Preliminary and Incomplete – Please Do Not Quote

Digestible Microfoundations: Buffer Stock Saving in a Krusell–Smith World

April 30, 2012

Christopher D. Carroll[1] JHU

Jiri Slacalek[2] ECB

Kiichi Tokuoka[3] IMF

Abstract

Krusell and Smith (1998) showed that it is possible to construct rational expectations macroeconomic models with serious microfoundations. We argue that three modifications to their framework are required to fulfill its promise. First, we replace their assumption about household income dynamics with a process that matches microeconomic data. Second, our agents have finite lifetimes *a la* Blanchard (1985), which has both substantive and technical benefits. Finally, we calibrate heterogeneity in time preference rates so that the model matches the observed degree of inequality in the wealth distribution. Our model has substantially different, and considerably more plausible, implications for macroeconomic questions like the aggregate marginal propensity to consume out of an economic 'stimulus' program.

Keywords Microfoundations, Wealth Inequality, Marginal Propensity to Consume

JEL codes D12, D31, D91, E21

PDF: http://econ.jhu.edu/people/ccarroll/papers/BSinKS.pdf
Web: http://econ.jhu.edu/people/ccarroll/papers/BSinKS/
Archive: http://econ.jhu.edu/people/ccarroll/papers/BSinKS.zip

[1]Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD, http://econ.jhu.edu/people/ccarroll/, ccarroll@jhu.edu [2]Slacalek: European Central Bank, Frankfurt am Main, Germany, http://www.slacalek.com/, jiri.slacalek@ecb.europa.eu [3]Tokuoka: International Monetary Fund, Washington, DC, ktokuoka@imf.org.

The views presented in this paper are those of the authors, and should not be attributed to the International Monetary Fund, its Executive Board, or management, or to the European Central Bank.

1 Introduction

Macroeconomists have sought credible microfoundations since the dawn of our discipline. Keynes, his critics, and subsequent generations through Lucas (1976) and beyond have agreed on this, if little else.

Since Keynes's time, consumption modeling has been a battleground between two microfoundational camps. 'Bottom up' modelers (e.g. Modigliani and Brumberg (1954); Friedman (1957)) drew wisdom from microeconomic data and argued that macro models should be constructed by aggregation from microeconomic models that matched robust micro facts. 'Top down' modelers (e.g., Samuelson (1958); Diamond (1965); Hall (1978)) treated aggregate consumption as reflecting the optimizing decisions of representative agents; with only one such agent (or, at most, one per generation), these models had 'microfoundations' under a generous interpretation of the word.

The tractability of representative agent models has made them appealing for business cycle analysis. But such models have never been easy to reconcile with either macroeconomic[1] or microeconomic[2] evidence on consumption dynamics, nor with microeconomic theory which implies that people who differ from each other (in age, preferences, wealth, liquidity constraints, taxes, and other dimensions) should respond differently to any given shock. If any of these differences matter (and it is hard to see how they could not),[3] the aggregate size of a shock is not a sufficient statistic to calculate the aggregate response; information about how the shock is distributed is indispensable.

Bottom-up models, however, also have their problems. Even judged by a sympathetic standard that asks how well they can match measured wealth heterogeneity, bottom-up models have not been as successful as their champions might have initially hoped. For example, bottom-up models calibrated to match the wealth holdings of the median household generally fail to match the large size of the aggregate capital stock, because they seriously underpredict the upper parts of the wealth distribution (Carroll (2000b); Cagetti (2003)). Alternatively, models calibrated to match the aggregate level of wealth greatly overpredict wealth at the median (Hubbard, Skinner, and Zeldes (1994); Carroll (2000b)). A further problem is that (at least until Krusell and Smith (1998)) there has been no common answer to the question of how to analyze systematic macroeconomic fluctuations (business cycles) in bottom-up models.

This paper aims to reconcile the camps. We construct a workhorse model that answers the main objections to both kinds of models by making three modifications to the well-known Krusell–Smith ('KS') framework.[4] First, we replace KS's highly stylized assumptions about the nature of idiosyncratic income shocks with a microeconomic

[1] See, e.g., Campbell and Mankiw (1991) and the vast related literature following Hall (1978). A newer literature attempts to fix the problems identified in that literature by introducing habit formation (see, e.g., Fuhrer (2000)); but this is at the cost of intensifying the conflict with microeconomic evidence (see the next footnote).

[2] A large microeconomic literature, for example, has found average values of the marginal propensity to consume out of transitory income much greater than the 3–5 percent implied by representative agent models; see Table 13 in the Appendix E.

[3] Gorman (1953) shows that the essential requirement is that marginal propensities to consume be identical for all consumers. See Kirman (1992) and Solow (2003) for discussions of the deficiencies of representative agent models.

[4] In practice, we use the slightly modified version of that model presented in the recent JEDC volume referenced below.

labor income process that captures the essentials of the empirical consensus from the labor economics literature about actual income dynamics in micro data (with credibly calibrated transitory and permanent shocks).[5] Second, agents in our model have finite lifetimes *a la* Blanchard (1985), permitting a kind of primitive life cycle analysis and also solving some technical problems created by the incorporation of permanent shocks. Finally, we obtain a necessary extra boost to wealth inequality by calibrating a simple measure of heterogeneity in 'impatience.'[6]

The resulting framework differs sharply from the benchmark KS model in its implications for important microeconomic and macroeconomic questions. A timely macroeconomic example is the response of aggregate consumption to an 'economic stimulus payment,' interpreted here as a one-time lump sum transfer to households. In response to a $1-per-capita payment, the baseline version of the KS model implies that the annual marginal propensity to consume (MPC) is about 0.05,[7] almost irrespective of how the cash is distributed across households. In contrast, a version of our model that matches the distribution of liquid financial wealth implies that if the entire tax cut were directed at households in the bottom half of the liquid financial-wealth-to-income distribution, the MPC would be 0.83, which counts as a big improvement in realism, given the vast body of microeconomic evidence that consistently finds MPCs much greater than the 3–5 percent figure that characterizes representative agent models.[8] Furthermore, the model's differences with the representative agent framework are not peculiar to unusual events like a stimulus payment; to the extent that different kinds of macroeconomic shocks tend systematically to be differently distributed across the population (for example, labor income shocks may affect a less wealthy set of households than capital income shocks), this improvement in realism may also matter for general questions of macroeconomic dynamics.

Section 2 of the paper begins building the model's structure by adding microeconomic modeling elements to a benchmark representative agent model. Using this model (without macroeconomic dynamics), the section closes by estimating the degree of heterogeneity in impatience necessary to match the degree of inequality in the U.S. wealth distribution; we find that relatively small differences in impatience substantially affect the model's fit to the wealth data. Section 3 builds up the full version of the model by adding aggregate shocks of the KS type, and presents detailed comparisons of our model with theirs. Section 4 further improves the model by introducing an aggregate income process that is analytically simpler than the KS 'toy' aggregate process, that we believe is more empirically plausible as well, and that simplifies model solution and

[5]See, e.g., Hryshko (2010) for a recent contribution to, and overview of, the empirical literature.

[6]The word is in quotes because we refer here not to the pure time preference rate but instead to a relation between time preference rate, the interest rate, relative risk aversion, the intertemporal elasticity of substitution, the magnitude of risk, and expected income growth. All of these surely vary in the population, but we argue that few if any important macroeconomic questions depend on which particular kinds of heterogeneity are most responsible for the heterogeneity in wealth-to-income ratios. See subsection 2.4 for a fuller discussion.

[7]That is, if a dollar were given to every household in the economy, over the subsequent year household spending would be higher by about $0.05. For simplicity, we assume that the higher taxes needed to finance the tax cut would be imposed on unborn generations, though the result would not change much if there were an immediate increase in the tax rate to defray the stimulus program's cost.

[8]See the work summarized in Table 13 below.

simulation considerably. We offer this final, simpler version of the model as our preferred jumping-off point for future macroeconomic research.

2 The Model without Aggregate Uncertainty

2.1 The Perfect Foresight Representative Agent Model

To establish notation and a transparent benchmark, we begin by briefly sketching a standard perfect foresight representative agent model.

The aggregate production function is

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

where Z_t is aggregate productivity in period t, $\boldsymbol{K}_t$ is capital, ℓ is time worked per employee, and $\boldsymbol{L}_t$ is employment. The representative agent's goal is to maximize discounted utility from consumption

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

for a CRRA utility function $\mathrm{u}(\bullet) = \bullet^{1-\rho}/(1-\rho)$.[9] The representative agent's state at the time of the consumption decision is defined by two variables: $\boldsymbol{M}_t$ is market resources, and Z_t is aggregate productivity.

The transition process for $\boldsymbol{M}_t$ is broken up, for clarity of analysis and consistency with later notation, into three steps. Assets at the end of the period are market resources minus consumption, equal to

$$\boldsymbol{A}_t = \boldsymbol{M}_t - \boldsymbol{C}_t,$$

while next period's capital is determined from this period's assets via

$$\boldsymbol{K}_{t+1} = \boldsymbol{A}_t.$$

The final step can be conceived as 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:

$$\boldsymbol{M}_{t+1} \quad = \quad \daleth \boldsymbol{K}_{t+1} + \underbrace{Z_{t+1} \boldsymbol{K}_{t+1}^{\alpha} (\ell \boldsymbol{L}_{t+1})^{1-\alpha}}_{\boldsymbol{K}_{t+1}\mathsf{r}_{t+1} + (\ell \boldsymbol{L}_{t+1})\mathsf{W}_{t+1}}$$

where r_{t+1} is the interest rate,[10] W_{t+1} is the wage rate,[11] and $\daleth = 1-\delta$ is the depreciation factor for capital.

After normalizing by the productivity factor $\mathrm{Z}_t = Z_t^{1/(1-\alpha)}(\ell \boldsymbol{L}_t)$,[12] the representative agent's problem is

[9] Substitute $\mathrm{u}(\bullet) = \log \bullet$ for the case where $\rho = 1$.

[10] Equal to the marginal product of capital, $\alpha Z_{t+1} \boldsymbol{K}_{t+1}^{\alpha-1} (\ell \boldsymbol{L}_{t+1})^{1-\alpha}$.

[11] Equal to the marginal product of labor, $(1-\alpha) Z_{t+1} \boldsymbol{K}_{t+1}^{\alpha} (\ell \boldsymbol{L}_{t+1})^{-\alpha}$.

[12] Details of this normalization are discussed in Carroll (2000a).

Table 1 Parameter Values and Steady State

Description	Parameter	Value	Source
Representative agent model			
Time discount factor	β	0.99	JEDC (2010)
Coef 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			
Unempl insurance payment	μ	0.15	JEDC (2010)
Unemployment rate	u	0.07	Mean in JEDC (2010)
Probability of death	D	0.005	Yields 50 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	0.016/4	Carroll (1992); median in Table 3

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

$$\mathrm{V}(M_t, Z_t) = \max_{\mathrm{C}_t} \ \mathrm{u}(\mathrm{C}_t) + \beta\,\mathbb{E}_t\left[\Gamma_{t+1}^{1-\rho}\mathrm{V}(M_{t+1}, Z_{t+1})\right] \quad (2)$$

s.t.

$$A_t = M_t - C_t \quad (3)$$

$$K_{t+1} = A_t/\Gamma_{t+1} \quad (4)$$

$$M_{t+1} = \daleth K_{t+1} + K_{t+1}^{\alpha}, \quad (5)$$

where the non-bold variables are the corresponding bold variables divided by Z_t (e.g., $A_t = \boldsymbol{A}_t/\mathrm{Z}_t$, $M_t = \boldsymbol{M}_t/\mathrm{Z}_t$); $\Gamma_{t+1} = \mathrm{Z}_{t+1}/\mathrm{Z}_t$; and the expectations operator $\mathbb{E}_t$ here signifies the perfection of the agent's foresight (but will have the usual interpretation when uncertainty is introduced below).

