NORTHWESTERN UNIVERSITY

Heterogeneous Agent Dynamics across the Business Cycle

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Brian Boyd O'Quinn

EVANSTON, ILLINOIS

December 2017

Abstract

Heterogeneous Agent Dynamics across the Business Cycle

Brian Boyd O'Quinn

This dissertation consists of two papers united by a common element: they both study the behavior of heterogeneous agents across the business cycle.

In Chapter 1, I consider: 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.

In Chapter 2, I address the question: what is the effect of plant entry and exit on productivity throughout the business cycle? According to Schumpeter's theory of creative destruction, recessions should cleanse the economy of unproductive plants. I also consider the hypothesis that economic booms should force less productive plants to close due to increased competition for

inputs. Using plant-level data from Chile, 1979–96, I estimate productivity using two contemporary methods and develop metrics to isolate the change in average productivity due solely to plant entry and exit. The results support both propositions. I find that entry–exit behavior during a recession improved productivity by 2.4 percentage points per year over periods of moderate economic growth. Similarly, entry–exit behavior during economic booms improved productivity by 1.9 percentage points per year over periods of moderate economic growth.

Acknowledgements

I am grateful to Giorgio Primiceri for his counsel and steadfast support throughout my studies.

I appreciate the valuable suggestions and comments of Guido Lorenzoni and Larry Christiano regarding the first chapter of this dissertation. I also thank Mirko Wiederholt and Jonathan Parker for their guidance in the early stages of that work.

I am grateful for the advice of Gaston Illanes on the development of the second chapter. The data set used in that chapter was supplied, along with the basic code for their estimator, by Amit Gandhi, Salvador Navarro, and David Rivers. They received the data set from David Greenstreet, so he should be credited for that. For my early work on the second chapter, I used a data set given to me by Aviv Nevo.

I want to recognize Northwestern University for providing me ample computational capability. In particular, I found the Social Sciences Computing Cluster to be an excellent resource.

Lastly, I thank my sisters: Shannon O'Quinn, who proofread the dissertation, and Erin O'Quinn, who formatted two of the figures.

Dedication

To my parents—Guy and Spring O'Quinn—in gratitude for their boundless love

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Unemployment and Credit Constraints in a Heterogeneous Agent Model

1.1. 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 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.

[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.

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 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 1.3. In Section 1.4, I calibrate most of the parameters for the model and examine the baseline case where the credit constraint is constant over time. Section 1.5 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 1.6, 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 1.7 concludes.

1.2. 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.

1.3. 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 1.3.2.

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.

$$V(e_{i,t}, b_{i,t-1}; \Gamma_t) = \max_{c_{i,t}, b_{i,t}} u(c_{i,t}) + \beta E_t V(e_{i,t+1}, b_{i,t}; \Gamma_{t+1}) \tag{1.3.1}$$

subject to:

$$p_t c_{i,t} + q_t b_{i,t} = w_t(1-\tau_t)e_{i,t} + \eta_t(1-e_{i,t}) + b_{i,t-1} \tag{1.3.2}$$

$$b_{i,t} \geq \bar{b}_t \tag{1.3.3}$$

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.

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

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

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).

$$\rho = \frac{\eta_t}{w(1-\tau_t)} \tag{1.3.4}$$

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.

$$w\tau_t L_t = \eta_t(1-L_t) + (1-q_t)B \tag{1.3.5}$$

By way of Equations 1.3.4 and 1.3.5, 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.

$$r_t = \max\left(r^* + \psi(L_t - L^*), 0\right) \tag{1.3.6}$$

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 $\Gamma_t = \Gamma_{t+1}$.

[2] This condition is mathematically implied by the two other market-clearing conditions and Equation 1.3.2. It is included for the sake of completeness.

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 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.

1.4. Baseline calibration and results

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

Calibration

Table 1.1 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

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 1.1. Baseline parameter values

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 1.3.6, steady-state employment will

[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.

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 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.

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 1.1 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

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

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

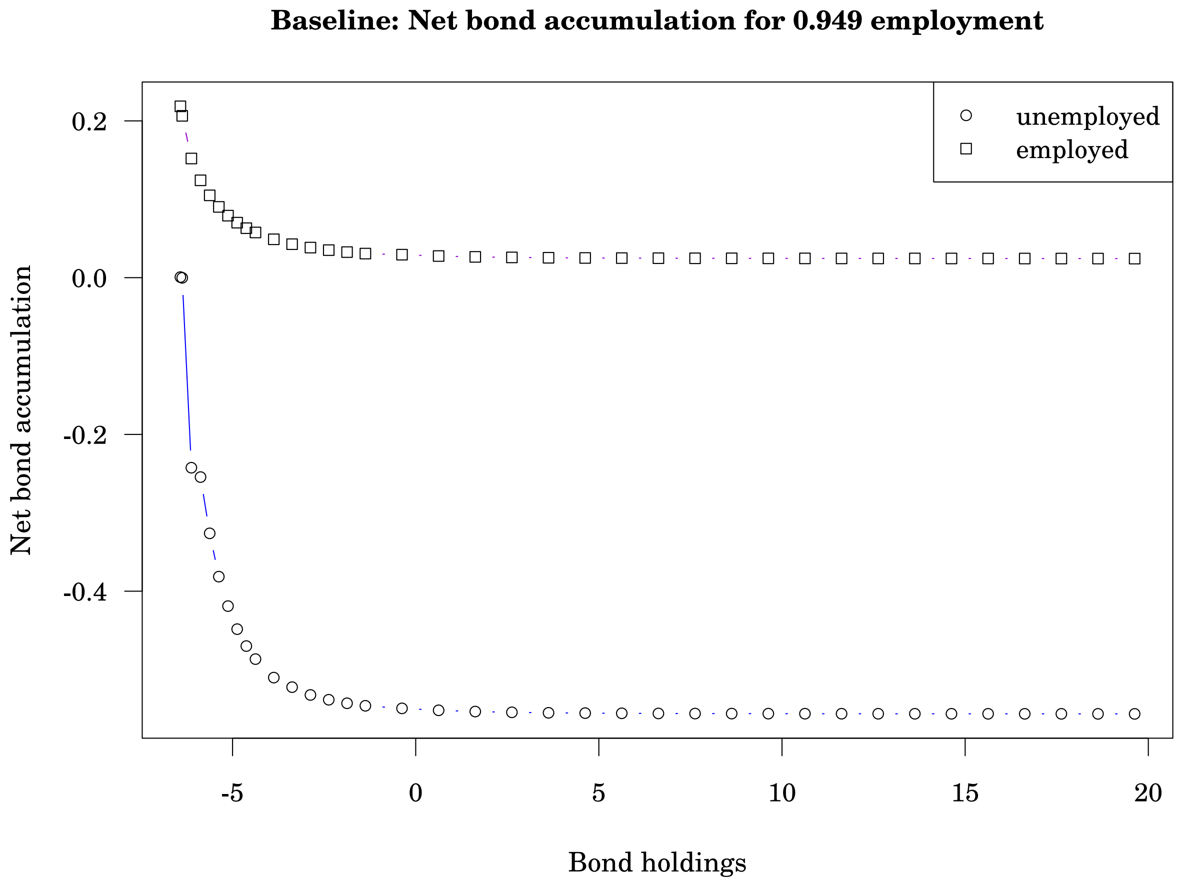

Figure 1.1. Net bond accumulation for 94.9% employment in the baseline model

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.

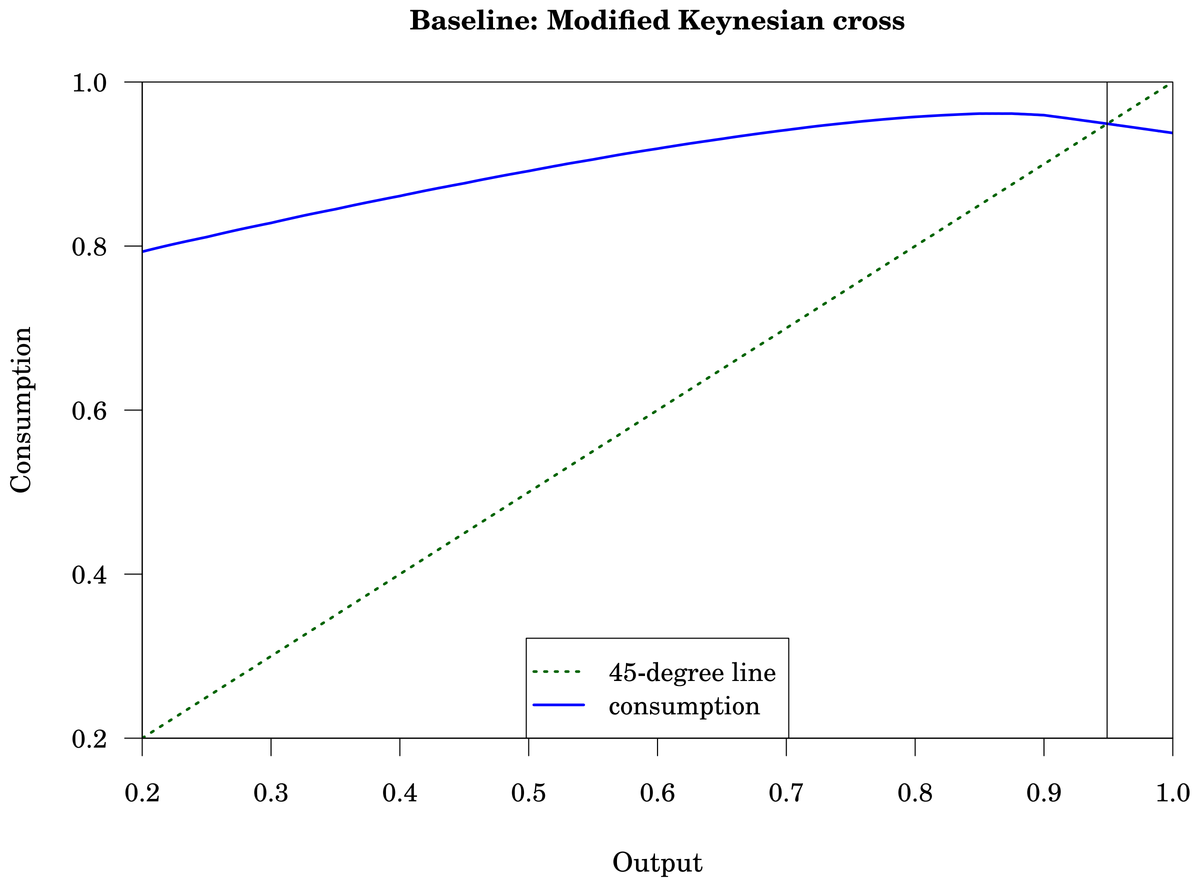

Figure 1.2. Modified Keynesian cross for the baseline model

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 1.2 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.

1.5. 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 1.3.3 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 1.3.1, 1.3.2, and Inequality 1.5.1.

$$b_{i,t} \geq \min\left(\bar{b}_t, \frac{\phi\left(w_t(1-\tau_t)e_{i,t} + \eta_t(1-e_{i,t})\right) + b_{i,t-1}}{q_t}\right) \tag{1.5.1}$$

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 1.3. 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.

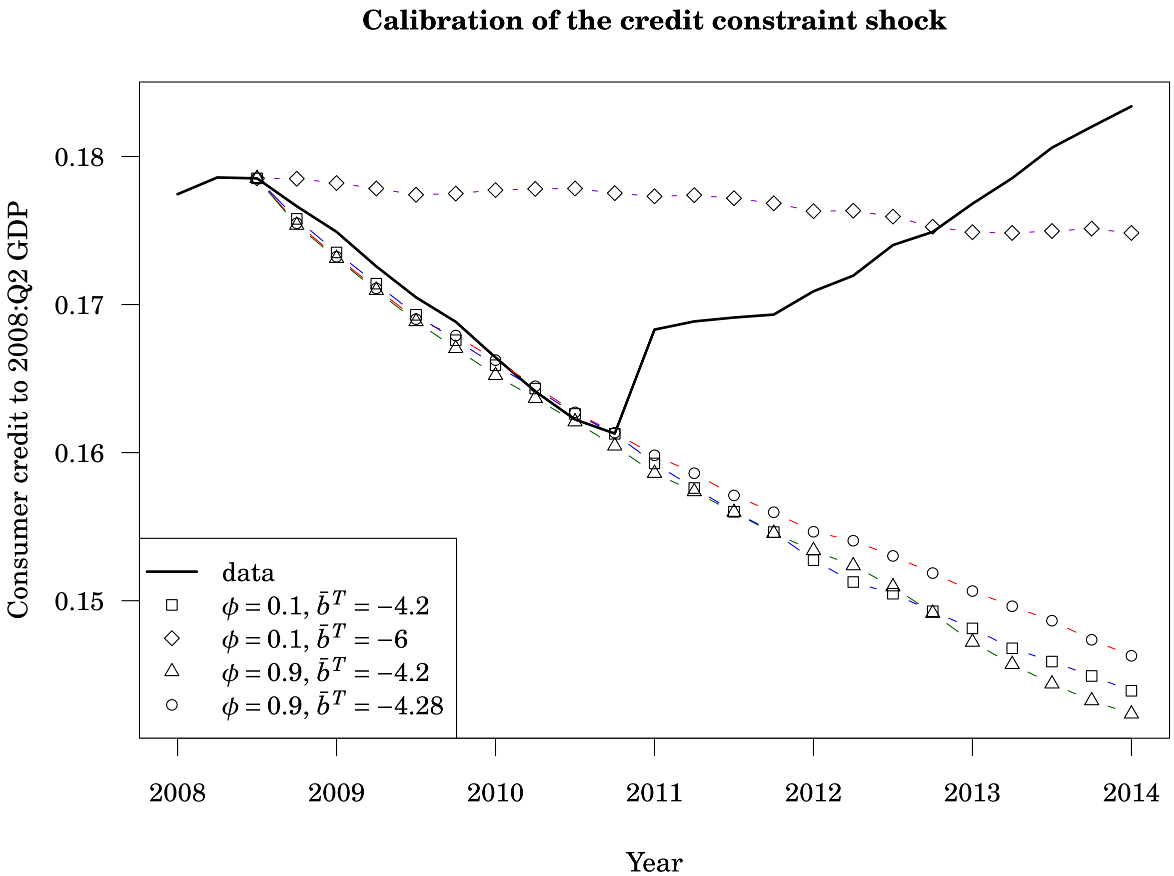

Figure 1.3. Illustration of the calibration of the credit constraint shock and various alternative calibrations

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 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]

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 1.3, 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 1.5.1. The indirect effect is that all households are now closer to their

[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.

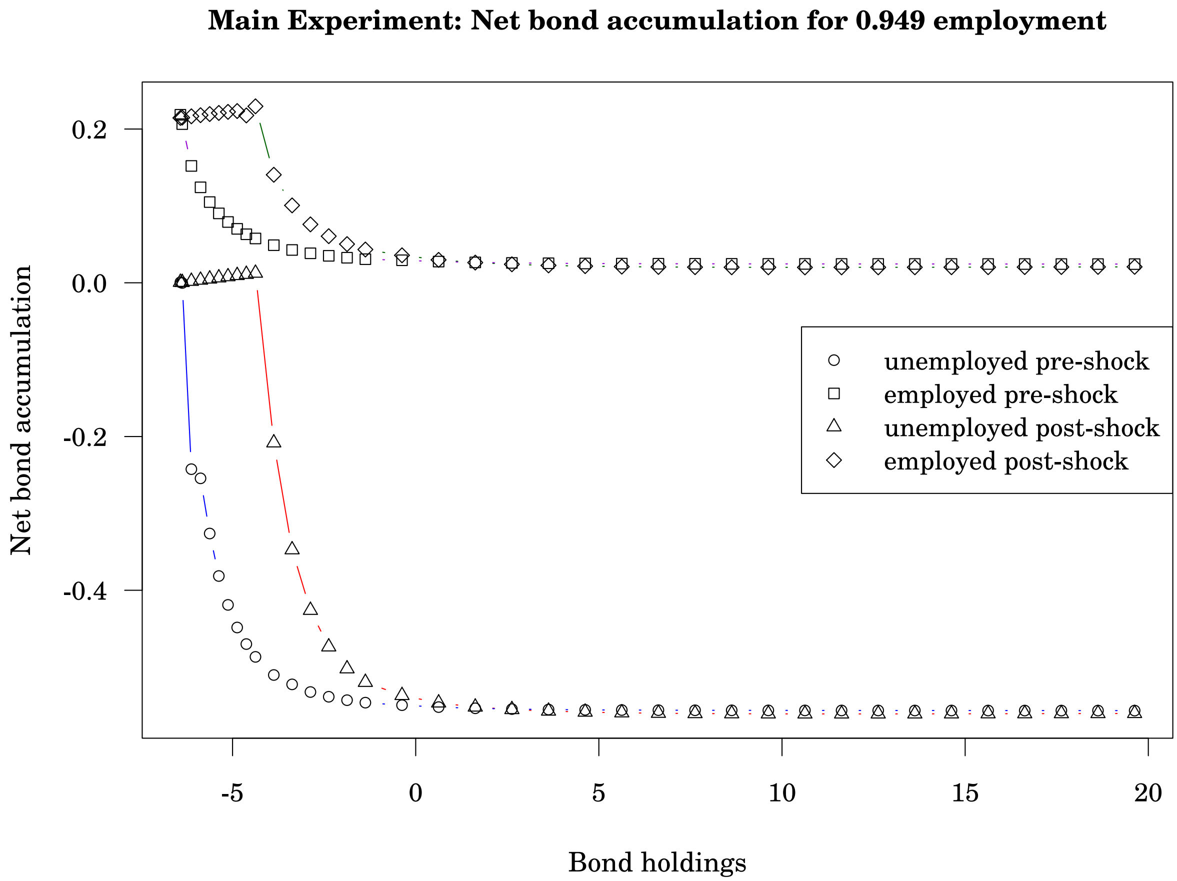

Figure 1.4. Net bond accumulation for 94.9% employment in the credit constraint shock model

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 1.4.[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

[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 1.5.1) 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.

the 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 1.4, 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 1.5. The solid line, representing consumption pre-shock, 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 1.6 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 1.2.

