

# The Distribution of Wealth and the Marginal Propensity to Consume

May 7, 2013

Christopher D. Carroll<sup>1</sup>  
JHU

Jiri Slacalek<sup>2</sup>  
ECB

Kiichi Tokuoka<sup>3</sup>  
MoF, Japan

---

## Abstract

We present a macroeconomic model calibrated to match both microeconomic and macroeconomic evidence on household income dynamics, as well as standard empirical measures of the degree of wealth inequality. We show that our model's predictions (unlike those of competing models) are consistent with the substantial body of microeconomic evidence which suggests that the marginal propensity to consume is much larger than the 0.02–0.05 range implied by commonly-used macroeconomic models. Our model also (plausibly) predicts that the aggregate MPC can differ greatly depending on how the shock is distributed across categories of households (e.g., employed or unemployed).

---

**Keywords** Microfoundations, Wealth Inequality, Marginal Propensity to Consume

**JEL codes** D12, D31, D91, E21

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

---

<sup>1</sup>Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD, <http://econ.jhu.edu/people/ccarroll/>, [ccarroll@jhu.edu](mailto:ccarroll@jhu.edu)   <sup>2</sup>Slacalek: European Central Bank, Frankfurt am Main, Germany, <http://www.slacalek.com/>, [jiri.slacalek@ecb.europa.eu](mailto:jiri.slacalek@ecb.europa.eu)   <sup>3</sup>Tokuoka: Ministry of Finance, Tokyo, Japan, [kiichi.tokuoka@mof.go.jp](mailto:kiichi.tokuoka@mof.go.jp)

We thank Michael Ehrmann, Jonathan Parker and numerous seminar audiences for helpful comments. The views presented in this paper are those of the authors, and should not be attributed to the European Central Bank or the Japanese Ministry of Finance.

# 1 Introduction

In developed economies, wealth is very unevenly distributed. In the United States, for example, the top 1 percent of households own roughly one third of aggregate wealth, while the bottom 60 percent of households own very little, on net.<sup>1</sup>

We modify a standard Krusell and Smith (1998)-style macroeconomic model so that it matches empirical evidence on the degree of wealth inequality and on the nature and magnitude of microeconomic income shocks. We then show that this new model has very different implications than standard models (including both representative agent models and the original Krusell and Smith (1998) model) for the aggregate marginal propensity to consume (MPC) out of a common transitory shock that hits all households equally. Furthermore, if the shock does not hit all households equally, our model implies that its aggregate effect should depend upon how it is distributed across households. In particular, shocks that disproportionately affect poor households should have a disproportionately larger immediate effect on aggregate consumption, because theory implies that poorer households have a higher marginal propensity to consume. Both of these implications find support in the microeconomic evidence.

We make three principal modifications to Krusell and Smith (1998) in order to construct an economy that matches the empirical degree of wealth heterogeneity.

First, we endow households with an income process that matches the microeconomic data reasonably well. Following the verbal description of Friedman (1957) and empirical work of many others, we model income as an amalgam of a permanent component and a transitory component (both of which further consist of an idiosyncratic and an aggregate part). The income process we introduce is considerably more tractable than the KS process and is also a much closer match to both household-level and aggregate data.

Second, to ensure that stable wealth and income distributions exist, we assume that households face an exogenous risk of death *à la* Blanchard (1985). Consequently, each household has a finite lifetime. The wealth of those who die is distributed among survivors; newborns start life earning the mean level of (permanent) income.

Finally, to match the observed inequality in the distribution of wealth we allow for differences in time preference rates. We find that a modest degree of heterogeneity in impatience is sufficient to capture the extreme skewness in the empirical distribution of net wealth: it is enough that all households have a (quarterly) discount factor roughly between 0.98 and 0.99.

The baseline Krusell and Smith (1998) model uses an income process with low persistence (very far from the microeconomic evidence). This process generates little heterogeneity around the aggregate level of wealth which it is calibrated to match,

---

<sup>1</sup>The 2010 wave of the Survey of Consumer Finances found that the wealthiest 1 percent of households owned 35.3 percent of total net wealth, while the bottom 60 percent of households held 2.4 percent of wealth.

which is much higher than the median level of wealth in microeconomic data. Most households living in the Krusell–Smith world thus inhabit the flat region of the microeconomic consumption function and consequently the KS economy inherits a key feature of the corresponding representative-agent model: Households’ large levels of wealth serve as highly effective buffers, so aggregate household spending reacts only negligibly to transitory income shocks.<sup>2</sup>

In our model, because many households are slightly impatient and therefore hold little wealth, they are not able to insulate their spending even from transitory shocks very well. In our model, when households in the bottom half of the wealth distribution receive a one-off \$1 in income, they consume up to 50 cents of this windfall in the first year, ten times as much as the corresponding annual MPC in the baseline Krusell–Smith model. For the population as a whole, the aggregate annual MPC out of a common transitory shock ranges between 0.2 and 0.4, depending on whether we target our model to match the empirical distribution of net worth or of liquid assets.<sup>3</sup>

While these MPCs are roughly an order of magnitude larger than those implied by off-the-shelf representative agent models (about 0.02 to 0.05), they are in line with the large and growing empirical literature estimating the marginal propensity to consume summarized in Table 1. Various authors have estimated the MPC using quite different household-level datasets, in different countries, using alternative measures of consumption and diverse episodes of transitory income shocks; our reading of the literature is that most estimates of the aggregate MPC range between 0.2 and 0.6,<sup>4</sup> considerably exceeding the low values implied by representative agent models or the standard framework of Krusell and Smith (1998).

Our work also supplies a rigorous rationale for the conventional wisdom that the effects of an economic stimulus are particularly strong if it is targeted to poor individuals and to the unemployed. For example, our simulations imply that a stimulus targeted on the bottom half of the wealth distribution or the unemployed is 2–3 times more effective in increasing aggregate spending than a stimulus of the same size concentrated on the rest of the population. This finding is in line with the recent estimates of Blundell, Pistaferri, and Preston (2008), Broda and Parker (2012), Kreiner, Lassen, and Leth-Petersen (2012) and Jappelli and Pistaferri (2013), who report that households with little liquid wealth and without high past income react particularly strongly to an economic stimulus.<sup>5</sup>

---

<sup>2</sup>KS also explore a version of their model in which preferences change stochastically; we discuss the implications of this variation of their model at the appropriate point below.

<sup>3</sup>Because the holdings of liquid financial and retirement assets are substantially more heavily concentrated close to zero than holdings of net worth, the aggregate MPC in an economy calibrated to the former measure of wealth is considerably higher.