Except where otherwise noted, our parametric assumptions match those of the papers in the special issue of the *Journal of Economic Dynamics and Control* (2010, Volume 34, Issue 1, edited by den Haan, Judd, and Julliard) devoted to comparing solution methods for the KS model (the parameters are reproduced for convenience in the top panel of Table 1).[13] The model is calibrated at the quarterly frequency. When aggregate shocks are shut down ($Z_t = 1$ and $\boldsymbol{L}_t = \boldsymbol{L}$), the model has a steady-state solution with a

[13] Examples of such authors include Young (2010) and Algan, Allais, and Den Haan (2008).

constant ratio of capital to output and constant (gross) interest and wage factors, which we write without time subscript as r and W and which are reflected in Table 1.[14]

Henceforth, we refer to the version of the model solved by the papers in the special JEDC volume as the 'KS-JEDC' model, while we call the original KS model solved in Krusell and Smith (1998) 'KS-Orig' model. (The only effective difference between the two is the introduction (for realism) of unemployment insurance in the KS-JEDC version, which does not matter much for any substantive results.[15,16])

2.2 The Household Income Process

For our purposes, the principal conclusion of the large literature on microeconomic labor income dynamics is that household income can be reasonably well described as follows. The idiosyncratic permanent component of labor income p evolves according to

$$p_{t+1} = \mathsf{G}_{t+1} p_t \psi_{t+1} \tag{6}$$

where G_{t+1} captures the predictable low-frequency (e.g., life-cycle and demographic) components of income growth, and the Greek letter psi mnemonically indicates the *p*ermanent *s*hock to *i*ncome. Actual income is the product of permanent income, a mean-one transitory shock, and the wage rate:

$$\boldsymbol{y}_{t+1} = p_{t+1} \xi_{t+1} \mathsf{W}_{t+1}.$$

After taking logarithms, this income process is strikingly similar to Friedman (1957)'s characterization of income as having permanent and transitory components. Because this process has been used widely in the literature on buffer stock saving, and though similar to Friedman's formulation is not identical to it, we henceforth refer to it as the Friedman/Buffer Stock (or 'FBS') process.[17]

Table 2 summarizes the annual variances of log permanent shocks (σ^2_ψ) and log transitory shocks (σ^2_ξ) estimated by a selection of papers from the extensive literature.[18] Some authors have used a process of this kind to describe the labor income or wage process for an individual worker (top panel) while others have used it to describe the process for overall household income (bottom panel); it seems to work reasonably well in both cases (though, obviously, with different estimates of the variances). (Recent work by Sabelhaus and Song (2010) using newly available data from Social Security

[14]In the steady state, $\boldsymbol{K}_t/(\ell \boldsymbol{L}_t) = \bar{k} = (\alpha\beta/(1-\beta\daleth))^{1/(1-\alpha)} = 38.0$, r (gross interest rate) $= \alpha\bar{k}^{\alpha-1}$, and $\mathsf{W} = (1-\alpha)\bar{k}^{\alpha}$.

[15]To be very precise, another difference is the introduction of ℓ (time worked per employee) in the KS-JEDC model, but this does not have a real impact.

[16]Details about the unemployment insurance scheme are described later in the paper.

[17]Guvenen (2007) refers to a process like this one as a 'restricted income process' (RIP) as distinguished from a process that he proposes which is similar but which allows each individual to have a distinct idiosyncratic mean growth rate. Guvenen's argument that each household has its own growth rate is intuitively plausible (indeed, it occurred to earlier authors who tested and rejected it), but Hryshko (2010) argues that there is no evidence that the Guvenen income process describes the data better (in a quantitatively meaningful way) than the restricted income process. Since incorporation of Guvenen's income process introduces serious modeling difficulties, it seems prudent to avoid using it unless the evidence for idiosyncratic growth factors becomes much more compelling.

[18]All the authors cited above used U.S. data. Nielsen and Vissing-Jorgensen (2006) used Danish data and estimated $\sigma^2_\psi = 0.005$ and $\sigma^2_\xi = 0.015$. It would be reasonable to interpret their estimates as the lower bounds for the U.S., given that their administrative data is well-measured and but that Danish welfare is more generous than the U.S. system.

earnings files finds that the variances of both transitory and permanent shocks have declined during the "Great Moderation" period at all ages; they also find distinct life cycle patterns of shocks by age, with young people experiencing higher levels of both kinds of shocks than the middle-aged).

The second-to-last line of the table shows what labor economists would have found, when estimating a process like the one above, if the empirical data were generated by households who experienced an income process like the one assumed by the KS-JEDC model.[19] This row of the table makes our point forcefully: The empirical procedures that have actually been applied to empirical micro data, if used to measure the income process households experience in a KS economy, would have produced estimates of σ_{ψ}^2 and σ_{ξ}^2 that are orders of magnitude different from what the actual empirical literature finds in actual data. This discrepancy naturally makes one wonder whether the KS-JEDC model's well-known difficulty in matching the degree of wealth inequality is largely explained by its highly unrealistic assumption about the income process.[20]

2.3 Finite Lifetimes and the Finite Variance of Permanent Income in the Cross-Section

One might wish to use the FBS income process specified in subsection 2.2 as a complete characterization of household income dynamics, but that idea has a problem: 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; that is, there is no ergodic distribution of permanent income in the population.

This problem and several others can be addressed by assuming that the model's agents 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 $t+1$. (The probability of NOT dying is the cancelation of the probability of dying: $\not{\mathsf{D}} = 1 - \mathsf{D}$). 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 an amount exactly corresponding to the mortality rate.

In order to maintain a constant population (of mass one, uniformly distributed on the unit interval), we assume that dying households are replaced by an equal number of newborns; we write the population-mean operator as $\mathbb{M}[\bullet_t] = \int_0^1 \bullet_{t,\iota} d\iota$. Newborns, we assume, begin life with a level of idiosyncratic permanent income equal to the mean level of idiosyncratic permanent income in the population as a whole. Conveniently, our

[19] First, we generated income draws according to the income process in the KS-JEDC model. Then, following the method in Carroll and Samwick (1997), we estimated the variances under the assumption that these income draws were produced by the process $\boldsymbol{y}_t = p_t \xi_t$ where $p_t = p_{t-1}\psi_t$. In doing so, as in Carroll and Samwick (1997), the draws of $\boldsymbol{y}_t$ are excluded when $\boldsymbol{y}_t$ is very low relative to its mean (see Carroll and Samwick (1997) for details about this restriction).

[20] The final line reports the variances estimated using income draws generated by the process assumed in Castaneda, Diaz-Gimenez, and Rios-Rull (2003), who were able to reproduce the skewness of the U.S. wealth distribution by reverse-engineering the income-process assumptions required to allow a Markov income process to generate the observed degree of wealth inequality. This process, too, bears little resemblance to the observable micro data on income dynamics.

Table 2 Estimates of Annual Variances of Log Income, Earning and Wage Shocks

Authors	Permanent σ_ψ^2	Transitory σ_ξ^2
Individual data		
MaCurdy (1982)[‡]	0.013	0.031
Topel (1990)	0.013	0.017
Topel and Ward (1992)	0.017	0.013
Moffitt and Gottschalk (1995)[⋆]	0.001	0.180
Meghir and Pistaferri (2004)[⋄]	0.031	0.032
Low, Meghir, and Pistaferri (2005)	0.011	–
Jensen and Shore (2008)[⋄]	0.054	0.171
Hryshko (2010)[⋄]	0.038	0.118
Guvenen (2009)	0.015	0.061
Household data		
Carroll (1992)	0.016	0.027
Carroll and Samwick (1997)	0.022	0.044
Storesletten, Telmer, and Yaron (2004a)	0.017	0.063
Storesletten, Telmer, and Yaron (2004b)	0.008–0.026	0.316
Blundell, Pistaferri, and Preston (2008)[⋄]	0.010–0.030	0.029–0.055
Implied by KS-JEDC	0.000	0.038
Implied by Castaneda, Diaz-Gimenez, and Rios-Rull (2003)	0.029	0.005

Notes: ‡ : MaCurdy (1982) did not explicitly separate ψ_t and ξ_t, but we have extracted σ_ψ^2 and σ_ξ^2 as implications of statistics that his paper reports. First, we calculate $\text{var}(\log \boldsymbol{y}_{t+d} - \log \boldsymbol{y}_t)$ and $\text{var}(\log \boldsymbol{y}_{t+d-1} - \log \boldsymbol{y}_t)$ using his estimate (we set $d = 5$). Then, following Carroll and Samwick (1997) we obtain the values of σ_ψ^2 and σ_ξ^2 which can match these statistics, assuming that the income process is $\boldsymbol{y}_t = p_t \xi_t$ and $p_t = p_{t-1}\psi_t$ (i.e., we solve $\text{var}(\log \boldsymbol{y}_{t+d} - \log \boldsymbol{y}_t) = d\sigma_\psi^2 + 2\sigma_\xi^2$ and $\text{var}(\log \boldsymbol{y}_{t+d-1} - \log \boldsymbol{y}_t) = (d-1)\sigma_\psi^2 + 2\sigma_\xi^2$). ⋆ : Moffitt and Gottschalk (1995) estimated the income process with random walk plus ARMA. Using income draws generated by their estimated process and following Carroll and Samwick (1997), we have estimated the variances under the assumption that these income draws were produced by the process $\boldsymbol{y}_t = p_t \xi_t$ where $p_t = p_{t-1}\psi_t$. ⋄ : Meghir and Pistaferri (2004), Jensen and Shore (2008), Hryshko (2010), and Blundell, Pistaferri, and Preston (2008) assume that the transitory component is serially correlated (an MA process), and report the variance of a subelement of the transitory component. For example, Meghir and Pistaferri (2004) and Blundell, Pistaferri, and Preston (2008) assume an MA(1) process $\log \xi_t = v_t + \vartheta v_{t-1}$ and obtain estimates (σ_v^2, ϑ)=$(0.0300, -0.2566)$ and $(0.0286\text{-}0.0544, 0.1132)$, respectively. σ_ξ^2 for these four articles reported in this table are calculated by $(1 + \vartheta^2)\sigma_v^2$ using their estimates.

definition of the permanent shock implies that in a large population, mean idiosyncratic permanent income will remain fixed at $\mathbb{M}[p] = 1$ forever, while the mean of p^2 is given by[21]

$$\mathbb{M}[p^2] \;=\; \frac{\mathsf{D}}{1 - \not{D}\mathbb{E}[\psi^2]} \tag{7}$$

and the variance of p by

$$\sigma_p^2 \;=\; \mathbb{M}[p^2] - 1.$$

Of course for all of this to be valid, it is necessary to impose the parametric restriction $\not{D}\mathbb{E}[\psi^2] < 1$ (a requirement that does not do violence to the data, as we shall see). Intuitively, the requirement is that, among surviving consumers, income does not spread out so quickly as to overwhelm the compression of the permanent income distribution that arises because of the equalizing force of death and replacement.

Since our goal here is to produce a realistic distribution of permanent income across the members of the (simulated) population, we measure the empirical distribution of permanent income in the cross section using data from the Survey of Consumer Finances (SCF), which conveniently includes a question asking respondents whether their income in the survey year was about 'normal' for them, and if not, asks the level of 'normal' income.[22] This corresponds well with our (and Friedman (1957)'s) definition of permanent income p (and Kennickell (1995) shows that the answers people give to this question can be reasonably interpreted as reflecting their perceptions of their permanent income), so we calculate the variance of $p^i \equiv p^i/\mathbb{M}[p^i]$ among such households.[23]

The results from this exercise are reported in Table 3 (with a final row that makes the point that both the KS-Orig and KS-JEDC models assume that permanent shocks did not exist). Substituting these estimates for σ_p^2 into (7) and (8), we obtain estimates of the variance of ψ. Reassuringly, we can interpret the variances of ψ thus obtained as being easily in the range of the estimated variances of $\log(\psi) = \sigma_\psi^2$ in Table 2.[24] Such a correspondence, across two quite different methods of measurement, suggests there is considerable robustness to the measurement of the size of permanent shocks. (Below, we will choose a calibration for σ_ψ^2 that is in the middle range of estimates from either method.)

2.4 The Wealth Distribution with Transitory and Permanent Shocks

We now examine how wealth would be distributed in the steady-state equilibrium of an economy with wage rates and interest rates fixed at the steady state values calibrated in Table 1 of subsection 2.1, an income process like the one described in subsection 2.2, and finite lifetimes per subsection 2.3.

[21] See Appendix A for the derivation.

[22] SCF1992 only asked whether the income level was about 'normal' or not.

[23] We restrict the sample to households between the ages of 25 and 60, because the interpretation of the question becomes problematic for retired households.

