

Should I Stay or Should I Go? Inter-state Mobility and Earnings Gains of Young College Graduates

Andrew Glover and José Mustre-del-Río

September 2024

RWP 24-08

<http://doi.org/10.18651/RWP2024-08>

FEDERAL RESERVE BANK *of* KANSAS CITY



Should I Stay or Should I Go? Inter-state Mobility and Earnings Gains of Young College Graduates*

Andrew Glover[†] José Mustre-del-Río[‡]

September 2, 2024

Abstract

In the late 1990s, nearly 7 percent of young college graduates moved across state lines every year. These workers enjoyed 30 percent higher earnings three years after moving relative to similar stayers, but their gains were not immediate, amounting to only 7 percent in the first year post-move. By the mid-2010s, mobility fell by more than half, and average earnings gains among movers fell and became more front-loaded. At the same time, debt increased among all young college graduates. We propose a model of geographic mobility with incomplete markets, where moving to a new state can deliver earnings gains that are either front- or back-loaded. Incomplete markets and high interest rates on debt reduce workers' acceptance of back-loaded opportunities, even if they have the same present-discounted increase in earnings as front-loaded opportunities. We find that lower potential gains account for most of the decline in mobility across periods, but that the lower initial wealth of young college graduates also reduced their mobility. The wealth effect on mobility is especially strong for poor individuals, so wealth changes generate an endogenous increase in income inequality later in the life cycle.

JEL classification: D60, E21, E44.

Keywords: incomplete markets, geographic mobility, wealth, debt forgiveness.

*The views expressed herein are those of the authors and should not be attributed to the Federal Reserve Bank of Kansas City (FRB-KC) or the Federal Reserve System. The authors thank detailed comments from Ayşegül Şahin. They also thank Daniele Coen-Pirani, Jordan Rappaport, and participants at the 2022 Labor Markets and Macroeconomic Outcomes Conference in Santa Barbara for helpful comments and suggestions. The authors thank the Center for Advancement of Data and Research in Economics (CADRE) at the FRB-KC for providing essential computational resources that contributed to the results reported in this paper.

[†]Federal Reserve Bank of Kansas City; e-mail: andrew.glover@kc.frb.org

[‡]Federal Reserve Bank of Kansas City; e-mail: jose.mustre-del-rio@kc.frb.org.

1 Introduction

In this paper we investigate how the relationship between inter-state mobility, labor market returns from moving, and initial wealth has evolved over time. We document that mobility is historically greatest for young, college educated workers, who moved across state lines at a 6.6 percent annual rate in the late 1990's. For these workers, a move was associated with significant earnings gains after three years, of over 30 percent relative to similar workers who didn't move. However, their earnings gains were accumulated gradually, amounting to less than 10 percent in the first year post move. That the earnings gains from moving were back-loaded is important because young, college educated workers often have significant debt and may need to finance a move by borrowing at high interest rates in order to pay upfront costs.¹

While young college-educated workers were the most mobile 30 years ago, they moved across state lines significantly less by the late 2010's. By 2016, mobility among young college graduates declined to 4.2 percent at the same time as their labor market prospects and financial conditions upon entering the labor market worsened. First, post-move earnings growth became much weaker and the moves that did occur were more front-loaded: a typical move in the mid 2010's generated only 15 percent higher earnings in the first year post-move, declining to just 5 percent two years later.² Second, and key for our analysis, we also observe a substantial increase in the the debt burden of wealth-poor young college graduates. For example, while the 10th percentile of net wealth-to-income among this group was -1.4 in the late 1990s, it decreased to -3.4 in the latter period, suggesting young college graduates entered the labor market with greater debt burdens (relative to income) in the mid 2010's.

Motivated by these observations, we build an incomplete markets model in which earnings are endogenous to a workers' decision to move to a new state for job opportunities. A move idiosyncratically delivers either front-loaded or back-loaded earnings gains, but in order to realize these gains a worker must pay an immediate financial cost.³ Uninsurable income risk and financial market frictions affect the payoff from moving differently across the wealth and income distribution. Low income or wealth workers who have to borrow in order to finance a move effectively discount future consumption gains more heavily, since interest rates on debt are significantly higher than those on savings. On the other hand, workers who enter with high wealth or income are able to finance a move without incurring debt and therefore more likely to move to opportunities.

We first use our model to account for the decline in inter-state mobility between the late 1990's and the mid 2010's. Four factors affect mobility: changes in the present value of earnings from moving, changes in the timing of earnings gains, changes in wealth upon entering the labor market, and changes in the disutility from moving. The decline in present value of earnings gains from moving is by far the biggest contributor to lower mobility. However, the observed decline in financial wealth for young workers over this period also

¹Consistently, we document that revolving credit spikes significantly around a move and remains higher for at least three years.

²Earnings growth is relative to similar workers who do not move. So a decline from 15 percent to 5 percent does not imply a decline in the level of earnings, but that similar stayers enjoy faster earnings growth.

³We also assume that movers incur a utility cost, as is common in the literature and necessary to match mobility rates for realistically calibrated financial costs.

has a significant negative effect on mobility.

The decline in mobility due to lower initial wealth leads to persistently lower income for workers who enter the labor force poor. Our model predicts that workers who entered the mid-2010's labor force in the bottom wealth quartile had nearly 3 percent lower cumulative income through age 38 than wealth-poor workers who entered in the mid-1990's. On the other hand, workers in the top quartile of initial wealth saw an increase that helped them move earlier and thereby accumulate 1.5 percent higher cumulative earnings. In short, the relative decline in wealth for poor young workers caused a 4.5 percent endogenous increase in earnings inequality.

These findings suggest that recently proposed student loan forgiveness policies may increase mobility, raise earnings, and reduce income inequality.⁴ To test this claim, we use the model to evaluate such policies. We find that they increase mobility and early-career earnings growth for low-wealth workers. This generates a persistent increase in income over the life cycle, which has positive effects on welfare. Furthermore, the tax burden of financing debt forgiveness is relatively small, so the negative effect on other workers (including those with high initial wealth) is also small. Therefore, debt forgiveness raises welfare, on average.

The most directly related papers to ours are those that study the economic motives for interstate mobility. Classics in the literature include Kennan and Walker (2011) and Borjas et al. (1992). More recently, Kaplan and Schulhofer-Wohl (2017) and Karahan and Rhee (2019) study the secular decline in interstate mobility. Kaplan and Schulhofer-Wohl (2017) explain the decline in mobility through the decline in the variance of earnings across states. We complement Kaplan and Schulhofer-Wohl (2017) by noting how the gains to interstate mobility were most important among the young in the late 1990s, when overall mobility was higher, and declined most significantly in the late 2010s, when mobility was lower. Additionally, we highlight how the gains (when operative) accrue over time after moving, particularly among the young. Karahan and Rhee (2019) uses the decline in household wealth following the housing bust to help explain the decline in mobility during the Great Recession. We highlight the importance of credit constraints for mobility.

Also related is the work of Coen-Pirani (2010) and Coen-Pirani (2021). Coen-Pirani (2010) focuses on gross migration flows across U.S. states and documents that states that experience relative losses of population do so mainly because of low inflows rather than because of especially high outflows. He also documents that workers moving into a state tend to be observationally similar to workers migrating out of it. This finding supports our modeling assumption that idiosyncratic matching effects are important for understanding mobility. Coen-Pirani (2021) builds a rich quantitative internal migration model and finds that a higher degree of income tax progressivity tends to reduce migration rates. Our identification strategy, which relies on the frequency of interstate migration, is analogous to his.

