



**On the Distribution of the Welfare Losses
of Large Recessions**

Dirk Krueger

University of Pennsylvania,
CEPR, CFS, NBER, and Netspar

Kurt Mitman

IIES, Stockholm University, and CEPR

Fabrizio Perri

Federal Reserve Bank of Minneapolis
and CEPR

Staff Report 532

July 2016

Keywords: Great Recession; Wealth inequality; Social insurance; Welfare loss from recessions

JEL classification: E21, E32, J65

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Federal Reserve Bank of Minneapolis • 90 Hennepin Avenue • Minneapolis, MN 55480-0291

<https://www.minneapolisfed.org/research/>

On the Distribution of the Welfare Losses of Large Recessions*

July 2016

Dirk Krueger

*University of Pennsylvania,
CEPR, CFS, NBER and Netspar*

Kurt Mitman

*IIES, Stockholm University and
CEPR*

Fabrizio Perri

*Federal Reserve Bank of
Minneapolis and CEPR*

Abstract

How big are the welfare losses from severe economic downturns, such as the U.S. Great Recession? How are those losses distributed across the population? In this paper we answer these questions using a canonical business cycle model featuring household income and wealth heterogeneity that matches micro data from the Panel Study of Income Dynamics (PSID). We document how these losses are distributed across households and how they are affected by social insurance policies. We find that the welfare cost of losing one's job in a severe recession ranges from 2% of lifetime consumption for the wealthiest households to 5% for low-wealth households. The cost increases to approximately 8% for low-wealth households if unemployment insurance benefits are cut from 50% to 10%. The fact that welfare losses fall with wealth, and that in our model (as in the data) a large fraction of households has very low wealth, implies that the impact of a severe recession, once aggregated across all households, is very significant (2.2% of lifetime consumption). We finally show that a more generous unemployment insurance system unequivocally helps low-wealth job losers, but hurts households that keep their job, even in a version of the model in which output is partly demand determined, and therefore unemployment insurance stabilizes aggregate demand and output.

Keywords: Great Recession, Wealth Inequality, Social Insurance

JEL Classifications: E21, E32, J65

*Thanks to Tony Smith for an excellent discussion and to Simone Civalo for outstanding research assistance. Krueger gratefully acknowledges financial support from the NSF under grants SES 1123547 and SES 1326781. Mitman gratefully acknowledges financial support from the Mannerfelt Foundation. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The objective of this paper is to quantify the distribution of welfare losses across households induced by a severe economic downturn of the magnitude of the U.S. Great Recession. As a laboratory for our analysis, we use an augmented version of the canonical Krusell-Smith (1998) real business cycle model with household income and wealth heterogeneity, as presented in Krueger, Mitman and Perri (2016). The model features business cycles driven by productivity shocks in an economy populated by agents that face different types of idiosyncratic income shocks and accumulate wealth in order to finance consumption during retirement and to self-insure against idiosyncratic income risk. In this framework, a recession affects households in several ways. First, in an economic downturn more households find themselves without a job, and a job loss reduces their lifetime income, consumption and welfare. Households that do not lose their job at the onset of the recession also suffer welfare losses, as their wages will fall and they have a higher probability of losing their job in the future, during the course of the long-lasting recession. The main focus of this paper is to document the size and heterogeneity of these losses across the income and wealth distribution, as well as to investigate the extent to which social insurance policies, such as stylized unemployment insurance, affect these losses.

In order to render this analysis empirically meaningful, we first measure, using micro data from the Panel Study of Income Dynamics (PSID), the extent of income, wealth, and consumption inequality in 2006, prior to the Great Recession. After calibrating the model to macroeconomic and microeconomic data, we first evaluate how well it captures the stylized cross-sectional facts, and then use the model as a laboratory for characterizing and quantifying the distribution of welfare losses from experiencing a severe aggregate economic downturn in which unemployment risk rises and household incomes (even conditional on not losing a job) fall.

Our first set of results shows that the welfare losses vary strongly with characteristics of households such as wealth, income, employment status, and household preferences (their patience, more specifically). In our benchmark economy, the welfare losses range from over 5%, for a low-wealth household that loses its job at the onset of the recession, to less than 2% for a high-wealth household, that the onset of the recession keeps its job.

Our second set of results uses the equilibrium distribution of households over the relevant individual household states to integrate welfare losses across households, and obtains a measure of the aggregate welfare losses.¹ We find that in our benchmark economy the aggregate (under the veil of ignorance) welfare cost of the recession is 2.2% and that almost 2% of the working-age households experience a loss exceeding

¹ As we will argue below, this measure can alternatively be interpreted as the expected welfare loss, under the veil of ignorance of where in the distribution a household is placed just prior to the onset of a severe recession.

3% of lifetime consumption.

Our final set of results studies how the welfare costs discussed above change as we vary features of the benchmark economy. One important lesson is that the shape of wealth distribution, especially at the bottom, has a large impact on the aggregate welfare impact of a recession as well as its distribution. In order to demonstrate this, we repeat our welfare calculations with an economy studied by Krusell and Smith (1998), which displays significantly less wealth inequality than our benchmark, and, in particular, displays a much smaller fraction of households at the bottom of the wealth distribution. In that economy, we find that the aggregate welfare cost the recession falls to 1.6% (versus 2.2% in the benchmark) and that virtually no households experience a loss that exceeds 3% of lifetime consumption.

We then study the impact of a stylized public unemployment insurance system and find that reducing the generosity of unemployment insurance matters significantly for welfare, even when the wealth distribution is allowed to change in response to the policy change. In particular, we first reduce the unemployment insurance replacement rate from 50% (our benchmark value) to 10% and let the wealth distribution change in response to the lower replacement rate. In this case, the aggregate welfare costs from a recession are similar to the benchmark, but the fraction of households that experience losses bigger than 3% increases significantly from 2% to 3.1%. When we reduce unemployment insurance without allowing the wealth distribution to respond immediately (an experiment we label "unemployment insurance shock", which can be interpreted as an unexpected expiration of unemployment insurance benefits), we find that the aggregate welfare costs increase only moderately, but the fraction of households that loses more than 3% of lifetime consumption jumps dramatically, to 7.5% of the population.

We finally study how unemployment insurance affects the welfare costs from a recession in an economy in which output is partially determined by aggregate demand.² In such an economy, one might think that unemployment insurance, by stabilizing aggregate demand, would stabilize output and thus reduce overall welfare losses. However, it turns out that even though more generous unemployment insurance indeed mitigates the collapse in aggregate consumption, output, and wages in the short run and reduces the welfare losses of the job losers (just as it did in the benchmark economy), larger unemployment insurance benefits increase the welfare losses of those who do not lose their job in a recession. The key intuition behind this finding is that more generous unemployment insurance reduces the incentives for precautionary saving in the recession. In the economy without the consumption externality, the reduction in saving (relative to the low unemployment insurance scenario) has a relatively small impact on factor prices and thus on overall welfare. In the economy with the consumption externality, however, the reduction in saving, by reducing the capital stock

² We achieve this by introducing in the model an externality from aggregate consumption to Total Factor Productivity (TFP)

and consumption in the medium run, affects negatively the evolution of TFP in the medium and long run. We find that, quantitatively, this effect can be quite large, and outweigh the benefits that come from short run demand stabilization for all but the most impatient households. This finding suggests, more broadly, that even if output is demand determined and unemployment insurance stabilizes the economy in the short run, it is not necessarily the case that such a policy helps both the unemployed and the employed (even though the latter might well become unemployed during the course of a severe and long-lasting recession).

This work is part of a broader research agenda 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 welfare costs of macroeconomic fluctuations) 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").³

Our work is also related to a large body of research that studies the welfare impact of aggregate fluctuations. In a seminal contribution, Lucas (1987) calculates an upper bound on the welfare costs of business cycle fluctuations in a representative household economy. His hypothetical thought experiment compares an aggregate consumption path with empirically observed volatility to one with the same mean but without fluctuations and asks what percentage of lifetime consumption households would be willing to give up to have the random consumption process replaced by its deterministic mean process. This gain is, in that context, the most a costless and perfectly effective stabilization policy can hope to achieve. Lucas finds that the so-calculated welfare loss from business cycles is minuscule, on the order of less than 1/100 of 1%. The conclusion from this calculation is that aggregate stabilization policy, whether fiscal or monetary policy, has only very limited potential for improving welfare of the representative household, even if it were perfectly effective and costless.

A substantial literature has revisited the Lucas thought experiment in a variety of models, including those with household heterogeneity, in which actors in the economy face both idiosyncratic and aggregate risk.⁴ Representative contributions include Imrohorglu (1989), Krusell and Smith (1999), Chatterjee and Corbae (2007), Krebs (2007), and Krusell, Mukoyama, Sahin, and Smith (2009) and has concluded that the welfare costs of business cycles are a) strongly heterogeneous across the population, with wealth-poor households suffering the most, and b) for these households, the costs can be large at least one order of magnitude (i.e. ten times) as large as the original Lucas numbers.

In this paper, we ask a related, but conceptually different question. Rather than assess-

³ For a more systematic review of these literatures see Krueger, Mitman and Perri, 2016

⁴ See Lucas (2003) for an early summary of this literature.

ing the benefits of removing all cycles from an *ex ante* point of view, we document the welfare losses of different households from actually experiencing a recession in the current period.⁵ Conceptually we follow the papers by Glover, Heathcote, Krueger and Rios-Rull (2014) as well as Hur (2014) in quantifying the welfare losses of *experiencing* a recession in this way. In the same spirit as Lucas (1987) we do not actually study concrete policies that would have prevented the Great Recession, but rather assess how valuable such a hypothetical policy would be if it existed and could be implemented costlessly, thereby quantifying how painful it is to experience a severe recession.⁶

The paper is organized as follows. The next section summarizes the cross-sectional inequality facts from the PSID that motivate and discipline our model-based analysis. Section 3 sets out the model and summarizes its calibration.⁷ In section 4, we assess how well the model fits the cross-sectional facts, and compare its cross-sectional distribution with that of the original Krusell and Smith (1998) economy. Section 5 contains our main welfare results; it measures how large, and how dispersed the welfare losses from a Great Recession-like economic downturn are, and how the magnitudes of these losses depend on the presence and size of a publicly provided unemployment insurance system. Section 6 concludes and technical details about the welfare cost measures and the definition of a recursive competitive equilibrium are contained in the appendix.

