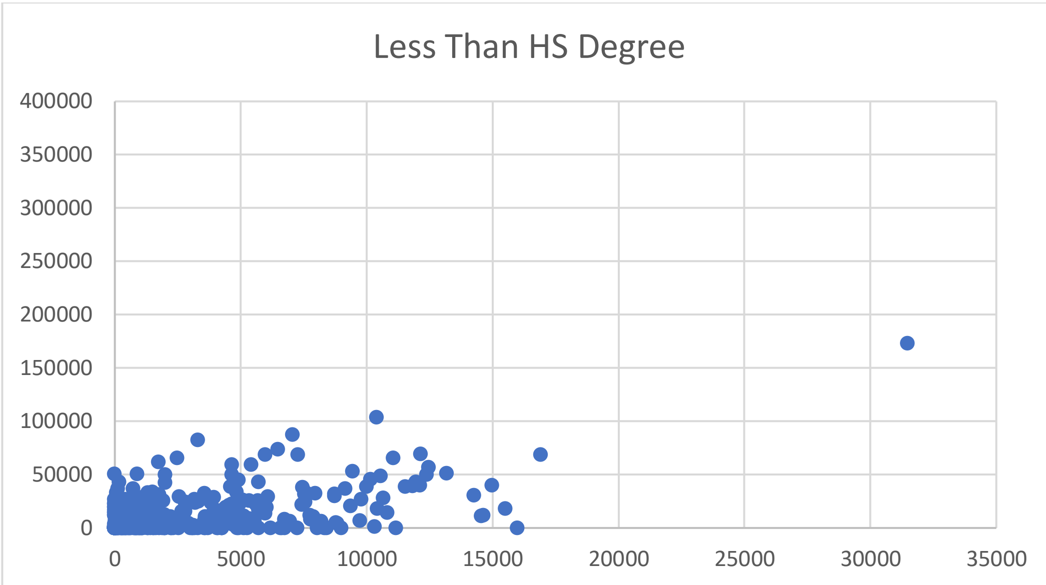

Figure 7: Average initial income vs average final income just as in Figure 3, but only including individuals who did not graduate high school are included

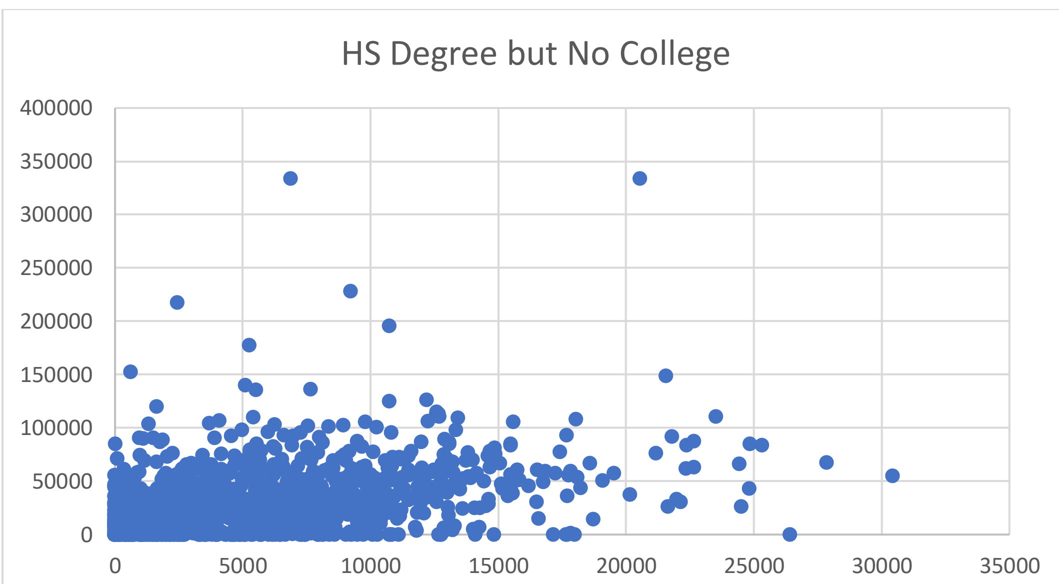

Figure 8: The subset of respondents who graduated high school but did not complete any years of college