[24] So long as the variance of the permanent shocks is small, these two measures should be approximately the same.

Table 3 Variance of Permanent Income

Dataset	var(p)	$\mathbb{E}[\psi^2]$	σ^2_ψ
SCF1992	2.5	1.015	0.015
SCF1995	7.5	1.018	0.018
SCF1998	3.1	1.015	0.015
SCF2001	3.6	1.016	0.016
SCF2004	5.2	1.017	0.017
KS-Orig or KS-JEDC	0	1	0

The process of noncapital income of each household follows

$$\boldsymbol{y}_t = p_t \xi_t \mathsf{W}_t \tag{8}$$

$$p_t = p_{t-1}\psi_t \tag{9}$$

$$\mathsf{W}_t = (1-\alpha)Z_t(\boldsymbol{K}_t/\ell \boldsymbol{L}_t)^\alpha, \tag{10}$$

where $\boldsymbol{y}_t$ is noncapital income for the household in period t, equal to the permanent component of noncapital income p_t multiplied by a transitory income shock factor ξ_t and wage rate W_t; the permanent component of noncapital income in period t is equal to its previous value, multiplied by a mean-one iid shock ψ_t, $\mathbb{E}_t[\psi_{t+n}] = 1$ for all $n \geq 1$. $\boldsymbol{K}_t$ is capital and $\boldsymbol{L}_t = 1 - u_t$ is the employment rate (because u_t is the unemployment rate). Since there is no aggregate shock, Z_t, $\boldsymbol{K}_t$, $\boldsymbol{L}_t$, and W_t are constant ($Z_t = Z = 1$, $\boldsymbol{K}_t = \boldsymbol{K}$, $\boldsymbol{L}_t = \boldsymbol{L}$, and $\mathsf{W}_t = \mathsf{W} = (1-\alpha)(\boldsymbol{K}/\ell\boldsymbol{L})^\alpha$).

Following the assumptions in the JEDC volume, the distribution of ξ_t is:

$$\xi_t = \mu \text{ with probability } u_t \tag{11}$$

$$= (1-\tau_t)\ell\theta_t \text{ with probability } 1-u_t, \tag{12}$$

where $\mu > 0$ is the unemployment insurance payment when unemployed and $\tau_t = \mu u_t/\ell\boldsymbol{L}_t$ is the rate of tax collected to pay unemployment benefits (see Table 1 for parameter values).[25] The probability of unemployment is constant ($u_t = u$); later we allow u to vary over time.

The decision problem for the household in period t can be written using normalized variables; the consumer's objective is to choose a series of consumption functions c between now and the end of the horizon that satisfy:

[25]The KS-Orig model assumed no unemployment insurance ($\mu = 0$).

$$\begin{aligned} \mathrm{v}(m_t) &= \max_{\mathrm{c}_t} \ \mathrm{u}(\mathrm{c}_t) + \beta \not{D} \mathbb{E}_t \left[\psi_{t+1}^{1-\rho} \mathrm{v}(m_{t+1}) \right] && (13) \\ & \text{s.t.} \\ a_t &= m_t - c_t \\ a_t &\geq 0 \\ k_{t+1} &= a_t / (\not{D} \psi_{t+1}) && (14) \\ m_{t+1} &= (\daleth + r) k_{t+1} + \xi_{t+1} && (15) \end{aligned}$$

where the non-bold ratio variables are defined as the bold (level) variables divided by the level of permanent income $\boldsymbol{p}_t = p_t \mathsf{W}$ (e.g., $m_t = \boldsymbol{m}_t/(p_t\mathsf{W})$). The only state variable is (normalized) cash-on-hand m_t. The household's employment status is *not* a state variable, unlike in the KS-JEDC model, where tomorrow's employment status depends on today's status. This substantially simplifies the analysis (which is useful for computational and analytical purposes), arguably without too much sacrifice of realism (except possibly for detailed studies of the behavior of households during extended unemployment spells).

Since households die with a constant probability D between periods, the effective discount factor is $\beta\not{D}$ (in (13)); the effective interest rate is $(\daleth + r)/\not{D}$ (combining (14) and (15)).[26]

Aside from heterogeneity in impatience (introduced below), three parameters characterize our modifications to the KS-JEDC model: D, σ_θ^2, and σ_ψ^2. $\mathsf{D} = 0.005$ implies the average length of working life is $1/0.005 = 200$ quarters $= 50$ years (dating from entry into the labor force at, say, age 25). The variance of log transitory income shocks σ_θ^2 is the value advocated in Carroll (1992) (based on the *Panel Study of Income Dynamics* (PSID) data),[27] as is $\sigma_\psi^2 = 0.016$ (but note that this value also matches the median value in Table 3).[28,29] Other parameter values (ρ, α, δ, and ℓ) are from the JEDC volume (Table 1).

The one remaining unspecified parameter is the time preference factor. As a preliminary theoretical consideration, note that Carroll (2011) (generalizing Deaton (1991) and Bewley (1977)) has shown that models of this kind do not have a well-defined solution unless the condition holds:

$$\frac{(\mathsf{R}\beta)^{1/\rho}}{\acute{\Gamma}} < 1 \qquad (16)$$

[26] $(\daleth + r)$ is scaled by $1/\not{D}$ due to the Blanchardian mutual insurance scheme as described in the previous subsection.

[27] This paper assumes that each period corresponds to a *quarter*, while $\sigma_\theta^2 = 0.010$ from Carroll (1992) is the value on an annual basis. Therefore, following Carroll, Slacalek, and Tokuoka (2008), 0.010 needs to be multiplied by 4 since the variance of log transitory income shocks of *quarterly* data should be four times as large as that of annual data. (Note further that Carroll (1992)'s calibration of $\sigma_\theta^2 = 0.010$ was considerably lower than his raw empirical estimate of 0.027, on the grounds that a substantial portion of the changes in measured income is likely to come from measurement error).

[28] Since σ_ψ^2 in Table 3 (0.016) is estimated using annual data, it needs to be *divided* by 4, following Carroll, Slacalek, and Tokuoka (2008) (recall that our model is calibrated quarterly).

[29] Using *quarterly* income draws generated by this section's income process with these parameter values, we have estimated the *annual* ARMA process for $\log(\xi_t)$ assumed in Moffitt and Gottschalk (1995): $\log(\xi_t) = a_1 \log(\xi_{t-1}) + v_t + m_1 v_{t-1}$. The estimates of a_1 and m_1 are positive and negative, respectively, in line with the coefficients estimated by Moffitt and Gottschalk (1995). This suggests that Moffitt and Gottschalk's findings are qualitatively consistent with the other papers in this literature, and with our own calibration of the income process. See Appendix B for details.

where

$$\acute{\Gamma} \;=\; \left(\mathbb{E}[\psi^{-\rho}]\right)^{-1/\rho}\Gamma.$$

Carroll (2011) dubs this inequality the 'Growth Impatience Condition' because it guarantees that consumers are sufficiently impatient to prevent the indefinite increase in the *ratio* of net worth to permanent income when income is growing (see also Szeidl (2006)). This condition is an amalgam of the pure time preference factor, expected growth, the relative risk aversion coefficient, and the real interest factor. Thus, a consumer can be 'impatient' in the required sense even if $\beta = 1$, so long as expected income growth is positive.[30]

We begin by searching for the time preference factor $\grave{\beta}$ such that if all households had an identical $\beta = \grave{\beta}$ the steady-state value of the capital-to-output ratio ($\boldsymbol{K}/\boldsymbol{Y}$) would match the value that characterized the steady-state of the perfect foresight model.[31] $\grave{\beta}$ turns out to be 0.9888 (recall that this is at a quarterly, not an annual, rate).

We now ask whether this model with realistically calibrated income and finite lifetimes (henceforth, the model is referred to as the 'β-Point' model) can reproduce the degree of wealth inequality evident in the micro data. An improvement in the model's ability to match the data is to be expected, since 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 vital explanation for variation in the level of target wealth (and, therefore, on average, actual wealth) across households.

Table 4 shows that compared to the distribution of net worth implied by our solution of the KS-JEDC model solved without an aggregate shock (or the results of the KS-Orig model from Krusell and Smith (1998)),[32] our β-Point model does indeed yield a substantial improvement (compare the first, third and fourth columns to the last column).[33] For example, in our β-Point model, the fraction of total net worth held by the top 1 percent is about 10 percent, while the corresponding statistic is only 3 percent in our solution of the KS-JEDC model (or the KS-Orig model).

The KS-JEDC model's failure to match the wealth distribution is not confined to the top. In fact, perhaps a bigger problem is that the model generates a distribution of wealth in which most households' wealth levels are 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 all households in the KS-JEDC model have net worth between 0.5 times mean net worth and 1.5 times mean net worth; in the SCF data from 1992–2004, the corresponding fraction ranges from only 20 to 25 percent.

---

[30] This near-equivalence explains why we do not bother to include a growth term in the process for noncapital income in (8)–(10) despite the presence of such a term in (6); inclusion of the income growth term should mostly just result in an offsetting effect on our estimated time preference rate, and would complicate our simulations unnecessarily.

[31] Output is the sum of noncapital and capital income.

[32] Our solution of the KS-JEDC model is similar to the results of the KS-Orig model in terms of wealth distribution; what small differences do exist reflect the minor difference in the assumption about unemployment insurance (discussed earlier) as well as the fact that the KS-Orig model was solved with aggregate shocks turned on.

[33] Throughout this paper, we will examine the distribution of net worth (not financial or gross assets), unless otherwise noted.

Table 4 Proportion of Net Worth by Percentile (in percent)

	Micro Income Process					
	Friedman/Buffer Stock		KS-JEDC	KS-Orig$^{\diamond}$		
	Point Discount Factor‡ (β-Point)	Uniformly Distributed Discount Factors* (β-Dist)	Our Solution		Hetero	U.S. Data*
Top 1%	10.3	24.9	3.0	3.0	24.0	29.6
Top 10%	38.6	65.6	23.0	19.0	73.0	66.1
Top 20%	54.9	81.0	40.1	35.0	88.0	79.5
Top 40%	75.7	93.1	66.0			92.9
Top 60%	88.9	97.4	84.0			98.7
Top 80%	97.0	99.3	95.2			100.4

Notes: $\boldsymbol{K}_t/\boldsymbol{Y}_t = 10.3$. ‡ : $\grave{\beta} = 0.9888$. * : $(\grave{\beta}, \nabla) = (0.9869, 0.0052)$, which implies $\{\grave{\beta} - \nabla, \grave{\beta} + \nabla\} = \{0.9816, 0.9921\}$. $^{\diamond}$: The results are from Krusell and Smith (1998) who solved the models with aggregate shocks. * : U.S. data is the SCF reported in Castaneda, Diaz-Gimenez, and Rios-Rull (2003). Bold quantiles are targeted.

But while the β-Point model fits the data better than the original KS model, it still falls short of matching the empirical degree of wealth inequality. The proportion of net worth held by households in the top 1 percent of the distribution is three times smaller in the model than in the data (compare the first and last columns in the table). This failure reflects the fact that, empirically, the distribution of wealth is considerably more unequal than the distribution of permanent income.

