

NBER WORKING PAPER SERIES

MACROECONOMICS AND HOUSEHOLD HETEROGENEITY

Dirk Krueger
Kurt Mitman
Fabrizio Perri

Working Paper 22319
<http://www.nber.org/papers/w22319>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2016

We wish to thank the editors John Taylor and Harald Uhlig, discussants of the chapter, Carlos Thomas, Claudio Michelacci, Tony Smith and Chris Tonetti, as well as seminar participants at several conferences and institutions for many useful comments and Simone Civalè for outstanding research assistance. Krueger gratefully acknowledges financial support from the NSF under grants SES 1123547 and SES 1326781. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Dirk Krueger, Kurt Mitman, and Fabrizio Perri. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Macroeconomics and Household Heterogeneity
Dirk Krueger, Kurt Mitman, and Fabrizio Perri
NBER Working Paper No. 22319
June 2016
JEL No. E21,E32,J65

ABSTRACT

The goal of this chapter is to study how, and by how much, household income, wealth, and preference heterogeneity amplify and propagate a macroeconomic shock. We focus on the U.S. Great Recession of 2007-2009 and proceed in two steps. First, using data from the Panel Study of Income Dynamics, we document the patterns of household income, consumption and wealth inequality before and during the Great Recession. We then investigate how households in different segments of the wealth distribution were affected by income declines, and how they changed their expenditures differentially during the aggregate downturn. Motivated by this evidence, we study several variants of a standard heterogeneous household model with aggregate shocks and an endogenous cross-sectional wealth distribution. Our key finding is that wealth inequality can significantly amplify the impact of an aggregate shock, and it does so if the distribution features a sufficiently large fraction of households with very little net worth that sharply increase their saving (i.e. they are not hand-to mouth) as the recession hits. We document that both these features are observed in the PSID. We also investigate the role that social insurance policies, such as unemployment insurance, play in shaping the cross-sectional income and wealth distribution, and through it, the dynamics of business cycles.

Dirk Krueger
Economics Department
University of Pennsylvania
160 McNeil Building
3718 Locust Walk
Philadelphia, PA 19104
and NBER
dkrueger@econ.upenn.edu

Fabrizio Perri
Research Department
Federal Reserve Bank of Minneapolis
90 Hennepin Avenue
Minneapolis MN 55401
fperri@umn.edu

Kurt Mitman
Institute for International Economic Studies
Stockholm University
106 91 Stockholm
SWEDEN
kurt.mitman@iies.su.se

1 Introduction

How important is household heterogeneity for the amplification and propagation of macroeconomic shocks? The objective of this chapter is to give a quantitative answer to a narrower version of this broad question.¹ Specifically, we narrow the focus of this question along two dimensions. First, we mainly focus on a specific macroeconomic event, namely the U.S. Great Recession of 2007-2009.² Second, we focus on specific dimensions of household heterogeneity, namely that in earnings, wealth, and household preferences, and their associated correlations with, and consequences for, the cross-sectional inequality in disposable income and consumption expenditures.³

The Great Recession was the largest negative macroeconomic downturn the United States has experienced since World War II. The initial decline in economic activity was deep and had an impact on all macroeconomic aggregates—notably private aggregate consumption and employment—and the recovery has been slow. Is the cross-sectional distribution of wealth an important determinant of the dynamics of the initial downturn and the ensuing recovery? That is, does household heterogeneity matter in terms of aggregate economic activity (as measured by output and labor input), its composition between consumption and investment, and, eventually, the cross-sectional distribution of consumption and welfare?

To address these questions *empirically*, we make use of recent waves of the Panel Study of Income Dynamics (PSID), which provides household-level panel data on earnings, income, consumption expenditures, and wealth for the United States. To answer these questions *theoretically and quantitatively*, we then study various versions of the canonical real business cycle model with aggregate technology shocks and ex-ante household heterogeneity in preferences and ex-post household income heterogeneity induced by

¹ In this chapter we focus on household heterogeneity. A sizeable literature has investigated similar questions in models with firm heterogeneity. Representative contributions from this literature include Khan and Thomas (2008) and Bachmann, Caballero and Engel (2013). We abstract from firm heterogeneity in this chapter, but note that the methodological challenges in computing these classes of models are very similar to the ones encountered here.

² By focusing on a business cycle event, and macroeconomic fluctuations more generally, we also abstract from the interaction between income or wealth inequality and aggregate income growth rates in the long run. See Kuznets (1955), Benabou (2002) or Piketty (2014) for important contributions to this large literature.

³ Excellent earlier surveys of different aspects of the literature on macroeconomics with microeconomic heterogeneity are contained in Deaton (1992), Attanasio (1999), Krusell and Smith (2006), Heathcote, Storesletten and Violante (2009), Attanasio and Weber (2010), Guvenen (2011) as well as Quadrini and Rios-Rull (2015).

the realization of uninsurable idiosyncratic labor earnings shocks, as in Krusell and Smith (1998). In the model, a recession is associated with lower aggregate wages and higher unemployment (i.e. a larger share of households with low labor income). The main empirical and model-based focus of the chapter is on the dynamics of macroeconomic variables—specifically, aggregate consumption, investment, and output—in response to such a business cycle shock. Specifically, we investigate the conditions under which the degree of wealth inequality plays a quantitatively important role for shaping this response. We also study how a stylized unemployment insurance program shapes the cross-sectional distribution of wealth and welfare, and how it affects the recovery of the aggregate economy after a Great Recession-like event.

We proceed in four steps: First, we make use of the PSID earnings, income, consumption and wealth data to document three sets of facts related to cross-sectional inequality. We summarize the key features of the joint distribution of income, wealth, and consumption prior to the Great Recession (i.e. for the year 2006). Next, we show how this joint distribution changed during the recession—over the 2006–2010 period—exploiting the panel dimension of the data to investigate how individual households fared and adjusted their consumption-savings behavior. The purpose of this empirical analysis is two fold. First, we believe the facts are interesting in their own right, as they characterize the distributional consequences of the Great Recession. Second, the facts serve as important moments for the evaluation of the different versions of the quantitative heterogeneous household model we study next.

In the second step, then, we construct, calibrate, and compute various versions of the canonical Krusell-Smith (1998) model and study its cross-sectional and dynamic properties. We first revisit the well-known finding that idiosyncratic unemployment risk and incomplete financial markets alone are insufficient to generate a sufficiently dispersed model-based cross-sectional wealth distribution. The problem is two fold: in the model, the very wealthy are not nearly wealthy enough, and the poor hold far too much wealth relative to the data. We argue that it is the discrepancy at bottom of the distribution that implies that the model generates an aggregate consumption response to a negative technology shock that is essentially identical to the response in a representative agent model.

We then study extensions of the model in which preference heterogeneity, idiosyncratic labor productivity risk conditional on employment, and a stylized life-cycle structure interact with the presence of unemployment insurance and social security to deliver a wealth distribution that is consistent with the data. In these economies, the decline in aggregate consumption is substantially larger than in the representative agent economy, by approximately 0.5 percentage points. This finding is primarily due to these economies now being populated by more wealth-poor households whose consumption responds strongly to the aggregate shock, both for those households that experience a transition from employment to unemployment, but also for households that have not lost their job, but understand they are facing a potentially long-lasting reces-

sion with elevated unemployment *risk*. We also stress that data and theory show that these wealth poor households do not behave as hand-to-mouth consumers, but are the group that reduces their expenditure rates strongly as their recession hits. This behavior implies that our benchmark model has quantitatively very different implications relative to a model where a large fraction of households is exogenously assumed to be hand-to-mouth consumers.

The more severe consumption declines in economies with larger wealth inequality imply a smaller collapse in investment, and thus a faster recovery from the recession, although this last effect is quantitatively small.

In light of the previous finding that larger wealth inequality—specifically, the importance of a large fraction of wealth-poor households—is an important contributor to an aggregate consumption collapse in the Great Recession, in the third step we determine whether public unemployment insurance is important for the dynamics of the economy in response to an aggregate shock. The answer to this question depends crucially on whether the distribution of household wealth has had a chance to respond to changes in the policy. In the short run, an unexpected cut or expiration of unemployment insurance benefits induces a significantly larger negative consumption response. These dynamics are explained by forward-looking households responding to lower public insurance by increasing their precautionary savings. The increased investment generates a medium-run boost to output, at the cost of a slow recovery of consumption.

In the long run, the new ergodic distribution of wealth features fewer people with zero or few assets. The consumption dynamics in response to a negative technology shock under this rightward shift in the wealth distribution are less severe than in they are in response to an unexpected shock, but still larger than in the economy with high unemployment insurance. Thus, for a *given wealth distribution* a cut in social insurance will result in a larger aggregate consumption drop. However, since social insurance policies themselves shape the ergodic distribution of wealth, and especially influence the share of households with zero or close to zero net worth, the aggregate consumption response across different economies is partially offset by these distributional shifts.

In the models considered thus far, the wealth distribution has had a potentially large effect on the *division* of aggregate output between consumption and investment, but not on output itself. In the final step, we therefore study an economy with a New Keynesian flavor—we introduce an aggregate demand externality that makes output partially demand-determined and generates an endogenous feedback effect from private consumption to total factor productivity, and thus output. In this model, social insurance policies might not only be beneficial in providing public insurance, but can also serve a potentially positive role for stabilizing aggregate output. We find that the output decline with an unemployment insurance benefit replacement rate of 50% to a Great Recession-like shock is 1 percentage point smaller on impact than in an economy with a replacement rate of 10%.

This work is part of a broader research agenda (and aims to partially synthesize it) that seeks to explore the importance of micro heterogeneity in general, and household income and wealth heterogeneity in particular, for classic macroeconomic questions (such as the impact of a particular aggregate shock) that have traditionally been answered within the representative agent paradigm (i.e. goes from micro to macro). It also builds upon, and contributes to, the related but distinct literature that studies the distributional consequences of macroeconomic shocks (i.e. goes from macro to micro).

The chapter is organized as follows. Section 2 documents key dimensions of heterogeneity among U.S. households, prior to and during the Great Recession. Sections 3 and 4 present our benchmark real business cycle model with household heterogeneity and discuss how we calibrate it. Section 5 studies to what extent the benchmark model is consistent with the cross-sectional facts presented in section 2, and section 6 studies how the aggregate consumption response to a large shock depends on the cross-sectional wealth distribution. In section 7 we augment the model with demand externalities in order to investigate the importance of cross-sectional wealth heterogeneity for the dynamics of aggregate output. Section 8 concludes, and the appendix contains details about the construction of the empirical facts, about the theory, and the computational algorithm used.

2 The Great Recession: a Heterogeneous Household Perspective

In this section we present the basic facts about the cross-sectional distribution of earnings, income, consumption, and wealth before and during the Great Recession. The main data set we employ is the Panel Survey of Income Dynamics (PSID) for the years 2004, 2006, 2008, and 2010. This data set has two key advantages for the purpose of this study. First, it contains information about household earnings, income, a broad and comprehensive measure of consumption expenditures, and net worth for a sample of households representative of the U.S. population. Second, it has a panel dimension so we can, in the same data set, both measure the key dimensions of cross-sectional household heterogeneity as well as investigate how different groups in the income and wealth distribution changed their consumption expenditure patterns during the Great Recession.⁴

The purpose of this empirical section is to provide simple and direct evidence for the importance of household heterogeneity for macroeconomic questions. It complements

⁴ Empirical analyses of the joint wealth, income, and consumption distribution using the same panel data set are also contained in Fisher, Johnson, Smeeding and Thompson (2015) for the United States, and in Krueger and Perri (2011) for Italy. See Skinner (1987), Blundell, Pistaferri and Preston (2008) and Smith and Tonetti (2014) for an alternative method for constructing an income-consumption panel using both the PSID and the Consumer Expenditure Survey (CE).

the large empirical literature documenting inequality trends in income, consumption and wealth in the United States and around the world.⁵ If, as we will document, there are significant differences in behavior (for example, along the consumption and savings margin) across different groups of the earnings and wealth distribution during the Great Recession, then keeping track of the cross-sectional earnings and wealth distribution and understanding their dynamics is likely important for analyzing the unfolding of the Great Recession from a macroeconomic and distributional perspective.

2.1 Aggregates

We start our analysis by comparing the evolution of basic U.S. macroeconomic aggregates from the National Income and Product Accounts (NIPA) with the aggregates for the same variables obtained from the PSID. In Figure 1 we compare trends in aggregate per capita disposable income (panel A) and per capita consumption expenditures (panel B) from the Bureau of Economic Analysis (BEA) with the corresponding series obtained by aggregating household level in the PSID, for the years 2004 through 2010, the last available data point for the PSID.⁶

The main conclusion we draw from figure 1 is that both the NIPA and the PSID paint the same qualitative picture of the U.S. macroeconomy over the period 2004-2010. Both disposable income and consumption expenditures experience a slowdown, which is somewhat more pronounced in the PSID. Furthermore, PSID consumption expenditure data also display a much weaker aggregate recovery than what is observed in the NIPA data.⁷

2.2 Inequality before the Great Recession

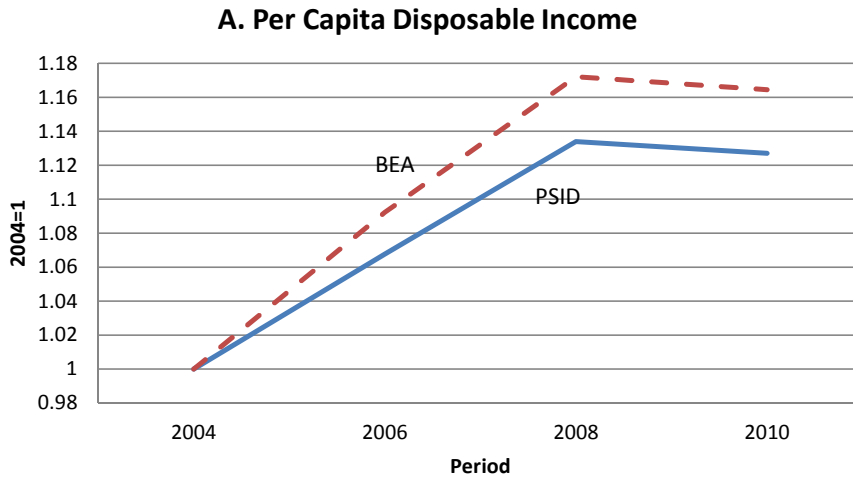
In this section we document basic inequality facts in the United States for the year 2006, just before the Great Recession hit the economy. Since the Great Recession greatly affected households in the labor market, and our models below focus on labor earnings risk, we restrict attention to households with heads between ages 25 and 60, which in 2006 represents slightly less than 80% of total households in the PSID. Table 1 reports statistics that characterize, for this group of households, the distributions of four key variables: earnings, disposable income, consumption expenditures, and net worth. Our definition of earnings captures income sources that we will model as ex-

⁵ For representative contributions, see e.g. Piketty and Saez, 2003, Krueger and Perri, 2006, Krueger, Perri, Pistaferri and Violante 2010, Piketty, 2014, Aguiar and Bils, 2015, Atkinson and Bourguignon, 2015 and Kuhn and Rios-Rull, 2015.

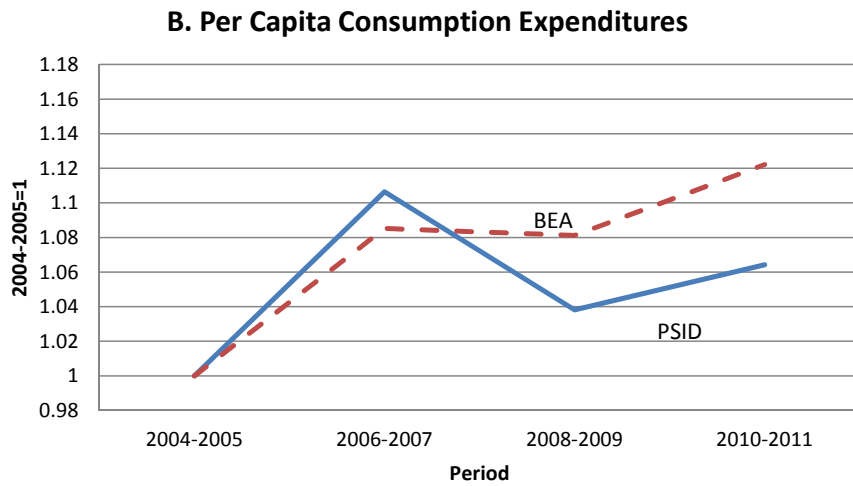
⁶ In section A.1 in the data appendix, we describe in detail how these series are constructed.

⁷ As Heathcote, Perri and Violante (2010) document, this discrepancy between macro data and aggregated micro data is also observed in previous recoveries from U.S. recessions.

Figure 1: The Great Recession in the NIPA and in the PSID Data



Note: In 2004 the per capita level in PSID is \$21364, in BEA is \$24120



Note: In 2004-05 the per capita level in PSID is \$15084, in BEA is \$18705

ogenous to household choices; they include all sources of labor income plus transfers (but not including unemployment benefits) minus tax liabilities.⁸ Disposable income includes earnings plus unemployment benefits, plus income from capital, including rental equivalent income of the main residence of the household. Consumption expenditures include all expenditure categories reported by the PSID i.e. cars and other vehicles purchases, food at home and away, clothing and apparel, housing including rent and imputed rental services for owners, household equipment, utilities and transportation expenses. Finally, net worth includes the value of the sums of households'

⁸ During the Great Recession, transfers and taxes have played an important role in affecting household income dynamics. See, for example, Perri and Steinberg, 2012.

assets minus liabilities.⁹

Table 1: Means and Marginal Distributions in 2006

Source	<i>Variable</i>						
	Earn.	Disp Y		Cons. Exp		Net Worth	
	PSID	PSID	CPS	PSID	CE	PSID	SCF(2007)
Mean (2006\$)	54,349	64,834	60,032	42,787	47,563	324,951	538,265
% Share by:							
Q1	3.6	4.5	4.4	5.6	6.5	-0.9	-0.2
Q2	9.9	9.9	10.5	10.7	11.4	0.8	1.2
Q3	15.3	15.3	15.9	15.6	16.4	4.4	4.6
Q4	22.7	22.8	23.1	22.4	23.3	13.0	11.9
Q5	48.5	47.5	46.0	45.6	42.4	82.7	82.5
90 – 95	10.9	10.8	10.1	10.3	10.2	13.7	11.1
95 – 99	13.1	12.8	12.8	11.3	11.1	22.8	25.3
Top 1%	8.0	8.0	7.2	8.2	5.1	30.9	33.5
Gini	0.43	0.42	0.40	0.40	0.36	0.77	0.78
Sample Size	6,232	6,232	54,518	6,232	4,908	6,232	2,910

Table 1 reports, for each variable (earnings, disposable income, consumption expenditures, and net worth), the cross-sectional average (in 2006 dollars), as well as the share of the total value held by each of the five quintiles of the corresponding distribution. At the bottom of the table, we also report the share held by the households between the 90th and 95th percentile, between the 95th and 99th percentile, by those in the top 1% of the respective distribution, and the Gini index of concentration. All statistics are computed from PSID data, but for disposable income, consumption expenditures, and net worth we also compare the statistics from the PSID with the same statistics computed from alternative micro data sets. In particular, for disposable income we use households from the 2006 Current Population Survey (CPS), which is a much larger sample often used to compute income inequality statistics. For consumption expenditures we use household data from the 2006 Consumer Expenditure Survey (CE). Finally, for net worth we use the 2007 Survey of Consumer Finances (SCF), which is the most commonly used dataset for studying the U.S. wealth distribution.

⁹ Assets include the value of farms and of any businesses owned by the household, the value of checking/saving accounts, the value of stocks or bonds owned, the value of primary residence and of other real estate assets, the value of vehicles, and the value of individual retirement accounts. Liabilities include any form of debt including mortgages on the primary residence or on other real estate, vehicle debt, student loans, medical debt, and credit card debt.

The table reveals features that are typical of distributions of resources across households in developed economies. Earnings and disposable income are both quite concentrated, with the bottom quintiles of the respective distributions holding shares smaller than 5% (3.6% and 4.5% to be exact) and the top quintiles holding almost 50% (48.5% and 47.5% to be precise). The distributions of earnings and disposable income look quite similar, since for the households in our sample (ages 25 to 60), capital income is a fairly small share of total disposable income (constituting only roughly 1/6 of disposable income).¹⁰ Note also that the distributions of disposable income in PSID and CPS look quite similar.¹¹

The table also shows that consumption expenditures are less unequally distributed than earnings or income, with the bottom quintile accounting for a bigger fraction (5.6%) of total expenditures. The distributions of consumption expenditures in the PSID and the CE are also fairly comparable.

Finally, net worth is by far the most concentrated variable, especially at the top of the distribution. The bottom 40% of households hold essentially no net worth at all, whereas the top quintile owns 83% of all wealth, and the top 10% holds around 70% of total wealth. Comparing the last two columns demonstrates that, although the average level of wealth in the PSID is substantially lower than in the SCF, the distribution of wealth across the five quintiles lines up quite closely between the two data sets, suggesting that the potential underreporting or mismeasurement of wealth in the PSID might affect the overall amount of wealth measured in this data set, but not the cross-sectional distribution too significantly, which is remarkably comparable to that in the SCF.

Although the marginal distributions of earnings, income, and wealth are interesting in their own right, the more relevant object for our purposes is the joint distribution of wealth, earnings, disposable income, and consumption expenditures.¹² To document the salient features of this joint distribution, we divide the households in our 2006 PSID sample into net worth quintiles, and then for each *net worth quintile* we report, in Table 2, key differences across these wealth groups.

The table shows two important features of the data. The first is that, perhaps not sur-

¹⁰ Recall that our definition of earnings is net of taxes and it already includes government transfers.

¹¹ The CPS income has a lower mean because it does not include the rental equivalent from the main residence. Notice also that both distributions are much less concentrated at the top than are income distributions computed by using tax returns, as in Piketty and Saez, (2008). Two reasons account for this difference. The first is that Piketty-Saez focus on income measures before taxes and transfers, whereas here we restrict attention to after-tax and after-transfers income, which is less concentrated; the second is that they focus on tax units, which is a unit of analysis different than households. See Burkhauser et al. (2012) for more on this distinction.

¹² The class of models we will construct below will have wealth – in addition to current earnings – as the crucial state variable, and thus we stress the correlation of net worth with earnings, income, and especially consumption here.

Table 2: PSID Households across the net worth distribution: 2006

NW Q	% Share of:			% Expend. Rate		Head's	
	Earn.	Disp Y	Expend.	Earn.	Disp Y	Age	Edu (yrs)
Q1	9.8	8.7	11.3	95.1	90.0	39.2	12
Q2	12.9	11.2	12.4	79.3	76.4	40.3	12
Q3	18.0	16.7	16.8	77.5	69.8	42.3	12.4
Q4	22.3	22.1	22.4	82.3	69.6	46.2	12.7
Q5	37.0	41.2	37.2	83.0	62.5	48.8	13.9
	Correlation with net worth						
	0.26	0.42	0.20				

prisingly, households with higher net worth tend to have higher earnings and higher disposable incomes. The last row of the table shows more precisely the extent to which earnings and disposable income are positively correlated with net worth. One simple explanation for this is that wealthier households tend to be older and more educated, as confirmed by the last two columns of the table. The second observation is that consumption expenditures are also positively correlated with net worth, but less so than the two income variables. The reason is that, as can be seen in the last two columns of the table, the lower is net worth, the higher the consumption rate. We measure the consumption rate by computing total consumption expenditures for a specific wealth quintile and then dividing it by total earnings (or disposable income) in that wealth quintile. The differences in the consumption rates across wealth quintiles are economically significant: for example, between the bottom and the top wealth quintile, the differences in the consumption rates range between 20% and 30%.

Another way to look at the same issue is to notice that the households in the bottom two net worth quintiles, basically hold no wealth (see table 1 above), but are responsible for $11.3\% + 12.4\% = 23.7\%$ of total consumption expenditures (see table 2), making this group quantitatively consequential for aggregate consumption dynamics. The differences across groups delineated by wealth constitute *prima facie* evidence that the shape of the wealth distribution *could* matter for the aggregate consumption response to macroeconomic shocks such as the ones responsible for the Great Recession.

In the next subsection, we will go beyond household heterogeneity at a given point in time and empirically evaluate how, during the Great Recession, expenditures and saving behavior changed differentially for households across the wealth distribution.

2.3 The Great Recession across the Income and Wealth Distributions

In table 3 we report for all households, and for households in each of the five quintiles of the net worth distribution, the changes (both percentages and absolute) in net worth, percentage changes in disposable income, and consumption expenditures and change in consumption expenditure rates (in percentage points).¹³ For each variable we first establish a benchmark (the growth rate in a non-recession period) by reporting the change or growth rate for the 2004-2006 period, and then report the same variable for the 2006-2010 period, which covers the whole recession. To make the two measures comparable, all changes are annualized.¹⁴

Table 3: Annualized Changes in Selected Variables across PSID Net Worth

	Net Worth*		Disp Y (%)		Cons. Exp.(%)		Exp. Rate (pp)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)		
	04-06	06-10	04-06	06-10	04-06	06-10	04-06	06-10		
All	15.7	44.6	-3.0	-10	4.1	1.2	5.6	-1.3	0.9	-1.6
NW Q										
Q1	NA	12.9	NA	6.6	7.4	6.7	7.1	0.6	-0.2	-4.2
Q2	121.9	19.5	24.4	3.7	6.7	4.1	7.2	2	0.3	-1.3
Q3	32.9	23.6	4.3	3.3	5.1	1.8	9	0	2.3	-1.1
Q4	17.0	34.7	1.7	3.8	5.0	1.7	5.9	-1.5	0.5	-2
Q5	11.6	132.2	-4.9	-68.4	1.8	-1.2	2.7	-3.5	0.5	-1.4

*The first figure is the percentage change (growth rate), the second is the change in 000's of dollars

Table 3 reveals a number of interesting facts that we want to highlight. From the first four columns of the table, notice that all groups of households experienced increases in net worth between 2004 and 2006, likely mainly because of the rapid growth in asset prices (stock prices and especially real estate prices) during this period, with low-wealth households experiencing the strongest percentage growth in wealth (but of course starting from very low levels: see again table 1). Turning to disposable income (second variable of table 3), we observe that households originally at the bottom of the wealth distribution experience faster disposable income growth than those in the highest wealth quintile (7.4% v/s 1.8%). This is most likely due to mean reversion

¹³ To construct these changes, we keep the identity of the households fixed; for example, to compute the 2004-2006 change in net worth for Q1 of the net worth distribution, we select all households in the bottom quintile of the wealth distribution in 2004, compute their average net worth (or income or consumption) in 2004 and 2006, and then calculate the percent difference between the two averages. For the consumption expenditure rates, we report percentage point differences.

