WEALTH AND VOLATILITY

Jonathan Heathcote
Fabrizio Perri

ABSTRACT

Periods of low household wealth in United States macroeconomic history have also been periods of
high business cycle volatility. This paper develops a simple model that can exhibit self-fulfilling fluctuations
in the expected path for unemployment. The novel feature is that the scope for sunspot-driven volatility
depends on the level of household wealth. When wealth is high, consumer demand is largely insensitive
to unemployment expectations and the economy is robust to confidence crises. When wealth is low,
a stronger precautionary motive makes demand more sensitive to unemployment expectations, and
the economy becomes vulnerable to confidence-driven fluctuations. In this case, there is a potential
role for public policies to stabilize demand. Microeconomic evidence is consistent with the key model
mechanism: during the Great Recession, households with relatively low wealth, ceteris paribus, cut
expenditures more sharply.

Jonathan Heathcote
Federal Reserve Bank of Minneapolis
Research Department
90 Hennepin Ave.
Minneapolis, MN 55401
heathcote@minneapolisfed.org

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

1 Introduction

Over the past 10 years 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
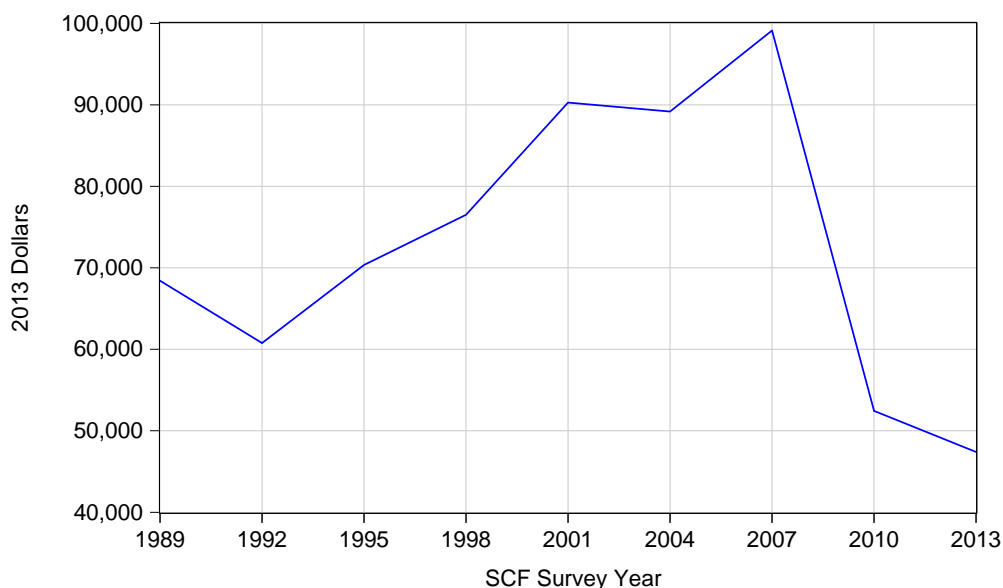
Note: Sample includes households with heads between ages 22 and 60.

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. Figures 2 and 3 provide some motivating evidence for our claim.

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

1

characterized by deep recessions and elevated macroeconomic volatility.[1]

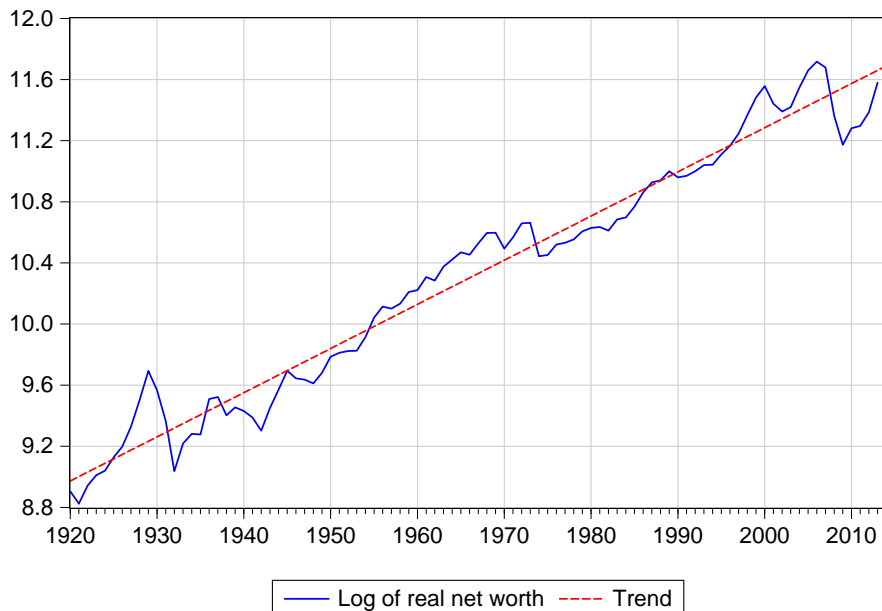Figure 2: Household net worth since 1920

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 rate 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 average deviation from trend (the difference between the blue and red lines in Figure 2) over the same 10-year window. The figure reveals that periods when wealth is high relative to trend, reflecting high prices for housing or stocks (or both), 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, wealth is well above trend and volatility is low.

Why should wealth affect volatility? We develop a microfounded dynamic equilibrium model in

---

[1]In 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

Note: Standard deviations of GDP growth are computed over 40-quarter rolling windows. Observations for net worth are averages over the same windows.
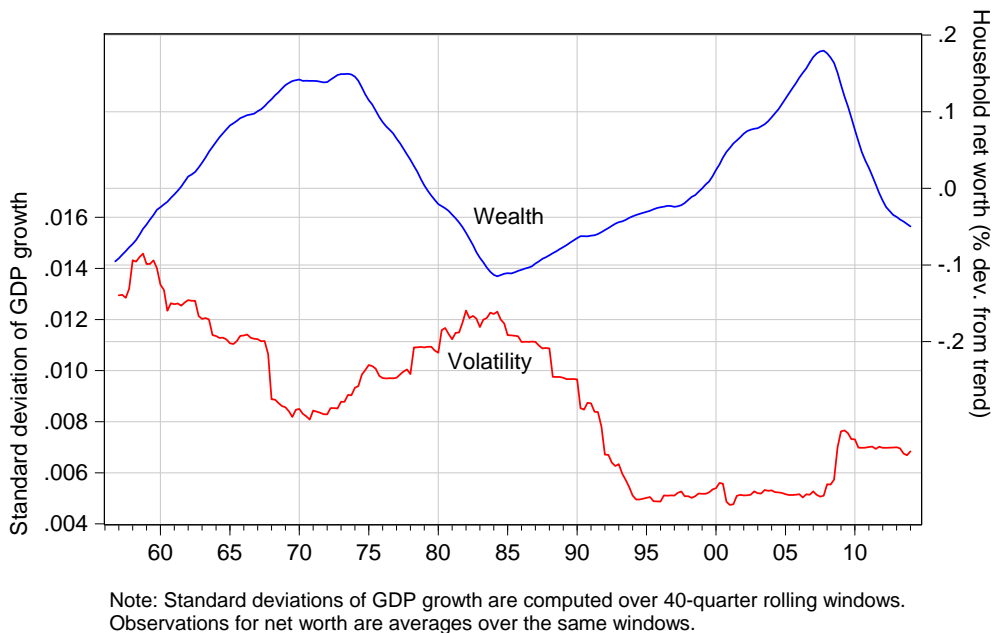
which economic fluctuations are driven by fluctuations in household optimism or pessimism. The novel feature is that the scope for equilibrium fluctuations due to "animal spirits" depends crucially on the value of wealth in the economy, which is in turn determined by fundamentals. When the fundamentals are such that wealth is high, the economy has a unique equilibrium and is not subject to confidence-driven fluctuations. When wealth is low, there are multiple possible equilibria, and the economy is vulnerable to confidence-driven fluctuations that are a source of macroeconomic volatility. In this case, there is a potential role for public policies to stabilize demand.

The model features a decentralized frictional labor market in which individual workers are either employed or unemployed. There is no explicit unemployment insurance, but savings can be used to smooth consumption in the event of an unemployment spell. Thus, positive unemployment risk generates a precautionary motive to save. 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

---

[2]Challe 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.

precautionary motive, which is the hallmark of incomplete markets models. We will heavily exploit one property of the model: higher household wealth makes precautionary saving demand (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.

Firms in the model are willing to hire up to the point at which they anticipate being able to sell the resulting output. Thus, the equilibrium employment rate is determined by consumption demand rather than by desired labor supply. Positive unemployment does not trigger expansionary price and quantity adjustments, since in our frictional environment workers cannot increase their probability of being matched with potential employers by signaling a willingness to accept lower wages – and once matched they have no incentive to do so.

Our theoretical analysis emphasizes the key role of the level of household wealth in determining the nature of equilibrium unemployment dynamics. When wealth is high, the precautionary motive to save is weak irrespective of the unemployment rate. Thus, high wealth rules out a confidence-driven collapse in demand and output. When wealth is low, on the other hand, the precautionary motive to save is strong and increases sharply with the expected unemployment rate. Thus, a recession driven by a self-fulfilling wave of pessimism becomes possible: if agents collectively expect higher unemployment, they all simultaneously reduce demand, leading to a fall in hiring and rationalizing the expected unemployment.

In addition to establishing that low asset values are a prerequisite for confidence-driven fluctuations, we develop other predictions of the theory that are useful for evaluating the confidence-driven interpretation of business cycles. First, we show that confidence-driven fluctuations must be persistent in expected terms, since a rapid expected recovery would cut against the incentive to cut current consumption. Second, similar logic implies that confidence-driven recessions that are especially deep are likely to be especially persistent. Third, confidence-driven fluctuations tend to be larger the lower are asset values.

These predictions square nicely with the view that the Great Depression and the Great Recession were both driven by a confidence-driven decline in demand. Both of these events involved large and sharp declines in output, followed by very sluggish recoveries. Both followed sharp declines in household wealth. Thus, in the context of our simple model, it seems plausible that a persistent collapse in confidence played an important role in these episodes.

Recessions in our model are associated with a desire to increase precautionary savings, which leads to an increase in the risk of unemployment, which in turn rationalizes the stronger precautionary motive. If this precautionary channel played an important role during the Great Recession, then one should expect low wealth households to have reduced consumption especially sharply, since their precautionary savings should be most sensitive to increased risk. We use micro data from the Consumer Expenditure Survey (CES) and the Panel Study of Income Dynamics (PSID) to test this implication of the theory. In both data sets we find that low net worth households systematically increased savings rates by much more than high net worth households around the onset of the recession.

Finally, we use the model to evaluate two specific policies designed to counteract the confidence-driven decline in demand that can fuel a model recession. Introducing a lump-sum unemployment benefit, financed by a tax on workers, is an effective policy. Unemployment benefits make unemployment less painful and thereby reduce the sensitivity of demand to the expected unemployment rate. A sufficiently generous benefit rules out sunspot-driven fluctuations and ensures full employment. Raising government consumption, again financed by taxing workers, turns out to be an ineffective policy. Higher government spending might be expected to make aggregate (private plus public) demand less sensitive to expectations, but it has an unintended general equilibrium effect. In particular, taxation reduces asset values, which reduces self-insurance against unemployment risk and thereby increases the sensitivity of demand to perceived unemployment risk. The net effect is that the set of steady state unemployment rates is insensitive to the level of government purchases.