2.5 Heterogeneous Impatience

As the simplest method to address this defect, we introduce heterogeneity in impatience: Each household is now assumed to have an idiosyncratic (but fixed) time preference factor.[34] We do not think of this assumption as only capturing actual variation in pure rates of time preference across people (though such variation surely exists). Instead, we view discount-factor heterogeneity as a shortcut that captures the essential consequences of many other kinds of heterogeneity (e.g., heterogeneity in age, income growth expectations, investment opportunities, tax schedules) as well. To be more concrete, take the example of age. A robust pattern in most countries is that income grows much faster for young people than for older people. According to (16), young people should therefore tend to act, financially, in a more 'impatient' fashion than older people. In particular, we should expect them to have lower target wealth-to-income ratios. Thus, what we are capturing by allowing heterogeneity in time preference factors is probably

[34] This differs from KS's experiment with heterogeneity, in which a household's discount factor could change suddenly; they interpreted such a change as reflecting a dynastic transition.

also some portion of the difference in behavior that (in truth) reflects differences in age instead of in time preference factors, and that would be introduced into the model if we had a more complex specification of the life cycle that allowed for different growth rates for households of different ages.[35]

One way of gauging a model's predictions for wealth inequality is to ask how well it is able to match the proportion of total net worth held by the wealthiest 20, 40, 60, and 80 percent of the population. Because these statistics have been targeted by other papers (e.g., Castaneda, Diaz-Gimenez, and Rios-Rull (2003)), we adopt a goal of matching them.[36]

We replace the assumption that all households have the same time preference factor with an assumption that, for some ∇, time preference factors are distributed uniformly in the population between $\grave{\beta} - \nabla$ and $\grave{\beta} + \nabla$ (for this reason, the model is referred to as the 'β-Dist' model below). Then, using simulations, we search for the values of $\grave{\beta}$ and ∇ for which the model best matches the fraction of net worth held by the top 20, 40, 60, and 80 percent of the population, while at the same time matching the aggregate capital-to-output ratio from the perfect foresight model. Specifically, defining w_i and ω_i as the proportion of total aggregate net worth held by the top i percent in our model and in the data, respectively, we solve the following minimization problem:

$$\min_{\grave{\beta},\nabla} \sum_{i=20,40,60,80} (w_i - \omega_i)^2 \tag{17}$$

subject to the constraint that the aggregate wealth (net worth)-to-output ratio in the model matches the aggregate capital-to-output ratio from the perfect foresight model ($\boldsymbol{K}_{PF}/\boldsymbol{Y}_{PF}$):[37]

$$\boldsymbol{K}/\boldsymbol{Y} \;=\; \boldsymbol{K}_{PF}/\boldsymbol{Y}_{PF} \tag{18}$$

The solution to this problem is $(\grave{\beta}, \nabla) = (0.9869, 0.0052)$.

The introduction of even such a relatively modest amount of time preference heterogeneity sharply improves the model's fit to the targeted proportions of wealth holdings (second column of the table). The ability of the model to match the targeted moments does not, of course, constitute a formal test, except in the loose sense that a model with such strong structure might have been unable to get nearly so close to four target wealth points with only one free parameter.[38] But the model also sharply improves the fit to locations in the wealth distribution that were not explicitly targeted; for example, the

[35]We could of course model age effects directly, but it is precisely the inclusion of such realism that has made OLG models unpopular for business cycle modeling; they are too unwieldy to use for many practical research purposes and (perhaps more important) it is too difficult distill their mechanics into readily communicable insights. Our view is that, for business cycle analysis purposes, the only thing of substance that is gained in exchange for many different kinds of extra complexity is a widening of the distribution of wealth-to-income ratios. We achieve such a widening transparently and parsimoniously by incorporating discount factor heterogeneity.

[36]Castaneda, Diaz-Gimenez, and Rios-Rull (2003) targeted various wealth and income distribution statistics, including net worth held by the top 1, 5, 10, 20, 40, 60, 80 percent, and the Gini coefficient.

[37]In estimating these parameter values, we approximate the uniform distribution by seven points ($\grave{\beta} - 3\nabla/3.5$, $\grave{\beta} - 2\nabla/3.5$, $\grave{\beta} - \nabla/3.5$, $\grave{\beta}$, $\grave{\beta} + \nabla/3.5$, $\grave{\beta} + 2\nabla/3.5$, $\grave{\beta} + 3\nabla/3.5$). Increasing the number of points further does not notably change the results below.

[38]As is clear from the minimization problem above, we are estimating two parameters ($\grave{\beta}$ and ∇). However, that

net worth shares of the top 10 percent and the top 1 percent are also included in the table, and the model performs reasonably well in matching them.

Of course, Krusell and Smith (1998) were well aware that their baseline model match the wealth distribution well. They, too, examined whether inclusion of a form of discount rate heterogeneity could improve the model's match to the data. Specifically, they assumed that the discount factor takes one of the three values (0.9858, 0.9894, and 0.9930), and that agents anticipate that their discount factor might change between these values according to a Markov process. As they showed, the model with this simple form of heterogeneity (henceforth 'KS-Orig Hetero' model) did improve the model's ability to match the wealth holdings of the top percentiles (see KS-Orig Hetero column in the table). Indeed, as inspection of the long-dashing locus in Figure 1 shows, their model of heterogeneity went a bit too far: It concentrated almost all of the net worth in the top 20 percent of the population (though rather evenly among that top 20 percent). By comparison, the figure shows that our model does a notably better job matching the data across the entire span of wealth percentiles.

The reader might wonder why we do not simply adopt the KS specification of heterogeneity in time preference factors, rather than introducing our own novel form of heterogeneity. The principal answer is that our purpose here is to define a method of explicitly matching the model to the data via statistical estimation of a parameter of the distribution of heterogeneity, letting the data speak flexibly to the question of the extent of the heterogeneity required to match model to data. A second point is that, having introduced finite horizons in order to yield an ergodic distribution of permanent income, it would be peculiar to layer on top of the stochastic death probability a stochastic probability of changing one's time preference factor within the lifetime; Krusell and Smith motivated their differing time preference factors as reflecting different preferences of alternative generations of a dynasty, but with our finite horizons assumption we have eliminated the dynastic interpretation of the model. Having said all of this, the common point across the two papers is that a key requirement to make the model fit the data is a form of heterogeneity that leads different households to have different target levels of wealth.

3 KS Aggregate Shocks

In this section, we examine a model with an FBS household income process that also incorporates KS aggregate shocks, and investigate the model's performance in replicating aggregate statistics. Krusell and Smith (1998) assumed that the level of aggregate productivity alternates between $Z_t = 1 + \triangle^Z$ if the aggregate state is good and $Z_t = 1 - \triangle^Z$ if it is bad; similarly, $\boldsymbol{L}_t = 1 - u_t$ where $u_t = u^g$ if the state is good and $u_t = u^b$ if bad. (For reference, we reproduce their assumed parameter values in Table 5.)

The decision problem for an individual household in period t can be written using normalized variables and the employment status ι_t:

estimation is subject to a constraint (matching the targeted aggregate net worth-to-output ratio) that effectively pins down one of the parameters ($\grave{\beta}$), so effectively only ∇ works to match the four wealth target points.

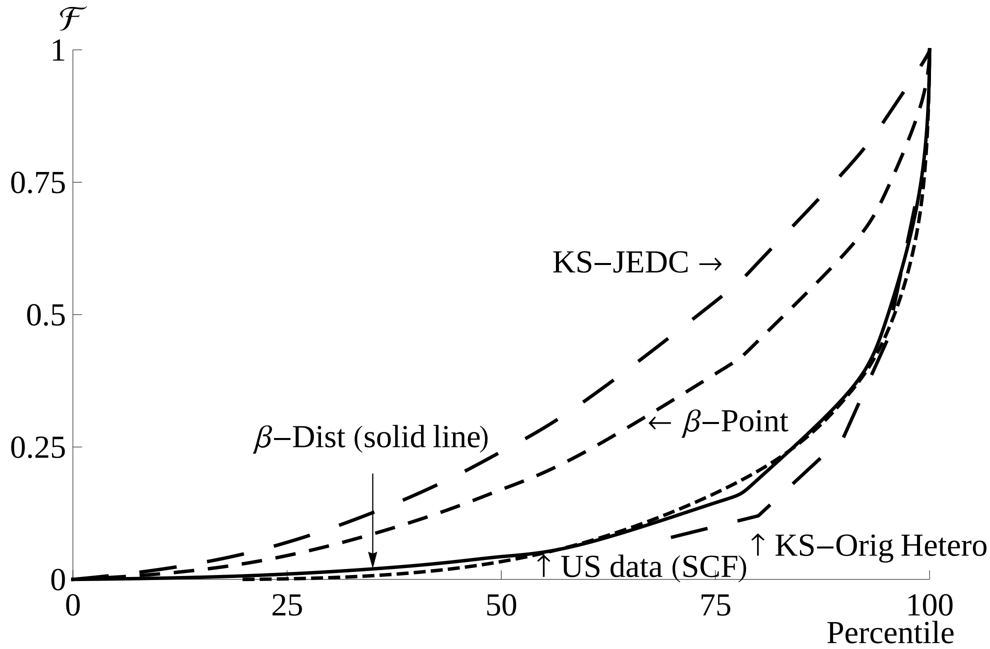

Figure 1 Cumulative Distribution of Net Worth

Table 5 Parameter Values for Aggregate Shocks

Description	Parameter	Value	Source
KS shocks			
Shock to productivity	$\triangle^Z$	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)
FBS shocks			
Variance of log Ψ_t	σ^2_Ψ	0.00004	Carroll, Slacalek, and Tokuoka (2008)
Variance of log Ξ_t	σ^2_Ξ	0.00001	Carroll, Slacalek, and Tokuoka (2008)

$$
\begin{aligned}
\mathrm{v}(m_t, \iota_t; \boldsymbol{K}_t, Z_t) &= \max_{c_t} \ \mathrm{u}(c_t) + \beta \not{\mathsf{D}} \mathbb{E}_t \left[(\Gamma_{t+1}\psi_{t+1})^{1-\rho} \mathrm{v}(m_{t+1}, \iota_t; \boldsymbol{K}_{t+1}, Z_{t+1}) \right] \\
&\text{s.t.} \\
a_t &= m_t - c_t \\
a_t &\geq 0 \\
k_{t+1} &= a_t / (\not{\mathsf{D}} \Gamma_{t+1}\psi_{t+1}) \\
m_{t+1} &= (\daleth + r_{t+1}) k_{t+1} + y_{t+1} \\
r_{t+1} &= \alpha Z_{t+1} (\boldsymbol{K}_{t+1} / \ell \boldsymbol{L}_{t+1})^{\alpha - 1}, \qquad (19)
\end{aligned}
$$

where

- the non-bold *individual* variables (lower-case variables except for ι_t and ψ_t) are the bold (level) variables divided by $\mathrm{Z}_t \boldsymbol{p}_t$ (e.g., $a_t = \boldsymbol{a}_t / \mathrm{Z}_t \boldsymbol{p}_t$, $m_t = \boldsymbol{m}_t / \mathrm{Z}_t \boldsymbol{p}_t$),
- $\Gamma_{t+1} = \mathrm{Z}_{t+1} / \mathrm{Z}_t$,
- $\boldsymbol{L}_t = 1 - u_t$, and
- the income process is the same as in (8)–(12) but the employment transition process follows KS-JEDC.

There are more state variables in this version of the model than in the model with no aggregate shock: The aggregate variables Z_t and $\boldsymbol{K}_t$, and the household's employment status ι_t whose transition process depends on the aggregate state. Solving the full version of the model above with both aggregate and idiosyncratic shocks is not straightforward; the basic idea for the solution method is the key insight of Krusell and Smith (1998). See Appendix C for details about our solution method.

We now report the results of simulations, both for the model in which all households have the same time preference factor (β-Point model) and for the version with a uniform distribution of time preference factors (β-Dist model). While the β-Point model uses $\grave{\beta}$ estimated in Section 2, the β-Dist model uses parameter values reestimated by solving the minimization problem (17) with the KS aggregate shocks ($(\grave{\beta}, \nabla) = (0.9851, 0.0074)$). Results using our solution of the KS-JEDC model (with the KS aggregate shocks, $\theta_t = 1$, $\psi_t = 1$ for all t, and no death ($\mathsf{D} = 0$)) are also reported for comparison.

3.1 Some Macroeconomic Statistics

Table 6 shows some aggregate statistics that we think are useful for macroeconomic analysis: The serial correlation of consumption growth, and correlation between consumption growth, income growth, and interest rates at several frequencies. The results are generally similar across the β-Point, β-Dist, and KS-JEDC models. They all produce positive $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$, and high correlation of consumption growth with (current) income growth or (current) interest rates. The serial correlation of consumption growth in our solution of the KS-JEDC model is similar to that reported by Maliar, Maliar, and Valli (2008) who also solved the KS-JEDC model (fourth column

Table 6 Aggregate Statistics with KS Aggregate Shocks

	Micro Income Process				
	Friedman/Buffer Stock		KS-JEDC		None
	β-Point	β-Dist	Our Solution	Maliar et al. (2008)	Rep Agent Model
$corr(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$	0.14	0.05	0.23	0.28	0.24
$corr(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_t)$	0.91	0.97	0.86		0.84
$corr(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_{t-1})$	0.12	0.06	0.15		0.15
$corr(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_{t-2})$	0.09	0.04	0.13		0.11
$corr(\Delta \log \boldsymbol{C}_t, \mathsf{r}_t)$	0.78	0.66	0.85		0.86
$corr(\Delta \log \boldsymbol{C}_t, \mathsf{r}_{t-1})$	0.13	−0.05	0.26		0.28
$corr(\Delta_4 \log \boldsymbol{C}_t, \Delta_4 \log \boldsymbol{Y}_t)$	0.80	0.92	0.70		0.67
$corr(\Delta_8 \log \boldsymbol{C}_t, \Delta_8 \log \boldsymbol{Y}_t)$	0.74	0.90	0.63		0.61

Notes: Δ_4 and Δ_8 are one-year and two-year growth rates, respectively.

of the table).[39,40] We also report results for the representative agent model with the KS aggregate income shock parameters (last column), the results of which are very close to those of our solution of the KS-JEDC model.

