

Interpreting Life-Cycle Inequality Patterns as an Efficient Allocation: Mission Impossible?

Alejandro Badel* Mark Huggett†

October 26, 2007

Abstract

Data on consumption, earnings, wages and hours dispersion over the life cycle is commonly viewed as incompatible with a Pareto efficient allocation. We show that a model with preference and wage shocks and full insurance produces the rise in consumption, wages and hours dispersion over the life cycle found in U.S. data. The efficient allocation model requires an increasing preference shifter dispersion profile to account for an increasing consumption dispersion profile. We examine U.S. data and find support for the view that the dispersion in preference shifters increases with age.

JEL Classification: E21, D91, D52

Keywords: Life-Cycle Inequality, Efficient Allocation, Preference Shocks

*Georgetown University. email: ab377@georgetown.edu.

†Georgetown University. email: mh5@georgetown.edu.

1 Introduction

A common finding from U.S. repeated cross-section data is that dispersion in consumption, wages and earnings increases with age. Figure 1 displays dispersion patterns based upon U.S. Consumer Expenditure Survey (CEX) data over the period 1981 to 2003.¹

Insert Figure 1 Here

The standard explanation for the pattern in Figure 1 relies upon idiosyncratic shocks that are not completely insured. It is now routine to posit that idiosyncratic shocks impact wages or earnings and that they are transmitted to consumption as some friction makes all or part of these shocks effectively uninsured.²

The literature contains virtually no attempts to interpret the patterns in Figure 1 within a model with full risk sharing. The paper by Storesletten, Telmer and Yaron (2001) is, to our knowledge, the only exception. They ask if a model with non-separable preferences between consumption and leisure can generate hours and consumption patterns similar to those in Figure 1. They find that there are no utility function parameters for which their model produces the observed rise in consumption dispersion over the life cycle together with only a “small” rise in hours dispersion.³ They conclude that this finding casts doubt on the full risk-sharing model and, thus, strengthens the argument in favor of models with uninsurable, idiosyncratic risk.

We show that the dispersion patterns in Figure 1 can be reconciled with an efficient allocation under full information. First, we present a simple model of efficient allocations under full information that exactly replicates the linearized profiles in Figure 1. This holds for a standard separable utility function with elasticity parameters within the range of those estimated in the literature. Our model has wage shocks as well as preference shifter shocks.

¹The methodology employed to produce Figure 1 follows Deaton and Paxson (1994) in using a time and age dummy variable regression. Our data and methodology are described in the Appendix.

²See Huggett (1996), Storesletten, Telmer and Yaron (2004), Heathcote, Storesletten and Violante (2004), Krueger and Perri (2006), Guvenen (2006), Huggett, Ventura and Yaron (2007) and Kaplan (2007), among others, for work based upon this type of mechanism.

³They require that the “small” rise in the variance of log hours over the life cycle is .08 or less. The rise in Figure 1, for the linearized hours profile, from age 28 to 62 is .02.

Within this model, preference shocks produce all of the rise in consumption dispersion, whereas wage shocks produce all of the rise in hours dispersion. Second, we examine the same micro-level data on consumption, wages and hours worked on which Figure 1 is based to determine whether or not there is support for the proposition that the dispersion of preference shifters increases with age. We present evidence that the variance of preference shifters increases with age. These findings suggest that the efficient allocation view is actually up to the standard of matching the patterns of Figure 1.

We focus on a specific set of facts. This implies that our results should not be taken as suggestive of a likely success of theories featuring full insurance in interpreting all aspects of the joint distribution of household consumption, wages and hours.⁴ Rather, our results simply show that the facts in Figure 1, related to the rise in cross-section inequality, are consistent with a theory featuring full insurance.

Our model requires a mild increase in the dispersion of hours worked with age to match the facts. We view this as consistent with the fact that the hours dispersion profile over the life cycle is basically flat. However, depending on sample selection criteria, this profile has been found to be either slightly increasing, slightly decreasing or U-shaped - see the Appendix. While a decreasing hours dispersion profile cannot be produced by our simple model, such a profile is also problematic for incomplete insurance models - see Kaplan (2007).

Section 2 outlines the model and discusses its implications for interpreting inequality over the life cycle. Section 3 estimates the life cycle dispersion profile of preference shifters. Data sources, sample selection criteria and empirical methods are described in an appendix.

2 Model

We analyze an overlapping generations economy. In this economy agents face idiosyncratic productivity and preference shocks. A continuum of agents is born at each time t . The size of each birth cohort is denoted by N_t . Agents are characterized by their age j , their year of birth b , and their own shock

⁴Cochrane (1991) and Attanasio and Davis (1996) examine data on household consumption growth, and find it hard to square with an efficient allocation. On the other hand, Altug and Miller (1990) consider several moments of the data, and find support for an efficient allocation model.

history $s^j = (s_0, s_1, \dots, s_j)$. At any age $j = 0, 1, \dots, J$ there is a finite number of possible shock histories s^j for the agent that occur with probability $P(s^j)$. An agent's productivity $w(s^j)$ and preference shift $z(s^j)$ are exogenously determined by the shock history.

