# The Early Career and the Sources of Life-Cycle Hours Growth*

Paulo Lins
Clemson University

Xiaonan Ma

This version:
March, 2026

## Abstract

Hours worked rise sharply in the first decade after labor-market entry. We show that this growth is primarily an intensive-margin phenomenon driven by human-capital incentives: about 60% of the increase in cumulative hours reflects longer workweeks among employed workers, not additional weeks worked. To reach this conclusion, we estimate a life-cycle model with learning-by-doing and on-the-job search using quarter-since-entry profiles constructed from the NLSY79 and NLSY97. The model replicates the joint dynamics of hours, wages, employment, and job mobility from the moment workers leave school. In the estimated model, the return to current hours operates not only through future wages but also through improved job stability and access to better outside offers. A decomposition of wage inequality shows that the growing covariance between human capital and match quality accounts for the majority of the rise in wage variance over the first 20 years after entry.

**Keywords:** Life-Cycle Models, Labor Supply, Human Capital Accumulation

**JEL - Classification:** D15, J22, J24

*We thank Alex Bick, Mark Bils, Adam Blandin, Nataliya Gimpelson, Eugenio González Flores, Chris Herrington, Ronni Pavan, Joseph Pedtke, Juan Sánchez, Minchul Yum, and Ruilin Zhou for their suggestions, and the participants of seminars in Virginia Commonwealth University and Stockman Conference at University of Rochester. For their helpful comments on an earlier draft, we are also grateful to Roman Merga and Gustavo Serra.

# 1 Introduction

Hours worked rise sharply over the first decade of working life. This pattern is hard to reconcile with standard intertemporal labor-supply logic: if wages are expected to rise with experience, workers should postpone labor supply. A common interpretation is that rising hours mainly reflect improving labor-market opportunities as workers move out of nonemployment and into better jobs (Topel and Ward, 1992; Neal, 1999). An alternative is that hours rise because work itself raises future productivity through human-capital accumulation (Heckman, 1976; Imai and Keane, 2004). Distinguishing among these forces is important for understanding the sources of lifetime earnings inequality (Huggett, Ventura, and Yaron, 2011; Bick, Blandin, and Rogerson, 2024) and the effects of policies that alter work incentives over the life cycle (Fan, Seshadri, and Taber, 2024).

This paper argues that life-cycle hours growth is primarily an intensive-margin phenomenon driven by human-capital incentives rather than by improving match quality. In the data, about 60% of the increase in cumulative hours over the first 20 years after labor-market entry comes from longer workweeks among employed workers, not from additional weeks worked. An estimated life-cycle model with learning-by-doing and on-the-job search replicates this decomposition and implies that rising hours are largely a human-capital investment decision.

Our key empirical contribution is to study the life cycle from labor-market entry rather than from an arbitrary age cutoff. Using weekly work histories from the NLSY79 and NLSY97, we construct quarter-since-entry profiles of employment, hours, wages, and transition rates. This timing matters because the first years after entry are precisely when job shopping is most intense and when workers accumulate skills most rapidly. Starting observation later in life misses much of this adjustment and changes the measured decomposition of both hours growth and wage inequality. For example, for a worker who finishes high school at age 18, beginning measurement at age 25 omits seven years of early-career adjustment.

To identify the mechanism behind hours growth, we build a life-cycle model around two state variables: human capital and match quality. Human capital is accumulated through learning-by-doing, so current hours directly raise future productivity. Match quality evolves through random search: workers draw new matches from an exogenous distribution and face a convex switching cost that reflects the idea that human capital becomes increasingly match-specific over time. This creates a key tension – workers with poor matches anticipate future transitions and therefore face a lower return to current hours, since switching destroys part of their stock. As experience accumulates, the cost of leaving rises, generating endogenous lock-in that reduces mobility even as outside offers continue to arrive. The model also features preference heterogeneity in the disutility of work, which Bick et al. (2024) identify as an important source of hours dispersion, and borrowing subject to constraints.

Wages, job-finding rates, separation rates, and the mean of the outside-offer distribution all depend on both state variables. As a result, both human capital and match quality shape the return to current hours through multiple paths: human capital raises current wages directly and also improves future job stability and outside options, while match quality also affects current wages and influences the likelihood of future transitions. The relative importance of each channel is not assumed but estimated. We target a broad set of life-cycle profiles – wages, hours, employment, transition rates, and dispersions – using indirect inference, and the joint behavior of these profiles disciplines the contribution of each force. In particular, no single profile is sufficient to separate the incentives created by skill accumulation from those created by changes in job opportunities.

With the estimated model in hand, we decompose wages and earnings over the life cycle into contributions from human capital, match quality, and their interaction. Wage inequality rises early in the career not only because workers accumulate different amounts of skill, but also because human capital and match quality increasingly sort together: workers who accumulate more skill also find and keep better matches. This growing covariance accounts for the majority of the increase in wage variance over the first 20 years after entry. The same mechanism also implies that the return to hours is not limited to future wages. In the estimated model, working more early in the career improves job stability and access to better outside offers, making these channels quantitatively important incentives for labor supply beyond those emphasized in standard human-capital models such as Imai and Keane (2004).[1]

**Related Literature** This paper brings together two literatures that have largely evolved separately. One studies life-cycle labor supply through human capital accumulation, emphasizing that learning-by-doing creates an incentive to work more early in life because current hours raise future wages (Heckman, 1976; Weiss, 1986;

[1]A counterfactual in which human capital is fully transferable across jobs generates counterfactually steep wage growth and eliminates early-career hours heterogeneity, implying that match-specificity, captured by the adjustment cost, is essential for disciplining both investment and mobility.

Imai and Keane, 2004). Recent quantitative work in this tradition decomposes lifetime earnings inequality into different components (Huggett et al., 2011; Bick et al., 2024) and studies how policy shapes skill-accumulation incentives over the life cycle (Fan et al., 2024). But these papers typically begin observation many years after the school-to-work transition and abstract from job search. A second literature studies wage growth and mobility in frictional labor markets, emphasizing movement toward better matches (Topel and Ward, 1992; Neal, 1999). Recent contributions combine human capital and match quality to decompose wage growth (Bagger, Fontaine, Postel-Vinay, and Robin, 2014; Ozkan, Song, and Karahan, 2023) or to quantify the sources of displacement costs (Burdett, Carrillo-Tudela, and Coles, 2020; Audoly, De Pace, and Fella, 2022). In this literature, however, hours are usually treated as exogenous or secondary. We combine these two views by studying a model in which hours are endogenous and respond jointly to human-capital incentives and labor-market frictions, and by estimating that model on labor-market profiles observed from entry.

The closest papers are Kaplan (2012), Erosa, Fuster, and Kambourov (2016), and Bick et al. (2024). Kaplan (2012) studies young workers and shows that involuntary unemployment is central to early-career hours inequality, but his framework does not include human capital accumulation. Erosa et al. (2016) develop a rich life-cycle model with intensive and extensive margins, preference heterogeneity, and fixed costs of work, but abstract from human capital accumulation and on-the-job search. Our framework differs by combining learning-by-doing with an explicit search process over employment, unemployment, and job-to-job transitions, so that hours affect future wages while match quality shapes both wage growth and mobility. In a different vein, Bick et al. (2024) document large differences in lifetime hours and identify their sources, pointing to the importance of preference heterogeneity. We complement their analysis by studying how hours, human capital, and job search interact from labor-market entry, also allowing for preference heterogeneity.

Empirically, we align data and model at labor-market entry by constructing quarter-since-entry profiles from weekly work histories, and we decompose hours growth into intensive and extensive margins following Baily, Hulten, and Campbell (1992) and Blundell, Bozio, and Laroque (2011). This timing is important because it captures the phase in which human-capital accumulation, nonemployment transitions, and job shopping are all most active. Related work on the school-to-work transition studies how schooling, occupational choice, and credit constraints shape labor-market entry (Keane and Wolpin, 1997; Lochner and Monge-Naranjo, 2011;

Flinn and Mullins, 2015); our focus is instead on the dynamics after entry, when hours, mobility, and wage growth are jointly determined.

### Roadmap

Section 2 describes the life-cycle model with human capital accumulation and job search frictions. Section 3 describes the data and our measurement choices, including labor market entry and the construction of quarters-since-graduation profiles. Section 4 describes our estimation strategy. We also use a statistical decomposition of the growth into its components. Section 5 reports the quantitative results. Section 6 explores quantitative implications. Section 7 concludes.

## 2 Model

The model is built to answer one question: why do hours rise early in the life cycle? Our mechanism is that hours are an investment. Working more today raises future human capital, and the return to that investment depends on both current match quality and expected mobility. A worker in a good match earns more today, but also has less reason to move. A worker in a poor match expects future job changes, which lowers the value of accumulating match-specific human capital. Hours, therefore, respond not only to the current wage but also to the effect of current work on future wages, job stability, and access to better outside offers.

To capture this mechanism, the model combines learning-by-doing with a frictional labor market. Human capital grows through work experience. Workers search both in and out of employment, and jobs differ in match quality. Human capital and match quality jointly shape wages, offer arrival rates, and separation risk, so both state variables affect the return to current hours through several margins at once. The model also allows for saving, borrowing constraints, and heterogeneity in the disutility of work.

### 2.1 Setup

Time is discrete and measured in months to capture the dynamics of job transitions and separations.

**Preferences** Workers live for a finite horizon $t = 0, 1, \ldots, T$. In each period, they choose consumption $c_t$ and, when employed, hours worked $h_t \in [0, 1]$. Lifetime utility is

$$E_0 \sum_{t=0}^{T} \beta^t \left( \ln(c_t) - \xi_t \frac{h_t^{1+\frac{1}{\gamma}}}{1 + \frac{1}{\gamma}} \right), \qquad \beta \in (0, 1).$$

The parameter $\xi_t > 0$ governs the disutility of work, while $\gamma > 0$ controls its curvature. We allow $\xi_t$ to vary across workers and over the life cycle. This heterogeneity helps the model match dispersion in hours, but it is not the main force behind rising average hours.

**Labor Market** At each date, a worker is either unemployed or employed in a match of quality $m_t$. Unemployed workers receive benefits $\bar{b}$ and search for jobs. Employed workers choose hours and may either remain in their current match or move to a new one when an outside offer arrives. Search is frictional in both states.

All transition rates use logistic functions of human capital and, where relevant, match quality, so that they remain in $[0, 1]$ for all states. This guarantees that continuation terms are well-defined expected values while allowing state variables to shift rates smoothly.

Each period, an unemployed worker receives a match offer with arrival rate $\lambda^U(s_t)$ that depends on his stock of human capital $s_t$,

$$\lambda^U(s_t) = \frac{1}{1 + \exp\left(\lambda_0^U - \rho_{\lambda s} \ln(s_t)\right)}.$$

The parameter $\lambda_0^U$ determines the baseline job-arrival rate, while $\rho_{\lambda s}$ determines how job-finding varies with human capital. Thus, more skilled workers receive job offers more frequently.

Match quality offers are drawn from an exogenous distribution. The cumulative distribution function (CDF) of match offers, $F_{q^U}(s)(m)$, is log-normal:

$$F_{q^U}(s)(m) = \Phi\left(\frac{\log(m) - \mu^U(s)}{\sigma^U(s)}\right),$$

where $\mu^U(s) = \mu_0^U + \mu_1^U s$ and $\ln \sigma^U(s) = \sigma_0^U + \sigma_1^U s$ are the mean and log standard deviation of the underlying normal distribution, and $\Phi(\cdot)$ denotes the standard normal CDF. If a worker rejects the offer, he remains unemployed; if he accepts it, he begins the next period employed at the accepted match. Human capital shifts both the mean and the spread of the offer distribution, so more skilled workers draw better matches on average. While unemployed, human capital depreciates at rate $\delta_u$ per period.

An employed worker is characterized by current human capital $s_t$ and match quality $m_t$. Labor income equals the wage per unit of labor times hours worked, where the wage rate is

$$w(s_t, m_t) = w_0 s_t^{\rho_{ws}} m_t^{\rho_{wm}}.$$

The parameter $w_0$ determines the wage level, while $\rho_{wm}$ and $\rho_{ws}$ govern the sensitivity of wages to match quality and human capital.

Each period, an employed worker receives an outside offer with arrival rate $\lambda^E(s_t, m_t)$,

$$\lambda^E(s_t, m_t) = \frac{1}{1 + \exp\left(\lambda_0^E - \rho_{\lambda s} \ln(s_t) - \rho_{\lambda m} \ln(m_t)\right)}.$$

Here, $\lambda_0^E$ determines the baseline arrival rate, while $\rho_{\lambda s}$ and $\rho_{\lambda m}$ determine how it varies with human capital and match quality. Match offers are drawn from an exogenous distribution distinct from that faced by unemployed workers. The CDF of match quality, $F_{q^E}(s)(m)$, is log-normal:

$$F_{q^E}(s)(m) = \Phi\left(\frac{\log(m) - \mu^E(s)}{\sigma^E(s)}\right),$$

where $\mu^E(s) = \mu_0^E + \mu_1^E s$ and $\ln \sigma^E(s) = \sigma_0^E + \sigma_1^E s$ are the mean and log standard deviation of the underlying normal distribution, and $\Phi(\cdot)$ denotes the standard normal CDF.

In addition, an employed worker faces exogenous separation risk $\pi(s_t)$, which declines with human capital:

$$\pi(s_t) = \frac{1}{1 + \exp\left(\pi_0 - \rho_{\pi s} \ln(s_t)\right)}.$$

The parameter $\pi_0$ determines the baseline separation rate, while $\rho_{\pi s} < 0$ governs how separation varies with human capital. Thus, more human capital reduces the risk of job loss. When separated, workers begin the next period unemployed. While employed, human capital evolves according to the law of motion described below.

**Human Capital** Employed workers accumulate human capital through a learning-by-doing process that depends on current human capital $s_t$, match quality $m_t$, and hours worked $h_t$,

$$\mathcal{H}(s_t, m_t, h_t, \alpha_t) = \alpha_t s_t^{\rho_{ss}} m_t^{\rho_{sm}} h_t^{\rho_{sh}} \ .$$

The parameter $\alpha_t > 0$ determines the efficiency of learning by doing, while $\rho_{ss}$, $\rho_{sm}$, and $\rho_{sh}$ are the elasticities of human capital accumulation with respect to current human capital, match quality, and hours worked. We allow workers to differ in learning ability, as described in more detail below.

To capture the idea that part of accumulated human capital is tied to the current job, we assume that workers lose human capital when they change jobs, either voluntarily or through separation. We model this loss as a one-time adjustment cost that is convex in current human capital,

$$\phi(s_t) = B s_t^{\nu} \ .$$

Here, $B > 0$ determines the level of the loss and $\nu > 1$ its curvature. The convexity implies that workers with more accumulated human capital lose more when leaving the current match.

The law of motion for human capital is therefore

$$s_{t+1} = \begin{cases} (1-\delta_e)s_t + \mathcal{H}(s_t, m_t, h_t, \alpha_t), & \text{if no job change ,} \\ (1-\delta_e)s_t + \mathcal{H}(s_t, m_t, h_t, \alpha_t) - \phi(s_t), & \text{if job change or separation ,} \end{cases}$$

where $\delta_e$ is the depreciation rate while employed. Thus, hours raise future human capital through learning by doing, but job changes reduce the amount of accumulated capital carried into the next period. For unemployed workers, human capital depreciates at rate $\delta_u$, with no new accumulation.

**Budget Constraint** Both employed and unemployed workers can save and borrow using a one-period risk-free asset with net return $r$. Borrowing is subject to the constraint $a_{t+1} \geq \bar{a}_t$, where $\bar{a}_t$ reflects a natural debt limit based on the present value of future income while unemployed and receiving unemployment benefits. For quantitative reasons, we allow the borrowing limit to be tighter than this natural debt limit. The per-period budget constraint is

$$c_t + a_{t+1} = \begin{cases} w(s_t, m_t)h_t + (1+r)a_t & \text{if employed,} \\ \bar{b} + (1+r)a_t & \text{if unemployed.} \end{cases}$$

**Retirement** Workers retire at age $T_R < T$. During retirement, they no longer supply labor and do not receive labor income. Instead, they receive common retirement income $\mathrm{Ret}$. The budget constraint then becomes

$$c_t + a_{t+1} = \mathrm{Ret} + (1+r)a_t, \qquad \text{for } t \geq T_R.$$

Retirees cannot borrow against future retirement income, so the borrowing constraint tightens to $a_t \geq 0$ for $t \geq T_R$. The terminal condition is $a_{T+1} = 0$.

**Initial Conditions and Type Heterogeneity** Workers begin life with initial assets, human capital, match quality, and employment status. We assume that assets, human capital, and match quality follow a jointly log-normal distribution, while initial employment status is drawn from a Bernoulli distribution with unemployment probability $p_{u_0}$.

Workers also differ in their disutility of labor, $\xi_t$, and in the efficiency of human capital accumulation, $\alpha_t$. The disutility parameter evolves over the life cycle according to an AR(1) process in logs,

$$\log \xi_{t+1} = (1-\rho_\xi)\mu_\xi + \rho_\xi \log \xi_t + \varepsilon_{j,t}, \qquad \varepsilon_{j,t} \sim N(0, \sigma_\varepsilon^2),$$

where $\mu_\xi$ is the unconditional mean, $\rho_\xi \in (0,1)$ is the persistence parameter, and $\sigma_\varepsilon = \sigma_\xi \sqrt{1-\rho_\xi^2}$ is the innovation standard deviation implied by the unconditional variance $\sigma_\xi^2$. The initial distribution of $\xi_t$ is log-normal with mean $\mu_\xi$ and standard deviation $\sigma_{\xi,0} < \sigma_\xi$, so dispersion increases over the life cycle as workers receive idiosyncratic shocks.