Our emphasis on the intertemporal trade-off between staying and moving is related to Bilal and Rossi-Hansberg (2021). Bilal and Rossi-Hansberg (2021) argue that the location decision can be understood as an asset investment decision. Consistently, in our model the mobility decision is effectively an investment in

⁴The Biden Administration has announced plans to cancel student debt under the *Higher Education Act*. This act would cancel up to \$20,000 in interest for all borrowers who have accrued or capitalized interest on their loans since entering repayment. For borrowers in an income-driven repayment (IDR) plan meeting certain income requirements it would also cancel all principal and interest above the principal and interest balance at the time their federal student loans entered repayment. See <https://studentaid.gov/manage-loans/forgiveness-cancellation/debt-relief-info>.

one's future earnings, which can be bolstered by moving.

Lastly, our finding that the labor earning benefit of moving is persists even a few years after moving is related to Roca and Puga (2017). Using rich administrative data for Spain, they find that workers in bigger cities (relative to smaller ones) obtain an immediate static premium and accumulate more valuable experience, the latter which persists even after leaving the location.

In the next section we describe in greater detail the empirical regularities documented at the onset. Sections 3 and 4 outline the model used for quantitative analysis and describe its parametrization. Section 5 presents the model's accounting of the decline in interstate mobility among young college graduates, while Section 6 presents the model's policy implications. Concluding remarks are discussed in Section 7.

2 Empirical Patterns

In this section we describe in greater detail how inter-state mobility, its labor market returns, and wealth among young college graduates have evolved since the late 1990's. First, we show that in the late 1990's, inter-state migration displayed a stark life-cycle pattern, with mobility being highest among the young. Second, we show that the returns of inter-state migration in terms of higher labor earnings after moving were persistent or back-loaded, particularly among the young. Then we show how these patterns changed in the mid 2010's, with mobility declining throughout the life cycle and the returns to mobility declining and becoming more front-loaded. Lastly, to motivate the importance of wealth and credit constraints in the migration decision we show that wealth among young college graduates declined between the two periods, and how debt changes around the time that somebody moves across states.

2.1 Mobility and Earnings

Our analysis of mobility and earnings draws on data from the Survey of Income and Program Participation (SIPP) from the U.S. Census Bureau. The SIPP is useful for our analysis as it is based on a nationally representative sample of individuals and records key demographics and labor market history measures including earnings. For the purposes of our analysis the SIPP is particularly useful as it tracks individuals over time even when they move to another state.

Our analysis sample is based on the 1996 and 2014 panels of the SIPP. We focus on these panels for a few reasons. First, to document how the labor earnings gains of inter-state mobility persist even a few years out requires data with a long panel dimension, which precludes the usage of the 2001 panel, which lasted three years. Second, since we want to document secular changes in mobility, our latter period of analysis should be far enough from the first period to detect any changes. Hence, the exclusion of the 2004 and 2008 panels. However, we note that our findings are robust to the incorporation of the other panels.

We aggregate observations to an annual frequency as this is the most natural measure of time given our interest in inter-state mobility and its consequences for earnings dynamics. We restrict our attention to college educated white males between the ages of 25 to 60, a sample for which our facts are most salient.

These restrictions lead to samples of 12,510 observations from roughly 3,000 individuals in the 1996 panel and 17,649 observations from nearly 4,400 individuals in the 2014 panel.

Starting with our first key observation, Figure 1 displays how inter-state mobility declines with age and has changed over time. Focusing on the blue bars, which are based on the 1996 panel, between the ages of 25 to 34, the annual rate of inter-state migration averages nearly 7 percent. However, by the ages of 45 to 59 the inter-state mobility rate averages a little less than 2 percent.

The orange bars, based on the 2014 panel, reveal that two decades later mobility rates were systematically lower across the life cycle. Among the youngest group mobility declined to 4.2 percent, while among the oldest mobility declined to 1 percent. Importantly, for both the youngest and oldest we can reject the null hypothesis of no change in mobility rates across time periods.




Figure 1: Inter-state mobility over the life-cycle and across time
Source: Survey of Income and Program Participation, U.S. Census Bureau.

In an attempt to understand what might be driving these life cycle and time patterns we next look at how labor income changes when individuals move to another state. Unlike previous work, we not only examine earnings changes upon moving, but also earnings changes 2 and 3 years after the migration decision. In the figures that follow, earnings changes are measured as the difference in (residual) earnings relative to pre-move earnings. Additionally, all earnings differences are calculated relative to earnings changes of corresponding stayers, which helps net out life cycle related earnings growth that would have occurred independent of the mobility decision. To fix ideas, suppose individuals who move between period t and $t + 1$ see (residual) earnings increase by 6 percent. Meanwhile, similar individuals who do not move see earnings increase by only 2 percent. In this case, our 1-year (relative) earnings change for inter-state migrants is 4 percent, which is the excess earnings growth we ascribe to inter-state migration. Next, suppose that the same set of individuals who moved between t and $t + 1$ see an earnings change of 10 percent between t and $t + 2$, while the same set of individuals who stayed between t and $t + 1$ see an earnings change of 4 percent between t and $t + 2$. Then, our 2-year (relative) earnings change for inter-state migrants is 6 percent as this is the excess earnings growth we ascribe from migrating, again between t and $t + 1$. The 3-year earnings change follows immediately.

The next series of figures demonstrate the second and (we believe) most novel finding of this paper—the benefits from inter-state migration were persistent or back-loaded in the late 1990s and later became more front-loaded in the mid 2010s. This observation is notable among young college graduates.

Figure 2 highlights the persistence of earnings gains upon migrating to another state among white male college graduates in the 1996 panel. The far left set of bars of this graph show the timing of earnings gains for young (ages 25 to 34) college graduates. The blue bar reveals that a year after moving earnings for this group are 7.5 percent higher than similar non-movers, though this difference is not statistically different from zero. However, if we look at these individuals two years after moving (the orange bar) their earnings are now about 18 percent higher, and this difference is highly significant. Lastly, the height of the green bar shows that three years after moving, earnings are roughly 30 percent higher, and still statistically different from zero. The middle set of bars of this graph shows that among college graduates between the ages of 35 and 44, earnings continue to increase two and three years out, though to a lesser degree than their younger counterparts. Importantly, only the three-year change is statistically different from zero for this age group. Similarly, while the far right bars show older college graduates (ages 45 to 54) experience earnings losses the year after moving, which reverse in subsequent years, though none of these earnings changes are statistically significant.




Figure 2: Earnings changes for inter-state movers relative to stayers, college graduates 1996 panel
 Source: Survey of Income and Program Participation, U.S. Census Bureau.

Figure 3 highlights how earnings gains upon migrating became smaller and front-loaded two decades later. The far left set of bars of this graph show that young college graduates in the 2014 panel experienced earnings gains of 15 percent in the first year after moving, followed by more moderate gains of 6 and 5 percent in the subsequent two years. Importantly, only the first-year gain is statistically different from zero. Comparing these bars to the corresponding ones in Figure 2 reveals how different the size and timing of gains became. Across the other two age groups we do not observe systematic or statistically significant gains upon moving.




Figure 3: Earnings changes for inter-state movers relative to stayers, college graduates 2014 panel
Source: Survey of Income and Program Participation, U.S. Census Bureau.

2.2 Changes in Initial Wealth

As noted in the introduction, aside from observing a change in the frequency and labor market returns to geographic mobility between the late 1990s and mid 2010s, we also observe a significant shift in wealth among young college graduates as they enter the labor market. Table 1 documents how the initial net wealth to labor income distribution for young college graduates shifted between these two periods. This table is based on data from the Survey of Consumer Finances (SCF) and focuses on the 1998 and 2016 surveys, which most closely align with our SIPP samples. We restrict our attention to individuals in the SCF ages 23 to 26 and with a college degree or more to get a sense of initial wealth when entering the labor market among young college graduates.

Table 1: Net Wealth to Labor Income Quantiles

Year	5	10	25	50	75	90	95
1998	-3.49	-1.36	-0.74	0.26	1.20	4.16	8.35
2016	-18.29	-3.36	-0.81	0.21	0.82	2.41	4.50

Source: Survey of Consumer Finances (SCF) 1998 and 2016.

