The Distribution of Wealth and the Marginal Propensity to Consume

November 7, 2014

Christopher Carroll[1]	Jiri Slacalek[2]	Kiichi Tokuoka[3]	Matthew N. White[4]
JHU	ECB	MoF, Japan	U of Delaware

Abstract

We present a macroeconomic model calibrated to match both microeconomic and macroeconomic evidence on household income dynamics. When the model is extended so that it can match empirical measures of wealth inequality in the U.S., we show that its predictions are consistent with extensive microeconomic evidence that suggests that the annual marginal propensity to consume (MPC) is much larger than the figure of roughly 0.04 implied by commonly-used macroeconomic models (even ones including some heterogeneity). Our model also plausibly predicts that the aggregate MPC can differ greatly depending on how the shock is distributed across households (e.g., low-wealth versus high-wealth, or employed versus unemployed).

Keywords Microfoundations, Wealth Inequality, Marginal Propensity to Consume

JEL codes D12, D31, D91, E21

PDF: http://econ.jhu.edu/people/ccarroll/papers/cstwMPC.pdf
Slides: http://econ.jhu.edu/people/ccarroll/papers/cstwMPC-Slides.pdf
Web: http://econ.jhu.edu/people/ccarroll/papers/cstwMPC/
Archive: http://econ.jhu.edu/people/ccarroll/papers/cstwMPC.zip
(Contains data and estimation code producing paper's results)

[1] Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD 21218-2685, http://econ.jhu.edu/people/ccarroll/, email: ccarroll@jhu.edu, phone: 410-381-1813 [2] Slacalek: DG Research, European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany, http://www.slacalek.com/, email: jiri.slacalek@ecb.europa.eu, phone: +49-69-1344-5047 [3] Tokuoka: Ministry of Finance, 3-1-1 Kasumigaseki, Chiyoda-ku, Tokyo 100-8940, Japan, email: kiichi.tokuoka@mof.go.jp, phone: +81-3-3581-4111 [4] White: Department of Economics, University of Delaware, Newark, DE 19702, http://www.lerner.udel.edu/faculty-staff/matthew-n-white/, email: mnwecon@udel.edu, phone: 302-831-1906

We thank Michael Ehrmann, Dirk Krueger, Jonathan Parker, Giorgio Primiceri, Gianluca Violante and numerous

1 Introduction

In capitalist economies, wealth is unevenly distributed. Recent waves of the triennial U.S. *Survey of Consumer Finances*, for example, have consistently found the top 1 percent of households holding about a third of total wealth, with the bottom 60 percent owning essentially no net wealth.[1]

Such inequality could matter for macroeconomics if households with different amounts of wealth respond differently to the same aggregate shock. Indeed, microeconomic studies (reviewed in section 2.2) have often found that the annual marginal propensity to consume out of one-time income shocks (henceforth, 'the MPC') is substantially larger for low-wealth than for high-wealth households. In the presence of such microeconomic heterogeneity, the aggregate size of, say, a fiscal shock is not sufficient to compute the shock's effect on spending; that effect will depend on how the shock is distributed across categories of households with different MPCs.

We began this project with the intuition that it might be possible to explain both wealth heterogeneity and MPC heterogeneity with a single mechanism: A description of household income dynamics that incorporated fully permanent shocks to household-specific income, calibrated using evidence from the existing large empirical microeconomics literature (along with correspondingly calibrated transitory shocks).[2,3]

In the presence of both transitory and permanent shocks, "buffer stock" models in which consumers have long horizons imply that decision makers aim to achieve a target *ratio* of wealth to permanent income. In such a framework, we thought it might be possible to explain the inequality in wealth as stemming mostly from inequality in permanent income (with any remaining wealth inequality reflecting the influence of appropriately calibrated transitory shocks). Furthermore, the optimal consumption function in such models is concave (that is, the MPC is higher for households with lower wealth ratios), just as the microeconomic evidence suggests.

[1] More specifically, in the 1998–2007 waves of the *Survey of Consumer Finances* the fraction of total net wealth owned by the wealthiest 1 percent of households ranges between 32.4 and 34.4 percent, while the bottom 60 percent of households held roughly 2–3 percent of wealth. The statistics from the 2010 SCF show even somewhat greater concentration, but may partly reflect temporary asset price movements associated with the Great Recession. Corresponding statistics from the recently released *Household Finance and Consumption Survey* show that similar (though sometimes a bit lower) degree of wealth inequality holds also across many European countries (see Carroll, Slacalek, and Tokuoka (2014)).

[2] Of course, we are not the first to have solved a model with transitory and permanent shocks; nor the first to attempt to model the MPC; see below for a literature review. Our paper's joint focus on the distribution of wealth and the MPC, however, is novel (so far as we know).

[3] The empirical literature typically finds that highly persistent (and possibly truly permanent) shocks account for a large proportion of the variation in income across households. For an extensive literature review, see Carroll, Slacalek, and Tokuoka (2013).

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. This paper is a revision of `this` one; a new section of the paper extends the original analysis to the case of a life cycle model, and Matthew White has joined as a coauthor.

In our calibrated model the degree of wealth inequality is indeed similar to the degree of permanent income inequality. And our results confirm that a model calibrated to match empirical data on income dynamics can reproduce the level of observed permanent income inequality in data from the *Survey of Consumer Finances*. But that same data shows that inequality in measured wealth is much greater than inequality in measured permanent income. Thus, while our initial model does better in matching wealth inequality than some competing models, its baseline version is not capable of explaining the observed wealth inequality in the U.S. as merely a consequence of permanent income inequality.

Furthermore, while the concavity of the consumption function in our baseline model does imply that low wealth households have a higher MPC, the size of the difference in MPCs across wealth groups is not as large as the empirical evidence suggests. And the model's implied aggregate MPC remains well below what we perceive to be typical in the empirical literature: 0.2–0.6 (see the literature survey below).

All of these problems turn out to be easy to fix. If we modify the model to allow a modest degree of heterogeneity in impatience across households, the modified model is able to match the distribution of wealth remarkably well. Moreover, the aggregate MPC implied by that modified model falls within the range of what we view as the most credible empirical estimates of the MPC (though at the low end).

In a further experiment, we recalibrate the model so that it matches the degree of inequality in *liquid financial assets,* rather than total net worth. Because the holdings of liquid financial assets are substantially more heavily concentrated close to zero than holdings of net worth, the model's implied aggregate MPC then increases to roughly 0.4, well into the middle of the range of empirical estimates of the MPC. Consequently, the aggregate MPC in our models is an order of magnitude larger than in models in which households are well-insured and react negligibly to transitory income shocks, having MPCs of 0.02–0.04.

We also compare the business-cycle implications of two alternative modeling treatments of aggregate shocks. In the simpler version, aggregate shocks follow the Friedmanesque structure of our microeconomic shocks: All shocks are either fully permanent or fully transitory. We show that the aggregate MPC in this setup essentially does not vary over the business cycle because aggregate shocks are small and uncorrelated with idiosyncratic shocks. Next, we present a version of the model where the aggregate economy alternates between periods of boom and bust, as in Krusell and Smith (1998). Intuition suggests that this model has more potential to exhibit cyclical fluctuations in the MPC, because aggregate shocks are correlated with idiosyncratic shocks. Finally, we confirm that the quantitative conclusions about the size of the MPC hold when we adopt a framework with overlapping generations of households with realistically calibrated life cycles.

In the models with aggregate shocks, we can explicitly ask questions like "how does the aggregate MPC differ in a recession compared to an expansion" or even more

complicated questions like "does the MPC for poor households change more than for rich households over the business cycle?" The answer is that neither the mean value of the MPC nor the distribution of MPCs changes much when the economy switches from one state to the other. To the extent that this feature of the model is a correct description of reality, the result is encouraging because it provides reason to hope that microeconomic empirical evidence about the MPC obtained during normal, nonrecessionary times may still provide a good guide to the effects of stimulus programs for policymakers confronting extreme circumstances like those of the Great Recession.[4]

The rest of the paper is structured as follows. The next section explains the relation of our paper's modeling strategy to (some of) the vast related literature. Section 3 presents the income process we propose, consisting of idiosyncratic and aggregate shocks, each having a transitory and a permanent component. Section 4 lays out two variants of the baseline, infinite horizon 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 5 compares the marginal propensities in these models to those in the Krusell and Smith (1998) model and investigates how the aggregate MPC varies over the business cycle. Section 6 shows that the quantitative conclusions about the MPC carry over into the setup with overlapping generations. Section 7 concludes.

2 Relation to the Literature

2.1 Theory

Our modeling framework builds on the heterogeneous-agents model of Krusell and Smith (1997, 1998), with the modification that we aim to accommodate transitory-and-permanent-shocks microeconomic income process that is a modern implementation of ideas dating back to Friedman (1957) (see section 3). However, directly adding permanent shocks to income would produce an ever-widening cross-sectional distribution of permanent income, which is problematic because satisfactory analysis typically requires that models of this kind have stable (ideally, invariant) distributions for the key variables, so that appropriate calibrations for the model's parameters that match empirical facts can be chosen.