The learning-efficiency parameter $\alpha_t$ is allowed to vary systematically with labor disutility. Specifically,

$$\log \alpha_t = C_0^\alpha + C_1^\alpha \log \xi_t,$$

where $C_0^\alpha$ and $C_1^\alpha$ govern the relationship between learning efficiency and labor disutility. We assume this relationship is negative, so workers who dislike work more are also less efficient at accumulating human capital. This generates a reinforcing mechanism: workers with high disutility supply less labor, accumulate less human capital through experience, and learn less efficiently on the job.

## 2.2 Workers' Problem and Value Functions

An employed worker at time $t$ chooses consumption $c_t$, savings $a_{t+1}$, hours worked $h_t$, and, when an outside offer arrives, whether to remain in the current job or move to a new employer. The value function depends on age $t$, assets $a_t$, human capital $s_t$, current match quality $m_t$, and current labor-disutility type $\xi_t$:

$$
\begin{aligned}
&V_t^W(a_t, s_t, m_t, \xi_t) = \\
&\quad \max_{c_t, a_{t+1}, h_t} \Bigg\{ u(c_t) - v(h_t) \\
&\quad + \beta\pi(s_t)\, \mathbb{E}_{\xi'} U_{t+1}(a_{t+1}, s'_{t+1}, \xi') \\
&\quad + \beta\Big(1 - \pi(s_t)\Big)\lambda^E(s_t, m_t)\mathbb{E}_{m',\xi'}\left[\max\left\{V_{t+1}^W(a_{t+1}, s_{t+1}, m_t, \xi'), V_{t+1}^W(a_{t+1}, s'_{t+1}, m', \xi')\right\}\right] \\
&\quad + \beta\Big(1 - \pi(s_t)\Big)\Big(1 - \lambda^E(s_t, m_t)\Big)\mathbb{E}_{\xi'} V_{t+1}^W(a_{t+1}, s_{t+1}, m_t, \xi')\Bigg\},
\end{aligned}
$$

subject to the budget constraint, the borrowing constraint, the hours restriction $0 < h_t < 1$, and the law of motion for human capital. Here, $s_{t+1}$ and $s'_{t+1}$ denote next-period human capital without and with a job transition, respectively. We omit the explicit functional forms of these laws of motion because they follow directly from the previous subsection.

Expectations are taken over next-period labor disutility, which also determines next-period learning efficiency. When an outside offer arrives, expectations also integrate over the match-offer distribution. An employed worker accepts a new offer when the value of the new match exceeds the value of remaining in the current one, taking into account the human-capital loss associated with switching. The optimal hours choice reflects not only the within-period tradeoff between consumption and leisure, but also the effect of current hours on future earnings, offer arrival rates, separation risk, and the stock of human capital carried into the future.

An unemployed worker at time $t$ chooses consumption $c_t$, savings $a_{t+1}$, and whether to accept an offer or remain unemployed. The value function depends on age $t$, assets $a_t$, human capital $s_t$, and current labor-disutility type $\xi_t$:

$$
\begin{aligned}
&V_t^U(a_t, s_t, \xi_t) = \\
&\quad \max_{c_t, a_{t+1}} \Bigg\{ u(c_t) - v(0) \\
&\quad + \beta \lambda^U(s_t) \mathbb{E}_{m', \xi'} \left[ \max \left\{ V_{t+1}^U(a_{t+1}, s_{t+1}, \xi'),\ V_{t+1}^W(a_{t+1}, s_{t+1}, m', \xi') \right\} \right] \\
&\quad + \beta \Big(1 - \lambda^U(s_t)\Big) \mathbb{E}_{\xi'} V_{t+1}^U(a_{t+1}, s_{t+1}, \xi') \Bigg\},
\end{aligned}
$$

subject to the budget constraint, the borrowing constraint, and the law of motion for human capital. As usual, accepting or rejecting an offer depends on whether the current draw exceeds the reservation match quality.

A retired worker at time $t \geq T_R$ chooses consumption $c_t$ and savings $a_{t+1}$. The value function depends on age $t$ and assets $a_t$:

$$
V_t^R(a_t) = \max_{c_t, a_{t+1}} \left\{ u(c_t) - v(0) + \beta V_{t+1}^R(a_{t+1}) \right\},
$$

subject to the retirement budget constraint and the no-borrowing constraint. The terminal condition is $V_{T+1}^R(a_{T+1}) = 0$.

Within each period, workers make their choices before shocks are realized. Shocks arrive in the following order: first, employed workers face a separation shock; second, both employed and unemployed workers receive match offers; third, labor-disutility shocks are realized.

The model contains two sources of uncertainty: shocks that arrive over the life cycle and heterogeneity in initial conditions. Separation and offer arrival events are Bernoulli, with probabilities given by the transition functions defined above. Match-quality draws are log-normal, and innovations to labor disutility are also log-normal. At labor-market entry, employment status is drawn from a Bernoulli distribution with unemployment probability $p_{u_0}$. Conditional on employment status, initial assets, human capital, and match quality are drawn from a joint log-normal distribution. Initial labor disutility is drawn from a log-normal distribution with mean $\mu_\xi$ and standard deviation $\sigma_{\xi,0}$, while initial learning efficiency is pinned down by the mapping described above.

### 2.3 Discussion

Human capital and match quality are the key state variables in the model. Human capital raises wages, improves labor market prospects, and lowers separation risk. Match quality affects current pay and the value of remaining in the current job relative to moving elsewhere. Together, these two variables shape the joint evolution of wages, job mobility, and labor supply over the life cycle. Preference heterogeneity plays a different role. The disutility parameter $\xi_t$ shifts labor supply directly and, through its correlation with learning efficiency, contributes to the dispersion of human capital over time.

The human capital channel resolves a tension that arises in standard intertemporal labor supply models. Intertemporal substitution predicts that workers should work more when wages are high and less when wages are expected to rise in the future. For young workers facing upward-sloping wage profiles, this logic implies low hours early in the career and high hours later, the opposite of the pattern observed in the data (Card, 1994). Human capital accumulation reverses this implication. Because hours today raise future earnings, working more when young is an investment, as in Heckman (1976) and Imai and Keane (2004). In this sense, expected wage growth is not taken as given; it is partly created by current labor supply.

Job search and switching costs introduce a second force through the interaction between mobility and accumulation. Workers with little experience and poor matches expect more future job transitions. Because changing jobs destroys part of the human capital accumulated in the current match, the return to current hours is lower when a worker expects to leave soon. This force dampens labor supply early in the career, when match quality is low and transition rates are high. As workers accumulate experience, they have more to lose from switching, which generates endogenous lock-in and reduces mobility with age. Together, these mechanisms produce rising hours and declining job-to-job transitions over the life cycle.

The model abstracts from several features present in richer labor market frameworks. Separations are exogenous, wages are given by a reduced-form wage function rather than bargaining, and the model omits firm heterogeneity and employer learning. Retirement is also modeled in a reduced-form way. These simplifications keep the framework tractable while preserving the mechanisms central to the paper: human capital accumulation, job search, and the loss of human capital at job transitions.

# 3 Measurement: Labor Market Entry and Profiles

In this section, we describe the sample construction, the identification of labor market entry, and the construction of quarter-since-graduation profiles. We use these profiles to estimate the model.

## 3.1 Data and Sample Restrictions

We use data from the National Longitudinal Survey of Youth 1979 and 1997 (NLSY79 and NLSY97, respectively) to construct profiles of many labor-market outcomes, particularly hours worked, wages, and transition rates. The NLSY79 is a nationally representative sample of 12,686 individuals born between 1957 and 1964 who were aged 14–22 in 1979. The NLSY97 is a nationally representative sample of 8,984 individuals born between 1980 and 1984 who were aged 12–17 in 1997. In addition to their panel structure, the main advantage of these datasets is the extensive information they provide on respondents' labor-market and educational histories.

Our empirical analysis focuses on individuals with a high school diploma as their highest educational attainment, since their labor market entry is clearly defined.[2] We use the NLSY employment histories to construct a quarterly panel. Variables such as weekly employment status and employer identifiers are reported for each week. Other variables, including wages, hours, and job characteristics, are reported only once per survey wave for each job. These are assigned to the survey year and, by convention, mapped to all calendar weeks within that survey year. We then aggregate the weekly data to a quarterly frequency, which is appropriate for analyzing the extensive margin and for capturing the rapidly changing labor-market outcomes of young workers.

After defining the target population and analysis frequency, we impose additional sample restrictions. Table 1 lists all restrictions and their impact on sample size. We sequentially apply the following criteria: (i) completion of the 12th grade; (ii) reporting a valid high school graduation date; (iii) graduation between ages 16 and 24 to restict the sample to a typical life-cycle trajectory; (iv) no military service; (v) at least 10 years of panel participation in the NLSY79 or 5 years in the NLSY97; and (vi) observations available for either the first or second year following gradua-

[2] For individuals who did not complete high school, entry is ambiguous–occurring either after dropping out or later. Similarly, for those with education beyond high school, it is unclear whether entry occurred after high school graduation, college graduation, or the completion of professional or postgraduate degrees.

tion. After applying all restrictions, the final sample consists of 3,341 respondents from the original 12,686 individuals in the NLSY79 and 1,563 respondents from the original 8,984 individuals in the NLSY97.[3]

| | NLSY79 | | NLSY97 | |
|---|---|---|---|---|
| | Ind. | % Sample | Ind. | % Sample |
| Initial Sample | 12,686 | 100.0% | 8,984 | 100.0% |
| Final Sample | 3,341 | 26.3% | 1,563 | 17.4% |
| *Restrictions applied* | | | | |
| Completed 12th grade | 6,009 | 47.4% | 2,386 | 26.6% |
| Reported a valid graduation date | 5,521 | 43.5% | 2,112 | 23.5% |
| Graduated between ages 16 and 24 | 5,072 | 40.0% | 2,024 | 22.5% |
| Did not serve in the military | 3,906 | 30.8% | 1,911 | 21.3% |
| Met panel duration requirements | 3,563 | 28.1% | 1,803 | 20.1% |
| Observed 1st or 2nd year post-graduation | 3,341 | 26.3% | 1,563 | 17.4% |

Table 1: Sample Selection Criteria and Final Sample Size

When computing life-cycle profiles to calibrate the model, we restrict the sample to individuals with a minimum level of labor-market attachment, defined as 520-5,200 annual hours. This restriction ensures that our profiles describe the outcomes of workers with sustained attachment, rather than being driven by individuals with little or no engagement in the labor market. A similar restriction is also used by Kaplan (2012) and Bick et al. (2024) when constructing their profiles. Importantly, most of our profiles are constructed conditional on employment, while the extensive margin is captured separately in the employment rates. We also use a statistical decomposition to quantify the contribution of entry and exit to the hours profile.

## 3.2 Identifying Labor Market Entry

We identify labor market entry using the high school graduation date. The NLSY79 and NLSY97 are unique for this purpose: in addition to their long panel structure and detailed labor market questions, they collect information on each diploma received, including the date it was awarded. Graduation dates are available in the NLSY97 but need to be created in the NLSY79. To do so, we rely on two sets of ques-

[3]Appendix A.1 reports descriptive statistics to evaluate potential selection concerns arising from our sample restrictions. The table compares three groups: the full NLSY79 cohort, all high school graduates, and the final analysis sample. As expected, high school graduates differ substantially from the full cohort, but the final sample closely resembles the broader population of high school graduates.

tions. First, respondents report whether they received a diploma, which diploma, and the date awarded. Second, they report the highest degree attained and the date it was obtained. When conflicting dates are reported for the same diploma, we manually cross-check responses using additional information, such as the highest degree attained up to that survey wave and the reason for leaving school.

Having identified labor-market entry, we construct quarters-since-graduation profiles. These profiles are preferable to the standard birth-age profiles for two reasons. First, within a given year bin, workers in our profiles have the same potential years of experience. By contrast, in birth-age profiles, workers are heterogeneous: at the same age, some may have just entered the labor market while others already have several years of experience. Second, quarters-since-graduation profiles align more closely with life-cycle models, which typically abstract from schooling and assume that workers are "born" when they enter the labor force.

Figure 1, Panels (a) and (b), shows the distribution of labor market entry in the NLSY79 and NLSY97, respectively. Each panel presents two histograms: the distribution for individuals with a high school diploma as their highest educational attainment and, in shaded bars, the distribution for the entire sample, including those with education above or below high school. These distributions are presented before applying any sample restrictions.

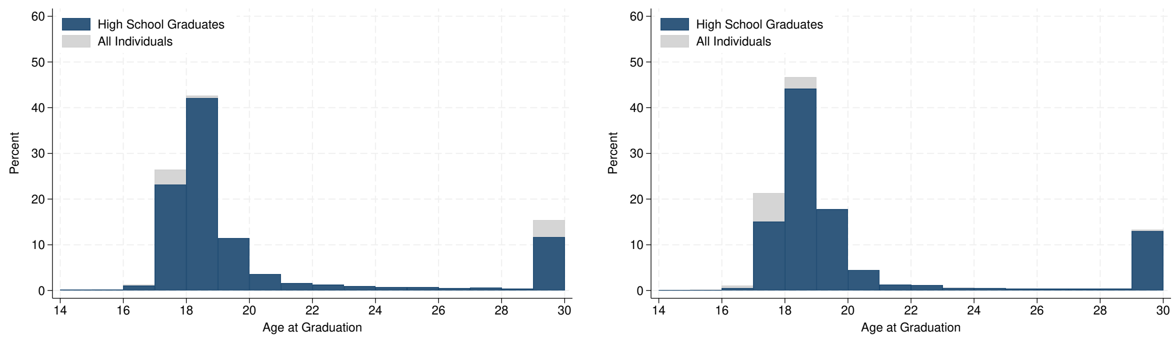


(a) NLSY79 (b) NLSY97

Figure 1: High School Graduation Distribution

Note: The figure plots the age distribution at high school graduation for individuals in the NLSY79 and NLSY97 samples. The lighter bars represent all individuals, while the darker bars highlight those whose highest educational attainment is a high school diploma. Ages below 14 and above 30 are grouped. For the NLSY79, it is not possible to distinguish between GED certificates and high school diplomas; both are treated equivalently. For the NLSY97, when both GED and high school graduation dates are available, the earlier date is used.

Two patterns stand out. First, the distribution of high school graduation is dispersed. In the NLSY79, the mean, median, and standard deviation of labor market entry are 19.64, 18.34, and 4.40 years, respectively; in the NLSY97, they are 19.14, 18.52, and 2.39 years. Second, by ages 21 and 25, individuals in the NLSY79 have, on average, 2.13 and 5.51 years of labor market experience, respectively, while those in the NLSY97 have 1.97 and 5.34 years, respectively. These patterns underscore why (i) using age as a proxy for labor market experience and (ii) restricting analysis to workers over 25 are problematic when studying young workers. Our approach avoids these issues by identifying entry based on high school graduation dates and constructing quarters-since-graduation profiles.

### 3.3 Measurement

We use eleven quarterly profiles to estimate the model. They comprise four mean outcomes – wages, employment rates, hours worked, and cumulative number of jobs – their corresponding cross-sectional standard deviations, and three labor-market transition rates.

**Mean Profiles** We construct four mean profiles by aggregating weekly observations to quarterly frequency. Let $i$ index individuals and $q$ index quarters since high school graduation.

- *Employment rate*: The fraction of weeks in quarter $q$ that individual $i$ spends employed, averaged across individuals. By construction, this measure lies between zero and one. In the data, non-employment includes both unemployment and non-participation.
- *Hours worked*: Total hours worked in quarter $q$, conditional on positive employment. We focus on hours in the main job, though respondents may hold multiple jobs within a week. We explore the difference between hours worked in the main job and a broader measure that accounts for hours worked across all jobs in an appendix. Hours are divided by 168 (the maximum weekly hours multiplied by 13) to bound them between 0 and 1.
- *Wages*: The average real hourly wage across all employed weeks in quarter $q$, conditional on employment. We construct real wages using the CPI. Wages are normalized to their first-quarter values, so the wage profile begins at 1.
- *Cumulative jobs*: The count of distinct jobs held by individual $i$ up to the end of quarter $q$. Unlike the other profiles, this is a stock variable measured at

quarter's end rather than a flow averaged over the quarter. The cumulative jobs profile is shifted so that its initial value equals one.

**Standard Deviation Profiles** For each mean profile, we compute the corresponding cross-sectional standard deviation. The standard deviations of hours and wages are computed conditional on positive employment and are adjusted for differences in units using logarithms. The standard deviation of employment captures dispersion in the fraction of weeks worked across individuals within a quarter.

**Transition Rates** We construct three transition rates from week-to-week changes in labor market status: job-to-job, job-to-unemployment, and unemployment-to-job. For each transition type, we count transitions occurring between consecutive weeks and divide by the number of individuals at risk in the origin state. We define a job-to-job transition as a direct change in job identifier between two consecutive employed weeks. A separation is a transition from employment to non-employment (unemployment or out of the labor force). Job finding is the transition from non-employment to employment.

Weekly hazard rates are converted to quarterly rates by compounding over thirteen weeks:

$$\lambda^{\text{quarter}} = 1 - (1 - \lambda^{\text{week}})^{13}.$$

We then apply a four-quarter moving average to smooth the series. Directly measured transition rates may be inconsistent with observed employment dynamics due to time-aggregation bias or measurement error. To ensure internal consistency, we adjust the separation rate following Shimer (2012). In particular, we use measured employment stocks to back out implied transition rates that satisfy an accounting identity. The employment identity $e_t = (1 - s_t)e_{t-1} + f_t(1 - e_{t-1})$ implies a separation rate

$$s_t = \frac{e_{t-1} + f_t(1 - e_{t-1}) - e_t}{e_{t-1}},$$

where $e_t$ is the employment rate and $f_t$ is the job-finding rate. We use this implied separation rate in estimation rather than the directly measured rate.