Agents care about expected utility derived from consumption and labor. The functions $c_b(s^j)$ and $l_b(s^j)$ denote age j consumption and labor in history s^j for an agent born in year b . Expected utility takes the following form, where u is a period utility function, β is a discount factor and φ_j is the probability of surviving up to age j :

$$E \left[\sum_{j=0}^J \varphi_j \beta^j u(c_b(s^j), l_b(s^j), z(s^j)) \right]$$

2.1 Planning Problem

At time $t = 1$ a planner maximizes the weighted sum of individual expected utilities subject to a resource constraint. The planner's objective includes the remaining expected utility of agents born before $t = 1$ and lifetime expected utility of agents born at time $t \geq 1$. The planner assigns weight $\gamma_b > 0$ to the agents of cohort b . The objective in the Planning Problem can be conveniently stated by noting that time t , birth year b and the age j of an agent are related via the identity $b = t - j$. The resource constraint says that total consumption equals total output at each time period. We assume that an agent's productivity $w(s^j)$ is always strictly positive.

$$\max \sum_{t=1}^{\infty} E \left[\sum_{j=0}^J \gamma_{t-j} N_{t-j} \varphi_j \beta^j u(c_{t-j}(s^j), l_{t-j}(s^j), z(s^j)) \right]$$

subject to

$$\sum_{j=0}^J E \left[c_{t-j}(s^j) - l_{t-j}(s^j) w(s^j) \right] N_{t-j} \varphi_j = 0, \forall t \geq 1$$

We make the following assumptions on preferences, population growth rates and planning weights:

A1: The period utility function u is additively separable, strictly concave, continuously differentiable, strictly increasing in consumption and decreasing

in labor. Furthermore, the period utility function satisfies the Inada conditions so that the range of the marginal utilities of consumption and labor are $(0, \infty]$ and $[0, -\infty)$ respectively.

A2: $N_t = (1 + n)^t$ and $\beta(1 + n) < 1$.

A3: The Planning weights are set to $\gamma_b = \beta^b, \forall b$.

We rewrite the Planner's objective below, making use of assumptions A2 and A3. This highlights the fact that the Planner faces effectively a sequence of static maximization problems with the same period objective function and the same resource constraint.

$$\sum_{t=1}^{\infty} [\beta(1+n)]^t E \left[\sum_{j=0}^J \frac{\varphi_j}{(1+n)^j} u(c_{t-j}(s^j), l_{t-j}(s^j), z(s^j)) \right]$$

Theorem 1: Assume A1 - A3.

- (i) *There exists a unique allocation $(c^*(s^j), l^*(s^j))$ that solves the problem of maximizing $E \left[\sum_{j=0}^J \frac{\varphi_j}{(1+n)^j} u(c(s^j), l(s^j), z(s^j)) \right]$ subject to the period resource constraint $\sum_{j=0}^J E [c(s^j) - l(s^j)w(s^j)] \frac{\varphi_j}{(1+n)^j} = 0$.*
- (ii) *The time-invariant allocation $(c^*(s^j), l^*(s^j))$ is the unique solution to the Planning Problem.*

Proof: See the Appendix.

2.2 Reinterpreting the Dispersion Profiles

We now relate the solution of the Planning Problem to the empirical profiles in Figure 1. To do so, we assume that the period utility function is of the class in assumption A1'. The parameters (ϕ, ρ) of this utility function have been the focus of a large literature which is reviewed and summarized by Browning, Hansen, and Heckman (1999). In this literature ϕ is interpreted as the inverse of the Frisch elasticity of labor supply, whereas ρ is interpreted both as the coefficient of relative risk aversion and the inverse of the elasticity of intertemporal substitution.

A1': *The period utility function is $u(c_j, l_j, z_j) = \exp(z_j) \frac{c_j^{1-\rho}}{1-\rho} - \frac{l_j^{1+\phi}}{1+\phi}$.*

Theorem 2: Assume A1', A2 and A3. Furthermore, assume wages and the preference shifter have a variance that rises linearly with age so that $\text{var}(\log w_j) = g_w + v_w j$ and $\text{var}(z_j) = g_z + v_z j$, where (g_w, v_w, g_z, v_z) are positive constants. Then the unique solution to the Planning Problem $(c^(s^j), l^*(s^j))$ has the property that within a birth cohort or in cross section, the variance of log hours and the variance of log consumption rise linearly with age:*

$$\Delta \text{var}(\log l_j^*) = (1/\phi^2) \Delta \text{var}(\log w_j) = (1/\phi^2) v_w$$

$$\Delta \text{var}(\log c_j^*) = (1/\rho^2) \Delta \text{var}(z_j) = (1/\rho^2) v_z$$