1.1 Related Literature

Our paper is related to several strands of literature.

First and most important, there are other 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 will argue that a precondition for a confidence-driven recession is a low level of household wealth.

In Guerrieri and Lorenzoni (2009), risk-averse agents trade in a decentralized fashion and face idiosyncratic risk. Like us, they emphasize the importance of wealth and credit access, but in their model an endogenous increase in precaution amplifies a fundamental aggregate productivity shock, whereas in ours it is a self-fulfilling prophecy. 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 the 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 is similar to that of Barro and Grossman (1971): we assume a simple environment in which all workers would like to work, but in which firms take demand as effectively fixed and hire only up to the point at which desired demand is filled. Alternative approaches to making labor supply irrelevant would work equally well.

A key force in our model is the precautionary motive to save in the face of unemployment risk. A large literature documents the empirical importance of the precautionary motive. 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) investigate the Great Recession and 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 this episode. Mody et al. (2012) similarly conclude that the global decline in consumption

during the Great Recession 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 increases in precautionary saving driven by larger idiosyncratic shocks, versus the direct effects of tighter credit, and conclude that a time-varying precautionary motive plays the dominant role (tighter credit, in their model, mostly affects the young, whereas all age groups increased saving). Finally, 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.

2 Model

There are two goods in the economy: a perishable consumption good, produced by a continuum of identical competitive firms using labor, and a durable asset, which is in fixed supply and which we label housing. There are two types of households in the model and a continuum of identical households of each type. These types share common preferences but differ with respect to the risk they face: income for the first type is risky, whereas income for the second is not. There are measure zero households of the "riskless" type and their only role in the model is to establish a floor for asset prices.

Each household of the first "risky" type contains a continuum of measure one of individuals. The measure of firms is equal to the measure of risky households. Thus, we can envision a representative firm interacting with mass one members of a representative risky household. The price of the consumption good is normalized to one in each period. The quantity of housing is normalized to one. The economy is closed.

Let s_t denote the current state of the economy and s^t denote the history up to date t. In each period, households of the risky type send out members to buy consumption and to look for jobs. Employment opportunities are randomly allocated across household members, but assets must be allocated across members before labor market outcomes are realized. It is therefore optimal to give each member an equal fraction $h(s^{t-1})$ of the assets the household carries in the period. The household can give its members consumption and savings instructions that are contingent on their labor market outcomes. The fraction $1 - u(s^t)$ of household members who find a job are paid a wage $w(s^t)$ and use wage income and asset holdings to pay for consumption $c^w(s^t)$. The fraction

$u(s^t)$ who are unemployed can only use wealth and (potentially) unemployment benefits to pay for consumption $c^u(s^t)$.

At the end of the period, the household regroups and pools resources, which determines the quantity of the asset carried into the next period $h(s^t)$. 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.

At the start of each period t, households observe s_t, update s^t, and assign probabilities to future sequences $\{s_\tau\}_{\tau=t+1}^{\infty}$. We assume that all households form the same expectations.

Preferences for a household are given by

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

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.

The household budget constraints for a risky household have the form

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

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

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

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

The left-hand side of eq. (1) captures total household consumption and the cost of net asset 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 aggregate variables $u(s^t)$, $w(s^t)$, and $p(s^t)$, the choices for $c^w(s^t)$ and $c^u(s^t)$ implicitly define the quantity of wealth carried into the next period $h(s^t)$. Equation (2) is the constraint that limits consumption of unemployed members to wealth plus unemployment benefits. Equation (3) is the analogous constraint for workers.

The budget constraint for the riskless household is identical, except that unemployment and transfers for this type are equal to zero.

The government balances its budget period by period, using payroll taxes on workers to finance unemployment benefits and (possibly) government spending G:

$$\left[1 - u(s^t)\right] T(s^t) = u(s^t)b + G.$$

2.1 Household's Problem

Consider the problem for the type that faces unemployment risk. 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 $c^u(s^t)$ are, respectively,

$$p(s^t)\mu(s^t) = \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t) \left[p(s^{t+1})\mu(s^{t+1}) + p(s^{t+1})\lambda(s^{t+1})\right] + \beta\phi,$$

$$\frac{1}{c^w(s^t)} = \mu(s^t),$$

$$\frac{u(s^t)}{c^u(s^t)} = \mu(s^t)u(s^t) + \lambda(s^t).$$

Combining these gives

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

where

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

This intertemporal condition can alternatively be written as

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

(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 $(\beta\phi)$ 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.

The analogous first-order condition for the type that does not face unemployment risk is

$$\frac{p(s^t)}{\hat{c}^w(s^t)} \geq \beta \sum_{s_{t+1}} \pi(s_{t+1}|s^t) \left[\frac{p(s^{t+1})}{\hat{c}^w(s^{t+1})} \right] + \beta\phi,$$

where hats denote allocations for this type. The inequality here reflects the fact that, given the preferences we will assume below, the type facing no unemployment risk will be at a corner in equilibrium, with zero housing.

2.2 Production and Labor Markets

We now describe how workers and firms meet, how production takes place, and how the equilibrium unemployment rate is determined.

Households observe the aggregate state s_t, and the household head then gives its members a reservation wage $w^*(s^t)$ specifying what wages to accept, along with contingent consumption instructions, $c^w(s^t)$ and $c^u(s^t)$. Each of the measure one of potential workers from a particular household is then randomly matched with firms across the economy, so each firm ends up matched with measure one of potential workers.

Each period consists of a continuous unit interval of time. Production opportunities at each representative firm arrive smoothly through the interval, with potential workers arriving sequen-

tially, and – if hired – producing instantaneously, one worker at a time. The production technology is linear: hiring measure $n(s^t)$ workers produces $y(s^t)$ units of output:

$$y(s^t) = n(s^t).$$

Potential workers matched with a particular firm are allocated a random position in the production queue indexed by $i \in [0, 1]$. Each firm decides whether or not to hire each successive worker in the queue. Firms have no control over the arrival rate of potential workers and thus have no control over their maximum production rate. Thus, the optimal strategy for the firm is to employ a worker as long as the worker's reservation wage $w^*(s^t)$ is less than or equal to the worker's marginal product (which is equal to one) and as long as the firm has not already produced sufficient output to satisfy anticipated demand $d(s^t)$.

Understanding the firms' incentives, a representative household head will optimally assign all of its members a reservation wage $w^*(s^t) = 1$. The key assumption underlying this result is that any individual worker's production queue position is exogenous, and thus the household cannot push its members to the front of the queue – and increase their employment probability – by signaling a lower reservation wage. At the same time, a higher reservation wage would guarantee non-employment.

Firms will hire workers and produce continuously until they have produced $d(s^t)$ units of output. Thus, those workers with queue index $i \leq d(s^t)$ end up employed, whereas those with $i > d(s^t)$ are unemployed. Thus, the unemployment rate is demand determined: $u(s^t) = 1 - d(s^t)$.

Once production is completed, at the instant of time that is fraction $d(s^t)$ into the period, all output is immediately sold in a competitive centralized market at a price normalized to one. Firms have no interest in hiring workers with queue positions $i > d(s^t)$ because they anticipate no additional demand within the period, because workers cannot consume the output of the firm at which they produce, and because output cannot be stored and carried into the next period.

Note that if orders fall short of potential output (i.e., if $d(s^t) < 1$), then labor supply will exceed labor demand, in the sense that all measure one of workers in each household are willing to work at any positive wage, whereas employment is determined by labor demand $n(s^t) = d(s^t) < 1$. However, as in Barro and Grossman (1971), involuntary unemployment does not engender the standard Walrasian adjustment process that ultimately equates labor demand and labor supply in

models with frictionless labor markets.[3]

Note that households make consumption decisions and firms make production plans given the same information set and identical expectations.

2.3 Equilibrium

A symmetric equilibrium in this model is a pair of policy parameters (G, b), a process for the state s_t (which in some examples will be a sunspot), and associated decision rules and prices $n(s^t)$, $u(s^t)$, $d(s^t)$, $c^w(s^t)$, $c^u(s^t)$, $w(s^t)$, $h(s^t)$, $p(s^t)$, $T(s^t)$ that satisfy, for all t and for all s_t, the following:

$$d(s^t) = [1 - u(s^t)]c^w(s^t) + u(s^t)c^u(s^t) + G \tag{5}$$

$$u(s^t) = 1 - n(s^t) \tag{6}$$

$$d(s^t) = n(s^t) \tag{7}$$

$$h(s^t) = 1 \tag{8}$$

$$\left[1 - u(s^t)\right] T(s^t) = u(s^t)b + G \tag{9}$$

$$w(s^t) = w^*(s^t) = 1 \tag{10}$$

$$c^u(s^t) = \min\left\{c^w(s^t), p(s^t)h(s^{t-1}) + b\right\} \tag{11}$$

$$\frac{p(s^t)}{c^w(s^t)} = \beta \sum_{s_{t+1}} \pi(s_{t+1}|s^t) \left[\frac{p(s^{t+1})}{c^w(s^{t+1})}\left(1 + \frac{u(s^{t+1})\max\left\{c^w(s^{t+1}) - \left[p(s^{t+1})h(s^t) + b\right], 0\right\}}{p(s^{t+1})h(s^t) + b}\right)\right] + \beta\phi \tag{12}$$

$$\frac{p(s^t)}{1 - T(s^t)} \geq \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t)\frac{p(s^{t+1})}{1 - T(s^{t+1})} + \beta\phi. \tag{13}$$

These equations, respectively, define aggregate demand and the unemployment rate, impose

---

[3]The reason why involuntary unemployment can persist is not simply that the equilibrium real wage $w(s^t)$ is always equal to one. Suppose the government were to legislate a reduction in the real wage. Firms would then make positive profits, but given unchanged demand would not want to hire additional workers, since they would not be able to sell additional output. At the same time, a lower wage would not change household behavior either, since lower wage income within the period would simply imply correspondingly larger profit income at the end of the period, leaving the budget constraints (1) and (2) unaffected.

Barro and Grossman (1971) make the same point: "The conclusion is that too high a real wage was not the cause of the lower employment, and a reduction in the real wage is only a superficial cure. The real cause of the problem was the fall in commodity demand, and only a reflation of commodity demand can restore employment to the proper level" (pp. 86-87).

goods market clearing, asset market clearing, and the government budget constraint, and impose optimal reservation wage setting, optimal intratemporal resource allocation within the household, and optimal intertemporal savings behavior. Equation (13) indicates that the presence of the riskless type whose consumption is $\hat{c}(s^t) = 1 - T(s^t)$ puts a floor under house prices.[4]

3   Steady States

Most of the analysis that follows focuses on a simple version of the model in which the government plays no role, so that $b = G = T(s^t) = 0$. We will return to consider various policy interventions in Section 7.