<sup>4</sup>Here and henceforth, when we use the term MPC without a timeframe, we are referring to the *annual* MPC; that is, the amount by which consumption is higher over the year following a transitory shock to income. This corresponds to the original usage by Keynes (1936) and Friedman (1957).

<sup>5</sup> Similar results are reported in Johnson, Parker, and Souleles (2006) and Agarwal, Liu, and Souleles (2007). Blundell, Pistaferri, and Saporta-Eksten (2012) estimate that older, wealthier households tend to use their assets more

**Table 1** Empirical Estimates of the Marginal Propensity to Consume (MPC) out of Transitory Income

Authors	Consumption Measure				
	Nondurables	Durables	Total PCE	Horizon*	Event/Sample
Blundell, Pistaferri, and Preston (2008) <sup>†</sup>	0.05				Estimation Sample: 1980–92
Coronado, Lupton, and Sheiner (2005)			0.36	1 Year	2003 Tax Cut
Hausman (2012)			0.6–0.75	1 Year	1936 Veterans’ Bonus
Jappelli and Pistaferri (2013)	0.48				Italy, 2010
Johnson, Parker, and Souleles (2009)	~ 0.25			3 Months	2003 Child Tax Credit
Lusardi (1996) <sup>‡</sup>	0.2–0.5				Estimation Sample: 1980–87
Parker (1999)	0.2			3 Months	Estimation Sample: 1980–93
Parker, Souleles, Johnson, and McClelland (2011)	0.12–0.30		0.50–0.90	3 Months	2008 Economic Stimulus
Sahm, Shapiro, and Slemrod (2010)			~ 1/3	1 Year	2008 Economic Stimulus
Shapiro and Slemrod (2009)			~ 1/3	1 Year	2008 Economic Stimulus
Souleles (1999)	0.045–0.09	0.29–0.54	0.34–0.64	3 Months	Estimation Sample: 1980–91
Souleles (2002)	0.6–0.9			1 Year	The Reagan Tax Cuts of the Early 1980s

Notes: \*: The horizon for which consumption response is calculated is 3 months or 1 year. The papers which estimate consumption response over the horizon of 3 months typically suggest that the response thereafter is only modest, so that the implied cumulative MPC over the full year is not much higher than over the first three months. †: elasticity.

Broda and Parker (2012) report the five-month cumulative MPC of 0.0836–0.1724 for the consumption goods in their dataset. However, the Homescan/NCP data they use only covers a subset of total PCE, in particular grocery and items bought in supercenters and warehouse clubs. We do not include the studies of the 2001 tax rebates, because our interpretation of that event is that it reflected a permanent tax cut that was not perceived by many households until the tax rebate checks were received. While several studies have examined this episode, e.g., Shapiro and Slemrod (2003), Johnson, Parker, and Souleles (2006), Agarwal, Liu, and Souleles (2007) and Misra and Surico (2011), in the absence of evidence about the extent to which the rebates were perceived as news about a permanent versus a transitory tax cut, any value of the MPC between zero and one could be justified as a plausible interpretation of the implication of a reasonable version of economic theory (that accounts for delays in perception of the kind that undoubtedly occur).

Recent work by Kaplan and Violante (2011) models an economy with households who choose between a liquid and an illiquid asset, which is subject to substantial transaction costs. Their economy features a substantial fraction of wealthy hand-to-mouth consumers, and consequently—like ours—responds strongly to a fiscal stimulus. In many ways their analysis is complementary to ours. While our setup does not model the choice between liquid and illiquid assets, theirs does not include transitory idiosyncratic (or aggregate) income shocks. A prior literature (all the way back to Deaton (1991, 1992)) has shown that the presence of transitory shocks can have a very substantial impact on the MPC (a result that shows up in our model), and the empirical literature cited below (including the well-measured tax data in DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013)) finds that such transitory shocks are quite large. Economic stimulus payments (like those studied by Broda and Parker (2012)) are precisely the kind of transitory shock to which we are interested in households’ responses, and so arguably a model (like ours) that explicitly includes transitory shocks (calibrated to micro evidence on their magnitude) is likely to yield more plausible estimates of the MPC when a shock of the kind explicitly incorporated in the model comes along (per Broda and Parker (2012)).

A further advantage of our framework is that it is consistent with the evidence which suggests that the MPC is higher for low-net-worth households. In the KV framework, among households of a given age, the MPC will vary strongly with the degree to which a household’s assets are held in liquid versus illiquid forms, but the relationship of the MPC to the household’s total net worth is less clear.

Finally, our model is a full rational expectations dynamic macroeconomic model, while their model does not incorporate aggregate shocks. Our framework is therefore likely to prove more adaptable to general-purpose macroeconomic modeling purposes.

On the other hand, given the substantial differences we find in MPC’s when we calibrate our model to match liquid financial assets versus when we calibrate it to match total net worth (reported below), the differences in our results across differing degrees of wealth liquidity would be more satisfying if we were able to explain them in a formal model of liquidity choice. For technical reasons not worth explicating here, the KV model of liquidity is not appropriate to our problem; given the lack of agreement in the profession about how to model liquidity, we leave that goal for future work (though preliminary experiments with modeling liquidity have persuaded us that the tractability of our model will make it a good platform for further exploration of this question).

The paper is structured as follows. Section 2 presents the income process we propose, consisting of idiosyncratic and aggregate shocks, each having a transitory

---

extensively to smooth spending. However, much of the empirical work (e.g., Souleles (2002), Misra and Surico (2011) or Parker, Souleles, Johnson, and McClelland (2011)) does not find that the consumption response of low-wealth or liquidity constrained households is statistically significantly higher, possibly because of measurement issues regarding credit constraints/liquid wealth and lack of statistical power.

and a permanent component. Section 3 lays out two variants of the baseline model—without and with heterogeneity in the rate of time preference—and explores how these models perform in capturing the degree of wealth inequality in the data. Section 4 compares the marginal propensities in these models to those in the Krusell and Smith (1998) model. Section 5 concludes.

## 2 The ‘Friedman/Buffer Stock’ Income Process

A key feature of our model is the labor income process, which closely resembles the verbal description of Friedman (1957) and which has been used extensively in the literature on buffer stock saving;<sup>6</sup> we therefore refer to it as the Friedman/Buffer Stock (or ‘FBS’) process.

