Buffer-Stock Saving in a Krusell–Smith World

January 22, 2015

Christopher D. Carroll[1] — JHU
Jiri Slacalek[2] — ECB
Kiichi Tokuoka[3] — MoF, Japan

Abstract

A large body of evidence supports Friedman (1957)'s proposition that household income at the microeconomic level can be reasonably well described as having both transitory and permanent components. We show how to modify the widely-used macroeconomic model of Krusell and Smith (1998) to accommodate a Friedmanesque income process. Our incorporation of appropriately calibrated permanent income shocks helps our model to explain a substantial part of the large degree of empirical wealth heterogeneity that is unexplained in the baseline Krusell and Smith (1998) model.

Keywords Household Income Process, Wealth Inequality

JEL codes D12, D31, D91, E21

PDF: http://econ.jhu.edu/people/ccarroll/papers/cstKS.pdf
Slides: http://econ.jhu.edu/people/ccarroll/papers/cstKS-Slides.pdf
Web: http://econ.jhu.edu/people/ccarroll/papers/cstKS/
Archive: http://econ.jhu.edu/people/ccarroll/papers/cstKS.zip

[1]Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD, http://econ.jhu.edu/people/ccarroll/, ccarroll@jhu.edu; and United States Consumer Financial Protection Bureau [2]Slacalek: European Central Bank, Frankfurt am Main, Germany, http://www.slacalek.com/, jiri.slacalek@ecb.europa.eu [3]Tokuoka: Ministry of Finance, Tokyo, Japan, kiichi.tokuoka@mof.go.jp

The views presented are those of the authors, and should not be attributed to the European Central Bank or to the Japanese Ministry of Finance. We thank Tamas Briglevics, Luigi Pistaferri, numerous seminar audiences, the editor and the referee for useful comments, and Ivan Vidangos for sharing with us the data from DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013), Figure IV(a). The long version of the paper is available as Carroll, Slacalek, and Tokuoka (2014), http://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1633.pdf.

1 Introduction

The Krusell and Smith (1998) ('KS') method for incorporating uninsurable household-level risk into macroeconomic models is a workhorse in macroeconomic modeling. However, the stochastic process that KS use to characterize household income dynamics is strongly inconsistent with microeconomic evidence. Since the point of the KS analysis was to examine *quantitative* implications of idiosyncratic uncertainty, it is hard to be confident about any of their (quantitative) conclusions if their calibration of the idiosyncratic risk is (quantitatively) implausible.

A large empirical literature, with a history traceable all the way back to Friedman (1957), finds that a simple income process consisting of a permanent and a transitory component—what we call the "Friedman/Buffer Stock" (FBS) process—captures the key features of the microeconomic data well.

In this paper, we solve a modified version of the KS model in which the microeconomic household income process has the FBS structure.

The microeconomic evidence indicates that permanent shocks are large. Consequently, even if households were to start life with the same levels of permanent income, long-lived households will end up differing greatly in their levels of permanent income. Because the model implies the existence of a target *ratio* of assets to permanent income, the model implies that the equilibrium cross-sectional wealth distribution is roughly as unequal as the distribution of permanent income.

Among the many ways to quantify the importance of our modification, perhaps the most interesting is to note that in our simulations, the top 1 percent of households are almost three times richer than the top 1 percent in the baseline KS model (although even our model falls short of the degree of inequality found in the data).

2 The Income Process

Friedman (1957) famously characterized income as having permanent and transitory components. A large subsequent empirical literature has confirmed the essential accuracy of this description (see e.g., Meghir and Pistaferri (2011) for a survey of the (vast) literature).[1]

Motivated by Friedman (1957) our household income process consists of permanent income p, a transitory shock ξ, and the wage rate W:

$$\boldsymbol{y}_t \quad = \quad p_t \xi_t \mathsf{W}_t. \tag{1}$$

[1] There is some evidence that the component Friedman labeled "permanent" may be merely highly persistent, with a serial correlation coefficient of, say, 0.98 or 0.99. For our purposes, we will stick with the original Friedman specification, in which the serial correlation coefficient is 1.0. Little of substance flows from this modest discrepancy, but models where the coefficient is not exactly one are harder to solve and to understand.

The permanent component follows a random walk:

$$p_t = p_{t-1}\psi_t, \tag{2}$$

where the Greek letter psi mnemonically indicates the **p**ermanent **s**hock to **i**ncome. Because the form (1)–(2) has also been used widely in the literature on buffer stock saving, we call it the "Friedman/Buffer Stock" (FBS) process.