Steady states are constant values (c^w, c^u, u, p) that satisfy equations (7) (goods market clearing), (11) (the budget constraint for unemployed members), (12) (the risky household's FOC), and (13) (the pricing floor established by the riskless type). These equations can be written, respectively, as

$$(1 - u)\, c^w + u c^u = 1 - u$$

$$c^u = \min\{c^w, p\}$$

$$\frac{p}{c^w} = \beta \frac{p}{c^w}\left(1 + \frac{u \max\{c^w - p, 0\}}{p}\right) + \beta\phi$$

$$p \geq \underline{p} \equiv \frac{\beta}{1 - \beta}\phi.$$

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 $c^w = c^u$:

$$p_F(u) = \frac{\beta\phi}{1 - \beta}(1 - u) \leq \underline{p} \tag{14}$$

with strict inequality for $u > 0$.

Proposition 1: Any steady state with positive unemployment must feature limited risk sharing: $u > 0 \implies c^w > c^u$.

---

[4]Note that with $h(s^t) = 1$, the intertemporal first-order condition for the risky household type (equation 12) would be identical if preferences were given by $u(c, h, \phi) = \log c + \phi \log h$. Thus, it is sufficient to assume linearity in preferences for the riskless type.

Proof: See Appendix A.

The logic for this result is that in any steady state with positive unemployment, expected individual consumption is less than one. If each household member consumed expected individual consumption, the price of housing would equal $p_F(u)$, which is below the price the riskless household (whose expected consumption is higher) is willing to pay, namely, $\underline{p} = p_F(0)$. It therefore follows that in steady states with positive unemployment, housing must have additional value as a source of liquidity for the risky household type. This in turn implies that in steady state, unemployed agents must be consuming less than employed households – the term labeled "liquidity value of wealth" in eq. (4) must be positive – so that the additional liquidity associated with housing wealth is priced.

Proposition 2: Let $\tilde{\phi} \equiv \frac{1}{2}\sqrt{\left(\frac{4}{\beta} - 3\right)} - \frac{1}{2}$. If $\phi \geq \tilde{\phi}$, then the only possible steady state is $p = \underline{p}$, $u = 0$. If $\phi < \tilde{\phi}$, then there exists a continuum of steady states in which the unemployment rate ranges from $u = 0$ to $u = u^+$ where

$$u^+ \equiv 1 - \frac{\beta}{1-\beta}\phi(1+\phi). \tag{15}$$

For each unemployment rate $u \in [0, u^+]$, the corresponding steady state house price is given by

$$p(u) = \frac{\beta(u+\phi)}{(1-\beta) + \frac{\beta u(1+\phi)}{1-u}}. \tag{16}$$

Proof: See Appendix A.

Note that the function $p(u)$ is concave and $p(0) = p(u^+) = \underline{p}$. A corollary follows:

Corollary: There is a range of values for $p \geq \underline{p}$ such that for any p in this range, there are two distinct steady state values for u. For $p > \underline{p}$, both of these steady states feature positive unemployment.

We define the liquidity value of housing, given a steady state unemployment rate u, as the equilibrium price (eq. 16) minus the fundamental component (eq. 14):

$$p_L(u) = \frac{\beta(u+\phi)}{(1-\beta) + \frac{\beta u(1+\phi)}{1-u}} - \frac{\beta\phi}{1-\beta}(1-u). \tag{17}$$

15

At $u = 0$, the liquidity value for housing is increasing in the unemployment rate, given $\phi \leq \tilde{\phi}$. The liquidity value shrinks to zero as $u \to 1$.

The nature of steady states, and the decomposition of housing value into fundamental and liquidity components, are best understood graphically. For the purposes of plotting numerical examples, we need to parameterize the model.

3.1 Parameterization

The model has only two parameters, β and ϕ. Thinking of a period length as a year, we set $\beta = (1 + 0.05)^{-1}$.

We then set our baseline value for ϕ so that the floor on house prices, $\underline{p} = \beta\phi/(1-\phi)$, is equal to 0.75, which implies $\phi = 0.0375$. This choice has two important features. First, $\phi < \tilde{\phi}$, guaranteeing, by virtue of Proposition 2, that the model will exhibit a continuum of steady states. The highest possible steady state unemployment rate $u^+ = 0.222$. Second, given $\underline{p} = 0.75$, in steady states with unemployment, unemployed workers will consume $c^u = 0.75$, which is 25 percent less than full employment consumption. This is consistent with estimates of the size of the consumption loss for households that experience a job loss (Chodorow-Reich and Karabarbounis, 2014).

3.2 Understanding Steady State Multiplicity

Figure 4 shows steady state equilibria for our baseline parameterization. Specifically the hump-shaped solid black line plots $p(u)$ from eq. (16), which corresponds to the steady state price that the risky household is willing to pay as a function of the unemployment rate. Each point on this line is a steady state. The green horizontal line shows $\underline{p}$: the lower bound on house prices established by the riskless household.

Suppose we start in the steady state with $p = \underline{p}$ and $u = 0$ and consider how the steady state price $p(u)$ changes in response to a marginal increase in unemployment. On the one hand, higher unemployment reduces expected income, reducing fundamental housing demand 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, because there is a large gap between the consumption

Figure 4: Steady states for $\beta = (1 + 0.05)^{-1}$, $\phi = 0.0375$
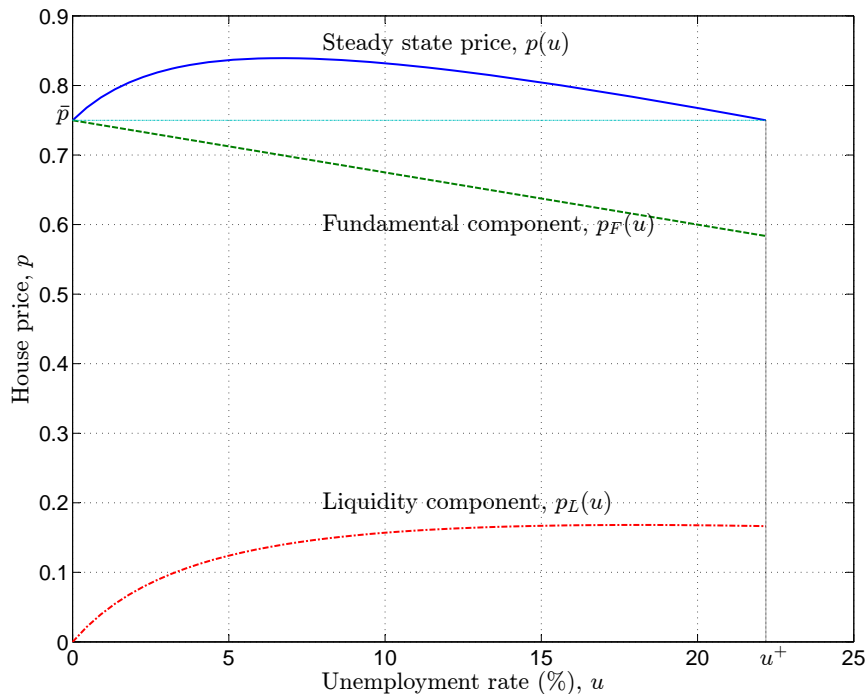
levels of employed and unemployed workers. This means that a marginal increase in unemployment (starting from $u = 0$) translates into an increase in the steady state asset price. But for high enough unemployment rates, the marginal negative impact on fundamental value comes to dominate, so that the steady state price becomes a declining function of u. Thus, there is a second equilibrium at $p = \underline{p}$ with $u = u^+$.

For each $p \geq \underline{p}$ in the range of the $p(u)$ function, there are two steady states, one with low and one with high unemployment. 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 equilibrium, 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 equilibrium, the fundamental share of house value is higher (and the liquidity share lower) than in the high unemployment equilibrium. Note that for $p > \bar{p}$, if steady states exist, the two steady state unemployment rates are closer together the larger is p. There are no steady states with $u > u^+$, because such high unemployment rates would imply values for p below the lower bound $\underline{p}$ established

by the marginal riskless household.

Figure 5 contrasts the baseline low ϕ parameterization described above to an alternative in which ϕ is larger and equal to $\tilde{\phi}$.

Figure 5: Steady states for $\phi = 0.0375$(black) and $\phi = \tilde{\phi}$ (red)
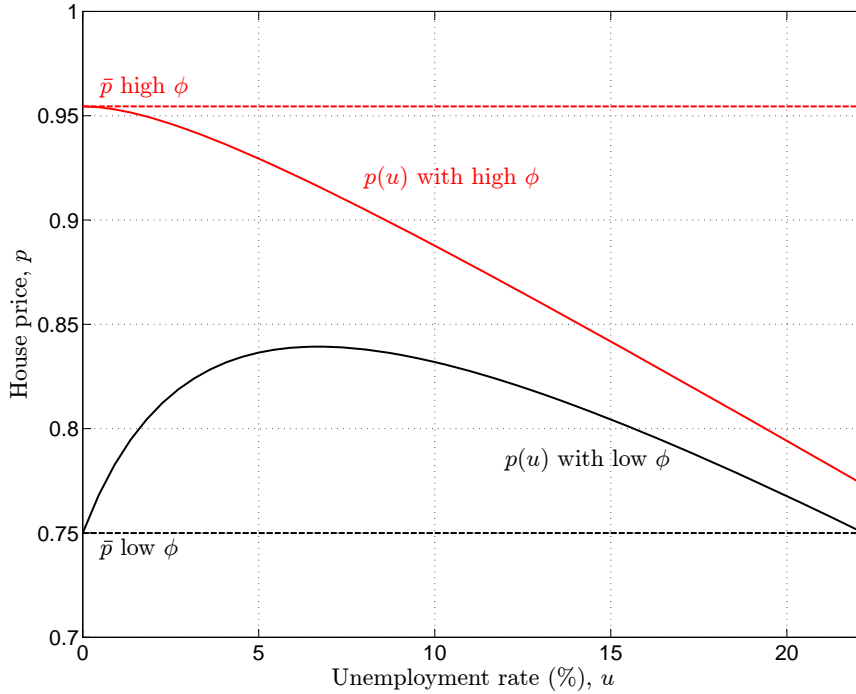
The top pair of red lines correspond to the case $\phi = \tilde{\phi}$, so that the taste for housing is high, while the bottom pair of black lines correspond to the baseline parameterization shown in the previous plot. In both cases, the solid lines depict the set of steady states, while the dashed lines show the respective price floors $\underline{p}$. Because the taste for housing is relatively strong with $\phi = \tilde{\phi}$, house prices are high, and as a consequence unemployed household members can afford consumption similar to employed household members. Thus, the liquidity value of housing is relatively low, and the fundamental component is the primary determinant of house value. As a consequence, the $p(u)$ curve is always (weakly) declining in u, and the equilibrium is therefore unique: $u = 0$ is the only steady state satisfying both $p = p(u)$ and $p \geq \underline{p}$. Zero is the only unemployment rate at which the steady state price $p_F(0)$ weakly exceeds the floor $\underline{p}$ established by the marginal riskless household.