Household income  $y_t$  consists of the aggregate wage rate  $W_t$  and two idiosyncratic components, a permanent component  $p_t$  and the transitory shock  $\xi_t$ :

$$y_t = p_t \xi_t W_t.$$

The permanent component follows a geometric random walk:

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

where the Greek letter psi mnemonically indicates the mean-one white noise permanent shock to income,  $\mathbb{E}_t[\psi_{t+n}] = 1 \ \forall \ n > 0$ . The transitory component is:

$$\xi_t = \mu \text{ with probability } u_t, \tag{2}$$

$$= (1 - \tau_t) \ell \theta_t \text{ with probability } 1 - u_t, \tag{3}$$

where  $\mu > 0$  is the unemployment insurance payment when unemployed,  $\tau$  is the rate of tax collected to pay unemployment benefits,  $\ell$  is time worked per employee and  $\theta$  is white noise.

The aggregate wage rate

$$W_t = (1 - \alpha) Z_t (K_t / \ell L_t)^\alpha, \tag{4}$$

is determined by productivity  $Z_t$  ( $= 1$ ), capital  $K_t$ , and the aggregate supply of effective labor  $L_t$ . The latter is again driven by two aggregate shocks:

$$L_t = P_t \Xi_t, \tag{5}$$

$$P_t = P_{t-1} \Psi_t, \tag{6}$$

where  $P_t$  is aggregate permanent productivity,  $\Psi_t$  is the aggregate permanent shock

---

<sup>6</sup>A large empirical literature has found that variants of this specification capture well the key features of actual household-level income processes; see Topel (1991), Carroll (1992), Moffitt and Gottschalk (2011), Storesletten, Telmer, and Yaron (2004), Low, Meghir, and Pistaferri (2010), DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013), and many others (see Table 1 in Carroll, Slacalek, and Tokuoka (2013) for a summary).

and  $\Xi_t$  is the aggregate transitory shock.<sup>7</sup> Like  $\psi_t$  and  $\theta_t$ , both  $\Psi_t$  and  $\Xi_t$  are assumed to be iid log-normally distributed with mean one.

Alternative specifications have been estimated in the extensive literature, and some authors argue that a better description of income dynamics is obtained by allowing for an MA(1) or MA(2) component in the transitory shocks, and by substituting AR(1) shocks for Friedman’s “permanent” shocks. The relevant AR and MA coefficients have recently been estimated in a new paper using a much higher-quality (and larger) data source than any previously available for the U.S.: IRS tax records. The authors’ point estimate for the size of the AR(1) coefficient is 0.98 (that is, very close to 1). Our view is that nothing of great substantive consequence hinges on whether the coefficient is 0.98 or 1.<sup>8,9</sup>

For modeling purposes, however, our task is considerably simpler both technically and to communicate to readers when we assume that the “persistent” shocks are in fact permanent. So that is what we do here.

## Comparison to the KS Income Process

Krusell and Smith (1998) assume that the level of aggregate productivity has a first-order Markov structure, alternating between two states:  $Z_t = 1 + \Delta^Z$  if the aggregate state is good and  $Z_t = 1 - \Delta^Z$  if it is bad; similarly,  $L_t = 1 - u_t$  (unemployment rate) where  $u_t = u^g$  if the state is good and  $u_t = u^b$  if bad. The idiosyncratic and aggregate shocks are thus correlated; the law of large numbers implies that the number of unemployed individuals is  $u^g$  and  $u^b$  in good and bad times, respectively.

The KS process for aggregate productivity shocks has little empirical foundation because the two-state Markov process is not flexible enough to match the empirical dynamics of unemployment or aggregate income growth well. In addition, the KS process—unlike income measured in the data—has low persistence. Indeed, the KS process 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 contrast, our assumption that the structure of aggregate shocks resembles the structure of idiosyncratic shocks is valuable not only because it matches the data well, but also because it makes the model easier to solve. In particular, the elimination of the ‘good’ and ‘bad’ aggregate states reduces the number of state variables to two (individual market resources  $m_t$  and aggregate capital  $K_t$ ) after normalizing the model appropriately. Employment status is not a state variable (in eliminating the

---

<sup>7</sup>Note that  $\Psi$  is the capitalized version of the Greek letter  $\psi$  used for the idiosyncratic permanent shock; similarly (though less obviously),  $\Xi$  is the capitalized  $\xi$ .

<sup>8</sup>Simulations have also convinced us that even if the true coefficient is 1, a coefficient of 0.98 might be estimated as a consequence of the bottom censorship of the tax data caused by the fact that those whose income falls below a certain threshold do not owe any tax.

<sup>9</sup>See Carroll, Slacalek, and Tokunaka (2013) for further discussion of these issues.



aggregate states, we also shut down unemployment persistence, which depends on the aggregate state in the KS model). As a result, given parameter values, solving the model with the FBS aggregate shocks is much faster than solving the model with the KS aggregate shocks.<sup>10</sup>

### 3 Modeling Wealth Heterogeneity: The Role of Shocks and Preferences

This section describes the key features of the framework in the absence of aggregate uncertainty.<sup>11</sup> Here, we allow for and for heterogeneity in time preference rates, and estimate the extent of such heterogeneity by matching the model-implied distribution of wealth to the observed distribution.<sup>12,13</sup>

#### 3.1 Homogeneous Impatience: The ‘ $\beta$ -Point Model’

The economy consists of a continuum of households of mass one distributed on the unit interval, each of which maximizes expected discounted utility from consumption,

$$\max \mathbb{E}_t \sum_{n=0}^{\infty} \beta^n u(c_{t+n})$$

for a CRRA utility function  $u(\bullet) = \bullet^{1-\rho}/(1-\rho)$ .<sup>14</sup> The household consumption functions  $\{c_{t+n}\}_{n=0}^{\infty}$  satisfy:

$$v(m_t) = \max_{c_t} u(c_t) + \beta \mathbb{E}_t [\psi_{t+1}^{1-\rho} v(m_{t+1})] \quad (7)$$

s.t.

$$a_t = m_t - c_t, \quad (8)$$

$$k_{t+1} = a_t / (\psi_{t+1}), \quad (9)$$

$$m_{t+1} = (\bar{\gamma} + r)k_{t+1} + \xi_{t+1}, \quad (10)$$

$$a_t \geq 0, \quad (11)$$

---

<sup>10</sup>As before, the main thing the household needs to know is the law of motion of aggregate capital, which can be obtained by following essentially the same solution method as in Krusell and Smith (1998) (see Appendix D of Carroll, Slacalek, and Tokuoka (2013) for details).

<sup>11</sup>Carroll, Slacalek, and Tokuoka (2013) provides further technical details of the setup.