For our baseline, infinite horizon model we solve this problem, essentially, by killing off agents in our model stochastically using the perpetual-youth mechanism of

[4]This is an interesting point because during the episode of the Great Recession there was some speculation that even if empirical evidence suggested high MPCs out of transitory shocks during normal times, tax cuts might be ineffective in stimulating spending because prudence might diminish the MPC of even taxpayers who would normally respond to transitory income shocks with substantial extra spending. While that hypothesis could still be true, it is not consistent with the results of our models.

Blanchard (1985): Dying agents are replaced with newborns whose permanent income is equal to the mean level of permanent income in the population, so that a set of agents with dispersed values of permanent income is replaced with newborns with the same (population-mean) permanent income. When the distribution-compressing force of deaths outweighs the distribution-expanding influence from permanent shocks to income, this mechanism ensures that the distribution of permanent income has a finite variance.

A large literature starting with Zeldes (1989) has studied life cycle models in which agents face permanent (or highly persistent) and transitory shocks; a recent example that reflects the state of the art is Kaplan (2012). For the most part, that literature has been focused on microeconomic questions like the patterns of consumption and saving (or, recently, inequality) over the life cycle, rather than traditional macroeconomic questions like the average MPC (though very recent work by Kaplan and Violante (2011), discussed in detail below, does grapple with the MPC). Such models are formidably complex, which probably explains why they have not (to the best of our knowledge) yet been embedded in a dynamic general equilibrium context like that of the Krusell and Smith (1998) type, which would permit the study of questions like how the MPC changes over the business cycle. However, in section 6 we present a life cycle model, which documents that our quantitative conclusions about the size of the MPC and its distribution across households continue to hold in a framework with overlapping generations.

Perhaps closest to our paper in modeling structure is the work of Castaneda, Diaz-Gimenez, and Rios-Rull (2003). That paper constructs a microeconomic income process with a degree of serial correlation and a structure to the transitory (but persistent) income shocks engineered to match some key facts about the cross-sectional distributions of income and wealth in microeconomic data. But the income process that those authors calibrated does not resemble the microeconomic evidence on income dynamics very closely because the extremely rich households are assumed to face unrealistically high probability (roughly 10 percent) of a very bad and persistent income shock. Further, Castaneda, Diaz-Gimenez, and Rios-Rull (2003) did not examine the implications of their model for the aggregate MPC, perhaps because the MPC in their setup depends on the distribution of the deviation of households' actual incomes from their (identical) stationary level. That distribution, however, does not have an easily measurable empirical counterpart.[5]

One important difference between the benchmark version of our model and most of the prior literature is our incorporation of heterogeneous time preference rates as a way of matching the portion of wealth inequality that cannot be matched by the dispersion in permanent income. A first point to emphasize here is that we find

[5] Heathcote (2005) uses an income process similar to Castaneda, Diaz-Gimenez, and Rios-Rull (2003) to calibrate an economy which matches the empirical wealth heterogeneity and has the aggregate MPC of 0.29, also thanks to households which are credit-constrained.

that quite a modest degree of heterogeneity in impatience is sufficient to let the model capture the extreme dispersion 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.

Furthermore, our interpretation is that our framework parsimoniously captures in a single parameter (the time preference rate) a host of deeper kinds of heterogeneity that are undoubtedly important in the data (for example, heterogeneity in expectations of income growth associated with the pronounced age structure of income in life cycle models). The sense in which our model 'captures' these forms of heterogeneity is that, *for the purposes of our question* about the aggregate MPC, the crucial implication of many forms of heterogeneity is simply that they will lead households to hold different wealth positions which are associated with different MPCs. Since our baseline model captures the distribution of wealth and the distribution of permanent income already, it is not clear that *for the purposes of computing MPCs*, anything would be gained by the additional realism obtained by generating wealth heterogeneity from a much more complicated structure (like a fully realistic specification of the life cycle). We confirm this point quantitatively in the life cycle framework of section 6. Similarly, it is plausible that differences in preferences aside from time preference rates (for example, attitudes toward risk, or intrinsic degrees of optimism or pessimism) might influence wealth holdings separately from either age/life cycle factors or pure time preference rates. Again, though, to the extent that those forms of heterogeneity affect MPCs by leading different households to end up at different levels of wealth, we would argue that our model captures the key outcome (the wealth distribution) that is needed for deriving implications about the MPC.

2.2 Empirics

In our preferred models, 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 that 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 about 0.2 and about 0.4, depending on whether we target our model to match the empirical distribution of net worth or of liquid assets.

While the MPCs from our models are roughly an order of magnitude larger than those implied by off-the-shelf representative agent models (about 0.02 to 0.04), they are in line with the large and growing empirical literature estimating the marginal propensity to consume summarized in Table 1 and reviewed extensively in Jappelli

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

Authors	Consumption Measure			Horizon⋆	Event/Sample
	Nondurables	Durables	Total PCE		
Agarwal and Quian (2014)			0.90	10 Months	Growth Dividend Program Singapore 2011
Blundell, Pistaferri, and Preston (2008)‡	0.05				Estimation Sample: 1980–92
Browning and Collado (2001)			~ 0		Spanish ECPF Data, 1985–95
Coronado, Lupton, and Sheiner (2005)			0.36	1 Year	2003 Tax Cut
Hausman (2012)			0.6–0.75	1 Year	1936 Veterans' Bonus
Hsieh (2003)‡	~ 0		0.6–0.75		CEX, 1980–2001
Jappelli and Pistaferri (2013)	0.48				Italy, 2010
Johnson, Parker, and Souleles (2009)	~ 0.25			3 Months	2003 Child Tax Credit
Lusardi (1996)‡	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)			$\sim 1/3$	1 Year	2008 Economic Stimulus
Shapiro and Slemrod (2009)			$\sim 1/3$	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 typically 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).

and Pistaferri (2010).[6] 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 while a couple of papers find MPCs near zero, most estimates of the aggregate MPC range between 0.2 and 0.6,[7] considerably exceeding the low values implied by representative agent models or the standard framework of Krusell and Smith (1997, 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 tax-or-transfer 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.[8]

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 significant 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

[6] See also Pistaferri and Saporta-Eksten (2012).

[7] 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).

[8] 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 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.

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 that 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 infinite horizon 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.

On the other hand, given the substantial differences we find in MPCs 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).

3 The 'Friedman/Buffer Stock' Income Process

A key component 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;[9] we therefore refer to it as the Friedman/Buffer Stock (or 'FBS') process.

Household income $\boldsymbol{y}_t$ is determined by the interaction of the aggregate wage rate W_t and two idiosyncratic components, the permanent component p_t and the transitory shock ξ_t:

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

The permanent component follows a geometric random walk:

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

[9] 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).

where the Greek letter *psi* mnemonically indicates the mean-one white noise *p*ermanent *s*hock to *i*ncome, $\mathbb{E}_t[\psi_{t+n}] = 1 \; \forall \; n > 0$. The transitory component is:

$$\begin{aligned} \xi_t &= \mu \text{ with probability } u_t, & (2) \\ &= (1-\tau_t)\ell\theta_t \text{ with probability } 1-u_t, & (3) \end{aligned}$$

where $\mu > 0$ is the unemployment insurance payment when unemployed, τ_t is the rate of tax collected to pay unemployment benefits, ℓ is time worked per employee and θ_t is white noise. (This specification of the unemployment insurance system is taken from the special issue of the the Journal of Economic Dynamics and Control (2010) on solution methods for the Krusell–Smith model.)

In our preferred version of the model, the aggregate wage rate

$$\mathsf{W}_t = (1-\alpha)Z_t(\boldsymbol{K}_t/\ell\boldsymbol{L}_t)^{\alpha}, \quad (4)$$

is determined by productivity Z_t $(= 1)$, capital $\boldsymbol{K}_t$, and the aggregate supply of effective labor $\boldsymbol{L}_t$. The latter is again driven by two aggregate shocks:

$$\begin{aligned} \boldsymbol{L}_t &= P_t\Xi_t, & (5) \\ P_t &= P_{t-1}\Psi_t, & (6) \end{aligned}$$

where P_t is aggregate permanent productivity, Ψ_t is the aggregate permanent shock and Ξ_t is the aggregate transitory shock.[10] Like ψ_t and θ_t, both Ψ_t and Ξ_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 of DeBacker, Heim, Panousi, Ramnath, and Vidangos (2013) 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.[11,12] 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.

This FBS aggregate income process differs substantially from that in the seminal paper of Krusell and Smith (1998), which assumes that the level of aggregate productivity has a first-order Markov structure, alternating between two states: $Z_t = 1+\triangle^Z$

[10] 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.

[11] 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.