Following the assumptions in the the special volume of the *Journal of Economic Dynamics and Control (2010)* devoted to comparing solution methods for the KS model, the distribution of ξ_t is:

$$\begin{aligned} \xi_t &= \mu \text{ with probability } u_t, \\ &= (1-\tau_t)\ell\theta_t \text{ with probability } 1-u_t, \end{aligned}$$

where $\mu > 0$ is the unemployment insurance payment when unemployed, $\tau_t = \mu u_t/\ell \boldsymbol{L}_t$ is the rate of tax collected to pay unemployment benefits, and ℓ is time worked per employee.

3 The Model

Households maximize discounted utility from consumption

$$\max \sum_{n=0}^{\infty} \beta^n \mathrm{u}(\boldsymbol{c}_{t+n})$$

for a CRRA utility function $\mathrm{u}(\bullet) = \bullet^{1-\rho}/(1-\rho)$. The aggregate production function is

$$a_t \boldsymbol{K}_t^{\alpha}(\ell \boldsymbol{L}_t)^{1-\alpha}, \tag{3}$$

where a_t is aggregate productivity, $\boldsymbol{K}_t$ is capital, and $\boldsymbol{L}_t$ is aggregate employment..

The decision problem for the household in period t can be written using variables normalized with the level of permanent income $\boldsymbol{p}_t = p_t \mathsf{W}_t$ (e.g., $c_t = \boldsymbol{c}_t/\boldsymbol{p}_t$). The consumption functions $\mathrm{c}_t(m_t)$ produce levels of c_t that satisfy:

$$\begin{aligned} \mathrm{v}(m_t) &= \max_{c_t} \; \mathrm{u}(c_t) + \beta \not{\mathsf{D}} \mathbb{E}_t \left[\psi_{t+1}^{1-\rho} \mathrm{v}(m_{t+1})\right] && (4) \\ & \text{s.t.} \\ a_t &= m_t - c_t, && (5) \\ k_{t+1} &= a_t/(\not{\mathsf{D}}\psi_{t+1}), && (6) \\ m_{t+1} &= \left(\daleth + \mathsf{r}_{t+1}\right)k_{t+1} + \xi_{t+1}, && (7) \\ a_t &\geq 0, \end{aligned}$$

where $\not{\mathsf{D}} = 1 - \mathsf{D}$ denotes the probability of *not* dying between periods (see below),

$\daleth$ is the depreciation factor for capital, and the gross interest rate r_{t+1} is calculated as marginal aggregate output from capital.

The transition process for market resources m_t is broken up, for clarity of analysis, into three steps (5)–(7). Equation (5) indicates that assets at the end of the period are equal to market resources minus consumption, while next period's capital is determined from this period's assets via equation (6). The final step (7) describes the transition from the beginning of period $t+1$ when capital has not yet been used to produce output, to the middle of that period, when output has been produced and incorporated into resources but has not yet been consumed. These steps can of course be combined into a single transition equation, as is usual in solving problems like this:

$$m_{t+1} \quad = \quad \left(\daleth + \mathsf{r}_{t+1}\right)\left((m_t - c_t)/(\not{D}\psi_{t+1})\right) + \xi_{t+1}, \tag{8}$$

whose impenetrability motivates our disarticulation into bite-sized pieces in (5)-(7).

If the probability of death is $\mathsf{D} = 0$, the FBS income process implies there is no ergodic distribution of permanent income in the population: Since each household accumulates a permanent shock in every period, the cross-sectional distribution of idiosyncratic permanent income becomes wider and wider indefinitely as the simulation progresses.

We address this problem by assuming that households have finite lifetimes *a la* Blanchard (1985). Death follows a Poisson process, so that every agent alive at date t has an equal probability D of dying before the beginning of period $t+1$. Households engage in a Blanchardian mutual insurance scheme: Survivors share the estates of those who die. Assuming a zero profit condition for the insurance industry, the insurance scheme's ultimate effect is simply to boost the rate of return (for survivors) by the mortality rate.

To maintain a constant population, dying households are replaced by an equal number of newborns who begin life with idiosyncratic permanent income equal to the population mean. Mean idiosyncratic permanent income will thus remain fixed at 1 forever, while the variance of p is finite as long as $\not{D}\mathbb{E}[\psi^2] < 1$ (a requirement that does not do violence to the data). Intuitively, permanent income among survivors does not spread out so quickly as to overwhelm the compression of distribution due to death and replacement.

