

Wealth and Volatility

Jonathan Heathcote
Minneapolis Fed and CEPR

Fabrizio Perri
University of Minnesota, Minneapolis Fed, NBER and CEPR

PRELIMINARY AND INCOMPLETE

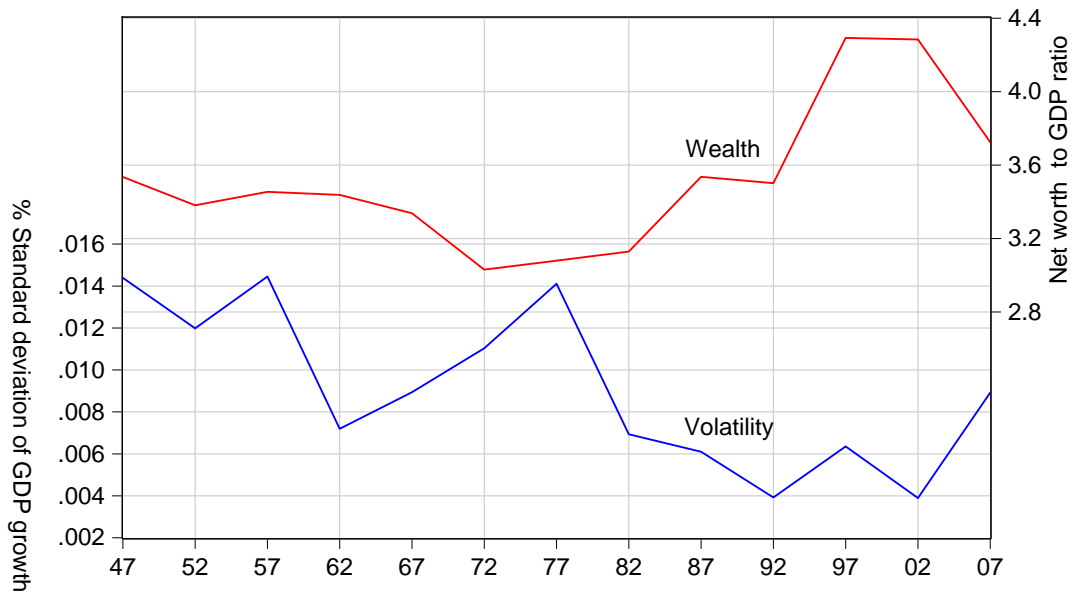
April 2012

1 Introduction

Figure 1 documents the relationship between household wealth and aggregate volatility in the United States. Volatility is measured as the standard deviation of the quarterly real GDP growth rate over a 5 years period (where the x axis in the graph denotes the start of that period), while wealth is measured as net worth of households and non profit organization over GDP in the middle of the same 5 years period. The figure reveals that periods when net worth is high, reflecting high prices for housing and/or stocks, tend to be periods of low volatility in aggregate output (and hence employment and consumption). Conversely, periods in which asset values are low tend to be periods of high macroeconomic volatility. For example during the periods of late 70s and early 1980s wealth is at its historical low and volatility is at its historical high, while during the late 1990s and early 2000s wealth is at its highest and volatility is at its lowest. Periods in which wealth is at intermediate levels, such as the 1962-1966 and 2007-2011 are also periods in which volatility is at intermediate levels.

In this paper we argue that (exogenously given) fluctuations in wealth are an important factor in determining the volatility of business cycles. We make the argument by developing a micro-founded dynamic equilibrium model that contains elements of a traditional Keynesian framework in 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. When wealth is high the economy has a unique equilibrium and behaves neoclassically. When wealth is low, the economy becomes vulnerable to additional confidence-driven fluctuations and hence is, in general, more volatile. In this case, there is a potential role for public policies to stabilize demand.

Our model economy is populated by a large number of identical households, each of which comprises many members. Competitive firms operate a linear technology that transforms employed workers into a perishable consumption good. In addition to non-durable consumption, households



Note: Standard deviations of GDP growth are computed over 20 quarters periods. Observations for net worth to GDP ratio refer to the middle of those same periods. The x axis denote the start of those periods

Figure 1: Wealth and volatility

enjoy utility from durable housing, which is in fixed supply. Households commit to consumption orders for their members. Then jobs are allocated by random matching in a decentralized labor market, in which firms hire sufficient workers to fulfill their orders. The decentralized nature of the labor market is important because it means that equilibrium output will be determined by desired consumption demand, rather than by desired labor supply, and because it allows for the possibility of equilibrium unemployment.

Agents must finance their consumption orders using wage income, home equity, or non-collateralized borrowing. Unemployment expectations matters not only because the expected unemployment rate determines expected income, but also because if house prices are low the household anticipates that members who do not find jobs will have to use expensive credit. The fact that consumption for each member is committed prior to realization of the member's idiosyncratic employment status increases the precautionary motive associated with perceived unemployment risk.

Our first result is to show that this environment allows for multiple equilibria in which households can collectively co-ordinate on a range of expectations about unemployment, each of which turns out to be self-fulfilling. In particular, there is a range of values for house prices in which the model has two steady states (with the same house price in each). In the optimistic steady state, households expect low unemployment, are therefore not too concerned about credit costs for unemployed workers, and set consumption demand high. Facing high demand, firms employ a large fraction of workers, and the expectation of low unemployment is rationalized. In the pessimistic steady state, households expect high unemployment. Because they do not want to commit to high consumption given high idiosyncratic unemployment risk and costly credit, they set consumption demand low. Facing low demand, firms hire few workers, and unemployment is in fact high, as expected.

Precautionary saving in housing offers a way for households to self-insure against unemployment risk. The less wealth a household has, the more reliant the household will be on costly credit in the event of unemployment. Thus the lower is household wealth, the more sensitive is consumption demand to the expected unemployment rate. This increased sensitivity of demand to expectations increases the range of unemployment rates that can be supported in a rational expectations equilibrium. To see this consider the extreme case in which households have no wealth. Then, if they are maximally pessimistic and expect 100% unemployment they will set consumption to zero, and 100% unemployment can occur in equilibrium. But if households have positive wealth they will choose positive consumption even if they expect 100% unemployment. Firms must then hire a positive fraction of workers to fill these orders, and so 100% unemployment cannot be an equilibrium.

Thus the maximum magnitude of sun-spot driven equilibrium fluctuations will vary inversely with the level of house prices.

Home values in the model are endogenous equilibrium objects, and reflect both the fundamental flow utility from home ownership and the liquidity value of being able to finance consumption out of home equity in the event of unemployment. Because this liquidity value is tied to the level of unemployment, house prices themselves are indeterminate, and like the unemployment rate, can potentially fluctuate in response to changes in expectations. However, for most of our analysis we will explore what dynamics for unemployment are possible for alternative constant values for equilibrium house prices. We do this because in the data a large share of house price volatility is at lower frequencies than typical business cycle fluctuations.

We further limit the set of values for house prices we consider by introducing as an additional model element a fringe group of households that does not face unemployment risk. The presence of this group establishes a lower bound on housing demand and thus on equilibrium house prices. If this lower bound is sufficiently high, unemployed workers can finance consumption entirely out of home equity, and full employment is the only possible equilibrium.

We use the model to offer an interpretation of recent macroeconomic history. The Great Moderation was a time in which US house and stock prices were very high by historical standards. High household wealth levels in this period meant that the economy was robust in the sense that it was not subject to large recessions induced by declines in confidence. However, the sharp declines in house and stock prices between mid 2007 and mid 2009 left the economy fragile, and vulnerable to a confidence-driven recession. One possible “trigger” for a collective loss of confidence was the collapse of Lehman Brothers in the Fall of 2008. We show that following a transitory sunspot-induced recession, equilibrium dynamics imply a slow recovery, even though fundamentals remain unchanged. Of course, fluctuations in consumer confidence are only one source of business cycles, and over a longer history economic cycles in the United States likely have a number of causes above and beyond fluctuations in animal spirits.

The model has policy implications. We evaluate two specific policies. The first is a lump-sum unemployment benefit, financed by a proportional income tax. This policy makes unemployment less painful, and thereby reduces the sensitivity of demand to the expected unemployment rate. A sufficiently generous benefit rules out sunspot-driven fluctuations and ensures full employment. The second policy we consider is government consumption financed by lump-sum taxation, in the spirit of the 2009 stimulus plan. This policy also makes aggregate (private plus public) demand

less sensitive to expectations, and thereby rules out equilibria with very high unemployment rates. However, taxation also reduces asset values, which increases the sensitivity of demand to perceived unemployment risk.

1.1 Related Literature

From a theoretical viewpoint, our paper is related to recent work by Guerrieri and Lorenzoni (2009), Chamley (2010), and Farmer (2010). In both our model and Farmer’s output is demand-determined, and asset prices play a critical role in determining demand. One difference is that agents in our model face idiosyncratic risk and thus have a precautionary motive for saving, which is central to delivering a connection between the level of asset prices and the volatility of output. Guerrieri and Lorenzoni develop a model in which risk-averse agents trade in a decentralized fashion and face idiosyncratic risk. Their model recessions feature an increase in precautionary saving, as do ours, but in their model an endogenous increase in precaution amplifies a fundamental aggregate productivity shock, while in ours it is a self-fulfilling prophesy.