[12] See Carroll, Slacalek, and Tokuoka (2013) for further discussion of these issues.

if the aggregate state is good and $Z_t = 1 - \triangle^Z$ if it is bad; similarly, $\boldsymbol{L}_t = 1 - u_t$ (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 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.[13]

Because of its familiarity in the literature, we present below (in section 5.3) comparisons of the results obtained using both alternative descriptions of the aggregate income process. Nevertheless, our preference is for the FBS process, not only because it yields a much more tractable model but also because it much more closely replicates empirical aggregate dynamics that have been targeted by a large applied literature.

4 Modeling Wealth Heterogeneity: The Role of Shocks and Preferences

This section describes the key features of our infinite horizon modeling framework.[14] Here, we allow for heterogeneity in time preference rates, and estimate the extent

[13] 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).

[14] Carroll, Slacalek, and Tokuoka (2013) provides further technical details of the setup.

of such heterogeneity by matching the model-implied distribution of wealth to the observed distribution.[15,16]

4.1 Homogeneous Impatience: The 'β-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,

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

for a CRRA utility function $\mathrm{u}(\bullet) = \bullet^{1-\rho}/(1-\rho)$.[17] The household consumption functions $\{\mathrm{c}_{t+n}\}_{n=0}^{\infty}$ satisfy:

$$\begin{aligned}
\mathrm{v}(m_t) &= \max_{\mathrm{c}_t} \ \mathrm{u}(\mathrm{c}_t) + \beta \mathbb{D} \mathbb{E}_t \left[\psi_{t+1}^{1-\rho} \mathrm{v}(m_{t+1})\right] && (7)\\
&\text{s.t.} \\
a_t &= m_t - c_t, && (8)\\
k_{t+1} &= a_t/(\mathbb{D}\psi_{t+1}), && (9)\\
m_{t+1} &= (\mathsf{R} + r)k_{t+1} + \xi_{t+1}, && (10)\\
a_t &\geq 0, && (11)
\end{aligned}$$

where the variables are divided by the level of permanent income $\boldsymbol{p}_t = p_t \mathsf{W}$, so that when aggregate shocks are shut down the only state variable is (normalized) cash-on-hand m_t.[18]

Households die with a constant probability $\mathsf{D} \equiv 1 - \mathbb{D}$ between periods.[19] Consequently, the effective discount factor is $\beta\mathbb{D}$ (in (7)). The effective interest rate is $(\mathsf{R} + r)/\mathbb{D}$, where $\mathsf{R} = 1 - \delta$ denotes the depreciation factor for capital and r is the interest rate (which here is time-invariant and thus has no time subscript).[20] The production function is Cobb–Douglas:

$$Z\boldsymbol{K}^{\alpha}(\ell \boldsymbol{L})^{1-\alpha}, \tag{12}$$

[15] 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.

[16] Terminologically, in the first setup (called 'β-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 'β-Dist' below) households are heterogeneous both *ex ante* (due to different discount factors) and *ex post* (due to different discount factors and different shocks).

[17] Substitute $\mathrm{u}(\bullet) = \log(\bullet)$ for $\rho = 1$.

[18] Again see Carroll, Slacalek, and Tokuoka (2013) for details.

[19] 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 $\mathbb{D}\mathbb{E}[\psi^2] < 1$.

[20] 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 Z is aggregate productivity, $\boldsymbol{K}$ is capital, ℓ is time worked per employee and $\boldsymbol{L}$ 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 are 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

$$k_{t+1} = a_t/(\cancel{\mathbb{D}}\psi_t).$$

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:

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

Solving the maximization (7)–(11) gives the optimal consumption rule. A target wealth-to-permanent-income ratio exists if a death-modified version of Carroll (2011)'s 'Growth Impatience Condition' holds (see Appendix C of Carroll, Slacalek, and Tokuoka (2013) for derivation):

$$\frac{(\mathsf{R}\beta)^{1/\rho}\,\mathbb{E}[\psi^{-1}]\cancel{\mathbb{D}}}{\Gamma} < 1, \tag{13}$$

where $\mathsf{R} = \cancel{\gimel} + r$, and Γ is labor productivity growth (the growth rate of permanent income).

4.2 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 the 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

Table 2 Parameter Values and Steady State of the Infinite Horizon Models

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/(quarterly) output ratio	$\boldsymbol{K/Y}$	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)
Probability of death	D	0.00625	Yields 40-year working life
FBS income shocks			
Variance of log $\theta_{t,i}$	σ_θ^2	0.010×4	Carroll (1992), Carroll, Slacalek, and Tokuoka (2013)
Variance of log $\psi_{t,i}$	σ_ψ^2	0.010/4	Carroll (1992), DeBacker et al. (2013), Carroll, Slacalek, and Tokuoka (2013)
Unemployment rate	u	0.07	Mean in JEDC (2010)
Variance of log Ξ_t	σ_Ξ^2	0.00001	Authors' calculations
Variance of log Ψ_t	σ_Ψ^2	0.00004	Authors' calculations
KS income shocks			
Aggregate 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)

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

(1992). The variances of idiosyncratic components lie in the upper part of the range spanned by empirical estimates.[21] However, we believe our values are reasonable also because the standard model omits expenditure shocks (such as a sudden shock to household's medical expenses or durable goods).[22, 23]

The variances of the aggregate component of the FBS income process were estimated as follows, using U.S. 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.[24,25] Differencing equation (5) and expressing the second moments yields

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

Given $\text{var}\big(\Delta \log \boldsymbol{L}_t\big)$ and ς we identify $\sigma_\Xi^2 = \text{var}\big(\Delta \log \boldsymbol{L}_t\big)/(\varsigma+2)$ 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 yields income dynamics that match the same aggregate statstics that are matched by 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 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).

4.3 Wealth Distribution in the 'β-Point' Model

To finish calibrating the model, we assume (for now) that all households have an identical time preference factor $\beta = \grave{\beta}$ (corresponding to a point distribution of β) and henceforth call this specification the 'β-Point' model. With no aggregate uncertainty, we follow the procedure of the papers in the JEDC volume by backing out the value

[21] For a fuller survey, see Carroll, Slacalek, and Tokuoka (2013), which documents that the income process described in section 3 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.

[22] When we alternatively set the quarterly standard deviation of transitory shocks to 0.1 (instead of the value of 0.2 implied by Table 2), the results below change only little (e.g., under the FBS aggregate income process, the average MPC for the economy calibrated to liquid assets is 0.4 (instead of 0.42).

[23] Table 2 calibrates variances of idiosyncratic income components based on annual data, as we have not been able to find any literature that models income dynamics at a frequency higher than annual and simultaneously matches the annual data that are the object of most scholarly study.

[24] This calibration allows for transitory aggregate shocks, although the results below hold even in a model without transitory aggregate shocks, i.e., for $\sigma_\Xi^2 = 0$.

[25] 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.

of $\grave{\beta}$ for which the steady-state value of the capital-to-output ratio ($\boldsymbol{K/Y}$) matches the value that characterized the steady-state of the perfect foresight version of the model; $\grave{\beta}$ turns out to be 0.9899 (at a quarterly rate).[26]

Carroll, Slacalek, and Tokuoka (2013) show that the β-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.'[27] For example, while the top 1 percent households living in the KS-JEDC model own only 3 percent of total wealth,[28] those living in the β-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,*[29] to the dashed curve labeled β-Point.

However, the wealth heterogeneity in the β-Point model essentially just replicates heterogeneity in permanent income (which accounts for most of the heterogeneity in total income); for example the Gini coefficient for permanent income measured in the *Survey of Consumer Finances* of roughly 0.5 is similar to that for wealth generated in the β-Point model. Since the empirical distribution of wealth (which has the Gini coefficient of around 0.8) 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.

4.4 Heterogeneous Impatience: The 'β-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

[26] Our calibration of $\rho = 1$ follows JEDC. We find, as previous work has found, that ρ and β are not sharply identifiable using methods of the kind we employ here. Our approach therefore is to set a value of one parameter (ρ) and estimate the other conditional on the assumed value of the first.

[27] 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 4.1 and the KS-JEDC model is the income process. In addition, households in the KS-JEDC model do not die.

[28] See below for a discussion of the extension of their model in which households experience stochastic changes to their time preference rates; that version implies more wealth at the top.

[29] For the empirical measures of wealth we target the data from 2004 (and include only households with positive net worth). The wealth distribution in the data was stable until 2004 or so, although it has been shifting during the housing boom and the Great Recession; the effects of these shifts on our estimates of β and ∇ are negligible.

