Using differences across US states to think about consumption

Charles Grant

September 2001

A Thesis Submitted for the Degree of Doctor of Philosophy to University College London, the University of London.

Contents

List of Tables

List of Figures

Abstract

This thesis investigates some contemporary issues in consumption using household data. It exploits state of residence information available over several years using the Consumer Expenditure Survey: a US survey of spending, and a variety of other household characteristics. The thesis contains three distinct studies. The first looks at how consumer bankruptcy rules affects the debt holdings, and consumption behaviour of US households. Harsher punishment results in more debt but less smoothing. The second study looks at how differences in state taxes translates into differences in the ability of agents to share the idiosyncratic component of their income shocks, finding that making taxes more re-distributive reduces agents ability to insure risks. The last study accepts that some agents are credit-constrained, and recovers estimates of the supply of, and the demand for, credit. This leads to estimates of the proportion of agents credit constrained, 28%, of how agents differ, and of how much more agents constrained agents wish to borrow.

Acknowledgements

I would like to thank the many people whoe comments and advise helped to improve this work. These included Orazio Attanasio, Erich Battistin, Gabriella Berloffa, Nick Bloom, Chris Carroll, Vassilis Hajivassiliou, Christos Koulovatianos, Tullio Jappelli, Nick Souleles, Mario Padula, Luigi Pistaferri, Ian Preston, Oved Yosha as well as several seminar participants. All remaining errors, are, of course my own.

Declaration

No part of this thesis has previously been presented to any University for any degree.

Charles Grant

Chapter 1

Introduction:

This thesis explores some contemporary issues in consumption. Modern work on consumption started with the seminal paper by Hall (1978).[1] This paper assumed quadratic and time separable utility functions and explored the implications of the permanent income hypothesis. He developed an Euler equation that he then tested, and rejected, using US aggregate data. Consequently, much of the next 15 years was spent testing various developments of this basic formulation. Nevertheless, subsequent studies confirmed Hall's result: changes in income are over-sensitive to predictable changes in income, the excess sensitivity highlighted by Flavin (1981), and changes too little in response to unpredictable income shocks, the excess smoothness of, for instance Campbell (1987). Simple life-cycle models of consumption have also been rejected: Carroll and Summers (1991) shows how even over the whole life-cycle that income and consumption seem to follow each other.

A number of explanations have been offered. Many studies have suggested alternative utility functions, such as iso-elastic or log utility, and have allowed life-cycle or preference parameters to enter the utility function. If such utility functions are chosen, then aggregate data is no longer appropriate. Attanasio and Weber (1993) showed that using micro-data

[1] A much fuller introduction to consumption is contained in Deaton (1992).

and explicitly controlling for taste shifters at the household level could account for much of the variation in consumption. Nevertheless, it is not believed that tastes-shifters, such as family size, can itself explain all of the rejection of the Euler equation first highlighted by Hall. Instead two other possible explanations have been popular. One idea is that if the agent is risk-averse, then there will be a precautionary motive to saving. Indeed, if agents are risk-averse, prudent and cautious, in the sense described by Zeldes (1989), then in a remarkable paper Carroll (1997) demonstrated that this itself *could* cause income and consumption to track together over the life-cycle. The second popular explanation is that at least some households are sometimes credit-constrained. Households would like to borrow against their future income, and thus smooth consumption, but for some reason they are not allowed to borrow as much as they would like at the prevailing interest rate. In practise, this different explanations are difficult to convincingly distinguish, at least with the kind of data consumption economists have hitherto used. The problem is that the different explanations have few, if any, observational implications that differ between the explanations.

Much of the more recent consumption literature has instead looked at risk-sharing across states of nature. One implication the complete market hypothesis is that agents should be able to pool their risks, and thus only aggregate risk, and not idiosyncratic risk, should enter into changes in consumption. This was first tested by Mace (1991), and, as might be expected, was rejected. Nevertheless, the study of risk-sharing has given some important insights. In particular, the theoretical literature, in such papers as Kehoe and Levine (1993) or Kocherlakota (1997) has stressed that risk-sharing is limited since agents will only credibly commit to sharing their income when they have good draws from nature if they are sufficiently compensated, or face an environment that is sufficiently risky, when they have bad income draws. Otherwise, they will not lend regardless of the interest rate. In other words, such models endogenously create credit-constraints, and such constraints are intimately related to the punishment a debtor faces if he defaults on any debt contract. This motivates the first study in the thesis. This chapter investigates how the default rules, that is the rule about how the debtors assets are shared between the borrower and lender when the debtor defaults, affects the amount of debt that debtors hold. Intuitively, one would expect that harsher rules will allow more debt to be held in equilibrium: a result confirmed by this study. One problem is that it then seems optimal to punish default with an arbitrarily large punishment. One possible explanation of arbitrarily large punishments are not observed might be that allowing default in *ex post* bad income realisations may insure the agent against bad income shocks. Hence this study also looks at the insurance effect of the default rules, and finds weak evidence to support the proposition that harsher rules provide less insurance since more of the risk of bad income realisations is borne by the debtor in this case.

The problem of risk-sharing is more directly analyzed in the fourth chapter. So far the empirical literature has noted that full insurance can be rejected by the data. However, full insurance seems a wholly implausible hypothesis, and it is hardly surprising that this model is rejected. However, the tests that have been constructed are not fully constructive: they do not help us to understand either how much insurance is available to agents, or what factors may contribute to the overall level of insurance. Chapter 4 concentrates on exactly this problem. It will formulate a measure of risk-sharing that can be related to some policy instrument. The chapter attempts to explain how the amount of extra insurance a particular policy instrument gives. The policy instrument that the chapter investigates is the tax system: the study measures how much extra risk-sharing occurs when taxes become more redistributive, i.e. for a fixed total tax revenue, how the distribution of tax liabilities between rich and poor changes the extent to which agents can smooth consumption against states of nature. As will be seen, the study reaches a surprising conclusion.

Chapter 5 takes as given that at least some agents are credit constrained. Chapter 3 found that borrowing was related to the punishment for default, hence this might not be too unreasonable. If some agents are credit constrained then observed borrowing is the minimum of desired borrowing (demand) and the maximum level of borrowing that any lender will allow (supply). Only this minimum is generally observed and it is not usually known whether the agent wished to borrow more than this observed amount. Chapter 4 explains how to separately recover both supply and demand as long as exclusion restrictions can be made: things that enter supply and not demand, and things that enter demand and not supply. The plausibility of the results depends crucially on the ability to find appropriate exclusion restrictions. The approach, as will be explained, builds on the literature on disequilibrium models.

The main body of work is contained in chapters 3 to 5. Most consumption literature has focused on changes over time, either looking at aggregate time series, or looking at synthetic cohorts in which identification is through a large number of time series. Instead this thesis will exploit differences across the environment in which agents reside. In fact, it is difficult to know how the issues addressed in this thesis could convincingly be studied in any other way. The difference that is exploited is that different agents reside in different US states. These US states differ in the their regulations on bankruptcy regulations (exploited in chapter 3), in their tax systems (exploited in chapter 4), and in their banking regulations (exploited in chapter 5). However, unlike a comparison across countries, states and the individuals in the states are still fairly homogeneous. For instance, the law regarding bankruptcy only differs between states in regulating which assets may be kept when an agent becomes bankrupt: all other regulations regarding bankruptcy are fixed across states. It also seems reasonable to suppose attitudes towards risk, and towards default are similar across US states in a way that they might not be across, say, OECD countries.

This thesis is able to exploit differences across states since it uses the Consumer Expenditure Survey. State data is available for 1980-1998 using this data set, something that is fairly unique for data. Chapter 2 starts with a brief description of this data set. Chapter 3 studies bankruptcy regulations and their effect on borrowing and consumption smoothing. Chapter

4 looks and risk-sharing in more detail, and investigates whether the tax system helps agents to smooth against their idiosyncratic risks. Chapter 5 looks at credit constraints, and the work is briefly summarised and discussed in chapter 6.

Chapter 2

Data Description:

Since the empirical research contained in the different chapters of this thesis exploits the US Consumer Expenditure Survey (henceforth the CEX), this chapter gives a brief explanation, and discussion of those aspects of the data that are common between the different chapters.[1] The CEX is a consumer survey of households that is available on a continuous basis since 1980. The survey is conducted by the Bureau of Labor Statistics in the US and was originally designed to construct a measure of inflation. As a result a large number of households, around 7,000 each quarter, are asked to respond to extremely detailed questions about current spending, as well as being asked about a variety of demographic and other household characteristics. Households are interviewed 5 times; the first being a contact interview (from which no information is disclosed); and four subsequent interviews at 3 monthly periods. The survey is constructed as a rotating panel, hence each quarter, one fifth of the households having reached their fifth interview drop out and are replaced. In each interview households recall their household expenditures in each of the previous three months. However, since many households, for many expenditure items, have assigned to them the same expenditure to each of the three months, only the last of these three months will ever be considered.

[1] Attanasio and Weber (1993) also contains a detailed description of the data.

Information is also recorded on income, although (unless the householder has changed employment) this information is only asked in the second and fifth interview. A large number of demographic and other household characteristics are recorded, including state of residence, although for reasons of confidentiality this information is sometimes suppressed. The fact that state information is available will be consistently exploited in this thesis.

From 1988 the survey has also included additional information on the households debts in the second and fifth interviews. This information is used to construct the total unsecured debts held by the household, including debts held in revolving credit accounts (including store, gasoline, and general purpose credit cards), in installment credit accounts, credit at banks or savings and loan companies, in credit unions, at finance companies, unpaid medical bills, and other credit sources. It also includes negative balances held in checking or brokerage accounts. Excluded from the total are mortgage, and other secured debts. This data will be used in two of the chapters of this thesis.

Throughout this work, income, debt and exemption values will be deflated by the CEX price index so that they are in real terms. The price index is constructed as a Stone-Geary price index for individual households. This work will consistently exclude farming and self-employed households. Also excluded are large households with eight or more members, and households in which the respondent answers that they have received no education. Large households are excluded because as the household becomes larger it becomes increasingly problematic to describe the household's characteristics by the characteristics of the household head. Those households with heads who did not go to school may have poorer responses to the survey. In any case, the number of households in these two categories was small. Other excluded observations will be highlighted in the chapter for which the exclusion is relevant.

Using the CEX has a number of advantages vis-a-vis other possible data sources. Unlike the PSID, for instance, a large part of the households consumption can be reliably calculated. The PSID, in contrast, only measures food consumption. This can be a problem if food consumption varies much less than total non-durable consumption in response to income shocks. Furthermore, the CEX surveys many more households. The disadvantage is that the furthest apart any two observations can be is 9 months. A second problem is that household responses are based on recall data, rather than a diary (the construction of the British Family Expenditure Survey, for instance). Never-the-less, it is not believed that this is too serious a drawback.

The last and most important advantage is that the CEX contains information on state of residence. Coupled with the large sample size, this has allowed this work to exploit the differences between US states to measure some of the consumption effects that could not be assessed in any other way. The only other data source for much of the information used in this thesis is the Survey of Consumer Finances (SCF) which contains information on household assets and debts that will be exploited in two of the chapters of this thesis. One problem with this survey is that state information is only contained in the 1983 wave of the data. A further problem is that that data set has many fewer observations. Overall, the CEX seems to be the dataset that can best be used to explore some of the issues that I hope to address in the following chapters.

Chapter 3

Bankruptcy Law, Credit Constraints and Insurance: Some Empirics

Abstract

Bankruptcy (defaulting on one's debts) acts as insurance if it allows default in cases of negative income shocks. However, whether bankruptcy provides insurance depends on the bankruptcy rules (the punishment for default) that is enforced. Bankruptcy rules can instead cause the consumer to be credit constrained. If debts are not fully enforceable, a rational lender may limit how much debt any borrowers are allowed to hold. This limit increases as the punishment for defaulting increases. The US provides a natural test of the theory since rules about which assets may be kept by the debtor, the state exemptions, when filing for bankruptcy differ dramatically across the different states. Regressions show that increasing the level of these exemptions causes less debt to be held by consumers. The chapter also tests the theory more indirectly by regressing changes in the level, and in the variance, of consumption, which suggests that bankruptcy provides insurance.

3.1 Introduction

In recent years, a great deal of attention by consumption economists has been devoted to the observation that consumption and income seem to follow each other, both over the comparatively short intervals[1] and over the whole lifecycle. Consumers seem to consume more in the middle of their life, in their 40's and 50's, than either at the beginning or at the end of their life. Carroll and Summers (1991) have shown how income and consumption seem to track each other over the life-cycle. Several explanations have been suggested in the literature, two of the most popular are: (1) households are risk averse, prudent and impatient in the sense of Zeldes (1989) and Carroll (1997); and (2) households are credit constrained and can not borrow, Deaton (1991). However, it has been difficult to distinguish the relative importance of these different explanations. This chapter will use bankruptcy legislation as an instrument that can shed light on these theories.

In a parallel literature, see Jackson (1986), bankruptcy rules have been motivated, particularly among lawyers, as a device that creates insurance when agents face uncertainty about the future. At the same time a theoretical literature has attempted to explain the fact that consumers can not fully insure all their idiosyncratic risk. Papers such as those by Kehoe and Levine (1993) and Kocherlakota (1996) argue that this limited insurance is due to the fact that debts can not be fully enforced.[2] In such a framework, as shall be shown in section 3.2, the presence of a bankruptcy law which limits the punishment for default, may instead create credit-constraints. One aim of this chapter is to bring together these two literatures, and to empirically test what effect bankruptcy rules may have.

In the literature on risk-sharing, any mean preserving action that reduced uncertainty

[1] See, for example, Hall (1978)

[2] For an earlier literature on credit constraints see, for instance, Stiglitz and Weiss (1981). In that literature credit constraints arise since agent's types are imperfectly observed, and different types have different propensities to default.

will be welfare improving. Bankruptcy legislation can reduce uncertainty if the consumer can default on his debt when his income is low. For bankruptcy legislation to act as insurance, actually defaulting must be negatively correlated with income.[3] Bankruptcy legislation can have very different or even perverse effects if this is not the case. Section 3 starts with a very simple discussion of how the penalty[4], or sharing rule (how much the creditor and the debtor each receive when the debtor defaults) affects the debtor's incentives to default.

The punishment for debt differs across the different states of the United States quite substantially, as, when defaulting, different levels of assets can be kept in different states. Borrowers are assumed to be otherwise identical, and lenders face no constraints as to which state they will lend in. An identifying assumption of the paper is that any other differences in the operation of credit markets across the different states is orthogonal to the bankruptcy exemptions. This allows the theory to be tested by comparing the level of debt held by households in the different US states. The level of debt should be systematically related to the level of assets that may be kept in bankruptcy. The empirical section investigates some of the implications. The first part of this section uses an approach similar to that of Gropp, Scholz and White (1997). However, their study is limited to a single cross section as, they use the Survey of Consumer Finances for which state data is only available in 1983. In contrast, this chapter is able to exploit data changes over time as well as across states. This allows us to potentially control for state specific effects that might be correlated with the bankruptcy legislation.

This section also reports results for consumption growth, which is a more direct test of consumption smoothing, at least for the ability to smooth relatively high frequency events, and for the change in the variance of consumption, which, as will be explained, is a direct

[3] Or, more generally, whatever the consumer faces uncertainty about.

[4] In general the punishment could be losing a portion of their assets; being denied any credit for a period afterwards; and perhaps losing (or having garnished) some of their future income. There may also be a social stigma attached to default.

test for the extra insurance induced by the bankruptcy rules. This part directly tests the claim that bankruptcy rules are providing insurance.

The chapter is organized in the following way. Section 2 expounds the theory stated above. In section 3 a brief account of the rules in personal bankruptcy as they pertain to the United States is given. Section 4 contains a description of the data. In section 5 there is a description the regression results, and the chapter concludes in section 6.

3.2 Theory

One of the suggested explanations for why consumption follows income over the life-cycle, is that consumer are risk-averse, impatient and cautious in the sense outlined by Zeldes (1989). If agents were risk-averse then anything that reduced uncertainty would be welfare improving: this could motivate bankruptcy legislation. If, for some reason, a contingent claims market in which consumer could insure themselves against bad income draws did not exist, then a bankruptcy rule could imitate some of the useful features of such a market. Bankruptcy legislation can act as insurance since it allows consumers with low income draws to default on their debt. To illustrate these ideas consider the following discussion.

Suppose the consumer lives for two periods, but second period income is uncertain and drawn from some distribution $y_2 \in \Gamma$. (Suppose that the moments of y are bounded and the utility function is strictly increasing and strictly concave in all its arguments and continuously differentiable.) Then uncertainty about future income causes the consumer to reduce consumption in period 1 and we can write (ignoring higher moments):

$$c_1 = c_1 \left[y_1, E\left(y_2\right), var\left(y_2\right) \right] \tag{3.1}$$

Consumption in period 1 is increasing in the first two arguments and, if agents are risk-averse, falling in the third. Increasing the variance of period 2 income reduces period 1 consumption, and thus also the level of borrowing at the end of period 1 since assets evolve according to the equation:

$$A_2 = (1+r)(y_1 - c_1) \tag{3.2}$$

Suppose the consumer could default on his debt if it were larger than some critical level. If the bank operates in a competitive environment, then it will make zero profits. The banks zero profit condition is:

$$\int_{default} q(y_t, A_t)\ dy_t + \int_{no-default} \frac{1+r}{1+r^f} A_t\ dy_t = A_t \tag{3.3}$$

Here r^f is the risk free rate and $q(\cdot)$ is the 'punishment' in the event of default: it is the amount that the bank can make the consumer pay when he defaults on his debt.[5] In this model, assuming the interest rate is small, the extra interest rate paid $r - r^f$ is exactly that needed to offset the loss the bank makes when the consumer defaults. It is conceptually the same as an insurance premium. If at least some debt will be held, so that $A_2 < 0$, then second period wealth, allowing for default, can be defined as:

$$W_2 = \begin{cases} y_2 - q & \text{default} \\ y_2 + \frac{1+r}{1+r^f} A_2 & \text{no default} \end{cases} \tag{3.4}$$

Define $\widehat{y}$ in the following way:

$$\widehat{y_2} = \begin{cases} y_2 - q - A_2 & \text{default} \\ y_2 + \frac{r-r^f}{1+r^f} A_2 & \text{no default} \end{cases} \tag{3.5}$$

Clearly $q(\cdot) \in [0, y_2]$, while it is optimal for the consumer to default if and only if $q < -A_2$. The consumer would be indifferent between receiving y with default allowed, or receiving $\widehat{y}$ with default not allowed. Remembering $A_2 < 0$, when default occurs $\widehat{y} > y$ while $\widehat{y} < y$ when

[5] This formulation implicitly imposes that the bank is risk neutral. More generally, the qualitative arguments hold as long as the bank is less risk-averse than the consumer. For simplicity it will also be assumed that there is no deadweight loss: the bank receives what the consumer pays.

the consumer does not default. If default happens when income is low then $var\,(\widehat{y}) < var\,(y)$ and so allowing default acts in the same way as compressing the distribution of income.[6] This will increase both period 1 consumption and the level of debt (A_2 falls). In period two, consumption is higher when default occurs, and lower when it does not. Overall, allowing default is unambiguously welfare improving since expected lifetime utility has increased.

The possibility of default acts as insurance since in low income states the consumer does not have to repay any debts. The bank bears the risk of low income realisations rather than the consumer. Crucial to this argument is that default occurs when income is low as insurance only happens when default is negatively correlated with income. If this is not true then any bankruptcy rule will not act as insurance. It is essentially trivial to devise rules where this is true. However, consider the following simple example where this is not true.

Example

Consider a consumer who lives for two periods and maximises utility over two goods; a durable d, that depreciates at rate α and a non-durable good c. The price of the non-durable good is normalised to one, while the price of the durable good is p. Income and consumption are as before. The consumer (uniquely) chooses his first period consumption bundle (c_1, d_1) which also defines his level of assets at the beginning of period 2:

$$A_2 = (1 + r)\,(y_1 - c_1 - p\,d_1) \tag{3.6}$$

In the second period the consumer realises y_2 which defines his second period consumption bundle (c_2, d_2). That is, in the second period, period two wealth W_2 is distributed over the two goods. Now consider the following bankruptcy rule. Suppose the punishment consisted of having the durable good, in excess of some exempt level E, seized and sold. Once the

[6]The banks no-profit condition ensures that $E(\widehat{y}) = E(y)$. Implicit in this statement is that income is exogenous, and that there are no moral hazard issues.

debt has been repaid in full, the consumer can retain any remaining value of the durable good. That is:

$$q_2 = \min\left[A_2, \max\left(\alpha\, p\, d_1 - E, 0\right)\right] \tag{3.7}$$

Given this framework, it is optimal to default if $W_2\,(default) > W_2\,(repay)$. Thus the consumer will default if the following holds:

$$y_2 + E > y_2 + \alpha\, p\, d_1 + A_2 \tag{3.8}$$

Clearly it does not make sense to default if the debt can be fully enforced, or if $A_2 > 0$, so assume that neither of these is true. In which case the consumer will default if:

$$-A_2 > \alpha\, p\, d_1 - E \tag{3.9}$$

That is, the consumer will default whenever second period debts can not be fully enforced. The important point here is that the decision to default is *independent* of the realisation of second period income. No matter what income the consumer receives in the second period, he will default as long as his debt is sufficiently large. If $\alpha\, p\, d_1 < E$ then the consumer will always default whenever he holds any debt. Since default is independent of income, bankruptcy can not insure consumers against low income draws.