¹⁴ Table A2 in the data appendix reports bootstrap standard errors for all figures in table 3. In tables A3 and A4, we separately report the changes for the 2006-2008 and 2008-2010 time periods.

in income: low-wealth households are also low-income households, and on average low income households experience faster income growth. Finally, expenditure growth roughly tracked the growth of income variables between 2004 and 2006, and as a result the consumption rates of each group remained roughly constant, perhaps with the exception of households initially in the middle quintile who experienced strong consumption expenditure growth, and thus their consumption rate displays a marked rise.

Now we turn to the dynamics in income, consumption, and wealth during the Great Recession. The columns labeled 06-10 display very significant changes in the dynamics of household income, consumption, and net worth throughout the wealth distribution, relative to the previous time period. Growth in net worth slowed down substantially for all households (it actually turned negative, from +15.6% to -3%) and most significantly so at the top of the wealth distribution. In fact, for households initially (that is, in 2006) in the top wealth quintile net worth fell 4.9% per year over the period 2006-2010. Income growth also slowed down, although not uniformly across the wealth distribution. Table 3 shows that the slowdown in income growth is modest at the bottom of the wealth distribution (from 7.4% to 6.7%), whereas the middle and top quintiles experience a more substantial slowdown. For example, the fourth wealth quintile went from annual disposable income growth of 5% between 2004 and 2006 to a growth rate of 1.7% between 2006 and 2010.

Most important for our purposes is the change in consumption expenditures at different points in the wealth distribution, especially in relation to the magnitude of the associated earnings and disposable income changes (as evident in the movement of the consumption rates over time). The first fact we want to highlight is that, overall, PSID households cut the growth in expenditures from +5.6% to -1.3%. Although the decline in the growth rate of consumption expenditures is sizeable across all quintiles, the fall is most pronounced at the bottom of the wealth distribution. To highlight the starkest differences across the wealth distribution, focus on the difference between the top and the bottom wealth quintile. Between 2004 and 2006 the households in both the bottom and the top wealth quintiles display small (less than 0.5 percentage point) changes in the consumption rate (out of disposable income). By contrast, between 2006 and 2010, households at the bottom end of the 2006 wealth distribution reduced the change in their consumption rate by 4 percentage points (from -0.2% to -4.2%), whereas the top quintile's change in the consumption rate declined by only 1.9 percentage points (from 0.5% to -1.4%). In other words, during the Great Recession saving rates increased across the wealth distribution, but more strongly so at the bottom of the wealth distribution.¹⁵

To investigate the sources of the decline in expenditures growth across the wealth distribution in greater detail, we now decompose the difference in consumption growth

¹⁵ Heathcote and Perri (2015) also document a similar pattern using data from the Consumer Expenditure Survey.

across the two periods as follows:

$$g_{c,it} - g_{c,it-1} \simeq g_{y,it} - g_{y,it-1} + \frac{\rho_{it} - \rho_{it-1}}{\rho_{it-1}} - \frac{\rho_{it-1} - \rho_{it-2}}{\rho_{it-2}}, \quad (1)$$

where $g_{c,it} = \frac{C_{it} - C_{it-1}}{C_{it-1}}$ is the growth rate of consumption expenditure for group i (for example households in the first wealth quintile in period $t - 1$) across periods t and $t - 1$, $g_{y,it}$ is the same measure for disposable income, and $\rho_{it} = \frac{C_{it}}{Y_{it}}$ is the consumption rate out of disposable income for group i in period t .

The first column of table 4 reports the changes in consumption growth rates for all households and for each group, i.e. the term $g_{c,it} - g_{c,it-1}$, which is the difference between column (6) and column (5) in table 3. The second and third columns of the table report the two right-hand-side terms from equation (1): the first term, labeled as change in disposable income growth Y , and the second term, labeled as change in the growth of the expenditure rate C/Y . Intuitively, if we see group i 's consumption growth slowing down, it could be because its income growth is slowing down, i.e. $g_{y,it} - g_{y,it-1}$ falls, or because, keeping income growth fixed, the growth in its expenditure rates, i.e. $\frac{\rho_{it} - \rho_{it-1}}{\rho_{it-1}}$, falls. The numbers in parentheses in the table represent the relative contribution of each term.¹⁶

Overall table 4 portrays a clear message. Households in the PSID reduce their expenditure growth significantly more than the slowdown in their disposable income alone would suggest (-6.9% v/s 2.9%). This implies that, overall, households increase their saving rate. However, the increase in saving rates, although present among all wealth quintiles, is quantitatively most potent for the first quintile, i.e. for those households entering the recession with the lowest net worth. Indeed, for these households the increase in the saving rate accounts for over two-thirds (69%) of the consumption growth decline, whereas for the other wealth groups consumption expenditure growth fell because both income growth slowed down and saving increased. We believe this fact is especially interesting, since it suggests that the decline in consumption at the bottom of the wealth distribution is not simply explained by standard hand-to-mouth behavior (i.e. the decline in income of these households), but primarily by changes in consumption behavior though a decline in expenditure rates.

Having documented the salient features of the joint wealth, income, and consumption distribution in the United States prior to the Great Recession and their dynamics over the course of the downturn, we now proceed with a quantitative evaluation of how well standard economic theory, in the form of the canonical heterogeneous household business cycle model with uninsurable idiosyncratic earnings risk, can explain these patterns. We then use this model as a quantitative laboratory to assess the importance of cross-sectional household heterogeneity for aggregate business cycles.

¹⁶ The relative contributions do not sum to 1 as the decomposition in 1 is not exact, and it excludes terms that involve the product of growth rates.

Table 4: Decomposing changes in expenditure growth

	Change C Growth	Change Y Growth	Change C/Y Growth
	$g_{c,t} - g_{c,t-1}$	$g_{y,t} - g_{y,t-1}$	$\frac{\rho_{it} - \rho_{it-1}}{\rho_{it-1}} - \frac{\rho_{it-1} - \rho_{it-2}}{\rho_{it-2}}$
All	-6.9	-2.9 (42%)	-3.8 (55%)
NW Q			
Q1	-6.5	-0.7 (11%)	-4.5 (69%)
Q2	-5.2	-2.6 (50%)	-2.3 (44%)
Q3	-9.0	-3.3 (37%)	-5.2 (58%)
Q4	-7.4	-3.3 (48%)	-3.8 (55%)
Q5	-6.2	-3.0 (42%)	-3.4 (55%)

3 A Canonical Business Cycle Model with Household Heterogeneity

In this section we lay out the benchmark model on which this chapter is built. The model is a slightly modified version of the original Krusell and Smith (1998) real business cycle model with household wealth and preference heterogeneity¹⁷ and shares many features of the model recently studied by Carroll, Slacalek, Tokunaka and White (2015).

3.1 Technology

In the spirit of real business cycle theory, aggregate shocks take the form of productivity shocks to the aggregate production function

$$Y = Z^* F(K, N). \quad (2)$$

Total factor productivity Z^* in turn is given by

$$Z^* = ZC^\omega, \quad (3)$$

where the exogenous part of technology Z follows a first-order Markov process with transition matrix $\pi(Z'|Z)$. Here C is aggregate consumption and the parameter $\omega \geq 0$ measures the importance of an aggregate demand externality. In the benchmark model, we consider the case of $\omega = 0$ in which case total factor productivity is exogenous and determined by the stochastic process for Z (and in which case we do not

¹⁷ Krusell and Smith (1998) in turn build on stationary versions of the model with household wealth heterogeneity, and thus on Bewley (1986), Imrohoroglu (1989), Huggett (1993, 1997) and Aiyagari (1994). See Deaton (1991) and Carroll (1992, 1997) for important early partial equilibrium treatments.

distinguish between Z and Z^*). In section 7 we consider a situation with $\omega > 0$. In that case current TFP and thus output is partially determined by demand (aggregate consumption).

In either case, in order to aid the interpretation of the results, we will mainly focus on a situation in which the exogenous technology Z can take two values, $Z \in Z_l, Z_h$. We then interpret Z_l as a severe recession and Z_h as normal economic times.

Finally, we assume that capital depreciates at a constant rate $\delta \in [0, 1]$.

3.2 Household Demographics, Endowments and Preferences

3.2.1 Demographics and the Life Cycle

In each period a measure 1 of potentially infinitely lived households populates the economy. Households are either young, working households (denoted by W) and participate in the labor market or are old and retired (and denoted by R). We denote a household's age by $j \in \{W, R\}$. Young households have a constant probability of retiring $1 - \theta \in [0, 1]$, and old households have a constant probability of dying $1 - \nu \in [0, 1]$. Deceased households are replaced by new young households. Given these assumptions, the distribution of the population across the two ages is given by

$$\begin{aligned}\Pi_W &= \frac{1 - \theta}{(1 - \theta) + (1 - \nu)} \\ \Pi_R &= \frac{1 - \nu}{(1 - \theta) + (1 - \nu)}.\end{aligned}$$

This simple structure captures the life cycle of households and thus their life-cycle savings behavior in a parsimonious way.

3.2.2 Preferences

Households do not value leisure, but have preferences defined over stochastic consumption streams, determined by a period utility function $u(c)$ with the standard concavity and differentiability properties, as well as a time discount factor β that may be heterogeneous across households (but is fixed over time for a given household). Denote by B the finite set of possible time discount factors.

3.2.3 Endowments

Since households do not value leisure in the utility function, young households supply their entire time endowment (which is normalized to 1) to the market. However, they

face idiosyncratic labor productivity and thus earnings risk. This earnings risk comes from two sources. First, households are subject to unemployment risk. We denote by $s \in S = \{u, e\}$ the current employment status of a household, with $s = u$ indicating unemployment. Employment follows a first-order Markov chain with transitions $\pi(s'|s, Z', Z)$ that depend on the aggregate state of the world. This permits the dependence of unemployment-employment transitions on the state of the aggregate business cycle.

In addition, conditional on being employed, a household's labor productivity $y \in Y$ is stochastic and follows a first order Markov chain; denote by $\pi(y'|y) > 0$ the conditional probability of transiting from state y today to y' tomorrow, and by $\Pi(y)$ the associated (unique) invariant distribution. In the benchmark model we assume that, conditional on being employed, transitions of labor productivity are independent of the aggregate state of the world.¹⁸

For both idiosyncratic shocks (s, y) we assume a law of large numbers, so that idiosyncratic risk averages out, and only aggregate risk determines the number of agents in a specific idiosyncratic state $(s, y) \in S \times Y$. Furthermore, we assume that the share of households in a given idiosyncratic employment state s only depends on the current aggregate state¹⁹ Z , and thus denote by $\Pi_Z(s)$ the deterministic fraction of households with idiosyncratic unemployment state s if the aggregate state of the economy is given by Z . We denote the cross-sectional distribution over labor productivity by $\Pi(y)$; by assumption this distribution does not depend on the aggregate state Z .

Households can save (but not borrow)²⁰ by accumulating (moderately risky) physical capital²¹ and have access to perfect annuity markets.²² We denote by $a \in A$ the asset holdings of an individual household and by A the set of all possible asset holdings. Households are born with zero initial wealth, draw their unemployment status according to $\Pi_Z(s)$ and their initial labor productivity from $\Pi(y)$. The cross-sectional population distribution of employment status s , labor productivity y , asset holdings a , and discount factors β is denoted as Φ and summarizes, together with the aggregate shock Z , the aggregate state of the economy at any given point in time.

¹⁸ Even for the unemployed, the potential labor productivity y evolves in the background and determines the productivity upon finding a job, as well as unemployment benefits while being unemployed, as described below.

¹⁹ This assumption imposes consistency restrictions on the transition matrix $\pi(s'|s, Z', Z)$. By assumption, the cross-sectional distribution over y is independent of Z to start with.

²⁰ We therefore abstract from uncollateralized household debt, as modeled in Chatterjee et al. (2007) and Livshits, MacGee and Tertilt (2007). Herkenhoff (2015) provides an investigation of the impact of increased access to consumer credit on the U.S. business cycle.

²¹ We therefore abstract from household portfolio choice. See Cocco, Gomes and Maenhout (2005) for the analysis of portfolio choice in a canonical partial equilibrium model with idiosyncratic risk, and Krusell and Smith (1997) and Storesletten, Telmer and Yaron (2007) for general equilibrium treatments.

²² Thus the capital of the deceased is used to pay an extra return on capital $\frac{1}{v}$ of the retired survivors.

3.3 Government Policy

3.3.1 Unemployment Insurance

The government implements a balanced budget unemployment insurance system whose size is parametrized by a replacement rate $\rho = \frac{b(y, Z, \Phi)}{w(Z, \Phi)y}$ that gives benefits b as a fraction of potential earnings wy of a household, with $\rho = 0$ signifying the absence of public social insurance against unemployment risk.²³ These benefits are paid to households in the unemployment state $s = u$ and financed by proportional taxes on labor earnings with tax rate $\tau(Z, \Phi)$. Taxes are levied on both labor earnings and unemployment benefits.

Recall that by assumption the number of unemployed $\Pi_Z(u)$ only depends on the current aggregate state. The budget constraint of the unemployment insurance system then reads as

$$\Pi_Z(u) \sum_y \Pi(y) b(y, Z, \Phi) = \tau(Z, \Phi) \left[\sum_y \Pi(y) [\Pi_Z(u) b(y, Z, \Phi) + (1 - \Pi_Z(u)) w(Z, \Phi) y] \right].$$

Exploiting the fact that $b(y, Z, \Phi) = \rho w(Z, \Phi) y$ and that the cross-sectional distribution over y is identical among the employed and unemployed we can simply write

$$\Pi_Z(u) \rho = \tau(Z, \Phi) [\Pi_Z(u) \rho + (1 - \Pi_Z(u))]$$

and conclude that the tax rate needed to balance the budget satisfies

$$\tau(Z, \Phi; \rho) = \left(\frac{\Pi_Z(u) \rho}{1 - \Pi_Z(u) + \Pi_Z(u) \rho} \right) = \left(\frac{1}{1 + \frac{1 - \Pi_Z(u)}{\Pi_Z(u) \rho}} \right) = \tau(Z; \rho) \in (0, 1). \quad (4)$$

That is, the tax rate $\tau(Z; \rho)$ only depends (positively) on the exogenous policy parameter ρ measuring the size of the unemployment system as well as (negatively) on the exogenous ratio of employed to unemployed $\frac{1 - \Pi_Z(u)}{\Pi_Z(u)}$, which in turn varies over the business cycle.

3.3.2 Social Security

The government runs a balanced budget PAYGO system whose size is determined by a constant payroll tax rate τ_{SS} (that applies only to labor earnings). Social security benefits $b_{SS}(Z, \Phi)$ of retirees are assumed to be independent of past contributions, but because of fluctuations in the aggregate tax base will vary with the aggregate state

²³ Recall that even unemployed households carry with them the idiosyncratic state y even though it does not affect their current labor earnings since they are unemployed.

of the economy Z . The budget constraint then determines the relationship between benefits and the tax rate according to

$$b_{SS}(Z, \Phi)\Pi_R = \tau_{SS}\Pi_W \left[\sum_y \Pi(y) (1 - \Pi_Z(u)) w(Z, \Phi)y \right],$$

Note that in the absence of unemployment (and with average labor productivity of working people equal to 1), we have

$$\tau_{SS} = \frac{b_{SS}(Z, \Phi) \Pi_R}{w(Z, \Phi) \Pi_W}$$

In this case the social security tax rate is simply equal to the average replacement rate $\frac{b_{SS}(Z, \Phi)}{w(Z, \Phi)}$ times the old age dependency ratio $\frac{\Pi_R}{\Pi_W}$.

3.4 Recursive Competitive Equilibrium

As is wellknown, the state space in this economy includes the entire cross-sectional distribution Φ of individual characteristics,²⁴ (j, s, y, a, β) . Since the dynamic programming problems of young, working age households and retired households differ significantly from each other (in terms of both individual state variables as well the budget constraint) it makes notation easier to separate age $j \in \{W, R\}$ from the other state variables. The dynamic programming problem of retired households then reads as

$$v_R(a, \beta; Z, \Phi) = \max_{c, a' \geq 0} \left\{ u(c) + v\beta \sum_{Z' \in Z} \pi(Z'|Z) v_R(a', \beta; Z', \Phi') \right\}$$

subject to

$$\begin{aligned} c + a' &= b_{SS}(Z, \Phi) + (1 + r(Z, \Phi) - \delta)a/v \\ \Phi' &= H(Z, \Phi', Z') \end{aligned}$$

For working household households, the decision problem is given by

$$v_W(s, y, a, \beta; Z, \Phi) = \left\{ \max_{c, a' \geq 0} u(c) + \beta \sum_{(Z', s', y') \in (Z, S, Y)} \pi(Z'|Z) \pi(s'|s, Z', Z) \pi(y'|y) \right\}$$

²⁴ In order to make the computation of a recursive competitive equilibrium feasible, we follow Krusell and Smith (1998), and many others since, and define and characterize a recursive competitive equilibrium with boundedly rational households who use only a small number of moments (and concretely here, just the mean) of the wealth distribution to forecast future prices. For a discussion of the various alternatives in computing equilibria in this class of models, see the January 2010 special issue of the *Journal of Economic Dynamics and Control*.

$$\times [\theta v_W(s', y', a', \beta; Z', \Phi') + (1 - \theta)v_R(a', \beta; Z', \Phi')]$$

subject to

$$\begin{aligned} c + a' &= (1 - \tau(Z; \rho) - \tau_{SS})w(Z, \Phi)y[1 - (1 - \rho)1_{s=u}] + (1 + r(Z, \Phi) - \delta)a \\ \Phi' &= H(Z, \Phi', Z'), \end{aligned}$$

where $1_{s=u}$ is the indicator function that takes the value 1 if the household is unemployed, and thus labor earnings equal unemployment benefits $b(y, Z, \Phi) = \rho w(Z, \Phi)y$.

Definition 1 *A recursive competitive equilibrium is given by value and policy functions of working and retired households, v_j, c_j, a'_j , pricing functions r, w , and an aggregate law of motion H such that*

1. *Given the pricing functions r, w , the tax rate given in equation (4), and the aggregate law of motion H , the value function v solves the household Bellman equation above and c, a' are the associated policy functions.*
2. *Factor prices are given by*

$$\begin{aligned} w(Z, \Phi) &= ZF_N(K(Z, \Phi), N(Z, \Phi)) \\ r(Z, \Phi) &= ZF_K(K(Z, \Phi), N(Z, \Phi)). \end{aligned}$$

3. *Budget balance in the unemployment system: equation (4) is satisfied*
4. *Market clearing*

$$\begin{aligned} N(Z, \Phi) &= (1 - \Pi_Z(u)) \sum_{y \in Y} y \Pi(y) \\ K(Z, \Phi) &= \int ad\Phi. \end{aligned}$$

5. *The aggregate law of motion H is induced by the exogenous stochastic processes for idiosyncratic and aggregate risk as well as the optimal policy function a' for assets.²⁵*

3.5 A Taxonomy of Different Versions of the Model

Table 5 summarizes the different versions of the model we will study in this chapter, including the section of the chapter in which it will appear. We start with a version of the model in which total factor productivity is exogenous. The only source of propagation of the aggregate shocks is the capital stock, which is predetermined in the short run

²⁵ We give the explicit statement of the law of motion H in appendix B.

(and thus output is exogenous), but responds in the medium run to technology shocks and/or reforms of the social insurance system. We study two versions of the model, the original Krusell-Smith (1998) economy without preference heterogeneity (which we will alternatively refer to as the KS economy, the low–wealth inequality economy, or the homogeneous discount factor economy), and a model with permanent discount factor heterogeneity (which we refer to as the high–wealth inequality economy, the heterogeneous discount factor economy, or simply the benchmark economy). The latter economy also features an unemployment insurance system whose size is consistent with U.S. data. In section 5.1 we discuss the extent to which both versions of this model match the empirically observed U.S. cross-sectional wealth distribution, and in section 6.1 we trace out the model-implied aggregate consumption, investment, and output dynamics in response to a Great Recession type shock.

Table 5: Taxonomy of Different Versions of the Model Used in the Paper

Name	Discounting	Techn.	Soc. Ins.	Section
KS	$\beta = \bar{\beta}$	$\omega = 0$	$\rho = 1\%$	Sec. 6.1
Het. β	$\beta \in [\bar{\beta} - \epsilon \bar{\beta} + \epsilon]$	$\omega = 0$	$\rho = 50\%$	Sec. 6.1
Het. β	$\beta \in [\bar{\beta} - \epsilon \bar{\beta} + \epsilon]$	$\omega = 0$	$\rho = 10\%$	Sec. 6.3
Dem. Ext.	$\beta \in [\bar{\beta} - \epsilon \bar{\beta} + \epsilon]$	$\omega > 0$	$\rho = 50\%$	Sec. 7

In order to assess the interaction of wealth inequality and social insurance policies for aggregate macro dynamics, in section 6.3 we study a version of the heterogeneous discount factor economy with smaller unemployment insurance. In section 7 the assumption of exogenous TFP is relaxed, and we present a version of the model in which TFP and thus output is partially demand-determined. In this version of the model, household heterogeneity has a potential impact not only on the size of the consumption recession, but also on the magnitude of the output decline, and by stabilizing individual consumption demand, unemployment insurance may act as a quantitatively important source of macroeconomic stabilization.

4 Calibration of the Benchmark Economy

In this section we describe how we map our economy to the data. Since we want to address business cycles and transitions into and out of unemployment, we calibrate the model to *quarterly* data.

4.1 Technology and Aggregate Productivity Risk

Following Krusell and Smith (1998), we assume that output is produced according to a Cobb-Douglas production function

$$Y = ZK^\alpha N^{1-\alpha}. \quad (5)$$

We set the capital share to $\alpha = 36\%$ and assume a depreciation rate of $\delta = 2.5\%$ per quarter. For the aggregate technology process, we assume that aggregate productivity Z can take two values $Z \in \{Z_l, Z_h\}$, where we interpret Z_l as a potentially severe recession. The aggregate technology process is assumed to follow a first-order Markov chain with transitions

$$\pi = \begin{pmatrix} \rho_l & 1 - \rho_l \\ 1 - \rho_h & \rho_h \end{pmatrix}.$$

The stationary distribution associated with this Markov chain satisfies

$$\begin{aligned} \Pi_l &= \frac{1 - \rho_h}{2 - \rho_l - \rho_h} \\ \Pi_h &= \frac{1 - \rho_l}{2 - \rho_l - \rho_h} \end{aligned}$$

With the normalization that $E(Z) = 1$, the aggregate productivity process is fully determined by the two persistence parameters ρ_l, ρ_h and the dispersion of aggregate productivity, as measured by Z_l/Z_h .

For the calibration of the aggregate productivity process, we think of a $Z = Z_l$ realization as a severe recession such as the Great Recession or the double-dip recession of the early 1980s (and a realization of $Z = Z_h$ as normal times). In this interpretation of the model, by choice of the parameters $\rho_l, \rho_h, Z_l/Z_h$ we want the model to be consistent with the fraction of time periods spent in severe recessions, their expected length conditional on slipping into one, and the decline in GDP per capita associated with severe recessions.²⁶

For this we note that with the productivity process set out above, the fraction of time spent in severe recessions is Π_l , whereas, conditional on falling into one, the expected length is given by

$$EL_l = 1 \times 1 - \rho_l + 2 \times \rho_l (1 - \rho_l) + \dots = \frac{1}{1 - \rho_l}. \quad (6)$$

²⁶ This chapter shares the focus on rare but large economic crises with the body of work on rare disasters, see e.g. Rietz (1988), Barro (2006) and Gourio (2013).

This suggests the following calibration strategy:

1. Choose ρ_l to match the average length of a severe recession EL_l . This is a measure of the persistence of recessions.
2. Given ρ_l , choose ρ_h to match the fraction of time the economy is in a severe recession, Π_l .
3. Choose $\frac{Z_l}{Z_h}$ to match the decline in GDP per capita in severe recessions relative to normal times.

In order to measure the empirical counterparts of these entities in the data, we need an operational definition of a severe recession. This definition could be based on GDP per capita, total factor productivity, or unemployment rates, given the model assumption that the aggregate unemployment rate $\Pi_Z(y_u)$ is only a function of the aggregate state of the economy Z .

We chose the latter and define a severe recession to be one where the unemployment rate rises above 9% at least for one quarter and determine the length of the recession to be the period for which the unemployment rate remains above 7%. Using this definition over period from 1948.I to 2014.III we identify two severe recession periods: from 1980.II to 1986.II and from 2009.I to 2013.III. This delivers a frequency of severe recessions of $\Pi_l = 16.48\%$ with expected length of 22 quarters. The average unemployment rate in these severe recession periods is $u(Z_l) = 8.39\%$ and the average unemployment rate in normal times is $u(Z_h) = 5.33\%$. The implied Markov transition matrix that delivers this frequency and length of severe recessions has $\rho_l = 0.9545$ and $\rho_h = 0.9910$ and thus is given by

$$\pi = \begin{pmatrix} 0.9545 & 0.0455 \\ 0.0090 & 0.9910 \end{pmatrix}.$$

For the ratio $\frac{Z_l}{Z_h}$ we target a value of $\frac{Y_l}{Y_h} = 0.9298$, that is, a drop in GDP per capita of 7% relative to normal times.²⁷ With average labor productivity if employed equal to 1 and if unemployed equal to zero, unemployment rates in normal and recession states equal to $u(Z_l) = 8.39\%$ and $u(Z_h) = 5.33\%$, and a capital share $\alpha = 0.36$, this requires $\frac{Z_l}{Z_h} = 0.9614$, which, together with the normalization

$$Z_l\Pi_l + Z_h\Pi_h = 1.$$

determines the levels of Z as $Z_l = 0.9676$, $Z_h = 1.0064$. Note that because of endogenous dynamics of the capital stock which falls significantly during the recession, the

²⁷ This is the decline in real GDP per capita during the two recession periods we identified, after GDP per capita is linearly detrended. The exact magnitude of the real GDP per capita decline is not crucial for our results, but it is important that severe recessions are deeper and (especially) more persistent than regular business cycle fluctuations.

dispersion in total factor productivity is smaller than what would be needed to engineer a drop in output by 7% only through TFP and increased unemployment (which is the drop in output on impact, given that the capital stock is predetermined).²⁸

4.2 Idiosyncratic Earnings Risk

Recall that households face two types of idiosyncratic risks: countercyclical unemployment risk described by the transition matrices $\pi(s'|s, Z', Z)$ and, conditional on being employed, acyclical earnings risk determined by $\pi(y'|y)$. We describe both components in turn.