Figure 1 Distribution of Net Worth (Lorenz Curve)

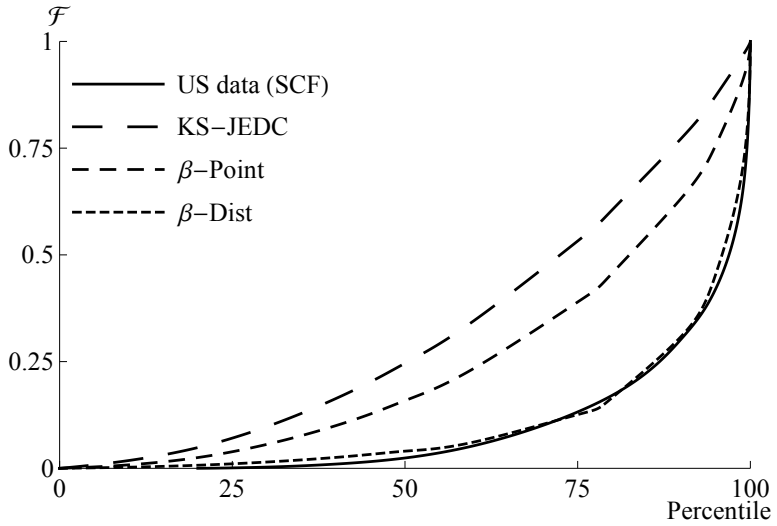

Notes: The solid curve shows the distribution of net worth in the 2004 *Survey of Consumer Finances*.

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 "death-modified 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 β 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. We make this point quantitatively in section 6 below, which solves the 'β-Dist' model in a realistic life cycle framework.

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. We follow other papers (in particular Castaneda, Diaz-Gimenez, and Rios-Rull (2003)) in matching these statistics.[30]

Our specific approach is to replace the assumption that all households have the same time preference factor with an assumption that, for some dispersion ∇, 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). 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:

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

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 ($\boldsymbol{K}_{PF}/\boldsymbol{Y}_{PF}$):[31]

$$\boldsymbol{K}/\boldsymbol{Y} \;=\; \boldsymbol{K}_{PF}/\boldsymbol{Y}_{PF}. \tag{15}$$

The solution to this problem is $\{\grave{\beta}, \nabla\} = \{0.9876, 0.0060\}$, so that the discount factors are evenly spread roughly between 0.98 and 0.99.[32]

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.[33] 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

[30] 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.

[31] In estimating these parameter values, we approximate the uniform distribution with the following seven points (each with the mass of 1/7): $\{\grave{\beta} - 3\nabla/3.5,\ \grave{\beta} - 2\nabla/3.5,\ \grave{\beta} - \nabla/3.5,\ \grave{\beta},\ \grave{\beta} + \nabla/3.5,\ \grave{\beta} + 2\nabla/3.5,\ \grave{\beta} + 3\nabla/3.5\}$. Increasing the number of points further does not notably change the results below. When solving the problem (14)–(15) for the FBS specification we shut down the aggregate shocks (practically, this does not affect the estimates given their small size).

[32] With these estimates, even the most patient consumers with $\beta = \grave{\beta} + 3\nabla/3.5$ (see footnote 31) satisfy the death-modified 'Growth Impatience Condition' of (13) (a sufficient condition for stationarity of the wealth distribution), derived in Appendix C of Carroll, Slacalek, and Tokuoka (2013).

[33] Because the constraint (15) effectively pins down the discount factor $\grave{\beta}$ estimated in the minimization problem (14), only the dispersion ∇ works to match the four wealth target points.

the top 10 percent and the top 1 percent are also shown in the figure, and the model performs reasonably well in matching them.[34]

Of course, Krusell and Smith (1997, 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, their results show that 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 β-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; estimation of their framework would have required searching for more than one parameter, and possibly as many as three of four. Indeed, had they intended to estimate parameters, they might have chosen a method more 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.

[34]We have examined the results for alternative calibrations of ρ; unsurprisingly, for larger calibrations of ρ, ∇ is larger. For example, for $\rho = 2$, ∇ is a bit more than twice as large. However, implications for the MPC are roughly similar.

5 The MPC in an Infinite Horizon Model

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-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 (financed by a tax on unborn future generations).[35] Our question is: *By how much will aggregate consumption increase?*

5.1 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 β-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 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 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 β-Point model is 0.1 in annual terms (second column of Table 3),[36] 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 β-Dist model (third column of the table) produces an even higher average MPC (0.23), since in the β-Dist model there are more households who possess less wealth,

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

[36] 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.

Figure 2 Empirical Wealth Distribution and Consumption Functions of the β-Point and β-Dist Models

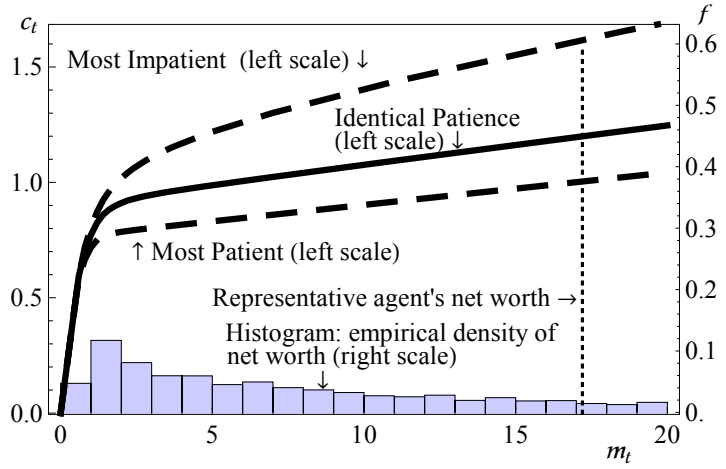

Notes: The solid curve shows the consumption function for β-Point model. The dashed curves show the consumption functions for the most patient and the least patient consumers for β-Dist model. The histogram shows the empirical distribution of net worth (m_t) in the *Survey of Consumer Finances* of 2004.

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).[37]

Kaplan, Violante, and Weidner (forthcoming) estimate that roughly a third of U.S. households are hand-to-mouth (in that they spend all their income in every pay-period). Of these households, roughly two thirds are wealthy—they own an illiquid asset—and the rest are poor. Because a state variable in our model is the ratio of wealth to permanent income, it can well be that households with low wealth–permanent income ratios own relatively high wealth (if their permanent income is high). In fact, a tabulation of the one third of households with the highest MPCs in the β-Dist model reveals that these households have quite diverse wealth holdings: half of them are in the bottom wealth quintile, one-third are in the second quintile and about 15 percent are in the third quintile.

Comparison of the third and fifth 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

[37]The MPCs calculated in Table 3 are 'theoretical', i.e., based on the slope of the consumption function. Alternatively, we have also calculated the following 'discrete' MPCs based on an increase in spending over the next four quarters after the household received an unexpected $ 1,000 extra in income. The implied MPCs for such calculation are slightly lower than the ones we report, e.g., for the aggregate MPC in the infinite horizon β-Dist model we get a value of 0.17 (instead of 0.2 reported in the second last column of Table 3).

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

	Krusell–Smith (KS) Aggregate Process				Friedman/Buffer Stock (FBS) Aggregate Process		Life Cycle Model		
Model	KS-JEDC Our Solution	β-Point	β-Dist	β-Dist	β-Dist	β-Dist	β-Point	β-Dist	β-Dist
Wealth Measure		Net Worth	Net Worth	Liquid Assets	Net Worth	Liquid Assets	Net Worth	Net Worth	Liquid Assets
Overall average	0.05	0.10	0.23	0.43	0.20	0.42	0.11	0.29	0.42
By wealth/permanent income ratio									
Top 1%	0.04	0.06	0.05	0.12	0.05	0.12	0.08	0.07	0.07
Top 10%	0.04	0.06	0.06	0.12	0.06	0.12	0.08	0.07	0.07
Top 20%	0.04	0.06	0.06	0.13	0.06	0.13	0.09	0.07	0.07
Top 40%	0.04	0.06	0.08	0.20	0.06	0.17	0.08	0.07	0.11
Top 50%	0.05	0.07	0.09	0.23	0.06	0.22	0.08	0.08	0.15
Top 60%	0.04	0.07	0.12	0.28	0.09	0.24	0.08	0.10	0.20
Bottom 50%	0.05	0.13	0.35	0.59	0.32	0.58	0.13	0.49	0.70
By income									
Top 1%	0.05	0.08	0.12	0.17	0.16	0.36	0.07	0.23	0.35
Top 10%	0.05	0.08	0.15	0.27	0.17	0.36	0.08	0.24	0.37
Top 20%	0.05	0.09	0.16	0.31	0.17	0.37	0.08	0.25	0.38
Top 40%	0.05	0.10	0.18	0.34	0.19	0.38	0.09	0.26	0.39
Top 50%	0.05	0.11	0.19	0.35	0.18	0.39	0.09	0.26	0.39
Top 60%	0.05	0.10	0.19	0.37	0.20	0.39	0.09	0.27	0.40
Bottom 50%	0.05	0.09	0.27	0.50	0.22	0.45	0.12	0.31	0.46
By employment status									
Employed	0.05	0.09	0.20	0.39	0.19	0.39	0.10	0.28	0.42
Unemployed	0.06	0.23	0.54	0.80	0.41	0.73	0.13	0.39	0.56
Time preference parameters‡									
$\grave{\beta}$		0.9899	0.9849	0.9573	0.9876	0.9636	0.9948	0.9888	0.9804
∇			0.0094	0.0206	0.0060	0.0133		0.0163	0.0242

Notes: Annual MPC is calculated by $1-(1-\text{quarterly MPC})^4$. "Liquid Assets" refers to liquid financial plus retirement assets. ‡: Discount factors are uniformly distributed over the interval $[\grave{\beta}-\nabla, \grave{\beta}+\nabla]$.