Note that without the riskless type, there would be a continuum of steady states with unemployment rates between zero and one, with each unemployment rate corresponding to a different

steady state asset price as given by eq. (14) (see Farmer, 2013). The presence of the riskless type puts a floor on the asset price, which in turn establishes a floor for steady state consumption demand and output.

To summarize our steady state analysis, with strong demand for housing (high ϕ), the fundamental component of house prices is large, which translates into a weak precautionary motive and a relatively small liquidity component to house values. This in turn implies a unique full employment steady state. With weak demand for housing (low ϕ), house values are lower but are initially increasing in the unemployment rate, reflecting a high value for additional liquidity. This implies the existence of two possible steady state unemployment rates for the same price level.

4 Dynamics

We now introduce dynamics. Our focus will primarily be on constructing equilibria in which the unemployment changes over time while asset prices are constant, so that $p(s^t) = p \geq \underline{p}$. We start by considering unemployment dynamics in the perfect foresight case. We then show that one can introduce sunspots in the model and thereby generate confidence-driven fluctuations in economic activity. The ultimate goals of this section are twofold. The first goal is to characterize some general features of fluctuations driven by nonfundamental changes in expectations. The second is to show that the model can be used to interpret the time path for the unemployment rate in the United States over the course of the Great Recession.

We will maintain – for the sake of simplifying the exposition – the assumptions $b = G = T(s^t) = 0$.

4.1 Deterministic Dynamics

Imposing $p(s^t) = p$ and the equilibrium market-clearing condition $h(s^t) = 1$, the perfect-foresight version of the intertemporal FOC for the risky household type (eq. 12) is

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

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

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

These two equations can be used to plot the implied dynamics for the unemployment rate. To do so, we use the same parameter values as before ($\phi = 0.0375$) and set $p = \underline{p} = 0.75$.
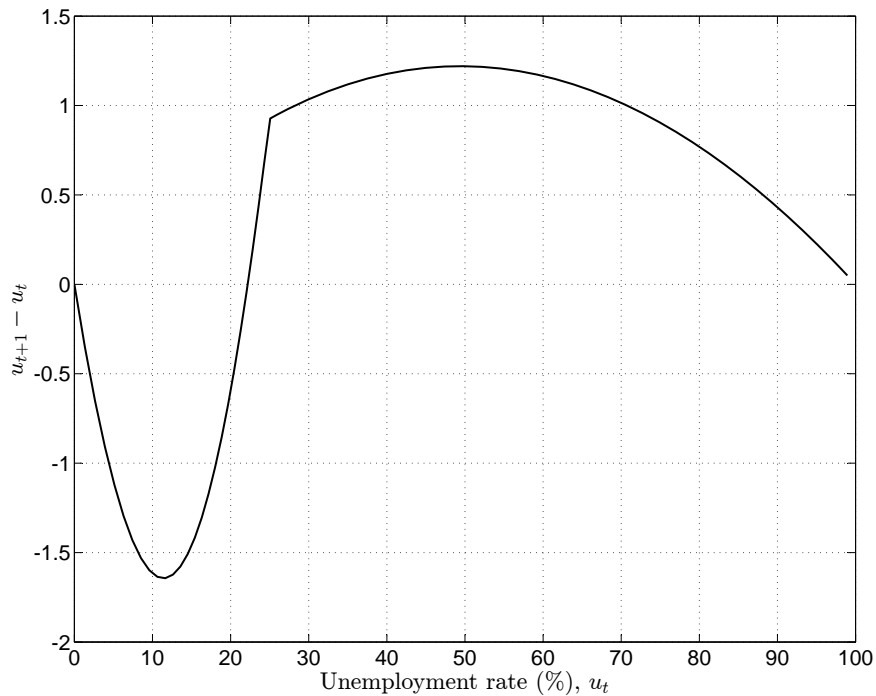
Figure 6: Deterministic dynamics

Figure 6 plots the change in the unemployment rate $u_{t+1} - u_t$ against the unemployment rate u_t. The two points at which the change in the unemployment rate is zero correspond to the two steady state unemployment rates at $p = 0.75$. Denote these rates u_L and u_H. The figure indicates that for any initial unemployment rate below u_H, unemployment will gradually converge, over time, to u_L. Thus, the low unemployment steady state is locally dynamically stable: if unemployment starts out below u_L, unemployment will rise, whereas if it starts above u_L (but below u_H), unemployment will fall. The fact that this steady state is dynamically stable will later allow us to introduce sunspot shocks that generate fluctuations in the neighborhood of u_L.

The high unemployment steady state is not stable. If unemployment starts above u_H, it will increase toward maximum unemployment, in expected terms. Note that any such paths are not

20

equilibria, because in the limit they imply that households will end up with zero income and consumption, which cannot be optimal given positive wealth equal to p.

4.2 The Great Recession

We now show that our model can generate dynamics for house prices and unemployment that are qualitatively similar to those experienced by the United States over the course of the Great Recession. Panel A of Figure 7 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 price.[5] 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 7: House prices and unemployment

A. Data   B. Model
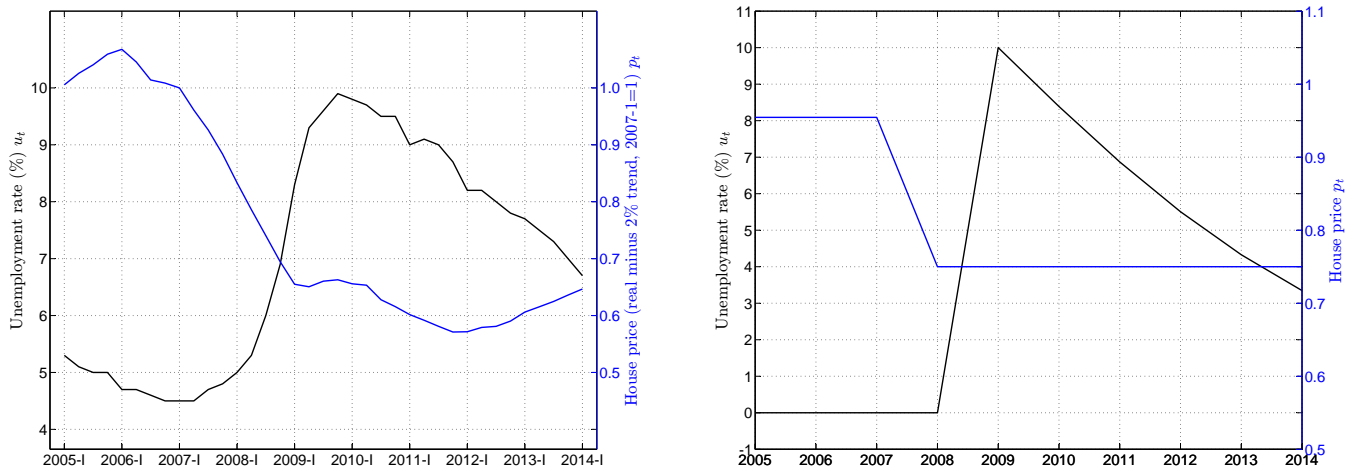
The sequence of model events that generates the times series plotted in Panel B of Figure 7 is as follows. Initially, the fundamental demand for housing is strong, so that there is a unique full

---

[5]This 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).

21

employment steady state. This corresponds to the period known as the Great Moderation during which U.S. house (and stock) prices were high by historical standards. Then, between 2007 and 2008, there is a permanent unanticipated decline in the taste for housing, which reduces ϕ from $\phi = \tilde{\phi}$ to $\phi = 0.0375$ (our baseline value). Agents initially believe that this will simply lead to a permanent change in house prices to the implied new value for $\underline{p}$, but that the unemployment rate will remain at zero. Given these beliefs, there is no immediate change in household demand – notwithstanding the decline in household wealth – and thus no immediate change in the unemployment rate.

With ϕ below $\tilde{\phi}$, however, the model now has multiple steady states, and the zero unemployment steady state is locally stable. Thus, the economy is now vulnerable to a confidence-driven recession. We assume that the economy is hit by an unanticipated shock to the expected path for the unemployment rate from 2009 onward 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, the economy is hit by no further fundamental shocks to preferences or to expectations, so households enjoy perfect foresight over the evolution of the unemployment rate. The economy converges toward the low unemployment steady state according to the dynamics described in Figure 6.

Although the model we have developed is very simple, it can replicate some key features of the Great Recession: (i) a decline in asset values that precedes the decline in real economic activity, (ii) a rapid contraction, and (iii) a slow recovery. In the next section, we will argue that there are good economic reasons for why confidence-driven recessions will typically be persistent in nature. A key part of the intuition is that a rapid expected recovery would imply a strong intertemporal motive to borrow and spend in the near term and therefore make it difficult to engineer a demand-driven recession in the first place.

Note that the model simulation just described relies on two zero probability shocks: one to preferences and one to expectations. We now move to construct sunspot equilibria in which agents take as given a positive probability of switching between boom and recession states.

22

4.3 Two-State Sunspot Equilibria

We now construct equilibria with sunspots. We start with perhaps the simplest possible equilibrium of this type: a two-state Markov sunspot equilibrium, in which asset prices are constant and in which 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. Let λ denote the probability that $s_{t+1} = L$ (H) given that $s_t = L$ (H). This is now a three-parameter model, where the parameters are β, ϕ, and λ.

The unemployment rates in the two states are $u(L) = 0$ and $u(H) > 0$. Assuming that unemployed workers are constrained, $c^u(H) = p$. From the resource constraint, consumption of workers is then given by

$$
\begin{aligned}
c^w(L) &= 1 \\
c^w(H) &= 1 - \frac{u(H)p}{[1 - u(H)]}.
\end{aligned}
$$

The intertemporal first-order conditions in the zero and positive unemployment rate states are (again assuming that unemployed workers are constrained)

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

Existence of Equilibria with Sunspots: We first ask, when does a sunspot equilibrium of the type described exist? An equilibrium of the type described exists if and only if there is a solution $\{u(H), c^w(H), c^w(L), p\}$ to the previous four equations that satisfies (i) $u(H) \in (0,1]$, (ii) $p \geq \underline{p}$, and (iii) $c^w(H) > p$. We now provide a partition of the parameter space into a region in which a solution with these properties exists and a region in which there is no such solution.

Proposition 3: A two-state sunspot equilibrium of the type described exists if and only if $\lambda \geq \underline{\lambda}$ where

$$
\underline{\lambda} = \frac{1}{2}\frac{(2+\rho)\underline{p}^2 - \underline{p} + 1}{\underline{p}^2 - \underline{p} + 1}
$$

and where $\underline{p} = \frac{\beta\phi}{1-\beta}$ is the floor on house prices established by the riskless household type.

Proof: See Appendix A.

From this proposition, it is immediate that the larger is ϕ, the larger is $\underline{\lambda}$ and thus the more persistent changes in the unemployment rate must be in any sunspot-driven model of fluctuations. Consider some special cases.