4.2.1 Unemployment Risk

Idiosyncratic unemployment risk is completely determined by the four 2 by 2 transition matrices $\pi(s'|s, Z', Z)$ summarizing the probabilities of transiting in and out of unemployment for each (Z, Z') combination. Thus $\pi(s'|s, Z', Z)$ has the form

$$\begin{bmatrix} \pi_{u,u}^{Z,Z'} & \pi_{u,e}^{Z,Z'} \\ \pi_{e,u}^{Z,Z'} & \pi_{e,e}^{Z,Z'} \end{bmatrix}, \quad (7)$$

where, for example, $\pi_{e,u}^{Z,Z'}$ is the probability that an unemployed individual finds a job between one period and the next, when aggregate productivity transits from Z to Z' . Evidently each row of this matrix has to sum to 1. Note that, in addition, the restriction that the aggregate unemployment rate only depends on the aggregate state of the economy imposes one additional restriction on each of these 2 by 2 matrices, of the form

$$\Pi_{Z'}(u) = \pi_{u,u}^{Z,Z'} \times \Pi_Z(u) + \pi_{e,u}^{Z,Z'} \times (1 - \Pi_Z(u)). \quad (8)$$

Thus, conditional on targeted unemployment rates in recessions and expansions, (Π_l, Π_h) this equation imposes a joint restriction on $(\pi_{u,u}^{Z,Z'}, \pi_{e,u}^{Z,Z'})$ for each (Z, Z') pair. With

²⁸ In the short run,

$$\frac{Y_l}{Y_h} = \frac{Z_l}{Z_h} \left(\frac{1 - u(Z_l)}{1 - u(Z_h)} \right)^{0.64}$$

so that in order to generate a drop in output of 7% in the short run would require

$$\frac{Z_l}{Z_h} = \frac{0.9298}{\left(\frac{0.9161}{0.9467} \right)^{0.64}} = 0.9496.$$

these restrictions, the idiosyncratic transition matrices are uniquely pinned down by $\pi_{u,e}^{Z,Z'}$, i.e. the job-finding rates.²⁹

We compute the job finding rate for a quarter as follows. We consider an individual that starts the quarter as unemployed and compute the probability that at the end of the quarter that individual is still unemployed. The possible ways that this can happen are (denoting as f_1, f_2, f_3 and as s_1, s_2, s_3 the job-finding and job-separation rates in months 1, 2, and 3 of the quarter):

1. Does not find a job in month 1, 2, or 3, with probability $(1 - f_1) \times (1 - f_2) \times (1 - f_3)$.
2. Finds a job in month 1, loses it in month 2, does not find in month 3, with probability $f_1 \times s_2 \times (1 - f_3)$.
3. Finds a job in month 1, keeps it in month 2, loses it in month 3, with probability $f_1 \times (1 - s_2) \times s_3$.
4. Finds a job in month 2, loses it in month 3, with probability $(1 - f_1) \times f_2 \times s_3$.

Thus the probability that someone that was unemployed at the beginning of the quarter is not unemployed at the end of the quarter is:

$$f = 1 - ((1 - f_1)(1 - f_2)(1 - f_3) + f_1 s_2 (1 - f_3) + f_1 (1 - s_2) s_3 + (1 - f_1) f_2 s_3) \quad (9)$$

We follow Shimer (2005) to measure the job-finding and separation rates from CPS data as averages for periods corresponding to specific Z, Z' transitions.³⁰ Equating these with $\pi_{u,e}^{Z,Z'}$ delivers the following employment-unemployment transition matrices:

- Aggregate economy is and remains in a recession: $Z = Z_l, Z' = Z_l$

$$\begin{pmatrix} 0.3378 & 0.6622 \\ 0.0606 & 0.9394 \end{pmatrix} \quad (10)$$

²⁹ One could alternatively use job-separation rates $\pi_{e,u}^{Z,Z'}$.

³⁰ Let u_t = unemployment rate and u_t^S = short-term unemployment rate (people who are unemployed this month, but were not unemployed last month). Then we can define the monthly job-finding rate as $1 - (u_{t+1} - u_{t+1}^S)/u_t$ and the separation rate as $u_{t+1}^S/(1 - u_t)$. The series we use from the CPS are the unemployment level (UNEMPLOY), the short-term unemployment level (UNEMPLT5) and civilian employment (CE16OV). There was a change in CPS coding starting in February 1994 (inclusive), so UNEMPLT5 in every month starting with February 1994 is replaced by $UEMPL5 \times 1.1549$.

- Aggregate economy is and remains in normal times: $Z = Z_h \cdot Z' = Z_h$

$$\begin{pmatrix} 0.1890 & 0.8110 \\ 0.0457 & 0.9543 \end{pmatrix} \quad (11)$$

- Aggregate economy slips into recession: $Z = Z_h \cdot Z' = Z_l$

$$\begin{pmatrix} 0.3382 & 0.6618 \\ 0.0696 & 0.9304 \end{pmatrix} \quad (12)$$

- Aggregate economy emerges from recession: $Z = Z_l \cdot Z' = Z_h$

$$\begin{pmatrix} 0.2220 & 0.7780 \\ 0.0378 & 0.9622 \end{pmatrix} \quad (13)$$

We observe that the resulting matrices make intuitive sense. One possible (but quantitatively minor) exception is that the job-finding rate is higher if the economy remains in normal times than if it emerges from a recession. On the other hand, the lower job-finding rate is consistent with the experience during the Great Recession per our definition, as job-finding rates did not recover until well into 2014, whereas by our calibration the recession ended in 2013.

4.2.2 Earnings Risk Conditional on Employment

In addition to unemployment risk, we add to the model earnings risk, conditional on being employed. This allows us to obtain a more empirically plausible earnings distribution and makes earnings risk a more potent determinant of wealth dispersion (and thus reduces the importance of preference heterogeneity for this purpose). We assume that, conditional on being employed, log-labor earnings of households follow a process with both transitory and persistent shocks.³¹ The process is specified as

$$\log(y') = p + \epsilon \quad (14)$$

$$p' = \phi p + \eta \quad (15)$$

with persistence ϕ and innovations of the persistent and transitory shocks (η, ϵ) , respectively.³² The associated variances of the shocks are denoted by $(\sigma_\eta^2, \sigma_\epsilon^2)$, and therefore the entire process is characterized by the parameters $(\phi, \sigma_\eta^2, \sigma_\epsilon^2)$. We estimate

³¹ The formulation of log-earnings or log-income as a stochastic process with transitory and persistent (or fully permanent) shocks follows a large empirical literature in labor economics. See Meghir and Pistaferri (2004), Storesletten, Telmer and Yaron (2004a), Guvenen (2009) and the many references discussed therein.

³² Note that we assume that the variance and persistence of this process are independent of the state of the business cycle. Earnings risk in the data *is* countercyclical, as stressed by Storesletten, Telmer and Yaron (2004b, 2007) and Guvenen, Ozkan and Song (2014); in our benchmark model earning risk is also countercyclical but only because of countercyclical unemployment risk.

this process for household labor earnings after taxes (after first removing age, education and time effects) from *annual* PSID data and find estimates of $\phi, \sigma_{\eta}^2, \sigma_{\epsilon}^2$ equal to 0.9695, 0.0384 and 0.0522 respectively.³³ Next we translate these estimates into a quarterly persistence and variance.³⁴ We then use the Rouwenhorst procedure to discretize the persistent part of the process into a seven-state Markov chain.³⁵ The *iid* shock only enters the computation of the expectation on the right-hand side of the Euler equation.³⁶ We approximate the integral calculating the expectation using a Gauss-Hermite quadrature scheme with three nodes. Thus, we effectively approximate the continuous state space process by a discrete Markov chain with $7 \times 3 = 21$ states.³⁷

4.3 Preferences and the Life Cycle

In the benchmark economy we assume that the period utility function over current consumption is given by a constant relative risk aversion utility function with parameter $\sigma = 1$. As described above, we study two versions of the model: the original Krusell-Smith (1998) economy in which households have identical time discount factors, and a model in which households, as in Carroll, Slacalek, Tokuoka and White (2015) have permanently different time discount factors (and die with positive probability, in order to ensure a bounded wealth distribution).

For the model with preference heterogeneity, we assume that households at the beginning of their life draw their permanent discount factor β from a uniform distribu-

³³ For the exact definition of the labor earnings after taxes, sample selection criteria and estimation method, please see appendix A.

³⁴ In order to ensure that quarterly log-earnings has the same persistence as annual log-earnings, we choose the persistence of the quarterly AR(1) to be $\phi = \hat{\phi}^{\frac{1}{4}}$. For the variances, we note that the main purpose of the earnings shocks is to help deliver a plausible cross-sectional distribution of labor income. Therefore we aim to maintain the same cross-sectional distribution of earnings at the quarterly frequency as we estimate at the annual frequency. Choosing a quarterly transitory variance equal to its annual counterpart and

$$\frac{\sigma_{\eta}^2}{1 - \phi^2} = \frac{\hat{\sigma}_{\eta}^2}{1 - \hat{\phi}^2}$$

achieves this goal.

³⁵ See Kopecky and Suen (2010) for a detailed description and evaluation of the Rouwenhorst method.

³⁶ This is because we use cash at hand and the persistent income state as state variables in the individual household dynamic programming problem.

³⁷ For the computation of the distributional statistics we simulate a panel of households. In this simulation, realizations of the persistent shock remain on the grid, but the transitory shock is drawn from a normal distribution and thus is not restricted to fall on one of the quadrature points.

tion³⁸ with support $[\bar{\beta} - \epsilon, \bar{\beta} + \epsilon]$ and choose $(\bar{\beta}, \epsilon)$ so that the model wealth distribution (with an unemployment insurance replacement rate of 50%) has a Gini coefficient for the working age population of 77% as in the data and a quarterly wealth-to-output ratio of 10.26 (as in Carroll, Slacalek, Tokuoka and White, 2015) This requires $(\bar{\beta} = 0.9864, \epsilon = 0.0053)$ and implies that annual time discount factors in this economy range from $\beta = 0.9265$ to $\beta = 0.9672$. Finally, households in the working stage of their life cycle face a constant probability $1 - \theta$ of retiring, and retired households face a constant probability $1 - \nu$ of dying. For our quarterly model we choose $1 - \theta = 1/160$, implying an expected work life of 40 years, and $1 - \nu = 1/60$, with a resulting retirement phase of 15 years in expectation.

For the original Krusell-Smith economy, we choose the common quarterly discount factor $\beta = 0.9899$ to ensure that the capital-output ratio in this economy (again at quarterly frequency) equals that in the heterogeneous β economy. In this economy households neither retire nor die.

4.4 Government Unemployment Insurance Policy

The size of the social insurance (or unemployment insurance, more concretely) system is determined by the replacement rate ρ . For the benchmark economy that we assume $\rho = 50\%$ (see, e.g. Gruber (1994)). We will also consider a lower value of $\rho = 10\%$, motivated by the observation that many households qualifying for unemployment insurance benefits fail to claim them (see, e.g. Blank and Card, 1991 or Chodorow-Reich and Karabarbounis, 2016).

Finally, the payroll tax rate for social security is set to $\tau_{SS} = 15.3\%$. This choice implies an average (over the business cycle) and empirically plausible replacement rate of the social security system of approximately 40%. In the KS economy, in order to avoid numerical problems with zero consumption, we include a minimal unemployment insurance system with a replacement rate of $\rho = 1\%$.

³⁸ In practice we discretize this distribution and assume that each household draws one of five possible β 's with equal probability; thus $B = \{\beta_1, \dots, \beta_5\}$ and $\Pi(\beta) = 1/5$. We also experimented with stochastic β 's as in Krusell and Smith (1998) but found that the formulation we adopt enhances the model's ability to generate sufficiently many wealth-poor households. The results for the stochastic β economy generally lie in between those obtained in the original Krusell and Smith (1998) economy documented in detail in this chapter, and the results obtained in the model with permanent β heterogeneity, also documented in great detail below.

5 Evaluating the Benchmark Economy

5.1 The Joint Distribution of Earnings, Income, Wealth and Consumption in the Benchmark Economy

In this section we evaluate the extent to which our benchmark model is consistent with the main empirical facts characterizing the joint distribution of wealth, income, and consumption expenditures, as well as the changes in this distribution when the economy is subjected to a large negative aggregate shock.

5.1.1 Wealth Inequality in the Benchmark Economy

We have argued in the introduction that a model-implied cross-sectional wealth distribution that is consistent with the empirically observed concentration, and especially with a share of wealth of the bottom 40% of close to zero, is crucial when using the model as a laboratory for studying aggregate fluctuations. We now document that our benchmark economy has this property, whereas an economy akin to the one studied in Krusell and Smith's (1998) original work in which wealth inequality is entirely driven by idiosyncratic unemployment shocks and incomplete financial markets does not.³⁹

Table 6: Net Worth Distributions: Data v/s Models

% Share held by:	Data		Models	
	PSID, 06	SCF, 07	Bench	KS
Q1	-0.9	-0.2	0.3	6.9
Q2	0.8	1.2	1.2	11.7
Q3	4.4	4.6	4.7	16.0
Q4	13.0	11.9	16.0	22.3
Q5	82.7	82.5	77.8	43.0
90 – 95	13.7	11.1	17.9	10.5
95 – 99	22.8	25.3	26.0	11.8
T1%	30.9	33.5	14.2	5.0
Gini	0.77	0.78	0.77	0.35

³⁹ We retain *our* calibration of idiosyncratic unemployment risk, and thus the cross-sectional wealth distribution in our version of the Krusell-Smith economy differs from their original numbers, but not in a magnitude substantial enough to change any of the conclusions below.

Table 6 reports selected statics for the wealth distribution, those computed from the data (PSID and SCF) as well as those from two model economies, the original Krusell-Smith (1998) economy and our benchmark model with idiosyncratic income risk, incomplete markets, a rudimentary life cycle structure, unemployment insurance, and heterogeneous discount factors.⁴⁰ As indicated in the calibration section, through appropriate choice of the time discount factor(s), both economies have the same average (over the business cycle) capital-output ratio, and the benchmark economy displays a wealth Gini coefficient in line with the micro data from the PSID. All other moments of the empirical cross-sectional wealth distribution were not targeted in the calibration of the models.

From the table we note that, overall, the benchmark model fits the empirical wealth distribution in the data quite well (albeit not perfectly), especially at the bottom of the distribution. Specifically, it captures the fact that households constituting the bottom two quintiles of the wealth distribution hardly have any wealth, but also that the top wealth quintile holds approximately 80% of all net worth in the U.S. economy. We also acknowledge that the benchmark model makes the wealth upper middle class (quintile 4 and also the bottom part of quintile 5) somewhat too wealthy. For example, households between the 90th and 99th percentiles of the net worth distribution account for about 36% of wealth in the data, but 44% in the model. Most problematically, the benchmark model still misses the wealth concentration at the *very top* of the distribution significantly. In the data the top 1% wealth holders account for over 30% of overall net worth in the economy, whereas the corresponding figure in the model is only 14.0%. A histogram of the model-implied wealth distribution can be found in figure 10 below.⁴¹

Finally, table 6 reproduces the well-known—since Krusell and Smith (1998)—result that transitory unemployment risk and incomplete financial markets alone are incapable of generating sufficient wealth dispersion. The problem relative to the data is two-fold: households at the top of the wealth distribution are not nearly wealthy enough, and, as we will argue, more importantly for the results to follow, households at the bottom of the distribution hold significantly too much wealth in the model. Relative to SCF or PSID micro data, in the model the bottom 40% own about 19% of net worth in the economy, whereas in the data that share is approximately 0. As a summary measure of wealth inequality, whereas the wealth Gini in the data is well above 0.7, the original Krusell-Smith model delivers a number of only 0.35.

⁴⁰ Recall that in the data, we restrict attention to working-age households. Consequently, when we report cross-sectional statistics from the benchmark model (which includes a retirement phase), we restrict attention to households in the working stages of their life.

⁴¹ Although this is clearly a shortcoming, note that in this range of wealth levels, the consumption function is essentially linear (as we will display below) and thus mechanically reshuffling wealth between the top 1% and the top 20% through top 1% would not alter aggregate consumption responses to shocks significantly. We will return to this point in section 6.2.

In the next section we now decompose which model elements in the benchmark economy are responsible for generating a more realistic wealth distribution than in the original Krusell-Smith economy. We then turn to an evaluation of the benchmark model's success in reproducing the empirical *joint* distribution of earnings, income, consumption, and wealth in the data.

5.2 Inspecting the Mechanism I: What Accounts for Wealth Inequality in the Benchmark Economy?

A substantial literature, recently surveyed in De Nardi (2015), De Nardi, Fella and Yang (2015) and Benhabib and Bisin (2016), explores alternative mechanisms for generating the empirically observed high wealth concentration in the data.⁴² These mechanisms include the use of very large but transient income realizations that the PSID misses out on (as in Castaneda, Diaz-Gimenez and Rios-Rull, 2003, Kindermann and Krueger (2015) or Brueggemann and Yoo, 2015), large uninsured or only partially insured medical expenditure shocks in old age (see e.g. De Nardi, French and Jones, 2010, or Ameriks, Briggs, Caplin, Shapiro and Tonetti, 2015), the intergenerational transmission of wealth through accidental and intended bequests (as e.g. in De Nardi, 2004), the interaction between wealth accumulation and entrepreneurship (see Quadrini, 2000, Cagetti and De Nardi 2006 and Buera 2009) or idiosyncratic shocks to investment opportunities or its returns, as in Benhabib, Bisin and Zhu (2011).

In our benchmark model, we instead follow the sizeable literature that has explored the potential importance of empirically realistic, highly persistent earnings risk (conditional on employment) as well as preference heterogeneity in general, and cross-sectional dispersion in patience specifically, for generating an empirically plausible cross-sectional wealth distribution. Household heterogeneity in time discount factors had already been explored by the original Krusell and Smith (1998) paper, and has been further analyzed by Hendricks (2007) and Carroll, Slacalek, Tokunaka and White (2015); the latter also incorporates a stochastic earnings process in the analysis.

In the previous section, we argued that preference heterogeneity, when combined with idiosyncratic unemployment and earnings shocks as well as rudimentary life cycle elements⁴³ and social insurance policies, generates a wealth distribution that resembles the data in 2006 well, both at the bottom and at the top of the distribution. In table 7

⁴² Gabaix, Lasry, Lions and Moll (2014) evaluate whether the existing theories discussed there are consistent with the secular rise in the share of income and wealth accruing to the top 1% of households, and argue that only theories embedding "superstar" phenomena are capable of reproducing the facts at the very top of these distributions.

⁴³ The literature on quantitative studies of the cross-sectional wealth distributions in general equilibrium life-cycle economies with uninsurable idiosyncratic income risk starts with Huggett (1996).

we now show precisely which model elements are responsible for this finding.⁴⁴

Table 7: Net Worth Distributions and Consumption Decline: Different Versions of the Model

% Share:	Models*				
	KS	$+\sigma(y)$	+Ret.	$+\sigma(\beta)$	+UI
Q1	6.9	0.7	0.7	0.7	0.3
Q2	11.7	2.2	2.4	2.0	1.2
Q3	16.0	6.1	6.7	5.3	4.7
Q4	22.3	17.8	19.0	15.9	16.0
Q5	43.0	73.3	71.1	76.1	77.8
90 – 95	10.5	17.5	17.1	17.5	17.9
95 – 99	11.8	23.7	22.6	25.4	26.0
T1%	5.0	11.2	10.7	13.9	14.2
Wealth Gini	0.350	0.699	0.703	0.745	0.767

*The KS model only has unemployment risk and incomplete markets, and thus the first column repeats information from table 6. The column $+\sigma(y)$ adds idiosyncratic earnings shocks (transitory and permanent) while employed. The column +Ret. adds the basic life cycle structure (positive probability of retirement and positive probability of death, plus social security in retirement). The column $+\sigma(\beta)$ incorporates preference heterogeneity into the model, and finally the column +UI raises the replacement of the unemployment insurance system from 1% to 50%; the resulting model is therefore the benchmark model, with results already documented in table 6. In all models, the (mean) discount factor is calibrated so that all versions have the same capital-output ratio.

The table (which partially repeats information from table 6 to facilitate comparisons across different model economies) displays the share of net worth held by the five wealth quintiles, the wealth Gini, and more detailed information about the top of the net worth distribution, in the data and in a sequence of models, ranging from the original Krusell-Smith (1998) economy to our benchmark economy in the last column.

The table contains several important quantitative lessons. First, comparing the first and the second model columns, we see that the inclusion of highly persistent earnings risk, in addition to unemployment risk, increases wealth dispersion very significantly, relative to the economy with *only* unemployment risk. Consistent with a sizeable literature estimating stochastic labor earnings or income processes (see e.g. Storesletten, Telmer and Yaron, 2004a) we find that the persistent component is indeed very persistent, with an annual autocorrelation (conditional on remaining employed) of

⁴⁴ Castaneda, Diaz-Gimenez and Rios-Rull (1998) provide a decomposition similar in spirit, but focus on the evolution of the cross-sectional income distribution over the cycle.

0.97. Thus, the economy contains a share of households with close to permanently low earnings, even in the absence of unemployment. These households, located predominantly in the lowest wealth quintile, have had no opportunity to accumulate significant wealth.⁴⁵ Consequently the share of wealth held by the poorest households shrinks to fairly close to zero with idiosyncratic income risk, as observed in the data. At the same time, the top wealth quintile is populated with households with high earnings realizations for whom the risk of a persistent fall in earnings provides motivation to accumulate substantial wealth. As a result, the wealth Gini doubles in the economy with earnings risk, relative to the original Krusell-Smith unemployment-only model.

Second, adding a more explicit life-cycle structure does not change the wealth distribution (of the working-age population) much, but as we will see in the next section, will imply a more plausible *joint* wealth-consumption distribution, by adding a life-cycle savings for retirement motive to the precautionary saving motive. It also somewhat reduces wealth concentration at the top of the distribution, since earnings risk ceases with retirement and thus trims the precautionary motive of the wealth-rich.⁴⁶

Third, as the examination of the very top of the wealth distribution in the first three columns of table 7 reveals, income risk and life-cycle elements alone are insufficient to generate the very high wealth concentration observed in the data. This is where the discount factor heterogeneity in the benchmark model plays a crucial role. It creates a class of households that are patient and have a high propensity to save, and the fact that in addition to a precautionary saving motive, they also save for retirement (a phase they value highly because of their patience) ensures that they do not start to decumulate wealth even at high wealth levels. As table 7 displays (comparing the last two columns), the model with both features (the life cycle and preference heterogeneity) is able to generate the wealth concentration at the top quintile of the distribution close to what is observed in U.S. data (albeit not at the very top of the distribution).

Finally, inserting an unemployment insurance system into the model further reduces the wealth held by the bottom two quintiles of the distribution, since now losing a job with little net worth is not nearly as harmful. In Krueger, Mitman and Perri (2016) we argue that the size of the unemployment insurance system not only crucially shapes the bottom of the wealth distribution, but also has a strong impact on the welfare losses from severe recessions in the class of heterogeneous household macro models we study in this chapter.

⁴⁵ And if an unemployment insurance system with replacement rate of $\rho = 50\%$ is in place, as in the benchmark economy, they have no strong motive, either.

⁴⁶ Our model imposes substantial structure on the link between idiosyncratic income shocks and consumption over the life cycle. In methodologically complementary work, Arellano, Blundell and Bonhomme (2015) estimate a more flexible nonlinear empirical model of household earnings and consumption over the life cycle.

5.2.1 Income and Consumption at Different Points of the Wealth Distribution

In this section we evaluate the ability of the benchmark model to reproduce key features of the joint distribution of income, consumption, and wealth in the PSID data. To do so, table 8 reports the share of earnings, disposable income, consumption expenditures, and the expenditure rates for the five quintiles of the wealth distribution, both for the data (as already contained in table 2) and for the benchmark model.

Table 8: Selected Variables by Net Worth: Data v/s Models

NW Q	% Share of:						% Expend. Rate			
	Earnings		Disp Y		Expend.		Earnings		Disp Y	
	Data	Mod	Data	Mod	Data	Mod	Data	Mod	Data	Mod
Q1	9.8	6.5	8.7	6.0	11.3	6.6	95.1	96.5	90.0	90.4
Q2	12.9	11.8	11.2	10.5	12.4	11.3	79.3	90.3	76.4	86.9
Q3	18.0	18.2	16.7	16.6	16.8	16.6	77.5	86.0	69.8	81.1
Q4	22.3	25.5	22.1	24.3	22.4	23.6	82.3	87.3	69.6	78.5
Q5	37.0	38.0	41.2	42.7	37.2	42.0	83.0	104.5	62.5	79.6
	Correlation with net worth									
	0.26	0.46	0.42	0.67	0.20	0.76				

On the positive side, first, the model is consistent with the significantly positive correlation between net worth on the one hand, and earnings, disposable income and consumption expenditures on the other. The shares of the latter three variables are all increasing with the net worth quintiles. Second, as in the data, disposable income (which includes capital income) displays a higher correlation with net worth than with labor earnings. Third, the model reproduces the crucial fact that the bottom two wealth quintiles, while accounting for essentially zero net worth, contribute a very significant share to aggregate consumption expenditures. In the data, that share is 23.7%, and in the model it is still highly significant at 17.9%. Since, as we will show below, this low-wealth group has the largest declines in their consumption, the fact that it accounts for a substantial part of aggregate consumption to start with is in turn crucial for the macro responses to an aggregate shock in the model. Fourth, turning to the consumption expenditure rates, the model is broadly consistent with the levels found in the data, and is broadly consistent with the empirical finding in the data that these rates decline with net worth. However, the wealth gradient is not quite as steep in the model as it is in the data, and in the model the top wealth quintile has expenditure rates that are higher than the fourth quintile (very slightly so in relation to disposable income, much more strongly so in relation to labor earnings).

For this last finding, the inclusion of a retirement phase and thus a life cycle savings motive into the model is absolutely crucial. A pure infinite horizon version of the model, even with idiosyncratic income shocks and preference heterogeneity, dis-

plays expenditure rates that are significantly too high—averaging 100% across wealth quintiles—and implies expenditure rates that are U-shaped with respect to net worth. Absent the life-cycle savings motive, households accumulate wealth exclusively for the purpose of smoothing out negative income fluctuations, and thus individuals in the fourth and fifth wealth quintiles, having accumulated enough net worth for this purpose, display very high expenditure rates (in fact, significantly larger than 100%)—especially with respect to labor earnings. Preference heterogeneity mitigates this effect somewhat, but the resulting model still displays grossly counterfactual expenditure rates, whereas the version of the model with stochastic retirement brings the implications of the model much closer to their empirical counterpart, and is our primary justification for the presence of this model element.

We would also like to flag another dimension along which the model is not fully successful in capturing the empirical facts. First, although the model does generate consumption expenditure shares that are strongly increasing with wealth, not only are the wealth-poor too consumption-poor in the model (as already discussed above), but also the wealth rich (quintile 5) consume too much in the model (42% relative to 37.2% in the data). This is true even though the model captures the earnings and income share of this group of households quite well. This problem of the model is summarized by the fact that the correlation between net worth and consumption expenditures is positive in the model, as it is in the data, but is much larger than it is in the PSID.

We conclude this section with the overall assessment that the benchmark model captures well many qualitative features of the cross-sectional joint distribution of net worth, earnings, income and consumption expenditures, but fails to quantitatively match the joint distribution of net worth and expenditures, with the wealth poor consuming too little, and the wealth rich consuming too much, relative to the data.