<sup>12</sup>The key differences between Carroll, Slacalek, and Tokuoka (2013) and this paper are that the former includes neither aggregate FBS shocks nor heterogeneity in impatience. Also, Carroll, Slacalek, and Tokuoka (2013) does not investigate the implications of various models for the marginal propensity to consume.

<sup>13</sup>Terminologically, in the first setup (called ‘ $\beta$ -Point’ below) households have *ex ante* the same preferences and differ *ex post* only because they get hit with different shocks; in the second setup (called ‘ $\beta$ -Dist’ below) households are heterogeneous both *ex ante* (due to different discount factors) and *ex post* (due to different discount factors and different shocks).

<sup>14</sup>Substitute  $u(\bullet) = \log(\bullet)$  for  $\rho = 1$ .



where the variables are divided by the level of permanent income  $\mathbf{p}_t = p_t W$ , so that when aggregate shocks are shut down the only state variable is (normalized) cash-on-hand  $m_t$ .<sup>15</sup>

Households die with a constant probability  $D \equiv 1 - \varnothing$  between periods.<sup>16</sup> Consequently, the effective discount factor is  $\beta\varnothing$  (in (7)). The effective interest rate is  $(\Upsilon + r)/\varnothing$ , where  $\Upsilon = 1 - \delta$  denotes the depreciation factor for capital and  $r$  is the interest rate (which here is time-invariant and thus has no time subscript).<sup>17</sup> The production function is Cobb–Douglas:

$$ZK^\alpha(\ell L)^{1-\alpha}, \quad (12)$$

where  $Z$  is aggregate productivity,  $K$  is capital,  $\ell$  is time worked per employee and  $L$  is employment. The wage rate and the interest rate are equal to the marginal product of labor and capital, respectively.

As shown in (8)–(10), the evolution of household’s market resources  $m_t$  can be broken up into three steps:

1. Assets at the end of the period equal to market resources minus consumption:

$$a_t = m_t - c_t.$$

2. Next period’s capital is determined from this period’s assets via

$$k_{t+1} = a_t/(\varnothing\psi_t).$$

3. Finally, 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 is:

$$m_{t+1} = (\Upsilon + r)k_{t+1} + \xi_{t+1}.$$

Solving maximization (7)–(11) gives the optimal consumption rule. A target wealth-to-permanent-income ratio exists if a mortality-modified version of Carroll (2011)’s ‘Growth Impatience Condition’ holds:

$$\frac{(\mathbf{R}\beta)^{1/\rho} \mathbb{E}[\psi^{-1}]}{\Gamma} < 1, \quad (13)$$

---

<sup>15</sup>Again see Carroll, Slacalek, and Tokuoka (2013) for details.

<sup>16</sup>Following Blanchard (1985), the wealth of those who die is distributed among survivors; newborns start earning the mean level of income. Carroll, Slacalek, and Tokuoka (2013) show that a stable cross-sectional distribution of wealth exists if  $\varnothing\mathbb{E}[\psi^2] < 1$ .

<sup>17</sup>Below we allow time-varying interest rates implied by the aggregate dynamics of  $K_t$ ; for simplicity, the reader can think of the model here as a small open economy, and the model below as a closed economy.

where  $R = \Upsilon + r$ , and  $\Gamma$  is labor productivity growth (the growth rate of permanent income).<sup>18</sup>

## 3.2 Calibration

We calibrate the standard elements of the model using the parameter values used for the papers in the special issue of the *Journal of Economic Dynamics and Control* (2010) devoted to comparing solution methods for the KS model (the parameters are reproduced for convenience in Table 2). The model is calibrated at the quarterly frequency.

We calibrate the FBS income process as follows. The variances of idiosyncratic components are taken from Carroll (1992) because those numbers are representative of the large subsequent empirical literature all the way through a new paper by DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013) whose point estimate of the variance of the permanent shock almost exactly matches the calibration in Carroll (1992).<sup>19</sup>

The variances of the aggregate income process were estimated as follows, using US NIPA labor income, constructed as wages and salaries plus transfers minus personal contributions for social insurance. We first calibrate the signal-to-noise ratio  $\varsigma \equiv \sigma_{\Psi}^2 / \sigma_{\Xi}^2$  so that the first autocorrelation of the process, generated using the logged versions of equations (5)–(6), is 0.96.<sup>20,21</sup> Differencing equation (5) and expressing the second moments yields

$$\begin{aligned} \text{var}(\Delta \log L_t) &= \sigma_{\Psi}^2 + 2\sigma_{\Xi}^2, \\ &= (\varsigma + 2)\sigma_{\Xi}^2. \end{aligned}$$

Given  $\text{var}(\Delta \log L_t)$  and  $\varsigma$  we identify  $\sigma_{\Xi}^2 = \text{var}(\Delta \log L_t) / (\varsigma + 2)$  and  $\sigma_{\Psi}^2 = \varsigma \sigma_{\Xi}^2$ . The strategy yields the following estimates:  $\varsigma = 4$ ,  $\sigma_{\Psi}^2 = 4.29 \times 10^{-5}$  and  $\sigma_{\Xi}^2 = 1.07 \times 10^{-5}$  (given in Table 2).

This parametrization of the aggregate income process is consistent with assumptions in standard exercises in the real business cycle literature including Jermann (1998), Boldrin, Christiano, and Fisher (2001), and Chari, Kehoe, and McGrattan (2005). It also fits well the broad conclusion of the large literature on unit roots

---

<sup>18</sup>Appendix C of Carroll, Slacalek, and Tokuoka (2013) generalizes the ‘Growth Impatience Condition’ for a model with death,  $\beta \neq 1$ , like the one exposed here:  $\frac{(R\beta)^{1/\rho} \mathbb{E}[\psi^{-1}]}{\Gamma} \beta < 1$ .

<sup>19</sup>For a fuller survey, see Carroll, Slacalek, and Tokuoka (2013), which documents that the income process described in section 2 fits cross-sectional variance in the data much better than alternative processes which do not include a permanent, or at least a highly persistent, component.

<sup>20</sup>This calibration allows for transitory aggregate shocks, although the results below hold even in a model without transitory aggregate shocks, i.e., for  $\sigma_{\Xi}^2 = 0$ .

<sup>21</sup>We generate 10,000 replications of a process with 180 observations, which corresponds to 45 years of quarterly observations. The mean and median first autocorrelations (across replications) of such a process with  $\varsigma = 4$  are 0.956 and 0.965, respectively. In comparison, the mean and median of sample first autocorrelations of a pure random walk are 0.970 and 0.977 (with 180 observations), respectively.