Proof: The uniqueness and time invariance of the solution follows from Theorem 1. The linear rise in the variance follows from plugging the utility function into the necessary conditions for an interior solution to the Planning Problem, taking logs, and then taking variances. These conditions are that the two equations below hold for some positive values $\Lambda_t, \forall t \geq 1$:

$$u_c(c_{t-j}(s^j), l_{t-j}(s^j), z(s^j)) = \frac{\Lambda_t}{\beta^t}$$

$$-u_l(c_{t-j}(s^j), l_{t-j}(s^j), z(s^j)) = \frac{\Lambda_t w(s^j)}{\beta^t}$$

◇

We now determine if the solution to the Planning Problem can match the linear rise in the variance profiles in Figure 1 for parameter values that are consistent with those estimated in the literature. Directly from Figure 1 we calculate that the empirical slopes of the wage profile v_w , labor profile $slope_l$ and consumption profile $slope_c$ are $(v_w, slope_l, slope_c) = (.0035, .0006, .0013)$. Plugging these values into the conclusion of Theorem 2 gives the two equations below.

$$.0006 = slope_l = (1/\phi^2) v_w = (1/\phi^2) .0035$$

$$.0013 = slope_c = (1/\rho^2) v_z$$

These equations tell us two things. The first equation tells us that $\frac{1}{\phi} = .41$ so that the Frisch elasticity of labor supply is .41. This value is consistent with the range of estimates from the literature. Specifically, Browning, Hansen, and Heckman (1999, Table 3.3) note that estimates of this parameter for male labor hours vary from nearly zero to about one half. Given that the parameter v_z governing the variance profile of the preference shifters is not pinned down by the evidence in Figure 1, the second equation tells us that a solution to the Planning Problem can match the rise in consumption dispersion in Figure 1 for any positive value of the parameter ρ .

We conclude that it is straightforward to reconcile the empirical rise in the dispersion of log wages, log hours and log consumption in the U.S., as documented in Figure 1, with an efficient allocation. This follows from a standard period utility function using standard estimates of these preference parameters.⁵ The key to interpreting Figure 1 in this way is that preference shifters have an increasing variance with age. To determine whether or not this interpretation of Figure 1 can be easily dismissed, one would need to examine the empirical nature of preference variation. We do so next.

3 Dispersion of Preference Shifters

We start our analysis of the empirical nature of preference shifters with a first order necessary condition of the planning problem. This is the first equation below. Assumption A1' implies that the preference shifter satisfies the second equation. Taking variances of both sides of this restriction among age j agents produces the third equation.

$$-\frac{u_l(c_j, l_j, z_j)}{u_c(c_j, l_j, z_j)} = w_j$$

$$z_j = \rho \log c_j + \phi \log l_j - \log w_j$$

⁵We have conducted a similar analysis that differs in that the information displayed in Figure 1 is constructed by controlling either for cohort effects or for year effects and the number of recessionary quarters experienced by a cohort. We use the NBER definition of a recession. While the slopes differ from the benchmark analysis in Figure 1, the Frisch elasticity of labor supply that is required to make sense of hours and wage data, respectively .21 and .35, remain within the range of estimates from the literature.

$$var(z_j) = var(\rho \log c_j + \phi \log l_j - \log w_j)$$

The variance equation suggests the outlines of a strategy for determining the nature of the variance of the preference shifter. Simply, the variance of the preference shifter within a cohort could be measured by the right-hand-side of this equation. Such an analysis infers the preference shifter variance conditional on values of the preference parameters (ρ, ϕ) and, importantly, assuming no measurement error. We now discuss how to deal with measurement error.

3.1 Preference Shifters and Measurement Error

Assume that consumption, hours and wages are measured with error. Let a measured (log) variable, say $\log \hat{x}$, be equal to the true value plus a measurement error so that $\log \hat{x} = \log x + \epsilon^x$, for $x = c, l, w$. Then, based upon the reasoning from the previous section, the preference shifter is related to measured variables as in the first equation below. Assuming that the combined measurement error term (i.e $\epsilon_j \equiv -\rho\epsilon_j^c - \phi\epsilon_j^l + \epsilon_j^w$) is independent of the preference shifter, the second equation follows from the first.

$$z_j = \rho(\log \hat{c}_j - \epsilon_j^c) + \phi(\log \hat{l}_j - \epsilon_j^l) - (\log \hat{w}_j - \epsilon_j^w)$$

$$var(z_j) + var(\epsilon_j) = var(\rho \log \hat{c}_j + \phi \log \hat{l}_j - \log \hat{w}_j)$$