Figure 3.1 shows the level of debt for which the consumer is just indifferent between default and repayment when the utility function takes the simple form $u\,(c_1, d_1, c_2, d_2) = \ln c_1 + \ln d_1 + \beta \ln c_2 + \beta \ln d_2$. It shows that more debt can be held, before defaulting, as the level of the exemption E increases. This suggests that the consumer's optimal strategy is to borrow an arbitrarily large amount and default in the second period. A rational lender can anticipate this, and will never lend more than $\alpha\, p\, d_1 - E$. Further, for any level of assets above the default level, repayment is certain regardless of income, and there is no interest rate premium. Backward induction implies this would hold for any finitely lived consumer.

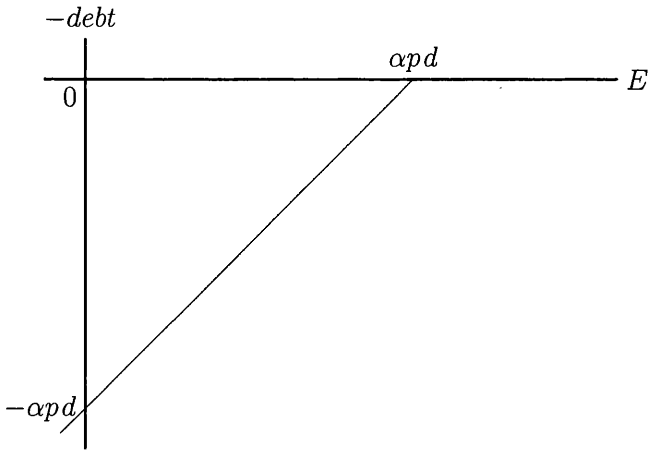

Figure 3.1: The feasible region for debt holdings when the utility function has the form $u(c_1, d_1, c_2, d_2) = \ln c_1 + \ln d_1 + \beta \ln c_2 + \beta \ln d_2$. Default is assured below the line.

In this framework the ability to default has created a credit constraint, and there has been no reduction in uncertainty about second period income. Even if there is no uncertainty about second period income in the example above, the consumer will still be denied credit even though second period income will cover his debt. This limited enforceability unambiguously reduces welfare, in contrast to the case in which debts are fully enforceable.

Kocherlakota (1996) and Kehoe and Levine (1993) among others have considered models in which there are a large number of *ex ante* identical and infinitely lived consumers, and a single non-durable good. Default is punished by being denied access to credit. Since the consumer is infinitely lived the backward induction reasoning can no longer hold. In general, there are many subgame perfect Nash equilibria to this problem, including the belief that no debt will ever be honoured, and no debt is ever allowed. However, these papers asked what is the highest level of lending that can be supported as a subgame perfect Nash equilibrium. This solution will obviously entail that default is punished by being permanently excluded from the credit market. The exact solution depends on the income process, which is assumed to be bounded and drawn from a stochastic Markov process. Since the income process is mean-reverting, default is most valuable when access to credit in the future has the least value, which is when current income is high.

Recall that a possible motivation for bankruptcy legislation is that it reduced uncertainty about future income. The consumer receives $\widehat{y}$ rather than y. Here we have a model where default occurs when income is high. The banks zero profit condition still holds and so the consumer pays extra in low income realisations, and pays nothing in high income realisations: here $var\left(\widehat{y}\right)$ exceeds $var\left(y\right)$. The bankruptcy rule, rather than compressing the distribution of period 2 outcomes, widens the distribution. In the model presented by Kocherlakota (1996) and Kehoe and Levine (1993), default is never allowed. Indeed, not allowing bankruptcy gives the equilibrium that generates the most welfare. Not allowing bankruptcy will place a limit on the amount of debt that is allowed, since consumers will never be allowed to hold enough debt for it to be optimal for them to default. This is another model that endogenously derives credit constraints.

Table 3.1 shows what happens as the level of exemption increases. Suppose that income is exogenous,[7] and bankruptcy provided insurance. As the exemption becomes more generous, the punishment falls. Since default occurs if $q < -A_t$, while repayment takes place if $q < A_2$, then reducing $q\left(\cdot\right)$ will reduce the level of default. Further, as long as default is negatively correlated with income, increasing the level of the exemption will further compress the distribution of second period outcomes, and will provide more insurance. That is, the consumer will want to hold more debt. Lastly, if the bank's zero profit condition holds, a simple application of Leibniz's rule shows that the level of the exemption will raise the

[7]This is important since it rules out moral hazard problems. If income is a function of the punishment then, despite the bank's zero profit condition holding $\frac{\partial E(y)}{\partial E} \neq 0$ and thus it becomes more problematic to describe how borrowing behaviour changes as the punishment for default changes.

interest rate.

The implications of our example or of the model of Kehoe and Levine (1993) or Kocherlakota (1996) are different. Ruling out default means that reducing the punishment will reduce the level of debt that the consumer will be allowed to hold. It will have no effect on the default rate, since default is never allowed. Interest rates will not change either, all consumers will pay the riskless rate r^f. There is no interest rate premium as default never happens.

Table 3.1: Expected effect of increasing the punishment for default.

	Credit Constraints	Insurance
Borrowing	increases	falls
Defaults	no default	fall in the level
Interest rate	no change	increases
Optimal Punishment	very high	very low

In section 4 these ideas about holdings of debt are tested using data. A consumer could be observed in any period of his life, and, in any given period, it is not known whether the consumer is credit constrained.

3.3 Personal Bankruptcy in the United States:

The United States contains some of the most generous bankruptcy regulations for default on debt in the world. The Federal Bankruptcy Act of 1978 specified individuals could choose to file for personal bankruptcy under either Chapter 7 or under Chapter 13, in cases which were not deemed a 'substantial abuse' of the bankruptcy regulations.[8] Chapter 7 was limited to

[8]In practise this meant that bankruptcy would not be allowed if the money had been borrowed with no intention of repaying the money; in cases where the debtor could reasonably repay their debts without

those with assets of less than $750,000 and the aim of the act was to allow those genuinely unable to repay their debts the chance to have a fresh start. Under the act, the debtor had his debts expunged, in return for surrendering all his assets except those deemed by the court necessary for him to make his fresh start: the federal exemptions are shown in table 3.2.[9] Under Chapter 13, the debtor agreed a repayment schedule for part or all of the debt: in practise a ceiling to how much was going to be repaid under Chapter 13 was set by the amount that the debtor could be forced to surrender under Chapter 7. Many courts preferred the debtor to file under chapter 13, but enforced purely nominal repayment schedules. Around 70% of personal bankruptcy cases resulted in a filing for Chapter 7, with the remainder under Chapter 13.

Where the value of the property was in excess of the exemption, the asset would be sold and the amount in excess of the exemption went to satisfy the debt. Cash up to the value of the exemption is retained by the debtor. In some cases the courts insisted that the money had to be re-invested in an exempt asset within a certain amount of time.

3.3.1 State Exemptions:

Since bankruptcy had traditionally been regulated by the individual states, the 1978 act allowed debtors to choose between the exemption allowed by the state and the exemption

resulting in substantial hardship; and in cases where the debtor had changed jurisdiction in order to take advantage of more generous exemptions in the new regime. However, the meaning of substantial abuse did not extend to the ability to repay out of current income, even in cases where current income was high.

[9]Case law has created an obligation for these exemptions to be 'liberally construed' by the courts. These exempt assets would only be surrendered if a valid lien had been created for them: in practise this meant they would only be surrendered if the lender had lent the money specifically to purchase the assets. The exemptions also did not apply to debts arising from state and federal taxes, fines issued by the courts, alimony or child support. The act specifically disallowed the creation of liens that were not related to the purchase of the asset.

Table 3.2: Federal exemptions for Chapter 7 bankruptcy.

Description	Amount $	Comments
Current exemptions:		
1. House	15,000	
2. Car	2,400	
3. Household Goods	8,000	$400 each item (furnishings, goods, clothes, appliances, books, animals, musical instruments) for personal use only.
4. Jewelry	1,000	personal use only.
5. Other Property	800	+ $7,500 of (1) that is unused.
6. Tools of Trade	1,500	Items needed for job.
Prior to 1994:		
1. House	7,500	
2. Car	1,200	
3. Household Goods	4,000	$200 each item.
4. Jewelry	500	
5. Other Property	400	+ $3.750 of (1) that is unused.
6. Tools of Trade	750	
Prior to 1984:		
3. Household Goods		no limit on aggregate amount that can be claimed under this category.
5. Other Property		Allowed all of unclaimed exemption from (1).

Source: Title, 11, Section 522(d) of the annotated federal code. While not recorded, the federal legislation also allowed (with some limits) insurance policies, pensions and annuities, social security payments, and awards adjudicated by the courts to be exempted.

set by the federal government.[10] It also allowed each state to refuse to allow the federal exemptions: the states that have enacted such legislation has been given in table 3.3 below. In the survey used in this chapter, roughly 18% of people are better off claiming the federal exemption rather than the state exemption.

Naturally, in cases where he had the option, the debtor would choose the larger of the state and the federal exemption. The chapter will exploit the differences in the level of the exemption to assess how the punishment in bankruptcy affects the level of debt and the amount of consumption smoothing. This chapter is able to exploit changes in the level in two dimensions; differences across the different states at a point in time, and changes over time.

Table 3.3 shows which states have opted out of the federally set bankruptcy exemptions.[11] As the table shows, most states have disallowed the federal exemptions, and in most cases where the state has not opted out, the state has enacted its own exemptions which may be chosen instead of the federal exemption: in these cases the state exemptions are usually more generous than the exemptions contained in the federal legislation.[12]

As for the federal exemptions, each state has set a variety of things that are exempt from seizure or forced sale for the satisfaction of a debt. The federal law demanded that the state exemptions should act in the same way as the federal exemptions, except in regard to what was exempt, and to what value. In many cases the courts have chosen to interpret legislation in slightly different ways. For example, all states have allowed tools and equipment needed for work to be exempted, up to a limit. However, some jurisdictions have chosen to allow

[10]The source for all the legislation, and legal comments, is derived from the Annotated State Codes published by Westlaw.

[11]Since residents of Montana, North Dakota, Rhode Island, and Wyoming are not sampled in the CEX survey, these states have been excluded from the analysis below.

[12]In two cases, Arkansas and New Hampshire, the state later reversed legislation that refused the federal exemption, while in Illinois, the state opted out of the federal exemptions in 1981, only for the courts to rule that this opt-out, was illegal causing fresh legislation to be re-enacted in the following year.

a car used to drive to work to fall under this definition, while other jurisdictions have not allowed this. The courts have also allowed debtors substantial room for manoeuvre in fully exploiting all the exemptions available: in most cases they have allowed the debtor to re-arrange his portfolio of assets prior to default and substitute exempt assets for non-exempt assets (some limit is placed on the ability to re-arrange assets by 'abuse/fraud' provisions).

Since there is considerable scope for substituting between assets when filing for bankruptcy, the exemptions have been added together, to arrive at a total money value of the exemption for each state. This chapter has summed the exemption on the homestead to the exemption on other assets but it has excluded the exemption on 'tools of trade'. The 'tools of trade' exemption has been excluded since, for the most part, they do not give rise directly to consumption and thus directly enter the utility function. In any case, including these items does not substantively change any of the results. As already stated the calculated exemption value differs between states and across time. It can also differ across subgroups of the population within the state: many states increase the value of exemptions for older, disabled, or married people, or if the debtor has other dependents. In cases where the federal exemption is allowed, the state and federal exemption has been compared and the household has been assigned the larger of the two exemptions.[13] In each case it is the overall household's exemption that has been calculated rather than the individuals in the household. This household exemption will depend on the marital status and age of the household head, on the number of dependents (both children and old people) and on whether the household head, or his spouse, is disabled. The exemption will also depend on the date at which the household is observed, since the exemptions evolve over time.

In calculating the level of exemptions a number of simplifications had to be made. The homestead exemption is that stated in the state legislation. In cases where the homestead exemption was unlimited[14], then a dummy was included in the regressions and the value of

[13]In California, the household was assigned the larger of the two state exemptions.

[14]Note that in these cases there were acreage limits for the property which placed limits on how much

Table 3.3: Whether, and in which year, the state passed legislation to not allow the federal exemptions to be claimed.

Alabama	1980	Mississippi	1982
Alaska	1982	Missouri	1982
Arizona	1980	Nebraska	1980
Arkansas	1981-1991	Nevada	1983
California	1984	New Hampshire	no
Colorado	1981	New Jersey	no
Connecticut	no	New Mexico	no
Delaware	1981	New York	1982
District of Columbia	no	North Carolina	1981
Florida	1979	Ohio	< 1991
Georgia	1981	Oklahoma	1978
Hawaii	no	Oregon	1981
Idaho	1983	Pennsylvania	no
Illinois	1982	South Carolina	1980
Indiana	1980	South Dakota	1980
Iowa	1981	Tennessee	1980
Kansas	1980	Texas	no
Kentucky	1980	Utah	1981
Louisiana	1979	Vermont	no
Maine	1981	Virginia	1979
Maryland	1982	Washington	< 1988
Massachusetts	no	West Virginia	1981
Michigan	no	Wisconsin	no
Minnesota	no		

Source: Westlaw (various) annotated state codes.

the continuous exemption was set at the value of the exemption on other items. In cases where no specific monetary limit was put on a particularly category of goods (for instance some states had an allowance for "all necessary wearing apparel") a value was assigned to the exemption of the good. The following values were adopted: clothes are assigned a value of $1000, books $1000, pictures $1000, other personal possessions $500, jewellery (including watches and wedding rings) $1500, home furnishings $5000, and fuel and provisions $500.

The final issue is to consider what happens when either the state or the federal exemption changes, due to local or national legislation. Most states changed the level of exemptions at least once (if preferred to the federal exemption), and the federal exemptions also changed in this period. While most states only made one or two changes during the period, Minnesota changed the exemption a remarkable seven times.[15] In cases where the month in which the legislation was passed is known (to me), then any observation that is within three months of this legislation has been removed. In cases where the month in which the legislation is not known (the year always is) then all observations for that year have been removed.[16]

Table 3.4 shows the level of exemptions and how they evolve over time. In most states, the exemptions rarely change (observe that the quartiles do not change much) but in most years at least one state changes its level of exemptions (notice how the means change). The homestead exemption is typically much larger than the total exemptions for other property (excluding the 'tools of trade' exemption) and this in turn is usually larger than the 'tools

could be claimed.

[15]Arizona, Colorado, Connecticut, Washington DC, Florida, Hawaii, Illinois, Iowa, Maine, Michigan, Missippi, New Jersey, New Mexico, North Carolina, Oklahoma, Oregon, Pennsylvania, South Carolina, South Dakota, Tennessee, Texas, Utah, and Virginia had one change; Alaska, Arkansas, Idaho, Nevada, and Vermont had two changes, California, New Hampshire, and Washington had three changes, Minnesota had seven changes, while Alaska, Georgia, Indiana, Kansas, Kentucky, Louisiana, Maryland, Massachusetts, Missouri, New York and Wisconsin did not change.

[16]In Ohio, since I have not been able to date the legislation, I have only included observations from 1991, since I was able to confirm that there has been no change in the level of the exemption after this date.

of trade' exemption. The level of the exemption is growing over time, and there is evidence of the distribution being skewed to the left, as the mean is larger than the median in all the cases shown above.

Table 3.4: The level of exemptions (in dollars) over the sample period.

Year	mean	25%	50%	75%
homestead*:				
1988	25,824	8,000	20,000	45,000
1992	28,543	8,000	20,000	100,000
1996	39,821	10,000	30,000	100,000
other assets				
1988	9,507	5,400	7,400	12,700
1992	11,276	5,400	7,400	12,700
1996	14,901	5,825	11,500	19,500
'tools of trade':				
1988	2,389	0	750	5,000
1992	2,504	0	750	5,000
1996	3,053	0	1,000	5,000

* *In calculating the mean for the homestead exemptions, the unlimited homestead exemptions have been omitted.*

As an example of how much the legislation can differ, it is instructive to compare the most, and one of the least generous jurisdictions. In West Virginia a bankrupt has a homestead exemption of up to $5,000 and can also keep up to $1,000 of other personal property. In contrasts Texas, the most generous state, allows the home to be exempt from seizure, no matter what the value of the house, as well as allowing individuals to keep $15,000 of other assets (which could include two cars) while other types of households could keep $30,000. In May 1991, these limits were doubled.

Both table 3.4 and the comparison between Texas and West Virginia show that there is considerable heterogeneity among states with regard to the level of exemptions that may be claimed as exempt in bankruptcy. It is precisely this heterogeneity that will be exploited in this chapter. States also differ in rules concerning garnishment: court orders that take a proportion of wage income directly from employers to lenders. However, since bankruptcy overrides garnishment, filing for bankruptcy tends to be higher in states which allow garnishment, but may not reflect differences in default (less than a quarter of defaults result in a filing for bankruptcy). Usury limits also differ across states, but by 1988 these rules had mostly been repealed. Other possible differences are differences in stigma and in welfare rules. A clear assumption is that omitted state heterogeneity is orthogonal to the state bankruptcy exemptions.[17]

3.4 Data Description:

This work uses the Consumer Expenditure Survey, which is described earlier. The chapter also constructs the total unsecured debts held by the household, including debts held in revolving credit accounts (including store, gasoline, and general purpose credit cards), in installment credit accounts, credit at banks or savings and loan companies, in credit unions, at finance companies, unpaid medical bills, and other credit sources. It also includes negative balances held in checking or brokerage accounts. Excluded from the total are mortgage, and other secured debts. This contrasts with the approach taken by Gropp *et. al.* (1997). Hynes and Berkowitz (1998) argue that the impact of bankruptcy exemptions on secured and unsecured debt ought to be very different, and in their study they consider mortgage debt.

[17]A further problem arises if debtors can move state to take advantage of the more generous exemptions within the new state. In practise the courts were reluctant to allow people who had recently moved to take advantage of the new rules, and indeed some states specified a qualifying period before movers were entitled to the new exemptions.

While mortgage (and other secured) debt is also likely to be important for the household, the creditor has an additional claim to such assets in the event of bankruptcy and can always claim the house (or other security) if the debtor defaults. The housing, or other exemption will not affect the creditors rights in this case, and hence it does not make sense to include such debts in the analysis. Other secured debts (for instance on cars) have also been excluded.

Table 3.5: Summary statistics for different exemption quartiles.

	all	quartile			
		1	2	3	4
total debt ($)	2123	2085	1974	2291	2096
holds debt (%)	62.63	63.15	58.99	64.16	64.36
interest rate (%)	14.70	15.45	14.85	14.34	14.27
interest rate (%) if owe $1,000+	19.01	19.60	19.08	19.03	18.03
defaults per 1000	8.58	11.39	8.19	7.11	8.09

Table 3.5 gives a brief summary of the data, and compares the different exemption quartiles of the state exemptions. It shows that the level of debt changes from quartile 1 (in which the lowest level of assets may be kept) to quartile 4. The average level of debt held is around $2,100 (the median is $385, while the 75th percentile is $2,250) but there is no strong pattern to the level of debt. It is also difficult to see a pattern to the number of people holding at least some debt in the sample. In all cases around 60% of people hold debt. However, looking at the interest rate suggests laxer rules imply a higher interest rate. The interest rate is constructed as the reported costs divided by the reported level of debt. The interest rate is thus the average interest rate on all debts rather than the marginal interest rate, which is what motivates the marginal borrowing decision. This pattern of interest rates falling as the level of the exemptions increases remains if larger debts only are looked at.

These results are significant in themselves (at the 10% level) if a one-sided rank-order test is done. The rate of default, calculated from aggregate data as the ratio of the number of bankruptcy filings, divided by the number of households (rather than individuals) resident in the state, is much higher for the first quartile for which the highest exemptions, but otherwise there does not seem to be a clear pattern to the defaults. As might be expected, the pattern for defaults and the interest rate is similar, but does not match completely: perhaps because the interest rate not only reflects the probability of default, but also the cost to the lender of default.

3.5 Regressions:

According to the theory outlined above, debtors will hold debt up to some maximum amount. By comparing the level of debt that individuals hold across states, the impact of the state exemptions can be assessed. Since debts are bounded at zero, a simple tobit model, in which the level of debt is regressed on a set of household characteristics, and the bankruptcy exemption to which the household is eligible, can be used. The key assumptions here are that household characteristics, and the size of the exemption, are exogenous. Further assumptions are that the household's state of residence is also exogenous, and that any changes in the level of exemptions over time are unexpected. In reality, household's decisions about education, residence and fertility may well be related to the ability to smooth consumption: at some level all economic decisions are endogenous. However, for this discussion it is assumed that these issues are of secondary importance, and they shall be ignored.

3.5.1 Results:

In table 3.6 the results of the tobit are displayed. They show that increasing the exemptions reduces the amount of debt that is held by households. The regressions are for the level of debt, and the level of income: recall that example 1 implied that there should be a linear relationship. The first regression shows the coefficients on all the control variables, without including the exemption variables or income. These variables will partly account for preferences, and partly account for income. The regression includes age and cohort effects, which means that time is excluded (age, cohort and time are collinear). This implicitly assumes all changes over time in the population is due to individual cohorts aging, and new cohorts replacing old cohorts. If year effects are important, then this will show up in the age and cohort coefficients. However, the paper does not attempt to interpret these coefficients. The interest rate that is included is the municipal bond rate deflated by the inflation rate. The regression has 10,418 observations: the small number is due to the fact that only the second interview for those households with full state information are included. Furthermore, households who are very close to a *change* in their exemption level are also excluded (within three months if the month is known and in the same year if it is not). The reason for excluding these households, is it is not clear whether one should use the existing exemption, or the exemption that may rationally anticipated shortly in the future.