A challenge in constructing models in which demand-side factors play an important role in that many forces that tend to reduce desired consumption demand (e.g. lower asset values, greater idiosyncratic risk) also tend to increase desired household labor supply. Hall (2005), Michaillat (2010) and Shimer (2012) have developed models with decentralized labor markets in which they assume that the real wage does not fall (much) in response to a negative productivity shock, leading to a large fall in vacancy posting and a surge in involuntary unemployment. These models lack clear microfoundations for wage formation. On the other hand, they offer a natural resolution to the longstanding discrepancy between small microeconomic estimates for labor supply elasticities, and the large macro elasticities implicit in the large movements in aggregate hours over the business cycle (see, for example, Chetty et al., 2011). The resolution is simply that large falls in aggregate hours during recessions do not reflect an unwillingness to supply labor, but instead indicate that wages are “stuck” at too-high levels. In our model, the household is assumed to first choose a reservation wage, and then the firm only gets to decide whether or not to accept or reject a potential match. Workers then always extract all the surplus from a firm-worker match and, given a linear technology, firms are always indifferent about how many workers to hire. The level of demand then effectively selects a particular employment rate.

We are not the first to argue for a link between asset values and volatility, but our mechanism reverses the usual direction of causation. Others (see Lettau et al. 2008) have pointed 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.

2 Model

There are two goods in the economy: a perishable consumption good, produced by a continuum of identical competitive firms using labor, and an 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, while income for the second is not.

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 riskless type of household is measure zero, but its presence will establish a floor for asset prices. 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 individuals, and the consumption order must be specified before this allocation is realized. Thus, the optimal strategy is to send each member out with the same order $c(s^t)$ and an equal fraction $h(s^{t-1})$ of the assets the household carries in the period. 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 clear their consumption orders. The fraction $u(s^t)$ who are unemployed pay for as much of their consumption order $c(s^t)$ as they can given assets on hand, and if necessary borrow to pay the rest. Borrowing incurs intermediation costs which, for simplicity, we model as a tax that depends on the amount borrowed. These taxes are rebated to households within the period as lump-sum transfers denoted $T(s^t)$. At the end of the period the household regroups and pools resources, which determines the quantity of the asset carried into the next period $h(s^t)$.

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 (exploiting the fact that each household member enjoys the same consumption level) are given by

$$\sum_{t=0}^{\infty} \beta^t \sum_{s^t} \pi(s^t) u(c(s^t), h(s^{t-1}), \phi(s^t)).$$

where β is the discount factor, $\pi(s^t)$ is the probability of history s^t as of date 0, and $\phi(s^t)$ is a potentially stochastic preference weight.

The household budget constraint for a risky household has the form:

$$\begin{aligned} c(s^t) + p(s^t) [h(s^t) - h(s^{t-1})] &\leq [1 - u(s^t)] w(s^t) - \frac{\psi}{2\bar{z}} u(s^t) \min \{ [p(s^t)h(s^{t-1}) - d - c(s^t)], 0 \}^2 + T(s^t) \\ c(s^t), h(s^t) &\geq 0 \end{aligned}$$

The left hand side of the budget constraint captures consumption and the cost of net asset purchases. The first term on the right hand side is household earnings, while the second is the cost of borrowing for unemployed workers who use credit to pay for consumption. Note that this cost is quadratic and only applies to the fraction $u(s^t)$ of workers who are unemployed, and only if the household sets the consumption order above the value of home equity. Home equity is the value of housing owned by the household, $p(s^t)h(s^{t-1})$, minus an amount d that we think of as mortgage debt. The size of credit costs is determined by the parameter ψ relative to average productivity \bar{z} .

Note that $h(s^{t-1})$ was effectively chosen in the previous period. In the current period, given aggregate variables $u(s^t)$, $w(s^t)$, $p(s^t)$ and $T(s^t)$, the choice for $c(s^t)$ implicitly defines the quantity of wealth carried into the next period $h(s^t)$.

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

2.1 Household's problem

Consider the problem for the type that faces unemployment risk. Let $\mu(s^t)$ be the multiplier on the budget constraint at history s^t . Assuming the non-negativity constraints on $c(s^t)$ and $h(s^t)$ are non-binding, the first order conditions with respect to $c(s^t)$ and $h(s^t)$ are:

$$\beta^t \pi(s^t) u_c(s^t) = \mu(s^t) \left(1 + \frac{\psi}{\bar{z}} u(s^t) b(s^t) \right) \quad (1)$$

$$\mu(s^t)p(s^t) = \sum_{s^{t+1}} \pi(s^{t+1}|s^t) \left(\beta^{t+1} u_h(s^{t+1}) + \mu(s^{t+1}) \left(p(s^{t+1}) + \frac{\psi}{\bar{z}} u(s^{t+1}) p(s^{t+1}) b(s^{t+1}) \right) \right) \quad (2)$$

where $b(s^t)$ denotes borrowing of the unemployed:

$$b(s^t) = \max \{ [c(s^t) - (p(s^t)h(s^{t-1}) - d)], 0 \}.$$

Substituting (1) into (2)

$$\frac{u_c(s^t)p(s^t)}{1 + \frac{\psi}{\bar{z}} u(s^t)b(s^t)} = \beta E_{|s^t} [u_h(s^{t+1}) + u_c(s^{t+1})p(s^{t+1})] \quad (3)$$

This looks like a standard inter-temporal first order condition for a consumption-savings problem, except the denominator of the left hand side indicates an additional motivation for saving when unemployment necessitates borrowing: saving one additional unit of the asset is really cheaper than the price $p(s^t)$ because reducing current consumption reduces expected borrowing costs. Credit costs are isomorphic to a tax on current (but not future) consumption at rate $1 + \frac{\psi}{\bar{z}} u(s^t)b(s^t)$. Under this interpretation, choosing lower $c(s^t)$ and thus lower $b(s^t)$ reduces the effective tax rate.

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

$$u_{\hat{c}}(s^t)p(s^t) \leq \beta E_{|s^t} [u_{\hat{h}}(s^{t+1}) + u_{\hat{c}}(s^{t+1})p(s^{t+1})]$$

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

Each representative firm produces according the following linear technology:

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

where $n(s^t)$ is the mass of workers employed by the representative firm, and $z(s^t)$ is potentially stochastic productivity. In equilibrium $u(s^t) = 1 - n(s^t)$. We now describe how equilibrium employment is determined.

Households first observe the aggregate state s_t and then give workers instructions about what wages to accept, i.e. they specify a reservation wage $w^*(s^t)$. Firms and workers meet in a decentralized labor market. A unit mass of workers meets each firm, where these meetings occur in a

random sequence throughout the period. Each firm takes as given the wage $w^*(s^t)$, the price at which it can sell output (normalized to one), and decides whether or not to hire each successive worker it meets. When a firm hires a worker it produces and sells the resulting output immediately.

The optimal strategy for the firm in this environment is to employ a worker if and only if the worker's reservation wage $w^*(s^t)$ is less than or equal to the worker's marginal product $z(s^t)$. Understanding the firms' incentives, a representative household will optimally assign its members a reservation wage $w^*(s^t) = z(s^t)$. Recall that a lower reservation wage does not increase the probability that a given household member will find a job, while a higher reservation wage would guarantee non-employment.

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

We assume that $z(s^t)$ follows a first-order Markov process, with mean $\bar{z} = 1$.

2.3 Equilibrium

In some versions of the model, labor productivity z_t and the preference for housing ϕ_t will be sufficient statistics for all aggregate variables. In other versions, sunspot-driven fluctuations in expectations will generate non-fundamental driven movements in consumption and house prices. Thus we define the current aggregate state of the economy $s_t = (z_t, \phi_t, c_t, p_t)$. [REVISIT THIS] The distribution of housing wealth does not appear because we will focus on equilibria in which all housing is owned by the risky household type, and in which each risky household is representative. (In Section XX we will consider an example in which the distribution of wealth between types is allowed to vary). A symmetric equilibrium in this model is a process for s_t and associated decision rules and prices $n(s_t)$, $u(s_t)$, $w(s_t)$, $h(s_t)$, $T(s_t)$ that satisfy, for all t and for all s_t :

1.

$$w(s^t) = w^*(s^t) = z(s^t)$$

2.

$$n(s^t) = 1 - u(s^t)$$

3.

$$h(s^t) = 1$$

4.

$$c(s^t) = z(s^t)(1 - u(s^t))$$

5.

$$T(s^t) = \frac{\psi}{2\bar{z}} u(s^t) \max \{ [c(s^t) - (p(s^t) - d)], 0 \}^2$$

6.

$$\frac{u_c(s^t)p(s^t)}{1 + \frac{\psi}{\bar{z}} u(s^t) \max \{ [c(s^t) - (p(s^t) - d)], 0 \}} = \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t) [u_h(s^{t+1}) + u_c(s^{t+1})p(s^{t+1})]$$

7.

$$u_{\hat{c}}(s^t)p(s^t) \geq \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t) [u_{\hat{h}}(s^{t+1}) + u_{\hat{c}}(s^{t+1})p(s^{t+1})]$$

2.4 Discussion

In a symmetric equilibrium, each firm employs sufficient workers to satisfy demand: $n(s^t) = c(s^t)/z(s^t)$. Thus, in this environment, the consumption order $c(s^t)$ determines employment $n(s^t)$ and unemployment $u(s^t) = 1 - n(s^t)$. If orders fall short of potential output, i.e., if $c(s^t) < z(s^t)$, then labor supply will exceed labor demand, in the sense that all measure 1 of workers in each household are willing to work at any positive wage, while employment is determined by labor demand $n(s^t) = c(s^t)/z(s^t) < 1$.