I compare the distribution of bond holdings (wealth) for households before the credit constraint shock and nine quarters following the shock in Figure 1.7, which is plotted using a twenty-bin histogram. The 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

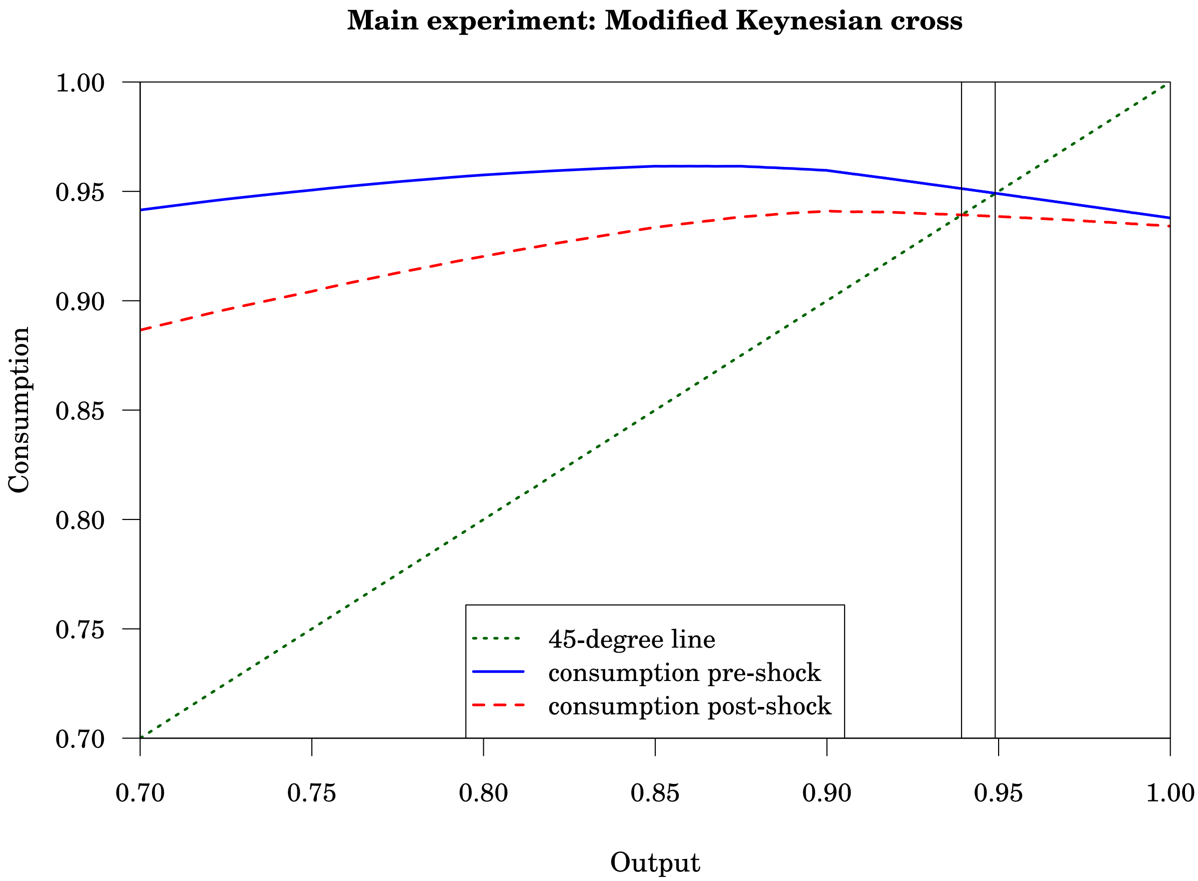

Figure 1.5. Modified Keynesian cross for the credit constraint shock model

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.

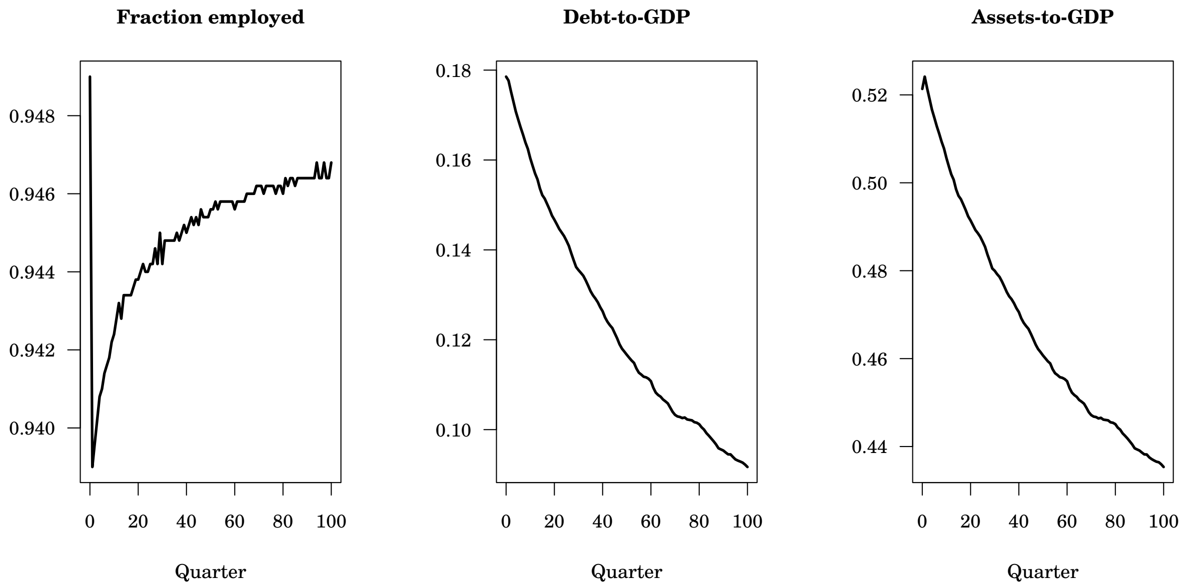

Figure 1.6. Impulse responses to the credit constraint shock

	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 1.2. Results of the credit constraint shock model

1.6. 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

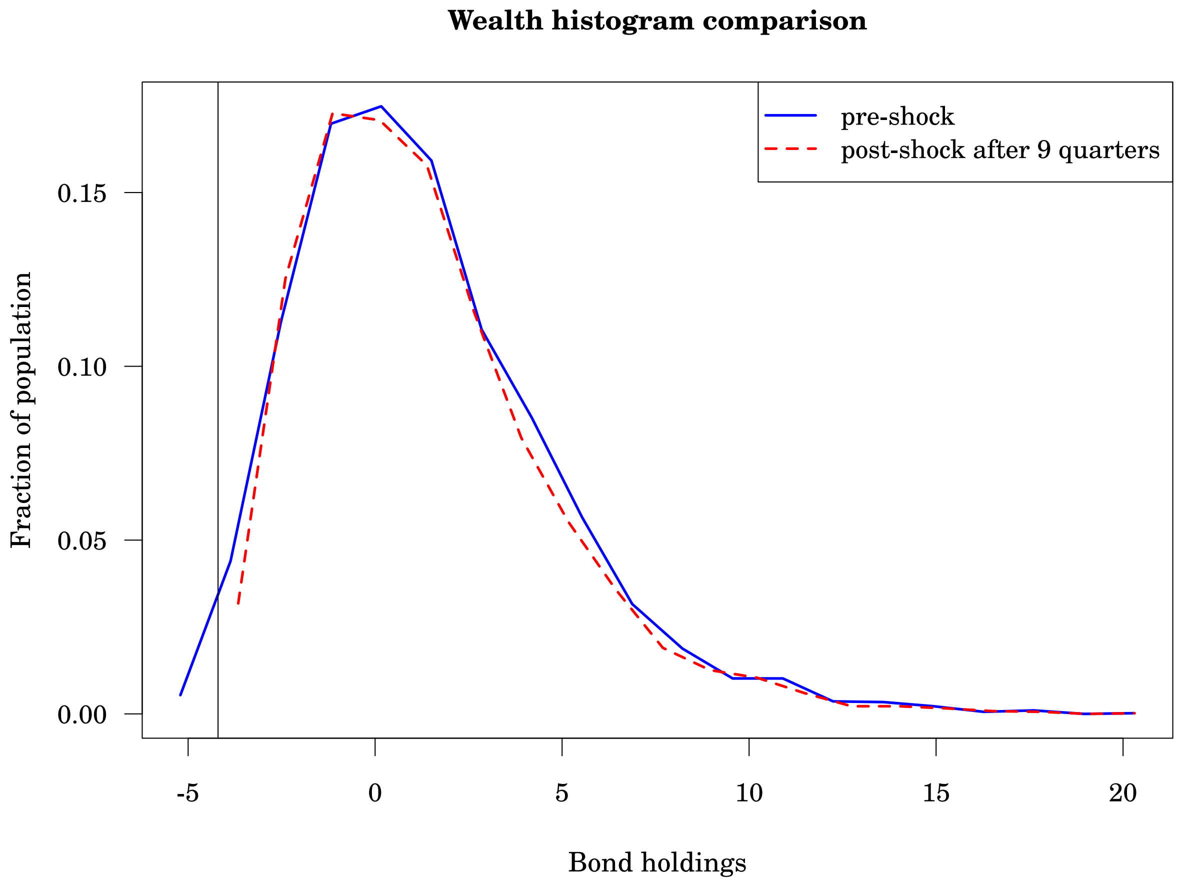

Figure 1.7. Wealth histograms of households before and after the credit constraint shock

have been in 2010:Q3 had the linear trend continued. I show this in Figure 1.8. 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 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 1.4. Prior to the shock, the interest rate is 2.45%. Immediately after the shock, it falls 69 basis points to 1.76%.

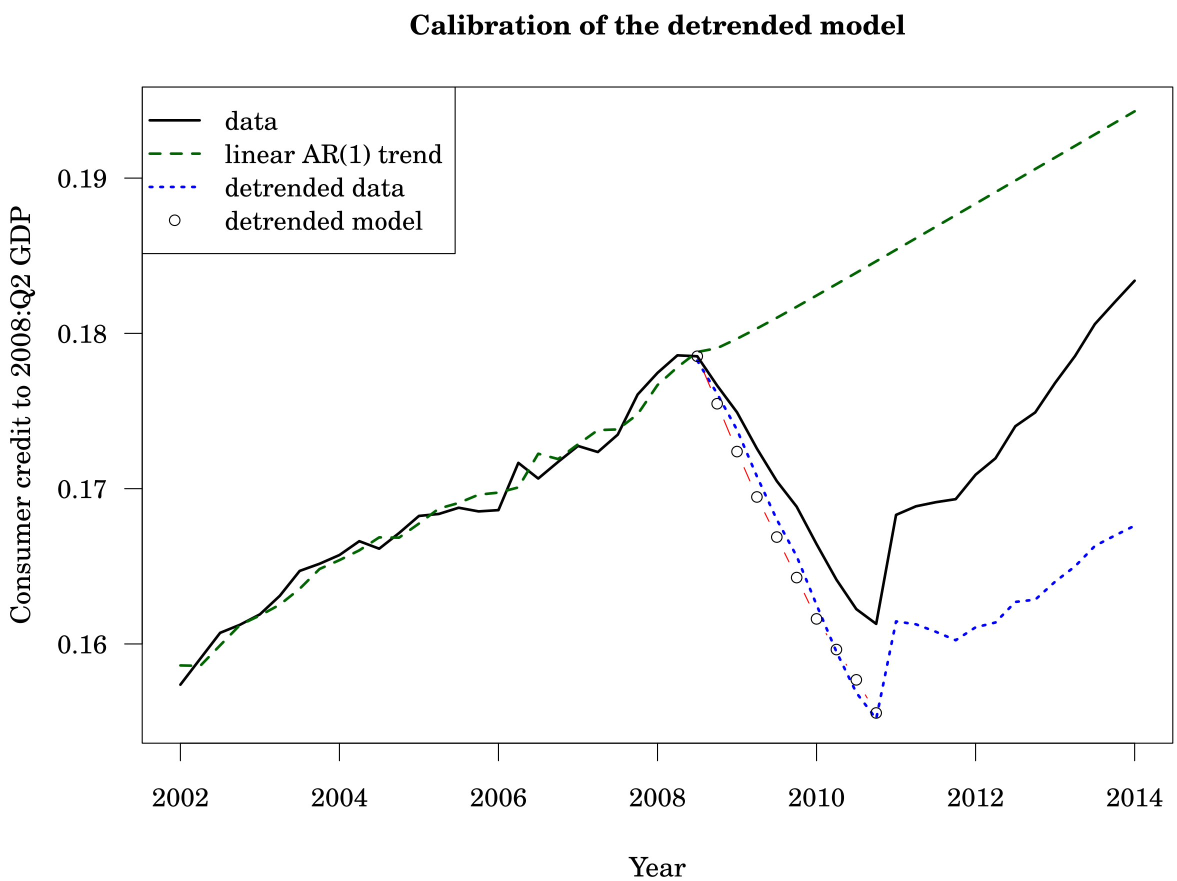

Figure 1.8. Illustration of the calibration of the detrended model

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 1.3 provides the parameter values for this new calibration.

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 1.3. Parameter values for the calibration with $\psi = 0.25$

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 1.3.6. 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 1.3.6. 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 1.4 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

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 1.4. Unemployment rates for select times for all models

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]

Alternatively, consider an economy at steady state with 94.9% employment, matching the original $\psi = 0.5$ calibration in Section 1.5. The steady-state interest rate is 2.45% (see Equation 1.3.6). 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 1.6.1.

$$r_t = \max\left(r^* + \psi(L_t - L^*), 0.0245\right) \tag{1.6.1}$$

[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.

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 1.5: $\bar{b}^T = -4.2$.

For computational purposes, when $\bar{b}_t = \bar{b}^L$, the model uses Equation 1.3.6, with the interest rate floor at 0%. When $\bar{b}_t = \bar{b}^T$, the model uses Equation 1.6.1, 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 1.4 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 1.5, the unemployment rate rises from 5.1% to 6.1%, for this experiment, it jumps from 5.1% to 10.46% after the shock.

1.7. Conclusion

I have examined this model in four different parametrizations, summarized in Table 1.4. 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.

CHAPTER 2

The Effect of Plant Entry and Exit on Productivity across the Business Cycle

2.1. Introduction

Prior to Keynes, economists did not seek to alleviate recessions because they were thought to have an important function: to cleanse the economy of inefficiency. Of economists that held this view, Schumpeter advanced it most famously, and it is encapsulated by his concept of "creative destruction." This paper seeks to study the cleansing effect of recessions in the particular context that relatively unproductive plants will cease to operate, or exit, during a recession.[1] This paper asks the question: how much is average productivity improved by the exit of these inefficient plants? If it is significant, it could have policy implications regarding fiscal stimulus, corporate bailouts, and protectionism.

Melitz (2003) builds a model which predicts the opposite of Schumpeter's view of recessions. The model implies that, during an economic boom, there is increased competition for scarce inputs/factors, and only the most productive plants will survive. Recessions are therefore "sullying:" unproductive plants can enter because input prices are low. Between these two opposing views, the key is competitive pressure. Schumpeter's theory would suggest competitive pressure is higher during a recession, due to low demand for output in a demand-shock recession, or high prices for inputs in a supply-shock recession. Melitz's model would suggest it is higher during a boom, due to high demand for inputs.

[1]A recession could also remove inefficient production lines within a plant or result in a poor manager being replaced, among other things. However, this paper focuses on plant exit.

Kehrig (2011) puts these two theories in opposition and, using U.S. data, concludes that Melitz's view is correct and Schumpeter's is not. However, I suggest that these theories need not be completely at odds. Perhaps both economic booms and recessions are cleansing, and only periods of moderate economic growth are sullying. Therefore, this paper also considers the proposition that booms are cleansing and assesses the improvement in average productivity caused by plants exiting during a boom.

These theories naturally apply to plant entry as well. For example, in Schumpeter's view, only highly productive plants would enter during a recession. Thus, this paper additionally evaluates the effect of plant entry on productivity during recessions and booms.

I employ a common data set used in the production function literature: plant-level data from Chile for the years 1979–96. During that period, Chile experienced a recession in 1982 and 1983. To estimate productivity, this paper uses two modern production function estimation methods. The primary method follows Ackerberg, Caves, and Frazer (2015) with an intermediate step to correct for selection. To check the robustness of the results, I use the estimation technique developed by Gandhi, Navarro, and Rivers (2016).

These methods result in an estimate for the distribution of total factor productivity across plants over time. I develop metrics to isolate the change in average productivity due solely to plant entry and exit. Then I examine those metrics during the recession, periods of moderate growth, and booms. These are operationalized to be periods of real GDP growth less than 0%, between 0% and 10%, and greater than 10%, respectively.