The classic reference point for consumption growth measurement is the random walk model of Hall (1978), and the large literature that rejects the random walk proposition in favor of models that either contain some 'rule-of-thumb' consumers who set spending equal to income in every period (Campbell and Mankiw (1989)) or, more popular recently, models with habit formation or 'sticky expectations' (Carroll, Slacalek, and Tokuoka (2008)) that imply serial correlation in consumption growth (see Carroll, Sommer, and Slacalek (2011) for evidence).

The KS-JEDC model produces a relatively high correlation coefficient $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$, which is closer to the U.S. data (where the statistic is about one-third) than that produced by standard consumption models stemming from Hall (1978).[41] As noted already, our β-Point and β-Dist models also imply positive $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$, although not as high as that predicted by the KS-JEDC model. At first blush, it seems puzzling that the KS-JEDC model, which includes neither habits nor sticky expectations, generates a substantial violation of the random walk proposition. This puzzle does not seem to have been noticed in the previous literature on the KS-JEDC model, but after some investigation we determined that the KS-JEDC model's sticky consumption growth is produced by the high degree of serial correlation in interest

[39] The difference between the results in Maliar, Maliar, and Valli (2008) and ours reflects approximation error in solving the consumption function.

[40] Although not reported here, our solution of the KS-JEDC model closely matches theirs in other aggregate statistics as well (e.g., variance of aggregate consumption (level), correlation between income and consumption levels).

[41] However, it should be noted that the serial correlation coefficient for consumption growth calculated using the U.S. data may be significantly underestimated because of measurement error and some other factors (Carroll, Sommer, and Slacalek (2011)). This would imply that the models above do not reproduce stickiness in aggregate consumption growth well.

rates in the model, which results from the assumption about the process of aggregate productivity shocks (see Appendix D for details). The interesting questions, in a model with time-varying interest rates, are, first, whether one can reliably estimate an intertemporal elasticity of substitution (IES) from the coefficient in a regression of consumption on the predictable component of interest rates (as Hall (1988) attempts to do), and, second, whether consumption growth is serially correlated *after accounting for* the predictable component related to interest rates (no random walk).[42]

3.2 The Aggregate Marginal Propensity to Consume

A macroeconomic question of perhaps even greater interest is whether a model that manages to match the distribution of wealth has similar, or different, implications from the KS-JEDC or representative agent models for the reaction of aggregate consumption to an economic 'stimulus' payment.

Specifically, we pose the question as follows. The economy has been in its steady-state equilibrium leading up to date t. Before the consumption decision is made in that period, the government announces the following plan: Effective immediately, every household in the economy will receive a 'stimulus check' worth some modest amount $\$x$ (financed by a tax on unborn future generations).[43]

In theory, the distribution of wealth across recipients of the stimulus checks has important implications for aggregate MPC out of transitory shocks to income. To see why, the solid line of Figure 2 plots our β-Point model's individual consumption function in the good (aggregate) state, with the horizontal axis being cash on hand normalized by the level of (quarterly) permanent income. Because the households with less normalized cash have higher MPC,[44] the average MPC is higher when a larger fraction of households has less (normalized) cash on hand.

There are many more households with little wealth in our β-Point model than in the KS-JEDC model, as illustrated by comparison of the short-dashing and the long-dashing lines in Figure 1. The greater concentration of wealth at the bottom in the β-Point model, which is the case in the data (see the histogram in Figure 2), should produce a higher average MPC, given the concave consumption function.

Indeed, the average MPC out of the transitory income ('stimulus check') in our β-Point model is 0.09 in annual terms (first column of Table 7),[45] about double the value in the KS-JEDC model (0.05) (the fourth column of the table) or the perfect foresight partial equilibrium model (0.04). Our β-Dist model (second column of the table) produces an even higher average MPC (0.19), since in the β-Dist model there are more households who possess less wealth, are more impatient, and have higher MPCs (Figure 1 and

[42] See Appendix D for the analysis of the KS-JEDC model.

[43] This financing scheme, along with the lack of a bequest motive, eliminates any Ricardian offset that might otherwise occur.

[44] Consumption functions of the KS-JEDC and β-Dist models have a similar form.

[45] The MPCs that we calculate are annual MPCs given by $1-(1-\text{quarterly MPC})^4$ (recall again that the models in this paper are calibrated quarterly). We make this choice because earlier influential studies (e.g., Souleles (1999); Johnson, Parker, and Souleles (2006)) attempted to estimate long-term MPCs, which refers to the amount of extra spending that has occurred over the course of a year or 9 months in response to a one unit increase in resources. Henceforth, the casual usage of the term 'the MPC' refers to annual MPC.

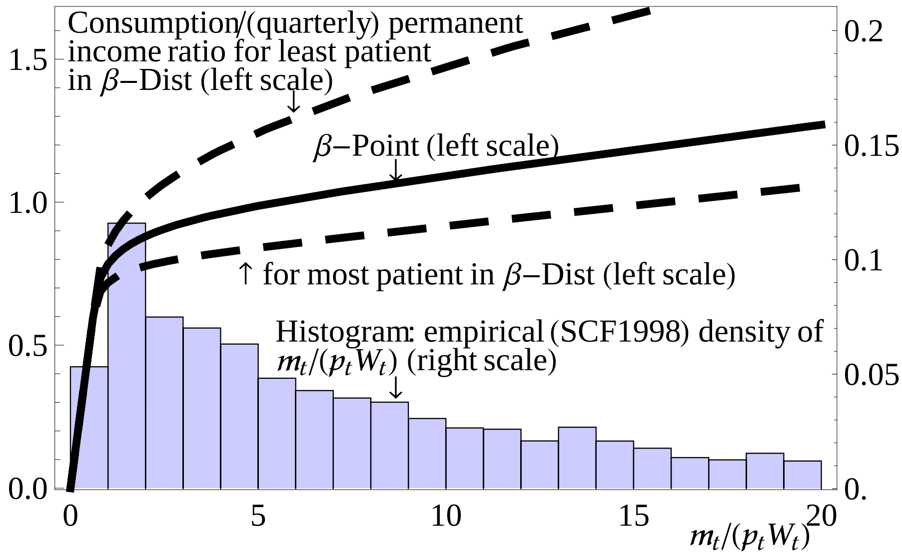

Figure 2 Consumption Functions of β-Point and β-Dist Models

dashed lines in Figure 2).[46] However, this is still at best only at the lower bound of typical empirical MPC estimates which are typically between 0.2–0.5 or even higher (see Table 13 in the Appendix E).

Thus far, we have been using total household net worth as our measure of wealth. Implicitly, this assumes that all of the household's debt and asset positions are perfectly liquid and that, say, a household with home equity of $50,000 and bank balances of $2,000 (and no other balance sheet items) will behave in every respect similarly to a household with home equity of $10,000 and bank balances of $42,000. This seems implausible. The home equity is more illiquid (tapping it requires, at the very least, obtaining a home equity line of credit, which requires an appraisal of the house and some paperwork).

Otsuka (2003) formally analyzes the optimization problem of a consumer with a FBS income process who can invest in an illiquid but higher-return asset (think housing), or a liquid but lower-return asset (cash), and shows, unsurprisingly, that the marginal propensity to consume out of shocks to liquid assets is higher than the MPC out of shocks to illiquid assets. Her results would presumably be even stronger if she had allowed that households hold so much of their wealth in illiquid forms (housing, pension savings), for example, as a mechanism to overcome self-control problems (see Laibson (1997) and many others).[47]

These considerations suggest that it may be more plausible, for purposes of extracting a predictions about the MPC out of stimulus checks, to focus on matching the distribu-

[46]The results are similar with aggregate shocks turned off.

[47]Indeed, using a model with both a low-return liquid asset and a high-return illiquid asset, Kaplan and Violante (2011) have replicated high MPCs observed in the data.

Table 7 Average Marginal Propensity to Consume in Annual Terms

	KS Aggregate Process					Friedman/Buffer Stock Aggregate Process	
Model	β-Point	β-Dist	β-Dist	KS-JEDC Our Solution	KS-Orig Hetero Our Solution	β-Dist	β-Dist
Wealth Measure	Net Worth	Net Worth	Liquid Financial Assets			Net Worth	Liquid Financial Assets
Overall average	0.09	0.19	0.68	0.05	0.09	0.18	0.69
By wealth/permanent income ratio							
Top 1%	0.06	0.05	0.23	0.04	0.04	0.06	0.24
Top 10%	0.06	0.06	0.24	0.04	0.04	0.06	0.24
Top 20%	0.06	0.06	0.28	0.04	0.04	0.06	0.27
Top 40%	0.06	0.07	0.39	0.04	0.05	0.06	0.41
Top 60%	0.07	0.09	0.50	0.04	0.06	0.08	0.52
Bottom 1/2	0.12	0.28	0.83	0.05	0.13	0.28	0.84
By income							
Top 1%	0.08	0.13	0.36	0.05	0.04	0.15	0.64
Top 10%	0.08	0.14	0.48	0.05	0.04	0.15	0.64
Top 20%	0.09	0.14	0.52	0.05	0.04	0.16	0.64
Top 40%	0.10	0.16	0.57	0.05	0.05	0.17	0.65
Top 60%	0.10	0.16	0.61	0.05	0.06	0.18	0.66
Bottom 1/2	0.08	0.21	0.76	0.05	0.13	0.18	0.73
By employment status							
Employed	0.08	0.16	0.65	0.05	0.09	0.16	0.66
Unemployed	0.20	0.44	0.95	0.06	0.18	0.35	0.96
Time preference parameters‡							
$\grave{\beta}$	0.9888	0.9851	0.9037			0.9869	0.9111
∇		0.0074	0.0424			0.0052	0.0336

Notes: Annual MPC is calculated by $1-(1-\text{quarterly MPC})^4$. ‡: Discount factors are uniformly distributed over the interval $(\grave{\beta}-\nabla, \grave{\beta}+\nabla)$.

Table 8 Proportion of Wealth Held by Percentile (in percent)

	Liquid Financial Assets					Net Worth		
	1992	1995	1998	2001	2004	1992‡	1998	2004
Top 1%	42.2	52.7	47.6	49.6	50.6	29.6	34.4	33.9
Top 10%	79.4	84.8	83.2	85.2	86.1	66.1	68.9	69.7
Top 20%	90.2	92.8	92.5	93.4	93.8	79.5	82.1	82.9
Top 40%	97.4	98.1	98.1	98.3	98.6	92.9	94.3	94.7
Top 60%	99.4	99.6	99.6	99.6	99.7	98.7	99.1	99.0
Top 80%	100.0	100.0	100.0	100.0	100.0	100.4	100.4	100.2

Notes: Survey of Consumer Finances, ‡: From Castaneda, Diaz-Gimenez, and Rios-Rull (2003).

tion of liquid financial assets across households (that is, assets which are of the same kind as represented by the stimulus check, once it has been deposited into a bank account).

When we ask the model to estimate the time preference factors that allow it to best match the distribution of liquid financial assets (instead of net worth),[48] estimated parameter values are $(\grave{\beta}, \nabla) = (0.9037, 0.0424)$ and the average MPC is 0.68 (third column of the table), which lies in the upper part of the range typically reported in the literature (see Table 13), and is considerably higher than when we match the distribution of net worth. This reflects the fact that matching the more skewed distribution of liquid financial assets than that of net worth (Table 8) requires a wider distribution of the time preference factors, which produces even more households with little wealth. The estimated distribution of discount factors lies below that obtained by matching net worth and is considerably more dispersed because of substantially lower median and more unevenly distributed liquid financial wealth (compared to net worth).