Looking across the two rows of this table shows a meaningful decrease in net wealth for young college graduates across these two periods, particularly at the bottom end of the wealth distribution. Looking at the bottom 5th percentile, young college graduates in 1998 on net had a debt burden equal to roughly 3.5 times labor income. By 2016, individuals in this same percentile had a much larger net debt burden of slightly over 18 times their labor income. Looking at the 10th and 25th percentiles also shows an increase in net debt burdens across the two time periods, though of somewhat smaller magnitudes.

Given that debt forgiveness will be an important part of our analysis in subsequent sections, we also highlight how student loan debt, in particular, changed across these two periods. Using the same sample, Table 2 shows how student debt evolved across the wealth distribution. Focusing on the last column shows that among young college graduates in the bottom 25th percentile of the wealth distribution, student loan debt relative to labor income increased nearly two-fold from 1.7 to 3 between 1998 and 2016. Thus, a good portion of the deterioration in net wealth for this group over time is accounted by higher student loan debt. Indeed, the second to last column of this table shows that net wealth to labor income for those in the bottom 25th percentile on average also saw a two-fold decline from -1.4 to -2.7.

Table 2: Composition of Net Wealth

Percentile of Net Worth/Labor Income	Av. Labor* Income	Av. Net* Worth	Av. Student* Loan Debt	Net Worth/ Labor Income	Student Loan Debt/ Labor Income
1998					
<25th	35.86	-51.27	62.45	-1.43	1.74
25-75th	65.36	25.85	14.70	0.40	0.22
>75th	52.62	428.87	74.03	8.15	1.41
2016					
<25th	23.22	-62.82	68.84	-2.71	2.96
25-75th	57.11	13.86	13.86	0.24	0.24
>75th	40.75	60.76	23.71	1.49	0.58

Notes: * in thousands of real dollars.

Source: Survey of Consumer Finances (SCF) 1998 and 2016.

The middle row of this table highlights an important aspect of the increase in student loan debt across these two periods. Specifically, it shows that the driver of the greater student loan to income ratio wasn't higher student loan debt per se, but rather declining labor incomes. Focusing on the those in the bottom 25th percentile, the results from this middle column show that across the two time periods average student loan debt increased only slightly. Instead, average labor income fell dramatically leading to an increased debt burden.

2.3 Credit Dynamics and Mobility

We also document credit dynamics around the time of a move using the NYFed/Equifax Consumer Credit Panel. We split our data into two balanced panels covering the periods 2000-2006 and 2013-2019 in order to avoid the Great Recession. We then split each panel into three age groups, containing people aged 25-34, 35-44, and 45-54 in the initial year of each panel. Table 3 reports summary statistics for each of these groups in each sample period, further split by the number of inter-state moves completed by the borrower.

Looking at the first rows of each panel shows that the level of debt (revolving or total) is increasing with age regardless of the time period in consideration. Indeed, average revolving debt among the oldest group is roughly 1.8 times larger than average revolving debt for the youngest group. Similarly, total debt of the oldest group is roughly 1.6 times larger than that of the youngest group.

In contrast, we observe the opposite pattern when looking at growth in debt (revolving or total), with the young group experiencing the fastest growth and the oldest group the slowest. This pattern also holds true across time periods, but is starkest in the latter. Focusing on the 2013-2019 period, the average growth rate of revolving debt among young individuals was roughly 10.5 percent per year. In contrast, individuals in the 45-54 group saw their revolving debt contract by roughly 1 percent per year. Looking at total debt paints a similar picture: young individuals saw total debt increase by roughly 9 percent per year, whereas older individuals saw their debt decrease by roughly 1.4 percent per year.

Separating individuals into those who engage in interstate migration versus those who don't reveals key differences in the dynamics debt across age groups. Within the young group, we see higher growth rates of revolving and total debt for those who move than those who never move. For the middle group, we see that total debt grows faster for people who move, but that revolving debt does not. It may therefore be more important for a young person to have access to revolving debt when deciding whether to move, since she would not have other sources of liquidity at hand. We therefore look at the dynamics of revolving debt for each age group around the move date.

Table 3: Summary Statistics of Consumer Credit

		2000-2006					2013-2019				
		Mean	P25	P50	P75	Growth	Mean	P25	P50	P75	Growth
Revolving Debt	Age 25-34	7,513	394	2,574	8,889	10.02	4,024	0	1,034	4,292	10.54
	0 Moves	7,557	371	2,477	8,840	9.60	4,046	0	1,028	4,332	10.39
	1 Move	6,714	333	2,495	8,208	10.87	3,398	0	852	3,405	12.22
	2+ Move	7,982	650	3,298	9,832	11.83	4,748	114	1,432	5,473	9.91
	Age 35-44	10,866	616	3,492	12,398	7.42	9,159	113	2,077	8,818	2.83
	0 Moves	10,654	599	3,450	12,221	7.60	9,157	130	2,109	8,899	2.84
	1 Move	10,969	605	3,312	12,627	5.98	8,624	0	1,548	7,220	2.88
	2+ Moves	13,467	859	4,395	14,989	6.85	10,181	220	2,495	10,178	2.59
	Age 45-54	13,379	688	4,284	14,961	3.73	12,885	279	2,821	12,141	-1.18
	0 Moves	12,984	662	4,161	14,487	4.04	12,814	291	2,821	12,141	-1.22
	1 Move	12,844	717	3,885	15,173	2.94	13,853	40	2,453	11,271	-1.21
	2+ Moves	19,518	1,258	7,133	20,531	1.27	12,995	466	3,556	13,911	0.02
Total Debt	Age 25-34	41,315	2,310	14,609	57,738	16.37	53,309	3,570	18,890	68,970	8.73
	0 Moves	42,150	2,294	15,184	61,902	14.82	54,023	3,563	18,832	71,665	7.79
	1 Move	37,430	1,700	10,841	44,733	20.94	46,747	2,299	15,339	51,811	12.12
	2+ Moves	39,496	3,212	14,812	40,836	21.69	56,312	6,588	24,747	66,483	12.19
	Age 35-44	64,501	4,739	32,401	95,419	8.55	94,382	6,077	39,558	135,471	1.65
	0 Moves	63,928	4,937	33,069	94,308	8.16	95,183	6,267	41,810	137,146	1.24
	1 Move	64,949	3,079	24,807	102,110	10.38	88,227	4,246	26,838	118,933	4.40
	2+ Moves	71,335	4,968	31,272	104,438	10.94	89,698	6,871	29,407	125,867	4.62
	Age 45-54	65,717	5,496	33,690	91,396	5.31	94,410	5,840	40,299	128,078	-1.37
	0 Moves	63,496	5,445	32,967	88,489	5.30	93,090	5,862	40,882	126,982	-1.55
	1 Move	69,949	4,217	33,532	103,206	5.09	98,853	4,720	29,696	127,893	-0.04
	2+ Moves	92,102	9,103	48,225	127,894	5.55	125,540	8,111	49,808	163,872	0.47

Summary statistics by age group, time period, and number of interstate moves. Dollar amounts are in 2020 levels. Growth is annualized average percentage increase.

Source: FRBNY Consumer Credit Panel/Equifax.

We next use the Federal Reserve Bank of New York Consumer Credit/Equifax Panel to document how

different forms of consumer credit change around the time that a person moves to a new state. We look at revolving credit (such as credit cards) in the four quarters before and twelve quarters after a move occurs by estimating a distributed lag regression model. The inclusion of time, age, and individual fixed effects means that we are comparing debts for people who move at period 0 to people who do not move at that time. Each coefficient is the difference between these two people at that time, relative to the average difference between those two groups in all periods more than five quarters before the move occurs.

Figure 4 plots the dynamics of revolving debt for a person who moves across state lines, by age and period. To avoid the Great Recession, the top rows use data from 2000-2006 and the bottom rows use data from 2013-2019. The first column looks at a sample of people aged 25-34 in the initial year of the sample, and the next two columns look at successively older groups.

