

Wealth and Volatility*

Jonathan Heathcote
Federal Reserve Bank of Minneapolis and CEPR

Fabrizio Perri
Federal Reserve Bank of Minneapolis and CEPR

February 2016

Abstract

Between 2007 and 2013, U.S. households experienced a large and persistent decline in net worth. The objective of this paper is to study the business cycle implications of such a decline. We argue that low wealth opens the door to self-fulfilling fluctuations. If wealth-poor households expect high unemployment, they have a strong precautionary incentive to cut spending. In an environment in which output is demand-driven, this makes the expectation of high unemployment a reality. We first document that poor U.S. households cut spending during the Great Recession much more sharply than richer households, consistent with a precautionary interpretation. We then develop a macroeconomic model featuring unemployment risk, and study the link between asset values and the volatility and persistence of precaution-driven fluctuations..

Keywords: Business cycles; Aggregate demand; Precautionary saving; Multiple equilibria, Self-fulfilling crises


JEL classification codes: E12, E21

*We thank seminar participants at several institutions and conferences, and Mark Aguiar, Christophe Chamley, Andrew Gimber, Franck Portier, Emiliano Santoro and Immo Schott for insightful comments. Also thanks to Joe Steinberg for excellent research assistance. Perri thanks the European Research Council for financial support under Grant 313671 RESOCONBUCY. 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

Between 2007 and 2013, a large fraction of U.S. households experienced a large and persistent decline in net worth. Figure 1 plots median real net worth from the Survey of Consumer Finances (SCF), for the period 1989-2013, for households with heads between ages 22 and 60. Since 2007 median net worth for this group has roughly halved and shows no sign of recovery through 2013. In relation to income, the decline is equally dramatic: the median value for the net worth to income ratio fell from 1.58 in 2007 to 0.92 in 2013.

Figure 1: Median household net worth in the United States



Note: Sample includes households with heads between ages 22 and 60.

The objective of this paper is to study the business cycle implications of such a large and widespread fall in wealth. We will argue that falls in household wealth (driven by falls in asset prices) leave the economy more susceptible to confidence shocks that can increase macroeconomic volatility. Thus, policymakers should view low levels of household wealth as presenting a threat to macroeconomic stability. Figures 2 and 3 provide some motivating evidence for this message.

Figure 2 shows a series for the log of total real household net worth in the United States from 1920 to 2013, together with its linear trend. The figure shows that over this period there have been three large and persistent declines in household net worth: one in the early 1930s, one in the early 1970s, and the one that started in 2007. All three declines have marked the start of periods

characterized by deep recessions and elevated macroeconomic volatility.¹

Figure 2: Household net worth since 1920





Figure 3 focuses on the postwar period, for which we can obtain a consistent measure of macroeconomic volatility. We measure volatility as the standard deviation of quarterly real GDP growth over a 10-year window. The figure plots this measure of volatility for overlapping windows starting in 1947.1 (the values on the x -axis correspond to the end of the window), together with wealth, measured as the deviation from trend (the difference between the solid and dashed lines in Figure 2) averaged over the same 10-year window. The figure reveals that periods when wealth is high relative to trend, reflecting high prices for housing and/or stocks, tend to be periods of low volatility in aggregate output (and hence employment and consumption). Conversely, periods in which net worth is below trend tend to be periods of high macroeconomic volatility. For example, during windows ending in the late 1950s and early 1980s, wealth is well below trend, and volatility peaks; conversely, in windows ending in the early 2000s and late 1960s, wealth is well above trend and volatility is low.

Why should wealth affect volatility? The novel idea of this paper is that the value of wealth in an economy determines whether or not the economy is vulnerable to economic fluctuations driven

¹In order to construct a consistent series for net worth, we focus on three categories of net worth for which we can obtain consistent data throughout the sample: real estate wealth (net of mortgages), corporate securities, and government treasuries. See Appendix B for details on the construction of the series.

Figure 3: Wealth and volatility



by changes in household optimism or pessimism (animal spirits). When wealth is low, consumers are poorly equipped to self-insure against unemployment risk, and hence have a precautionary saving motive which is highly sensitive to unemployment expectations. Suppose households come to expect high unemployment. With low wealth, the precautionary motive to save will be strong, and households will cut desired expenditure sharply. In an environment in which demand affects output, this decline in spending rationalizes high expected unemployment. Suppose, instead, that households in the same low wealth environment expect low unemployment. In this case, because perceived unemployment risk is low, the precautionary motive will be weak, consumption demand will be relatively strong, and hence equilibrium unemployment will be low. Thus, when asset values are low, economic fluctuations can arise due to self-fulfilling changes in expected unemployment risk.

In contrast, when the fundamentals are such that asset values are high, consumers can use wealth to keep their consumption smooth through unemployment spells, and thus the precautionary motive to save is weak irrespective of the expected unemployment rate. Thus, high wealth rules out a confidence-driven collapse in demand and output.

In order to articulate our point we proceed in two steps.

First, we use micro data from the Consumer Expenditure Survey (CES) and the Panel Study

of Income Dynamics (PSID) to document that, around the onset of the Great Recession, low net worth households increased their saving rates (i.e., cut their expenditures) by significantly more than high net worth households. This pattern is especially remarkable when considered alongside a second finding, which is that low wealth households suffered much smaller wealth losses during the recession. This new evidence indicates that the precautionary motive, in the context of sharply eroded home equity wealth and rising unemployment risk, was a key driver of consumption dynamics during the recession.

Second, we develop a simple macro model in which precaution-driven changes in consumer demand can generate self-fulfilling aggregate fluctuations. The model contains two key ingredients. First, labor markets are frictional, so that unemployment can arise in equilibrium. Specifically, we assume a sticky wage, so that consumer demand, rather than desired labor supply, determines the unemployment rate. Second, unemployment risk is imperfectly insurable, so that anticipated unemployment generates a precautionary motive to save.

It is important that idiosyncratic unemployment risk in the model is imperfectly insurable, so that a precautionary motive is active and consumption demand is sensitive to perceived unemployment risk. We therefore rule out explicit unemployment insurance, but assume that households own an asset (housing) that provides services and can also be used to smooth consumption in the event of an unemployment spell. We avoid the numerical complexity associated with standard incomplete markets models (e.g., Huggett 1993 or Aiyagari 1994) by assuming that individuals belong to large representative households.² However, the household cannot reshuffle resources from working to unemployed household members within the period. This preserves the precautionary motive, which is the hallmark of incomplete markets models. We will heavily exploit one property of the model: higher household wealth (i.e., high house prices) makes desired precautionary saving (and thus consumption demand) less sensitive to the level of unemployment risk. The intuition is simply that higher wealth permits higher consumption for unemployed household members, and thus better within-household risk sharing.

After describing the structure of model, we characterize equilibria with fixed house prices. When fundamentals are such that house prices (and hence net worth) are high, then no self-fulfilling fluctuations in unemployment are possible (conditional on prices staying high). When fundamentals are such that house prices are low, in contrast, the unemployment rate can fluctuate. Suppose the economy starts in a “good” state in which agents expect low unemployment. Because

²Challe and Ragot (2012) show that an alternative way to preserve a low-dimensional cross-sectional wealth distribution while still admitting a precautionary motive is to assume that utility is linear above a certain consumption threshold.

unemployment is expected to be low, agents have a weak precautionary motive and expenditure is high (supporting low unemployment), while the precautionary demand for saving is low (supporting low house prices). However, the economy is fragile, and a collective loss of confidence can trigger a switch to a “bad” state in which agents expect high unemployment. If this switch occurs, then, because of low asset prices, the precautionary motive to save strengthens, and the associated fall in expenditure fully rationalizes the rise in expected unemployment. By itself, increased precautionary saving will increase housing demand. This additional liquidity demand for housing offsets the reduction in housing demand arising from higher unemployment, and explains why a recession can happen without a collapse in house prices.

It is important at this point to distinguish our theory from an alternative theory of self-fulfilling fluctuations in which declines in output are coincident with declines in asset values (see, e.g., Farmer 2013). In fluctuations of this type, the narrative goes as follows. Households expect asset prices to collapse, which makes them feel poorer, and a standard wealth effect channel induces them to cut spending. Low demand and the associated fall in output then rationalizes the expected fall in asset prices.

Although the two theories are related, there is a crucial difference between them. In our theory, the primary factor that drives the reduction in spending, and hence the recession, is the expected increase in risk. In the alternative theory, it is the expected fall in asset prices. We do not view the two theories as mutually exclusive, but we note that if the main driver of reduced spending during the Great Recession was falling asset prices, then high wealth households (who suffered the largest wealth losses) should have exhibited the largest spending declines. Instead, the data show that low wealth households were the ones who reduced consumption most, suggesting an important role for our precautionary mechanism.

After characterizing equilibria in the model, we show that the theory can be applied to help us better understand features of business cycles, and to design policies to ward off future confidence crises.

First, all confidence-driven fluctuations must be persistent in expected terms, since a rapid expected recovery would cut against the incentive to cut current consumption. Similar logic implies that confidence-driven recessions that are especially deep are likely to be especially persistent. We also show that confidence-driven fluctuations tend to be larger the lower are asset values, consistent with the evidence presented in Figures 2 and 3. All these predictions are consistent with the view that both the Great Depression and the Great Recession were driven, at least in part, by a confidence-driven decline in demand. Both events followed sharp declines in household wealth, both

involved large and sharp declines in output, and both were followed by very sluggish recoveries. We add to the case for such an interpretation of the Great Recession by showing that a simply parameterized version of the model can generate dynamics for unemployment that are similar to those experienced by the United States between 2007 and 2014.

Finally, we use the model to evaluate the role of public unemployment insurance as a stabilization tool designed to counteract confidence-driven fluctuations. Recall that a confidence-driven recession is ignited by an increase in perceived labor income risk. Unemployment benefits, by partially insuring such risk, can prevent this type of recession from getting started. Obviously our model is too stylized to argue that unemployment benefits should be more generous; rather the result highlights a novel positive feature of generous benefits, and one that should be weighed against the costs of such a policy.

The paper is organized as follows. Section 2 presents the evidence on households' expenditures and wealth during the Great Recession. Sections 3 and 4 contain our theory and its applications. Section 5 discusses the related literature, and Section 6 concludes.

2 Microeconomic Evidence

Our idea is that when wealth is low, desired household consumption becomes sensitive to perceived unemployment risk. When this risk rises, wealth-poor households reduce consumption sharply, which translates into lower employment, and rationalizes ex-post the fear of high unemployment. Is this mechanism empirically relevant for understanding the decline in aggregate consumption during the Great Recession? An alternative – possibly complementary – hypothesis is that negative wealth effects associated with sharp declines in asset prices played the dominant role. In this Section we offer microeconomic evidence that can help discriminate between these two hypotheses. The key insight is that the precautionary mechanism should be quantitatively more important for the consumption behavior of low wealth households, while declines in asset values should matter more for high wealth households. Thus, if the precautionary mechanism is more important in aggregate we should expect to see relatively sharp declines in consumption for low wealth households, while if traditional wealth effects played the dominant role, we should expect to see wealthier households cutting consumption disproportionately more than poorer ones.