This paper finds support for both Schumpeter's theory and Melitz's model; both recessions and booms are cleansing. The process of creative destruction, in which unproductive plants exit and productive plants enter, is generally at work, improving average productivity. However, during recessions and booms, this process improves average productivity more than in periods of moderate growth. During the recession, average productivity is improved by about 3.6 percentage points per year due to plant entry and exit. During years of moderate GDP growth, this number is 1.2 percentage points, and during boom years, 3.1 percentage points. These results are mostly driven by the exit of unproductive plants for recessions and the entry of productive plants for booms.

This is not to say that recessions are good. Furthermore, on the whole, productivity falls during the recession. It is tempered, however, by plants selecting whether to enter or exit. That is, had more productive plants not entered and less productive plants not exited, the decline in productivity during the recession would likely have been greater.

In the next section, I summarize the related literature, including papers on estimating production functions, the papers cited above, and others. In Section 2.3, I discuss the Chilean data set and examine plant entry and exit rates. Section 2.4 outlines the primary production function estimation method. I develop the Entry and Exit Metrics in Section 2.5 and in Section 2.6 present the main results. Section 2.7 considers weighting the metrics by plant size, and Section 2.8 examines the robustness of the results to an alternative estimation method. In Section 2.9, I extend the concepts discussed above from macroeconomic business cycles to industry-specific cycles. Section 2.10 concludes.

2.2. Related work

The possibility that recessions may cleanse the economy of unproductive means of production has been studied by others besides Schumpeter. Caballero and Hammour (1994) examine this, and in particular, consider the extent to which a recession increases the rate at which production units close versus decreasing the rate at which they open.[2] For example, it is theoretically possible that a recession's impact would be absorbed entirely by a reduction in the opening of new production

[2]A production unit could be a production line, a plant, or an entire firm.

units, allowing older production units to close at normal rates. In their paper, they build a structural vintage capital model and calibrate it to job creation and destruction numbers in the United States from 1972 to 1983. The model assumes that older capital is less productive than newer capital, that job destruction means old production units are being closed, and job creation means new units are being opened. Given that job destruction is more responsive to recessions than job creation, they conclude that recessions are cleansing.

In contrast to Caballero and Hammour's work, as well as similar work by other economists using strictly labor data, this paper estimates productivity at the plant level using data on plant inputs and outputs. This paper primarily uses a modern version of a "proxy-variable" production function estimator.

Proxy-variable estimators were first developed in 1996 by Olley and Pakes, hereafter "OP." The main purpose of the proxy-variable estimation technique is to overcome the issue of simultaneity, also called "transmission bias." In the model, a plant chooses an input, such as labor, partially based on a plant-specific productivity level unobservable to the researcher. When plant-specific productivity is high, the plant uses more of the input. This causes the estimate for the effect of the input to be biased away from zero when estimated using ordinary least squares, and this is called transmission bias. To overcome this, OP used a proxy variable, investment, to control for changes in the unobserved plant-specific productivity level. They also considered the issue of selection, or "survival bias," but found it to have little to no effect. Thus, correction for survival bias has been excluded from most papers applying proxy-variable estimators. However, since this paper studies plant survival in particular, I will correct for the selection issue using an analogous method to the one used in OP.

Levinsohn and Petrin (2003), hereafter "LP," developed the next generation of the proxy-variable estimator. Whereas OP used investment as their proxy variable, LP introduced the use of intermediate inputs as the proxy variable. They showed that its use requires fewer assumptions than investment. Furthermore, investment is often zero for a plant, which makes it unattractive for use as a proxy for plant-specific productivity. When the proxy variable is instead intermediate inputs, which are mostly/generally material inputs, this is not such a problem. Additionally, LP reformulated the estimator to be partly a generalized method of moments (GMM) estimator, whereas OP was solvable using non-linear least squares.

In this paper, I use the latest proxy variable estimator, specified by Ackerberg, Caves, and Frazer (2015). They, "ACF," address an identification issue with LP and relax the required timing assumptions regarding the plant's choice of inputs and innovations in plant-specific productivity. While in their 2015 paper, they deal with simulated data to test LP and their estimator's applicability to various data-generating processes, their 2006 working paper used part of the same Chilean data set I use. They used data from 1979 to 1986, whereas I use data from 1979 to 1996.

To check the robustness of my results to a different production function estimator, I apply the method developed by Gandhi, Navarro, and Rivers (2016), hereafter "GNR." Their paper illustrates potential pitfalls with the use of value-added production functions, and they promote instead the use of gross-output production functions. However, they also show that proxy-variable estimators are unidentified for typical specifications of gross-output production functions, and thus they provide an alternative estimation method. Proxy-variable estimators use a monotonicity condition: that as a plant's productivity, known to its operators, increases, then the proxy variable, which is typically intermediate inputs, increases. In GNR's estimator, this monotonicity condition is replaced by the plant's first order condition on intermediate inputs. Consequently, the first stage of the GNR estimator involves "share regression," whereupon the ratio between intermediate-input costs and revenue is regressed on the plant's inputs. This forms a partial derivative of the production function with respect to intermediate inputs. This partial derivative is then integrated to find the production function up to a constant term, and the second stage of the estimator finds the constant of integration.

Liu (1992) was the first to develop and use the Chilean plant-level data set for 1979 through 1986. Using estimation methods based on a fixed-effects model, she finds that exiting plants have lower productivity than plants that remain. She also tracks the productivity of entering plants over time. That competitive pressures select against low productivity plants is a common result between this paper and hers. To some extent, this paper is a modern update: I use more data and apply more modern production function estimators which account for the transmission bias addressed in OP. She defines exit to mean plants permanently exiting the data set, whereas I consider exit as a plant simply closing. Furthermore, she does not address the recession in particular. Finally, she does not account for entering plants when no capital data is observed and lacks an estimate for the aggregate effect of entry and exit on productivity.

Foster, Haltiwanger, and Syverson (2005) consider the issue that firms exit based on profitability, not productivity. They focus on plants from industries that have no differentiation in output and for which quantity and price are known, such as gasoline. Their data comes from a U.S. survey that takes place every five years and thus is unsuitable for studying the effect of recessions. Fortunately, they find a high degree of correlation between measures of profitability and productivity; therefore, little concern is warranted when it comes to considering plants exiting based on productivity instead of profitability.

There is some question as to the effect of a recession on the dispersion of productivity. Kehrig (2011), using LP's estimator, is one of the few papers to address this. Kehrig suggests there are two mutually exclusive effects. The first effect is that recessions distress firms, and some of the least productive firms are forced to close due to diminished demand for their output (in a demand-driven recession) or increased cost of their inputs (in a supply-driven recession). This view is congruent with Schumpeter's theory of creative destruction and predicts that recessions will decrease the dispersion of productivity.

The second effect focuses on competition for scarce inputs such as labor, raw materials, etc. During an economic boom, demand for inputs increases, driving up their prices. Only the most productive firms will be able to compete for the costly inputs, and less productive firms are forced to exit. Inversely, during a recession, less productive firms enter to take advantage of the weak demand for inputs. As these less productive firms enter, the dispersion of productivity during a recession increases. This is the effect predicted by the model presented in Melitz (2003). Kehrig, using U.S. data from 1972 to 2005, finds that dispersion increases during recessions and that the least productive firms see greater declines in productivity during a recession than more productive firms. Thus he finds support for Melitz's model over Schumpeter's theory.

On the other hand, Faggio, Salvanes, and Van Reenen (2010) find opposite results using U.K. data from 1984 to 2001. They define productivity as value added per worker, in contrast to this paper and Kehrig's, which both use TFP estimates from a proxy-variable estimator. Faggio, Salvanes, and Van Reenen assert that their productivity measure follows very closely with a TFP measure derived from average cost shares which itself gives similar results to TFP estimates from more sophisticated estimators. They find that productivity dispersion decreased during the recession of the early 1990s and that the left tail of the productivity distribution was truncated. They consider these results consistent with Schumpeter's theory. Note that these two results, regarding overall dispersion and the behavior of the left tail, are just the opposite of what Kehrig found.

However, examining productivity dispersion is not a good way to test between Schumpeter's and Melitz's propositions. There are plausible reasons to think that dispersion would increase during a recession regardless of plants entering or exiting. Depending on how well inputs are measured, it is likely that in estimating the production function, labor and capital intensity are not captured. During a boom, labor and capital are likely used to their fullest extent. However, during a recession, some plants may fire workers, whereas others may hoard them. Some plants may leave capital idle; others may sell their capital. If these differences are not captured, it will appear as if the plant productivity distribution is more disperse.

Moreover, if there are any adjustment frictions that may differ across firms, we should expect an increase in the dispersion of measured productivity during a recession. Recessions are generally large shocks. If plants re-optimize at different rates, which will be the case if some plants are locked into certain prices and others are not, the recessionary shock will increase dispersion in estimated productivity. On the other hand, periods of high growth generally do not come as large shocks, but are eased into as several periods of accelerating growth. As it is gradual, differences in plants' abilities to optimize will be less important.

Both Kehrig and Faggio et al. used plant- or firm-level data. Thus, it is observable when a plant or firm enters or exits. Therefore, there is no need to assert that differences in dispersion across the business cycle come from entry or exit, which both papers must do in order to use dispersion to evidence the views of Schumpeter or Melitz. Instead of using productivity dispersion to examine these ideas, this paper uses metrics that directly rely on information in the data set regarding plant entry and exit.

All of that said, I find that productivity dispersion increases during the recession. However, I do not have information about labor and capital intensity, so this fact is explainable as discussed above.

2.3. Data

This paper's data set is a panel of plants in Chile from the year 1979 to 1996. The original data source is Chile's annual census on manufacturing, *Encuesta Nacional Industrial Annual*. The census data was first organized as a data set, documented in English, and examined by Liu (1992), covering 1979 through 1986. This paper uses a more recent version of the data set, prepared by Greenstreet (2007). In its various versions, it is a common data set for production function analysis, used in LP, ACF (2006), and GNR.

From 1974 to 1979, Chile's government liberalized its trade policy, privatized state-run firms, and deregulated markets. This set in motion a period of transition for the Chilean economy, which is captured in the first years of the data set. During 1982 and 1983, Chile experienced a recession due to the Latin American debt crisis of 1981 combined with a highly leveraged financial sector. This recession, and its effect on plant entry and exit, is the focus of this paper.

To look for the effect predicted by Melitz (2003), I consider years with real GDP growth in excess of 10% as an economic boom. Thus, I classify 1989, 1992, and 1995 as economic boom years, with real GDP growing 10.6%, 12.3%, and 10.8% respectively. I selected 10% as the threshold because that limited the period of study to three boom years, which is comparable with the two years of recession. Additionally, there is a reasonable gap between the boom year with the least growth, 10.6%, and the year with the next highest growth, which is 1991 with 8.1% growth.[3]

Like GNR, I examine the five largest industries in the data set as determined by three-digit ISIC (International Standard Industrial Classification) Revision 2 codes. These industries are food products (311), textiles (321), apparel (322), wood products (331), and metal products (381). I restrict my analysis to these five industries. Plants that change industry are dropped from the panel: this is important for when I study industry-specific growth rates in Section 2.9.

In order to discuss plant entry and exit, some definitions are in order. I consider a plant open for the year if it is open at least one day. I define a plant to have entered in year t if it is open in year t and not open in year $t-1$. A plant has exited in year t if the plant is not open in year t and is open in year $t-1$. Finally, a plant is persisting in year t if it is open in both year t and $t-1$.

[3] Real GDP grew at 8.3% in 1979. However, while that year is in the panel, it is excluded from the entry–exit analysis described subsequently. It is impossible to infer whether a plant entered or exited in 1979.

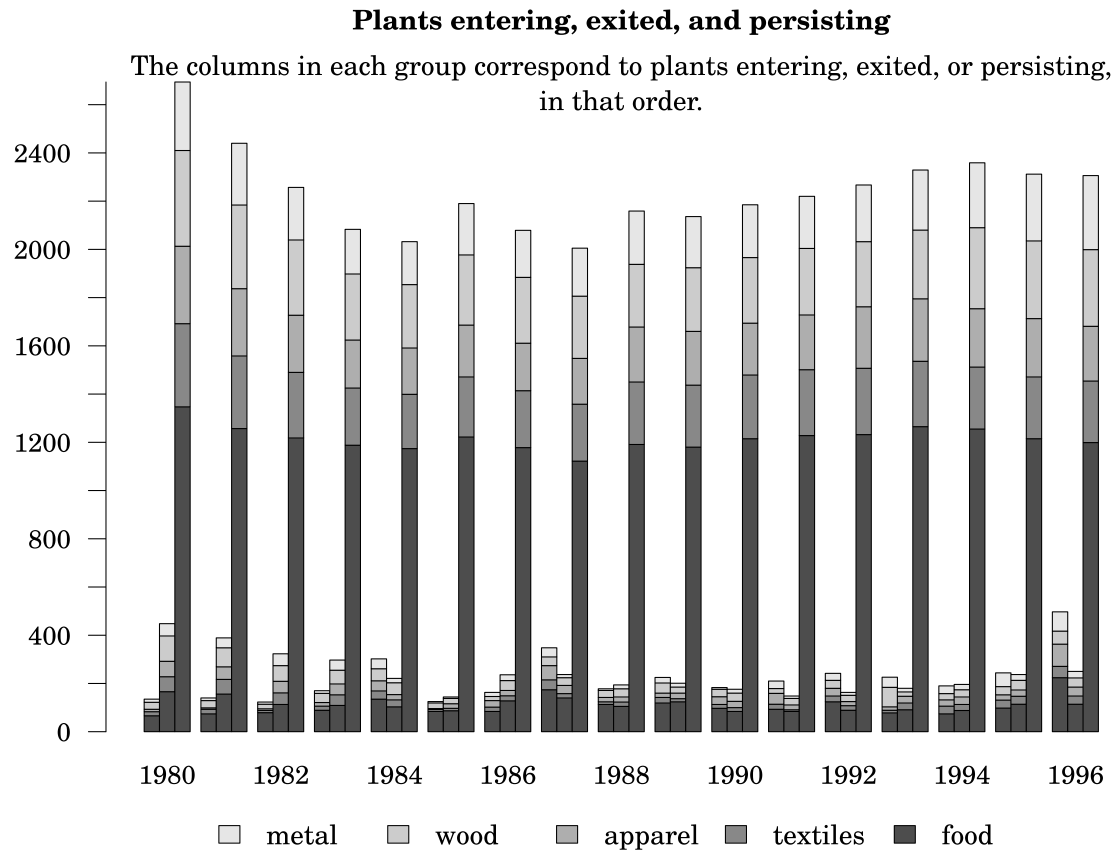

Figure 2.1. Number of plants entering, exited, and persisting over time

Given these definitions, there is no way to ascertain the status of plants during that first year of the data set, 1979. Figure 2.1 shows the number of plants in each state over time.

Note that the number of plants exiting during the recession years 1982–83 is not much different than in the prior years 1980–81. This is contrary to what one would expect to see: that the recession should cause a large increase in the number of plants exiting. As previously mentioned though, the Chilean economy was in a state of transition during these years, and previously protected plants were being forced to exit. This is a mildly unfortunate feature of the data: that there is only one recession to study and that the recession occurred at a time of already naturally high exit numbers. I will address this issue in Section 2.9.

Figure 2.2 further illustrates this effect in terms of rates. Each circle represents a particular industry at a particular time, and the area thereof is proportional to the number of operating plants. The largest industry, by number of plants, is the food industry, and it has the largest circles. There is a downward trend in exit rates from 1980 to 1991. However, while exit rates were also high in 1980 and 1981, the recession years of 1982–83 saw exit rates slightly above the trend. So it is likely the recession increased the rate of exit at least a small amount.

Without the largest industry, food, there would have been an increase in the exit rate during the recession. Unlike the other four industries, the food industry produces a consumer staple / nondurable good. Food consumption is less income elastic than consumption for the products of the other industries; therefore, the food industry was subject to less competitive pressure during the recession. Furthermore, both the food and wood industries export significant amounts of their production. Unlike the wood industry though, the food industry was able to increase sales to the external sector in the face of decreased domestic demand. Appendix C provides more details for both of these effects that uniquely diminish the recession's impact on the food industry.

An expected feature of the recession is the high entry rate in the year following it, 1984. While this is partly due to the re-opening of some plants that exited during the recession, the majority of the plants are new. Another expected feature is that exit rates are low in the years 1984 and 1985. This is likely due to the fact that the recession had already removed relatively unproductive plants.

Considering the graphs in the right column of Figure 2.2, exit rates tend to be higher during recessions than in years of positive GDP growth. Furthermore, years of real GDP growth greater than 10%, classified as booms, have lower exit rates than years of moderate real GDP growth. The opposite is true for entry rates. This is congruent with Schumpeter's theory and makes the predictions of Melitz's model more doubtful.