2 The Wealth, Income and Consumption Distribution Prior to the Great Recession

In this section, our aim is two-fold. First we document the extent of income, wealth and consumption inequality at the eve of the Great Recession. A plausible model aimed at measuring the distribution of the welfare losses from a large aggregate economic downturn needs to be consistent with these pre-recession cross-sectional facts. Second, our data set of choice, the PSID contains comprehensive information on household earnings, income, consumption expenditures and wealth. These rich data allow us to empirically characterize the *joint* distribution of these variables, which is in turn informative about the plausibility of competing mechanisms that are suitable for gen-

⁵ We share the focus on idiosyncratic unemployment risk in rare but severe recessions with Chatterjee and Corbae (2007).

⁶ Krueger and Perri (2004) use a similar approach to the one adopted here to quantify the welfare losses from the observed increase in *earnings inequality* in the last 25 years.

⁷ In the interest of space, the details of the calibration are in Krueger, Mitman and Perri (2016). That paper shares the broad theme (the interaction between the dynamics of macroeconomic variables and household heterogeneity in earnings, income, consumption and wealth), but focuses on the impact of micro heterogeneity for the dynamics of aggregate consumption, investment and output in the Great Recession, rather than studying the normative question of the distribution of welfare losses from the Great Recession.

erating the empirically observed wealth dispersion.

2.1 The Cross-Sectional Distributions of Earnings, Disposable Income, Consumption, and Wealth

The main focus of our welfare analysis is on households that face labor income risk more generally, and unemployment risk specifically. Therefore in our empirical analysis we restrict attention to households with heads between ages 25 and 60. For these households, table 1 reports, for each variable of interest in the PSID (earnings, disposable income, consumption expenditures and net worth), the 2006 cross-sectional average as well as the share of the total value held by each of the five quintiles of the corresponding distribution. In addition, to obtain a more precise picture of earnings, income and wealth concentration, the table also reports the Gini index as well as the share held by the households between the 90th and 95th percentile, between the 95th and 99th percentile, and by those in the top 1% of the respective distribution.⁸ Finally, and for comparison, the table contains information about the income, wealth and consumption distribution from alternative data sets.⁹

The marginal distributions of earnings, disposable income, consumption expenditures and net worth have properties that are well-known from other studies and data sets. Namely, all distributions display a fairly high degree of concentration, with the top 20% of households controlling close to 50% of earnings, income and consumption, and more than 80% of net worth, whereas the bottom quintile accounts for no wealth, less than 5% of earnings (income) and about 6% of consumption expenditures. Comparing across variables, we see that the distributions of earnings and disposable income for our sample look fairly similar, since capital income constitutes only about 1/6 of disposable income for these households (ages 25 to 60).

In the class of heterogeneous-agent models we will study in the next section, net worth is the key endogenous state variable. Thus, any model aimed at capturing the cross-sectional distribution of welfare losses from economic downturns should feature an empirically plausible wealth distribution. As table 1 shows, this requires generating a wealth distribution that is substantially more concentrated than that of earnings, in-

⁸ The earnings variable includes all sources of labor income plus non-UI transfers minus tax liabilities. Disposable income encompasses 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 PSID plus the rental equivalent of the main residence. Finally wealth is measured as net worth, the value of all financial and real assets (including owner-occupied houses) net of all household liabilities.

⁹ For disposable income from the 2006 Current Population Survey (CPS), for consumption from the 2006 Consumer Expenditure Survey (CE) and for net worth from the 2007 Survey of Consumer Finances (SCF), which, because of oversampling of rich households, paints a more accurate empirical picture for the very top of the distribution of net worth in the data.

Table 1: Means and Marginal Distributions in 2006

Source	Variable						
	Earn.	Disp Y		Cons. Exp		Net Worth	
	PSID	PSID	CPS	PSID	CE	PSID	SCF(07)
Mean (06\$)	52,783	62,549	60,032	43,980	47,563	291,616	497,747
% 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	0.42	0.42	0.40	0.40	0.36	0.77	0.78
Sample	6,232	6232	54,518	6,232	4,908	6,232	2,910

come and consumption – a distribution in which the bottom 40% of household hold virtually no wealth and more than 80% is held by the top quintile. Comparing the last two columns of the table, it is noteworthy that the net worth distribution from the PSID lines up fairly well with that obtained from the SCF (although it misses a significant share of wealth at the very top of the wealth distribution, resulting in lower average wealth per household in the PSID, relative to the SCF).

2.2 Income and Consumption by Wealth Quintiles

The previous section ranked households separately by income, consumption, and wealth and characterized the empirical marginal distributions. Although useful for descriptive purposes, economic theory imposes restrictions on the joint distribution of these variables. In recursive formulations of heterogeneous household models wealth – in addition to current earnings – acts as the crucial individual state variable, and thus in table 2 we highlight the correlation of net worth with earnings, income and especially consumption at the individual household level. We divide the 2006 PSID sample into net worth quintiles, and then report the share of earnings, income and consumption accruing to each *net worth quintile*. We also calculate, in the last two rows of the table, the *expenditure rates* of the different wealth quintiles, by dividing consumption expen-

ditures of the respective wealth quintile by earnings (income) of the wealth quintile.

Table 2: Earnings, Disposable Income and Expenditures by Net Worth in 2006

Quintile(NetW)	% Share of:			% Expend. Rate	
	Earn.	Disp Y	Expend.	Earn.	Disp Y
Q1	9.8	8.7	11.3	95.1	90.0
Q2	12.9	11.2	12.4	79.3	76.4
Q3	18.0	16.7	16.8	77.5	69.8
Q4	22.3	22.1	22.4	82.3	69.6
Q5	37.0	41.2	37.2	83.0	62.5
	Correlation with net worth				
	0.26	0.42	0.20		

From table 2 we observe that earnings, income and consumption at the household level are all substantially, but far from perfectly, correlated with net worth (see the last row for the correlation coefficients). Note, however, that the correlation of consumption expenditures with net worth is weaker than for earnings or income. As the last two columns show, low-wealth households have systematically higher consumption rates than high-wealth households. These differences in consumption expenditure behavior across the wealth distribution are economically significant: the gap between the savings rate (measured as 1 minus the consumption rate) of the top and the bottom wealth quintiles amounts to 15 to 30 percentage points.

Two key facts from tables 1 and 2 emerge: wealth is highly concentrated, with the bottom two quintiles of the wealth distribution accounting for very little of aggregate wealth; at the same time, these households hold about 20% of aggregate labor earnings and income, and are responsible for $11.3\% + 12.4\% = 23.7\%$ of total consumption expenditures. In Krueger, Mitman and Perri (2016) we argue that these observations about the cross-sectional wealth distribution and the joint distribution of wealth and consumption are key for the argument that household heterogeneity matters, in a quantitatively significant way, for the dynamics of aggregate consumption in a severe recession.

In this paper, we will show that it is precisely the households at the bottom of the wealth distribution with low savings rates and high propensities to consume out of current income that suffer the largest welfare losses from a severe recession. Further, these losses are much more severe than those sustained by the "average" household in society. In order to make this argument, we first develop a macroeconomic business cycle model with household heterogeneity and ensure that the model is consistent with the two key stylized facts from the data discussed above. In section 3 we lay out the model and briefly discuss how we calibrate it, in section 4 we discuss the fit of the model vis-a-vis the PSID micro data, and section 5 then contains our welfare results.

3 Household Heterogeneity and the Macro Economy: Model and Calibration

In this section we lay out the canonical real business cycle model with household income and wealth heterogeneity that we use as a measurement tool and laboratory for counterfactual policy analysis. It builds on the seminal heterogeneous-household general-equilibrium models of Bewley (1986), Imrohoroglu (1989), Huggett (1993, 1997), Aiyagari (1994) and especially Krusell and Smith (1998), but modifies the latter along several important dimensions. The objective of these modifications is two-fold. First, in order to obtain a better fit of the cross-sectional household income and wealth distributions, we augment the model's idiosyncratic unemployment shock process with a stochastic process for earnings, conditional on being employed. Further, we allow for preference heterogeneity, and we incorporate basic life cycle elements and social insurance policies into the model.¹⁰ Second, in order for these policies to potentially have some aggregate demand stabilizing effects, and to make contact with business cycle models stressing demand-determined aggregate fluctuations more broadly, we also consider a variant of the model with an aggregate consumption demand externality.

3.1 Aggregate Risk

As in Krusell and Smith (1998) and the real business cycle literature more generally (see e.g. Kydland and Prescott, 1982), the ultimate source of aggregate fluctuations are exogenous stochastic movements in total factor productivity, denoted as Z^* . Total factor productivity Z^* is determined by

$$Z^* = ZC^\omega \tag{1}$$

where the exogenous part of TFP, Z follows a first order Markov process with state space $Z \in Z_l, Z_h$ and transition matrix $\pi(Z'|Z)$. When bringing the model to the data we will interpret Z_h as normal times and Z_l as a severe recession (such as the Great Recession). In equation (1) the term C is aggregate consumption, and the parameter $\omega \geq 0$ governs the strength of the aggregate demand externality. For much of the paper, and in our benchmark model calibration, we will study the case of $\omega = 0$. In that case, as in the standard Real Business Cycle literature, TFP simply follows an exogenous process and there is no distinction between Z and Z^*). If instead $\omega > 0$, current TFP, and thus output, is partially demand-determined, and, by stabilizing aggregate consumption, demand social insurance policies might mitigate the severity of a recession, and thus moderate the welfare losses stemming from it.

¹⁰ On the household side, the model shares its key features with the recent study by Carroll, Slacalek, Tokunaka and White (2015).

3.2 Technology

Given the process for TFP, output in the economy is produced according to a standard neoclassical aggregate production function:

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

Capital used in production in turn depreciates at a constant rate $\delta \in [0, 1]$.

3.3 Household Demographics, Endowments and Preferences

3.3.1 Demographics and the Life Cycle

The economy is populated by a unit mass of potentially infinitely lived households that differ by age, preferences, unemployment status, earnings and wealth. We model the life cycle of households in a parsimonious way, by assuming that households are either of working age (indexed by W) and participate in the labor market or are retired (denoted by R). A household's age is denoted by $j \in \{W, R\}$. Young households transit into retirement with constant probability $1 - \theta \in [0, 1]$. Old households die with constant probability $1 - \nu \in [0, 1]$ and are replaced by new young households. By assumption, the realizations of the retirement and the mortality shock are orthogonal to all other household characteristics, and thus the parameters θ, ν completely determine the time-invariant distribution of households across the two stages of the life cycle, given by

$$\Pi_W = \frac{1 - \theta}{(1 - \theta) + (1 - \nu)}$$
$$\Pi_R = \frac{1 - \nu}{(1 - \theta) + (1 - \nu)}$$

We explicitly model a life cycle (albeit rudimentary) of households to induce a retirement savings motive, which in turn, as we will argue below, helps the model to deliver a more plausible consumption-savings behavior across the wealth distribution that eludes a model where savings are accumulated for purely precautionary motives.