**Table 2** Parameter Values and Steady State

Description	Parameter	Value	Source
Representative agent model			
Time discount factor	$\beta$	0.99	JEDC (2010)
Coef of relative risk aversion	$\rho$	1	JEDC (2010)
Capital share	$\alpha$	0.36	JEDC (2010)
Depreciation rate	$\delta$	0.025	JEDC (2010)
Time worked per employee	$\ell$	1/0.9	JEDC (2010)
Steady state			
Capital/(quarterly output) ratio	$K/Y$	10.26	JEDC (2010)
Effective interest rate	$r - \delta$	0.01	JEDC (2010)
Wage rate	$W$	2.37	JEDC (2010)
Heterogenous agents models			
Unempl insurance payment	$\mu$	0.15	JEDC (2010)
Probability of death	$D$	0.00625	Yields 40-year working life
FBS income shocks			
Variance of $\log \theta_{t,i}$	$\sigma_\theta^2$	$0.010 \times 4$	Carroll (1992), Carroll, Slacalek, and Tokuoka (2013)
Variance of $\log \psi_{t,i}$	$\sigma_\psi^2$	$0.010/4$	Carroll (1992), Carroll, Slacalek, and Tokuoka (2013)
Unemployment rate	$u$	0.07	Mean in JEDC (2010)
Variance of $\log \Xi_t$	$\sigma_\Xi^2$	0.00001	Authors' calculations
Variance of $\log \Psi_t$	$\sigma_\Psi^2$	0.00004	Authors' calculations
KS income shocks			
Aggregate shock to productivity	$\Delta^Z$	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)

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

of the 1980s, which found that it is virtually impossible to reject the existence of a permanent component in aggregate income series (see Stock (1986) for a review).

### 3.3 Wealth Distribution in the ‘ $\beta$ -Point’ Model

To finish calibrating the model, we assume (for now) that all households have an identical time preference factor  $\beta = \hat{\beta}$  (corresponding to a point distribution of  $\beta$ ) and henceforth call this specification the ‘ $\beta$ -Point’ model. With no aggregate uncertainty, we back out the value of  $\hat{\beta}$  for which the steady-state value of the capital-to-output ratio ( $K/Y$ ) matches the value that characterized the steady-state of the perfect foresight version of the model;  $\hat{\beta}$  turns out to be 0.9899 (at the quarterly rate).

Carroll, Slacalek, and Tokuoka (2013) show that the  $\beta$ -Point model matches the empirical wealth distribution substantially better than the version of the Krusell and Smith (1998) model analyzed in the Journal of Economic Dynamics and Control (2010) volume, which we call ‘KS-JEDC.’<sup>22</sup> For example, while the top 1 percent households living in the KS-JEDC model own only 3 percent of total wealth, those living in the  $\beta$ -Point are much richer, holding roughly 10 percent of total wealth. This improvement is driven by the presence of the permanent shock to income, which generates heterogeneity in the level of wealth because, while all households have the same target wealth/permanent income *ratio*, the equilibrium dispersion in the level of permanent income leads to a corresponding equilibrium dispersion in the level of wealth.

Figure 1 illustrates these results by plotting the wealth Lorenz curves implied by alternative models. Introducing the FBS shocks into the framework makes the Lorenz curve for the KS-JEDC model move roughly one third of the distance toward the data from the 2004 Survey of Consumer Finances, to the dashed curve labeled  $\beta$ -Point.

The wealth heterogeneity in the  $\beta$ -Point model essentially just replicates heterogeneity in permanent income (which accounts for most of the heterogeneity in total income). Since the distribution of wealth is considerably more unequal than the distribution of income (or permanent income), the setup only captures part of the wealth heterogeneity in the data, especially at the top.


### 3.4 Heterogeneous Impatience: ‘ $\beta$ -Dist Model’

Because we want a modeling framework that matches the fact that wealth inequality substantially exceeds income inequality, we need to introduce an additional source of heterogeneity (beyond heterogeneity in permanent and transitory income). We

---

<sup>22</sup>The only notable difference between the KS-JEDC and the original Krusell and Smith (1998) model is the introduction of unemployment insurance in the KS-JEDC version, which does not matter much for any substantive results. The key difference between our model described in section 3.1 and the KS-JEDC model is the income process. In addition, households in the KS-JEDC model do not die.

**Figure 1** Distribution of Net Worth (Lorenz Curve)



Notes: The solid curve shows the distribution of net worth in the 2004 Survey of Consumer Finances.

accomplish this by introducing heterogeneity in impatience. Each household is now assumed to have an idiosyncratic (but fixed) time preference factor. We think of this assumption as reflecting not only actual variation in pure rates of time preference across people, but also as reflecting other differences (in age, income growth expectations, investment opportunities, tax schedules, risk aversion, and other variables) that are not explicitly incorporated into the model.

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. Our “growth impatience condition” (13) captures the intuition that people facing faster income growth tend to act, financially, in a more ‘impatient’ fashion than those facing lower growth. So we should expect young people to have lower target wealth-to-income ratios than older people. Thus, what we are capturing by allowing heterogeneity in time preference factors is probably also some portion of the difference in behavior that (in truth) reflects differences in age instead of in pure time preference factors. Some of what we achieve by allowing heterogeneity in  $\beta$  could alternatively be introduced into the model if we had a more complex specification of the life cycle that allowed for different income growth rates for households of different ages.<sup>23</sup>

One way of gauging a model’s predictions for wealth inequality is to ask how well

<sup>23</sup>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 to distill the mechanics of the many moving parts 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 target wealth-to-income ratios. We achieve such a widening transparently and parsimoniously by incorporating discount factor heterogeneity.

it is able to match the proportion of total net worth held by the wealthiest 20, 40, 60, and 80 percent of the population. We follow other papers (in particular Castaneda, Diaz-Gimenez, and Rios-Rull (2003)) in matching these statistics.<sup>24</sup>

Our specific approach is to replace the assumption that all households have the same time preference factor with an assumption that, for some dispersion  $\nabla$ , time preference factors are distributed uniformly in the population between  $\dot{\beta} - \nabla$  and  $\dot{\beta} + \nabla$  (for this reason, the model is referred to as the ‘ $\beta$ -Dist’ model). Then, using simulations, we search for the values of  $\dot{\beta}$  and  $\nabla$  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  $\omega_i$  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:

$$\{\dot{\beta}, \nabla\} = \arg \min_{\{\dot{\beta}, \nabla\}} \sum_{i=20, 40, 60, 80} (w_i(\dot{\beta}, \nabla) - \omega_i)^2 \quad (14)$$

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 ( $K_{PF}/Y_{PF}$ ):<sup>25</sup>

$$K/Y = K_{PF}/Y_{PF}. \quad (15)$$

The solution to this problem is  $\{\dot{\beta}, \nabla\} = \{0.9876, 0.0060\}$ , so that the discount factors are evenly spread roughly between 0.98 and 0.99.<sup>26</sup>

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, bringing it reasonably in line with the data (Figure 1). 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.<sup>27</sup> But the model also sharply improves the fit to locations in the wealth distribution that were not explicitly targeted; for example, the 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

---

<sup>24</sup>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.

