Modeling Intergenerational Economic Mobility Dynamics Via Calibrating Markov Models

Aneesh Tekulapally
Advisor: Professor Daniel Rigobon

Submitted in partial fulfillment
of the requirements for the degree of
Bachelor of Science in Engineering
Department of Operations Research and Financial Engineering
Princeton University

May 2024

I hereby declare that I am the sole author of this thesis.

I authorize Princeton University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Aneesh Tekulapally

I further authorize Princeton University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Aneesh Tekulapally

Abstract

Quantifying rates of intergenerational economic mobility, or the ability for children to achieve a higher standard of living than their parents (the "American Dream"), is a challenging empirical task. Previous studies have largely relied on measures such as log-log elasticities and rank-rank correlations to assess levels of mobility. However, these models are limited in their precise quantification of specific transition probabilities, which model the likelihood of a child transitioning to a particular socioeconomic group given their parents' data. In this work, we formulate a series of Markov transition matrices to model observed rates of intergenerational mobility over several decades using data from the Panel Study of Income Dynamics. Both relative and absolute mobility formulations are considered, where relative mobility is defined by income quintiles and absolute mobility is defined by discrete fixed income buckets. We demonstrate that rates of intergenerational relative mobility have remained remarkably stable between the 1968 and 1997 birth cohorts, with high levels of income persistence. These findings largely align with past literature even though income inequality has increased in subsequent decades. However, our absolute mobility formulation indicates statistically significant results of higher levels of upward mobility (almost double the likelihood of joining the top income bucket) when comparing the 1968 and 1997 birth cohorts. We note that the results from the absolute mobility parameterization are confined to this paper's specific parameterization and are likely structurally skewed towards greater perceived levels of intergenerational mobility. Furthermore, we acknowledge that our findings are limited in terms of broader inferences regarding intergenerational mobility due to the lack of significant robustness checks and limited historical data. Yet, the results are still interesting and useful for researchers in terms of providing a future model framework for better estimators of economic mobility.

Acknowledgements

Thank you, Professor Rigobon, for your invaluable support and guidance over the last several months. Particularly in the several weeks preceding the deadline, your mentorship and time were a calming presence that made the thesis experience enjoyable.

Additionally, I would like to thank my graduate student TA, Rajita, for leading our breakout room sessions throughout the year and providing useful insights for an inexperienced researcher like me.

Next, I would like to thank my family and friends who have endlessly encouraged and pushed me. Mom, Dad, Abhi, and friends have been the best role models, and I'm so appreciative for all that you've done.

Lastly, I need to thank Princeton University for the opportunity of learning and growing at this amazing institution.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The "American Dream" – the belief that each generation can achieve a better life than the previous one – has been a core American value since the nation's founding (Samuel, 2012). While intergenerational economic mobility has long been a central tenet of the American ethos, it was not until the 20th century that systematic studies on this topic began to appear. Early research was primarily rooted in sociological theory, focusing on the relationship between immigration and economic mobility or examining the role of education and family background in economic outcomes (Borjas, 1992; Hout, 1988).

In the 1960s, with President Lyndon B. Johnson's Great Society programs aimed at alleviating poverty and inequality, there was a growing interest in studying intergenerational economic mobility (Katz, 2013). The demand for empirical evidence of economic progress led to the aggregation of new, large-scale longitudinal datasets. This development enabled researchers to conduct more rigorous quantitative analyses on income mobility across generations.

Studying intergenerational economic mobility provides crucial insights into how effectively institutions and policies foster equal opportunities. By examining economic mobility trends over time, policymakers can gain a deeper understanding of recurring poverty cycles and analyze environments where an individual's economic progress is independent of their prevailing socioeconomic background (Solon, 1999).

Furthermore, investigating intergenerational economic mobility through quantitative analysis of socioeconomic group transitions has widespread implications. A better understanding of the likelihood of transitioning between economic subgroups can help researchers explore how factors like family background, education, and labor market conditions are linked to children's economic outcomes with greater certainty. Note that the complex nature of modeling economic mobility measures makes it difficult to establish direct causal links, yet analysis still offers valuable insights for all stakeholders involved.

Moreover, research on intergenerational mobility plays a vital role in public discourse around challenging systemic inequities and ensuring fair distribution of collective resources (Corak, 2013). Therefore, a deeper understanding of mobility can lead to more informed discussions about inequality and socioeconomic persistence. Findings from studies on intergenerational mobility can either affirm or challenge prevalent narratives concerning the "American Dream," potentially shaping public opinion and influencing political will in the allocation of government resources.

The primary objective of this paper is to analyze observed changes in intergenerational economic mobility rates across several decades, employing two distinct models: relative and absolute mobility. Relative mobility is defined as a child's ability to transition to a different income quintile compared to their parent's income quintile. In contrast, absolute mobility is defined as a child's likelihood of moving to a higher or lower income bucket than their parents, where income buckets are determined by fixed dollar thresholds rather than quintile rankings. Specifically, this paper aims to calibrate a series of Markov transition matrices for both relative and absolute mobility analyses, providing better estimators of individuals' transitions between economic groups over time. By doing so, we seek to contribute to the understanding of intergenerational mobility patterns and their evolution across generations.

The Markov mathematical framework is the preferred stochastic tool for modeling socioeconomic mobility as it represents the likelihood of transitioning from one state to another within a defined system. Markov processes can neatly represent transitions between social classes over time and can model the long-term stationary distribution of the likelihood of moving between economic ranks. Crucially, Markov chains are suitable for socioeconomic mobility analysis due to their memoryless property. This property means that the probability of transitioning to a future state depends solely on the current state rather than the sequence of preceding events. Thus, the underlying qualitative principle assumed by this model is that an individual's economic situation is primarily influenced by their parents' economic situation, rather than any residual effects from prior generations.

Recent studies on relative economic mobility have utilized a Markov transition framework, but this approach has been limited to analyzing a single column of the transition matrix. This paper aims to expand on prior research by comparing transition probabilities between two distinct birth cohorts on an element-by-element basis and plotting more comprehensive time series data of transition probabilities spanning multiple decades. The implementation presented in this paper enables a more extensive examination of how rates of relative mobility have evolved across generations.

Previous research has quantified absolute intergenerational income mobility by examining the proportion of children whose incomes surpass those of their parents. However, these studies often present a binary perspective, solely focusing on whether children earn more or less than their parents. This study aims to expand upon prior work by analyzing the extent of deviation in children's earnings relative to their parents'. Specifically, we introduce a Markov transition matrix to quantify the likelihood of a child belonging to a fixed absolute income bucket, conditional on the income bucket of their parents. Each income bucket serves as a proxy for various lifestyle categories, with all observed incomes adjusted to the first year in the dataset.

The following provides an overview of the subsequent chapters:

- Chapter 2: Introduces the mathematical architecture of Markov models for relative and absolute mobility. Discusses the mathematical features of Markov transition matrices, including the stationary distribution and mixing times. It then defines the delta matrix to compare transition probabilities, and lastly, discusses the model assumptions.

- Chapter 3: Overviews the datasets used to calibrate the Markov models while detailing the cleaning and preprocessing procedures.

- Chapter 4: Presents results for the relative mobility parametrization, including how to calibrate the transition matrices, comparing the first and last transition matrix, and time series plots of specific transition probabilities.

- Chapter 5: Similar to the previous chapter, this chapter provides results for the absolute mobility parametrization. It compares the first and last transition matrix, plots time series of specific transition probabilities, and discusses the stationary distribution and mixing times.

- Chapter 6: Reflects on the previous chapters and places the results in the context of existing literature. Lastly, it presents this paper's contributions, limitations, and areas for future work.

1.1 Literature Review

1.1.1 Relative Mobility

Following the proliferation of large-scale longitudinal datasets, Zimmerman (1992) employed several novel econometric strategies to estimate the elasticity between parents' and their children's lifetime earnings in the United States. By measuring parent-child earnings elasticity, this foundational paper delved into relative intergenerational economic mobility.

Zimmerman (1992) makes an important contribution to the literature on intergenerational mobility by challenging consensus views from previous studies (House, 1976; Behrman and Taubman, 1985). These papers indicate that children's earnings are not strongly correlated with their parent's earnings with an intergenerational earnings elasticity of roughly 0.2 or less (lower elasticity signals greater mobility and less correlation between parent-child incomes). Through the formal modeling of measurement error and transitory fluctuations in annual earnings measures, Zimmerman's statistical model suggests intergenerational earnings elasticity of around 0.4, indicating much less mobility than prior studies. Note that an intergenerational earnings elasticity of 1 signals a perfect correspondence between parent-child incomes.

However, more robust methods utilizing panel data on earnings histories would help reduce sensitivities to assumptions as the paper's results rely on a relatively small sample of under 900 father-son pairs. Additionally, Zimmerman (1992) focuses on understanding average intergenerational mobility between 1981 and 1996, as it lacks time series data on earnings elasticity.

Zimmerman (1992) also conceptualizes a Markov transition matrix as an alternative model to analyze intergenerational mobility. The paper defines a transition matrix that categorizes income ranks (such as quartiles) of father-son pairs, where each element of the transition matrix signifies the likelihood that a son will reach a particular quartile, given that their father is in a specified income interval. In this formulation, absolute immobility would be represented by an identity transition matrix, which indicates a direct and complete correlation between the father's and son's income ranks. Conversely, complete mobility would be depicted by a transition matrix with identical probabilities in each element, as the son's income rank would then be independent of their father's income rank.

Nonetheless, the paper underutilizes this Markov approach due to empirical challenges related to data availability at the time. Given the utility of transition matrices in quantifying intergenerational economic mobility, the increased collection of income data and the aggregation of datasets present avenues for further analysis.

Recent studies by Chetty et al. (2014) present new observations of relative intergenerational income mobility that build off Zimmerman (1992). Unlike Zimmerman's focus on average parent-child income elasticity over a time interval, (Chetty et al., 2014) focus on measuring intergenerational mobility over several decades via rank-rank correlations and quintile Markov transition matrices.

Chetty et al. (2014) addresses numerous pitfalls of Zimmerman (1992), through a more robust mathematical framework and larger more complete datasets. The paper indicates that utilizing a rank-based model rather than elasticity demonstrates greater resilience to lifecycle and attenuation bias from transitory income fluctuations, which are also prominent in Zimmerman (1992). Regarding data, Chetty et al. (2014) vastly expands upon the 900 father-son pairs sampled in Zimmerman (1992), using de-identified tax records to link parents' income to their children's income. For children born from 1980 onwards, the study constructs a parent-child sample using U.S. tax records from 1996-2012, linking approximately 95% of children in each birth cohort to their parents and yielding a sample of 3.7 million children per cohort. However, for birth cohorts prior to 1980, the study utilizes the Statistics of Income (SOI) annual cross-sections, a stratified random sample covering 0.1% of tax returns, to link children to parents based on dependent information available starting in 1987. By utilizing substantial sample sizes and more comprehensive income data, Chetty et al. (2014) establishes that rank-rank correlations of intergenerational mobility have remained remarkably stable for individuals born between 1971 and 1993.

Like Zimmerman (1992), Chetty et al. (2014) also supplements the paper's primary methodology—rank-rank correlations for Chetty et al. (2014) and parent-child elasticity for Zimmerman (1992)—by considering Markov transition matrices to directly measure the probability that a child reaches the top quintile of the income distribution. Chetty et al. (2014) formulates quintiles by ordering children in comparison to their peers within the same birth cohort, and similarly ranking parents based on their standing amongst other parents with children in the corresponding cohort.

While Zimmerman's analysis based on Markov chains was rudimentary in design, Chetty et al. (2014) plots the probabilities of children ascending to the highest income quintile within their birth cohort, based on their parents' income quintile (Chetty et al., 2014, see Figure 3 on pg. 21). Note that Chetty et al. (2014) does not reveal or discuss the complete quintile Markov transition matrix estimates; instead, it presents the vector corresponding to transitions to the highest child income quintile from the set of parent income quintiles. This vector, corresponding to the lowest parent income quintile to the highest, was estimated as 9%, 18%, 19%, 23%, 31% for the child birth cohort of 1971, with these probabilities exhibiting little to no directional trend for the subsequent 15 years. This supports the paper's primary rank-rank correlation analysis, indicating that intergenerational mobility remained stable throughout the 1970s and 1980s.

The partial implementation of the Markov transition matrix to directly quantify the probability of a child moving to the top quintile was a significant advancement in understanding relative intergenerational economic mobility. However, the Markov transition matrices can be further defined and explored for broader applicability in modeling mobility dynamics. By doing so, it's possible to generate time series probabilities of a child moving from any defined income rank to another, instead of merely the probability of reaching the top quintile. For example, we can plot the diagonal entries of the transition matrix that correspond with the probability of income persistence or remaining in the same income quintile as your parents. Additionally, conducting an element-wise statistical significance comparison of transition matrices offers greater robustness to the conclusions of Chetty et al. (2014).

1.1.2 Absolute Mobility

Prior studies of intergenerational economic mobility have focused on predominantly relative mobility, typically measured as income elasticities or rank-rank correlations (Zimmerman, 1992; Chetty et al., 2014). While many of the leading studies on relative mobility including Chetty et al. (2014) suggest that relative mobility has been fairly stable over recent decades, a newer paper also by Chetty et al. posits that relative and absolute mobility can diverge if income growth rates vary across the distribution (Chetty et al., 2017). This document introduces absolute mobility as an alternative metric to intergenerational economic mobility, specifically defining it as the proportion of children who have outearned their parents between the 1940s and the 1980s birth cohort.

Despite longstanding interest in absolute mobility studies, data limitations linking parents and children have impeded empirical analysis. Chetty et al. (2017) makes two methodological contributions to overcome the prevailing data limitations. First, Chetty et al. (2017) constructs marginal income distributions for parents and children by birth cohort, using Census and Current Population Surveys (CPS) cross-sectional data. Second, it estimates copulas between parents' and children's incomes by using tax data and assumes copula stability over time to project backward. By combining marginal income distributions with the copula—which provides probabilities for each child-parent income rank pairing—Chetty et al. (2017) can effectively estimate the fraction of children earning more than their parents.

Chetty et al. (2017) documents significant empirical findings that rates of absolute mobility have fallen sharply, from approximately 90% for the 1940 birth cohort to 50% for the 1980 birth cohort with the most severe declines amongst middle-class families. The novel combination of marginal income distributions with copulas between parent-child incomes has been pivotal in studying absolute intergenerational economic mobility.

However, further analysis could be conducted to better understand absolute intergenerational mobility over time. To extend the analysis from Chetty et al. (2017), we propose modeling absolute mobility as a Markov process, with a transition matrix representing the probability of a child belonging to a predefined income bucket conditional on their parents' bucket. This parameterization differs from that of relative mobility, as there is not an equal number of individuals in each subgroup, and income is not normalized to form quintiles for every birth cohort. Additionally, conceptualizing absolute mobility as a Markov process introduces useful mathematical features, such as the stationary distribution. This distribution reveals the long-term probability of individuals belonging to an income subgroup and is useful in understanding and comparing the implied long-term trends of current mobility dynamics over time.

Chapter 2

Markov Model Mathematical Formulation

This chapter first mathematically defines an arbitrary Markov model and then provides a parameterization framework for relative and absolute intergenerational mobility. Subsequently, mathematical features like the stationary distribution and mixing times of Markov matrices are discussed in the context of intergenerational mobility. Lastly, model assumptions are stated and examined in light of our specific parameterization.