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.[38]

5.2 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 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).[39]

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 in the U.S. 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),[40] estimated parameter values are $\{\grave{\beta}, \nabla\} = \{0.9573, 0.0206\}$ under the KS

[38] Of course, if one consequence of business cycles is to increase the magnitude of idiosyncratic shocks, as suggested for example by McKay and Papp (2011), 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.

[39] 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.

[40] 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).

Figure 3 Empirical Distribution of Liquid Financial Assets + Retirement Assets and Consumption Functions of β-Dist Model

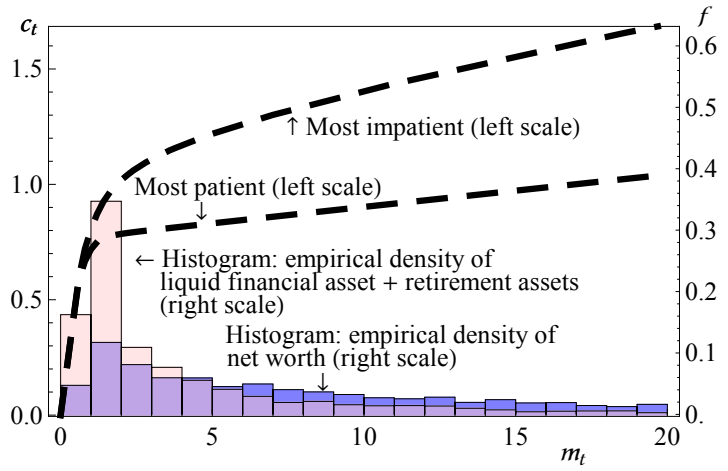

Notes: The dashed curves show the consumption functions for the most patient and the least patient consumers for β-Dist model. The blue (dark grey) and pink (light grey) histograms show the empirical distributions of net worth and liquid financial and retirement assets, respectively, in the *Survey of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

aggregate income process and the average MPC is 0.43 (fourth column of the table), 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 (see Figure 3) requires a wider distribution of the time preference factors, ranging between 0.94 and 0.98, which produces even more households with little wealth.[41] 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).[42]

Figure 4 shows the cumulative distribution functions of MPCs for the KS-JEDC model and the β-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.[43] The figure illustrates that the MPCs for KS-JEDC

[41]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.

[42]Our value of the survival probability $\mathcal{D} = 1 - 0.00625$ implies that 8 percent of households are older than 100 years. To keep the model consistent we keep them in the economy. However, the results essentially do not change—under the FBS aggregate shocks, the aggregate MPC is 0.43 instead of 0.42—if we alternatively replace the 100-year-olds with newborns (assuming they do not anticipate being replaced). This is reasonable given the small number of such households and given that the consumption function is almost linear at high levels of wealth.

[43]We have also solved a version of the model that matches only "very liquid assets" (excluding retirement and

Figure 4 Distribution of MPCs Across Households

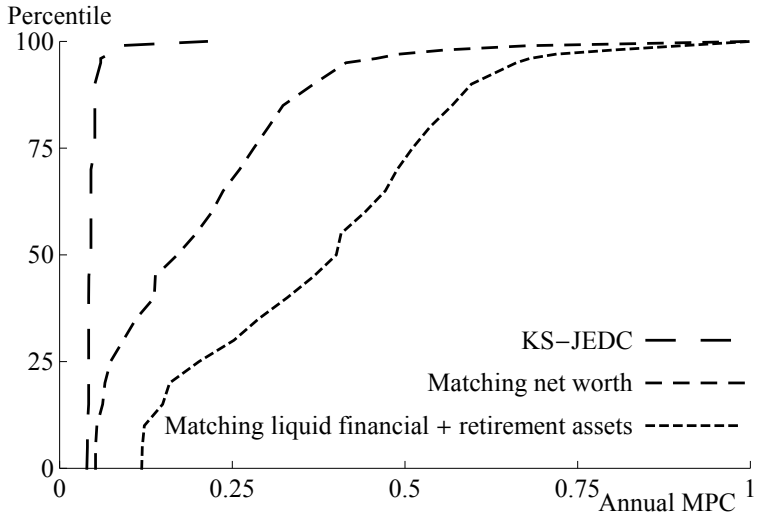

model are concentrated tightly around 0.05, which sharply contrasts with the results for the β-Dist models. Because the latter two models match the empirical wealth distribution, they imply that a substantial fraction of consumers has very little wealth.

Table 3 illustrates the distribution of MPCs by wealth, income, and employment status. In contrast to the KS-JEDC model, the β-Point and in particular β-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 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.[44]

5.3 MPC over the Business Cycle

Because our models include FBS or KS aggregate shocks, we can investigate how the economy's average MPC and its distribution across households varies over the business cycle. Table 4 reports the results for the following experiments with the β-Dist models calibrated to the net worth distribution (and compares them to the

other assets that might not be instantly accessible); as would be expected, that exercise produces an even higher average MPC.

[44]These studies include Blundell, Pistaferri, and Preston (2008), Broda and Parker (2012), Kreiner, Lassen, and Leth-Petersen (2012) and Jappelli and Pistaferri (2013).

baseline results from Table 3). For the model with KS aggregate shocks, in which recessions/expansions can be defined as bad/good realizations of the aggregate state:

1. 'Expansions vs. Recessions': $Z_t = 1 + \triangle^Z$ vs. $Z_t = 1 - \triangle^Z$.

2. 'Entering Recession': Bad realization of the aggregate state directly preceded by a good one: $Z_t = 1 - \triangle^Z$ for which $Z_{t-1} = 1 + \triangle^Z$.

For the model with FBS aggregate shocks, we consider large bad realizations of the aggregate shock:

1. 'Large Bad Permanent Aggregate Shock': bottom 1 percent of the distribution in the permanent aggregate shock

2. 'Large Bad Transitory Aggregate Shock': bottom 1 percent of the distribution in the transitory aggregate shock

In the KS setup, the aggregate MPC is countercyclical, ranging between 0.21 in expansions and 0.25 in recessions. The key reason for this business cycle variation lies in the fact that aggregate shocks are correlated with idiosyncratic shocks. The movements in the aggregate MPC are driven by the inadequately insured households at the bottom of the distributions of wealth and income. MPCs for rich and employed households essentially do not change over the business cycle. The scenario 'Entering Recession' documents that the length of the recession matters, so that initially the MPCs remain close to the baseline values, and increase only slowly as recession persists.

In the FBS setup, the distribution of the MPC displays very little cyclical variation for both transitory and permanent aggregate shocks. This fact is caused because the precautionary behavior of households is driven essentially exclusively by idiosyncratic shocks, as these shocks are two orders of magnitude larger (in terms of variance) and because they are uncorrelated with aggregate shocks.

Of course, these results are obtained under the assumptions that the parameters and expectations in the models are constant, and that the wealth distribution is exogenous. These assumptions are likely counterfactual in events like the Great Recession, during which objects such as expectations about the future income growth or the extent of uncertainty may well have changed.

As Figure 2 suggests, the aggregate MPC in our models is a result of an (interrelated) interaction between two objects: The distribution of wealth and the consumption function(s). During the Great Recession, the distribution of net worth has shifted very substantially downward. Specifically, Bricker, Kennickell, Moore, and Sabelhaus (2012) document that over the 2007–2010 period median net worth fell 38.8 percent (in real terms).[45] *Ceteris paribus,* these dynamics resulted an increase

[45]The *Survey of Consumer Finances* also documents that net worth decreased considerably relative to income; for example, the median net worth-to-income ratio declined from 8.5 in 2007 to 5.6 in 2010.