## 4 Estimation

Our estimation strategy proceeds in two steps. First, we assign values to some parameters that are well established in the literature or directly observable in the data. Second, we estimate the remaining parameters by matching empirical and model-based profiles.

### 4.1 Mapping Model Variables to the Data

Before describing the estimation procedure, we clarify how model objects correspond to their empirical counterparts. The wage function $w(s_t, m_t)$, which depends on human capital and match quality, maps to the hourly wage rate in the NLSY. Hours worked in the model, $h_t \in [0, 1]$, correspond to the average workweek in the data, which we normalize by 168 hours to ensure the same bounds. The employment rate in the model is computed as the fraction of months worked in a quarter, consistent with the data definition based on weeks worked. In the model, non-employment encompasses both unemployment and non-participation; given our sample restriction to workers with 520 to 5,200 annual hours, this approximation is reasonable. Finally, cumulative jobs in the model are measured by counting transitions – both separations to unemployment and job-to-job moves – and correspond to the number of distinct jobs held in the data, both recorded as end-of-period stocks.

### 4.2 Externally Parameters

We calibrate 15 parameters externally, as summarized in Table 2.

**Demographics and Institutions.** We set the total number of periods to $T = 840$, corresponding to a 70-year life cycle at monthly frequency. Retirement occurs at period $T_R = 540$, implying 45 years of active working life followed by 25 years of retirement. This corresponds to a worker who enters the labor market at age 20 and retires at age 65. The monthly interest rate is set to $r = 0.0025$, equivalent to approximately 3 percent per year. The wage scale $w_0$ is normalized to unity.

Retirement income is determined by a replacement rate $\kappa = 0.18$. This value reflects the product of three components: a 60 percent replacement ratio for high school graduates, consistent with U.S. Social Security; the approximate fraction of total weekly hours devoted to work (roughly 0.3); and our wage normalization. Unemployment benefits follow the same logic: $\bar{b} = 0.12 = 0.4 \times 0.3 \times 1.0$, where the first term is the unemployment insurance replacement rate, approximately 40 percent of prior earnings for this population.

| Parameter | Description | Value |
|---|---|---|
| *Demographics and Institutional* | | |
| $T$ | Total life periods (70 years, monthly) | 840 |
| $T_R$ | Retirement period (45 years of work life) | 540 |
| $r$ | Interest rate (monthly, 3% annual) | 0.0025 |
| $w_0$ | Wage scale normalization | 1.00 |
| $\kappa$ | Replacement rate for retirement income | 0.18 |
| $\bar{b}$ | Unemployment benefit | 0.12 |
| *Preferences* | | |
| $\beta$ | Discount factor (monthly) | 0.994 |
| $\gamma$ | Frisch elasticity of labor supply | 0.55 |
| *Initial Asset Conditions* | | |
| $p_{A_0}$ | Probability of zero initial assets | 0.16 |
| $\mu_{A_0}$ | Mean of log initial assets | 1.26 |
| $\sigma_{A_0}$ | Std. dev. of log initial assets | 0.64 |
| $\bar{a}$ | Borrowing limit | 4.54 |
| *Labor Market* | | |
| $p_{e_0}$ | Initial employment rate | 0.75 |
| $\delta_e$ | Human capital depreciation (employed) | 0.0025 |
| $\rho_{\lambda m}$ | Offer rate elasticity w.r.t. match quality | 0.00 |

Table 2: Externally Calibrated Parameters

**Preferences** The utility of consumption follows a logarithmic specification. The curvature parameter for the disutility of labor is set to $\gamma = 0.55$, consistent with micro estimates of the Frisch elasticity around 0.5 for high school-educated workers.[4] The discount factor is set to $\beta = 0.994$ per month, implying an annual discount factor of approximately $0.93$. This relatively low value implies impatient workers who tend to save less and discount the future heavily.

**Initial Asset Conditions** Initial asset holdings are calibrated from the 1983 wave of the Survey of Consumer Finances (SCF), restricting the sample to high school-educated household heads aged 25 or younger. Net worth is expressed in 1993 dollars and scaled by average monthly entry earnings from the NLSY79, computed as mean hourly wages multiplied by $0.3 \times 168 \times 4.3$, where $0.3 \times 168$ reflects the average weekly hours in the NLSY sample and $4.3$ is the number of weeks per month. The borrowing limit $\bar{a}$ is set to the first percentile of the scaled net worth distribution, yielding $\bar{a} = 4.54$. The initial cross-sectional distribution of assets is

[4] A low value of $\gamma$ implies a highly convex disutility of labor, which helps stabilize hours worked over the life cycle after an initial period of growth. By contrast, other studies estimating life-cycle models with endogenous labor supply and human capital accumulation obtain higher values: Imai and Keane (2004) report $\gamma \approx 3.82$, while Wallenius (2011) find $\gamma \approx 1.13$.

specified as a two-component mixture: a point mass at zero, reflecting the empirical share of households with no assets ($p_{A_0} = 0.16$), and a lognormal distribution for strictly positive holdings with parameters $(\mu_{A_0}, \sigma_{A_0}) = (1.26, 0.64)$ estimated by matching the mean and variance of log assets among positive observations.

**Labor Market** The initial employment rate $p_{e_0}$ is set to match the average employment rate observed in the NLSY79 during the first quarter after graduation. A significant share of respondents report working while still enrolled in high school, so calibrating this rate is necessary to capture the gradual transition into full labor market attachment. Human capital depreciates at a rate $\delta_e = 0.0025$ per month while employed, roughly 3% annually.

We set the elasticity of offer arrival with respect to match quality to $\rho_{\lambda m} = 0$, so that the on-the-job arrival rate depends only on human capital. The parameters governing the baseline separation rate $\pi_0$ and the offer arrival rates $\lambda^U$ and $\lambda^E$ are calibrated internally within the estimation routine to match the separation rate, job-finding rate, and job-to-job transition rate in the first quarter after graduation, as measured in the NLSY. This calibration ensures that the model replicates the initial labor market dynamics before the estimated elasticity and age-decay parameters shape the subsequent life-cycle profiles. We report these parameters in the next table.

### 4.3 Internally Calibrated Parameters

Table 3 reports the 32 parameters estimated using our indirect inference procedure. We estimate the model parameters by minimizing the distance between simulated and empirical moments using a multi-stage global optimization procedure inspired by Arnoud, Guvenen, and Kleineberg (2019). In the first stage, we evaluate the objective function at $2^{12}$ Sobol quasi-random points spanning the parameter space, which provides systematic coverage of the high-dimensional domain. We then select the $2^4$ best-performing points and use each as a starting value for local optimization via a derivative-free algorithm. In a final refinement stage, we form convex combinations of the current best estimate with each remaining local minimum and re-run local optimization from these combined starting points. This strategy guards against convergence to local minima by repeatedly probing the region between competing candidates, and we take the overall minimum across all stages as our point estimate $\hat{\theta}$. When computing the distance between data and model moments for estimation purposes, we use an identity matrix as the weight matrix.

We compute standard errors using a bootstrap-based sandwich estimator. On the data side, we draw 100 bootstrap samples by resampling individuals with replacement within each cohort and recompute all empirical moments for each draw. The covariance matrix of these bootstrap moment vectors yields $\hat{\Omega}$, our estimate of the sampling variance of the data moments. On the model side, let $g(\theta) = \hat{m} - f(\theta)$ denote the vector of moment deviations, where $\hat{m}$ collects the empirical moments and $f(\theta)$ their model-implied counterparts. We approximate the Jacobian $\hat{J} = \partial f(\theta)/\partial\theta'$ by central finite differences, perturbing each parameter by $\pm 10^{-6}$ around $\hat{\theta}$ and re-solving the model at each perturbation. Because we minimize the unweighted sum of squared moment deviations (i.e. $W = I$), the asymptotic variance-covariance matrix of the estimator is given by the classical minimum distance sandwich formula,

$$\widehat{\mathrm{Var}}(\hat{\theta}) \;=\; (\hat{J}'\hat{J})^{-1}\,\hat{J}'\,\hat{\Omega}\,\hat{J}\,(\hat{J}'\hat{J})^{-1},$$

and we report the square roots of its diagonal as standard errors.[5]

[5]The table also reports the transition rate levels $\lambda^U$, $\lambda^E$, and $\pi_0$, which are calibrated within the estimation routine to match the first-quarter transition rates. When computing their standard errors, we use the full sample.

| Parameter | Description | Value | Std. Error |
|---|---|---|---|
| *Initial Distributions* | | | |
| $\mu_{s_0}$ | Mean of log initial human capital | 0.072 | 0.012 |
| $\sigma_{s_0}$ | Log std. dev. of initial human capital | −0.567 | 0.011 |
| $\mu_{m_0}$ | Mean of log initial match quality | 0.415 | 0.005 |
| $\sigma_{m_0}$ | Log std. dev. of initial match quality | −0.469 | 0.021 |
| $\rho_{sm,0}$ | Correlation of initial $\log s_0$ and $\log m_0$ | −0.558 | 0.041 |
| *Labor Market Parameters* | | | |
| $\lambda^U$ | Offer arrival rate (unemployed) | 0.615 | 0.000 |
| $\lambda^E$ | Offer arrival rate (employed) | 1.909 | 0.000 |
| $\pi_0$ | Separation rate intercept | 2.637 | 0.000 |
| $\rho_{\lambda s}$ | Offer rate elasticity w.r.t. human capital | 0.050 | 0.000 |
| $\rho_{\pi s}$ | Separation rate elasticity w.r.t. human capital | 2.258 | 0.000 |
| $\rho_{\pi m}$ | Separation rate elasticity w.r.t. match quality | 0.908 | 0.000 |
| $\rho_{ws}$ | Wage elasticity w.r.t. human capital | 0.301 | 0.000 |
| $\rho_{wm}$ | Wage elasticity w.r.t. match quality | 0.750 | 0.000 |
| *Human Capital Parameters* | | | |
| $\rho_{ss}$ | HC accumulation elasticity w.r.t. human capital | 0.003 | 0.000 |
| $\rho_{sm}$ | HC accumulation elasticity w.r.t. match quality | 0.009 | 0.000 |
| $\rho_{sh}$ | HC accumulation elasticity w.r.t. hours | 0.353 | 0.000 |
| $\delta_u$ | HC depreciation rate (unemployed) | 0.006 | 0.000 |
| $B$ | Scale parameter in $\phi(s)$ | 0.010 | 0.000 |
| $\nu$ | Curvature parameter in $\phi(s)$ | 4.859 | 0.000 |
| *Match Quality Distribution* | | | |
| $\mu_0^E$ | Mean intercept | −0.035 | 0.000 |
| $\mu_1^E$ | Mean slope w.r.t. human capital | 1.102 | 0.000 |
| $\sigma_0^E$ | Log std. dev. intercept | −1.548 | 0.000 |
| $\sigma_1^E$ | Log std. dev. slope w.r.t. human capital | 0.021 | 0.000 |
| $\mu_0^U$ | Mean intercept | 0.153 | 0.000 |
| $\mu_1^U$ | Mean slope w.r.t. human capital | −0.097 | 0.000 |
| $\sigma_0^U$ | Log std. dev. intercept | −2.397 | 0.000 |
| $\sigma_1^U$ | Log std. dev. slope w.r.t. human capital | −0.409 | 0.000 |
| *Type Heterogeneity* | | | |
| $\mu_\xi$ | Mean labor disutility (levels) | 78.02 | 0.000 |
| $\sigma_\xi$ | Log std. dev. of stationary $\log \xi$ | −0.351 | 0.000 |
| $\rho_\xi$ | Persistence of $\log \xi$ | 0.996 | 0.000 |
| $\sigma_{\xi,0}$ | Log std. dev. of initial $\log \xi$ | −1.029 | *n.i.* |
| $C_0^\alpha$ | Mean $\alpha$ in levels | 1.869 | 0.000 |
| $C_1^\alpha$ | Log std. dev. of $\alpha$ | −1.327 | 0.000 |
| *Age Effects* | | | |
| $\zeta_\xi$ | Age decay in labor disutility | −0.0003 | 0.000 |
| $\zeta_\lambda$ | Age decay in offer arrival rate | 0.0010 | 0.000 |
| $\zeta_\pi$ | Age decay in separation rate | −0.0010 | 0.000 |
| $\zeta_w$ | Age decay in wages | −0.0002 | 0.000 |
| $\zeta_\alpha$ | Age decay in HC accumulation | 0.0000 | 0.000 |

Table 3: Internally Calibrated Parameters

**Initial Distributions** The parameters $\mu_{s_0}$, $\sigma_{s_0}$, $\mu_{m_0}$, $\sigma_{m_0}$, and $\rho_{sm,0}$ characterize the joint distribution of human capital and match quality at labor market entry. The negative correlation ($\rho_{sm,0} = -0.56$) implies that workers with high human capital tend to enter lower-quality matches. These parameters are mostly identified by the shapes of the first observations from the wage and hours profiles. Importantly, the negative correlation is identified by the high log-workweek dispersion observed early in workers' careers.

**Labor Market Parameters** The intercepts $\lambda^U$, $\lambda^E$, and $\pi_0$ are calibrated within the estimation routine to match the first-quarter transition rates observed in the data. They need to be calibrated together with other parameters, since they depend on the initial distributions of match quality and the human capital stock, as well as the elasticity parameters.

The elasticities $\rho_{\lambda s}$, $\rho_{\pi s}$, and $\rho_{\pi m}$ govern how transition rates respond to human capital and match quality. The separation rate is highly sensitive to human capital ($\rho_{\pi s} = 2.26$) and match quality ($\rho_{\pi m} = 0.91$), implying that more experienced workers and better matches face substantially lower separation risk. The offer arrival rate is almost independent of human capital ($\rho_{\lambda s} = 0.05$). These parameters are mostly identified by the shape of the transition profiles.

The elasticities $\rho_{ws}$ and $\rho_{wm}$ govern how wages respond to human capital and match quality. Wages load on both human capital ($\rho_{ws} = 0.30$) and match quality ($\rho_{wm} = 0.75$), with match quality playing a larger role. These parameters are mostly identified by the shapes of the profiles of mean wages and cross-sectional dispersion of wages.

**Human Capital Accumulation** The elasticities $\rho_{ss}$, $\rho_{sm}$, and $\rho_{sh}$ shape the human capital production function. The elasticity with respect to the stock ($\rho_{ss} = 0.00$) generates almost no increasing returns of human capital accumulation as workers gain experience. The positive elasticity with respect to hours ($\rho_{sh} = 0.35$) implies that working more accelerates skill acquisition, providing an investment motive for early-career labor supply. Match quality plays a minor role ($\rho_{sm} = 0.01$). The depreciation rate while unemployed ($\delta_u = 0.006$) exceeds that while employed, penalizing non-employment spells. These parameters are mostly identified by the wage profile, but they also influence the hours profiles.

The scale $B$ and curvature $\nu$ in the loss function $\phi(s) = Bs^\nu$ determine how much human capital is forfeited upon job transitions. The high curvature ($\nu = 4.86$)

implies that losses are convex in the stock: workers with more accumulated experience have more to lose. This generates endogenous lock-in, reducing job mobility as tenure grows and helping explain why job-to-job transitions decline with age. These parameters are mostly identified by the job-to-job transition rates and the number of jobs held.

**Match Quality Distributions** The parameters governing match quality offers differ between employed and unemployed workers. For employed workers, the mean of the offer distribution rises steeply with human capital ($\mu_1^E = 1.10$), so that skilled workers draw from a more favorable distribution. For unemployed workers, the dependence on human capital is now negative ($\mu_1^U = -0.10$), and the distribution is more compressed ($\sigma_0^U = -2.40$ versus $\sigma_0^E = -1.55$, both in logs). These parameters are mostly identified by the dispersion profiles.

**Type Heterogeneity** The disutility of labor $\xi_t$ is highly persistent ($\rho_\xi = 0.996$) and fans out over the life cycle: initial dispersion ($\sigma_{\xi,0} = -1.03$ in logs) is smaller than stationary dispersion ($\sigma_\xi = -0.35$ in logs). The parameters that map the accumulation of human capital imply that it is negatively correlated with the disutility of labor. These parameters are identified by the dispersion in hours.

**Age Effects** To improve the model's fit, we allow for age-decay parameters that capture residual life-cycle variation in transition rates not explained by human capital and match quality. The decay operates multiplicatively. For example, the transition rate at age $t$ is scaled by $(1+\zeta)^t$ relative to its baseline value. All parameters are small, implying small residual life-cycle variation not captured by the human capital and match quality.

### 4.4 Targeted Moments: Model Fit

Figure 2 compares the model-implied and data profiles for average outcomes. Panels (2a) and (2b) present the employment rate and average workweek. The data show a rise in employment rates and an increase in the average workweek during the initial years of employment, before reaching a plateau. The model matches both patterns relatively well, capturing the rise and the plateau after some years in both measures. The rise in workweek is steeper in the data than in the model. Panel (2c) presents the average wage profile. The data shows a continuous upward trend that the model successfully reproduces. The model overestimates wages later in the life cycle by 50 percentage points. Finally, panel (2d) presents the number of jobs held over time. The data shows high job mobility in the early years of the labor market, with frequent job changes occurring within the first 15 years. After this period, job mobility decreases as workers transition to longer-term positions. The model replicates this pattern, capturing the early-career job changes followed by a decline in job switching, but it also overestimates the number of jobs later in the life cycle.