The variance equation tells us that the variance of the measured preference shifter (the right-hand-side term) equals the variance of the true preference shifter plus the variance of the measurement error term. The variance equation is useful as we would like to know whether the variance of the preference shifter is increasing with age. This can be determined by examining how the variance of the measured preference shifter changes with age while controlling for year or cohort effects. We therefore assume that the variance of the measurement error term does not vary with age or year. Under this assumption, $\Delta_j var(z_j) = \Delta_j var(\rho \log \hat{c}_j + \phi \log \hat{l}_j - \log \hat{w}_j)$ holds within cross

section or within cohort. In the analysis that follows, we analyze how the variance of the measured preference shifter changes with age.

3.2 Empirical Assessment

We calculate the variance of the measured preference shifter within age groups for each year at particular values for the preference parameters (ρ, ϕ) . We fix ϕ at the value $\phi = 1.0/.41$. This value makes sense of the hours and wage variance profiles within the model as discussed in section 2.2. We then calculate the sample variance of the measured preference shifter within age-year cells for values of the parameter ρ on a fine grid on the interval $[1.0, 4.0]$.⁶ The construction of these cells is described in the Appendix. We regress these variances on dummy variables for age groups and time period as was done in the construction of Figure 1. We then focus on the age coefficients to highlight how preference dispersion changes with age, controlling for time effects.

Insert Figure 2 Here

The results of this analysis are contained in Figure 2. Figure 2 shows the change in the variance of the measured preference shifter across age groups for different values of the preference parameter ρ . For example, at age 40 we plot the difference between the age 40 dummy coefficient and the coefficient on the youngest age group, which is age 28. Figure 2 shows from age 28 to age 62 that the variance of the measured preference shifter rises for all values of ρ in the range considered. This rise is greater for higher values of the coefficient of relative risk aversion ρ .

The next issue to address is whether the rise in the variance of the preference shifter, for any value of ρ , is consistent with the rise in consumption dispersion over the life cycle in Figure 1. For this to hold within the model, it would need to be the case that the rise in the variance of log consumption over the life cycle equals $(1/\rho^2)$ times the empirical rise in the variance of the preference shifter. Of course, the rise in the variance of the preference shifter in Figure 2 depends upon (ρ, ϕ) .

We calculate, based upon the data in Figure 2, that there is a value of ρ which is consistent with the rise in consumption dispersion over the life

⁶Browning et. al. (1999, Table 3.1) show that estimated values of ρ are in the range of below 1 to 4, when period utility is assumed to be separable in consumption and labor.

cycle. Specifically, this holds for $\rho = 2.5$. The rise in preference dispersion is .29 at this value and, thus, $(1/\rho^2) \times .29$ approximately equals .0463, which is the rise in consumption dispersion over the life cycle from Figure 1.

This type of calculation can be done somewhat differently. One can linearize the empirical profiles in Figure 2 by fitting a linear trend. This yields an empirical slope value $v_z(\rho, \phi)$ for each linearized profile. One then asks whether there is any value of ρ , given $\frac{1}{\phi} = .41$, such that $.0013 = slope_c = (1/\rho^2)v_z(\rho, \phi)$ holds. This calculation follows from Theorem 2 and Figure 1. We find that the value $\rho = 2.25$ approximately satisfies this equation.

Thus, the upshot of this section is two fold. First, there is evidence that the variance of the preference shifters rises with age for a range of values of the preference parameters. This evidence is also relevant for research which uses the standard competitive labor supply model to interpret data because the evidence is based upon the assumption that the intratemporal marginal rate of substitution equals labor productivity each period.⁷ Second, there are values of the preference parameters (ρ, ϕ) , that lie comfortably within the range of estimates in the literature, for which the rise in the variance of preference shifters is quantitatively consistent with the rise in consumption dispersion documented in Figure 1.

⁷The results in Figure 2 do depend on a functional form assumption which is also commonly used in the labor supply literature.