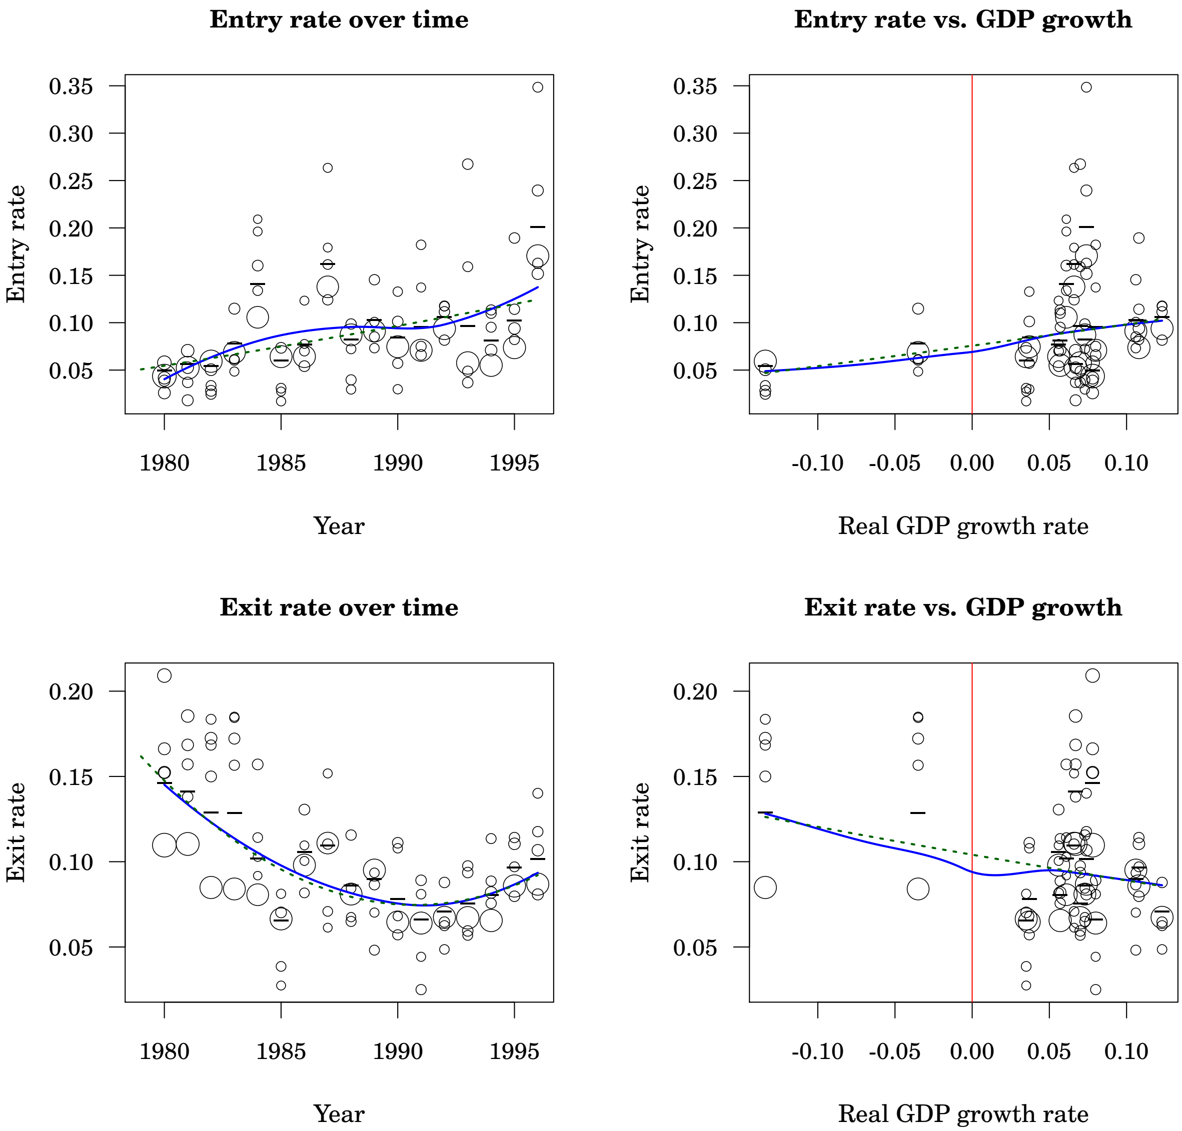

Figure 2.2. Entry and exit rates over time and versus real GDP growth rates. Each circle represents a particular industry at a particular time, and the area thereof is proportional to the number of operating plants. The dashes represent the weighted average for that year. The dashed line represents a fitted curve from a quadratic ordinary least squares regression. The solid line is a fitted curve from quadratic local regression.

The census is conducted only for plants with at least ten employees. This means that there is some risk of falsely identifying a plant as entering when in fact it operated in the previous year with fewer than ten employees and now operates with at least ten. The same risk holds for improperly identifying exit. Greenstreet (2007) addressed this issue by excluding plants that appear to enter with fewer than fifteen employees.[4] In Section 2.7, I will address the issue by assigning less weight to smaller plants in my analysis.

The data set includes a measure for double-deflated real value added, that is, deflated output minus deflated inputs. There are some observations for which real value added is negative. This is mostly due to using multiple different deflators for inputs and outputs. This will be addressed in Section 2.8.

In Liu's original data set, plants were only required to report measures for fixed assets in 1980 and 1981. The capital series is constructed using real investment and by assuming fixed depreciation rates for each class of assets (buildings, vehicles, and equipment). This leads to an issue where plants that enter after 1981 generally are missing a measure for their capital stock. I consider this issue in Appendix B, in which I examine a model similar to the one in Section 2.4 but without capital.

2.4. Production function estimation

The primary model this paper uses for estimating the production function follows ACF (2006) but includes an additional step to correct for survival bias / selection. There are three types of intermediate inputs: real materials, real energy, and real services. The sum of the real inputs is real intermediates, $\mathscr{M}$, and the logarithm of that is represented by μ. ACF (2006) use a value-added production function. Where Y_{it} is the real gross output of plant i at time t, real value added is:

$$V_{it} = Y_{it} - \mathscr{M}_{it}$$

Let L_{it} represent a measure for the number of employees, weighted by their compensation, and let K_{it} represent the real value of the plant's mid-year capital stock. The ACF production

[4]The estimation of his sequential learning model is particularly adversely affected by the risk of spurious entry, as opposed to both spurious entry and exit.

function has a Cobb-Douglas form as follows:

$$V_{it} = L_{it}^{\beta_l} K_{it}^{\beta_k} \exp(\omega_{it} + \varepsilon_{it})$$

Thus, the total factor productivity of the plant is $\exp(\omega_{it} + \varepsilon_{it})$. It is assumed that ω_{it} is observed by the plant's operators but not the researcher, and ε_{it} is unobserved entirely.

Where K'_{it} is the plant's end-of-year capital stock, I define the plant's information set at t as:

$$\mathscr{I}_{it} = \{(Y_{i\tau-1}, L_{i\tau}, K_{i\tau}, K'_{i\tau}, \mathscr{M}_{i\tau}, \omega_{i\tau}) \,|\, \tau \leq t\}$$

Given this definition for the information set, let idiosyncratic productivity, ω_{it}, be a first-order Markov process, and let the unobservable productivity shock, ε_{it}, have conditional mean zero.

$$\Pr(\omega_{it+1} \,|\, \mathscr{I}_{it} \cup \{Y_{it}\}) = \Pr(\omega_{it+1} \,|\, \omega_{it})$$

$$E[\varepsilon_{it} \,|\, \mathscr{I}_{it}] = 0$$

Letting lower-case letters denote the (natural) logarithms, the log production function is:

$$v_{it} = \beta_l l_{it} + \beta_k k_{it} + \omega_{it} + \varepsilon_{it}$$

While this is a linear equation, one cannot simply apply ordinary least squares at this stage due to the issue of transmission bias / simultaneity. As explained by OP, plants that observe high ω will choose to invest more and hire more. Thus, marginal increases in value added or output due to an increase in ω_{it} will seem to be caused only by increases in l_{it} or k_{it}, which will bias the estimates for β_l or β_k away from zero. Therefore, in order to estimate β_l and β_k without bias, a different estimation method must be used.

One class of estimation methods designed to address transmission bias are proxy-variable methods. Assuming all plants face identical prices, a plant's (conditional) demand for intermediate inputs can be written as:

$$\mu_{it} = h(l_{it}, k_{it}, \omega_{it})$$

Assuming that this function is strictly monotonic in ω_{it} for relevant values of l_{it} and k_{it}, this can be inverted to:

$$\omega_{it} = h^{-1}(l_{it}, k_{it}, \mu_{it})$$

Using this method, μ_{it} is called the "proxy variable." LP demonstrated how monotonicity holds under common regularity conditions on the plant's gross output production function and the plant's optimizing behavior. Substituting into the log production function yields:

$$\begin{aligned} v_{it} &= \beta_l l_{it} + \beta_k k_{it} + h^{-1}(l_{it}, k_{it}, \mu_{it}) + \varepsilon_{it} \\ &= \psi(l_{it}, k_{it}, \mu_{it}) + \varepsilon_{it} \end{aligned}$$

where the $\beta_l l_{it} + \beta_k k_{it}$ is subsumed into the ψ function, which is to be estimated nonparametrically. I estimate ψ with a cubic polynomial series/sieve estimator and define the fitted values of that function as $\hat{\psi}_{it}$.

Up to this point, I have followed ACF (2006). Now, I detour slightly to correct for survival bias with an intermediate stage, following a method similar to OP's work regarding selection correction.

The selection / survival bias issue arises because plants may choose to exit based on their idiosyncratic productivity and capital. The idea is that plants with high capital may be less willing to exit during times of low ω than plants with less capital. This may be due to greater costs associated with offloading a larger plant's assets or that larger plants have greater access to liquidity to withstand periods of low productivity. The implication is that, in the data, large plants may have lower average ω than smaller plants. Thus, without taking into account this selection issue, β_k will be negatively biased (as will β_l insofar as large plants hire many workers). The solution to this issue is to employ a Heckman-like index for use in the final stage of the procedure as OP did.

Before continuing, a few timing assumptions are in order.[5] I assume that at the beginning of the year, plants observe their idiosyncratic productivity and decide whether to exit according to a threshold rule, which itself is a function of the plant's beginning-of-year capital. That is, a plant will exit in year $t+1$ if ω_{it+1} is less than $\underline{\omega}(k'_{it})$.[6] If a plant does not exit, then at the beginning of the year, it will choose its levels of capital investment, labor, and intermediate inputs. This determines variables k_{it}, k'_{it}, l_{it}, and μ_{it}.

Let $o_{it} = 1$ if a plant operates in year t and $o_{it} = 0$ if the plant does not. The probability that a plant operates in period $t+1$ given the information it has at time t is therefore a function of l_{it}, k_{it}, μ_{it}, and k'_{it}.

$$\begin{aligned}
\Pr(o_{it+1} = 1 \mid \mathscr{I}_{it}) &= \Pr(\omega_{it+1} \geq \underline{\omega}(k'_{it}) \mid \mathscr{I}_{it}) \\
&= \Pr(\omega_{it+1} \geq \underline{\omega}(k'_{it}) \mid \omega_{it}, k'_{it}) \\
&= \Pr(\omega_{it+1} \geq \underline{\omega}(k'_{it}) \mid h^{-1}(l_{it}, k_{it}, \mu_{it}), k'_{it}) \\
&= p(l_{it}, k_{it}, \mu_{it}, k'_{it})
\end{aligned}$$

The function p is estimated nonparametrically. In particular, I estimate p using probit regression with a cubic polynomial in l_{it}, k_{it}, μ_{it}, and k'_{it}. I call the fitted values $\hat{p}_{it}$, and where Φ is

[5] Note that up to this point, I have not needed to make any nontrivial timing assumptions: capital and labor may be chosen concurrently with intermediate inputs. This is one of the contributions of ACF.

[6] Recall that K'_{it} represents plant i's end-of-year capital stock in time t. Therefore, K'_{it} is also the plant's beginning-of-year capital stock in time $t+1$.

the standard normal cumulative distribution function, the estimated mean function is given as:

$$\hat{p}_{it} = \Phi(\sum_{\alpha_l+\alpha_k+\alpha_\mu+\alpha_{k'}\leq 3} \gamma_{\alpha_l,\alpha_k,\alpha_\mu,\alpha_{k'}} l_{it}^{\alpha_l} k_{it}^{\alpha_k} \mu_{it}^{\alpha_\mu} {k'_{it}}^{\alpha_{k'}}) \text{ with } \alpha_l,\alpha_k,\alpha_\mu,\alpha_{k'} \geq 0$$

The final stage of the algorithm is to use GMM on moment conditions of the prediction error in ω. Define ξ_{it} as the prediction error in ω:

$$\xi_{it} = \omega_{it} - E[\omega_{it} | \mathscr{I}_{it-1}]$$

By the timing assumption, k'_{it-1} and l_{it-1} are determined in $t-1$. Consequently, they must be uncorrelated with prediction error ξ_{it}. Thus, they can be used as instruments in the following GMM moment conditions:

$$E[\xi_{it} | k'_{it-1}] = E[\xi_{it} k'_{it-1}] = 0$$

$$E[\xi_{it} | l_{it-1}] = E[\xi_{it} l_{it-1}] = 0$$

To utilize these moment conditions, I first need a way to calculate an estimate for ξ_{it}. Because $\hat{\psi}_{it}$ does not include ε_{it} , note that for some guessed parameters, $(\tilde{\beta}_l, \tilde{\beta}_k)$, the implied $\tilde{\omega}_{it}$ is:

$$\tilde{\omega}_{it} = \hat{\psi}_{it} - \tilde{\beta}_l l_{it} - \tilde{\beta}_k k_{it}$$

Let Ω represent the function that estimates $E[\omega_{it} | \mathscr{I}_{it-1}]$. Because ω is a first-order Markov process, the estimate for the expected value of ω_{it} would normally only be a function of ω_{it-1}. However, because of the selection issue, that estimate would be biased. Therefore, to adjust for that bias, I must include the selection index $\hat{p}_{it-1}$. I estimate Ω by regressing $\tilde{\omega}_{it}$ onto a cubic polynomial of $\tilde{\omega}_{it-1}$ and $\hat{p}_{it-1}$. The residuals of that regression represent the prediction error given the guessed parameters:

$$\tilde{\xi}_{it} = \tilde{\omega}_{it} - \tilde{\Omega}(\tilde{\omega}_{it-1}, \hat{p}_{it-1})$$

Then I can multiply $\tilde{\xi}_{it}$ by k'_{it-1} and l_{it-1} to find the value of the moment conditions for $(\tilde{\beta}_l, \tilde{\beta}_k)$. I thus search across the parameter space for the values $\hat{\beta}_l$ and $\hat{\beta}_k$ that best satisfy the sample analog of the moment conditions using continuously updating GMM.

2.5. Entry and Exit Metrics

Whereas OP and GNR define productivity as $\exp(\omega_{it} + \varepsilon_{it})$, I leave productivity in natural logarithms: simply $\omega_{it} + \varepsilon_{it}$. For the ACF estimation method, let the residuals $r_{it} = v_{it} - \hat{\beta}_l l_{it} + \hat{\beta}_k k_{it}$ represent the estimate for $\omega_{it} + \varepsilon_{it}$, plant i's productivity in year t.

The production function estimation routine can be thought to provide a series of plant productivity distributions over time. To assess the effect of plant entry and exit on aggregate productivity, I must isolate the changes in the productivity distribution due to time. To identify the effect of time between two years, I compare the productivity levels of plants that exist in both years.

I find it helpful to consider the problem graphically. Consider Figure 2.3. Each rectangle represents a plant's productivity level at a particular time. The gray rectangles represent plants that remain open (persist) throughout the sample. The black rectangle represents a plant that exits in time 3. The white one represents a plant that enters in time 2.

From time 1 to time 2, average productivity increased from 1.5 to 3. However, some of that change was due to a relatively productive plant entering; some of the change was just a general increase in productivity between the years. I identify the time effect as the change in average productivity of plants operating in both time 1 and time 2. This is the average pairwise difference, and in Figure 2.3, this is 1.

Thus, in order for the distribution in time 2 to be comparable to time 1, it must be shifted down by 1. I call this "adjusted productivity" and the adjusted productivity distribution is shown in Figure 2.4. The adjusted productivity distribution for time 1 is the same as the productivity distribution for time 1. For all subsequent times, the productivity distribution is shifted such that

Productivity distribution over time

Average:	1.5	3	1.75
Difference:	n/a	1.5	–1.25
Average pairwise difference:	n/a	1	–1.5

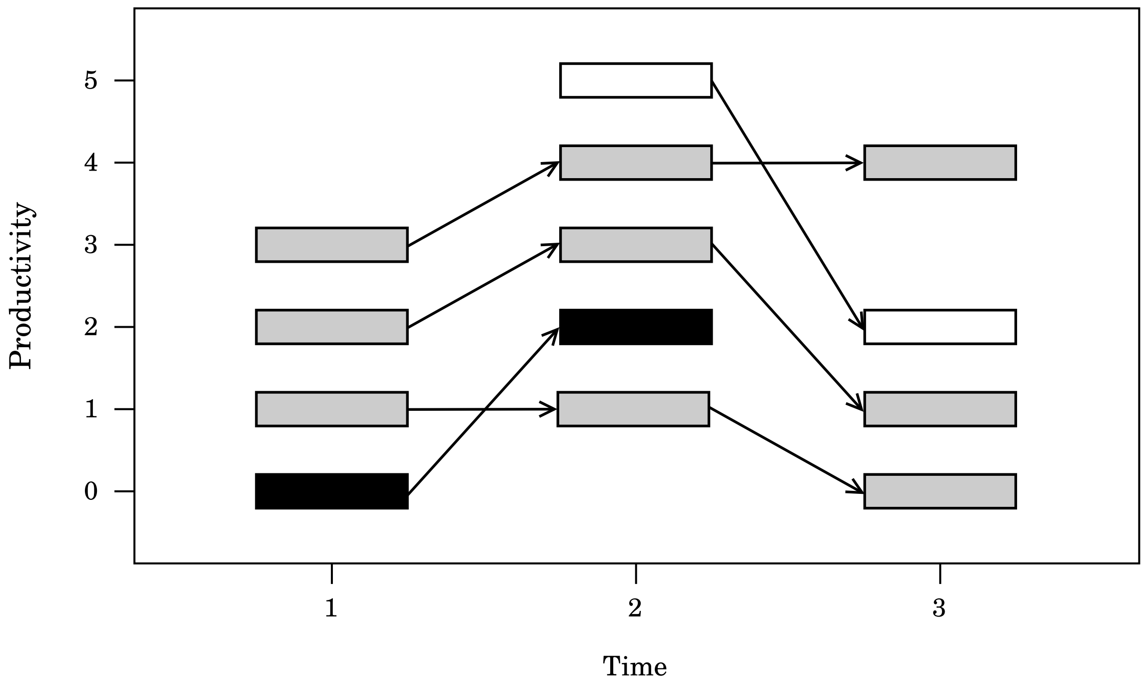

Figure 2.3. Example productivity distribution over time. This figure and the next one illustrate the concept of "adjusted productivity."

the average pairwise difference is 0. This is equivalent to minimizing the sum of squared pairwise differences.