When the level of the exemption is included, (and also dummy for unlimited homestead exemption,) we find that the coefficients are not significant level. A joint test of the level and including a dummy for the unlimited homestead exemption is also not significant. Although the negative coefficient is consistent with the simple theory of credit constraints expounded earlier. Other things to note are that households headed by females or non-white people seem to hold lower levels of debt (which may partly reflect the greater chance they have of being turned down, see Hajivassiliou and Ioannides, 2002c). Better educated people also hold higher levels of debt as well.

Table 3.7 shows the effect of including state specific dummies. Including these state specific effects ought to control for other state specific effects that are not included in the regression. Thus it measures the effect of changes over time. When these dummies are included, the control variables do not change substantially. However, the effect on the exemption coefficient is substantial: the results are now highly significant, as shown by the F-test. This time increasing the level of the exemption from the 25th centile to the 75th centile entails a reduction in about $1,500. This is large figure, but not implausible, recall table 3.5 showed the average level of debt is around $2,100 dollars.[18] The true effect is likely to be under-estimated. The correct regression to run is a tobit which is truncated at zero and at the point where the credit constraints bite, which, however, is unknown. Unfortunately, it is not even known if the consumer is credit constrained. The level of debt that the consumer will hold will only change for the higher level of exemptions, if the consumer is credit constrained, and he is able to borrow more money at the lower level of exemptions (where the punishment for default is bigger). For households that are not credit-constrained, there will be no change in the level of debt that they hold. Thus the amount calculated from the table will under-estimate the true effect. It would have been nice to have included time dummies, to exploit purely the cross-sectional variance, but this is not possible since the regression already includes age and cohort effects. Clearly age, cohort, and time effects are not all separately identifiable.[19]

A second feature of table 3.7 is that including income in the regression does not substan-

[18]The distribution of observed debt may have fatter tails than would be implied by the normality assumption used in the tobit regression, which may affect the results. A second potentially serious econometric problem is the presence of at least one regressor that contains substantial measurement error. This is an endemic problem with no fully satisfactory solution, consequently all that this paper can do is acknowledge the problem.

[19]Another possibility would be to first difference the data, and regress the change in the exemptions on the change in the level of debt. However, the exemptions change relatively rarely, and hence the results would be identified by a comparatively small number of observations. Furthermore, these changes take time to be legislated, and, at least over the period of two or three months, can be anticipated by lenders. This would cause the level of debt to already reflect the new exemptions immediately before the change (something that in the reported regressions motivated the exclusion of such observations).

tially change the results. Included in the regression is the current level of income. This will include both temporary and permanent components. If the temporary component is high then this will reduce the level of borrowing in the current period, while if the permanent component is high, then the effects would be a little more ambiguous. Suppose individual i's income, denoted y, follows the following process:

$$y_{it} = \theta x_{it} + v_i + \varepsilon_{it} \tag{3.10}$$

where x is a set of other explanatory factors (that evolve over the lifetime), v_i can be thought of as permanent income, and ε_{it} is temporary income. The permanent effect will unambiguously raise consumption, and it will raise debt in periods where θx_{it} is unusually low. This is indeed what the regressions find: increasing income does raise the level of debt that the individual holds.

The effect of the exemptions on the interest rate that is charged is reported in tables 3.8 and 3.9. The interest rate is the self reported interest rate from the 5th interview and it is only calculated for those who hold at least some debt. This explains why the sample size is much smaller than in the other regression. Again, the identifying assumption is that the interest rate charged is independent (in a statistical sense) of whether any debt will be held: we are not just selecting the low interest rate people. This may not be a particularly appealing assumption in this case. The results suggest that perhaps better educated people face lower interest rates, although the effects are small. In table 3.9 neither the level of the exemptions nor the level of income enter significantly into the results. This is disappointing given table 3.5, where there is a clear monotonic relationship between the interest rate and the exemption quartiles. These results could be due to the small sample size and the fact that self reported interest rates are likely to be measured extremely inaccurately. However, while this can explain the insignificance of the results in table 3.8 it can not explain the sign (measurement error in the left-hand side does not bias the point estimates). Table 3.9 reports estimates when state specific dummies are included, and again the results are not significant. The identifying assumption may also be causing these results. As the interest rate increases, some households would decide not to hold debt, thus downward biasing the results if the sample is restricted to those holding any debt.[20]

3.5.2 Consumption:

So far these equations have been couched in terms of the level of debt that is held by the household. It is also interesting to think more directly about consumption. For instance, consider the standard Euler equation for consumption growth that has been estimated in the literature:[21]

$$\Delta \ln c_{it} = \gamma^{-1} (r - \delta) + \beta X_{it} + u_{it} \tag{3.11}$$

This framework implies an iso-elastic utility function with relative risk aversion parameter γ, δ is the discount rate, while X_{it} represents observed taste shifters, such as family size. According to the theory, nothing else should enter the regression. However, the literature has consistently rejected this: current and future income both seem to enter significantly.[22] Two of the most popular explanations for this can be interpreted as having implications for the error term u_{it}. Writing

$$u_{it} = \frac{\gamma}{2} var(c_{it}) + \phi ln(1 + \psi_{t-1}) + \varepsilon_{it}$$

then if the relative risk aversion parameter γ is non-zero, there is a precautionary motive to saving, and the variance of permanent income c_{it} will enter the equation. Alternatively, if some consumers are credit-constrained, the kuhn-tucker condition has an associated multiplier ψ, which will be positive when credit constraints are binding. The rejection of equation

[20]Charles (2000) also reports results for a probit on whether any debt is held and on the probability of default, although only the first of these was significant.

[21]See, for instance, Deaton (1992)

[22]See for instance, Hall and Mishkin (1982), or Flavin (1981).

3.11 can be thought of as an omitted variable problem, as is well known in the literature. However, bankruptcy constraints can account for part of this error term. Suppose the bankruptcy exemptions were included in equation 3.11 as an additional explanatory variable, then if the bankruptcy exemptions are providing insurance, this will reduce the variance of consumption, and thus the parameter on the bankruptcy rules will be negative. If instead bankruptcy rules created credit constraints, then the estimated coefficient in equation 3.11 should be positive.

A second and additional approach is due to Deaton and Paxson (1994) and will again help to test whether bankruptcy rules provide insurance. Consider a consumer who faces both temporary and permanent shocks to his (log-)income (here v_{it} denotes the permanent shock while ε_{it} denotes the temporary shock).

$$y_{it} = y_{it-1} + v_{it} + \varepsilon_{it} - \varepsilon_{it}$$

Suppose he can smooth some proportion $1-\phi_j$ of each shock. Then there will be the following relationship between (log-)income and (log-)consumption (assuming the interest rate equals the discount rate and ignoring preference shocks).

$$c_{it} = c_{it-1} + \phi_1 v_{it} + \phi_2 \varepsilon_{it}$$

Full insurance implies $\phi_1 = \phi_2 = 0$, while the permanent income hypothesis argues $\phi_1 = 1$, $\phi_2 = 0$ (that is, all temporary shocks are fully insured, but permanent shocks are uninsured). More generally, one may believe both types of shocks are partly insured. An implication of the model is that the cross-sectional variance of consumption should be growing over time, and that the variance of the change in consumption is related in a very simple way to the level of insurance.[23]

$$var\left(c_{it} - c_{it-1}\right) = \phi_1^2 var\left(v_{it}\right) + 2\phi_2^2 var\left(\varepsilon_{it}\right)$$

[23] One could equally use the change in the variance in consumption, although, anticipating table 3.12, the results, while similar, were not as clear cut.

If the variance of the income shocks is identical across all states, but bankruptcy rules generate insurance, then this should cause the variance the change in consumption to be negatively correlated with these bankruptcy rules: insurance means that the variance of the change in consumption is lower. In either case the sign of the effect can be identified, but the size of the coefficient can not be.

Results:

The regressions relating to equation 3.11 augmented by the state exemptions are on state-quarter averages, and are displayed in tables 3.10 and 3.11. The first shows that the level of exemptions seems to be of negligible importance when considering consumption growth. Including state specific effects increases the size of the estimated effect dramatically. Although the $F-$test continues to reject the significance of the results, the results are negative, and close to being significant: the point estimate suggests that the bankruptcy exemptions provide some insurance. These results seem to contrast the earlier results on debt holding, where increasing the level of exemptions reduced the level of debt. These results might suggest that bankruptcy is making it easier to smooth consumption over relatively short time intervals, but cause problems in smoothing consumption over lower frequency events. Recall that the earlier regressions investigated level effects while the regressions in this section investigate the effect of changes in the levels.

The results for the variance of the change in consumption are displayed in table 3.12.In these regressions, the dependent variable was constructed as the difference in the variance of c_{it} where the variance was constructed for each state-year cell. At time t household heads were aged between 30 and 55 while in time $t+1$ household heads were aged between 31 and 56. The fact that the constant is both positive and significant rules out complete insurance. The coefficient on the level of the exemption is small and far from significant, perhaps due to the low power of the test. By contrast, the coefficient on the dummy for an unlimited housing exemption is both negative, and significant (at the 10 percent level). The coefficients change little if year, or state and year dummies are included, although for the latter, the variance increases substantially. The results from this section suggest that bankruptcy exemptions are providing an insurance role.

3.6 Conclusion:

The results show that as the level of the exemptions increases, households hold less debt. Including state specific dummies dramatically increases the estimated effect. The fact that the tobit regressions showed that the exemptions were negatively related to the level of debt held suggests that credit constraints are important. This result can not be interpreted as resulting from either uncertainty, or from the fact that consumption changes with household needs. The insurance argument would seem to imply the opposite effect. This could be a way of comparing the comparative importance of these two arguments, although a much more realistic model would fully interact the two effects. The argument for credit constraints seems to be an incomplete argument, not least because casual observation shows that a great many people default in the United States. The fact that default is observed would seem to support the view that incomplete insurance is still an important additional factor that helps to explain the inability of consumers to fully smooth consumption.

In contrast to the regressions on the level of debt held, the consumption growth equations show weak evidence that the bankruptcy laws were providing insurance to consumers. One way to reconcile these results with the levels results is to argue that bankruptcy rules help consumers to smooth relatively high frequency income shocks, but at the cost of making it more difficult to smooth income over the lifecycle. The regressions of the variance of the growth of consumption (rather than the levels) seemed to suggest that bankruptcy rules reduces variance of consumption, which supports the idea that the rules do providing insurance.

These results have suggested an indirect way of testing for credit constraints and for incomplete insurance. They offer testable implications for the way that increasing the level of the bankruptcy exemptions will affect the level of borrowing, the interest rate, and the rate of bankruptcy. The fact that credit constraints and insurance suggest that the effect of increasing the bankruptcy exemptions have different effects on the level of borrowing can potentially offer a test of their relative importance. Gropp *et.al.* (1997) showed results that suggested that for low asset people, the increasing the state exemptions reduced the level of borrowing, while the results were reversed for high asset people. This is what might be expected: for low asset people credit constraints dominate; while for high asset people insurance dominates.

3.7 Appendix

Table 3.6: Results of a tobit regression on household debt (probability in parentheses).

parameter	(1)	(2)	(3)
age	1990.10	1947.56	1882.81
	(0.548)	(0.555)	(0.566)
age-squared	-208.41	-200.26	-220.73
	(0.776)	(0.784)	(0.761)
age-cubed	-2.05	-2.62	1.09
	(0.968)	(0.958)	(0.983)
completed high school	367.81	370.18	205.27
	(0.025)	(0.026)	(0.173)
some college	1014.52	1012.12	771.22
	(0.000)	(0.000)	(0.001)
college graduate	1302.04	1302.98	936.59
	(0.000)	(0.000)	(0.001)
Black	-171.54	-160.91	-64.40
	(0.457)	(0.489)	(0.789)
Asian	-480.67	-483.46	-354.46
	(0.020)	(0.023)	(0.063)
Native American	-468.75	-473.45	-450.52
	(0.061)	(0.059)	(0.071)
female household head	-482.36	-483.23	-391.28
	(0.008)	(0.008)	(0.031)
not married	-198.38	-195.24	-7.86
	(0.447)	(0.453)	(0.977)
family-size	205.62	203.91	157.15
	(0.137)	(0.143)	(0.249)
family-size squared	59.81	59.34	83.14
	(0.255)	(0.258)	(0.104)
interest rate	307.02	302.18	316.56
	(0.203)	(0.210)	(0.189)
exemption × 1000	-	-1.31	-1.58
		(0.739)	(0.695)
unlimited homestead exemption	-	-65.35	-115.24
		(0.746)	(0.585)
income × 1000	-	-	45.05
			(0.000)
F-test*	-	0.07	0.16
		(0.93)	(0.85)
no. of observations	10,418	10,418	10,418

*The F-test is a joint test for the exemption and the dummy for unlimited homestead exemption. All regressions include a constant, month and cohort dummies. Standard errors allow for clustering by state.

Table 3.7: Results of a tobit regression on household debt, including state dummies (probability in parentheses).

parameter	(4)	(5)	(6)
age	1675.67	1831.95	1755.01
	(0.682)	(0.654)	(0.667)
age-squared	-141.50	-170.82	-188.57
	(0.877)	(0.851)	(0.836)
age-cubed	-6.42	-4.21	-0.89
	(0.921)	(0.948)	(0.989)
completed high school	385.49	375.21	223.41
	(0.036)	(0.041)	(0.230)
some college	981.67	965.02	746.13
	(0.000)	(0.000)	(0.000)
college graduate	1287.783	1269.22	935.79
	(0.000)	(0.000)	(0.000)
Black	-166.87	-167.94	-72.74
	(0.376)	(0.373)	(0.701)
Asian	-532.52	-518.78	-408.66
	(0.532)	(0.543)	(0.631)
Native American	-760.27	-765.27	-706.52
	(0.023)	(0.022)	(0.035)
female household head	-494.83	-479.83	-396.79
	(0.001)	(0.001)	(0.006)
not married	-189.88	-281.31	-88.79
	(0.291)	(0.124)	(0.634)
family-size	228.53	250.07	202.40
	(0.004)	(0.002)	(0.011)
family-size squared	56.74	47.66	70.95
	(0.349)	(0.432)	(0.243)
interest rate	232.12	260.24	263.49
	(0.474)	(0.423)	(0.417)
exemption × 1000	-	-26.39	-21.47
		(0.004)	(0.018)
unlimited homestead exemption	-	-5013.23	-4266.27
		(0.002)	(0.010)
income × 1000	-	-	42.18
			(0.000)
F-test*	-	4.75	3.33
		(0.008)	(0.035)
no. of observations	10,418	10,418	10,418

*The F-test is a joint test for the exemption and the dummy for unlimited homestead exemption. All regressions include a constant, cohort dummies, and state dummies.

Table 3.8: Results of a linear regression on the interest rate that the household pays (probability in parenthesis).

parameter	(1)	(2)	(3)
age/10	0.007	0.008	0.008
	(0.598)	(0.578)	(0.556)
age-squared/100	-0.007	-0.007	-0.008
	(0.420)	(0.398)	(0.319)
age-cubed/1000	0.003	0.003	0.002
	(0.627)	(0.630)	(0.675)
completed high school	-0.024	-0.024	-0.022
	(0.126)	(0.119)	(0.151)
some college	-0.004	-0.004	-0.001
	(0.779)	(0.779)	(0.925)
college graduate	-0.038	-0.039	-0.033
	(0.016)	(0.014)	(0.044)
Black	0.025	0.023	0.022
	(0.078)	(0.103)	(0.115)
Asian	-0.034	-0.034	-0.034
	(0.521)	(0.520)	(0.517)
Native American	0.034	0.033	0.033
	(0.151)	(0.170)	(0.168)
not married	-0.001	-0.001	-0.002
	(0.741)	(0.708)	(0.581)
female head	0.014	0.016	0.013
	(0.140)	(0.132)	(0.174)
family-size	-0.008	-0.008	-0.006
	(0.165)	(0.169)	(0.256)
family-size squared	-0.006	-0.006	-0.007
	(0.121)	(0.169)	(0.093)
exemption $\times$ 1000	-	-2.99e-04	-1.95e-04
		(0.331)	(0.330)
unlimited homestead exemption	-	-0.020	-0.021
		(0.157)	(0.123)
income $\times$ 1000	-	-	-3.68e-04
			(0.123)
F-test*	-	1.24	1.34
		(0.288)	(0.262)
No. of observations	6262	6262	6262

*The F-test is a joint test for the significance of the exemption level and the dummy for unlimited homestead exemption. All regressions include a constant and cohort dummies.

Table 3.9: Results of a linear regression on the interest rate that the household pays, including state dummies. (probability in parenthesis).

parameter	(1)	(2)	(3)
age/10	0.003	0.006	0.007
	(0.822)	(0.682)	(0.641)
age-squared/100	-0.007	-0.007	-0.008
	(0.394)	(0.392)	(0.312)
age-cubed/1000	0.003	0.003	0.002
	(0.600)	(0.610)	(0.655)
completed high school	-0.021	-0.021	-0.019
	(0.178)	(0.181)	(0.230)
some college	-0.002	-0.002	0.000
	(0.854)	(0.860)	(0.981)
college graduate	-0.036	-0.036	-0.030
	(0.021)	(0.022)	(0.067)
Black	0.022	0.022	0.021
	(0.119)	(0.119)	(0.137)
Asian	-0.026	-0.026	-0.026
	(0.620)	(0.617)	(0.613)
Native American	0.059	0.059	0.058
	(0.031)	(0.033)	(0.035)
not married	-0.001	-0.001	-0.002
	(0.884)	(0.755)	(0.606)
female head	0.013	0.013	0.012
	(0.160)	(0.162)	(0.220)
family-size	-0.007	-0.007	-0.005
	(0.188)	(0.226)	(0.346)
family-size squared	-0.007	-0.007	-0.008
	(0.110)	(0.098)	(0.070)
exemption $\times$ 1000	-	-0.001	-0.001
		(0.271)	(0.231)
unlimited homestead exemption	-	-0.093	-0.103
		(0.418)	(0.370)
income $\times$ 1000	-	-	-3.99e-04
			(0.236)
F-test*	-	0.66	0.77
		(0.516)	(0.463)
No. of observations	6262	6262	6262

*The F-test is a joint test for the significance of the exemption level and the dummy for unlimited homestead exemption. All regressions include a constant, cohort dummies and state dummies.

Table 3.10: Results of a linear regression of the growth rate of consumption (probability in parentheses).

parameter	(1)	(2)	(3)	(4)
age/10	0.013	0.017	0.018	-0.007
	(0.641)	(0.540)	(0.524)	(0.760)
age-squared/100	-0.031	-0.272	-0.023	-0.019
	(0.160)	(0.220)	(0.293)	(0.318)
age-cubed/1000	-0.008	-0.007	-0.003	-0.017
	(0.742)	(0.774)	(0.882)	(0.428)
completed high school	0.036	0.056	0.016	0.226
	(0.674)	(0.533)	(0.861)	(0.004)
some college	0.093	0.124	0.070	0.316
	(0.285)	(0.168)	(0.450)	(0.000)
college graduate	0.103	0.140	0.050	0.329
	(0.200)	(0.099)	(0.594)	(0.000)
Non-white	0.019	0.009	-0.004	0.067
	(0.578)	(0.301)	(0.906)	(0.048)
female head	-0.072	-0.071	-0.063	-0.084
	(0.244)	(0.249)	(0.306)	(0.114)
not married	-0.111	-0.109	-0.049	0.181
	(0.183)	(0.194)	(0.577)	(0.013)
family-size	0.597	0.683	0.066	0.025
	(0.135)	(0.093)	(0.102)	(0.476)
family-size squared	0.001	0.006	0.006	-0.018
	(0.949)	(0.839)	(0.825)	(0.464)
real interest rate	-0.006	-0.006	-0.006	-0.016
	(0.322)	(0.322)	(0.290)	(0.004)
exemption × 1000	-	-0.0002	-0.0002	-0.0001
		(0.213)	(0.158)	(0.388)
unlimited homestead exemption	-	-0.002	-0.008	-0.022
		(0.840)	(0.572)	(0.081)
income × 1000	-	-	-0.078	-
			(0.034)	
income growth × 1000	-	-	-	0.352
				(0.000)
F-test*	-	1.01	1.60	1.57
		(0.365)	(0.203)	(0.208)

*The regressions are on the state average in each quarter for the variables (and only includes those cells with at least 50 observations). The F-test is a joint test for the exemption and the dummy for unlimited homestead exemption. All regressions include a constant and seasonal dummies.

