Unemployment and Credit Constraints in a Heterogeneous Agent Model

Brian O'Quinn*

December 8, 2017

Abstract

What is the link between the drop in consumer credit during the Great Recession and increased unemployment? I build a heterogeneous household model with endogenous idiosyncratic risk of unemployment, incomplete insurance, sticky wages, and a central bank that follows a predetermined interest rate rule. After a shock to their credit constraints, households try to save more and thereby reduce their spending. This results in job rationing because prices are rigid. With a typical interest rate rule, I find that a tightening in credit constraints that matches the decline in consumer credit between 2008:Q2 and 2010:Q3 can explain about a 1 percentage point increase in unemployment. Without an interest rate decrease, my model exhibits a 5.36 percentage point increase in unemployment.

1 Introduction

In the years prior to the U.S. recession in 2008–09, households had been steadily taking on more and more debt. When the recession hit, this trend was abruptly broken, and households began to reduce their debt levels very quickly over the next several quarters. Real consumer credit per capita declined

*Department of Economics, Northwestern University (e-mail: b-oquinn@northwestern.edu)
I would like to thank Giorgio Primiceri for his guidance and counsel. I also appreciate Larry Christiano and Guido Lorenzoni for their helpful suggestions and comments. Finally, I am grateful to Mirko Wiederholt and Jonathan Parker for their input in the early stages of this paper.

to 2002 levels over the course of nine quarters. From its peak in 2008:Q2, real consumer credit per capita fell 9.65% to a trough in 2010:Q3. This time period was also marked by a significant increase in the unemployment rate. Prior to the crisis, the unemployment rate was around 5%. During the recession, it rose to a peak of 10%.

A common thread in the stories offered for this severe recession is an adverse positive feedback loop.[1] One such loop, centered on consumer saving behavior, is the paradox of thrift. As the paradox classically goes, as a result of some impetus, households choose to save more of their income. Thus they decrease demand for consumption goods, which causes economic output to fall. This, in turn, lowers the income of households and further reduces demand for consumption. In most formulations of the paradox, aggregate saving remains unchanged or even shrinks.

The initial shock that triggers this cycle varies across narratives. Explanations for this shock include popping asset bubbles (and the resulting wealth effect), confidence shocks, inflation expectations shocks, preference shocks, and credit constraint shocks.

The purpose of this paper is to present a simple model that captures this positive feedback loop of the paradox of thrift. However, I include an additional feedback channel: the motivation to self-insure against unemployment. In the model, a tightening of credit constraints serves as the initial shock. As a consequence of the cutback in consumer credit, households face increased incentive to save and curtail demand for consumption. This results in a decline in output and employment. As the risk of unemployment rises, households want to increase precautionary saving and will reduce consumption demand further, thereby perpetuating the cycle.

The main question that I address is a quantitative one: What portion of the observed increase in unemployment during the recession can be attributed to the tightening of households' credit constraints and the consequent fall in consumer credit?

My model is built on four components: heterogeneous households, a credit constraint shock, nominal wage rigidity, and an interest rate rule. I require heterogeneous households because an economic

[1] A positive feedback loop is one in which the effect of a process tends to amplify the cause. These have also been called, in various conceptualizations, "vicious cycles" and "death spirals." They contrast with negative feedback loops which damp themselves.

model with one representative household generally cannot describe debt. Some households must have positive asset holdings and some must have negative asset holdings (debt). I introduce heterogeneity by assuming binary employment outcomes and incomplete insurance. A household is either employed or unemployed. There is no intensive margin for labor such as varying hours worked. This is the starting point for creating differences between households. Employed households earn a wage, and unemployed households earn an unemployment benefit which is less than the wage. Therefore, households will differ by their employment history.

Incomplete insurance is required to maintain these differences between households. Otherwise, households would agree to an income-sharing arrangement which would make their employment histories irrelevant to their budget constraint.

With these two assumptions, households will differ by their levels of wealth. Households that have been employed for a while will generally have higher wealth than households that are in a spell of unemployment. These two assumptions make for a heterogeneous agent model that I have to solve computationally.

The next component is an exogenous credit constraint for all households. The economy will initially be in a steady state with loose credit and some small amount of unemployment. The credit constraint is then tightened by an unanticipated and permanent shock.

After the credit constraint shock, households will increase demand for saving for two reasons. First, a fraction of households will be below the new credit constraint and will have to save more to come into compliance with the new, tighter constraint. Second, all households now have a smaller wealth buffer from which to draw in case they are unemployed, and they will want to save more to restore that buffer.

This increase in demand for saving could be negated by an appropriate decrease in the real interest rate. While those households below the new credit constraint still need to save more, a decrease in the real interest rate will cause wealthier households to reduce saving. If the real interest rate fell sufficiently, the macroeconomic effect of the increased demand for saving would be nullified. Therefore, in order for the increase in demand for saving to influence the economy as a whole, I need to restrict the movement of the real interest rate.

The real interest rate is a function of the inflation rate and the nominal interest rate. This leads me to the last two components of the model: nominal wage rigidity and an interest rate rule for the central bank. Wage rigidity is a common assumption in New Keynesian models. In my model, nominal wage rigidity restricts the inflation rate, and the central bank's interest rate rule determines the nominal interest rate. After the credit constraint shock, the real interest rate will not decrease enough to negate the shock's effect, and there will be an increase in demand for saving.

Firms are perfectly competitive and their only input is labor. Thus, since nominal wages are rigid, so too is the price of output. This implies real wages are also rigid, and that results in job rationing after a negative demand shock. As households try to save more and demand for consumption goods falls, prices are unable to adjust. Therefore, firms will lay off workers, and the unemployment rate will rise. In this way, my model captures the paradox of thrift.

There are two channels through which the positive feedback cycle continues. The first is the typical Keynesian multiplier effect, or income effect. Diminished output implies diminished household income which causes demand for consumption goods to fall even further. The second is a precautionary saving effect. The rise in the unemployment rate implies that the probability of a household losing its job has increased and the expected duration of an unemployment spell has lengthened. In light of this heightened risk of unemployment, households will increase demand for saving even further.

My model allows me to quantify how much the increase in unemployment can be attributed to the credit constraint shock. If I calibrate the shock to match the fall in consumer credit between 2008:Q2 and 2010:Q3, the model can explain approximately a 1 percentage point increase in the unemployment rate. Adjusting for the upward trend in consumer credit prior to the shock, I can explain a 1.38 percentage point increase in the unemployment rate. This is in light of the fact that consumer credit makes up only 17.9% of household liabilities. Furthermore, in the model, the central bank responds in the same quarter as the shock by lowering the interest rate: there are no delays.

Moreover, my quantitative model can demonstrate the importance of the interest rate responding to changes in unemployment through the interest rate rule. I consider a situation where the interest rate does not fall after the credit constraint shock. This is to proxy for the case wherein the interest rate is at the zero lower bound. Alternatively, it could be that the central bank is simply unresponsive. Either way, after the credit constraint shock, the unemployment response is much greater when the interest rate is unable to fall. In this situation, after an identical shock to credit constraints, the unemployment rate will rise 5.36 percentage points.

These two results lead me to conclude that it is unlikely that the rapid drop in consumer credit alone explains the levels of unemployment observed in the U.S. during the Great Recession. While the unemployment rate increased about 5 percentage points, consumer credit alone explains only 1 to 1.38 percentage points. There must have been other factors at work: these factors likely brought the nominal interest rate to the zero lower bound. In that case, the shock to consumer credit can explain the 5 percentage point increase in the unemployment rate.

In the next section, I summarize the related literature. I set up the assumptions and equations of the model in Section 3. In Section 4, I calibrate most of the parameters for the model and examine the baseline case where the credit constraint is constant over time. Section 5 features the main experiment of the paper. I calibrate the shock to the credit constraint and present a number of figures as to the shock's effects on the economy. In Section 6, I provide a set of alternative calibrations and experiments, including my proxy for an economy with the nominal interest rate at the zero lower bound. Section 7 concludes.

2 Related work

Since my model is a heterogeneous agent model, I have solved it using the computational method illustrated in Krusell and Smith (1998). In their model, households differ by their level of capital holdings, and employment outcomes are determined by an exogenous first-order Markov process. In their model, as well as in mine, households differ by their wealth levels. However, in my model, the probability of a household being employed is based on their employment state in the previous period and the employment rate as determined by aggregate demand. Thus, in this paper, employment is endogenously determined.

Eggertsson and Krugman (2012) build a model with patient and impatient agents. That is, some of their agents have high time discount factors (β's) and some have low. All impatient agents are exactly at their credit constraint: the credit constraint binds with equality for them. Given this, they explore how a credit constraint shock would affect the economy. All impatient agents are forced to move, in one period, out of violation of the new tighter credit constraint. My model differs in that all households have identical preferences. Furthermore, in my model, only a small fraction of households will ever be in direct violation of the credit constraint after it is shocked. Additionally, I allow those households to gradually come into compliance with the new credit constraint instead of forcing them to do so in one period. That said, Eggertsson and Krugman's model is analytically solvable due to the fact that there are just two types of households.

Guerrieri and Lorenzoni (2017) also explore tightening credit constraints in a heterogeneous agent model. In their paper, the decrease in output following the shock to the credit constraint is due to low productivity workers working more while high productivity workers work less. They also extend their model to include sticky prices and consider different paths for the interest rate after the shock. Since the decline in output is due to a fall in average productivity, they provide more of a supply-side explanation for the recession than a demand-side one. In my model, all workers have identical productivity, and the decrease in output will be due to diminished demand for consumption goods.

Hall (2011) illustrates how a fixed interest rate can lead to unemployment. He shows that inflation is mostly exogenous over time: prices do not necessarily fall when there is high unemployment. Given this, when the nominal interest rate is bound by the zero lower bound, the real interest rate is also constrained. If this bounded real interest rate does not match the real interest rate that would be implied by inter-temporal preferences and production technology, there will be unemployment. Like Eggertsson and Krugman (2012), he has two types of households: some households are always at their credit constraint and others are not. Among other components, his model features sticky real wages and exogenous inflation. He demonstrates that when the interest rate is pinned, there is unemployment.