In this environment, firms are on their labor demand curve. However, no single atomistic household has an incentive to choose a lower reservation wage, because a lower wage will not increase the probability of its members forming successful matches. Thus unemployment does not exert downward pressure on wages, breaking the standard Walrasian adjustment process that ultimately equates labor demand and labor supply in models with frictionless labor markets.

The model of the household we adopt is a simple way to introduce idiosyncratic risk and a precautionary motive, without having to keep track of the distribution of wealth. We assumed that the household must choose commit to a consumption choice for each household member before members' idiosyncratic employment status is revealed. Had we instead allowed the household to send members out with two contingent consumption orders, higher expected unemployment would still strengthen the precautionary motive to save, but the effect would be much weaker, because the household could reduce the cost of unemployment spells by instructing unemployed workers to reduce spending rather than increase costly borrowing. In reality, large categories of consumption

are either impossible or very costly to reduce quickly in response to a negative income shock. It is possible to eat at home rather than go out for dinner, but it is difficult to return an iPhone to the Apple Store, and it is difficult to reduce housing related expenses without down-sizing to a smaller home.

The economic logic for the quadratic credit cost is that in the event that an unemployed worker is forced to borrow, he will exhaust cheap sources of credit first before turning to more expensive sources - thus the marginal cost of credit should be increasing in the amount borrowed. For our purposes a quadratic cost function is particularly tractable, but any exponent larger than unity delivers qualitatively similar results.¹

2.5 Preferences

We will assume the utility function is of the following separable quasi-linear form

$$u(c, h, \phi) = \log c + \phi h$$

Given this utility function coupled with $\hat{h}(s^t) = 0$ for all s^t , the inter-temporal first order condition for the riskless type simplifies to

$$p(s^t) \geq \beta \sum_{s^{t+1}} \pi(s^{t+1}|s^t) \left[\frac{\phi(s^{t+1}) + z(s^{t+1})^{-1} p(s^{t+1})}{z(s^t)^{-1}} \right] \quad (4)$$

Thus the presence of this type puts a floor under house prices.²

For the rest of the paper we will assume $z(s^t) = \bar{z} = 1$ for all s^t . We will also assume $\phi(s^t) = \phi$ for all s^t . However, we will compare equilibria with alternative constant values for ϕ .

3 Steady States

In steady state, the expression for the asset price floor above simplifies to

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

Steady state consumption is given by

$$c = 1 - u.$$

¹A linear cost function would introduce a kink in the agent's budget set.

²Note that with $h(s^t) = 1$ the inter-temporal first order condition for the risky household type 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.

Thus a sufficient condition for unemployed agents to be able to finance consumption out of home equity at any unemployment rate in steady state is

$$\underline{p} - d \geq 1$$

which is satisfied if

$$\phi \geq \bar{\phi} = \frac{1 - \beta}{\beta}(1 + d)$$

3.1 Full employment steady state with high house prices

Claim: If $\phi \geq \bar{\phi}$, the only possible steady state is $p = \underline{p}$, $u = 0$.

Proof: The steady state version of the risky type's intertemporal first order condition is

$$\begin{aligned} \frac{p}{(1 - u)} &= \beta\phi + \beta\frac{p}{(1 - u)} \\ p &= \frac{\beta\phi}{(1 - \beta)}(1 - u) \end{aligned} \tag{5}$$

Given the price floor

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

the only possible steady state is $p = \underline{p}$, $u = 0$.

Note that this uniqueness result 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. 5 (see Farmer 2010). 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.

3.2 Multiple steady states with low house prices

Now consider the case $\phi < \bar{\phi}$, so that demand for housing by the riskless type is not strong enough to drive house prices to a level at which borrowing is never necessary.

Now steady states are solutions (c, u, p) to the following equations

$$\begin{aligned} \frac{pc^{-1}}{1 + \psi u \max\{[c - (p - d)], 0\}} &= \beta(\phi + pc^{-1}) \\ c &= 1 - u \end{aligned} \tag{6}$$

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

Claim: Any steady state with positive unemployment must feature costly credit: $u > 0 \implies c > p - d$.

Proof: Suppose, contrary to the claim, that $u > 0$ and $c \leq p - d$. Then the price that solves the inter-temporal FOC would be

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

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.

The logic for this result is that in any steady state with positive unemployment, households facing unemployment risk have lower expected income than households who do not face that risk. For the risky households to be nonetheless willing to pay more for housing, it must be that housing has additional value as a source of liquidity for the risky type. This in turn implies that in steady state unemployed agents must be using costly credit, so that the additional liquidity associated with home equity is priced.

Claim: There exists a steady state with positive unemployment if and only if

$$\phi < \bar{\phi} \text{ and } \psi \geq \bar{\psi} = \frac{(1-\beta)^2}{(1-\beta)(1+d) - \beta\phi}$$

Proof: [TO BE ADDED]

The logic for the threshold $\bar{\psi}$ is as follows. Suppose we start with $p = \underline{p}$ and $u = 0$, and consider how the steady state asset price 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 for housing. At $\psi = \bar{\psi}$ the two effects exactly offset, and a marginal increase in unemployment is consistent with the same steady state asset price. For $\psi > \bar{\psi}$, the liquidity effect dominates, and a marginal increase in unemployment necessitates an increase in the steady state asset price: i.e. $\frac{\partial p}{\partial u} > 0$ at $p = \underline{p}$ and $u = 0$. For higher unemployment rates, the fundamental effect must eventually dominate, and by continuity, there must be at least one additional equilibrium at $p = \underline{p}$ and $u > 0$.

3.2.1 Simple example with $\phi = 0$

For $\phi = 0$, the set of solutions to these equations is particularly simple. This is a model in which housing offers no utility, and no financial return, and is only held because it offers liquidity and reduces borrowing costs. From 6 steady state consumption demand is given by

$$c = \frac{\rho}{\psi u} + p - d$$

where $\rho = \frac{1-\beta}{\beta}$ is the rate of time preference. Thus steady state demand is increasing in home equity $p - d$, and decreasing in unemployment risk u . Note that positive wealth establishes a positive floor on consumption demand, even if the expected unemployment rate is 100%. As the unemployment risk u is reduced, a higher steady state borrowing level for unemployed workers $c - (p - d)$ is required for agents to maintain constant consumption over time.

In any steady state, the inter-temporal motive to borrow plus the precautionary motive to save add up to a constant desired level of asset holdings equal to p . Equivalently steady state demand must equal steady state supply:

$$\frac{\rho}{\psi u} + p - d = 1 - u$$

This is a quadratic equation with potentially two interior solutions:

$$u = \frac{1}{2} \left((1 - (p - d)) \pm \sqrt{(1 - (p - d))^2 - 4 \frac{\rho}{\psi}} \right)$$

This equation has two interior solutions as long as $\psi > \bar{\psi} = \frac{1-\beta}{1+d}$. At both steady state unemployment rates, the household chooses to set $c = 1 - u$. Note that the ratio of wealth to expected income differs across the two steady states. In both steady states, wealth is the same, and equal to p , while expected income differs. In the high unemployment steady state, one can interpret the higher steady state wealth to income ratio as reflecting a greater precautionary demand for saving in the face of greater unemployment risk.