Table 4 Marginal Propensity to Consume over the Business Cycle

Model	Krusell–Smith (KS): β-Dist				Friedman/Buffer Stock (FBS): β-Dist		
Scenario	Baseline	Recession	Expansion	Entering Recession	Baseline	Large Bad Permanent Aggregate Shock	Large Bad Transitory Aggregate Shock
Overall average	0.23	0.25	0.21	0.23	0.20	0.20	0.21
By wealth/permanent income ratio							
Top 1%	0.05	0.05	0.05	0.05	0.05	0.05	0.05
Top 10%	0.06	0.06	0.06	0.06	0.06	0.06	0.06
Top 20%	0.06	0.06	0.06	0.06	0.06	0.06	0.06
Top 40%	0.08	0.08	0.08	0.08	0.06	0.06	0.06
Top 50%	0.09	0.10	0.09	0.09	0.06	0.06	0.09
Top 60%	0.12	0.12	0.11	0.12	0.09	0.09	0.09
Bottom 50%	0.35	0.38	0.32	0.35	0.32	0.32	0.32
By income							
Top 1%	0.12	0.13	0.11	0.13	0.16	0.16	0.17
Top 10%	0.15	0.16	0.15	0.16	0.17	0.17	0.17
Top 20%	0.16	0.17	0.16	0.17	0.17	0.17	0.18
Top 40%	0.18	0.19	0.18	0.18	0.19	0.19	0.19
Top 50%	0.19	0.20	0.18	0.19	0.18	0.19	0.20
Top 60%	0.19	0.20	0.19	0.20	0.20	0.20	0.20
Bottom 50%	0.27	0.30	0.24	0.27	0.22	0.21	0.22
By employment status							
Employed	0.20	0.20	0.20	0.20	0.19	0.19	0.19
Unemployed	0.54	0.56	0.51	0.47	0.41	0.41	0.41

Notes: Annual MPC is calculated by $1-(1-\text{quarterly MPC})^4$. The scenarios are calculated for the β-Dist models calibrated to the net worth distribution. For the KS aggregate shocks, the results are obtained by running the simulation over 1,000 periods, and the scenarios are defined as (i) 'Recessions/Expansions': bad/good realization of the aggregate state, $1-\triangle^Z/1+\triangle^Z$; (ii) 'Entering Recession': bad realization of the aggregate state directly preceded by a good one: $Z_t = 1-\triangle^Z$ for which $Z_{t-1} = 1+\triangle^Z$. The 'baseline' KS results are reproduced from column 3 of Table 3. For the FBS aggregate shocks, the results are averages over 1,000 simulations, and the scenarios are defined as (i) 'Large Bad Permanent Aggregate Shock': bottom 1 percent of the distribution in the permanent aggregate shock; (ii) 'Large Bad Transitory Aggregate Shock': bottom 1 percent of the distribution in the transitory aggregate shock. The 'baseline' FBS results are reproduced from column 5 of Table 3.

in the aggregate MPC, as the fraction of wealth-poor, high-MPC households rose substantially.

It is also likely that the second object, the consumption function, changed as many of its determinants (such as the magnitude of income shocks[46]) have not remained unaffected by the recession. And, of course, once parameters are allowed to vary, one needs to address the question about how households form expectations about these parameters. These factors make it quite complex to investigate adequately the numerous interactions potentially relevant for the dynamics of the MPC over the business cycle. Consequently, we leave the questions about the extent of cyclicality of the MPC in more complicated settings for future research.

6 The MPC in a Life Cycle Model

For ease of exposition and tractability of the aggregate shock processes, the models used in previous sections assume that households have infinite horizons, with no difference between "old" and "young" agents. Our qualitative results hold even when households are instead assumed to live out a finite life cycle, with more realistic assumptions about changes in the income process and mortality as the household ages. This section discusses the assumptions used in an overlapping generations life cycle specification and presents analogous results corresponding to the analysis in section 5 by re-estimating the β-Point and β-Dist models. In this environment, wealth heterogeneity emerges not only from shocks to permanent and transitory income and differences in discount factors, but also through demographic differences in age and education, via differential mortality and income growth expectations. While these latter factors were abstracted into time preference heterogeneity in our benchmark model, here we model them explicitly to demonstrate the robustness of our results to the simplifying assumptions.

6.1 Life Cycle of a Household

The economy consists of a continuum of expected utility maximizing households with a common CRRA utility function over consumption, $\mathrm{u}(\bullet) = \bullet^{1-\rho}/(1-\rho)$; each household has a time discount factor β. A household enters the economy at time t aged 24 years, endowed with an education level $e \in \{D, HS, C\}$ (for dropout, high school, and college, respectively), an initial permanent income level $\boldsymbol{p}_0$, and a stock of capital $\boldsymbol{k}_0$. Each quarter, the household receives (after tax) income, chooses how much

[46] See, e.g., Guvenen, Ozkan, and Song (2012) and Blundell, Low, and Preston (2013), and the literature on the 'scarring' effect of deep recessions on workers' lifetime income profiles.
Carroll, Slacalek, and Tokuoka (2014) document that an increase in the variance of transitory income shocks makes the consumption function steeper close to the origin.

of their market resources to consume and how much to save, and then transitions to the next quarter by facing shocks to mortality and income.

The FBS income process of section 3 translates into the life cycle framework as follows. A household receives a permanent shock to income when transitioning into period t, denoted by ψ_t (along with the age–education-specific average growth factor $\overline{\psi}_{es}$), as well as an after tax transitory shock ξ_t. The life cycle variant of the income process can be summarized by:

$$\begin{aligned} \boldsymbol{y}_t &= \xi_t \boldsymbol{p}_t = (1-\tau)\theta_t \boldsymbol{p}_t, \\ \boldsymbol{p}_t &= \psi_t \overline{\psi}_{es} \boldsymbol{p}_{t-1}. \end{aligned}$$

Households that have already lived for s periods have permanent shocks drawn from a lognormal distribution with mean 1 and variance $\sigma^2_{\psi s}$, and transitory shocks drawn from a lognormal distribution with mean $1/\mathscr{u}$ and variance $\sigma^2_{\theta s}$ with probability $\mathscr{u} = (1-u)$ and a degenerate distribution at μ with probability u. The prospect of unemployment (at rate u) is a completely transitory event: unemployment in period t has no effect on the probability of unemployment in period $t+1$. The non-zero transitory shock when unemployed represents a welfare benefit funded by income taxes, as discussed below. When transitioning from one period to the next, a household with education e that has already lived for s periods faces a D_{es} probability of death. In the main specification, the assets of a household that dies are completely taxed by the government to fund activities outside the model.[47]

The household's permanent income level will be factored out from the problem, so that the only state variable that affects the choice of optimal consumption is normalized market resources m_t. After this normalization, the household's budget transition functions can be described by:

$$\begin{aligned} a_t &= m_t - c_t, && (16) \\ k_{t+1} &= a_t/\psi_{t+1}, && (17) \\ m_{t+1} &= (\beth + r)k_{t+1} + \xi_{t+1}, && (18) \\ a_t &\geq 0. && (19) \end{aligned}$$

These transition constraints are identical to the infinite horizon model except that capital owned by surviving households does not grow with the inverse survival probability, and income is taxed at a marginal rate τ depending on the household's age and employment status.

Starting from some terminal age $\overline{s}$ at which $\mathsf{D}_{e\overline{s}} = 1$, a household's problem can be solved by backward induction until $s = 0$. At age $\overline{s}$, the household will consume all market resources, generating a consumption function of $\boldsymbol{c}_{e\overline{s}}(\boldsymbol{m}_t, \boldsymbol{p}_t) = \boldsymbol{m}_t = m_t \boldsymbol{p}_t$ and

[47] As a further robustness check, we also estimate versions in which the assets of the newly deceased are distributed to a random household, with varying preferences for bequests. In an online appendix we show that under a wide range of parameters governing preferences over bequests, both the overall aggregate marginal propensity to consume and its decompositions by wealth and income are little changed from the original specification.

a value function of $V_{e\bar{s}}(\boldsymbol{m}_t, \boldsymbol{p}_t) = \mathrm{u}(\boldsymbol{m}_t) = \boldsymbol{p}_t^{1-\rho}\mathrm{u}(m_t)$. At any earlier age, the value function is recursively defined by:

$$V_{es}(\boldsymbol{m}_t, \boldsymbol{p}_t) = \max_{\boldsymbol{c}_t} \mathrm{u}(\boldsymbol{c}_t) + \beta \not{\mathrm{D}}_{es}\, \mathbb{E}_t \left[V_{es+1}(\boldsymbol{m}_{t+1}, \boldsymbol{p}_{t+1})\right] \text{ s.t. (16)–(19).} \tag{20}$$

To eliminate the permanent income level as a state variable, further define the normalized consumption function as $\mathrm{c}_{es}(m_t) = \boldsymbol{c}_{es}(\boldsymbol{m}_t, \boldsymbol{p}_t)/\boldsymbol{p}_t$ and the normalized value function as $\mathrm{v}_{es}(m_t) = V_{es}(\boldsymbol{m}_t, \boldsymbol{p}_t)/\boldsymbol{p}_t^{1-\rho}$. Dividing (20) by $\boldsymbol{p}_t^{1-\rho}$, the problem is reduced to a single state dimension and can be expressed as:

$$\begin{aligned} \mathrm{v}_{es}(m_t) &= \max_{c_t} \mathrm{u}(c_t) + \beta \not{\mathrm{D}}_{es}\, \mathbb{E}_t \left[\psi_{t+1}^{1-\rho}\mathrm{v}_{es+1}(m_{t+1})\right] \text{ s.t. (16)–(19),} \\ \mathrm{c}_{es}(m_t) &= \arg\max_{c_t} \mathrm{u}(c_t) + \beta \not{\mathrm{D}}_{es}\, \mathbb{E}_t \left[\psi_{t+1}^{1-\rho}\mathrm{v}_{es+1}(m_{t+1})\right] \text{ s.t. (16)–(19).} \end{aligned} \tag{21}$$

A standard envelope condition applies in this model, so that $\mathrm{v}'_{es}(m_t) = \mathrm{u}'(\mathrm{c}_{es}(m_t))$, and the first order condition for the solution to (21) is:

$$c_t^{-\rho} = (\daleth + r)\beta \not{\mathrm{D}}\, \mathbb{E}_t \left[(\psi_{t+1}c_{t+1})^{-\rho}\right]. \tag{22}$$

In this way, the value function need not be tracked or recorded during the solution process, as the age-dependent consumption functions are sufficient.[48]

6.2 Macroeconomic Dynamics

The analysis in section 5 demonstrated that while there is considerable variation in the marginal propensity to consume across income, wealth, and employment status, the MPC does not appreciably change depending on the structure of aggregate shocks to the economy nor to the current macroeconomic state. Moreover, for reasons previously discussed, it is fairly difficult to account for macroeconomic state variables in an overlapping generations model. Rather than expend significant energy on a feature that would yield little of interest, we do not model aggregate shocks in this section but instead focus on the effects of idiosyncratic shocks and household-level dynamics. However, there are some additional macroeconomic features of the model that warrant discussion.