2.1 Markov Model Definition

Let $X = \{X_t, t \in \mathbb{N}\}$ be an arbitrary Markov chain with state space $\mathcal{D} \in \{d_1, d_2, \ldots, d_k\}$. Throughout this paper, we use the notation:

$$p_{i,j} := \mathbb{P}\{X_{t+1} = j \mid X_t = i\}, \quad i, j \in \mathcal{D} \tag{2.1}$$

for all $n \geq 0$. We call $p_{i,j}$ the transition probability from i to j, and the matrix $P = [p_{i,j}]_{i,j \in \mathcal{D}}$ is called the transition probability matrix defined below:

$$P = \begin{bmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,k} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ p_{k,1} & p_{k,2} & \cdots & p_{k,k} \end{bmatrix}. \tag{2.2}$$

The transition probability matrix P satisfies the following properties:

$$0 \leq p_{i,j} \leq 1 \tag{2.3}$$

$$\sum_{j \in D} p_{i,j} = \sum_{j \in D} \mathbb{P}\{X_{t+1} = j \mid X_t = i\} = 1. \tag{2.4}$$

In addition to the transition probability matrix, the initial state distribution of a Markov chain can be defined as

$$\mathbb{P}_{x_0}\{\cdot\} := \mathbb{P}\{\cdot | X_0 = x_0\}. \tag{2.5}$$

2.2 Model Parameterization

Intergenerational mobility can be conceptualized by defining a simplified Markov model for each child birth cohort and analyzing changes in transition probabilities over time. A Markovian transition framework is valuable for modeling mobility, as it provides insights into the likelihood of a child belonging to a particular subgroup, given information about their parents' subgroup. The mathematical formulation is defined as follows:

Time Period (t)

To model transition probabilities over a generational shift from parent to child, a Markov parameterization only necessitates one time step. Therefore, $t \in \{0, 1\}$ with

$t = 0$ representing the parent generation and $t = 1$ representing the child generation.

State Space ($\mathcal{D}$)

The state space is defined by income subgroups that divide the distribution of income in a particular birth cohort. Note that these subgroups are uniquely parameterized for relative and absolute mobility. For relative mobility, the income distribution is subdivided into quintiles, while for absolute mobility, it is subdivided into fixed income buckets, for example, individuals making between $2,700 and $4,800. The specific parameterization of the state space for relative mobility is found in Section 4.1.1, whereas the parameterization for absolute mobility is detailed in Section 5.1.1.

Although the construction of these subgroups varies fundamentally between relative and absolute mobility, both parameterizations subdivide the income distribution into five distinct subgroups. These five subgroups can be defined as follows:

$$\mathcal{D} = \{\sigma_1, \sigma_2, \ldots, \sigma_5\} \tag{2.6}$$

Here, $\sigma_i \ \forall \ i \in 1, 2, \ldots, 5$ represents the five distinct income subgroups, with σ_1 corresponding to the lowest income grouping and σ_5 to the highest income grouping.

Transition Probability Matrix (P_c)

The Markov transition matrix encapsulates the probability of a child belonging to a specific income subgroup, given their parents' income subgroup. Note that this subgroup is defined by the state space of the relative or absolute mobility parameterization. Specifically, each entry $p_{i,j}$ represents the probability of transitioning from income subgroup i in the parent generation to income subgroup j in the child generation where $i, j \in \{1, 2, \ldots, 5\}$.

While state spaces differ between parameterizations, a more abstract transition matrix P_c can be defined below:

$$P_c = \begin{bmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,5} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,5} \\ \vdots & \vdots & \ddots & \vdots \\ p_{5,1} & p_{5,2} & \cdots & p_{5,5} \end{bmatrix}. \tag{2.7}$$

The subscript c corresponds with the transition matrix for a specific child birth cohort. For example, P_{1968} is the transition matrix for children born in 1968. Additionally, note that each row transition probability vector sums to one within each birth cohort such that Equations 2.3 and 2.4 are satisfied.

Figure 2.1: Visual Representation of Markov Transition Matrix (P_c)

Each element of the state space, σ_i, has an edge leading to every other element, including itself, since a child can transition to any possible income subgroup.

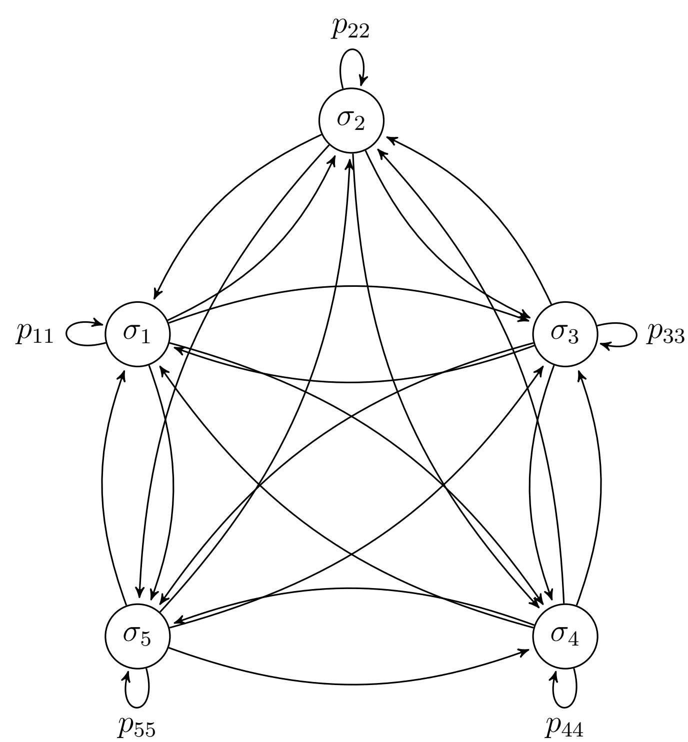

Initial State Distribution ($\mathbb{P}_{x_0}\{\cdot\}$)

The initial state distribution vector is deterministic as the income subgrouping of the parent generation in both the relative and absolute mobility parameterization is known from empirical data. In effect, there's no randomness in the state space $\mathcal{D}$ when $t = 0$ (the parent generation).

2.3 Markov Mathematical Properties

Markov transition matrices possess intrinsic characteristics that hold qualitative significance when modeling intergenerational mobility. Two noteworthy characteristics are the stationary distribution and mixing time, defined below:

Stationary Distribution

The stationary distribution represents the long-run probability distribution of the states in a Markov chain. Mathematically, the stationary distribution can be expressed as the vector π that satisfies the condition $\pi P = \pi$, where P is the Markov transition matrix. This condition implies that π is an eigenvector of P, as its span remains unchanged after successive transformations by P. However, it is important to note that π corresponds to the largest eigenvector of the transition matrix with a corresponding eigenvalue of 1. This is because the stationary distribution for a stochastic matrix must remain a valid probability vector, thereby satisfying Equations 2.3 and 2.4. The steps below outline the computational process:

1. Construct the transition probability matrix (P) for the Markov chain.
2. Find the eigenvalues (λ) and corresponding eigenvectors (v) of the transition probability matrix (P) by solving $Pv = \lambda v$.
3. Identify the eigenvalue $\lambda = 1$ and find its corresponding eigenvector (v).
4. Normalize the eigenvector (v) by taking the Euclidean norm such that $\pi = \frac{v}{|v|}$.

In the context of intergenerational mobility, the stationary distribution indicates the long-term equilibrium probability of an individual belonging to a specific income subgroup. Note that while our parameterization focuses on a single generational leap, the stationary distribution reflects the equilibrium implied by repeated generational shifts within a given transition matrix.

Since transition matrices (P_c) are calibrated per birth cohort, the stationary distribution can be analyzed over time to observe changes in the equilibrium distribution. This analysis offers insights into the long-term effects of current patterns in intergenerational mobility if they were to persist. By examining the stationary distribution, we can gain a better understanding of absolute mobility, especially since the predefined income buckets highlight long-term implied changes in the distribution of these income categories. However, the stationary distribution has limited qualitative use in analyzing relative mobility by definition, as it is defined by quintiles where 20% of the population is always in each grouping.

Mixing Times

The mixing time of a Markov process is a measure of how quickly the process converges to its stationary distribution. The stationary distribution of a Markov process with transition matrix P_c is the normalized eigenvector corresponding to the largest eigenvalue. However, the rate of convergence to the stationary distribution is determined by the second-largest eigenvalue of the transition matrix. Specifically, the mixing time (t_{mix}) for a stochastic matrix is a function of the spectral gap ($\Delta_{Gap} = \lambda_1 - \lambda_2$). Since λ_1 corresponding to the largest eigenvector is 1, t_{mix} is proportional to Δ_{Gap}:

$$t_{mix} = \frac{1}{\Delta_{GAP}} = \frac{1}{1 - \lambda_2} \tag{2.8}$$

In the analysis of mixing times, the specific value of the mixing time computed does not hold significant relevance within our mathematical formulation. Rather, the significance lies in observing the relative increase or decrease through time series analysis. A relatively shorter mixing time suggests that the Markov process converges more rapidly to its stationary distribution, indicating higher intergenerational mobility. This implies that it requires fewer generations for the distribution of individuals or families across socioeconomic states to become independent of their initial state. Conversely, a relatively longer mixing time denotes a slower rate of convergence towards the stationary distribution, which implies lower intergenerational mobility and greater persistence. This scenario indicates that individuals are more likely to remain within their socioeconomic states over multiple generations, experiencing limited opportunities for upward or downward mobility.

2.4 Delta Matrix

To quantify differences in transition probabilities across birth cohorts, we can define an arbitrary matrix that encapsulates observed differences between two transition matrices (P^b, P^a):

$$\Delta_P = P^b - P^a = \begin{bmatrix} p^b_{1,1} - p^a_{1,1} & p^b_{1,2} - p^a_{1,2} & \cdots & p^b_{1,5} - p^a_{1,5} \\ p^b_{2,1} - p^a_{2,1} & p^b_{2,2} - p^a_{2,2} & \cdots & p^b_{2,5} - p^a_{2,5} \\ \vdots & \vdots & \ddots & \vdots \\ p^b_{5,1} - p^a_{5,1} & p^b_{5,2} - p^a_{5,2} & \cdots & p^b_{5,5} - p^a_{5,5} \end{bmatrix} \tag{2.9}$$

The Δ_P matrix enhances the robustness of mobility models by directly comparing the differences in transition probabilities across specific birth cohorts. It serves a critical function by enabling the analysis of observed changes in key mobility measures. For example, it examines the likelihood that an individual remains in the lowest income subgroup if their parents were also in the lowest income subgroup ($p^b_{1,1} - p^a_{1,1}$).

2.5 Model Assumptions

To calibrate mobility dynamics via a Markov process the following are key assumptions:

Markov Property

The Markov property states that the probability of transitioning to the next state depends only on the current state. Mathematically this property can be denoted as:

$$P(X_n = x \mid X_{n-1} = x', \ldots, X_0 = x_0) = P(X_n = x \mid X_{n-1} = x') \tag{2.10}$$

In models of intergenerational mobility, the assumption is often made that income depends primarily on the income subgroup of the immediate prior generation rather than on those further back. This assumption aligns with the intuitive understanding that parents' economic circumstances, resources, and behaviors exert a stronger and more direct influence on their children's economic outcomes than extended family history.

Discrete Time Steps

In discrete-time models, we examine a system's state at designated points in time, tracking its progression through distinct steps rather than continuous change. In the formulation of the Markov transition matrix presented in this section, the time step represents a generation change from parent to child. Specifically, $t = 0$ corresponds to observations of parental income data while $t = 1$ corresponds to observations of child income data.

Discrete State Space

The Markov process is assumed to operate within a finite state space defined by five income subgroupings ($\mathcal{D}$) for each birth cohort. These income subgroupings form the basis for the relative and absolute mobility formulations and are both discrete but distinctively defined.

Other

In this section, the parameterization of the Markov model obviates the need for the conventional assumptions often required to calculate the stationary distribution. These assumptions include time-homogeneity, where transition probabilities between time steps are invariant over time; irreducibility, denoting the possibility of transitioning from any state to any other state within the system; and aperiodicity, which refers to the absence of cyclical patterns or loops within the process.

This parameterization does not require these assumptions because the transition matrices are calibrated on a single generational leap from parent to child. Such an approach emphasizes transitions from one state to another without having to account for the passage of time beyond a single generational change or the long-term behavior of the system.

Chapter 3

Data

This study utilizes panel data extracted from publicly accessible income datasets to calibrate Markov models for analyzing relative and absolute income mobility. This chapter offers an overview of the data aggregation and cleaning processes necessary to ensure an effective exploration of intergenerational mobility dynamics.

3.1 Sourcing

The foremost publicly accessible repository of longitudinal income data is the Panel Study of Income Dynamics (PSID), an ongoing study administered by the University of Michigan's Institute for Social Research (Survey Research Center, 2023). This paper leverages the comprehensive scope of the dataset, spanning its inaugural year of 1968 to the present.

The PSID began in 1968 as a nationally representative sample of over 18,000 individuals within the United States. Data on employment, income, wealth, marriage, education, and other factors have been repeatedly collected for each individual, with the same data collected on subsequent descendants. Throughout the last five decades, the dataset has expanded to reflect more than 85,000 individuals. The table below presents an overview of the downloaded variables along with descriptors.

Table 3.1: Summary of PSID Variables

These variables were extracted from the PSID database. All other variable options were deselected, with default settings retained. Note that these variable codes are specific to the 1968 dataset; codes for equivalent data in subsequent years will differ.

Category	**Variable**	**Description**
Identification	ER30000	Release number
	ER30001	Interview number
	ER30002	Person number
	ER32052	Individual's cohort year
Family	ER30003	Relationship to family head
	ER32009	Mother's ID
	ER32010	Mother's Person Number
	ER32011	Mother's Birth Year
	ER32016	Father's ID
	ER32017	Father's Person Number
	ER32018	Father's Birth Year
Personal	ER32000	Sex
	ER32004	Age
	ER30005	Marital Status
Income	ER30012	Total individual income

The resulting output is a dataset beginning in 1968 with subsequent yearly data corresponding to the variables in Table 3.1. Note that after 1997, the PSID datasets were not consistently released every year (just a few years with no data collected). To manage this issue, we retrospectively compute necessary values by using previously gathered data on those individuals and filling in parameters for the missing year. For any missing income variables, we update the prior year's income figure by implementing an adjustment factor that mirrors the income growth experienced during the relevant period. This adjustment is based on the individual's position within the income distribution (their income percentile) and the rate at which income changed for that specific percentile. Data for calculating income growth at the percentile level is derived from the Integrated Public Use Microdata Series (IPUMS), specifically the Current Population Survey (CPS) which provides yearly 1% samples of the population (Flood et al., 2023).

3.2 Cleaning

In preparation for analysis, income data within yearly PSID datasets is adjusted, repaired, and aggregated for balance and consistency. The table below details the cleaning methods conducted on the raw PSID files.

Table 3.2: Documentation for PSID Clean-Up

These are the two primary functions to reorganize the data from the PSID database.