3.2.2 Example with $\phi > 0$

For $\phi > 0$ solving for steady state unemployment rates in closed-form is slightly more complicated, because the equation defining steady states is now a cubic in u . However, it is straightforward to characterize the set of steady states numerically. We now plot the set of steady states for $\psi = 1$, $\beta = 0.9$, $\phi = 0.05$ and $d = 0.1$. Note that this parameter configuration satisfies $\phi < \bar{\phi} = 0.1$ and $\psi > \bar{\psi} = 0.15$, so there are steady states in which $u > 0$ and $p > \bar{p} = 0.45$. Figure 2 shows the

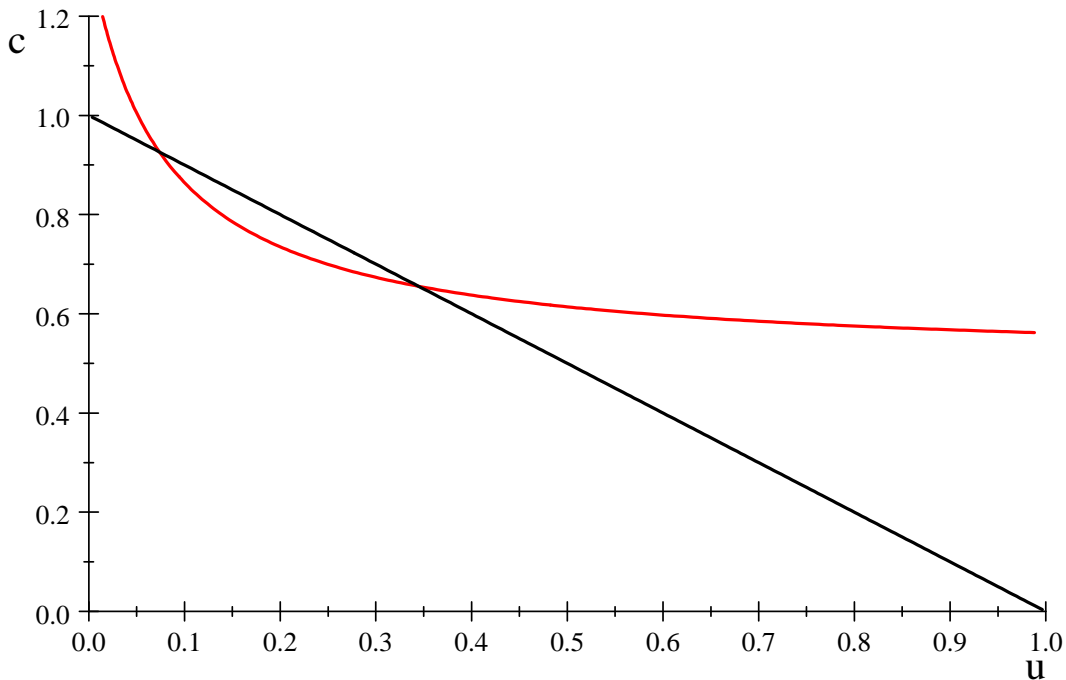


Figure 2: Steady states for $p = 0.6$.

equations defining steady state consumption demand (curved red line) and supply (straight black line) given $p = 0.6$.

It is clear that there are two steady states here, one with low, and one with higher unemployment. In the low unemployment equilibrium wealth is low relative to consumption, but the household does not want to increase saving because unemployment risk is low. In the high unemployment equilibrium, unemployment risk is high, but the household does not want to increase saving because wealth is already high relative to consumption.

Figure 2 shows the set of steady states for a particular price p while Figure 3 shows the set of steady states for all $p \geq \underline{p}$.

The green horizontal line here is $p = \bar{p}$, and the hump-shaped black line shows unemployment rates that satisfy the inter-temporal first order condition and market clearing at prices $p \geq 0$, i.e. solutions to

$$\frac{p(1-u)^{-1}}{1 + \psi u [(1-u) - (p-d)]} = \beta [\phi + p(1-u)^{-1}]. \quad (7)$$

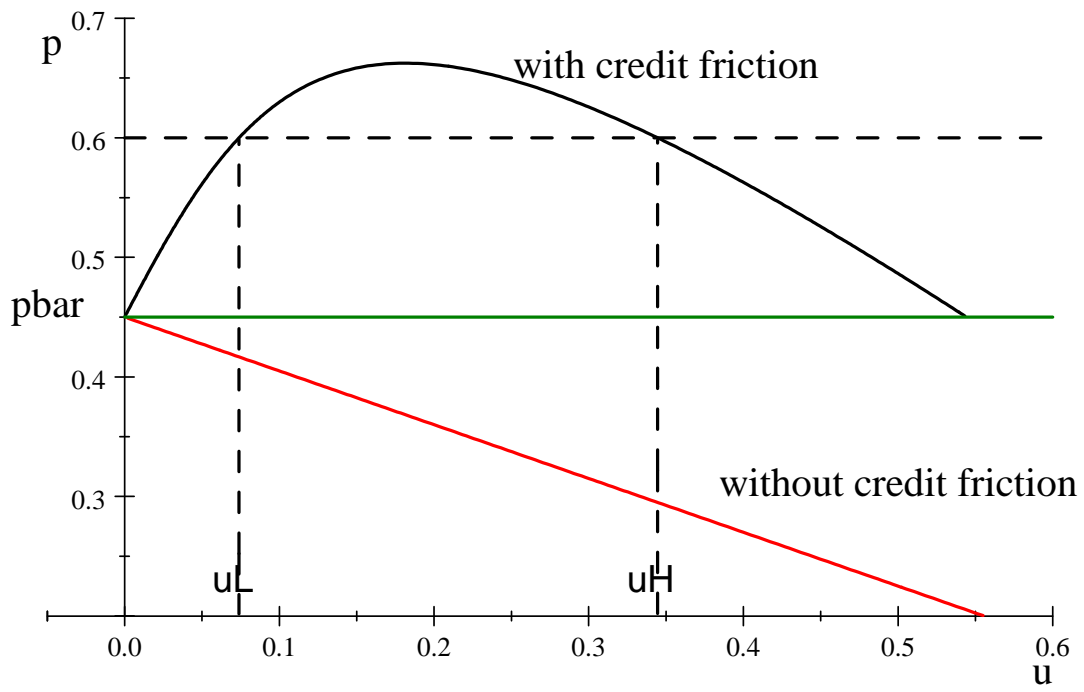


Figure 3: Steady states for $p = 0.6$.

It is clear that for all p such that a steady state exists, there are two steady state unemployment rates given this particular parameter configuration. The distance between the two steady state unemployment rates is decreasing in p .

The red line in the picture shows the price the household would be willing to pay absent the liquidity value for housing, i.e. the solutions to

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

Thus for a given steady state unemployment rate, the red line shows the “fundamental” asset value $p_F(u)$, while the vertical distance between the black and red lines measures the “liquidity” value of housing (which is zero when $u = 0$).

The figure also shows the two steady states for $p = 0.6$, the steady state price assumed for the previous figure. In the low unemployment equilibrium (u_L) the fundamental share of house value is high relative to the liquidity value, while in the high unemployment steady state (u_H) the opposite is true.

4 Dynamics

We now want to consider dynamics. Suppose that $p(s^t) = p > \underline{p}$. The inter-temporal FOC is (using the resource constraint to substitute out consumption)

$$\frac{(1-u(s^t))^{-1}}{1 + \psi u(s^t) \max \{[(1-u(s^t)) - (p-d)], 0\}} = \beta \frac{\phi}{p} + \beta E_{|s^t} \left[(1-u(s^{t+1}))^{-1} \right].$$

Let $x(s^t) = (1-u(s^t))^{-1}$, so $u(s^t) = 1 - x(s^t)^{-1}$. Then we can rewrite this as

$$E_{|s^t} [x(s^{t+1})] = \frac{x(s^t)}{\beta [1 + \psi (1 - x(s^t)^{-1}) \max \{[x(s^t)^{-1} - (p-d)], 0\}]} - \frac{\phi}{p}$$

or, more compactly, as

$$E_{|s^t} [x(s^{t+1})] = F(x(s^t)) \tag{8}$$

Suppose we now introduce a “sunspot”, v_t , where v_{t+1} (part of the date $t+1$ state, s_{t+1}) is symmetrically distributed with mean zero with a support $[-\varepsilon_{t+1}, +\varepsilon_{t+1}]$ defined by $\varepsilon_{t+1} = G(x(s^t))$, and where

$$x(s^{t+1}) = E_{|s^t} [x(s^{t+1})] + v_{t+1}(s_{t+1})$$

The assumption that v_{t+1} has mean zero is a pre-requisite for rational expectations. The assumption that the distribution for v_{t+1} is symmetric could be relaxed.

Given this specification for a sunspot shock, we can re-write the inter-temporal FOC as

$$x(s^{t+1}) = F(x(s^t)) + v_{t+1}(s_{t+1})$$

Thus the F function tells us the expected part of $x(s^{t+1})$ while the sunspot v_{t+1} tells us the impact of the impact of the sunspot shock.

Absent any fundamental shocks (shocks to productivity or preferences) we can define equilibrium recursively, with the current state being x_t , and the shock being v_{t+1} . (The current sunspot shock v_t is redundant as a state because x_t is a sufficient statistic for current unemployment, and v_t does not help forecast v_{t+1} given that the sunspot shock is iid over time.)

4.1 No sunspot shocks

Suppose first that $G(x(s^t)) = \varepsilon_{t+1} = 0$, so that $x(s^{t+1}) = E_{|s^t} [x(s^{t+1})]$.

It is straightforward to compute the dynamics for x_t given an initial x_0 , and thus to plot the corresponding dynamics for unemployment:

$$\Delta u = u' - u = 1 - \left(\frac{(1-u)^{-1}}{\beta(1+\psi u \max\{[(1-u) - (p-d)], 0\})} - \frac{\phi}{p} \right)^{-1} - u$$

Figure 4 plots Δu as a function of u , assuming $p = 0.6$.

Let u_L and u_H denote the low and high unemployment steady states. The plot indicates that the low unemployment steady state is locally dynamically stable: if unemployment starts out below u_L , unemployment will rise, while if it starts above u_L (but below u_H) unemployment will fall. Because this steady state is dynamically stable, we can introduce sunspot shocks that will 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 towards maximum unemployment, in expected terms (the kink in the path reflects the point at which the household stops using costly credit). Note that any such paths are not equilibria, because in the limit they imply that households will end up with zero income and consumption, which cannot be optimal given positive wealth.