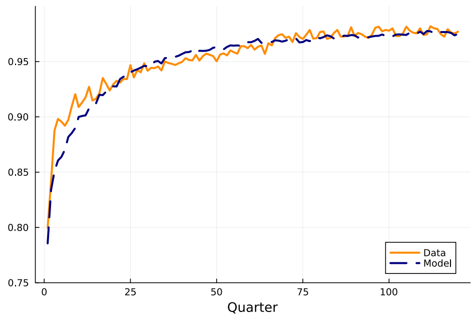


(a) Employment Rate

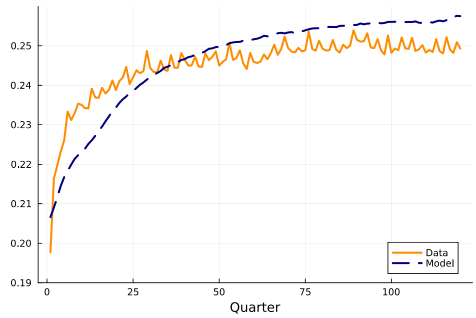


(b) Average Workweek

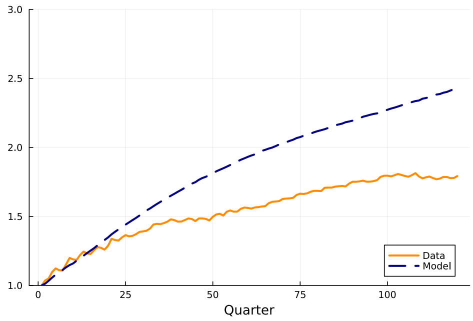


(c) Average Wage

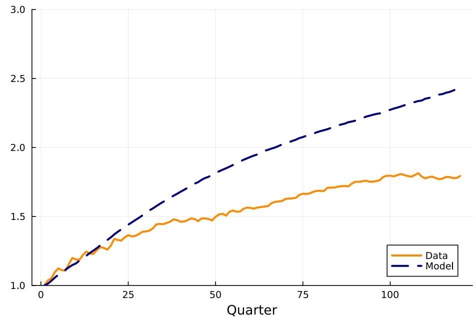


(d) Number of Jobs

Figure 2: Comparison of Model-Implied and Data Profiles

Figure 3 compares the model-implied and data profiles for dispersion (standard deviations) of outcomes. Panels (3a) and (3b) report the cross-sectional standard deviation of the employment rate and log hours, respectively. In the data, dispersion for employment rates is larger in the early labor-market years and then stabilizes as workers sort into more persistent employment. However, the workweek displays a U-shaped pattern, as noted by Kaplan (2012). It is high early on, it decreases markedly in the first years, and starts to rise again. Importantly, the heterogeneity in preferences and its dynamics are crucial for generating this moment. Without it, hours dispersion would decrease markedly over the life-cycle.

Panel (3c) shows the standard deviation of wages, which increases over the life cycle in the data. The model also produces rising wage dispersion, consistent with mechanisms of human capital accumulation and job changes. Finally, panel

(3d) reports dispersion in the number of jobs held. The data exhibit substantial heterogeneity in early-career mobility: some workers switch frequently while others settle quickly. The dispersion in the number of jobs held continues to increase over the life cycle.

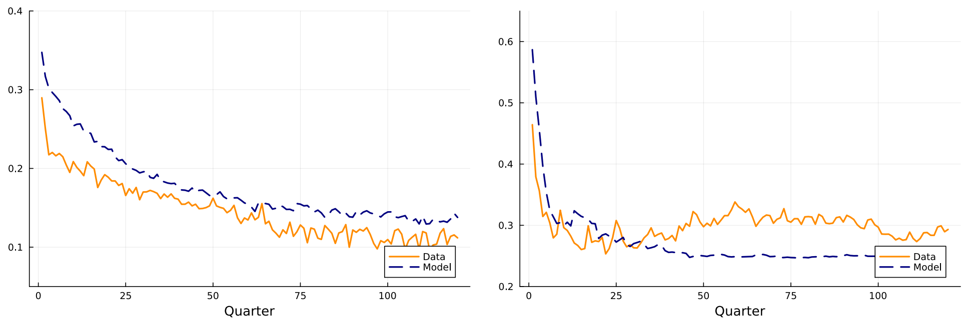


(a) Employment Rate (Std) (b) Average Workweek (Std)

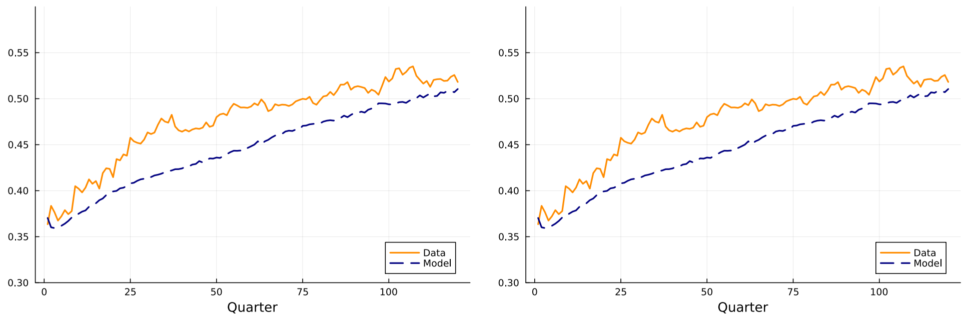


(c) Average Wage (Std) (d) Number of Jobs (Std)

Figure 3: Comparison of Model-Implied and Data Profiles (Standard Deviations)

Finally, figure 4 reports the life-cycle profiles of key labor-market transition rates. The job-to-job (EE) transition rate is highest early on and then falls steadily, consistent with rapid early-career reallocation followed by greater job stability; the model captures the downward slope reasonably well, although it understates the very high level of EE transitions at labor-market entry. The job-to-unemployment (EU) separation rate declines with experience in both the model and the data, and overall, the model tracks the observed decline in separations. The data exhibit some spikes due to adjustments that make them consistent with other series. Finally, the unemployment-to-job (UE) rate declines over time in the data and in the model, indicating that job finding becomes less likely later in the working life.

### 4.5 Untargeted Moments: Decomposing Hours into Components

Hours worked over the life cycle can rise either because individuals work more weeks per year (the extensive margin) or because they work longer weeks (the intensive margin). Many models of labor supply emphasize the extensive margin,

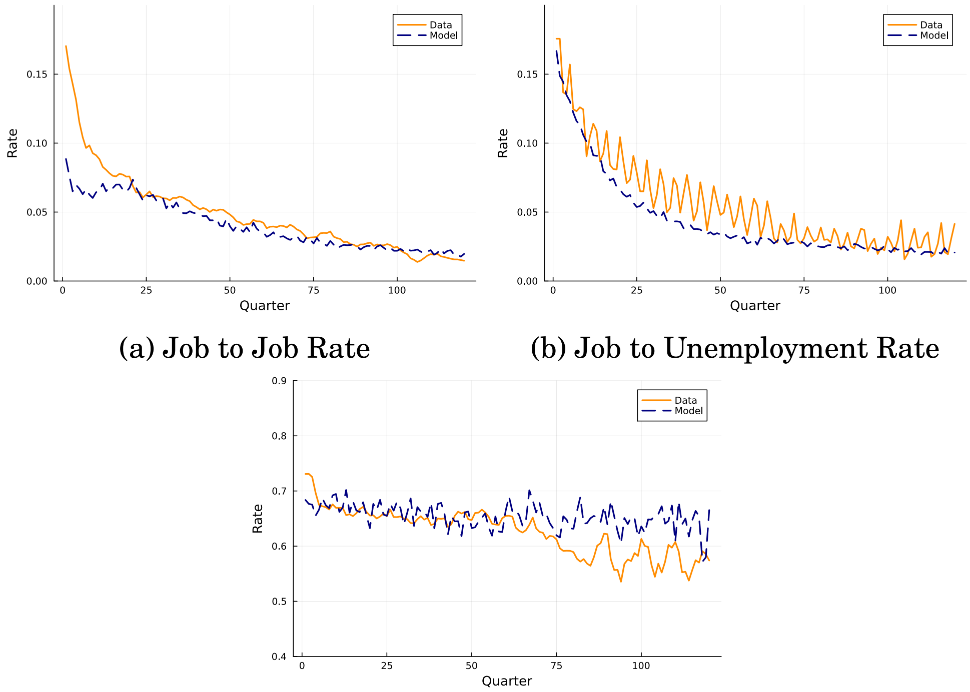


(a) Job to Job Rate (b) Job to Unemployment Rate

(c) Unemployment to Job Rate

Figure 4: Comparison of Model-Implied and Data Profiles (Standard Deviations)

focusing on participation as the main source of hours dynamics. In contrast, we show that the intensive margin dominates: most of the increase in hours comes from longer workweeks rather than additional weeks of employment. We apply our proposed decomposition to our quantitative life-cycle model and to the NLSY data. Importantly, this decomposition is treated as an untarget moment.

**Intensive vs. Extensive Margins** Let $H_t$ denote the mean hours worked $t$ quarters after labor-market entry. With $S_t$ the set of individuals observed in $t$ and $N_t = |S_t|$ the number of individuals in $S_t$, then

$$H_t = \frac{1}{N_t} \sum_{i \in S_t} H_{it} \ . \tag{1}$$

Equivalently, equation (1) can be written as a weighted average,

$$H_t = \sum_{i \in S_t} \omega_{it} H_{it} \ , \quad \omega_{it} = \frac{1}{N_t} \ .$$

Individual hours are the product of the average workweek and the number of weeks worked:

$$H_{it} = h_{it} \, n_{it} \ ,$$

where $h_{it}$ is average hours per week and $n_{it}$ is weeks worked. We compute $h_{it} = H_{it}/n_{it}$ if $n_{it} > 0$ and $h_{it} = 0$ otherwise. Following the literature, we include only workers with positive hours worked.

The life-cycle hours profile $\{H_t\}_{t=0}^{T}$ is the sequence of mean hours worked from entry through period $T$. To quantify the contributions of the intensive and extensive margins, we decompose changes in adjacent mean hours using a method similar to that of Baily et al. (1992) and Blundell et al. (2011). Specifically, the change between $t-1$ and $t$ can be written as

$$\begin{aligned} H_t - H_{t-1} = & \underbrace{\sum_{i \in S_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1}}_{\text{intensive margin}} \\ & + \underbrace{\sum_{i \in S_{t \cap t-1}} \omega_{it} h_{it-1} \left( n_{it} - n_{it-1} \right)}_{\text{extensive margin (1): within-continuers}} \\ & + \underbrace{\sum_{i \in S_{t/t \cap t-1}} \omega_{it} h_{it} n_{it} - \sum_{i \in S_{t-1/t \cap t-1}} \omega_{it-1} h_{it-1} n_{it-1} + \sum_{i \in S_{t \cap t-1}} \left( \omega_{it} - \omega_{it-1} \right) h_{it-1} n_{it-1}}_{\text{extensive margin (2): composition effect}} \\ & + \underbrace{\sum_{i \in S_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) \left( n_{it} - n_{it-1} \right)}_{\text{quadratic term}} . \end{aligned} \tag{2}$$

The first, second, and fourth terms use individuals observed in both periods, which we call *continuers* $S_{t \cap t-1} = \{i \mid i \in S_t \cap S_{t-1}\}$. The third term reflects sample turnover and uses *entrants* and *leavers*: individuals who appear in $t$ but not in $t-1$, and individuals who appear in $t-1$ but not in $t$. These groups are represented by $S_{t-1/t \cap t-1}$ and $S_{t/t \cap t-1}$, respectively. Appendix B.1 provides the step-by-step derivation.[6]

The first term captures the intensive margin: the change in total hours due to variation in the workweek, holding weeks worked fixed at the $t-1$ level. The second term captures the within-continuers component of the extensive margin: the change in total hours due to variation in weeks worked, holding the workweek fixed at the $t-1$ level. The third term captures the composition effect: the change in total hours due to entry and exit from the sample and the resulting reweighting of continuers. The fourth term captures the interaction effect: the change in total

[6]Our decomposition can be interpreted as a discrete-time approximation to the total derivative, measuring how total hours would change if either the intensive or extensive margin were held constant. This approach is analogous to a Laspeyres index, in which changes are evaluated relative to a base period. Because it is a discrete approximation, second-order interaction terms do not cancel and must be explicitly included.

hours due to simultaneous changes in weeks worked and the workweek.

This decomposition is important because it directly reveals the relative importance of the extensive and intensive margins and thereby informs modeling choices. If the extensive margin were the primary driver of life-cycle hours growth, the intensive margin contribution would be negligible, and models of labor supply could focus exclusively on participation and weeks worked. Fan et al. (2024) is a recent paper that takes this route. By contrast, if the intensive margin is central, then explicitly modeling variation in the workweek becomes essential.

Finally, we track how hours accumulate over the life cycle by adding up the changes between each period:

$$H_T - H_0 = \sum_{t=1}^{T} (H_t - H_{t-1}) \ . \tag{3}$$

This gives the total increase in hours worked from entry to period $T$. We can also express this as a growth rate:

$$\frac{H_T - H_0}{H_0} = \frac{\sum_{t=1}^{T} (H_t - H_{t-1})}{H_0} \ . \tag{4}$$

Each component of the decomposition (2) – intensive, extensive, composition, and interaction – can be summed in the same way. Adding them up over all periods shows how much each channel contributes to the overall rise in hours worked across the life cycle.

**Implementation of the Decomposition** Figure 5 summarizes the main finding. It plots the contributions of the intensive, extensive (weeks), extensive (composition), and quadratic margins in green, orange, blue, and purple, respectively. The intensive margin accounts for nearly 100 additional hours over the ten years after graduation, while the extensive margin accounts for about 60 hours. The quadratic term's contribution is negligible. The model reproduces the same pattern, even though it misses some of the hours growth due to missing hours and to early-on employment rate growth in the labor market. Interestingly, the model captures the negligible impact of the quadratic term and the composition effect. The results show that life-cycle hours growth is driven primarily by longer workweeks rather than by greater participation in the data or the model. We investigate the mechanism behind the growth in the workweek in the next section.[7]

[7] All results use the same sample as in Section A.2, pooling the NLSY79 and NLSY97 cohorts.

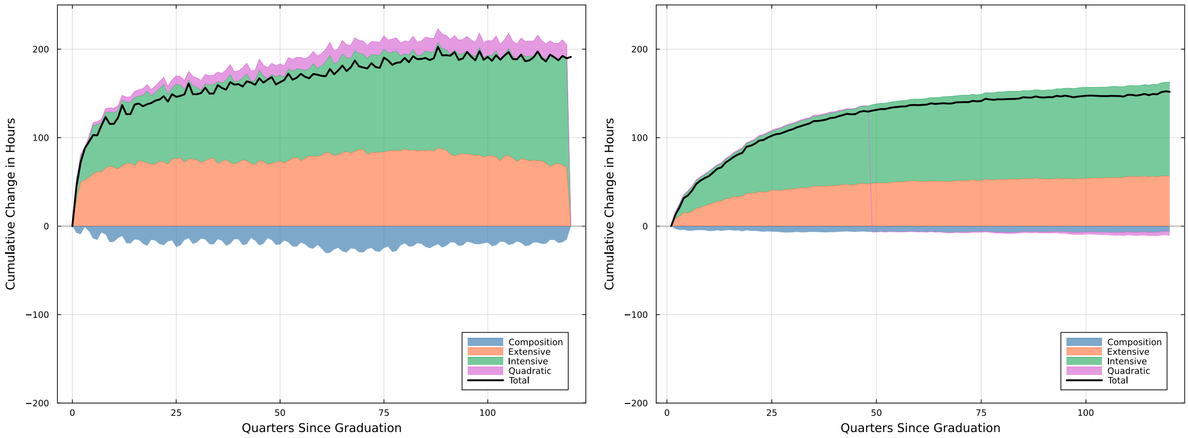


(a) Data Decomposition

(b) Model Decomposition

Figure 5: Intensive vs. Extensive Margin Decomposition

Table 4 helps see the magnitude. It shows the percentage contribution of each margin to cumulative growth in hours worked. We present results across different horizons: 5, 10, 15, 20, and 25 years after graduation. Row 1 reports that, five years after graduation, the intensive margin accounts for 57% of the increase in hours worked. The intensive margin is measured by the increase in the average workweek. By contrast, the extensive margin accounts for 38%, split into a 49% contribution from additional weeks worked among continuers and a –12% contribution from changes in the composition of workers in the sample. The quadratic term plays only a minor role, contributing less than –5%. Observe that, even though it is not a target, the model replicates the relative importance of each margin well. The contribution of the intensive margin increases with years of labor market experience. Fifteen years after graduation, the intensive margin accounts for 65% of the cumulative increase in hours worked, while the extensive margin accounts for 35%. Again, this dominance of the intensive margin implies that longer workweeks, not greater participation, drive the rise in hours worked after graduation.