Adjusted productivity isolates the effects of entry and exit on productivity. If there was no entry or exit, then average adjusted productivity would be constant over time. From time 1 to time 2, average adjusted productivity increased from 1.5 to 2. Thus, the white plant's entry caused average productivity to increase by 0.5.

Returning to Figure 2.3, similar operations are applied for moving from time 2 to time 3. Between those times, average productivity fell by 1.25, but productivity fell by 1.5 on average for

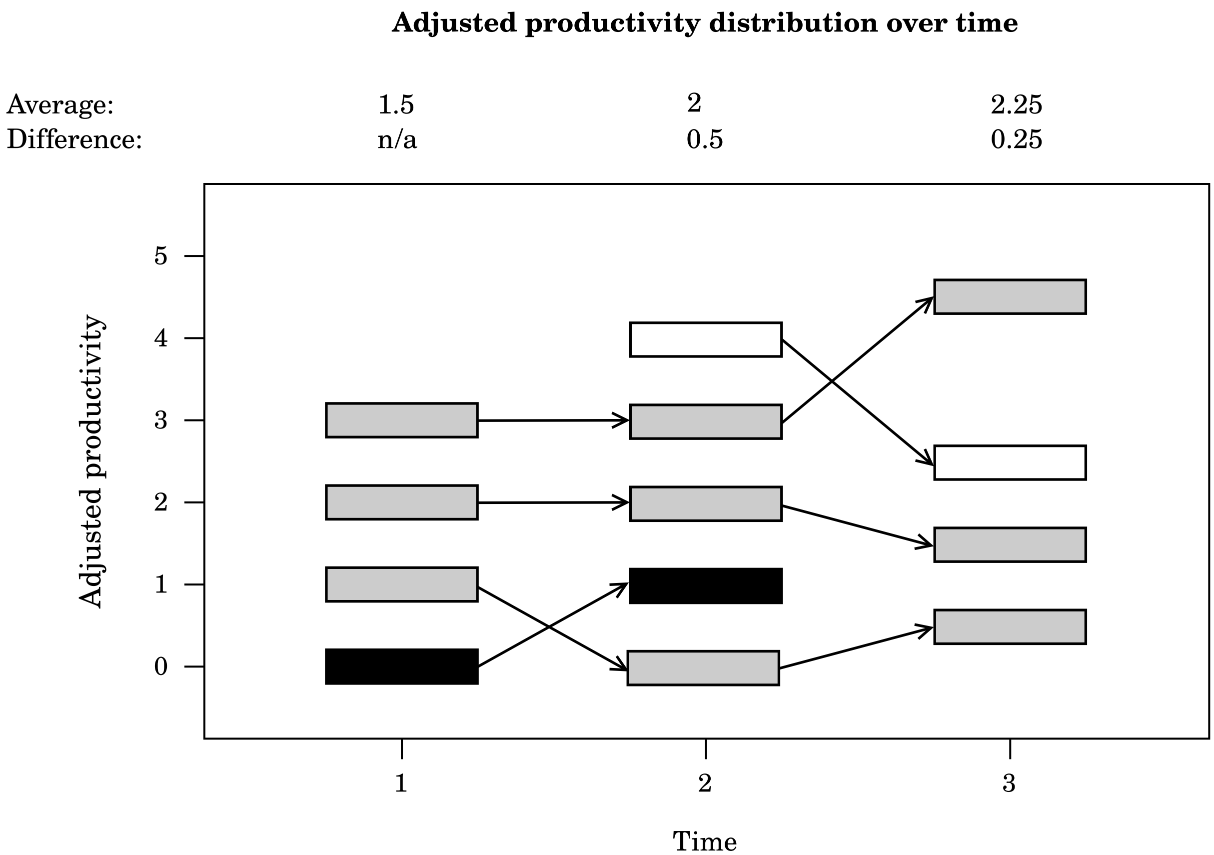

Figure 2.4. Example adjusted productivity distribution over time. This figure and the previous one illustrate the concept of "adjusted productivity."

plants operating in both times. Adjusted productivity for time 3 is equal to productivity in time 2 minus the cumulative sum of the average pairwise differences. The cumulative average pairwise difference for time 3 is $1 + -1.5 = -0.5$, so adjusted productivity is equal to productivity plus 0.5. Thus, when one compares Figure 2.3 to Figure 2.4, it is apparent that, for time 3, adjusted productivity is productivity shifted up by 0.5.

Average adjusted productivity increased from 2 to 2.25 between times 2 and 3; therefore, the effect of the plant exiting was to increase average productivity by 0.25.

Using this intuition, the mathematical formula for these concepts follows. Let r_{it} represent the estimated productivity for plant i at time t. Let $o_{it} = 1$ if plant i is operating in time t, and

$o_{it} = 0$ otherwise. Suppose the first observation time is t_1. Then let $\dot{r}_t$ represent the average pairwise difference in productivity:

$$\dot{r}_t = \begin{cases} \frac{\sum_i r_{it} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}} - \frac{\sum_i r_{it-1} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}} & t > t_1 \\ 0 & t = t_1 \end{cases}$$

I define $\tilde{r}_{it}$, the adjusted productivity of plant i in time t, as the plant's productivity r_{it} minus cumulative average pairwise differences.

$$\tilde{r}_{it} = r_{it} - \sum_{\tau=t_1}^{t} \dot{r}_\tau$$

Additionally, let $\tilde{r}_{\cdot t}$ represent average adjusted productivity in time t.[7]

$$\tilde{r}_{\cdot t} = \frac{\sum_i \tilde{r}_{it} o_{it}}{\sum_i o_{it}}$$

While the change in average adjusted productivity captures the effect of entry and exit as discussed regarding Figure 2.4, it would be good to separate the effect of entry from exit. For this purpose, I define the "Entry Metric" as the cumulative increase in average productivity due to plant entry, and the "Exit Metric" similarly for plant exit. I construct these metrics iteratively, such that:

$$\text{Entry Metric}_t = \begin{cases} \sum_{\tau=t_1+1}^{t} \Delta\,\text{Entry Metric}_\tau & t > t_1 \\ 0 & t = t_1 \end{cases}$$

$$\text{Exit Metric}_t = \begin{cases} \sum_{\tau=t_1+1}^{t} \Delta\,\text{Exit Metric}_\tau & t > t_1 \\ 0 & t = t_1 \end{cases}$$

[7] I retain the o_{it} in the numerator so as to emphasize the number of non-zero elements being summed, which is illustrative in the decomposition that follows.

To determine the ΔEntry Metric and ΔExit Metric, I decompose the change in average adjusted productivity into the sum of two addends, one particular to entry and one particular to exit. For $t > t_1$, let:

$$\begin{aligned}
\Delta\tilde{r}_{\cdot t} &= \tilde{r}_{\cdot t} - \tilde{r}_{\cdot t-1} \\
&= \frac{\sum_i \tilde{r}_{it} o_{it}}{\sum_i o_{it}} - \frac{\sum_i \tilde{r}_{it-1} o_{it-1}}{\sum_i o_{it-1}} \\
&= \frac{\sum_i [r_{it} - \sum_{\tau=t_1}^{t} \dot{r}_\tau] o_{it}}{\sum_i o_{it}} - \frac{\sum_i [r_{it-1} - \sum_{\tau=t_1}^{t-1} \dot{r}_\tau] o_{it-1}}{\sum_i o_{it-1}} \\
&= \frac{\sum_i r_{it} o_{it}}{\sum_i o_{it}} - \sum_{\tau=t_1}^{t} \dot{r}_\tau - \frac{\sum_i r_{it-1} o_{it-1}}{\sum_i o_{it-1}} + \sum_{\tau=t_1}^{t-1} \dot{r}_\tau \\
&= \frac{\sum_i r_{it} o_{it}}{\sum_i o_{it}} - \frac{\sum_i r_{it-1} o_{it-1}}{\sum_i o_{it-1}} - \dot{r}_t \\
&= \underbrace{\frac{\sum_i r_{it} o_{it}}{\sum_i o_{it}} - \frac{\sum_i r_{it} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}}}_{\Delta \text{Entry Metric}_t} + \underbrace{\frac{\sum_i r_{it-1} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}} - \frac{\sum_i r_{it-1} o_{it-1}}{\sum_i o_{it-1}}}_{\Delta \text{Exit Metric}_t}
\end{aligned}$$

Note that ΔEntry Metric and ΔExit Metric are defined only for $t > t_1$. They both require the use of information regarding the operating status of plants in the previous period which is necessary to identify the persisting plants.

Recursively formulated, the Entry Metric is:

$$\text{Entry Metric}_t - \text{Entry Metric}_{t-1} = \Delta\text{Entry Metric}_t = \frac{\sum_i r_{it} o_{it}}{\sum_i o_{it}} - \frac{\sum_i r_{it} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}}$$

It is compelling that the ΔEntry Metric$_t$, while derived from adjusted productivity, is the moments estimator for the following simple difference in conditional expectations:

$$E[\omega_{it} + \varepsilon_{it} \,|\, o_{it} = 1] - E[\omega_{it} + \varepsilon_{it} \,|\, o_{it} = 1 \wedge o_{it-1} = 1]$$

I interpret the ΔEntry Metric$_t$ to measure the increase in average productivity due to plants entering in year t. Behind this interpretation is the implicit counterfactual assumption that had the entering plants not entered, the average productivity of the persisting plants would not have been different.

Similarly, the Exit Metric is recursively formulated as:

$$\text{Exit Metric}_t - \text{Exit Metric}_{t-1} = \Delta\text{Exit Metric}_t = \frac{\sum_i r_{it-1} o_{it} o_{it-1}}{\sum_i o_{it} o_{it-1}} - \frac{\sum_i r_{it-1} o_{it-1}}{\sum_i o_{it-1}}$$

Also, the ΔExit Metric$_t$ is the moments estimator for:

$$E[\omega_{it-1} + \varepsilon_{it-1} | o_{it} = 1 \wedge o_{it-1} = 1] - E[\omega_{it-1} + \varepsilon_{it-1} | o_{it-1} = 1]$$

The ΔExit Metric$_t$ is defined using r_{it-1}, which is important since plants that exit in time t are missing r_{it}. I interpret the ΔExit Metric$_t$ to measure the increase in average productivity due to plants exiting in t. Unlike the Entry Metric, this interpretation requires two counterfactual assumptions. The first counterfactual assumption is the same as that for the Entry Metric: that the average productivity of persisting plants would have been the same had the exiting plants not exited.

The second counterfactual assumption is that had the exiting plants remained, they would have maintained their relative position in the productivity distribution from the previous year. This second assumption is necessary to move from the productivity estimate of $t-1$ to a counterfactual productivity level for t. Alternatively stated, this assumption is that counterfactual productivity, r^*_{it}, for plant i that operated in $t-1$ but exited in t is given by:

$$r^*_{it} = r_{it-1} + \dot{r}_{\cdot t}$$

By studying the Entry and Exit Metrics, I can evaluate the effect of entry and exit on average productivity. This is something that previous studies, that relied solely on dispersion or quantile statistics, could not do.

In Appendix A, I address how these metrics are modified to handle missing data and discuss alternative counterfactual perspectives.

2.6. Results

I estimate a separate production function model for each of the five industries. Then I demean the residuals across models to make them cross-comparable. Unlike Kehrig (2011), I do not divide by the standard deviation estimate, since that would destroy the interpretation of the metrics described below.

Figure 2.5 shows the Entry and Exit Metrics for each industry, as well as the sum of the metrics. The areas of the circles are proportional to the number of plants for which I have a productivity estimate at that time in that industry. The trend lines are calculated by weighting the data points accordingly. By number of plants, the food industry is the largest, and it has the largest circles on the graphs. The Exit Metric increases over time: low productivity plants tend to exit, thereby bringing up the average level of productivity. The magnitude of this effect is remarkable: real productivity is about 25 (log)% higher in 1990 than it was in 1980 strictly due to plants exiting.

Unlike the Exit Metric, the Entry Metric is non-monotone, and the magnitude is much smaller compared to the Exit Metric. Prior to 1987, entering plants tended to improve the average level of productivity. Afterward, however, entering plants decreased it. It could be expected that entering plants would generally improve the average level of productivity as they would likely have newer capital and technology than older plants. One possible explanation for the weaker entry effect is that nascent plants are not likely at peak productivity. New plants may not have yet fully trained

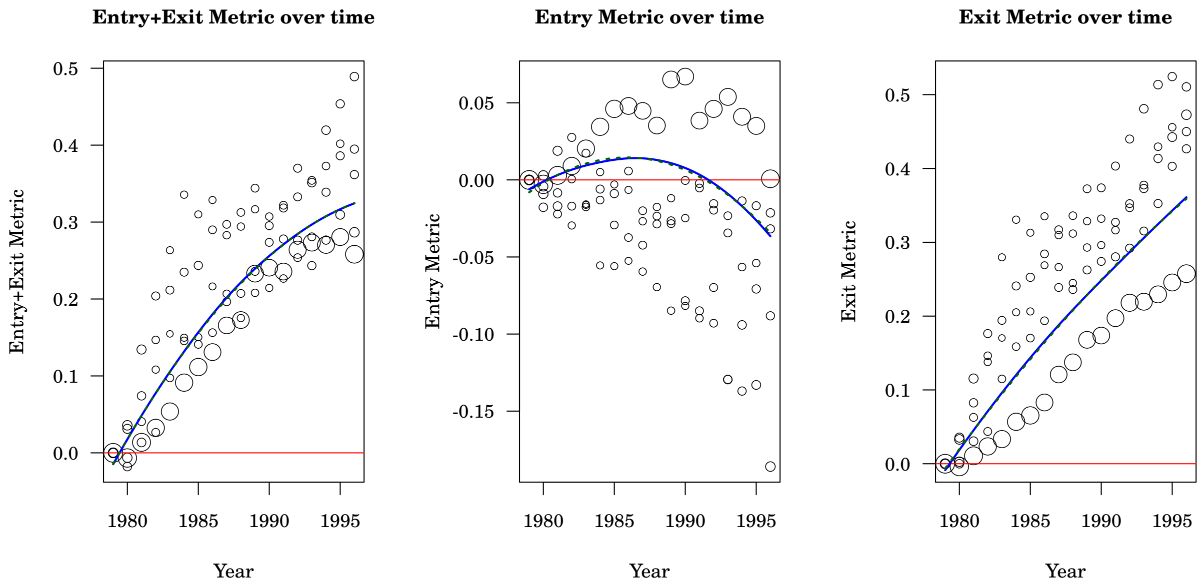

Figure 2.5. Entry and Exit Metrics over time

their workforce, optimized systems of production, or otherwise engaged in learning-by-doing. Furthermore, new plants, if they belong to new firms, may not have the market power to command prices similar to their more well-established competition. Regarding the Entry Metric, this effect explains why the magnitude is small and the slope is generally negative.

Why might the Entry Metric increase up through 1987 and decrease thereafter? One explanation for this would be that since the Chilean economy was in a state of flux in the early part of the data set, entering plants, backed by new foreign and domestic investment, were able to carve out niches in their industries at the expense of older plants that were previously protected by regulations. As time progressed, these niches were filled, and the old protected plants were driven out or made more efficient, and thus the nascent plant effect dominates.