Note that costly credit is being used at each point along any equilibrium path for unemployment corresponding to different initial unemployment rates $u_0 \in [0, u_H]$. The logic is as follows. First, we

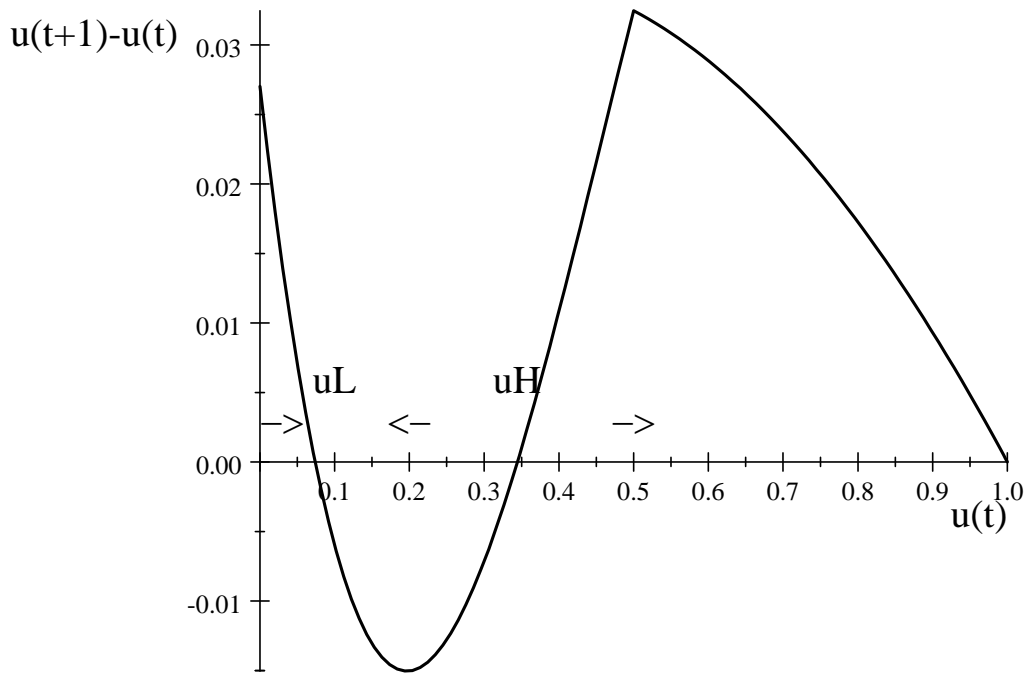


Figure 4: Unemployment Dynamics

have argued that the costly credit is used in all steady states with positive unemployment, and is thus being used at the high unemployment steady state u_H . Second, given that $(1 - u_H) > (p - d)$, the credit constraint must also be binding at all lower unemployment rates, since these correspond to higher consumption levels.

4.2 Introducing Sunspots

When we introduce sunspots, a range of paths become possible. For a particular current x_t ,

$$x_{t+1} \in [F(x_t) - G(x_t), F(x_t) + G(x_t)]$$

The bounds $G(x_t)$ have to be such that $x_{t+1} \leq (1 - u_H)^{-1}$ and $x_{t+1} \geq 0$, so that the unemployment rate remains between zero and u_H . The reason x_{t+1} can never exceed $(1 - u_H)^{-1}$ (equivalently u_{t+1} cannot exceed u_H) is that such realizations would make explosive paths possible (for example, $v_{t+\tau} = 0$ for all $\tau \geq 2$).

5 Review: Asset Prices and Volatility

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

1. First, if the taste for housing is strong ($\phi \geq \bar{\phi}$) or if credit is cheap ($\psi \leq \bar{\psi}$) then the only equilibrium is permanent full employment.
2. Second, if the taste for housing is weak ($\phi < \bar{\phi}$) so that home equity is inadequate to finance consumption of the unemployed and if, in addition, credit is expensive ($\psi > \bar{\psi}$) then the model has two steady states for the same house price p . The unemployment rate in the high unemployment steady state u_H is decreasing in p .
3. Third, given low house prices and expensive credit, and absent expectational errors, for any possible steady state price $p \geq \underline{p}$, the set of equilibria is as follows. First, for any initial unemployment rate $u_0 \in [0, u_H)$ there is an equilibrium path in which the unemployment rate steadily converges to u_L . Second, there is an equilibrium in which the unemployment rate is initially u_H and remains always at u_H .

4. Fourth, because the low unemployment steady state is dynamically stable, non-fundamental sunspot shocks can generate fluctuations in unemployment. Thus for any $p \geq \underline{p}$ and any $u_t \in [0, u_H)$ sunspot shocks can move the economy to any $u_{t+1} \in [0, u_H)$ (in expected terms unemployment evolves according to 8).
5. Finally, combining points 1, 2 and 4 and assuming costly credit, the range of equilibrium unemployment rates $[0, u_H)$ is decreasing in the steady state house price p . For high enough p (i.e. $\phi \geq \bar{\phi}$) the only possible unemployment rate is $u = 0$. For lower values for p (i.e. $\phi < \bar{\phi}$) unemployment can fluctuate between 0 and u_H . Moreover, since u_H is declining in p the range of equilibrium unemployment rates and thus the maximum size of sunspot shocks, is decreasing in p .

6 Microeconomic evidence

As discussed above, when wealth is low demand (and hence output) in our model can fall in response to a negative shock to expectations. Low wealth implies a strong precautionary motive and this leads to a large consumption fall in response to higher perceived unemployment risk, thereby making the expectation of higher unemployment self-fulfilling.

In this section we present some direct evidence for the first part of this mechanism using micro data.

Consider in particular two households, one with low wealth and the other with high wealth, which face the same increase in unemployment/income risk. Our theory suggests that the low wealth household should decrease its consumption more than the high wealth household. To test this prediction we use data from the Consumer Expenditure (CE) Survey, which is the only US data set that contains information on household-level wealth, expenditures and income. Households in the CE survey are interviewed for a maximum of four consecutive quarters. Households report consumption expenditures in all four interviews, report income information in the first and last interview, and report wealth information in the last interview only. We use CE data from the first quarter of 2005 to the first quarter of 2011 and we select households for which we have four consumption expenditure observations, two income observations and one wealth observation.

The goal is to compare changes in consumption expenditures during the course of the Great Recession for wealth rich versus poor households, controlling for potentially different income trends across the groups. It is important that we measure changes in expenditures using the same house-

holds, rather than comparing the average consumption of different sections of the wealth distribution at different points of time, as changes in moments of the cross sectional distribution might reflect changes in composition of the cross-sections and not actual changes in expenditures. To measure actual expenditure changes we exploit the limited panel dimension of the CE survey.

In particular, we first rank households in each survey quarter by net worth relative to average quarterly consumption expenditure, which we use as a proxy for permanent income. Net worth here includes net financial wealth plus housing wealth net of all mortgages (including home equity loans). We then construct two wealth groups in each month of the survey, using the median as the break-point. We compute aggregate nine-month consumption (income) growth for each group as the difference between aggregate 4th interview consumption (income) less aggregate 1st interview consumption (income). Note that we are using the exact same set of households to measure consumption and income growth. Finally we using a standard weighting scheme to convert these observations on monthly growth rates to a quarterly frequency, assuming constant growth between interview dates.

Before comparing changes in expenditures across groups, we compare aggregate consumption in our CE sample to a conceptually-similar definition of consumption from the National Income and Product Accounts (NIPA). In particular we take aggregate NIPA consumption expenditure minus the following categories: housing and utilities, health care, and financial services.³ Figure 5 compares the CE and NIPA series, both of which are real, per-capita measures. The timing and depth of the Great Recession are broadly similar across the two measures, although the extent of the recovery in expenditures is much weaker in the CE than in NIPA.

Table 1 reports some characteristics of the two groups. First the table indicates that with respect to demographics the wealth poor group appears to be composed by younger, less educated and poorer (in terms of income) households. The table also shows that the differences in average net worth across the different groups are large: the median per capita household wealth in the poorest group is close to 0 while the corresponding value for the rich group quartile exceeds 60000 2005 dollars.