Figure 3 shows the cumulative distribution functions of MPCs for the β-Dist models estimated to match the empirical distribution of net worth and liquid financial assets. The Figure illustrates the high values of implied MPCs obtained for both models, especially the latter.

MPCs are generally higher among low wealth/income households and the unemployed in both our β-Point and β-Dist models (rest of the rows in Table 7). These results provide the basis for a common piece of conventional wisdom about the effects of economic stimulus mentioned in our introduction: If the purpose of the stimulus payments is to stimulate consumption, it makes much more sense to target those payments to low-wealth households than to distribute them uniformly to the population as a whole.

[48] We define liquid financial assets as the sum of transaction accounts (deposits), CDs, bonds, stocks, and mutual funds. We take the same approach as before: we match the fraction of liquid financial assets held by the top 20, 40, 60, and 80 percent of the population (in the SCF1998), while at the same time matching the aggregate liquid financial assets-to-income ratio (which is 3.9 in the SCF1998).

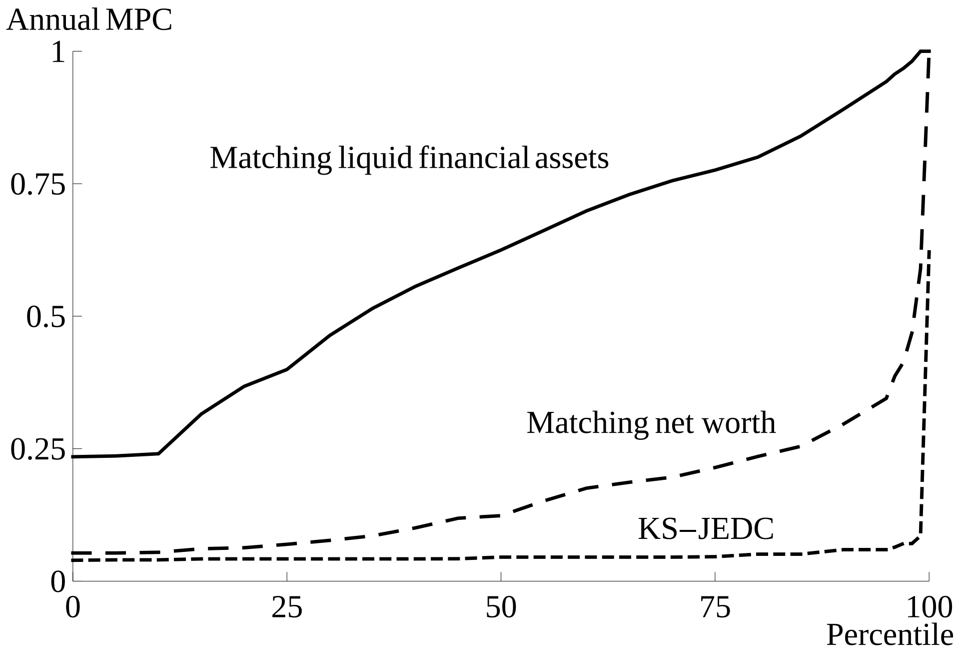

Figure 3 Distribution of MPCs across Households

4 A More Plausible and More Tractable 'Friedman/Buffer Stock' Aggregate Process

The KS process for aggregate productivity shocks has little empirical foundation; indeed, it appears to have been intended by the authors as an illustration of how one might incorporate business cycles in principle, rather than a serious candidate for an empirical description of actual aggregate dynamics. In this section, we introduce an aggregate income process that is considerably more tractable than the KS aggregate process and is also a much closer match to the aggregate data. We regard the version of our model with this new income process as the 'preferred' version for use as a starting point for future research.

The aggregate production function is the same as equation (1), but following Carroll, Slacalek, and Tokuoka (2008), the aggregate state (good or bad) no longer exists in this model ($Z_t = 1$). Aggregate productivity is instead captured by $\boldsymbol{L}_t$. Specifically, $\boldsymbol{L}_t = P_t\Xi_t$; P_t is aggregate permanent productivity, where $P_{t+1} = P_t\Psi_{t+1}$; Ψ_{t+1} is the aggregate permanent shock; and Ξ_t is the aggregate transitory shock (note that Ψ is the capitalized version of the Greek letter ψ used for the idiosyncratic permanent shock; similarly (though less obviously), Ξ is the capitalized ξ). Both Ψ_t and Ξ_t are assumed to be log normally distributed with mean one, and their log variances are from Carroll, Slacalek, and Tokuoka (2008), who have estimated them using U.S. data (Table 5).

The assumption that the structure of aggregate shocks resembles the structure of idiosyncratic shocks is valuable not only because it matches the data better, but also because it makes the model easier to solve. In particular, the elimination of the 'good'

Table 9 Aggregate Statistics in β-Dist Model under 'Plausible' Aggregate Process

	Friedman/Buffer Stock Aggregate Process		KS Aggregate Process
	Net Worth	Liquid Financial Assets	Net Worth
$\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$	0.10	0.24	0.05
$\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_t)$	0.82	0.83	0.97
$\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_{t-1})$	0.08	0.18	0.06
$\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{Y}_{t-2})$	0.05	0.13	0.04
$\varrho(\Delta \log \boldsymbol{C}_t, \mathsf{r}_t)$	0.46	0.66	0.66
$\varrho(\Delta \log \boldsymbol{C}_t, \mathsf{r}_{t-1})$	0.20	0.33	-0.05
$\varrho(\Delta_4 \log \boldsymbol{C}_t, \Delta_4 \log \boldsymbol{Y}_t)$	0.92	0.92	0.92
$\varrho(\Delta_8 \log \boldsymbol{C}_t, \Delta_8 \log \boldsymbol{Y}_t)$	0.95	0.95	0.90

and 'bad' aggregate states reduces the number of state variables to two (m_t and $\boldsymbol{K}_t$) after normalizing the model by $\boldsymbol{p}_t P_t$ (as elaborated in Carroll, Slacalek, and Tokuoka (2008)). As in Section 2, employment status is not a state variable (in eliminating the aggregate states, we also shut down unemployment persistence, which depends on the aggregate state in the KS-JEDC or KS-Orig model). As before, the main thing the household needs to know is the law of motion of $\boldsymbol{K}_t$, which can be obtained by following essentially the same method as described in the Appendix C.

When matching the distribution of net worth, aggregate statistics produced by the β-Dist model with our preferred (Friedman/Buffer Stock) aggregate process are relatively similar to those under the KS aggregate process, despite the difference in the aggregate process (first and third columns of Table 9). Given that there is no aggregate state in the economy, we are using $\grave{\beta}$ and ∇ estimated in Section 2 and assuming that the unemployment rate u_t is fixed at 0.07 (same as in Section 2). Our preferred version of the β-Dist model maintains positive $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$, and high correlation of consumption growth with income growth or interest rates. We have obtained similar results by matching the distribution of liquid financial assets (second column of the table).

More importantly, the preferred version of the β-Dist model can produce high MPCs. For example, in the net worth case, the average MPC is 0.18, which is very close to the estimate under the KS aggregate process (compare second and last but one columns of Table 7). In the liquid financial assets case, the average MPC is higher at 0.69 (last column).[49]

[49] As in the net worth case, parameter value are estimated with aggregate shocks turned off $((\grave{\beta}, \nabla) = (0.9111, 0.0336))$.

5 Conclusion

This paper found that the performance of a KS-type model in replicating wealth distribution can be improved significantly by introducing i) a microfounded income process, ii) finite lifetimes, and iii) heterogeneity in time preference factors. Moreover, such modifications improve macroeconomic characteristics of the model by substantially boosting the MPC out of transitory income.

Appendix

A Derivation of Variance of Permanent Income

The evolution of the square of p is given by

$$\begin{aligned} p_{t+1,i} &= p_{t,i}\psi_{t+1,i}(1-\mathsf{d}_{t+1,i})+\mathsf{d}_{t+1,i} \\ p_{t+1,i}^2 &= (p_{t,i}\psi_{t+1,i}(1-\mathsf{d}_{t+1,i}))^2+2p_{t,i}\psi_{t+1,i}\underbrace{\mathsf{d}_{t+1,i}(1-\mathsf{d}_{t+1,i})}_{=0}+\mathsf{d}_{t+1,i}^2, \end{aligned}$$

where $\mathsf{d}_{t+1,i}=1$ if household i dies.

Because $\mathbb{E}_t[(1-\mathsf{d}_{t+1,i})^2]=1-\mathsf{D}$ and $\mathbb{E}_t[\mathsf{d}_{t+1,i}^2]=\mathsf{D}$, we have

$$\begin{aligned} \mathbb{E}_t[p_{t+1,i}^2] &= \mathbb{E}_t[(p_{t,i}\psi_{t+1,i}(1-\mathsf{d}_{t+1,i}))^2]+\mathsf{D} \\ &= p_{t,i}^2\not\mathsf{D}\,\mathbb{E}[\psi^2]+\mathsf{D}, \end{aligned}$$

and

$$\mathbb{M}\left[p_{t+1}^2\right] = \mathbb{M}[p_t^2]\not\mathsf{D}\,\mathbb{E}[\psi^2]+\mathsf{D}.$$

Finally, the steady state expected level of $\mathbb{M}[p^2]\equiv\lim_{t\to\infty}\mathbb{M}[p_t^2]$ can be found from

$$\begin{aligned} \mathbb{M}[p^2] &= \mathsf{D}+\not\mathsf{D}\,\mathbb{E}[\psi^2]\mathbb{M}[p^2] \\ \mathbb{M}[p^2] &= \frac{\mathsf{D}}{1-\not\mathsf{D}\,\mathbb{E}[\psi^2]}. \end{aligned}$$

B Estimating Moffitt and Gottschalk Income Process

This appendix estimates the *annual* income process à la Moffitt and Gottschalk (1995) using *quarterly* income draws generated by our income process (Section 2) with parameter values from Table 1. Moffitt and Gottschalk (1995) assume log permanent income $\log(p_t)$ follows a random walk and log transitory income $\log(\xi_t)$) an ARMA process:

$$\begin{aligned} \boldsymbol{y}_t &= p_t\xi_t, \\ \log(p_t) &= \log(p_{t-1})+\log(\psi_t), \\ \log(\xi_t) &= a_1\log(\xi_{t-1})+v_t+m_1v_{t-1} \end{aligned}$$

Like Moffitt and Gottschalk (1995), we match the covariance matrix of the annual income draws, and obtain estimate with the same signs as theirs obtained using the PSID data; see Table 10, confirming that our calibration is qualitatively consistent with Moffitt and Gottschalk's.

Interestingly, even though our true quarterly transitory shock process is just white noise, if we estimate the process on an annual basis we obtain positive AR (a_1) and negative MA (m_1) coefficients, reflecting time aggregation. This suggests that the positive a_1 and negative m_1 reported in Moffitt and Gottschalk (1995) may be (at least) partly due to time aggregation.

Table 10 Estimates of Moffitt and Gottschalk Annual Income Process

	σ^2_ψ	σ^2_v	a_1	m_1
Our estimates	0.015	0.025	0.504	−0.521
Moffitt and Gottschalk (1995)	0.00159	0.169	0.622	−0.344

C Solution Method to Obtain Law of Motion

C.1 Solution Methods

Broadly speaking, the literature takes one of the following two approaches in solving the KS problem in Section 3:

1. Relying on simulation to obtain the law of motion of per capita capital
2. (In principle) not relying on simulation to obtain the law of motion of per capita capital

Table 11 lists some existing articles that solve the KS-JEDC model according to this categorization. All articles in the table except Kim, Kim, and Kollmann (2010) solve the exact KS-JEDC model using various methods.[50]

The advantage of the first approach is that simulation performed to obtain the law of motion generates micro data, which can be used directly to investigate issues such as wealth distribution. The disadvantage is that this approach is generally subject to cross-sectional sampling variation, because this approach typically performs simulation using a finite number of households. Young (2010) and Den Haan (2010b)'s approaches can also be categorized in the first approach but avoid cross-sectional sampling variation by running *nonstochastic* simulation that approximates the density of wealth with a histogram.