Figure 2.6 plots the change in the Entry and Exit Metrics against the annual growth rate of real GDP. The ΔExit Metric takes on a convex shape. Years of negative GDP growth are associated with an increase in average productivity due to plant exit. This is evidence for Schumpeter's theory that recessions are periods of intensified creative destruction. However, periods of high

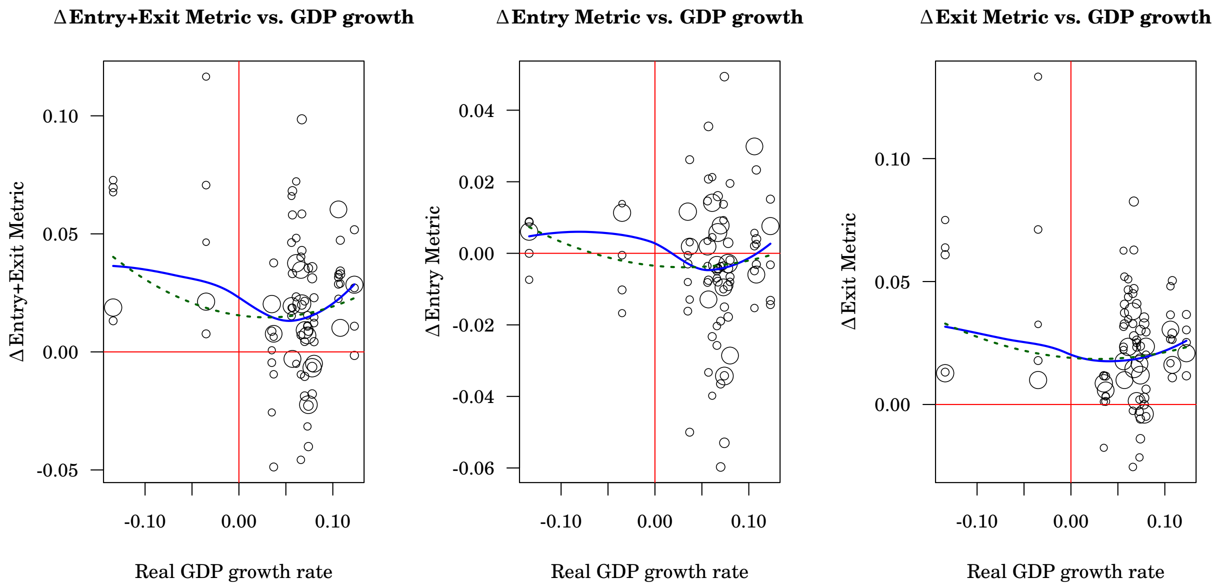

Figure 2.6. Change in the Entry and Exit Metrics versus real GDP growth

GDP growth are also associated with an increase in average productivity due to plant exit. This is evidence for Melitz's idea that increased competition for inputs during a boom will cause productivity gains from exiting plants. Thus, the views of Schumpeter and Melitz are not exclusive.

Reflecting back to Figure 2.2, during economic booms, exit rates are low and entry rates are high. I suggested these facts cast doubt on Melitz's model; thus, the result that productivity is improved by exiting plants during a boom is remarkable. Since the number of exiting plants is low, yet average productivity improves with their exit, the productivity of plants that exit during a boom must be particularly low.

Note that for the Δ Exit Metric, the largest industry, food, is pulling down the average change in the Exit Metric during the recessionary years. Had I excluded that industry from the analysis, the graph in Figure 2.6 would have been more convex. As discussed in Section 2.3, food is a consumer staple and thus is subject to a smaller demand shock during the recession than the other industries. There is less competitive pressure forcing unproductive food plants to exit; therefore, the change in the Exit Metric is smaller.

Category			1982	1983	1982–83	$g \geq 0$	$0 \leq g < 10$
category avg. growth	entry rate exit rate plants open	mean bootstrap mean std. error	difference = column − row standard error percentile p-value (one-sided)				
1982 −13.4%	0.0545 0.1252 2380	0.035033 0.035164 0.00080284					
1983 −3.5%	0.0816 0.1248 2253	0.036865 0.036956 0.000909	−0.00179 0.000712 0.991				
1982–83 −8.5%	0.0675 0.1250 4633	0.035925 0.036037 0.00077867	−0.000874 0.000347 0.991	0.000919 0.000365 0.009			
$g \geq 0$ 7.3%	0.1011 0.0921 37121	0.016085 0.016082 0.00023206	0.0191 0.000764 0	0.0209 0.000918 0	0.02 0.000764 0		
$0 \leq g < 10$ 6.3%	0.0999 0.0945 29695	0.012435 0.012429 0.00019455	0.0227 0.000757 0	0.0245 0.000895 0	0.0236 0.000746 0	0.00365 5.92×10^{-5} 0	
$g \geq 10$ 11.2%	0.1059 0.0821 7426	0.031022 0.03103 0.00043865	0.00413 0.000838 0	0.00593 0.00104 0	0.00501 0.000874 0	−0.0149 0.000242 1	−0.0186 0.000302 1

Table 2.1. For the ACF estimator, the average ΔEntry + Exit Metric for separate periods and the differences between periods. Variable g represents the percent real GDP growth rate.

The change in the Entry Metric during the recessionary years is greater than the average change during growth years. So, this is evidence that only highly productive plants could possibly enter during a recession.

How significant are these effects? Using a non-parametric block bootstrap, run for 999 iterations, I can establish the results in Table 2.1 for the ΔEntry + Exit Metric.[8] The average increase in productivity per year due to plant entry and exit during the recession was 2 percentage points

[8] The table requires some way to aggregate the ΔEntry + Exit Metric across years and industries. The question is how to combine the ΔEntry + Exit Metric for any given industry with the other industries. Furthermore, for the rows of the table that are not "1982" and "1983," there is a question of how to aggregate across years. I take an average weighted according to the number of extant residuals for that year in that industry, which corresponds to the areas of the circles in Figures 2.5 and 2.6.

higher than in years of positive GDP growth. Therefore, the recession years saw greater improvement in average productivity due to entry and exit than the average positive growth year.

The effect is slightly more pronounced if one compares 1982 and 1983 versus years of moderate GDP growth, when the growth rate was between 0% and 10%. Then the difference is 2.36 percentage points per year.

The years of economic boom saw an average 1.86 percentage point increase in productivity due to plant entry and exit over years of moderate growth. This evidences the implication of Melitz's model. In these exceptional growth years, the entry rate is higher than the exit rate, and entry contributed more than exit relative to the recessionary years.

2.7. Weighted average productivity

Up to this point, when discussing changes in average productivity, the average has simply been calculated across numbers of plants. No account was made for the size of the plants. So one cannot really say that economy-wide productivity increases by the aforementioned amounts due to entry and exit. It could very well be that these changes are insignificant if the size of the plants entering and exiting is small. Thus, I consider weighting the Entry and Exit Metrics by w_{it}, a measure for the size of plant i in year t:

$$\Delta\,\text{Weighted Entry Metric}_t = \frac{\sum_i w_{it} r_{it} o_{it}}{\sum_i w_{it} o_{it}} - \frac{\sum_i w_{it} r_{it} o_{it} o_{it-1}}{\sum_i w_{it} o_{it} o_{it-1}}$$

$$\Delta\,\text{Weighted Exit Metric}_t = \frac{\sum_i w_{it-1} r_{it-1} o_{it} o_{it-1}}{\sum_i w_{it-1} o_{it} o_{it-1}} - \frac{\sum_i w_{it-1} r_{it-1} o_{it-1}}{\sum_i w_{it-1} o_{it-1}}$$

For weighting by plant size, a natural choice for weights would be real value added or gross output. However, for those particular weighting schemes, outliers are exaggerated and diminished asymmetrically. Consider a plant with implausibly high real value added relative to its capital and labor input. Such a plant would have very high estimated productivity, and the weight of that plant would be very high. Assigning a large weight to such a plant is exactly the opposite of what a statistician would generally do to an observation that is already an outlier bordering on the realm of credibility. This is not an issue for plants with very low value added and estimated productivity, which would be given very low weight. Therefore, there exists an inherent asymmetry.

Admittedly, there is supposed to be an asymmetry with the weights: larger plants should be weighted more. My concern is that measurement error in value added, which is estimated as productivity, will improperly emphasize positive outliers. Since this is a study of plants with very low productivity exiting during periods of high competitive pressure, I cannot simply exclude observations with extreme productivity estimates.

This issue exists because productivity is correlated with real value added and gross output. However, there is another suitable measure for plant size: its use of inputs. Because the ACF estimation method is "close" to ordinary least squares, inputs are fairly uncorrelated with estimated productivity. Instead of choosing one particular input (labor or capital) as the weight, I have opted to use a mix, the fitted values for real value added:

$$w_{it} = \exp(v_{it} - r_{it}) = \exp(\hat{\beta}_l l_{it} + \hat{\beta}_k k_{it})$$

Whereas the correlation between the productivity estimate and real value added is 0.20, the correlation between the productivity estimate and w_{it} is -0.017.[9] Furthermore, not only does input usage provide an uncorrelated measure for plant size, it also provides a nice interpretation as to the extent that resources are being used efficiently. As large unproductive plants exit, they free up labor and capital for use in more productive plants. At least, this is true insofar as said resources are simply employed or organized inefficiently as opposed to being inherently unproductive.

[9]The Spearman correlations are 0.58 for real value added and 0.025 for w_{it}.

Category			1982	1983	1982–83	$g \geq 0$	$0 \leq g < 10$
category	entry rate	mean	difference = column – row				
avg. growth	exit rate	bootstrap mean	standard error				
	plants open	std. error	percentile p-value (one-sided)				
1982	0.0545	0.033858					
−13.4%	0.1252	0.033998					
	2380	0.00053069					
1983	0.0816	0.0075657	0.0264				
−3.5%	0.1248	0.0076027	0.000575				
	2253	0.00050964	0				
1982–83	0.0675	0.020249	0.0137	−0.0127			
−8.5%	0.1250	0.020336	0.000298	0.000278			
	4633	0.00043315	0	1			
$g \geq 0$	0.1011	0.001398	0.0326	0.00625	0.019		
7.3%	0.0921	0.0013522	0.000582	0.000605	0.00052		
	37121	0.0001524	0	0	0		
$0 \leq g < 10$	0.0999	-0.0035166	0.0376	0.0112	0.0239	0.00491	
6.3%	0.0945	-0.0035559	0.000584	0.000594	0.000514	5.94×10^{-5}	
	29695	0.00014118	0	0	0	0	
$g \geq 10$	0.1059	0.017149	0.0169	−0.00948	0.00325	−0.0157	−0.0206
11.2%	0.0821	0.017083	0.000616	0.000676	0.00058	0.00019	0.00025
	7426	0.00028511	0	1	0	1	1

Table 2.2. For the ACF estimator, the average ΔWeighted Entry + Exit Metric for separate periods and the differences between periods. Variable g represents the percent real GDP growth rate.

As discussed in Section 2.3, there exists a possible issue with improperly identifying a plant as having exited when in fact it operated with fewer than ten employees. A similar risk holds for spurious identification of entry. By weighting plants based on their input usage, less weight is given to these plants for which there is a greater risk of spurious entry or exit.

Table 2.2 shows the results with this weighting scheme.[10] Once again, there is strong evidence for Schumpeter's theory as the recessionary years show a greater improvement in the Entry and Exit Metric per year than growth years by 1.9 percentage points. The prediction of

[10]For this table only, I average across time and industry according to the sum of the weights of the plants with extant residuals for that year in that industry.

Melitz's model still holds as well, with periods of exceptional growth improving productivity due to entry and exit by 2.06 percentage points per annum over periods of moderate growth.

2.8. Robustness of the estimation method

An alternative method for the estimation of production functions has been developed by GNR. In their paper, they argue against the use of structural value-added production functions as used by ACF (2006) and other papers. They focus on estimating the gross output function (which includes intermediate inputs). Furthermore, instead of using a proxy variable, they build their estimation routine on a plant's first order condition for flexible inputs. Their standard model, which they use on this same Chilean data set, uses stronger timing restrictions than I employed with the ACF estimator above. In particular, they assume that both capital and labor are predetermined. That is, they assume that the only input over which a plant has any control in year t is the intermediate input; both labor and capital are determined in the previous year.

Since GNR and ACF are both contemporary estimators, I present the GNR results here. For this estimator, I adopt their stronger assumptions regarding capital and labor timing. Additionally, in keeping with the standard model GNR present in their paper, I make no adjustment for survival bias. I follow the setup presented in their paper exactly, except I include an interaction term for labor, capital, and intermediate inputs in the polynomial sieve estimator for the share regression, which would otherwise be purely quadratic.[11]

As mentioned in Section 2.3, for a number of observations, the measure for real value added is negative. By estimating a gross output production function by way of GNR's method, this issue is sidestepped as the number of observations with negative real gross output is very small. The ACF estimator, corresponding to Table 2.1, returns 37,513 productivity estimates. By comparison, the GNR estimator returns 38,500 productivity estimates. However, there is little change in the

[11]This matches the computer code that GNR have made available that implements their estimator.

Category			1982	1983	1982–83	$g \geq 0$	$0 \leq g < 10$
category avg. growth	entry rate exit rate plants open	mean bootstrap mean std. error	difference = column – row standard error percentile p-value (one-sided)				
1982 −13.4%	0.0545 0.1252 2380	0.01583 0.017208 0.0027267					
1983 −3.5%	0.0816 0.1248 2253	0.023271 0.024829 0.0029831	−0.00762 0.000817 1				
1982–83 −8.5%	0.0675 0.1250 4633	0.019417 0.020882 0.0028238	−0.00367 0.000394 1	0.00395 0.000423 0			
$g \geq 0$ 7.3%	0.1011 0.0921 37121	0.0072007 0.0082027 0.0012465	0.009 0.00165 0	0.0166 0.00189 0	0.0127 0.00172 0		
$0 \leq g < 10$ 6.3%	0.0999 0.0945 29695	0.0057459 0.0065936 0.001035	0.0106 0.00183 0	0.0182 0.00207 0	0.0143 0.00191 0	0.00161 0.000241 0	
$g \geq 10$ 11.2%	0.1059 0.0821 7426	0.013179 0.014815 0.0021816	0.00239 0.00115 0.027	0.01 0.00133 0	0.00607 0.00117 0	−0.00661 0.000992 1	−0.00822 0.00123 1

Table 2.3. For the GNR estimator, the average ΔEntry + Exit Metric for separate periods and the differences between periods. Variable g represents the percent real GDP growth rate.

entry–exit analysis whether these approximately 1,000 productivity estimates are included or not.[12]

Table 2.3 shows the results for the GNR model. Once again, the recessionary years saw a greater increase in productivity due to entry and exit than growth years, by 1.27 percentage points per annum, evidencing Schumpeter's theory. Additionally, the economic boom years saw a greater increase in productivity due to entry and exit than moderate growth years by 0.8 percentage points per annum, which supports the prediction of Melitz's model. In their paper, GNR,

[12]That is, the differences between Tables 2.1 and 2.3 are due to the difference in the production function estimation method and not these approximately 1,000 additional observations.

with a gross output production function, find that the productivity distribution is much less disperse than ACF with a value-added production function. Therefore, the fact that these numbers are attenuated relative to the results in Table 2.1 is unsurprising given that productivity is less disperse in a GNR model than an ACF model.

2.9. Industry growth rates

Instead of focusing exclusively on the business cycle, I also consider industry-specific expansion and contraction phases. Schumpeter's theory, that competitive pressure increases during a recession and thus unproductive plants are forced to exit, can conceivably be extended to apply to industry-specific contractions. The hypothesis of Melitz's model can be likewise analogously extended.

Figure 2.7 shows the same data as Figure 2.6, except that the horizontal axis values have been replaced by industry-specific real growth rates of output. Comparing the two figures, the ΔEntry Metric has lost much of its shape. For the ΔExit Metric graph, the food industry's less variable growth rate pushes its more moderate values towards the center, and the smaller industries play a greater effect on the shape of the fitted curves.

Table 2.4 shows that periods of negative industry growth, on average, saw an increase in productivity by 1.62 percentage points per year over periods of positive industry growth, due only to plant entry and exit. This is about 38 basis points less than the analog for economy-wide growth. Economy-wide busts likely apply more competitive pressure to plants than sectoral busts. When only a single industry sees a large decline in demand, some plants may leave that industry, making room for the remaining plants to survive. During a recession, all industries are affected, and switching industries provides plants little to no relief. Furthermore, credit is likely more available during an industry-specific bust than a recession.

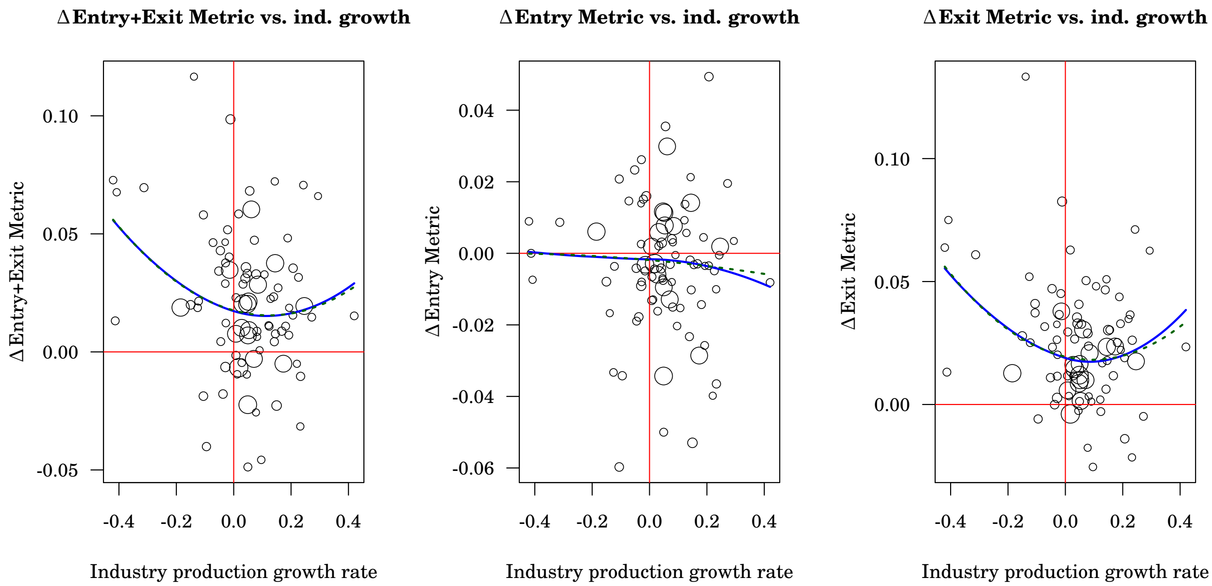

Figure 2.7. Change in the Entry and Exit Metrics versus the industry production growth rate

Periods of industry growth greater than 10% saw productivity increase by 0.49 percentage points per year over periods of industry growth between 0% and 10%. However, this is about 1.4 percentage points less than the analog for economy-wide growth. When the entire economy is booming, there is very high demand for common inputs, such as labor or electricity. However, when only a single industry is booming, the demand for common inputs does not increase as much, so there is less competitive pressure.