My model is similar in its assumptions about wages and inflation. In my model though, I do not have a fixed fraction of agents at the credit constraint. Also, his experiment, at the core, examines a pinned versus unpinned interest rate. My model takes a step back from that and considers a credit constraint shock which will endogenously affect the interest rate.

Michaillat (2012) lays down a framework for discussing job rationing and demonstrates how rationing can occur when wages are rigid and marginal product of labor is decreasing. My model features fixed wages, and in effect, diminishing marginal product of labor. While the production function has constant marginal product of labor, the marginal revenue product of labor is decreasing. Past a point, while a firm could hire another worker to produce another unit of output, the firm would be unable to sell that unit of production, even for its marginal cost. My model has no matching frictions and thus no frictional unemployment: all unemployment will be rationing/cyclical unemployment.

Ravn and Sterk (2017) feature a model with heterogeneous households that save for precautionary reasons. They include a job matching aspect with two pools of unemployed workers: the short-term unemployed and the long-term unemployed. They then shock the job matching component of the model and consider the effects. While the setup of the model is similar, the questions I consider are different. Whereas they are focused on shocks to the job matching aspect of the model, I study a change in credit constraints and do not need to include job matching.

Schmitt-Grohé and Uribe (2012) create a representative agent model to explain the recession, and in particular, the jobless recovery. They shock inflation expectations, and this causes the economy to fall into a liquidity trap, that is, a period where the zero lower bound on the nominal interest rate binds. Similar to my model, they assume households supply labor perfectly inelastically and include downward nominal wage rigidity. However, this paper and theirs do have differences regarding these two assumptions. In their model, the representative household can work for some or all of its time endowment. If the household does not work all its hours, they call that unemployment. In my model, households either work or do not, and the fraction of households that do not work is unemployment. Also, my nominal wage rigidity assumption is technically stronger than theirs.

There are two key differences between their model and mine. First, they shock inflation expectations, and I shock the credit constraint. Second, in their model, in order to get unemployment, the nominal interest rate must hit the zero lower bound. In my model, I can get a response of increased unemployment after the shock without the zero lower bound binding.

3 Model

Households

This is a quarterly model, and there are three types of actors: households, firms, and the government.

In the model, there is a continuum of households of measure $I = 1$, indexed by i. Households are infinitely-lived and risk averse. There is only one type of consumption good and households value it using an isoelastic (constant relative risk aversion) utility function. Households gain no utility from leisure and suffer no disutility from working. Therefore, their utility function, with coefficient of risk aversion γ, is as follows:

$$u(c_{i,t}) = \frac{c_{i,t}^{1-\gamma} - 1}{1-\gamma}$$

A household i starts any given period t with some level of bonds $b_{i,t-1}$, which it chose in the previous period. Households can choose to hold a negative level of bonds; this represents debt. If a household is employed, it will earn a wage w_t, which is taxed at a rate τ_t. If a household is unemployed, it collects an unemployment benefit from the government η_t and is not taxed. These things make up a household's budget, as shown in Equation 3.2.

A household will spend its entire budget on purchasing consumption goods, $c_{i,t}$, or bonds, $b_{i,t}$. Consumption goods are purchased at price p_t, and bonds are purchased at price q_t. All households are bound by a credit constraint $\bar{b}_t$: they must hold at least $\bar{b}_t$ bonds, where $\bar{b}_t$ is a negative number. This will be the source of the aggregate shock to the economy.

Given all this, each household solves a recursive maximization problem, discounting utility across time by discount factor β.

$$V(e_{i,t}, b_{i,t-1}; \Gamma_t) = \max_{c_{i,t}, b_{i,t}} u(c_{i,t}) + \beta E_t V(e_{i,t+1}, b_{i,t}; \Gamma_{t+1}) \tag{3.1}$$

subject to:

$$p_t c_{i,t} + q_t b_{i,t} = w_t(1-\tau_t)e_{i,t} + \eta_t(1-e_{i,t}) + b_{i,t-1} \tag{3.2}$$

$$b_{i,t} \geq \bar{b}_t \tag{3.3}$$

In period t, household i's idiosyncratic state variables are $e_{i,t}$, the employment state of the household, and $b_{i,t-1}$, the quantity of bonds purchased by the household last period. For any period t, the household is either employed, $e_{i,t} = 1$, or unemployed, $e_{i,t} = 0$. A household's employment status is not chosen by the household, but rather, is determined by aggregate demand for consumption goods.

Since households are idiosyncratically employed or unemployed, they will vary in their employment histories. Furthermore, since there is incomplete insurance as in Bewley (1977), households hold bonds, not just to earn a rate of return, but as a precaution against unemployment. Therefore, the idiosyncrasy in employment outcomes causes households to be heterogeneous in their levels of bond holdings. The joint distribution of employment states, $e_{i,t}$, and bond holdings $b_{i,t-1}$, is represented by the term Γ_t. In particular, note that households observe the current level of employment when they make their decisions.

Firms

There is an infinite number of perfectly competitive firms. Firms employ a measure of households L_t and produce consumption goods C_t according to a linear production function:

$$C_t = AL_t$$

Thus, firms maximize profits according to the following optimization problem:

$$\max_{L_t} p_t AL_t - w_t L_t$$

The first order condition implies a simple relationship between the price of the consumption good and the wage.

$$w_t = p_t A$$

Furthermore, firms earn zero profits.

Employment dynamics

At the start of the period, a fraction λ of employed households are separated from their jobs. Those employed households who avoid this separation will have jobs this period, assuming firms terminate no jobs. In this way, a household's employment state is somewhat persistent. Households employed in the previous period have a higher chance of employment than those unemployed in the previous period.

These factors combine to give the following probabilities, where $L_t \in [0,1]$ represents the fraction of households employed in period t, and $\tilde{L}_t = (1-\lambda)L_{t-1}$ represents the fraction of the households that survive the separation shock in period t.

$$\Pr(e_{i,t} = 1 \mid e_{i,t-1} = 1) \quad = \quad \begin{cases} (1-\lambda)\frac{L_t}{\tilde{L}_t} & L_t < \tilde{L}_t \\ (1-\lambda) + \lambda\frac{L_t - \tilde{L}_t}{1-\tilde{L}_t} & L_t \geq \tilde{L}_t \end{cases}$$

$$\Pr(e_{i,t} = 1 \mid e_{i,t-1} = 0) \quad = \quad \begin{cases} 0 & L_t < \tilde{L}_t \\ \frac{L_t - \tilde{L}_t}{1-\tilde{L}_t} & L_t \geq \tilde{L}_t \end{cases}$$

For example, if $L_t < \tilde{L}_t$, then firms want to further reduce their number of employees beyond those removed by the λ shock. The probability of a previously employed worker having a job then is the probability they survive the λ separation shock times the probability that they are of the L_t chosen among all of the workers in $\tilde{L}_t$.

Model assumptions

At this point, it is worthwhile to talk about why there is unemployment. Recall that households have no disutility from working and supply labor perfectly inelastically. Consider a situation in which some households are employed and others are unemployed but want to work at the current wage. The classical response is to predict that the real wage, w_t/p_t must fall until there are no unemployed households.

Since this is a model for recessions, and during the recession there is greater than normal unemployment, the nominal wage would only have a tendency to fall. However, many macroeconomic models assume that wages have nominal downward rigidities, such as Schmitt-Grohé and Uribe (2012). In the United States, this assumption is supported by survey data examined by Barattieri, Basu, and Gottschalk (2014). In my model, I assume that the nominal wage cannot decrease. And since the nominal wage would only decrease in the model, this implies that the nominal wage is constant. I normalize the nominal wage to $w_t = 1$ for all t. Despite fixing the nominal wage, the classical response could still work: the real wage, w_t/p_t, can be made to fall by increasing p_t.

However, since firms are perfectly competitive, they will price the consumption good at $p_t = \frac{w_t}{A}$. Therefore, with a fixed nominal wage and perfectly competitive firms, the real wage is fixed at $\frac{w_t}{p_t} = A$.

Note that there are three markets: labor, consumption goods, and bonds; therefore, there are three prices: w_t, p_t, and q_t. I have discussed how nominal wages and the price of consumption goods are set and are incapable of adjusting to bring the economy from a state with unemployed households to a state of full employment.

The last hope for full employment is in the price of bonds, q_t. By lowering the bond price, a benevolent social planner can encourage saving, and inversely, discourage consumption. If the bond price can take on any value, then by adjusting the bond price appropriately, the planner can target the level of consumption corresponding to full employment $C_t = A$. However, the government will follow an interest rate rule, and this determines the bond price exactly. Because the bond price is constrained, as well as the other two prices, the model can exhibit unemployment in the form of job rationing.

Throughout the rest of the paper, I drop the subscripts on wage, w, and the price of the consumption good, p, since they are constant over time.

Government

The government taxes wages and sets unemployment benefits according to some ratio ρ, which is the ratio between the unemployment benefit and the after-tax wage. For example, when $\rho = 0.4$, an unemployed household receives benefits equal to 40% of the after-tax income of an employed household (excluding interest income).

$$\rho = \frac{\eta_t}{w(1-\tau_t)} \tag{3.4}$$

The government maintains a fixed level of debt. Thus, households as a group always hold positive net assets (bonds). Let $B > 0$ represent this fixed level of government-held debt per capita. The government always runs a balanced budget: tax revenue equals transfer payments plus service on the debt.

$$w\tau_t L_t = \eta_t(1-L_t) + (1-q_t)B \tag{3.5}$$

By way of Equations 3.4 and 3.5, the unemployment benefit, η_t, and the wage tax, τ_t, are functions of L_t alone.