3.3.2 Preferences

Households have time separable preferences representable by von Neumann-Morgenstern expected utility. The period felicity function is given by $u(c)$ and is continuous, strictly increasing, strictly concave and satisfies the Inada conditions. Households discount the future with a time discount factor β that is fixed over time for a given household, but might vary across households. Since Krusell and Smith's (1998) original paper, preference heterogeneity in time discount factors is known to be an effective tool for

helping to generate a cross-sectional wealth distribution as dispersed as in the data. The set of possible time discount factors is given by $\beta \in B$, the finite set of possible time discount factors, and $\pi(\beta)$ denotes the share of households with time discount factor β . By assumption (of a law of large numbers), $\pi(\beta)$ is both the individual probability distribution from which a newborn household draws its β as well as the cross-sectional distribution of time discount factors in the population.¹¹

3.3.3 Endowments

In the macroeconomics literature that models household heterogeneity explicitly, the key source of such heterogeneity stems from labor income. Following much of this literature ever since Bewley (1986), we model household labor income as a stochastic process with idiosyncratic shocks. Earnings risk has two sources: unemployment risk and labor income risk conditional on being employed. We now describe both sources in turn, starting with unemployment risk. The inclusion of this risk not only is empirically relevant and helps us to connect our results to the original Krusell and Smith (1998) paper, but also allows us to discuss unemployment insurance as one example of public social insurance policies. Let $s \in S = \{u, e\}$ denote the current employment status of a household, where $s = u$ stands for unemployment. We assume that this employment status follows a Markov chain with transitions $\pi(s'|s, Z', Z)$. The dependence of unemployment-employment transitions on the aggregate state transitions for TFP captures the fact that job finding and job losing rates depend on the state of the business cycle.

For the idiosyncratic unemployment shock, we assume a law of large numbers, so that idiosyncratic risk averages out, and only aggregate risk determines the number of unemployed households. Furthermore, we assume that the share of households in a given idiosyncratic employment state s depends only on the *current* aggregate state¹² 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 .

Unemployment risk alone is, of course, insufficient to generate an empirically plausible labor earnings distribution. Thus, following much of the literature since Aiyagari (1994), we will estimate a stochastic labor income process, conditional on being employed, from the PSID data and embed it into the model as a second source of idiosyncratic income risk.¹³

¹¹ The fact that households die with positive probability ensures ergodicity of the cross-sectional wealth distribution even in the presence of permanent time discount factor heterogeneity across households.

¹² This assumption imposes consistency restrictions on the transition matrix $\pi(s'|s, Z', Z)$ and prevents the unemployment rate from becoming a continuous aggregate variable. The restrictions are discussed in detail in the original Krusell and Smith (1998) paper.

¹³ As described below, we discretize the continuous state space process into a finite Markov chain before incorporating it into the model.

Therefore, conditional on being employed, labor income is the product of an aggregate wage w per labor efficiency unit and an idiosyncratic part y best interpreted as individual labor productivity. The idiosyncratic component $y \in Y$ follows a first order Markov chain with transition matrix $\pi(y'|y) > 0$. Let $\Pi(y)$ denote the associated unique invariant distribution; by assumption, this distribution does not depend on the aggregate state Z .¹⁴ As with unemployment risk, we assume that idiosyncratic labor productivity risk washes out in the aggregate because of a law of large numbers.

3.3.4 Trading Opportunities and Initial Conditions

Households can save (but are prevented from borrowing) by accumulating (risky) physical capital and are assumed to have access to perfect annuity markets in retirement. The capital of the deceased is used to pay an extra return on capital $\frac{1}{v}$ of the survivors. 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 and 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.

3.4 Government Policy

In our economy the government organizes two social insurance programs: a pure pay as you go (PAYGO) social security system and an unemployment insurance system. The purpose of the presence of the former is simply to ensure that the private life cycle savings motive has a plausible magnitude. In contrast, the role of the stylized public unemployment insurance for the cross-sectional wealth distribution and for the welfare losses from great recessions is an important aspect of our analysis and will be investigated in greater detail below.

3.4.1 Unemployment Insurance

The government uses proportional taxes on labor income (and unemployment benefits) to finance unemployment benefits $b(y, Z, \Phi)$ that potentially depend on the aggregate state of the economy (Z, Φ) as well as the potential earnings wy of a household.¹⁵

¹⁴ Even for the unemployed, potential labor productivity evolves according to the same process, and it determines the productivity upon finding a job, as well as unemployment benefits while being unemployed, as described in the next section.

¹⁵ Recall that even unemployed households carry with them the idiosyncratic state y .

The size of the system is indexed by the replacement rate ρ , the key policy parameter that we will vary in our counterfactual policy experiments. Thus, benefits of a currently unemployed person are given by

$$b(y, Z, \Phi) = \rho w(Z, \Phi)y \quad (3)$$

Budget balance of the unemployment insurance system requires

$$\Pi_Z(u) \sum_y \Pi(y) b(y, Z, \Phi) = \tau(Z, \Phi) \left[\sum_y \Pi(y) [\Pi_Z(u) b(y, Z, \Phi) + (1 - \Pi_Z(u)) w(Z, \Phi)y] \right] \quad (4)$$

and using equation (3) and the fact that the distribution over y is identical among the employed and unemployed, we determine the required tax rate to be

$$\tau(Z, \Phi; \rho) = \left(\frac{\Pi_Z(u)\rho}{1 - \Pi_Z(u) + \Pi_Z(u)\rho} \right) = \left(\frac{1}{1 + \frac{1 - \Pi_Z(u)}{\Pi_Z(u)\rho}} \right) = \tau(Z; \rho) \in (0, 1) \quad (5)$$

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

3.4.2 Social Security

The government runs a balanced budget, pure PAYGO social security system characterized by a constant payroll tax rate τ_{SS} on labor income and unemployment benefits. Social security benefits $b_{SS}(Z, \Phi)$ of retired households are assumed to be independent of past contributions, reflecting in a stylized way the highly redistributive nature of the current U.S. system, but also avoiding the need to track past earnings or contributions as additional continuous state variables in the household maximization problem. Given these assumptions, the budget constraint of the social security system then reads as

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

The social security replacement rate is an intuitive function of the payroll tax rate τ_{SS} , the old age dependency ratio $\frac{\Pi_W}{\Pi_R}$, and average labor productivity in the economy:

$$\frac{b_{SS}(Z, \Phi)}{w(Z, \Phi)} = \tau_{SS} \frac{\Pi_W}{\Pi_R} [\Pi_Z(u)\rho + (1 - \Pi_Z(u))]$$

Note that both social security benefits and wages, and therefore the replacement rate, will fluctuate with the aggregate state of the economy Z , although the quantitative importance of these fluctuations is minor.

3.5 Recursive Competitive Equilibrium

In this section, we formally define a recursive competitive equilibrium for this economy with cross-sectional household heterogeneity and aggregate fluctuations. The aggregate state of the economy at every point in time is characterized by the aggregate productivity shock Z and the cross-sectional distribution Φ over individual household characteristics (s, y, a, β) . This high-dimensional state space is the source of computational challenges when computing recursive competitive equilibria, and we follow a large part of the literature in approximating the cross-sectional distribution Φ by a small number of moments.¹⁶

Given the distinction between working-age and retired households, it is helpful to separate age $j \in \{W, R\}$ from the other state variables of the household and to state the individual dynamic programming problem of both types of households separately. The problem of retired households is given by

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

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 young, working households, the decision problem is given by

$$\begin{aligned} v_W(s, y, a, \beta; Z, \Phi) &= \left\{ \max_{c, a' \geq 0} u(c) + \beta \sum_{(Z', s', y') \in (Z, S, Y)} \pi(Z'|Z) \pi(s'|s, Z', Z) \pi(y'|y) \right. \\ &\quad \left. \times [\theta v_W(s', y', a', \beta; Z', \Phi') + (1 - \theta) v_R(a', \beta; Z', \Phi')] \right\} \end{aligned}$$

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*

¹⁶ See our companion paper, Krueger, Mitman and Perri (2016), for a more detailed discussion of the performance of the modified Krusell and Smith (1998) algorithm we use to solve the model.

1. Given the pricing functions r, w , the tax rate given in equation (5) 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 (5) 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. See appendix A for an explicit statement of the law of motion for this economy.

3.6 Calibration of the Benchmark Economy

We now briefly describe how we parameterize our business cycle model with household heterogeneity.¹⁷ Since the welfare losses from severe recessions induced in part by more frequent transitions into unemployment, we want to calibrate the model at a *quarterly* frequency.

3.6.1 Aggregate Risk

In the benchmark economy, we set $\omega = 0$; in this case TFP follows a fully exogenous process and there is no distinction between the exogenous process for Z and TFP Z^* . We model aggregate risk as a two-state process with $Z \in \{Z_l, Z_h\}$ where the state Z_l denotes a potentially severe recession and Z_h indicates normal times (which in our calibration strategy will include periods of shallow recessions). The Markov chain describing aggregate state transitions is given by

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

¹⁷ Krueger, Mitman and Perri (2016) contains a more detailed discussion and justification of these calibration choices.

Normalizing average productivity $E(Z) = 1$, the process is characterized by three parameters, the persistence of both states ρ_l, ρ_h and the dispersion of aggregate productivity, Z_l/Z_h . In Krueger, Mitman and Perri (2016), we define, empirically, a severe recession as an economic downturn in which the unemployment rate rises above 9% at least for one quarter and measure its length by the time period for which the unemployment rate remains above 7%. Using this operational definition of a severe recession, we identify two such downturns in U.S. postwar data (1948 to 2014.III): a period encompassing the double-dip recession of the early 1980s (1980.II-1986.II) and the Great Recession (according to our operational definition lasting from (2009.I-2013.III)). Matching the average length of a severe recession (22 quarters) and the frequency with which these recessions occur in our model (16.5% of all quarters) requires $\rho_l = 0.9545$ and $\rho_h = 0.9910$ and thus the transition matrix is given by:

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

Finally, we choose the ratio $\frac{Z_l}{Z_h}$ such that on average (across severe recessions), GDP per capita falls by 7% relative to normal trend growth. Given that during severe recessions in the model unemployment rises and thus labor input falls as well (as does capital, in the medium run), this requires $\frac{Z_l}{Z_h} = 0.9614$, which in turn, together with the normalization $E(Z) = 1$, pins down $Z_l = 0.9676, Z_h = 1.0064$.