5.3 The Dynamics of Income, Consumption and Wealth in Normal Times and in a Recession

The previous section studied the joint distribution of the key economic variables at a given point in time (2006 in the data, a period after a long sequence of normal macroeconomic performance in the model). We now put the model to a more ambitious (and to our knowledge novel) test and assess whether the *dynamics* of wealth, income and consumption implied by the model can match those observed in the data. We ask this question both for a period of macroeconomic stability (in section 5.3.1), and then, in section 5.3.2, for a period characterized by a severe macroeconomic crisis. Note that none of the empirical moments along this dimension were targeted in the calibration of the model.

5.3.1 Normal Times: 2004-2006

In the data we are somewhat limited in our choices by the sparse time series dimension of the PSID (for which comprehensive consumption data are available). We take *normal times* in the data to be the period from 2004 to 2006; we map this period into the model by studying an episode of eight quarters of good productivity, $Z = Z_h$, which in turn followed a long sequence of good aggregate shocks so that aggregates and distributions have settled down prior to this episode.

Table 9: Annualized Changes in Selected Variables by Net Worth in Normal Times (2004-2006): Data v/s Model

NW Q	Net Worth (%)		Disp Y (%)		Expend (%)		Exp. Rate (pp)	
	Data	Model	Data	Model	Data	Model	Data	Model
Q1	NaN	44	7.4	7.2	7.1	6.7	-0.2	-0.4
Q2	122	33	6.7	3.1	7.2	3.6	0.3	0.5
Q3	33	20	5.1	1.6	9	2.5	2.3	0.8
Q4	17	9	5	0.5	5.9	1.7	0.5	1.2
Q5	12	3	1.8	-1.0	2.7	0.5	0.5	1.4
All	16	5	4.1	0.7	5.6	1.8	0.9	0.7

Table 9 reports the statistics for the data (and thus repeats the information from table 3 above) together with the model.⁴⁷ Recall from the description of table 9 that for a given variable x (wealth, income, and consumption) and each wealth quintile we compute the quintile average for x in 2004 and the average x for the *same* households⁴⁸ in 2006 and then report the annualized percentage difference between the two figures. For the expenditure rates, which are already in percentage units, we compute the annualized percentage point differences between 2004 and 2006.

For net worth, the model captures the fact that in good economic times, wealth-poor households accumulate wealth at a faster rate than wealth-poor households. The percentage increase in wealth for all groups is lower in the model than in the data. We should note that in the data, the 2004-2006 period was one of rapid appreciation of house prices and financial asset valuations, whereas in our model the relative price of wealth (capital) is constant at one, and thus an increase in net worth during normal times in the model has to come from net capital accumulation of households.⁴⁹

⁴⁷ Since in tables 9 and 10 the statistics for earnings and disposable income are quite similar, we only report those for disposable income.

⁴⁸ These households would typically not be in the same wealth quintile in 2006 as they were in 2004.

⁴⁹ In a model without retirement and thus without life-cycle saving, generating positive changes in net worth for *all* wealth quintiles is of course very difficult; justifying again the inclusion of a basic life cycle element into the economy.

In terms of earnings (not reported) and disposable income, the model displays the substantial mean reversion built into the estimates of the idiosyncratic unemployment and earnings process, with income of the lowest wealth quintile rising fast (7.2%) and income of the highest wealth group actually falling (by 1%) even though aggregate incomes do not. This is because low wealth households tend to be low labor earnings and thus low income households with income. As we saw earlier, this is qualitatively consistent with the data, but quantitatively the model implies differences in income growth between the top and the bottom of the wealth distribution that are too large. In other words, the model implies slightly too much downward and upward mobility in incomes when households are ranked by wealth.⁵⁰

Finally, for changes in consumption expenditures, table 10 reveals that during normal times, as in the data (and as for disposable income), consumption growth is strongest at the low end of the wealth distribution. The wealth gradient of the consumption growth rates (again, as for disposable income), is somewhat steeper in the model. As in the data, the expansion of consumption for households in the lowest (in 2004) wealth quintile falls short of their income growth and thus the expenditure rate of this group falls during normal times. The opposite is true for the wealthiest group of households in the population: as in the data, the expenditure rate of this group expands as the macro economy remains in normal times. The reason for this differential behavior in expenditure rates between the wealth-poor and the wealth-rich is intuitive from the perspective of the model: low wealth households have had, on average, unfortunate earnings realizations and their wealth is below their target wealth. Therefore, these households cut their expenditure to re-build their wealth buffers. The opposite logic applies to households at the top of the wealth distribution. This implication of the model matches the data, although quantitatively, the difference in changes in expenditure rates between the top and the bottom wealth quintiles is a bit larger in the model than in the data.

5.3.2 A Great Recession

After documenting the dynamics of wealth, income, and consumption (ordered by wealth) in normal times, table 10 displays the same *model* statistics during a period in which the macro economy undergoes a large recession, induced by a transition of aggregate TFP from $Z = Z_h$ to $Z = Z_l$.⁵¹ To facilitate comparisons between the two tables, we display the difference in the growth rates between the recession period and normal times in table 11.

⁵⁰ Ranking households by earnings or income would make this statement even stronger.

⁵¹ In the model the Great Recession hits in Q.I, 2009, consistent with our calibration. In that quarter, Z switches from $Z = Z_h$ to $Z = Z_l$ and remains there until Q.III, 2013. The statistics are based on comparing the average of the four 2010 quarters to the average of the four 2008 quarters. In the data, as discussed in section 2, we consider the period from 2006 to 2010 because of the timing of the income and consumption data. Note that in the data, changes are all annualized.

Again, first focusing on net worth, the key endogenous state variable in our model that underlies the dynamics of all other economic variables, we observe that as in normal times (as in the data), the growth rate of net worth is declining in the level of net worth. And as in the data, the Great Recession significantly slows down the pace of wealth accumulation across all quintiles, and turns it negative for the wealthiest households, although the reduction predicted by the model is smaller than in the data. In the model, the wealth of the top net worth quintile declines by 1%, relative to the 3% growth in normal times. For the same quintile, annual wealth growth in the data slows down from 12% to -5% over a two-year period. As discussed above, in the data a large part of this reduction in wealth at the top of the distribution is likely the consequence of asset *price* movements which are, by construction, absent in the one-asset model studied here.⁵²

Table 10: Annualized Changes in Selected Variables by Net Worth in a Severe Recession: Data v/s Model

NW Q	Net Worth (%)		Disp Y (%)		Expend.(%)		Exp. Rate (pp)	
	Data	Model	Data	Model	Data	Model	Data	Model
Q1	NaN	24	6.7	4.9	0.6	4.5	-4.2	-0.4
Q2	24	15	4.1	0.3	2.0	1.2	-1.3	0.8
Q3	4	8	1.8	-2.4	0.8	0.0	-1.1	2.2
Q4	2	4	1.7	-4.0	-1.7	-1.5	-2.0	3.2
Q5	-5	-1	-1.2	-6.4	-3.7	-3.5	-1.4	4.6
All	-3	1	1.2	-3.7	-1.3	-0.8	-1.6	2.0

The two other empirical facts we have documented in section 2.3 were that income declines in the recession hit the top wealth quintiles more than the bottom quintiles, and that households in the bottom quintiles cut expenditure rates more than households in the top quintiles. Comparing disposable income growth rates in tables 9 and 10, we observe that the first fact is captured well by the model, at least qualitatively. In the model, the decline in the income growth rate is 2.3 percentage points for the lowest wealth quintile, but 5.4 percentage points for the highest wealth quintile (and the decline is monotonically increasing in wealth in between these two extreme wealth quintiles). In the data, the wealth-poorest 20% of the working-age population see their income growth rate slow down by 0.7 percentage point, whereas for the wealthiest households, income growth slows down by 3.0 percentage points.

In contrast, the performance of the model with respect to the changes in consumption rates is more mixed. In the model, in the recession households all increase consumption by more, or cut consumption by less, than disposable income, resulting in a rise

⁵² Huo and Rios-Rull (2016) and Kaplan, Mitman, and Violante (2016) investigate the role of price movements in housing in explaining aggregate consumption dynamics in the Great Recession.

Table 11: Difference in Annualized Growth Rates between Recession Period and Normal Times: Data and Model

NW Q	Net Worth (%)		Disp Y (%)		Expend.(%)		Exp. Rate (pp)	
	Data	Model	Data	Model	Data	Model	Data	Model
Q1	<i>NaN</i>	-20	-0.7	-2.3	-6.5	-2.2	-4.0	0.0
Q2	-98	-18	-2.6	-2.8	-5.2	-2.4	-1.6	0.3
Q3	-29	-12	-3.3	-4.0	-9.0	-2.7	-3.4	1.4
Q4	-15	-5	-3.3	-4.5	-7.4	-2.8	-2.5	2.0
Q5	-17	-4	-3.0	-5.4	-6.2	-2.9	-1.9	3.2
All	-19	-4	-2.9	-4.4	-6.9	-2.6	-2.5	1.3

in consumption rates, with the increase in consumption rates being smallest at the low end of the wealth distribution. In the data, all groups instead cut their consumption rates, the more so the less wealthy they are. Thus, although the model is consistent with the relative movement (in the recession vis-à-vis normal times) in consumption rates across wealth levels, with the wealth-poor decreasing consumption rates the most—in the data—or increasing them the least—in the model, the latter overstates consumption growth in the recession and thus underpredicts the decline in expenditure rates evident in the data.

In the model, when the recession hits and thus incomes decline (or grow less) relative to normal times, households have strong incentives to use their wealth to smooth consumption. This is especially true for those falling into unemployment. On the other hand, since the recession is long-lasting and comes with elevated unemployment risk, the motive to engage in precautionary saving against future unemployment spells increases, especially among those with little wealth coming into the recession. For high wealth households, the first motive dominates and the consumption rates of these households increase in the recession, whereas for low-wealth households both motives roughly balance out, leaving consumption rates roughly unchanged across the two time periods. We will show below that in an economy with less generous unemployment insurance, the precautionary savings motive becomes more potent, especially at the low end of the wealth distribution, and low-wealth households indeed cut their consumption rates during recessions, as is the case in the data.

We conclude this section by briefly summarizing the strengths and shortcomings of our baseline model when confronted with the PSID earnings, income, consumption, and wealth data. The model succeeds in replicating the observed cross-sectional wealth distribution (except at the very top) and does well in capturing the salient features of the joint distribution of wealth, income, and expenditures. It also replicates the relative movements of expenditure rates by wealth as the economy falls into a recession. However, it fails to predict the *decline* in consumption expenditure rates during recessions and fails to capture the large movements in wealth we see in the data during the years

2006-2010, since it abstracts from asset price movements.

In the next section we use the benchmark model and some of its variants to quantify the extent to which wealth inequality is important in determining the magnitude of aggregate consumption movements in response to a Great Recession type business cycle shock in TFP.

6 Cross-Sectional Household Heterogeneity and the Aggregate Dynamics of Consumption and Investment in a Severe Crisis

In this section we argue that the cross-sectional distribution of households across individual characteristics (primarily in wealth and impatience) is a crucial determinant of the aggregate consumption and investment response to a negative business cycle shock. In addition, we show that in the presence of such significant household heterogeneity, the generosity of social insurance policies strongly affects the dynamics of macroeconomic aggregates.

Our focus on the impact of household heterogeneity in wealth for the aggregate consumption dynamics during large recession is shared with a number of recent studies, including Guerrieri and Lorenzoni (2012), Glover, Heathcote, Krueger and Rios-Rull (2014), Heathcote and Perri (2015) as well as Berger and Vavra (2015).

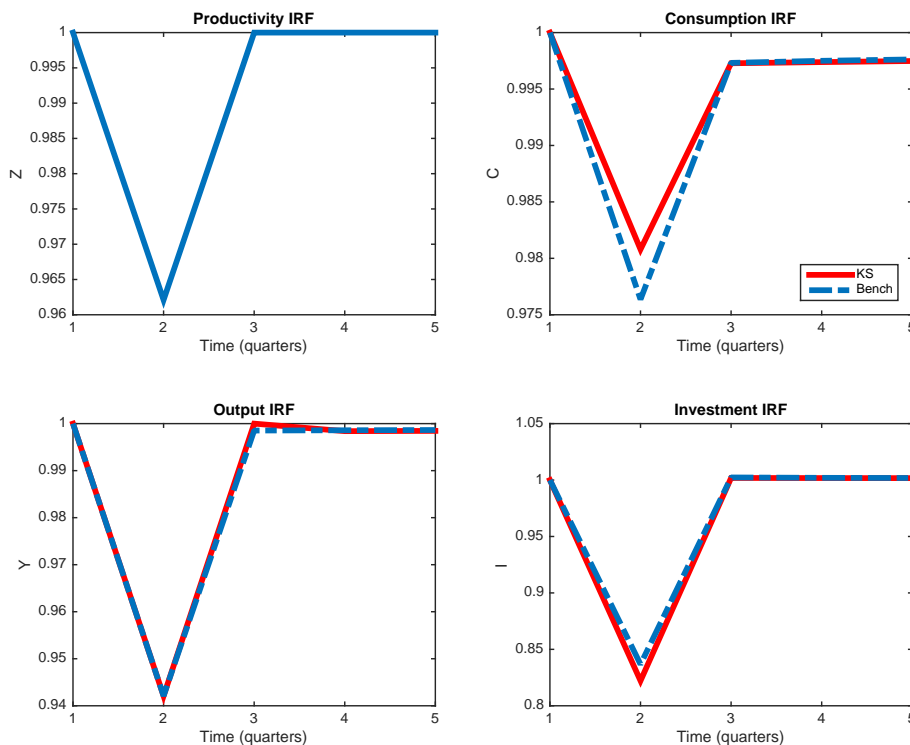
When exploring the role that social insurance policies can play in shaping the aggregate consumption (and, in the next section, output) response to adverse business cycle shocks in economies with household heterogeneity we build, on the work by Krusell and Smith (2006), which also focuses on income insurance programs, and more concretely, unemployment insurance.⁵³ Our work is also related to McKay and Reis (2016), who conduct a comprehensive study of automatic stabilization programs on business cycle dynamics, to Heathcote (2005), Kaplan and Violante (2014) and Jappelli and Pistaferri (2014), who study the role of discretionary changes in income taxation on aggregate consumption, and Brinca, Holter, Krusell and Malafry (2016), who investigate the magnitude of aggregate fiscal multipliers in this class of heterogeneous agent models.

⁵³ As we do, Auclert (2014), Auclert and Rognlie (2016) and Kekre (2016) also stress the importance of the heterogeneity in the marginal propensity to consume across households for the dynamics of aggregate demand and the impact of redistributive policies. Wong (2015) stresses the heterogeneity in age across households for the transmission of monetary policy shocks to aggregate consumption.

6.1 Benchmark Results

We consider two thought experiments, both of which take as an initial condition the wealth distribution after a long sequence of good shocks so that the cross-sectional distribution has settled down. Then a severe recession hits. In the first thought experiment, productivity returns to the normal state $Z = Z_h$ after one quarter (and remains there forever after). Although this thought experiment is not a good depiction of the actual Great Recession because of the short duration of the downturn, it displays the mechanics of the model recession most clearly.⁵⁴ In the second thought experiment we plot the responses of the economy to a Great Recession of typical length (according to our calibration) that lasts for 5.5 years (22 quarters). In both cases we trace out the impulse response functions (henceforth IRF) for the key macroeconomic aggregates. The main focus of interest is on the extent to which the aggregate consumption and investment responses differ across two economies that differ fundamentally in their extent of household heterogeneity.

Figure 2: Impulse Response to Aggregate Technology Shock in two Economies: One time Technology Shock



⁵⁴ Of course, households form expectations and make decisions based on the persistent Markov chain for Z driving the model even in this thought experiment.

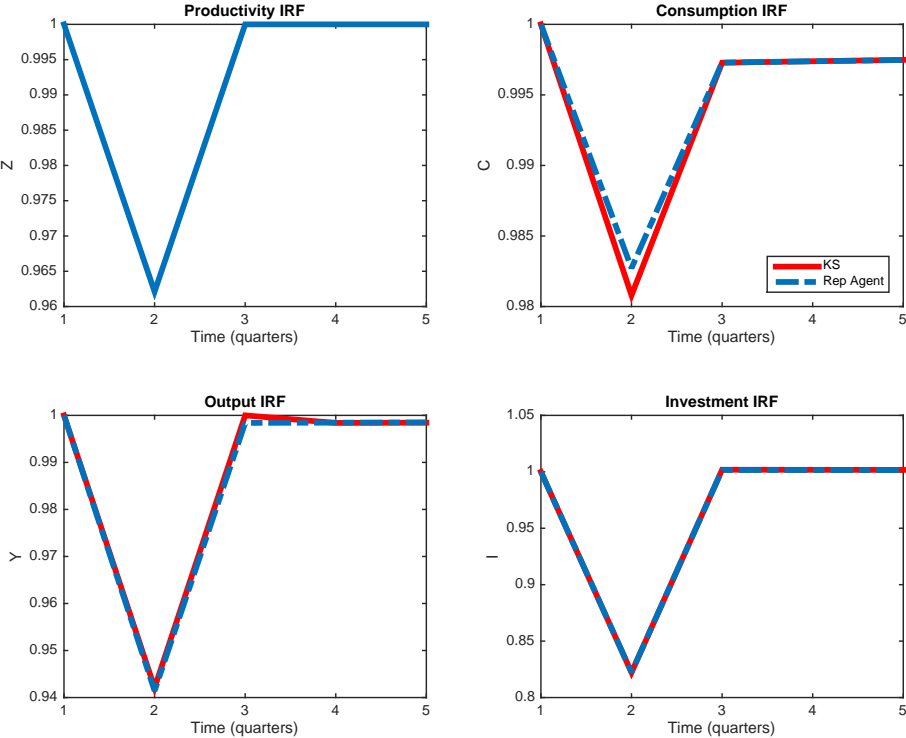
To make our main point, we perform both experiments for two model economies: the original Krusell-Smith economy without preference heterogeneity, life-cycle structure, and only modest unemployment insurance, and our benchmark model that includes these features and therefore, as documented above, provides a model wealth distribution that matches its empirical counterpart very well. We will also show that the aggregate consumption and investment behavior over the business cycle in the KS economy approximates an economy with representative agents (RA) very well (as already noted in the original Krusell and Smith (1998) paper), and thus as far as macroeconomic aggregates are concerned, the KS and the RA economy can be treated as quantitatively equivalent.

In figure 2 we plot the model impulse response to a onetime negative technology shock in which Z switches to Z_l after a long spell of good realizations Z_h . The upper left panel plots the time series of TFP Z fed into the model, and the remaining sub-plots show the model-implied dynamics of aggregate consumption, investment, and output induced by the Great Recession type TFP shock. By construction the time paths of exogenous TFP Z are identical in both economies in the short run; for output they are identical on impact and virtually identical over time. Since TFP and labor supply are exogenous in both economies and follow the same time path, capital is predetermined on impact, and the one time shock is not sufficient to trigger a substantially different dynamics of the capital stock, the time path of output is virtually identical in both economies. Thus, the key distinction between both economies is the extent to which a very similar decline and recovery in output is reflected in lower aggregate consumption rather than aggregate investment.

The key observation we want to highlight is that the aggregate consumption (and thus investment) response to the negative productivity shock differs substantially between the two economies. In the benchmark model, consumption falls by 2.4% in response to a technology shock that induces a decline in output by 6% on impact. The same fall in output triggers a decline of only 1.9% in the original Krusell-Smith (labeled as KS) economy. Thus the impact of the recession on aggregate consumption increases by 0.5% percentage points more in the economy with empirically plausible wealth heterogeneity. Given that output is exogenous in the short run, and is used for consumption and investment only in this closed economy, the investment impulse response necessarily shows the reverse pattern: the decline in investment is much weaker in the high wealth inequality economy. This in turn triggers a less significant decline and more rapid recovery of the macro economy once the recession has ended. However, given that new investment is only a small fraction of the capital stock, these differential effects on capital, and thus output, are quantitatively minor, at least in the case in which the recession is short-lived.⁵⁵

⁵⁵ In section C.4 we argue that the fact that the wealth distribution is quantitatively important for the current aggregate consumption response to a TFP shock does not imply that higher moments of the wealth distribution are needed to accurately forecast *future* wages and interest rates.

Figure 3: Impulse Response Functions (IRF) to Aggregate Technology Shock in KS and RA Economies

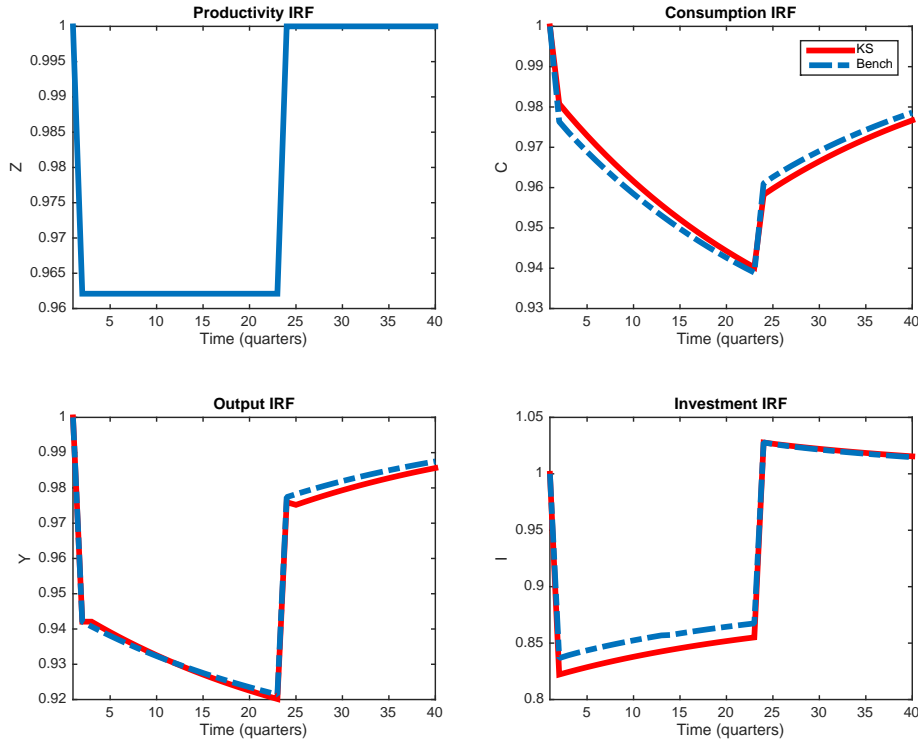


Note that for all practical purposes, in what follows the KS economy displays aggregate consumption-investment dynamics that are very close to those in a representative agent (RA) economy. Figure 3 shows this fact by displaying impulse responses to a one-period recession shock in the KS and RA economies. Although not identical, the impulse responses are quantitatively very close. For example, the aggregate consumption decline in the RA economy amounts to 1.78%, relative to a fall in aggregate consumption of 1.9% in the KS economy.

In figure 4, we display the dynamics of macroeconomic aggregates in a prolonged and severe recession, with a length of 22 quarters, under our operational definition of a severe recession. It demonstrates that in a Great Recession lasting several years, the differences in capital and output dynamics across the low-wealth inequality KS economy and the high inequality benchmark are now more noticeable, especially toward the end of the recession. As a result, the recovery after TFP has turned back up again is substantially stronger in the benchmark economy, by approximately 1 percentage point for capital and 0.3 percentage point for output in the period in which the recession ends.

Since the KS economy and the benchmark differ along several model dimensions, in

Figure 4: Impulse Response to Aggregate Technology Shock in Two Economies: "Typical" Severe Recession Technology Shock



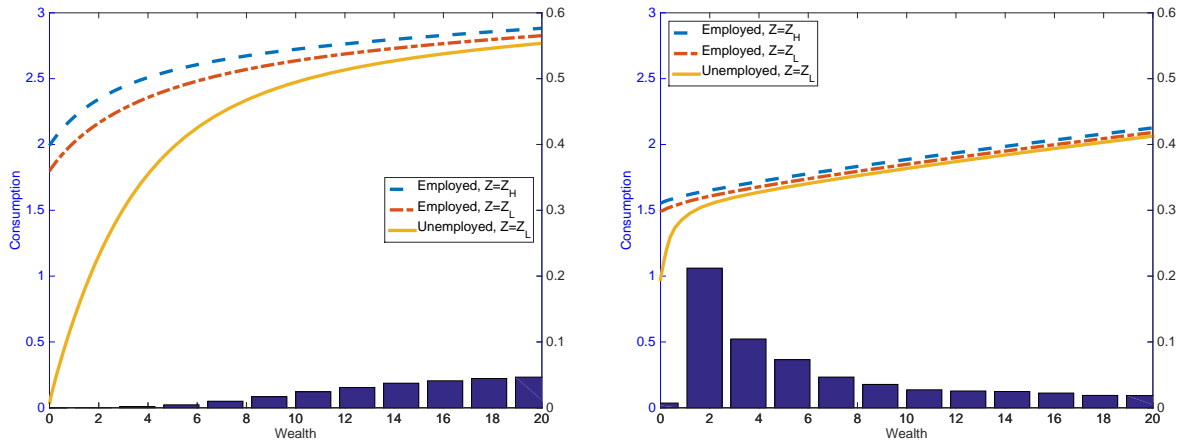
the next section we break down the reasons for the differential aggregate consumption response, again focusing on the interaction between the aggregate movement in consumption in a Great Recession and the cross-sectional wealth distribution prior to it.

6.2 Inspecting the Mechanism II: What Accounts for the Size of the Aggregate Consumption Recession

The key finding from the last section is that the aggregate consumption recession in our benchmark economy with preference and realistic wealth heterogeneity is more than twice as deep as it is in the corresponding RA economy (which in turn displays aggregate time series that are very close to those in the original KS economy). In this section, we dissect the reasons behind this finding. To start, in figure 5, we display the consumption functions and wealth distributions for both the KS and the benchmark economy. The left panel shows the consumption functions (plotted against individual wealth on the x -axis) in the original KS economy for three combinations of idiosyncratic employment and aggregate productivity states. For a given wealth level, the

vertical difference between the consumption functions for the employed in aggregate state $Z = Z_h$ (blue dashed line) and the employed in aggregate state $Z = Z_l$ (red dot-dashed line) gives the consumption drop in the Great Recession, conditional on not losing a job. In the same way, the vertical distance between the blue-dashed consumption function and the orange solid consumption function (for the unemployed in the recession) gives the consumption decline for those households that lose their jobs in a recession. The figure also contains the pre-recession wealth distribution, displayed as a histogram, with the mass of a particular wealth bin being measured on the right y -axis.⁵⁶ The right panel displays the same information, but for our benchmark economy, for working-age households with median earnings state y and mean discount factor β .