The government also sets the nominal interest rate according to an interest rate rule involving the level of employment. Since inflation is zero, as discussed above, no inflation term appears in this interest rate rule. Furthermore, since inflation is zero, the real interest rate, r_t, is always the same as the nominal interest rate.

$$r_t = \max\left(r^* + \psi(L_t - L^*), 0\right) \tag{3.6}$$

The term r^* is the target real interest rate. The term L^* is the target employment rate, and ψ is the government's responsiveness to the employment gap. In reality, households can always hold currency and earn a nominal rate of return of zero. Thus, in the model, I assume a zero lower bound on the nominal interest rate. Since the nominal and real interest rates are always equal, the real interest rate cannot be less than zero.

Again, this is a quarterly model, and I specify interest rates in annual terms. Therefore, the bond price q_t is related to the interest rate according to the following equation:

$$q_t = \frac{1}{(1+r_t)^{1/4}}$$

Equilibrium

For any given period t, an intra-temporal equilibrium is defined by a joint distribution of saving decisions and employment outcomes, Γ_t, where all of the above equations hold and the three markets are cleared:

- The measure of households working equals the measure of households employed:

$$L_t = \int e_{i,t} \mathrm{d}\Gamma_t$$

- The number of bonds sold by the government equals the net measure of bonds saved:

$$B = \int b_{i,t} \mathrm{d}\Gamma_t$$

- Total production equals the measure of total consumption:[2]

$$C_t = \int c_{i,t} \mathrm{d}\Gamma_t$$

Equilibrium is defined as the sequence of intra-temporal equilibria where the evolution of Γ_t is consistent with the household policy function. The economy is in steady state in period t if $\Gamma_t = \Gamma_{t+1}$.

Solution method

My solution method is similar to Krusell and Smith (1998). To briefly summarize it, the first step is to guess laws of motion for various moments of the Γ_t distribution. Given these laws of motion, I estimate the households' value and policy functions by value function iteration. I then simulate the economy given the households' policy function. After that, I estimate new laws of motion for the moments of Γ_t based on the simulation's results. If the estimated laws of motion are similar to the guessed ones, then I have found valid laws of motion and policy functions. If the estimated laws of

[2] This condition is mathematically implied by the two other market-clearing conditions and Equation 3.2. It is included for the sake of completeness.

Name	Variable	Value
Utility discount factor	β	$\frac{1}{1.025^{1/4}} \approx 0.9938$
Constant of relative risk aversion	γ	4
Nominal wage	w	1
Total factor productivity	A	1
Job separation shock	λ	0.1
Unemployment benefit to post-tax wages	ρ	0.4
Target annual interest rate	r^*	0.025
Interest rate response to employment gap	ψ	0.5
Target employment rate	L^*	0.95
Permanent government debt	B	1.30145
Household credit constraint	$\bar{b}$	−6.374

Table 1: Baseline parameter values

motion differ from the guessed ones, then I repeat the process with new estimated laws of motion serving as the guessed laws of motion.

Computationally, the only moment required of Γ_t is the mean of the marginal distribution of employment. That is, households only need the current employment rate and its law of motion to make their consumption/saving decisions. This is similar to a result in Krusell and Smith (1998), where only the mean of the capital distribution is required, while the variance and other statistics are extraneous.

4 Baseline calibration and results

I first calibrate the model without the credit constraint shock and examine the steady state.

Calibration

Table 1 describes the values of the parameters in the baseline calibration.

I want the household discount rate, β, to correspond to an annual interest rate of 2.5% if this was a model with perfect insurance (a representative agent model). Note that β is provided in quarterly terms by the formula in the table. The nominal wage is fixed by assumption. The choice of w is a decision about how to scale the other variables. As previously mentioned, I set $w = 1$. Similarly, total factor productivity just scales things, so I fix $A = 1$. Shimer (2005) determines the quarterly job separation rate to be $\lambda = 0.1$ and has the unemployed earn 40% of what the employed earn, so $\rho = 0.4$.

I make $\psi = 0.5$, which is a typical coefficient for the output gap in an interest rate rule.[3] I set $L^* = 0.95$ and $r^* = 0.025$. This will imply the unemployment rate is about 5%. Of course, these two parameters are not uniquely determined: any L^* and r^* satisfying the following equation would be equivalent:

$$r^* - 0.5L^* = 0.025 - 0.5(0.95) = -0.45$$

Setting $r^* = 0.025 = \left(\frac{1}{\beta}\right)^4 - 1$ is convenient, based on how β was chosen. However, it should be expected that the steady-state interest rate will be less than r^* and the steady-state employment rate will be less than L^*. In a world with perfect insurance, the annual interest rate would equal $r = r^* = \left(\frac{1}{\beta}\right)^4 - 1$. However, without perfect insurance, saving provides not only a rate of return from interest, but also insurance against unemployment. Therefore, the required rate of return on saving will be less than $\left(\frac{1}{\beta}\right)^4 - 1$. Thus, by way of Equation 3.6, steady-state employment will be less than L^*. Of course, this should not be alarming: equivalent r^* and L^* exist such that they are equal to their steady-state values.

Targeting around 5% unemployment is, of course, a bit wasteful; there is nothing in the model to prevent the government from targeting 0% unemployment. I chose the 5% unemployment target to reflect the reality of natural frictional/structural unemployment, and this simply models that. In fact, an earlier specification of the model included a reduced form of frictional unemployment, where a random 5% of households were made unemployed and unemployable for the period, and thus the maximum employment rate was 95%. That feature was dropped for simplicity's sake, since, for the most part, it only justified targeting 5% unemployment.

To calibrate B and $\bar{b}$, I chose to look at U.S. consumer credit and exclude other household liabilities. Naturally, home mortgages make up the majority of household liabilities. However, I suspect that the

[3]Note, however, that it is actually being applied to a difference in employment rates or output, and not an output gap percentage. This is a fairly trivial distinction though.

model would need to be expanded to include homes and mortgages explicitly if I wanted to properly include home mortgages.

Similarly, from the balance sheet for U.S. households, I chose to examine only the most liquid assets: currency, deposits, and money market funds, to the exclusion of other household assets. Again, real estate makes up a large portion of household assets, but I have excluded it for reasons already mentioned. I have also excluded corporate equity holdings, savings bonds, and other financial instruments. I believe that these types of assets are saved mostly to fund college education or retirement. Since the model does not simulate either college expenses or retirement, to include those assets would require that model households hold a relatively large amount of assets for precautionary saving motives. Furthermore, with such high levels of mean asset holding, it is very hard to have indebted households in the model.

Explained another way, when the household wealth distribution's mean is so far to the right, it is very hard to get the tail of the distribution left of zero such as to have sufficient household debt compared to the consumer credit data. The only way to increase the variance of the distribution is to lower the wealth floor $\bar{b}$, but that ceases being effective at increasing variance past a point. The idiosyncratic unemployment shocks only provide so much variance to household wealth. When the economy has so much in assets, households never have to enter debt.

The consumer credit series and the "currency, deposits, and money market funds" series are drawn from Table B-100 of the Federal Reserve's Z.1 report, March 6, 2014. I adjust the series for seasonality using X-13ARIMA-SEATS.[4] I combine this with the seasonally adjusted GDP series from the U.S. Bureau of Economic Analysis's million-dollar National Income and Product Accounts tables to get debt-to-GDP and assets-to-GDP.

The variables B and $\bar{b}$ are calibrated such that debt-to-GDP and assets-to-GDP match 2008:Q2, the quarter before the noticeable drop in consumer credit. These ratios are 0.1785 and 0.5214 respectively.

[4] I also tried other methods of seasonal adjustment and found similar results.

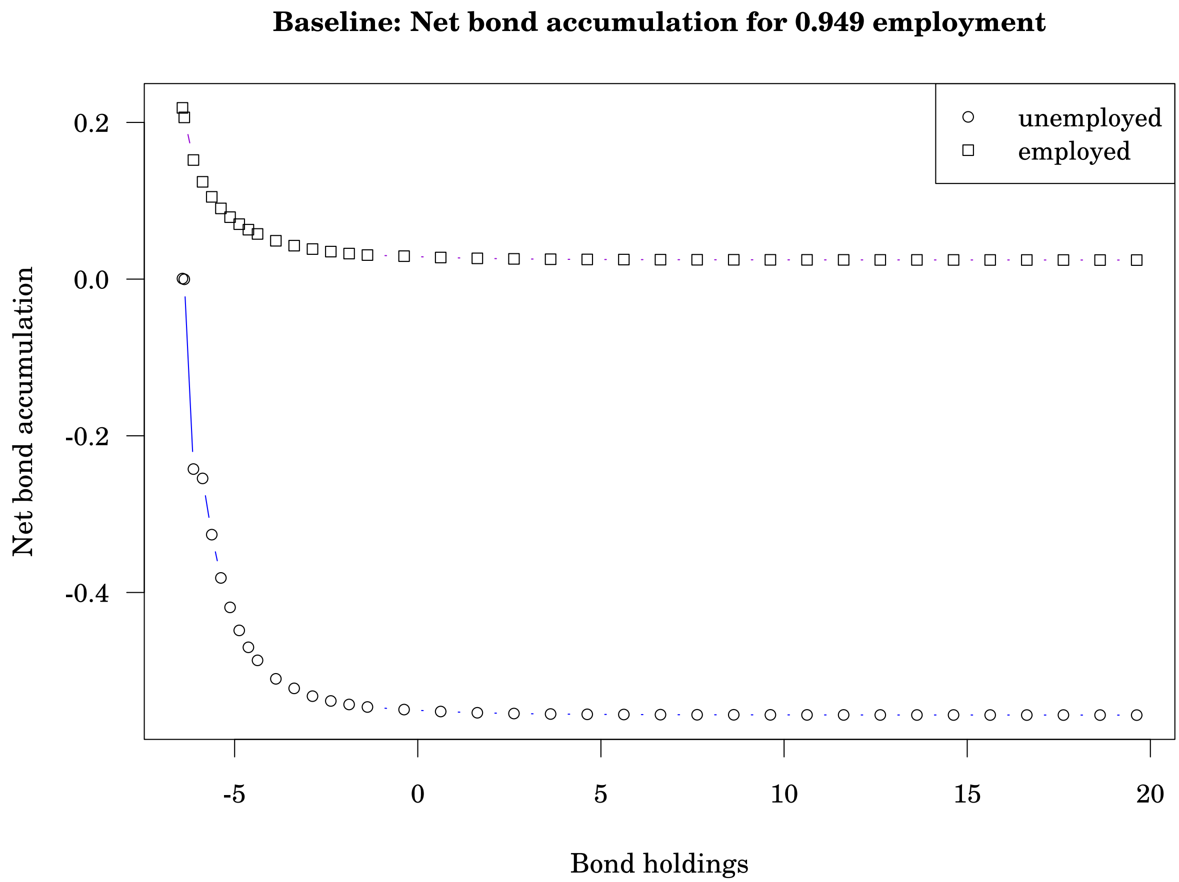

Figure 1: Net bond accumulation for 94.9% employment in the baseline model

Results

In this baseline model without a credit constraint shock, it is informative to consider the policies of households. In steady state, the economy has an unemployment rate of 5.1%, so I will examine how much households want to save at this level of employment. Figure 1 shows net bond accumulation for employed and unemployed households at various bond holding levels.