Function	**Description**
Clean	Processes the downloaded PSID data file for a specified year. Addresses missing values and inconsistencies, generates additional variables, and unifies formats across reporting years to account for variable additions.
Reshape	Consolidates PSID files into a unified dataset and structures individual entries into family units, facilitating the streamlined retrieval of parental data.

The following subsections offer a comprehensive overview of each cleaning process function. Note that this same dataset is used for both the relative and absolute mobility Markov calibration.

3.2.1 Clean

The cleaning function initiates an iterative process on the downloaded PSID data, starting from the initial year of 1968. First the function identifies and removes rows that have missing data in key columns related to parental identification, sex, age, and income. It is important to note that rows with insufficient parental linking data are also excluded at this stage. This ensures that the subsequent reshaping analysis is conducted with complete information. Furthermore, the function creates a new variable for unique identification (`ID`), which merges the release year, interview number, and individual person number. This unique identifier is vital for accurately tracking individuals over time. The `ID` field allows for the inclusion of new individuals through birth, adoption, or marriage, facilitating a comprehensive longitudinal analysis.

Furthermore, the function addresses inconsistencies in income reporting across different years of the PSID. It ensures that when additional income variables are introduced in later years, they are appropriately standardized. This standardization process adjusts income figures to reflect a consistent metric, specifically aggregate taxable income, in contrast to the post-tax income reported in certain years.

As a consequence of these cleaning steps, the number of observations in each data file decreases. This decrease is more pronounced in earlier years (approximately a 30% reduction) and less substantial in later years (roughly a 5% decrease). This results in a standardized dataset with approximately 12,000 observations per year until 1982, followed by an increase to approximately 18,000 observations per year until 1990, and roughly 26,000 observations per year afterward.

3.2.2 Reshape

The reshaping function reorganizes cleaned PSID data files into an aggregated dataset structured around family units. We note that creating family units based on data from a single survey year presents challenges, primarily due to the possibility of children being part of more than one household or family. This situation often arises from the merging of families that necessitate modifications to the dataset.

To address these challenges and ensure the precise construction of family units, this study implements several steps:

1. A foundational family unit, benchmarked against the 1968 cohort (marking the inaugural year of the dataset), is established through the assignment of a distinct `FAMID` to a primary member within the family structure. For the purposes of this analysis, the primary member is determined to be the mother to accommodate any single-parent households.
2. After the initial identification process, income and age data for the same set of individuals are integrated. This integration uses subsequent PSID datasets and incorporates the data into the established family units, utilizing the unique `ID` markers generated during the data cleansing phase. Note that with each additional dataset integrated, a verification process is conducted to incorporate new family members into the existing family units. For instance, if a child's birth occurs five years after the 1968 baseline, the respective child's data is added to the pertinent family unit.
3. To expand the dataset, family entities identified in post-1968 surveys, which were not part of the original dataset, are also assigned a `FAMID` to represent a new family unit. This procedure also applies to families that have undergone splits.

The resulting output is a restructured income file that is organized by family units with longitudinal income data.

3.3 Pre-processing

Income data from PSID must be adjusted for inflation to accurately reflect differences in purchasing power across dataset years. This adjustment is critical for analyzing both relative and absolute mobility. For the relative mobility bootstrap analysis (Section 4.2), inflation adjustment ensures statistical significance by adjusting income in different years to the same reference year. Table 4.2 details the methodologies and applications of this inflation adjustment for the statistical significance tests of relative mobility. However, inflation adjustment is generally unnecessary for analyses of relative mobility, except for statistical significance tests. This is because relative mobility depends on quintile rankings within a cohort's income distribution, which are unaffected by the uniform scaling of income. Conversely, analyses of absolute mobility, which compare nominal income values across generations, require inflation adjustment.

Given that all data is in US dollars, to create an apples-to-apples comparison, we can adjust dollars for inflation using the Consumer Price Index (CPI). The CPI tracks the average change over time in the prices paid by urban consumers for a market basket of consumer goods and services. In this paper, historical CPI values are retrieved by referencing the Federal Reserve Economic Data (FRED) publicly available datasets (U.S. Bureau of Labor Statistics, 2023).

Relative Mobility

To adjust older nominal dollar values to current prices, we multiply the original prices by the current CPI (as of the end of 2023 in our case) and then normalize by the original CPI corresponding to the year of the original price. This is defined in Equation 3.1 below:

$$AV_C = \frac{A_O * CPI_C}{CPI_O} \tag{3.1}$$

Here, AV_C represents the adjusted value in current dollars (2023 in this case), A_O denotes the amount in original dollars (dependent on the year being adjusted), CPI_C is the current Consumer Price Index, and CPI_O signifies the original CPI (dependent on the year being adjusted).

Absolute Mobility

To adjust newer nominal dollar values to previous prices, we multiply the amount in current dollars by the old CPI we are benchmarking to and then normalize by the current CPI corresponding to the year of calculation. This is defined in Equation 3.2 below:

$$AV_O = \frac{A_C \times CPI_O}{CPI_C} \tag{3.2}$$

Here, AV_O represents the adjusted value back to the original reference year (1968 in the absolute mobility parametrization), A_C denotes the amount in current dollars (dependent on the year being adjusted), and CPI_C and CPI_O carry the same meanings as before.

Chapter 4

Relative Mobility Markov Models

This chapter presents an overview of the mathematical and computational processes involved in parameterizing relative mobility models. It then compares the transition matrices for the 1968 and 1997 birth cohorts, testing for statistical significance. The chapter concludes by providing time series outputs of the diagonal and outer edges of the transition matrices, which correspond to levels of income persistence and mobility at the extremes. The results indicate that intergenerational mobility, when modeled on a relative basis using income quintiles, has remained consistent between 1968 and 1997 birth cohorts.

4.1 Parameterization

4.1.1 Model

To parameterize the Markov transition matrix initially defined in Section 2.2 for relative mobility, the state space ($\mathcal{D}$) needs to be explicitly defined to create five distinct income subgroupings. The time period (t), transition probability matrix (P_c), and initial state distribution ($\mathbb{P}_{x_0}\cdot$) can be directly parameterized from the definitions in Section 2.2.

For relative mobility, the state space is defined as income quintiles, such that $\mathcal{D} = \{\sigma_1, \sigma_2, \ldots, \sigma_5\}$, where σ_1 is the lowest income quintile, and σ_5 is the highest income quintile. Thus, the Markov transition matrix (P_c) for relative mobility models a child's probability of belonging to a certain income quintile given information on the parental income quintile for a birth cohort c. Each entry $p_{i,j}$ represents the probability of transitioning from income quintile i in the parent generation to income quintile j in the child generation, where $i, j \in \{1, 2, \ldots, 5\}$. It is important to note that both parent and child income are measured at 26 years of age in this model.

4.1.2 Data

To formulate the census data for the relative mobility parameterization outlined above, several data processing steps are conducted. The table below synthesizes the various methods to create the transition probability matrix P_c for relative mobility.

Table 4.1: Relative Mobility Data Parameterization Overview

These are the three primary functions to create the Markov transition matrix for each child birth cohort. Note that the pivotal assumption is that transition matrices are made from parent-child income data when they were both 26 years old.

Function	**Description**
Quintile Ranks	Calculates quintile ranks for each data entry based on individual income. Note that quintiles are computed by comparing individuals of the same age within a birth cohort's dataset.
Family Linking	Links quintile information for parents to their respective children's records using the family unit structures created in Section 3.2.2. For children who have quintile information available for both parents, the parent with the higher quintile rank is selected.
Matrix Calculation	Computes the transition probability estimates for each element of the Markov transition matrices using the matched parent-child quintile data over the entire distribution of data for each child birth cohort.

The resulting output given the datasets available is an array containing a P_c

transition matrix corresponding to a birth cohort $c \in \{1968, 1969, \ldots, 1997\}$.

4.2 Results

The first (P_{1968}) and last (P_{1997}) transition matrices are shown below, spanning 29 birth years apart:

$$P_{1968} = \begin{bmatrix} 0.300 & 0.218 & 0.189 & 0.214 & 0.079 \\ 0.189 & 0.209 & 0.237 & 0.218 & 0.146 \\ 0.138 & 0.190 & 0.197 & 0.272 & 0.203 \\ 0.081 & 0.144 & 0.202 & 0.263 & 0.309 \\ 0.099 & 0.096 & 0.171 & 0.222 & 0.412 \end{bmatrix} \tag{4.1}$$

$$P_{1997} = \begin{bmatrix} 0.313 & 0.190 & 0.197 & 0.215 & 0.084 \\ 0.170 & 0.200 & 0.249 & 0.231 & 0.150 \\ 0.140 & 0.202 & 0.188 & 0.270 & 0.200 \\ 0.087 & 0.142 & 0.207 & 0.261 & 0.302 \\ 0.088 & 0.128 & 0.159 & 0.243 & 0.382 \end{bmatrix} \tag{4.2}$$

Both P_{1968} and P_{1997} transition matrices suggest intergenerational income persistence, with partial diagonal dominance indicating a relative "stickiness" of income quintile status from parents to children. In both matrices, the highest probabilities are $p_{1,1}$ and $p_{5,5}$, suggesting a high likelihood of income persistence at the extreme quintiles. The lowest probabilities are $p_{1,5}$ and $p_{5,1}$, indicating limited mobility at the extremes of the income distribution.

For transition probabilities off the diagonal, both matrices generally show higher probabilities adjacent to the diagonal elements than farther away within the same row. This indicates that when mobility does occur, it is more likely to involve moving to a neighboring quintile rather than a significant jump in relative income standing.

4.2.1 Delta Matrix

To analyze how intergenerational income mobility patterns have changed over the dataset, we can construct a delta matrix (Δ_P) conceptualized in Section 2.4 by subtracting the transition probability matrix of an earlier birth cohort from that of a later birth cohort. Here, Δ_P is explicitly defined as:

$$\Delta_P = P_{1997} - P_{1968} = \begin{bmatrix} 0.013 & -0.028 & 0.008 & 0.001 & 0.005 \\ -0.019 & -0.009 & 0.012 & 0.013 & 0.004 \\ 0.002 & 0.012 & -0.009 & -0.002 & -0.003 \\ 0.006 & -0.002 & 0.005 & -0.002 & -0.007 \\ -0.011 & 0.032 & -0.012 & 0.021 & -0.030 \end{bmatrix} \tag{4.3}$$

Given that P_{1997} and P_{1968} have the same dimensionality, we can compute an element-wise division of matrix elements to observe the percentage differences relative to P_{1968}, for illustrative purposes. We define matrix C as the full matrix of these element-wise divisions in percentages as $\left[c_{ij} = \left(\frac{a_{ij}}{b_{ij}} - 1\right) \times 100\right]$, where c_{ij} is the comparison element in percent, a_{ij} is the element in P_{1997}, and b_{ij} is the corresponding element in P_{1968}. The computed matrix C is shown below:

$$C = \begin{bmatrix} 4.3 & -12.8 & 4.2 & 0.5 & 6.3 \\ -10.1 & -4.3 & 5.1 & 6.0 & 2.7 \\ 1.4 & 6.3 & -4.6 & -0.7 & -1.5 \\ 7.4 & -1.4 & 2.5 & -0.8 & -2.3 \\ -11.1 & 33.3 & -7.0 & 9.5 & -7.3 \end{bmatrix} \tag{4.4}$$

An element-wise comparison of the P_{1968} and P_{1997} transition matrices reveals that the probabilities of transitioning between different income quintiles across parent-child generations are quite similar. Qualitatively, Δ_P suggests that the broad dynamics of intergenerational income mobility have not undergone substantial changes between the 1968 and 1997 birth cohorts.

Transition Probability Statistical Significance

To assess whether the observed differences in the probabilities of transitioning from parent to child income quintiles between P_{1968} and P_{1997} are statistically significant, we can employ a bootstrap test. Traditional statistical tests make assumptions about the underlying data distribution, which may not hold for income mobility data where the income distribution is typically right-skewed and variable across years. The table below describes the steps involved in using bootstrapping to test whether the differences in transition probabilities between P_{1968} and P_{1997} are statistically significant:

Table 4.2: Bootstrap Process Overview for Statistical Significance

These are the primary steps to conduct a bootstrapping analysis on the income dataset. Note that the process is repeated 5,000 times ($B = 5,000$) to create a distribution for each element of Δ_P, where each element represents a specific transition probability.

Function	**Description**
Income Adjustment	Scales the income values for both chosen birth cohort years to be equivalent to 2023 income levels (relative mobility analysis requires income to be scaled to a common reference year for quintile calculation). The inflation adjustment process for relative mobility is described by Equation 3.1.
Random Selection	Resamples the data for both chosen birth cohorts (with replacement) and maintains the original sample size. Note for each birth year, data is resampled from the combined dataset irrespective of birth year to test if segmenting by birth year is statistically significant from splitting the data randomly.
Matrix Calculation	Computes income quintile ranks for the two newly generated random samples (one for each birth cohort) using matched parent-child data. Then estimates transition probabilities for each element of the transition matrix, as illustrated in Table 4.1.2. Note the process is carried out separately for the two random samples representing the different birth cohorts.
Calculating Test Statistic	Evaluates the Δ matrix, which measures the difference in transition probability matrices between two birth cohorts. Here, the two transition probability matrices are created through bootstrapping and under the assumption that the birth cohorts are identical in terms of outcomes for children.

Here, the null hypothesis (H_0) posits no variation in transition probabilities between the 1968 and 1997 birth cohorts, while the alternative hypothesis (H_1) asserts such a difference exists. Note that several hypothesis tests are conducted concurrently for each element in the delta matrix.

Table 4.3: Summary of Statistical Significance of Delta Matrix Elements

All observed differences in transition probabilities in Δ_P defined in Equation 4.3 are not statistically significant at $\alpha = 0.05$ when resampling the combined dataset irrespective of birth year ($B = 5,000$). Note that the P value is calculated as the proportion of generated results that are as extreme or more extreme than the observed value.

Transition Probability	Observed Value	P value
$p_{1,1}$	0.013	0.252
$p_{1,2}$	−0.028	0.141
$p_{1,3}$	0.008	0.384
$p_{1,4}$	0.001	0.894
$p_{1,5}$	0.005	0.521
$p_{2,1}$	−0.019	0.167
$p_{2,2}$	−0.009	0.313
$p_{2,3}$	0.012	0.234
$p_{2,4}$	0.013	0.224
$p_{2,5}$	0.004	0.573
$p_{3,1}$	0.002	0.761
$p_{3,2}$	0.012	0.234
$p_{3,3}$	−0.009	0.294
$p_{3,4}$	−0.002	0.775
$p_{3,5}$	−0.003	0.640
$p_{4,1}$	0.006	0.457
$p_{4,2}$	−0.002	0.759
$p_{4,3}$	0.005	0.477
$p_{4,4}$	−0.002	0.756
$p_{4,5}$	−0.007	0.372
$p_{5,1}$	−0.011	0.296
$p_{5,2}$	0.032	0.172
$p_{5,3}$	−0.012	0.276
$p_{5,4}$	0.021	0.182
$p_{5,5}$	−0.030	0.136

From the bootstrap analysis conducted in Table 4.3, none of the observed differences in transition probabilities are statistically significant at the $\alpha = 0.05$ level, which fails to reject H_0. However, certain transition probabilities, specifically $p_{1,2}$, $p_{2,1}$, and $p_{5,5}$, with P values of 0.141, 0.167, and 0.136, respectively are statistically significant at the $\alpha = 0.15$ level. The lower P values suggest that the corresponding observed differences in transition probabilities have the strongest evidence towards being different between the two cohorts, although not statistically significant at typical levels of 5%. Notably, $p_{5,5}$ with an observed difference of 0.03 and a P value of 0.136 is the most interesting, as it represents a decrease in income persistence for children of the highest income quintile between the 1968 and 1997 birth cohorts.