References

- Altug, S. and R. Miller (1990), Household Choices in Equilibrium, *Econometrica*, 58, 543-70.
- Attanasio, O. and S. Davis (1996), Relative Wage Movements and the Distribution of Consumption, *Journal of Political Economy*, 104, 1227- 62.
- Browning, M., Hansen, L. and J. Heckman (1999), Micro Data and General Equilibrium Models, in Handbook of Macroeconomics, ed. J.B. Taylor and M. Woodford, (Elsevier Science B.V, Amsterdam).
- Cochrane, J. (1991), A Simple Test of Consumption Insurance, *Journal of Political Economy*, 99, 957- 76.
- Deaton, A. and C. Paxson (1994), Intertemporal Choice and Inequality, *Journal of Political Economy*, 102, 437-67.
- Guvenen, F. (2006), Learning your Earning: Are Labor Income Shocks Really Very Persistent?, Institute for Empirical Macroeconomics, Discussion Paper 145.
- Heathcote, J., Storesletten, K. and G. Violante (2004), The Cross-Sectional Implications of Rising Wage Inequality in the United States, manuscript.
- Huggett, M. (1996), Wealth Distribution in Life-Cycle Economies, *Journal of Monetary Economics*, 38, 469 - 94.
- Huggett, M., Ventura, G. and A. Yaron (2007), Sources of Lifetime Inequality, manuscript.
- Kaplan, G. (2007), Inequality and the Lifecycle, manuscript.
- Krueger, D. and F. Perri (2006), Does Income Inequality Lead to Consumption Inequality? Evidence and Theory, *Review of Economic Studies*, 73, 163-93.
- Storesletten, K., Telmer, C. and A. Yaron (2001), How Important Are Idiosyncratic Shocks? Evidence from Labor Supply, *American Economic Review Papers and Proceedings*, 91, 413- 17.
- Storesletten, K., Telmer, C. and A. Yaron (2004), Consumption and Risk Sharing Over the Life Cycle, *Journal of Monetary Economics*, 51, 609- 33.

A Appendix

A.1 Data

We employ the publicly available micro data of Krueger and Perri (2006). Their dataset comes from the CEX and contains household level measures of consumption, male hours and male wages for the 1981-2003 period. We use the "ND" consumption measure, which includes nondurable components of household expenditures. To control for family size and composition, we normalize household consumption using a standard adult equivalent scale.⁸ Annual consumption is computed as the sum of quarterly flows, restricting our sample to households completing all four interviews of the CEX. Annual hours of work of the head of household "reference person" in the CEX are computed as the product of reported "hours usually worked per week" and reported "number of weeks worked by reference person full or part time in last 12 months, including paid vacation and paid sick leave". We measure hourly wages as total annual labor income plus transfers minus taxes, as defined in Krueger and Perri (2006), divided by annual hours of work.

Our basic sample selection criteria are the same as in the "benchmark sample" of Krueger and Perri (2006). We additionally restrict the sample to households headed by a white male that had not retired, was not self employed, working for the army, or in the forestry/fishing sector. As in Heathcote et. al. (2006), we also restrict annual hours of work to be between 520 and 5096, reflecting high attachment to the labor force. In order to reduce measurement error in hours, we further restrict the sample to heads of household who work 24 hours or more per week and 16 weeks or more per year.

A.2 Estimating Life Cycle Dispersion Profiles

The life cycle dispersion profiles displayed in Figure 1 and in Figure 2 are constructed in two steps. Step 1 is to calculate variances of log consumption, log hours, log wages and measured log preference shifters for households grouped by age and year. The construction of the preference shifter profiles in Figure 2 requires the simultaneous observation of consumption, hours, and wages for each household. This requirement considerably reduces the sample sizes within an age-year cell. To obtain reasonable sample sizes, we define a household to belong to cell (a, y) if the interview year is in $[y - 1, y + 1]$ and the age of the head of household is in

⁸The number of adult equivalents is given by $\sqrt{a + 0.5c}$, where a is the number of people over 15 years of age and c is the number of people 15 years of age or younger living in the household.

$[a - 2, a + 2]$. This means, for example, that the relevant variance for the age=28 and year=1985 cell was calculated using data on all households of ages 26 to 30 interviewed in years 1984 to 1986.

The Krueger-Perri dataset contains interviews from the first quarter of 1980 to the first quarter of 2004. Since we consider only households with four complete quarterly interviews, our dataset starts in the last quarter of 1980. Also, as in Krueger and Perri (2006), we consider households whose fourth interview was conducted in the first quarter of any year t as part of year $t - 1$'s data. We drop the 1980 observations because of cell size considerations. Our sample therefore contains data for ages $a = 25, 26, \dots, 65$ and years $y = 1981, 1982, \dots, 2003$. Consistent with our definition of age-year groups, the resulting cells are available for $a = 27, 28, \dots, 63$ and $y = 1982, 1983, \dots, 2002$. The resulting cell sizes range from 66 to 545 observations, with an average of 335 observations.

Step 2 in the construction of dispersion profiles is to run an ordinary least squares regression of the relevant variance in cell (y, a) on age dummy variables and year dummy variables. The regression equation is additive in age and year and there are no cross age-year terms. The age coefficients in Figure 1 are normalized so that the variance at age 40 equals the unconditional sample variance of each variable at age 40.