In this section we use provide novel evidence, based on data from the Consumer Expenditure Survey (CES) and the Panel Study of Income Dynamics (PSID) to show that at the onset of the recession lower wealth households exhibit systematically larger declines in their consumption rates.

This evidence is broadly consistent with Mian et al. (2013), who find that zip codes in the United States with poorer and more levered households experienced the sharpest consumption declines during the Great Recession. Collectively, this evidence lends support to demand-driven theories of the Great Recession, and to the importance of wealth in understanding demand dynamics.

2.1 Empirical Strategy

Our goal is to compare changes in consumption rates during the course of the Great Recession for wealth rich versus wealth poor households. In each data set we rank households by net worth and compute changes over time in the consumption rates of households in different groups of net worth distribution. It is important that the set of households in each wealth group is fixed when we measure the change in the consumption rate between t and $t + 1$, so that the change in the measured consumption rate reflects a true change in savings behavior, and is not an artifact of a change in the composition of the groups. Fortunately both the PSID and CES data sets have a panel dimension: in the PSID, households are re-interviewed every two years, while in the CES they are interviewed for four consecutive quarters, and are asked about income in their first and last interviews.


2.2 Aggregates

Before contrasting consumption behavior across wealth groups, we first explore the dynamics of aggregate consumption, income and wealth in our cross sectional data, in order to verify that the micro data captures the broad contours of the Great Recession. Panel A of Figure 4 shows the dynamics of average per capita expenditures in the PSID and the CES against the equivalent measure in the National Income and Product Accounts (NIPA). Panel B shows average per capita disposable income in the PSID and the CES versus NIPA personal disposable income. Panel C shows median household net worth in the PSID and the CES versus median net worth in the Survey of Consumer Finances (SCF). Our consumption concept includes all categories except expenditure on housing and on health. Net worth includes net financial wealth plus housing wealth net of all mortgages (including home equity loans).³ The key message from the figure is that the dynamics of consumption, income and wealth are broadly comparable across data sets. In particular, both micro data sets exhibit a marked reduction in consumption expenditure during the recession.⁴

³Appendix B reports more details of how we measure each variable. We do not impose any sample selection when constructing the PSID, CES and SCF series in Figure 4.

⁴One discrepancy is that consumption expenditures decline somewhat earlier in the PSID than in the CES or the NIPA. Note, however, that due to the bi-annual nature of the PSID we have no observation for 2007. In addition, it is

Figure 4: Comparing aggregates across micro data sets



2.3 Measurement

We now describe precisely how we define and compute changes in consumption rates for rich and poor households in the PSID. The procedure for the CES is very similar, adapted to the slightly different panel structure of the survey (see Appendix B).

First, for any year t , we construct the sample we use to measure changes in consumption rates between year t and year $t + 2$. We select all households with a head or spouse aged between 22 and 60, and which report income, consumption and wealth in both the t and $t + 2$ waves. We focus on households of working age, since unemployment risk is most relevant for this group.

difficult to date consumption precisely in the PSID because some of the survey questions ask explicitly about spending in the previous year – the year to which we attribute consumption – while others ask about current consumption. In Appendix B we discuss how excluding the latter consumption categories reduces the difference in dynamics between the PSID and the other two sources.

Second, we rank households by net worth in year t relative to the average of consumption expenditures in years t and $t + 2$. We then divide the sample into two equal size subgroups, rich and poor, where the dividing line is (weighted) median net worth relative to consumption. We measure household wealth relative to consumption since the strength of a household’s precautionary motive to save is likely to be more closely connected to wealth relative to permanent income (for which average consumption is a proxy), rather than to absolute wealth.

Third, for each group we compute consumption rates in years t and $t + 2$, where the consumption rate is defined as the average consumption of the group divided by the average disposable income of the group. The change in the consumption rate between t and $t + 2$ for each group is simply the $t + 2$ rate minus the t rate.

We then move to compute the change in the consumption rate from $t + 2$ to $t + 4$. This involves constructing a new sample, following the same procedure described in the first step, ranking households in the new set to construct new rich and poor groups, and constructing new measures of consumption rates for $t + 2$ and $t + 4$.

2.4 Descriptive Statistics

Table 1 reports characteristics of the rich and poor groups in both the PSID and the CES, for the year 2006. Differences between rich and poor are very similar across the two data sets. With respect to demographics, the wealth poor group tends to be younger and less educated. The most striking difference between the rich and poor groups, not surprisingly, is in terms of wealth. Median net worth for the poor group is near zero, while for the rich group it is around \$265,000 in the PSID and around \$187,000 in the CES.⁵ This dramatic difference suggests that the precautionary saving motive for the poor group should be much stronger than for the rich group, while the rich group likely experiences much larger capital losses when housing and stock prices fall. The wealth-poor group has a little more than half the average income of the rich group, but has a much higher consumption rate, so that differences in consumption between the two groups are quite small.

⁵One reason for the difference between the two data sets is that the measure of net worth in the PSID includes more assets, such as individual retirement accounts, vehicles, and family businesses. See Appendix B for details.

TABLE 1. CHARACTERISTICS OF THE WEALTH RICH AND THE WEALTH POOR, 2006

	PSID		CES	
	Poor	Rich	Poor	Rich
Sample size	3446	2523	1915	1960
Mean age of head	37.9 (0.21)	47.1 (0.21)	40.2 (0.25)	46.4 (0.24)
Heads with college (%)	21.3 (0.86)	36.5 (1.1)	24.8 (1.1)	39.4 (1.2)
Mean household size	2.45 (0.04)	2.72 (0.03)	2.84 (0.04)	2.79 (0.04)
Mean household net worth (current \$)	11,931 (879)	619,831 (49,388)	11,967 (1,155)	338,535 (12,644)
Median household net worth	5,000 (476)	265,000 (6,602)	1,800 (294)	187,102 (4,893)
Per capita disposable income	15,028 (256)	28,475 (667)	18,739 (334)	30,184 (593)
Per capita consumption expenditure	9,831 (177)	13,101 (250)	9,185 (232)	10,858 (188)
Consumption rate (%)	65.8 (0.90)	46.0 (0.86)	49.0 (1.18)	36.0 (0.66)

Note: Bootstrapped standard errors are in parentheses.

2.5 Changes in Consumption Rates: Rich versus Poor Households


Figure 5 contains the key findings of this section. The figure plots changes in consumption rates computed from the PSID (top panels) and from the CES (bottom panels). The two panels on the left (A and C) show the evolution over time of the changes in consumption rates of the bottom 50% of the net worth distribution (labeled “Poor”) against the same changes for the top 50% (labeled “Rich”). Both panels show that in “normal” times, there are no significant differences in consumption rate changes between the two wealth groups, while at the onset of the recession low net worth households reduce consumption rates significantly more than high net worth households.⁶

The two panels on the right (B and D) disaggregate more finely by net worth, and report the changes in consumption rates for the five quintiles of the net worth distribution. Here we focus on just two intervals: a pre-recession period, and the period when the economy enters recession. The plots show that as the economy moves into recession, consumption rates fall for all wealth quintiles, but the consumption rate drop is a declining function of initial net worth, with the bottom quintiles

⁶Consumption rate declines in the CES appear to be smaller than in the PSID. We conjecture that this primarily reflects the fact that the CES consumption rate changes are computed over 9-month intervals, while the PSID changes are recorded over 2-year intervals.

experiencing the largest consumption rate drops, and the top quintiles experiencing much smaller falls.

Figure 5: Changes in consumption rates for rich and poor



The disproportionate expenditure decline for wealth-poor households during the recession points toward an important precautionary motive in the face of rising unemployment risk. Moreover, the observed differential consumption decline is especially remarkable considering that it occurs at a time of sharply falling asset prices which, *ceteris paribus*, should disproportionately hurt the rich and lead them to reduce consumption more than the poor. To quantify the differential impact of asset price declines, we exploit the fact that households in the PSID report wealth at each interview, allowing us to compute changes in net worth over time for both the rich and poor groups. Table 2 reports the changes in consumption rates plotted in Figure 5 for each wealth group as defined above (lines 1 and 4), alongside changes in wealth (as a fraction of group-specific average income) over the same period (lines 2 and 5). Consider the period 2006–2008. Line 5 of the table shows that

wealth-rich households experience a decline in net worth equivalent to 137 percent of their annual income over this period, while wealth-poor households actually see net worth increase (line 2). A conventional wealth effects story would therefore predict larger declines in consumption rates for the rich, contrary to the pattern observed. Note also that households in the poor group increase their net worth by 83 percent of their income during the recession. This provides independent evidence that this group really is increasing saving. In addition, the net worth data confirm the popular narrative of deleveraging during the Great Recession: low wealth households cut expenditure to reduce debts and rebuild eroded balance sheets.

The long panel dimension of the PSID also allows us to compute future income growth for the two groups, which we take as a measure of income prospects. Table 2 reports the growth rates of income over the two-year period following the measured consumption rate changes (lines 3 and 6).⁷ As the economy enters the recession, expected future income growth declines for both the rich and poor groups, perhaps accounting for some of the decline in the aggregate consumption rate. However, the change in expected income growth is similar across wealth groups: the poor income growth changed by -8.8% (5.6-14.4) while the rich income growth changed by -7.6% (-0.2-7.4), suggesting that differential income prospects are not the primary factor behind the especially sharp decline in the consumption rate of the poor.

TABLE 2. CONSUMPTION, WEALTH & FUTURE INCOME, PSID

		2004-2006	2006-2008	2008-2010
POOR				
1.	Δ consumption rate (pp)	-2.0 (1.2)	-9.3 (1.2)	-0.0 (1.2)
2.	Δ net worth (% of income)	113 (15)	83 (32)	61 (11)
3.	Δ future income (%)	14.4 (1.6)	5.6 (1.3)	
RICH				
4.	Δ consumption rate (pp)	-0.8 (1.8)	-4.6 (1.1)	2.6 (1.1)
5.	Δ net worth (% of income)	189 (71)	-137 (50)	-22 (23)
6.	Δ future income (%)	7.4 (1.8)	-0.2 (1.9)	

Note: Bootstrapped standard errors are in parentheses.

⁷Thus, future income growth for 2004-2006 refers to income growth of the group between 2006 and 2008, and analogously for the 2006-2008 period. We cannot compute future income growth for the 2008-2010 period because data for 2012 are not yet available.