Nevertheless, the bootstrapping results indicate that the observed variations in transition probabilities, as represented by the Δ_P matrix, can be attributed to random fluctuations at the 5% significance level. Even when considering a significance level of $\alpha = 0.15$, there are minimal discernible changes in transition probabilities. This supports the initial observation from the element-wise comparison, suggesting that levels of intergenerational income mobility have remained consistent between the 1968 and 1997 birth cohorts.

Delta Matrix Statistical Significance

Although individual transition probabilities are not statistically significant, we can repeat the bootstrapping procedure outlined in Table 4.2 to determine whether in aggregate the observed differences in the transition matrices P_{1968} and P_{1997} are statistically significant. Here, instead of comparing individual elements of the transition probability matrix as before, we compare the Frobenius norm of the observed Δ_P matrix with the Frobenius norms of bootstrapped Δ matrices that are randomly sampled from the combined dataset regardless of birth year. This approach emphasizes the total magnitude of change in the transition matrices rather than on individual probability shifts.

The Frobenius norm is a matrix norm that represents the square root of the sum of the absolute squares of the matrix elements. It provides a single scalar value that captures the overall magnitude of the matrix. Notably due to the squaring of each element before summing, the Frobenius norm is always positive regardless of the individual signs of the matrix elements. The Frobenius norm for the observed Δ_P in Equation 4.3 is defined below:

$$\|\Delta_P\|_F = \sqrt{\sum_{i=1}^{5}\sum_{j=1}^{5} |\delta_{ij}|^2} = 0.068 \tag{4.5}$$

Here, H_0 posits that there is no aggregate difference in the transition probability matrices P_{1968} and P_{1997}, whereas H_1 suggests that an aggregate difference exists between the transition probability matrices.

Table 4.4: Summary of Statistical Significance of Aggregate Delta Matrix
The observed Frobenius norm of Δ_P defined in Equation 4.3 is not statistically significant at $\alpha = 0.05$ when resampling the combined dataset irrespective of birth year ($B = 5,000$). Note that the P value is calculated as the proportion of generated results that are as extreme or more extreme than the observed value.

Observed Value	***P* value**
0.068	0.112

The bootstrap analysis conducted in Table 4.4 on the aggregate observed differences in transition matrices P_{1968} and P_{1997} indicates that the aggregate observed differences are not statistically significant at the $\alpha = 0.05$ level, which fails to reject the null hypothesis H_0. Even though the P value of the aggregate bootstrap analysis is lower than all the element-wise analyses, a similar result of intergenerational income mobility remaining stable when comparing the 1968 and 1997 birth cohorts can be inferred.

However, the observed differences in transition matrices at an aggregate level are significant at the 15% level. Yet, this hypothetical statistical significance provides minimal qualitative insight. The Frobenius norm, being inherently positive and scalar by definition, obscures information about the directionality of mobility, as well as the specific transition probabilities that represent real changes in income mobility.

4.2.2 Time Series

The following figures present a temporal comparison of select transition probabilities from P_c where $c \in \{1968, 1969, \ldots, 1997\}$. As a reminder, the first income quintile corresponds to the lowest income subgroup, and the fifth quintile corresponds to the highest. Additionally, note that parent-child income is measured when both groups were 26 years old.

Figure 4.1: Probability of Child Income Quintile Persistence

The likelihood that a child remains in the same income quintile as their parents appears unchanged for children born between 1968 and 1997. Notably, the highest probabilities correspond to remaining in the highest or lowest quintiles.

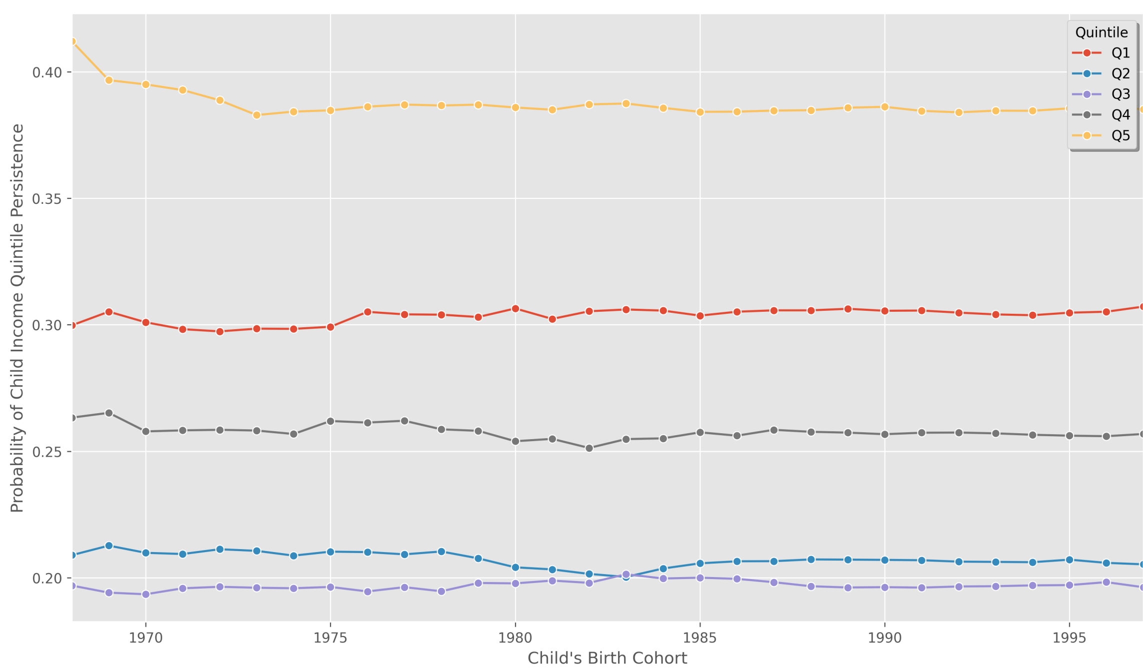

Figure 4.2: Probability of Child Quintile from Bottom Parental Quintile

Conditional on having a parent in the lowest-income quintile, a child is most likely to remain in the bottom quintile and least likely to reach the top quintile, with the remaining quintiles having a roughly equal probability. These probabilities remain steady across birth cohorts from 1968 to 1997.

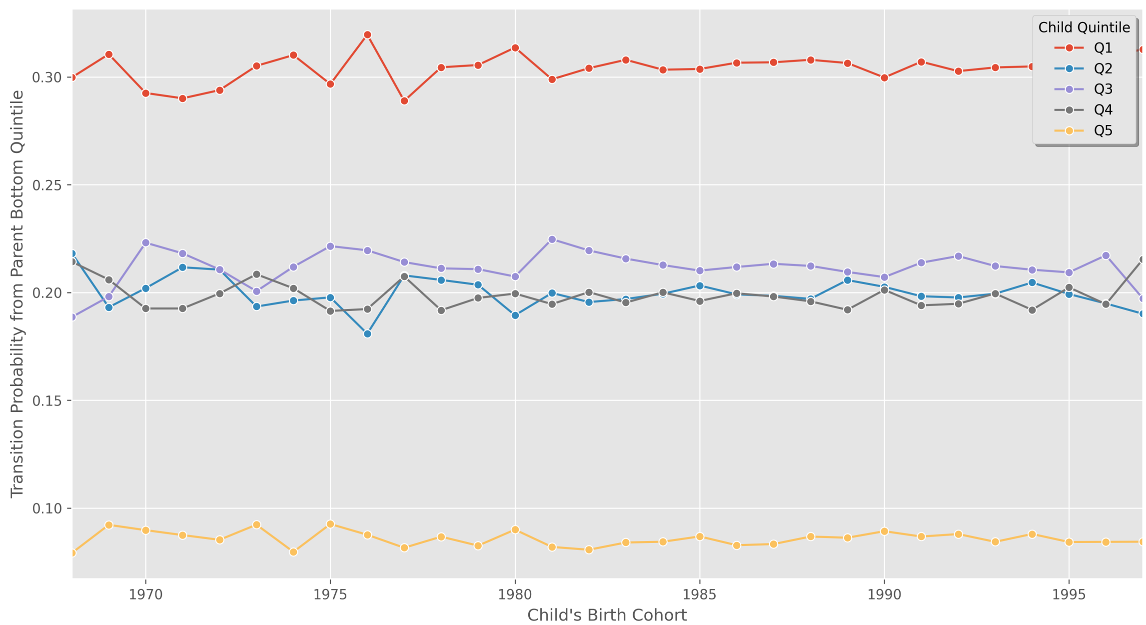

Figure 4.3: Probability of Child Quintile from Top Parental Quintile

Conditional on having a parent in the highest-income quintile, a child is most likely to remain in the top quintile, followed by the fourth, third, second, and lowest quintiles, respectively. These rates exhibit little trend across birth cohorts.

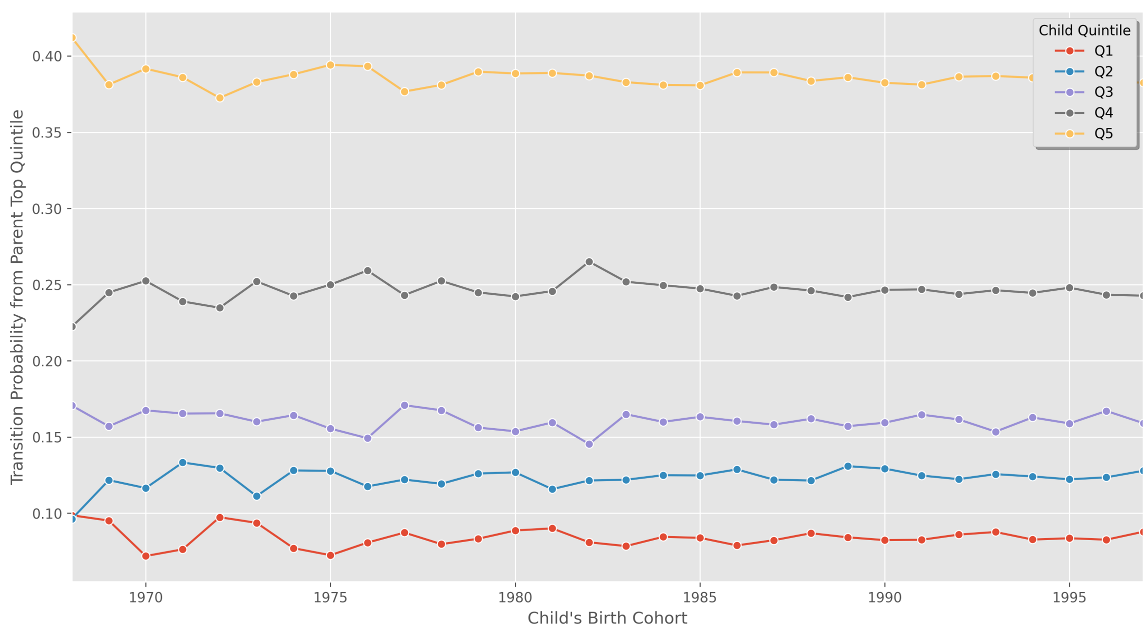

Figure 4.4: Probability of Child Reaching Bottom Quintile

The probability of a child belonging to the bottom quintile decreases as their parental income quintile increases. The likelihoods remained consistent between the 1968 and 1997 birth cohorts.

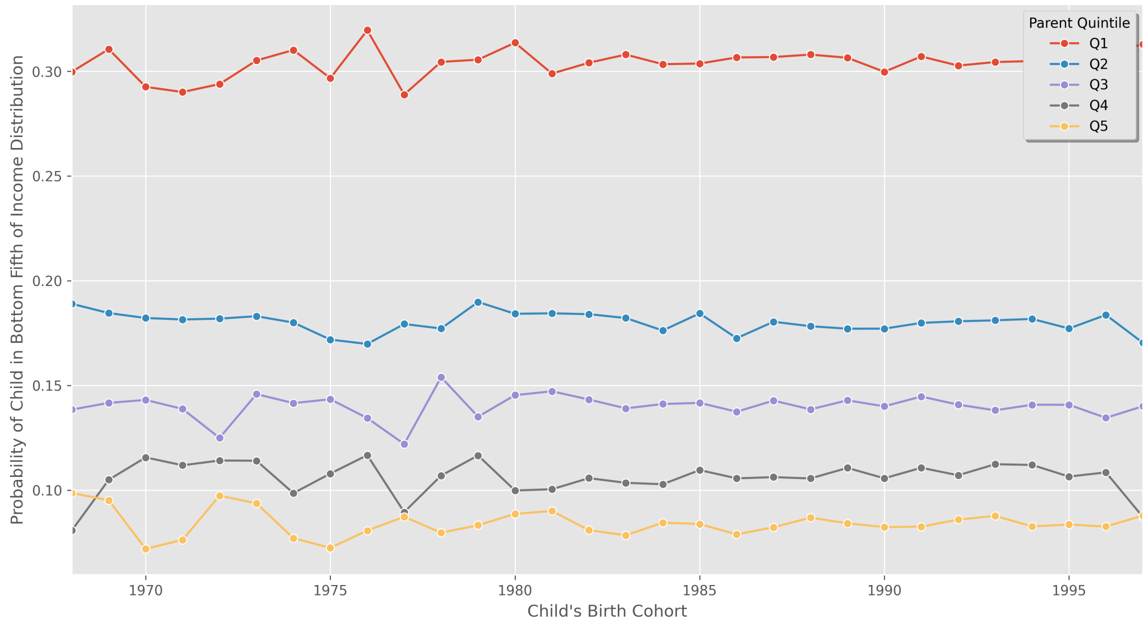

Figure 4.5: Probability of Child Reaching Top Quintile

The probability of a child belonging to the top quintile increases as their parental income quintile increases. The likelihoods remained consistent between the 1968 and 1997 birth cohorts.