Figure 5: Consumption Function and Wealth Distribution: Krusell-Smith (left panel), Benchmark (right panel)



The first observation we make is that, for a given level of wealth, the drop in individual consumption as the KS economy falls into a Great Recession is substantially *larger* than in our benchmark economy.⁵⁷ This is especially true for households with little wealth that lose their jobs at the onset of the recession, because of the virtual absence of unemployment insurance.

The observation of larger individual consumption declines in the KS economy would suggest that the aggregate consumption recession is actually larger than it is in the benchmark economy, in contrast to the result documented in the previous section. However, as figure 5 (and table 6 above) display clearly, the cross-sectional wealth distribution places almost no mass on households with very little net worth, exactly the

⁵⁶ The aggregate capital stock associated with these plots is the pre-recession capital stock; note that both economies, by virtue of the calibration, have the same average (over the cycle) capital stock.

⁵⁷ Figure 5 displays the consumption functions in the benchmark economy for individuals with median (y, β) , but the same statement applies, qualitatively, to the consumption functions for households with other (y, β) characteristics. Recall that there is no (y, β) heterogeneity in the original KS economy.

households with the largest consumption declines. In contrast, the benchmark model with realistic wealth inequality places substantial probability mass at zero or close to zero wealth where the individual consumption losses are significant, especially (but not only) for newly unemployed households.⁵⁸ Note that average net worth is the same in both economies: we truncate the plots at net worth twenty times average income in order to make the individual consumption declines at the low end more clearly visible, but the benchmark economy has a fat right-tailed wealth distribution that is well approximated by a Pareto distribution (as in the data, see e.g. Benhabib and Bisin, 2016), whereas the original KS economy displays a wealth distribution whose right tail more closely resembles that of a log-normal distribution. Thus, both distributions have the same mean even though, as clearly visible from the figure, the benchmark economy has substantially more mass of households at low levels of net worth.

As we will see in subsection 6.3, public social insurance programs will affect both the determinants of the aggregate consumption dynamics—the consumption response to aggregate shocks for a given wealth level—and the wealth distribution itself. Both components are crucial when determining the overall impact of unemployment insurance policies on the macro economy over the business cycle. Before turning to this point, we first further explore the precise reasons behind the significant differences in aggregate and distributional characteristics between the original KS economy and our benchmark, thereby pinpointing precisely which model elements (and their interaction) are responsible for the differences in aggregate consumption dynamics across different economies.

Recall that relative to the KS model, our benchmark includes idiosyncratic earnings shocks, a rudimentary life cycle structure with social security system, permanent preference heterogeneity as well as a more generous unemployment insurance system.

In table 12 we repeat the information from table 7 on the wealth distribution in different versions of the model, but now we also document the magnitude of the aggregate consumption response on impact in a Great Recession. Figure 6 displays the associated impulse responses. From the table and figure we observe that the introduction of persistent idiosyncratic income risk on top of unemployment risk significantly amplifies the aggregate consumption response above that of the original KS model. In fact, the magnitude of the aggregate consumption response is larger than that obtained in the benchmark (the second to last column in the table). This is perhaps not surprising given our arguments thus far, as this version of the model generates significantly larger wealth inequality—and importantly—the two lowest wealth quintiles that hold very little net worth.⁵⁹

⁵⁸ The wealth distribution in the right panel of figure 5 is for the entire working-age population, rather than conditioning on the specific (y, β) types for which the consumption functions are displayed.

⁵⁹ Note, however, that this mechanism is insufficient to generate the very high wealth concentration, as the examination of the wealth share very top of the wealth distribution reveals.

Figure 6: Consumption Recessions in Various Versions of the Model

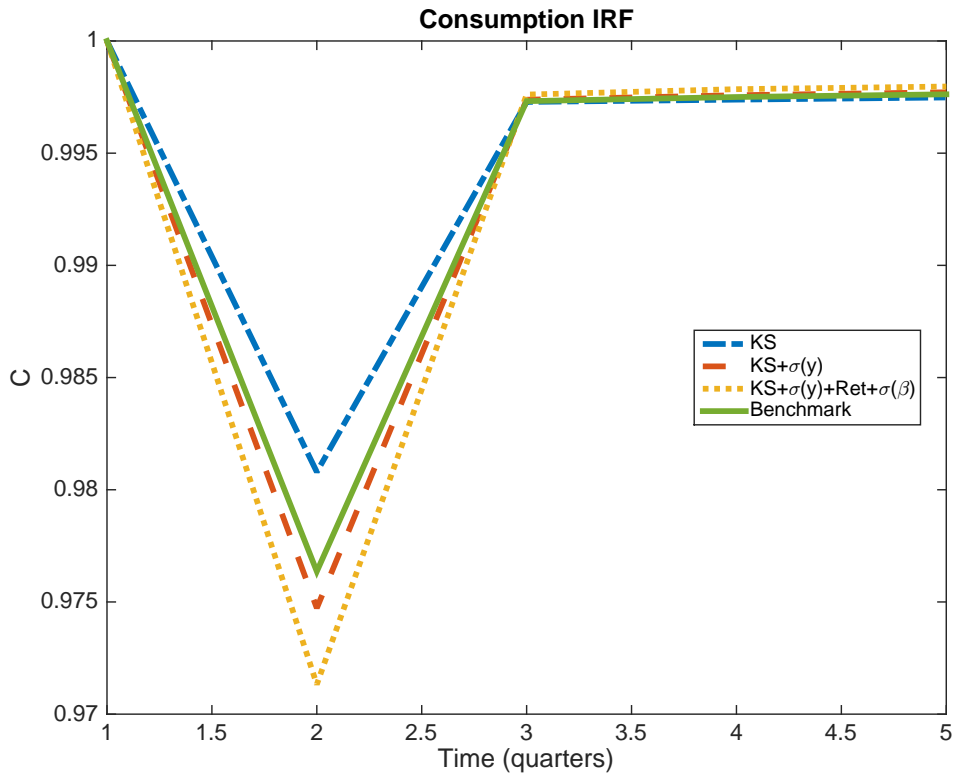


Figure 7: Consumption Function and Wealth Distribution: KS (left panel) and KS w/Income Risk (right panel)

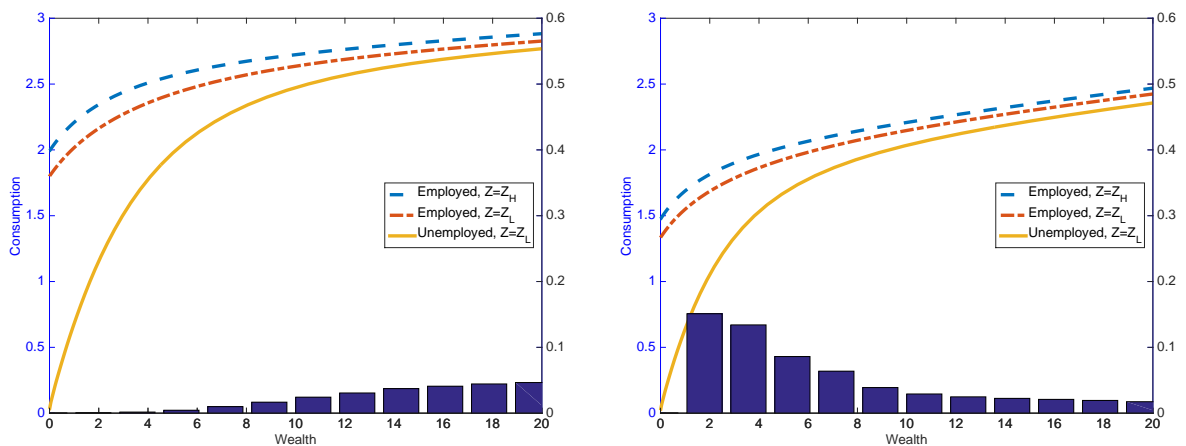


Figure 7 compares the consumption functions and equilibrium wealth distributions in the KS economy and the KS economy with just persistent earnings shocks added. In the latter, the policy functions are displayed for the median y realization. Whereas

Table 12: Net Worth Distributions and Consumption Decline: Different Versions of the Model

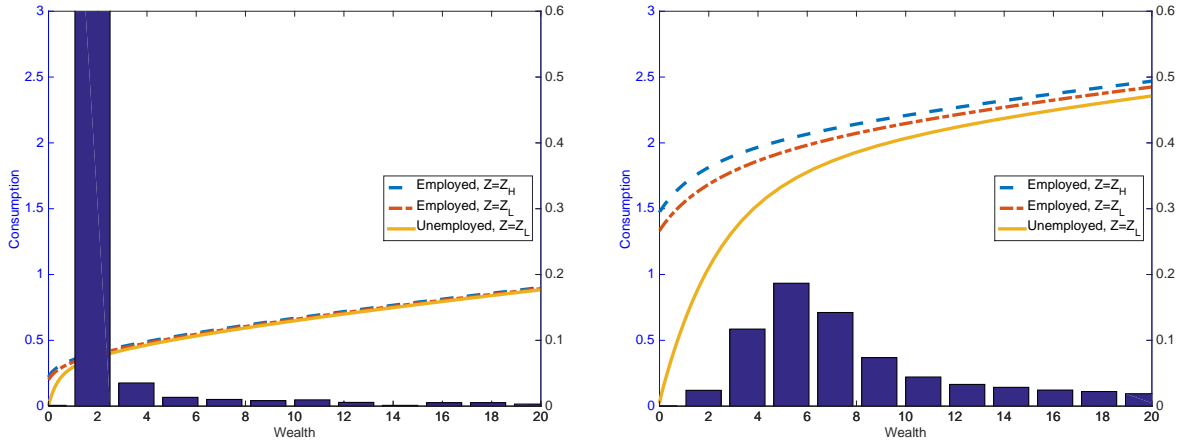
% Share:	Models*					KS + Top 1%
	KS	$+\sigma(y)$	+Ret.	$+\sigma(\beta)$	+UI	
Q1	6.9	0.7	0.7	0.7	0.3	5.0
Q2	11.7	2.2	2.4	2.0	1.2	8.6
Q3	16.0	6.1	6.7	5.3	4.7	11.9
Q4	22.3	17.8	19.0	15.9	16.0	16.5
Q5	43.0	73.3	71.1	76.1	77.8	57.9
90 – 95	10.5	17.5	17.1	17.5	17.9	7.4
95 – 99	11.8	23.7	22.6	25.4	26.0	8.8
T1%	5.0	11.2	10.7	13.9	14.2	30.4
Wealth Gini	0.350	0.699	0.703	0.745	0.767	0.525
ΔC	-1.9%	-2.5%	-2.6%	-2.9%	-2.4%	-2.0%

the consumption policy functions look broadly similar in both economies, the mass of households with low wealth and thus a large consumption response to the recession shock increases very substantially relative to the original KS economy. In this variant, the wealth distribution at the *bottom* looks already quite similar to the benchmark economy, although the absence of significant unemployment insurance implies that the mass of households at exactly zero wealth is negligible. On the other hand, because of the absence of unemployment insurance, the consumption drop of the wealth-poor for a given wealth level is comparable in magnitude to that in the original KS economy.

Figure 8, which displays the consumption functions and wealth distributions for two different types households in the $KS + \sigma(y)$ economy, clarifies the interaction between earnings inequality and wealth inequality. Households with low current (and very persistent) income realizations are highly concentrated at the low end of the wealth distribution. But even among households with contemporaneous median income, there is significantly more mass in the wealth region where consumption falls substantially upon unemployment.

Moving to the third column of table 12, we see that although the introduction of life-cycle elements is crucial for delivering joint income-consumption distributions, their impact on the dynamics of aggregate consumption in the recession is limited. In contrast, adding preference heterogeneity to the model helps to amplify the consumption drop. Crucially, now the economy is populated by a share of highly impatient house-

Figure 8: Consumption Function and Wealth Distribution: KS Low Income (left panel) and KS Median Income (right panel)



holds at the bottom of the wealth distribution. In normal times, unemployment risk is low and these households consume at a high rate because of their impatience, ending up with little or no wealth. When the economy falls into the recession, idiosyncratic unemployment risk goes up significantly for the "foreseeable future" from the point of view of impatient households. Faced with the elevated chance of becoming unemployed, impatient households who have not yet lost their jobs and have currently medium to high income realizations start to save more for precautionary reasons.⁶⁰

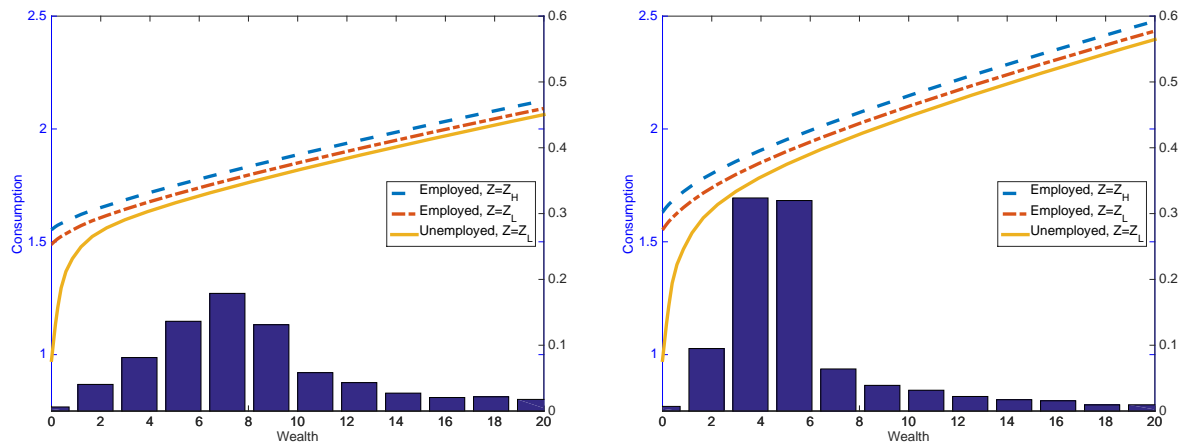
For more patient employed households, the increase in precautionary saving and resulting drop in consumption at the onset of the recession is not quite as severe. These households were already saving a larger fraction of their income even in good times, since their patience makes them more focused on the long horizon. Because the persistent idiosyncratic income component is more persistent than the recession, patient households with high current income expect to have high income even when exiting the recession, so the short-run possibility of increased unemployment is not as big of a concern to them.

Figure 9 displays the consumption policy functions for patient and impatient households, as well as the wealth distribution among these households. The key observation is that consumption falls more pronouncedly for impatient households when the aggregate state turns bad, even conditional on *not* losing a job. Also, not unexpectedly, among impatient households wealth levels tend to be lower, as the group-specific wealth distributions underneath the consumption functions in figure 9 show. As a broad summary measure of this differential effect, the contribution to the aggre-

⁶⁰ The small share of impatient, low-wealth households that do in fact lose their jobs at the onset of the recession behave like hand-to-mouth consumers instead, cutting their consumption one for one with income, and consume whatever little wealth they might have at the beginning of the recession.

gate decline in consumption is more than twice as large for the most impatient group of households than for the most patient group, even though that they constitute equal shares of the population.

Figure 9: Consumption Function and Wealth Distribution: Patient Households (left panel) and Impatient Households (right panel)



In the aggregate, the decline in aggregate consumption in the economy with income and preference heterogeneity amounts to 2.9%, and is thus a full 1 percentage point larger than in the KS economy, and 1.11% larger than in the representative agent economy. Both dimensions of heterogeneity are quantitatively important for the magnitude of the aggregate fluctuations, and so is their interaction, as the previous discussion of the importance of the impatient, employed with high income has indicated.

Finally, the second to last column of table 12 raises the unemployment insurance replacement rate to our benchmark value of 50%. As we discuss and quantify in the next part of the chapter, section 6.3, this change in the generosity of social insurance has a two-fold impact on the economy: for a given wealth level it softens the decline in household consumption in the recession, but it also shifts the wealth distribution toward wealth levels that imply a large decline in consumption and thus make the recession more costly in welfare terms. The first effect reduces the aggregate consumption response to the Great Recession shock, the second magnifies it. As table 12 shows, the net effect is a reduction of aggregate consumption volatility (with a decline of 2.4%), bringing the implications of the benchmark economy closer to that of the RA and KS economies with absent or limited wealth heterogeneity.

To summarize the main lessons from this section, the key aspects of the benchmark model that make its implied consumption dynamics different from its RA counterpart in a quantitatively meaningful way are a) an equilibrium wealth distribution that makes the wealth-poor poor enough and has them cut consumption more significantly than the average household when the recession hits; and b) that these wealth-poor

households make up a significant share of aggregate consumption. These requirements are achieved through highly persistent income shocks that generate a set of households that are born wealth-poor and never accumulate much wealth, and are compounded by the presence of impatient households that do not want to accumulate much wealth. If these households do not have access to generous unemployment insurance, their consumption falls a lot more than that of the representative household in a recession, either because they have in fact lost their jobs (and the incidence of job loss is higher in recessions), or because they have not lost their job, but have cut consumption to hedge against a now more likely job loss in the future.

Preference heterogeneity produces not only impatient households with the characteristics discussed thus far, but also patient households that find it optimal to accumulate large amounts of wealth, thereby contributing significantly to wealth inequality. However, it is the lack of wealth at the bottom, as opposed to significant concentration at the very top, that is crucial for explaining aggregate consumption dynamics. To make this point sharply, we consider a version of the model that is identical to the original KS model but adds limited preference heterogeneity. Specifically, it constructs a model in which 99% of the population has a lower time discount factor β_l than the remaining 1% of the population. The two discount factors are chosen to match the capital-output ratio in the benchmark economy (which essentially pins down β_l) and the share of wealth held by the top 1%–30%—as in the PSID data (whereas in the benchmark economy, we match the capital-output ratio and the wealth Gini). This pins down the time discount factor β_h of the remaining 1% of the population.

The purpose of this economy is to evaluate the importance of the wealth concentration at the very top of the distribution for the aggregate consumption decline in a Great Recession (and to demonstrate that it is straightforward, with appropriate preference heterogeneity in time discount factors, to generate a wealth distribution as concentrated at the top as it is in the data). The wealth distribution and aggregate consumption decline from this version of the model are reported in the last column of table 12. Since consumption functions are approximately linear for households with above-median wealth, and the individual consumption *drop* in a recession is roughly invariant to net worth at that level, it does not matter much *for aggregate consumption dynamics* if the top of the wealth distribution is populated by 1% of astronomically wealthy households, or by 20% of merely super rich households. Consequently, the consumption response is roughly the same in this variant of the model and in the original KS economy (and the RA economy for that matter).

6.2.1 The Importance of Precautionary Saving vs. "Hand-to-Mouth" Consumers

Given the importance we assigned to households with *little net worth* in our discussion above, in this subsection we briefly ask whether a model with a *fixed* fraction of households κ that always have zero wealth and thus simply consume their income in every

period has the same implications for the consumption dynamics as our benchmark model.⁶¹

We have resolved our model under the assumption that the bottom $\kappa = 40\%$ of the wealth distribution in model period $t - 1$ just consumes their earnings and unemployment benefits (if applicable) from period t on, whereas the remainder of the distribution (in period $t - 1$) continues to follow the intertemporally optimal decision rules from the benchmark economy.

The drop of consumption in a one-period Great Recession now amounts to 2%, relative to the decline in the benchmark economy of 2.4%. The drop is larger in the benchmark economy since households at the bottom of the wealth distribution on average (and especially those not currently unemployed) find it optimal to *reduce* consumption rates for precautionary reasons: the Great Recession is expected to last a long time, and those not yet affected by a job loss try to build a buffer to hedge against the increased risk of being laid off in the future.⁶² This precautionary saving motive in the face of increased idiosyncratic risk in recessions, also discussed lucidly in a recent paper by McKay (2015), is absent among households that follow a mechanical hand-to-mouth consumption rule and is responsible for the deeper recession in the benchmark economy.⁶³ We will return to this point in the next section, where we study the impact of the generosity of unemployment insurance on our results, and will show that with less generous unemployment insurance benefits, the additional precautionary savings motive from elevated unemployment risk is more potent, and the divergence between the class of models studied here and hand-to-mouth consumer models is even more significant.

It is important to note that in our formulation the share of households that behave as hand-to-mouth consumers is exogenous. In recent work Kaplan and Violante (2014) and Bayer, Lüticke, Pham-Dao and Tjaden (2015) construct models with wealthy hand-to-mouth consumers where a share of households endogenously choose to behave like hand to mouth consumers despite having non-trivial net worth. However, since their

⁶¹ This question is interesting from a modeling perspective since a model in which a fixed fraction of hand-to-mouth households and the remaining fraction employs permanent income consumption and savings functions (which are linear in wealth with identical marginal propensities to consume out of wealth, given our model) would give rise to easy aggregation.

⁶² In the versions of the model studied here, labor supply is exogenous (but its productivity fluctuating over the cycle), and thus saving is the only possible household response to hedge against higher idiosyncratic risk. In models with endogenous labor supply choice, such as the ones studied in Chang and Kim (2007) and Athreya, Owens and Schwartzman (2015), households have another margin of adjustment and thus the impact of elevated risk on precautionary saving will be smaller. For a model that combines household precautionary saving and *frictional* labor markets, see Krusell, Mukoyama and Sahin (2010).

⁶³ Obviously, the magnitude of this effect depends on the share of hand-to-mouth consumers κ . In the limit, as $\kappa = 0$ we are back in the benchmark economy. For $\kappa = 20\%$ the fall in aggregate consumption is 2.1%, about halfway between the RA economy and the benchmark.

net worth is primarily in the form of assets that are costly to liquidate (think of owner-occupied real estate and tax-favored retirement accounts), the consumption behavior of this group of households approximates that of the hand-to-mouth consumers modeled here, especially for income shocks of moderate magnitude.

6.3 The Impact of Social Insurance Policies

In this section we ask how the presence of public social insurance programs affects the response of the macro economy to aggregate shocks in a world with household heterogeneity.⁶⁴ We focus specifically on the effects of government-provided, and tax-financed unemployment insurance. We will argue that the impact of this policy is two-fold: it changes the consumption-savings response of a household with a *given* wealth level to income shocks, and it changes the cross-sectional wealth distribution in society, at least in the medium to long run. In order to decompose the overall impact of social insurance into these two effects, we consider two thought experiments. In the first, we simply compare the dynamics of macroeconomic aggregates of the benchmark economy with that of an identical economy that has a lower unemployment insurance replacement rate of $\rho = 10\%$. We interpret the latter economy as providing basic social insurance (as embedded in basic welfare programs), or alternatively, as a world where a significant share of households do not claim unemployment benefits despite being entitled to it.⁶⁵ This thought experiment will encompass the effect of unemployment insurance both on individual consumption behavior as well as on the equilibrium wealth distribution. To isolate the former effect, we will also consider an economy with low unemployment insurance, but entering the recession with the *same pre-recession wealth distribution* as in the benchmark economy.⁶⁶

In the left panel of figure 10, we plot, against wealth, the consumption functions (for the unemployed in the low and the employed in the high aggregate shock, with the mean discount factor) as well as the wealth histogram in the benchmark economy (with a replacement rate of 50%). This was the right panel of figure 5. The right panel of figure 10 does the same for an economy with an unemployment insurance system of only

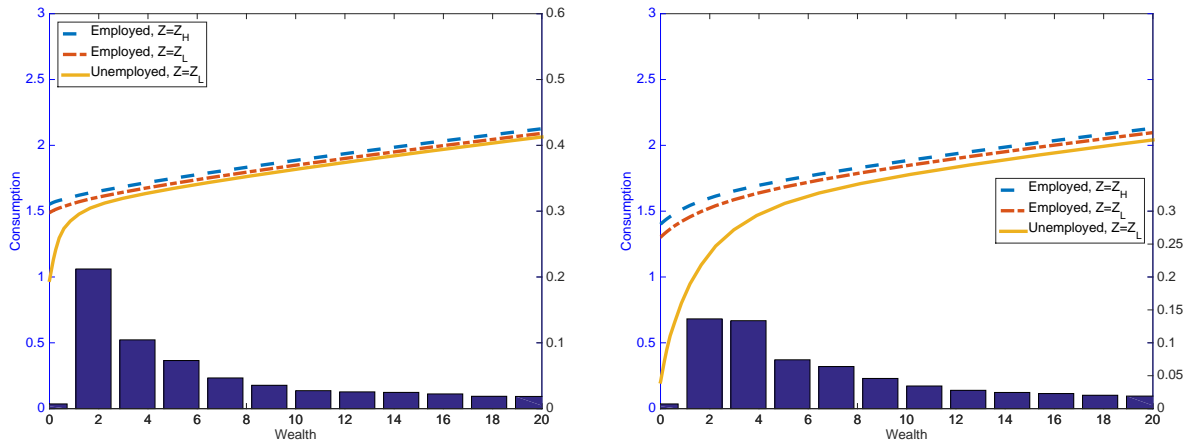
⁶⁴ The purpose of this analysis is purely positive in nature, and limited in scope by the assumption that transitions between employment and unemployment are exogenous and thus policy-invariant. See Hagedorn, Karahan, Manovskii and Mitman (2013) and Hagedorn, Manovskii and Mitman (2015) for an analysis of the effects of unemployment benefit extensions on vacancy creation and employment.

⁶⁵ We prefer to model a replacement rate of $\rho = 10\%$ rather than $\rho = 1\%$ as in the original Krusell-Smith economy studied in the previous subsection, since we think $\rho = 10\%$ is a more empirically relevant case. The resulting macro effects will lie right in between that of the benchmark economy, and the economy with a replacement rate of $\rho = 1\%$ displayed in the forth column (the $\sigma(\beta)$ economy) of table 12.

⁶⁶ One can interpret this thought experiment as a surprise permanent removal (or a surprise failure of extension) of unemployment benefits exactly in the period in which the recession hits.

10%. We chose to display the consumption function for the employed in an expansion and the unemployed in a recession because this helps us to best to understand what drives the aggregate consumption impulse response below.⁶⁷

Figure 10: Consumption Function and Wealth Distribution: Benchmark (left panel) and Low UI (right panel)



We want to highlight three observations. First, in the high unemployment insurance economy, households with low wealth consume much more than in the economy with small unemployment insurance. Second, and relatedly, the decline in consumption for low-wealth households from experiencing a recession with job loss is much more severe in the low-benefit economy. Third, the size of the social insurance system however, by affecting the extent to which households engage in precautionary saving, is a crucial determinant of the equilibrium wealth distribution. In the benchmark economy (as in the data), a sizeable mass of households has little or no wealth, whereas in the no-benefit economy this share of the population declines notably. Specifically, average assets increase by 0.5% relative to the benchmark economy, and only 0.9% of the population holds exactly zero assets, relative to 3.1% in the benchmark economy.