Table 3.11: Results of a linear regression of the growth rate of consumption when state effects are included (probability in parentheses).

parameter	(1)	(2)	(3)	(3)
age/10	0.013	0.017	0.016	0.007
	(0.641)	(0.578)	(0.602)	(0.765)
age-squared/100	-0.031	-0.031	-0.014	-0.029
	(0.160)	(0.202)	(0.556)	(0.136)
age-cubed/1000	-0.008	-0.011	-0.005	-0.018
	(0.742)	(0.685)	(0.842)	(0.416)
completed high school	0.036	0.057	0.041	0.040
	(0.674)	(0.532)	(0.724)	(0.673)
some college	0.093	0.071	0.018	-0.008
	(0.285)	(0.581)	(0.885)	(0.935)
college graduate	0.103	0.166	0.096	0.0.065
	(0.200)	(0.165)	(0.418)	(0.496)
Non-white	0.019	0.104	0.156	0.025
	(0.578)	(0.315)	(0.132)	(0.761)
female head	-0.072	-0.072	-0.030	-0.042
	(0.244)	(0.384)	(0.709)	(0.529)
not married	-0.111	-0.172	-0.049	0.291
	(0.183)	(0.098)	(0.363)	(0.001)
family-size	0.597	0.642	0.044	0.033
	(0.135)	(0.162)	(0.328)	(0.363)
family-size squared	0.001	0.004	0.001	0.008
	(0.949)	(0.894)	(0.966)	(0.754)
real interest rate	-0.006	-0.007	-0.007	-0.018
	(0.322)	(0.264)	(0.239)	(0.001)
exemption × 1000	-	-0.002	-0.002	-0.002
		(0.056)	(0.070)	(0.025)
unlimited homestead exemption	-	-0.337	-0.315	-0.334
		(0.086)	(0.104)	(0.033)
income × 1000	-	-	-0.173	-
			(0.000)	
income growth × 1000	-	-	-	0.494
				(0.000)
F-test*	-	2.07	1.84	2.57
		(0.127)	(0.157)	(0.077)

*The regressions are on the state average in each quarter for the variables (and only includes those cells with at least 50 observations). The F-test is a joint test for the exemption and the dummy for unlimited homestead exemption. All regressions include a constant and seasonal dummies.

Table 3.12: Regressing $\Delta var\,(c_{it})$ against the bankruptcy exemptions.

	coeff.	s.d.(robust)	prob.
exemption (× 1000)	-0.343	0.468	0.466
unlimited housing exemption	-0.063	0.032	0.054
constant	0.492	0.013	0.000
+ year dummies:			
exemption (× 1000)	-0.453	0.509	0.376
unlimited housing exemption	-0.061	0.029	0.039
constant	0.445	0.020	0.000
+ year and state dummies:			
exemption (× 1000)	-0.323	2.332	0.890
unlimited housing exemption	-0.009	0.315	0.975
constant	0.475	0.393	0.229

The number of observations was 115, while the minimum cell size was 75.

Chapter 4

Risk Sharing and the Tax System

Abstract

Several papers have documented that US consumers can not fully insure themselves against all their idiosyncratic risks, but little is understood about which mechanisms provide insurance. We investigate whether, as some suggest, progressive taxes provide additional insurance. The methodology distinguishes insurance from redistribution, and can by applied to testing any potential insurance mechanism. Using repeated cross-sections from the US consumer expenditure survey (CEX), we relate changes in consumption inequality to several measures of tax progressivity. Identification exploits the variation in taxes both across states and over time. Our results suggest, under weak assumptions, that progressive taxes induce insurance, while stronger assumptions quantify this effect.

4.1 Introduction

Over the last twenty years or more, there has been considerable interest, particularly among consumption economists, about the extent to which consumers can smooth out income shocks. While the earlier literature considered consumers smoothing over time, much of the more recent literature has looked at smoothing across states of nature, asking whether households can fully insure their idiosyncratic risks. An implication of the full insurance hypothesis is that neither temporary nor permanent shocks to income will have any impact on consumption behaviour. While this hypothesis has consistently been rejected by the literature (for US consumers, see for instance Mace 1991, Cochrane 1991, or Attanasio and Davis 1996), it seems unlikely that agents are completely autarkic: more likely, they have some, but partial, ability to smooth shocks. A natural corollary is to ask what mechanisms provide insurance? This has important welfare implications, since providing additional insurance, other things remaining equal, will make consumers better off. Thus we investigate what changes the total amount of insurance available to consumers, explicitly recognizing that any potential mechanism might merely crowd out, or substitute for, other risk-sharing mechanisms.

One particularly popular candidate instrument that might provide insurance is the tax and benefit system. This is explicitly recognized in some tax regimes: for instance, in Britain part of the tax system is labeled national insurance. Moreover, some economists, such as Varian (1980), have argued that this is a more important motivation for progressive taxes than redistribution. Several papers have studied this problem, including Attanasio and Rios-Rull (2000), and Kruegar and Perri (1999), by simulating over some plausible parameter estimates. Never-the-less, whether the tax system really does provide insurance, or merely crowds out, or substitutes for, other forms of insurance is an open issue.

In an interesting and much discussed paper, Asdrubali, Sørensen and Yosha (1996) investigate how state level shocks are insured, or smoothed, across US states. Their exercise decomposes insurance into its various constituent parts, by regressing changes in consumption against changes in the components of income. They find that federal taxes smooth 13% of income variation, but their results suggest that taxes are only half as important as credit markets, and a third as important as capital markets (e.g. holding a balanced portfolio of assets), in smoothing state level consumption. Instead, this chapter looks at individual level shocks, and investigates the level of insurance between households within the state. Our approach has a number of advantages: using individual data we can separate insurance from redistribution, something they concede they can not distinguish. We also want to know the marginal effect. We have in mind the following policy experiment: suppose the tax system was changed so that it became more redistributive; will the level of insurance change, or will consumers merely substitute from alternative insurance mechanisms. In the decomposition of Asdrubali *et. al.* (1996) this effect is not captured.

The methodology of this chapter builds and develops the earlier work of Deaton and Paxson (1994). We exploit differences across US states, as well as time, to identify what effect, if any, the tax system has on the total amount of insurance available to consumers: identification follows from the fact that different US states have different tax regimes. We use household level consumption data available from the CEX for the years 1980-1998. The aim is not to reject full insurance, but rather to test whether, and how much, a progressive tax system provides insurance. For this, as will be explained, it is not necessary to explicitly state what is observed in the full insurance case, unlike the paper by Deaton and Paxson (1994). While quite general assumptions can be used to identify whether taxes provide insurance, tighter assumptions are needed to quantify how much is provided. We will find that the tax system really does help agents to insure themselves against income shocks.

The chapter is organised so that first it will discuss how a progressive tax system might provide insurance. Then it discusses the literature on how to test for (complete) insurance, before explaining how the literature can be used to test for the marginal, or extra, amount of insurance that some policy instrument, such as the tax system, might provide. We then briefly discuss our data, and how the tax system works in the US, before we present the results of our regressions. Finally after discussing our results, we briefly conclude.

4.2 Taxes and Transfers

It has been traditional in the public economics literature to view taxation as a method of reducing cross-sectional inequality. The idea is that redistributing income from high to low income households may be a 'good thing'. Obviously, unless agents were altruistic, such a plan will be supported by low income people and opposed by high income people. If, on the other hand, agents were altruistic, there would seem to be no motivation for a public scheme of re-distribution, since agents would be privately motivated to redistribute their income. However, this argument ignores another (possible) important motivation for a redistributive tax.[1] If future income is uncertain, then, in the absence of alternative insurance mechanisms, a redistributive tax system can provide insurance against this uncertainty. Varian (1980), for instance, argued that this ought to motivate extremely high marginal tax rates.

Suppose individual i's current income y_{it} has a permanent y^p_{it} and a temporary component[2]:

$$\begin{aligned} y^p_{it} &= y^p_{it-1} + f_{it} \\ y_{it} &= y^p_{it} + \varepsilon_{it} \end{aligned} \tag{4.1}$$

This periods income is subject to permanent f_{it} and temporary u_{it} shocks, which are assumed to be mean zero processes. This framework implies income can be written:

$$y_{it} = f_{i0} + \sum_{s=1}^{t} f_{is} + \varepsilon_{it} \tag{4.2}$$

[1]Here a poll tax is deemed non-redistributive, and any tax that increases as income increases is deemed to be redistributive, or progressive.

[2]This analysis is similar to that presented in Deaton et. al. (2000)

Taxes are typically a function of current income: suppose taxes are levied at some flat rate level τ. In which case expected *per capita* tax revenue (where the expectation is over the i individuals) is:

$$\tau E_i \left[f_{i0} + \sum_{s=1}^{t} f_{is} + \varepsilon_{it} \right] = \tau E_i f_{i0} = \tau \bar{f}_{i0}$$

Suppose these taxes are re-distributed back to all agents as a lump sum transfer.[3] After tax and transfer income, $\tilde{y}_{it}$, would be:

$$\tilde{y}_{it} = f_{i0} - \tau \left(f_{i0} - \bar{f}_0 \right) + (1-\tau) \left[\sum_{s=1}^{t} f_{is} + \varepsilon_{it} \right] \tag{4.3}$$

and the change in income can be decomposed into an insurance component and a re-distribution component.

$$y_{it} - \tilde{y}_{it} = \underbrace{\tau \left(f_{i0} - \bar{f}_0 \right)}_{redistribution} + \underbrace{\tau \left[\sum f_{is} + \varepsilon_{it} \right]}_{insurance} \tag{4.4}$$

Re-distribution is on the initial distribution of income, while insurance acts on any income shocks. We concentrate on this second feature of the tax system. Whether, in practise, the tax system provides insurance depends on what mechanisms the agent has to smooth consumption in the presence of income shocks. If no other insurance mechanisms were available, then clearly a progressive tax system would provide insurance. On the other hand, the agent may already have alternative risk-sharing arrangements available to him, in which case the tax system may instead interrupt the agents private incentives to participate in these arrangements. In this case, a progressive tax system may provide no additional insurance, but may instead merely crowd out, or substitute for, existing arrangements, a point highlighted by Attanasio and Rios-Rull (2000).

[3]It would be simple to adjust this equation if some fixed proportion of $\bar{f}_{i0}$ is spent by the government on public goods, rather than being redistributed. This adjustment has no effect on the analysis.

We propose using the different tax systems in the different states of the US to examine how the level of insurance changes as the degree of re-distribution changes. However, rather than attempting to simulate over the possible values of the parameters in a fully structural model, an approach taken in Attanasio and Rios-Rull (2000) or Deaton, Gourinchas and Paxson (2000) among others, we will instead estimate a reduced form, and see what insights can be gained from that approach. While the structural approach can give valuable insights, we do not believe that it can successfully answer the question that we are interested in: it would require knowledge of preferences, the income process, and *all* insurance mechanisms available to the agent. This is not necessary in the reduced form approach that we take in this chapter. While it can not describe the mechanism that provides insurance, it can say whether insurance is provided. This is often what a policy maker wants to know. The aim is to compare different tax systems to see how the overall level of insurance changes across regimes.

4.3 Empirical Framework

A number of approaches have been used to test for full insurance. The earliest tests are based on exclusion restrictions. Suppose each individual i's income y_{it} was subject to aggregate η_t and idiosyncratic u_{it} shocks each period t so that

$$y_{it} = y_{it-1} + \eta_t + u_{it}$$

By construction, the idiosyncratic component u_{it} sums to zero over the i individuals. Full insurance implies that this term does not affect changes in consumption, hence a valid test for full insurance is to regress changes in consumption Δc_{it} on the idiosyncratic component of the income shock. Mace (1991), Cochrane (1991), and Attanasio and Davis (1996) all apply this test to US households, and decisively reject this implication of the full insurance hypothesis. A second approach, due to Jappelli and Pistaferri (1999), instead tests the ordering of agents.

Under full insurance, if agents, in any time period, are ordered by their level of consumption, then this ordering (after controlling for taste-shifters) will not change over time. They reject full insurance, and argue that neither measurement error, nor taste-shifters, could explain this rejection.

A third implication was investigated by Deaton and Paxson (1994). For the moment ignore taste-shifters, then, with quadratic utility, and if the interest rate r equals the discount rate δ, consumption follows a martingale:

$$c_{it} = c_{it-1} + \varepsilon_{it} \tag{4.5}$$

in which case the cross-sectional variance of consumption for a fixed membership group j evolves according to the relationship:

$$var^j(c_{it}) - var^j(c_{it-1}) = var^j(\varepsilon_{it}) + 2cov^j(c_{it-1}, \varepsilon_{it})$$

Deaton and Paxson (1994) showed that if lagged aggregate consumption is in each agent's information set then $cov^j(c_{it-1}, \varepsilon_{it}) = 0$ (at least on average over a large enough number of time periods). If income follows the stylized process:

$$\begin{aligned} y_{it}^p &= y_{it-1}^p + f_{it} \\ y_{it} &= y_{it}^p + v_{it} \end{aligned} \tag{4.6}$$

then, for quadratic utility, the change in consumption is due to the permanent shock, and the annuity value of the transitory shock, i.e.:

$$\varepsilon_{it} = f_{it} + \frac{r}{1+r} v_{it}$$

Blundell and Preston (1998) used this to show, in their proposition 3, that:

$$\Delta var^j(c_{it}) = var^j(f_{it}) + \left(\frac{r}{1+r}\right)^2 var^j(v_{it}) \tag{4.7}$$

If the permanent and temporary shocks are fully insured, then neither should affect consumption, which would imply:

$$\Delta var^j (c_{it}) = 0$$

This implication of full insurance was rejected by Deaton and Paxson (1994). An advantage of their approach, compared to the other two, lies in the weaker data requirements. In order to test for full insurance, researchers need only observe the cross-sectional moments of the distribution of consumption. In contrast, the first approach requires controlling for idiosyncratic versus aggregate income shocks, which unlike in the Deaton and Paxson case, may not be feasible with a time-series of cross-sections, while the second approach requires the availability of panel data, and for all taste-shifters to be fully parametrized. For these reasons we build from the approach of Deaton and Paxson (1994).[4]

While full insurance has been consistently rejected by all of the approaches discussed above, the tests are not fully constructive, in the sense that they do not help us to understand either how much insurance is available to agents, nor what mechanisms might change the overall level of insurance. The next two subsections discuss incomplete insurance under quadratic and alternative utility specifications.

4.3.1 Quadratic Utility

Under quadratic utility, when $r = \delta$, the permanent income hypothesis implies:

$$c_{it} = c_{it-1} + f_{it} + \frac{r}{1+r} v_{it}$$

[4] A further reason for not using the mean is that consumption growth not only depends on uncertainty (and hence the level of insurance), but also on the tax rate. Since we will look at the effect of the tax system on changes in consumption, and taxes affect the after tax interest rate, then it will also effect the expected growth rate of consumption. It will not, however, affect the higher moments of consumption.

If insurance is incomplete, so that the agent can insure the proportions ϕ of his permanent shock, and ψ of the temporary shock then:

$$\begin{aligned} c_{it} &= c_{it-1} + (1-\phi) f_{it} + (1-\psi) \left(\tfrac{r}{1+r}\right) v_{it} \\ \Rightarrow \quad \Delta var^j (c_{it}) &= (1-\phi)^2 var^j (f_{it}) + (1-\psi)^2 \left(\tfrac{r}{1+r}\right)^2 var^j (v_{it}) \end{aligned}$$

where the variance is taken for a fixed membership group j. In this framework, ϕ and ψ are determined by the characteristics of the economy in which the agent resides. They reflect *all* the mechanisms available to the agent that can insure idiosyncratic risk. Two cases are of particular interest: if $\phi = \psi = 1$ then this implies full insurance, while $\phi = \psi = 0$ reflects autarky, where none of the agents idiosyncratic risk can by insured. Further, if r is sufficiently small, or if the temporary shock can be fully insured ($\psi = 1$), then[5]:

$$\Delta var^j (c_{it}) = (1-\phi)^2 var^j (f_{it}) \tag{4.8}$$

From this last equation two things are immediately apparent. First, unless $\phi = 1$, the full insurance case, the variance (or standard deviation) of consumption will be increasing over time. Secondly, as ϕ increases (the amount of insurance increases) the growth in the variance of consumption falls.

4.3.2 Alternative Utility Specifications

For more complicated utility functions the simple relationship in equation 4.8 no longer holds. One of the simplest extensions is discussed by Attanasio and Jappelli (2001). Maintaining quadratic utility, but relaxing $r = \delta$, results in the Euler equation:

$$c_{it} = \left(\frac{1+\delta}{1+r}\right) c_{it-1} + \varepsilon_{it} \tag{4.9}$$

[5]This assumption may not be too unreasonable. Attanasio and Davis (1996) offer evidence that temporary shocks can be smoothed by US consumers, but permanent shocks can not.

where, under the assumptions made above on the income process, consumption innovations ε_{it} are equal to a linear transformation of income shocks, i.e.:

$$\varepsilon_{it} = f_{it} + \frac{r}{1+r} v_{it}$$

And under full insurance $var^j(c_{it})$ may be growing or declining depending on whether $r > \delta$ or $r < \delta$. With incomplete insurance and under the assumptions made above on income shocks the evolution of the variances for group j takes the form:

$$var^j(c_{it}) \quad = \left(\tfrac{1+\delta}{1+r}\right)^2 var^j(c_{it-1}) + (1-\phi)^2 \, var^j(f_{it}) \tag{4.10}$$

where the assumption that idiosyncratic shocks are all smoothed away is maintained.

Attanasio and Jappelli (2001) also discuss the case of isoelastic utility functions with relative risk aversion parameter σ. In this case, consumption innovations are not a simple function of income shocks, as in the quadratic utility case. If the distribution of consumption growth was log-normal, then the Euler equation takes the form:

$$\ln c_{it} = \ln c_{it-1} + \sigma(r-\delta) + \frac{1}{2\sigma} var(\Delta \ln c_{it}) + \varepsilon_{it}$$

If the fixed membership group j is sufficiently homogeneous, so that all members faced the same uncertainty about consumption growth, then the variance of consumptions evolves according:

$$var^j(\ln c_{it}) = var^j(\ln c_{it-1}) + (1-\phi)^2 \, var^j(\varepsilon_{it}) \tag{4.11}$$

Notice that the last term of the RHS of equation 4.11 is the variance for group j of the *consumption* innovation and that full insurance would now imply that there is no change in the variance of log-consumption over time.

Blundell and Preston (1998) discuss the cases of constant absolute risk aversion (CARA) and constant relative risk aversion (CRRA) preferences. For the first they use a relationship derived by Caballero (1990) who finds that when income is generated by a random walk and income shocks, ε_{it}, are log-normally distributed the growth of consumption is:

$$\Delta \ln c_{it} = \Gamma_{it} + \Lambda_{it} + \varepsilon_{it}$$

where Γ_{it} is the slope of the consumption path and Λ_{it} is a term that accounts for revisions to variance forecasts. Again, if the fixed membership group j is sufficiently homogeneous, so that Γ_{it} and Λ_{it} are constant within each group, then under incomplete insurance, the variance of the growth rate of consumption evolves according:

$$var^j (\Delta \ln c_{it}) = (1 - \phi)^2 \, var^j (\varepsilon_{it}) \tag{4.12}$$

Blundell and Preston (1998) also showed that a relationship of the same form as equation 4.12 hold when preferences are described by a CRRA functions and the income process is given by:

$$\ln y_{it} = \ln y_{it}^p + v_{it}$$
$$\ln y_{it}^p = \ln y_{it-1}^p + f_{it}$$

Namely, it holds that:

$$var^j (\Delta \ln c_{it}) = (1 - \phi)^2 \, var^j (f_{it}) + (1 - \psi)^2 \left(\frac{r}{1+r} \right)^2 var^j (u_{it}) \tag{4.13}$$

If, as before, we assume that either the interest rate is sufficiently small, or that $\psi \approx 1$, equation 4.13 boils down to equation 4.12.

Lastly, it is worth considering completely arbitrary preferences, that nevertheless maintain that the real interest rate and the discount rate are equal. Most generally, changes in marginal utility must obey the relationship:

$$\lambda (c_{it}) = \lambda (c_{it-1}) + \zeta_{it}$$

where $\lambda(\cdot)$ is the marginal utility of consumption and ζ_{it} is its innovation at time t. Writing λ_{it-1} in place of $\lambda(c_{it-1})$, the variance of the marginal utility consumption for group j under incomplete insurance evolves according to:

$$var^j(\lambda_{it}) = var^j(\lambda_{it-1}) + (1-\phi)^2 var^j(\zeta_{it}) \tag{4.14}$$

which implies that if $\phi \neq 1$ the variance of the marginal utility of consumption is increasing over time. There are two problems in bringing 4.14 to the data. First, the variance of the marginal utility can be computed only if one makes specific assumptions on preferences. Second, the variance of the innovation in the marginal utility of consumption is not observable.