Overall, the evidence in Table 2 supports our hypothesis that the differential consumption rate changes shown in Figure 5 reflect a strong precautionary motive to save on the part of wealth-poor households, in the face of rising unemployment risk, and, more broadly, that the precautionary motive played a major role in accounting for the decline in aggregate consumption during the Great Recession. The magnitudes of the observed changes in consumption rates are economically relevant. For example, in the PSID over the period 2006-2008, poor households reduced their consumption rate by about 4 percentage points more than rich households. If we attribute this difference entirely to a stronger precautionary motive, then given that the poor account for about 1/3 of total disposable income (see Table 1), we can conclude that increased precautionary saving by the poor reduced aggregate consumption by $1/3 \times 4\% \simeq 1.3\%$ of aggregate disposable income.

This evidence shows that a rising precautionary motive to save may have played an important role in reducing consumer demand during the Great Recession. In the next section, we develop a simple macro model with a precautionary motive, and show that when wealth is sufficiently low, changes in precautionary saving can make perceived changes in unemployment risk self-fulfilling.

3 Theory

The model we will develop is stylized, and has two key ingredients. First, unemployment risk is assumed to be imperfectly insurable. This implies that when wealth is low, fear of high unemployment generates a precautionary motive to save, and thus a reduction in consumer demand. Second, wages are assumed to be rigid. This implies that a reduction in consumer demand will drive a fall in output and an increase in unemployment, validating the initial fear of unemployment. These two ingredients are central to the implication of the theory that will be our focus: when the level of household wealth is low, expectations of high unemployment can be self-fulfilling.

There are two goods in the economy: a perishable consumption good, c , produced by a continuum of identical competitive firms using labor, and housing, h , which is durable and in fixed supply normalized to one. There is a continuum of identical households, each of which contains a continuum of measure one of individuals. Households and firms share the same information set and have identical expectations. The price of the consumption good is normalized to one in each period. We focus on equilibria in which the price of housing relative to consumption, denoted by p , is fixed.

Let s_t denote the current state of the economy and s^t denote the history up to date t . Following each possible history s^t , the benevolent household head sends out its members to look for jobs in

the labor market, in which the probability of any given member finding a job – exogenous from the household’s perspective – is $1 - u(s^t)$. The household head can give its members consumption and savings instructions that are contingent on their idiosyncratic labor market outcomes, but housing wealth must be allocated across members before these outcomes are realized. The fraction $1 - u(s^t)$ of household members who find a job are paid a wage $w(s^t)$ and can use wage income and housing wealth (via direct sales or home equity borrowing) to pay for consumption $c^w(s^t)$. The fraction $u(s^t)$ who do not find a job and are unemployed can only use wealth and (potentially) unemployment benefits b to pay for consumption $c^u(s^t)$. At the end of the period, the household regroups and pools resources, which determines the value of housing wealth carried into the next period $h(s^t)$. Thus, while within household insurance is limited within the period, it is perfect between periods. This model of the household is a simple way to introduce idiosyncratic risk and a precautionary motive, without having to keep track of the cross-sectional distribution of wealth.

Preferences for a household are given by

$$\sum_{t=0}^{\infty} \beta^t \sum_{s^t} \pi(s^t) \{ [1 - u(s^t)] \log c^w(s^t) + u(s^t) \log c^u(s^t) + \phi \log h(s^{t-1}) \},$$

where β is the discount factor, $\pi(s^t)$ is the probability of history s^t as of date 0, and ϕ is a parameter determining the utility from housing. Note that utility is effectively Cobb-Douglas between housing and non-housing consumption, a specification consistent with Davis and Ortalo-Magne (2001).

Because employment opportunities are randomly allocated, it is optimal to give each household member an equal fraction $ph(s^{t-1})$ of the assets the household carries into the period. The household budget constraints therefore take the form

$$[1 - u(s^t)] c^w(s^t) + u(s^t) c^u(s^t) + p [h(s^t) - h(s^{t-1})] \leq [1 - u(s^t)] [w(s^t) - T(s^t)] + u(s^t) b \quad (1)$$

$$c^u(s^t) \leq ph(s^{t-1}) + b \quad (2)$$

$$c^w(s^t) \leq ph(s^{t-1}) + w(s^t) - T(s^t) \quad (3)$$

$$c^w(s^t), c^u(s^t), h(s^t) \geq 0.$$

The left-hand side of eq. (1) captures total household consumption and the cost of net housing purchases. The first term on the right-hand side is earnings for workers $w(s^t)$ less payroll taxes $T(s^t)$, and the second is unemployment benefits b for the fraction $u(s^t)$ of members who do not find a job. Note that $h(s^{t-1})$ was effectively chosen in the previous period. In the current period, given $u(s^t)$, $w(s^t)$, and p , the choices for $c^w(s^t)$ and $c^u(s^t)$ implicitly define the value of wealth carried into the next period, $ph(s^t)$. Equation (2) is the constraint that limits consumption of unemployed

members to wealth plus unemployment benefits. Equation (3) is the analogous constraint for workers.

3.1 Household's Problem

Let $\mu(s^t)$ denote the multiplier on (1) and let $\lambda(s^t)$ denote the multiplier on (2). We conjecture and later verify that the other constraints do not bind in equilibrium.

The first order conditions (FOCs) for $h(s^t)$, $c^w(s^t)$, and $c^u(s^t)$ can be combined to give

$$\frac{p}{c^w(s^t)} = \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t) \left[(1 - u(s^{t+1})) \frac{p}{c^w(s^{t+1})} + u(s^{t+1}) \frac{p}{c^u(s^{t+1})} \right] + \beta \frac{\phi}{h(s^t)},$$

where

$$c^u(s^{t+1}) = \begin{cases} c^w(s^{t+1}) & \text{if } c^w(s^{t+1}) \leq ph(s^t) + b \\ ph(s^t) + b & \text{if } c^w(s^{t+1}) > ph(s^t) + b. \end{cases}$$

This intertemporal condition can alternatively be written as

$$\frac{p}{c^w(s^t)} = \beta \sum_{s_{t+1}} \pi(s_{t+1}|s^t) \left[\frac{p}{c^w(s^{t+1})} \left(1 + \underbrace{\frac{u(s^{t+1}) \max \{ c^w(s^{t+1}) - [ph(s^t) + b], 0 \}}{ph(s^t) + b}}_{\text{liquidity value of wealth}} \right) \right] + \beta \frac{\phi}{h(s^t)}. \quad (4)$$

This first-order condition can be interpreted as follows. The utility cost of buying an additional unit of housing is the price times the marginal utility of consumption for a worker. The return is the discounted utility flow $-\beta\phi/h(s^t)$ – plus the next period price times next period marginal utility for workers plus an additional term that regulates the liquidity value of additional wealth in the next period. This liquidity value is proportional to the unemployment rate – which captures the number of household members who will value extra liquidity – times the difference in consumption for employed versus unemployed workers – which captures the value of being able to allocate consumption more evenly across household members. When either the unemployment rate is zero or the borrowing constraint is nonbinding for unemployed workers – so that employed and unemployed members enjoy equal consumption – this liquidity term drops out, and the intertemporal first-order condition takes the usual representative agent form. Conversely, when there is a positive probability of unemployment at $t + 1$ and when workers consume more than the unemployed, there will be a precautionary motive to save that will be larger the higher are expected unemployment rates. Further, and most important, the precautionary motive to save will be stronger the lower are expected house prices, since lower asset prices will imply a higher marginal utility of consumption

for unemployed household members.

3.2 Production and Labor Markets

Our stylized model of the labor market is constructed to ensure that output is determined by desired aggregate goods demand, rather than by the desired level of labor supply.⁸ We assume that the wage $w(s^t)$ is exogenously fixed at 1, which is a worker's physical marginal product. Given this wage, all workers would like to work, and firms are indifferent about how many workers to hire. The equilibrium employment level, $1 - u(s^t)$, will be the level at which demand for goods, $[1 - u(s^t)] c^w(s^t) + u(s^t) c^u(s^t)$, is equal to supply, $1 - u(s^t)$. In Appendix A we show how our fixed wage assumption can be micro-founded as a special case of the standard Mortensen and Pissarides framework.

Note that absent a sticky wage or other labor market frictions, labor market clearing would dictate a unique equilibrium at full employment. Thus, the rigid wage assumption is important for introducing the possibility of multiple values for equilibrium unemployment, as we will shortly describe.

3.3 Equilibrium

In general, this class of models has many equilibria, where different levels of output correspond to different values for asset prices. As we explained earlier, we are interested in exploring whether we can generate economic fluctuations driven by changes in risk, rather than by changes in asset prices. We therefore focus on equilibria in which the price of housing is fixed. A fixed price equilibrium is a constant house price p , an unemployment benefit b , a process for the state s_t (which later will be a sunspot), and associated quantities $u(s^t)$, $c^w(s^t)$, $c^u(s^t)$, $h(s^t)$, $T(s^t)$ that satisfy, for all t and for all s_t , the following:

$$[1 - u(s^t)] c^w(s^t) + u(s^t) c^u(s^t) = 1 - u(s^t) \quad (5)$$

$$h(s^t) = 1 \quad (6)$$

$$[1 - u(s^t)] T(s^t) = u(s^t) b \quad (7)$$

$$c^u(s^t) = \min \{c^w(s^t), ph(s^{t-1}) + b\} \quad (8)$$

⁸This is sometimes referred to as the "old-Keynesian" model. See, for example, Cochrane (2013).

$$\frac{p}{c^w(s^t)} = \beta \sum_{s_{t+1}} \pi(s_{t+1}|s^t) \left[\frac{p}{c^w(s^{t+1})} \left(1 + \frac{u(s^{t+1}) \max \{c^w(s^{t+1}) - [ph(s^t) + b], 0\}}{ph(s^t) + b} \right) \right] + \beta \frac{\phi}{h(s^t)} \quad (9)$$

Equation (5) is goods market clearing, (6) is housing market clearing, (7) is the government budget constraint, and (8) and (9) describe optimal consumption and savings behavior.

We now go on to show that if house prices are relatively high – reflecting a relatively high value for the taste parameter ϕ – then the model implies a monotonic steady state relationship between unemployment and house prices. Thus, for a given level of asset prices multiple equilibria are not possible and an increase in expected unemployment risk cannot generate a recession. In contrast, if housing prices are relatively low, then the same house price is consistent with two different steady state levels for unemployment. Second, we will show that, in the latter low-price case, the low unemployment steady state is locally stable, which introduces the possibility of confidence-driven fluctuations in unemployment. The mechanism is that when wealth is low, a jump in the expected path for unemployment implies a strong response of desired precautionary savings, which makes the prophecy of higher unemployment self-fulfilling.

Most of the analysis that follows focuses on a simple version of the model in which the government plays no role, so that $b = T(s^t) = 0$. We will return to consider unemployment insurance in subsection 4.4.