<sup>25</sup>In estimating these parameter values, we approximate the uniform distribution with the following seven points (each with the mass of 1/7):  $\{\dot{\beta} - 3\nabla/3.5, \dot{\beta} - 2\nabla/3.5, \dot{\beta} - \nabla/3.5, \dot{\beta}, \dot{\beta} + \nabla/3.5, \dot{\beta} + 2\nabla/3.5, \dot{\beta} + 3\nabla/3.5\}$ . Increasing the number of points further does not notably change the results below.

<sup>26</sup>With these estimates, even the most patient consumers with  $\beta = \dot{\beta} + 3\nabla/3.5$  (see footnote 25) satisfy the death-modified ‘Growth Impatience Condition’ of footnote 18 (a sufficient condition for stationarity of the wealth distribution), derived in Appendix C of Carroll, Slacalek, and Tokunaka (2013).

<sup>27</sup>Because the constraint (15) effectively pins down the discount factor  $\dot{\beta}$  estimated in the minimization problem (14), only the dispersion  $\nabla$  works to match the four wealth target points.

(1998) were well aware that their baseline model provides a poor match to the wealth distribution. In response, they 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 did improve the model’s ability to match the wealth holdings of the top percentiles. Indeed, unpublished results kindly provided by the authors show 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. By comparison, our model  $\beta$ -Dist 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 (though simple) 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. Krusell and Smith were not estimating a distribution in this manner; had they done so, they might have chosen a method like ours. 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 alternating 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 wealth data is a form of heterogeneity that leads different households to have different target levels of wealth.


## 4 The Aggregate Marginal Propensity to Consume

Having constructed a model with a realistic household income process which is able to reproduce steady-state wealth heterogeneity in the data, we now turn on aggregate shocks and investigate the model’s implications about relevant macroeconomic questions. In particular, we ask 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-



**Figure 2** Empirical Wealth Distribution and Consumption Functions of  $\beta$ -Point and  $\beta$ -Dist Models



Notes: The solid curve shows the consumption function for  $\beta$ -Point model. The dashed curves show the consumption functions for the most patient and the least patient consumers for  $\beta$ -Dist model. The histogram shows the empirical distribution of net worth ( $m_t$ ) in the Survey of Consumer Finances of 2004.

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 one-off ‘stimulus check’ worth some modest amount  $\$x$  (financed by a tax on unborn future generations).<sup>28</sup> Our question is: *By how much will aggregate consumption increase?*

#### 4.1 Matching Net Worth

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  $\beta$ -Point model’s individual consumption function using the FBS aggregate income process, 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 MPCs, 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  $\beta$ -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

<sup>28</sup>This financing scheme, along with the lack of a bequest motive, eliminates any Ricardian offset that might otherwise occur.

the  $\beta$ -Point model, which mirrors 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  $\beta$ -Point model is 0.1 in annual terms (second column of Table 3),<sup>29</sup> about double the value in the KS-JEDC model (0.05) (first column of the table) or the perfect foresight partial equilibrium model with parameters matching our baseline calibration (0.04). Our  $\beta$ -Dist model (third column of the table) produces an even higher average MPC (0.23), since in the  $\beta$ -Dist model there are more households who possess less wealth, are more impatient, and have higher MPCs (Figure 1 and dashed lines in Figure 2). However, this is still at best only at the lower bound of empirical MPC estimates, which are typically between 0.2–0.6 or even higher (see Table 1).

Comparison of columns 3 and 5 of Table 3 makes it clear that for the purpose of backing out the aggregate MPC, the particular form of the aggregate income process is not essential; both in qualitative and in quantitative terms the aggregate MPC and its breakdowns for the KS and the FBS aggregate income specification lie close to each other. This finding is in line with a large literature sparked by Lucas (1985) about the modest welfare cost of the aggregate fluctuations associated with business cycles and with the calibration of Table 2, in which variance of aggregate shocks is roughly two orders of magnitude smaller than variance of idiosyncratic shocks. (Of course, if one consequence of business cycles is to increase the magnitude of idiosyncratic shocks, as suggested for example by Guvenen, Ozkan, and Song (2012) and Blundell, Low, and Preston (2013), the costs of business cycles could be much larger than in traditional calculations that examine only the consequences of aggregate shocks.)

## 4.2 Matching Liquid Assets

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, with the attendant inconvenience and expense of 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

---


<sup>29</sup>The casual usage of the term ‘the MPC’ refers to annual MPC given by  $1 - (1 - \text{quarterly MPC})^4$  (recall again that the models in this paper are calibrated quarterly). We make this choice because existing influential empirical studies (e.g., Souleles (1999); Johnson, Parker, and Souleles (2006)) estimate longer-term MPCs for 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.

**Table 3** Average (Aggregate) Marginal Propensity to Consume in Annual Terms

Model	Krusell–Smith (KS) Aggregate Process			Friedman/Buffer Stock (FBS) Aggregate Process		
	KS-JEDC Our Solution	$\beta$ -Point	$\beta$ -Dist	$\beta$ -Dist	$\beta$ -Dist	$\beta$ -Dist
Wealth Measure		Net Worth	Net Worth	Liquid Financial and Retirement Assets	Net Worth	Liquid Financial and Retirement Assets
Overall average	0.05	0.10	0.23	0.44	0.20	0.42
By wealth/permanent income ratio						
Top 1%	0.04	0.06	0.05	0.12	0.05	0.12
Top 10%	0.04	0.06	0.06	0.12	0.06	0.12
Top 20%	0.04	0.06	0.06	0.13	0.06	0.13
Top 40%	0.04	0.06	0.08	0.20	0.06	0.17
Top 50%	0.05	0.07	0.09	0.25	0.06	0.22
Top 60%	0.04	0.07	0.12	0.28	0.09	0.24
Bottom 50%	0.05	0.13	0.35	0.59	0.32	0.58
By income						
Top 1%	0.05	0.07	0.12	0.17	0.16	0.36
Top 10%	0.05	0.08	0.15	0.28	0.17	0.36
Top 20%	0.05	0.09	0.16	0.31	0.17	0.36
Top 40%	0.05	0.10	0.18	0.35	0.19	0.38
Top 50%	0.05	0.11	0.19	0.36	0.18	0.39
Top 60%	0.05	0.10	0.19	0.37	0.20	0.39
Bottom 50%	0.05	0.09	0.27	0.51	0.22	0.45
By employment status						
Employed	0.05	0.09	0.20	0.40	0.19	0.39
Unemployed	0.06	0.23	0.53	0.81	0.41	0.73
Time preference parameters <sup>‡</sup>						
$\beta$		0.9899	0.9848	0.9570	0.9876	0.9636
$\nabla$			0.0094	0.0210	0.0060	0.0133