In the following, we show how to tackle the first problem, while the next subsection shows how our identification procedure tackles the second problem. The Euler equation can be used to characterize the *changes* in the variance the marginal utility of consumption. The first order Taylor expansion of it around c_{it-1} gives:

$$\lambda_{it} = \lambda_{it-1} + \lambda'_{it-1}(c_{it} - c_{it-1}) + o(c_{it} - c_{it-1}) \tag{4.15}$$

Hence the variance of (4.15) is:

$$\begin{aligned} var^j(\lambda_{it}) &= var^j(\lambda_{it-1}) + E^j(\lambda'_{it-1})^2 var^j(\Delta c_{it}) \\ &\quad + var^j(\lambda'_{it-1}) E^j(\Delta c_{it})^2 + var^j(\lambda'_{it-1}) var^j(\Delta c_{it}) \end{aligned} \tag{4.16}$$

where $E^j(\cdot)$ is the expected value for fixed membership group j. From equation (4.16) one can see that if the variance of the change in consumption increases then the change in the variance of the marginal utility of consumption also increases. This implies that those factors that cause the variance of the change in consumption to decrease cause the change of the marginal utility of consumption to decrease too, i.e. the overall amount of insurance available to individuals to increase. This allows us to identify what provides insurance at the margin, even without making the assumptions about the functional form of the utility function. In other words, there is a monotone mapping between the changes in the variance of marginal utility of consumption and the variance of changes.[6]

4.3.3 The regression

The discussion above highlights that the sign of the effect of a policy instrument, as captured by ϕ, can be always identified, while the coefficient is interpretable only by assuming quadratic utility. In the discussion on quadratic utility above, we considered $\Delta var^j(c_{it})$, but if, in equation 4.5, we first difference before taking variances, we obtain:

$$var^j(\Delta c_{it}) = var^j(\varepsilon_{it}) \tag{4.17}$$

and the implications are the same. However, in this approach, no assumption about the $cov^j(c_{it-1}, \varepsilon_{it})$ is necessary, although the disadvantage is that each household must be observed at least twice. The discussion on alternative utility specifications, culminating in equation 4.16, also suggests using the variance of changes rather than changes in variances. The third reason for doing this is it allows us to easily measure the insurance effect of our policy instrument. This last follows since equation 4.8 becomes:

$$var^j(\Delta c_{it}) = (1-\phi)^2 var^j(f_{it}) \tag{4.18}$$

which implies that if instead the standard deviations were taken on each side, then:

$$s.d.^j(\Delta c_{it}) = (1-\phi)\, s.d.^j(f_{it}) \tag{4.19}$$

Recall that ϕ is the amount of insurance available to households within each regime. This amount may depend on some policy instrument z^j which varies exogenously across regimes.

[6]One could obtain an entirely similar derivation using logs instead of levels.

In which case, estimation can recover ϕ as a function of z^j. Variation between regimes allows this function to be recovered. Writing:

$$\phi\left(z^j\right) = \alpha_0 + \alpha_1 z^j + \xi^j \tag{4.20}$$

where the error term ξ^j captures any differences in the level of insurance that are not modeled. Full insurance implies that both $\alpha_0 = 1$ and $\alpha_1 = 0$; incomplete insurance implies $\alpha_0 < 0$; while z_j provides additional insurance if $\alpha_1 > 0$. If instead z^j merely crowded out, or substituted, for other insurance mechanisms, then this would imply $\alpha_1 = 0$, while a negative coefficient would imply that there was more than complete crowding out (that is the overall level of insurance is reduced). Substituting into equation 5.1 results in:

$$s.d.^j\left(\Delta c_{it}\right) = \left(1 - \alpha_0 - \alpha_1 z^j + \xi^j\right) s.d.^j\left(f_{it}\right) \tag{4.21}$$

The regression that is run takes the form:

$$s.d.^j\left(\Delta c_{it}\right) = \beta_0 + \beta_1 z^j + error \tag{4.22}$$

and identifying terms implies:

$$\begin{aligned} \beta_0 &= \left(1 - \alpha_0\right) s.d.\left(f_{it}\right) \\ \beta_1 &= -\alpha_1 s.d.\left(f_{it}\right) \\ error &= \xi^j s.d.\left(f_{it}\right) \end{aligned}$$

Hence if the variance of the permanent shock were known and constant across all groups j, the level of insurance, and how it changes with the policy instrument z^j, can be recovered.

For departures from the simple cases, the coefficient is no longer so easily identified but, the sign can always be interpreted. This is illustrated in figure 4.1. The top panel shows the simple quadratic case where $r = \delta$. Full insurance implies that $var^j\left(c_{it}\right)$ is time-invariant, while full risk-sharing is rejected if the variance is growing over time. This is longer true for more general preferences: instead full risk-sharing may imply that $var^j\left(c_{it}\right)$ is either increasing or decreasing, depending on the parameters of the model. The diagram shows the variance decreasing under full insurance. However, our aim is to compare the different policy regimes, indexed by z^j, hence we compare the $var^j(c_{it})$ across the different regimes: in those regimes where this variance is growing more slowly, there must be more risk-sharing or insurance. This remains true even if the full insurance case can not be fully specified.

4.4 Data

This chapter uses the Consumer Expenditure Survey (CEX) which has been described earlier. Information on state of residence will be exploited. Comparing these states will enable us to test whether the tax (and transfer) system can provide insurance as different states have different tax policies. This chapter uses a sample of around 100,000 households taken from the years 1980 to 1998 for which we have full state information. The data was also restricted to those households headed by individuals who were between 30 and 50 in the first year in which they are sampled.[7]

In the US taxes are raised at a variety of levels; those entitled to levy taxes include the federal and state legislators, county administrations, and school boards. Taxes include income taxes, sales taxes, property taxes and duty. This chapter will concentrate on income tax, which is raised both federally and by states. While the consumer expenditure survey includes questions on the amount of taxes that the household pays, we do not believe that the answers that households give are particularly accurate or reliable. Nevertheless, some results are included using these variables. However, we have also constructed a tax liability figure using the TAXSIM programme developed by Feenberg (see Feenberg and Coutts, 1993 for details) and provided by the NBER. Using a variety of household variables, including husband's and wife's salary income, taxes and other costs on property, interest, dividends

[7]This means if we are taking the *k-th* difference they were aged 30+k to 50+k in the last year in which they were sampled.

Figure 4.1: The evolution of $var^j(c_{it})$.

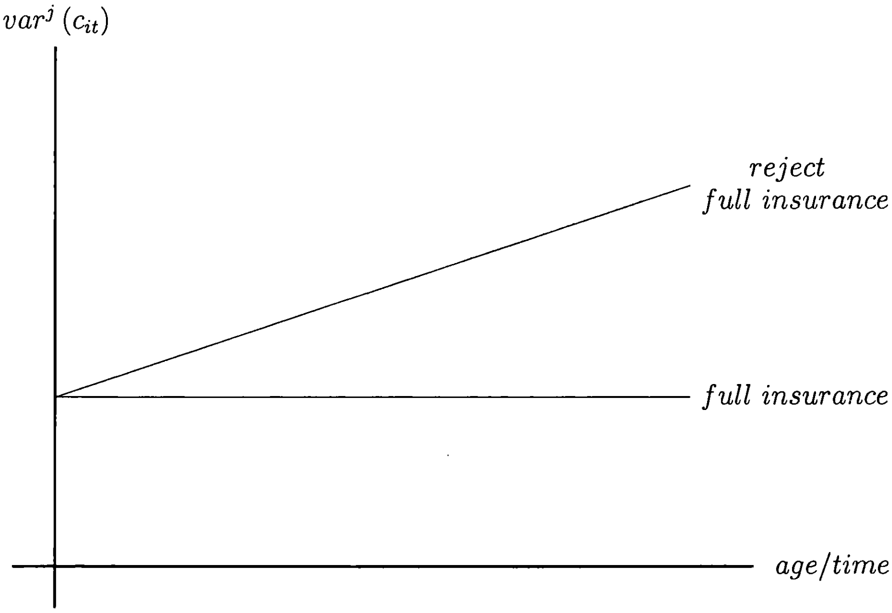

(a) Quadratic preferences where $r = \delta$.

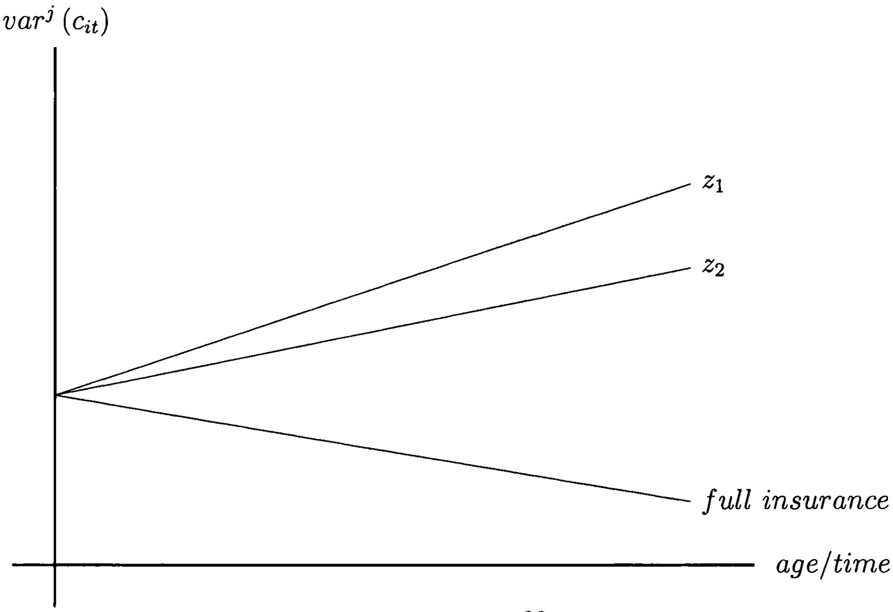

(b) A more general case, comparing regimes z_1 and z_2.

etc., and details about the household's characteristics (such as number of dependent children) as well as the state of residence, the programme constructs both the state and the federal tax bracket, tax liability, and marginal tax rate for each household in the sample. The calculation specifically allows for the fact that a variety of allowances are allowed, and also includes an assessment of households entitlement to the Earned Income Tax Credit[8](whic h results in some households net tax liability being negative). Tables 4.1-4.5 summarize some of the main features of the data.

The current federal marginal tax rate (see table 4.1) varies from 15% for those whose income is less than $26,250 (for single people, for married couples the threshold is $43,850) up to 39.6% for income over $288,350. Table 4.2 summarizes the proportion of people in each tax bracket in our sample. It also highlights that the brackets themselves have varied over the years. For the period 1982-1986 a large number of tax brackets were applicable: too many to give anything other than fairly broad summary statistics. However, the table demonstrates that around 15% (depending on the year) did not pay any tax, while the median household was in the 23% tax bracket. In these years the highest tax bracket for federal taxes was set at 50%. In 1987 the number of brackets was greatly reduced, and from this year every tax bracket is recorded in table 4.2. There was a further reduction in the following year. Between 1987 and 1996 the proportion of households who did not pay tax gradually increased to 19%. Other features are the introduction of a 31% tax bracket, for the top 5% of earners in 1991, the introduction of a 36% bracket in 1993, of a 39.6% bracket in 1996, and the abolition of the 0% bracket in 1997.

States taxes can differ quite widely among the different US states. Table 4.3 shows the current tax rates applicable in different US states. From this we can see 8 states, including Texas and Florida, do not to levy any income tax on their residents: state revenue in these states comes mostly from sales taxes. In addition New Hampshire and Tennessee only charge

[8]See Scholz (1996) for an explanation of the EITC.

state income tax against dividend and interest income. The other states have a variety of income tax bands and exemptions (or tax credits) that are applicable. Although some states have a flat rate tax, in most states, the marginal tax rate increases with income, and there are a variety of tax allowances to which households are entitled. Table 4.3 provides a summary description of the state tax systems in the U.S.

The chapter also exploits the transfers that typically poorer agents receive. Such transfers include social security and railroad retirement income, supplementary security income, unemployment compensation, worker's compensation and veterans payments, public assistance or welfare, pension income, and the value of food stamps received: the CEX includes questions on all these transfers. Table 4.4 shows that the average amount of transfer, over the whole sampled population, is \$871, but that only 18.6% of households receive a transfer. Conditional on receiving at least something, households receive an average of \$4,680. This is not a substantial amount, when the average household salary income in the survey is \$34,574.

4.4.1 Measuring Tax Progressivity

In order to assess if tax systems are affecting the level of insurance available to households we need some measure of the progessivity[9] of the tax system in each state. If the marginal tax rate were the same for all households, then this could be used in the regressions. As it increases we should expect the cross-sectional variance of consumption to increase. The intuition is that higher marginal tax rates cause the after taxes income distribution to be more concentrated. However, from the previous discussion, the marginal tax rate that consumers face is an increasing function of income, and furthermore, households differ in the allowances and exemptions they can claim. Under such circumstances, no completely sat-

[9]Recall that we define a tax to be progressive if the tax liability increases as taxes increase, thus a flat rate income tax is progressive.

isfactory measure of tax progressivity exists. We choose to use two measures that capture progressivity. The first is to take the average of the households marginal tax rate, or tax bracket, within each year-state-cohort j. From table 4.5 the average federal bracket is 20.2%, and the average of the reported marginal tax rates (which accounts for various allowances) is 19.2%. The state rates vary from zero in Texas and Florida, which charge no taxes, to an average marginal tax rate of 7.4% in New York.

One problem with using this measure is that it ignores any heterogeneity in tax rates across households. For instance, a given mean marginal tax rate of 20% could be due to all households paying a marginal (and average) tax rate of 20%, from only the top half of the income distribution paying, but paying 40%, or from only the bottom half paying 40%. These situations differ strongly in the degree to which taxes are re-distributive. A related problem, is that the mean marginal tax rate is not invariant to the overall level of taxes that are taken. Our interest is in the degree to which taxes re-distributes income, hence we wish to construct a measure that is not affected by the overall tax level. This motivates a second measure of how much the tax system redistributes income, constructed as:

$$\tau = 1 - \sqrt{\frac{var^j\left(\tilde{y}_{it}\right)}{var^j\left(y_{it}\right)}} \tag{4.23}$$

that is, as one minus the square-root of the ratio of the variance of income after taxes compared to before taxes for each group j. If all households faced the same marginal tax rate, and there were no allowances, then this constructed τ would exactly equal the marginal tax rate. Moreover, a larger τ implies more re-distribution. Table 4.5 displays the constructed values for τ for the whole of the US and for the 6 largest US states. It shows that this measure averages to 32.0% over the US, but that states can differ from 27.6% in Florida (where there is no income tax), to 36.8% in New York, traditionally viewed as one of the more progressive states. This measured τ can be regressed on how much the variance of consumption changes over time. A negative coefficient implies that the variance is growing more slowly, which is interpreted as meaning more insurance is being provided by the tax system. That is, a negative coefficient means that as taxes become more re-distributive, agents are better able to smooth against uncertain income shocks. In contrast, a positive coefficient implies that taxes more than crowd out alternative insurance mechanisms, and the overall level of insurance from the tax system is reduced.

4.5 Results

Section 4 motivated regressing the change in the variance of consumption (or the variance of thc change) against a measure of tax progressivity. The regression takes the form:

$$\Delta^k var^j (c_{it}) = \beta_0 + \beta_1 \tau_t^j + \varepsilon^j$$

where j represents the state-time combination, and we make explicit that the difference is over k periods, where a period is one year. The number of groups that could reasonably be defined (ensuring that the cell size was at least 75[10]) was around 180, but depended on the regression. Obviously, this process meant that several smaller states were never included in the analysis, since fewer observations came from these states. However, there were sufficient observations in most of the larger states. The regression also assumes that states do not endogenously choose their tax rates.

The results of the regressions are recorded in tables 4.6-4.9 below. In each case the basic regression in column 1 takes households who are 25-55 years of age for the years 1982-1997. However the variance of the growth rate of households within this group is likely to display considerable predictable heterogeneity.[11] Hence, column 2 further controls for a full set of predictable, and predetermined set of taste-shifters.[12] That is, instead of using $s.d. (\Delta c_{it})$,

[10]experimenting with different cell sizes did not qualitatively change the results.

[11]The basic regression only controls for seasonality.

[12]The set will also control for predictable income changes, if predictable income changes impact on changes in current consumption.

column 2 uses $s.d.\left[\Delta c_{it} - E\left(\Delta c_{it}|X_{it}\right)\right]$ The set of control variables X includes age, time, education, sex, race, marital status, and changes in family size.[13]

A second important issue is that we wish to observe the standard deviation of consumption, and for the measure of compression for the whole population in each group, but we only observe a small sample of the population. Thus errors are introduced into both the left-hand side and into the right-hand side of each regression. That is, we have

$$\Delta c_{it} + \varepsilon_{jt} = \beta_0 + \beta_1\left(\tau_{jt} + \varsigma_{jt}\right) + u_{jt} \tag{4.24}$$

Within each state one would expect past values of τ to be correlated its current value, but measurement error induced by the small sampling sizes to be uncorrelated with the past, hence instrumenting with lagged values of τ will remove any downward bias. A further problem is that ε_{jt} and ς_{jt} may be correlated, hence for some of the regressions τ twice lagged will be the instrument.

Included are regressions for the level and the log-level of consumption. The regression constructs cells based on state and cohort, where the cohorts are defined as those who are 25-40 and those who are 41-55. The second give slightly nicer results, and are the ones we should include in the chapter. In each case I have run regressions for the 'basic' regression, which merely controls for seasonality, consumption controlling for a full set of controls, and some instrumented regressions, where the instruments are τ (or the marginal tax rate) lagged once, lagged once or twice, or lagged twice. The significance of the results is starred once for significant at the 10% level, twice for significant at the 5% level and three times for significant at the 1% level. For the regressions using the mean tax rate mean marginal tax rate, and the tax brackets, generated by TAXSIM, and includes both federal and state taxes. The regressions with our constructed measure of tax progressivity, uses the before and after

[13]One issue is that this set may not control for all predictable taste-shifters. Also, there may be errors in measuring $E\left(\Delta c_{it}|X_{it}\right)$.

tax income reported by households in the CEX, taxes only calculated from the TAXSIM, and finally taxes from the TAXSIM programme together with reported transfers.

Table 4.6 shows that in the basic regression, and in the regression including a set of control variables, that the coefficient on the mean marginal tax rate, and on the tax bracket, is significant at the 1% level. However, perhaps surprisingly, the coefficient is positive. This suggests that other important insurance mechanisms are being crowded out by the introduction of a more re-distributive tax system. When the instrumented results are included, the mean marginal tax rate remains significant, at the 5% level, although this is not true for the mean tax bracket. When the change in the variance of log-consumption was considered in table 4.7 the coefficient on the constant becomes highly significant (at the 1% level), and the marginal tax rate, and the tax bracket both remain significant. This remains true when τ is instrumented by its lag, and even in the last column, where it is lagged twice. The results in this last column are significant at the 1% level.

Tables 4.8 and 4.9 consider the constructed value of re-distribute that was created using before and after tax income. The first results refer to levels, and show, except in the less reliable case where we have used the reported tax liabilities in the CEX, that τ is positive and significant at the 5% level. However, when τ is instrumented by past values of the variable, the results are more ambiguous, and only remain significant for the middle panel that ignores transfer income. When log-consumption is considered the results are significant at the 1% level, with and without the controls and for both τ and the constant. The significance of the results remain with the IV-estimate. The last column, our most preferred regression, is also significant.

The fact that the constant was significant and positive in the regressions, at least for the log-consumption regressions, confirms the results in Deaton and Paxson (1994). This rejects full insurance, although it is open to the criticism made in Attanasio and Jappelli (2001). The results for the coefficient for the measure of tax-progressivity have a number of interpretations. At the most basic, the coefficient on τ can be treated as an exclusion restriction: under full insurance the growth in the variance of consumption, regardless of the utility function, should be uncorrelated with anything except changes in tastes. Hence the fact that τ enters significantly rejects the full insurance hypothesis. A second interpretation says that while the coefficient is not interpretable, the fact that the coefficient is positive, and significant, means that making the tax system more re-distributive is reducing the amount of insurance that the agents have. There must be other mechanisms that operate privately, that are being disrupted by the imposition of the tax system: agents incentives to participate in these private insurance mechanisms is reduced since they have access to a public mechanism. However, how these private mechanisms operate is not calculated.

The final interpretation allows us to quantify how much dis-insurance the tax system provides. The theory section highlighted when this case arises and the terms can be identified:

$$\begin{aligned} \beta_0 &= (1-\alpha_0)\, s.d.\,(f_{it}) \\ \beta_1 &= -\alpha_1 s.d.\,(f_{it}) \\ error &= \xi^j s.d.\,(f_{it}) \end{aligned}$$

For identification it is necessary to know, or estimate, the standard deviation of the permanent shock. For log-income, this has been done by MaCurdy (1982) among others.[14] His paper estimated the variance of the permanent shock to be 0.34: the implications for the estimated dis-insurance that is thus generated by a more re-distributive tax system is tabulated in table 4.10. From the table, both the basic, and the control regression suggests, for either the marginal tax rate or the tax bracket, that households can insure roughly one third of their permanent income shock. Each one percent increase in the marginal tax rate reduces the level of insurance by 0.7%. Instrumenting both increases the estimated basic level of insurance and the estimated reduction in insurance induced by the tax system. Those

[14]We are not aware of any papers that have estimated this parameter when income has been measured in levels, and hence concentrate on the log-level case.

regressions that use the constructed τ proposed in equation 4.5 show similar results for the basic and the control regression. When instruments are used the results that used the estimated taxes from the TAXSIM routine described above show that households can insure roughly 45% of their permanent shock, but that each 1% increase in the measure of how re-distributive the tax system is reduces the amount of insurance by 1%. To put this in perspective, these results imply that if we examine table 4.5, if a household was moved from California to Florida, then the total amount of the permanent shock that the agent could insure would increase from around 9% to around 18%, that is it would roughly double.

4.6 Conclusions

In this chapter we tried to relate the amount of insurance available to households to the system of taxes and transfers. A large strand of literature has documented the absence of full insurance, at least for the U.S. economy. However, it is likely that households *can* smooth some of their idiosyncratic shocks. Several mechanisms might provide insurance: some of them, such as financial markets, are related to formal economic interactions; others, such as family networks, relate to informal interactions. The tax system is the mechanism this chapter explores.

