Inequality and the Lifecycle

Author: Greg Kaplan

Presentation: Dan Greenwald

November 13, 2012

Motivation

- Literature has been largely successful in accounting for the evolution of both average consumption and its variance across agents in an incomplete markets, lifecycle setting.
- However, these models have not similarly addressed inequality in labor supply, which is important for two reasons:
    1. Modeling the labor supply decision accurately can help us to understand the structure of labor markets.
    2. Labor supply enters the utility function directly, and so is essential for measuring welfare.
- Standard incomplete markets models with frictionless labor supply decisions cannot account for both cross-sectional data on consumption and hours — a robust result.
- Instead, this paper will incorporate involuntary unemployment and nonseparability between disutility of the extensive and intensive margins of work to match the data.

Demographics

- Continuum of households, indexed by i.
- Agents work until T^{ret} and then retire.
- Mortality risk: unconditional probability S_t of surviving until age t, certain death at age T.
- No bequest motive.
- Perfect annuity markets to hedge against survival risk (i.e. higher interest rate on savings).

Preferences

- Time-separable expected utility preferences over annual consumption c_{it} and hours worked H_{it}:

$$\mathbb{E} \sum_{t=1}^{T} \beta^{t-1} S_t \left[\frac{c_{it}^{1-\gamma}}{1-\gamma} - \varphi_i \frac{H_{it}^{1+\sigma}}{1+\sigma} \right]$$

- φ_i is allowed to vary across households (but is fixed over time).

Technology

- Each household consists of a single worker with a fixed time endowment (1) that is divided between work and leisure.
- Worker receives labor income $y_{it} = w_{it} H_{it}$, where w_{it} is an individual-specific exogenous wage process given by

$$\log w_{it} = \kappa_t + \alpha_i + z_{it} + \varepsilon_{it}$$
$$z_{it} = \rho z_{it} + \eta_{it}$$

- Asset markets are limited to a risk-free security a_{it} that pays interest rate R. Holdings can be negative up to an exogenous borrowing limit $\underline{a}$.

Government

- Taxes:
  1. Progressive tax on labor income, $\zeta(y)$.
  2. Proportional tax on capital income τ_a.
  3. Proportional consumption tax τ_c.
- Progressive social security system, where benefits are functions of α_i, the fixed component in wages.

Data

- Data on wages, hours, and wealth from the PSID, with long panel dimension.
- Consumption data (cross-sectional) from CEX, also contains data on earnings and hours for constructing joint moments.
- Allows for the possibility that log hours, earnings, and consumption contain classical measurement error.
- Focus of the paper is on residual inequality, so first-stage regression is performed to purge effects of experience, education, year/cohort, and race.
- Year/cohort effects can not be separately identified along with lifecycle effects, so separate estimation is performed controlling for each.

Empirical Patterns

- First moments:
  1. Mean log consumption increases roughly linearly during working years.
  2. Mean log hours has an inverse U shape: increases over the first 10 years, then flattens, and decreases before retirement.
- Second moments:
  1. Variance of log consumption increases by less than the variance of log wages.
  2. Variance of log hours is strongly U shaped: sharp decrease over first 15 years.
     - This decrease is a robust finding across other data sets, and is not driven by observable characteristics or compositional effects.
     - Will cause problems for the benchmark model.

Estimation with First Moments

- Estimation performed using simulated method of moments.
- Targeted moments are mean residual log hours and log (normalized) consumption at 36 experience levels during working life only.
- Estimated parameters are $(\gamma, \sigma, \underline{a}, \{\phi_i\})$.
- Internally calibrated parameters: $\mathbb{E}[\varphi_i]$ (to match mean log hours), β (to match 75th percentile of wealth at retirement).
- Externally calibrated parameters: $R = 3\%$, $\tau_a = 0.4$, $\tau_c = 0.08$, $\zeta(\cdot)$, distribution of a_{i0}, stochastic process for individual wages.