Most important though, for our purposes, is the comparison of income and consumption growth for the groups. The table shows that although the wealth poor experience a slightly stronger income growth than the wealth rich, they experience a substantially stronger reduction in consumption expenditures (-5.6% v/s -3.1%). We view this as *prima-facie* evidence that the level of wealth is a

³We exclude the categories for which the matching of the CE with NIPA is particularly poor

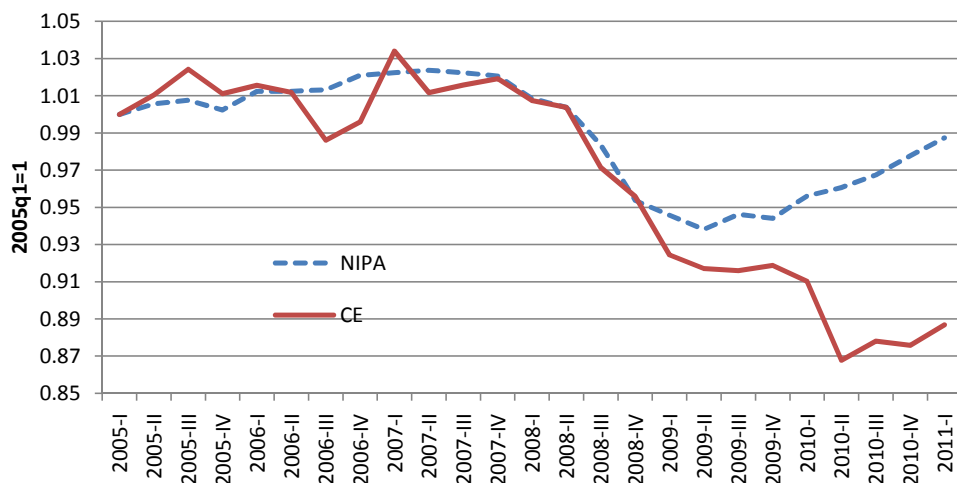


Figure 5: Aggregate real consumption expenditures per capita

major factor, independent from income, in determining demand response in a turbulent time.

Table 1. Characteristics of two wealth groups, 2005.1-2011.1

	Wealth Group	
	0-50	50-100
Sample size	8864	8873
Average age of head	41.4	46.9
Percent of heads with college	25.7	40.5
Average household size	2.9	2.8
Per capita real net wealth (2005\$)		
Mean	1498	119796
Median	238	63162
Average per capita after tax income (2005\$)		
Level	22117	32811
Growth rate	-0.3%	-1.0%
Average per capita consumption expenditures (2005\$)		
Level	9353	11252
Growth rate	-5.6%	-3.1%

We analyze the issue in more detail as figure 6 plots (annualized) growth rates of per capita real consumption for each of our four wealth groups from the 4th quarter of 2005 through the 1th quarter of 2011 while figure 7 plots the change in consumption rates (i.e. consumption expenditure over after tax total income) for the two groups. These 2 figures tell a clear story. The first suggests that the the wealth-poor group reduced consumption sooner and most periods more sharply than the wealth rich. The second suggests that the relative reduction in consumption expenditures of

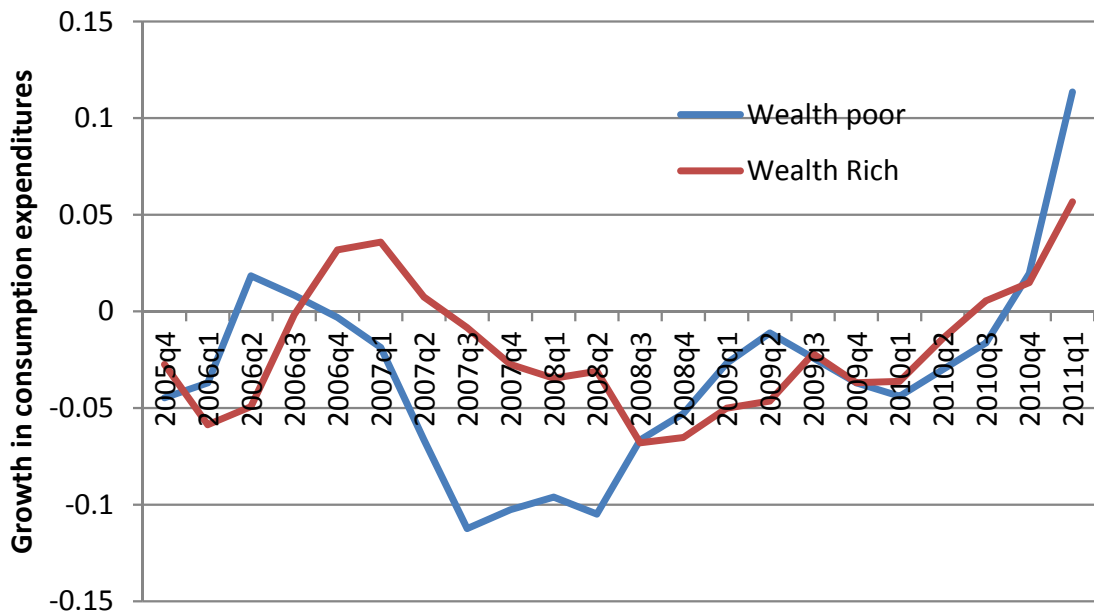


Figure 6: Growth in consumption expenditures for rich and poor

the wealth poor throughout the recession is not just driven by a poor relative income growth by the wealth poor but rather by an increasing saving (reduced consumption rate). Our theory suggests that this increase in relative saving rate of the wealth poor is the result of the strengthening of their precautionary motive to save, in response to higher unemployment risk.

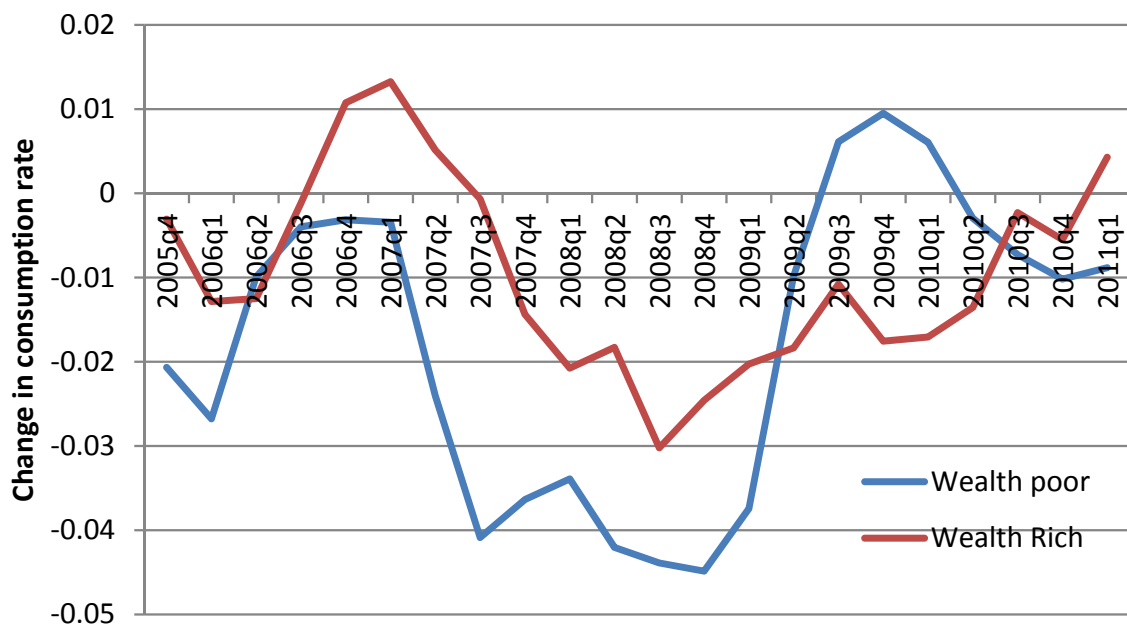


Figure 7: Changes in consumption rates for rich and poor

7 The Great Recession

The key characteristics of the Great Recession were a sharp fall in asset prices, accompanied by a sharp fall in spending, and a rise in unemployment. Since the recession officially ended in mid 2009, the subsequent recovery has been slow.

Figure 8 shows timepaths for (i) the price of net worth, (ii) durable consumption, and (iii) the unemployment rate over a five year period between the first quarter of 2007 and the last quarter of 2011.⁴ Asset prices fall 30% relative to trend in the first two years of the sample, while durables consumption falls around 20%. Note that the fall in asset prices precedes the fall in consumer spending, which in turn leads the fall in the unemployment rate. Figure 9 shows price series for the two key components of net worth: house prices and stock prices. This figure shows that throughout 2007 house prices were already falling steadily relative to trend. In contrast, the fall in stock prices

⁴A nominal net return series for net worth is constructed from the Flow of Funds as household and non-profit sector holding gains on assets stated at market value (Table R100, line 9) divided by previous quarter net worth (Table B100, line 43). A price index is created from this series by deflating by the GDP deflator. The series plotted for the real price of net worth and real per-capita durable consumption are deviations from a constant 2.1% trend, the average growth rate of GDP per capita in the 1947-2007 period.