Taxes and transfers re-distribute income across households, but also might smooth some of the income shocks households incur. This in turn affects how consumption changes over time. If the changes in consumption respond to idiosyncratic shocks, we should expect this response to be weaker as more insurance becomes available to households. Moreover, if the cross-sectional variance of consumption trends up, we should expect this trend to be flatter, the more households are able to insure their consumption against idiosyncratic income shocks. Around these two intuitions, the empirical results of the chapter are organized.

We test if changes in consumption respond to the degree of compression of the income distribution induced by taxes and transfers. We measure the degree of distribution induced by taxes and transfers with (a) the marginal tax rate, and the tax bracket; (b) a constructed measure of how much taxes compress the distribution of current income. Our regressions show that increasing the degree to which taxes re-distribute income is negatively correlated with the increase in the variance of consumption growth. The chapter suggests three interpretations: (i) this can be thought of as an exclusion restriction that confirms the rejection of full insurance found by previous authors; (ii) the fact that the coefficient is positive suggests that increasing the degree to which taxes re-distribute income reduces the total amount of insurance available to households; and (iii) under stronger assumptions, the coefficients themselves can be interpreted, and the results suggest that moving from a highly re-distributive state, such as California, to one that re-distributes less, such as Florida, can double the amount of insurance that is available to households.

The overall, and tentative conclusion of this chapter, is that the tax system seems to be rather a poor mechanism for providing insurance against the idiosyncratic income shocks that households suffer. Instead, the tax system crowds out other mechanisms which we do not attempt to describe, although the methodology described in this chapter can be extended to any potential insurance mechanism: we merely need to observe sufficient variation in our sample. This result suggests that any defense of a re-distributive tax system that appeals to the self-interest of the agents involved must argue on the grounds that *ex ante* inequality is itself a 'bad thing' since the results do not support the suggestion made by Varian (1980) that progressive taxes can be motivated by agents insurance incentives, at least for the ranges of taxes that are observed in the US.

Lastly, one could speculate on what might affect the level of insurance in the economy. This paper finds that the tax system, at the margin, does not have much effect. It seems that households are able to co-operate and self insure most of the risks they face. Possible mechanisms may be transfers made between neighbours or close relatives, transfer from parents to children (poorer children getting more than richer children) or perhaps some other mechanism. This paper gives a methodology in which these other mechanisms can be investigated.

Table 4.1: Thresholds for current federal tax brackets

Tax Rate	Tax Bracket		
(%)	single	married jointly	married seperately
15	0	0	0
28	26,250	43,850	21,925
31	63,550	105,950	52,975
36	132,660	161,450	80,725
39.6	288,350	288,350	144,175

Table 4.2: Proportion paying at each marginal federal tax rate for 1982-1998

Year	Tax Bracket	% Paying Rate	Year	Tax Bracket	% Paying Rate
1982	0	13.4			
	10-20	19.5			
	21-30	30.8			
	31-40	25.2			
	41-50	11.1			
1983	0	14.6	1984	0	15.8
	10-20	26.8		10-20	28.2
	21-30	31.6		21-30	29.0
	31-40	21.9		31-40	19.7
	41-50	5.2		41-50	7.4
1985	0	15.2	1986	0	15.5
	10-20	26.3		10-20	26.6
	21-30	28.4		21-30	28.4
	31-40	22.5		31-40	22.3
	41-50	7.7		41-50	7.3
1987	0	17.1	1988	0	17.0
	11	2.9			
	15	38.0		15	42.4
	28	23.3		28	40.5
	35	15.9			
	38.5	2.6			
1989	0	17.2	1990	0	17.4
	15	42.4		15	42.3
	28	40.2		28	40.2
1991	0	17.1	1992	0	18.6
	15	43.9		15	43.0
	28	33.2		28	31.5
	31	5.6		31	6.7
1993	0	18.0	1994	0	17.8
	15	44.9		15	43.7
	28	30.8		28	32.6
	31	5.5		31	5.1
	36	0.5		36	0.6
1995	0	17.6	1996	0	19.4
	15	43.2		15	42.1
	28	33.3		28	32.6
	31	5.1		31	4.3
	36	0.5		36	1.2
				39.6	0.1
1997	15	61.1	1998	15	58.2
	28	32.1		28	34.2
	31	4.9		31	5.2
	36	1.4		36	1.8
	39.6	0.2		39.6	0.3

Table 4.3: State Individual Income Tax Rates in the US in 1999

State	Tax Rates		Exemptions		
	low	high	single	married	dependents
Alabama	2.0	5.0	1,500	3,000	300
Alaska	no state tax				
Arizona	2.87	5.04	2,100	4,200	2,300
Arkansas	1.0	7.0	20*	40*	20*
California	1.0	9.3	72*	142*	227*
Colorado	4.63	4.63		none	
Connecticut	3.0	4.5	12,000	24,000	0
Delaware	2.2	5.95	110*	220*	110*
Florida	no state tax				
Georgia	1.0	6.0	2,700	5,400	2,700
Hawaii	1.5	8.5	1,040	2,080	1,040
Idaho	2.0	8.2	2,900	5,800	2,900
Illinois	3.0	3.0	2,000	4,000	2,000
Indiana	3.4	3.4	1,000	2,000	1,000
Iowa	0.36	8.98	40*	80*	40*
Kansas	3.5	6.45	2,250	4,500	2,250
Kentucky	2.0	6.0	20*	40*	20*
Louisiana	2.0	6.0	4,500	9,000	1,000
Maine	2.0	8.5	2,850	5,700	2,850
Maryland	2.0	4.75	1,850	3,700	1,850
Massachusetts	5.6	5.6	4,400	8,800	1,000
Michigan	4.2	4.2	2,800	5,600	2,800
Minnesota	5.35	7.85	2,900	5,800	2,900
Mississippi	3.0	5.0	6,000	12,000	1,000
Missouri	1.5	6.0	2,100	4,200	2,100
Montana	2.0	11.0	1,610	3,220	1,610

*Tax Credits.

Table 4.3: (cont.) State Individual Income Tax Rates in the US

State	Tax Rates		Exemptions		
	low	high	single	married	dependents
Nebraska	2.51	6.68	91*	182*	91*
Nevada	no state tax				
New Hampshire	taxes unearned income only				
New Jersey	1.4	6.37	1,000	2,000	1,500
New Mexico	1.7	8.2	2,900	5,800	2,900
New York	4.0	6.85	-	-	1,000
North Carolina	6.0	7.75	2,500	5,000	2,500
North Dakota	2.67	12.0	2,900	5,800	2,900
Ohio	0.691	6.98	1,050	2,100	1,050
Oklahoma	0.5	6.75	1,000	2,000	1,000
Oregon	5.0	9.0	132*	264*	132*
Pennsylvania	2.8	2.8		none	
Rhode Island	25.5% of federal taxes				
South Carolina	2.5	7.0	2,900	5,800	2,900
South Dakota	no state tax				
Tennessee	taxes unearned income only				
Texas	no state tax				
Utah	2.3	7.0	2,175	4,350	2,174
Vermont	24% of federal taxes				
Virginia	2.0	5.75	800	1,600	800
Washington	no state tax				
West Virginia	3.0	6.5	2,000	4,000	2,000
Wisconsin	4.6	6.75	700	1,400	400
Wyoming	no state tax				
Dist. Columbia	5.0	9.0	1,370	2,740	1,370

*Tax Credits.

Table 4.4: The level of transfers in the US 1982-1998

transfer	average	average if received	% receive
social security	247	6,710	3.6
supplementary security income	74	3,328	2.2
unemployment compensation	160	2,439	6.5
worker's compensation	117	4,484	2.6
welfare	169	3,768	4.4
pension	266	8,815	3.0
food stamps	101	1,918	5.3
total *	871	4,680	18.6

*Excluding pension income

Table 4.5: Measuring tax progressivity

	mean marginal rate	mean tax bracket	τ
Federal	19.2	20.2	
State:			
Overall	3.7	4.2	32.0
California	5.0	5.3	36.4
Florida	-	-	27.6
New York	6.3	7.4	36.8
Ohio	3.8	4.0	33.0
Pennsylvania	2.2	2.4	29.8
Texas	-	-	28.9

*The results are for the whole sample 1982-1998.

Table 4.6: Regressing $\Delta^k sd^j(c_{it})$ against the mean tax rate (standard errors in parentheses).

	basic	control	Instrument τ_{t-1}	Instrument τ_{t-1} and τ_{t-2}	Instrument τ_{t-2}
Marginal Tax Rate					
τ	1.633***	1.636***	1.590**	2.146**	2.126**
	(0.590)	(0.583)	(0.775)	(0.854)	(0.846)
constant	0.177	0.174	0.189	0.081	0.086
	(0.127)	(0.125)	(0.172)	(0.184)	(0.182)
Tax Bracket					
τ	1.036**	1.054**	0.805	1.163	1.294*
	(0.521)	(0.515)	(0.668)	(0.754)	(0.708)
constant	0.294**	0.288**	0.349**	0.277	0.247
	(0.126)	(0.124)	(0.163)	(0.179)	(0.167)
N	185	185	176	166	166

Table 4.7: Regressing $\Delta sd^j(\ln c_{it})$ against the mean tax rate (standard errors in parentheses).

	basic	control	Instrument τ_{t-1}	Instrument τ_{t-1} and τ_{t-2}	Instrument τ_{t-2}
Marginal Tax Rate					
τ	0.396*	0.402**	0.481**	0.574**	0.678***
	(0.198)	(0.194)	(0.207)	(0.232)	(0.233)
constant	0.390***	0.379***	0.358***	0.339***	0.317***
	(0.042)	(0.041)	(0.045)	(0.049)	(0.050)
Tax Bracket					
τ	0.391**	0.428**	0.377**	0.423**	0.512***
	(0.172)	(0.168)	(0.148)	(0.172)	(0.175)
constant	0.379***	0.367***	0.376***	0.366***	0.346***
	(0.039)	(0.038)	(0.035)	(0.039)	(0.040)
N	185	185	176	166	166

Table 4.8: Regressing $\Delta^k sd^j\,(c_{it})$ against the τ^j (standard errors in parentheses).

	basic	control	Instrument		
			τ_{t-1}	τ_{t-1} and τ_{t-2}	τ_{t-2}
Taxes from CEX					
τ	1.309	1.322	1.575	2.422	2.781
	(0.845)	(0.843)	(1.494)	(1.715)	(1.928)
constant	0.229	0.226	0.169	-0.021	-0.106
	(0.203)	(0.202)	(0.359)	(0.409)	(0.460)
Taxes only					
τ	1.910**	1.912**	1.851	2.812**	2.995**
	(0.829)	(0.826)	(1.382)	(1.368)	(1.414)
constant	0.098	0.097	0.117	-0.091	-0.133
	(0.190)	(0.189)	(0.323)	(0.314)	(0.322)
Taxes and Transfers					
τ	1.459**	1.459**	1.452	1.838	1.629
	(0.724)	(0.721)	(1.368)	(1.345)	(1.351)
constant	0.166	0.166	0.172	0.084	0.137
	(0.185)	(0.184)	(0.356)	(0.347)	(0.349)
N	188	188	179	169	169

Table 4.9: Regressing $\Delta^k sd^j(\ln c_{it})$ against the τ^j (standard errors in parentheses).

	basic	control	Instrument τ_{t-1}	Instrument τ_{t-1} and τ_{t-2}	Instrument τ_{t-2}
Taxes from CEX					
τ	0.374**	0.395**	0.619**	0.627**	1.161**
	(0.181)	(0.182)	(0.279)	(0.279)	(0.485)
constant	0.381***	0.376***	0.319***	0.317***	0.192*
	(0.042)	(0.042)	(0.067)	(0.067)	(0.115)
Taxes only					
τ	0.520***	0.519***	0.589**	0.624**	0.633*
	(0.184)	(0.184)	(0.279)	(0.287)	(0.356)
constant	0.350***	0.350***	0.331***	0.324***	0.322***
	(0.041)	(0.041)	(0.065)	(0.065)	(0.081)
Taxes and Transfers					
τ	0.493***	0.490***	0.634***	0.653***	0.611**
	(0.185)	(0.185)	(0.243)	(0.251)	(0.305)
constant	0.344***	0.344***	0.304***	0.300***	0.311***
	(0.046)	(0.046)	(0.062)	(0.063)	(0.077)
N	186	186	176	166	166

Table 4.10: Estimated parameters in the insurance function $\phi(\tau) = 1 - \alpha_0 - \alpha_1 \tau$ from tables 4.7 and 4.9.

	basic	control	Instrument		
			τ_{t-1}	τ_{t-1} and τ_{t-2}	τ_{t-2}
Marginal tax rate					
α_0	0.33	0.35	0.38	0.41	0.45
α_1	-0.67	-0.69	-0.82	-0.98	-1.16
Tax Bracket					
α_0	0.35	0.37	0.35	0.37	0.40
α_1	-0.67	-0.73	-0.64	-0.72	-0.87
Incomes from the CEX					
α_0	0.34	0.35	0.45	0.45	0.67
α_1	-0.64	-0.67	-1.06	-1.07	-1.99
Taxes only using TAXSIM					
α_0	0.39	0.39	0.43	0.44	0.44
α_1	-0.89	-0.89	-1.01	-1.07	-1.08
Taxes and Transfers using TAXSIM					
α_0	0.41	0.41	0.47	0.48	0.46
α_1	-0.84	-0.84	-1.08	-1.11	-1.04

Chapter 5

Estimating Credit Constraints among US households.

abstract

It is often argued that credit constraints are pervasive: agents can not borrow as much as they would like. This chapter shows how to separately identify the demand for debt, and the maximum amount that agents are allowed to borrow, even if it is not known which consumers are credit constrained. Identification requires making appropriate exclusion restrictions. Using this approach, an estimate of the prevalence of credit constraints is constructed, and of differences across sub-groups. This chapter also estimates how much more such households would like to borrow. However, it is not possible to test this model against models in which credit constraints are never binding. The methodology is implemented using US household data (the CEX survey), finding that between 26% and 31% of households are credit constrained: the poorly educated and ethnic minority and men groups are less often credit constrained, while credit constraints decline gently with age.

5.1 Introduction

It is widely believed by economists, in almost all branches, that at least some agents are credit-constrained. It is offered as one of the most important explanations for a wide-variety of phenomena that are observed in economics, and also implicitly informs the policy debate, not only at the macro-economic level in motivating fiscal and monetary policy, but it also motivates micro policies such as subsidising university education for under-graduates, or supporting small business investment. While credit constraints are widely seen as pervasive, little is known about its incidence or importance. This chapter attempts to look at this issue while particularly investigating consumer borrowing behaviour.

When considering how much consumers borrow, several questions may be of interest. (i) Are there restrictions on the amount that consumers can borrow, that is, are at least some consumers credit constrained? (ii) How many consumers are credit constrained? (iii) How do credit constraints differ with household characteristics? (iv) How much more would these consumers borrow if unconstrained?

The approach taken in this chapter is to briefly characterize what it means for a consumer to be credit-constrained, and then to use a very simple form of credit-constraints that the literature motivates to start attempting to answer some of the questions that are raised above. As will be seen, a simple way to characterize debt holdings is to think of actual debt as the minimum of the amount of debt that the household wishes to hold (the demand for debt), and the amount that any lender is prepared to lend (the supply of debt). While in this framework, for reasons to be discussed below, it is not possible to test the first of the questions that we might want to answer, the other questions can be answered.

The chapter starts with a brief review of the literature in section 2. This discussion is restricted to the literature on consumption, and will help to motivate the rest of the chapter. Section 3 proposes an estimation strategy, and explains that identification requires exclusion restrictions on the parameters that enter the demand and supply equation respectively.

Estimation must cope with two problems. Credit-constraints are not directly observed, but must be replaced with some proxy variable that will be a function of observed characteristics. Further, there is a selectivity issue since even if it was observed which households were credit-constrained, demand (or supply) conditional on the household being unconstrained does not equal the unconditional demand for debt. Having estimated these equations, section 4 will then recover the estimated incidence of credit-constraints among US households, and how this differs as household characteristics change. This section will also discuss how to recover a measure of how much more credit constrained households wanted to borrow. The selection issue involved in estimation will also be important here as it requires the construction of the difference between demand and supply conditional on being credit-constrained and the former can not be replaced by, for instance, with the unconditional demand for debt without downward biasing the results.

The data used is a sample of around 10,000 households from the US consumer expenditure survey for the years 1988-1994. Preliminary results show that between 26% and 31% of households are credit-constrained. This figure is higher than several previous studies. It also shows that credit constraints are less serious for married couples, and for white households. However, credit-constraints decline only gently with age while state banking regulations that are designed to encourage competition have little effect.

5.2 The Literature On Credit Constraints.

Ever since the seminal paper by Hall (1978) consumption economists have questioned simple versions of the life-cycle/permanent income hypothesis (PIH). That chapter, and an enormous number of succeeding papers, have rejected the Euler equation formulation for consumption. Another strand of the literature showed how consumption tracked income over the life-cycle, which again rejects simple versions of the PIH.[1] A number of explanations for this have been suggested in the literature: one of the most popular is that at least some consumers face binding credit constraints. These consumers would like to borrow more in order to increase their level of consumption, such consumption is compatible with their life-cycle budget constraint, but for some reason they are not able to borrow as much as they would like at the 'market clearing' interest rate. For instance, Hayashi (1987) defined consumers as credit constraints if either (i) "they face some quantity constraint on the amount of borrowing", or (ii) "the loan rate available to them is higher than the rate at which they could lend". The first is often called *credit rationing*, and there is a large literature, going back to Jaffee and Russell (1976), or Stiglitz and Weiss (1981), which motivates such credit rationing by lenders as due to the fact that some consumers default on their loans, and there is imperfect information as to which agents will default. The decision to default is not modelled. Nevertheless, such models show that it can be optimal to restrict lending to consumers. Such models imply (see figure 5.1) that lending takes place in discrete jumps; that is, there are a countable number of ordered points $(0, b_1, b_2, \ldots)$ between which lending takes place at a constant marginal rate of interest. At each of these points b_i, there is a jump in the marginal rate of interest charged (perhaps to infinity, in which case no lending occurs beyond b_i).

A more recent literature has attempted to explicitly model the decision to default by consumers. This literature aimed to explain the limited ability of consumers to pool risk, and includes papers by Kehoe and Levine (1993) and Kocherlakota (1996). In these papers, the standard model of an infinitely lived, utility maximizing consumer subject to a life-cycle budget constraint is augmented by an additional constraint on the consumer's behaviour. This additional constraint explicitly accounts for the fact that *ex post* the consumer may wish to default on his debt, and suffer any penalty which default causes. The punishment for

[1] See, for instance Carroll and Summers (1991)

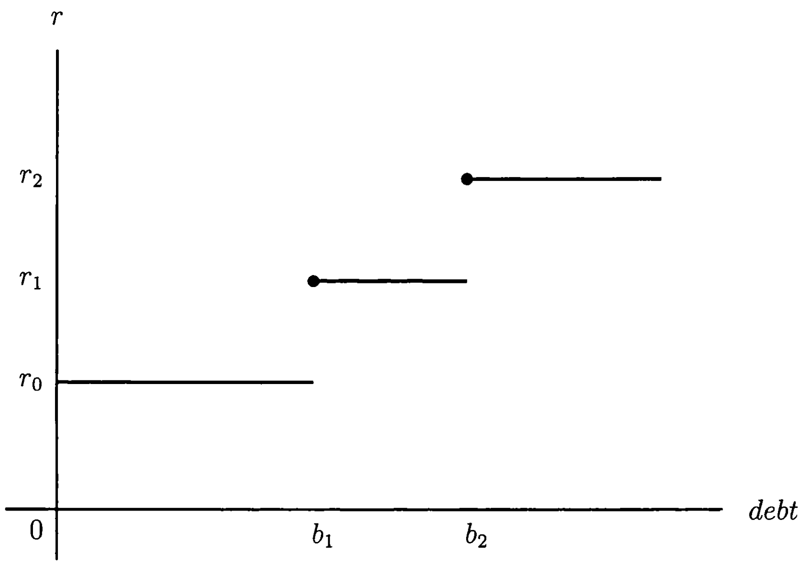

Figure 5.1: The marginal rate of interest as debt increases.

default could take many forms but these papers concentrate on default resulting in autarky, in which the consumer is permanently excluded from the credit markets after default. By solving these models for a decentralised market economy these models can endogenously create credit rationing in which the ability to borrow is restricted to some maximum level which depends on the parameters of the model. Above this maximum level default is assured, and hence it is never rational for lenders to allow borrowing beyond this limit. Such models differ from the earlier literature in that information is perfect, and in that there is only one interest rate at which lending occurs (if it occurs at all).[2]

Several papers have estimated the proportion of households that are credit constrained. The simplest approach is taken by Hall and Mishkin (1982), who attributed the rejection of the PIH model in Euler equations to a fixed fraction of the population simply spending their current income: this fraction was estimated to be some 20% of the population. Hajivassiliou and Ioannides (1996) formalize how the Euler equation is affected by credit constraints, and show how these constraints can motivate switching regression econometric models. See also Hajivassiliou and Ioannides (2001), Hajivassiliou and Ioannides (2002a) and Hajivassiliou and Ioannides (2002b). Mariger (1987) instead tried to estimate the effective time horizon in the Euler equation and concluded that around 19% of households were constrained. One problem is that the rejection of the PIH in Euler equations could instead be due to mis-specification of the Euler equation, a point that is well known in the literature.