## 5 Quantitative Experiments

In this section, we investigate the roles of human capital accumulation, job mobility, and preference heterogeneity in shaping wages and hours over the life cycle. In each exercise, we shut down a specific mechanism while keeping all others active. This approach quantifies how much of the baseline outcomes is driven by each channel. We address two questions: (i) how much of wage and hours growth is explained by each channel and their interaction? And (ii) how important is match-specific

Table 4: Contributions of Intensive and Extensive Margins: Data vs. Model

| | | Data | | |
|---|---|---|---|---|
| Years | Intensive | Extensive (1) | Extensive (2) | Quadratic |
| 5 | 57.0% | 49.4% | -11.6% | 5.2% |
| 10 | 57.3% | 45.0% | -9.4% | 7.1% |
| 15 | 59.6% | 47.1% | -14.9% | 8.2% |
| 20 | 56.5% | 45.9% | -10.2% | 7.8% |
| | | Model | | |
| 5 | 60.3% | 38.6% | -5.6% | 1.0% |
| 10 | 63.7% | 35.4% | -6.3% | 0.8% |
| 15 | 64.9% | 35.4% | -4.0% | -0.2% |
| 20 | 65.7% | 35.2% | -4.6% | -0.9% |

Note: Extensive (1) refers to within-continuers margin. Extensive (2) refers to composition margin. All values represent percentage contributions to cumulative growth up to year X.

human capital?

### 5.1 Human Capital, Job Mobility, Preference Heterogeneity

In our first exercise, we compare four parametrizations of our economy: (i) the baseline one, (ii) one without human capital, (iii) one without job search, and (iv) one without preference heterogeneity. In the figures, the baseline is the blue dashed line; the parametrization without human capital is the dotted red line; the parametrization without job mobility is the dashed-dotted green line; and the parametrization without preference heterogeneity is the solid purple line. We restrict ourselves to showing only the wage and total hours profiles, specifically their means and standard deviations. Total hours is the product of the number of weeks worked and the workweek.

**No Job Mobility** We parametrize an economy in which workers can accumulate human capital and exhibit preference heterogeneity, but cannot change jobs. For that, we set the initial employment rate to 1, the offer arrival rates to essentially 0 ($\lambda^U = \lambda^E \to 0$), and the separation rate to 0 ($\pi \to 0$). All elasticities linking job flows to human capital and match quality are also set to 0 ($\rho_{\lambda s} = \rho_{\lambda m} = \rho_{\pi s} = \rho_{\pi m} = 0$). In this economy, workers are permanently employed in their initial match and can only increase wages through human capital accumulation.

**No Human Capital** We parametrize an economy in which workers search for better matches and exhibit preference heterogeneity, but cannot accumulate human capital. For that, we set the human capital learning efficiency to 0 ($C_0^{\alpha} \to \infty$ and $C_1^{\alpha} = 0$), the accumulation elasticity to hours, human capital, and match quality to 0 ($\rho_{sh} = \rho_{ss} = \rho_{sm} = 0$), the depreciation rates to 0 ($\delta_e = \delta_u = 0$), and the switching costs to 0 ($B = \nu = 0$). In this economy, initial human capital is actually permanent, and workers can only increase their earnings through job search. We do not zero out the dependence of labor market flows on human capital and match quality, so that, in this economy, workers will experience differences in arrival and separation rates due to their initial, permanent human capital and current match quality.

**No Preference Heterogeneity** We parametrize an economy in which workers have common preferences. Workers are still accumulating human capital, making transitions, and losing jobs, but they do not differ in their preferences. For that, we set the initial and stationary variance of $\log \xi$ to 0 ($\sigma_{\xi} \to -\infty$ and $\sigma_{\xi,0} \to -\infty$, parameters representing the log of std. dev.) and the persistence of $\log \xi$ to 1 ($\rho_{\xi} \to 1$).

**Results** Figure 6 presents the results. We begin by examining wage dynamics. Panel (a) shows that the baseline parameterization yields the largest wage growth, while the parameterization without human capital accumulation yields the lowest. Human capital is thus essential to generate the observed wage growth profile. The intermediate position of the no-mobility parameterization reveals that job-to-job transitions and the effect of human capital in the match distribution amplify wage growth significantly.

Panel (c) examines wage dispersion. Both counterfactuals generate rapid dispersion growth, but for different reasons. Without human capital accumulation, workers face no cost to climbing the job ladder, leading to fast sorting across match qualities. Without job mobility, workers face no loss of firm-specific capital from transitions, leading to fast skill accumulation without the baseline model's friction. The baseline parameterization yields slower dispersion growth because the prospect of accumulating human capital that is partially lost upon transition slows both job ladder climbing and skill investment. Preference heterogeneity becomes increasingly important for wage dispersion among mid-career workers.

We now turn to the dynamics of hours. Panel (b) shows that the baseline parameterization and the no-mobility parameterization yield nearly identical hours

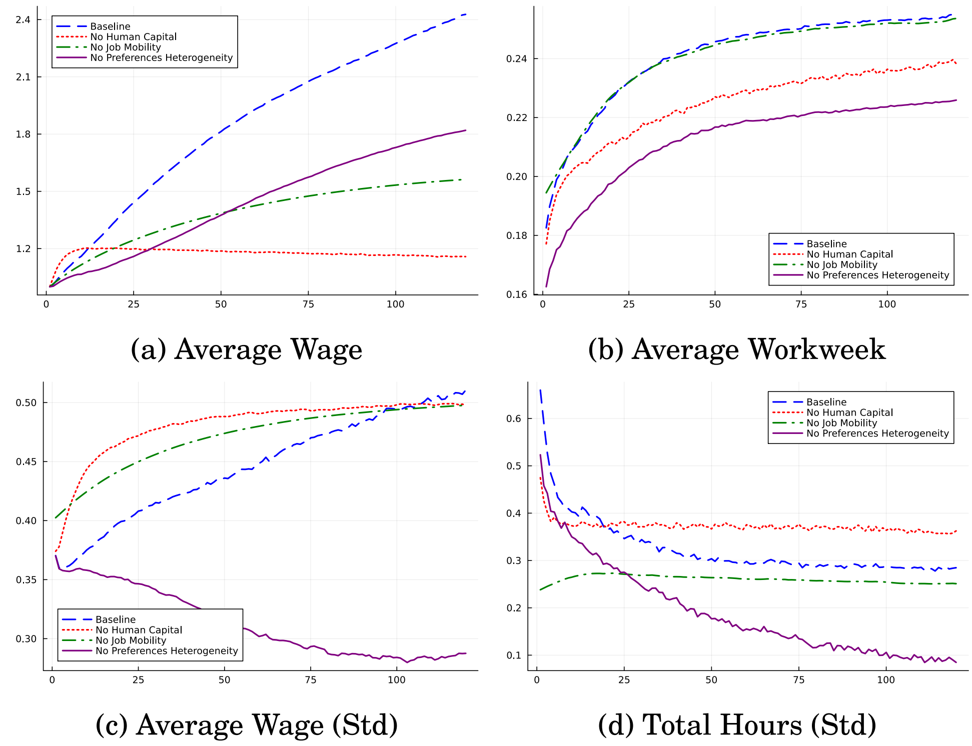


(a) Average Wage (b) Average Workweek

(c) Average Wage (Std) (d) Total Hours (Std)

Figure 6: Comparison of Human Capital and Job Search Models

profiles, with one key exception: in the first years after labor market entry, the no-mobility model predicts higher labor supply. This difference reflects the baseline model's forward-looking investment behavior, where workers who anticipate future job transitions reduce their early-career hours to avoid overinvesting in firm-specific human capital they will partially lose. The no-human-capital parameterization predicts uniformly lower hours since labor supply lacks an investment component.

Panel (d) examines hours dispersion. The early-career spike in hours dispersion is driven by mobility frictions: when we eliminate job mobility, this early dispersion disappears. This result is consistent with the early career wedges documented by Kaplan (2012). The baseline model generates this pattern because workers with good early matches invest heavily in hours, while those with poor matches invest less in anticipation of a transition. As with wages, preference heterogeneity becomes increasingly important for the dispersion of hours among mid-career workers.

**Decomposing Labor Supply Incentives** To examine how human capital, job mobility, and preference heterogeneity shape labor supply incentives over the life cycle, we compare the ratio of the marginal rate of substitution (MRS) between leisure and consumption to the observed wage across all model parametrizations.

For each employed agent, we compute

$$\frac{\text{MRS}_{it}}{w_{it}} = \frac{\xi_{it} \cdot h_{it}^{1/\gamma} \cdot c_{it}}{w(s_{it}, m_{it})},$$

and report cross-sectional averages by age. This ratio captures the full wedge between the observed wage and the shadow value of time. In a standard life-cycle model without frictions, this ratio equals one: workers equate their marginal disutility of labor to the wage. When human capital accumulation is present, the ratio exceeds one because working today raises future earnings, so agents are willing to supply labor even when the current wage falls short of their MRS. This mechanism is stressed by Imai and Keane (2004). However, since our model also features borrowing constraints, the ratio additionally reflects a consumption-smoothing wedge: when constraints bind, consumption is suppressed below the unconstrained path, which lowers the measured MRS relative to the human capital wedge alone. An increasing MRS/wage profile signals that borrowing constraints tighten over the life cycle relative to income growth, or equivalently, that agents would prefer to borrow against future income gains but cannot.

First, Figure 7 shows that the ratio exceeds one in all parametrizations, indicating that agents face significant wedges between the current wage and the shadow value of time. All profiles also exhibit an increasing region, indicating that borrowing constraints bind. In our model, the borrowing limit equals the present value of unemployment insurance, which is tight for agents with high expected wage growth.

Second, the ranking across parametrizations is informative. The baseline model displays the highest ratio, followed by the model without job mobility, then the model without human capital. This ordering confirms that human capital is an important amplifier of the value of working: by supplying labor today, workers not only raise future wages but also improve their future labor market outcomes. Comparing the baseline to the no-mobility economy isolates this latter channel. In the baseline, working brings job stability and better matches in addition to wage growth, whereas in the no-mobility economy, working raises wages only and has no effect on the current match, which is fixed. The gap between these two curves measures the contribution of endogenous job mobility to labor supply incentives. This is an additional channel absent from Imai and Keane (2004), in which the return to working operates exclusively through wages. Finally, the model without preference heterogeneity closely tracks the baseline, suggesting that heterogeneous preferences play a limited role in shaping the average MRS/wage profile.

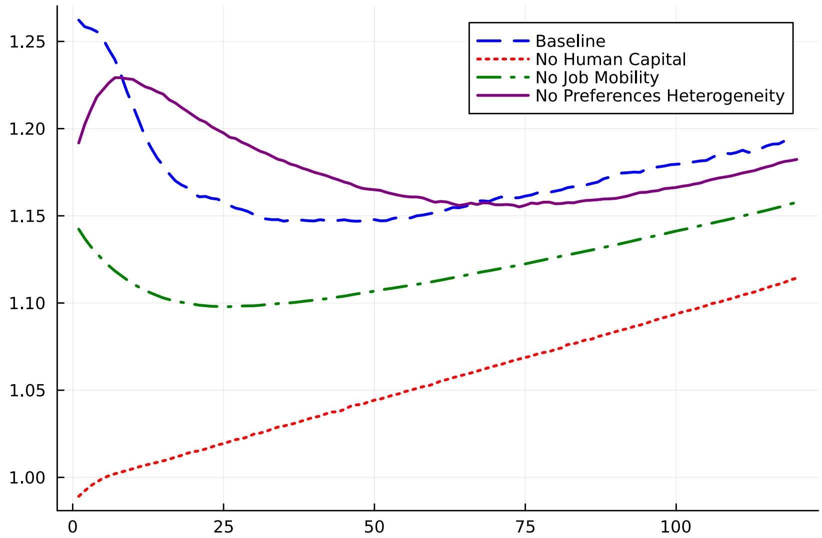


Figure 7: Marginal Rate of Substitution across Models

Note: The figure plots the marginal rate of substitution (MRS) between leisure and consumption against wages for three models. The baseline model (blue) includes both human capital accumulation and job mobility. The human capital only model (red) shuts down job mobility. The job mobility only model (green) shuts down human capital accumulation. A declining MRS indicates that workers supply labor beyond what current wages justify, reflecting the dynamic investment motive.

## 5.2 The Role of Match-Specific Human Capital

Our second exercise evaluates the importance of match-specific human capital. In the baseline model, workers lose a fraction of their accumulated skills when changing jobs, captured by the switching cost $\phi(s) = Bs^{\nu}$ with $\nu > 1$. We simulate a counterfactual economy in which all human capital is general and fully transferable across jobs by setting $B = \nu = 0$.

Figure 8 presents the results. Eliminating switching costs fundamentally alters the investment-mobility trade-off. In the baseline model, workers moderate their skill investment in poor matches to avoid accumulating capital they will lose upon transition. With general human capital, this friction disappears: workers invest aggressively in all matches and transition freely up the job ladder, carrying their full accumulated skills to each new employer.

This change produces dramatically counterfactual wage dynamics. Panel (a) shows that general human capital generates wage growth nearly three times as large as in the baseline parameterization. Panel (c) shows the consequences for wage dispersion: the standard deviation of log wages rises by 0.3 points above the baseline by year 20, as workers sort rapidly across the match quality distribution without the friction of potential human capital loss constraining their transitions.

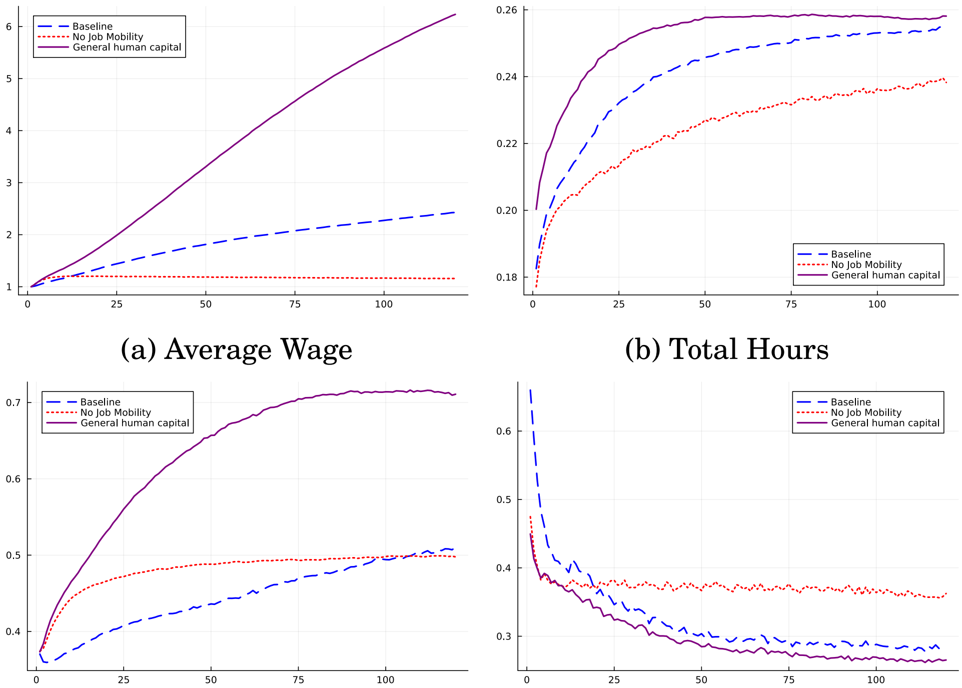


(a) Average Wage (b) Total Hours

(c) Average Wage (Std) (d) Total Hours (Std)

Figure 8: Comparison of Human Capital and Job Search Models

The hours dynamics tell a complementary story. Panel (b) shows that general human capital predicts uniformly higher hours throughout the career. Workers invest more aggressively because the return to skill accumulation is higher. Panel (d) shows that the early-career spike in hours dispersion largely disappears. In the baseline model, this spike arises because workers in poor matches invest less, anticipating separation and human capital loss. With general human capital, all workers invest heavily regardless of current match quality, compressing the hours distribution.

These results demonstrate that match-specific human capital is essential for matching the observed wage and hours dynamics. The friction created by imperfect transferability disciplines both investment behavior and mobility decisions, generating the gradual wage growth and early-career hours heterogeneity observed in the data.

## 6 Implications for Wage and Earnings Inequality

We use the quantitative model to decompose the cross-sectional variance of log wages and log earnings into contributions from human capital, match quality, and hours worked. We then construct the percentage contribution of each component to total variance and to cumulative variance growth. An important exercise we perform is computing this decomposition for different base periods, and we show that the decomposition results differ when measurement begins years after labor market entry, stressing why the early-career period cannot be ignored. This result reinforces our message that the common practice of excluding this period obscures the mechanisms driving life-cycle dynamics.

Importantly, this decomposition was not targeted in our calibration procedure, making it a demanding test of the model's ability to replicate the sources of earnings inequality observed in the data.

**Earnings Variance Decomposition** We first decompose log earnings into contributions from wages and hours worked, which is particularly convenient since it can also be done in the data. In the data, we construct earnings by computing, for each individual, labor earnings as hours worked times the wage rate, with the former reflecting both worked weeks and the average workweek. In the model, we construct earnings analogously. Then, for each year $t$, we compute the cross-sectional variance and covariance of log wages and log hours among the employed.

Formally, earnings are the product of wages and hours worked, so log earnings among employed workers satisfy:

$$\log e_t = \log w_t + \log h_t.$$

The variance of log earnings decomposes as:

$$\mathbf{Var}(\log e_t) = \mathbf{Var}(\log w_t) + \mathbf{Var}(\log h_t) + 2\,\mathbf{Cov}(\log w_t, \log h_t).$$

This separates the contribution of wage inequality from the contribution of hours inequality. The covariance term captures whether high-wage workers also work longer hours, which would amplify earnings inequality.

Table 5 presents the decomposition results for both the model and the data. Despite not targeting these moments, the model captures the quantitative patterns remarkably well. In the data, wage dispersion is the primary driver, accounting for 65-73% of earnings variance over the first 20 years. The model generates wage contributions of 47 to 55 percent, somewhat lower but of the same order of magnitude. Hours dispersion accounts for 20 to 27 percent of the data and 15 to 33 percent of the model, with both showing a declining pattern over the life cycle.