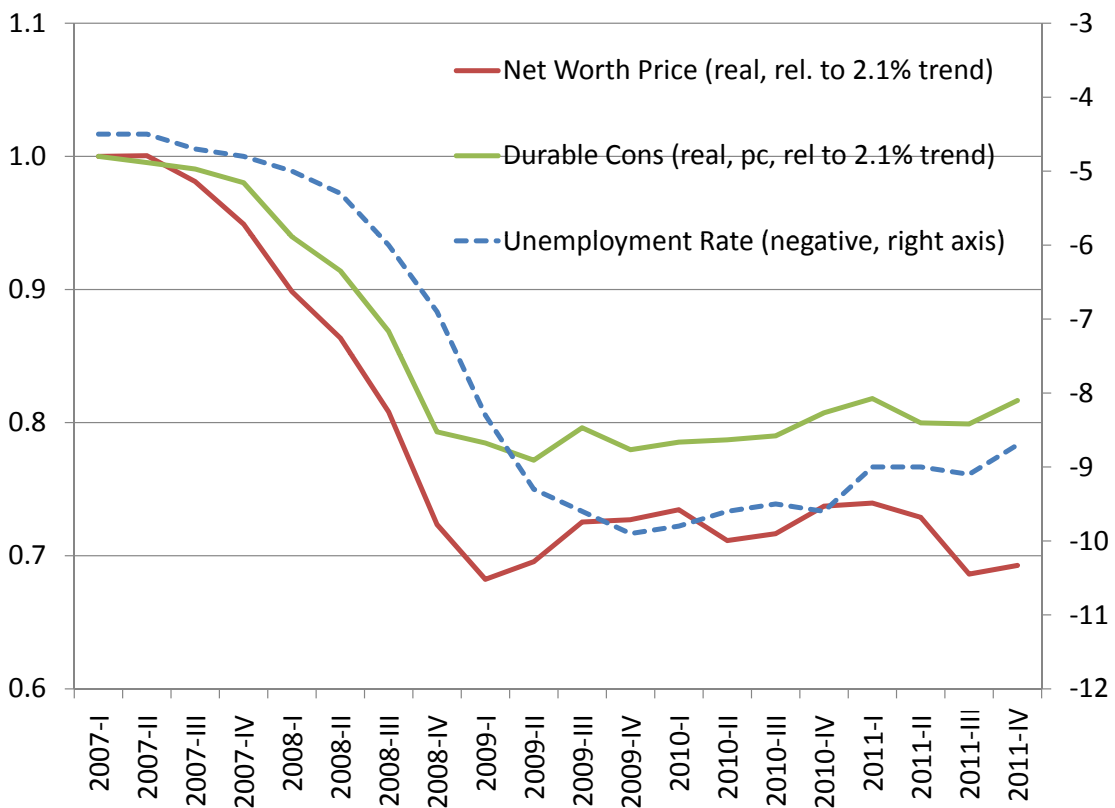


Figure 8: Asset prices, Consumption and Unemployment

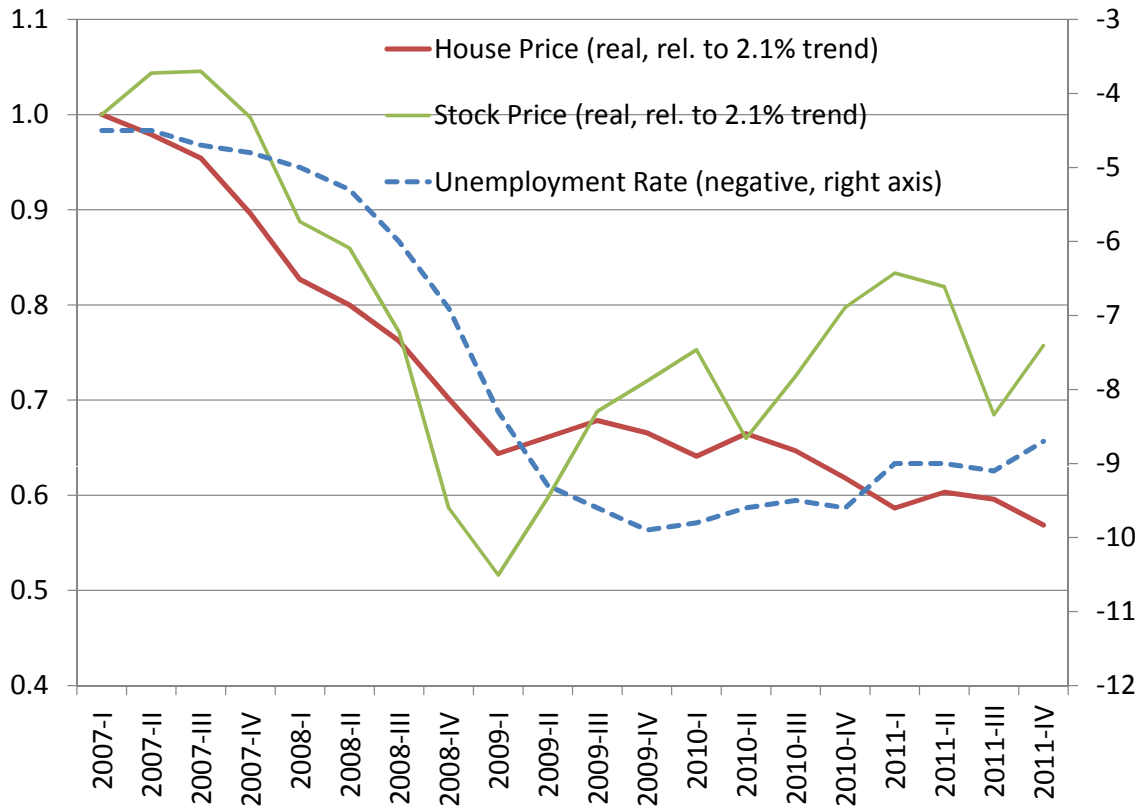


Figure 9: Housing prices, stock prices and unemployment

was concentrated in 2008, and the rise in unemployment between mid 2008 and mid 2009. Thus a sharp fall in house prices precedes the falls in consumer spending and the rise in unemployment. This fall in house prices drove the ratio of per capita home equity to per capita annual disposable income from a post-War high of 1.61 in the first quarter of 2006 to 1.19 by the fourth quarter of 2007, and to a post-War low of 0.70 by the end of 2011.

We now show that our model can generate dynamics that qualitatively resemble those described above. In the context of the model we interpret the events of the Great Recession as follows. Before 2006 the economy was robust in that the demand for housing was strong ($\phi > \bar{\phi}$) and thus there was no scope for expectations-driven cycles. Then, during 2006 and 2007, there was a fall in ϕ , the preference weight on housing, that left the economy fragile and vulnerable to a confidence-driven recession. In 2008 a negative aggregate confidence shock, perhaps triggered by the collapse of Lehman Brothers then led to a severe recession.

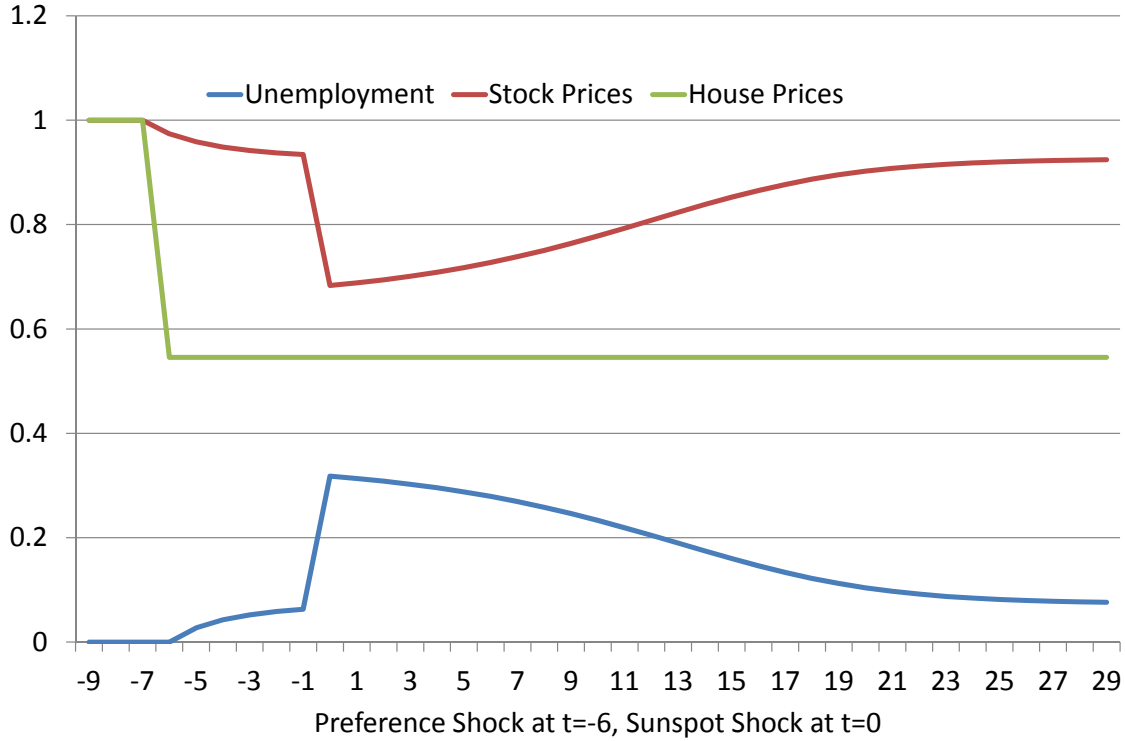


Figure 10: Model's prediction

Figure 9 shows the dynamics for a simulation of the model along the lines just described. At date $t = -6$ there is a permanent shock to preferences that reduces ϕ from its initial high value $\phi_0 > \bar{\phi}$ to a new lower value $\phi_1 = 0.05 < \bar{\phi}$. All other parameters are set to the same values as in the previous figures. From $t = -6$ onwards, we assume the equilibrium house price is constant at $p = 0.6$. Recall that at this price, instead of a unique full employment steady state, the model has two steady states, both with positive unemployment. To construct the figure, we assume that prior to date $t = 0$, the economy converges, with no surprises, towards the low unemployment steady state. Then, at $t = 0$, a one-off expectational shock generates a sharp fall in spending and an associated sharp rise in unemployment. Fearing high unemployment, households expect higher borrowing costs, and cut consumption orders to reduce credit costs. In aggregate this decline in demand does indeed translate into higher unemployment. After $t = 0$ the economy is hit by no further shocks, and the path for unemployment (the blue line) traces out the dynamics implied by Figure 4 below.