The advantages of the second approach are: i) there is no cross-sectional sampling variation; ii) it is generally faster than the first approach. Using the second approach, Algan, Allais, and Den Haan (2008) and Reiter (2010) find a wealth distribution function of various moments,[51] while Reiter (2010) calculates a matrix for the transition probabilities of individual wealth (see Appendix **??** for details about his technique). Kim, Kim, and Kollmann (2010) use a perturbation method that linearizes the system. Although they are not able to solve the exact same KS-JEDC model and thus modify the form of the utility function, they can solve a related problem very quickly.

We use the first approach because it directly generates various micro data (e.g., individual wealth and MPC), which can be used to examine wealth distribution and the aggregate MPC. Details about our algorithm are in the next subsection.

[50] Kim, Kim, and Kollmann (2010) modified the form of the utility function.

[51] Simulation plays a part in Algan, Allais, and Den Haan (2008)'s method (they use simulation to find the function).

Table 11 Methods of Solving KS-JEDC Model

Authors	Description
Relying on simulation	
Young (2010)	Grid-based method
Den Haan (2010b)	Grid-based method
(In principle) not relying on simulation	
Algan, Allais, and Den Haan (2008)	Parametrization method
Reiter (2010)	Parametrization method
Kim, Kim, and Kollmann (2010)	Perturbation method
Den Haan and Rendahl (2010)	Explicit aggregation method

C.2 Our Algorithm

In solving the problem in section 3 we closely follow the stochastic simulation method of Krusell and Smith (1998). Krusell and Smith find that per capita capital today ($\boldsymbol{K}_t$) is sufficient to predict per capita capital tomorrow ($\boldsymbol{K}_{t+1}$). Our procedure is as follows:

1. Solve for the optimal individual decision rules given some 'beliefs' π that determine the (expected) law of motion of per capita capital. The law of motion is takes the log-linear form given by $\pi = (\pi_0, \pi_1, \pi'_0, \pi'_1)$:
$$\log \boldsymbol{K}_{t+1} = \pi_0 + \pi_1 \log \boldsymbol{K}_t$$
if the aggregate state in period t is good ($Z_t = 1 + \triangle^Z$), and
$$\log \boldsymbol{K}_{t+1} = \pi'_0 + \pi'_1 \log \boldsymbol{K}_t$$
if the aggregate state is bad ($Z_t = 1 - \triangle^Z$).

2. Simulate the economy populated by $7,000$ households (which experiments determined is enough to suppress idiosyncratic noise) for $1,100$ periods (following Maliar, Maliar, and Valli (2010)). When starting a simulation, $p_{t,i} = 1$ for all i, the distribution of $m_{t,i}$ is generated assuming $k_{t,i}$ is equal to its steady state level (38.0) for all i, and $Z_t = 1 + \triangle^Z$ (the aggregate state is good). (The steady state level of $k_{t,i}$ is $\bar{k} = (\alpha\beta/(1-\beta\daleth))^{1/(1-\alpha)}$. With $k_{t,i} = 38.0$ for all i, $\boldsymbol{k}_{t,i} = \boldsymbol{K}_t = 41.2$.) The newborn households start life with $p_{t,i} = 1$ and $k_{t,i} = 0$.

3. Estimate $\tilde{\pi}$, which determines the law of motion of per capita capital, using the last $1,000$ periods of data generated by the simulation (we discard the first 100 periods).

4. Compute an improved vector for the next iteration by $\hat{\pi} = (1-\eta)\tilde{\pi} + \eta\pi$. $\eta = 1/2$ is used for the β-Dist model. (Our experiments found that we can reach the solution faster with $\eta = 1/2$.)

We repeat this process until $\hat{\pi} = \pi$ with a given degree of precision.[52]

From the second iteration and thereafter, we use the terminal distribution of wealth (and permanent component of income (p)) in the previous iteration as the initial one. For the case of the β-Dist model, the number of households is multiplied by 10 in the final two (or three) iterations to reduce cross-sectional simulation error.[53]

While we can eventually obtain some solution whatever the initial π is, we use π obtained using the representative agent model as the starting point. This can significantly reduce the time needed to obtain the solution.

Parameter values to solve the model are from Table 1 (except for the unemployment rate u_t) and Table 5. The time preference factors are imposed to be those estimated in Section 2.

C.3 Tricks to Reduce Simulation Errors

In obtaining the aggregate law, we introduce the following tricks to reduce simulation errors (or to speed up the solution given a degree of estimate precision):

- **Death:** When death is concentrated among households at the very top of the wealth distribution, per capita capital would be at a lower than normal level. To alleviate simulation errors from this source, each period we: i) sort households by wealth level, ii) construct groups, the size of which is the inverse of the death probability (under our parameter choice, the size of each group is 200 and the first group contains households from the wealthiest to the 200th), and iii) pick one household that dies within each group.
- **Permanent income shocks:** In our methodology, permanent shocks to income are approximated by n discrete points. Similarly to the death element, after sorting we set up groups each of size n. We randomize shocks within each group subject to the constraint that each shock point is experienced by one of the group members every period, making the group mean of the shocks equal to the theoretical mean.[54]

C.4 Estimated Laws of Motion

The estimated laws of motions for β-Point, β-Dist and KS-JEDC models are given in Table 12. The fit measured with R^2 in all specifications exceeds 0.9999.[55]

[52]In our analysis below, the process is iterated until the difference between each estimate (π_0, π_1, π_0', or π_1') and its previous value is smaller than 1 percent.

[53]This is enough to ensure that the maximum deviation of each estimate of π_0, π_1, π_0' and π_1' from its previous value is less than 1 percent.

[54]This idea is motivated by Braun, Li, and Stachurski (2009), who proposed the estimation of densities with smaller simulation errors by calculating conditional densities given simulated data.

[55]Note that, as pointed out by Den Haan (2010a), R^2 only measures in-sample fit and should be interpreted with caution.

Table 12 Estimated Laws of Motion

$\log \boldsymbol{K}_{t+1} = \pi_0 + \pi_1 \log \boldsymbol{K}_t + \epsilon_{t+1}$

Model	β-Point		β-Dist		KS-JEDC	
State	Good	Bad	Good	Bad	Good	Bad
π_0	0.140	0.127	0.154	0.144	0.138	0.122
π_1	0.963	0.965	0.959	0.960	0.963	0.966

Notes: The coefficients for the KS-JEDC model are very close to those estimated in Maliar, Maliar, and Valli (2010).

D Experiment to Understand Sticky Consumption Growth in KS-JEDC Model

Although $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$ reported in subsection 3.1 may not be high enough relative to that observed in the U.S. data, it is still not clear why simulations produce such a high value.

Previous studies on KS type models have not investigated this issue. Using the KS-JEDC model, we performed an experiment to understand the phenomenon better. In this experiment we assume that the aggregate state switches from good to bad (or from bad to good) every eight quarters.[56]

Figure 4 plots $\Delta \log \boldsymbol{C}_t$ 24 quarters of simulated observations (the state is bad for the first eight quarters, good for the next eight quarters, and bad for the final eight quarters). The figure shows that $\Delta \log \boldsymbol{C}_t$ is very persistent (it is negative in the bad state and positive in the good state), resulting in a relatively high $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$.

It is easy to understand that $\Delta \log \boldsymbol{C}_t$ is higher when the state is good (and vice versa) given the following facts:

- A first order approximation of the Euler equation yields:

$$\Delta \log \boldsymbol{C}_t \approx b_0 + b_1 r_t, \tag{20}$$

  where $b_0 \equiv -\rho^{-1}(1-\beta+\delta)$, $b_1 \equiv \rho^{-1}$, ρ is the coefficient of relative risk aversion, r_t is the interest rate, β is the time preference factor, and δ is the depreciation rate. Indeed, when we conduct an IV regression of equation (20) using r_{t-1} as the instrument,[57] which effectively means estimating $\Delta \log \boldsymbol{C}_t = b_0 + b_1 \mathbb{E}_{t-1}[r_t] + \varepsilon_t$, the estimate of $b_1 \equiv \rho^{-1}$ is 0.95 (with a standard deviation of 0.08) and relatively close to the actual value of ρ^{-1} $(=1)$. This suggests that using the predictable component of interest rates ($\mathbb{E}_{t-1}[r_t]$), we can obtain a reasonable estimate of intertemporal elasticity of substitution.

- When the state is good, $r_t = \alpha Z_t(\boldsymbol{K}_t/\ell \boldsymbol{L}_t)^{\alpha-1}$ (from (19)) is higher because Z_t

[56] Because one state switches to another with a probability of 0.125, the average length of each state is eight quarters in typical simulation.

[57] The data that produced Table 6 are used.

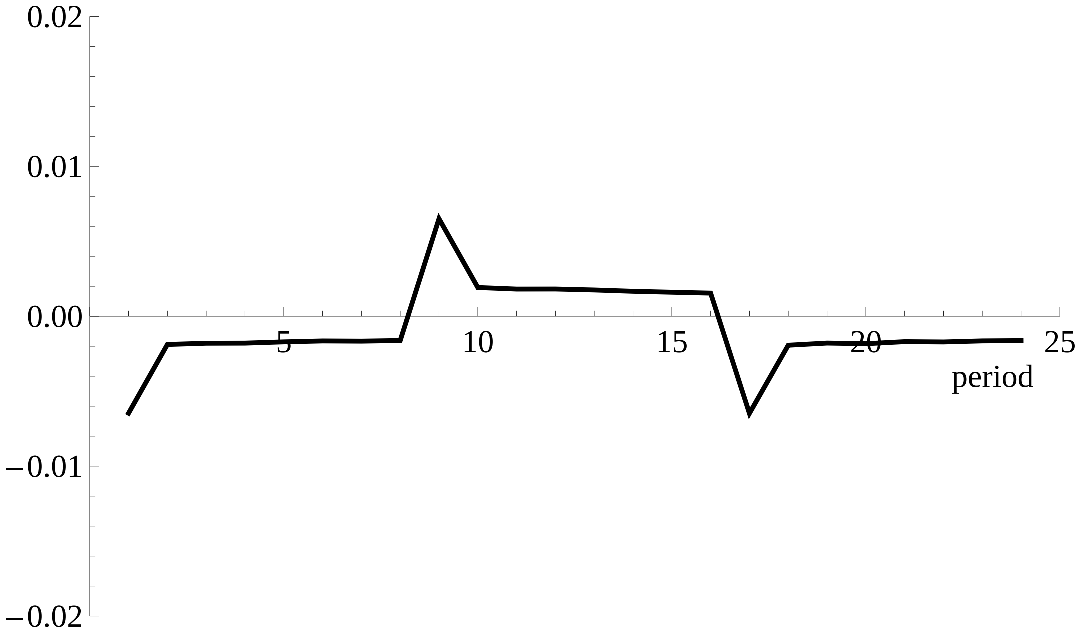

Figure 4 Dynamics of $\Delta \log \boldsymbol{C}_t$ in KS-JEDC Model

(aggregate productivity) is higher, as can be seen in Figure 5, which plots the dynamics of r_t for the 24 quarters.

While in typical simulation one state does not generally last for exactly eight quarters, we observe sticky aggregate consumption growth (and a relatively high $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$) because the same mechanisms are at work as in the experiment above.

In sum, a relatively high $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$ in the KS-JEDC model can be interpreted as a consequence of the persistent behavior of the interest rate r_t. Indeed, denoting $\varepsilon_t = \Delta \log \boldsymbol{C}_t - b_0 - b_1 \mathbb{E}_{t-1}[r_t]$ the residual after controlling for the predictable component of consumption growth related to interest rates, we find that $\varrho(\varepsilon_t, \varepsilon_{t-1}) = 0.02$ is much lower than $\varrho(\Delta \log \boldsymbol{C}_t, \Delta \log \boldsymbol{C}_{t-1})$.[58]

E Empirical Estimates of MPCs

Table 13 summarizes the estimates of MPCs obtained using household-level data on various recent fiscal stimulus measures in the U.S.

[58]Estimating an AR(1) process on ε_t produces a small and statistically insignificant coefficient on lagged ε_{t-1}.