- When $\phi = 0$, so that housing has no fundamental value, $\underline{\lambda} = 0.5$, and aggregate fluctuations can be generated by a sunspot process that is iid over time.

- For $\phi > 0$, $\underline{\lambda} > \frac{1}{2}$, and thus confidence-driven fluctuations must be persistent. At the baseline parameter values $\underline{\lambda} = 0.8635$, so the expected duration of the unemployment state must be at least $1/(1 - 0.8635) = 7.3$ years.

- As $\phi \to \tilde{\phi}$, $\underline{\lambda} \to 1$, and thus regimes of zero or positive unemployment must be expected to be near permanent.

- For $\phi \geq \tilde{\phi}$, there are no sunspot equilibria of this type. Thus, confidence-driven fluctuations are only possible if the taste for housing, and hence the fundamental component of housing value, is sufficiently low.

To summarize, sunspot-driven fluctuations can only arise when two conditions are satisfied: (i) asset values must be sufficiently low (precisely, ϕ must be sufficiently low), and (ii) fluctuations must be sufficiently persistent in expected terms (λ must be sufficiently high).

The logic for the first result is straightforward. For households to be willing to pay the same price for housing in the positive unemployment / low consumption state as in the zero unemployment / high consumption 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.

The logic for why confidence-driven recessions must be persistent is as follows. The reason households reduce spending when the sunspot shock flips the economy into the positive unemployment state is that they anticipate a high likelihood that the unemployment rate will be high in the next period, and thus they have a strong precautionary motive to save today. Iterating forward, expecting that the unemployment rate will be high in the next period (and thus that consumption

24

will be low) only makes sense if the unemployment rate is likely to be high two periods into the future.[6]

Persistence and Volatility: We now turn to a second question. For values for $\phi < \tilde{\phi}$ (and thus $\underline{\lambda} < 1$) and for $\lambda \in [\underline{\lambda}, 1]$, what is the relationship between λ on the one hand and the unemployment rate in the recession state $u(H)$ on the other?

Figure 8 shows $u(H)$ for a set of alternative model parameterizations in which we vary the persistence parameter λ. In each case, β and ϕ are fixed at their baseline values. Corresponding to each different λ, there is a λ-specific constant house price p.

Figure 8: Larger recessions are more persistent
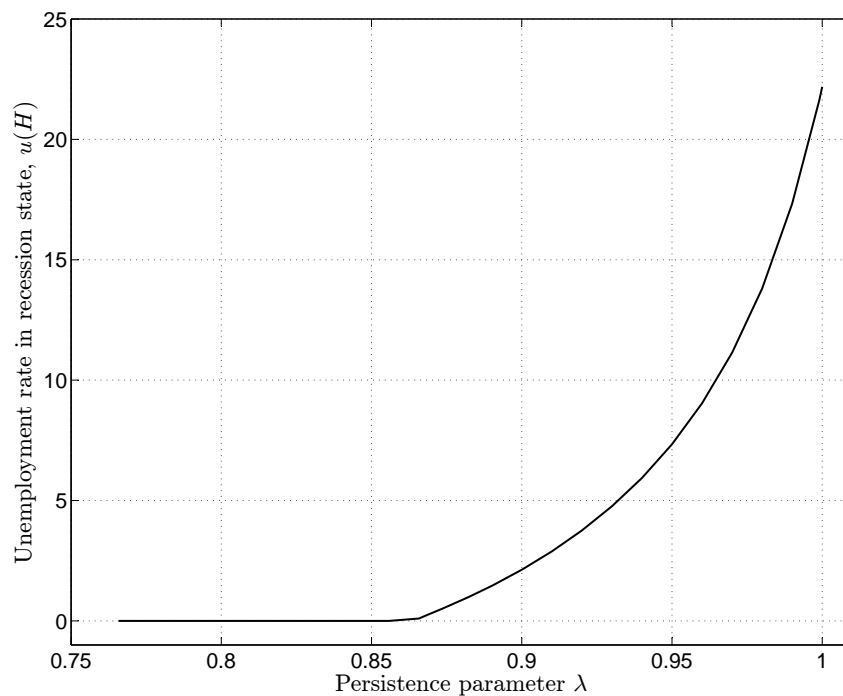
The key takeaway from the figure is that, in this class of equilibria, more persistent fluctuations and larger fluctuations go hand in hand. As $\lambda \to \underline{\lambda}$, the unemployment rate in the recession state $u(H)$ converges to zero, and fluctuations vanish. As we increase λ, unemployment in the recession

---

[6]Put 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 a strong intertemporal motive to use wealth to support consumption in the high unemployment state. But then this consumption would translate into additional demand and employment, and the conjectured equilibrium would unravel.

state rises. In the limit $\lambda \to 1$, the expected duration of either state becomes infinite, and the unemployment rate in the recession state converges to highest possible value that can arise in steady state, namely, u^+.[7]

Why are more persistent fluctuations also larger in magnitude? At a basic level, standard permanent income logic suggests that if the recession state is expected to last a long time, households will reduce spending sharply when the economy enters the recession state. More precisely, the strength of the household's intertemporal savings motive is closely connected to expected income growth, which in turn depends (negatively) on both the persistence of the sunspot process and (positively) on the level difference between output in the normal and recession states. The more persistent are shocks, the larger must be this level difference in order to maintain similar expected income growth in the recession state and to thereby make these fluctuations consistent with optimal intertemporal consumption choices.

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).

Recall that p is an endogenous equilibrium object in the model. But there exist sunspot equilibria with different constant values for p, just as there exist steady states with different constant values for p (see Figure 4). In fact, one can roughly think of the exercise here as constructing equilibria in which the economy bounces between the two steady states.[8]

In Figure 9 we fix $\lambda = 0.99$ and plot $u(H)$ (in red) and $u(L)$ (in blue) against p.[9] 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. 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

---

[7]In both the limit $\lambda \to \underline{\lambda}$, and the limit $\lambda \to 1$, the equilibrium house price p converges to the lower bound $\underline{p}$. The steady state price is a hump-shaped function of λ for intermediate values for λ.

[8]As $\lambda \to 1$, this analogy is exact.

[9]The dashed lines show the same objects for a lower $\lambda = 0.98$.

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.

One way to understand these theoretical predictions is that in order to support a price considerably above the zero unemployment fundamental ($\underline{p} = 0.75$), 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 $u(H)$.

Figure 9: Wealth and volatility
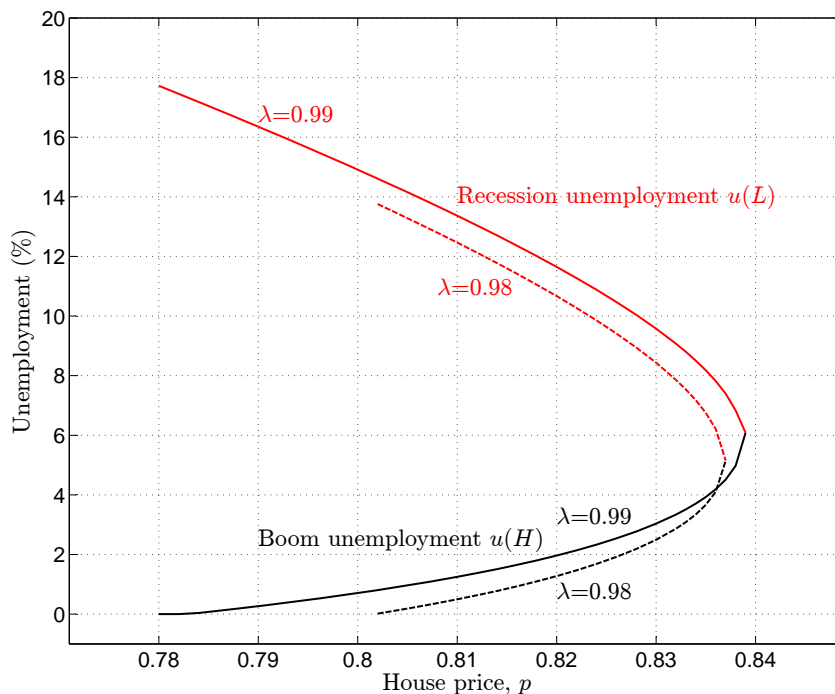
Note that in this class of sunspot equilibria, for a fixed λ, the equilibrium with the highest constant value for asset prices $\tilde{p}$ implies a constant positive unemployment rate. This $(\tilde{p}, u)$ combination corresponds to one of the steady states illustrated in Figure 4. For each different value for λ, the corresponding highest price equilibrium corresponds to a different steady state $(\tilde{p}, u)$ combination. Recall that for $\phi \geq \tilde{\phi}$, the highest price in any steady state corresponds to zero unemployment and is equal to $\underline{p}$. Combining these observations, it follows that $\phi \geq \tilde{\phi}$ rules out the existence of any sunspot equilibria of this type; in any such equilibrium, we would have $p \leq \tilde{p}$ (by definition) and $\tilde{p} < \underline{p}$ (by virtue of Proposition 2), thereby contradicting the requirement $p \geq \underline{p}$.

5 Review: Asset Prices and Volatility

Recall that the starting point for this paper is the strong positive empirical correlation between the level of U.S. household wealth and U.S. macroeconomic volatility. We now collect the various theoretical results we have established that relate to the relationship between asset prices and output volatility.

First, if the taste for housing is sufficiently strong ($\phi \geq \tilde{\phi}$), the model has a unique steady state with full employment. Furthermore, there do not exist sunspot equilibria in which the unemployment rate fluctuates over time.

Second, if the taste for housing is weak ($\phi < \tilde{\phi}$) so that the fundamental component of home value is small, then the economy has multiple steady states. In particular, there is a range of values for $p \geq \underline{p}$ such that for any p in that range, the model has two steady state unemployment rates. The low unemployment steady state is locally stable, introducing the possibility of sunspot-driven fluctuations, in which changes in expectations unrelated to fundamentals translate into macroeconomic volatility.

Third, when a low fundamental component to home value makes possible fluctuations driven by self-fulfilling changes in expectations, the theory imposes restrictions on the nature of those fluctuations. In particular, fluctuations must be expected to be persistent in order for changes in the expected unemployment rate to be self-fulfilling. The larger is the fundamental component of house value – and thus the less sensitive is demand to expectations – the more persistent must be changes in the equilibrium unemployment rate. Thus, the theory raises an interesting contrast between fundamentals-driven versus confidence-driven fluctuations: the persistence of a fundamentals-driven recession will be tightly linked to the persistence of the underlying fundamental shock, whereas a confidence-driven recession will necessarily be persistent.

Fourth, when sunspot driven fluctuations are possible, the model can be used to investigate the conditions under which equilibrium fluctuations are large in magnitude. We found, first, that fluctuations tend to be larger in magnitude the more persistent is the underlying sunspot process. This implication of the theory 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. Our second result is that, holding fixed structural parameters (including the persistence of the sunspot process), fluctuations are larger the lower is the level of asset prices.

Taking these observations in combination indicates that the empirical evidence we presented earlier on the link between the level of wealth and the volatility of macro aggregates 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 larger the lower is the level of asset prices.