Unlike the infinite horizon perpetual youth model, the economy is now perpetually growing, with each new cohort larger than the last and ongoing technological progress. The expected permanent income growth for a household $\overline{\psi}_{es}$ comprises the household's own effective labor supply growth plus technological growth. When aggregating wealth, the contribution of a household that has already lived for s quarters is thus discounted by a factor of $(1+\Gamma)^{-s}$ relative to the youngest cohort, where Γ is the technological growth rate. Moreover, older households were born into smaller cohorts

[48] In practice, we use the method of endogenous gridpoints, as originally described in Carroll (2006), to discretize the state space and approximate consumption functions at each age and education level.

relative to the newest generation, so our population weighting scheme scales their contribution by the population growth rate N.

As mentioned in section 6.1, households are subject to a tax rate of τ depending on their age and employment status. Households are assumed to retire at age 65 (i.e. when $s = 164$), captured in the model with an expected permanent growth factor well below 1 at this age.[49] Income before retirement is earned through labor, while income after retirement is provided by a pay-as-you-go social security system funded by taxes on the employed. The social security tax rate is calculated as the rate that balances outlays to retired households and tax revenues from the working population:

$$\tau_{SS} = \frac{\sum_{e\in\{D,HS,C\}} \left[\theta_e \bar{\boldsymbol{p}}_{e0} \sum_{t=164}^{384} \left(((1+\Gamma)(1+N))^{-t} \prod_{s=0}^{t} (\overline{\psi}_{es} \not{\mathcal{D}}_{es})\right)\right]}{\sum_{e\in\{D,HS,C\}} \left[\theta_e \bar{\boldsymbol{p}}_{e0} \sum_{t=0}^{163} \left(((1+\Gamma)(1+N))^{-t} \prod_{s=0}^{t} (\overline{\psi}_{es} \not{\mathcal{D}}_{es})\right)\right]}.$$

Here, θ_e is the proportion of each new generation with education level e, and $\bar{\boldsymbol{p}}_{e0}$ is the average permanent income of that education type when they enter the economy at age 24. Note that neither permanent nor transitory shocks are relevant, as they average to unity across a cohort. The tax to fund unemployment benefits is simply the product of the unemployment rate and the benefit replacement rate: $\tau_U = u\mu$. Employed households pay a total income tax rate of $\tau = \tau_{SS} + \tau_U$, while unemployed and retired households have $\tau = 0$.

6.3 Calibration

Calibrations of the distributional parameters are taken from related estimates in the literature. Average permanent income growth rates $\overline{\psi}_{es}$ are calculated using the same trajectories as in Cagetti (2003) for those with less than a high school education, a high school degree, and four or more years of college. The permanent and transitory shock variances are approximated from the results of Sabelhaus and Song (2010), with extrapolation for ages 55–64.[50] Households are assumed to retire at age 65, withdrawing from the labor force and only receiving income from a pay-as-you-go social security system financed by taxes on the working population. Baseline mortality rates at each age are taken from the Social Security Administration's 2010 Actuarial Life Table,[51] then adjusted by education level using estimates by Brown, Liebman, and Pollett (2002) and converted to quarterly probabilities;[52] households die with certainty if they reach age 120. The unemployment benefit μ is set to 0.15 to match

[49] The drop in permanent income at retirement depends on the household's education: dropouts' income fall by 44%, high school graduates by 56%, and college graduates by 69%.

[50] We assume that $\sigma^2_\psi = \sigma^2_\theta = 0$ in retirement, so there is no income risk.

[51] Following the bulk of related literature, we use women's mortality rates to allow us to simulate households living past the husband's death.

[52] For ages 101–120, we use the adjustment for age 100, as this table does not extend to very late ages to which very few people live.

Table 5 Parameter Values in the Life Cycle Model

Description	Parameter	Value
Coefficient of relative risk aversion	ρ	1
Effective interest rate	$(r-\delta)$	0.01
Population growth rate	N	0.0025
Technological growth rate	Γ	0.0037
Rate of high school dropouts	θ_D	0.11
Rate of high school graduates	θ_{HS}	0.55
Rate of college graduates	θ_C	0.34
Average initial permanent income, dropout	$\overline{\boldsymbol{p}}_{D0}$	5000
Average initial permanent income, high school	$\overline{\boldsymbol{p}}_{HS0}$	7500
Average initial permanent income, college	$\overline{\boldsymbol{p}}_{C0}$	12000
Unemployment insurance payment	μ	0.15
Unemployment rate	u	0.07
Labor income tax rate	τ	0.0942

Cagetti (2003), while the unemployment probability is $u = 7\%$, the average rate in the infinite horizon model.

We assume that the population grows at a rate of 1% annually, while total factor productivity grows at a 1.5% annual rate; these approximately match long run rates in the United States. Educational attainment rates are set to be fairly consistent with U.S. educational rates over the past twenty years, and average initial permanent (quarterly) income at age 24 for each educational group are roughly calibrated to recent data.[53] Each simulated household is given an initial lognormal shock to permanent income with standard deviation 0.4, approximately matching the total variance of income among young households in the SCF 2004 data. Households begin with a very low wealth to permanent income ratio, drawn uniformly from $\{0.17, 0.50, 0.83\}$. Other basic parameters are set to match the values used in the infinite horizon model. A summary of the model parameters is provided in Table 5.

6.4 Aggregate Marginal Propensity to Consume

Following the same procedure as in the benchmark infinite horizon model, we first assume that all households have the same time preference factor $\grave{\beta}$ as in the β-Point model. Seeking the value of $\grave{\beta}$ at which the aggregate capital to income ratio matches

[53] Precision here is unimportant: after several years of simulation, the initial permanent income differences between types matters much less than their income growth trajectories and idiosyncratic shocks.

that of the perfect foresight version of the infinite horizon model ($\boldsymbol{K/Y} = 10.26$), we find $\grave{\beta} = 0.9948$. As before, the simulated distribution of wealth in the β-Point life cycle model matches the empirical distribution from the 2004 Survey of Consumer Finances considerably better than the KS-JEDC model; indeed, the life cycle model has a somewhat better fit than the infinite horizon model, moving about two thirds of the way from the KS-JEDC's Lorenz curve to the empirical distribution, rather than one third. The additional wealth heterogeneity arises through differences in households' expectations of the future that were suppressed in the infinite horizon model: income growth rates vary with both education and age (particularly the timing of retirement), while the increasing probability of death plays a key role in older households' target wealth-to-income ratio.

To better fit the empirical distribution of wealth, we again estimate the β-Dist model by assuming that each household has its own discount factor drawn from a uniform distribution, $\beta \in [\grave{\beta} - \nabla, \grave{\beta} + \nabla]$, approximated by a ten point discrete distribution.[54] As in section 4, the object of the estimation is to find parameters $\grave{\beta}$ and ∇ that minimize the distance between empirical and simulated shares of wealth held by the richest 20%, 40%, 60%, and 80%, as in (14), subject to the constraint that the simulated $\boldsymbol{K/Y}$ ratio equals the target ratio of 10.26, as in the β-Point model. Estimation reveals that the optimal parameters are $\{\grave{\beta}, \nabla\} = \{0.9889, 0.0162\}$, quite similar to the infinite horizon estimates of $\{0.9849, 0.0094\}$ in the infinite horizon model with Krusell–Smith macroeconomic shocks.