3.6.2 Technology

With the exogenous process for TFP Z determined as discussed above, output is produced according to a standard Cobb-Douglas production function:

$$Y = ZK^\alpha N^{1-\alpha}. \quad (6)$$

The capital share is chosen to be $\alpha = 36\%$ and the quarterly depreciation rate to $\delta = 2.5\%$.

3.6.3 Household Demographics, Preferences, and Endowments

On the household side, we need to specify the basic demographic structure of the economy, preferences of households (with discount factor heterogeneity being its crucial element), and their stochastic labor income endowments, including the process for unemployment spells.

Demographics and the Life Cycle 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.

Preferences All households are assumed to have instantaneous logarithmic utility, but display permanent heterogeneity with respect to their impatience. Following Carroll et al. (2015) we assume that households draw their discount factors β from a uniform distribution with support $[\bar{\beta} - \epsilon, \bar{\beta} + \epsilon]$. In our numerical implementation we discretize this uniform 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 choose the parameters $(\bar{\beta}, \epsilon)$ such that the benchmark economy has a wealth Gini as in the data and a quarterly wealth-to-output ratio of 10.26 (as targeted by Carroll et al., 2015) The resulting parameters are $(\bar{\beta} = 0.9864, \epsilon = 0.0053)$; recall that the model is calibrated at a quarterly frequency, and thus, as a reference the annualized discount factors of households range from $\beta = 0.9265$ to $\beta = 0.9672$.

Endowments: Idiosyncratic Earnings Risk Households face two sources of idiosyncratic labor earnings risk: countercyclical unemployment risk described by the transition matrices $\pi(s'|s, Z', Z)$ and labor productivity and thus earnings risk, conditional on being employed, driven by the Markov transition matrix $\pi(y'|y)$.

Unemployment Risk Households can either be employed (in which case their labor earnings are determined by the labor productivity process y described below), or unemployed (in which case they receive unemployment benefits). Since employment-unemployment transition probabilities are allowed to depend on aggregate state transitions from Z to Z' , unemployment risk is governed by four 2 by 2 transition matrices $\pi(s'|s, Z', Z)$. Each of these matrices has two free parameters, but the assumption that the aggregate unemployment rate depends only on the aggregate state imposes, for each Z, Z' pair, the restriction that

$$\Pi_{Z'}(u) = \pi(u|u, Z', Z) \times \Pi_Z(u) + \pi(u|e, Z', Z) \times (1 - \Pi_Z(u)) \quad (7)$$

With these restrictions, each matrix is uniquely pinned down by quarterly job finding rates $\pi(e|u, Z', Z)$. In Krueger, Mitman and Perri (2016), we describe how we compute quarterly job finding rates (as a function of aggregate state transitions) from Current Population Survey (CPS) data. The resulting rates are

$$\begin{pmatrix} \pi(e|u, Z_l, Z_l) \\ \pi(e|u, Z_h, Z_l) \\ \pi(e|u, Z_l, Z_h) \\ \pi(e|u, Z_h, Z_h) \end{pmatrix} = \begin{pmatrix} 0.6622 \\ 0.7780 \\ 0.6618 \\ 0.8110 \end{pmatrix} \quad (8)$$

In general and as expected, job finding rates are high (and job losing rates low) when tomorrow's aggregate state of the world is $Z' = Z_h$ and low if the economy falls into a large recession, $Z' = Z_l$.

Earnings Risk Conditional on Employment A plausible earnings distribution requires a second source of earnings risk (or at least earnings heterogeneity). Following a large empirical literature, we specify a statistical model of log-earnings composed of a transitory and a persistent (potentially permanent component):

$$\log(y') = p + \epsilon \quad (9)$$

$$p' = \phi p + \eta \quad (10)$$

with persistence ϕ and persistent and transitory shocks (η, ϵ) with variances $(\sigma_\eta^2, \sigma_\epsilon^2)$. In Krueger, Mitman, and Perri (2016) we estimate this process for household labor earnings after taxes (after first removing age, education and time effects) from *annual* PSID data and find estimates of $(\hat{\phi}, \hat{\sigma}_\eta^2, \hat{\sigma}_\epsilon^2) = (0.9695, 0.0384, 0.0522)$. Next we translate these estimates into a quarterly persistence and variance¹⁸ and discretize the process into a finite state Markov chain.¹⁹

3.6.4 Government Policy

Unemployment Insurance Our model has two policy parameters: the replacement rate ρ of the unemployment insurance system, and the payroll tax rate τ_{SS} of the social

¹⁸ We map the estimated annual persistence into quarterly persistence by setting $\phi = \hat{\phi}^{\frac{1}{4}}$. Our main objective when choosing quarterly variances is that the resulting process delivers 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. This is achieved by setting the quarterly transitory variance equal to its annual counterpart and

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

¹⁹ We discretize the persistent component into a 7 point Markov chain using Rouwenhorst's method. When taking expectations with respect to the transitory shock, we perform the integration of the shock through a Gauss-Hermite quadrature with 3 nodes. Thus, essentially, we discretize the continuous AR(1) plus iid process into a discrete process with $7 \times 3 = 21$ states.

security system. We choose $\rho = 50\%$ in the benchmark economy, but also experiment with lower numbers to assess the importance of the generosity of the social insurance system for the size of the welfare losses and their distribution across households with different levels of wealth.

Social Security We choose the payroll tax rate for social security as $\tau_{SS} = 15.3\%$, which, given our demographics and other parameters of the model, implies a replacement rate of approximately 40% on average (over the cycle).

4 The Cross-Sectional Earnings, Income, Wealth and Consumption Distribution in the Benchmark Economy

Prior to turning to the welfare results, we first want to ensure that the model is indeed broadly consistent with the cross-sectional facts discussed in section 2.

4.1 Wealth Inequality in the Benchmark Economy

A large literature within macroeconomics with household heterogeneity proposes a variety of mechanisms to generate an empirically plausible cross-sectional distribution of net worth.²⁰ For example, Castaneda, Diaz-Gimenez and Rios-Rull (2003) and Kindermann and Krueger (2015) propose the presence of very large but transient income realizations not captured in PSID data. De Nardi, French and Jones (2010) and Ameriks, Briggs, Caplin, Shapiro and Tonetti (2015) stress the role of large uninsured, or only partially insured medical expenditure shocks in old age, De Nardi (2004) studies wealth concentration induced by the intergenerational transmission of wealth through accidental and intended bequests. Quadrini (2000), Cagetti and De Nardi (2006), and Buera (2009) propose models of entrepreneurs and study their importance for the right tail of the wealth distribution, and relatedly, Benhabib, Bisin and Zhu (2011) propose idiosyncratic shocks to investment opportunities as a main driver of wealth concentration.

In our benchmark model instead, dispersion in wealth is driven by uninsurable (because of incomplete markets) idiosyncratic unemployment and income shocks and preference heterogeneity, as already proposed by Krusell and Smith (1998), and further analysed by Hendricks (2007) and Carroll et al. (2015). These features interact with a rudimentary life cycle structure and a publicly provided unemployment insurance system. In Krueger, Mitman and Perri (2016) we argue that, broadly, persistent

²⁰ This literature is extensively surveyed in De Nardi (2015), De Nardi, Fella and Yang (2015) and Benhabib and Bisin (2016)

idiosyncratic income shocks, the public provision of unemployment insurance and a life cycle where individuals are born with little wealth are key for generating a substantial share of working-age households with little or no net worth, whereas preference heterogeneity is crucial for obtaining a model net worth distribution with a fat right tail.

Table 3 displays the resulting wealth distribution for our benchmark model with idiosyncratic income risk, incomplete markets, a rudimentary life cycle structure, unemployment insurance and heterogeneous discount factors, as well as the distribution of net worth implied by the original Krusell-Smith (1998) model (with our calibration of aggregate and idiosyncratic employment risk), as well as the empirical cross-sectional wealth distribution from section 2.²¹ Through appropriate choice of the time discount factor(s) both model 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. All other moments of the wealth distribution were not targeted in the calibration.

We observe that our model fits the empirical wealth distribution in the data fairly well, both at the bottom and at the top of the distribution. It captures the fact that households constituting the bottom two quintiles of the wealth distribution hardly have any wealth, and that the top wealth quintile holds more than 80% of all net worth in the U.S. economy. We also acknowledge that the model makes the wealth upper middle class (quintile 4) somewhat too wealthy and somewhat misses the wealth concentration at the *very top* of the distribution. Specifically, 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 14%. On the other hand, households between the 90th and the 99th percentile of the net worth distribution account for about 37% of wealth in the data, but 44% in the model.

Finally, table 3 reproduces the well-known result from Krusell and Smith's (1998) original paper that unemployment risk and incomplete 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 households at the bottom of the distribution hold significantly too much wealth in the model, relative to SCF or PSID micro data.²² As a summary measure, whereas the wealth Gini in the data is close to 0.8 in the working age-population, it is only 0.35 in

²¹ Recall that in the data, we restrict attention to households with at least one member of working age. 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.

²² In Krueger, Mitman and Perri (2016), we argue that it is the latter feature (the wealth holdings of the wealth-poor) that is responsible for the finding that in our benchmark model the decline in *aggregate* consumption in the Great Recession is significantly more pronounced than in the original Krusell and Smith (1998) model, which in turn approximates the aggregate consumption dynamics in a representative agent model quite closely.

Table 3: 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

the original Krusell-Smith model.²³

4.2 Income and Consumption at Different Points along the Wealth Distribution

The benchmark model delivers a plausible cross-sectional wealth distribution. In this section, we assess whether it is also consistent with the basic empirical facts concerning the *joint* income-wealth-consumption distribution. Table 4 sorts household by net worth and reports the share of earnings, disposable income, consumption expenditures and the expenditure rates for the five quintiles of the wealth distribution, for both the data (as already contained in table 2) and the benchmark model.

Overall, the model fares quite well in replicating the joint distributions of these variables. As in the data, households in the lowest two wealth quintiles hold a significant share of aggregate earnings, income and (crucially for welfare) consumption (close to 18%), although the share is still somewhat too low relative to the data, where the consumption share of the bottom 40% of the wealth distribution amounts to 24%.²⁴

²³ For the entire population, the wealth Gini is approximately 3-5 percentage points higher, in the data (and in both data sets) as well as in the benchmark model.

²⁴ In the Krusell-Smith economy with only earnings risk, which also delivers a plausible share of income and earnings held by the lowest wealth quintiles, the consumption share of the bottom 40% is only approximately 16%. Heterogeneity in time discount factor generates impatient households that end up holding little wealth (and thus are located in the lowest wealth quintile) but consume at a high rate, because of their impatience.