3.4 Steady States

Steady states are constant values (c^w, c^u, u, p) that satisfy equations (5), (8), and (9). Imposing $w(s^t) = 1$ and $h(s^t) = 1$, these equations can be written, respectively, as

$$\begin{aligned} (1 - u) c^w + u c^u &= 1 - u \\ c^u &= \min \{c^w, p\} \\ \frac{p}{c^w} &= \beta \frac{p}{c^w} \left(1 + \frac{u \max \{c^w - p, 0\}}{p} \right) + \beta \phi. \end{aligned}$$

Let $p_F(u)$ denote the fundamental price of housing: the price households would be willing to pay in steady state if there was perfect risk sharing within the household, so that $c^w = c^u$:

$$p_F(u) = \frac{\beta \phi}{1 - \beta} (1 - u). \quad (10)$$

The model has a continuum of steady states in which the unemployment rate ranges from $u = 0$

to $u = 1$, and the corresponding house price ranges from $p = p_F(0)$ at $u = 0$ to $p = 0$ at $u = 1$.

To characterize the relation between steady state u and p more sharply, it is useful to divide the parameter space into regions according to the value for ϕ .⁹ There are three regions of interest:

1. High house prices: $\phi \geq \phi_H$, where $\phi_H = \frac{(1-\beta)}{\beta}$. In this region, the steady state house price is a linear decreasing function of the steady state unemployment rate.
2. Moderate house prices: $\phi \in [\phi_L, \phi_H]$, where

$$\phi_L = \frac{1}{2} \sqrt{\left(\frac{4}{\beta} - 3\right)} - \frac{1}{2} < \phi_H. \quad (11)$$

In this region, the house price is a concave decreasing function of the unemployment rate.

3. Low house prices: $\phi < \phi_L$. Here, the house price is a hump-shaped function of the unemployment rate.

Figure 6 plots the set of steady states for three different values for ϕ , corresponding to each of these three different regions of the parameter space. The solid line corresponds to $\phi \geq \phi_H$ (high prices), the dashed line to $\phi \in [\phi_L, \phi_H]$ (moderate prices), and the dotted line to $\phi < \phi_L$ (low prices). For each value for ϕ , each combination (p, u) on the corresponding line constitutes a steady state.¹⁰

We now offer some intuition for each case.


First, if $\phi \geq \phi_H$ then housing has sufficient fundamental value to enable the household to equate consumption between employed and unemployed members, so $c^w(u) = c^u(u) = 1 - u$. House prices are therefore given by their fundamental value, $p(u) = p_F(u)$, which implies a linear inverse relationship between unemployment and house prices, as in the model analyzed by Farmer (2013). The steady state house price is declining in u because higher unemployment implies lower consumption, a higher marginal utility of consumption, and a lower willingness to exchange consumption for housing.

If $\phi \leq \phi_H$, then the fundamental component of house value is not sufficient to support full risk sharing within the household, and unemployed members are liquidity constrained: $c^u(u) = p(u)$ and $c^w(u) = 1 - \frac{u}{1-u}p(u) > c^u(u)$. In this case, the steady state relationship between unemployment

⁹See Appendix A for derivations of the following results.

¹⁰The actual parameter values used to construct the plot are $\beta = (1 + 0.05)^{-1}$ and $\phi = \phi_H$ (solid) $\phi = \phi_L$ (dashed) and $\phi = 0.0375$ (dotted).

Figure 6: Steady states for $\phi \geq \phi_H$ (solid), $\phi \in [\phi_L, \phi_H]$ (dashed), and $\phi < \phi_L$ (dotted)



and house prices is given by

$$p(u) = p_M(u) = \frac{\beta(u + \phi)}{(1 - \beta) + \frac{\beta u(1 + \phi)}{1 - u}} > p_F(u). \quad (12)$$

Housing now has additional value as a source of liquidity – the term labeled “liquidity value of wealth” in eq. (4) is positive. This liquidity component to house value can be defined as the equilibrium price (eq. 12) minus the fundamental component (eq. 10):

$$p_L(u) = p_M(u) - p_F(u). \quad (13)$$

At $u = 0$, the liquidity component $p_L(u)$ is increasing in the unemployment rate, while it shrinks to zero as $u \rightarrow 1$.

If $\phi \in [\phi_L, \phi_H]$, moderate house prices mean that the consumption differential between workers and the unemployed is modest. Thus the liquidity component housing value is relatively small, and the fundamental component remains the primary determinant of house value. The price function

$p_M(u)$ is a concave function of u , but is still declining in u . Thus, any p in the range of $p_M(u)$ corresponds to a unique steady state unemployment rate u .

If $\phi < \phi_L$, then house prices are lower still. The steady state price / unemployment relationship is still given by $p(u) = p_M(u)$. However, lower house prices now mean worse consumption insurance within the household, and a larger liquidity component to home value. This liquidity component is now so large that steady state house prices are increasing in the unemployment rate at $u = 0$.

Figure 7: Steady state price decomposition, $\beta = (1 + 0.05)^{-1}$, $\phi = 0.0375$




Figure 7 decomposes house prices $p_M(u)$ under our baseline (solid line) parameterization into fundamental and liquidity components. Suppose we start in the steady state with $u = 0$ and $p = p_M(0)$ and consider how the steady state price changes in response to a marginal increase in steady state unemployment. On the one hand, higher unemployment reduces expected income, and the fundamental component of the price $p_F(u)$. On the other hand, increasing unemployment raises the liquidity value of housing, $p_L(u)$, since the household has a stronger incentive to accumulate housing as an asset that members can use to smooth consumption through unemployment spells. The marginal liquidity value is initially strong (given $\phi < \phi_L$), because there is a large gap between the

consumption levels of employed and unemployed workers. This explains why the steady state house price is initially increasing in the unemployment rate. As unemployment rises further, however, the optimal consumption of workers declines, narrowing the consumption differential relative to the unemployed, and shrinking the liquidity component of house value.

The key feature of the low price regime is that for each $p \geq p_M(0)$ in the range of the $p_M(u)$ function, there are two steady state unemployment rates, which we denote $u_L(p)$ and $u_H(p)$, where $u_H(p) > u_L(p)$. In the low unemployment steady state, wealth is low relative to per capita consumption, but the household does not want to increase saving because there is low unemployment risk – and thus a modest precautionary motive to save. In the high unemployment steady state, unemployment risk is high, but the household does not want to increase saving further because wealth is already high relative to consumption. Thus, in the low unemployment steady state, the fundamental share of house value is higher (and the liquidity share lower) than in the high unemployment steady state.

3.5 Deterministic Dynamics and Sunspots

We have established that when fundamentals dictate low house values, two distinct steady state unemployment rates are consistent with the same price. We now show that, given the same fundamentals, there also exist equilibria with rich unemployment dynamics. To rule out changes in demand driven by wealth effects, and to isolate and emphasize our precautionary mechanism, we continue to focus on constructing equilibria in which the unemployment changes over time while asset prices are constant.

We start by considering unemployment dynamics in the perfect foresight case. We then show when and how one can introduce sunspots into the model to generate confidence-driven fluctuations in economic activity.

Imposing $p(s^t) = p$ and the housing market-clearing condition $h(s^t) = 1$, the perfect-foresight version of the intertemporal FOC (eq. 9) is

$$\frac{p}{c_t^w} = \beta \frac{p}{c_{t+1}^w} \left(1 + \frac{u_{t+1} \max \{c_{t+1}^w - p, 0\}}{p} \right) + \beta \phi, \quad (14)$$

where the consumption of the representative worker and the unemployment rate are linked via the resource constraint:


$$(1 - u_t) c_t^w + u_t \min \{c_t^w, p\} = 1 - u_t. \quad (15)$$

These two equations can be used to derive the dynamics for the unemployment rate for any initial unemployment rate u_0 .

In Appendix A we show that for a steady state pair (p^*, u^*) on the upward-sloping portion of the $p_M(u)$ function, a small deviation of the unemployment rate from u^* induces dynamics governed by eq. 14 that gradually return the economy to u^* . In contrast, for any steady state pair (p^*, u^*) on the downward-sloping portion of the $p_M(u)$ function, the dynamics are locally unstable. Recall that if $\phi \geq \phi_L$, then the $p_M(u)$ function is monotonically declining, and thus all steady states are locally unstable. In fact, if $\phi \geq \phi_L$, then for any constant price p , the only perfect foresight equilibrium features constant unemployment at the steady state rate corresponding to that price. We will therefore focus on $\phi < \phi_L$, in which case the $p_M(u)$ function is increasing in u at low unemployment levels, so that low unemployment steady states are locally stable.

To more fully characterize the global dynamics of the system, we plot a parametric example, in which $\beta = (1 + 0.05)^{-1}$, $\phi = 0.0375 < \phi_L$, and the house price is fixed at $p^* = p_M(0) = 0.75$. Figure 8 uses eqs. (14) and (15) to plot the change in the unemployment rate $u_{t+1} - u_t$ against the unemployment rate u_t .

Figure 8: Deterministic dynamics



The two points at which the change in the unemployment rate is zero correspond to the low and high unemployment steady states, $u_L(p^*)$ and $u_H(p^*)$. Because $u_L(p^*)$ is locally stable, there is a continuum of perfect foresight equilibria, indexed by an initial unemployment rate in the interval $[0, u_H(p^*)]$, each of which features a path for the unemployment rate that converges to $u_L(p^*)$.

In contrast, the high unemployment steady state $u_H(p^*)$ has the saddle-path property, so that if the unemployment rate starts slightly below $u_H(p^*)$ then unemployment will tend to fall, while if it starts above $u_H(p^*)$ it will rise. The unemployment paths that start out below $u_H(p^*)$ eventually converge to $u_L(p^*)$ and are valid perfect foresight equilibria. Paths that begin above $u_H(p^*)$ converge to maximum unemployment. These paths are not equilibria, because zero consumption cannot be optimal given positive wealth equal to p^* .

The fact that the system is locally stable in the neighborhood of $u_L(p^*)$ means that we can also construct sunspot equilibria in which agents take as given a positive probability of switching between boom and recession states. We will focus on perhaps the simplest equilibrium of this type: a two-state sunspot equilibrium, in which asset prices are constant and the unemployment rate bounces between zero and a positive value, with symmetric Markov transition probabilities. Let L and H denote the zero and positive unemployment states and let λ denote the probability that $s_{t+1} = L$ (H) given $s_t = L$ (H). This is now a three-parameter model, where the parameters are β , ϕ , and λ .