The covariance between wages and hours is positive in both the model and the data, indicating that high-wage workers tend to work longer hours, thereby amplifying earnings inequality. In the data, this covariance accounts for 8 to 9 percent of total variance, while the model generates a larger contribution of 20 to 30 percent. The model thus overstates the strength of the wage-hours complementarity, though it correctly captures the qualitative pattern.

Table 5: Earnings Variance Decomposition: Levels

| | Model | | | | Data (NLSY79) | | | |
|---|---|---|---|---|---|---|---|---|
| Years | Wages | Hours | Cov$(w, h)$ | Total | Wages | Hours | Cov$(w, h)$ | Total |
| 0 | 47.1% | 33.0% | 20.0% | 100% | 64.6% | 26.6% | 8.8% | 100% |
| 5 | 49.5% | 20.9% | 29.7% | 100% | 72.8% | 19.7% | 7.5% | 100% |
| 10 | 51.9% | 17.2% | 30.9% | 100% | 72.2% | 19.8% | 8.0% | 100% |
| 20 | 55.3% | 14.8% | 29.9% | 100% | 69.3% | 22.1% | 8.7% | 100% |

*Note:* The table reports percentage contributions to the cross-sectional variance of log earnings among employed workers. Earnings variance is decomposed as Var$(\log e)$ = Var$(\log w)$ + Var$(\log h)$ + 2Cov$(\log w, \log h)$. Data results are from NLSY79 annual data.

Table 6 shows the decomposition of earnings variance growth, measured as changes relative to the base year. In the data, wage dispersion accounts for an extremely large share of the growth in variance early in the career (229% at year 5), declining to 94% by year 20. The model generates more moderate but still substantial wage contributions of 59 to 70% that rise over time. Hours dispersion contributes negatively in both model and data, indicating that hours variance actually declines over the career. In the data, this negative contribution is substantial early on (- 111% at year 5), before converging toward zero. The model generates smaller but persistent negative contributions (-18 to -30%), correctly capturing the qualitative pattern.

The covariance term shows the largest discrepancy between the model and the data. In the data, the covariance initially contributes negatively (-18% at year 5), then turns slightly positive (6%) by year 20. In contrast, the model generates large positive contributions throughout (48-70%). This reflects the model's mechanism, in which the wage-hours relationship becomes increasingly positive as careers progress: high-wage workers work more hours while low-wage workers work fewer, amplifying earnings inequality. It is also important that the model needs a negative correlation between initial human capital and initial match quality.

Table 6: Earnings Variance Decomposition: Changes

| | Model | | | | Data (NLSY79) | | | |
|---|---|---|---|---|---|---|---|---|
| Years | Wages | Hours | Cov$(w,h)$ | Total | Wages | Hours | Cov$(w,h)$ | Total |
| 5 | 59.4% | −29.4% | 70.0% | 100% | 229.4% | −111.3% | −18.1% | 100% |
| 10 | 66.4% | −29.7% | 63.3% | 100% | 139.9% | −38.2% | −1.7% | 100% |
| 20 | 70.4% | −18.3% | 47.9% | 100% | 93.5% | 0.9% | 5.6% | 100% |

*Note:* The table reports percentage contributions to the change in cross-sectional variance of log earnings relative to the base year (year 0 for model, year 1 for data). The decomposition is $\Delta\text{Var}(\log e) = \Delta\text{Var}(\log w) + \Delta\text{Var}(\log h) + 2\Delta\text{Cov}(\log w, \log h)$. Data results are from NLSY79 annual data.

**Wages Variance Decomposition** Given the importance of wages for earnings inequality, we proceed by decomposing the cross-sectional variance of log wages into contributions from human capital and match quality. The wage function $w(s_t, m_t) = w_0\, s_t^{\rho_{ws}}\, m_t^{\rho_{wm}}$ implies that log wages are linear in log human capital and log match quality:

$$\log w_t = \log w_0 + \rho_{ws} \log s_t + \rho_{wm} \log m_t.$$

Taking cross-sectional variances among employed workers yields the decomposition:

$$\text{Var}(\log w_t) = \rho_{ws}^2\, \text{Var}(\log s_t) + \rho_{wm}^2\, \text{Var}(\log m_t) + 2\rho_{ws}\rho_{wm}\, \text{Cov}(\log s_t, \log m_t).$$

The first term captures the contribution of human capital dispersion, the second captures the contribution of match quality dispersion, and the third captures their interaction. We again express each component as a share of total variance and track how these shares evolve over the life cycle.

Table 7 reports the decomposition of wage variance at different horizons after labor market entry. The model generates quantitatively reasonable patterns across all components. Match quality is the dominant source of wage dispersion at labor market entry, accounting for 74% of wage variance. However, its relative importance declines over time as human capital dispersion grows: by 20 years, match quality accounts for 44 percent, while human capital accounts for 22%. The positive and growing contribution of the covariance term (13 to 34%) reflects endogenous sorting: workers who accumulate more human capital also tend to hold better matches, amplifying wage inequality. This sorting intensifies over the course of the life cycle as workers climb the job ladder and accumulate skills simultaneously.

Table 7: Wage Variance Decomposition: Levels

| | Wage Variance: Var$(\log w_t)$ | | | |
|---|---|---|---|---|
| Years | Human Capital | Match Quality | Cov$(s, m)$ | Total |
| 0 | 13.9% | 73.7% | 12.7% | 100% |
| 5 | 10.7% | 60.3% | 29.1% | 100% |
| 10 | 15.4% | 52.1% | 32.5% | 100% |
| 20 | 22.4% | 43.5% | 34.2% | 100% |

*Note:* The table reports percentage contributions to the cross-sectional variance of log wages among employed workers. Wage variance is decomposed as Var$(\log w) = \rho_{ws}^2$Var$(\log s) + \rho_{wm}^2$Var$(\log m) + 2\rho_{ws}\rho_{wm}$Cov$(\log s, \log m)$.

To understand the sources of rising wage inequality over the life cycle, we decompose changes in variance relative to labor market entry. Table 8 shows that the growth in wage inequality is driven primarily by the intensification of sorting between human capital and match quality, with the covariance term contributing 60 to 83 percent of wage variance growth. This large positive contribution indicates that the relationship between skills and match quality becomes increasingly positive as careers progress: workers who accumulate more human capital systematically move into higher-quality matches, and this sorting amplifies wage inequality beyond what the dispersion in either component alone would generate.

Human capital dispersion also contributes positively to wage variance growth, accounting for essentially zero at year 5 but rising to 33 percent by year 20. This growing contribution reflects heterogeneity in the efficiency of human capital accumulation and the taste of work. Notably, match quality dispersion contributes very little to wage variance growth (7 to 16

**The Importance of Entry Timing** An important question is whether the decomposition of inequality growth depends on when the observation begins. To address this, we compare two measurement approaches: one that tracks workers from true labor market entry, and another that begins observation five years after entry, mimicking studies that restrict attention to workers with some labor market experience. Table 9 presents the comparison. The left columns measure changes from labor market entry; the right columns measure changes from five years after entry. The results reveal striking differences in the decomposition patterns.

When measured from entry, the wage variance decomposition yields clear and interpretable patterns. The covariance between human capital and match qual-

Table 8: Wage Variance Decomposition: Changes

| | Change in Wage Variance: $\Delta\text{Var}(\log w_t)$ | | | |
|---|---|---|---|---|
| Years | Human Capital | Match Quality | Cov$(s, m)$ | Total |
| 5 | 0.1% | 16.2% | 82.9% | 100% |
| 10 | 18.6% | 6.6% | 74.4% | 100% |
| 20 | 32.7% | 6.9% | 60.3% | 100% |

*Note:* The table reports percentage contributions to the change in cross-sectional wage variance relative to labor market entry (year 0). The decomposition is $\Delta\text{Var}(\log w) = \rho_{ws}^2 \Delta\text{Var}(\log s) + \rho_{wm}^2 \Delta\text{Var}(\log m) + 2\rho_{ws}\rho_{wm}\Delta\text{Cov}(\log s, \log m)$. The large positive contribution of Cov$(s, m)$ reflects endogenous sorting: workers who accumulate more human capital systematically move into higher-quality matches, amplifying wage inequality growth. Model results from annual simulations. This decomposition was not targeted in the calibration.

Table 9: How Entry Timing Affects Inequality Decomposition

| | Change in Wage Variance: $\Delta\text{Var}(\log w)$ | | | | | |
|---|---|---|---|---|---|---|
| | From Entry ($t = 0$) | | | From 5 Years ($t = 5$) | | |
| Years | HC | Match | Cov$(s, m)$ | HC | Match | Cov$(s, m)$ |
| 5 | 0.1% | 16.2% | 82.9% | 45.6% | -4.7% | 60.6% |
| 10 | 18.6% | 6.6% | 74.4% | 48.6% | -0.2% | 51.8% |
| 20 | 32.7% | 6.9% | 60.3% | 46.7% | 9.3% | 44.6% |
| | **Change in Earnings Variance: $\Delta\text{Var}(\log e)$** | | | | | |
| | From Entry ($t = 0$) | | | From 5 Years ($t = 5$) | | |
| Years | Wages | Hours | Cov$(w, h)$ | Wages | Hours | Cov$(w, h)$ |
| 5 | 59.4% | -29.4% | 70.0% | 65.6% | -27.1% | 61.5% |
| 10 | 66.4% | -29.7% | 63.3% | 69.6% | -19.0% | 49.3% |
| 20 | 70.4% | -18.3% | 47.9% | 79.9% | -11.4% | 31.4% |

*Note:* The table compares the decomposition of variance changes when measured from labor market entry versus from five years after entry. The left panel computes changes relative to year 0 (labor market entry). The right panel computes changes relative to year 5. “Years” in the right panel refers to years elapsed from the new baseline (e.g., Year 5 in the right panel corresponds to year 10 from true entry).

ity is the dominant driver of inequality growth, contributing 60 to 83% across all horizons. Match quality dispersion contributes modestly (7 to 16%), while human capital dispersion starts near zero and rises to 33% by year 20.

When measurement begins five years after entry, the decomposition changes substantially. Human capital dispersion now appears to be the dominant force, accounting for 46 to 49% of wage variance growth. Match quality dispersion contributes almost nothing or even a slight negative amount (-5 to 9%). The covariance term remains important (45 to 61%) but is no longer dominant. This altered decomposition arises because the delayed baseline misses the first five years when the most rapid sorting occurs and match quality dispersion grows fastest. Starting observation later effectively conditions on this initial sorting, leaving human capital accumulation as the primary source of subsequent inequality growth.

The earnings variance decomposition shows similar patterns but smaller differences across measurement approaches. The earnings decomposition is more stable across baselines because both wage and hours dynamics continue throughout the career, unlike the match quality dispersion, which is concentrated early on.

These findings highlight the importance of measuring outcomes at labor market entry rather than later in the career. Studies that begin observation after workers have accumulated some experience implicitly condition on initial sorting and skill accumulation, potentially leading to different conclusions about the relative importance of various inequality-generating mechanisms. The early-career period contains crucial dynamics that cannot be recovered by observing workers only after they have settled into the labor market.

## 7 Conclusion

Life-cycle hours growth is primarily an intensive-margin phenomenon driven by human-capital incentives. In the estimated model, the return to current hours operates not only through future wages, but also through job stability and access to better outside offers. These channels interact because human capital and match quality jointly determine labor-market outcomes. The same framework also implies that wage inequality rises early in the career largely through the growing covariance between skill and match quality, a pattern that is obscured when measurement begins several years after labor-market entry.

Two extensions seem especially relevant. First, our model abstracts from heterogeneity in education and family structure. Incorporating these dimensions would allow the framework to speak more directly to differences in life-cycle labor supply and wage dynamics across demographic groups. Second, we model hours as a continuous choice, while, in practice, workers often face fixed-hour contracts, part-time arrangements, and other institutional constraints on adjustment along the intensive margin. Gimpelson (2024) shows that modeling labor-market statuses as stochastic hours constraints can account for a substantial share of hours variation. Extending our framework in this direction would help clarify how much of early-career hours growth reflects human-capital incentives versus institutional constraints on hours adjustment.

## References


ARNOUD, A., F. GUVENEN, AND T. KLEINEBERG (2019): "Benchmarking global optimizers," *Working Paper*.

AUDOLY, R., F. DE PACE, AND G. FELLA (2022): "Job ladder, human capital, and the cost of job loss," *FRB of New York Staff Report*.

BAGGER, J., F. FONTAINE, F. POSTEL-VINAY, AND J.-M. ROBIN (2014): "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics," *American Economic Review*, 104, 1551–1596.

BAILY, M. N., C. HULTEN, AND D. CAMPBELL (1992): "Productivity dynamics in manufacturing plants," *Brookings Papers on Economic Activity: Microeconomics*, 1992, 187–267.

BICK, A., A. BLANDIN, AND R. ROGERSON (2024): "Hours Worked and Lifetime Earnings Inequality," *Working paper*.

BLUNDELL, R., A. BOZIO, AND G. LAROQUE (2011): "Labor supply and the extensive margin," *American Economic Review*, 101, 482–486.

BURDETT, K., C. CARRILLO-TUDELA, AND M. COLES (2020): "The cost of job loss," *The Review of Economic Studies*, 87, 1757–1798.

CARD, D. (1994): "Intertemporal Labor Supply: An Assessment," in *Advances in Econometrics, Sixth World Congress*, ed. by C. Sims, New York: Cambridge University Press, chap. 2, 49–78.

EROSA, A., L. FUSTER, AND G. KAMBOUROV (2016): "Towards a micro-founded theory of aggregate labour supply," *The Review of Economic Studies*, 83, 1001–1039.

Fan, X., A. Seshadri, and C. Taber (2024): "Estimation of a life-cycle model with human capital, labor supply, and retirement," *Journal of Political Economy*, 132, 48–95.

Flinn, C. and J. Mullins (2015): "Labor market search and schooling investment," *International Economic Review*, 56, 359–398.

Gimpelson, N. (2024): "Disentangling the Roles of Preferences and Shocks in Labor Supply," *Working paper*.

Heckman, J. J. (1976): "A life-cycle model of earnings, learning, and consumption," *Journal of political economy*, 84, S9–S44.

Huggett, M., G. Ventura, and A. Yaron (2011): "Sources of lifetime inequality," *American Economic Review*, 101, 2923–54.

Imai, S. and M. P. Keane (2004): "Intertemporal labor supply and human capital accumulation," *International Economic Review*, 45, 601–641.

Kaplan, G. (2012): "Inequality and the life cycle," *Quantitative Economics*, 3, 471–525.

Keane, M. P. and K. I. Wolpin (1997): "The career decisions of young men," *Journal of Political Economy*, 105, 473–522.

Lochner, L. J. and A. Monge-Naranjo (2011): "The nature of credit constraints and human capital," *American Economic Review*, 101, 2487–2529.

Neal, D. (1999): "The complexity of job mobility among young men," *Journal of Labor Economics*, 17, 237–261.

Ozkan, S., J. Song, and F. Karahan (2023): "Anatomy of lifetime earnings inequality: Heterogeneity in job-ladder risk versus human capital," *Journal of Political Economy Macroeconomics*, 1, 506–550.

Shimer, R. (2012): "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 15, 127–148.

Topel, R. H. and M. P. Ward (1992): "Job mobility and the careers of young men," *The Quarterly Journal of Economics*, 107, 439–479.

Wallenius, J. (2011): "Human capital accumulation and the intertemporal elasticity of substitution of labor: How large is the bias?" *Review of Economic Dynamics*, 14, 577–591.

Weiss, Y. (1986): "The determination of life cycle earnings: A survey," *Handbook of Labor Economics*, 1, 603–640.

# A Additional Empirical Results

## A.1 Sample Description

Table A1: Descriptive Statistics

| | Full Cohort | HS Graduates | Final Sample |
|---|---|---|---|
| Age (quarters) | 133.92 | 137.94 | 147.35 |
| | (50.29) | (51.24) | (52.09) |
| Sex | 1.51 | 1.46 | 1.47 |
| | (0.50) | (0.50) | (0.50) |
| Highest grade ever | 14.01 | 11.96 | 11.99 |
| | (3.15) | (0.19) | (0.09) |
| Total hours worked | 377.29 | 361.03 | 373.73 |
| | (1351.91) | (1170.26) | (1228.60) |
| Real hourly wage | 12.56 | 10.74 | 11.02 |
| | (8.44) | (6.36) | (6.46) |
| Fraction unemployed | 0.05 | 0.06 | 0.06 |
| | (0.19) | (0.21) | (0.21) |
| Fraction employed | 0.70 | 0.66 | 0.67 |
| | (0.44) | (0.45) | (0.45) |

## A.2 Quarters-Since-Graduation Profiles

Our first set of results presents mean and standard deviation profiles by quarters since graduation. Both cohorts are shown in all figures to assess the consistency of the empirical patterns, with the NLSY79 represented by solid lines and the NLSY97 by dashed lines. At the end of this section, we also present analogous profiles based on potential experience to assess what is missed when using experience rather than quarters since graduation.