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

**Figure 3** Empirical Distribution of Liquid Financial Assets + Retirement Assets and Consumption Functions of  $\beta$ -Dist Model



Notes: The dashed curves show the consumption functions for the most patient and the least patient consumers for  $\beta$ -Dist model. The blue (dark grey) and pink (light grey) histograms show the empirical distributions of net worth and liquid financial and retirement assets, respectively, in the Survey of Consumer Finances of 2004.

the annual 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 permitted households to hold 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).<sup>30</sup>

These considerations suggest that it may be more plausible, for purposes of extracting predictions about the MPC out of stimulus checks, to focus on matching the distribution of liquid financial and retirement assets across households. The inclusion of retirement assets is arguable, but a case for inclusion can be made because retirement assets, such as IRA's and 401(k)'s, can be liquidated under a fairly clear rule (e.g., a penalty of 10 percent of the balance liquidated).

When we ask the model to estimate the time preference factors that allow it to best match the distribution of liquid financial and retirement assets (instead of net worth),<sup>31</sup> estimated parameter values are  $\{\beta, \nabla\} = \{0.9570, 0.0210\}$  under the KS aggregate income process and the average MPC is 0.44 (fourth column of the table),

<sup>30</sup>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.

<sup>31</sup>We define liquid financial and retirement assets as the sum of transaction accounts (deposits), CDs, bonds, stocks, mutual funds, and retirement assets. We take the same approach as before: we match the fraction of liquid financial and retirement assets held by the top 20, 40, 60, and 80 percent of the population (in the SCF 2004), while at the same time matching the aggregate liquid financial and retirement assets-to-income ratio (which is 6.6 in the SCF 2004).

**Table 4** Proportion of Wealth Held by Percentile (in Percent)

	Net Worth	Liquid Financial and Retirement Assets
Top 1%	33.9	34.6
Top 10%	69.7	75.3
Top 20%	82.9	88.3
Top 40%	94.7	97.5
Top 60%	99.0	99.6
Top 80%	100.2	100.0

Notes: The data source is the 2004 Survey of Consumer Finances.

which lies at the middle of the range typically reported in the literature (see Table 1), 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 and retirement assets than that of net worth (Table 4 and Figure 3) requires a wider distribution of the time preference factors, ranging between 0.94 and 0.975, which produces even more households with little wealth.<sup>32</sup> 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 and retirement assets (compared to net worth).

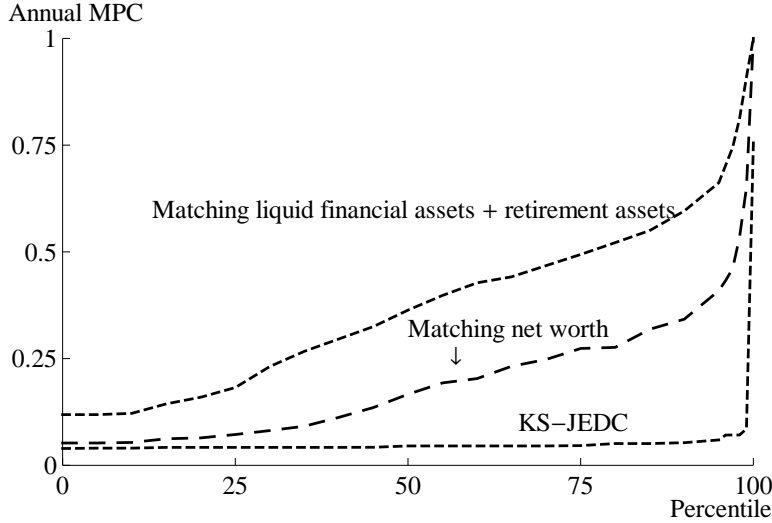
Figure 4 shows the cumulative distribution functions of MPCs for the KS-JEDC model and the  $\beta$ -Dist models (under the KS aggregate income shocks) estimated to match, first, the empirical distribution of net worth and, alternatively, of liquid financial and retirement assets.<sup>33</sup> The figure illustrates that the MPCs for KS-JEDC model are concentrated tightly around 0.05, which sharply contrasts with the results for the  $\beta$ -Dist models. Because the latter two models match the empirical wealth distribution, they imply that a substantial fraction of consumers has very little wealth.

Table 3 illustrates the distribution of MPCs by wealth, income, and employment status. In contrast to the KS-JEDC model, the  $\beta$ -Point and in particular  $\beta$ -Dist models generate a wide distribution of marginal propensities. Given the considerable concavity of the theoretical consumption function in the relevant region, these results indicate that the aggregate response to a stimulus program will depend greatly upon which households receive the stimulus payments. Furthermore, unlike the results from

<sup>32</sup>Table 4 also illustrates that the distribution of liquid financial and retirement assets is more concentrated close to zero than the distribution of net worth, e.g., the top 10 percent of households hold 75 percent of liquid assets and 70 percent of net worth.

<sup>33</sup>We have also solved a version of the model that matches only “very liquid assets” (excluding retirement and other assets that might not be instantly accessible); as would be expected, that exercise produces an even higher average MPC.