An interesting theoretical result is that such a sunspot equilibrium exists only if the sunspot process is sufficiently persistent. The lower bound for λ is given by¹¹

$$\bar{\lambda} = \frac{1}{2} \left(1 + \frac{1}{\beta} \left(\frac{\left(\frac{\beta\phi}{1-\beta} \right)^2}{\left(\frac{\beta\phi}{1-\beta} \right)^2 - \left(\frac{\beta\phi}{1-\beta} \right) + 1} \right) \right). \quad (16)$$

The logic for why sunspot fluctuations can only arise when λ is sufficiently high is as follows. Households reduce spending when the sunspot shock flips the economy into the positive unemployment state at date t because they anticipate a high likelihood that the unemployment rate will be high in $t + 1$, and thus they have a strong precautionary motive to save today. Iterating forward, expecting that the unemployment rate will likely be high at $t + 1$ (and thus that consumption will likely be low at $t + 1$) can only be rational if the unemployment rate is also likely to be high at $t + 2$. This logic dictates persistence in the unemployment process.¹²

¹¹See Appendix A for the derivation.

¹²Put differently, suppose one were to try to construct sunspot-driven cycles in which the sunspot process were iid. Households would then have no differential precautionary motive to save in the two states, but they would have an intertemporal motive to use wealth to support consumption in the high unemployment state. But then

From eq. (16), it is immediate that the larger is ϕ , the larger is $\bar{\lambda}$. Consider some special cases.

- For $\phi > 0$, $\bar{\lambda} > \frac{1}{2}$, and thus confidence-driven fluctuations must be persistent.
- As $\phi \rightarrow \phi_L$, $\bar{\lambda} \rightarrow 1$, and thus the zero and positive unemployment regimes must be near permanent.
- For $\phi \geq \phi_L$, there are no sunspot equilibria of this type. Thus, confidence-driven fluctuations are only possible if the fundamental component of housing value is sufficiently low.

The logic for why fixed-price sunspot fluctuations can only arise for low enough ϕ is closely related to the logic for while low ϕ is required for steady state multiplicity. In particular, for households to be willing to pay the same price for housing in the positive unemployment state as in the zero unemployment state, it must be that housing has sufficient additional liquidity value in the positive unemployment state. Housing only has significant liquidity value when the fundamental component of housing value is low.

4 Applications

We now apply the theory to business cycle fluctuations. We start in subsection 4.1 by using the model to interpret the time path for the unemployment rate in the United States over the course of the Great Recession. In subsection 4.2 we show that the theory implies a natural link between the depth of a recession and the speed of the subsequent recovery. In subsection 4.3 we revisit the link between wealth and volatility that we discussed in the Introduction. Finally, in subsection 4.4 we use the model to explore the potential role of unemployment insurance as a stabilization tool.


4.1 The Great Recession

Our model can generate dynamics for unemployment that are qualitatively similar to those experienced by the United States over the course of the Great Recession. Panel A of Figure 9 shows time paths for the unemployment rate and for house prices between the first quarter of 2005 and the first quarter of 2014. The house price series plotted is the Case-Shiller U.S. National Home Price Index, deflated by the GDP deflator, and relative to a 2 percent trend growth rate for the real

this consumption would translate into additional demand and employment, and the conjectured equilibrium would unravel.

price.¹³ Between the start of 2007 and the end of 2008, house prices fell by 30 percent relative to trend, largely accounting for the sharp fall in median net worth documented in Figure 1. The rise in the unemployment rate was concentrated in the second half of 2008 and the first half of 2009. Thus, the fall in house prices began well before the most severe portion of the recession.

Figure 9: House prices and unemployment



The starting point for the model simulation shown in Panel B of Figure 9 is a situation in which $\phi < \phi_L$, and thus the economy is in the region of the parameter space where house prices are relatively low. We assume that prior to 2009, the economy is in the low unemployment steady state, and set $p = p_M(0)$, so that this steady state features zero unemployment. We then compute the unemployment dynamics that arise due to an unanticipated shock to the expected path for unemployment, assuming no changes in the equilibrium house price.¹⁴

Thinking of a period length as a year, we set $\beta = (1 + 0.05)^{-1}$. We set $\phi = 0.0375$, which implies a full employment price $p_M(0) = 0.75$, and hypothetical consumption values at full employment of $c^w(0) = 1$ and $c^u(0) = 0.75$. Thus, our baseline choice for ϕ is consistent with Chodorow-Reich and Karabarbounis (2014), who document that households that experience a job loss experience a decline in household consumption of around 25 percent. The fall in consumption associated with

¹³This is the average growth rate for real GDP per capita between 1947 and 2007. It is also close to the average growth rate for real house prices between 1975 and 2006 (see Figure 1 in Davis and Heathcote, 2007).

¹⁴We focus on an equilibrium with constant house prices to emphasize that fluctuations in unemployment in our model do not require shocks to asset values. It would be straightforward to replicate the decline in house prices observed throughout 2007 and 2008 by introducing an earlier phase in the model simulation with a higher value for ϕ and thus higher house prices. Then a negative unanticipated shock to ϕ would drive the model decline in house prices.

unemployment is a natural calibration target, since the key mechanism we want to evaluate is the response of precautionary saving to changes in perceived unemployment risk.

In 2009 we hit the model economy with a one-time unanticipated shock to the expected unemployment path such that the unemployment rate jumps immediately to 10 percent. One possible trigger for this collective loss of confidence is the collapse of Lehman Brothers in the fall of 2008. Households cut back consumption – thereby rationalizing the surge in unemployment – because they now expect persistently high unemployment and therefore have a strong precautionary motive to save. From this point onward, households enjoy perfect foresight over the evolution of the unemployment rate, and the economy converges toward $u_L = 0$ according to the dynamics described in Figure 8.

Although this model is simple, it can replicate some key features of the Great Recession. The negative shock to expectations generates a deep and rapid contraction, followed by a very slow recovery. Moreover, these unemployment dynamics occur in the context of low but stable house prices: house prices were in fact quite stable between the start of 2009 and the start of 2014.

Why does the model predict a slow recovery? Part of the intuition is that a rapid expected recovery would imply a strong intertemporal motive to borrow and spend, which would cut against the precautionary motive to save in the near term. We expand on the logic for why confidence-driven fluctuations will typically be persistent in the next subsection.

Why does rising consumption during the recovery not drive up the equilibrium house price? The reason is that as the unemployment rate declines, so does the liquidity value of housing, and this declining liquidity value exactly offsets housing’s rising fundamental value.

Note that the possibility for a confidence-driven recession in the model hinges on ϕ being relatively low, so that house prices in the full employment steady state are relatively low, and the precautionary motive to save becomes strong when unemployment rises. The decline in house prices that occurred throughout 2007 and 2008 may therefore have played an important role in creating the pre-conditions for the Great Recession.


4.2 Deep Recessions and Slow Recoveries

When a low fundamental component to home value makes fixed price expectations-driven fluctuations possible, such fluctuations will tend to be larger in magnitude the more persistent is the underlying sunspot process. This theoretical link between persistence and amplitude can perhaps shed light on why the United States experienced such slow recoveries from the Great Depression

and the Great Recession – the two deepest recessions in the last century.

Figure 10 illustrates how the equilibrium value for $u(H)$ varies with the persistence parameter λ . In each case, $\beta = (1 + 0.05)^{-1}$ and $\phi = 0.0375$.

Figure 10: Larger recessions are more persistent



The figure indicates that more persistent fluctuations and larger fluctuations go hand in hand. At our calibrated values for β and ϕ the smallest value for the sunspot persistence parameter λ such a fixed price sunspot equilibrium exists is $\bar{\lambda} = 0.8635$ (see eq. 16). Thus, the expected duration of the unemployment state must be at least $1/(1 - 0.8635) = 7.3$ years. As $\lambda \rightarrow \bar{\lambda}$, the unemployment rate in the recession state $u(H)$ converges to the low unemployment steady state value, $u_L(p_F(0)) = 0$, and fluctuations vanish. As we increase λ , unemployment in the recession state rises: in the limit $\lambda \rightarrow 1$, $u(H)$ converges to the high unemployment steady state value, $u_H(p_F(0))$, and fluctuations are maximized.¹⁵

Why are more persistent fluctuations also larger in magnitude? At a basic level, standard

¹⁵In both limit cases ($\lambda \rightarrow \bar{\lambda}$ and $\lambda \rightarrow 1$) the equilibrium house price p converges to $p_F(0)$, while p is a hump-shaped function of λ for intermediate values for λ .

permanent income logic suggests that the longer the recession state is expected to last, the more sharply will households cut spending when the economy enters the recession state. In addition, when recessions are more persistent, equilibrium expected consumption growth is weaker. Thus, increasing persistence dampens the intertemporal motive to borrow and spend during a recession, thereby making it easier to support a deep precautionary-savings-driven recession.

4.3 Wealth and Volatility

We next construct a slightly different set of sunspot equilibria. We again focus on two-state Markov equilibria in which asset prices are constant. However, instead of focusing on equilibria in which $u(L) = 0$ and exploring the effects of varying λ , we instead fix λ and explore how the unemployment rates $u(L)$ and $u(H)$ vary with different choices for p . The goal is to explore how macroeconomic volatility – the difference between $u(H)$ and $u(L)$ – varies with wealth (conditional on a parameterization in which this sort of sunspot equilibrium exists).

In Figure 11 we fix $\lambda = 0.99$ and plot $u(H)$ (in red) and $u(L)$ (in blue) against p .¹⁶ The key point is that the larger is p , the smaller is the gap between the unemployment rates in the two states. Thus, higher asset prices imply less macroeconomic volatility, consistent with our characterization of U.S. macroeconomic history in Section 1. One way to understand this prediction is that in order to support a price considerably above the zero unemployment fundamental ($p_F(0) = 0.75$), the fundamental and liquidity components of home value must both be large in each state. A large liquidity component in the low unemployment state implies a relatively high value for $u(L)$, whereas a high fundamental component in the high unemployment state implies a relatively low value for $u(H)$.¹⁷


A second finding is that higher asset prices also translate into lower average unemployment rates over the cycle: in the equilibrium with maximum volatility ($p \approx 0.78$), the average unemployment rate is 9 percent, whereas in the equilibrium with least volatility ($p \approx 0.84$), the average unemployment rate is 6 percent. Thus, the model suggests that high *levels* of unemployment should go hand in hand with high *volatility* of unemployment.

To summarize, we have shown that sunspot-driven fluctuations with fixed asset prices are only possible when the taste for housing ϕ is sufficiently low. When such fluctuations can arise, fluctua-

¹⁶The dashed lines show the same objects for a lower persistence parameterization, with $\lambda = 0.98$.

¹⁷Note that in this class of sunspot equilibria, the equilibrium with the highest constant value for asset prices implies a constant positive unemployment rate. This (p, u) combination corresponds to one of the steady states illustrated in Figure 7. For each different value for λ , the corresponding highest price equilibrium corresponds to a different steady state (p, u) combination.