	(1)	(2)	(3)
Estimated parameters			
γ	1.60	1.86	1.65
	(1.22,1.98)	(1.57,2.68)	(1.36,2.15)
σ	2.00	2.48	1.94
	(1.31,3.13)	(1.37,5.86)	(1.31,2.67)
CV[φ]	1.27	0.98	1.16
	(0.65, 142.0)	(0.55,3.40)	(0.72,9.77)
$\underline{a}$	-0.30		-0.27
	(-6.90,-0.10)		(-0.63,0.09)
$\sigma^2_{\epsilon,y}$			0.00
			(0.00,0.05)
$\sigma^2_{\epsilon,h}$			0.03
			(0.00,0.05)
$\sigma^2_{\epsilon,c}$			0.05
			(0.03,0.08)
Test of over-identifying restrictions			
p-value	0.45	0.00	0.00
Calibrated parameters			
β	0.990	0.988	0.990
E[φ]	29.4	67.2	30.3
Targeted moments			
E[$\log c$]	X	X	X
E[$\log h$]	X	X	X
V[$\log c$]			X
V[$\log h$]			X

Table 1: Parameter estimates

Estimation Results

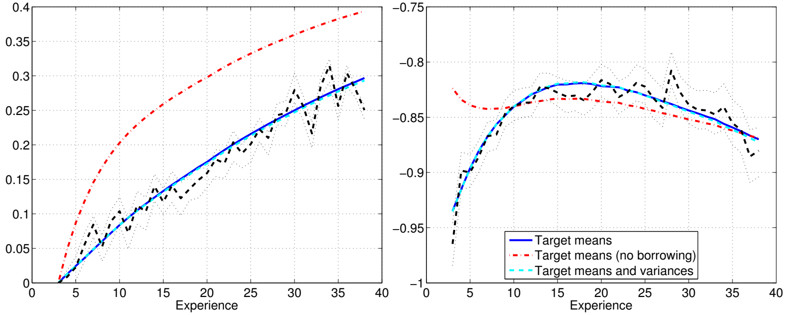

- Benchmark model does well in fitting first moments.
- Ability to borrow (esp. for young agents) is critical to obtaining this fit.

Estimation with Second Moments

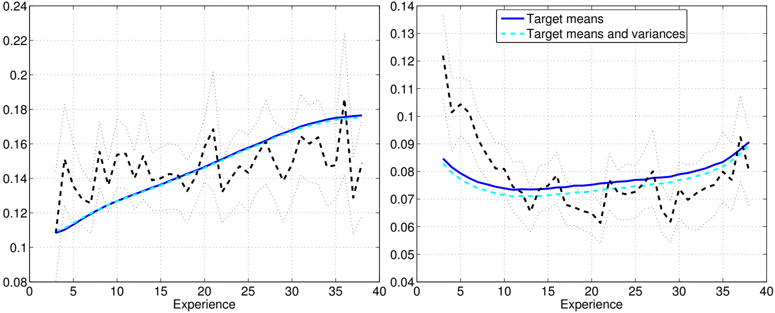

(a) Variance log consumption (b) Variance log hours

- However, the model fails to fit second moments, even if they are targeted in the estimation.
- In particular, fails to capture the decline in hours inequality early in the lifecycle.
- Test of overidentifying restrictions including second moment conditions is overwhelmingly rejected.

Problems with the Benchmark Model

- First order condition for the choice of hours:

$$w_{it} c_{it}^{-\gamma} = \varphi_i H_{it}^{\sigma}$$

- This implies

$$\sigma^2 \Delta \text{Var}(\log H_{it}) = \Delta \text{Var}(\log w_{it}) + \gamma^2 \Delta \text{Var}(\log c_{it}) - 2\gamma \text{Cov}(\log c_{it}, \log w_{it}) + 2\gamma \text{Cov}(\Delta \log c_{it}, \log \varphi_i)$$

- From the data, the LHS is negative at young ages, but the variance of wages and consumption are both increasing in the data, and so the RHS in the data is positive for any plausible γ.
- Note that this does not depend on the specific wage process chosen, only on moments, so this is robust across all specifications with this FOC.