Note that in response to the negative sunspot shock, households cut back consumption, even though they expect positive income and consumption growth looking forward. Moreover, house-

holds have not experienced any loss in wealth (by assumption p is held fixed). The reason they nonetheless cut consumption is that higher unemployment risk generates a stronger precautionary motive to save. In fact this precautionary motive is so strong that positive expected consumption growth and an associated inter-temporal motive to dis-save is required to prevent a rise in house prices. After the sharp initial contraction, the economy gradually converges back towards the low unemployment steady state. In this recovery phase, the precautionary motive to save gradually declines, and so does expected income and consumption growth, such that the desired amount of household saving remains constant and equal to p .

8 Policy

TO BE COMPLETED

9 Extension to incorporate a centralized labor market

The crucial model assumption that opens the door to multiple equilibria is the decentralized labor market. With a conventional centralized labor market, absent a utility cost of working, all agents would choose to work at any positive wage, and full employment would be the only equilibrium.

The goal of this section is to present a model somewhere in between the baseline model, in which labor supply plays no role and the real wage does not move in response to changes in the unemployment rate, and the alternative flexible centralized labor market model, which cannot generate unemployment.

Suppose there are two sectors producing two different goods, c_1 and c_2 , using linear technologies. Good c_1 is produced by firms who hire workers in a centralized labor market, while good c_2 is produced by firms who hire in the decentralized fashion described in the baseline model.

Suppose preferences are

$$u(c_1, c_2, h, \phi) = \lambda \log c_1 + (1 - \lambda) \log c_2 + \phi h$$

The two goods are sold on a centralized market at prices p_1 and p_2 , where we normalize the latter to one, as in the baseline model. The household chooses the fraction κ of its members to send to work in the centralized market. Wages in the two markets are w_1 and w_2 .

All workers sent to the sector 1 find jobs that pay w_1 , while fraction $1 - u$ of the workers sent to sector 2 find jobs paying w_2 . In both sectors the sector-specific value of the marginal product is equal to the sector wage. In equilibrium the household must be indifferent between sending an additional worker to either sector.

The household budget constraint is now

$$(p_1 c_1 + c_2) + p(h' - h) = \kappa w_1 + (1 - \kappa)(1 - u)w_2 - \frac{\psi}{2}(1 - \kappa)u \max\{[(p_1 c_1 + c_2) - (ph - d)], 0\}^2 + T.$$

Note that from the perspective of the budget constraint the two types of consumption are treated symmetrically. Note also that workers sent to sector 2 may end up unemployed, which adds to potential credit costs.

Taking first-order conditions with respect to c_1 , c_2 , and κ gives

$$\frac{c_2}{c_1} = \frac{p_1(1 - \lambda)}{\lambda}$$

and

$$w_2 = \frac{w_1}{(1 - u)} + \frac{\psi}{2} \frac{u}{(1 - u)} \max\{[(p_1 c_1 + c_2) - (ph - d)], 0\}^2 \quad (9)$$

This is quite intuitive. It says that there must be an equilibrium wage premium in sector 2 that compensates for both the difference in job finding rates - the first term - and the extra credit risk from working in sector 2 - the second term. An interesting implication is that the wage differential between the two sectors must be increasing in the unemployment rate. This fits with the notion that relative wages for centralized labor market jobs like pizza delivery and fruit picking fall in recessions.

Suppose the production functions are

$$c_1 = \kappa \quad (10)$$

$$c_2 = (1 - \kappa)(1 - u)$$

The respective sector-specific wages are then

$$w_2 = p_2 = 1 \quad (11)$$

and

$$w_1 = p_1 = \frac{c_2}{c_1} \frac{\lambda}{1 - \lambda} = \frac{(1 - \kappa)(1 - u)}{\kappa} \frac{\lambda}{1 - \lambda}. \quad (12)$$

This last equation says that the relative wage in the centralized sector will be higher: (1) the stronger is the relative taste for good 1, (2) the smaller is the fraction of workers sent to the centralized sector, and (3) the lower is the unemployment rate in the decentralized sector.

We can now go back to the labor-sector-choice-indifference condition to solve for the equilibrium value for κ , given u and p . Substituting 11 and 12 into 9 along with 10 gives

$$1 = \frac{(1 - \kappa)}{\kappa} \frac{\lambda}{1 - \lambda} + \frac{\psi}{2} \frac{u}{(1 - u)} \max \left\{ \left[\frac{(1 - \kappa)(1 - u)}{1 - \lambda} - (p - d) \right], 0 \right\}^2 \quad (13)$$

Given u and p , this equation implicitly defines a solution for κ . Absent credit costs ($\psi = 0$) the solution is simply $\kappa = \lambda$. Why is the unemployment rate irrelevant in this case? The intuition is that given Cobb-Douglas preferences, holding fixed κ (at $\kappa = \lambda$) a higher unemployment rate and lower associated output of good 2 translates automatically into a lower relative price of good 1 and thus a lower relative wage in the centralized sector. This decline in the wage in the “safe” sector exactly offsets the decline in the job finding probability in the “risky” sector, and the household optimally maintains the same division of job seekers between the two sectors.

When $\psi > 0$ and unemployed agents are using credit, the unemployment rate factors into the optimal choice for κ . In particular, the fraction of household members in the centralized market κ will be increasing in the unemployment rate. Intuitively a greater relative supply of good 1, the centralized market good, is required to increase the equilibrium wage differential between the two sectors, which compensates for the extra credit costs that unemployed workers in the decentralized sector will pay.

We have not yet addressed the main question of interest, which is whether the introduction of a centralized labor market fundamentally changes the nature of equilibrium dynamics that can arise in the model. The answer is no.

The inter-temporal first-order condition for housing is essentially unchanged relative to the baseline model. Specifically, assuming a constant house price, it is now

$$\frac{c_2^{-1}}{1 + \psi(1 - \kappa)u \max \{[(p_1 c_1 + c_2) - (p - d)], 0\}} = \beta \frac{\phi}{p} + \beta E [c_2^{-1}]$$

where $c_2 = (1 - \kappa)(1 - u)$ and κ is the implicit function of u and p described in 13. Thus the key equilibrium condition defining the equilibrium dynamics for unemployment is very similar to the one for the baseline model, the only difference being the presence of κ .

The nice thing about the model is that the parameter λ controls the overall size of labor market frictions: as λ moves from zero to one, the economy moves from our original model to a Walrasian

one.

10 Conclusions

TO BE ADDED

References

- [1] Benito, A. (2006) “Does Job Insecurity Affect Household Consumption?”, *Oxford Economic Papers* 56.
- [2] Chamley, C. (2010) “Self-fulfilling Traps with Money as a Medium of Exchange” Working paper.
- [3] Chetty, R., A. Guren, D. Manoli, and A. Weber (2011) “Are Micro and Macro Labor Supply Elasticities Consistent? A Meta-Analysis of Extensive Margin Elasticities” NBER Working Paper 16729.
- [4] Cooper, R., and A. John (1988), “Coordinating coordination failures in Keynesian models.” *Quarterly Journal of Economics* 103 (3), 441-63
- [5] Corsetti, G., A. Meier and G. Muller (2010), “What determines government spending multipliers?”, Working paper
- [6] Farmer, R. (2010), “Confidence Crashes and Animal Spirits”, Working paper
- [7] Guerrieri V. and G. Lorenzoni (2009) “Liquidity and Trading Dynamics”, *Econometrica* 77(6).
- [8] Hall, R.E. (2005), “Employment Fluctuations with Equilibrium Wage Stickiness”, *American Economic Review* 95(1), 50-65.
- [9] Kocherlakota, N. (2009), “Bursting Bubbles: Consequences and Cures”, Working paper
- [10] Lettau, M. Ludvigson, S., and J. Wachter (2008), “The Declining Equity Premium: What Role Does Macroeconomic Risk Play?” *Review of Financial Studies*, 21, 1653-1687
- [11] Mian, A., Rao, K., and A. Sufi (2011), “Household Balance Sheets, Consumption, and the Economic Slump”, Working paper.

- [12] Michaillat, P. (2012), “Do Matching Frictions Explain Unemployment? Not in Bad Times”, *American Economic Review*
- [13] Shimer, R. (2012), “Wage Rigidities and Jobless Recoveries”, *Journal of Monetary Economics*