Figure 11: Wealth and volatility



tions are larger the lower are asset values. Thus, the empirical evidence we presented earlier on the link between the level of wealth and macroeconomic volatility can be interpreted on two levels: (i) when asset values are sufficiently high, confidence-driven fluctuations cannot arise, and (ii) when asset values are lower, the amplitude of confidence-driven fluctuations is inversely related to the level of asset prices.

4.4 Unemployment Insurance

We have shown that when household wealth is low, the economy is vulnerable to fluctuations driven by self-fulfilling changes in perceived unemployment risk. The key mechanism is that when agents anticipate higher unemployment, their precautionary motive to save becomes stronger and they reduce consumption. A natural policy to consider in this context is unemployment insurance, since unemployment benefits shrink the gap between the consumption levels of workers and the unemployed, and should therefore dampen the sensitivity of precautionary savings to changes in unemployment risk. We therefore introduce an unemployment benefit b for unemployed workers financed by a tax $T(s^t)$ on workers. The government budget constraint is eq. (7).

We find that unemployment benefits have a large impact on the shape of the steady state

relationship between unemployment and house prices, because they shrink the liquidity component of house value. The larger is b , the smaller is the value of ϕ below which the model has multiple steady state unemployment rates for the same price, so that fluctuations driven solely by changes in unemployment risk are possible. In fact, with positive unemployment benefits, the threshold becomes $\tilde{\phi}_L = \phi_L - \left(\frac{1-\beta}{\beta}\right)b$, where ϕ_L is defined in eq. (11).¹⁸

Conversely, holding fixed ϕ , a sufficiently generous unemployment benefit $b \geq \tilde{b}$ rules out the sort of sunspot-driven fluctuations we have studied, where

$$\tilde{b} = \frac{\beta}{1-\beta} \left(\frac{1}{2} \sqrt{\frac{4}{\beta}} - 3 - \frac{1}{2} - \phi \right).$$

Given our baseline parameter values, $\tilde{b} = 0.204$.

We conclude that our stylized model highlights an important feature of unemployment insurance, namely its ability to prevent the emergence of fluctuations arising from self-fulfilling changes in expected unemployment risk.¹⁹

5 Related Literature

On the empirical side, our model is related to a large literature which relates individual expenditures to labor income risk and to wealth, in order to assess the importance of the precautionary motive for consumption dynamics. Using British micro data, Benito (2006) finds that more job insecurity (using both model-based and self-reported measures of risk) translates into lower consumption. Importantly for the mechanism in our model, he finds that this effect is stronger for groups that have little household net worth. Engen and Gruber (2001) exploit state variation in unemployment insurance (UI) benefit schedules and estimate that reducing the UI benefit replacement rate by 50 percent for the average worker increases gross financial asset holdings by 14 percent. Carroll (1992) argues that cyclical variation in the precautionary savings motive explains a large fraction of cyclical variation in the savings rate.

Carroll et al. (2012) find that increased unemployment risk and direct wealth effects played the dominant roles in accounting for the rise in the U.S. savings rate during the Great Recession.

¹⁸This condition can be derived by substituting $c^w(u) = 1 - \frac{u}{1-u}p - \frac{ub}{1-u}$ and $c^u(u) = p + b$ into the steady state version of the household's first order condition, computing $\frac{dp}{du}$ by implicit differentiation, and solving for the value $\tilde{\phi}_L$ such that at $\phi = \tilde{\phi}_L$, $\frac{dp}{du} = 0$ at $u = 0$.

¹⁹A complete analysis of unemployment insurance would obviously need to consider other costs of such a policy (i.e. incentives) and would ideally micro-found why private unemployment insurance markets do not exist. We have simply assumed that the government (but not households) can access a technology that allows it to reshuffle resources between unemployed and employed household members within the period.

Mody et al. (2012) similarly conclude that the global decline in consumption was largely due to an increase in precautionary saving. Alan et al. (2012) exploit age variation in savings responses in U.K. data to discriminate between increased precautionary saving driven by larger idiosyncratic shocks versus the direct effects of tighter credit. They conclude that a time-varying precautionary motive was the key factor: tighter credit, in their model, mostly affects the young, whereas all age groups increased saving. Mian and Sufi (2010 and 2015) and Baker (2015) use, respectively, county and household-level data to show that consumption declines during the Great Recession were larger for units with lower initial net worth, evidence again consistent with a heightened precautionary motive. Kaplan et al. (2014) argue that the number of households for whom the precautionary motive is strong might be much larger than would be suggested by conventional measures of net worth, since there is a large group of households with highly illiquid wealth.

On the theory side, there is a long tradition of models in which self-fulfilling changes in expectations generate fluctuations in aggregate economic activity (see Cooper and John, 1988, for an overview). A classic early contribution is Diamond (1982), who constructs a model in which the expected presence of more trading partners makes trade easier, thereby stimulating production and generating the existence of more trading partners. In Farmer (2013, 2014) the labor market features search and matching frictions. Rather than assuming Nash bargaining over wages, he assumes that households form expectations – tied to asset prices – about the level of output, and that wages then adjust to support the associated level of hiring. Chamley (2014) constructs a model in which different equilibria are supported by differences in the strength of the precautionary motive to save, as in our model. In the low output equilibrium, individuals are reluctant to buy goods because they are pessimistic about their future opportunities to sell goods and because credit is restricted. In Kaplan and Menzio (2014), multiplicity is driven by a shopping externality: when more people are employed, the average shopper is less price sensitive, thereby increasing firms’ profits and spurring vacancy creation. Bacchetta and Van Wincoop (2013) note that with strong international trade linkages, expectations-driven fluctuations will necessarily tend to be global in nature.

Perhaps the most important difference between all these papers and ours is that we focus on the role of household wealth in determining when self-fulfilling fluctuations can arise. In most models that admit nonfundamental-driven fluctuations, the theory has little to say about when fluctuations should occur. In contrast, we have argued that a precondition for a confidence-driven recession is a low level of household wealth.

Guerrieri and Lorenzoni (2009), Challe et al. (2015), den Haan et al. (2016), and Ravn and Sterk (2014) all emphasize the role of precautionary savings as a mechanism that amplifies

fundamental shocks, but none of these papers consider the possibility of self-fulfilling precaution-driven fluctuations. In Beaudry et al. (2014) the precautionary savings channel amplifies a negative demand shock – via higher unemployment risk – but in their model, the impetus to low demand is excessively high past wealth accumulation, whereas we emphasize vulnerability when wealth is low.

Our emphasis on the role of asset values in shaping the set of possible equilibrium outcomes is shared by the literature on bubbles in production economies. Martin and Ventura (2014) consider an environment in which credit is limited by the value of collateral. Alternative market expectations can give rise to credit bubbles, which increase the credit available for entrepreneurs and therefore generate a boom (see also Kocherlakota, 2009). Hintermaier and Koeniger (2013) link the level of wealth to the scope for equilibrium multiplicity in an environment in which sunspot-driven fluctuations correspond to changes in the equilibrium price of collateral against consumer borrowing.

In other papers that emphasize a link between asset values and volatility, causation generally runs from volatility to asset prices. For example, Lettau et al. (2008) point out that higher aggregate risk should drive up the risk premium on risky assets relative to safe assets. Lower prices for risky assets like housing and equity then just reflect higher expected future returns on these assets. In our model, asset prices are the primitive, and the level of asset prices determines the possible range of equilibrium output fluctuations (i.e., macroeconomic volatility).

Our emphasis on the role of confidence is also a feature of Angeletos and La’O (2013), in which sentiment shocks (i.e., shocks to expectations about other agents behavior) can lead to aggregate fluctuations. Angeletos et al. (2014) develop a quantitative dynamic business cycle model that builds on similar ideas.

A general challenge in constructing models in which a notion of demand plays an important role in driving fluctuations is that many forces that tend to reduce desired consumption (such as lower asset values or greater idiosyncratic risk) also tend to increase desired labor supply. For this reason, models that emphasize the demand channel – including ours – need to ensure that increased desired labor supply does not automatically increase equilibrium output. Hall (2005), Farmer (2013), Michaillat (2012), and Shimer (2012) all exploit the fact that there are different possible ways to split the match surplus in search matching models, and negative aggregate shocks to the economy can be amplified if wages do not decline much in response. An alternative is to simply assume that real wages are sticky (see, for example, Midrigan and Philippon, 2011). Our approach blends both approaches: wages do not respond to unemployment in equilibrium precisely because the labor market is frictional.

6 Conclusions

The message of this paper is that when household wealth is low, a decline in consumer confidence can be self-fulfilling, because with low wealth, higher unemployment risk implies a large increase in the precautionary motive to save, rationalizing low equilibrium consumption and output. We argued that the decline in U.S. house prices in 2007 and 2008 reduced U.S. household net worth and left the U.S. economy vulnerable to just such a self-fulfilling wave of pessimism.

We developed a simple model to characterize the conditions under which confidence-driven fluctuations can arise and to better understand the link between the level of wealth on the one hand and the volatility and persistence of fluctuations on the other. Precautionary motives that vary with wealth and unemployment risk play the key role in these links. These precautionary motives are central in standard intertemporal consumption theory and are consistent with consumption patterns observed in micro data during the course of the Great Recession.

An obvious project for future work would be to develop a richer, more quantitative version of the model. One interesting extension would be to introduce household heterogeneity in net worth, to better understand the implications for aggregate precautionary demand of the extremely uneven distribution of net worth in the United States. A richer model of labor markets, in which desired labor supply plays a role in the long-run adjustment process, is another direction for future research. Finally, it would be interesting to introduce capital in the environment, since an increase in perceived unemployment risk would both increase the supply of saving and reduce the demand for investment, thereby changing the contours of the macroeconomic response.

References

- [1] Aiyagari, S.R. (1994) “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659-684.
- [2] Alan, S., T.F. Crossley, and H. Low (2012) “Saving on a Rainy Day, Borrowing for a Rainy Day,” Institute for Fiscal Studies Working Paper W12/11.
- [3] Angeletos, G.M., and J. La’O (2013) “Sentiments,” *Econometrica*, 81(2), 739-779.
- [4] Angeletos, G.M., F. Collard, and H. Dellas (2014) “Quantifying Confidence,” Working Paper, MIT.
- [5] Bacchetta, P., and E. Van Wincoop (2013) “The Great Recession: A Self-Fulfilling Global Panic,” Working Paper, University of Virginia.
- [6] Baker, S. (2015) “Debt and the Consumption Response to Household Income Shocks,” Working Paper, Northwestern University.