Revolving debt does not change much in the year prior to a move, but it rises immediately at the time of move and only gradually returns to a similar level as a year before the move. This general pattern is present in all age groups and periods, although we see that the young group's debt rises more significantly, both economically and statistically, and remains above the pre-move level for longer than the other groups. These dynamics suggest that access to unsecured credit is useful during a move, especially for younger people.

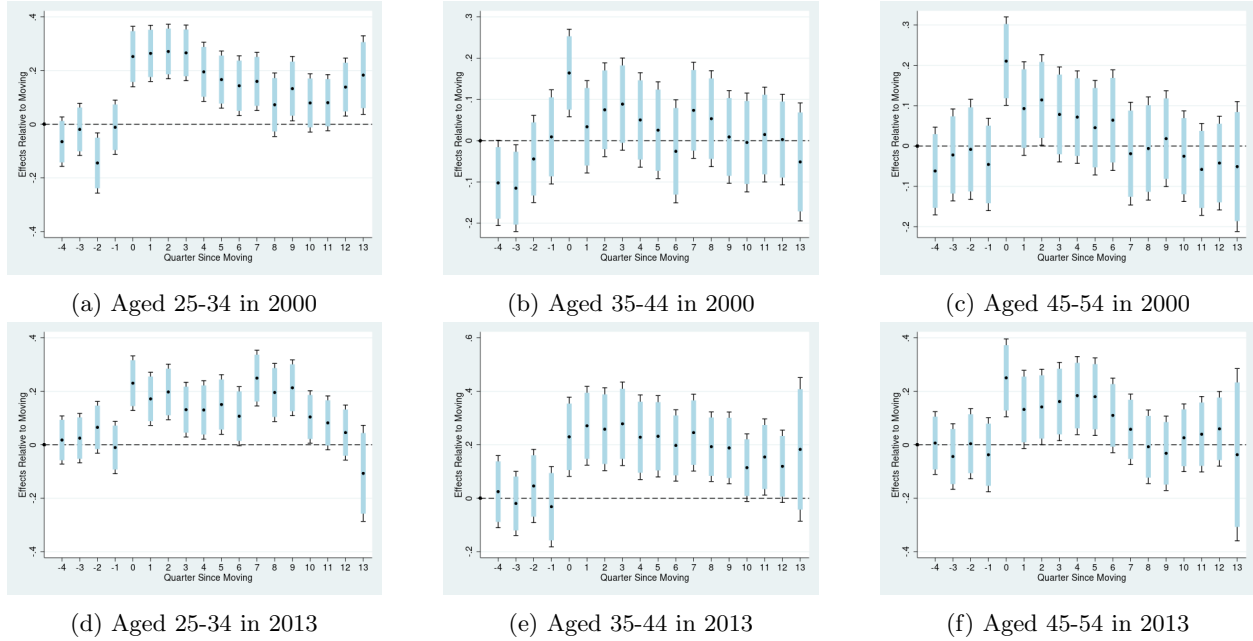


Figure 4: Dynamics of Revolving Credit Around Move

Notes: Coefficients on distributed lag regression of the form

$$y_{i,t} = \alpha_t + \gamma_i + \sum_{j=-4}^{12} \beta_j \text{MOVE}_{i,t-j} + \beta_{13} \text{MOVE}_{i,13+} + \varepsilon_{i,t},$$

in which the dependent variable $y_{i,t}$ is the inverse-hyperbolic sine of revolving debt and the independent variable $\text{MOVE}_{i,t}$ is an indicator variable that takes the value one if a person moves across state lines in quarter t and $\text{MOVE}_{i,13+}$ indicates that the person moved thirteen or more quarters ago. The coefficient for thirteen quarters corresponds to thirteen or more quarters since the move. Regression includes time, age, and individual fixed effects. Standard errors clustered at the state level.

Source: FRBNY Consumer Credit Panel/Equifax.

In summary, the results from this section highlight several important facts about inter-state migration over the life-cycle and across time. First, interstate migration declines sharply over the life-cycle. Second, the labor earnings benefits from moving to another state are persistent, with gains occurring even two to three years after moving. Third, the decline in interstate migration between the late 90s and mid 2010s was associated with a decline in the level and persistence of earnings gains upon moving. We also find that revolving debt increases significantly upon moving across states. This suggests that credit access (or lack thereof due to financial market incompleteness) is an important ingredient for understanding the migration decision. In the next section we link these observations in a life-cycle model of consumption, savings, and mobility under incomplete markets.

3 A Model of Consumption, Savings, and Mobility

We now develop a model of consumption, savings/borrowing, and inter-state mobility. Before going into the full quantitative model, we consider the choice of moving or staying with exogenous consumption processes to develop intuition for how wealth affects the effective costs and benefits of moving.

The costs and benefits of moving. Consider a worker who is offered a one-time choice between two consumption processes, $(c_j^s, c_j^m)_{j=0}^J$ with $c_0^m = c_0^s - f$ and $c_j^m = c_j^s + \Delta_j^c$ for $j \geq 1$. The worker discounts future consumption using factor β and has felicity $u(c)$ with $u'(c) > 0$ and $u''(c) < 0$. In this setting, the worker chooses to move if and only if

$$\mathbb{E}_0 \left[u(c_0^s - f) - \nu + \beta \sum_{j=0}^{J-1} \beta^j u(c_{j+1}^m) \right] > \mathbb{E}_0 \left[u(c_0^s) + \beta \sum_{j=0}^{J-1} \beta^j u(c_{j+1}^s) \right]. \quad (1)$$

Approximate each side of this expression around $c_j^m = c_j^s$. We can then say that the worker moves if and only if

$$\sum_{j=1}^J \mathbb{E}_0 \left[\frac{\beta^j u'(c_j^s)}{u'(c_0^s)} \right] \Delta_j^c > f. \quad (2)$$

This expression gives is intuition for how wealth and the timing of consumption can affect mobility, both in the cross-section and over time.

Suppose that the consumption processes arise from borrowing and savings decisions by households with different levels of initial wealth and that borrowing incurs a higher interest rate than saving earns ($r^b > r$). First consider a household who is always a net saver and receives a real interest rate of $r = 2\%$. If this worker moves, their effective discounted value of consumption gains at age j are $\mathbb{E}_0 \left[\frac{\beta^j u'(c_j^s)}{u'(c_0^s)} \Delta_j^c \right] \approx 0.98^j \Delta_j^c$. This worker would therefore put a high value on expected consumption gains from moving, even if they accrued many years in the future. Now consider a worker who is initially in debt and expects to borrow for the first five years of their working life at a real rate of $r^b = 15\%$. This worker would value future benefits more than five years out at $0.85^5 = 0.44$. This difference in discounting suggests that the high-wealth worker will move at a higher frequency and therefore enjoy higher realized earnings throughout their working life.

The intuition from this analysis will carry over into our full model, both in terms of low wealth workers being less likely to move to opportunities in the cross section and lower wealth reducing mobility for poor households over time. If we could observe the consumption processes directly, we could take these closed form expressions directly to the data. In the absence of such data, we use the structure of a standard consumption-savings model to evaluate the gains and costs of moving.

Environment. We study a life cycle incomplete markets model close to that of Kaplan and Violante (2010). As they do, we model the choices of an individual that lives up to T periods and works until age $W \leq T$. At any time a unit mass of people are born.

Agents survive to the next period with probability ϱ_j , which depends on age j . Preferences are given by the expected value of the discounted sum of momentary utility

$$\mathbb{E}_0 \left[\sum_{j=1}^T (\beta \varrho_n)^j u(c_j) \right],$$

where c_j is consumption at age j . The utility function u is strictly increasing, strictly concave, and twice differentiable.