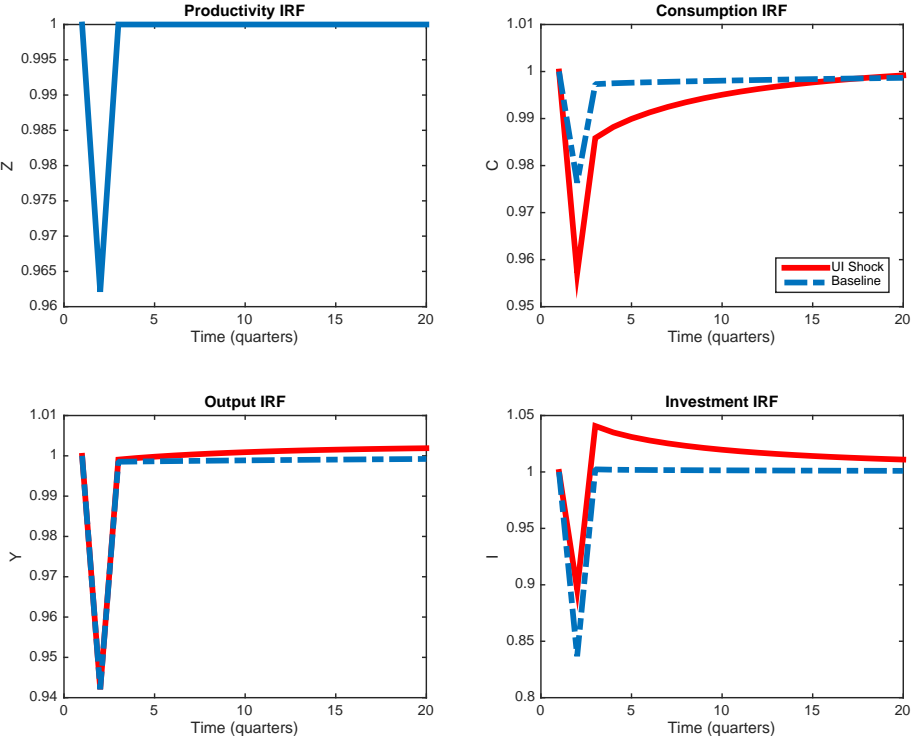
The difference in the consumption decline in a recession across the two economies can then be decomposed into the differential consumption response of households, integrated with respect to the *same* cross-sectional wealth distribution (which is a counterfactual distribution for one of the two economies), and the effect on the consumption response stemming from a policy-induced difference in the wealth distribution coming into the recession. As it turns out, both effects (the change in the consumption functions and the change in the wealth distribution) are quantitatively large, but partially offset each other.

In order to isolate the first effect, we now plot, in figure 11, the recession impulse

⁶⁷ Setting $\rho = 0$ would create the problem of zero consumption in some of the decomposition analyses we conduct later on.

response for the benchmark economy and the economy with low unemployment insurance, but starting at the *same pre-recession wealth distribution* as in the benchmark economy. Under this fixed wealth distribution scenario, the consumption response in both cases is given by the difference in the consumption functions (in both panels) integrated with the wealth distribution of the high UE insurance economy. We find that consumption declines much more substantially in the economy with a low replacement rate, by 4.6%, relative to 2.4% in the benchmark economy. This is, of course, exactly what the consumption functions in figure 10 predict.

Figure 11: Impulse Response to Aggregate Technology Shock with and without Generous Unemployment Insurance, Fixed Wealth Distribution: Onetime Technology Shock



To further quantify what drives this differential magnitude in the consumption response, in table 13 we display the fall in consumption for four groups in the population that differ in their transitions between idiosyncratic employment states as the aggregate economy slips into a recession. The share of households undergoing a specific transition is exogenous and the same across both economies, and is given in the second column of the table. Most (88.1%) households retain their jobs even though the aggregate economy turns bad. In contrast, the fraction of households making the transition from employment to unemployment is only 6.6% (and 3.5% of households make the reverse transition), but based on the consumption functions we expect them to display the largest decline in individual consumption.

Table 13: Consumption Response by Group in three Economies: Share of Total Decline

Transitions	Pop. Share	$\rho = 50\%, \Phi^{\rho=0.5}$	$\rho = 10\%, \Phi^{\rho=0.5}$	$\rho = 10\%, \Phi^{\rho=0.1}$
$s = e, s' = e$	88.1%	79.8%	72.8%	71.6%
$s = e, s' = u$	6.6%	13.8%	18.5%	21.8%
$s = u, s' = e$	3.5%	2.5%	2.9%	0.3%
$s = u, s' = u$	1.8%	3.8%	5.8%	6.3%
Total Decline	100%	-2.4%	-4.6%	-2.7%

The aggregate consumption decline documented in the last row of table 13 corresponds to the impulse responses of figure 11. The rows above give the share of the consumption decline accounted for by each of the four groups, so that the sum of the rows adds up to 100%. Similarly, table 14 summarizes the percentage consumption decline of each of the four groups and gives, in the second column, the pre-recession population shares of each of these four groups.

Table 14: Consumption Response by Group in 3 Economies: Consumption Growth Rates of Different Groups

Transitions	Pop. Share	$\rho = 50\%, \Phi^{\rho=0.5}$	$\rho = 10\%, \Phi^{\rho=0.5}$	$\rho = 10\%, \Phi^{\rho=0.1}$
$s = e, s' = e$	88.1%	-1.5%	-2.3%	-1.5%
$s = e, s' = u$	6.6%	-3.5%	-7.6%	-6.1%
$s = u, s' = e$	3.5%	-1.2%	-2.3%	-0.0%
$s = u, s' = u$	1.8%	-3.5%	-8.8%	-6.8%
Total Decline	100%	-2.4%	-4.6%	-2.7%

From both tables we observe that, even though the share of households that become newly unemployed (6.6% of the population, $s = e, s' = u$) and remain unemployed (1.8% of the population, $s = u, s' = u$) is relatively small, these groups account for a disproportionately large fraction of the overall consumption collapse in both the economy with generous, and in the economy with modest unemployment insurance.⁶⁸ See columns 3 and 4 of table 13 (which are based on the same pre-recession wealth distribution).

These two groups of households make up 8.4% of the population, but in the benchmark economy (column 3, table 13) account for 17.6% of the consumption drop. Carrying out the same decomposition for the economy with a small unemployment insurance system (column 4, table 13) we observe that the total drop in consumption is about

⁶⁸ For a recent empirical study on the link between unemployment and consumption expenditures, see Ganong and Noel (2015), who find reductions in consumption expenditures that are quantitatively similar to the ones our model with low unemployment insurance predicts.

twice as large now, as already displayed in the impulse response plot. Now the (newly and existing) unemployed have significantly larger percentage consumption drops (see the fourth column of table 14) and the share of the (now larger) consumption drop rises to 24.3%. Of course, the more pronounced consumption drop of the unemployed in a low UI benefit environment (and holding the wealth distribution fixed) is exactly what one would expect, and is already apparent in the policy functions of figure 10.

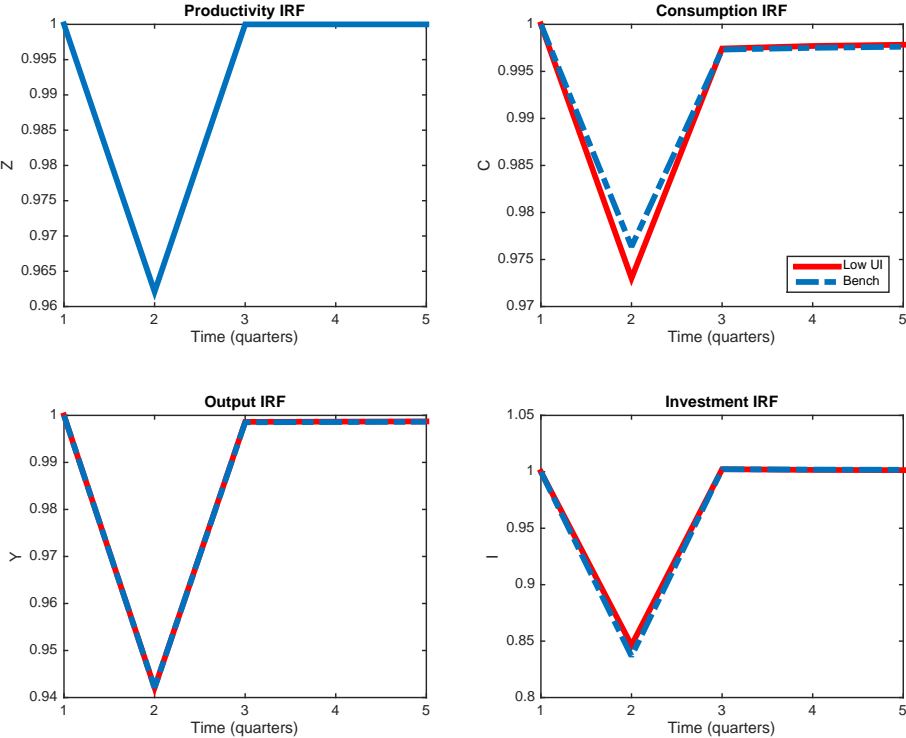
Table 14 contains a second important observation that we wish to stress. Looking at the magnitude of the consumption drops of households that have *not yet* lost their jobs as the economy falls into the recession (households with the idiosyncratic state transitions $s = u, s' = e$ and $s = e, s' = e$), we observe that these households, which constitute the vast majority of the population, also cut their consumption much more significantly in the (surprise) low-benefit economy, again comparing columns 3 and 4 of table 14. This is true even though these groups in both economies start with the same wealth distribution (by construction of the thought experiment) and experience the same income loss coming from a modest decline in aggregate wages. The lower UI benefits do not have an immediate impact on the earnings of these households, since they are currently employed even though the macro economy is doing poorly. The larger cuts in consumption of these groups instead emerge because future unemployment risk has gone up for these households as the economy falls into the highly persistent recession, and the potential future income losses from unemployment are larger in the economy with low unemployment insurance. Employed households, especially those with little net worth to start with, respond by elevating their saving and cutting their consumption rates, and since employed households make up 91.6% of the population, the extra fall in consumption of about 1 percentage point (in the economy with low UI, relative to the economy with high UI) is an important contributor to the overall larger decline of aggregate consumption in the low UI economy.

Finally, we document what happens if the wealth distribution is determined endogenously and responds to the absence of an unemployment insurance system. Figure 12 displays the impulse responses for the benchmark economy (again) and the no-benefits economy with a pre-recession wealth distribution that emerges in *that economy* after a long period of economic prosperity.⁶⁹ Column 5 of tables 13 and 14 breaks down the consumption response by subgroups. Overall we observe that the endogenous shift in the wealth distribution to the right that is due to the less generous unemployment insurance partially offsets the larger individual consumption declines in the no-benefits economy for a given wealth level.

To see this more precisely, compare the third and fifth columns of table 14. The aggregate consumption decline in the economy with little unemployment insurance is somewhat larger than in the benchmark economy (by 0.3 percentage point). But very notably, in this economy the unemployed (both newly and already existing ones) account for a substantially larger share of the reduction in consumption, even though this

⁶⁹ That wealth distribution was displayed in the right panel of figure 10.

Figure 12: Impulse Response to Aggregate Technology Shock with and without Generous Unemployment Insurance: Onetime Technology Shock



group understands the possibility of a Great Recession and has access to self-insurance opportunities to prepare for it. This is primarily because the employed, now fully aware of the fact that unemployment benefits will be low if they happen to become unemployed in the recession, enter the recession with larger wealth levels and do not cut their consumption as much as when they were surprised by the expiration of their benefits (compare columns 4 and 5 in table 14 for the employed, $s' = e$). Thus, all of the larger magnitude of the aggregate consumption decline with low UI benefits is driven by the small group of unemployed (compare columns 3 and 5 of table 14). The end effect is an aggregate consumption decline of 2.7% that is somewhat larger, but broadly consistent with that in the benchmark economy even though *individual consumption* responses to the crisis differ markedly across the two economies for the unemployed.

6.3.1 Revisiting the Importance of “Hand-to-Mouth” Consumers

In the absence of a generous unemployment insurance system, not only is the decline in aggregate consumption larger, as the previous section has argued, but the wealth-poor, not yet unemployed households have a greater incentive to save for now more

likely unemployment spells. As such, our economy with low replacement rate responds to aggregate shocks more strongly, relative to an economy with hand-to-mouth consumers, than the benchmark economy with $\rho = 50\%$. Recall that with $\rho = 50\%$ the aggregate consumption decline was 2.4%, relative to a fall of 2% in an economy with 40% hand-to-mouth consumers. With $\rho = 10\%$, the fall amounts to 2.7% in our economy and 2.1% in the hand-to-mouth consumer economy, and thus the divergence between the two models becomes stronger, on account of the elevated importance of the precautionary savings behavior of the wealth-poor, which is absent in models with exogenously given fixed shares of hand-to-mouth consumers. The recent papers by Ravn and Sterk (2013), McKay (2015) and Den Haan, Rehdal and Riegler (2016) are important examples that have stressed the importance of precautionary savings in the face of increased idiosyncratic risk for the dynamics of macro aggregates.⁷⁰

7 Inequality and Aggregate Economic Activity

In the model studied so far, the wealth distribution did potentially have an important impact on the dynamics of aggregate consumption and investment, but—by construction—only a fairly negligible effect on aggregate economic activity. Output depends on capital, labor input, and aggregate TFP, and in the previous model the latter two are exogenously given. The capital stock is predetermined in the short run, and even in the medium run only responds to net investment, which is a small fraction of the overall capital stock. So the output response to a negative productivity shock is exogenous on impact and, to a first approximation, exogenous (to the wealth distribution and to social insurance policies) even in the medium run. That is why in the previous section we focused on the distribution of the output decline between aggregate consumption and investment.

In the models discussed so far, aggregate demand played no independent role in shaping business cycle dynamics and, by construction, government demand management is ineffective. We now present a version of the model in which the output response to a negative shock is endogenous even in the short run, and thus potentially depends on the wealth distribution in the economy as well as policies that shape this distribution. The model retains the focus on real, as opposed to nominal, factors.⁷¹

The aggregate production function continues to be given by

⁷⁰ In related work Harmenberg and Oberg (2016) analyze the dynamics of consumption expenditures on durables in the presence of time-varying income risk.

⁷¹ In this chapter we abstract completely from nominal frictions that make output partially demand-determined. Representative papers that contain a lucid discussion of the demand- and supply-side determinants of aggregate *output* fluctuations in heterogeneous agent New Keynesian models are Gornemann, Kuester and Nakajima (2012), Challe, Matheron, Ragot and Rubio-Ramirez (2015) and Kaplan, Moll and Violante (2016).

$$Y = Z^*F(K, N)$$

with $Z^* = ZC^\omega$ and $\omega > 0$,

but now consider a world in which $\omega > 0$ and thus TFP $Z^* = ZC^\omega$ endogenously responds to the level of aggregate demand. A decline in aggregate consumption triggered by a fall in Z and an ensuing reduction of aggregate wages and household incomes endogenously reduces TFP and thus output further. This model with aggregate demand externalities is in the spirit of Bai, Rios-Rull and Storesletten (2012), Huo and Rios-Rull (2013) and Kaplan and Menzio (forthcoming), who provide micro foundations for the aggregate productivity process we are assuming here.⁷²

Since in this model a reduction in aggregate consumption C (say, induced by a negative Z shock) feeds back into lower TFP and thus lower output, government "demand management" might be called for, even in the absence of incomplete insurance markets against idiosyncratic risk. A social insurance program that stabilizes consumption demand of those adversely affected by idiosyncratic shocks in a crisis might be desirable not just from a distributional and insurance perspective, but also from an aggregate point of view. In the model with consumption externalities, in addition to providing consumption insurance it increases productivity and accelerates the recovery.⁷³

We now first discuss the calibration of the extended model before documenting how the presence of the demand externality affects our benchmark results.

7.1 Calibration Strategy

We retain all model parameters governing the idiosyncratic shock processes (s, y) , but recalibrate the *exogenous* part of aggregate productivity Z . In addition we need to specify the strength of the externality ω . Our basic approach is to use direct observations on TFP to calibrate the exogenous process Z and then choose the magnitude of the externality ω such that the demand externality model displays the same volatility of output as the benchmark model which (as the reader might recall) was calibrated to match the severity of the two severe recession episodes we identified in the data.⁷⁴

⁷² We are certainly not claiming that our and their formulations are isomorphic on the aggregate level; rather, their work provides the fully micro-founded motivation for the reduced form approach we are taking in this section.

⁷³ We think of this model as the simplest structure embedding a channel through which redistribution affects output directly and in the short run.

⁷⁴ An alternative approach would have been to retain the original calibration of the Z process, choose a variety of ω values, and document how much amplification, relative to the benchmark model, the externality generates. The drawback of this strategy is that output is counterfactually volatile in these thought experiments unless $\omega = 0$.

7.1.1 Exogenous TFP Process Z

For comparability with the benchmark results, we retain the transition matrix $\pi(Z'|Z)$ but recalibrate the states (Z_l, Z_h) of the process. To do so, we HP-filter the Fernald (2012) data for total factor productivity, identify as severe recessions the empirical episodes with high unemployment as in the benchmark analysis, and then compute average TFP (average percentage deviations relative to the HP-trend) in the severe recession periods, identified from unemployment data, as well as in normal times. This delivers

$$\frac{Z_l}{Z_h} = \frac{1 - 1.84\%}{1 + 0.36\%} = 0.9781.$$

Thus, the newly calibrated exogenous TFP process is significantly less volatile than in the benchmark economy, where the corresponding dispersion of TFP was given by $\frac{Z_l}{Z_h} = 0.9614$.

7.1.2 Size of the Spillover ω

Given the exogenous TFP process, we now choose ω such that the externality economy has exactly the same output volatility as the benchmark economy. This requires $\omega = 0.30$.

7.2 Results

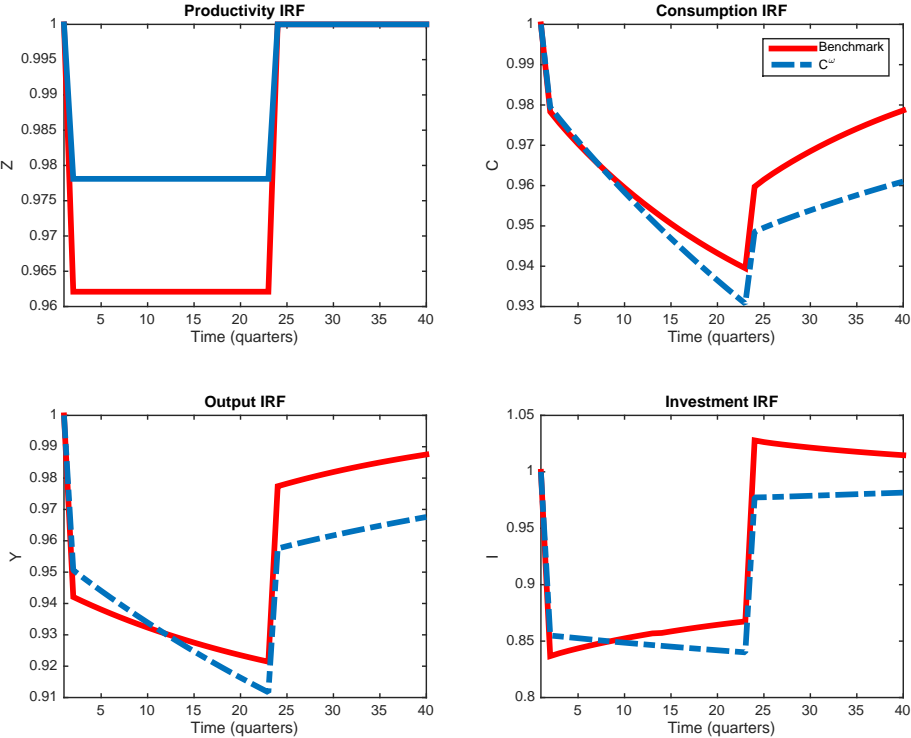
7.2.1 Aggregate Dynamics

In figure 13 we display the dynamics of a typical Great Recession (22 quarters of low TFP) in both the baseline economy and the demand externality economy (labeled C^ω).⁷⁵ The upper left panel shows that, as determined in the calibration section, a significantly smaller exogenous shock (2.2% as opposed to a 3.9% fall in TFP) is needed in the externality economy to generate a decline in output (and thus consumption and investment) of a given size. The impulse response functions are qualitatively similar in both economies, but with important quantitative differences.

First, the average decline in output in a Great Recession is the same across both economies since this is how $\frac{Z_l}{Z_h}$ was calibrated in the externality economy. However, since aggregate consumption declines during the course of a Great Recession and aggregate consumption demand impacts productivity, the decline in output is more pronounced and the recovery slower in the externality economy. Thus, the consumption externality

⁷⁵ The figure for a one-quarter Great Recession is qualitatively similar, but less useful in highlighting the differences between both economies.

Figure 13: Impulse Response to Aggregate Technology Shock: Comparison between Benchmark and Demand Externality Economy

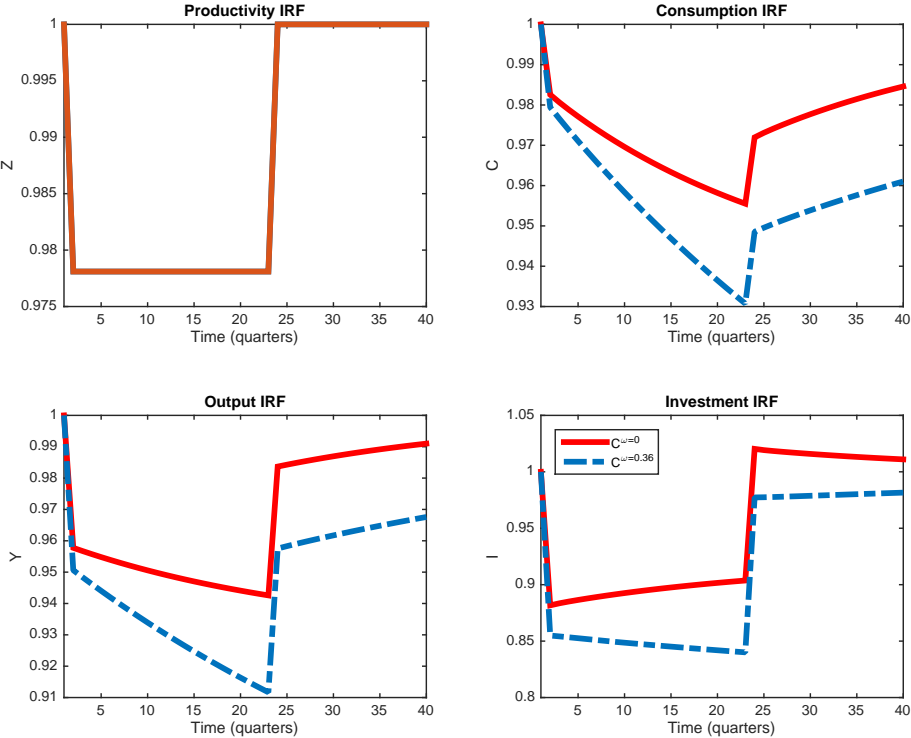


adds endogenous persistence to the model, over and above the channel already present through endogenous capital accumulation.

Of course, the demand externality mechanism also adds endogenous volatility to the model, but the fact that, via calibration, both models have the same output volatility obscures this fact. In figure 14 we display the magnitude of this amplification by comparing the impulse responses in two economies with the *same* exogenous TFP process (the one recalibrated for the demand externality economy), but with varying degrees of the externality ($\omega = 0$ and $\omega = 0.30$).

In contrast to figure 13, now the differences in the dynamics of the time series are purely driven by the presence of the demand externality. The amplification of the exogenous shock is economically important: the initial fall in output, consumption and investment is substantially larger (5.0%, 2.1% and 14.5% versus 4.2%, 1.7% and 11.8%, respectively). In addition, and consistent with figure 13, these larger output and consumption losses are more persistent in the economy with negative feedback effects from aggregate demand on productivity and thus production.

Figure 14: Impulse Response to Identical Aggregate Technology Shock: Comparison between Economies with and without Demand Externality

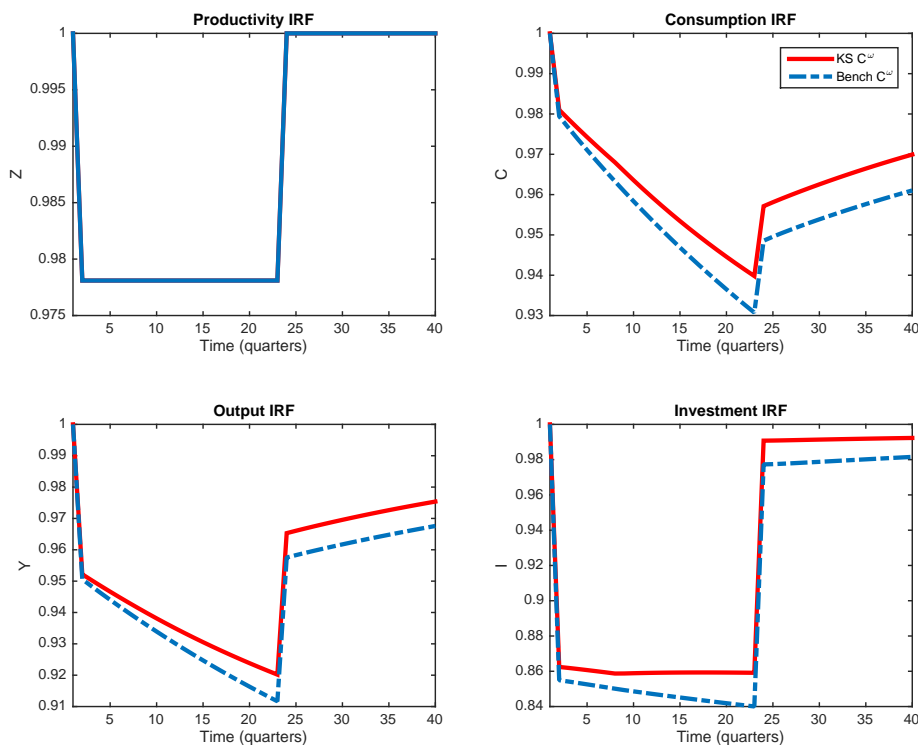


7.2.2 On the Importance of the Wealth Distribution when Output is Partially Demand-Determined

In principle, the previous results measuring the importance of aggregate consumption demand for output fluctuations did not require household heterogeneity at all. However, in the previous part of the chapter, we argued that the wealth distribution is a crucial determinant of aggregate consumption fluctuations, so it stands to reason the same is true with *output* fluctuations in economies where GDP is demand-determined. In figure 15 we verify this point, by displaying the aggregate impulse responses to a Great Recession in both the externality economy with plausible wealth heterogeneity and a version of the original Krusell-Smith economy, but also including the demand externality. The underlying exogenous TFP process is identical in both economies (and the same as in figure 14), and to display the differences between the models most clearly, we display the dynamics of the macro economy through a 22-quarter Great Recession.