- [7] Barro, R.J., and H.I. Grossman (1971) "A General Disequilibrium Model of Income and Employment," *American Economic Review* 61(1), 82-93.
- [8] Beaudry, P., D. Galizia, and F. Portier (2014) "Reconciling Hayek's and Keynes' Views of Recessions," NBER Working Paper 20101.
- [9] Benito, A. (2006) "Does Job Insecurity Affect Household Consumption?" *Oxford Economic Papers* 58(1), 157-181.
- [10] Carroll, C.D. (1992) "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence" *Brookings Papers on Economic Activity* 23(2), 61-156.
- [11] Carroll, C.D., J. Slacalek, and M. Sommer (2012) "Dissecting Saving Dynamics: Measuring Wealth, Precautionary, and Credit Effects," IMF Working Paper 12/219.
- [12] Challe, E., J. Mathéron, X. Ragot, and J. Rubio-Ramirez (2015) "Precautionary Saving and Aggregate Demand," Banque de France Working Paper No. 535.
- [13] Challe, E., and X. Ragot (2012) "Precautionary Saving Over the Business Cycle," Working Paper, Paris School of Economics.
- [14] Chamley, C. (2014) "When Demand Creates Its Own Supply: Saving Traps," *Review of Economic Studies* 81(2), 651-680.
- [15] Chodorow-Reich, G., and L. Karabarbounis (2014) "The Cyclicalities of the Opportunity Cost of Employment," Working Paper, University of Chicago.
- [16] Cochrane, J. (2013) "New vs. Old Keynesian Stimulus," The Grumpy Economist Blog, November 8, 2013.
- [17] Cooper, R., and A. John (1988) "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics* 103(3), 441-463.
- [18] Davis, M.A., and J. Heathcote (2007) "The Price and Quantity of Residential Land in the United States," *Journal of Monetary Economics* 54(8), 2595-2620.
- [19] Davis, M.A., and F. Ortalo-Magne (2011) "Household Expenditures, Wages, Rents," *Review of Economic Dynamics*, 14(2), 248-261.
- [20] Den Haan, W.J., P. Rendahl, and M. Riegler (2016) "Unemployment (Fears) and Deflationary Spirals," Working Paper, London Business School.
- [21] Diamond, P.A. (1982) "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90(5), 881-894.
- [22] Engen, E.M., and J. Gruber (2001) "Unemployment Insurance and Precautionary Saving," *Journal of Monetary Economics* 47(3), 545-579.
- [23] Farmer, R.E.A. (2013) "Animal Spirits, Financial Crises and Persistent Unemployment," *Economic Journal* 123(568), 317-340.
- [24] Farmer, R.E.A. (2014) "The Evolution of Endogenous Business Cycles," *Macroeconomic Dynamics*, 1-14.

- [25] Grebler L., D.M. Blank, and L. Winnick (1956) *Capital Formation in Residential Real Estate: Trends and Prospects*, Princeton, N.J.: Princeton University Press.
- [26] Guerrieri, V., and G. Lorenzoni (2009) "Liquidity and Trading Dynamics," *Econometrica* 77(6), 1751-1790.
- [27] Hall, R.E. (2005) "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review* 95(1), 50-65.
- [28] Hintermaier, T., and W. Koeniger (2013) "Collateral Constraints and Self-Fulfilling Macroeconomic Fluctuations," Working Paper, University of Bonn.
- [29] Huggett, M. (1993) "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control* 17(5-6), 953-969.
- [30] Kaplan, G., and G. Menzio (2014) "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," Working Paper, University of Pennsylvania.
- [31] Kaplan, G., G.L. Violante, and J. Weidner (2014) "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* (Spring).
- [32] Kocherlakota, N. (2009) "Bursting Bubbles: Consequences and Cures," Mimeo, Federal Reserve Bank of Minneapolis.
- [33] Lettau, M., S.C. Ludvigson, and J.A. Wachter (2008), "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies*, 21(4), 1653-1687.
- [34] Martin, A., and J. Ventura (2014) "Managing Credit Bubbles," Working Paper, CREI.
- [35] Mian, A., K. Rao, and A. Sufi (2013) "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128(4), 1687-1726.
- [36] Mian, A., and A. Sufi (2010) "Household Leverage and the Recession of 2007 to 2009," *IMF Economic Review*, 58(1), 74-117.
- [37] Mian, A., and A. Sufi (2015) "Who Bears the Cost of Recessions? The Role of House Prices and Household Debt," Working Paper, Princeton University.
- [38] Michaillat, P. (2012) "Do Matching Frictions Explain Unemployment? Not in Bad Times," *American Economic Review* 102(4), 1721-1750.
- [39] Midrigan, V., and T. Philippon (2011) "Household Leverage and the Recession," Working Paper, New York University.
- [40] Mody, A., F. Ohnsorge, and D. Sandri (2012) "Precautionary Savings in the Great Recession," IMF Working Paper 12/42.
- [41] Ravn M., and V. Sterk (2014) "Job Uncertainty and Deep Recessions," Working Paper, University College London.
- [42] Shimer, R. (2012) "Wage Rigidities and Jobless Recoveries," *Journal of Monetary Economics* 59, S65-S77.

Appendix A.

Microfoundations for Fixed Wages

In the model described in the text, we simply assumed a fixed wage, $w(s^t) = 1$. We now describe one possible way to micro-found this assumption.

Workers and firms meet in a decentralized search and matching environment, à la Mortensen and Pissarides. Households send out workers to search for jobs, while firms decide how many vacancies $v(s^t)$ to post. The number of matches formed is the minimum of the number of vacancies posted, $v(s^t)$, and the number of searching workers, which will be of measure one in equilibrium. Once matched, a worker-firm pair Nash-bargain over the wage, $w(s^t)$. If they agree on a mutually satisfactory wage, the worker-firm pair produces one unit of the consumption good. Thus total output is equal to the employment level: $y(s^t) = 1 - u(s^t)$. All employment relationships are then dissolved, and all output is sold on a centralized market.

We assume that it is costless to create vacancies, and that the worker is endowed with 100% of the bargaining power at the wage bargaining stage. Given these assumptions, there is an equilibrium with the following properties. At the start of the period, all households and firms co-ordinate expectations on anticipated period employment $1 - u(s^t)$ and demand $d(s^t) \equiv [1 - u(s^t)]c^w(s^t) + u(s^t)c^u(s^t)$. Firms then post $v(s^t) = d(s^t)$ vacancies, each of which generates an employment opportunity. Each such opportunity leads to a successful match, so that $1 - u(s^t) = v(s^t)$, and the bargained wage in each match is $w(s^t) = 1$, which is the full match surplus given that each worker's physical marginal product is equal to one and that output can be sold at a price of one. Note that given positive anticipated wages, it is optimal for each household to send each potential worker out to search. At the same time, given an anticipated wage equal to one and costless vacancy creation, firms are indifferent about hiring any number of workers, including the quantity that clears the goods market, $1 - u(s^t) = d(s^t)$.

With a centralized labor market and flexible wages, any incipient involuntary unemployment would set in motion a Walrasian adjustment process that would equate labor demand and labor supply at the full employment level. No such process occurs in this frictional decentralized model. In particular, as in Barro and Grossman (1971), there is nothing to rule out situations in which anticipated output falls short of potential, and some workers end up involuntary unemployed (recall that all workers are willing to work at any positive wage). Wages do not fall in such a situation because of a hold-up problem: once matched, workers demand the full match surplus.

Steady State (u, p) Relationship

Perfect Risk Sharing with $\phi \geq \phi_H$ We first guess and verify that if $\phi \geq \phi_H$, then at any unemployment rate $u \in [0, 1]$ there is a steady state in which $p(u) = p_F(u) = \frac{\beta\phi(1-u)}{1-\beta}$, and $c^w(u) = c^u(u) = 1 - u$.

If $c^w(u) = c^u(u) = 1 - u$, then risk sharing is perfect, and $p(u) = p_F(u)$ is the house price that satisfies the household's inter-temporal first-order condition. It only remains to verify that perfect risk sharing is feasible, i.e., that

$$p_F(u) \geq c^w(u) = 1 - u.$$

But this condition is satisfied iff $\phi \geq \phi_H$.

Imperfect Risk Sharing with $\phi < \phi_H$ We now guess and verify that if $\phi < \phi_H$ then at

any unemployment rate $u \in [0, 1]$ there is a steady state in which $p(u) = p_M(u) = \frac{\beta(u+\phi)}{(1-\beta) + \frac{\beta u(1+\phi)}{1-u}}$, and $c^u(u) = p_M(u)$, $c^w(u) = 1 - \frac{u}{1-u}p_M(u)$.

Substituting the expressions for $c^u(u)$ and $c^w(u)$ into the household's first-order condition, and solving for the house price gives $p(u) = p_M(u)$. It only remains to verify that perfect risk-sharing is not feasible, i.e., that

$$p_M(u) < c^w(u) = 1 - \frac{u}{1-u}p_M(u).$$

This inequality is satisfied as long as

$$p_M(u) < (1-u),$$

which in turn is satisfied iff $\phi < \phi_H$.

Note that since $p_M(u) \geq p_F(u)$ for any u , there are no steady states with $\phi \geq \phi_H$ in which risk sharing is imperfect, and there are no steady states with $\phi < \phi_H$ in which risk sharing is perfect.

Concavity of $p_M(u)$ with $\phi < \phi_H$ Given $\phi \leq \phi_H$,

$$\frac{\partial^2 p_M(u)}{\partial u^2} \propto \frac{\frac{\beta}{1-\beta}\phi - 1}{1 - u + u\frac{\beta}{1-\beta}(1+\phi)}.$$

The numerator is negative by virtue of $\phi \leq \phi_H$. A sufficient condition for denominator to be positive (and thus for the price function to be concave) is $\beta \geq 0.5$.

Monotonicity of $p_M(u)$ with $\phi \geq \phi_L$

If $\phi \leq \phi_H$, then $p(u) = p_M(u)$ and

$$\frac{\partial p_M(u)}{\partial u} \propto \kappa_1 = -(2u + \beta - 2u\beta + \beta\phi + 2u^2\beta + \beta\phi^2 - u^2 + u^2\beta\phi - 1).$$

It is immediate that at $u = 0$

$$\left(\frac{\partial p_M(u)}{\partial u} \right)_{|u=0} = \begin{cases} > 0 & \text{for } \phi < \phi_L \\ 0 & \text{at } \phi = \phi_L \\ < 0 & \text{for } \phi > \phi_L. \end{cases}$$

Now, if $\phi \in [\phi_L, \phi_M]$, then because $p_M(u)$ is concave and declining in u at $u = 0$, $p_M(u)$ must be monotonically declining in u .

Inverse U-Shape of $p_M(u)$ with $\phi < \phi_L$ If instead $\phi < \phi_L$, then $p_M(u)$ is increasing in u at $u = 0$. Since $p_M(u)$ is a continuous and concave function of u , and since $p = 0$ at $u = 1$, the $p_M(u)$ function must be hump-shaped. Thus, there exists a set of prices above $p_M(0)$ such that any price p in the set corresponds to two distinct steady state values for u .