A.3 Hours Dispersion over the Life Cycle

Our model implies a positive increase in log hours dispersion over the life cycle given an increase in wage dispersion over the life cycle - see Theorem 2. The empirical literature contains at least four different versions of the graph relating age to log hours dispersion, in addition to ours: Storesletten et. al. (2001) and Heathcote et. al. (2004) use PSID data, whereas Kaplan (2007) uses both CEX and PSID data. Examination of these profiles and comparison to ours, leads us to conclude that, varying sample selection criteria and methodology, one can obtain increasing, decreasing, and U-shaped profiles from the same basic dataset.

A.3.1 Profiles Based on PSID Data

Storesletten et. al. (2001) present a profile similar to ours in shape but with a larger rise in dispersion. They focus on ages 24 to 60 and find that dispersion increases by about .08 across these ages compared to .03 in our work. Most of the rise in dispersion is concentrated after age 50.

Heathcote et. al. (2004) focus on ages 25 to 60 and find a U-shaped profile, where the variance of log hours decreases by around .03 between age 25 and age 40, and then rises by a similar amount until age 60.

Kaplan (2007) analyzes hours dispersion with respect to potential work experience rather than age. In PSID data, he finds that dispersion decreases with potential experience by around .05 over the experience cycle.⁹ His profile includes males ages 20 to 60 but with potential experience levels between 3 and 32 years.

All three profiles are based upon PSID data and upon the same age and time dummy variable regression. Thus, the differences must be accounted for by differences in sample selection and fine points related to methodology.

A.3.2 Profiles Based on CEX Data

The CEX based profile in Kaplan (2007) directly contrasts with ours. Hours dispersion in his work decreases with experience by around .05 over the experience cycle. Households aged 20 to 60 are included in the profile's computation provided that potential experience is between 3 to 32 years. Most of the decline occurs in the first 10 years of experience (around .03) with the rest of the experience cycle characterized by ups and downs of around .02 in magnitude.

To trace the sources of differences, we recalculate our age profiles under several variations of the sample selection criteria and extend our methodology to control for additional demographic variables. The main results are shown in Figure A1 and our conclusions are listed below:

(i) Since we do not restrict the sample based on experience, our sample includes a larger fraction of the older workers compared to Kaplan's. This explains why our profile *rises* at the end of the life cycle while Kaplan's does not.¹⁰ If we take Kaplan's sample selection criteria and apply our methodology, we obtain a profile that is decreasing, but ends early in the life cycle (Figure A1, Profile 1).

(ii) A criterion in our sample selection procedure, not present in Kaplan's, excludes households with less than 24 hours worked per week or less than 16 weeks worked per year. If we relax this criterion, our profile rises early in the life cycle by a similar magnitude as it does if we employ Kaplan's sample selection criteria (Figure A1 Profiles 2 and 3). Although the fraction of households in the excluded set is very small, this additional selection criterion accounts for the fact that our profile is nearly *flat* at the start of the life cycle.¹¹ All remaining differences in

⁹Kaplan defines *potential experience* = *age-years of education-6*.

¹⁰The profile in Kaplan (2007) includes workers with 3 to 32 years of potential and 20 to 60 years of age. Following his sample selection criteria as closely as possible, we find that only 53% of workers older than 45, only 30% of workers older than 49, and only 6% of workers older than 54, have less than 33 years of potential experience.

¹¹Out of all households with a male head between 20 and 65 years of age, working between 520 and 5090 hours per year, only 2.3% work less than 24 hours per week and/or less than 14 weeks per year.

sample selection criteria have a small impact on results.¹²

(iii) Controlling for demographics (i.e. race and education) has a minimal impact on our results.¹³

A.4 Theorem 1

Theorem 1: Assume A1 - A3.

- (i) *There exists a unique allocation $(c^*(s^j), l^*(s^j))$ that solves the problem of maximizing $E \left[\sum_{j=0}^J \frac{\varphi_j}{(1+n)^j} u(c(s^j), l(s^j), z(s^j)) \right]$ subject to the period resource constraint $\sum_{j=0}^J E [c(s^j) - l(s^j)w(s^j)] \frac{\varphi_j}{(1+n)^j} = 0$.*
- (ii) *The time-invariant allocation $(c^*(s^j), l^*(s^j))$ is the unique solution to the Planning Problem.*

Proof:

- (i) By A1, the following two functions are well defined for $\Lambda > 0$:

$$c(s^j; \Lambda) \equiv u_c(\cdot, z(s^j))^{-1}(\Lambda)$$

$$l_j(s^j; \Lambda) \equiv u_l(\cdot, z(s^j))^{-1}(-w(s^j)\Lambda)$$

These functions are continuous in Λ by the continuous differentiability of the period utility function and are strictly decreasing and increasing, respectively, by concavity. The resource constraint is strictly decreasing in Λ . Assumption A1 implies that there are values Λ for which the constraint is strictly positive and different values Λ for which the constraint is strictly negative. The Intermediate Value Theorem then implies that there is a positive value Λ^* at which the resource constraint holds with equality.