6   Microeconomic Evidence

As discussed above, when wealth is low, demand (and hence output) in our model can fall in response to an increase in perceived unemployment risk. The mechanism is that a perceived increase in unemployment risk rationalizes a large precautionary fall in desired consumption, thereby making the expectation of higher unemployment self-fulfilling.

At the aggregate level, this mechanism is hard to test, because consumption and unemployment are mechanically inversely related: higher unemployment means lower output which necessarily corresponds to lower aggregate consumption. Thus it is hard to say whether changes in demand are driving output, or whether changes in output are driving demand.

At the microeconomic level, in contrast, we can isolate precautionary demand effects by comparing the consumption responses of households who face the same increase in aggregate unemployment risk, but who differ with respect to initial wealth. In particular, we can explore whether low wealth households – for whom the precautionary motive should be especially large – reduce their consumption rate disproportionately sharply as unemployment risk rises during the Great Recession.[10] The fact that there is enormous cross-sectional heterogeneity in wealth helps identify how the strength of the precautionary motive increases as wealth falls, which is a key element of our narrative.

In this section we use micro 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 did in fact exhibit systematically larger declines in their consumption rates. We believe we are the first to provide such evidence on the importance of the precautionary motive using micro data for U.S. households. This evidence as 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

---

[10]See Appendix B for a discussion of this effect in the model.

support to demand-driven theories of the Great Recession.

6.1 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. The PSID and the CES are the only U.S. data sets with household-level data on consumption, income and wealth. In each data set we rank households by net worth and compute changes over time in the consumption rates of households in the the top and bottom halves of the 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.

6.2 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 10 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 healthcare. Net worth includes net financial wealth plus housing wealth net of all mortgages (including home equity loans).[11] 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.[12]

---

[11]Appendix B reports more details of how we measure each variable. We do not use any age restrictions when constructing the PSID and CES series in Figure 10.

[12]One 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 difficult to date consumption precisely in the PSID because some of the survey questions ask explicitly about spending

Figure 10: Comparing aggregates across micro data sets
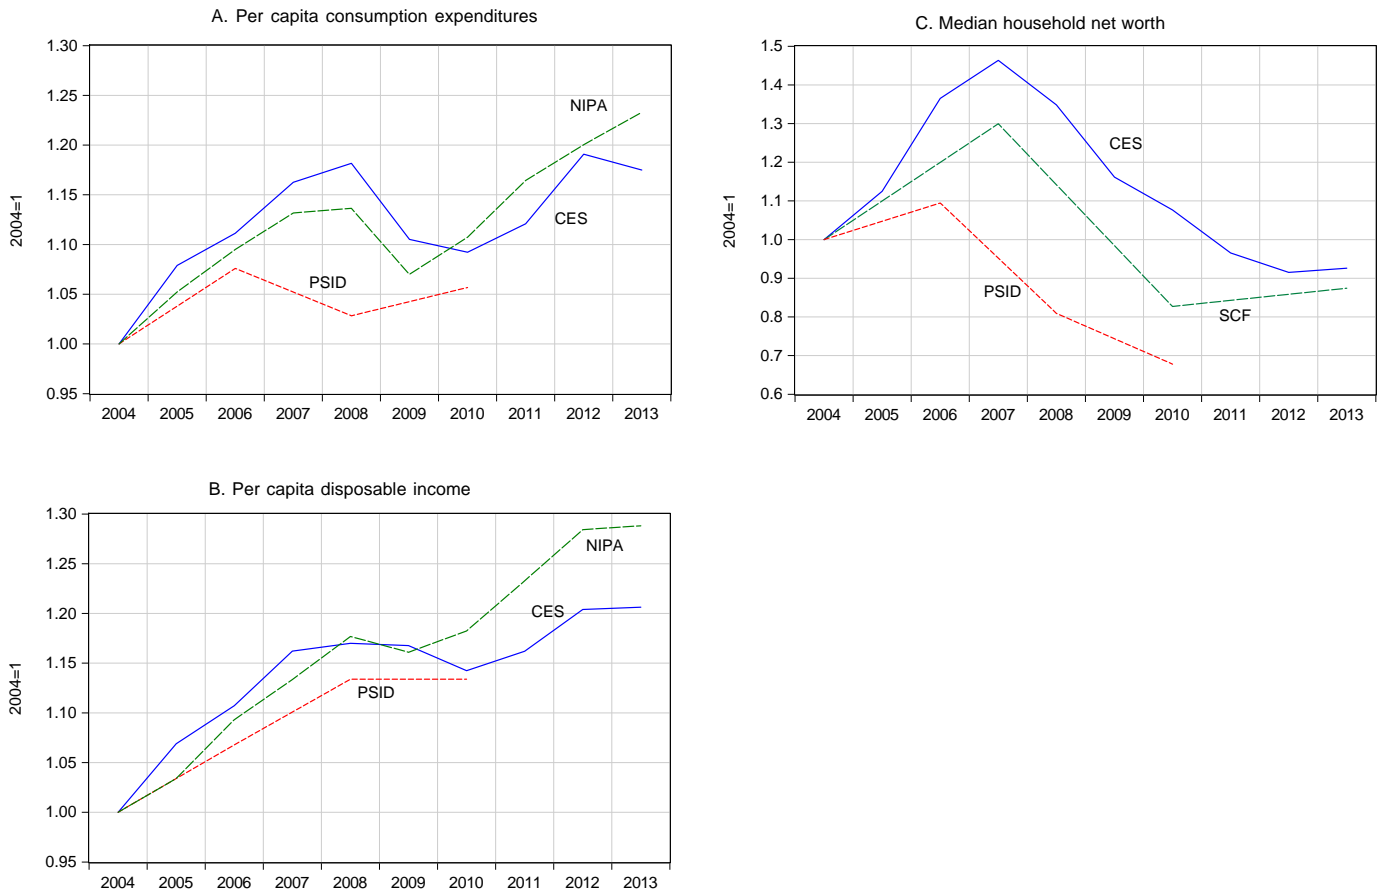
6.3 Measurement

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.

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.

31

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 $t + 4$.

6.4   Descriptive Statistics

Table 1 reports characteristics of the rich and poor groups in both the PSID and the CES, for the year 2006. Boot-strapped standard errors are in parentheses. 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 PSID and around 187,000 in CES.[13] This dramatic difference suggests that the precautionary saving motive for the poor group should be much stronger than for the rich group. 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.

---

[13]One reason for the difference between the two dataset is that the measure of net worth in PSID includes the value more assets such us 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	47.1	40.2	46.4
	(0.21)	(0.21)	(0.25)	(0.24)
Heads with college (%)	21.3	36.5	24.8	39.4
	(0.86)	(1.1)	(1.1)	(1.2)
Mean household size	2.45	2.72	2.84	2.79
	(0.04)	(0.03)	(0.04)	(0.04)
Mean household net worth (current $)	11,931	619,831	11,967	338,535
	(879)	(49,388)	(1,155)	(12,644)
Median household net worth	5,000	265,000	1,800	187,102
	(476)	(6,602)	(294)	(4,893)
Per capita disposable income	15,028	28,475	18,739	30,184
	(256)	(667)	(334)	(593)
Per capita consumption expenditure	9,831	13,101	9,185	10,858
	(177)	(250)	(232)	(188)
Consumption rate (%)	65.8	46.0	49.0	36.0
	(0.90)	(0.86)	(1.18)	(0.66)

Note: Bootstrapped standard errors are in parentheses.

6.5 Changes in Consumption Rates: Rich versus Poor Households

Figure 11 contains the key finding of this section. The figure plots changes in consumption rates in both the PSID (Panel A) and the CES (Panel B). Around the onset of the recession both data sets reveal a decline in the consumption rate of the poor that is significantly larger than the corresponding decline for the rich.[14]

Before concluding that the large fall in the consumption rate of the poor (relative to the rich) is due to the poor having a stronger precautionary motive, we consider two alternative possible explanations. The first is that the poor cut their consumption more because they suffered larger wealth losses. The second is that the poor cut their consumption more because their income prospects deteriorated more than those of the rich.

To evaluate the first alternative explanation we exploit the fact that households in the PSID

---

[14]Consumption rate declines in the CES (Panel B) 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.

Figure 11: Changes in consumption rates for rich and poor

A. PSID

B. CES
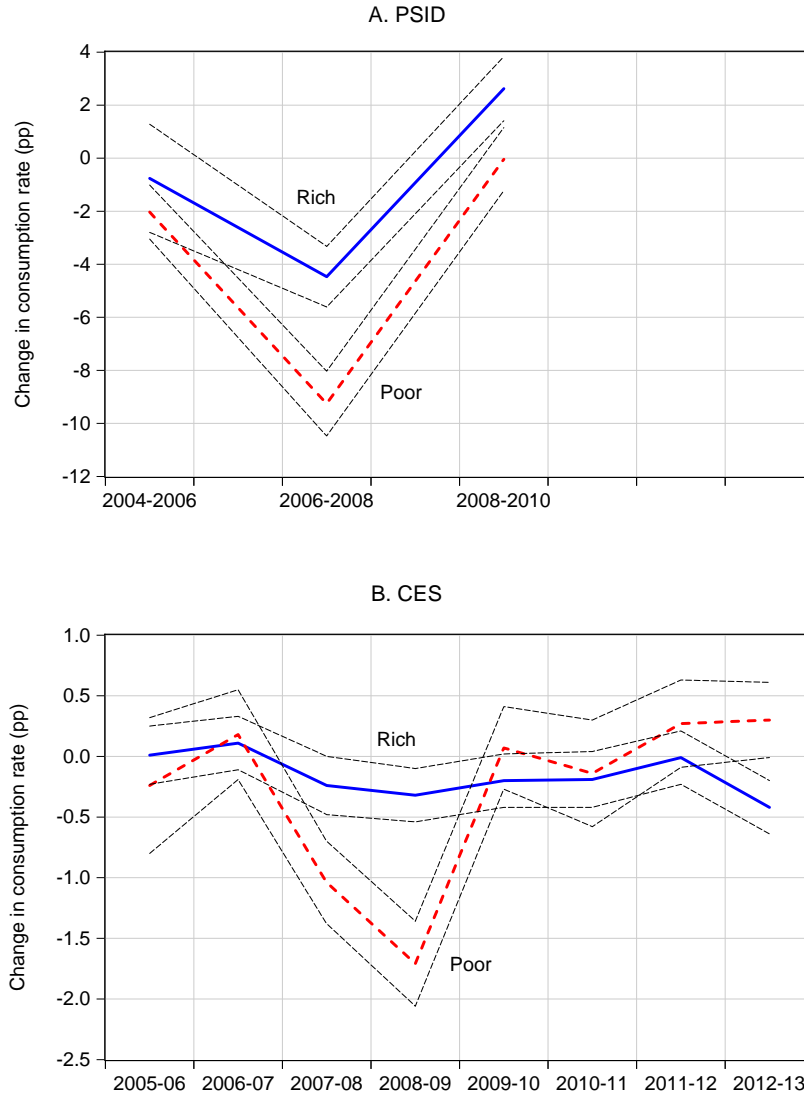
report wealth at each interview, which allows us to compute changes in net worth for both the rich and poor groups. To evaluate the second alternative we take advantage of the long panel dimension of the PSID, which allows us to compute future income growth for the two groups, which we take as a measure of income prospects. Table 2 reports the changes in consumption rates plotted in Figure 11 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), and the growth rates of income over the two year period following the measured consumption rate changes (lines 3 and 6).[15]

_____
[15]Thus, future income growth for 2004-2006 refers to income growth of the group between 2006 and 2008, and

TABLE 2. CONSUMPTION, WEALTH & FUTURE INCOME, PSID

		2004-2006	2006-2008	2008-2010
POOR				
1.	Δ consumption rate (pp)	-2.03 (1.2)	-9.25 (1.2)	-0.04 (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)	-0.76 (1.8)	-4.47 (1.1)	2.62 (1.1)
5.	Δ net worth (% of income)	189 (71)	-137 (50)	-22 (23)
6.	Δ future income (%)	7.39 (1.8)	-0.2 (1.9)	

Note: Bootstrapped standard errors are in parentheses.