The benchmark model also displays consumption rates (out of disposable income) that are strongly decreasing with net worth up until the fourth wealth quintile, although the wealth gradient is not quite as steep as the model. For this success of the model, capturing the life cycle at least in a rudimentary form as we have done is crucial, since the retirement saving motive slows down the consumption spending of households at the top of the wealth distribution.²⁵ However, in a model where asset accumulation is fundamentally driven by precautionary motives, eventually, as households become very wealth-rich, this motive subsides and they consume a large(r) share of their disposable income, in contrast to the data.

Table 4: 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				

After having verified that the model is indeed consistent with the main cross-sectional facts from the PSID data, we now use it as a measurement tool to quantify the welfare losses from a great recession.

5 The Welfare Cost of Great Recessions

Given the heterogeneity in the consumption response to the aggregate downturn documented above, it is plausible to conjecture that the welfare losses from this adverse macroeconomic event are very unevenly distributed as well. In this section, we document that this is indeed the case.

²⁵ In an infinite horizon version of the model without a retirement phase, the fifth wealth quintile consistently has the highest consumption rate out of earnings and disposable income. In this version, households accumulate wealth exclusively to smooth bad income realizations, and once they have accumulated significant wealth for this purpose they consume at a high rate, especially when earnings and thus incomes are low.

5.1 Measurement of the Welfare Cost: A Useful Decomposition

5.1.1 Household-Specific Welfare Losses

In general we calculate the welfare losses from a particular transition (aggregate or idiosyncratic) as the permanent percentage increase in consumption that a working-age household would require so that its welfare in the transition is the same as the welfare when the transition does not happen. In particular let $g_{ss',ZZ'}(y, a, \beta)$ be the required percentage consumption compensation for a household of type (y, a, β) for an aggregate transition from Z to Z' and at the same time an idiosyncratic transition from s to s' . For a given current aggregate capital stock K prior to the recession²⁶ this quantity is given by:

$$g_{ss',ZZ'}(y, a, \beta) = 100 * \left[\exp \left\{ \left(\frac{(1 - \theta\beta)(1 - \nu\beta)}{1 - \nu\beta + \beta(1 - \theta)} \right) [v_W(s, y, a, \beta; Z, K) - v_W(s', y, a, \beta, Z', K)] \right\} - 1 \right] \quad (11)$$

The welfare cost is clearly individual state specific, and our primary interest is in the extent to which its size differs across households, as well as in the dimensions of household heterogeneity that determine this size of the loss.²⁷ Of specific interest is the loss for individuals of type $(u, y, a \approx 0)$, that is, the welfare loss from a recession for an unemployed household with close to zero wealth. We are interested in transitions from normal times, $Z = Z_h$ to a severe recession, $Z' = Z_l$. In the aggregate, a larger share of households are unemployed in a recession, and thus it is instructive to measure the welfare losses of those households that lose their job as the economy transits into a recession. This loss of moving from $s = e$ to $s' = u$ when the aggregate economy transits from $Z = Z_h$ to $Z' = Z_l$ is then given by g_{eu,Z_hZ_l} , using the notation developed above.

Note that this welfare cost captures the fact (by using the value functions and thus the underlying transition matrices with positive persistence) that conditional on falling into a recession the economy is likely to remain there for an extended period of time (in expectation, 22 quarters), and that, conditional on not experiencing a recession today, the economy is also unlikely to experience one tomorrow (since the good aggregate state is highly persistent as well).

5.1.2 An Aggregate Welfare Loss Measure

Since the welfare loss depends on the individual characteristics of the households, it is also informative to aggregate these losses in some form, in order to arrive at an ag-

²⁶ Recall that we approximate, in the computational algorithm, the cross-sectional wealth distribution by its first moment. We choose the capital stock prevailing in the economy after a long sequence of good Z realizations. Our results are not sensitive to choosing different values of K in the ergodic set.

²⁷ See appendix B for the derivation of equation 11

gregate measure of the welfare losses. Suppose households are randomly placed into the pre-recession cross-sectional distribution over individual characteristics. Under the veil of ignorance of not knowing where in the distribution one would be placed, we ask by what percentage would lifetime consumption of everyone need to be increased to be compensated from the loss of falling into a recession. Such a measure holds the cross-sectional distribution constant and only changes the aggregate state. However, since the fraction of unemployed households increases in the aggregate recession state, expected lifetime utility in the recession event has to be calculated under a new cross-sectional distribution that scales up the share of households in the unemployment state.

Denote by $\Phi_{W,h}$ the pre-recession working-age cross-sectional distribution over individual characteristics, and by $\Phi_{W,h,l}$ the same distribution, but with the mass of employed households scaled down, and the mass of unemployed households scaled up, so as to be consistent with the increase in the aggregate unemployment rate in the great recession.²⁸ Using the same calculations as above we can then derive an aggregate measure of the welfare losses from the Great Recession as:

$$\bar{g} = 100 * \left[\exp \left\{ \left(\frac{(1 - \theta\beta)(1 - \nu\beta)}{1 - \nu\beta + \beta(1 - \theta)} \right) \left[\int v_W(s, Z_h) d\Phi_{W,h} - \int v_W(s, Z_l) d\Phi_{W,h,l} \right] \right\} - 1 \right] \quad (12)$$

5.1.3 Decomposition of Household-Specific Welfare Losses

To aid with the interpretation of the results and the sources by which they emerge, the following decomposition of individual welfare losses into an aggregate and an idiosyncratic component is useful:

$$1 + g_{eu,Z_h Z_l}(y, a, \beta) = (1 + g_{ee,Z_h Z_l}(y, a, \beta)) \times (1 + g_{eu,Z_l Z_l}(y, a, \beta)) \quad (13)$$

or (taking logs and approximating $\log(1 + g) \approx g$)

$$g_{eu,Z_h Z_l}(y, a, \beta) \approx g_{ee,Z_h Z_l}(y, a, \beta) + g_{eu,Z_l Z_l}(y, a, \beta) \quad (14)$$

That is, the welfare loss from losing a job as the economy turns bad is (approximately) the sum of the welfare loss of an aggregate downturn for a person that remains employed and the welfare loss of becoming unemployed in bad times.²⁹ For households

²⁸ Recall that the incidence of unemployment in the model, by assumption, is orthogonal to all other household characteristics.

²⁹ An alternative decomposition leading to the same qualitative results as described below is given by

$$g_{eu,Z_h Z_l}(y, a, \beta) \approx g_{eu,Z_h Z_h}(y, a, \beta) + g_{uu,Z_h Z_l}(y, a, \beta). \quad (15)$$

not changing their idiosyncratic employment status the welfare loss from the recession is simply the aggregate component, as defined above. Also note that on net, changes in unemployment status are small: only approximately 3% more households are unemployed in the recession than in normal times, and thus the aggregate welfare measure behind the veil of ignorance defined in equation (12) are expected to closely mirror the aggregate component of the individual welfare losses in the decomposition above (which will turn out to not vary too much across households with different characteristics).

5.2 Benchmark Results


In figure 1 we plot the welfare losses from the recession $g_{eu, Z_h Z_l}(y, a, \beta)$ against assets for four different (y, β) combinations: the most impatient households with discount factor β_1 , households with the median discount factor β_3 , households with lowest income y_1 and households with median income y_4 .³⁰

We make the following observations: first, experiencing a Great Recession *and* a concurrent job loss is very painful for many households, with welfare losses ranging from 2% to 5% of lifetime consumption. Second, these losses are very unequally distributed between wealthy and poor households. Holding other household characteristics (such as impatience and permanent income) constant, the additional losses wealth-poor households sustain relative to wealth-rich households is on the order of 2 percentage points, and the wealth gradient is quite steep at the low end of the wealth distribution.

Third, losing a job as the economy slips into a recession is significantly more painful if the job one held was a good one: households with higher current y suffer larger losses. The welfare losses are also distributed unequally across households that differ in their discount factors. For the same level of income, the welfare loss is significantly higher for more impatient households ($\beta = 1$ in figure 1, corresponding to the most impatient households in the economy) relative to households with the average discount factor ($\beta = 3$ in the figure). This can be explained by the high persistence of severe recessions. Since these recessions are very persistent, both the contemporaneous drop in income from being unemployed, and the fact that the households expect to face lower income and increased unemployment risk while the recession persists, disproportionately affect more impatient households, who value less the higher expected income when the economy emerges from the recession. Broadly, for impatient households the short run business cycle is the main determinant of lifetime utility, whereas more patient households place more weight on the remaining lifetime that, in expectation, is

³⁰ Recall that we discretize the persistent income state process into a 7 state Markov chain. The plots shown pertain to households that do not experience a (positive or negative) transitory shock in the current period. Results for the most patient households are very similar to the ones with β_3 .

Figure 1: Welfare Losses $g_{eu,Z_h,Z_l}(y, a, \beta)$ from Great Recession by Asset Holdings




spent in normal times.

In figure 2 we now decompose, using equation (14) the welfare losses into an aggregate and an idiosyncratic component, both for households with the median earnings shock y_4 and median $\beta = \beta_3$ and for the most impatient households with the lowest earnings shock, ($y = y_1, \beta = \beta_1$). Abstracting from approximation error (which is nontrivial only for households with close to zero assets), the overall welfare loss is the sum of both components. We observe that both components are large and decline with asset holdings, but much more so for the idiosyncratic component.

The figure also shows that in the benchmark economy, the aggregate component of the welfare losses (i.e. the welfare losses of households that do not lose their job at the start of the recession) is dominant and amounts to approximately 2% of consumption, with fairly modest variation across asset holdings and other household characteristics.

Why is the aggregate component so large? Recall that g_{ee,Z_h,Z_l} is the welfare loss from falling into a severe recession conditional on the household not losing its job. This loss partially comes from lower aggregate wages (and lower returns on capital for those with positive assets), but is to a large degree the result of higher *future* unemployment risk. Recall that we study very persistent recessions (lasting on average 22 quarters)

Figure 2: Decomposition of Welfare Losses into $g_{eu,Z_h,Z_l}(y, a, \beta)$ and $g_{ee,Z_h,Z_l}(y, a, \beta)$



and that the risk of entering unemployment during a recession is substantially higher than in normal times. Thus, a big part of the aggregate component of the welfare losses stems from higher future idiosyncratic risk, which also affects households that do not lose their job at the onset of the recession.