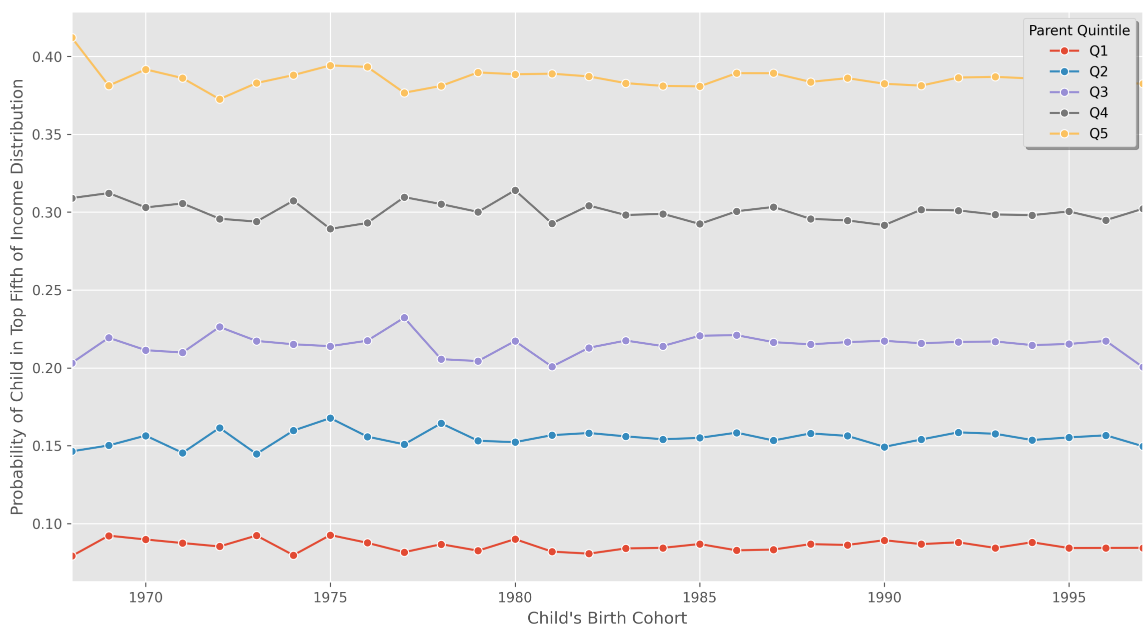

The time series plots of specific transition probabilities within Markov transition matrices P_c, where $c \in \{1968, 1969, \ldots, 1997\}$ suggest that intergenerational income mobility on a relative basis remained stable between 1968 and 1997 birth cohorts at the income age of 26. Consistently, higher variance in probabilities appears before 1980, which can likely be explained by the smaller sample size pre-1980 rather than an underlying year-to-year variation. Ultimately, the time series analysis of specific transition probabilities supports the conclusions of the delta matrix analysis conducted on the endpoints of the dataset (P_{1968} and P_{1997}), indicating that levels of intergenerational mobility on a relative basis are consistent across the 29-year period. It is important to note, however, that this stability pertains specifically to this paper's parametrization of intergenerational mobility defined by relative income quintiles.

Chapter 5

Absolute Mobility Markov Models

This chapter provides an overview of the mathematical and computational processes involved in parametrizing absolute mobility models. It then compares transition matrices corresponding to the 1968 and 1997 birth cohorts and tests for statistical significance. The chapter concludes with time series outputs for the diagonal and outer edge entries of the transition matrix, as well as the stationary distribution and mixing times. Results indicate that intergenerational mobility modeled by absolute mobility in predefined income buckets has changed over time. For the birth cohorts from 1968 to 1997, there is a discernible trend showing that children are increasingly likely to move into higher income buckets compared to their chances of remaining in middle-income buckets, while transitions to the lowest income bucket remain relatively unchanged.

5.1 Parameterization

5.1.1 Model

To parametrize the Markov transition matrix defined in Section 2.2 for absolute mobility, the state space ($\mathcal{D}$) needs to be explicitly defined in five distinct income subgroups.

Similar to relative mobility, the time period (t), transition probability matrix (P_c), and initial state distribution ($\mathbb{P}_{x_0}\cdot$) can be directly parameterized from the definitions in Section 2.2.

For absolute mobility, the state space is defined as five income buckets such that $\mathcal{D} = \{\sigma_1, \sigma_2, \ldots, \sigma_5\}$, where σ_1 is the lowest income bucket, and σ_5 is the highest income bucket. The income buckets are defined by the quintiles of the first year in the dataset, which is 1968 in our case. Hence, the state space for absolute mobility can be explicitly defined in 1968 dollar quintiles as follows:

$$\mathcal{D} = \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \sigma_3 \\ \sigma_4 \\ \sigma_5 \end{bmatrix} = \begin{bmatrix} [0,\ \ 780) \\ [780,\ \ 2{,}700) \\ [2{,}700,\ \ 4{,}800) \\ [4{,}800,\ \ 7{,}000) \\ [7{,}000,\ \ \infty) \end{bmatrix} \tag{5.1}$$

In this study, the number of individuals representing the parent generation of 1968 remains fixed across income buckets by definition. However, subsequent birth cohorts exhibit varying population sizes within each bucket, reflecting shifts in the income distribution specific to their respective years.

The decision to base the income bucket cutoffs on the 1968 income quintiles is a significant assumption. This approach enables tracking changes in transition probabilities between fixed parent-child income buckets over time, revealing shifts in absolute mobility. Since the buckets remain constant, it allows for direct observation of how income transitions have changed compared to the 1968 quintile cutoffs. Furthermore, it quantifies the proportion of children experiencing upward or downward income mobility in absolute dollar terms, compared to their parents' income bucket.

The Markov transition matrix P_c models the absolute income mobility for a birth cohort c. It represents a child's probability of belonging to an absolute dollar income bucket, given information about the parental income bucket. Specifically, each entry

$p_{i,j}$ in the matrix represents the probability of transitioning from income bucket i in the parent generation to income bucket j in the child generation, where $i, j \in \{1, 2, \ldots, 5\}$. It is important to note that both parent and child incomes are measured in absolute dollar terms at age 26.

5.1.2 Data

To formulate the data for the absolute mobility parameterization outlined above, several data processing steps are conducted similarly to the relative mobility steps outlined in Table 4.1. The difference lies in the treatment of income data. Instead of calculating quintile ranks, all income data throughout the dataset is inflation-adjusted to 1968 dollars using the process outlined in Equation 3.2, and then placed into the income buckets defined in Equation 5.1.

For children with income bucket information on both parents, the parent with the higher income bucket is selected. Finally, using the data on paired parent-child income buckets for each birth cohort, transition matrices are calculated. The output, given the available datasets, is an array containing a P_c transition matrix corresponding to each birth cohort $c \in \{1968, 1969, \ldots, 1997\}$.

5.2 Results

The first (P_{1968}) and last (P_{1997}) absolute transition matrices are shown below, spanning 29 birth years apart:

$$P_{1968} = \begin{bmatrix} 0.173 & 0.363 & 0.286 & 0.133 & 0.045 \\ 0.103 & 0.345 & 0.300 & 0.189 & 0.063 \\ 0.114 & 0.262 & 0.355 & 0.152 & 0.117 \\ 0.060 & 0.253 & 0.295 & 0.233 & 0.159 \\ 0.078 & 0.189 & 0.240 & 0.253 & 0.240 \end{bmatrix} \tag{5.2}$$

$$P_{1997} = \begin{bmatrix} 0.145 & 0.261 & 0.304 & 0.184 & 0.106 \\ 0.111 & 0.242 & 0.358 & 0.157 & 0.131 \\ 0.067 & 0.194 & 0.254 & 0.261 & 0.223 \\ 0.041 & 0.187 & 0.283 & 0.237 & 0.251 \\ 0.035 & 0.178 & 0.188 & 0.223 & 0.376 \end{bmatrix} \tag{5.3}$$

The P_{1968} transition matrix reveals distinct mobility characteristics. The higher probabilities along the diagonal suggest a significant tendency for individuals to remain within their parents' income bucket, especially at the extremes of the income spectrum. There is a notable "stickiness" at $p_{1,1}$ and $p_{5,5}$, with a 17% probability of those in the lowest income bucket remaining there, and a 24% chance for those in the highest bucket to stay put. In contrast, $p_{1,5}$ with an 11% probability indicates a slim chance of moving from the lowest parental income bucket to the highest income bucket, and $p_{5,1}$ indicates a 4% probability of a child with parents in the highest income bucket falling to the lowest income bucket. The middle income buckets show more fluidity, with the second and third rows depicting a higher likelihood of movement into adjacent income categories rather than remaining stationary. This suggests a somewhat dynamic middle class with the potential for both upward and downward mobility.

In contrast, the P_{1997} transition matrix indicates shifts in mobility dynamics given the predefined income buckets in Equation 5.1. There appears to be a greater bifurcation at the extremes of the distribution, with $p_{1,1}$ decreasing, pointing to an increase in upward mobility in the lowest class. However, $p_{5,5}$ shows a marked increase in persistence, with a probability of 38%, which may reflect further consolidation of wealth. Furthermore, P_{1997} exhibits higher diagonal elements for the upper income buckets compared to P_{1968}, suggesting even greater "stickiness" or entrapment at the top of the income distribution for the later cohort. Conversely, the diagonal elements for lower income buckets are smaller in 1997, suggesting improved absolute mobility out of poverty or low-income levels, relative to the earlier cohort. However, in both P_{1968} and P_{1997}, there are generally higher probabilities adjacent to the diagonal in the same row, indicating that when mobility occurs, it is likely to be to neighboring income buckets.

Additionally, there appears to be a higher magnitude of transition probabilities below the diagonal for P_{1968}, representing downward mobility, compared to the corresponding upward mobility elements in P_{1997}. This indicates that for the 1968 birth cohort there was a greater chance of declining or steady income levels. In comparison, the 1997 birth cohort exhibits a greater chance of income outperformance relative to the prior generation on an absolute dollar basis.

5.2.1 Delta Matrix

Similar to the relative mobility analysis, a delta matrix (Δ_P) conceptualized in Section 2.4 can be constructed to directly compare transition probabilities on a per element basis. Here, Δ_P is explicitly defined as:

$$\Delta_P = P_{1997} - P_{1968} = \begin{bmatrix} -0.028 & -0.102 & 0.018 & 0.051 & 0.061 \\ 0.008 & -0.103 & 0.058 & -0.032 & 0.068 \\ -0.047 & -0.068 & -0.101 & 0.109 & 0.106 \\ -0.019 & -0.066 & -0.012 & 0.004 & 0.092 \\ -0.043 & -0.011 & -0.052 & -0.030 & 0.136 \end{bmatrix} \tag{5.4}$$

The Δ_P can also be expressed as the percentage difference between the transition probabilities P_{1997} and P_{1968}, relative to P_{1968}, for illustrative purposes. We define C as the matrix of element-wise divisions corresponding to this percentage difference (see the mathematical definition for C in Section 4.2.1):

$$C = \begin{bmatrix} -16.2 & -28.1 & 6.3 & 38.3 & 135.6 \\ 7.8 & -29.9 & 19.3 & -16.9 & 107.9 \\ -41.2 & -26.0 & -28.5 & 71.7 & 90.6 \\ -31.7 & -26.1 & -4.1 & 1.7 & 57.9 \\ -55.1 & -5.8 & -21.7 & -11.9 & 56.7 \end{bmatrix} \tag{5.5}$$

An element-wise comparison between P_{1968} and P_{1997} reveals a shift away from downward mobility and towards greater upward mobility over the 29-year period. This trend is evidenced by negative values below the diagonal entries, indicating the 1968 birth cohort was more likely to transition towards lower income buckets than their parents. Conversely, positive values above the diagonal suggest that the 1997 birth cohort was more likely to outperform their parents' income on an absolute dollar basis. Overall, Δ_P suggests that intergenerational economic dynamics may have undergone substantial changes between the 1968 and 1997 birth cohorts.

Transition Probability Statistical Significance

To assess the statistical significance of observed changes between transition matrices P_{1968} and P_{1997}, we can apply the same bootstrap methodology used for determining statistically significant differences in the relative mobility parameterization. The precise process is described in Table 4.2. However, since all income is already deflated to 1968 dollars, the income adjustment step is unnecessary. Similar to relative mobility, the null hypothesis (H_0) posits no difference in transition probabilities between children born in 1968 and 1997. The alternative hypothesis (H_1) asserts that a difference exists. Note that we conduct multiple hypothesis tests simultaneously, one for each element of the transition matrix.

Table 5.1: Summary of Statistical Significance of Delta Matrix Elements

Of the observed differences in transition probabilities in Δ_P defined in Equation 5.4, 11 are statistically significant at $\alpha = 0.05$ (*). Several of these differences are also significant at $\alpha = 0.01$ (**) and $\alpha = 0.001$ (***) levels when resampling the combined dataset irrespective of birth year ($B = 5,000$). Note that the P value is calculated as the proportion of generated results that are as extreme or more extreme than the observed value.

Transition Probability	Observed Value	P value
$p_{1,1}$	−0.028	0.499
$p_{1,2}$	−0.102	0.023*
$p_{1,3}$	0.018	0.694
$p_{1,4}$	0.051	0.141
$p_{1,5}$	0.061	0.023*
$p_{2,1}$	0.008	0.672
$p_{2,2}$	−0.103	0.000***
$p_{2,3}$	0.058	0.072
$p_{2,4}$	−0.032	0.202
$p_{2,5}$	0.068	0.002**
$p_{3,1}$	−0.047	0.016*
$p_{3,2}$	−0.068	0.016*
$p_{3,3}$	−0.101	0.003**
$p_{3,4}$	0.109	0.002**
$p_{3,5}$	0.106	0.000***
$p_{4,1}$	−0.019	0.414
$p_{4,2}$	−0.066	0.054
$p_{4,3}$	−0.012	0.775
$p_{4,4}$	0.004	0.926
$p_{4,5}$	0.092	0.038*
$p_{5,1}$	−0.043	0.056
$p_{5,2}$	−0.011	0.714
$p_{5,3}$	−0.052	0.127
$p_{5,4}$	−0.030	0.410
$p_{5,5}$	0.136	0.000***

The bootstrap analysis conducted in Table 5.1 reveals several statistically significant transition probabilities. Notably, $p_{2,2}$, $p_{3,5}$, and $p_{5,5}$ are significant at the $\alpha = 0.001$ level. Consequently, the null hypothesis can be rejected in favor of the alternative hypothesis for particular elements in the Δ_P matrix. The probabilities exhibiting the most substantial observed differences are those in the last column, corresponding to children being in the highest income bucket, and the third row, representing child outcomes contingent on parents belonging in the middle income bucket.

The observed differences in the last column suggest an increase in the probability of ending up in the highest income bucket for all parental income buckets between the P_{1997} and P_{1968} cohorts. This broadly corresponds to children outperforming their parents in absolute dollars or performing in line if parents were already in the highest bucket. Conversely, the observed differences within the third row ($p_{3,1}$, $p_{3,2}$, and $p_{3,3}$) indicate a decrease in the probability of a child remaining in the same or lower income bucket given a parent in the third income bucket, with observed values of -0.047, -0.068, and -0.101, respectively. The resulting gains for $p_{3,4}$ and $p_{3,5}$ within the same row suggest that, given parents in the third income bucket, the probability of outperforming parents in absolute dollars has increased between the 1968 and 1997 birth cohorts.

Relaxing the $\alpha = 0.05$ level reveals that the observed differences for $p_{2,3}$, $p_{4,2}$, $p_{5,1}$, and $p_{5,3}$ are close to statistical significance, with P values of 0.072, 0.054, 0.056, and 0.127, respectively. Notably, $p_{5,1}$ exhibits an observed difference of -0.043 between P_{1997} and P_{1968}, suggesting greater income persistence as the probability of joining the first income bucket with the highest parental income bucket has decreased. Broadly, the statistical significance of the last column and the third row supports the element-wise comparison, indicating a broader shift towards greater intergenerational upward mobility on an absolute dollar basis between the 1968 and 1997 birth cohorts.

Delta Matrix Statistical Significance

To assess the statistical significance of the observed differences in transition matrices P_{1997} and P_{1968} for absolute mobility, a bootstrapping process can be conducted on the delta matrix for absolute mobility, akin to the relative mobility analysis. Instead of comparing individual elements, the Frobenius norm of the observed Δ_P is compared with the distribution of Δ_P under the null hypothesis. Here, H_0 posits that there is no aggregate difference in the transition probability matrices P_{1997} and P_{1968}, whereas H_1 proposes than an aggregate difference exists between the transition probability matrices.