Recall that Chile's economy is in transition during the early years of the data set. As shown in Figure 2.2, exit rates are very high in those years. In the previous sections, where I focused on the business cycle, I was forced to include those years since the recession occurred in 1982 and 1983. However, by changing my focus to industry-specific growth rates, I can discard those years and study exclusively 1984 through 1996. This means that the Entry and Exit Metrics are now set to zero for the year 1984, and the first non-zero period for both will be 1985. Therefore, Figure 2.8 is the same as Figure 2.7 without points corresponding to 1979 through 1984.

Range of g			$(-\infty,-10)$	$[-10,0)$	$(-\infty,0)$	$[0,\infty)$	$[0,10)$
$g \in \ldots$ avg. growth	entry rate exit rate plants open	mean bootstrap mean std. error	difference = column − row standard error percentile p-value (one-sided)				
$(-\infty,-10)$ −21.3%	0.0816 0.1226 4110	0.033688 0.033822 0.00090812					
$[-10,0)$ −3.3%	0.1122 0.1203 5431	0.028537 0.028544 0.00050206	0.00528 0.000992 0				
$(-\infty,0)$ −11.1%	0.0988 0.1213 9541	0.030769 0.030831 0.00050859	0.00299 0.000562 0	−0.00229 0.00043 1			
$[0,\infty)$ 9.4%	0.0968 0.0882 32213	0.01462 0.014614 0.00021225	0.0192 0.00086 0	0.0139 0.000472 0	0.0162 0.000453 0		
$[0,10)$ 4.7%	0.0881 0.0878 21450	0.012977 0.012995 0.00018334	0.0208 0.000876 0	0.0155 0.000492 0	0.0178 0.000478 0	0.00162 9.74×10^{-5} 0	
$[10,\infty)$ 19.0%	0.1145 0.0890 10763	0.017903 0.01785 0.00035321	0.016 0.000859 0	0.0107 0.00049 0	0.013 0.000463 0	−0.00324 0.000195 1	−0.00485 0.000292 1

Table 2.4. For the ACF estimator, the average ΔEntry + Exit Metric by industry growth ranges and the differences between ranges. Variable g represents the percent industry growth rate.

Table 2.5 shows that, due to entry and exit of plants, periods of negative industry growth saw an increase in productivity by 1.37 percentage points per year over periods of positive growth. Since the recession is not included, it is reasonable that this is 25 basis points less than when the full temporal range of the data set is used. The figure and the table both show that the effect theorized by Schumpeter can be extended beyond just recessions.

While Figure 2.8 seems to show evidence for the implication of Melitz's model, I do have a concern that it is being partially driven by the leverage point at 42% industry production growth, which corresponds to the textile industry in 1986. As shown in Table 2.5, the average increase

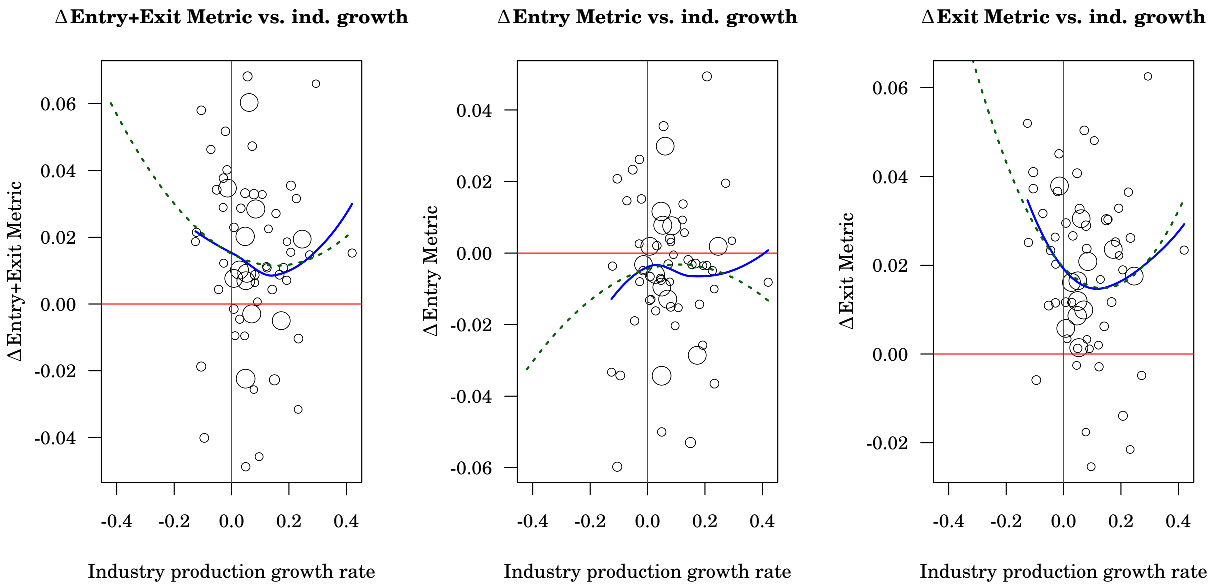

Figure 2.8. Change in the Entry and Exit Metrics versus the industry production growth rate for 1985–96

in productivity per year for industry growth rates greater than 10% is less than the average for growth rates between 0% and 10%.

2.10. Conclusion

There is robust evidence for Schumpeter's theory that recessions are periods of intensified "creative destruction" which cleanse the economy of less productive plants. Particular to the Chilean 1982–83 recession, entry and exit behavior is estimated to have improved average productivity by about 1.4 to 2.4 percentage points per annum over years of moderate economic growth. Outside of the recession, this paper also finds evidence for analogous behavior causing improvements in productivity simply during downturns in specific industries.

Melitz's (2003) model predicts that economic booms will similarly cleanse the economy of less productive plants due to increased competition for inputs. In the three nonconsecutive years Chile experienced real GDP growth in excess of 10%, entry and exit behavior improved average

Range of g			$(-\infty,-10)$	$[-10,0)$	$(-\infty,0)$	$[0,\infty)$	$[0,10)$
$g \in \ldots$ avg. growth	entry rate exit rate plants open	mean bootstrap mean std. error	difference = column − row standard error percentile p-value (one-sided)				
$(-\infty,-10)$ −11.4%	0.1467 0.0891 1196	0.018637 0.01875 0.0015885					
$[-10,0)$ −3.4%	0.1368 0.0999 3914	0.027164 0.02706 0.00074917	−0.00831 0.00163 1				
$(-\infty,0)$ −5.2%	0.1391 0.0974 5110	0.025168 0.025115 0.00073757	−0.00637 0.00125 1	0.00194 0.000382 0			
$[0,\infty)$ 10.4%	0.1041 0.0801 26602	0.011437 0.011432 0.00023834	0.00732 0.00149 0	0.0156 0.000704 0	0.0137 0.000649 0		
$[0,10)$ 5.2%	0.0775 0.0515 16405	0.011668 0.011679 0.00022282	0.00707 0.00156 0	0.0154 0.00078 0	0.0134 0.000749 0	−0.000247 0.000114 0.985	
$[10,\infty)$ 19.0%	0.1149 0.0844 10197	0.010959 0.010923 0.00039096	0.00783 0.00135 0	0.0161 0.000586 0	0.0142 0.000462 0	0.000509 0.000234 0.015	0.000756 0.000348 0.015

Table 2.5. For the ACF estimator, the average ΔEntry + Exit Metric by industry growth ranges and the differences between ranges for 1985–96. Variable g represents the percent industry growth rate.

productivity by about 0.8 to 1.9 percentage points per annum over years of moderate economic growth. The evidence for an analogous effect during industry-specific booms appears a bit lacking.

While the recession's improvement in productivity through entry–exit behavior is higher than the boom years', it is not clear that the effect posited by Schumpeter is stronger than the one predicted by Melitz's model. One must note that the recession's average annual real GDP growth was −8.5%, and it is being compared to moderate economic growth at 6.3%. On the other hand, the average growth during a boom year was 11.2%, and it is being compared to 6.3%, which corresponds to a smaller absolute difference in growth rates.

Regardless, this paper finds evidence for the predictions of both Schumpeter and Melitz, and the dilemma presented by Kehrig (2011) is a false one. Furthermore, the use of dispersion and quantile statistics to assess entry–exit behavior is unnecessary and is likely confounded by other effects as discussed at the end of Section 2.2.

Throughout the specifications of the models of this paper, plant exit has primarily driven the results. This is likely due to the fact that nascent plants, while likely equipped with the latest technology and new equipment, still must train a new workforce, develop routines, and generally experience a degree of learning-by-doing. Further research into the effects of entry and exit behavior could include evaluating the productivity of these plants over a few years, adjusting for survival bias, to better ascertain the effect of their entry.

References

References for Chapter 1

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.

References for Chapter 2

Ackerberg, Daniel, Kevin Caves, and Garth Frazer. 2006. "Structural Identification of Production Functions." Working paper.

Ackerberg, Daniel, Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica,* 83(6): 2411–2451.

Caballero, Ricardo J. and Mohamad L. Hammour. 1994. "The Cleansing Effect of Recessions." *American Economic Review,* 84(5): 1350–1368.

Faggio, Guilia, Kjell G. Salvanes, and John Van Reenen. 2010. "The Evolution of Inequality in Productivity and Wages: Panel Data Evidence." *Industrial and Corporate Change,* 19(6): 1919–1951.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2005. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" NBER Working Paper.

Gandhi, Amit, Salvador Navarro, and David Rivers. 2016. "On the Identification of Production Functions: How Heterogeneous is Productivity?" Working paper.

Greenstreet, David. 2007. "Exploiting Sequential Learning to Estimate Establishment-Level Productivity Dynamics and Decision Rules." Department of Economics, University of Oxford, Discussion Paper.

Kehrig, Matthias. 2011. "The Cyclicality of Productivity Dispersion." Center for Economics Studies, U.S. Census Bureau, Discussion Paper.

Levinsohn, James and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies,* 70(2): 317–341.

Liu, Lili. 1993. "Entry–Exit, Learning, and Productivity Change: Evidence from Chile." *Journal of Development Economics,* 42: 217–242.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica,* 71(6): 1695–1725.

Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica,* 64(6): 1263–1297.

APPENDIX A

More on the Entry and Exit Metrics

A.1. Missing data considerations

As discussed in Section 2.3, there are a number of missing values for several variables, such as real value added and capital. For those observations for which such values are missing, the production function estimation routine cannot provide an estimate for productivity.

Once again, let r_{it} represent the productivity estimate of plant i at time t. Also as before, let $o_{it} = 1$ if the plant operates during time t and $o_{it} = 0$ otherwise. Now, define a new indicator variable for the existence of the productivity estimate. Let $e_{it} = 1$ if the residual exists and $e_{it} = 0$ if it is missing. Naturally, if plant i has a productivity estimate for time t then it operated during time t. That is, $e_{it} = 1 \implies o_{it} = 1$. And similarly, if plant i did not operate during time t, then its productivity estimate is missing: $o_{it} = 0 \implies e_{it} = 0$.

As before, I want to identify the time effect using the change in productivity for persisting plants, which is called the average pairwise difference.

$$\dot{r}_t = \begin{cases} \frac{\sum_i r_{it} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}} - \frac{\sum_i r_{it-1} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}} & t > t_1 \\ 0 & t = t_1 \end{cases}$$

For this new definition, I have swapped o's for e's. To properly capture the time effect, it is important to use only plants that have productivity estimates in both times. Technically, if the calculation was restricted only to plants that operated in both times, the difference may include plants that have productivity estimates in one time and not the other.[1]

[1]Then the word "pairwise" in "average pairwise difference" would not be an appropriate description at all.

Adjusted productivity can be defined as before, and average adjusted productivity simply replaces o's for e's:

$$\tilde{r}_{it} = r_{it} - \sum_{\tau=t_1}^{t} \dot{r}_\tau$$

$$\tilde{r}_{\cdot t} = \frac{\sum_i \tilde{r}_{it} e_{it}}{\sum_i e_{it}}$$

However, now differences in adjusted productivity will not only capture the effect of plant entry and exit, but also the effect of persisting plants switching between having productivity estimates and not. To see this, decompose the difference in average adjusted productivity as before:

$$\Delta \tilde{r}_{\cdot t} \quad = \quad \frac{\sum_i r_{it} e_{it}}{\sum_i e_{it}} - \frac{\sum_i r_{it} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}} + \frac{\sum_i r_{it-1} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}} - \frac{\sum_i r_{it-1} e_{it-1}}{\sum_i e_{it-1}}$$

Note that the first term can be split into three parts: plants that persist and have productivity estimates in $t-1$, plants that persist but are missing productivity estimates in $t-1$, and plants that are entering.

$$\frac{\sum_i r_{it} e_{it}}{\sum_i e_{it}} = \frac{\overbrace{\sum_i r_{it} e_{it} e_{it-1}}^{\text{persist \& estimates in } t-1} + \overbrace{\sum_i r_{it} e_{it} o_{it-1}(1-e_{it-1})}^{\text{persist \& no estimates in } t-1} + \overbrace{\sum_i r_{it} e_{it}(1-o_{it-1})}^{\text{entering}}}{\sum_i e_{it} e_{it-1} + \sum_i e_{it} o_{it-1}(1-e_{it-1}) + \sum_i e_{it}(1-o_{it-1})}$$

An analogous equation can be found for the term containing the productivity estimates of the plants that exit.

Since the difference in average adjusted productivity contains unwanted terms, there is no need to use it exactly. However, now that the desired terms have been identified, I define the ΔEntry Metric as:

$$\Delta \text{Entry Metric}_t = \frac{\sum_i r_{it} e_{it} e_{it-1} + \sum_i r_{it} e_{it}(1-o_{it-1})}{\sum_i e_{it} e_{it-1} + \sum_i e_{it}(1-o_{it-1})} - \frac{\sum_i r_{it} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}}$$

This is the moments estimator for the following difference in conditional expectations:

$$E[\omega_{it} + \varepsilon_{it} \,|\, e_{it} = 1 \wedge (e_{it-1} = 1 \vee o_{it-1} = 0)] - E[\omega_{it} + \varepsilon_{it} \,|\, e_{it} = 1 \wedge e_{it-1} = 1]$$

Similarly, the ΔExit Metric is defined as:

$$\Delta\,\text{Exit Metric}_t = \frac{\sum_i r_{it-1} e_{it} e_{it-1}}{\sum_i e_{it} e_{it-1}} - \frac{\sum_i r_{it-1} e_{it} e_{it-1} + \sum_i r_{it-1} e_{it-1}(1 - o_{it})}{\sum_i e_{it} e_{it-1} + \sum_i e_{it-1}(1 - o_{it})}$$

It is the moments estimator for:

$$E[\omega_{it-1} + \varepsilon_{it-1} \,|\, e_{it} = 1 \wedge e_{it-1} = 1] - E[\omega_{it-1} + \varepsilon_{it-1} \,|\, e_{it-1} = 1 \wedge (e_{it} = 1 \vee o_{it} = 0)]$$

Because the Entry Metric and the Exit Metric are measures of difference in average productivity relative to the same set of persisting plants, they can be naturally added together while retaining their meaning.[2]

A.2. Alternative counterfactual assumptions

One of the counterfactual assumptions I use is that exiting plants, had they not exited, would have productivity equal to their previous productivity plus the average pairwise difference. The other counterfactual assumption is that when a plant enters or exits, the productivity average of the plants that persisted is unchanged.

It might be noted that there is something of an asymmetry here: exiting plants have an additional assumption tied to them that entering plants do not. There are, in fact, four basic counterfactual scenarios that could be considered:

[2] If I had defined the metrics as

$$\Delta\,\text{Entry Metric}_t = \frac{\sum_i r_{it} e_{it}}{\sum_i e_{it}} - \frac{\sum_i r_{it} e_{it} o_{it-1}}{\sum_i e_{it} o_{it-1}} \qquad \Delta\,\text{Exit Metric}_t = \frac{\sum_i r_{it-1} e_{it-1} o_{it}}{\sum_i e_{it-1} o_{it}} - \frac{\sum_i r_{it-1} e_{it-1}}{\sum_i e_{it-1}}$$

while they would individually still have reasonable interpretations, they would not be comparable as the ΔEntry Metric's subtrahend does not generally equal the ΔExit Metric's minuend.

(1) Entering plants do not enter and exiting plants do not exit.

(2) Entering plants have always existed and exiting plants do not exit.

(3) Entering plants do not enter and exiting plants never existed.[3]

(4) Entering plants have always existed and exiting plants have never existed.

My study uses the first counterfactual scenario. In it, I need to assign counterfactual productivity levels to plants that exit, hence the need for the singular assumption regarding exiting plants.

There are a number of alternatives to the particular assumption regarding the counterfactual productivity of plants that exited. For example:

- Instead of using all available observations to construct the pairwise differences, I could restrict myself to using only observations of plants that persist throughout the entire sample to construct the pairwise differences used to define adjusted productivity.
- Instead of using the average pairwise difference to construct adjusted productivity, I could use the median pairwise difference. This would be equivalent to minimizing the sum of absolute pairwise differences.
- Instead of using pairwise differences, I could assign counterfactual productivity levels to exiting plants by looking at persisting plants with similar labor, capital, and/or productivity levels. This would involve regressing next-period productivity on current labor, capital, and productivity and using that regression to predict counterfactual next-year productivity for exiting plants.