Local Stability / Instability of Steady States

Assuming the borrowing constraint for the unemployed is always binding, so that $c^u(u_t) = p$ and $c^w(u_t) = 1 - \frac{u_t}{1-u_t}p$, we can compute the dynamics for unemployment by implicitly differentiating the household's FOC (eq. 14):

$$\frac{\partial u_{t+1}}{\partial u_t} = \frac{p^2 (u_{t+1} + pu_{t+1} - 1)^2}{\beta (u_t + pu_t - 1)^2 (p^2 + pu_{t+1}^2 - p + u_{t+1}^2 - 2u_{t+1} + 1)}.$$

Now we explore whether or not the unemployment dynamics are locally stable at a steady state u . Imposing $u_{t+1} = u_t = u$, the expression for $\frac{\partial u_{t+1}}{\partial u_t}$ simplifies to

$$\left(\frac{\partial u_{t+1}}{\partial u_t} \right)_{|u_t=u} = \frac{p^2}{\beta (p^2 + pu^2 - p + u^2 - 2u + 1)}.$$

Unemployment is locally stable iff $\frac{\partial u_{t+1}}{\partial u_t} < 1 \Leftrightarrow 1 - \frac{\partial u_{t+1}}{\partial u_t} > 0$. Substituting in $p = p_M(u)$,

$$\begin{aligned} 1 - \left(\frac{\partial u_{t+1}}{\partial u_t} \right)_{|u_t=u} &\propto \kappa_2 \\ \kappa_2 &= \left(\frac{\beta(u + \phi)}{(1 - \beta) + \frac{\beta u(1 + \phi)}{1 - u}} \right)^2 \left(1 - \frac{1}{\beta} \right) + \left(\frac{\beta(u + \phi)}{(1 - \beta) + \frac{\beta u(1 + \phi)}{1 - u}} \right) u^2 - \left(\frac{\beta(u + \phi)}{(1 - \beta) + \frac{\beta u(1 + \phi)}{1 - u}} \right) + (1 - u)^2 \end{aligned}$$

But recall now that the condition for $p_M(u)$ to be increasing in u (i.e., for $\frac{\partial p_M(u)}{\partial u} > 0$) is

$$\kappa_1 = - (2u + \beta - 2u\beta + \beta\phi + 2u^2\beta + \beta\phi^2 - u^2 + u^2\beta\phi - 1) > 0.$$

Note that

$$\kappa_2 = \frac{1 - \beta(1 - u)}{\left(\frac{u}{1 - u}\beta(1 + \phi) + (1 - \beta) \right)^2} \kappa_1.$$

Since $\frac{1 - \beta(1 - u)}{\left(\frac{u}{1 - u}\beta(1 + \phi) + (1 - \beta) \right)^2} > 0$, the conditions are the same, i.e.,

$$1 - \left(\frac{\partial u_{t+1}}{\partial u_t} \right)_{|u_t=u} > 0 \Leftrightarrow \frac{\partial p_M(u)}{\partial u} > 0.$$

Derivation of Expression for $\bar{\lambda}$

The unemployment rates in the two states are $u(L) = 0$ and $u(H) > 0$. Assuming that unemployed workers are constrained, $c^u(H) = p$. The intertemporal FOCs in the zero and positive unemployment rate states are (again assuming that unemployed workers are constrained)

$$\begin{aligned} \frac{p}{c^w(L)} &= \beta(1 - \lambda)p \left([1 - u(H)] \frac{1}{c^w(H)} + u(H) \frac{1}{p} \right) + \beta\lambda \frac{p}{c^w(L)} + \beta\phi, \\ \frac{p}{c^w(H)} &= \beta\lambda p \left([1 - u(H)] \frac{1}{c^w(H)} + u(H) \frac{1}{p} \right) + \beta(1 - \lambda) \frac{p}{c^w(L)} + \beta\phi, \end{aligned}$$

where, from the resource constraint, consumption of workers is given by $c^w(L) = 1$ and $c^w(H) = 1 - u(H)p/[1 - u(H)]$.

A two state fixed price sunspot equilibrium exists if and only if there is a solution $\{u(H), c^w(H), c^w(L), p\}$

to these four equations that satisfies (i) $u(H) \in (0, 1]$, and (ii) $c^w(H) > p$.

Substituting the expressions for $c^w(L)$ and $c^w(H)$ into the household's first-order conditions gives

$$\begin{aligned} 1 &= \beta(1 - \lambda) \left(\frac{(1 - u(H))^2}{(1 - u(H)) - u(H)p} + u(H)\frac{1}{p} \right) + \beta\lambda + \frac{\beta\phi}{p}, \\ \frac{(1 - u(H))}{(1 - u(H)) - u(H)p} &= \beta\lambda \left(\frac{(1 - u(H))^2}{(1 - u(H)) - u(H)p} + u(H)\frac{1}{p} \right) + \beta(1 - \lambda) + \frac{\beta\phi}{p}. \end{aligned}$$

At $u(H) = 0$, both equations collapse to give $p = \beta\phi/(1 - \beta)$ for any values for (β, ϕ, λ) . Now consider the sensitivity of p to $u(H)$ in the neighborhood of $u(H) = 0$. Implicitly differentiating both first order conditions and equating the resulting expressions for $\frac{\partial p}{\partial u(H)}$ at $u(H) = 0$ and $p = \frac{\beta\phi}{1 - \beta}$ gives the solution for $\bar{\lambda}$ in the text.

Appendix B. Empirical Analysis

Total household net worth

Total net worth of U.S. households is computed as the sum of the following components of the Financial Accounts of the United States (Z1 release): (i) Households and nonprofit organizations; real estate at market value minus home and commercial mortgages, (ii) Households and nonprofit organizations; corporate equities, (iii) Households and nonprofit organizations; treasury securities, including U.S. savings bonds. For the period 1920-1944, these series are not available from the source above so we extend them as follows. The value of real estate is backcast using the growth rate of the value of total residential non-farm wealth in Grebler et al. (1956). Home mortgages are backcast using the growth rate of nonfarm residential mortgage debt from Grebler et al. (1956) and commercial mortgages are backcast using the growth rate of nonfarm commercial mortgage debt from the same source. The value of Treasury securities is backcast using the growth rate of the amount of public debt outstanding from the Treasury Department, and finally the value of corporate equities is backcast using the historical growth of the S&P 500 price index.

Micro Data

For each data set our key variables are net worth, disposable income and consumption expenditures. Below we first briefly describe the data sets, and then discuss the construction of these variables. The micro data used for the analysis are available on the authors' websites.

The Panel Study of Income Dynamics (PSID) is a panel of U.S. households, selected to be representative of the U.S. population, collected (starting from 1997) at a bi-annual frequency. Starting in 2004 the PSID reports, for every household in the panel, comprehensive consumption expenditure information, alongside information on income and wealth. Our panel includes all households which have at least one member aged between 22 and 60, which report yearly consumption expenditure of at least \$1,000, and which are in the panel for at least two consecutive interviews.

The Consumer Expenditure Survey (CES) is a rotating panel of U.S. households, selected to be

representative of the U.S. population, collected at a quarterly frequency. Households in the CES report information for a maximum of four consecutive quarters. Households report consumption expenditures in all four interviews, income information in the first and last interview, and wealth information in the last interview only. We use CES data from the first quarter of 2004 to the last quarter of 2013, and include all households which have at least one member aged between 22 and 60, which report yearly consumption expenditure of at least \$1,000, and which report consumption and income in the first and last interviews.

The Survey of Consumer Finance (SCF) is triennial survey of US households. The survey collects information on household income but focuses primarily on detailed information about household financial and non-financial assets and debts.

Net Worth In all three data sets we construct net worth by summing all categories of financial wealth (i.e. bank accounts, bonds, stocks) plus real estate wealth minus the value of any household debt (including mortgages, home equity loans and other debts). The PSID and SCF have a more accurate record of wealth, and report also the values of individual retirement accounts (IRAs), of family businesses, and of vehicles. Our measure of net worth in PSID and SCF also includes these variables.

Disposable Income In both data sets we construct disposable income by summing all money income received by all members of the household, including transfers, and then subtracting taxes. In the PSID we compute taxes using the NBER TAXSIM utility, while in the CES we use taxes paid as reported by the household.

Consumption Expenditure In both data sets we construct expenditure by summing the value of the purchases of: new/used cars and other vehicles, household equipment (including major appliances), goods and services used for entertainment purposes, food and beverages (at home and out), clothing and apparel (including jewelry), transportation services (including gasoline and public transportation), household utilities (including communication services such as telephone and cable services), education, and child care services. The two major categories that are excluded from our analysis are health expenditures and housing services. We exclude these categories to enable better comparison with NIPA data. Our key result regarding the differential behavior of consumption rates between rich and poor (shown in Figure 5) survives with consumption measures that include these two categories. We also experimented with a narrower consumption measure which excludes food, transportation services and utilities. The reason to consider excluding these categories is that households in the PSID are asked how much they spent on these in a typical week and not explicitly for the whole year (as for the other consumption categories). For this narrower consumption definition the discrepancy between aggregate consumption expenditures in the PSID and the other two data sources (Panel A of Figure 4) is much smaller. We have also reproduced Figure 5 with this narrower consumption measure, and found that the patterns of changes in consumption rates are very similar.

Measuring changes in consumption rates in the CES

We now outline the procedure used to produce Figure 5, Panel B.

1. Select households that (i) contain a head or spouse aged between 22 and 60, (ii) are interviewed

for the first time in year t (e.g. the first year in the sample), and (iii) report annual income and quarterly consumption in their first and last interviews, and report wealth in their last interview.

2. Rank these households by net worth in their last interview (the only time wealth is reported) relative to the average of consumption reported in the first and last interview, and divide the sample into two equal (weighted) size subgroups: the rich and the poor.
3. For each group, compute the consumption rate in the first interview (in year t) and the last interview. Note that because the first and last interviews are 9 months apart, for some households the last interview is at the end of year t while for the rest it is in year $t + 1$. The consumption rate is the average (annualized) consumption of the group in the quarter divided by the average disposable income of the group in the year.
4. Record for year t changes in consumption rates from the first to the last interview.
5. Move to year $t + 1$, and repeat Steps 1 through 4 to construct new rich and poor samples and to compute changes in consumption rates for year $t + 1$.

To produce Figure 5, Panel D, we follow the same procedure, except that to compute consumption rate changes over a two-year window, for a given wealth quintile, we add together the quintile-specific changes in the consumption rate over each of the two years in the window. Note that, in contrast to the PSID, the set of households used for a given quintile differs across the two years of a particular two-year window.