**Figure 4** Distribution of MPCs Across Households



the baseline KS-JEDC model or from a representative agent model, the results from these simulations are easily consistent with the empirical estimates of aggregate MPCs in Table 1 and the evidence that households with little liquid wealth and without high past income have high MPCs.<sup>34</sup>

## 5 Conclusion

We have shown that a model with a realistic microeconomic income process and modest heterogeneity in time preference rates is able to match the observed degree of inequality in the wealth distribution. Because many households in our model accumulate very little wealth, the aggregate marginal propensity to consume out of transitory income implied by our model, roughly 0.2–0.4 depending on the measure of wealth we ask our model to target, is consistent with most of the large estimates of the MPC reported in the microeconomic literature. Indeed, some of the dispersion in MPC estimates from the microeconomic literature (where estimates range up to 0.7 or higher) might be explainable by the model’s implication that there is no such thing as “the” MPC – the aggregate response to a transitory income shock should depend on details of the recipients of that shock in ways that the existing literature may not have been sensitive to (or may not have been able to measure). If some of the experiments reported in the literature reflected shocks that were concentrated in different regions of the wealth distribution than other experiments, considerable

<sup>34</sup>These studies include Blundell, Pistaferri, and Preston (2008), Broda and Parker (2012), Kreiner, Lassen, and Leth-Petersen (2012) and Jappelli and Pistaferri (2013).

variation in empirical MPCs would be an expected consequence of the differences in the experiments.

Additionally, our work provides researchers with an easier framework for solving, estimating, and simulating economies with heterogeneous agents and realistic income processes than has heretofore been available. Although benefiting from the important insights of [Krusell and Smith \(1998\)](#), our framework is faster and easier to solve than the KS model or many of its descendants, and thus can be used as a convenient building block for constructing micro-founded models for policy-relevant analysis.



## 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.
- BLANCHARD, OLIVIER J. (1985): “Debt, Deficits, and Finite Horizons,” *Journal of Political Economy*, 93(2), 223–247.
- BLUNDELL, RICHARD, HAMISH LOW, AND IAN PRESTON (2013): “Decomposing changes in income risk using consumption data,” *Quantitative Economics*, 4(1), 1–37.
- BLUNDELL, RICHARD, LUIGI PISTAFERRI, AND IAN PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5), 1887–1921.
- BLUNDELL, RICHARD, LUIGI PISTAFERRI, AND ITAY SAPORTA-EKSTEN (2012): “Consumption Inequality and Family Labor Supply,” NBER Working Papers 18445, National Bureau of Economic Research, Inc.
- BOLDRIN, MICHELE, LAWRENCE J. CHRISTIANO, AND JONAS D. FISHER (2001): “Habit Persistence, Asset Returns and the Business Cycle,” *American Economic Review*, 91(1), 149–66.
- BRODA, CHRISTIAN, AND JONATHAN A. PARKER (2012): “The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption,” mimeo, Northwestern University.
- 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>.
- (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., JIRI SLACALEK, AND KIICHI TOKUOKA (2013): “Buffer-Stock Saving in a Krusell–Smith World,” mimeo, Johns Hopkins University.
- 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.

- CHARI, V. V., PATRICK J. KEHOE, AND ELLEN R. MCGRATTAN (2005): “A Critique of Structural VARs Using Real Business Cycle Theory,” working paper 631, Federal Reserve Bank of Minneapolis.
- 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.
- (1992): *Understanding Consumption*. Oxford University Press, New York.
- DEBACKER, JASON, BRADLEY HEIM, VASIA PANOUSI, SHANTHI RAMNATH, AND IVAN VIDANGOS (2013): “Rising Inequality: Transitory or Permanent? New Evidence from a Panel of US Tax Returns,” *mimeo*.
- FRIEDMAN, MILTON A. (1957): *A Theory of the Consumption Function*. Princeton University Press.
- GUVENEN, FATIH, SERDAR OZKAN, AND JAE SONG (2012): “The Nature of Countercyclical Income Risk,” Working Paper 18035, National Bureau of Economic Research.
- HAUSMAN, JOSHUA K. (2012): “Fiscal Policy and Economic Recovery: The Case of the 1936 Veterans’ Bonus,” *mimeo*, University of California, Berkeley.
- JAPPELLI, TULLIO, AND LUIGI PISTAFERRI (2013): “Fiscal Policy and MPC Heterogeneity,” discussion paper 9333, CEPR.
- JERMANN, URBAN J. (1998): “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, 41(2), 257–75.
- 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.
- JOURNAL OF ECONOMIC DYNAMICS AND CONTROL (2010): “Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty,” *edited by Wouter J. Den Haan, Kenneth L. Judd, Michel Juillard*, 34(1), 1–100.

- KAPLAN, GREG, AND GIOVANNI L. VIOLANTE (2011): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *NBER Working Paper Number W17338*.
- KEYNES, JOHN MAYNARD (1936): *The General Theory of Employment, Interest, and Money*. Harcourt, Brace.
- KREINER, CLAUS THUSTRUP, DAVID DREYER LASSEN, AND SØREN LETH-PETERSEN (2012): “Heterogeneous Responses and Aggregate Impact of the 2001 Income Tax Rebates,” discussion paper 9161, CEPR.
- 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 (2010): “Wage Risk and Employment Over the Life Cycle,” *American Economic Review*, 100(4), 1432–1467.
- LUCAS, ROBERT E. (1985): *Models of Business Cycles, Yrjö Jahnsson Lectures*. Basil Blackwell, Oxford.
- 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.
- MISRA, KANISHA, AND PAOLO SURICO (2011): “Heterogeneous Responses and Aggregate Impact of the 2001 Income Tax Rebates,” discussion paper 8306, CEPR.
- MOFFITT, ROBERT, AND PETER GOTTSCHALK (2011): “Trends in the Covariance Structure of Earnings in the U.S.: 1969–1987,” *Journal of Economic Inequality*, 9, 439–459, doi: 10.1007/s10888-010-9154-z.
- 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*.

- SAHM, CLAUDIA R., MATTHEW D. SHAPIRO, AND JOEL B. SLEMROD (2010): “Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications,” *Tax Policy and the Economy*, 24, 69–110.
- SHAPIRO, MATTHEW D., AND JOEL SLEMROD (2003): “Consumer Response to Tax Rebates,” *American Economic Review*, 93(1), 381–396.
- SHAPIRO, MATTHEW W., AND JOEL B. SLEMROD (2009): “Did the 2008 Tax Rebates Stimulate Spending?,” *American Economic Review*, 99(2), 374–379.
- 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.
- STOCK, JAMES H. (1986): “Unit roots, structural breaks and trends,” in *Handbook of Econometrics*, ed. by R. F. Engle, and D. McFadden, vol. 4 of *Handbook of Econometrics*, chap. 46, pp. 2739–2841. Elsevier.
- STORESLETTEN, KJETIL, CHRIS I. TELMER, AND AMIR YARON (2004): “Consumption and Risk Sharing Over the Life Cycle,” *Journal of Monetary Economics*, 51(3), 609–633.
- TOPEL, ROBERT H. (1991): “Specific Capital, Mobility and Wages: Wages Rise with Job Seniority,” *Journal of Political Economy*, 99, 145–176.