As the figure clearly indicates, in the economy with realistic wealth inequality, the output recession is significantly greater, with output losses of 5.0% on impact and 8.8% at the end of the recession, compared with declines of 4.8% and 8.0% in the original KS

Figure 15: Impulse Response to Identical Aggregate Technology Shock: Comparison between Economies with High and Low Wealth Inequality



economy (but with demand externality).⁷⁶ In table 15 we summarize the consumption and output declines (on impact, and at the end of a Great Recession) for both the original KS and the benchmark economy, both with and without consumption externality.⁷⁷ It reconfirms the main message of figure 15: larger wealth dispersion, and especially lower wealth at the bottom of the wealth distribution, amplifies aggregate consumption recessions, as well as aggregate output recessions if the level of production is partially demand-determined. In the latter case, lower output in turn feeds back into an even more severe consumption recession. The magnitude of the differences is quantitatively significant, amounting to an additional drop of aggregate (and thus per capita) consumption of 0.9% at the end of the recession, because of larger wealth inequality induced by more realistic household heterogeneity (again comparing the benchmark model with the original KS economy).

⁷⁶ As in the economy without externality, the KS version of the model provides a very good approximation, as far as macroeconomic aggregates are concerned, for the corresponding representative agent economy.

⁷⁷ It is important to note that the results with $\omega = 0$ and $\omega = 0.3$ are not directly comparable, since in the economy with demand externality we feed in smaller TFP fluctuations, as described in the calibration section.

Table 15: Consumption and Output Declines in Four Economies

Economy	$\Delta_1 C$	$\Delta_1 Y$	$\Delta_{22} C$	$\Delta_{22} Y$
KS, $\omega = 0$	-1.9%	-5.8%	-6.0%	-8.0%
Bench., $\omega = 0$	-2.4%	-5.8%	-6.1%	-7.8%
KS, $\omega = 0.3$	-1.9%	-4.8%	-6.0%	-8.0%
Bench., $\omega = 0.3$	-2.1%	-5.0%	-6.9%	-8.8%

7.2.3 On the Interaction of Social Insurance and Wealth Inequality with Demand Externalities

In section 6.3 we demonstrated that the presence of social insurance policies has a strong impact on the aggregate consumption response to an adverse aggregate shock *for a given wealth distribution*, but also alters the long-run wealth distribution in the economy. With output partially demand-determined, these policies indirectly impact aggregate productivity and thus output. As the previous figures suggested, the effects are particularly important in the medium run due to the added persistence in the demand externality economy.

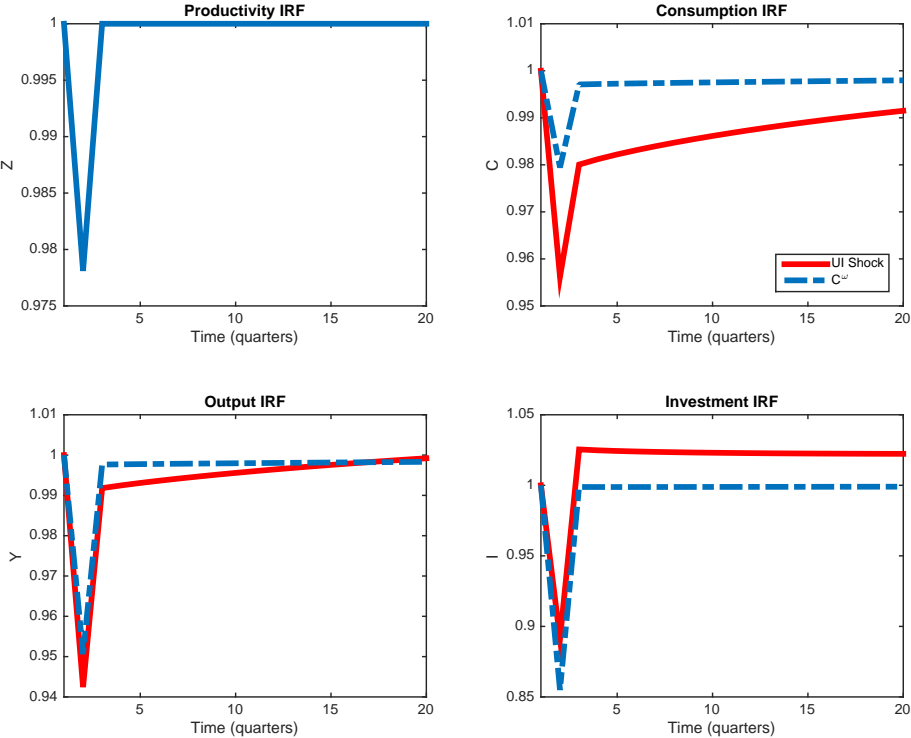
In figure 11 above we documented that, holding the wealth distribution fixed, the size of the social insurance system matters greatly for the aggregate consumption (and thus investment) response to an aggregate productivity shock. Figure 16 repeats the same thought experiment (an impulse response to a TFP shock in economies with $\rho = 50\%$ and $\rho = 10\%$ with the same pre-recession wealth distribution), but now for the consumption externality model.

The key observations from figure 16 are that now, in the consumption externality model, the size of the unemployment insurance system affects not only the magnitude of the aggregate consumption decline on impact, but also aggregate output, and the latter effect is quite persistent.

This can perhaps be more clearly seen in figure 17, which displays the *difference* in the impulse response functions for output and consumption between economies with $\rho = 50\%$ and $\rho = 10\%$, for both the benchmark model and the demand externality model. The presence of sizeable unemployment insurance stabilizes aggregate consumption more in the externality economy (the UI-induced reduction in the fall of C is 2.3% on impact and 1.3% after ten quarters of the initial shock in the externality economy, relative to 1.9% and 0.5% in the benchmark economy).

In addition, whereas in the benchmark economy more generous social insurance has no impact on output in the short run (by construction) and a moderately negative impact in the medium run (since investment recovers more slowly in the presence of more generous UI), with partially demand-determined output, UI stabilizes output significantly

Figure 16: Impulse Response to Aggregate Technology Shock with and without Generous Unemployment Insurance in Consumption Externality Model, Fixed Wealth Distribution

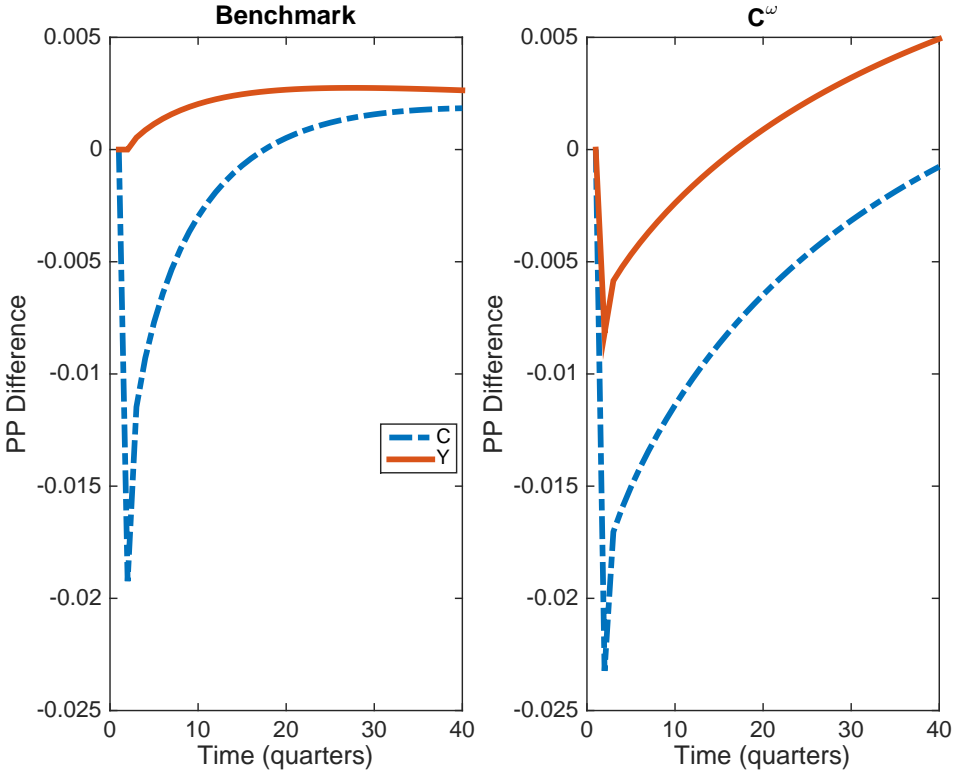


(close to 1% on impact, with the effect fading away only after 20 quarters—despite the fact that the shock itself only lasts for one quarter in this thought experiment.

Finally, we want to make a perhaps somewhat unexpected observation that turns out to be important for the calculation of the welfare losses of Great Recessions that we pursue in Krueger, Mitman and Perri (2016).⁷⁸ The surprise removal of unemployment benefits leaves households—especially those at the low end of the wealth distribution—with suboptimally small assets. These households start to save massively, especially in light of the elevated unemployment risk. Thus, in the medium run, wealth (the capital stock) and therefore aggregate consumption starts to rise. And since total factor productivity is linked to aggregate consumption demand (and since

⁷⁸ In that paper, we contribute to the very large literature that studies the normative consequences of social insurance policies (such as unemployment insurance, social security and progressive income taxation) in quantitative heterogeneous household models. See Domeij and Heathcote (2004), Cautt, Imrohorglu and Kumar (2006), Conesa, Kitao and Krueger (2009), Peterman (2013), Heathcote, Storesletten and Violante (2014), Mitman and Rabinovich (2015), Bakis, Kaymak and Poschke (2015), Karabarbounis (2016), Krebs, Kuhn and Wright (2015), and Krueger and Ludwig (2016) for recent representative contributions to this literature.

Figure 17: Difference in IRF between $\rho = 50\%$ and $\rho = 10\%$, with and without Consumption Externality



the capital stock in the economy increases), aggregate wages and output rise strongly in the medium run in the externality economy with low unemployment insurance benefits.⁷⁹ As long as households are sufficiently patient⁸⁰ and have *not* lost their job in the recession, the stronger recovery of the macro economy with low unemployment benefits might make these households prefer less generous unemployment insurance, despite the fact that unemployment insurance benefits act as effective aggregate demand stabilizers in the short run (again as figure 17 clarifies).

This last finding, discussed in much greater length in Krueger, Mitman and Perri (2016), leads us back to the main overall theme of this chapter: we have demonstrated that the extent of household heterogeneity with respect to income, wealth and preferences, in a canonical heterogeneous household business cycle model, determines the aggregate consumption and output dynamics over the business cycle in a quantitatively significant way. It gives social insurance policies that shape the income, consumption and wealth distributions a potentially important role in aggregate consump-

⁷⁹ Mitman and Rabinovich (2014) argue, reversely, that the *extension* of unemployment benefits goes a long way towards explaining recent slow recoveries in U.S. data.

⁸⁰ Recall that the population is heterogeneous with respect to the time discount factor.

tion and output stabilization and has (as we show in our companion work) welfare implications that vary strongly across households with different characteristics. Modelling microeconomic heterogeneity explicitly in the analysis of Great is therefore potentially quantitatively important, even if the object of research interest is purely aggregate in nature.

8 Conclusion

In this chapter we used PSID data on earnings, income, consumption, and wealth as well as different versions of a canonical business cycle model with household earnings and wealth heterogeneity to study the conditions under which the cross-sectional wealth distribution shapes the business cycle dynamics of aggregate output, consumption and investment in a quantitatively meaningful way. We argued that the low end of the wealth distribution is crucial for the answer to this question. We studied mechanisms that helped to generate close to 40% of households without significantly positive net worth, including highly persistent earnings shocks, preference heterogeneity and publicly provided social insurance programs. We showed that the decline in consumption of this group of wealth-poor households at the onset of the recession generates a significantly larger aggregate consumption drop than in a representative-household version of the neoclassical growth model. The same is true for output if it is partially demand-determined. We argued that the key mechanism underlying this result is increased precautionary savings against elevated unemployment risk, and we investigated the extent to which social insurance programs impact the strength of this channel.

Our work suggests that there are at least three important research directions that could yield new insights on the role of heterogeneity for macro outcomes.⁸¹

The first is the introduction of additional dimensions of household heterogeneity, so that the model can better capture the *joint* distribution of wealth, income and expenditure we observe in the data. A more accurate mapping between the model and household micro data might change our quantitative conclusions regarding the impact of household heterogeneity on macro dynamics.

The second dimension is the introduction of a richer model of the labor market, with elastic labor supply and other frictions impacting equilibrium hours and unemployment. Doing so would allow us to better understand the link between changes in aggregate consumption expenditures and changes in aggregate output, which in this chapter we have modeled in a very reduced form way.

The final direction for promising work is the explicit introduction of aggregate shocks to the net worth of households (which one may call financial shocks). The micro data

⁸¹ We fully acknowledge that exciting work in all these dimensions is already under way.

on the dynamics of household wealth have shown that during the Great Recession large changes in the net worth of households occurred, and the current model with only one asset does not capture these changes. Introducing a mechanism that can generate these fluctuations in the price of different assets could modify the mechanisms leading from the micro wealth distribution to aggregate consumption and output described in this chapter.

More generally, the emergence of new rich household and firm-level data sets, coupled with continuous theoretical and computational advances in the solutions of macro models with micro heterogeneity, as well as renewed scientific and popular interest in distributional questions, make the research field of quantitative heterogeneous agent macroeconomics an exciting area for future inquiry.

References

- [1] Aguiar, M., and M. Bilal (2015), "Has Consumption Inequality Mirrored Income Inequality?", *American Economic Review* 105: 2725-2756.
- [2] Aiyagari, S. R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics* 109: 659-684.
- [3] Ameriks, J., J. S. Briggs, A. Caplin, M. D. Shapiro and C. Tonetti (2015), "Long-Term Care Utility and Late in Life Saving", "NBER working paper No.20973.
- [4] Arellano, M., R. Blundell and S. Bonhomme (2015), "Household Earnings and Consumption: A Nonlinear Framework", *Working paper*.
- [5] Athreya, K., A. Owens and F. Schwartzman (2015), "Does Redistribution Increase Output? The Centrality of Labor Supply", *Working paper*.
- [6] Atkinson, A. B., and F. Bourguignon, eds. (2015), *Handbook of Income Distribution*, Vol. 2 North-Holland, Amsterdam.
- [7] Attanasio, O. (1999), "Consumption", in: J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Elsevier Science, Amsterdam.
- [8] Attanasio, O. P., and G. Weber (2010), "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy", *Journal of Economic Literature* 48: 693-751.
- [9] Auclert, A. (2014), "Monetary Policy and the Redistribution Channel" *Working paper*.
- [10] Auclert, A. and M. Rognlie (2016), "Inequality and Aggregate Demand", *Working paper*.
- [11] Bachmann, R., R. J. Caballero and E. Engel (2013), "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", *American Economic Journal: Macroeconomics* 5: 29-67.
- [12] Bai, Y., J.-V. Rios Rull and K. Storesletten (2012), "Demand Shocks as Productivity Shocks", *Working Paper*.
- [13] Bakis, O., B. Kaymak and M. Poschke (2015), "Transitional Dynamics and the Optimal Progressivity of Income Redistribution", *Review of Economic Dynamics*, 18: 679-693.
- [14] Barro, R. J. (2006), "Rare Disasters and Asset Markets in the Twentieth Century", *Quarterly Journal of Economics* 121: 823-866.

- [15] Bayer, C., R. Lütticke, L. Pham-Dao and V. Tjaden (2015), "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk", *Working paper*.
- [16] Benabou, R. (2002), "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?", *Econometrica* 70: 481-517.
- [17] Benhabib, J., and A. Bisin (2016), "Skewed Wealth Distributions: Theory and Empirics", NBER working paper No. 21924.
- [18] Benhabib, J., A. Bisin and S. Zhu (2011), "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents ", *Econometrica* 79: 123-157.
- [19] Berger, D., and J. Vavra (2015), "Consumption Dynamics During Recessions" *Econometrica* 83: 101-154.
- [20] Bewley, T. (1986), "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers", in: W. Hildenbrand and A. Mas-Colell, eds., *Contributions to Mathematical Economics in Honor of Gerard Debreu* North-Holland Amsterdam, 79-102.
- [21] Blank, R. M. and D. E. Card (1991), "Recent Trends in Insured and Uninsured Unemployment: Is There an Explanation?" *Quarterly Journal of Economics* 106:1157-1190.
- [22] Blundell, R., L. Pistaferri and I. Preston (2008), "Consumption Inequality and Partial Insurance" *American Economic Review* 98:1887-1921.
- [23] Brinca, P., H. A. Holter, P. Krusell and L. Malafry (2016), "Fiscal Multipliers in the 21st Century" *Journal of Monetary Economics* 77:53-69.
- [24] Brüggemann, B., and J. Yoo (2015), "Aggregate and Distributional Effects of Increasing Taxes on Top Income Earners ", *Working paper*.
- [25] Buera, F. (2009), "A Dynamic Model of Entrepreneurship with Borrowing Constraints: Theory and Evidence", *Annals of Finance* 5: 443-464.
- [26] Burkhauser, R. V., J. Larrimore and K. Simon (2012), "A 'Second Opinion' on the Economic Health of the American Middle Class ", *National Tax Journal* 65: 7-32
- [27] Cagetti, M., and De Nardi, M. (2006), "Entrepreneurship, Frictions, and Wealth", *Journal of Political Economy*, 114, 835-870.
- [28] Carroll, C. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence" , *Brookings Papers on Economic Activity* 1992: 61-135.
- [29] Carroll, C. (1997), "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1-55.

- [30] Carroll, C. D., J. Slacalek, K. Tokuoka and M. White (2015), "The Distribution of Wealth and the Marginal Propensity to Consume", *Working paper*.
- [31] Castaneda, A., J. Diaz-Gimenez and J.-V. Rios-Rull (1998), "Exploring the Income Distribution Business Cycle Dynamics", *Journal of Monetary Economics* 42: 93-130.
- [32] Castaneda, A., J. Diaz-Gimenez and J.-V. Rios-Rull (2003), "Accounting for the U.S. Earnings and Wealth Inequality", *Journal of Political Economy* 111:818-857.
- [33] Caucutt, E., S. Imrohoroglu and K. Kumar (2006), "Does the Progressivity of Income Taxes Matter for Human Capital and Growth?", *Journal of Public Economic Theory* 8: 95-118.
- [34] Challe, E., J. Matheron, X. Ragot and J. F. Rubio-Ramirez (2015), "Precautionary Saving and Aggregate Demand", *Working Paper*.
- [35] Chang, Y., and S.-B. Kim (2007), "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations", *American Economic Review* 97:1939-1956.
- [36] Chatterjee, S., D. Corbae, M. Nakajima and J.-V. Rios-Rull (2007), "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default" *Econometrica*, 75:1525-1590.
- [37] Chodorow-Reich G. and L. Karabarbounis (2016), "The Cyclicity of the Opportunity Cost of Employment", *Journal of Political Economy*, forthcoming
- [38] Cocco, J. F., F. J. Gomes and P. J. Maenhout (2005), "Consumption and Portfolio Choice over the Life-Cycle", *Review of Financial Studies* 18: 491-533.
- [39] Conesa, J. C., S. Kitao and D. Krueger (2009), "Taxing Capital? Not a Bad Idea after All!" *American Economic Review*, 99, 25-48.
- [40] Deaton, A. (1991), "Saving and Liquidity Constraints" *Econometrica* 59: 1221-1248.
- [41] Deaton, A. (1992), *Understanding Consumption*, Oxford University Press, New York.
- [42] De Nardi, M. (2004), "Wealth Inequality and Intergenerational Links", *Review of Economic Studies* 71: 743-768.
- [43] De Nardi, M. (2015), "Quantitative Models of Wealth Inequality: A Survey", *NBER working paper* No. 21106.
- [44] De Nardi, M., G. Fella and F. Yang (2015), "Piketty's Book and Macro Models of Wealth Inequality", *NBER working paper* No. 21730.
- [45] De Nardi, M., E. French and J. B. Jones (2010), "Why Do the Elderly Save? The Role of Medical Expenses", *Journal of Political Economy* 118: 39-75.

- [46] Den Haan, W. J. (2010), "Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents", *Journal of Economic Dynamics & Control* 34: 79-99.
- [47] Den Haan, W., P. Rendahl and M. Riegler (2016), "Unemployment (Fears), Precautionary Savings, and Aggregate Demand", *Working paper*.
- [48] Domeij, D. and J. Heathcote (2004), "On the Distributional Effects of Reducing Capital Taxes", *International Economic Review* 45: 523-554.
- [49] Fernald, J. (2012), "Productivity and Potential Output before, during, and after the Great Recession", Federal Reserve Bank of San Francisco Working Paper 2012-18.
- [50] Fisher, J., D. Johnson, T. Smeeding and J. Thompson (2015), "Inequality in 3D: Income, Consumption and Wealth", Presentation at 2015 NBER Summer Institute.
- [51] Gabaix, X., J. Lasry, P. Lions and B. Moll (2014), "The Dynamics of Inequality", *Working paper*.
- [52] Ganong, P., and Noel, P. (2015), "How Does Unemployment Affect Consumer Spending?", *Working paper*.
- [53] Glover, A., J. Heathcote, D. Krueger and J.-V. Rios-Rull (2014), "Inter-generational Redistribution in the Great Recession", *Working paper*.
- [54] Gomes, F., and A. Michaelides (2008), "Asset Pricing with Limited Risk Sharing and Heterogeneous Agents", *Review of Financial Studies* 21: 415-448.
- [55] Gornemann, N., K. Kuester and M. Nakajima (2012), "Monetary Policy with Heterogeneous Agents", *Working Paper*.
- [56] Gourio, F. (2013), "Credit Risk and Disaster Risk", *American Economic Journal: Macroeconomics*, 5: 1-34.
- [57] Gruber, J. (1994). "The Consumption Smoothing Benefits of Unemployment Insurance," *NBER working paper* No. 4750.
- [58] Guerrieri, V. and G. Lorenzoni (2012), "Credit Crises, Precautionary Savings, and the Liquidity Trap", *Working Paper*.
- [59] Guvenen, F. (2009), "An Empirical Investigation of Labor Income Processes", *Review of Economic Dynamics* 12: 58-79.
- [60] Guvenen, F. (2011), "Macroeconomics with Heterogeneity: A Practical Guide", *Economic Quarterly* 97: 255-326.
- [61] Guvenen, F., S. Ozkan and J. Song (2014), "The Nature of Countercyclical Income Risk", *Journal of Political Economy* 122: 621-660.

- [62] Hagedorn, M., F. Karahan, I. Manovskii and K. Mitman (2013), "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects", *NBER working paper* No. 19499.
- [63] Hagedorn, M., I. Manovskii and K. Mitman (2015), "The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?", *NBER working paper* No. 20884.
- [64] Harmenberg, K. and E. Oberg (2016), "Durable Expenditure Dynamics under Time-Varying Income Risk", *Working Paper*.
- [65] Heathcote, J. (2005), "Fiscal Policy with Heterogeneous Agents and Incomplete Markets", *Review of Economic Studies* 72: 161-188.
- [66] Heathcote, J., and F. Perri (2015), "Wealth and Volatility", *NBER working paper* No. 20994.
- [67] Heathcote, J., Perri, F., and Violante, G. (2010), "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006", *Review of Economic Dynamics* 13: 15-51.
- [68] Heathcote, J., K. Storesletten and G. Violante (2009), "Quantitative Macroeconomics with Heterogeneous Households", *Annual Review of Economics* 1: 319-354.
- [69] Heathcote, J., S. Storesletten and G. Violante (2014), "Optimal Tax Progressivity: An Analytical Framework," *Working paper*.
- [70] Hendricks, L. (2007), "How Important Is Discount Rate Heterogeneity for Wealth Inequality?", *Journal of Economic Dynamics and Control* 31: 3042-3068.
- [71] Herkenhoff, K. (2015), "The Impact of Consumer Credit Access on Unemployment," *Working paper*.
- [72] Huggett, M. (1993), "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies", *Journal of Economic Dynamics and Control* 17: 953-969.
- [73] Huggett, M. (1996), "Wealth Distribution in Life-Cycle Economies", *Journal of Monetary Economics* 38: 469-494
- [74] Huggett, M. (1997), "The One-Sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics", *Journal of Monetary Economics* 39:385-403.
- [75] Huo, Z. and J.-V. Rios-Rull (2013), "Paradox of Thrift Recessions ", *Working Paper*.
- [76] Huo, Z. and J.-V. Rios-Rull (2016), "Balance Sheet Recessions ", *Working Paper*.
- [77] Imrohoroglu, A. (1989), "Cost of Business Cycles with Indivisibilities and Liquidity Constraints", *Journal of Political Economy* 97: 1364-1383.

- [78] Jappelli, T., and L. Pistaferri (2014), "Fiscal Policy and MPC Heterogeneity", *American Economic Journal: Macroeconomics* 6: 107-136.
- [79] Kaplan, G. and G. Menzio (2014), "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations", *Journal of Political Economy*, forthcoming.
- [80] Kaplan, G., K. Mitman and G. Violante (2016), "Consumption and House Prices in the Great Recession: Model Meets Evidence", *Working paper*.
- [81] Kaplan, G., B. Moll and G. Violante (2016), "Monetary Policy According to HANK", *NBER working paper* No. 21897
- [82] Kaplan, G. and G. Violante (2014), "A Model of the Consumption Response to Fiscal Stimulus Payments", *Econometrica* 82: 1199-1239.
- [83] Karabarbounis, M. (2015), "A Road Map for Efficiently Taxing Heterogeneous Agents", *American Economic Journal: Macroeconomics*, forthcoming.
- [84] Kekre, R. (2015), "Unemployment Insurance in Macroeconomic Stabilization", *Working paper*
- [85] Khan, A. and J. K. Thomas (2008), "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics", *Econometrica* 76: 395-436.
- [86] Kopecky, K. and R. Suen (2010), "Finite State Markov-Chain Approximations to Highly Persistent Processes", *Review of Economic Dynamics* 13: 701-714.
- [87] Krebs, T., M. Kuhn and M. Wright (2015), "Human Capital Risk, Contract Enforcement and the Macroeconomy", *American Economic Review* 105: 3223-3272.
- [88] Krueger, D. and A. Ludwig (2016), "On the Optimal Provision of Social Insurance: Progressive Taxation versus Education Subsidies in General Equilibrium", *Journal of Monetary Economics*, forthcoming.
- [89] Krueger, D., K. Mitman and F. Perri (2016), "On the Distribution of the Welfare Losses of Large Recessions", *Proceedings of the Econometric Society World Congress 2016*, forthcoming.
- [90] Krueger, D., and F. Perri (2006), "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory", *Review of Economic Studies* 73: 163-193.
- [91] Krueger, D., and F. Perri (2011), "How Do Households Respond to Income Shocks?", *Working Paper*.
- [92] Krueger, D., F. Perri, L. Pistaferri and G. Violante (2010), "Cross-Sectional Facts for Macroeconomists", *Review of Economic Dynamics* 13: 1-14.
- [93] Krusell, P, T. Mukoyama and A. Sahin (2010), "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations", *Review of Economic Studies* 77: 1477-1507.