The idiosyncratic component captures the direct impact of losing one's job at the onset of the recession, triggering immediate earnings losses (of 50% given the size of the unemployment insurance system). For households with little or no wealth, these earnings losses translate directly into current consumption losses of similar magnitude, and thus the idiosyncratic component is more potent for households at the low end of the asset distribution. Note, however, that unemployment spells are expected to be short (certainly relative to the length of the recession) and thus the idiosyncratic component contributes at most half (for impatient households with little income and assets) of the total welfare losses.³¹

Now that we have characterized how the welfare losses are distributed across individual states, we can use the distribution of households across these states at the start of

³¹ In our set-up, households do not suffer from persistent earnings losses upon re-employment. Introducing this empirically plausible feature (see, for example, Davis and Von Wachter, 2011) into the model would magnify the idiosyncratic – but also the aggregate – component of the welfare losses.

the recession, together with the implied transition during recession, to calculate aggregate welfare losses, as defined in equation (12), and to calculate two summary statistics of the impact of the recession, that is the fraction of households that experience welfare losses bigger than 3% and 4% of lifetime consumption as a consequence of the recession. These statistics are reported on line 1 of table 5. Note that aggregate losses amount to 2.16% of lifetime consumption, under the veil of ignorance of the position in the cross-sectional wealth distribution, and thus, as expected, are similar to the welfare losses of households that do not lose their job in the recession (the lines labeled $g_{ee, Z_h, Z_l}(y, a, \beta)$ in figure 2). This is simply because these households represent, even in a recession, the majority of the population. Note, however, that there is a non trivial mass of households that experience significantly larger welfare losses, and those are the low wealth households that happen to lose their job during the recession. Line 1 of table 5 shows that almost 2% of households experience losses bigger than 3% of lifetime consumption and 0.3% of households have losses larger than 4% of lifetime consumption.

Overall, we conclude that in our baseline model, the welfare impact of a large downturn is significant for most households, but with heterogeneity in the magnitude of the losses. In the remaining of the paper we explore how the magnitude of these losses changes when we change features of the model economy.

Table 5: Summary Statistics for Welfare Losses from Great Recessions


	Aggregated Welf. Loss	% of hh's with loss	
	(% of lifetime cons.)	>3%	>4%
1. Benchmark	2.16	1.9%	0.3%
2. KS (low wealth heterog.)	1.60	0.0%	0.0%
3. Low UI	2.14	3.1%	1.7%
4. UI shock	2.24	7.5%	2.7%
5. Cons. Externality	2.51	6.0%	0.8%
6. Cons. Externality+UI Shock	2.10	9.3%	3.2%

5.3 Welfare Costs of the Great Recession and the Wealth Distribution

In this subsection, we argue that the cross-sectional wealth distribution is a crucial determinant of the aggregate welfare losses from the great recession. We do so by showing that these aggregate losses are significantly larger in the benchmark economy that features realistic wealth inequality than they are in the Krusell-Smith economy where wealth dispersion is much more limited.

In figure 3 we display the cross-sectional wealth distribution and the welfare losses (the aggregate component and the total loss) for the original Krusell-Smith economy

Figure 3: Welfare Losses and Wealth Distribution, Krusell-Smith (left panel) and Benchmark (right panel) Economies



(left panel) and our benchmark economy (right panel), for households with a median persistent income state and a median time discount factor. Thus, the lines in the right panel of this figure replicates the information contained in the left panel of figure 2.

The first observation we make, comparing the left and right panels, is that for a given amount of wealth, the welfare losses from the recession for those losing a job are actually substantially larger in the original Krusell-Smith economy than in our benchmark economy (compare the solid lines in both subpanels, with the units being given on the left y -axis).³² This is especially true for households with little or no wealth, and is primarily due to the fact that the KS economy has (virtually) no unemployment insurance whereas our benchmark economy has substantial unemployment insurance.

However, in the left panel, the wealth distribution (scale on the right y -axis) has virtually no mass at the low end where the welfare losses are largest. The right panel features lower welfare losses for wealth-poorer households, but a lot of probability mass at the wealth levels where the losses are the largest.³³ This in turn crucially shapes the size of the aggregated welfare losses from recession according to the measure constructed in equation (12): as noted above, it is 2.16% in the benchmark economy but "only" 1.60% in the Krusell-Smith economy (see line 2 of table 5). The additional loss

³² The aggregate component (dashed line), which measures the welfare losses from the recession for those not losing a job, is quite similar in both economies, but slightly larger for the most impatient households in the benchmark economy, for whom the recession period constitute a more important component of expected discounted lifetime utility.

³³ Both economies have the same wealth-to-income ratio. We truncate the plots at a wealth level 20 times average income, but the wealth distribution for the benchmark economy has fat tails (and is well approximated by a Pareto distribution) and extends far to the right, whereas the corresponding distribution for the KS economy is approximately log-normal in its right tail and thus has little mass beyond 20 times average income.

of more than half of 1% of lifetime consumption may serve as a quantitative summary measure for the importance of realistic household heterogeneity for the welfare loss of severe recession question we address in this paper.

Figure 4: Histogram of Welfare Losses, Krusell-Smith (left panel) and Benchmark (right panel) Economies




Figure 4 provides a different perspective on the same phenomenon. It depicts, again for both economies, the histogram of welfare losses.³⁴ It shows the much larger left tail of welfare consequences from the great recession in the benchmark economy, where more than two thirds of households lose at least 2%, close to 2% of households lose more than 3%, and 0.3% of working age households lose more than 4% of permanent consumption. In the KS economy (see line 2 of table 5), in contrast, the share of households experiencing welfare losses in excess of 3% is negligible (about 3/100 of 1%), exactly because in this economy, households save away from the low wealth levels entailing large welfare losses from economic downturns.

5.4 Welfare Costs of Great Recessions and Unemployment Insurance


The previous analysis was based on a model with generous unemployment insurance so that even for households that lose their jobs in a recession the immediate income losses constitute 50% of their potential earnings. In this section, we document how the impact of the size of the UI system on our findings. We are especially interested in the extent to which our previous result, that the aggregate component of the welfare losses

³⁴ Even in bad economic times, some households transit from unemployment to employment and thus would have, according to our measures, a negative idiosyncratic and total welfare loss from the recession. We show a histogram based on net flows, reflecting the aggregate component of the losses for 97% of the population and the sum of the idiosyncratic and the aggregate component for 3% of the population, where 3% is the difference in the unemployment rate between the normal and the recession aggregate state of the economy.

is quantitatively more important than the idiosyncratic component, hinges on the fact that generous UI benefits soften the latter.

To do so, we now compare our previous results with those in an economy with little unemployment insurance, concretely, with a replacement rate of 10%. Note that, as the previous two subsections have already indicated, the generosity of unemployment insurance will have a significant impact on the welfare losses from the recession for households with a given level of wealth. For the aggregate component of individual welfare as well as aggregated (across households) measures of welfare, it will also be crucially important how the presence of UI benefits affects the cross-sectional distribution of wealth in the economy itself. For this reason we perform two distinct experiments. The first (labeled low UI) simply recomputes the benchmark economy with $\rho = 10\%$. Note that in this case, the wealth distribution before the recession will be very different than the one in the benchmark economy, because households, anticipating the higher loss of income in case of unemployment will be more reluctant to have low wealth. Welfare losses for a given state (and for households with median time discount factor $\beta = \beta_3$) are depicted in the middle panel of figure 5, while aggregate statistics are reported on line 3 of table 5.

Figure 5: Welfare Decomposition in Benchmark, Low UI and UI shock



The first key observation from the middle panel of figure 5 is that the welfare losses from losing a job in a Great Recession roughly double for low wealth-households, and now approximate 8% rather than 4% of lifetime consumption (we display results for

median earnings and impatient households). For households with wealth exceeding average wealth in the economy, the losses are still substantially larger with little unemployment insurance, but the difference is not nearly as sizeable.

Second, whereas the welfare losses of households that do not lose their job at the onset of the recession remains quantitatively roughly unchanged, with small ρ unemployment spells themselves are much more costly with little social insurance against them, especially for households with little financial wealth coming into the recession. As a result, the idiosyncratic component of the welfare loss now dominates in the low ρ economy, especially for households with no or very little net worth. In terms of aggregate statistics, line 3 of table 5 reveals several noteworthy features. First, the aggregate welfare losses are virtually unchanged, despite the large change in the welfare losses state by state, and across the wealth distribution. The reason is that the two changes offset each other in the computation of aggregated welfare losses. The decline in UI implies that in some states (low wealth), the welfare losses will be much higher, but the endogenous response of households implies that there will be much fewer households in those states. Notice, though, that despite this, in the low UI economy more households experience large losses (i.e. 3.1% of households, v/s 1.9% of households in the benchmark, experience losses larger than 3% of lifetime consumption). This is because the accumulation of wealth can provide only partial insurance against unemployment shocks and thus, under the low UI regime, there some unlucky households will experience large income losses that will translate into larger welfare losses.

Our second experiment (labeled UI shock, see the right panel of figure 5) considers a scenario in which unemployment benefits are cut unexpectedly from 50% to 10% exactly as the economy enters the recession (and remain low from that point on). The difference relative to the first scenario is that both the benchmark and the counterfactual scenario share *the same* cross-sectional wealth distribution coming into the recession. The right panel of figure 5 displays the results from this thought experiment, and aggregate statistics are reported on line 4 of table 5. Since the two low-UI thought experiments that distinguish the middle and right panels in figure 5 differ only in the initial wealth distributions prior to the Great Recession (and thus in the evolution of the wealth distribution through the recession), they affect households only through differential dynamics of factor prices. Households that lose their job are less affected by factor prices (wages), hence their welfare losses (the lines labeled $g_{eu, Z_h Z_l}(y, a, \beta)$) are very similar across panels. Households that do not lose their job are more affected by factor prices, and thus their welfare losses (the lines labelled $g_{ee, Z_h Z_l}(y, a, \beta)$) differ

more noticeable across the two experiments.³⁵

Overall, the largest difference between the two thought experiments is the wealth distribution with which the economy enters the recession, and this distinction is primarily responsible for the different magnitude of the numbers in lines 3 and 4 of table 5. Note that in the UI shock experiment, a much larger fraction of households experience large welfare losses (7.5% of households in this case experience welfare losses larger than 3%), because these households are caught by surprise by the fall in unemployment insurance benefits and did not accumulate sufficient wealth to effectively smooth consumption in case of a job loss.