The Frobenius norm, defined in Equation 4.5, is a matrix norm that can be calculated for the absolute mobility delta matrix as follows:

$$\|\Delta_P\|_F = \sqrt{\sum_{i=1}^{5}\sum_{j=1}^{5}|\delta_{ij}|^2} = 0.339 \tag{5.6}$$

Table 5.2: Summary of Statistical Significance of Aggregate Delta Matrix
The observed Frobenius norm of Δ_P defined in Equation 5.4 is statistically significant at $\alpha = 0.001$ (***) when resampling the combined dataset irrespective of birth year ($B = 5,000$). Note that the P value is calculated as the proportion of generated results that are as extreme or more extreme than the observed value.

Observed Value	***P* value**
0.339	0.000***

The bootstrap analysis (Table 5.2) reveals a statistically significant difference (at $\alpha = 0.001$) in aggregate transition probabilities (P_{1997} and P_{1968}). This supports the rejection of H_0 in favor of H_1, implying a difference in transition probabilities between the 1968 and 1997 birth cohorts under the absolute mobility model defined by Equation 5.1. However, the Frobenius norm cannot determine the direction of change in intergenerational mobility due to its scalar and positive properties.

5.2.2 Time Series

The following figures present a temporal comparison of select transition probabilities of P_c, where $c \in \{1968, 1969, \ldots, 1997\}$. All figures in this subsection display a two-year rolling average to account for high year-to-year variation. Note that parent-child income is recorded when both groups were 26 years old.

- Figure 5.1: Depicts the persistence of income buckets from parents to children.
- Figure 5.2: Illustrates the probability of a child belonging to a higher income bucket than their parents, contingent upon the parental income bucket.
- Figure 5.3/5.4: Charts the bottom and top rows of the transition matrix, respectively, corresponding to the probability of a child's income bucket given the bottom and top parental income buckets.
- Figure 5.5/5.6: Displays the bottom and top columns of the transition matrix, showing the probability of a child attaining the bottom and top income buckets, respectively, dependent on their parents' income bucket.

Figure 5.1: Probability of Child Income Bucket Persistence

Income persistence is similar for the lowest and highest income buckets for children born between 1968-1997. The middle income buckets see less persistence over time, but the highest earners increasingly remain in their bucket. Notably, children from the lowest income families tend to move upward given the low income persistence.

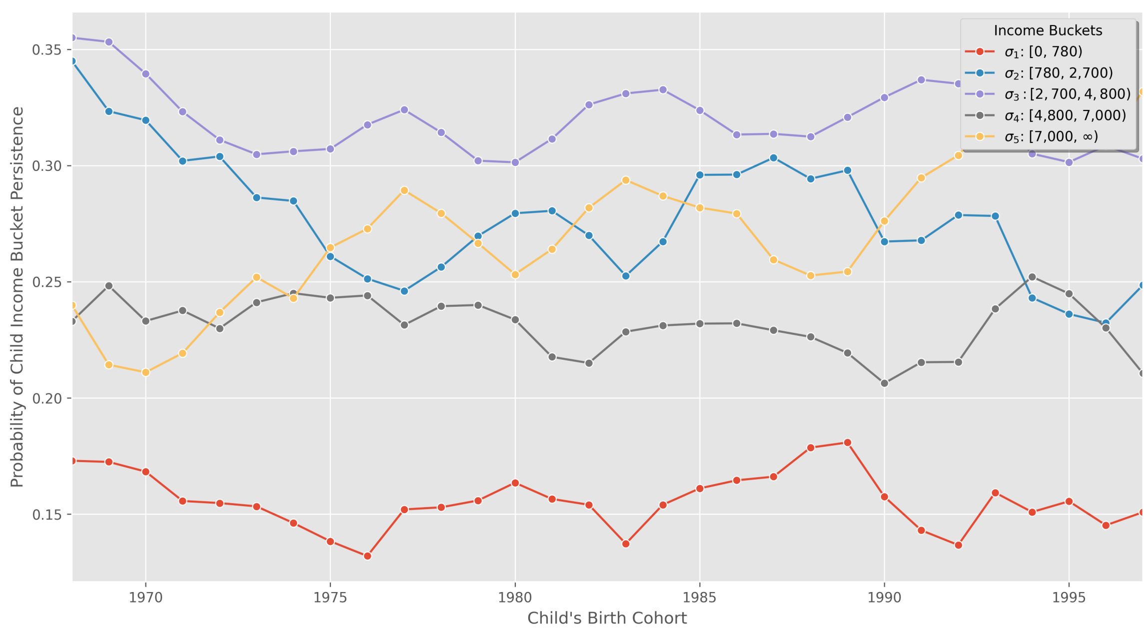

Figure 5.2: Probability of Child Outearning Parental Bucket

The probability of children attaining higher income buckets than their parents exhibits a slight upward trend from the 1968 to 1997 birth cohorts, indicating modestly greater absolute upward mobility across income buckets. Notably, these probabilities decline as parental income increases.

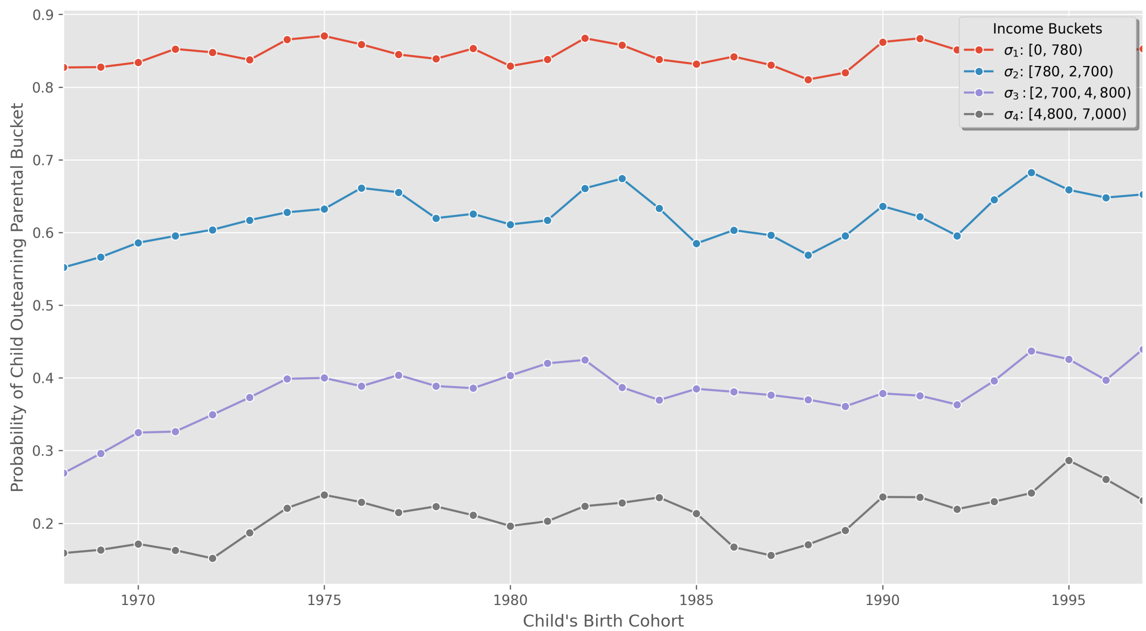

Figure 5.3: Probability of Child Bucket from Bottom Parental Bucket

Conditional on parents in the lowest income bucket, the highest probabilities are associated with the child attaining the second or third income buckets. Probabilities stayed relatively stable from the 1968 to 1997 birth cohorts, with a slight rise in the chance of entering the top income bucket.

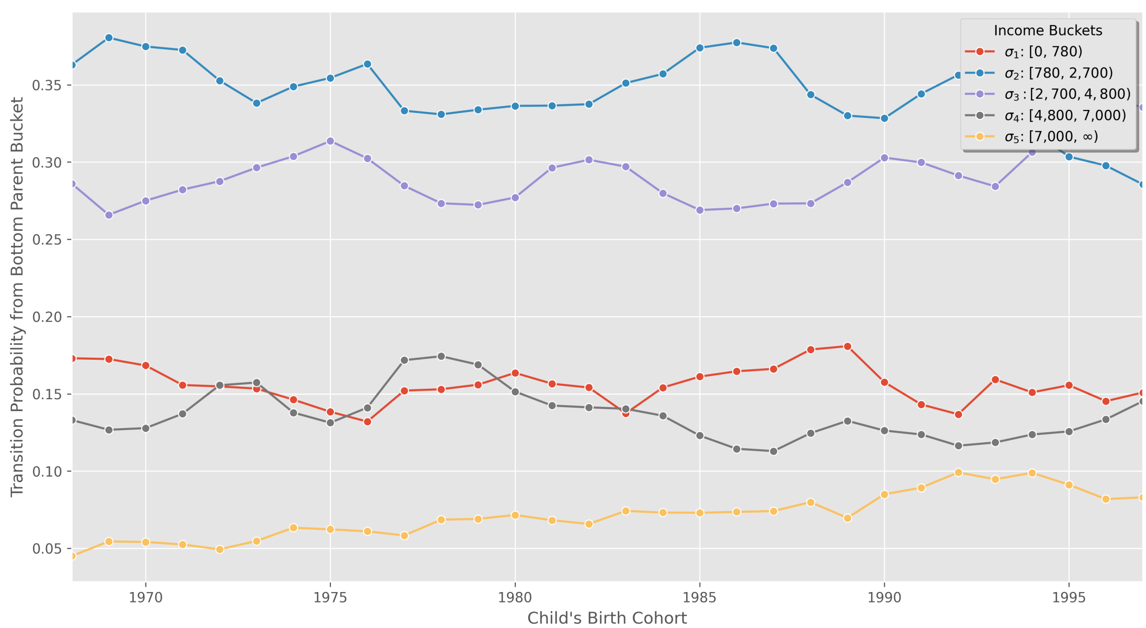

Figure 5.4: Probability of Child Bucket from Top Parental Bucket

Conditional on parents in the highest income bucket, the chances of remaining in the top bucket rose slightly, with equal chances of falling to middle buckets, and least likely to drop to the bottom. Transition probabilities, except to the top bucket, stayed consistent from 1968 to 1997 cohorts.

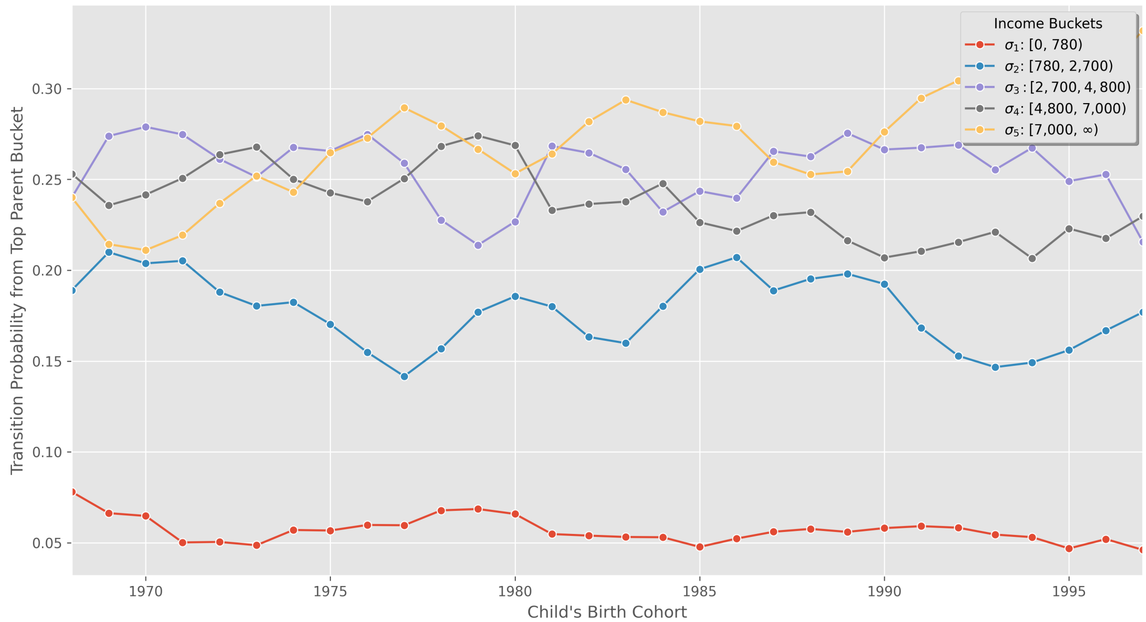

Figure 5.5: Probability of Child Reaching Bottom Bucket

The probability of a child belonging to the bottom bucket decreases as their parental bucket increases. The likelihoods exhibit a marginal trend between the 1968 and 1997 birth cohorts.

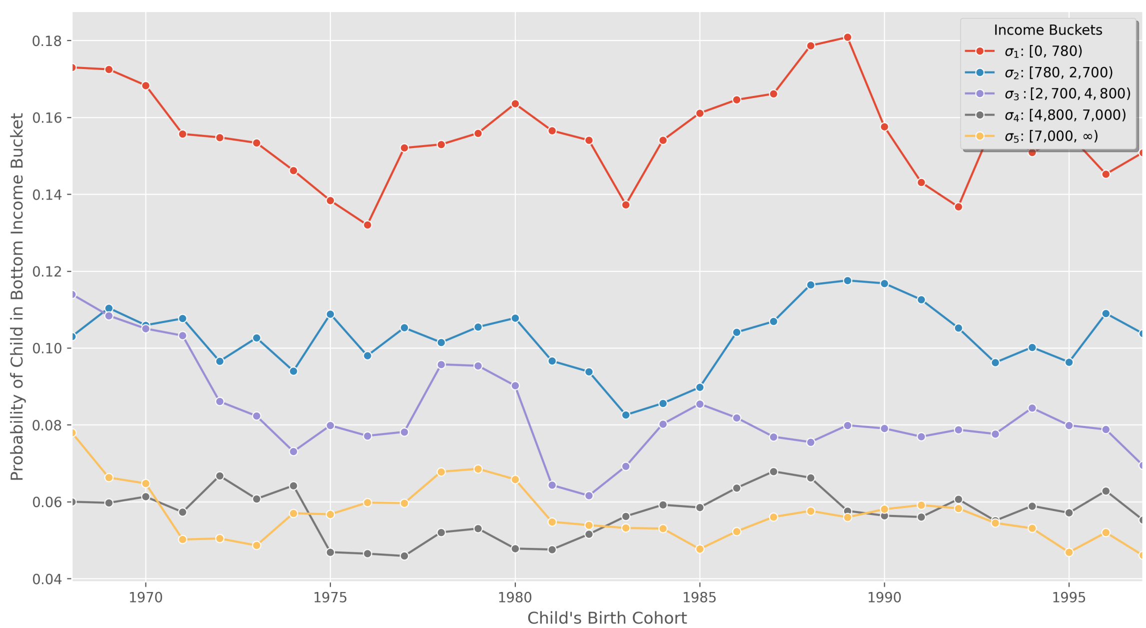

Figure 5.6: Probability of Child Reaching Top Bucket

The probability of a child belonging to the top bucket increases as their parent's bucket increases. All likelihoods trended upwards between the 1968 and 1997 birth cohorts, indicating greater upward absolute mobility irrespective of parental bucket.