The horizontal axis corresponds to the household's wealth or bond holdings. Very poor households are to the left, and rich households are to the right. Close to the credit constraint, households save a lot, and this decreases as household wealth increases. As expected, employed households save, increasing their bond holdings over time, and unemployed households dissave.

Intra-temporal equilibrium is found by looking for a fixed point of L, the measure of households employed. First I suppose some value of L^1, a guess at the level of employment.[5] By way of the central bank's interest rate rule, this provides r^1. Furthermore, with this L^1, I can determine which households are employed and which are unemployed based on the previous employment distribution. At this point, households have all the information they need to use their policy rule and report how many bonds they want to buy. They know the wage, the price of output, and the price of bonds (by way of r^1). They know their current wealth and whether they are employed or not. They finally know what the employment rate is today, which informs them as to what the employment rate is likely to be tomorrow through an estimated law of motion for employment rates. With this information, they indicate how many bonds they are willing to buy, or equivalently, how much output they want to buy, C^1. If $C^1 \neq AL^1$, then this is not an intra-temporal equilibrium. If $C^1 \neq AL^1$, then I can find a new $L^2 = C^1/A$, and repeat this process until $L^{k+1} = L^k$ for some natural number k.

The solution method helps explain how Figure 2 is constructed. On the horizontal axis is output AL. On the vertical axis is aggregate demand / consumption, C. The vertical line at 0.949 is the intra-temporal equilibrium level of output, which corresponds to the unemployment rate 5.1%. The 45-degree line represents the firms' side of the economy, $C = AL$. The consumption demand line represents the household and government's side. It plots the demand for consumption goods given a particular level of employment and output. This is not a typical Keynesian cross: consumption demand begins to decrease around 0.9 output or 90% employment. This is due to the interest rate rule of the central bank. To the left of 90% employment, the interest rate is zero. To the right, however, the interest rate is increasing. This causes households to divert from spending to saving. Since the interest rate is not constant, I call this graph a modified Keynesian cross.

If the interest rate was constant at zero throughout, then consumption demand would continue its mostly linear trend and intersect the 45-degree line beyond output = 1. That is, there would be no intra-temporal equilibrium. There would be excess demand for consumption goods and excess supply of bonds. At very high levels of employment, the precautionary motive to save is small, and households must earn a positive interest rate, near $1/\beta - 1$, in order to be enticed to buy all the bonds the government supplies.

[5] The superscript here does not indicate an exponent.

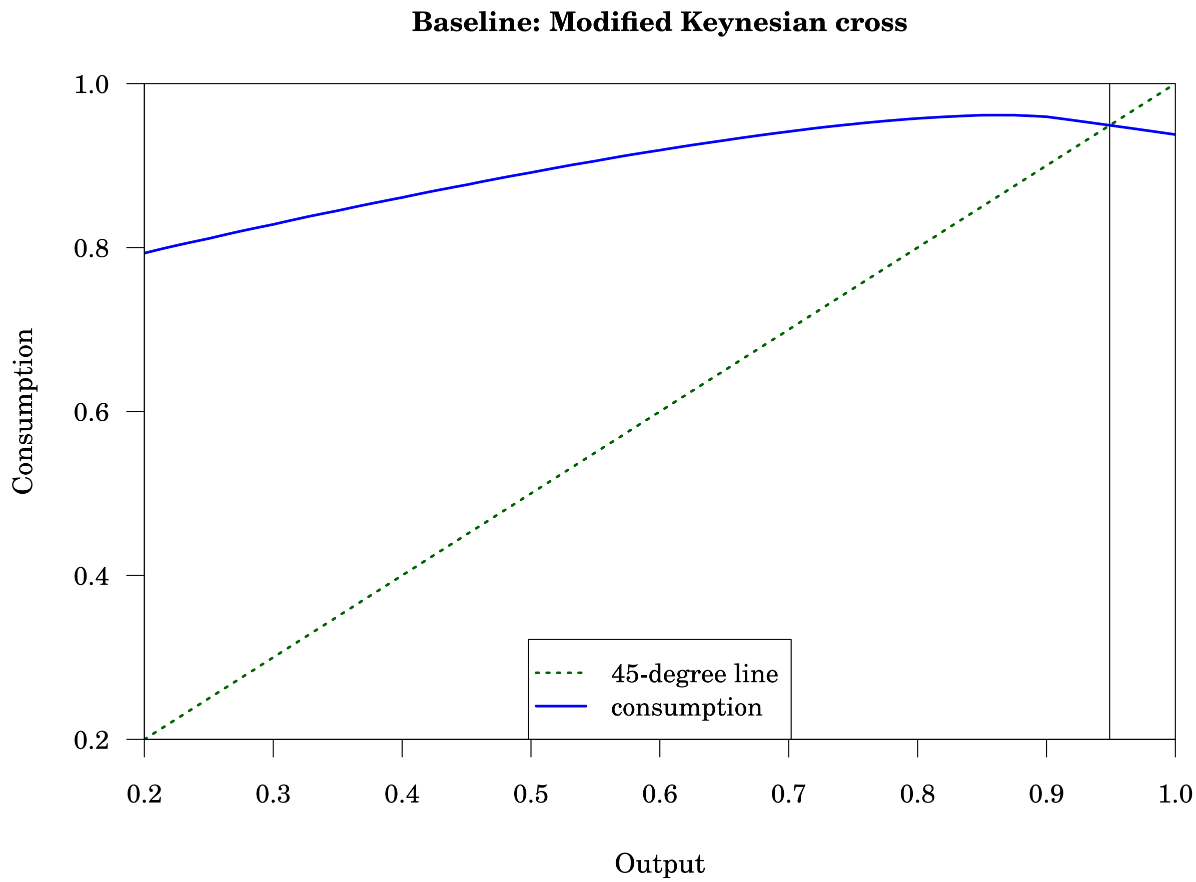

Figure 2: Modified Keynesian cross for the baseline model

5 Main experiment

Credit constraint shock

In the model presented in the previous section, the credit constraint faced by households, $\bar{b}$, had been constant over time. As an experiment, I now introduce a permanent, unforeseen shock to it. That is, at some point, all households' credit constraints will simultaneously become permanently tighter. Prior to the shock, households assign probability zero to the possibility that the credit constraint will change.

With this new feature, I have to deal with the possibility that households near the credit constraint prior the shock will find themselves in violation of it after the shock. To force such households to come into immediate compliance with the new, tighter credit constraint would force them to make a large cut to consumption. In fact, some households would have insufficient income to save to come into compliance even if they cut consumption to zero. To avoid complicating the model with the possibility that households go bankrupt, I need to permit these households to gradually come into compliance with the tighter credit constraint. Additionally, some consumer credit instruments have a term longer than one quarter, and a gradual adjustment to the credit constraint will better reflect this fact.

Guerrieri and Lorenzoni (2017) tighten the credit constraint linearly over the course of six quarters. There are seven different consumer credit constraint states: loose, tight, and five intermediate states. Since the credit constraint changes gradually, forcing households into compliance each period is not that harsh, and households never can go bankrupt. However, for computational reasons, I prefer just two consumer credit constraint states: loose and tight.

For households in violation of the new tight credit constraint, they are required to devote a fraction ϕ of their income to saving or come into immediate compliance with the new constraint. For these households with $b_{i,t-1} < \bar{b}_t$, that would choose to save $b_{i,t} < \bar{b}_t$, they are instead required to save according to the following inequality:

$$q_t b_{i,t} - b_{i,t-1} \geq \phi\left(w_t(1-\tau_t)e_{i,t} + \eta_t(1-e_{i,t})\right)$$

As such, Inequality 3.3 is replaced by a new credit constraint where households can buy and hold fewer bonds than $\bar{b}_t$. The household problem is now described by Equations 3.1, 3.2, and Inequality 5.1.

$$b_{i,t} \geq \min\left(\bar{b}_t, \frac{\phi\left(w_t(1-\tau_t)e_{i,t} + \eta_t(1-e_{i,t})\right) + b_{i,t-1}}{q_t}\right) \tag{5.1}$$

Calibration

Let $\bar{b}^L = -6.374$, the loose credit constraint from the previous section. There are two new variables to calibrate: $\bar{b}^T$, the tightened credit constraint, and ϕ, the fraction of income that a household must save if $b_{i,t}$ would be less than $\bar{b}^T$.

I am interested in the period 2008:Q2 through 2010:Q3, when consumer credit decreased. However, during this period, GDP slightly increased in 2008:Q3, but then fell and then rose, with a trough in 2009:Q2. This makes debt-to-GDP an unsatisfactory variable to use in calibration. Therefore, I calibrate $\bar{b}^T$ and ϕ using real consumer credit per capita.

I take the seasonally adjusted consumer credit series discussed earlier and convert it to real 2009 dollars using a chained GDP deflator. To find per-capita amounts, I simply divide real consumer credit by the U.S. population in the middle of the quarter. Throughout, I will call this just debt per capita.