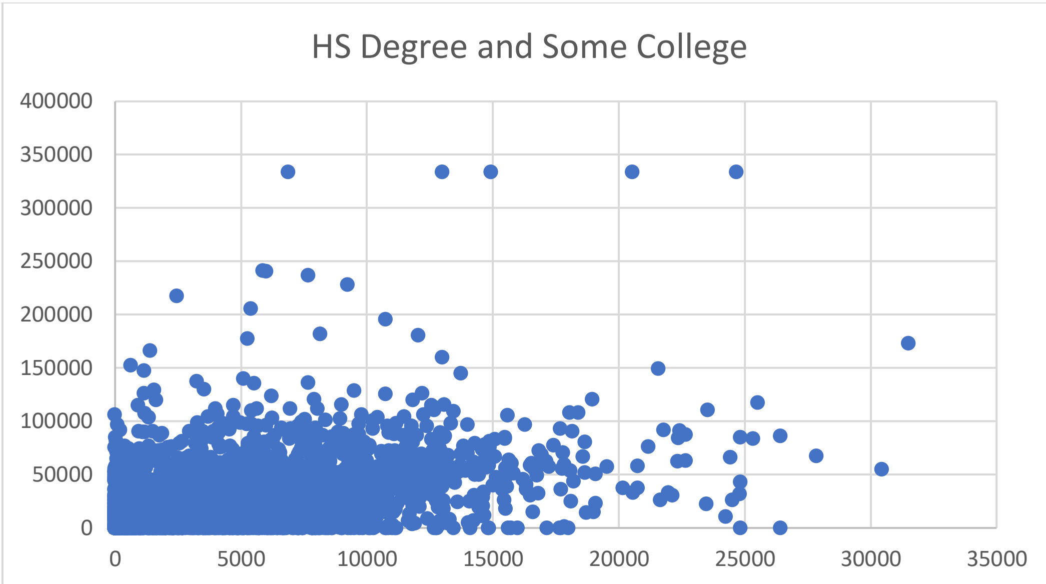

Figure 9: The subset of respondents who completed some years of college but did not graduate

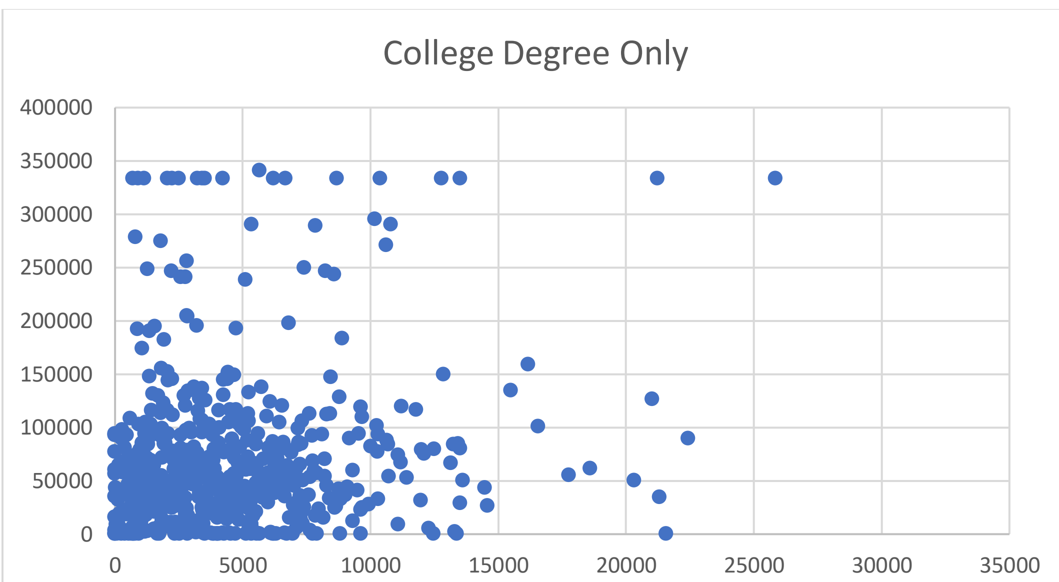

Figure 10: The subset of respondents with a college degree but no years of graduate school completed