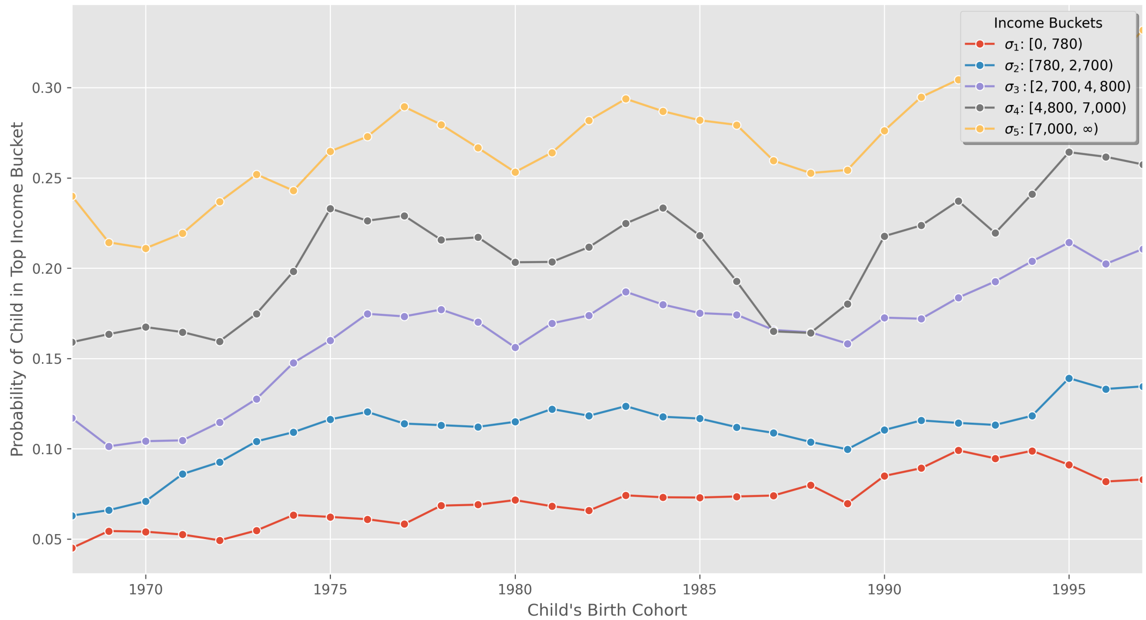

The time series plots of specific transition probabilities within the Markov transition matrix P_c, where $c \in \{1968, 1969, \ldots, 1997\}$ suggest that intergenerational income mobility on an absolute basis has skewed towards greater levels of upward mobility between 1968 and 1997 at the income age of 26. Figure 5.2 shows a gradual upward trend, indicating that children have a higher probability of belonging to an income bucket greater than their parents, suggesting increased upward mobility. Furthermore, Figure 5.6 indicates that the likelihood of a child reaching the top income bucket has increased steadily over time across all parental income buckets, suggesting that upward mobility has become a more likely outcome for children.

However, other transition probabilities have remained stable between the 1968 to 1997 birth cohorts, specifically relating to outcomes with parents in the lowest income bucket (Figure 5.3) and children ending up in the lowest income bucket (Figure 5.5). Given the gains over time in the probability of belonging to the higher income buckets, this implies (since rows sum to 1 in the transition matrix) that the likelihood of joining the second or third income bucket has likely slightly decreased in favor of a greater probability of joining the higher buckets. This suggests that while the difficulty of getting out of the lowest income bucket may remain similar between the 1968 and 1997 birth cohorts, there is a higher probability that if mobility does occur, it will be to join one of the buckets at the top of the income distribution.

Across all figures, there is a higher year-over-year variance, which likely corresponds to income being tracked for parent-child pairs at the age of 26. This variance is less prevalent in the relative mobility analysis, as the income subgrouping is based on quintiles that generally normalize for this effect across the population. Since the absolute mobility analysis is dependent on fixed absolute income buckets, this normalization effect does not occur, which creates greater variation year-over-year as the income distribution changes.

Holistically, the time series analysis of transition probabilities supports the trends observed in the delta matrix analysis conducted between P_{1968} and P_{1997}. It suggests that over time, there is a greater likelihood of children joining the highest income buckets; however, there are not significant changes in the likelihood of children ending up in the lowest income bucket. This represents a general shift in levels of intergenerational mobility measured on an absolute basis, defined specifically by Equation 5.1, across the 29-year period of the dataset.

5.2.3 Stationary Distribution & Mixing Times

As introduced in Section 2.3, the stationary distribution for absolute mobility represents the long-term equilibrium probability of an individual belonging to a certain income bucket set by Equation 5.1. Similarly, the mixing time defined in Section 2.3 models the time it takes for the Markov process to converge at the stationary distribution. The following figures present a temporal comparison of the implied stationary distribution and mixing times for the absolute mobility transition matrices.

Figure 5.7: Stationary Distribution Across Birth Cohorts

The stationary distribution remains relatively stable across the first, third, and fourth income buckets for the 1968 to 1997 birth cohorts. However, the distribution for the second and fifth income buckets switches over time, indicating a greater long-term likelihood of belonging to the highest absolute income bucket.

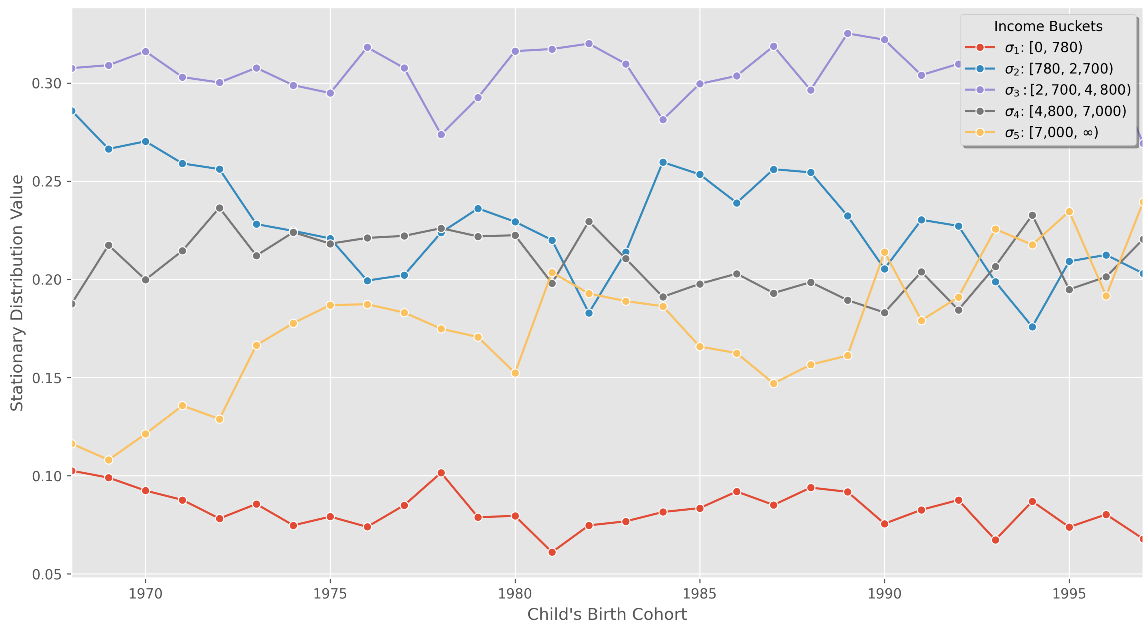

Figure 5.8: Mixing Times Across Birth Cohorts

The mixing times show little directional trend from the 1968 to 1997 birth cohorts.

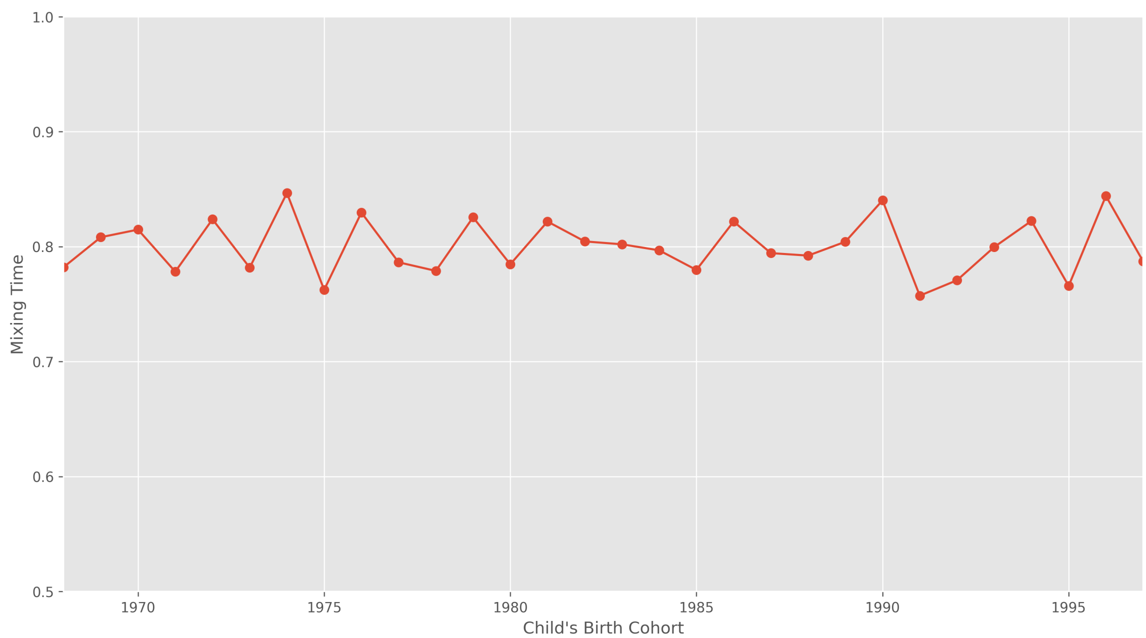

The time series plots of the stationary distribution exhibit patterns closely aligned with those depicted in Figure 5.1 which plots income persistence on absolute dollar terms. The overall trends in the stationary distribution suggest that, over the long term, the stable configuration is one in which more people are in the highest bucket over time between 1968 and 1997 birth cohorts. Nonetheless, a considerable proportion of children are anticipated to remain entrenched in lower income categories, reflecting the enduring nature of these segments when defined by absolute income levels.

Unfortunately, the variance observed in the mixing time of the stationary distribution is minimal, providing little qualitative insight into trends pertaining to intergenerational mobility. Remember that the magnitude of the mixing time holds no significant relevance given the mathematical formulation in Section 2.3, and that a relative increase in mixing time suggests greater levels of intergenerational income persistence.

Lastly, note that these observations regarding the stationary distribution and mixing time are specific to scenarios where absolute mobility is accurately characterized by the defined income buckets in Equation 5.1.

Chapter 6

Discussion & Conclusions

In this study, we calibrate a set of Markov transition matrices to model intergenerational mobility dynamics through the lenses of relative and absolute mobility. Relative mobility is parameterized by income quintiles recalculated for every birth cohort, while absolute mobility is defined by predefined income buckets that remain constant across birth cohorts.

Relative Mobility

Numerous studies have modeled relative mobility through log-log elasticities and rank-rank correlations. However, to our knowledge, none have parameterized and analyzed a series of Markov transition matrices over an extended time period. Notably, Chetty et al. (2014) formulates a Markov transition matrix but only presents findings related to the last column, representing the probability that a child reaches the top quintile given various parental income quintiles.

The results from this study's parameterization of relative mobility suggest no discernible trends in intergenerational mobility probabilities between the 1968 and 1997 birth cohorts based on time-series data. An element-wise comparison of P_{1968} and P_{1997}, corresponding to the first and last calibrated transition matrices derived from the dataset, suggests no observed differences in transition probabilities at the 5% statistical significance level. Even when relaxing the statistical significance levels to 15%, only a few transition probabilities exhibit evidence of differing between the two birth cohorts, and none of these differences were substantial.

The results from this study's parameterization reflect systematic inertia in socioeconomic mobility, where the quintile income ranking of parents significantly influences the quintile income of children, with these observations remaining consistent over the 29-year period of the dataset. This level of income persistence also coincides with a lack of substantial jumps in relative mobility, as transitions to neighboring quintiles are significantly more likely. Broadly, these results suggest that children have the same likelihood of upward mobility as their parents over the aggregated dataset.

This study's relative mobility results are holistically analogous to existing literature. Zimmerman (1992) models relative mobility via log-log elasticities and indicates an intergenerational earnings elasticity of 0.4, signaling a high correspondence between parent and children's incomes (elasticity of 1 represents perfect matching). This supports the higher transition probabilities found consistently across the diagonal entries of the transition matrices, with children most likely to remain within the same quintile as their parents. Furthermore, specific transition probabilities in this paper's parameterization are supported by the empirical values reported by Chetty et al. (2014). Specifically, the vector corresponding to transitions to the highest child income quintile from the set of parent quintile values in Chetty et al. (2014) is roughly the same as the normalized last column in our matrix formulation. Thus, Zimmerman (1992) and Chetty et al. (2014) generally support the conclusions from this paper's formulation that levels of intergenerational mobility have remained somewhat steady between the 1968 and 1997 birth cohorts.

Prior studies, including Lee and Solon (2009), suggest that intergenerational income mobility on a relative basis has remained stable even further back, including the 1950 and 1970 birth cohorts, based on a log-log elasticities approach measuring income at the age of 26. This paper's results, in conjunction with prior results, suggest that levels of intergenerational mobility may have remained consistent from the 1950 to 1997 birth cohorts given this paper's formulation of relative mobility.

The stability of intergenerational mobility rates in a relative mobility parametrization is surprising given higher rates of income inequality over the last several decades. The Gini coefficient measures the statistical dispersion of the income distribution, with the Census Bureau reporting an estimate of 0.481 for 2016, which has increased by about 20% from 1980 to 2016 (Habib and Perese, 2016). Note that the Gini coefficient is range-bounded from 0 to 1, with perfect equality corresponding to 0 and perfect inequality corresponding to 1. Krueger (2012) notes a "Great Gatsby Curve" where greater income inequality, measured by a country's higher Gini coefficient, coincides with lower economic mobility. Since the US Gini coefficient has increased, indicating higher income inequality, Krueger (2012) implies economic mobility levels should decrease. However, this paper's results suggest mobility levels have remained steady despite the jump in inequality. This suggests that while income quintile definitions have likely grown farther apart in absolute dollars due to increased inequality, the transition probabilities have not changed overall. The lack of significant observable changes in transition probabilities is likely due to the recalculation of quintiles across income data, which reduces noise and the effects of changes in the wider income distribution.

Absolute Mobility

Intergenerational absolute mobility has remained relatively unexplored compared to relative mobility due to challenges in obtaining suitable datasets. The seminal paper by Chetty et al. (2017) laid the foundation for absolute mobility research by plotting the proportion of children earning more than their parents. To our knowledge, no subsequent studies have parameterized a series of Markov transition matrices to conceptualize absolute mobility, where the state space is defined by predefined, constant income buckets.