In 2010:Q4, debt per capita increased by 4.35% (not annualized), breaking the decreasing trend it had been following from 2008:Q3 to 2010:Q3. Past that point, the series is increasing. I believe this suggests that the economy saw another aggregate shock at that point, bringing it out of the phase of declining consumer credit. I have not explored or calibrated the model for a subsequent shock. That said, a permanent tightening shock is prevalent in the literature, such as in Guerrieri and Lorenzoni (2017) and Eggertsson and Krugman (2012).

The term ϕ affects how quickly the debt per capita levels fall. The higher ϕ, the quicker debt per capita falls. In other words, as ϕ increases, the impulse response of debt per capita becomes more convex. For 2008:Q2 through 2010:Q3, though, debt per capita falls fairly linearly. As such, I choose ϕ to be fairly small: $\phi = 0.1$. Given this parameter, I choose $\bar{b}^T = -4.2$ to match the relative decrease in debt per capita from 2008:Q2 through 2010:Q3.

In the model, since the credit constraint shock is permanent, the economy will converge to some new steady-state levels of employment and debt per capita. However, this steady-state level of debt per capita will be much smaller than debt per capita in 2010:Q3. This can be seen in Figure 3. In the calibration with the parameters, $\phi = 0.1$ and $\bar{b}^T = -4.2$ (shown by the line with square markers), debt continues to decrease after 2010:Q3. Note that in this figure, I have scaled the vertical axis by dividing by 2008:Q2 GDP to make the numbers a bit more meaningful than just real consumer credit per capita.

If instead I targeted a steady-state value of debt to match the debt in 2010:Q3, then the shock would have to be minuscule. If $\bar{b}^T = -6$, then steady-state debt after the shock would match the debt

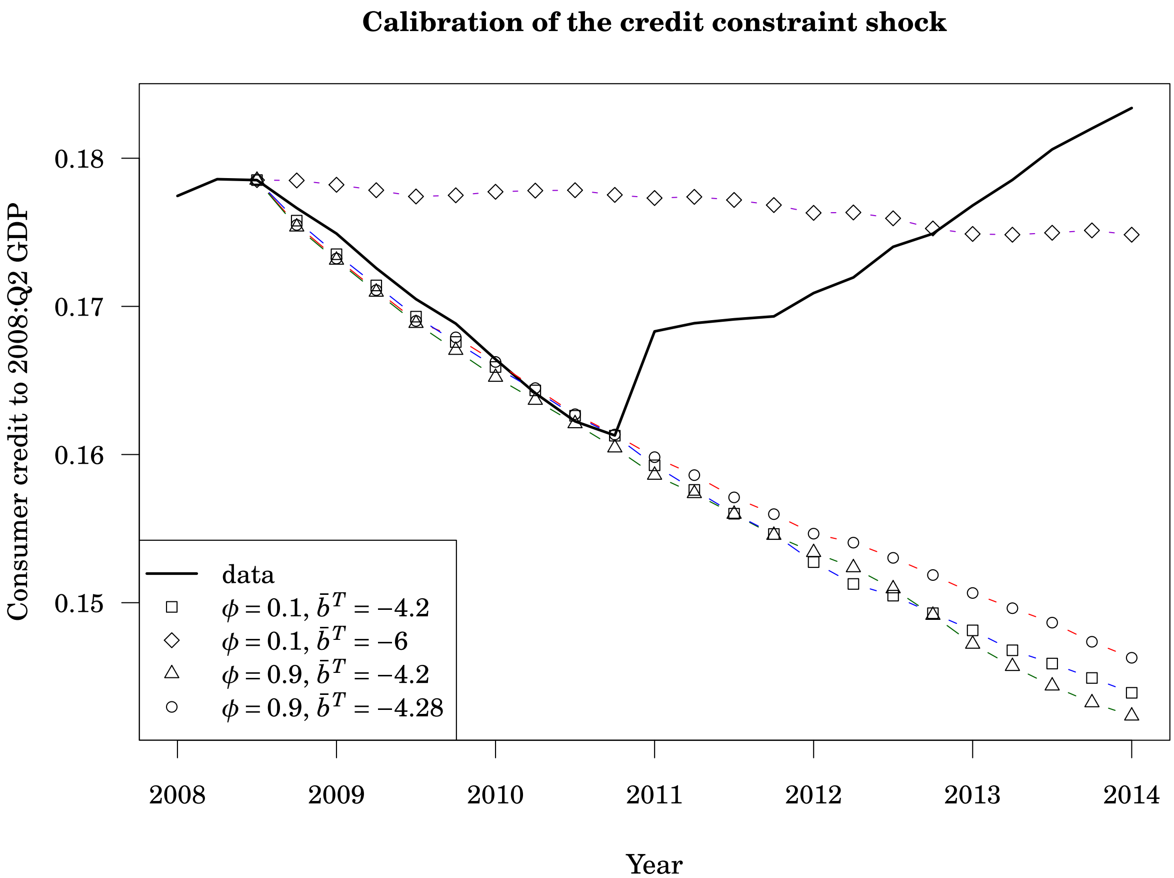

Figure 3: Illustration of the calibration of the credit constraint shock and various alternative calibrations

observed in the data in 2010:Q3. However, such a calibration cannot replicate the rapid drop in debt seen in the data. I show such a calibration with $\phi = 0.1$ and $\bar{b}^T = -6$ in the figure.

Therefore, in order to replicate the rapid decrease in debt per capita the U.S. economy experienced between 2008:Q2 and 2010:Q3, steady-state debt per capita has to be much less than 2010:Q3 debt per capita. As previously mentioned, I think there was another shock in 2010:Q4 that broke the decline in debt and changed expectations about levels of future consumer credit. I also think that the financial crisis and the aftermath was such a surprise and rarity that it is plausible that, during the crisis, people thought consumer credit was going to fall much more than it did, such that the expected steady state was below the 2010:Q3 trough.[6]

[6] Indeed, the fact that the model cannot explain the rapid decrease in consumer credit without very low steady-state consumer credit implies that, during the recession, people thought consumer credit was going to fall much more than it did.

I could consider an alternative calibration: what happens if I raise ϕ above 0.1 and choose to make the magnitude of the shock smaller, that is, choose a smaller $\bar{b}^T$? This will raise the steady-state level of debt. However, things change very little. Very few households are affected by the ϕ in the credit constraint; therefore, raising ϕ does not allow me to decrease the magnitude of the shock much. In Figure 3, I show what happens when I raise ϕ to 0.9. If I keep $\bar{b}^T = -4.2$, then I miss my target in 2010:Q3 by a little. This is shown by the line with triangle markers. If I reduce $\bar{b}^T$ to -4.28, then I hit the target. However, as can be inferred from the line with circle markers, steady-state debt per capita does not change much at all. In fact, if I select $\phi = 0.9$ and $\bar{b}^T = -4.28$ as my chosen parametrization, then that actually slightly increases the effect on unemployment; my selected parameters ($\phi = 0.1$ and $\bar{b}^T = -4.2$) are more conservative.

In conclusion, the only way to replicate such a drop in debt over just nine quarters is to have a shock of the magnitude I have calibrated. If I target a steady-state debt per capita value equal to debt per capita in 2010:Q3, then the shock must be minute and debt per capita will decrease very slowly. If I try to raise ϕ, that permits a shock of only slightly smaller magnitude and does not change the unemployment response or the steady-state debt per capita much at all.

Results

After the credit constraint shock, households want to save more, and there are two immediate effects. The direct effect is that households below the credit constraint are forced to move away from it by way of Inequality 5.1. The indirect effect is that all households are now closer to their credit constraint, which means that their buffer against unemployment has shrunk. To rebuild their buffers, all households will want to save more. This is illustrated in Figure 4.[7]

This figure illustrates a cross section of the household policy functions at an employment rate of 94.9%, or equivalently, an unemployment rate of 5.1%. For the most part, the curves shift to the

[7] There is a slight dip in the bond accumulation line for the employed post-shock for wealth levels around -4.5. This is an artifact of the solution method. During value function iteration, for these wealth levels, the program identifies that either the minimum wealth should be saved (as determined by Inequality 5.1) or something a bit more. As more value function iterations occur, households "discover" that they should save more than this minimum level of wealth, with the wealthier households discovering it first. This jump in their saving decision introduces a kink (a non-differentiable but continuous point) in the value function. A kink in the value function can cause the objective function to be bimodal in the subsequent iteration, which is ultimately responsible for this dip. I am unconcerned about this though: the number of households that are affected is very small.