- [94] Krusell, P. and A. Smith (1997), "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns", *Macroeconomic Dynamics* 1: 387-422.
- [95] Krusell, P., and A. A. Smith Jr. (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy* 106: 867-896.
- [96] Krusell, P., and A. A. Smith. Jr. (2006), "Quantitative Macroeconomic Models with Heterogeneous Agents", in: *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*.
- [97] Kuhn, M. and V. Rios-Rull (2015), "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomic Modelers", *Working Paper*.
- [98] Kuznets, S. (1955), "Economic Growth and Income Inequality", *American Economic Review* 45: 1-28.
- [99] Livshits, I., J. MacGee and M. Tertilt (2007), "Consumer Bankruptcy: A Fresh Start", *American Economic Review* 97: 402-418.
- [100] Maliar, L., S. Maliar and F. Valli (2010), "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm", *Journal of Economic Dynamics and Control* 34: 42-49
- [101] McKay, A. (2015), "Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics", *Working Paper*.
- [102] McKay, A., and R. Reis (2016), "The Role of Automatic Stabilizers in the U.S. Business Cycle", *Econometrica*, 84:141-194.
- [103] Meghir, C., and L. Pistaferri (2004), "Income Variance Dynamics and Heterogeneity", *Econometrica* 72: 1-32
- [104] Mitman, K., and S. Rabinovich (2014), "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?", *Working Paper*.
- [105] Mitman, K., and S. Rabinovich (2015), "Optimal Unemployment Insurance in an Equilibrium Business-Cycle Model", *Journal of Monetary Economics* 71: 99-118.
- [106] Perri F., and J. Steinberg (2012), "Inequality and Redistribution during the Great Recession", Federal Reserve Bank of Minneapolis, Economic Policy Paper.
- [107] Peterman, W. B. (2013), "Determining the Motives for a Positive Optimal Tax on Capital", *Journal of Economic Dynamics and Control* 37: 265-295.
- [108] Piketty, T. (2014), *Capital in the Twenty-First Century*, Belknap Press, Cambridge, MA.

- [109] Piketty, T., and E. Saez (2003), "Income Inequality in the United States, 1913-1998", *Quarterly Journal of Economics*, 118:1-41.
- [110] Quadrini, V. (2000), "Entrepreneurship, Saving, and Social Mobility", *Review of Economic Dynamics* 3: 1-40.
- [111] Quadrini, V., and J.-V. Rios-Rull (2015), "Inequality in: Macroeconomics", in A. B. Atkinson and F. J. Bourguignon, eds., *Handbook of Income Distribution*, Volume 2B, North Holland, Amsterdam 1229-1302.
- [112] Ravn, M. and V. Sterk (2013), "Job Uncertainty and Deep Recessions", *Working Paper*.
- [113] Rietz., T. (1988), "The Equity Risk Premium: A Solution", *Journal of Monetary Economics* 22:117-131.
- [114] Shimer, R. (2005), "The Cyclical Behavior of Equilibrium Unemployment and Vacancies", *American Economic Review* 95: 25-49.
- [115] Skinner, J. (1987), "A Superior Measure of Consumption from the Panel Study of Income Dynamics", *Economics Letters* 23: 213-216.
- [116] Smith, M. and C. Tonetti (2014), "A Bayesian Approach to Imputing a Consumption-Income Panel Using the PSID and CEX", *Working Paper*.
- [117] Storesletten, K. C. I. Telmer and A. Yaron (2004a), "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics* 51:609-633
- [118] Storesletten, K., C. I. Telmer and A. Yaron (2004b), "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy* 112: 695-717.
- [119] Storesletten, K., C. I. Telmer and A. Yaron (2007), "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," *Review of Economic Dynamics* 10:519-548.
- [120] Wong, A. (2015), "Population Aging and the Transmission of Monetary Policy to Consumption ", *Working Paper*.
- [121] Young, E. R. (2010), "Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control* 34: 36-41.

A Data and Estimation Appendix

A.1 Aggregates in PSID and BEA

The series for disposable income from the BEA is Disposable Personal Income minus Medicare and Medicaid transfers, which are not reported in the PSID. The disposable income series from the PSID is constructed by adding, for each household and from all members, wage and salary income, income from business and farm, income from assets (including the rental equivalent for the main residence for home owners), and all money transfers minus taxes (computed using the NBER TAXSIM calculator).

The series for consumption expenditures (from both the BEA and the PSID) includes the following expenditures categories: cars and other vehicles purchases, food (at home and away), clothing and apparel, housing (including rent and imputed rental services for owners), household equipment, utilities, transportation expenses (such as public transportation and gasoline), and recreation and accommodation services. In the PSID, imputed rental services from owners are computed using the value of the main residence times an interest rate of 4%. Total consumption expenditures are reported for a two-year period because of the timing of reporting in the PSID. In the PSID, some expenditures categories (food, utilities) are reported for the year of the interview, while others are reported for the year preceding the interview, so total expenditures span a two-year period. The measure of total consumption from the BEA is constructed by aggregating the different categories using PSID timing; so, for example, total expenditures in 2004-2005 include car purchases from 2004 and food expenditures from 2005. We have excluded health services because PSID only reports out-of-pocket expenditures and insurance premia. All PSID observations are aggregated using sample weights. Table A1 reports the 2004 levels of the per capita variables plotted in figure 1, alongside, for comparison purposes, the level of food expenditures from both sources and the total household personal consumption expenditures from the BEA.

Table A1. Per capita levels in 2004: BEA v/s PSID

	BEA	PSID
1. Disposable income	\$24120	\$21364
2. Personal Consumption (PSID aggregate)	\$18705	\$15889
3. Food Expenditures	\$3592	\$2707
4. Personal Consumption (Total)	\$27642	-

Table A1 suggests that the levels from the PSID and the BEA are not too far off, although there are differences. In particular, the aggregated PSID data are different from the aggregates from the BEA for two reasons. Comparing lines 2-3 across columns, we see that for a given category, the average from the PSID is different (typically lower) than that reported by the BEA. This discrepancy between aggregate and aggregate survey data has been widely documented before. The second reason is that some cate-

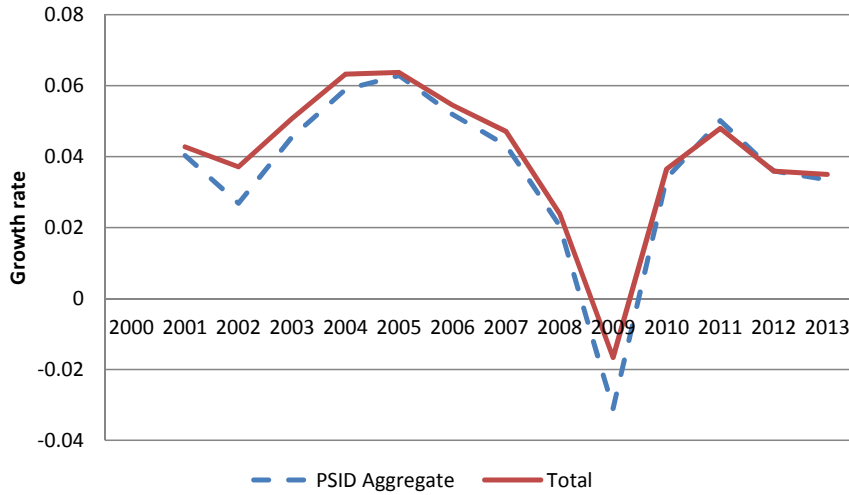


Figure A.1: BEA consumption growth for two different aggregates

gories are not included in our PSID aggregate, either because they are mismeasured in the PSID (e.g., Health expenditures) or because they are not reported by the PSID (e.g., expenditures in financial services). One might wonder whether these omitted categories matter for the aggregate pattern of expenditures. Figure A1 reports the growth rate of total household personal consumption expenditures from the BEA, along with the growth rate for the BEA consumption expenditures that are included in the PSID aggregate defined above. Table A1 suggests that categories included in the PSID aggregate cover only about 65% of the total consumption expenditures; Figure A.1 however, shows that the cyclical pattern of total expenditures is similar to the one in the PSID aggregate, suggesting that the missing consumption categories in the PSID aggregate should not make a big difference for our results.

A.2 Standard Errors and Additional Tables

Table A2. Annualized changes in variables across PSID net worth (2004-06 v/s 2006-10) with standard errors^a

	Net Worth ^b				Disp Y (%)		Cons. Exp.(%)		Exp. Rate (pp)	
	(1)		(2)		(3)	(4)	(5)	(6)	(7)	(8)
	04-06	06-10	04-06	06-10	04-06	06-10	04-06	06-10	04-06	06-10
All	15.7 (4.4)	44.6 (12.4)	-3.0 (1.6)	-10.2 (6.4)	4.1 (1.5)	1.2 (0.3)	5.6 (1.0)	-1.3 (0.5)	0.9 (0.9)	-1.6 (0.3)
Q1	NA (1.5)	12.9	NA (1.5)	6.6	7.4 (1.0)	6.7 (0.8)	7.1 (1.2)	0.6 (0.7)	-0.2 (0.9)	-4.2 (0.7)
Q2	121.9 (38.3)	19.5 (5.9)	24.4 (5.2)	3.7 (0.8)	6.7 (1.0)	4.1 (0.6)	7.2 (1.4)	2.0 (0.6)	0.3 (1.0)	-1.3 (0.4)
Q3	32.9 (3.7)	23.6 (3.1)	4.3 (1.5)	3.3 (1.1)	5.1 (0.7)	1.8 (0.4)	9.0 (4.1)	0.0 (0.7)	2.3 (2.6)	-1.1 (0.4)
Q4	17.0 (2.1)	34.7 (4.4)	1.7 (1.7)	3.8 (3.7)	5.0 (0.6)	1.7 (0.4)	5.9 (1.8)	-1.5 (0.5)	0.5 (1.1)	-2.0 (0.3)
Q5	11.6 (5.5)	132.2 (63.3)	-4.9 (1.7)	-68.4 (31.5)	1.8 (3.2)	-1.2 (0.6)	2.7 (1.7)	-3.5 (1.1)	0.5 (1.7)	-1.4 (0.8)

^aStandard errors (in parentheses) are computed using bootstrapping with 50 sample replications

^bThe first figure is the percentage change (growth rate), the second is the change in 000's of dollars. Standard errors for those figures are also in 000's of dollars

Table A3. Annualized changes in variables across PSID net worth (2006-08)

	Net Worth ^a		Disp Y (%)	Cons. Exp. (%)	Exp. Rate (pp)
All	-5.1	-17.3	2.5	-3.3	-3.6
Q1	NA	7.7	8.6	-0.7	-7.0
Q2	131.3	19.0	7.7	2.9	-3.1
Q3	18.5	13.8	3.4	-3.4	-4.2
Q4	10.4	23.0	3.2	-1.6	-3.0
Q5	-10.8	-150	-1.1	-7.3	-3.7

^aThe first figure is the percentage change (growth rate), the second is the change in 000's of dollars.

Table A4. Annualized changes in variables across PSID net worth (2008-10)

	Net Worth ^a		Disp Y (%)	Cons. Exp. (%)	Exp. Rate (pp)
All	0.5	1.3	-0.2	1.3	0.9
Q1	NA	14.7	5.4	1.8	-2.4
Q2	101.5	5.6	0.6	3.4	2.0
Q3	24.2	11.6	0.7	1.4	0.4
Q4	12.7	20.4	0.2	2.8	1.5
Q5	-4.2	-44.6	-2.6	-0.8	1.0

^aThe first figure is the percentage change (growth rate), the second is the change in 000's of dollars.

A.3 Estimation of Earnings Process for employed households

To estimate the income process for employed households, we use annual household data from the PSID from 1970 to 1997. (These are all the years the PSID survey was conducted annually and for which we can construct comparable data.) We select all households with a head between ages 25 to 60. For each household, we compute total household labor income as the sum of the labor income of the head, the labor income of the spouse, income from farm and business, plus transfers. We then compute tax liabilities for each household using the TAXSIM (ver. 9) tax calculator and subtract it from household labor income to construct household disposable labor income. We then deflate disposable income using the CPI and divide it by the number of members in the household to obtain a measure of per capita real disposable household income. We then exclude the household/years observations where the head of the household is unemployed and where the wage (computed as the head's labor income divided by the head's total hours worked) is below half the minimum wage for that year. On this sample, we regress the log of per capita real disposable income on age dummies, education dummies, interaction of age and education dummies, and year dummies. Before proceeding with estimation we exclude all household income sequences that are shorter than five years. This leaves us with our final sample of 3878 household/years sequences, of an average length of 13.1 years. On these data, we compute the first differences and then the autocovariance matrix of the first differences. We then estimate the stochastic process specified in the text using generalized method of moments, targeting the covariance matrix. The weighting matrix is the identity matrix. Many thanks to Chris Tonetti for providing the Matlab routines that perform the estimation.

B Theoretical Appendix

B.1 Explicit Statement of Aggregate Law of Motion for Distribution

Since the extent of heterogeneity and the choice problem of young and old households differ significantly, it is easiest to separate the cross-sectional probability measure Φ into two components (Φ_W, Φ_R) and note that the measures integrate to Π_W and Π_R , respectively. First define the Markov transition function, conditional on staying in the young age group $j = W$ as

$$Q_{W,(Z,\Phi,Z')}((s, y, a, \beta), (\mathcal{S}, \mathcal{Y}, \mathcal{A}, \mathcal{B})) = \sum_{s' \in \mathcal{S}, y' \in \mathcal{Y}} \begin{cases} \pi(s'|s, Z', Z)\pi(y'|y) : a'_W(s, y, a, \beta; Z, \Phi) \in \mathcal{A}, \beta \in \mathcal{B} \\ 0 \text{ else} \end{cases}$$

and for the old, retired age group, as

$$Q_{R,(Z,\Phi,Z')}((a, \beta), (\mathcal{A}, \mathcal{B})) = \begin{cases} 1 : a'_R(a, \beta; Z, \Phi) \in \mathcal{A}, \beta \in \mathcal{B} \\ 0 \text{ else} \end{cases}$$

For each Borel sets $(\mathcal{S}, \mathcal{Y}, \mathcal{A}, \mathcal{B}) \in P(\mathcal{S}) \times P(\mathcal{Y}) \times B(\mathcal{A}) \times P(\mathcal{B})$, the cross-sectional probability measures of the young and old tomorrow are then given by⁸²

$$\begin{aligned} H_W(Z, \Phi, Z')(\mathcal{S}, \mathcal{Y}, \mathcal{A}, \mathcal{B}) &= \theta \int Q_{W,(Z,\Phi,Z')}((s, y, a, \beta), (\mathcal{S}, \mathcal{Y}, \mathcal{A}, \mathcal{B})) d\Phi_W \\ &+ (1 - \nu) \mathbf{1}_{\{0 \in \mathcal{A}\}} \sum_{s' \in \mathcal{S}} \Pi_Z(s') \sum_{y' \in \mathcal{Y}} \Pi(y') \sum_{\beta' \in \mathcal{B}} \Pi(\beta') \end{aligned}$$

and

$$\begin{aligned} H_R(Z, \Phi, Z')(\mathcal{A}, \mathcal{B}) &= \nu \int Q_{R,(Z,\Phi,Z')}((a, \beta), (\mathcal{A}, \mathcal{B})) d\Phi_R \\ &+ (1 - \theta) \int Q_{W,(Z,\Phi,Z')}((s, y, a, \beta), (\mathcal{S}, \mathcal{Y}, \mathcal{A}, \mathcal{B})) d\Phi_W. \end{aligned}$$

⁸² These expressions capture the assumption that in each period, a measure $1 - nu$ of newborn households enter the economy as workers, with zero assets and with idiosyncratic productivities and discount factors drawn from the stationary distributions, and that a fraction $1 - \theta$ of working households retire, and that the retirement probability is independent of all other characteristics.

C Computational Appendix

The computational strategy follows the framework developed initially in Krusell and Smith (1998), which was further adapted by Storesletten, Telmer and Yaron (2007) and Gomes and Michaelides (2008). In particular, we employ the computational strategy outlined in Maliar, Maliar and Valli (2010), focusing on the nonstochastic simulation algorithm first introduced by Young (2010).

C.1 The individual problem

We approximate the true aggregate state ($S=(Z, \Phi)$) by \hat{S} , whose specific form depends on which version of the model we solve, which is detailed explicitly later. Thus, the household state is determined by $(s, y, a, \beta; \hat{S})$ in working life and $(a, \beta; \hat{S})$ when retired.

The solution method from Maliar, Maliar and Valli (2010) is an Euler equation algorithm that takes into account occasionally binding borrowing constraints. The problem to be solved is as follows:

Retired:

$$\begin{aligned} c_R(a, \beta; \hat{S})^{-\sigma} - \lambda &= \nu \beta \mathbb{E}[(1 - \delta + r'(\hat{S}'))c'_R(a'_R, \beta; \hat{S}')^{-\sigma}] \\ a'_R(a, \beta; \hat{S}) + c_R(a, \beta; \hat{S}) &= b_{SS}(\hat{S}) + (1 + r(\hat{S}) - \delta)a/\nu \\ a'_R(a, \beta; \hat{S}) &\geq 0 \\ \lambda &\geq 0, \quad \lambda a'_R(a, \beta; \hat{S}) = 0 \end{aligned}$$

Working:

$$\begin{aligned} c_W(s, y, a, \beta; \hat{S})^{-\sigma} - \lambda &= \theta \beta \mathbb{E}[(1 - \delta + r'(\hat{S}'))c'_W(s', y', a'_W, \beta; \hat{S}')^{-\sigma}] \\ &\quad + (1 - \theta) \beta \mathbb{E}[(1 - \delta + r'(\hat{S}'))c'_R(a'_R, \beta; \hat{S}')^{-\sigma}] \\ a'_W(s, y, a, \beta; \hat{S}) + c(s, y, a, \beta; \hat{S}) &= (1 - \tau(Z; \rho))w(\hat{S})y[1 - (1 - \rho)1_{s=u}] + (1 + r(\hat{S}) - \delta)a \\ a'_W(s, y, a, \beta; \hat{S}) &\geq 0 \\ \lambda &\geq 0, \quad \lambda a'_W(s, y, a, \beta; \hat{S}) = 0, \end{aligned}$$

where λ is the Lagrange multiplier on the borrowing constraint.

We eliminate consumption via the budget constraint and then guess a policy rule for $a'_W(s, y, a, \beta; \hat{S})$ and $a'_R(a, \beta; \hat{S})$. We then substitute the policy rule to compute $a''_W(s', y', a'_W, \beta; \hat{S}')$, $a''_R(a'_R, \beta; \hat{S}')$ and $a'_R(a'_R, \beta; \hat{S}')$, and use the Euler equation to back out the implied policy rule for a' . If the implied policy rule is the same as the conjectured policy rule, we have computed the optimal policy; if not, we update the guess and repeat.

C.2 The simulation algorithm

In order to simulate the model, we pick a grid on \mathcal{A} and fix a distribution of workers $\Phi_0 \in S \times Y \times A \times B$ space. We fix a long time series for the realization of the aggregate shock, Z . Using the realization Z_t and Φ_t , we can compute \hat{S}_t and then apply the policy rules from the individual problem and the Markov transition matrices associated with s and y to compute Φ_{t+1} by interpolating onto the grid points in \mathcal{A} .

C.3 Approximating the Aggregate Law of Motion

C.3.1 KS and Benchmark Economies

For the KS and benchmark economies, we approximate the true aggregate state with $\hat{S}_t = (Z, \bar{K}_t)$ where \bar{K}_t is the average capital in the economy. Agents need to forecast the evolution of the capital stock. We conjecture that the law of motion in capital depends only on the Z and \bar{K} :

$$\log(\bar{K}_{t+1}) = a_0(Z_t) + a_1(Z_t) \log(\bar{K}_t)$$

We conjecture coefficients a_0 and a_1 , solve the household problem, and simulate the economy. Then, using the realized sequence of \hat{S} , we perform the previous regression and check whether the implied coefficients are the same as the conjectured ones. If they are, we have found the law of motion; if not, we update our guess and repeat.

For the KS economy, the computed law of motion is as follows:

$$\begin{aligned} \log(\bar{K}_{t+1}) &= 0.1239 + 0.9652 \log(\bar{K}_t) && \text{if } Z_t = Z_l \\ \log(\bar{K}_{t+1}) &= 0.1334 + 0.9638 \log(\bar{K}_t) && \text{if } Z_t = Z_h. \end{aligned}$$

The R^2 for both regressions are in excess of 0.999999. Note, however, that Den Haan (2010) points out that despite having large R^2 values, the accuracy of the solution can still be poor, and suggests simulation of the capital stock under the policy rule and comparing it with the capital stock that is calculated by aggregating across the distribution. We do this for 3000 time periods. The average error between the implied law of motion from the forecast equations and the computed law of motion is 0.02%, with a maximum error of 0.10%.

For the benchmark economy, the computed law of motion is as follows:

$$\log(\bar{K}_{t+1}) = 0.0924 + 0.9716 \log(\bar{K}_t) \quad \text{if } Z_t = Z_l$$

$$\log(\bar{K}_{t+1}) = 0.0929 + 0.9723 \log(\bar{K}_t) \quad \text{if } Z_t = Z_h$$

The R^2 for both regressions are in excess of 0.99999. Similar to the previous computation, we check the accuracy of the law of motion. We find that the average error between the implied law of motion and the actual capital stock computed from the distribution is 0.01%, with a maximum error of 0.07%.

C.3.2 Consumption Externality Economy

In the economy with the aggregate consumption externality, we add contemporaneous consumption as a state variable in our approximation of the true aggregate state, $\hat{S} = (Z, \bar{K}, C)$. We therefore need an additional law of motion for how aggregate consumption evolves. We conjecture the same form of law of motion for the average capital stock; however, we allow the evolution of aggregate consumption to depend on both the average capital stock and aggregate consumption:

$$\begin{aligned} \log(\bar{K}_{t+1}) &= a_0(Z_t) + a_1(Z_t) \log(\bar{K}_t) \\ \log(C_{t+1}) &= b_0(Z_t, Z_{t+1}) + b_1(Z_t, Z_{t+1}) \log(\bar{K}_t) + b_2(Z_t, Z_{t+1}) \log(C_t). \end{aligned}$$

Note that because capital is predetermined in the current period, the forces rule for capital depends only on contemporaneous variables. Because aggregate consumption is an equilibrium outcome in the next period, we allow for the forecast to depend on the subsequent period's realization of the Z shock. Thus, there are four sets of coefficients to be estimated for the law of motion for consumption. The computed forecast equations are as follows:

$$\begin{aligned} \log(\bar{K}_{t+1}) &= 0.0872 + 0.9736 \log(\bar{K}_t) & \text{if } Z_t = Z_l \\ \log(\bar{K}_{t+1}) &= 0.0626 + 0.9816 \log(\bar{K}_t) & \text{if } Z_t = Z_h \end{aligned}$$

and

$$\begin{aligned} \log(C_{t+1}) &= -0.0205 + 0.0023 \log(\bar{K}_t) + 0.9675 \log(C_t) & \text{if } (Z, Z') = (Z_l, Z_l) \\ \log(C_{t+1}) &= -0.5061 + 0.2882 \log(\bar{K}_t) + 0.5297 \log(C_t) & \text{if } (Z, Z') = (Z_l, Z_h) \\ \log(C_{t+1}) &= -0.3560 + 0.1893 \log(\bar{K}_t) + 0.6626 \log(C_t) & \text{if } (Z, Z') = (Z_h, Z_l) \\ \log(C_{t+1}) &= -0.0506 + 0.0360 \log(\bar{K}_t) + 0.9295 \log(C_t) & \text{if } (Z, Z') = (Z_h, Z_h) \end{aligned}$$

with R^2 in excess of 0.9999, 0.9999999, 0.9999, 0.9999, 0.99999, 0.99999, respectively. As before, we check the accuracy of the two laws of motion. We find that the average error between the implied law of motion and the actual capital stock computed from the distribution is 0.02%, with a maximum error of 0.30%, and for the path of aggregate consumption the mean error is 0.02% with a maximum error of 0.24%. Although the externality economy has slightly larger forecast errors, the fit of the predicted aggregates is still excellent.

C.4 Digression: Why Quasi-Aggregation?

One of the implications of the results in the main text is that the wealth distribution (and especially the fraction of the population with little or no wealth) is quantitatively important for the macroeconomic consumption and investment response to an aggregate technology shock. This, however, does not imply that Krusell and Smith's (1998) original quasi-aggregation result fails.⁸³ Recall that this result states that only the mean of the current wealth distribution (as well as the current aggregate shock Z) is required to accurately predict the future capital stock and therefore future interest rates and wages.

The previous experiment compared consumption and investment dynamics in *two economies* that differed substantially in their wealth distributions. For a given economy, if the wealth distribution does not move significantly in response to aggregate shocks, then it would be irrelevant for predicting future aggregates and prices. However, in the high-wealth-inequality economy, the wealth distribution *does* move over the cycle. For example, the share of households at the borrowing constraint displays a coefficient of variation of 7%. However, what is really crucial for quasi-aggregation to occur is whether the movement, over the cycle, in the key features of the wealth distribution is explained well by movements in Z and K , the state variables in the forecast equations of households. We find that it is, even in the high-wealth-inequality economy.

For example, if we regress the fraction of people at the borrowing constraint tomorrow on Z in simulated data, we obtain an R^2 of around 0.8. Therefore, the vast majority of the variation in households at the borrowing limit is very well predicted by the aggregate state variables (Z, K). This finding is robust to alternative definitions of constrained households (households exactly at wealth 0, households who save less than 1%, less than 10%, or less than 25% of the quarterly wage) and alternative moments of the wealth distribution. It is this finding that makes quasi-aggregation hold, despite the strong impact of the wealth distribution on the aggregate consumption and investment response to aggregate technology shocks.

⁸³ In fact, our computational method that follows theirs rather closely relies on quasi-aggregation continuing to hold.

C.5 Recovering the Value Function

As we solve the model by exploiting the Euler equation, in order to perform the welfare calculations in section 6.4, we need to recover the value functions as a function of the idiosyncratic and aggregate states. To calculate them, we use policy function iteration. We make an initial guess for the value function, v^0 , then calculate v^1 by solving the recursive household decision problem (we need not perform the maximization, since we have already computed the optimal policy function). We approximate the value function with a cubic spline interpolation in assets, as well as in aggregate capital (and for the demand externality model, we also aggregate consumption). If v^1 is sufficiently close to v^0 (in the sup-norm sense), we stop; otherwise, we proceed to compute v^2 taking v^1 as the given value function. We proceed until convergence. For the economies with retirement, we first recover the value function for retired households, v_R , and then proceed to recover the value function for working-age households, v_W .