Some papers have instead tried to more directly estimate or test credit constraints. The problem is that, denoting $\pi_i = 1$ as a binary variable indicating that the agent is credit constrained, this variable π_i is not directly observed. In much of the literature, instead some proxy variable has been substituted for the unobserved latent variable. For instance, in a clas-

[2]Two papers that test the effect of punishing default are Gropp, Schulz, and White (1997) and Charles (2000). They exploited differences between state bankruptcy rules in the different US states. Charles (2000) showed that moving to a state with a stricter punishment for default results in debtors borrowing $400 more on average.

sic paper, Zeldes (1989) splits households by their level of assets: low asset households, with a gross assets to monthly income ratio of less than 2, are assumed to be credit constrained. Jappelli (1990) instead uses self-reported responses to a question about credit constraints contained in the Survey of Consumer Finances. The question asked if the agent had been rejected for a loan, or if he had failed to apply for a loan because he feared rejection.[3] In either case, having chosen the proxy variable z, the observations can be partitioned, and those who are thought to be credit constrained can be compared to those who are not. If the proxy variable z exists, then there is no need to estimate the incidence of credit constraints, however, different groups can still be compared. Jappelli (1990) found that about 12% of households are credit constrained, rising to 19% if discouraged borrowers are included. He also found that credit constraints are more often binding for low income, low asset, young, and black households.

However, it is often not clear what should be used as an appropriate proxy for credit constraints, a point made by Garcia, Lusardi, and Ng (1997). They use a switching regressions technique in the Euler equation, and note that agents should react differently to increases and decreases in income if they are at the margin of being constrained. Their technique allows for constraints to be a function of several variables, and they find that around 16% of agents are constrained.[4] Gross and Souleles (2002) look at credit card balances and limits, and note that consumers increased their borrowing in response to any raising of the credit limit: they interpret this as due to credit constraints and suggest the effect of credit constraints on consumer behaviour is substantial.

Early estimates of the extra amount that households wanted to borrow, such as Hayashi (1985) and Mariger (1986) estimated that credit constraints had little effect on debts holdings

[3] Jappelli, Pishke and Souleles (1998) combines the Euler equation approach taken by Zeldes and Jappelli's more reliable measure of being constrained.

[4] As in this chapter, the probability of being credit constrained must be bounded away from zero for the regression to be identified.

and consumption, by using an Euler equation approach. Cox and Jappelli (1993) look at a cross section of households and compare a group who are assumed to be credit-constrained, with a group who are not. They find that credit constrained households would like to hold over $8,000 more debt than they actually do. Two other papers worth mentioning are Perrudin and Sorensen (2000) and Duca and Rosenthal (1994). The first considers a two stage estimation of asset holding using the 1983 wave of the SCF, whereby a probit predicts which asset types are held, while the second stage predicts how much of each asset is held. They find that age, marital status, education, and sex all have substantial effects on both the type and quantity of assets. In contrast, Duca and Rosenthal (1994) look at how liquidity constraints affect the ability of households to enter the mortgage market again using Jappelli's 'turn-down' measure. Their model allows for selectivity by using a bivariate probit model for the housing choice and whether a household is credit constrained, finding borrowing constraints particularly affect younger households. They ask whether households borrow, but not how much.

One problem is that even if it were known which households were credit constrained, estimates of the demand, or other behavioural equations can still be biased. If we estimate the demand equation by selecting only those observations who are not credit constrained (for which $\pi_i = 0$), then this is likely to under-estimate the demand fro debt among credit-constrained households: those households with an unusually low level of demand, in the sense that they have low error draws in the demand equation, are less likely to be observed to be credit constrained. This selection problem must be accounted for if we want to recover a true estimate of how much more credit-constraints want to borrow than they are currently allowed. This problem is likely to be apparent in Cox and Jappelli (1993) in their estimate of how much more households wish to borrow, but is controlled for by Duca and Rosenthal (1994).

The next section is devoted to showing how these problems can be overcome. It directly models what credit constraints look like in an econometric model, and shows how such a model may be identified. In particular, credit constraints are not replaced by some proxy variable but instead they are modeled as arising from some equation $\pi_i = f(z, \varepsilon)$ where z is continuous and multi-dimensional. Furthermore, it is implicitly recognized that π_i is observed with error.

5.3 An Empirical Framework

The theory suggests (recall figure 5.1) that consumers can borrow any amount up until some limit. Both the amount that the consumer wishes to borrow, and the credit limit are functions of the household's characteristics. Denoting desired borrowing as "demand" and the credit limit as "supply" we can write:

$$\begin{aligned} y_{1i} &= f_1(X_{1i}, \varepsilon_{1i}) \qquad && \textit{demand} \\ y_{2i} &= f_2(X_{2i}, \varepsilon_{2i}) \qquad && \textit{supply} \end{aligned} \tag{5.1}$$

and the level of debt that is observed, y_i, is defined as $min(y_{1i}, y_{2i})$. These equations explicitly recognize that the econometrician does not observe all the characteristics that drive demand or supply. Further, one would expect any characteristic that is observed both by the lender and the borrower (but not the econometrician) to enter both equations, and thus, in general, one would expect the errors ε_{1i} and ε_{2i} to be correlated. The agent will be credit constrained if both demand exceeds supply and demand is positive. This formulation makes explicit that credit constraints can bind at some level other than zero, and that not all agents who do not borrow fail to borrow because they are credit constrained. An alternative and parsimonious way of representing the same result is to write

$$y_i = f_{1i} + \pi_i (f_{2i} - f_{1i}) + \varepsilon \tag{5.2}$$

This formulation will sometimes be useful later.

How estimation proceeds depends on what exactly is observed. The approach taken by Cox and Jappelli (1993) assumes that π_i is observed, or can be well approximated by some proxy variable z_i which partitions the data into those for whom supply is binding and those for whom it is not. From those who are credit constrained (where $\pi_i = 1$) the supply equation can be recovered, while the demand equation can be recovered from those who are not. In this case the extra amount that credit-constrained consumers wanted to borrow is the difference between demand and supply conditional on credit constraints being binding. The strategy highlighted replaces demand conditional on credit constraints by the estimated demands from those for whom the constraints are not binding. However, estimation must account for the selection problem: agents are only observed to be credit constrained (ignoring the zero observations for the time being) if demand exceeds supply. That is, any estimation strategy must explicitly recognize that

$$E(\varepsilon_{1i}|\pi_i = 0) \neq E(\varepsilon_{1i}|\pi_i = 1) \tag{5.3}$$

and similarly for ε_{2i}. Failure to account for this selection problem results in biased estimates of f_{1i} and f_{2i} and thus mis-estimates how much more credit-constrained consumers wish to borrow.

More often π_i is not observed, and there is no good proxy variable to replace it. In these cases the problem is that only y_i is observed, but it is not clear which of the underlying equations y_{1i} or y_{2i} has generated the observation. However, estimation can proceed by noting that what is observed is the minimum of supply and demand. That is:

$$y_i = \begin{cases} min\left(y_{1i}, y_{2i}\right) & y_{1i} > 0, y_{2i} > 0 \\ 0 & otherwise \end{cases} \tag{5.4}$$

This involves replacing our equation $\pi_i = z_i$ by $\pi_i = f(X_{1i}, X_{2i}, \nu_i)$. Two comments are worth making. First, while previously z_i was binary and one dimensional, this formulation explicitly recognizes that many variables can affect the incidence of credit constraints, and that such variables may be continuous. Secondly, it also recognizes that there will be heterogeneity across agents that is not captured through those variables X_{1i} and X_{2i} observed by the econometrician. This fact is captured by the addition of the error term ν_i. The aim is to replace π_i by a probability distribution which depends on the household's observable characteristics, rather than by some proxy variable so equation 5.2 becomes[5]:

$$y_i = f_{1i} + Pr\left(\pi_i = 1 | X_{1i}, X_{2i}, \varepsilon\right)\left(f_{2i} - f_{1i}\right) + \varepsilon_i \tag{5.5}$$

In our framework, this probability depends on the estimated parameters in the underlying supply and demand equations.

Throughout this chapter, the household is only observed once. If repeated observations of the same household are available then a variety of different techniques are available, that can exploit the panel structure of the data. This is essentially the approach taken by Hajivassiliou (1987) and Hajivassiliou (1994). These papers apply a switching regression framework, and exploit information about how the level of debt is changing to identify supply and demand.[6] However, given the that the data in this chapter is cross-sectional (since only the first interview for each household is used), rather than longitudinal, the same approach can not be taken here.

5.3.1 Estimation by Maximum Likelihood:

The framework discussed above, in which π is not observed, is very similar in form to standard canonical disequilibrium models, as discussed in Quandt (1988), and the discussion

[5]Again to aid clarity, the problem of the zero observations is temporarily ignored. This will be important in estimation, and is considered then.

[6]This chapter more closely follows Eaton and Gersovitz (1981), although that paper did not account for the fact that observations were repeated, nor did it allow for correlation between supply and demand. Furthermore, in that paper, all the agents (in their case countries) held at least some debt.

is similar to that contained there. Estimation can proceed by full maximum likelihood.[7] The likelihood of any observation y_i is thus

$$\begin{aligned}\mathcal{L}_i = \; & Pr\left(y_i|y_{1i}, y_{2i} \geq 0; y_{1i} < y_{2i}\right) Pr\left(y_{1i} < y_{2i}|y_{1i}, y_{2i} \geq 0\right) Pr\left(y_{1i}, y_{2i} \geq 0\right) \\ & + Pr\left(y_i|y_{1i}, y_{2i} \geq 0; y_{1i} \geq y_{2i}\right) Pr\left(y_{1i} \geq y_{2i}|y_{1i}, y_{2i} \geq 0\right) Pr\left(y_{1i}, y_{2i} \geq 0\right) \\ & + \left[1 - Pr\left(y_{1i}, y_{2i} \geq 0\right)\right]\end{aligned} \tag{5.6}$$

where it is implicitly recognized that we are also conditioning on X_{1i} and X_{2i}. If we impose that f_1 and f_2 are linear, and further assume that the error structure is bivariate normal with co-variance matrix Σ, so that

$$\begin{pmatrix} \varepsilon_{1i} \\ \varepsilon_{2i} \end{pmatrix} \sim N \left[0, \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix} \right] \tag{5.7}$$

and note that y_i is observed, then it is simple to construct the (log)-likelihood function for this problem. Although the exact form of the likelihood is given in the appendix, some technical remarks need to be made. Consistency of the estimator follows from the consistency of both the tobit model and of the canonical form of the disequilibrium model: for the later Hartley and Mallela (1977) highlighted a number of conditions that are needed to ensure the consistency of the estimator. The most important of which are (i) that the parameters to be estimated are in the interior of the parameter space; and (ii) there exist exclusion restrictions on the supply and demand equations. A third condition is that the fraction of observations falling within each regime approaches a strictly positive fraction as the number of observations approaches infinity. These conditions translate directly to the estimator above. The first of these conditions means that the estimated variances σ_1 and σ_2 must be bounded away from zero, and that the estimated correlation parameter ρ must be bounded away from ± 1. This is more onerous than might be thought: unlike in the straightforward

[7] Several other estimation strategies could instead be applied, such as SML, GMM, and MD or methods based on moment generating functions. This chapter does not claim any particular advantage in the estimation strategy chosen.

tobit model, it is not known, a priori, whether a non-zero observation obtains from the supply or the demand equation. If the regression includes a constant, then β_j can always be chosen so as to make $y_j = 0$ and then the likelihood is unbounded as $\sigma_j \longrightarrow 0$. In practise this problem translates into finding a suitable starting value that does not lie in the region in which the gradient points to the boundary of the parameter space.[8] The second condition means that there must be variables that enter the supply, and not the demand equation, and variables that enter the demand, and not the supply equation (e.g. $X_{1i} \neq X_{2i}$). This is the identification problem. Note however, that given such variables, other variables can freely enter both the supply and the demand equation, and the estimated effect on supply and/or demand of the variable can be separately identified.

Which exclusion restrictions should be made? The exclusions in this chapter are that quarter enters demand and not supply, while bank regulation enters supply an not demand. Are these restrictions reasonable? Using seasonal dummies only in the demand equation argues that lenders do not discriminate on the basis of which month borrowers ask for loans. Note that this does not mean that debtors observed, or unobserved characteristics do not change month by month, merely that banks do not use month in their assessment of whether to extend a loan to the potential debtor.[9] The question is whether this is a movement along the supply curve (demand only shifts), or the supply curve is shifting as well. If banks did discriminate on the basis of month then there would be incentives for borrowers to time their requests for debt at certain times in the year, and since it is as costly to request a loan in one month as another, it seems sensible to suppose there is a pooling equilibrium on

[8]This was done by using the estimated parameters from a tobit model as the initial starting values. Kooiman, van Dijk and Thurik (1985) suggest instead employing some penalty function for getting too close to the boundary, although it is not clear why this should cause convergence to the interior maximum: the solution which has a clear economic interpretation.

[9]Gross and Souleles (2002) note that credit limits on credit cards are more likely to rise around Christmas, and that lending on credit cards is much higher at this time of year.

month. A further argument, supported by Miron (1986), and Barsky and Miron (1989), is that the federal reserve explicitly tries to smooth out seasonal variations in supply through its market operations. On the other side, the banking regulations are due to state banking laws.[10] These rules, tabulated in 5.1 not only vary across states, but also across time, and it seems reasonable to suppose that they are unrelated to demand. The table shows when different states allowed intra-state branching, and highlights that there has been a gradual relaxation in banking regulations over the last 20 years. The timing of this deregulation differed from state to state, which can be exploited in the regressions that will be run.

5.4 Data Description:

The Consumer Expenditure Survey (CEX) has already been described Other surveys also ask questions on the level of household debt, but there are a number of advantages in using the CEX survey. The first is the large sample size: over 7000 interviews are conducted each quarter, and this chapter is able to exploit a sample size of 13918 observations. This is larger than most other data sets that are available: something that is important given the estimation strategy. A second advantage is that the CEX survey contains information on the state of residence of the household. This is crucial since the supply side instrument is the state level banking regulations. Few other surveys supply this kind of information. For instance, the Survey of Consumer Finances is perhaps a more natural data source for information on household debt but it only provides state information for 1983. This means that changes over time can not be captured. Furthermore, since the SCF is conducted over a month or two, month can not act as a demand instrument in that survey. The CEX survey,

[10]The rules used are whether intrastate branching through merger and acquisition is allowed (bank-law 1) and whether full intrastate branching was permitted (bank-law 2). Including whether interstate banking was also permitted made no difference to the regressions. These rules were obtained from Amel (1993) and Kroszner and Strahan (1999) where a much fuller discussion of these regulations is contained.

by contrast, surveys continuously throughout the year, and provides state information for all years that are surveyed (although, for confidentiality reasons, this information is suppressed for some states in some years). For the regressions the debt was deflated by a household specific Stone-Geary price index (where the prices relate to non-durable expenditure). The regressions will also include an interest rate, which will refer to the reported municipal bond rate (which has the advantage of being tax free).

The data used in this chapter are those observations from 1988 to 1993 for which full state information is available who were 'full income responders'. This differs from the other chapters, in that fewer years are being used. This chapter only uses data up to 1993 since the branch banking regulations exploited in this chapter had been almost completely repealed by this date, and there is no variation that can be exploited after 1993. To ensure a reasonably homogeneously defined group, sampling was restricted to single households or those headed by a couple, and those households whose head was between 25 and 55. Large households (with 7 or more members) were excluded, as well as households whose head received no education. Also excluded were self-employed households, those whose primary occupation as farming, and those households in which more adults other than the household head (and his or her partner) were working.

Table 5.2 summarises some of the features of the raw data, without conditioning on any observable characteristics of the households. It shows that the median debt holding in the whole sample is \$790 while 68.2% of households hold at least some debt. Conditional on holding debt, the average amount of debt held is \$4,008, a substantial amount. While this may seem large, other studies, such as Cox and Jappelli (1993) have found similarly large amounts. The table also compares the level of debt for each year. Except for the first year, the proportion holding any debt gradually declined over the period in question. However, the average size of the debt was increasing, and there were particularly dramatic increases in this quantity in 1990 and 1991. Comparing age groups shows that younger households are more likely to hold debt than older households, and that median debt holdings are roughly twice as large. However, on average, the level of debt was not much different for the 15-35 age group and for the 45-55 age group. The 35-45 age group hold around $300 more. The table also highlights some other features of the data. Childless households are less likely to hold debt, hold much less debt when they do by roughly $500. While there is little difference between households with either one or two children, having three or more children results in around $100 less being held, and 3% fewer households hold debt. The differences between education groups is dramatic; the most poorly educated group is 20% less likely to hold any debt, although when they hold debt, their holdings are only slightly less than the middle two education groups. The most educated group hold much greater amounts of debt, over $2,000 more than any other group. Other comparisons show that unmarried households, those headed by women, and ethnic minority households are all less likely to hold debts, and hold smaller debts when they do.

5.5 Results:

Table 5.3 displays the estimated supply and demand equations that were estimated. Omitted from the table are the coefficients on the year dummies and a set of regional dummies. Results are recorded for both the levels (in columns 2 and 3) and for the log-levels (in columns 4 and 5). For the log-level regression the left-hand side variable is the $\ln(1 + debt)$. The raw results are themselves not very easy to understand, and hence they will be discussed when a number of issues are discussed below. Recall that the introduction outlined a number of questions that might be interesting. The first question was to test whether $\pi_i \equiv 0$, and as discussed in the appendix, is not investigated in this chapter. The other questions can be addressed by constructing:

(ii) $E(\pi_i)$

(iii) $E(\pi_i|x_i)$

(iv) $E(y_{1i} - y_{2i}|\pi_i = 1)$

The $E(\pi_i)$ will give (assuming our observations are a random draw from the whole population) the fraction of people currently credit-constrained in the whole economy. We could also (but did not) construct the probability of being credit-constrained conditional on the current level of debt. By ranging over all possible y_i we can calculate the 'maximum' level of debt, y^* that a consumer is allowed to hold. Instead (iii) gives the probability of being credit-constrained conditional on the x-variates. By ranging over the x-variates different subgroups can be compared to see if there are significant observable differences across subgroups in their ability to borrow and smooth consumption. Lastly (iv) shows how much more such consumers would have borrowed in the absence of binding credit-constraints.

(a) *The proportion of households that are credit constrained.*

The proportion of households that are credit constrained is the unconditional expectation of π_i over all households. For a household to be credit constrained it is necessary for the demand for debt to exceed supply *and* for the demand for debt to be positive. These two conditions will not be independent, even in the case where the supply and demand equations have uncorrelated errors. This condition can be written as

$$\begin{aligned} E(\pi_i) &= Pr(y_{1i} \geq 0; y_{1i} \geq y_{2i}) \\ &= Pr(y_{1i} \geq y_{2i}; y_{1i}, y_{2i} \geq 0) + Pr(y_{1i} \geq 0; y_{2i} \leq 0) \\ &= Pr(y_{1i} \geq y_{2i}|y_{1i}, y_{2i} \geq 0)\ Pr(y_{1i}, y_{2i} > 0) \\ &\qquad + [1 - Pr(y_{1i}, y_{2i} > 0) - Pr(y_{1i} < 0)] \end{aligned} \tag{5.8}$$

Not all households that have no debt wish to hold debt, and at least some of the households who hold debt would like to hold more. This probability can easily be constructed and more details are given in the appendix. The variance can also be calculated, since the estimated $\hat{\pi}$ is distributed:

$$\sqrt{n}\left[\pi\left(\hat{\theta}\right) - \pi\left(\theta\right)\right] \sim N\left(0, \pi\prime\left(\hat{\theta}\right)\Sigma\,\pi\prime\left(\hat{\theta}\right)^{t}\right)$$

where $\theta^t = [\beta_1^t\ \beta_2^t\ \sigma_1\ \sigma_2\ \rho]$, Σ is the variance-covariance matrix of the parameters from the maximum likelihood function, while $\pi\prime$ is a vector of partial derivatives with respect to θ evaluated at estimated coefficients $\hat{\theta}$.

The results in table 5.4 suggest that the proportion of consumers who are credit-constrained is 26% when the levels equation is estimated, and 31% when the equation system is estimated in log-levels. These figures is slightly larger than is usually estimated For instance, Hall and Mishkin (1982) estimate around 20% of households are credit constrained, a figure that is close to the 19.4% estimated by Mariger (1987). Jappelli (1990) also estimates a figure of around 19%.[11] One explanation of why this study finds a higher figure than the self-reported responses contained in Jappelli (1990) is that this chapter will also include those households who were allowed to hold some debt, but not as much as they wish to hold. His definition only included those households who were rejected outright.

This number was constructed under the assumption that the errors u_{1i} and u_{2i} in equation 5.1 are uncorrelated ($\rho = 0$). This assumption seems implausible. One interpretation of the error term is that it is due to parameters being omitted from the regression, either because they are unmodelled, or because they are not observed. For the parameters in the regression to be identified it must be true that $E\left(u_{ji}|x_{ji}\right) = 0$ for $j = 1, 2$, the standard assumption in a regression. However, if $\rho = 0$ then this argues that any such omitted variable can only enter either the supply equation, or the demand equation, but not both. This seems unlikely. If instead an omitted variable increased (or decreased) both supply and demand, then this would cause a positive correlation between the errors, in which case $\rho > 0$, while

[11]The different results may partly reflect the later period that is studied here.

if the omitted variable entered the supply and demand equations with opposite signs, then $\rho < 0$. Unfortunately I have not been able to obtain results for the case of correlated errors because of convergence problems.