Problems with the Benchmark Model

- A mechanical way to generate a decline in Var(log H_{it}) would be to allow for time-varying individual shocks to the disutility of working, φ_{it}, but lacks underlying mechanism.
- From data, appears that declining variance of hours is not driven by observable characteristics (i.e. marraige, fertility, home ownership).
- Instead, will introduce age-varying labor market frictions, implied by observed age differences in job destruction rates.
- Most of the variance in log H_{it} is in weeks per year (65%), not hours per week (25%), and is concentrated in the bottom half of the distribution.

Unemployment Hypothesis

- Hypothesis: decline in inequality at the bottom of the hours distribution is due to age differences in the incidence of unemployment.
- Backed by PSID data: fraction of sample experiencing at least one week of unemployment during the calendar year drops from 25% to 5% over this period.
- Effect comes from incidence, not duration: duration of unemployment conditional on experiencing spell is very consistent with age.

Modeling Unemployment

- Extend the benchmark model so that each individual's time endowment is $\bar{h}_{it} \in [0, 1]$, so that $1 - \bar{h}_{it}$ is the fraction of the year spent unemployed.
- To prevent workers from unrealistically loading up on hours during time employed, model preferences over $(h_{it}, \bar{h}_{it})$:

$$v(h, \bar{h}) = \frac{(\bar{h}^{\chi} h)^{1+\sigma}}{1+\sigma}.$$

- With $\chi = 1$, individuals only care about total hours. With $\chi = 0$, individuals only care about intensity of hours when employed. With $\chi < 0$, individuals like working while employed less the longer they are unemployed.

Modeling Unemployment

- First order condition is now

$$w_{it} c_{it}^{-\gamma} = \bar{h}_{it}^{(\chi-1)(1+\sigma)} \varphi_i H_{it}^{\sigma}.$$

- This introduces a term

$$(\chi - 1)^2 (1 + \sigma)^2 \Delta \mathrm{Var}(\log \bar{h}_{it})$$

into the expression for $\sigma^2 \Delta \mathrm{Var}(\log H_{it})$. A decrease in this term in the data will allow for the desired fall in hours inequality.

- Also introduces a term

$$2(\chi - 1)(1 + \sigma) \Delta \mathrm{Cov}(\log \bar{h}_{it}, \log \varphi_i)$$

which can be used to model "lazy types" who are more likely to become unemployed. If this correlation is positive, and $\chi - 1 < 0$, then this can also have a negative effect.

Unemployment Process

- Treat unemployment as an exogenous shock, with an age-varying distribution.
- Conditional on receiving a shock, $\bar{h}_{it}$ takes one of two discrete values with equal probability, to match the mean and variance of time spent unemployed.
- Realizations of $\bar{h}_{it}$ are independent across time, which is backed up by data (autocorrelation after controlling for individual fixed effects is small).