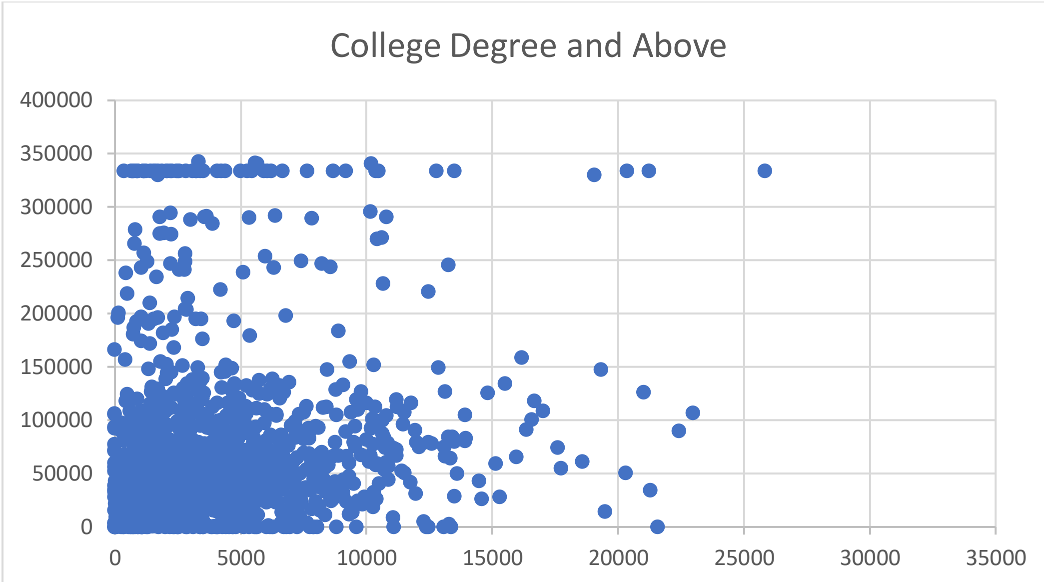

Figure 11: The subset of respondents with at least a bachelor's degree

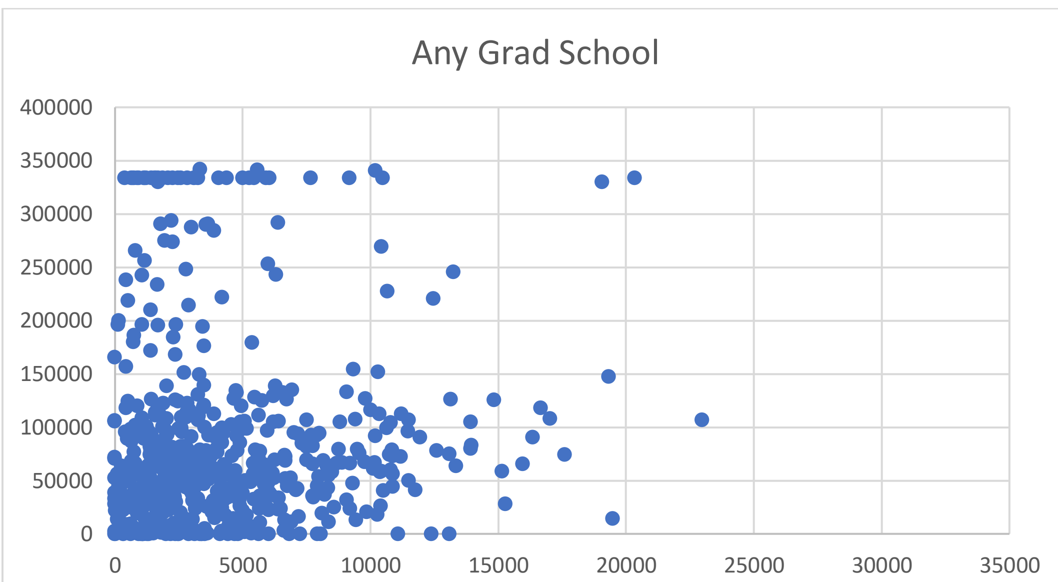

Figure 12: The subset of respondents who completed at least one year of graduate school

References

Mitra, Tapan, Ok, Efe. (1998). "The Measurement of Income Mobility: A Partial Ordering Approach." *Economic Theory* 12 (1): 77-102. https://www.jstor.org/stable/25055110.

Lee, Chun-Li, Solon, Gary. (2009). "Trends in Intergenerational Income Mobility." *The Review of Economics and Statistics* 91 (4): 766-772. https://www.jstor.org/stable/25651375.

Murtazashvili, Irina. (2012). "An Alternative Measure of Intergenerational Income Mobility Based on a Random Coefficient Model." *Journal of Applied Econometrics* 27 (7): 1161-1173. https://www.jstor.org/stable/23355929.

Smith, Patricia K. (1994). "Recent Patterns in Downward Income Mobility: Sinking Boats in a Rising Tide." *Social Indicators Research* 31 (3): 277-303. https://www.jstor.org/stable/27522748.

Solon, Gary. (1992). "Intergenerational Income Mobility in the United States." *The American Economic Review* 82 (3): 393-408. https://www.jstor.org/stable/2117312.

Tsui, Kai-yuen. (2009). "Measurement of Income Mobility: A Re-Examination." *Social Choice and Welfare* 33 (4): 629-645. https://www.jstor.org/stable/41108030.

Veum, Jonathan R. (1992). "Accounting for Income Mobility Changes in the United States." *Social Science Quarterly* 73 (4): 773-785. https://www.jstor.org/stable/42863119.