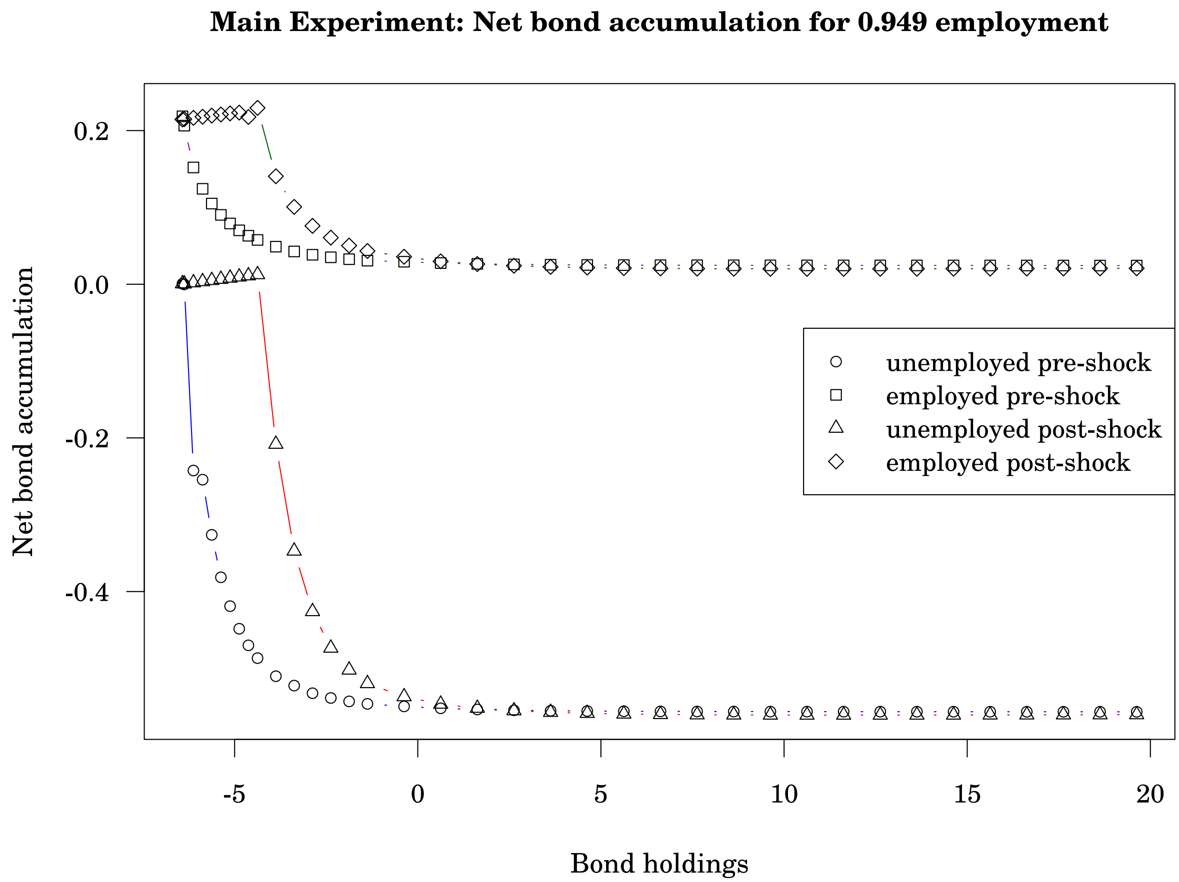

Figure 4: Net bond accumulation for 94.9% employment in the credit constraint shock model

right after the shock because the minimum level of bond holdings (the credit constraint) has shifted to the right. The increasing sections of the post-shock curves correspond to households in violation of the new credit constraint. They are required to save a fraction ϕ of their income.

One may note that in Figure 4, for high levels of bond holdings (greater than five or so), the post-shock curves are very slightly less than the pre-shock curves. This is because 94.9% employment is not a steady-state level of employment for the economy with tightened credit constraints. Households perceive that employment will fall (by way of the estimated law of motion for employment). Therefore, they understand that interest rates will fall in the future and choose to save a bit less than if they knew employment was going to stay at 94.9%.

As all households seek to save less and spend more, this pulls down the consumption demand line in the modified Keynesian cross diagram, Figure 5. The solid line, representing consumption pre-shock,

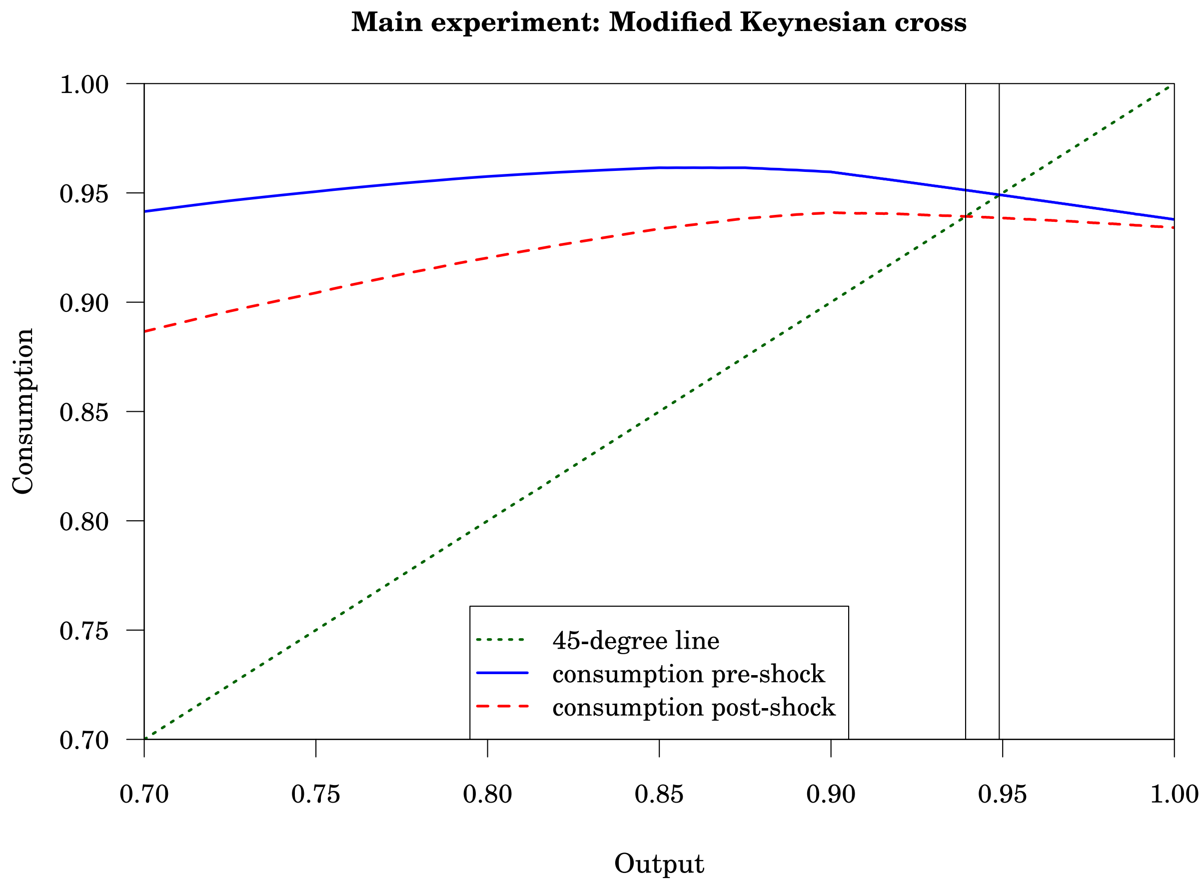

Figure 5: Modified Keynesian cross for the credit constraint shock model

is steady-state consumption demand with the loose credit constraint. The dashed line represents consumption demand for the quarter the credit constraint shock occurs. As the economy adjusts to the newer, tighter credit constraint, the consumption demand line gradually rises.

Figure 6 illustrates the effects of this permanent, unanticipated credit constraint shock on employment, debt-to-GDP, and assets-to-GDP. The initial level is marked as quarter 0, and the shock occurs in quarter 1. While this model is calibrated only to correspond to the nine-quarter drop in U.S. consumer credit, I have included 100 quarters in the graph. Even after the nine quarters have passed, employment continues to improve, and debt and asset levels continue to drop. I have provided the key numbers in Table 2.

I compare the distribution of bond holdings (wealth) for households before the credit constraint shock and nine quarters following the shock in Figure 7, which is plotted using a twenty-bin histogram. The

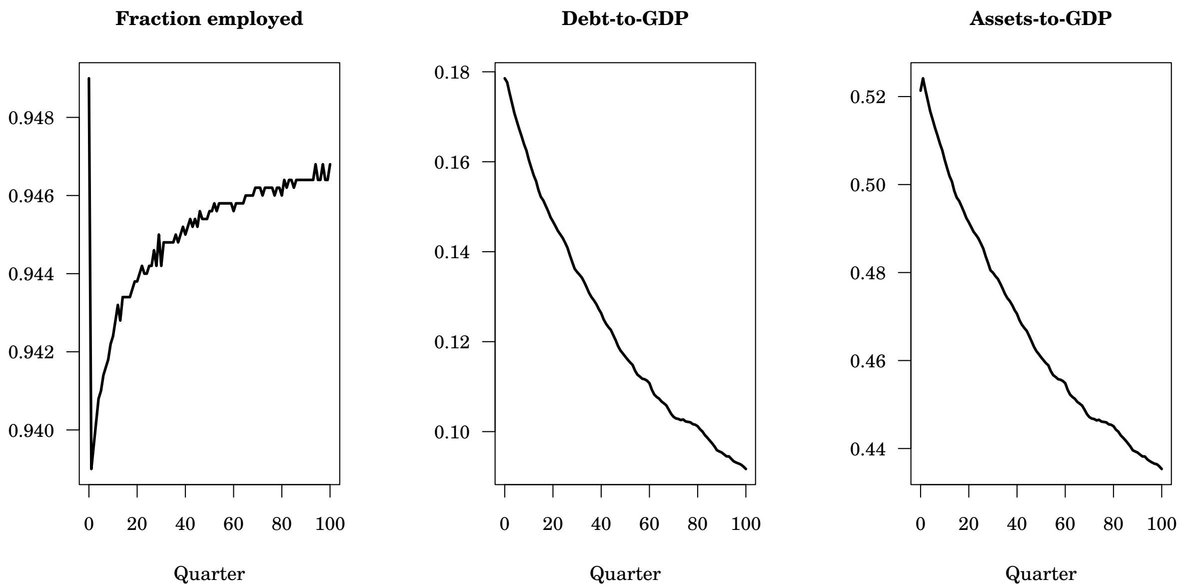

Figure 6: Impulse responses to the credit constraint shock

	Fraction employed	Interest rate	Debt-to-GDP	Assets-to-GDP
Pre-shock steady state	94.90%	2.45%	0.1785	0.5214
Post-shock 1 quarter	93.90%	1.95%	0.1777	0.6778
Post-shock 9 quarters	94.22%	2.11%	0.1625	0.5078
Post-shock steady state	94.80%	2.44%	0.0623	0.4055

Table 2: Results of the credit constraint shock model

vertical line represents the credit constraint $\bar{b}^T = -4.2$. Prior to the shock, 39.70% of households are in debt, and 1.18% of households hold less than $\bar{b}^T = -4.2$ bonds. That is, only 1.18% of households are going to be in violation of the credit constraint after the shock and will have to consider the ϕ saving constraint.

After nine quarters, 40.46% of households are in debt, and 0.04% hold less than −4.2 bonds. Additionally, the variance of the distribution decreases: households that are near the old credit constraint save more to move away from the new tighter credit constraint. Since the total number of bonds supplied by the government is fixed, this means that wealthier households reduce their bond holdings as the interest rate falls.