4 Calibration

We first calibrate the income process using the empirical literature (see Carroll, Slacalek, and Tokuoka (2014), Table 1, Meghir and Pistaferri (2011) and many others) and households' subjective estimates of their permanent income. In line with this evidence, we set the variance of the transitory and permanent income components at $\sigma_\theta^2 = 0.010 \times 4$ and $\sigma_\psi^2 = 0.010 \times 4/11$, specifically following the estimates of Carroll (1992).[2] In a sharp contrast, the income process used in the KS model implies estimates of σ_ψ^2 and σ_θ^2 that are orders of magnitude different from what the literature finds in actual data.

Figure 1 illustrates further the difficulties of the KS process. A key feature of the data is that the cross-sectional variance of the income profiles $\text{var}(\log \boldsymbol{y}_{t+r,i} - \log \boldsymbol{y}_{t,i})$ tends to grow linearly with the horizon r, with slope σ_ψ^2, mirroring closely the characteristics of the FBS process, where $\text{var}(\log \boldsymbol{y}_{t+r,i} - \log \boldsymbol{y}_{t,i}) = 2\sigma_\xi^2 + \sigma_\psi^2 \times r$. In contrast, the statistic for the KS process does not exhibit any trend, also reflecting the fact that the first autocorrelation of the KS income process is only roughly 0.2.

As a second check, we make sure that the parameters of the income process are in line with households' subjective estimates of their permanent income. The Survey of Consumer Finances (SCF) conveniently includes a question asking respondents whether their income in the survey year was about 'normal' for them, and if not, it asks the level of 'normal' income. The question corresponds well with our (and Friedman (1957)'s) definition of permanent income p. We calculate the cross-sectional variance σ_p^2 of $\boldsymbol{p}$ and, eventually, the implied variance of ψ. Reassuringly, the estimates of σ_ψ^2 from the 1992–2010 SCF data range between 0.015 and 0.018, closely in line with the estimates reported in Meghir and Pistaferri (2011). Such a correspondence, across two quite different methods of measurement, suggests there is considerable robustness to the measurement of the size of permanent shocks.

To calibrate the time preference factor $\grave{\beta}$ we shut down aggregate uncertainty and search for the value at which the steady-state capital-to-output ratio $\boldsymbol{K}/\boldsymbol{Y}$ matches the steady-state ratio in the perfect foresight model. Our remaining parameters match the values in the special JEDC (2010) volume; see Table 2 in the Appendix.

5 Matching the Wealth Distribution

We now ask whether our model with realistically calibrated income and finite lifetimes can reproduce the degree of wealth inequality evident in the micro data. An improvement in the model's ability to match the data (over the KS model) is to be expected, since (as noted above), in buffer stock models agents strive to achieve a target *ratio* of wealth to permanent income. By assuming no dispersion in the level of permanent income across households, KS's income process disables a potentially

[2] The empirical literature estimates variance of income shock using annual data. We back out the corresponding quarterly values as follows. For the variance of permanent shocks, σ_ψ^2, aggregating equation (2) to the annual frequency implies: $(p_{t+4} + p_{t+3} + p_{t+2} + p_{t+1})/4 = (p_t + p_{t-1} + p_{t-2} + p_{t-3})/4 + (\sum_{i=1}^4 \psi_{t+i} + \sum_{i=1}^4 \psi_{t-1+i} + \sum_{i=1}^4 \psi_{t-2+i} + \sum_{i=1}^4 \psi_{t-3+i})/4$. Consequently, the variance of the shocks on the right-hand side is: $\text{var}((\psi_{t+4} + 2\psi_{t+3} + 3\psi_{t+2} + 4\psi_{t+1} + 3\psi_t + 2\psi_{t-1} + \psi_{t-2})/4) = \text{var}(\psi)(1+4+9+16+9+4+1)/16 = 11/4 \times \text{var}(\psi)$. For transitory shocks, going from the quarterly frequency to annual implies dividing by 4 as we average over the four quarters.

vital explanation for variation in the level of target wealth (and, therefore, on average, actual wealth) across households.

Table 1 shows that the models with the FBS income process (in columns 3 through 6) do indeed yield a substantial improvement in matching the data (in the last column) over the distribution of wealth implied by our solution of the KS model (column 2), as parametrized in the JEDC (2010) volume (which we call 'KS-JEDC'). For example, in our model with FBS income dynamics and no aggregate uncertainty (column 3), the proportion of total net worth held by the top 1 percent is 11.5 percent, while the corresponding statistic in the KS-JEDC model is only 2.7 percent.