Line 5 of the table shows that wealth rich households experienced very large declines in net worth between 2006 and 2010, while changes in net worth for poor households over the same period (line 2) were positive. A conventional wealth effects story would therefore predict larger declines in consumption rates for the rich, which is the opposite of the pattern observed. Lines 3 and 6 indicate that as the economy entered the recession, expected future income growth declined 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 was similar across wealth groups, suggesting that differential income prospects were not the primary factor behind the especially sharp decline in the consumption rate of the poor.

Thus, the evidence in Table 2 points against two competitors to our hypothesis that the differential consumption rate changes shown in Figure 11 reflect a strengthening of the precautionary motive to save on the part of wealth poor households, in the face of rising unemployment risk. 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

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.

consumption by $1/3 \times 4\% \simeq 1.3\%$ of aggregate disposable income.

7   Policy

In the model we outlined in Section 2, agents derive no disutility from working and enjoy strictly increasing utility from consumption. Thus, equilibria with positive unemployment are inefficient. The economy can get stuck in those equilibria because agents cannot get together and agree to maintain consumption demand at the full employment level of output. Government policy can potentially address this coordination failure.

Here, we consider two policies.

The first is for the government to directly purchase goods and to finance the purchases by taxing workers. The government budget constraint is

$$[1 - u(s^t)]T(s^t) = G,$$

where G denotes nonvalued government purchases and $T(s^t)$ is the tax per worker.

The second is for the government to tax workers and transfer the revenues to unemployed workers in the form of a benefit b. The government budget constraint is then

$$[1 - u(s^t)]T(s^t) = u(s^t)b.$$

7.1   Tax and Spend

Given a credible government commitment to a fixed amount of purchases G, only the private portion of aggregate demand is now potentially sensitive to expectations. One might therefore guess that this policy should make aggregate demand less sensitive to the expected unemployment rate and therefore help to rule out equilibria with positive unemployment.[16] However, the problem with this policy is that the tax on workers required to finance expenditure reduces private consumption and therefore reduces the fundamental component of housing values. This reduces risk sharing within the household and makes desired saving more sensitive to perceived unemployment risk. The net effect is that the set of steady state unemployment rates ends up being completely insensitive to

---

[16]Corsetti et al. (2012) present empirical evidence suggesting that the government spending multiplier is relatively large during periods of financial crisis.

G.

Proposition 4: The set of steady state unemployment rates is equal to $[0, u^+]$ for any $G \in [0, 1 - u^+]$.

Proof: See Appendix A.

Figure 12 illustrates this result, showing the set of steady states for two different values for government purchases: $G = 0$ and $G = 0.1$. With the higher value for G, the floor on house prices established by the riskless household is reduced, since workers in the riskless household now pay taxes. Moreover, the floor is now declining in the steady state unemployment rate u, since per worker taxes $T = G/(1 - u)$ are increasing in u. The plot offers a concrete visual verification of Proposition 4: the maximum steady state unemployment rate is u^+ for both values for G. We conclude that the tax and spend policy is not an effective way to rule out an expectations-driven recession because the beneficial spending component of the policy is likely to be undermined by the negative impact of the associated taxation on asset prices.[17]

Figure 12: Tax and spend policy

---
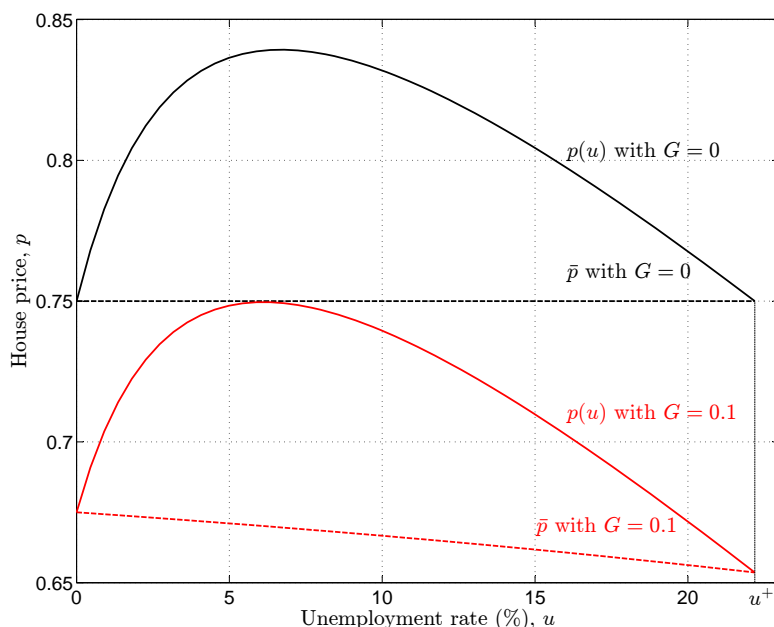
[17]The government could reduce the set of steady state unemployment rates by committing to extremely high levels of government purchases. Since private consumption cannot be negative, $u \leq 1 - G$. Thus, setting $G = 1$ would guarantee zero unemployment. Note, however, that such a policy would imply zero private consumption.

7.2 Unemployment Benefits

Now consider a policy of taxing workers to provide unemployment benefits. This policy is effective because it directly addresses the lack of within-household risk sharing that makes demand very sensitive to perceived unemployment risk. Once unemployment benefits are introduced, the gap between the consumption of workers and that of the unemployed is reduced, and the precautionary motive to save in the face of perceived unemployment risk therefore declines. As a result, the liquidity component to house prices is smaller, and high unemployment rates are no longer consistent with house prices that exceed the floor $\underline{p}$. Figure 13 shows the set of steady states for $b = 0$ and $b = 0.1$.

Figure 13: Unemployment insurance policy
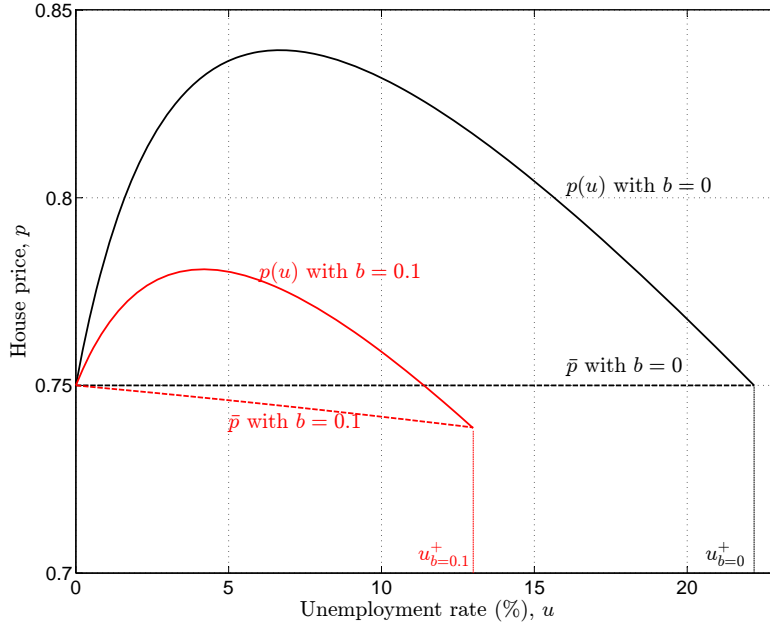
In fact, if

$$b \geq \underline{b} = \frac{1}{2\beta - 2}\left(\beta - \sqrt{4\beta - 3\beta^2} + 2\beta\phi\right),$$

then there is a unique steady state, with full employment.[18] Given our baseline parameter values, the critical threshold for benefits is $\underline{b} = 0.204$. Note that a policy of setting $b \geq \underline{b}$ is costless in equilibrium, since the mere anticipation of unemployment benefits is enough to prevent agents from

---

[18]This expression can be derived following a similar approach to the derivation of the threshold $\tilde{\phi}$ for the existence of multiple steady states absent unemployment benefits (see Proposition 2).

coordinating on an equilibrium with positive unemployment, and with full employment no taxes need be collected. In that sense, the policy is analogous to bank deposit insurance: the presence of sufficiently generous deposit insurance (unemployment insurance) rules out the equilibrium in which agents run on the bank (increase precautionary saving), and on the equilibrium path no tax revenue need be collected. Thus, in the context of this simple model, unemployment benefits are an attractive policy instrument. We recognize, however, that this simple model does not capture the potential costs of an unemployment insurance scheme in terms of disincentives to supply labor.[19]

8   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 important direction for future research.

---

[19]In addition, a more complete analysis of unemployment insurance would ideally microfound 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.

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] Barro, R.J., and H.I. Grossman (1971) "A General Disequilibrium Model of Income and Employment," *American Economic Review* 61(1), 82-93.

[7] Beaudry, P., D. Galizia, and F. Portier (2014) "Reconciling Hayek's and Keynes' Views of Recessions," NBER Working Paper 20101.

[8] Benito, A. (2006) "Does Job Insecurity Affect Household Consumption?" *Oxford Economic Papers* 58(1), 157-181.

[9] Carroll, C.D. (1992) "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence" *Brookings Papers on Economic Activity* 23(2), 61-156.

[10] Carroll, C.D., J. Slacalek, and M. Sommer (2012) "Dissecting Saving Dynamics: Measuring Wealth, Precautionary, and Credit Effects," IMF Working Paper 12/219.

[11] Challe, E., and X. Ragot (2012) "Precautionary Saving Over the Business Cycle," Working Paper, Paris School of Economics.

[12] Chamley, C. (2014) "When Demand Creates Its Own Supply: Saving Traps," *Review of Economic Studies* 81(2), 651-680.

[13] Chodorow-Reich, G., and L. Karabarbounis (2014) "The Cyclicality of the Opportunity Cost of Employment," Working paper, Chicago Booth.

[14] Cooper, R., and A. John (1988) "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics* 103(3), 441-463.

[15] Corsetti, G., A. Meier, and G. Muller (2012), "What Determines Government Spending Multipliers?," Working Paper, Cambridge.

[16] 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.