	(1)	(2)	(3)	(4)
Estimated parameters				
γ	2.33	2.47	2.24	2.49
	(1.86,3.48)	(1.82, 3.68)	(2.06, 2.92)	(2.03, 2.96)
σ	2.10	2.26	2.41	3.05
	(1.49, 3.69)	(1.56, 4.27)	(1.72, 3.95)	(1.75, 4.44)
CV[φ]	0.21	0.23	0.80	0.77
	(0.00, 0.52)	(0.08, 0.82)	(0.71, 0.94)	(0.71, 0.93)
$\underline{a}$	-0.22	-0.22	-0.17	-0.19
	(-0.44, -0.12)	(-0.45, -0.11)	(-0.41, -0.04)	(-0.39, -0.07)
$\sigma^2_{\epsilon,y}$	0.02	0.01	0.10	0.06
	(0.00, 0.20)	(0.00, 0.07)	(0.04, 62.32)	(0.04, 15.48)
$\sigma^2_{\epsilon,h}$	0.02	0.02	0.00	0.00
	(0.00, 0.04)	(0.00, 0.04)	(0.00, 0.01)	(0.00, 0.00)
$\sigma^2_{\epsilon,c}$	0.08	0.08	0.06	0.06
	(0.06, 0.09)	(0.06, 0.11)	(0.05, 0.07)	(0.05, 0.07)
χ	-0.36		-0.08	
	(-0.52,-0.23)		(-0.22, 0.44)	
$\rho_{\varphi,\bar{h}}$		0.74		0.24
		(-0.02, 1.09)		(-0.26, 0.58)
Test of over-identifying restrictions				
p-value	0.28	0.26	0.00	0.00
Calibrated parameters				
β	0.989	0.989	0.989	0.987
Eφ	66.4	88.4	115.1	305.0
Targeted moments				
E[$\log c$]	X	X	X	X
E[$\log h$]	X	X	X	X
V[$\log c$]	X	X	X	X
V[$\log h$]	X	X	X	X
COV[$\log w, \log h$]			X	X
COV[$\log c, \log h$]			X	X
COV[$\log w, \log c$]			X	X

Table 2: Parameter estimates with shocks to time endowment

Estimates

- χ is found to be slightly negative (-0.36).
- Negative χ implies a positive cross-sectional correlation between the fraction of the year employed and hours worked when employed, consistent with the data.
- Test of over-identifying restrictions no longer rejected even if variance moments are matched.
- A good fit can also be obtained with $\chi = 0$ and "lazy types" (i.e. positive covariance between $\bar{h}_i$ and φ_i).

Estimation Results

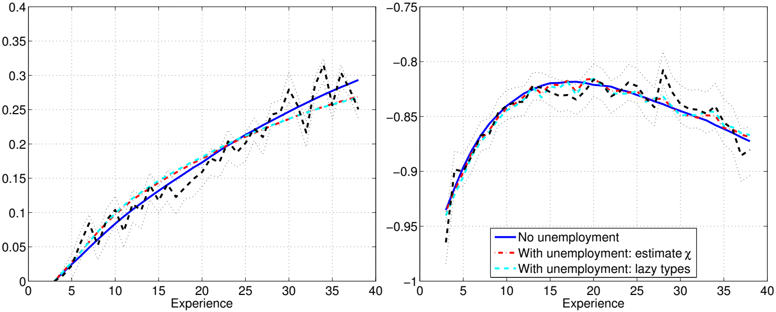

(a) Mean log consumption (b) Mean log hours

- As before, model is able to match first moments well.

Estimation Results

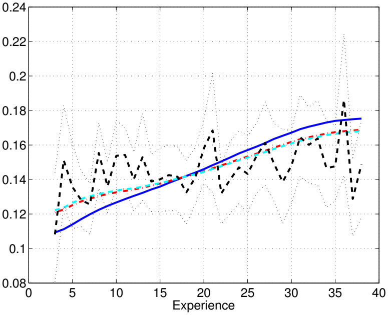

(a) Variance log consumption

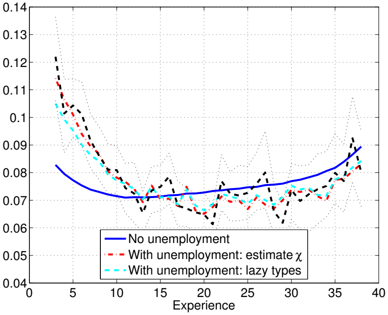

(b) Variance log hours

- However, model with unemployment is also able to get a very good fit of the second moments, including early fall in hours inequality.