We assume earnings follow a restricted income profile (RIP) process. Specifically, an agent i of age j receives income $y_{i,j}$. During working ages, income has a life cycle component that is common to all individuals, and a persistent component and i.i.d component that are idiosyncratic:

$$\log(y_{i,j}) = l_j + z_{i,j} + \varepsilon_{i,j},$$

where l_j denotes the life cycle component, $z_{i,j}$ is a persistent component, and $\varepsilon_{i,j}$ is a transitory component. We assume the persistent component $z_{i,j}$ follows an AR(1) process:

$$z_{i,j} = \rho z_{i,j-1} + e_{i,j}. \quad (3)$$

We assume $\varepsilon_{i,j}$, and $e_{i,j}$ are normally distributed with variances σ_ε^2 , and σ_e^2 , respectively. Once in retirement, the household receives a percentage of the last realization of the permanent component of its working-age income.

The effect of inter-state mobility on earnings is modeled in a parsimonious fashion via the persistent component of earnings. This effect lasts for three periods and can be either front-loaded (F) or back-loaded (B). We assume front-loaded offers arrive with probability θ while back-loaded offers arrive with complementary probability $1 - \theta$.

More precisely, the persistent component of an individual who migrates to another location is:

$$z_{i,j} = \alpha^{m,\ell} + \rho z_{i,j-1} + e_{i,j}. \quad (4)$$

where: $\alpha^{m,\ell}$ depends on the periods since the individual moved (m) and the type of move ($\ell \in \{F, B\}$). For a front-loaded F move, we assume:

$$\alpha^{m,F} = \begin{cases} 2\Delta_0 & \text{if } m = 1 \\ (2\Delta_0 - \Delta_1)(1+r) & \text{if } m = 2 \\ 2(\Delta_0 - \Delta_1)(1+r)^2 & \text{if } m = 3 \end{cases} \quad (5)$$

Meanwhile, for a back-loaded (B) move we assume:

$$\alpha^{m,B} = \begin{cases} 2(\Delta_0 - \Delta_1) & \text{if } m = 1 \\ (2\Delta_0 - \Delta_1)(1 + r) & \text{if } m = 2 \\ 2\Delta_0(1 + r)^2 & \text{if } m = 3 \end{cases} \quad (6)$$

Note that either type of move delivers the same present value as the rate on savings is used to discount future payments. Once a mover has reached $m = 3$, their earnings process reverts to (3) and they become “permanent stayers”; individuals who no longer receive migration offers. Lastly, as the focus of this paper is the mobility decisions of young individuals, we assume mobility offers arrive through age 34, after which no migration decisions are made.

While somewhat non-standard, this formulation is motivated by the empirical findings in Section 2. Data constraints limit how long we model the effect of mobility on earnings; i.e. why $m \leq 3$. Additionally, since repeat mobility in our sample is very low we impose that individuals can only move once, hence the “permanent stayer” assumption. Lastly, this specification captures in a flexible way the timing of earnings gains of movers both in the late 1990’s (when earnings gains appear more back-loaded) and the mid 2010’s (when earnings gains appear more front-loaded).

To close the model, as previously mentioned, markets are incomplete and thus agents are allowed to save using a single risk-free asset a , which pays a gross interest rate $1 + r = R$. We assume agents can borrow up to a limit \underline{a} at the gross borrowing rate $1 + r^b = R^b > R$.

Household Decision Problems

3.1 Stayer’s problem

An individual of age j who is a stayer and has yet to move solves:

$$\begin{aligned} V_j^s(a, z, \varepsilon) = & \max_{a' \geq \underline{a}} u(c) + \theta \beta \varrho_j \mathbb{E}[\max\{V_{j+1}^s(a', z', \varepsilon'), V_{j+1}^{F,1}(a' - \kappa(a'), \tilde{z}', \varepsilon') - \nu\} | z] \\ & + (1 - \theta) \beta \varrho_j \mathbb{E}[\max\{V_{j+1}^s(a', z', \varepsilon'), V_{j+1}^{B,1}(a' - \kappa(a'), \tilde{z}', \varepsilon') - \nu\} | z] \end{aligned} \quad (7)$$

$$\begin{aligned} \text{s.t.} \quad c &= \mathbb{I}_{a \geq 0} R a + \mathbb{I}_{a < 0} R^b a + \exp(l_j + z + \varepsilon) - a', \\ z' &= \rho z + e' \\ \tilde{z}' &= \alpha^{1,\ell} + \rho z + \tilde{e}' \text{ for } \ell \in \{F, B\} \end{aligned}$$

The continuation value takes into account that tomorrow this person may receive either an F or B type move. In either case, the individual can choose to accept the offer or reject it and remain a stayer. Additionally, regardless of move type, an individual must pay an upfront financial moving cost $\kappa(a')$, which

is a function of wealth.

We assume that when deciding to move the agent is subject to a utility cost μ_j^M and a utility shock ν , drawn from a logistic distribution. We normalize the variance of these shocks to be 1 and assume the distribution is the same regardless of move type. Note, both the mean utility cost and variance are age dependent - since we are focused on young movers, we will assume the utility cost is infinite for those over age 34.

3.2 Mover's problem

An individual of age j who is an ℓ type mover solves:

$$\begin{aligned}
V_j^{\ell,m}(a, z, \varepsilon) = \max_{a' \geq a} & \quad u(c) + \mathbb{I}_{m < 3} \beta \mathbb{E}[V_{j+1}^{\ell,m+1}(a', z', \varepsilon') | z] + \mathbb{I}_{m=3} \beta \mathbb{E}[V_{j+1}^{ps}(a', z', \varepsilon') | z] & (8) \\
\text{s.t.} & \quad c = \mathbb{I}_{a \geq 0} R a + \mathbb{I}_{a < 0} R^b a + \exp(l_j + z + \varepsilon) - a', \\
& \quad z' = \alpha^{m+1, \ell} + \rho z + e' \quad \text{for } \ell \in \{F, B\} \\
& \quad \hat{z}' = \rho z + \hat{e}'
\end{aligned}$$

The continuation value of this problem takes into account that individuals who have reached $m = 3$ will become permanent stayers the following period. As a result, their earnings evolve in the usual way. By contrast, if $m < 3$ these individuals continue to have their earnings influenced by their mobility status, which depends on move type (ℓ) and duration of mobility spell (m).

3.3 Permanent stayer problem

An individual of age j who is a permanent stayer solves the standard problem:

$$\begin{aligned}
V_j^{ps}(a, z, \varepsilon) = \max_{a' \geq a} & \quad u(c) + \beta \mathbb{E}[V_{j+1}^{ps}(a', z', \varepsilon') | z] & (9) \\
\text{s.t.} & \quad c = \mathbb{I}_{a \geq 0} R a + \mathbb{I}_{a < 0} R^b a + \exp(l_j + z + \varepsilon) - a', \\
& \quad z' = \rho z + e'
\end{aligned}$$

The continuation value of this problem takes into account that these individuals remain permanent stayers and their earnings evolve following the usual process.

4 Calibration

As we are interested in accounting for mobility patterns both in the late 1990's and mid 2010's, we calibrate our model separately for each period, though keeping certain parameters (e.g. those governing preferences)

constant across periods. For a subset of parameters we set values following external data. For the remaining parameters we select them by asking the model to replicate the previously outlined empirical facts of geographic mobility for young (25-34) college graduates in either the 1996 (early period) or 2014 (later period) samples of the SIPP.