Overall the main conclusion from this section is that changes in unemployment insurance significantly affect the distribution of welfare losses from a recession, much more so when households are caught by surprise by these changes. Finally, we wish to note that it is important to keep in mind two things when interpreting these welfare results. First, they do not represent a comprehensive normative assessment of the desirability of public unemployment insurance, but rather simply document how the welfare losses from recessions vary with the size of such a system. Second, and relatedly, given that employment-unemployment transitions are exogenous, the size of the unemployment insurance system does not impact individual incentives of seeking and keeping jobs. In addition, this analysis abstracts from the impact unemployment insurance has on the incentives of firms to create jobs. Recent research indicates that these distortions could potentially be large.³⁶

5.5 Welfare Costs of Severe Recessions when Output is Demand Determined

Krueger, Mitman, and Perri (2016) shows that in an economy in which aggregate output is partially determined because aggregate TFP depends on aggregate consumption, social insurance in general, and unemployment insurance specifically, stabilizes

³⁵ Note that especially in the economy with surprise UI cuts, the line $g_{ee, Z_h Z_l}(y, a, \beta)$ is U-shaped in household wealth (and exceeds the idiosyncratic component for wealthy households). These households derive most of their income from capital income, and the returns to capital fall in the recession because of the fall in TFP. This effect is present in both thought experiments. Furthermore, in the economy in which UI benefits unexpectedly fall (run out) and remain low, household wealth is suboptimally low (relative to the level of social insurance). Households respond with increased precautionary saving, pushing down rates of returns further. Thus, the welfare losses arising from aggregate factor price movements for wealthy households are more substantial (and the U-shaped nature of the aggregate component more severe) in the economy in which UI benefits run out by surprise.

³⁶ See, for example, Hagedorn et. al. (2013) and Hagedorn, Manovskii and Mitman (2015) who find large negative effects of unemployment benefit extensions on vacancy creation and employment, respectively, and Mitman and Rabinovich (2014) who argue that unemployment benefit extensions can explain the "jobless recovery" that followed the Great Recession.

the aggregate economy, in addition to providing social insurance against idiosyncratic unemployment shocks.³⁷ Thus, one would expect that in such an economy, differences in the welfare losses between an economy with low and one with high unemployment insurance replacement rate are especially large, and that the aggregate component of the welfare losses is responsible for this additional difference. We now investigate this conjecture in greater detail.

To do so, we extend the model to render aggregate output partially demand-determined. To achieve this, we assume that TFP Z^* is now a function of aggregate consumption:

$$Z^* = ZC^\omega, \quad (16)$$

where ω measures the strength of the aggregate demand externality. Recent work by Bai et al. (2012), Huo and Rios-Rull (2013), and Kaplan and Menzio (2016) provides micro foundations for an aggregate productivity process of this form.


As described in greater detail in Krueger, Mitman and Perri (2016), our calibration strategy is to use actual data on TFP by Fernald (2012) to calibrate the dispersion of the exogenous TFP process and then choose the externality parameter $\omega = 0.30$ such that output in the demand externality economy is as volatile as in the benchmark economy studied thus far. The resulting exogenous TFP fluctuations are significantly smaller, $\frac{Z_l}{Z_h} = 0.9781$, instead of $\frac{Z_l}{Z_h} = 0.9614$ as required by the benchmark economy. All other parameters are kept at their original values. The main purpose of the analysis there is to measure the additional benefit unemployment insurance has in the economy by stabilizing aggregate consumption demand and thus aggregate output.

Here, first we simply document the importance for the welfare losses of the amplification of aggregate shocks in the economy with partially demand-determined output. Line 5 in table 5 reports the aggregate welfare statistics for the economy with the consumption externality. Note that, even if the economy with consumption externality has, by construction, the same output volatility as the benchmark economy, the aggregate welfare losses in the former economy are larger (2.51% versus 2.16% in the benchmark). This economy has a stronger internal propagation mechanism and thus recessions are more persistent, and have more severe welfare consequences.

But the more interesting question that we wish to study with this version of the model is how the presence of the consumption externality (and thus the fact that output is partially consumption-demand determined) changes the impact of the generosity of unemployment insurance on the welfare losses from severe recessions. To analyze this issue, we now repeat, in the demand externality economy, the thought experiment of reducing, by surprise, the replacement rate of the unemployment insurance system from 50% to 10%. The results are shown in figure 6 and on line 6 of table 5.

³⁷ See McKay and Reis (2016) for a comprehensive analysis of the role of social insurance policies as automatic stabilizers of the macro economy.

Figure 6: Welfare Decomposition of Great Recession: Demand Externality Economy with UI Shock



In the benchmark economy, the welfare losses from the Great Recession increased sharply when unemployment benefits were cut, and the idiosyncratic part of the welfare losses started to dominate the overall welfare cost calculations, whereas the aggregate component remained broadly constant.


In the demand externality economy, the first finding persists. First, the welfare losses increase sharply for households with little net worth: they approximately double with a surprise cut in the replacement rate, and the increase is driven by an increase in the *idiosyncratic* component of the losses. This is reflected in the second and third columns of line 6 in table 5: when the economy is hit with the UI shock, the fraction of big welfare losers increases, just as it did when we introduce the UI shock in the benchmark economy. These big losers are households that lose their job and have low assets.

More surprising is that the welfare losses of households that do not lose their job (the lines $g_{ee,Z_h,Z_l}(y, a, \beta)$) become slightly *smaller* for median (y, β) households. This finding appears to run counter to the common intuition that in the demand externality economy, unemployment insurance stabilizes aggregate output, consumption, TFP and thus individual incomes. It also suggests that short-run demand stabilization, even if successful, need not have the intended beneficial welfare consequences. Indeed,

comparing line 5 and 6 of table 5 shows that in the consumption externality economy the UI shock *reduces* the aggregate welfare losses of the recession from 2.51% to 2.1% (whereas in the benchmark economy the UI shock makes aggregate losses larger).

Figure 7 provides the explanation for this seemingly contradictory finding. It plots, for both the benchmark and the externality economies, the *difference* in the evolution of wages and rates of return after a recession shock with high, relative to low unemployment insurance benefits. To interpret the units, consider the right panel pertaining to the externality economy. It shows that right after the recession hits, wages and rates of return are 0.75 basis points lower in the scenario with surprise unemployment insurance cuts than in the scenario without such cuts. However, as time passes, wages in the former scenario exceed that of the latter scenario, and quite strongly so. Why?

Figure 7: Impulse Responses of Factor Prices



When UI benefits by surprise run out when the recession hits, households (especially those at the low end of the wealth distribution) find themselves in a situation with sub-optimally low wealth levels, relative to the income risk they are facing, going forward. The behavioral response is to very substantially increase savings and cut consumption, especially for households that have not yet lost their job. On impact, as discussed above, aggregate consumption falls, and so do TFP and wages in the externality econ-

omy.³⁸ However, the precautionary saving translates into a large increase in the productive capital stock in the medium run, and thus a recovery in wages that is faster than in the economy without UI cuts. This effect is present in both the benchmark and the externality economy. However, in the externality economy, the medium run recovery in wages and household consumption driven by a larger capital stock translates into an increase in TFP and thus a further increase in wages, incomes and consumption. Thus the medium run recovery of the economy is more pronounced with low UI benefits, and is especially potent in the demand externality economy with low UI because of the endogenous feedback from aggregate wages and consumption to TFP. As a result, again comparing figures 5 and 6, the aggregate component of the welfare loss becomes *smaller* with a surprise cut in UI benefits in the demand externality economy, relative to the benchmark economy. Thus, even though UI benefits act as stabilizer of aggregate consumption and output in the short run, we conclude that the welfare consequences of this role of social insurance are more complex than appears at first sight, even in the absence of the negative incentive effects on worker search behavior and firm job creation incentives.³⁹


Implicit in the previous argument for a decline in the aggregate component of the welfare losses when UI falls is that the medium run gains from higher TFP and wages outweigh the lower wage and consumption losses in the short-term. The relative valuation of these effects depends crucially on a household's preference for the future. This is demonstrated by figure 8, which is identical to figure 6 but shows the welfare losses for the most impatient households in the economy (β_1 as opposed to β_3). For relatively patient households (Figure 6), the short-run effects of low wages and returns are dominated by the medium run movements in TFP and thus wages, and therefore the aggregate component of the welfare losses (the line g_{ee, Z_h, Z_l}) is noticeably *lower* with low unemployment benefits. For the most impatient households (Figure 8) the aggregate component of the welfare losses is instead roughly invariant to the size of the unemployment system (and in fact increasing non-trivially for the wealth-poorest when UI benefits are cut).

To summarize, we have established that more generous unemployment insurance is beneficial for households that lose their job during a severe recession, especially for those with low wealth. The effect of the size of the unemployment insurance system on households that are not (yet) laid off at the onset of the recession is instead ambiguous. In our benchmark economy, households that keep their job have to pay higher

³⁸ Note that this effect on TFP and wages is absent in the benchmark economy where TFP is fully exogenous.

³⁹ We wish to add two caveats to this discussion. First, this result depends on the assumption that the extra precautionary saving due to lower UI benefits flows into productive assets -the physical capital stock in this economy- that eventually raise wages and thus TFP. If the precautionary saving would instead be absorbed by unproductive assets or flow out of the economy, the described effect would be absent. Second, the additional potency of the effect in the demand externality economy clearly relies on the particular way we chose to model the relation between TFP and aggregate consumption.

Figure 8: Welfare Decomposition in Two Economies



taxes but also receive more benefits in case they become unemployed. Moreover, more generous unemployment insurance reduces precautionary saving and increases interest rates, benefiting especially high wealth households who keep their job. On net these effects are relatively small, though, and we conclude that households that get to keep their job during the recession are basically unaffected by the size of the unemployment insurance system.

More surprisingly, in a model in which TFP (and output) is demand determined, we find that households that retain their jobs suffer from more generous unemployment insurance. In the short run unemployment insurance increases aggregate demand and thus TFP, output and wages, which is beneficial for job keepers. The higher short run consumption and lower saving reduces capital, aggregate consumption, TFP and wages in the medium and long run, however. This negative effect quantitatively dominates the short run positive effect for welfare of most job keepers, and suggests that the majority of the population can actually be made worse off by more generous unemployment insurance at the onset of a severe recession.

This last finding brings us back to the broader point of this paper, and indeed the research agenda of which it is a part: the welfare losses from severe recessions can be large, are shaped in an important way by public social insurance, and perhaps most

importantly, vary fundamentally across a population displaying heterogeneity by income, wealth and preferences. Measuring this heterogeneity empirically and modelling it explicitly in general equilibrium is therefore crucial for normative analyses of the kind conducted here, in our view.⁴⁰

6 Conclusion

In this paper, we have used PSID data on earnings, income, consumption, and wealth as well as a canonical business cycle model with household earnings and wealth heterogeneity to measure the welfare losses of households across the wealth distribution, following a severe recession. We have argued that these losses are substantial, more than twice as high for wealth-poor relative to wealth-rich households, and that social insurance in the form of a publicly provided unemployment insurance program can effectively mitigate these costs, especially for households who lose their job and are at the low end of the wealth distribution.