¹²To compare all sample selection criteria, refer to our appendix A1 and Kaplan (2007, Appendix A and B).

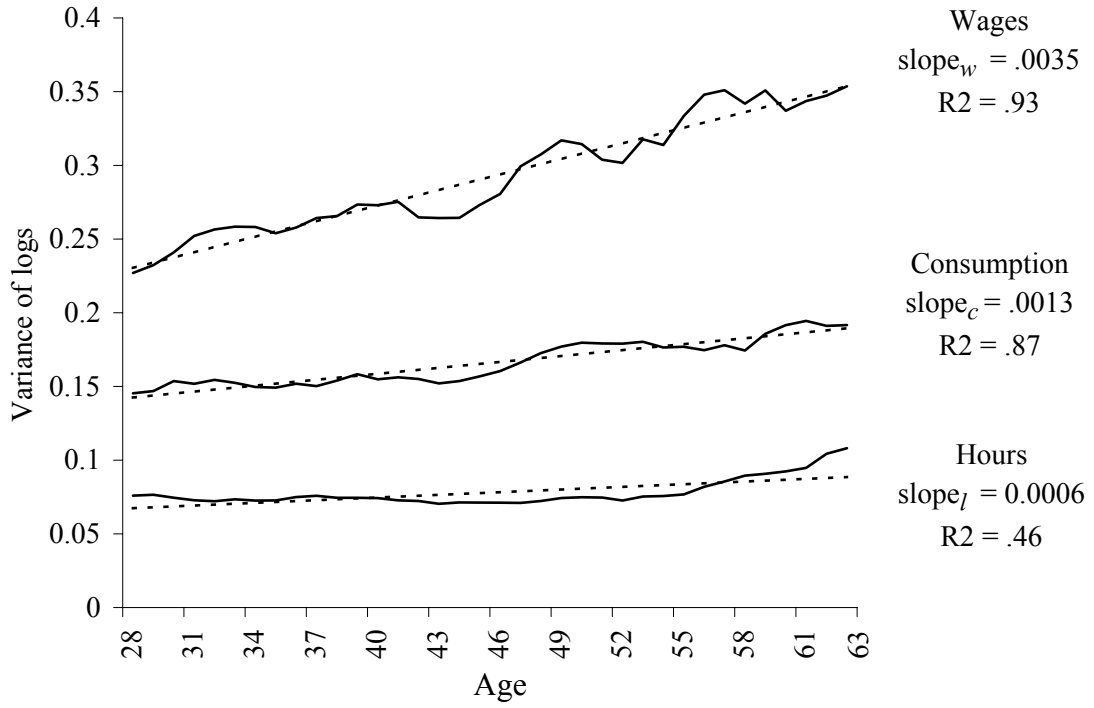
¹³We control for education and race in a separate exercise. To do this, we regress the raw data on education, race, year/race interaction, age, year, and age/year interaction dummies. Then we use those regression residuals to conduct a time-age dummy variable variance decomposition as described above. The results are almost identical to those shown in Figure A1. When using our sample selection criteria we do not include race dummies, since we allow only households with a white Reference Person.

The candidate allocation is $(c^*(s^j), l^*(s^j)) = (c(s^j; \Lambda^*), l(s^j; \Lambda^*))$. This allocation satisfies the Kuhn-Tucker conditions for a solution to this problem. As these conditions are sufficient conditions for finite-dimensional, concave maximization problems, the candidate allocation solves the problem. To establish uniqueness, note that if there were a different feasible allocation solving this problem, then a convex combination of the two solutions would be feasible, by the convexity of the resource constraint, and would increase the objective since the objective is strictly concave. Contradiction.

(ii) The time-invariant allocation $(c^*(s^j), l^*(s^j))$ from Theorem 1(i) satisfies the resource constraint to the Planning Problem each period and delivers a finite value for the objective function in the Planning Problem by assumptions A2 and A3. We now argue that there do not exist solutions which deliver an infinite value for the objective function. Suppose by way of contradiction that there is such a solution, then in some time period the value of the period objective must exceed the value implied by the solution constructed in Theorem 1(i). Contradiction.