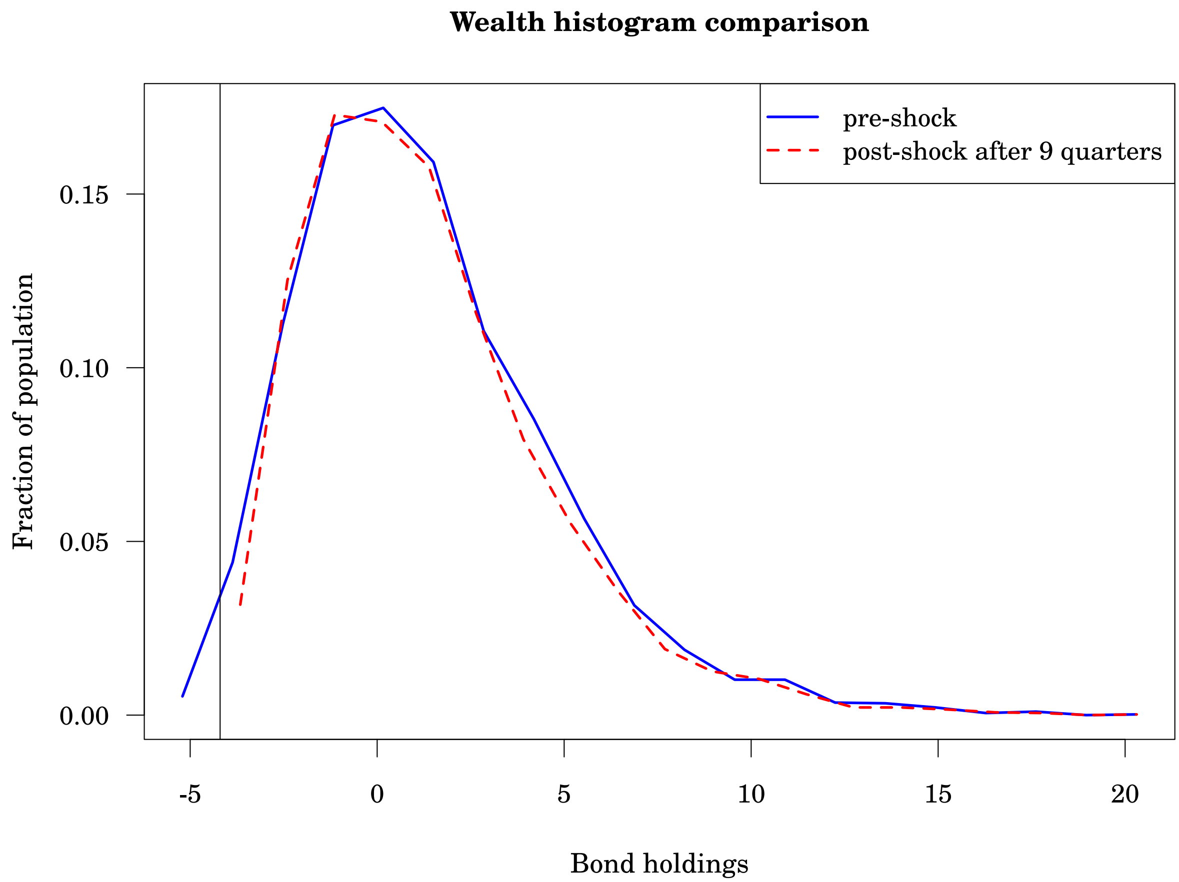

Figure 7: Wealth histograms of households before and after the credit constraint shock

6 Alternative experiments

Detrended calibration

From 2002 through 2008:Q2, real debt per capita had been increasing linearly. In the above section, I calibrated the shock's size using the difference between real debt per capita in 2010:Q3 and 2008:Q2. An alternative calibration could be considered where I calibrate the shock's size using the difference between real debt per capita in 2010:Q3 and what real debt per capita would have been in 2010:Q3 had the linear trend continued. I show this in Figure 8. I keep $\phi = 0.1$ and calibrate the new $\bar{b}^T$ to -3.9.

Naturally, a stronger credit constraint shock has a greater effect on unemployment. With the original calibration and $\bar{b}^T = -4.2$, unemployment increases 1 percentage point, from 5.10% to 6.10%. With

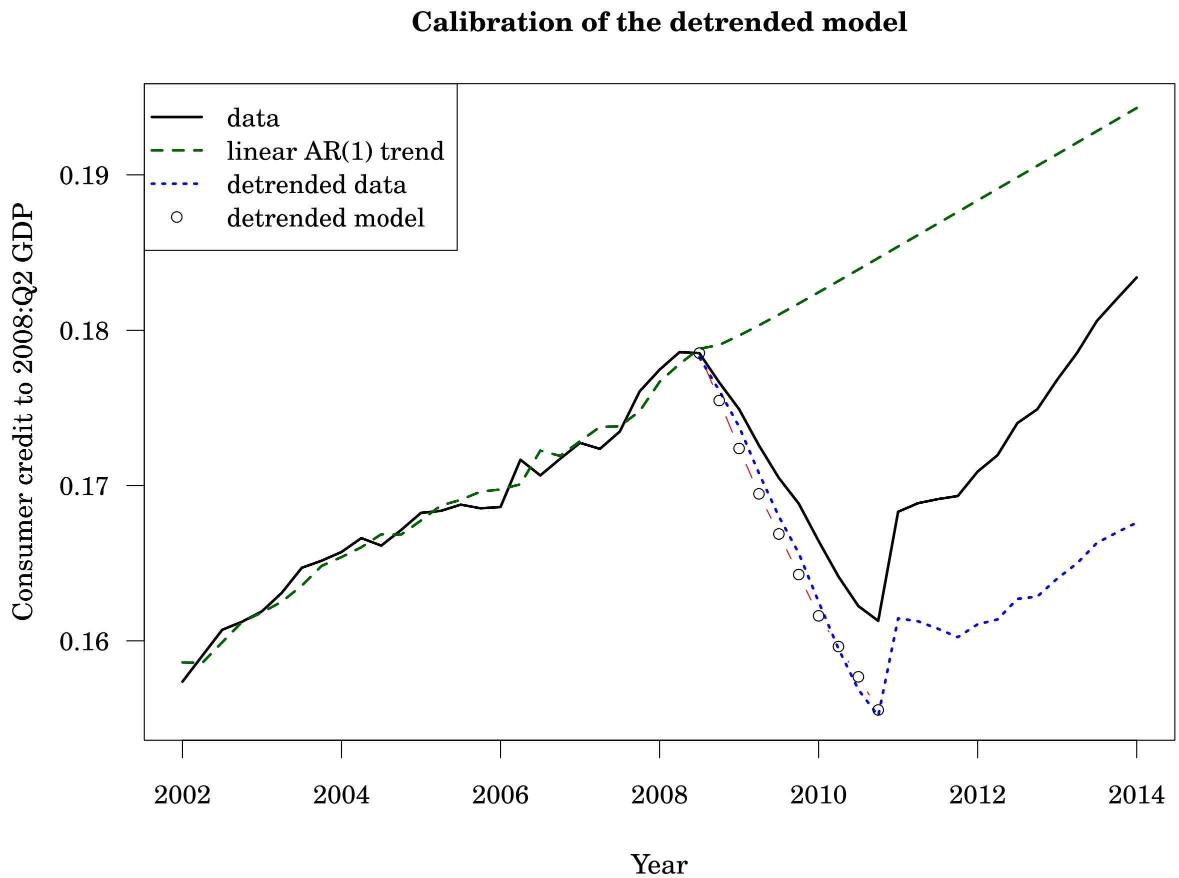

Figure 8: Illustration of the calibration of the detrended model

this detrended calibration and $\bar{b}^T = -3.9$, unemployment increases 1.38 percentage points, from 5.10% to 6.48%. The unemployment rates for this calibration are summarized in Table 4. Prior to the shock, the interest rate is 2.45%. Immediately after the shock, it falls 69 basis points to 1.76%.

Less responsive interest rate experiment

For a nominal interest rate rule, a typical coefficient for the output gap is around $\psi = 0.5$, which is what I studied in the previous sections. However, a typical coefficient for the employment gap is around $\psi = 0.25$. In my model, employment and output are linearly related, so either ψ is worthy of study as a plausible coefficient. For this experiment, I keep $L^* = 0.95$ and $r^* = 0.025$ unchanged from the baseline parameters. Then, I set ψ to 0.25 and recalibrate government bonds, B; the loose credit constraint, $\bar{b}^L$; and the tight credit constraint, $\bar{b}^T$. Table 3 provides the parameter values for

Name	Variable	Value
Interest rate response to employment gap	ψ	0.25
Permanent government debt	B	1.2998
Household loose credit constraint	$\bar{b}^L$	−6.435
Household tight credit constraint	$\bar{b}^T$	−4.25

Table 3: Parameter values for the calibration with $\psi = 0.25$

this new calibration.

In the previous calibration, where $\psi = 0.5$, before the shock, the steady-state employment rate is 94.9%, and the steady-state interest rate is 2.45% by way of Equation 3.6. In this calibration, the steady-state employment rate and interest rate are going to be a bit less. Suppose the steady-state interest rate when $\psi = 0.25$ was 2.45%. Then that would imply the steady-state employment rate should be 94.8% by Equation 3.6. However, this employment rate is slightly lower than the 94.9%, which means there is a slightly greater precautionary motive to save for this $\psi = 0.25$ calibration in steady state. Therefore, it should be expected that the steady-state interest rate will be slightly less than 2.45% because of the increased precautionary motive. Consequently, the employment rate will be slightly less than 94.8%; it will be 94.78%.

In this calibration, the pre-shock steady-state unemployment rate is 5.22%. Following the credit constraint shock, unemployment rises 1.38 percentage points to 6.60%. The interest rate falls 34.5 basis points, from 2.445% to 2.1%. Table 4 summarizes the unemployment rates for this experiment.