Looking forward, to us the future of the research agenda on macroeconomics with household heterogeneity looks promising. In the last ten years, inequality has re-emerged as a major issue in the academic, popular and public policy debate⁴¹ at the same time as the macro economy experienced the most severe recession since World War II. At the same time, administrative data sets are increasingly becoming available to researchers, allowing them to paint a much more comprehensive picture of cross-sectional distributions than was previously possible.⁴² The advancement of computational tools for studying heterogeneous household and/or firm models has progressed at a rapid speed in recent years.⁴³ Therefore, if we were still in graduate school, signing on to this research agenda would appear to be a very promising proposition.

⁴⁰ Krueger, Mitman and Perri (2016) argue that the same is true for positive questions surrounding the dynamics of aggregate consumption and output in the great recession.

⁴¹ see, for example, Piketty's (2014) best-selling book and the Occupy (Wall Street) Movement.

⁴² See e.g. Guvenen, Ozkan and Song (2014) for a recent example of using such administrative data for an important applied topic.

⁴³ For a collection of these advancements, see the special issues (January 2010 and February 2011) of the *Journal of Economic Dynamics and Control*.

References

- [1] Aiyagari, R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659-684.
- [2] Ameriks, J., J. Briggs, A. Caplin, M. Shapiro and C. Tonetti (2015), "Long-Term Care Utility and Late in Life Saving," *NBER Working paper* 20973.
- [3] Bai, Y., V. Ríos Rull and K. Storesletten (2012), "Demand Shocks as Productivity Shocks," *Working Paper*.
- [4] Benhabib, J. and A. Bisin (2016), "Skewed Wealth Distributions: Theory and Empirics," *NBER Working Paper* 21924.
- [5] Benhabib, J., A. Bisin, and S. Zhu, (2011), "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents ", *Econometrica*, 79(1), 123-157.
- [6] Bewley, T. (1986), "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in Werner Hildenbrand and Andreu Mas-Colell (eds.) *Contributions to Mathematical Economics in Honor of Gerard Debreu*, Amsterdam: North-Holland, 79-102.
- [7] Buera, F. (2009), "A Dynamic Model of Entrepreneurship with Borrowing Constraints: Theory and Evidence," *Annals of Finance*, 5, 443-464.
- [8] Cagetti, M., and De Nardi, M. (2006), "Entrepreneurship, Frictions, and Wealth," *Journal of Political Economy*, 114, 835-870.
- [9] Carroll, C., J. Slacalek, K. Tokuoka and M. White (2015), "The Distribution of Wealth and the Marginal Propensity to Consume," *Working paper*.
- [10] Castaneda, A., J. Díaz-Giménez and V. Ríos-Rull (2003), "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111, 818-857.
- [11] Chatterjee. S. and D. Corbae (2007), "On the Aggregate Welfare Cost of Great Depression Unemployment," *Journal of Monetary Economics*, 54, 1529-1544.
- [12] Davis S. and T. Von Wachter (2011), "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, 43(2), 1-72
- [13] De Nardi, M. (2004) "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743-768.
- [14] De Nardi, M. (2015) "Quantitative Models of Wealth Inequality: A Survey," *NBER Working Paper* 21106.
- [15] De Nardi, M., G. Fella and F. Yang (2015) "Piketty's Book and Macro Models of Wealth Inequality", *NBER Working Paper* 21730.

- [16] De Nardi, M., E. French and J. Jones (2010) "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39-75.
- [17] Fernald, J. (2012), "Productivity and Potential Output before, during, and after the Great Recession", *Federal Reserve Bank of San Francisco Working Paper* 2012-18.
- [18] Glover, A., J. Heathcote, D. Krueger and V. Rios-Rull (2014), "Inter-generational Redistribution in the Great Recession", *Working paper*.
- [19] Guvenen, F., S. Ozkan and J. Song (2014), "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122, 621-660.
- [20] 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* 19499.
- [21] Hagedorn, M., I. Manovskii and K. Mitman (2015), "The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?" *NBER Working Paper* 20884.
- [22] Hendricks, L. (2007), "How Important Is Discount Rate Heterogeneity for Wealth Inequality?" *Journal of Economic Dynamics and Control*, 31, 3042-68.
- [23] Huggett, M. (1993), "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953-969.
- [24] Huggett, M. (1997), "The One-Sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics," *Journal of Monetary Economics*, 39, 385-403.
- [25] Huo, Z. and V. Ríos-Rull (2013), "Paradox of Thrift Recessions," *NBER Working Paper* 19443.
- [26] Hur, S. (2014), "The Lost Generation of the Great Recession," *Working Paper*, University of Pittsburgh.
- [27] Imrohorglu, A. (1989), "Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97, 1364-1383.
- [28] Kaplan, G. and G. Menzio (2016), "Shopping Externalities And Self-Fulfilling Unemployment Fluctuations," *Journal Of Political Economy*, 124, 771-825.
- [29] Kindermann, F. and D. Krueger (2015), "High Marginal Tax Rates on the Top 1%? Lessons from a Life Cycle Model with Idiosyncratic Income Risk," *Working Paper*.
- [30] Krebs, T. (2007) "Job Displacement Risk and the Cost of Business Cycles," *American Economic Review*, 97, 664-686.
- [31] Krueger, D., K. Mitman and F. Perri (2016), "Macroeconomics and Household Heterogeneity, Including Inequality ", *Handbook of Macroeconomics, Volume 2*, forthcoming.

- [32] Krueger, D. and F. Perri (2004), "On the Welfare Consequences of the Increase in Inequality in the United States," in Mark Gertler and Kenneth Rogoff (eds) *NBER Macroeconomics Annual 2003*, vol. 18 Cambridge MA: MIT Press 83-138
- [33] Krusell, P., T. Mukoyama, A. Sahin and A. Smith (2009), "Revisiting the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 393-404.
- [34] Krusell, P. and A. Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867-896.
- [35] Krusell, P. and A. Smith (1999), "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics* 2, 245-272.
- [36] Lucas R. (1987), "Models of Business Cycles," New York: *Basil Blackwell*
- [37] Lucas R. (2003), "Macroeconomic Priorities," *American Economic Review*, 93, 1-14.
- [38] McKay, A. and R. Reis (2016), "The Role of Automatic Stabilizers in the U.S. Business Cycle," *Econometrica*, forthcoming.
- [39] Mitman, K. and S. Rabinovich (2014), "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?," *Working Paper*.
- [40] Piketty, T. (2014), "Capital in the Twenty-First Century," Cambridge, MA: *Belknap Press*.
- [41] Quadrini, V. (2000), "Entrepreneurship, Saving and Social Mobility," *Review of Economic Dynamics*, 3, 1-40.

is to compute how lifetime utility is scaled upward if consumption in every t , at every node of the event tree, is scaled up by a factor $1 + g$. Denoting lifetime utility of a working age individual with idiosyncratic state s and aggregate state S by $v_W(s, S)$ and the lifetime utility from the scaled-up consumption profile by $v_W(s, S; g)$ we have:

$$\begin{aligned}
v_W(s, S) &= E \sum_{t=0}^{\infty} (\beta\theta)^t \log(c_t) + \beta(1-\theta) E \sum_{t=1}^{\infty} (\beta\theta)^{t-1} \left\{ \sum_{\tau=t}^{\infty} (\beta\nu)^{\tau-t} \log(c_t) \right\} \\
&= E \sum_{t=0}^{\infty} (\beta\theta)^t \log(c_t) + \frac{\beta(1-\theta)}{1-\nu\beta} E \sum_{t=1}^{\infty} (\beta\theta)^{t-1} \log(c_t) \\
v_W(s, S; g) &= E \sum_{t=0}^{\infty} (\beta\theta)^t \log((1+g)c_t) + \frac{\beta(1-\theta)}{1-\nu\beta} E \sum_{t=1}^{\infty} (\beta\theta)^{t-1} \log((1+g)c_t) \\
&= \log(1+g) \left\{ \sum_{t=0}^{\infty} (\beta\theta)^t + \frac{\beta(1-\theta)}{1-\nu\beta} E \sum_{t=1}^{\infty} (\beta\theta)^{t-1} \right\} + \\
&\quad E \sum_{t=0}^{\infty} (\beta\theta)^t \log(c_t) + \frac{\beta(1-\theta)}{1-\nu\beta} E \sum_{t=1}^{\infty} (\beta\theta)^{t-1} \log(c_t) \\
&= \frac{\log(1+g)}{1-\theta\beta} \left(1 + \frac{\beta(1-\theta)}{1-\nu\beta} \right) + v(s, S) \\
&= \log(1+g) \left(\frac{1-\nu\beta + \beta(1-\theta)}{(1-\theta\beta)(1-\nu\beta)} \right) + v_W(s, S).
\end{aligned}$$

The scaling factor g that makes a working age household indifferent between remaining in normal times $Z = Z_h$ and experiencing a recession ($Z = Z_l$), with scaled-up consumption from that point on, then solves the equation:

$$v_W(s, Z_l; g) = v_W(s, Z_h) \quad (17)$$

which, given the result above, gives the scaling factor g as:

$$\log(1+g) \left(\frac{1-\nu\beta + \beta(1-\theta)}{(1-\theta\beta)(1-\nu\beta)} \right) + v_W(s, Z_l) = v_W(s, Z_h) \quad (18)$$

and thus:

$$g_{Z_h Z_l}(s) = 100 * \left[\exp \left\{ \left(\frac{(1-\theta\beta)(1-\nu\beta)}{1-\nu\beta + \beta(1-\theta)} \right) [v_W(s, Z_h) - v_W(s, Z_l)] \right\} - 1 \right] \quad (19)$$

and thus the welfare loss $g_{Z_h Z_l}(s)$ for a household of type s is positive as long as $v_W(s, Z_h) > v_W(s, Z_l)$. The same calculations apply if, in addition to a change in the aggregate state of the world the household experiences a change in the idiosyncratic state of the world (e.g. when losing a job at the onset of a severe recession), in which case part of the idiosyncratic state s in $v_W(s, Z_l)$ has to be replaced by the new state s' .

Note that if $\sigma \neq 1$, then the above calculations are simpler and deliver

$$g_{Z_h Z_l}(s) = 100 * \left(\frac{v_W(s, Z_h)}{v_W(s, Z_l)} \right)^{\frac{1}{1-\sigma}} > 0$$