4.1 First-stage parameters and initial conditions

A period in the model refers to a year. Individuals enter the model at age 25, retire at age 60, and die no later than at age 95. Survival rates are obtained from Kaplan and Violante (2014). We set the risk free interest rate to 3 percent and assume individuals have CRRA preferences over consumption setting $\sigma = 2$. We set the borrowing rate r^b equal to 15 percent and the borrowing limit \underline{a} so that the bottom 5 percent of the wealth distribution for new labor market entrants is just at the borrowing constraint. This allows us to externally set the initial distribution of wealth-to-earnings to match the distribution of wealth-to-earnings of 25-26 year-olds calculated from the 1998 SCF. The earnings process parameters $\rho, \sigma_e, \sigma_\epsilon$ and the life cycle component of income l_j are calibrated following Kaplan and Violante (2014). We set mobility cost κ to the equivalent of \$12,000 following Rappaport (2012). Lastly, while in retirement, individuals receive a fraction of the last realization of the persistent component of their working-age income. To be consistent with U.S. replacement ratios, we target an average replacement rate of 47 percent. Table 4 collects the parameters set using external information.

Table 4: Parameters Set Externally

Parameter	Value	Source
\bar{l}_j , common component of life-cycle	–	Kaplan and Violante (2014)
ρ , persistence of persistent shocks	1.0	Kaplan and Violante (2014)
σ_e^2 , Variance of persistent shocks	0.01	Kaplan and Violante (2014)
σ_ϵ^2 , Variance of transitory component	0.05	Kaplan and Violante (2014)
σ , Coefficient of relative risk aversion	2.00	Standard
\underline{a} , borrowing constraint	-3.66	SCF 1998
R , Gross risk free interest rate	1.03	Standard
R^b , Gross borrowing rate	1.13	Standard
W , Retirement age	60	Standard
T , last age of life	95	Standard
ρ_n , mortality profile	–	Kaplan and Violante (2014)

Unlike the aforementioned parameters, we allow the initial distribution of wealth to vary across periods. Specifically, we feed in the 5th, 10th, 25th, 50th, 75th, 90th, and 95th percentiles as shown in the top row of Table 1 into the model. When calibrating the model to match mobility patterns from the 1996 SIPP (early period) we use the 1998 SCF (top row). Similarly, when calibrating the model to match mobility patterns from the 2014 SIPP (late period) we use the 2016 SCF (bottom row).

4.2 Estimation

Next, we calibrate the remaining parameters “internally” by asking the model to match a few facts from Section 2. The remaining parameters are those that govern the present-value of earnings gains from moving (Δ_0, Δ_1) , the share of opportunities that are front-loaded (θ) , the average utility cost of moving μ_j^M , and the subjective discount factor β .

To calibrate the mobility parameters we target a few of the empirical regularities as presented in Section 2 for each period. Specifically, for each period we target the average mobility rate and the average earnings change of inter-state movers (relative to stayers) 1, 2, and 3 years after moving for the 25-34 age group. This set of moments helps identify $\Delta_0, \Delta_1, \theta$, and μ_j^M .

Finally, we use the average wealth to average income ratio from the 1998 SCF (for all ages) to pin down β . Unlike the other parameters, we treat this parameter as fixed across periods.

As is standard, we use a minimum distance estimator as in Chamberlain (1982, 1984), which minimizes a weighted squared sum of differences between model and data moments. The estimator solves the following problem:

$$\min_{\Theta} [\hat{g} - g(\Theta)]' W [\hat{g} - g(\Theta)], \tag{10}$$

where $g(\Theta)$ and \hat{g} are $(J \times 1)$ vectors of model-based and data-based moments, respectively; and Θ is an $N \times 1$ vector of structural parameters to be estimated. W is a $J \times J$ weighting matrix, which we assume to be the identity matrix following Altonji and Segal (1996). Table 5 shows that the model is able to match the early period, in particular, fairly well, while Table 6 displays the necessary parameter differences across periods to accomplish this fit.

Focusing on the early period, the model is very close to the data, which is accomplished by having mostly back-loaded moving opportunities $(\theta = 0.05)$ that growth over time. The implied parameters suggest an average gain reaching 16 percent (in a present value sense) in the third period following the move. This makes moving very attractive, which means that the average utility cost is quite high to match the average mobility rate observed in the early period.

Turning to the late period, the model fit is not as tight, but does replicate the different time profile of earnings gains upon moving combined with the lower mobility rate. Naturally, more moving offers are front-loaded $(\theta$ rises to 0.46). Interestingly, the model requires smaller values for both Δ_0 and Δ_1 to match the relative earnings gains of movers seen in the data. Viewed through the lens of the model, migration offers became more *similar* to the current location/offer. Indeed, if $\Delta_0 = \Delta_1 = 0$, then the earnings process for movers collapses to that of stayers. This quantitative finding echoes the results of Kaplan and Schulhofer-Wohl (2017), who explain the secular decline in interstate mobility with a decline in the variance of earnings across locations.

Table 5: Model Fit

	Early		Late	
	Data	Model	Data	Model
Avg. Mobility (in %)	6.67	6.51	4.21	4.22
Avg. Earnings Gain Movers after 1 yr (in %)	7.5	10.1	15.3	8.6
Avg. Earnings Gain Movers after 2 yrs (in %)	17.8	20.1	6.0	7.3
Avg. Earnings Gain Movers after 3 yrs (in %)	31.3	36.0	5.4	5.4
Mean nw/mean income	2.61	2.63		

Note: All mobility and earnings statistics are for young (25-34) college graduates.

Table 6: Implied Parameter Values

Parameter	Early	Late
Δ_0	0.08	0.01
Δ_1	0.06	0.03
θ	0.05	0.46
$\mu_{j \leq 34}^M$	4	1.6
β	0.975	0.975*

*We do not recalibrate β .

5 Accounting For the Decline in Mobility

With the model calibrated to both periods, we now explore the drivers of the decline in mobility between the late 1990's and mid 2010's. To do this, we change parameters or initial conditions from their early to late values to see how each affects mobility and earnings growth early in life. We find that the reduction in overall earnings gains is the most important factor for reducing mobility. However, the reduction in initial wealth across periods for young college educated workers also has a negative impact. In the second subsection, we show this is especially true for those who enter the labor force wealth poor.

5.1 Decomposing Changes From Early to Late on Average

Table 7: Accounting For The Decline in Mobility

	Early Calibration	Adding Late			
		Earnings Process Δ_0, Δ_1	Initial Wealth a_0	Share Front-Loaded θ	Utility Cost μ^M
Avg. Mobility (in %)	6.51	0.7	6.42	6.61	13.7
Avg. Earnings Gain Movers after 1 yr (in %)	10.1	8.9	10.2	15.9	5.2
Avg. Earnings Gain Movers after 2 yrs (in %)	20.1	8.0	20.2	26.0	12.9
Avg. Earnings Gain Movers after 3 yrs (in %)	36.0	9.8	35.8	36.5	28.1

All mobility and earnings statistics are for young (25-34) college graduates.

Table 7 highlights the key drivers of the decline in mobility between the early and late periods. For

reference, the first column in Table 7 restates the early-period calibration. The next column changes the earnings process (Δ_0, Δ_1) to the late period's, where the gains from moving fall sharply for either type of move. We see that overall mobility falls from 6.5 to 0.7 percent, well below the average mobility rate of 4.2 percent observed in the later period. This suggests that the overwhelming reason for lower mobility in the latter period is the smaller potential gains from moving, regardless of the type of move (front- or back-loaded).

However, lower earnings gains upon moving are not the only reason for the decline in mobility between periods. The next column of Table 7 examines the effect of changes in initial wealth across periods. Specifically, rather than feeding into the model the wealth to income percentiles from the early period (top row of Table 1), we instead feed into the model the corresponding percentiles for the late period (bottom row of Table 1). The third column of Table 7 shows that while the profile of earnings gains for movers remains essentially unchanged, measured mobility falls from 6.5 to 6.4 percent. While this decline appears modest we'll see in the next subsection that it masks a larger decline in mobility among the initially wealth poor.

The next column shows that a greater share of front-loaded moves (higher θ), all else equal, drives mobility up. This column shows that taking the early-period calibration and only changing the value of θ to the late period implies a small increase in mobility. Naturally, observed earnings gains in the first and second periods following the move rise, as more moves are associated with front-loaded gains. Again, we'll see in the next subsection that this small increase in average mobility masks a larger increase in mobility among the initially wealth poor, highlighting the interaction of market incompleteness, discounting, and the timing of earnings gains in shaping the mobility decision.