The β-Dist model is able to match the empirical Lorenz curve extremely well for the bottom 85% of the wealth distribution: the average difference between simulated and actual wealth shares at the levels of interest is less than 0.4%. Indeed, the life cycle model is able to match the extremely low asset holdings of the bottom half of the population significantly better than the infinite horizon model. However, the simulated share of wealth held by the very richest individuals is somewhat lower than the SCF data show; in contrast, the infinite horizon β-Dist model is able to match the Lorenz curve fairly well even in the top 10%. We surmise that this is due to differences in the distribution of income at the upper tail between the models. While about 8% of households in the infinite horizon model have working lives that exceed 100 years, allowing some of these households to attain extremely high permanent incomes through a series of fortuitous shocks,[55] individuals in the overlapping generations/finite horizon model retire with certainty after 41 working years, receiving a large negative shock to permanent income and experiencing no

[54] We implicitly assume that the distribution of discount factors is independent from education type. More realistically, individuals with higher discount factors are more likely to remain in school longer. As this would tend to make high income types retain even more assets for the future while low income types will save even less, we would need a narrower range of discount factors to match observed wealth inequality. For this reason and for the sake of simplicity, we ignore this complicating factor.

[55] While our results do not critically depend on the inclusion of these very long-lived households, the model requires a somewhat wider distribution of β when they are removed from the aggregation.

income shocks thereafter. Without the ability to generate extremely high income households, even the richest do not accumulate enough assets to match the top end of the empirical Lorenz curve. As the simulated distribution matches fairly well and we are concerned with the aggregate marginal propensity to consume, particularly among non-wealthy households, this is not a serious deficiency.

The right-hand panel of Table 3 shows that, across all households, the aggregate (annual) marginal propensity to consume in both the β-Point (0.11) and β-Dist (0.29) models is extremely similar to the corresponding averages in the infinite horizon model.[56,57] The relationship between wealth-to-permanent income and the MPC is nearly identical to the pattern in the infinite horizon case, with the MPC slowly rising with lower incomes among the wealthier half of the population, and spiking rapidly among the bottom half. However, the gradient of income to MPC is much shallower in the finite horizon model, with the wealthiest 1% of households' MPC only 20% less than the poorest half, rather than 50% less in the benchmark model. This is likely due to confounding effects from life cycle dynamics: income poor households in the overlapping generations model are made up of both the young (who have not had time to accumulate income growth) and the retired (whose cohorts began with lower initial permanent income and have experienced the large negative wage growth from retirement). [58]

Figure 5 presents the aggregate marginal propensity to consume by age for the entire population, as well as for the most patient and least patient types in the β-Dist model. After an initial drop as households build up a minimum buffer stock, the life cycle profile of the MPC takes an inverted U-shape for most β types: rising during the rapid income growth ages of 30–40 before falling as households anticipate their retirement and seek to retain assets to consume in old age. Post retirement, the MPC steadily grows as agents experience an ever increasing mortality risk. The most impatient households, with a quarterly discount factor of about $\beta = 0.9742$, have a significantly higher MPC throughout life as they disfavor saving—they begin saving for retirement less than ten years prior, and quickly deplete their assets if they live beyond age 75 (as evidenced by MPCs approaching 1 at these ages). In contrast, the

[56]When the annual marginal propensity to consume is calculated by simulating the change in consumption over four quarters resulting from an unexpected $1000 payment to each household, we find an aggregate value of 0.28, substantially the same and confirming the corresponding exercise in the infinite horizon model.

[57]Without much comment, we also present estimates of the β-Dist model when matching the empirical distribution of liquid financial and retirement assets rather than net worth, along with subpopulation average MPCs for these models. In each case, the results of the life cycle model align very well with the earlier findings in the infinite horizon setting.

[58] While the ratio of wealth to permanent income is a very strong determinant of the marginal propensity to consume in our model, the wide distribution of household incomes allows for even wealthy households to have high MPCs. Confirming a similar exercise in the benchmark model, we again find that among the one third of households with the highest MPCs, 51% are in the lowest wealth quintile, 32% are in the second wealth quintile, and 14% are in the middle wealth quintile. Even in a life cycle model in which wealth is highly correlated with both age and the marginal propensity to consume, there is still a significant fraction of "wealthy hand-to-mouth" households as found in Kaplan, Violante, and Weidner (forthcoming).

Figure 5 Aggregate Marginal Propensity to Consume by Age

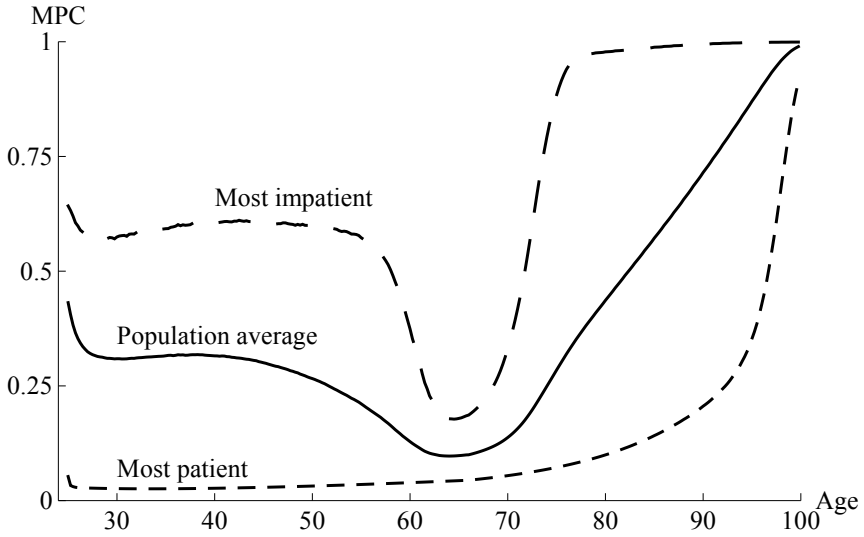

most patient households show an increasing marginal propensity to consume for their entire lives, though beginning from very low levels.

7 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 0.75 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

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 Reaction of Consumer Spending and Debt to Tax Rebates–Evidence from Consumer Credit Data," Discussion paper, National Bureau of Economic Research, http://www.nber.org/papers/w13694.

Agarwal, Sumit, and Wenlan Quian (2014): "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, forthcoming.

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.

Bricker, Jesse, Arthur B. Kennickell, Kevin B. Moore, and John Sabelhaus (2012): "Changes in U.S. Family Finances from 2007 to 2010: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, 98(2), 1–80.

Broda, Christian, and Jonathan A. Parker (2012): "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," mimeo, Northwestern University.

Brown, Jeffrey, Jeffrey B. Liebman, and Joshua Pollett (2002): "Estimating Life Tables That Reflect Socioeconomic Differences in Mortality," in *The Distributional Aspects of Social Security and Social Security Reform*, ed. by Martin Feldstein, and Jeffrey B. Liebman, pp. 447–457. University of Chicago Press.

Browning, Martin, and M. Dolores Collado (2001): "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates," *American Economic Review*, 91(3), 681–692.

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

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.

——— (2006): "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 91, 213–320.

——— (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.

——— (2014): "The Distribution of Wealth and the MPC: Implications of New European Data," *The American Economic Review*, 104(5), 107–111, At http://econ.jhu.edu/people/ccarroll/papers/cstMPCxc.

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, http://www.jstor.org/stable/2938366.

——— (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.

Heathcote, Jonathan (2005): "Fiscal Policy with Heterogeneous Agents and Incomplete Markets," *Review of Economic Studies*, 72(1), 161–188.

Hsieh, Chang-Tai (2003): "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund," *American Economic Review*, 99, 397–405.

Jappelli, Tullio, and Luigi Pistaferri (2010): "The Consumption Response to Income Changes," *The Annual Review of Economics*, 2, 479–506.

——— (2013): "Fiscal Policy and MPC Heterogeneity," discussion paper 9333, CEPR.

Jermann, Urban J. (1998): "Asset Pricing in Production Economies," *Journal of Monetary Economics*, 42(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 (2012): "Inequality and the Life Cycle," *Quantitative Economics*, 3(3), 471–525.

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

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner (forthcoming): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*.

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 (1997): "Income and Wealth Heterogeneity, Portfolio Choice and Equilibrium Asset Returns," *Macroeconomic Dynamics*, 1(2), 387–422.

——— (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, Yrjo 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.

McKay, Alisdair, and Tamas Papp (2011): "Accounting for Idiosyncratic Wage Risk Over the Business Cycle," working paper 28, Boston University.

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 (1999): "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, Vol 89(4), 959–73.

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*.

Pistaferri, Luigi, and Itay Saporta-Eksten (2012): "Changes in the Income Distribution and Aggregate Consumption," working paper 11, Ministry of Economy and Finance, Italy.

Sabelhaus, John, and Jae Song (2010): "The Great Moderation in Micro Labor Earnings," *Journal of Monetary Economics*, 57(4), 391–403.

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.

——— (2009): "Did the 2008 Tax Rebates Stimulate Spending?," *American Economic Review*, 99(2), 374–79.

Social Security Administration (2010): "Actuarial Life Table," available at http://www.ssa.gov/oact/STATS/table4c6.html.

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.

Zeldes, Stephen P. (1989): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104(2), 275–298.