Estimation Results

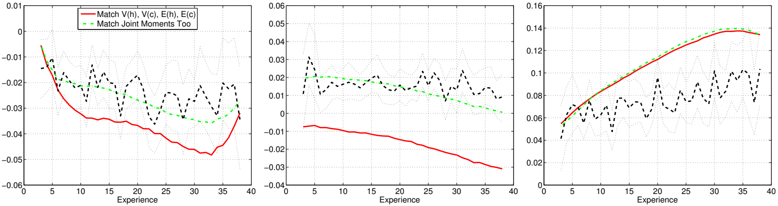

(a) Covariance log wages, log hours (b) Covariance log consumption, log hours (c) Covariance log wages, log consumption

- Unemployment model does less well matching covariance of wages, hours, consumption.
- However, if these covariances are targeted as well, the fit is much better, with the exception of (w, c) covariance.

Estimation Results

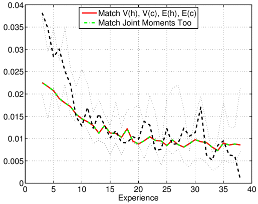

(a) Variance of $\log \bar{h}$

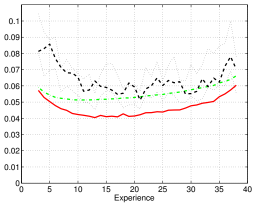

(b) Variance of $\log h$

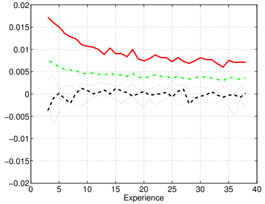

(c) Covariance $\log h$, $\log \bar{h}$

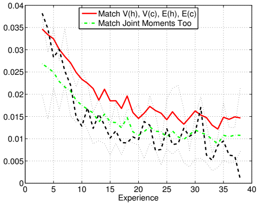

(d) Variance of $\log \bar{h}$

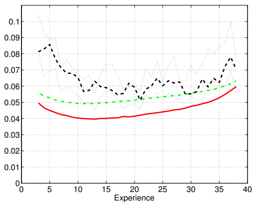

(e) Variance of $\log h$

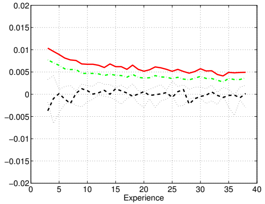

(f) Covariance $\log h$, $\log \bar{h}$

- Model does a pretty good job of matching intensive and extensive $(h, \bar{h})$ moments, especially if (w, c, h) covariances are targeted.

Estimates

- Model also matches key features of non-human wealth distribution: declining fraction of the population with non-positive wealth, and positive correlation between log wages and log wealth.
- Fit of non-human wealth inequality not as good.
    - Early in the lifecycle, model gini is too large (too many agents with nonpositive wealth).
    - Later in the lifecycle, model gini is too small (upper tail not fat enough).

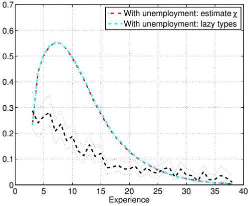

(a) Fraction with negative or zero assets

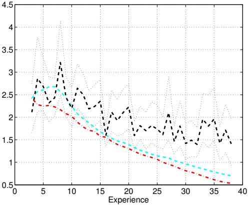

(b) Variance log wealth for wealth > 0

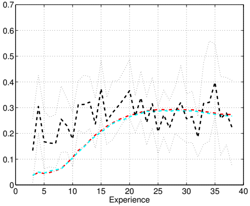

(c) Covariance log wealth, log wages

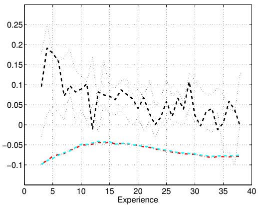

(d) Covariance log wealth, log hours

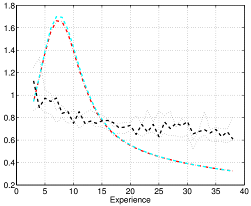

(e) Gini coefficient for wealth

Conclusion

- Standard model's intratemporal FOC produces very strong, counterfactual restrictions on joint distribution of consumption, wages, and hours, regardless of wage process.
- Modification to allow for exogenous unemployment and nonseparable preferences allows the model to correctly match the first and second moments of consumption, hours, and wages.