Unresponsive interest rate experiment

I have demonstrated the effect of changing ψ from 0.5 to 0.25 on unemployment. If the interest rate is less responsive to changes in unemployment, then the effect of the credit constraint shock on unemployment is greater. What is the effect on unemployment if the interest rate does not or can not fall at all after the shock? For example, what if the central bank's only mandate was to control inflation? If full employment is not a goal for the central bank, then $\psi = 0$, and the interest rate would remain unchanged after an increase in unemployment.[8]

[8] Another justification for $\psi = 0$ could be that the central bank suffers observation, decision, and implementation policy lags, and thus it does not respond contemporaneously to changes in the unemployment rate. However, this interpretation would only apply to the first quarter or two of this experiment.

Model	Unemployment rates			
	Pre-shock steady state	Post-shock, 1 quarter	Post-shock, 9 quarters	Post-shock steady state
Main experiment: $\bar{b}^T = -4.2$ and $\psi = 0.5$	5.10%	6.10%	5.78%	5.20%
Detrended calibration: $\bar{b}^T = -3.9$	5.10%	6.48%	6.04%	5.28%
Less responsive interest rate: $\psi = 0.25$	5.22%	6.60%	6.20%	5.54%
Unresponsive interest rate: rate floor = 2.45%	5.10%	10.46%	9.76%	8.63%

Table 4: Unemployment rates for select times for all models

Alternatively, consider an economy at steady state with 94.9% employment, matching the original $\psi = 0.5$ calibration in Section 5. The steady-state interest rate is 2.45% (see Equation 3.6). Next, suppose there is some unspecified shock to the economy that causes the nominal interest rate to fall to zero or some value near zero. Now the nominal interest rate cannot fall any further. The real rate cannot fall any further either, assuming inflation is constant. Suppose a credit constraint shock hit the economy in this state. The interest rate would be unresponsive to any subsequent increase in unemployment.

I am not going to model this unspecified shock that puts the economy at or near the zero lower bound on the nominal interest rate. However, as a proxy for this situation, I consider an interest rate floor at 2.45%. This will change the interest rate rule to Equation 6.1.

$$r_t = \max\left(r^* + \psi(L_t - L^*), 0.0245\right) \tag{6.1}$$

As a final note, I would suggest that the zero lower bound on the nominal interest rate applies to riskless overnight lending. For loans of a longer duration, like consumer credit instruments, if the yield curve is normal, a term premium is applied to the interest rate. Furthermore, this is a model of consumer credit. While default risk is not explicitly in my model, in reality, consumer credit is subject to default risk. Creditors would demand a default risk premium applied to the interest rate. So these two premia, when added to the zero lower bound for riskless overnight lending, could reasonably lead to a positive lower bound on the interest rate for consumer credit.

I wish to keep things comparable with the $\psi = 0.5$ calibration, where the steady-state interest rate is 2.45%. If I keep the tightened credit constraint the same at $\bar{b}^T = -4.2$, debt per capita falls a bit more than it did before. However, it misses its target by only 1.76%. Smaller values (more negative values) for $\bar{b}^T$, do not change the results much. Thus, for comparison purposes, I keep the tightened credit constraint the same as in Section 5: $\bar{b}^T = -4.2$.

For computational purposes, when $\bar{b}_t = \bar{b}^L$, the model uses Equation 3.6, with the interest rate floor at 0%. When $\bar{b}_t = \bar{b}^T$, the model uses Equation 6.1, with the interest rate floor at 2.45%. If there was an infinite number of households in my program, then I could use the 2.45% interest rate floor for all $\bar{b}_t$. However, the number of households I simulate is finite, and the exact fraction employed at any given period can vary slightly, or jitter, around the steady-state value for the loose credit constraint state.

If the interest rate can fall below 2.45% in the loose credit constraint state, then the jittering might make the employment rate for that period 94.88%. It would then recover to 94.9% in subsequent periods. However, if the interest rate has a floor at 2.45% in the loose credit constraint state, then the jittering might make the employment rate for that period 94.68%. Since the interest rate rule is flat in that region, the employment rate will not return to 94.9%.

It might be helpful to think of the jittering as miniature demand shocks. As will be demonstrated, the effect of the demand shock on the employment rate increases in magnitude when the interest rate is constant.

Additionally, I have tested the model with an interest rate floor at 2.4% for all $\bar{b}_t$ versus the model with an interest rate floor at 0% when $\bar{b}_t = \bar{b}^L$and 2.45% when $\bar{b}_t = \bar{b}^T$. The formulations behave absolutely identically. This is because in both formulations, the interest rate will fall in response to the negative miniature demand shock, and thus the employment rate only falls to 94.88%.

Thus, to keep this model's pre-shock employment rate comparable with the previous model's, I have the interest rate floor at 2.45% only after the shock.

Table 4 shows the unemployment rates for this experiment in which the interest rate does not fall below 2.45%. As can be seen, the effect on unemployment is greatly amplified. Whereas in Section 5, the unemployment rate rises from 5.1% to 6.1%, for this experiment, it jumps from 5.1% to 10.46% after the shock.

7 Conclusion

I have examined this model in four different parametrizations, summarized in Table 4. In the first three parametrizations, the unemployment rate increases 1 to 1.38 percentage points in response to the credit constraint shock. In the fourth parametrization, the unemployment rate increases 5.36 percentage points. To frame these numbers, during the Great Recession, the U.S. unemployment rate rose from about 5.3% to a peak of 10%, which corresponds to a 4.7 percentage point increase.

In 2008:Q2, consumer credit was only 17.9% of U.S. household liabilities; home loans make up most of household liabilities. However, this model is calibrated using consumer credit alone. Since consumer credit is a small fraction of household liabilities, a 1 to 1.38 percentage point increase in the unemployment rate is fairly reasonable if the interest rate is able to adjust. Since the unemployment rate rose by 4.7 percentage points in the data, consumer credit alone does not explain the observed increase in unemployment.

For further evidence, consider the fall in interest rates during the recession. In August 2008, the U.S. federal funds rate was 2.00%. At its lowest during the recession, the federal funds rate was 0.11%. Therefore, before the shock, the interest rate had at least 1.89 percentage points worth of room to fall. Of the three parametrizations where the interest rate could fall, the interest rate fell the most in the detrended calibration. In that calibration, the interest rate fell 0.69 percentage points. If the U.S. was subject to only a consumer credit shock, then the interest rate would have had enough room to adjust without reaching the zero lower bound.

If the interest rate cannot decrease after a credit constraint shock, then my model predicts a large increase in the unemployment rate: 5.36 percentage points. This increase in unemployment is much closer to what was observed during the recession. I conclude that there must have been some other factor that pushed the nominal interest rate down to zero or near zero. According to Hall (2011), inflation was nearly exogenous. Thus, neither the nominal nor real interest rate would have been able to fall further. With the interest rate unable to respond, the shock to consumer credit can explain the observed increase in the unemployment rate. Of course, the order of the two shocks does not matter: the unspecified shock could have occurred after the consumer credit shock or concurrent with it.

One of the causes that could be responsible for this unspecified shock is the asset price bubble. Households lost a lot of net worth in their homes and retirement accounts. Given this loss of wealth, households wanted to save more and spend (consume) less, due to a precautionary motive to save against the risk of unemployment and a desire to rebuild their portfolios for retirement. The subsequent effects would continue from there as I have explained above. My model could be extended to incorporate households owning homes which are subject to price bubbles.

Additionally, I can extend the model in two other ways. First, the distribution of wealth in my model is a bit too uniform. One way to introduce more variance in the current version of the model would be to decrease the degree of unemployment insurance ρ. Another way would be to make the employed more likely to keep jobs and the unemployed less likely to find them by decreasing the probability of separation λ. However, to alter these variables would run counter to Shimer's observations (2005). Thus, I am interested in following the method used by Krusell and Smith (1998) to change the shape of their wealth distribution. I will give households idiosyncratic discount rates, β's, which will evolve according to a first-order Markov process. Then by altering the transition probabilities and discount rates, I can match the wealth distribution to what is observed in the data.

Another extension is to give households an age. Households would be available to work for 45 years. After that, they would retire for 20 years, during which they would be unable to work. This would cause households to save not only for precautionary reasons but also for retirement. Calibrating the model with this extension, I could include stocks, corporate and government bonds, etc. as household assets. Since households would have another reason to save, shocking their credit constraint could result in a stronger effect on unemployment. This extension would also allow me to study the disparate impact of the credit constraint shock on households far from retirement, households near retirement, and households in retirement.

References

Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk. 2014. “Some Evidence on the Importance of Sticky Wages.” *American Economic Journal: Macroeconomics,* 6(1): 70–101.

Bewley, Truman. 1977. “The Permanent Income Hypothesis: A Theoretical Formulation.” *Journal of Economic Theory,* 16(2): 252–292.

Eggertsson, Gauti B. and Paul Krugman. 2012. “Debt, Deleveraging, and the Liquidity Trap: A Fisher–Minsky–Koo Approach.” *Quarterly Journal of Economics,* 127(3): 1469–1513.

Guerrieri, Veronica and Guido Lorenzoni. 2017. “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *Quarterly Journal of Economics,* 132(3): 1427–1467.

Hall, Robert E. 2011. “The Long Slump.” *American Economic Review,* 101(2): 431–469.

Krusell, Per and Anthony A. Smith, Jr. 1998. “Income and Wealth Heterogeneity in the Macroeconomy.” *Journal of Political Economy,* 106(5): 867–896.

Michaillat, Pascal. 2012. “Do Matching Frictions Explain Unemployment? Not in Bad Times.” *American Economic Review,* 102(4): 1721–1750.

Ravn, Morten O. and Vincent Sterk. 2017. “Job Uncertainty and Deep Recessions.” *Journal of Monetary Economics,* 90: 125–141.

Schmitt-Grohé, Stephanie and Martín Uribe. 2012. “The Making of a Great Contraction with a Liquidity Trap and a Jobless Recovery.” NBER Working Paper.

Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” *American Economic Review,* 95(1): 25–49.