The failure of the KS-JEDC model to match the wealth distribution is not confined to the top. In fact, perhaps a bigger problem is that most households in the model hold wealth levels not very far from the wealth target of a representative agent in the perfect foresight version of the model. For example, in steady state about 50 percent of households have wealth between 0.5 and 1.5 times the mean wealth; in the SCF data from 1992–2004, the corresponding proportion ranges from only 20 to 25 percent.

But while our model fits the data better than the KS-JEDC model, it still falls short of matching the empirical degree of wealth inequality. The proportion of wealth held by households in the top 1 percent is about three times smaller in the model than in the data (compare the third and last columns). This failure reflects the fact that, empirically, the distribution of wealth is considerably more unequal than the distribution of permanent income.[3]

A comparison of columns 3 and 4 shows that the FBS models without and with KS aggregate uncertainty do roughly equally well in matching the wealth distribution.[4] The fact that the specification of aggregate shock has little effect on the performance of the model in this respect is not surprising because it is well-known that aggregate shocks are much smaller than idiosyncratic shocks.

Columns 5 and 6 explore the effects of higher variances of income shocks. A significantly higher size of transitory shocks (setting $\sigma_\theta^2 = 0.15$, in line with the upper range of empirical estimates, DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013)) results in a stronger precautionary motive and a somewhat more even wealth distribution. In contrast, allowing for larger permanent shocks ($\sigma_\psi^2 = 0.030$) implies a less equal wealth distribution (because our model is scalable with permanent income).

Finally, statistics on aggregate dynamics (reported in Carroll, Slacalek, and Tokuoka (2014)) are generally similar for the KS and the FBS models, implying

[3] Carroll, Slacalek, and Tokuoka (2013) show that models with a modest amount of heterogeneity in impatience—allowing quarterly discount factors of households to range between 0.980 and 0.992—are able to match the wealth distribution very well.

[4] The match of the wealth distribution worsens slightly with aggregate uncertainty as households increase a bit their precautionary saving, responding to somewhat higher overall uncertainty.

a positive autocorrelation of consumption growth and a high contemporaneous correlation of consumption growth with income growth and interest rates.

6 Conclusion

This paper has two main results.

First, we have shown how to incorporate a (quantitatively) realistic microeconomic income process in a KS-type model.

Second, we have shown that while this modification substantially improves the fit between the model and the data, the empirically observed degree of inequality is even greater than that implied even by our modified model.

Thus, some mechanism other than shocks to permanent income seems necessary to bring this class of models into full alignment with the available data on wealth inequality (or, another class of models is needed).

References

Blanchard, Olivier J. (1985): "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 93(2), 223–247.

Carroll, Christopher D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992(2), 61–156, http://econ.jhu.edu/people/ccarroll/BufferStockBPEA.pdf.

Carroll, Christopher D., Jiri Slacalek, and Kiichi Tokuoka (2013): "The Distribution of Wealth and the Marginal Propensity to Consume," mimeo, Johns Hopkins University, At http://econ.jhu.edu/people/ccarroll/papers/cstMPC/.

——— (2014): "Buffer-Stock Saving in a Krusell–Smith World," working paper 1633, European Central Bank.

DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos (2013): "Rising Inequality: Transitory or Permanent? New Evidence from a Panel of US Tax Returns," *mimeo.*

Friedman, Milton A. (1957): *A Theory of the Consumption Function*. Princeton University Press.

Journal of Economic Dynamics and Control (2010): "Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty," *edited by Wouter J. Den Haan, Kenneth L. Judd, Michel Juillard*, 34(1), 1–100.

Krusell, Per, and Anthony A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

Meghir, Costas, and Luigi Pistaferri (2011): "Earnings, Consumption and Life Cycle Choices," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, chap. 9, pp. 773–854. Elsevier.