[17] Diamond, P.A. (1982) "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90(5), 881-894.

[18] Engen, E.M., and J. Gruber (2001) "Unemployment Insurance and Precautionary Saving," *Journal of Monetary Economics* 47(3), 545-579.

[19] Farmer, R.E.A. (2013) "Animal Spirits, Financial Crises and Persistent Unemployment," *Economic Journal* 123(568), 317-340.

[20] Farmer, R.E.A. (2014) "The Evolution of Endogenous Business Cycles," *Macroeconomic Dynamics,* 1-14.

[21] Grebler L., D.M. Blank, and L. Winnick (1956) *Capital Formation in Residential Real Estate: Trends and Prospects*, Princeton, N.J.: Princeton University Press.

[22] Guerrieri, V., and G. Lorenzoni (2009) "Liquidity and Trading Dynamics," *Econometrica* 77(6), 1751-1790.

[23] Hall, R.E. (2005) "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review* 95(1), 50-65.

[24] Hintermaier, T., and W. Koeniger (2013) "Collateral Constraints and Self-Fulfilling Macroeconomic Fluctuations," Working Paper, University of Bonn.

[25] Huggett, M. (1993) "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control* 17(5-6), 953-969.

[26] Kaplan, G., and G. Menzio (2014) "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," Working Paper, University of Pennsylvania.

41

[27] Kaplan, G., G.L. Violante, and J. Weidner (2014) "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* (Spring).

[28] Kocherlakota, N. (2009) "Bursting Bubbles: Consequences and Cures," Mimeo, FRB Minneapolis.

[29] 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.

[30] Martin, A., and J. Ventura (2014), "Managing Credit Bubbles," Working Paper, CREI.

[31] Mian, A., K. Rao, and A. Sufi (2013), "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 128(4), 1687-1726.

[32] Michaillat, P. (2012) "Do Matching Frictions Explain Unemployment? Not in Bad Times," *American Economic Review* 102(4), 1721-1750.

[33] Midrigan, V., and T. Philippon (2011), "Household Leverage and the Recession," Working Paper NYU.

[34] Mody, A., F. Ohnsorge, and D. Sandri (2012), "Precautionary Savings in the Great Recession," IMF Working Paper 12/42.

[35] Shimer, R. (2012), "Wage Rigidities and Jobless Recoveries," *Journal of Monetary Economics* 59, S65-S77.

Appendix A. Proofs

Proof of Proposition 1

Suppose, contrary to the claim, that $u > 0$ and $c^w = c^u$. Then the price that solves the intertemporal FOC would be

$$p = p_F(u) = \frac{\beta\phi}{1-\beta}(1-u) < \underline{p}, \tag{18}$$

where $p_F(u)$ is the fundamental steady state price given u. But $p < \underline{p}$ contradicts $p \geq \underline{p}$, which must hold in any steady state.

Proof of Proposition 2

From the pricing equation for the riskless type, $p \geq \underline{p} = \frac{\beta\phi}{1-\beta}$. Thus, there can be no steady states with $p < \underline{p}$.

From the previous proposition, if there is an equilibrium with $u > 0$, then unemployed agents are constrained. Thus, $c^u = p$ and $c^w = 1 - \frac{u}{1-u}p$.

The first-order condition for the risky household is therefore

$$\frac{p}{1-\frac{u}{1-u}p} = \beta\frac{p}{1-\frac{u}{1-u}p}\left[1 + \frac{u\left(1-\frac{u}{1-u}p-p\right)}{p}\right] + \beta\phi,$$

which simplifies to deliver the following equilibrium relation between steady state u and p :

$$p(u) = \frac{\beta(u+\phi)}{(1-\beta) + \frac{\beta u(1+\phi)}{1-u}}. \tag{19}$$

Note that at $u = 0$, this relation implies $p = \underline{p}$, so that is always a steady state.

Using the expression above, we can explore how p varies with u :

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

Now it is immediate that at $u = 0$

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

Now turn to the second derivative:

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

It is immediate that $\beta > 0.5$ is a sufficient condition for the second derivative to be negative and thus for p to be a concave function of u.

Combining these two results, it follows that for $\beta > 0.5$ and $\phi \geq \tilde{\phi}$, the value for p that satisfies the risky type's first-order condition is decreasing in u. Thus, $p \leq \underline{p}$.

But then for $\phi \geq \tilde{\phi}$, we have $p \geq \underline{p}$ (from the FOC for the riskless type) and $p \leq \underline{p}$ (from the FOC for the risky type). It follows that $p = \underline{p}$ and $u = 0$ is the only steady state.

If $\phi < \tilde{\phi}$, then $\frac{\partial p}{\partial u} > 0$ at $u = 0$. Since p is a continuous and concave function of u, and since $p = 0$ at $u = 1$, there must be a second steady state at $p = \underline{p}$ with $u > 0$.

This unemployment rate is given by

$$u^+ = 1 - \frac{\beta}{1-\beta}\phi(1+\phi). \tag{20}$$

By similar reasoning, there is a range of values for $p > \underline{p}$ such that given $\phi < \tilde{\phi}$, there are two steady states with positive unemployment.

Note that the uniqueness result with $\phi \geq \tilde{\phi}$ hinges on the presence of the riskless household type. Without this type, there would be a continuum of steady states with unemployment rates between zero and one, with each unemployment rate corresponding to a different steady state asset price as given by eq. 18 (see Farmer, 2013). The presence of the riskless type puts a floor on the asset price, which effectively establishes a floor for steady state consumption demand and output.

Proof of Proposition 3

Let $u(H) = \varepsilon$. By assumption $u(L) = 0$. From the budget constraints, $c^w(L) = 1$ and $c^w(H) = \frac{(1-\varepsilon)-\varepsilon p}{1-\varepsilon}$.

Thus, the FOCs for the household that faces risk in the L and H states can be rewritten, respectively, as

$$1 = \beta(1-\lambda)\left(\frac{(1-\varepsilon)^2}{(1-\varepsilon)-\varepsilon p} + \varepsilon\frac{1}{p}\right) + \beta\lambda + \frac{\beta\phi}{p}$$

$$\frac{(1-\varepsilon)}{(1-\varepsilon)-\varepsilon p} = \beta\lambda\left(\frac{(1-\varepsilon)^2}{(1-\varepsilon)-\varepsilon p} + \varepsilon\frac{1}{p}\right) + \beta(1-\lambda) + \frac{\beta\phi}{p}.$$

A first-order approximation to these conditions at $\varepsilon = 0$ gives

$$\beta + \frac{\beta\phi}{p} - 1 + \left(\frac{1}{p}\beta(1-\lambda)\left(p^2 - p + 1\right)\right)\varepsilon = 0$$

$$\beta + \frac{\beta\phi}{p} - 1 + \left(\frac{1}{p}\left(\beta\lambda - p^2 - p\beta\lambda + p^2\beta\lambda\right)\right)\varepsilon = 0.$$

To satisfy both FOCs, the coefficients on ε must be the same in both equations, which imposes the following restriction on the equilibrium relationship between p and λ in this class of equilibria:

$$\lambda = \frac{1}{2}\frac{\left(2 + \frac{1-\beta}{\beta}\right)p^2 - p + 1}{p^2 - p + 1}.$$

The right-hand side of this expression is increasing in p, and thus the lowest value for λ that can be supported corresponds to the lowest possible price for p, which is $\underline{p}$.

Proof of Proposition 4

The steady state equilibrium condition that must be satisfied for the household facing unemployment risk is

$$\frac{p}{1 - \frac{u}{1-u}p - T(u)} = \beta p \left(\frac{(1-u)}{1 - \frac{u}{1-u}p - T(u)} + \frac{u}{p} \right) + \beta\phi,$$

where equilibrium consumption values for workers and unemployed members are

$$c^w = 1 - \frac{u}{1-u}p - T(u)$$
$$c^u = p$$

and where taxes paid per worker are $T(u) = G/(1-u)$.

Let $\underline{p}(u)$ denote the floor on house prices established by the riskless household type:

$$\underline{p}(u) = \frac{\beta\phi(1 - T(u))}{(1-\beta)}.$$

Recall that at $u = 0$, the house price is equal to $\underline{p}(0)$. Let $u^{\#}$ denote the maximum possible steady state unemployment rate. At $u = u^{\#}$ the price is again equal to the floor established by the riskless household type: $p = \underline{p}(u^{\#})$. This implies the following relationship at $u = u^{\#}$ between the per-worker tax required to finance G and the house price:

$$T(u^{\#}) = 1 - \frac{(1-\beta)\underline{p}(u^{\#})}{\beta\phi}.$$

Substituting this into the first-order condition for the household facing risk gives

$$\frac{\underline{p}(u^{\#})}{\frac{(1-\beta)\underline{p}(u^{\#})}{\beta\phi} - \frac{u}{1-u}\underline{p}(u^{\#})} = \beta\underline{p}(u^{\#}) \left(\frac{(1-u^{\#})}{\frac{(1-\beta)\underline{p}(u^{\#})}{\beta\phi} - \frac{u}{1-u}\underline{p}(u^{\#})} + \frac{u^{\#}}{\underline{p}(u^{\#})} \right) + \beta\phi,$$

which simplifies to the following quadratic equation:

$$-u^{\#} \left(\beta\phi^2 + \beta\phi + u^{\#} + \beta - u^{\#}\beta - 1 \right) = 0.$$

One solution to this equation is $u^{\#} = 0$ – the minimum possible steady state unemployment rate. The other solution is

$$u^{\#} = 1 - \frac{\beta}{1-\beta}\phi(1+\phi) = u^{+}.$$

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.

Net Worth In both 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 has a more accurate record of wealth, and reports also the values of individual retirement accounts (IRAs), of family businesses, and of vehicles. Our measure of net worth in the PSID 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 11) 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 10) is much smaller. We have also reproduced Figure 11 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 11, 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 $t + 1$.

Wealth and the precautionary motive in the model

In the model we have laid out, all households are identical. To develop a theoretical link between household wealth on the one hand and the consumption response to a change in unemployment risk

on the other, we considered the following thought experiment. Take the baseline calibration of the model, and consider the two steady states at $p = \underline{p}$, one of which has $u = 0$ and the other of which has $u = u^+$. Now consider dropping a single atomistic household with arbitrary initial wealth h_0 into each of the two steady states. Figure 14 plots, for different values for h_0, the optimal values for initial consumption in each steady state. In the steady state with zero unemployment (the top line), the optimal choice is $c_0^w = 1$ irrespective of the initial value for h_0 (this is an implication of the quasi-linear preferences we assumed). In the steady state with positive unemployment, the optimal initial consumption of both workers and unemployed agents is increasing in h_0 : consumption of borrowing-constrained unemployed workers is mechanically increasing in h_0, whereas consumption of employed workers is increasing in h_0 because more initial wealth reduces the precautionary motive to save (see eq. 4). Thus, if we imagine the economy transiting unexpectedly from the full employment steady state to the positive unemployment steady state, the decline in household consumption will be larger the smaller is initial household wealth.

Figure 14: Wealth and the precautionary motive

48