Lastly, not surprisingly, lower utility costs of moving (lower μ^M) are also associated with higher mobility. The fifth column of Table 7 shows that taking the early-period calibration and only changing the value of μ^M to the late period implies nearly a two-fold increase in mobility. As all other parameters (including the earnings gains upon moving) are held fixed this increase in mobility leads to a decline in observed earnings gains upon moving at all horizons. This is due to the fact that by making the utility cost of moving lower, individuals at comparatively higher earnings levels move more than before. This mechanically drives down the average measured gain upon moving. To this point, average pre-move earnings among movers are 36 percent higher in this counterfactual than in the baseline economy.

Overall, the analysis of this subsection shows that while changes in the earnings process between the early and late periods help explain the majority of the observed decline in mobility, changes in initial wealth also matter. In the next subsection we focus our attention on those who are initially wealth poor to show that this effect is amplified.

5.2 Comparing Young Poor and Wealthy

While the change in average mobility due to changes in wealth documented in the previous subsection appears small, it masks substantial heterogeneity. Indeed, a key take-away from Table 1 is that college educated workers with initial wealth below the median became significantly poorer from the late 1990's to mid 2010's. Our model suggests this alone had a significant effect on their mobility decisions, which in turn

affected their earnings growth and therefore overall earnings inequality.

To see this more clearly, Table 8 presents a breakdown of average mobility and cumulative earnings through age 38 (when mobility no longer affects earnings) across the initial wealth distribution (i.e. from the late 1990's) for the early period along with the same counterfactuals presented in the previous section. Looking at the first column reveals that even the baseline calibration masks substantial heterogeneity in mobility across the initial wealth distribution. Indeed, the relationship between mobility and initial wealth is an inverted-U shape, with mobility being lower for the initially poor (who cannot easily move) and initially rich (who are well insured to begin with), but higher for those in the middle (who can afford and benefit from moving).

Table 8: Heterogeneity in Accounting For The Decline in Mobility

	Early Calibration	Adding Late			
		Earnings Process Δ_0, Δ_1	Initial Wealth a_0	Share Front-Loaded θ	Utility Cost μ^M
Avg. Mobility (in %)		ppt. chg. from early calibration			
below 25th pctile of a_0	6.48	-5.74	-0.98	0.19	6.47
between 25-75th pctile of a_0	7.12	-6.34	0.06	0.10	7.37
above 75th pctile of a_0	5.32	-4.70	0.52	0.02	7.69
average	6.51	-5.78	-0.09	0.10	7.22
Avg. Cumm. earnings through age 38		% chg. from early calibration			
below 25th pctile of a_0	32.99	-16.05	-2.96	1.13	17.05
between 25-75th pctile of a_0	32.32	-17.73	-0.01	0.46	19.86
above 75th pctile of a_0	34.26	-12.39	1.47	0.49	20.77
average	32.97	-15.92	-0.36	0.64	19.39

Notes: All percentiles are calculated based on the initial (age 25) distribution of wealth a_0 from the early calibration.

The second column of this table shows that changes in the earning process from the early to late period alone predict a sharp decline in mobility regardless of initial wealth. Mobility for those who enter the labor market with wealth between the 25th and 75th percentiles is reduced by 6.3 percentage points. With less earnings enhancing mobility early in the life cycle, cumulative earnings (through age 38) fall across the initial wealth distribution, though slightly more for the aforementioned group (who have the highest mobility rates in the baseline economy and see the largest decline in them).

Turning to the third column, we see that the previously documented effect on mobility of changing the initial wealth distribution from the early to late period is asymmetric between the wealth poor versus rich. Recall, while wealth on average fell between the early and late periods, it fell more dramatically for the bottom quartile. Focusing on the third column, the reduction in initial wealth among the bottom quartile alone leads to a nearly one percentage point *decline* in their mobility, much larger than the 0.1 percentage point decline in mobility documented in the previous section. By contrast, lower initial wealth among the

top quartile leads to a half percentage point *increase* in their mobility, as they seek to offset their decline in initial wealth with higher earnings through mobility. As consequence, cumulative earnings among the bottom quartile fall by nearly 3 percent, whereas they rise by 1.5 percent among the top quartile. All told, our model implies that changes in initial wealth between the two periods alone lead to an endogenous increase in income inequality (as measured by the gap in cumulative earnings through age 38 between the bottom and top 25th percentiles of initial wealth) of 4.5 percent.

Turning to the fourth column of this table we see more clearly how the interaction of wealth, market incompleteness, and the timing of earnings gains shapes mobility. Specifically, this column shows that a higher share of front-loaded mobility opportunities encourages mobility among the wealth poor, but has little impact among the wealth rich. Our model predicts that mobility among those in the bottom quartile of initial wealth increases by 19 basis points, but is essentially unchanged among those in the top quartile. Though modest, these asymmetric differences emphasize how debt-constrained individuals discount back-loaded opportunities (which are more common in the early calibration) and favor front-loaded ones (which are more common in the late calibration).

Lastly, the fifth column of this table echoes the results from the previous section and shows that a lower utility cost in the latter period helps drive up mobility. This is true across the initial wealth distribution, though slightly more so for the wealthy. While mobility increases by 6.5 percentage points among those in the bottom quartile of the initial wealth distribution, it increases by 7.7 percentage points among the top quartile. Naturally, the mobility decision of richer individuals is less dependent on wealth and more on the utility cost of moving. Thus, this group is more responsive to changes in the utility cost of moving.

Overall, a key take-away from this section is that the reduction in wealth that young college educated workers saw between the late 1990's and mid 2010's had significant effects on their mobility decisions. This in turn affected earnings growth early in their careers and subsequently widened earnings inequality. Our counterfactuals suggest that lower initial wealth alone lead to an endogenous increase of 4.5 percent in the cumulative earnings gap between the poor (bottom quartile of the initial wealth distribution) and wealthy (top quartile).

6 Policy Analysis

Motivated by the quantitative results from the previous two sections documenting the importance of initial wealth in shaping mobility decisions, we now consider policies that reduce debt for workers who enter the labor force with low wealth. In particular, we consider a version of recently proposed student loan forgiveness plans through the *Higher Education Act* and use our model as a laboratory to gauge the effects of student loan forgiveness on mobility.⁵

To operationalize the student debt forgiveness plan in our model, we proceed as follows. Using the results from Table 2, we calculate what fraction of initial net wealth in 2016 is student loan debt. We then subtract

⁵Additional details on the proposed legislation can be found at: <https://studentaid.gov/manage-loans/forgiveness-cancellation/debt-relief-info>

this amount from initial net wealth for each individual in our 2016 economy.⁶ In all experiments, we use lump sum taxes to finance the debt forgiveness policy.

Table 9: Student Debt Forgiveness and Mobility

	Late Calibration	Student Debt Forgiveness
Initial Net Worth		% chg. from late
below 25th pctile of a_0	-1.87	-174.8
between 25-75th pctile of a_0	0.53	47.7
above 75th pctile of a_0	3.63	18.3
average	0.71	155.8
Avg. Mobility (in %)		ppt. chg. from late
below 25th pctile of a_0	3.36	1.13
between 25-75th pctile of a_0	4.56	0.07
above 75th pctile of a_0	4.40	0.04
average	4.22	0.33
Avg. Cumm. earnings through age 38		% chg. from late
below 25th pctile	28.40	0.89
between 25-75th pctile	27.23	0.12
above 75th pctile	31.75	-0.03
average	28.66	0.27
Avg. Welfare Change (in %)		
below 25th pctile of a_0		261.2
between 25-75th pctile of a_0		-5.1
above 75th pctile of a_0		16.4
average		66.9
Lump sum tax/ Income		0.0081

Notes: All percentiles are calculated based on the initial (age 25) distribution of wealth a_0 from the late calibration. The welfare change is relative to laissez-faire for each period, measured as the amount of wealth that households would need to be indifferent between being born into the laissez-faire economy rather than the economy with debt forgiveness policies, measured as a percent of average income for new labor force entrants.