We begin with Figure A1, which illustrates the extensive margin of labor supply through profiles of employment, unemployment, and nonparticipation. Four facts emerge. First, employment rises sharply in the first few years after graduation, stabilizing at about 90–95 percent of weeks. Second, unemployment and nonparticipation both fall rapidly after entry, converging to levels below 5 percent. Third, the variability of all three statuses declines with age, reflecting increasing stability in labor market attachment over the life cycle. Finally, these patterns are remarkably consistent across the NLSY79 and NLSY97 cohorts.[8]

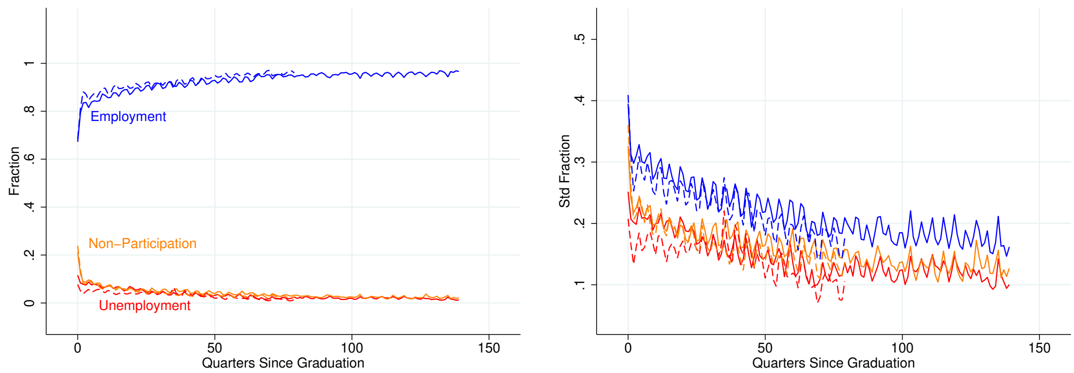


(a) Labor Market Statuses

(b) Std Labor Market Statuses

Figure A1: Labor Market Status Profiles

Note: The figure plots the fraction of weeks in each quarter that workers spend in employment, unemployment, or nonparticipation. These profiles are constructed using the NLSY's weekly work-history array; by construction, the fractions sum to one. The NLSY79 is represented by solid lines and the NLSY97 by dashed lines. We restrict the sample to individuals with a minimum level of labor market attachment, defined as between 520 and 5,200 annual hours. Appendix A.1 shows results using the full sample without dropping observations based on hours worked.

[8] Appendix Figure A5 shows analogous profiles without restricting the sample by hours worked. In that unrestricted sample: (i) employment starts at only 50 percent and never rises above 80 percent; (ii) nonparticipation begins around 35 percent, declines, but rises back to 25 percent with retirement; and (iii) unemployment starts at 14 percent and ends below 5 percent. By contrast, in our restricted sample: (i) employment quickly climbs above 85 percent and stabilizes near 90–95 percent, while nonparticipation and unemployment remain low throughout. Focusing on workers with meaningful labor market attachment allows us to capture the dynamics of labor supply choices of active workers, rather than shifts driven by structural shocks, such as skill obsolescence, or health shocks, such as disability.

Figure A2 presents the hours' profile and its components. The first row shows total (per-capita) hours worked per quarter, defined as the sum of hours across all weeks in the quarter. Increases in total hours can arise either from longer workweeks or from more weeks worked. The second row plots the average number of weeks worked per quarter, and the third row plots the average workweek. The left panels plot means, while the right panels plot standard deviations. Also, all profiles are measured among workers with positive hours in the quarter.

The first row shows that the mean total hours increase rapidly during the first five years of labor market experience, before slowing and plateauing after roughly ten years. For both cohorts, this initial phase adds more than 100 hours per quarter, which represents nearly one-third of the total lifetime increase. This steep early-career rise reflects simultaneous improvements in two margins. The second row shows that the number of weeks worked per quarter climbs from around 10.5 to nearly 13. The third row shows that the average workweek rises from around 32–33 hours to 38–39 hours. Both margins flatten after the first decade, leaving total hours largely stable for the remainder of the career. Importantly, the first five years are when the transition from low attachment and unstable early jobs to nearly full-time, full-year employment occurs.

The right panels reveal different patterns regarding second moments. The standard deviation of total hours follows a U-shape for the first two decades after graduation, consistent with Kaplan (2012). The source of this U-shape lies entirely in the workweek margin: the dispersion in weeks worked declines monotonically with experience, as most individuals converge to steady full-year employment. By contrast, dispersion in weekly hours first declines, as workers settle into stable jobs, and then rises again, reflecting growing heterogeneity in weekly labor supply at mid-career.

Lastly, comparing cohorts reveals systematic differences. For most outcomes, the NLSY79 profiles lie above those of the NLSY97 in both levels and dispersion. Mean total hours are higher in the NLSY79, and so is their dispersion. Weeks worked, by contrast, are lower in the NLSY79, with higher dispersion than in the NLSY97. The workweek is substantially higher for the early cohort, while its dispersion is similar across cohorts. These patterns suggest that the lower total hours of the NLSY97 cohort are driven primarily by shorter workweeks, whereas their weeks worked are more stable.

Figure A3, first row, presents the profile of real hourly wages, measured in 2007 dollars and deflated using the Consumer Price Index (CPI). Wages rise rapidly during the first five years after graduation, but, unlike the hours profile, they continue to grow at a steady pace rather than plateauing. The standard deviation of wages also increases gradually over the life cycle and does not display the U-shaped pattern observed for hours. Importantly, the wage profiles of the NLSY79 and NLSY97 cohorts are very similar in both levels and dispersion, suggesting robustness across cohorts.

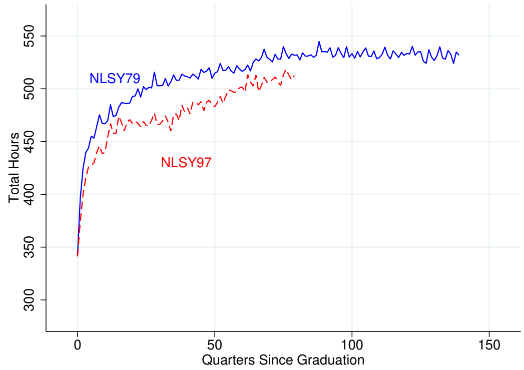


(a) Mean Hours Worked

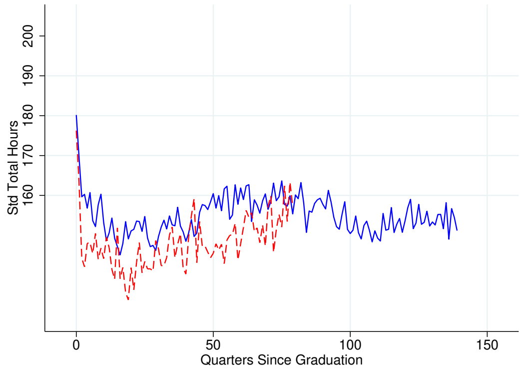


(b) Std Hours Worked

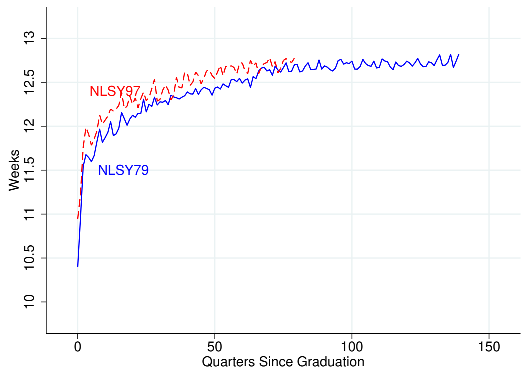


(c) Mean Weeks Worked in a Quarter

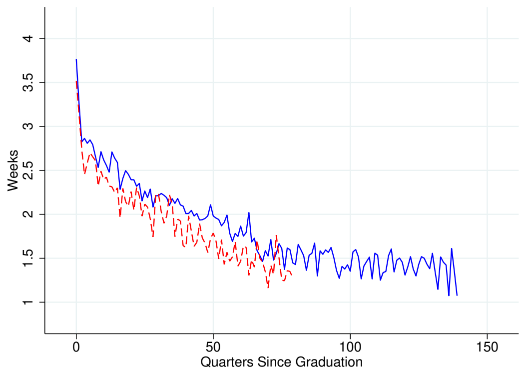


(d) Std Weeks Worked in a Quarter

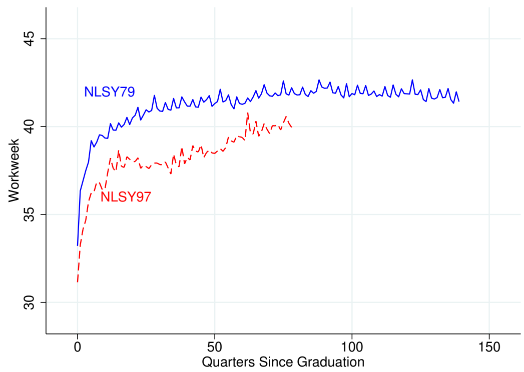


(e) Mean Workweek

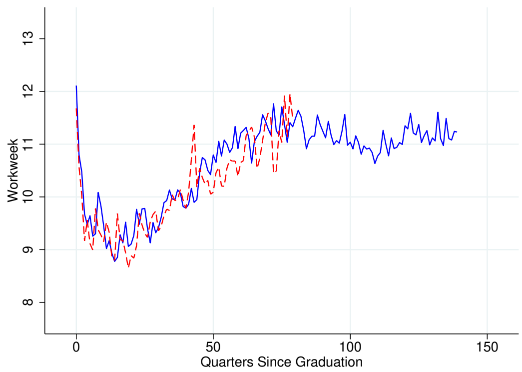


(f) Std Workweek

Figure A2: Total Hours and Components Profiles

Note: The figure plots total hours worked per quarter and its components: weeks worked and average workweek. These profiles are constructed using the NLSY's weekly work-history array, which records each respondent's labor market status, main job, and hours worked (if employed) for every week. We focus on hours worked in the main job, though respondents may hold multiple jobs within a given week. Including hours from all jobs does not change the early-career profile, but the gap between main-job hours and total hours across all jobs widens with experience, peaking about 20 years after graduation. The NLSY79 is represented by solid lines and the NLSY97 by dashed lines. The sample is restricted to individuals with a minimum level of labor market attachment, defined as between 520 and 5,200 annual hours.

Figure A3, second row, shows the mean and standard deviation of the cumulative number of jobs held over a worker's career. Workers tend to change jobs more frequently in the early years of their careers, with the average profile exhibiting a concave pattern. This concave pattern reflects the transition from unstable early-career matches to more stable relationships. The dispersion in cumulative jobs also rises concavely over time. Both profiles highlight the importance of job mobility in early-career wage growth, consistent with the findings of Topel and Ward (1992).


(a) Mean Real Wages

(b) Std Real Wages

(c) Number of Jobs Held

(d) Std Number of Jobs Held

Figure A3: Real Wages and Number of Jobs Held

Note: The figure plots profiles of real hourly wages and the cumulative number of jobs held. Both are constructed from the NLSY's weekly work-history array, which links each employed week to a main job with an associated hourly wage. For each quarter, we compute the average real wage across all workweeks, excluding weeks without employment. We focus on wages from the main job. The number of jobs held is measured as the cumulative count of distinct jobs a worker has reported up to each quarter, recorded at the end of the period. The NLSY79 is represented by solid lines and the NLSY97 by dashed lines. The sample is restricted to individuals with a minimum level of labor market attachment, defined as between 520 and 5,200 annual hours.

We now turn to profiles based on potential experience, defined as age minus years of education minus six. Because our sample is restricted to high school graduates, potential experience is mechanically determined by birth year. To ensure that most workers have completed formal schooling, we compute profiles beginning three years after potential labor market entry. Figure A4 plots the hours and wage potential-experience profiles. Although the mean and standard deviation of hours and wages resemble those from years-since-graduation profiles at later ages, two important features disappear. First, the sharp rise in mean hours during the first five years of labor market experience is muted. Second, the U-shaped pattern in the standard deviation of hours vanishes. This comparison highlights the importance of correctly identifying the timing of labor market entry and capturing the dynamics of young workers.

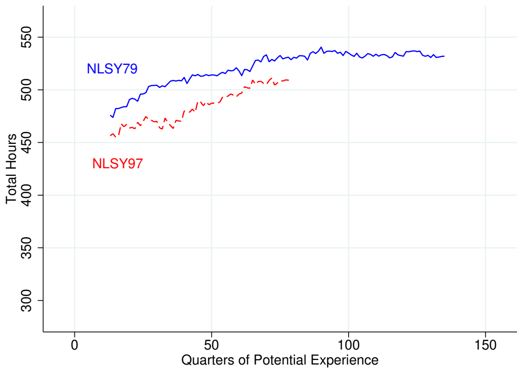


(a) Mean Hours Worked

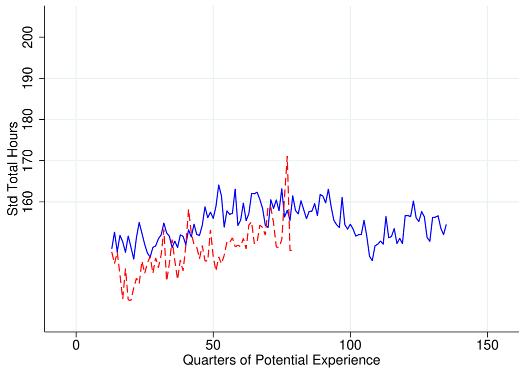


(b) Std Hours Worked

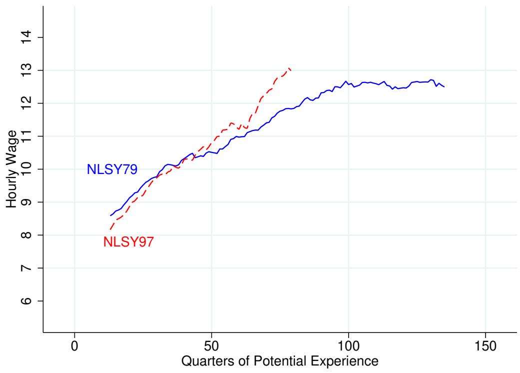


(c) Mean Real Wages

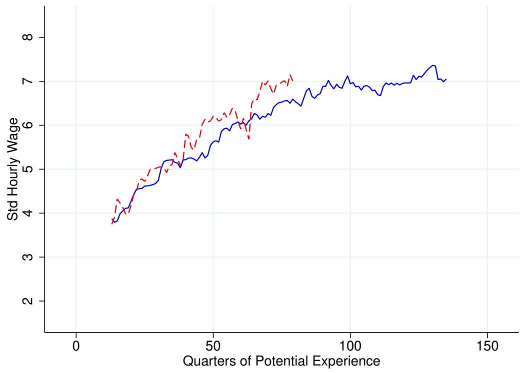


(d) Std Real Wages

Figure A4: Age-Based Hours and Wages Profiles

Note: The figure plots hours and wage profiles against potential experience, defined as age minus years of education minus six. Profiles start three years after potential entry to ensure schooling is completed. The NLSY79 is shown with solid lines and the NLSY97 with dashed lines.

## A.3 Labor Market Statuses Without Hours Restrictions

Figure A5 plots the fraction of weeks per quarter that individuals spend in employment, unemployment, or nonparticipation, when no restrictions are imposed on hours worked. Compared with the main text, where we restrict the sample to workers with meaningful labor market attachment, the profiles here reflect the entire population.

First, employment starts at about 50 percent immediately after graduation and increases gradually over the next two decades, but then begins to decline roughly 25 years after entry. By the end of the sample, the employment rate falls to about 70 percent. Second, nonparticipation shows the most pronounced changes. It starts high, near 35 percent, falls to roughly 16 percent by mid-career, and then rises again to about 26 percent toward the end of the observation window. This late-life increase reflects both voluntary and involuntary retirement. Finally, unemployment follows a U-shape: it is high at entry (around 14 percent), declines steadily during the early career years, and stabilizes at under 5 percent later in life.

Overall, the main difference relative to the restricted sample in the main text is that including weakly attached individuals lowers employment rates, raises nonparticipation, and highlights retirement dynamics. This confirms the value of focusing on workers with stronger labor market attachment when studying the determinants of hours and wages.

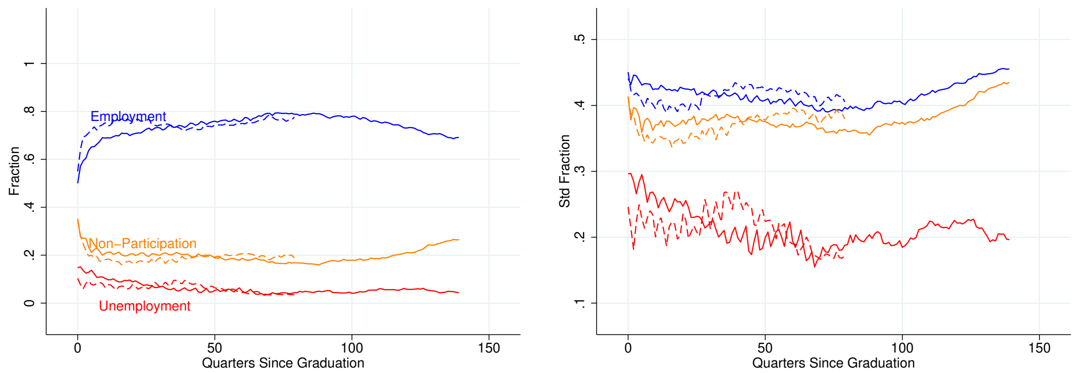


(a) Labor Market Statuses (b) Std Labor Market Statuses

Figure A5: Labor Market Statuses' Profiles

Note: The figure plots the fraction of weeks in each quarter that workers spend in employment, unemployment, or nonparticipation. These profiles are constructed using the NLSY's weekly work-history array; by construction, the fractions sum to one. The NLSY79 is represented by solid lines and the NLSY97 by dashed lines. We do not restrict the sample to individuals with a minimum level of labor market attachment, as we do in the main text.