This paper's particular parameterization of absolute mobility indicates that intergenerational mobility has shifted towards an increased probability of upward mobility. However, the data also suggests that the rates at which children fall into the lowest income bracket remain consistent between the birth cohorts of 1968 and 1997. Independent of the parental income category, results indicate a higher probability for children to ascend to an income bucket higher than that of their parents. Additionally, the probability of a child achieving the highest income category saw a progressive increase across successive birth cohorts, regardless of the parental income category. Conversely, the transition probabilities represented in the first column of the transition matrix, which estimated the likelihood of a child joining the first income bracket, showed no clear trend. This suggests that, due to the increased transition probabilities to higher income categories, the probability of a child advancing to the second or third income categories decreased correspondingly.

The trends in the time-series data are supported by comparing the P_{1968} and P_{1997} matrices, the first and last transition matrices from the dataset. This Δ_P comparison shows a shift toward upward mobility across income brackets. Specifically, the observed data in the last column, which indicates the likelihood of a child reaching the highest income bracket based on the parents' income, shows significant positive differences at the 5% level. The transition probabilities for P_{1997} relative to P_{1968} indicate a 136% increase in the chance of a child moving to the highest income bracket from the lowest parental income bracket. There is also a 57% increase in the chance of a child staying in the highest bracket when the parents are already in that bracket. However, differences related to a child ending up in the lowest income bracket were not significant at the 5% level, except when parents were in the middle income bracket.

Additionally, the gradual shift towards upward mobility is corroborated by trends in the stationary distribution vector. This vector suggests that, should the transition matrices continue to operate in the same manner over generations, there is a higher likelihood of individuals moving into the top income bucket as the birth cohorts progress from 1968 to 1997.

Furthermore, there also appears to be income inertia in intergenerational absolute mobility. The probability of remaining in the same income bucket as one's parents has remained consistent across cohorts, with the exception of the highest bucket, which experienced a notable increase. Moreover, the transition matrices showed no significant jumps in income buckets, indicating that movements were most likely to occur to adjacent income buckets. Holistically, these changes in transition probabilities suggest that over time, children are increasingly likely to surpass their parents' income bucket, particularly for those born in later cohorts.

Yet, the reliability of these findings is moderated by the high annual variation observed in the time series data and limited existing literature. Thus, it is crucial to interpret the results presented in Chapter 5 within the context of the specific parameterization assumptions made in this paper.

The two principal assumptions underpinning the parameterization of absolute mobility are: (1) the use of income data recorded at age 26 for both parents and offspring and (2) the creation of predefined income buckets, along with the corresponding adjustment for inflation. These assumptions are likely the primary drivers of the significant year-over-year variability and strong upward mobility trends observed in the time series analysis.

Using age 26 as a benchmark to compare incomes across generations introduces significant variability, complicating the analysis of economic mobility. Recent trends indicate a shift in career and education paths, as shown by Pew Research Center (2023) and Tamborini et al. (2015). Today, fewer young adults hold full-time jobs by

26, with many instead pursuing higher education. This leads to later entry into full-time work and a delay in reaching peak earnings. Thus, directly comparing the income of 26-year-olds from different generations might not accurately capture changes in economic mobility. This is because younger generations may be following a different economic trajectory influenced by increased educational attainment. Therefore, our methodology's failure to adjust the chosen age for income comparison may not fully reveal the true nature of upward absolute mobility amid evolving societal norms. Note that these dynamics have a lesser impact on relative mobility since quintiles are redefined each year, ensuring a more standardized comparison across time.

Moreover, pegging income buckets to quintiles of the 1968 income distribution and deflating subsequent earnings adds a degree of ambiguity. The median income in the United States was \$3,700 in 1968, \$3,600 in 1998, and \$3,900 in 2023, based on empirical data. In contrast, the 80% threshold rose from \$7,000 in 1968 to \$7,700 in 1998, and further to approximately \$8,600 by 2023. This suggests that while median income remained relatively steady, the income distribution widened above the median. As a result, the greater likelihood of upward mobility indicated by our absolute mobility model (defined by fixed buckets) may primarily reflect this widening income distribution rather than inherently greater levels of intergenerational mobility.

In light of the widening income distribution above the median, the findings in the absolute mobility delta matrix (Equation 5.4) warrant closer examination. Columns corresponding to children in higher income buckets show exclusively positive and statistically significant values. This phenomenon suggests that these results may be structurally skewed. A simple inflation adjustment likely does not adequately capture the true magnitude of changes in purchasing power over time, particularly as the distribution of income above the median shifts upward between the 1968 and 1997 birth cohorts. To better reflect these reductions in purchasing power, a different approach that considers both wage growth and inflation may be necessary to better model absolute intergenerational mobility.

Ultimately, these assumptions regarding absolute mobility place restrictions on broader inferences we can make about fundamental shifts in mobility patterns. Therefore, any conclusions about broader absolute mobility should be approached with caution.

6.1 Limitations

6.1.1 Data

This paper's implementation of Markov matrices to model intergenerational mobility is primarily limited by the underlying dataset. The transition matrix calibration would have benefited from a larger number of parent-child data points within each birth year to more accurately discern rates of mobility. Note that the current PSID dataset averages around 400 parent-child associations per birth cohort, where "associations" refer to incomes recorded at the age of 26 for both generations. These limited data points increase the yearly noise in our analysis and reduce the validity of potential takeaways. However, expanding the dataset proves to be formidable due to the PSID's methodology, which depends on annual interviews with a consistent pool of families.

A transition to a more expansive digitized dataset could significantly bolster the findings presented in this paper. The Statistics of Income (SOI) Data Master Files, maintained by the Internal Revenue Service (IRS), cover the entire population and contain variables that facilitate the creation of family units. Additionally, every individual's record in this dataset is tagged with a unique Taxpayer Identification Number (TIN), which streamlines the process of tracking children throughout the dataset. The availability of unique TIN numbers is a key advantage of the SOI datasets over other large datasets like IPUMS-CPS, where longitudinal analysis is impossible. However, the SOI dataset is only accessible to select researchers through solicited research projects awarded to winning applicants (Chetty et al. (2014) utilizes the SOI dataset). Thus, we reiterate that while the SOI dataset can offer more precise estimators than the manually curated PSID data, the PSID remains the most comprehensive and publicly available dataset for studying intergenerational mobility.

Moreover the analysis conducted in this paper is constrained by the limited breadth of available data, with the earliest data in the PSID corresponding to 1968. Thus, there are only 29 transition matrices total, encompassing the period between the 1968 to 1997 birth cohorts. This timeframe encompasses a single generational shift, as distinct human generational cohorts (e.g., baby boomers, Generation X) are typically demarcated in 25-year increments. These data limitations diminish the utility of the delta matrix, as it can only compare differences in intergenerational mobility across an abbreviated window. Ideally, this study would incorporate data from the beginning of the 20th century to examine how rates of mobility have evolved over a century.

6.1.2 Model Construction

The other significant limitation in this paper relates to the Markov transition matrix formulation. Specifically, the state space defined by segmenting the income distribution into five discrete income subgroups introduces bias. By categorizing income into lower and higher subgroups, the model implicitly biases transition probabilities, as the lowest income group is predisposed to experiencing a higher likelihood of upward mobility, and the highest income group is predisposed to a higher likelihood of downward mobility.

Furthermore, it is crucial to provide a broad caveat regarding the use of Markov processes in this paper and the corresponding mathematical properties, such as the stationary distribution and mixing time. The transition matrices calibrated in this model are calibrated over two time periods (parent and child generations), representing only one iteration of the transition matrix. This approach is atypical for Markov chain processes, as they are usually trained on a much wider set of observable time states. Consequently, although a stationary distribution can be calculated for every calibrated transition matrix per birth cohort, the stationary distribution is qualitatively infeasible to achieve, as the transition matrix corresponding to subsequent generational leaps is different to reflect changes in mobility. Thus, this limitation reduces the utility of the mixing times as well, as they are contingent on the stationary distribution.

6.2 Future Directions

6.2.1 Robustness

This paper makes key assumptions regarding the definition of the state space for relative and absolute mobility, as well as the age at which income is tracked for parents and children. These assumptions can be sensitized to test the robustness of the results.

State Space

The relative mobility state space is currently defined by income quintiles. In future work, this state space can be altered to include more granular states. For instance, instead of quintiles, 10% increments could be established. Alternatively, the income subgrouping could be set such that the variance within each subgroup is the same, better reflecting the right-skewed nature of the income distribution.

The absolute mobility state space is defined by fixed income buckets pegged to the 1968 income quintiles. This assumption, with income in future years deflated to 1968 levels for comparison, can be sensitized in several ways. The income cutoffs could be parameterized to consider a wider set of options, such as basing them off fixed variance levels in the 1968 income distribution or choosing a different reference year for bucketing and adjusting accordingly. Furthermore, as discussed in the absolute mobility analysis, the income distribution has changed over time, with income above the mean becoming more spread out. This structurally increases the rate at which upward mobility is achieved in the absolute parameterization. To account for these changes, a better purchasing power adjustment deflation could be considered besides just the Consumer Price Index (CPI) data. For example, income could be adjusted using the Producer Price Index (PPI), Gross Domestic Product (GDP), Employment Cost Index (ECI), or a weighted combination of these adjustments. In aggregate, sensitizing the results will likely reduce skews toward upward mobility and bring greater robustness to the absolute mobility analysis. This may allow for broader qualitative takeaways.

Lastly, instead of subgrouping based on income, in theory, future researchers could segment the population on other sociodemographic data like career or educational attainment to model intergenerational mobility.

Income Tracking Age

Another vector for sensitivity is the age at which parent and child income is tracked. This paper arbitrarily chooses the age of 26 as a way to extend the available time series. Sensitizing for this is important given changes in lifecycle earnings due to broader shifts in career and education paths discussed earlier. A simple robustness check could involve changing the income age to 30 or 35 and observing how the results change. In future studies, a more precise approach would be to formulate a metric that tracks the adjusted lifecycle age of earnings for each birth cohort and take income data for children at the corresponding adjusted age. Furthermore, to reduce year-over-year noise in observed data, particularly in the absolute mobility case, an average three or five-year window of income could be used instead of an individual data point.

Parent Income Subgrouping

Another assumption that can be addressed is using the lower parent income for relative and absolute mobility calculations, rather than the higher income parent. This assumption regarding dual-income families mainly affects later birth cohorts. While its impact likely applies consistently across income subgroups, dual-income households may be disproportionately present in certain subgroups, potentially impacting transition probabilities.

6.2.2 Model Construction

Future research can refine and adjust the Markov formulation presented in this study by incorporating additional income data spanning multiple generations. As income data encompassing future generations becomes available, it will be possible to extend the Markov formulation to consider an additional time period. In this case, the model would encompass two generational shifts, representing the transitions from parent to child and from child to grandchild. Within this extended modeling framework, the delta matrix analysis employed in the present study can be revisited to analyze differences in transition probabilities within the same familial lineage, allowing for a comparison of how transition probabilities differ between the parent-child and child-grandchild transitions.

Furthermore, the Markov formulation can be adapted more broadly by future researchers, particularly in the context of absolute mobility analyses. For instance, given information spanning at least three generations, future studies on absolute mobility could construct Markov transition matrices that reflect the likelihood of a child earning a certain percentage more than their parents, conditional on their parents having earned a certain percentage more than the child's grandparents. This extension would provide insights into upward mobility persistence across generations.

6.2.3 Event Studies

Further research can delve deeper into the qualitative underpinnings of the transition matrix results, given the prevalent role of intergenerational mobility in policy and individual decision-making. Potential event studies can create delta matrices that specifically compare two periods with different underlying tax regimes, geopolitical landscapes, or political leadership. For instance, a comparative analysis of transition matrices could show differences in intergenerational mobility between individuals eligible for the Vietnam War draft and those who were not. Note that observed differences between events are likely to be presented as stylized facts, as econometrically linking world events to differences in observed transition probabilities involves numerous confounding factors.

6.3 Implications

The work conducted in this paper has significant implications for stakeholders interested in understanding intergenerational mobility patterns and their evolution across generations. Specifically, the parameterization of mobility in terms of Markov transition matrices offers researchers a framework for analyzing rates of mobility over time, as well as the long-term implied stationary distribution of current mobility dynamics.

While publicly available datasets are limited in terms of historical data, the results of this paper can complement existing literature on intergenerational mobility and inform policymakers on recurring poverty cycles and rates of upward mobility. Furthermore, the more nuanced understanding of mobility presented in this study can provide future researchers with an avenue to explore how other factors, such as family background and education, are linked with children's economic outcomes.

Overall, research on intergenerational mobility plays a crucial role in public discourse, and the findings from this study, in conjunction with existing and future studies, can affirm or challenge narratives regarding the "American Dream".

Bibliography

Behrman, J. and Taubman, P. (1985). Intergenerational earnings mobility in the united states: Some estimates and a test of becker's intergenerational endowments model. *The Review of Economics and Statistics*, 67(1):144–151.

Borjas, G. J. (1992). Ethnic Capital and Intergenerational Mobility*. *The Quarterly Journal of Economics*, 107(1):123–150.

Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R., and Narang, J. (2017). The fading american dream: Trends in absolute income mobility since 1940. *Science*, 356(6336):398–406.

Chetty, R., Hendren, N., Kline, P., Saez, E., and Turner, N. (2014). Is the united states still a land of opportunity? recent trends in intergenerational mobility. *American Economic Review*, 104(5):141–47.

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives*, 27(3):79–102.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Warren, D., Chen, A., Cooper, G., Richards, S., Schouweiler, M., and Westberry, M. (2023). Ipums, current population survey: Version 11.0.

Habib, B. and Perese, K. (2016). *The Distribution of Household Income, 2016*.

House, J. S. (1976). *American Journal of Sociology*, 81(5):1236–1238.

Hout, M. (1988). More universalism, less structural mobility: The american occupational structure in the 1980s. *American Journal of Sociology*, 93(6):1358–1400.

Katz, M. (2013). *The Undeserving Poor: America's Enduring Confrontation with Poverty: Fully Updated and Revised*. OUP USA.

Krueger, A. (2012). The rise and consequences of inequality in the united states. Remarks by Council of Economic Advisers Chairman Alan Krueger at the Center of American Progress, Washington, D.C. [Accessed: 2024 04 09].

Lee, C.-I. and Solon, G. (2009). Trends in Intergenerational Income Mobility. *The Review of Economics and Statistics*, 91(4):766–772.

Pew Research Center, . (2023). Young adults in the u.s. are reaching key life milestones later than in the past.

Samuel, L. R. (2012). *Introduction*, pages 1–11. Syracuse University Press.

Solon, G. (1999). Chapter 29 - intergenerational mobility in the labor market. volume 3 of *Handbook of Labor Economics*, pages 1761–1800. Elsevier.

Survey Research Center, . (2023). Panel study of income dynamics.

Tamborini, C. R., Kim, C., and Sakamoto, A. (2015). Education and lifetime earnings in the united states. *Demography*, 52(4):1383–1407.

U.S. Bureau of Labor Statistics (2023). Consumer price index for all urban consumers: All items in u.s. city average [cpiaucsl]. Retrieved from FRED, Federal Reserve Bank of St. Louis. Accessed: March 5, 2024.

Zimmerman, D. J. (1992). Regression toward mediocrity in economic stature. *The American Economic Review*, 82(3):409–429.