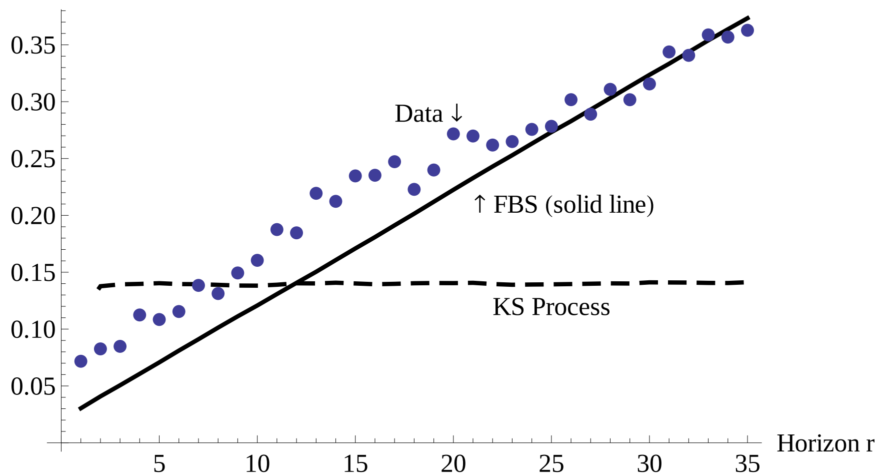

Figure 1 Cross-Sectional Variance of Income Processes and Data, $\mathrm{var}(\log \boldsymbol{y}_{t+r,i} - \log \boldsymbol{y}_{t,i})$

Notes: The data are based on DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013), Figure IV(a) and were normalized so that the variance for $r = 1$, $\mathrm{var}(\log \boldsymbol{y}_{t+1,i} - \log \boldsymbol{y}_{t,i})$ lie in the middle between the values for the KS and the FBS processes.

Table 1 Proportion of Net Worth by Percentile in Models and the Data (in Percent)

	Income Process					
	KS-JEDC	Friedman/ Buffer Stock‡				
	Our Solution	No Aggr Unc	KS Aggr Unc			
Percentile of Net Worth		$\sigma_\psi^2 = 0.01$ $\sigma_\theta^2 = 0.01$	$\sigma_\psi^2 = 0.01$ $\sigma_\theta^2 = 0.01$	$\sigma_\psi^2 = 0.01$ $\sigma_\theta^2 = 0.15$	$\sigma_\psi^2 = 0.03$ $\sigma_\theta^2 = 0.01$	Data*
Top 1%	2.7	11.5	9.1	8.8	15.0	33.9
Top 10%	20.2	38.9	35.9	35.3	44.8	69.7
Top 20%	35.6	55.3	52.4	51.9	60.0	82.9
Top 40%	60.0	76.5	74.1	74.0	78.4	94.7
Top 60%	78.5	89.7	88.2	88.2	89.8	99.0
Top 80%	92.1	97.4	96.8	96.9	97.0	100.2

Notes: $\boldsymbol{K}_t/\boldsymbol{Y}_t = 10.3$. ‡ : $\grave{\beta} = 0.9894$ when $(\sigma_\psi^2, \sigma_\theta^2) = (0.01, 0.01)$, $\grave{\beta} = 0.9888$ when $(\sigma_\psi^2, \sigma_\theta^2) = (0.01, 0.15)$ and $\grave{\beta} = 0.9862$ when $(\sigma_\psi^2, \sigma_\theta^2) = (0.03, 0.01)$. * : The data is the SCF 2004.

Appendix

Table 2 Parameter Values and Steady State

Description	Parameter	Value	Source
Representative agent model			
Time discount factor	β	0.99	JEDC (2010)
Coefficient of relative risk aversion	ρ	1	JEDC (2010)
Capital share	α	0.36	JEDC (2010)
Depreciation rate	δ	0.025	JEDC (2010)
Time worked per employee	ℓ	1/0.9	JEDC (2010)
Steady state			
Capital–output ratio	$\boldsymbol{K}/\boldsymbol{Y}$	10.26	JEDC (2010)
Effective interest rate	$\mathsf{r} - \delta$	0.01	JEDC (2010)
Wage rate	W	2.37	JEDC (2010)
Heterogenous agents models			
Unemployment insurance payment	μ	0.15	JEDC (2010)
Unemployment rate	u	0.07	Mean in JEDC (2010)
Probability of death	D	0.00625	Yields 40-year working life
Variance of log $\theta_{t,i}$	σ_θ^2	0.010×4	Carroll (1992)
Variance of log $\psi_{t,i}$	σ_ψ^2	$0.010 \times 4/11$	Carroll (1992)
			DeBacker et al. (2013)
KS aggregate shocks			
Shock to productivity	$\triangle^a$	0.01	Krusell and Smith (1998)
Unemployment (good state)	u^g	0.04	Krusell and Smith (1998)
Unemployment (bad state)	u^b	0.10	Krusell and Smith (1998)
Aggregate transition probability		0.125	Krusell and Smith (1998)

Notes: The models are calibrated at the quarterly frequency. The steady state values are calculated on a quarterly basis.