I chose the method I did because it uses all available observations to construct the average pairwise difference, it is mathematically parsimonious, and it makes the Entry and Exit Metrics more comparable and interpretable.

[3]This would just be a study of plants that only persist.

APPENDIX B

Model without Capital

The Chilean data set contains a potentially serious deficiency for the study of entering plants: capital is missing for many plants that enter after 1981. For 1979–86, the census only required the reporting of fixed asset values in 1980 and 1981. Starting from those fixed asset values, Liu (1992) recursively constructed the capital series of the data set using investment numbers and assumed depreciation rates. Thus, for plants exiting in 1980 and plants entering after 1981, there is often no capital data.

To address this issue, I consider an alternative model with energy usage in the place of capital. Energy consumption is correlated with capital, both statistically and theoretically. Greenstreet (2007) uses this idea to develop a capital services series.

Where M_{it} is real materials usage, and S_{it} is real services usage, let:

$$V_{it}^{ms} = Y_{it} - M_{it} - S_{it}$$

and

$$\mathscr{M}_{it}^{ms} = M_{it} + S_{it}$$

Then the capital-less value-added production function is, in log terms:

$$v_{it}^{ms} = \beta_l l_{it} + \beta_e e_{it} + \omega_{it} + \varepsilon_{it}$$

where e_{it} is the log of real energy usage. I apply the method described in Section 2.4, using $\mu_{it}^{ms} = \log(\mathscr{M}_{it}^{ms})$ as a proxy for ω_{it} and estimating β_l and β_e against instruments l_{it-1} and e_{it-1}.

Year	Entering			Exiting			Persisting (both years)		
	capital	energy	number	capital	energy	number	capital	energy	number
1980	0.7926	0.9481	135	0.5335	0.9754	448	0.8656	0.9621	2694
1981	0.4	0.9571	140	0.7866	0.9589	389	0.882	0.9611	2440
1982	0.3821	0.9919	123	0.8235	0.9721	323	0.8498	0.9459	2257
1983	0.4824	0.9647	170	0.7643	0.936	297	0.8296	0.9462	2083
1984	0.6126	0.9702	302	0.7059	0.9638	221	0.8504	0.9675	2032
1985	0.672	1	125	0.6875	0.9861	144	0.8489	0.9776	2190
1986	0.7055	0.908	163	0.7542	0.9915	236	0.8413	0.9567	2079

Table B.1. The fraction of extant residuals for entering, exiting, and persisting plants, for the model with capital and the model without

Note that the energy instrument has to be lagged. The model presented in Section 2.4 used capital at the beginning of the year t as an instrument. However, energy at time t cannot be assumed to be uncorrelated with the innovation in ω_t since energy usage is as flexible as materials and services usage.

Table B.1 shows the fraction of observations for which I can calculate residuals in both models. For example, in 1981, 140 plants entered. In the model that uses capital, I could calculate residuals for only 40% of those plants. With this energy-substitution model, I can calculate residuals for 95.71% of the plants. A large gain in the number of residuals that I can calculate is also seen for plants exiting in 1980.

The residuals of the energy-substitution model are highly correlated with residuals of the model with capital included, which suggests that this model is a reasonable replacement considering it swaps out one of two explanatory variables. Out of the five industries studied, the minimum Pearson correlation between the model with capital and the model without is 0.85. I considered a number of alternative formulations of a capital-less model, such as including services with energy instead of with materials, using v_{it} instead of v_{it}^{ms} as the dependent variable, and using services instead of energy. This model provided the reasonably best performance across the five industries as measured by Pearson, Spearman, and Kendall correlations. I considered maintaining the rank order of observations in the residual distribution very important, hence my use of Spearman and Kendall correlations, which as nonparametric statistics, consider rank alone.

Table B.2 shows the results of the model. There are some differences compared to the model with capital. For example, the change in the sum of the Entry and Exit Metrics in 1983 is smaller by about 1 percentage point. The recessionary years saw only a 1.2 percentage point per annum increase in productivity due to entry and exit over years of GDP growth. This is about 0.8 percentage points less than the model with capital.

However, the economic boom years saw productivity improve by 1.84 percentage points per annum due to entry and exit over years of moderate growth. This closely matches the result in the model with capital, with the effect diminished by only 0.02 percentage points.

Category			1982	1983	1982–83	$g \geq 0$	$0 \leq g < 10$
category	entry rate	mean	difference = column – row				
avg. growth	exit rate	bootstrap mean	standard error				
	plants open	std. error	percentile p-value (one-sided)				
1982	0.0545	0.035929					
−13.4%	0.1252	0.036127					
	2380	0.001615					
1983	0.0816	0.025563	0.00948				
−3.5%	0.1248	0.02665	0.00158				
	2253	0.0018392	0.001				
1982–83	0.0675	0.030839	0.00465	−0.00482			
−8.5%	0.1250	0.031474	0.000774	0.000802			
	4633	0.0015388	0.001	0.999			
$g > 0$	0.1011	0.019381	0.0168	0.00733	0.0122		
7.3%	0.0921	0.019321	0.00229	0.00244	0.00223		
	37121	0.0011363	0	0.001	0		
$0 < g < 10$	0.0999	0.015678	0.0205	0.011	0.0158	0.00369	
6.3%	0.0945	0.015631	0.0023	0.00248	0.00226	0.000122	
	29695	0.0012111	0	0	0	0	
$g > 10$	0.1059	0.034153	0.00208	−0.00739	−0.00257	−0.0147	−0.0184
11.2%	0.0821	0.034045	0.00232	0.00237	0.00221	0.000488	0.00061
	7426	0.00094264	0.155	0.994	0.949	1	1

Table B.2. For the ACF estimator without capital, the average ΔEntry + Exit Metric for separate periods and the differences between periods. Variable g represents the percent real GDP growth rate.

APPENDIX C

The Distinctiveness of the Food Industry

As seen in Figure 2.1, the food industry generally has about as many plants as the other four industries combined. This gives it a tremendous amount of weight in the calculation of average productivity for all the models in the paper except the weighted model in Section 2.7. However, that particular model is weighted by (expected) real value added, and the food industry generates about 32% more real value added (and 70% more real output) than the other industries combined. So in that model too, the food industry is weighted very heavily.

As opposed to the other four industries (textiles, apparel, wood products, and metal products), the food product industry saw a decline in the plant exit rate in 1982, the start of the recession. This is illustrated in Figure 2.2. These facts point to the need for a bit further study into the distinctiveness of the food industry.

Why was the food industry's exit rate unaffected by the recession? As shown at the top of Figure C.2, the food industry's production declined relatively less than other industries. One reason for this is that food is a consumer staple, consumption of which is less cyclical than the other more durable goods produced by the other four industries.[1] In Figure C.1, domestic food product consumption in 1982 fell relatively less than the other industries' products. Table C.1 shows the average income elasticity of consumption for the products of each industry. It also

[1]The consumption quantity is calculated as domestic production minus exports plus imports. The source of the export and import data is the Commodity Trade and Statistics Database (Comtrade), compiled by the United Nations Statistics Division. The early years of the Comtrade data for Chile were classified by SITC (Standard International Trade Classification) Revision 1. Starting from Revision 1 data for all years (1979–96), I converted the data to SITC Revision 2 using a conversion table produced by Robert Lipsey. I then converted from SITC Revision 2 to ISIC Revision 2 by way of Marc-Andreas Muendler's conversion table. However, those tables alone were insufficient to capture all the relevant Comtrade data; I had to make several modifications to them.

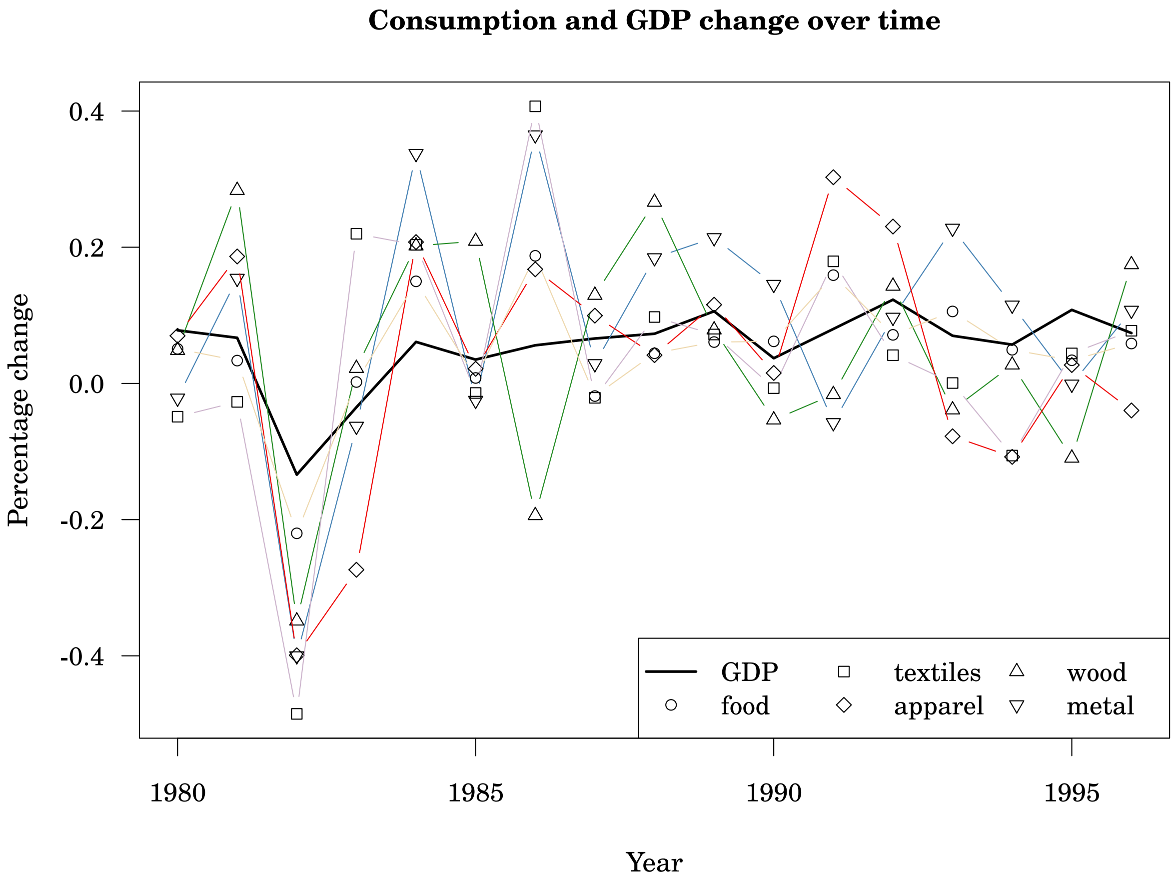

Figure C.1. Percent changes in consumption and GDP over time

presents the β's of a model similar to the Capital Asset Pricing Model (CAPM). Where j indexes the industries and $\%\Delta$ represents percentage change, this linear model is:

$$(\%\Delta\,\text{consumption}_{jt}) = \alpha_j + \beta_j(\%\Delta\,\text{GDP}_t) + \varepsilon_{jt}$$

Chile also exports a fair percentage of its food product production, as shown at the bottom of Figure C.2. During the recession, in the face of decreased domestic demand, Chile's food industry was able to increase exports, unlike most other industries. The only other industry to increase exports throughout the recession was the apparel industry, but exports made up a very small fraction of their total sales in the years around the recession.

Industry	food	textiles	apparel	wood	metal
income elasticity	1.02	0.60	1.61	1.16	1.99
CAPM-like β	1.06	1.56	2.38	1.46	1.92

Table C.1. The responsiveness of consumption to changes in GDP

Thus, the food industry was less affected by the recession than the other industries for two reasons. First, it produces a consumer staple, for which consumption is generally less elastic than the other industries. Second, it was able to partially make up for the decline in domestic consumption by increasing sales to the external sector, which helped insulate it from the increased competitive pressure felt by the other industries.

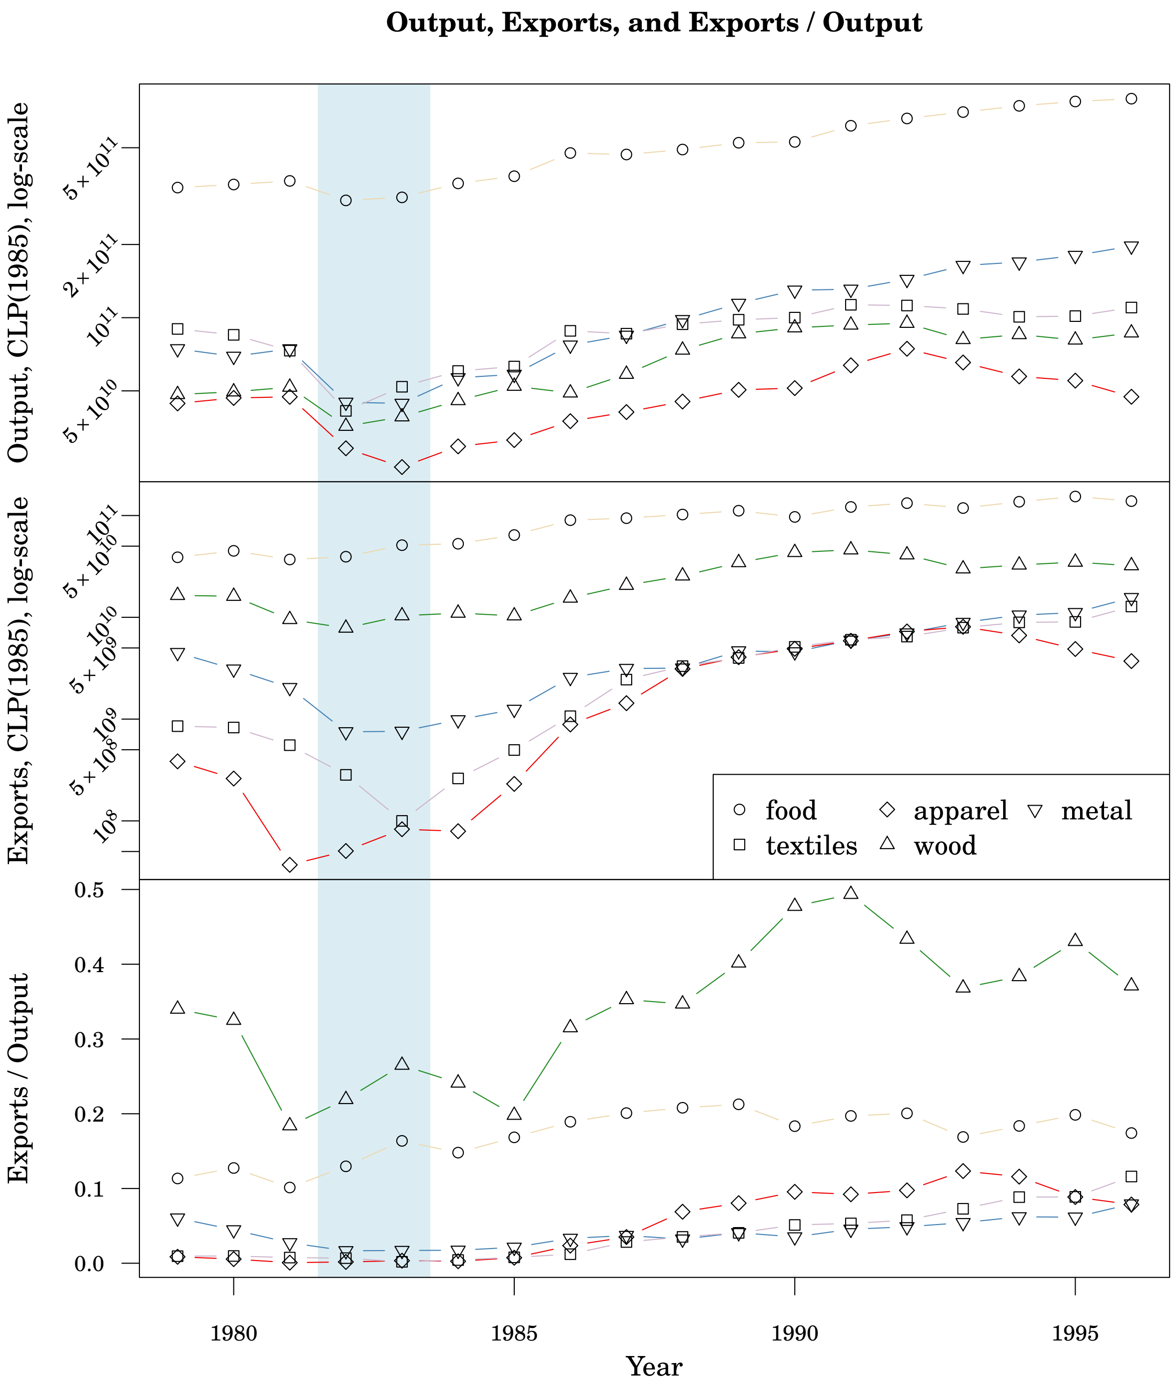

Figure C.2. A stack of three graphs. The top two are real output and real exports, measured in Chilean pesos (CLP) at 1985 prices, plotted on the common logarithm scale. The bottom graph is the fraction of real output exported. The recession in 1982 and 1983 is highlighted.