## A.4 Other Profiles Without Hours Restrictions

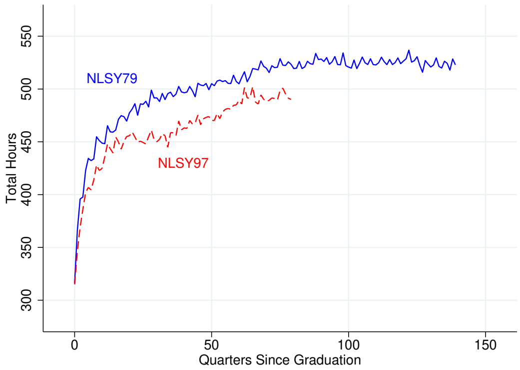


(a) Mean Hours Worked

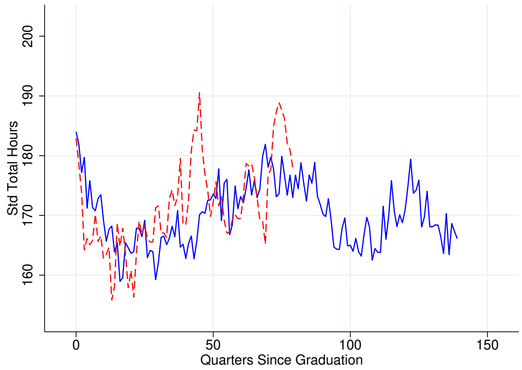


(b) Std Hours Worked

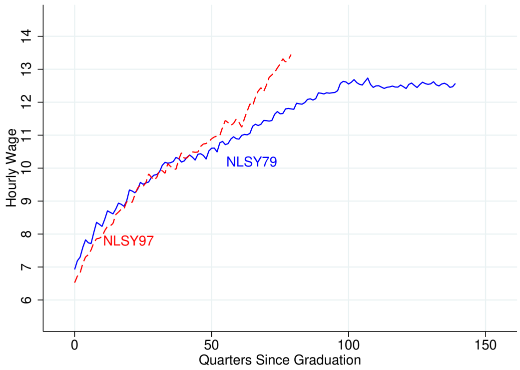


(c) Mean Real Wages

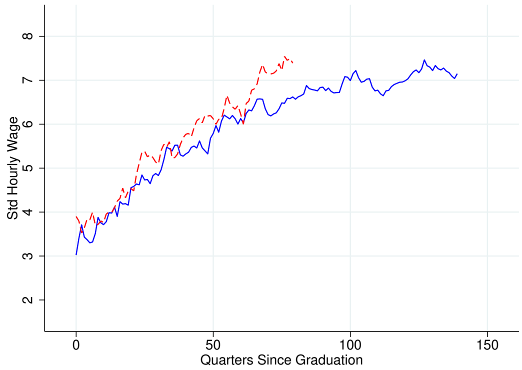


(d) Std Real Wages

Figure A6: Wages and Hours Profiles

Note: The figure plots hours and wage profiles against quarters-since-graduation. The NLSY79 is shown with solid lines and the NLSY97 with dashed lines. We do not restrict the sample to individuals with a minimum level of labor market attachment, as we do in the main text.

# B Hours Decomposition

## B.1 Details on the Hours Decomposition

In this appendix, we derive the statistical decomposition that we use to quantify the relative importance of both the average number of weeks worked annually and the average weekly hours in accounting for the increase in total hours.

**Goal:** Derive a decomposition of total hours worked into "intensive" and "extensive" margins. The intensive margin is measured by the workweek, while the extensive margin is measured by the number of weeks worked in a year and changes in the sample composition.

**Idea:** What would be the total hours worked if the intensive margin (or the extensive margin) is held constant? Perform a discrete-time approximation of the total derivative.

**Notation:**

- $H_t$ = the average total hours in period $t$, where period $t$ is measured in years since graduation. Its definition is: $H_t = \sum_{i \in E_t} \omega_{it} h_{it} n_{it}$.
- $E_{it}$ = the set of all workers $i$ employed in period $t$.
- $\omega_i t$ = the sample weight of worker $i$, initially taken to be equal to the inverse of the sample size in that period, $\omega_{it} = \frac{1}{N_t}$.
- $h_{it}$ = the average workweek in that period for worker $i$.
- $n_{it}$ = the number of weeks worked in that period by worker $i$.

**Implementation issues:**

- Because the sample is not a balanced panel, the sample weights change across periods. Changes in the sample composition, including workers' entry, exit, and weight changes, need to be tracked.
- By performing a discrete approximation of the total derivative, second-order terms do not cancel out. They need to be tracked.
- Also, because of the approximation of the total derivative, changes need to be evaluated in a base period. The Laspeyres index is used, i.e., period $t-1$ is the base period.

**Derivation Sketch:**

The derivation proceeds by following these steps:

- First, decompose $H_t - H_{t-1}$ into the intensive and extensive margins.
- Second, sum $H_t - H_{t-1}$ for $t \in \{1, \ldots, T\}$ to find the contributions of the intensive and extensive margins over an extended period.

**Formal Derivation:**

A change in average total hours between periods $t$ and $t-1$, $H_t - H_{t-1}$, can be decomposed as follows:

$$\begin{aligned}
H_t - H_{t-1} &= \sum_{i \in E_t} \omega_{it} h_{it} n_{it} - \sum_{i \in E_{t-1}} \omega_{it} h_{t-1} n_{t-1} \\
&= \sum_{i \in E_t} \omega_{it} h_{it} n_{it} - \sum_{i \in E_{t-1}} \omega_{it} h_{t-1} n_{t-1} \pm \sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it-1} n_{it-1} \\
&= \sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} n_{it} - h_{it-1} n_{it-1} \right) \\
&\quad + \underbrace{\sum_{i \in E_{t \cap t-1}} \left( \omega_{it} - \omega_{it} \right) h_{it-1} n_{it-1} + \sum_{i \in E_{t / t \cap t-1}} \omega_{it} h_{it} n_{it} - \sum_{i \in E_{t-1 / t \cap t-1}} \omega_{it} h_{it-1} n_{it-1}}_{\text{composition change}} .
\end{aligned}$$

The last terms capture the impact of weight changes, as well as entry and exit. We bunch them together in a *composition change* term. The set $E_{t \cap t-1}$ is composed of all workers that are present in the sample in both periods $t$ and $t-1$, or, more formally, $E_{t \cap t-1} = \{i \mid i \in E_t \cap E_{t-1}\}$.

Now working, continuing with the derivation, $H_t - H_{t-1}$ can be further decomposed as:

$$
\begin{aligned}
H_t - H_{t-1} &= \sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} n_{it} - h_{it-1} n_{it-1} \right) + \textit{comp. change} \\
&= \sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} n_{it} - h_{it-1} n_{it-1} \right) \pm \sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it} n_{it-1} + \textit{comp. change} \\
&= \sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1} + \sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it} \left( n_{it} - n_{it-1} \right) + \textit{comp. change} \\
&= \sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1} + \sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it} \left( n_{it} - n_{it-1} \right) \\
&\pm \sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it-1} \left( n_{it} - n_{it-1} \right) + \textit{comp. change} \\
&= \underbrace{\sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1}}_{\text{intensive margin}} + \underbrace{\sum_{i \in E_{t \cap t-1}} \omega_{it} h_{it-1} \left( n_{it} - n_{it-1} \right)}_{\text{extensive margin}} \\
&+ \underbrace{\sum_{i \in E_{t \cap t-1}} \omega_{it} \left( h_{it} - h_{it-1} \right) \left( n_{it} - n_{it-1} \right)}_{\text{quadratic term}} + \textit{comp. change} \ .
\end{aligned}
$$

The first term refers to the intensive margin. It answers how much the increase in the workweek between $t$ and $t-1$ increased total hours worked, holding the number of weeks worked in a year constant at its $t-1$ value. The second term refers to the extensive margin. Similarly, it answers how much the increase in weeks worked in a year between $t$ and $t-1$ increased total hours worked, holding the workweek constant at its $t-1$ value. Lastly, the last term captures changes in second-order terms. It is positive if the weeks worked and the workweek tend to move together.

The change in hours can be further adjusted to account for differences in the contributions of the intensive and extensive margins between job movers and job stayers. Because the change is computed between period $t$ and $t-1$, job-stayers are workers associated with the same employers in both periods, while job-movers are those associated with different employers.

The set of workers employed in both periods, $E_{t \cap t-1}$, is then split into two disjoint sets: the set of job-movers and job-stayers, $M_t$ and $S_t$, respectively. The full decomposition formula is

$$
\begin{aligned}
H_t - H_{t-1} = & \underbrace{\sum_{i \in M_t} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1}}_{\text{movers' intensive margin}} + \underbrace{\sum_{i \in M_t} \omega_{it} h_{it-1} \left( n_{it} - n_{it-1} \right)}_{\text{movers' extensive margin}} \\
& + \underbrace{\sum_{i \in M_t} \omega_{it} \left( h_{it} - h_{it-1} \right) \left( n_{it} - n_{it-1} \right)}_{\text{movers' quadratic term}} \\
& + \underbrace{\sum_{i \in S_t} \omega_{it} \left( h_{it} - h_{it-1} \right) n_{it-1}}_{\text{stayers' intensive margin}} + \underbrace{\sum_{i \in S_t} \omega_{it} h_{it-1} \left( n_{it} - n_{it-1} \right)}_{\text{stayers' extensive margin}} \\
& + \underbrace{\sum_{i \in S_t} \omega_{it} \left( h_{it} - h_{it-1} \right) \left( n_{it} - n_{it-1} \right)}_{\text{stayers' quadratic term}} \\
& + \textit{comp. change} \quad .
\end{aligned}
$$

The contribution of the intensive and extensive margin between periods $0$ and $T$ is easily found by summing their yearly contributions:

$$
H_T - H_0 = \sum_{t=1}^{T} \Big( H_t - H_{t-1} \Big)
$$

## B.2 Additional Decomposition Results

This appendix extends the hours decomposition from Section 6 by distinguishing between job movers and job stayers, and by further decomposing the extensive margin into composition effects. These additional results provide insight into whether hours growth is driven by stable employment relationships or by job mobility.

## B.3 Job Movers versus Job Stayers

We define job stayers as individuals who remain with the same employer between periods $t-1$ and $t$, and job movers as those who switch employers. The set of workers employed in both periods, $E_{t\cap t-1}$, can be partitioned as $E_{t\cap t-1} = J_t^S \cup J_t^M$, where $J_t^S$ and $J_t^M$ denote job stayers and job movers, respectively. This classification requires observing individuals in both periods, so it applies only to the intensive margin, the within-continuers component of the extensive margin, and the quadratic term.

For example, the intensive margin decomposes as:

$$\sum_{i \in E_{t\cap t-1}} \omega_{it}\left(h_{it} - h_{it-1}\right) n_{it-1} = \underbrace{\sum_{i \in J_t^M} \omega_{it}\left(h_{it} - h_{it-1}\right) n_{it-1}}_{\text{movers' intensive margin}} + \underbrace{\sum_{i \in J_t^S} \omega_{it}\left(h_{it} - h_{it-1}\right) n_{it-1}}_{\text{stayers' intensive margin}} . \quad (5)$$

We further refine the decomposition by splitting movers into job-to-job movers, who switch directly between employers, and job-to-nonemployment movers, who experience at least three weeks of nonemployment between jobs. This distinction separates voluntary or upward mobility from transitions associated with joblessness.

Table B1 reports the results for the data. Job stayers account for virtually all cumulative hours growth, contributing 95 to 118 percent across horizons when aggregating over all margins except composition. Job movers net to a small positive contribution at 5–10 years but turn negative by 15–20 years. Within the mover category, job-to-job transitions contribute positively (57 to 77 percent), but this is largely offset by the negative contribution of nonemployment movers ($-40$ to $-85$ percent).

On the intensive margin (Panel 2), both types of movers raise their workweek early in the career, but stayers dominate the aggregate contribution. On the extensive margin (Panel 3), job-to-job moves contribute strongly to weeks worked (39 to 59 percent), whereas nonemployment moves contribute sharply negatively ($-47$ to $-86$ percent by year 20), offsetting their intensive gains. The quadratic term (Panel 4) is small and positive throughout.

| Cumulative Contribution up to Year X | All Margins (except Composition) | | | | |
|---|---|---|---|---|---|
| | All Workers | Job Stayers | Job Movers | | |
| | | | All Movers | Job-to-Job | Non-Emp. |
| 5 | 111.6% | 95.1% | 16.5% | 56.2% | -39.7% |
| 10 | 109.4% | 102.2% | 7.2% | 70.6% | -63.5% |
| 15 | 114.9% | 119.1% | -4.2% | 77.6% | -81.8% |
| 20 | 110.2% | 117.9% | -7.7% | 76.8% | -84.5% |

| Cumulative Contribution up to Year X | Intensive Margin | | | | |
|---|---|---|---|---|---|
| | All Workers | Job Stayers | Job Movers | | |
| | | | All Movers | Job-to-Job | Non-Emp. |
| 5 | 57.0% | 36.4% | 20.6% | 15.2% | 5.4% |
| 10 | 57.3% | 38.7% | 18.6% | 16.5% | 2.1% |
| 15 | 59.6% | 43.7% | 15.9% | 16.4% | -0.5% |
| 20 | 56.5% | 41.5% | 15.0% | 15.6% | -0.6% |

| Cumulative Contribution up to Year X | Extensive: Within-Continuers | | | | |
|---|---|---|---|---|---|
| | All Workers | Job Stayers | Job Movers | | |
| | | | All Movers | Job-to-Job | Non-Emp. |
| 5 | 49.4% | 56.7% | -7.3% | 39.3% | -46.6% |
| 10 | 45.0% | 60.6% | -15.6% | 52.6% | -68.2% |
| 15 | 47.1% | 71.8% | -24.6% | 59.5% | -84.1% |
| 20 | 45.9% | 72.9% | -27.0% | 59.4% | -86.4% |

| Cumulative Contribution up to Year X | Quadratic Term | | | | |
|---|---|---|---|---|---|
| | All Workers | Job Stayers | Job Movers | | |
| | | | All Movers | Job-to-Job | Non-Emp. |
| 5 | 5.2% | 2.0% | 3.2% | 1.7% | 1.5% |
| 10 | 7.1% | 3.0% | 4.1% | 1.5% | 2.6% |
| 15 | 8.2% | 3.6% | 4.6% | 1.7% | 2.9% |
| 20 | 7.8% | 3.5% | 4.3% | 1.8% | 2.5% |

Table B1: Hours Decomposition by Job Movers and Stayers: Data

*Note:* Job stayers remain with the same employer between periods $t-1$ and $t$. Job movers switch employers. Job-to-Job movers transition directly between employers. Non-Emp. movers experience at least three weeks of nonemployment between jobs. All values represent percentage contributions to cumulative hours growth up to year X. Data from NLSY79.

### B.4 Composition Effects

The extensive margin can be separated into a within-continuers component (reported above) and a composition effect. The composition effect captures changes in total hours arising from net entry and exit as well as shifts in the relative weights of continuing workers. Table B2 reports the total composition contribution and its two subcomponents. Net entry is consistently negative and often large, as late entrants work fewer hours than those already in the sample. Shifts in continuers' weights partially offset this force, so the overall composition effect remains modest and slightly negative on average ($-9$ to $-15$ percent).

| Cumulative Contribution up to Year X | Composition | Net Entry | Changing Weights |
|---|---|---|---|
| 5 | -11.6% | -113.3% | 101.6% |
| 10 | -9.4% | -91.3% | 82.0% |
| 15 | -14.9% | -61.1% | 46.2% |
| 20 | -10.2% | 45.3% | -55.5% |

Table B2: Decomposition of Composition Effects: Data

*Note:* The composition effect reflects changes in total hours from net entry/exit and shifts in worker weights. Data from NLSY79.

### B.5 Model Comparison

Table B3 compares the mover-stayer decomposition in the data and model. In the data, job stayers account for 93 to 115 percent of hours growth, job-to-job movers contribute 55 to 75 percent, and nonemployment movers contribute $-41$ to $-87$ percent. The model overstates the role of stayers (343 to 525 percent) and the negative contribution of nonemployment movers ($-297$ to $-518$ percent), while matching the job-to-job contribution reasonably well (59 to 95 percent).

The model's overstatement of the stayer contribution reflects its learning-by-doing mechanism: workers accumulate human capital through hours worked at their current employer, creating strong incentives for stayers to work long hours and to increase labor supply as they acquire more human capital. A Ben-Porath-style model with transferable-skill investment would generate smaller differences between stayers and movers. We also believe that an inquiry into the extent of human capital wealth effects on labor supply would be interesting. We leave this extension for future work.

Table B3: Hours Decomposition by Job Movers and Stayers: Model versus Data

| Years | Job Stayers | Job-to-Job | Non-Emp. | Other |
|---|---|---|---|---|
| | | Data | | |
| 5 | 93.1% | 54.5% | -41.2% | -6.4% |
| 10 | 99.2% | 69.1% | -66.1% | -2.3% |
| 15 | 115.5% | 75.9% | -84.6% | -6.7% |
| 20 | 114.4% | 75.0% | -87.0% | -2.4% |
| | | Model | | |
| 5 | 342.7% | 58.7% | -296.6% | -4.8% |
| 10 | 409.3% | 73.9% | -379.5% | -3.6% |
| 15 | 461.0% | 82.9% | -442.7% | -1.2% |
| 20 | 525.2% | 94.7% | -518.2% | -1.7% |

*Note:* Job-to-Job refers to workers who change employers directly. Non-Emp. refers to workers who transition through nonemployment (at least three weeks). All values represent percentage contributions to cumulative hours growth up to year X. Data from NLSY79. Model results from annual simulations.