⁶Like the proposed plan, we assume that individuals earning more than \$120,000 are ineligible for student debt forgiveness. However, as we forgive debt at age 25, this constraint rarely binds in our experiments.

Table 9 shows how student debt forgiveness affects initial wealth, mobility, and cumulative earnings using our late calibration as a baseline. Looking at the top panel, by construction wealth of those who enter the labor force in the bottom quartile of the initial wealth distribution is impacted the most by the policy as they hold the largest amount of student debt both in absolute terms and relative to net worth or income. Indeed, our calculations imply these individuals switch from having negative to positive net wealth as a result of the student forgiveness program. By contrast, wealth at other points of the initial wealth distribution is impacted less as wealthier individuals tend to hold less student debt overall and relative to their net worth.

Consistent with these changes in initial wealth, the second panel of this table shows that average mobility of the poorest increases the most in response to the policy. Mobility among those who were originally in the bottom 25th percentile of initial net wealth increases by 1.1 percentage points relative to the baseline rate of 3.4 percent. By contrast, mobility of wealthier individuals is essentially unchanged.

Echoing the results from the previous sections, the third panel of this table shows that higher mobility is subsequently associated with a modest increase in cumulative earnings through age 38. Focusing again on those who were originally in the bottom 25th percentile of initial wealth, their cumulative earnings increase by nearly one percent thanks to the greater ability to move to higher earnings early in the life cycle.

The bottom panel of this panel shows that student loan forgiveness increases welfare, on average, and in particular, for those who enter the labor market wealth poor. Viewed through the lens of our model, the policy increases welfare on average by nearly 67 percent, or by 261 percent for those who were originally in the bottom quartile of initial net wealth.

An important issue with these welfare calculations is that they encompass both the welfare gains from higher earnings through mobility, but also the welfare gains from redistribution (e.g. forgiving debt through taxation). To measure the welfare gains associated with the latter effect, we also consider a simplified version of our model where there are no earnings gains from moving. Hence, when individuals move it is purely for non-economic reasons. We use this simplified model as a counterfactual baseline economy and then reconsider the effect of student loan forgiveness in this case.

Table 10 presents the results of this exercise and reveals that while redistribution is the most important driver of the welfare gains reported in the previous table, greater mobility (and the associated earnings gains) still matter. Looking at the top panel of this table, we see that with no economic gains associated with mobility, all else equal, mobility rates are lower across the initial wealth distribution. Compared to the “true” late calibration, the average mobility rate is two percentage points lower in this economy without earnings gains from moving.

When we forgive student debt in this counterfactual economy mobility still rises on average, and more so for the originally wealth poor. Indeed, the second column of this table shows that mobility rises by 71 basis points when originally wealth poor individuals enter the labor market with less debt. Comparing this number with the corresponding one in the previous table suggests debt forgiveness increases economic-related mobility by 42 basis points among the originally wealth poor.

Turning to the bottom panel of this panel, welfare still increases with the policy. Comparing the welfare numbers between this table and the previous one implies that greater earnings gains from higher mobility

Table 10: Student Debt Forgiveness and Mobility With No Earnings Gains

	Counterfactual Late Calibration	Student Debt Forgiveness
Avg. Mobility (in %)		ppt. chg. from counterfactual late
below 25th pctile of a_0	1.50	0.71
between 25-75th pctile of a_0	2.18	0.02
above 75th pctile of a_0	2.65	0.05
average	2.13	0.20
Avg. Welfare Change (in %)		
below 25th pctile of a_0		259.4
between 25-75th pctile of a_0		-7.1
above 75th pctile of a_0		14.7
average		65.0
Lump sum tax/ Income		0.0081

Notes: All percentiles are calculated based on the initial (age 25) distribution of wealth a_0 from the late calibration. The welfare change is relative to laissez-faire for each period, measured as the amount of wealth that households would need to be indifferent between being born into the laissez-faire economy rather than the economy with debt forgiveness policies, measured as a percent of average income for new labor force entrants.

(the main focus of this paper) increase welfare on average by nearly 2 percent.

7 Conclusion

We explore the importance of the timing and size of gains from moving across state lines to a job opportunity in a world with uninsurable income risk and financial market frictions. In the late 1990s the average annual inter-state mobility rate among college graduates was nearly 7 percent. Young inter-state migrants experienced large (roughly 30 percent) gains in their labor earnings within three years of moving, compared to similar non-migrants. By the mid 2010's these patterns changed dramatically with the average inter-state mobility rate declining dramatically and the gains from moving becoming much smaller and front-loaded.

Using a life cycle model of consumption, savings, and mobility we explore how changes in wealth (particularly among the young), credit constraints, and changes in the potential returns to moving explain these facts. We find that the main reason for declining mobility is the overall size of gains, but that the decline in wealth of workers upon entering the labor force also has a negative impact. The effect of lower wealth on mobility is especially negative for those who are poor, which means that changes in wealth endogenously increase earnings inequality through the mobility channel.

We then use the model to evaluate the potential effects of hypothetical tax-financed debt forgiveness programs for the young, similar to those proposed in recent student loan forgiveness programs. The model suggests that debt forgiveness could have a significant effect on the mobility of low-wealth workers. Greater mobility contributes to early-career earnings growth and thus welfare. As the model-implied tax burden of

forgiving debt among young workers is small, our analysis suggests that such a policy could have overall positive welfare implications.

References

- ALTONJI, J. G. AND L. M. SEGAL, “Small-Sample Bias In GMM Estimation Of Covariance Structures,” *Journal of Business and Economic Statistics* 14 (July 1996), 353–366.
- BILAL, A. AND E. ROSSI-HANSBERG, “Location as an Asset,” *Econometrica* 89 (September 2021), 2459–2495.
- BORJAS, G. J., S. G. BRONARS AND S. J. TREJO, “Self-selection and internal migration in the United States,” *Journal of urban Economics* 32 (1992), 159–185.
- CHAMBERLAIN, G., “Multivariate Regression Models for Panel Data,” *Journal of Econometrics* 18 (1982), 5–46.
- , “Panel Data,” in Z. Griliches and M. D. Intriligator, eds., *Handbook of Econometrics* volume 2 (North-Holland Press, 1984), 1247–1318.
- COEN-PIRANI, D., “Understanding Gross Worker Flows Across U.S. States,” *Journal of Monetary Economics* 57 (October 2010), 769–784.
- , “Geographic mobility and Redistribution,” *International Economic Review* 63 (August 2021), 921–952.
- KAPLAN, G. AND S. SCHULHOFER-WOHL, “Understanding the long-run decline in interstate migration,” *International Economic Review* 58 (2017), 57–94.
- KAPLAN, G. AND G. L. VIOLANTE, “How Much Consumption Insurance beyond Self-Insurance?,” *American Economic Journal: Macroeconomics* 2 (October 2010), 53–87.
- , “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica* 82 (July 2014), 1199–1239.
- KARAHAN, F. AND S. RHEE, “Geographic reallocation and unemployment during the Great Recession: The role of the housing bust,” *Journal of Economic Dynamics and Control* 100 (2019), 47–69.
- KENNAN, J. AND J. R. WALKER, “The effect of expected income on individual migration decisions,” *Econometrica* 79 (2011), 211–251.
- RAPPAPORT, J., “Why Does Unemployment Differ Persistently Across Metro Areas?,” *Economic Review* 2 (2012), 5–35.
- ROCA, J. D. L. AND D. PUGA, “Learning by Working in Big Cities,” *The Review of Economic Studies* 84 (January 2017), 106–142.