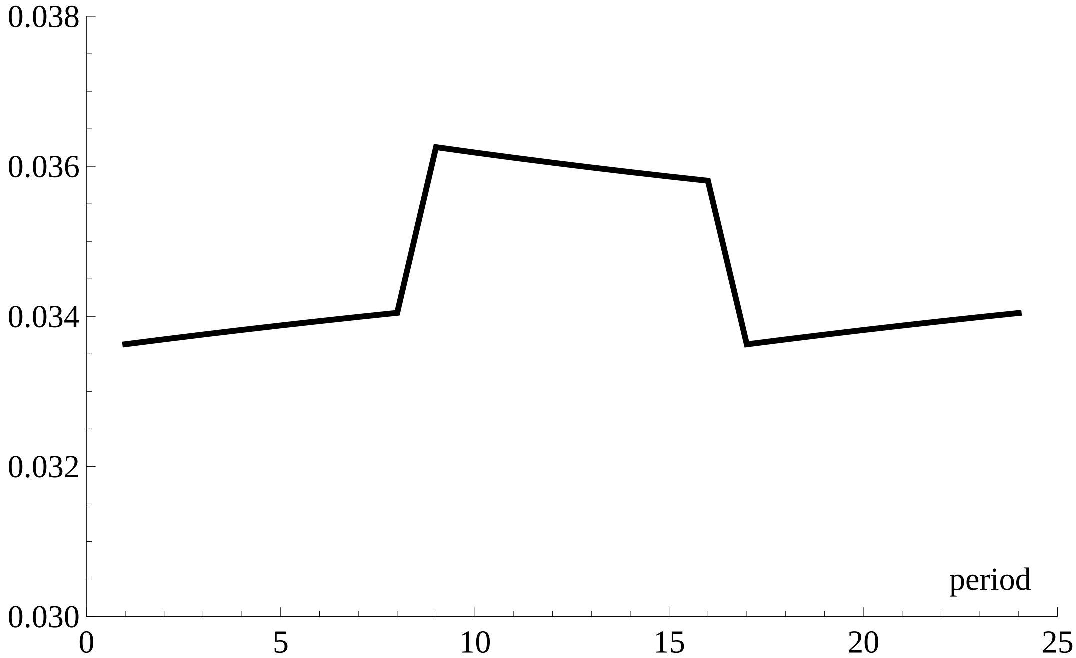

Figure 5 Dynamics of r_t in KS-JEDC Model

Table 13 Estimates of MPCs

	Consumption Measure		
Authors	Nondurables	Durables	Total PCE
Agarwal, Liu, and Souleles (2007)			0.4
Coronado, Lupton, and Sheiner (2005)			0.28–0.36
Johnson, Parker, and Souleles (2006)	0.12–0.30		0.50–0.90
Johnson, Parker, and Souleles (2009)	0.25		
Lusardi (1996)‡	0.2–0.5		
Parker (1999)	0.2		
Parker, Souleles, Johnson, and McClelland (2011)	0.12–0.30		
Sahm, Shapiro, and Slemrod (2009)			0.33
Shapiro and Slemrod (2009)			0.33
Souleles (1999)	0.09	0.54	0.64
Souleles (2002)	0.6–0.9		

Notes: ‡: elasticity.

References

Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles (2007): "The Response of Consumer Spending and Debt to Tax Rebates – Evidence from Consumer Credit Data," *Journal of Political Economy*, 115(6), 986–1019.

Algan, Yann, Oliver Allais, and Wouter J. Den Haan (2008): "Solving Heterogeneous-agent Models with Parameterized Cross-sectional Distributions," *Journal of Economic Dynamics and Control*, 32(3), 875–908.

Bewley, Truman (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252–292.

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

Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008): "Consumption Inequality and Partial Insurance," *Manuscript*.

Braun, Richard Anton, Huiyu Li, and John Stachurski (2009): "Computing Densities and Expectations in Stochastic Recursive Economies: Generalized Look-Ahead Techniques," *Manuscript*.

Cagetti, Marco (2003): "Wealth Accumulation Over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 21(3), 339–353.

Campbell, John Y., and N. Gregory Mankiw (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence," in *NBER Macroeconomics Annual, 1989*, ed. by Olivier J. Blanchard, and Stanley Fischer, pp. 185–216. MIT Press, Cambridge, MA, http://ideas.repec.org/p/nbr/nberwo/2924.html.

——— (1991): "The Response of Consumption to Income: A Cross-Country Investigation," *European Economic Review*, 35, 723–67.

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.

——— (2000a): "Requiem for the Representative Consumer? Aggregate Implications of Microeconomic Consumption Behavior," *American Economic Review*, 90(2), 110–115.

——— (2000b): "Why Do the Rich Save So Much?," in *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, ed. by Joel B. Slemrod. Harvard University Press, http://econ.jhu.edu/people/ccarroll/Why.pdf.

——— (2011): "Theoretical Foundations of Buffer Stock Saving," *Manuscript, Department of Economics, Johns Hopkins University*, http://econ.jhu.edu/people/ccarroll/papers/BufferStockTheory.

Carroll, Christopher D., and Andrew A. Samwick (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40(1), 41–71.

Carroll, Christopher D., Jiri Slacalek, and Kiichi Tokuoka (2008): "Sticky Expectations and Consumption Dynamics," *Manuscript*.

Carroll, Christopher D., Martin Sommer, and Jiri Slacalek (2011): "International Evidence on Sticky Consumption Growth," *Review of Economics and Statistics*, 93(4), 1135–1145, http://econ.jhu.edu/people/ccarroll/papers/cssIntlStickyC/.

Castaneda, Ana, Javier Diaz-Gimenez, and Jose-Victor Rios-Rull (2003): "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111(4), 818–857.

Coronado, Julia Lynn, Joseph P. Lupton, and Louise M. Sheiner (2005): "The Household Spending Response to the 2003 Tax Cut: Evidence from Survey Data," FEDS discussion paper 32, Federal Reserve Board.

Deaton, Angus S. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.

Den Haan, Wouter J. (2010a): "Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents," *Jouranl of Economic Dynamics and Control*, 34(1), 79–99.

——— (2010b): "Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty," *Journal of Economic Dynamics and Control*, 34(1), 4–27.

Den Haan, Wouter J., and Pontus Rendahl (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using Explicit Aggregation," *Journal of Economic Dynamics and Control*, 34(1), 69–78.

Diamond, Peter A. (1965): "National Debt in a Neoclassical Growth Model," *American Economic Review*, 55, 1126–1150.

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

Fuhrer, Jeffrey C. (2000): "Habit Formation in Consumption and its Implications for Monetary Policy Models," *American Economic Review*, 90(3), 367–390, http://www.jstor.org/stable/117334.

Gorman, William M. (1953): "Community Preference Fields," *Econometrica*, 21(1), 63–80.

Guvenen, Fatih (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, 97(3), 687–712.

——— (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12.

Hall, Robert E. (1978): "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 96, 971–87, Available at http://www.stanford.edu/~rehall/Stochastic-JPE-Dec-1978.pdf.

——— (1988): "Intertemporal Substitution in Consumption," *Journal of Political Economy*, XCVI, 339–357, Available at http://www.stanford.edu/~rehall/Intertemporal-JPE-April-1988.pdf.

Hryshko, Dmytro (2010): "RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles," *Manuscript, University of Alberta.*

Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes (1994): "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in *The Carnegie-Rochester Conference Series on Public Policy*, ed. by Allan H. Meltzer, and Charles I. Plosser, vol. 40, pp. 59–126.

Jensen, Shane T., and Stephen H. Shore (2008): "Changes in the Distribution of Income Volatility," *Manuscript.*

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006): "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 96(5), 1589–1610.

——— (2009): "The Response of Consumer Spending to Rebates During an Expansion: Evidence from the 2003 Child Tax Credit," working paper, The Wharton School.

Kaplan, Greg, and Giovanni L. Violante (2011): "A Model of the Consumption Response to Fiscal Stimulus Payments," *NBER Working Paper Number W17338.*

Kennickell, Arthur B. (1995): "Saving and Permanent Income: Evidence from the 1992 SCF," *Board of Governors of the Federal Reserve System.*

Kim, Jinill, Sunghyun Henry Kim, and Robert Kollmann (2010): "Solving the Incomplete Market Model with Aggregate Uncertainty Using a Perturbation Method," *Journal of Economic Dynamics and Control*, 34(1), 50–58.

Kirman, Alan P. (1992): "Whom or What Does the Representative Individual Represent?," *The Journal of Economic Perspectives*, 6(2), 117–136.

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

Laibson, David (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, CXII(2), 443–477.

Low, Hamish, Costas Meghir, and Luigi Pistaferri (2005): "Wage Risk and Employment Over the Life Cycle," *Manuscript, Stanford University.*

Lucas, Robert E. (1976): "Econometric Policy Evaluation: A Critique," in *The Phillips Curve and Labour Markets*, ed. by Karl Brunner, and Allan H. Meltzer, vol. 1, pp. 19–46. Journal of Monetary Economics (Supplement).

Lusardi, Annamaria (1996): "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets," *Journal of Business and Economic Statistics*, 14(1), 81–90.

MaCurdy, Thomas (1982): "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics*, 18(1), 83–114.

Maliar, Lilia, Serguei Maliar, and Fernando Valli (2008): "Comparing Numerical Solutions of Models with Heterogeneous Agents (Models B): a Grid-based Euler Equation Algorithm," *Manuscript.*

——— (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell Smith Algorithm," *Journal of Economic Dynamics and Control*, 34(1), 42–49.

Meghir, Costas, and Luigi Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Journal of Business and Economic Statistics*, 72(1), 1–32.

Modigliani, Franco, and Richard Brumberg (1954): "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in *Post-Keynesian Economics*, ed. by Kenneth K. Kurihara, pp. 388–436. Rutgers University Press, New Brunswick, N.J.

Moffitt, Robert A., and Peter Gottschalk (1995): "Trends in the Covariance Structure of Earnings in the U.S.: 1969-1987," *Manuscript.*

Nielsen, Helena Skyt, and Annette Vissing-Jorgensen (2006): "The Impact of Labor Income Risk on Educational Choices: Estimates and Implied Risk Aversion," *Manuscript.*

Otsuka, Misuzu (2003): "Household Portfolio Choice with Illiquid Assets," manuscript, Johns Hopkins University.

Parker, Jonathan A. (1999): "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 89(4), 959–973.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland (2011): "Consumer Spending and the Economic Stimulus Payments of 2008," *NBER Working Paper Number W16684.*

Reiter, Michael (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty by Backward Induction," *Journal of Economic Dynamics and Control*, 34(1), 28–35.

Sabelhaus, John, and Jae Song (2010): "The Great Moderation in Micro Labor Earnings," *Journal of Monetary Economics*, 57(4), 391–403, http://ideas.repec.org/a/eee/moneco/v57y2010i4p391-403.html.

Sahm, Claudia R., Matthew D. Shapiro, and Joel B. Slemrod (2009): "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," *NBER Working Paper Number W15421.*

Samuelson, Paul A. (1958): "An exact consumption loan model of interest with or without the social contrivance of money," *Journal of Political Economy*, 66, 467–482.

Shapiro, Matthew W., and Joel B. Slemrod (2009): "Did the 2008 Tax Rebates Stimulate Spending?," *American Economic Review*, 99(2), 374–379.

Solow, Robert M. (2003): "Dumb and Dumber in Macroeconomics," *Manuscript, MIT.*

Souleles, Nicholas S. (1999): "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 89(4), 947–958.

——— (2002): "Consumer Response to the Reagan Tax Cuts," *Journal of Public Economics*, 85, 99–120.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2004a): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51(3), 609–633.

——— (2004b): "Cyclical Dynamics in Idiosyncratic Labor-Market Risk," *Journal of Political Economy*, 112(3), 695–717.

Szeidl, Adam (2006): "Invariant Distribution in Buffer-Stock Saving and Stochastic Growth Models," *Manuscript, University of California at Berkeley.*

Topel, Robert H. (1990): "Specific Capital, Mobility and Wages: Wages Rise with Job Seniority," *NBER Working Paper Number W3294.*

Topel, Robert H., and Michael P. Ward (1992): "Job Mobility and the Careers of Young Men," *Quarterly Journal of Economics*, 107(2), 439–479.

Young, Eric R. (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control*, 34(1), 36–41.