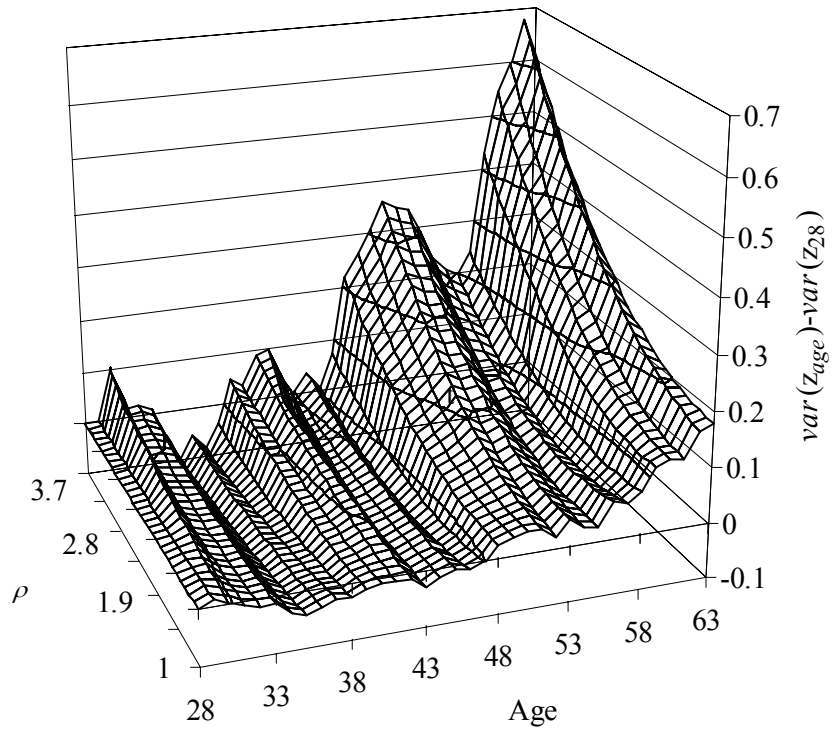
We now argue that $(c^*(s^j), l^*(s^j))$ is the unique solution to the Planning Problem. First, any feasible allocation leading to a greater value of the objective must produce a greater value in some period. By Theorem 1(i) this can not hold. Thus, $(c^*(s^j), l^*(s^j))$ solves the Planning Problem. Second, it is unique as any alternative feasible allocation must by Theorem 1(i) deliver strictly less utility at some time period. \diamond

Figure 1. Age Variance Profile of Wages, Consumption and Hours



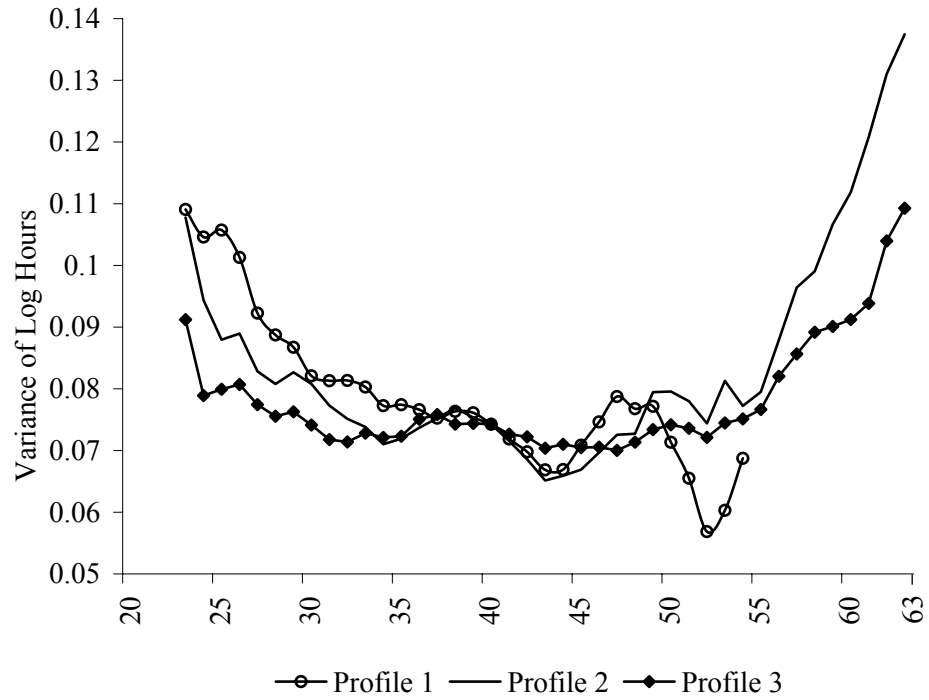
Source: Author's calculations based upon CEX data 1981-2003.

Figure 2. Change in the Variance of Preference Shifters



Source: Author's calculations based upon CEX data 1981-2003.

Figure A1. Dispersion of log Hours by Age and Sample Selection Criteria



Profile 1: Obtained by following the selection criteria in Kaplan (2007), and applying our methodology.

Profile 2: Obtained by applying our selection criteria but relaxing the condition that households work more than 23 hours per week and more than 13 weeks per year. Also, we have extended the sample to households aged 20 to 24, for comparison purposes.

Profile 3: Obtained by applying our benchmark sample selection criteria extended to include ages 20 to 24.

Note: The profiles are normalized so that variances at age 40 equal the unconditional sample variance at age 40 under our sample selection criteria.