(b) *Differences across consumers.*

As highlighted in the introduction, one of the advantages of the approach taken in this chapter is that differences across consumer types can be sensibly investigated. The first question that can be asked is are consumers with different observable characteristics differently credit constrained? However, the approach can go further than that: it can also investigate the *reason* for these differences across consumers. For instance, are these differences being driven by differences in the supply of loans, or by differences in the demand for loans?

Table 5.4 reports the average level of demand and the average level of supply conditional on the X-variates. The median demand (and supply) as are also calculated. For instance, when education was considered, the demand (or supply) were calculated as:

$$\bar{y}_1^j = \frac{1}{n}\sum_i E\left(y_{1i}^j | X_i^j\right)$$

where $\tilde{X}_{1i}^j$ equals X_{1i} except that the actual level of education is replaced by an arbitrary level of education.

In all cases the levels regression predicts a lower incidence of credit-constraints than the log-level regression. Perhaps more interesting in this table are the differences across agent types. The results show households headed by women are marginally more likely to be constrained than men. Table 5.5 highlights that this is because while the median demand for debt by men was lower by about \$470 for the level equation and about \$350 for the log-level equation, the supply of credit was higher for men, by around \$50 for the level equation, and by around \$140 for the log-level equation. Women are more likely to be credit constrained, by around 1.5% according to table 5.4 since they demand is higher but supply is lower. Much more important is being married. This makes the household nearly 5% less likely to be credit constrained for the level equaiton. For the log-level equation the difference is just over 1%. Table 5.5 highlights that this is mainly because for non-married households the demand for debt is very low, hence even though supply is also lower for unmarried households, they are less likely to be credit-constrained. For age table 5.4 shows that the incidence of credit constraints declines as the household ages. Again the effect is stronger in the levels equation, where there is a decline of nearly 4% between the ages of 25 and 55, compared to the log-level equation where the decline is a more modest 2.5%. As households age, the supply of credit falls, as shown in table aveds, but the demand for debt falls much more sharply.

As might be expected, poorly educated households, those who didn't complete high school, are much less likely to be credit-constrained, and there is still a small decrease for those households who started, or finished college. The main reason is that the demand for debt is higher for better educated households. A similar explanation holds for the much lower probability of being credit constrained among black (that is non-white) households. They are nearly 6% less likely to be credit constrained (3.7% in the log-level equation) than white households, even though the median supply of credit is similar for black and for white households.[12] Other features are that childless households are slightly more likely to be credit-constrained, although this is unlikely to be because of their increased mobility (which would suggest a supply side effect), but rather is due to their higher demand for credit.[13]

[12]Jappelli (1990) found blacks are more likely to be credit constrained, which contrasts with the results in this study. This is due to households who defined themselves as discouraged in his survey: he interpreted these households as credit-constrained but this study suggests that these households might have been given a loan if only they had asked for one; note the extremely low level of demand which is found in table 5.5.

[13]The framework in this paper assumes that application for loans are costless. The interpretation of the results is not so simple if applications are costly, for in that case some applicants may be discouraged from applying for a loan if they believe that there is a large chance of the application being turned down. This is likely to be more serious for those agents who ask for relatively smaller loans, such as black, female, or poorly educated households: see Hajivassiliou and Ionnides (2002c) for an attempt to address this issue.

(c) *The demand for debt among credit-constrained households*

The previous discussion investigated the extent of credit rationing among US households. However, a complete discussion will also consider how important rationing is for these consumers. In the absence of rationing how much more would households borrow? If the *extra* amount of debt that households wish to borrow is denoted Δ then the problem is to construct some estimate of[14]:

$$E(\Delta_i) = E(\Delta|\pi_i = 1)\, Pr(\pi_i = 1) + E(\Delta|\pi_i = 0)\, Pr(\pi_i = 0) \tag{5.9}$$

But by construction, the household can borrow as much as it likes if it is not credit constrained hence $\Delta_i = 0$ whenever $\pi_i = 0$, hence this term drops out. Thus

$$E(\Delta_i) = Pr(\pi_i = 1)\left[E(y_{1i}|\pi_i = 1) - y_i\right] \tag{5.10}$$

For the last term $y_{2i} = y_i$ whenever $\pi_i = 1$, and this is observed. The problem is to find some good proxy or estimate of the first. One approach, prevalent in the literature (see for instance Cox and Jappelli, 1993), is to replace $E(y_{1i}|\pi_1 = 1)$ by $E(y_{1i}|\pi_1 = 0)$. However, even if the estimate of $E(y_{1i}|\pi_1 = 0)$ is consistently estimated, using this in equation 5.9 will result in downward biased estimates of Δ. This is because (ignoring the zero observations).

$$\begin{aligned} E(y_{1i}|\pi_1 = 1) &= E(y_{1i}|y_{1i} > y_{2i}) \\ &> E(y_{1i}|y_{2i} > y_{1i}) \end{aligned}$$

A naive estimation strategy would thus under-estimate the true impact of credit-constraints on households. Instead construction of the difference entails allowing for the selectivity problem that this highlights. Given that f_1 and f_2 have both been estimated, then the expectation can be constructed:

$$E(\Delta_i) = \int_{y_i}^{\infty} y_{1i} Pr(y_{1i}|\pi_i = 1)\, dy_{1i}$$

[14]It should be understood that we are conditioning on all the observed characteristics of the agent, and on the observed level of borrowing, although the notation has suppressed this.

where

$$Pr\left(y_{1i}|\pi_i = 1\right) = Pr\left(y_{1i}|y_{1i} > y_{2i} > 0\right) + Pr\left(y_{1i}|y_{1i} > 0; y_{2i} < 0\right)$$

and in the case of uncorrelated errors this becomes:

$$Pr\left(y_{1i}|\pi_i = 1\right) = \frac{\phi\left(\frac{y_{1i} - X_{1i}\beta_1}{\sigma_1}\right)}{\Phi\left(\frac{y_{1i} - X_{1i}\beta_1}{\sigma_1}\right)}$$

which can be recovered. Construction of this results in an estimate of $E(\Delta)$ of $1,642 dollars for the level of debt, and $3,975 for the log-level regression. That is, in the absence of credit constraints, households would, on average, borrow over one and a half thousand more unsecured dollars than they currently do (or nearly $4,000 for the log-level equation). This number is large given that the average amount of debt households hold is $2,733. The results also show that this problem is particularly acute for highly educated households, for younger households, and for married households.

5.6 Conclusion:

The chapter described how it was possible to estimate a model in which some consumers were credit-constrained, even though it was not known which consumers were constrained. Estimation also addressed the selectivity problem: demand (or supply) conditional on being constrained does not equal the unconditional demand. Using the strategy proposed in the main body of the chapter, the demand for debt and the supply of debt were separately identified and estimated. From this we recovered the incidence of credit-constraints among households, estimated to be between about 26% and about 31% in the population as a whole. We found that marital status, education, and race dramatically changed the incidence of credit-constraints, as well as age. The main reason for the higher incidence of credit-constraints was that marginal groups were more reluctant to borrow.

Overall, the chapter shows how three of the questions that were raised in the introduction can be answered and, in the appendix, it explains why it is not practically possible to answer the first of the four questions. The results are perhaps a little surprising, and show that credit constraints may be more important than is commonly thought. Furthermore, the importance of credit constraints declined only moderately with age: it affects older households (those around 50) not much less than younger households.

Appendix

The likelihood function

$$ln\mathcal{L} = \sum_i d_i \; ln\left(h_{1i}h_{2i} + h_{3i}h_{4i}\right) + (1 - d_i)\, ln\left(h_{5i}\right)$$

where d_i is an indicator function for observing a non-zero debt holding, that is:

$$d_i = I\left(y_i > 0\right)$$

and

$$h_{1i} = \left(2\pi\sigma_1^2\right)^{-\frac{1}{2}} \exp\left[-\frac{\left(y_i - \beta_1' x_{1i}\right)^2}{2\sigma_1^2}\right]$$

$$h_{2i} = \frac{1}{2\pi\left(1-\rho^2\right)^{\frac{1}{2}}} \int_0^\infty \int_{y_1}^\infty exp\left[\frac{-1}{2\left(1-\rho^2\right)}\left(\xi_1^2 - 2\rho\xi_1\xi_2 + \xi_2^2\right)\right] d\xi_2 d\xi_1$$

$$h_{3i} = \left(2\pi\sigma_2^2\right)^{-\frac{1}{2}} \exp\left[-\frac{\left(y_i - \beta_2' x_{2i}\right)^2}{2\sigma_2^2}\right]$$

$$h_{4i} = \frac{1}{2\pi\left(1-\rho^2\right)^{\frac{1}{2}}} \int_0^\infty \int_{y_2}^\infty exp\left[\frac{-1}{2\left(1-\rho^2\right)}\left(\xi_1^2 - 2\rho\xi_1\xi_2 + \xi_2^2\right)\right] d\xi_1 d\xi_2$$

$$h_{5i} = 1 - \frac{1}{2\pi\left(1-\rho^2\right)^{\frac{1}{2}}} \int_{k_1}^\infty \int_{k_2}^\infty exp\left[\frac{-1}{2\left(1-\rho^2\right)}\left(\xi_1^2 - 2\rho\xi_1\xi_2 + \xi_2^2\right)\right] d\xi_2 d\xi_1$$

and where $k_1 = \frac{\beta_1' x_{1i}}{\sigma_1}$, $k_2 = \frac{\beta_2' x_{2i}}{\sigma_2}$, and $\Phi\left(\cdot\right)$ is the usual c.d.f. of the normal distribution and the correlation of the errors is defined as $\rho = \frac{\sigma_{12}}{\sigma_1\sigma_2}$.

The intuition is straightforward: non-zero observations may either have been generated by the demand equation or by the supply equation. If generated by the demand equation, then h_{1i} represents the p.d.f. of the observation, given that it is generated by the demand equation, while h_{2i} represents the probability that demand is exceeded by supply. The converse is represented by h_{3i} and h_{4i}. Finally, the term h_{5i} is the standard bivariate normal c.d.f. and represents the probability that either supply or demand is less than or equal to zero.[15] If the correlation between the errors is identically equal to zero (in which case $\rho \equiv 0$) then h_{2i} and h_{4i} simplify in the obvious way, while h_{5i} becomes:

$$h_{5i} = 1 - \Phi\left(\frac{\beta_1 x_{1i}}{\sigma_1}\right)\Phi\left(\frac{\beta_2 x_{2i}}{\sigma_2}\right)$$

Results will be presented for this simpler framework as well.

Testing for credit constraints

It would be useful to test the system of equations in our model against an appropriate market clearing model. In a market clearing model, the supply and demand equations jointly determine the level of borrowing, and there are no parameter restrictions implied by the model. Thus conventional tests are not appropriate, something that has been known at least since Hwang (1980).[16] To illustrate the argument, suppose, for the time being, the zero observations were ignored. As specified in equation 5.1 the model to be estimated takes the form[17]:

$$\begin{aligned} y_{1i} &= f_1(x_{1i}, \beta_1) + u_{1i} \qquad && i \in \Omega_1 \\ y_{2i} &= f_2(x_{2i}, \beta_1) + u_{2i} \qquad && i \in \Omega_2 \equiv \Omega \backslash \Omega_1 \end{aligned}$$

where Ω_1 and Ω_2 partition observations between regimes and $\Omega_i \in \Theta$ where Θ is the space of all possible partitions. That is, Ω_2 represents those agents who are credit constrained, while Θ represents *all* the possible ways of selecting credit constrained people. By making appropriate assumptions about the parameters this system of equations can be estimated in

[15] See Johnson and Kotz (1972) for a general discussion of the derivation of these equations.

[16] He suggested a cusum or cusum of squares test on the residuals from the market clearing model but noted the poor power of the test in large samples.

[17] For the purposes of this discussion we will ignore the observations in which no borrowing occurs.

a variety of ways, including FIML, SML, GMM, and MD. In the main part of the chapter, the equations have been linearized, and then estimated by maximum likelihood. Estimation comprises both estimating $\hat{f}_1$ and $\hat{f}_2$, and estimating $\hat{\Omega}_i$. Having derived some estimate of f_1 and f_2 the challenge is to test the estimated model against some alternative model in which nobody is ever credit constrained (or agents are always credit constrained). In such a model the system of equations reduces to only one equation. That is, all observations will fall into only one of the regimes, f_1 say. The null and alternative hypotheses can thus be written:

$$H_0 : \beta_i = \beta \qquad \forall i$$

$$H_1 : \beta_i = \begin{cases} \beta_1(\pi) & i \in \Omega_1 \\ \beta_2(\pi) & i \in \Omega_2 \end{cases}$$

Let π_i denote an indicator function for observation i belonging to the first regime[18]: $\pi_i = I[y_{2i} - y_{1i} \geq 0]$. (And let π be the stacked vector of π_i's.) In which case the system of equations can be re-written in the following way.

$$y_i = f_{1i} + \pi_i (f_{2i} - f_{1i}) + u_{1i} + \pi_i (u_{2i} - u_{1i}) \tag{5.11}$$

From this it is immediately apparent that there is a problem with testing the parameters of the model. Suppose that the aim was to test whether all observations were generated by the first equation (*i.e.* credit constraints are never binding). Equivalently this can be interpreted as either π_i identically equals zero, or $f_{2i} - f_{1i}$ identically equals zero. The problem is that if $\Omega_2 = \phi$ (the empty set) then π_i and $f_{2i} - f_{1i}$ are not separately identifiable. If the null hypothesis had generated the data, then it becomes problematic to test against the alternative hypothesis.

If π where known, then testing the model would simply be a matter of constructing Wald, LR-, or LM- statistics and comparing the test statistic against a standard chi-squared

[18] This definition can be reconciled with the earlier definition of π if we note that now we are ignoring the zero observations.

distribution with degrees of freedom equal to the dimension of β. However, π is not known. Since it is derived in a way that is dependent on the data, the Wald (similarly LR-, and LM-) statistic based on this estimated π will no longer have a standard distribution: instead the null hypothesis will be *over-rejected.* This could lead to the mistaken conclusion that some agents suffer binding credit constraints.

Table 5.1: When US states deregulated and relaxed state banking regulations.

State	intrastate branching through M & A	full intrastate branching	State	intrastate branching through M & A	full intrastate branching
Alabama	1981	1990	Minnesota	1993	-
Alaska	1970	1970	Miss.	1986	1989
Arizona	1970	1970	Missouri	1990	1990
Arkansas	1994	-	Nebraska	1985	-
California	1970	1970	Nevada	1970	1970
Colorado	1991	-	New Hamp.	1987	1987
Connect.	1980	1988	New Jersey	1977	-
Delaware	1970	1970	New Mexico	1991	1991
D.C	1970	1970	New York	1976	1976
Florida	1988	1988	N. Carolina	1970	1970
Georgia	1983	-	Ohio	1979	1989
Hawaii	1986	1986	Oklahoma	1988	-
Idaho	1970	1970	Oregon	1985	1985
Illinois	1988	1993	Penn.	1982	1990
Indiana	1989	1991	S. Carolina	1970	1970
Iowa	-	-	S. Dakota	1970	1970
Kansas	1987	1990	Tenn.	1985	1990
Kentucky	1990	-	Texas	1988	1988
Louisiana	1988	1988	Utah	1981	1981
Maine	1975	1975	Vermount	1970	1970
Maryland	1970	1970	Virginia	1978	1987
Mass.	1984	1984	Washington	1985	1985
Michigan	1987	1988	W. Virginia	1987	1987
			Wisconsin	1990	1990

Taken from Amel (1993) and Kroszner and Strahan (1999), while 1970 means 1970 or before.

Table 5.2: Some summary statistics on debt-holding among US households.

	median debt ($)	ratio holding debt (%)	mean debt* ($)
All	790	68.2	4,008
1988	615	69.0	3,496
1989	800	70.3	3,476
1990	800	69.7	3,962
1991	874	68.1	4,353
1992	795	67.4	4,325
1993	700	65.2	4,465
Age 25-35	900	71.0	3,955
Age 35-45	800	68.0	4,243
Age 45-55	451	62.4	3,874
No Children	570	65.3	2,539
1 Child	1,100	72.1	3,075
2 Children	1,055	72.4	3,017
3+ Children	991	69.1	2,909
some school	0	47.3	3,386
finished high school	700	67.3	3,431
some college	1,000	72.4	3,650
university degree	955	71.5	4,885
unmarried	400	62.3	3,621
married	1,000	71.1	4,218
white	850	70.0	4,045
black	231	56.0	3,983
male	800	68.3	4,177
female	670	67.7	3,624

*Conditional on holding at least some debt.

Table 5.3: Estimated debt equations, in $1,000's (standard errors in parentheses).

parameter	level		log-level	
	demand	supply	demand	supply
constant	-7.557	1.714	-0.180	0.787
	(0.813)	(0.135)	(0.130)	(0.101)
female head	0.676	-0.053	0.142	-0.045
	(0.278)	(0.045)	(0.046)	(0.032)
married	2.557	0.285	0.255	0.249
	(0.319)	(0.052)	(0.051)	(0.037)
(age-40)/10	-0.315	-0.053	-0.049	-0.032
	(0.319)	(0.052)	(0.052)	(0.038)
$(\text{age-40})^2/100$	-0.111	-0.007	0.004	-0.018
	(0.176)	(0.029)	(0.028)	(0.021)
$(\text{age-40})^3/1000$	-0.149	-0.006	-0.020	0.004
	(0.209)	(0.034)	(0.033)	(0.024)
Completed High school	4.089	0.016	0.496	0.054
	(0.403)	(0.071)	(0.065)	(0.055)
Some college	4.742	0.084	0.619	0.099
	(0.429)	(0.072)	(0.068)	(0.056)
College degree	5.004	0.053	0.539	0.142
	(0.426)	(0.073)	(0.067)	(0.058)
non-white	-2.584	0.037	-0.338	0.005
	(0.318)	(0.056)	(0.050)	(0.044)
one child	0.165	0.061	0.036	0.024
	(0.366)	(0.056)	(0.058)	(0.043)
two children	0.282	0.015	0.067	-0.039
	(0.352)	(0.055)	(0.057)	(0.041)
3+ children	-0.091	0.080	0.030	0.014
	(0.448)	(0.069)	(0.069)	(0.051)
interest rate	0.117	0.007	0.011	0.011
	(0.118)	(0.018)	(0.019)	(0.013)
q2	-0.237	-	-0.057	-
	(0.295)		(0.042)	
q3	-0.359	-	-0.073	-
	(0.297)		(0.042)	
q4	0.321	-	0.030	-
	(0.292)		(0.043)	
bank-law 1	-	-0.034	-	-0.006
		(0.063)		(0.040)
bank-law 2	-	-0.042	-	-0.040
		(0.047)		(0.031)

Table 5.4: The Proportion of credit constrained households.

Subgroup	level		log-level	
	constrained (%)	*s.e.*	constrained (%)	*s.e.*
All	26.15	0.32	31.61	0.30
Men	25.73	0.35	31.15	0.34
Women	27.40	0.60	32.85	0.58
Married	27.70	0.39	31.93	0.38
Unmarried	22.75	0.63	30.82	0.63
Age 25	27.41	0.92	32.87	0.77
Age 30	26.68	0.43	32.01	0.41
Age 35	26.31	0.45	31.54	0.44
Age 40	26.05	0.44	31.31	0.44
Age 45	25.67	0.56	31.15	0.52
Age 50	24.94	0.61	30.90	0.58
Age 55	23.60	1.31	30.35	1.19
Some High School	16.71	0.81	26.41	0.79
Completed High school	25.98	0.58	31.76	0.55
Some College	27.17	0.57	32.72	0.54
College degree	27.85	0.51	31.70	0.52
White	26.87	0.34	32.02	0.33
Black	20.95	0.71	28.34	0.69
1988	24.57	0.66	30.43	0.63
1989	25.95	0.65	32.30	0.62
1990	26.55	0.56	31.57	0.58
1991	27.29	0.64	31.99	0.61
1992	24.55	0.95	30.51	0.91
1993	24.47	0.96	30.32	0.93
No children	26.88	0.61	32.19	0.55
One child	25.43	0.68	31.20	0.60
Two children	26.01	0.71	31.53	0.64
Three-Four children	25.15	0.81	30.65	0.77
Bank Law 1 = no	26.07	0.35	31.58	0.34
Bank Law 1 = yes	26.17	0.32	31.62	0.33
Bank Law 2 = no	26.07	0.33	31.46	0.33
Bank Law 2 = yes	26.19	0.32	31.69	0.33

Table 5.5: Median demand and supply for different subgroups.

Group	level		log-level	
	demand	supply	demand	supply
All	1,557	2,065	1,410	2,164
Men	1,460	2,078	1,331	2,197
Women	1,934	2,025	1,681	2,056
Married	2,127	2,158	1,603	2,411
Unmarried	458	1,873	1,026	1,659
Age 25	2,015	2,137	1,725	2,156
Age 30	1,740	2,105	1,522	2,212
Age 35	1,595	2,079	1,410	2,217
Age 40	1,492	2,053	1,344	2,181
Age 45	1,355	2,023	1,284	2,117
Age 50	1,119	1,986	1,198	2,036
Age 55	758	1,936	1,058	1,949
Some High School	5	2,028	460	1,879
Completed High school	1,247	2,011	1,361	2,039
Some College	1,718	2,112	1,670	2,179
College degree	1,922	2,082	1,466	2,318
White	1,714	2,060	1,493	2,162
Black	303	2,098	791	2,178
1988	1,061	2,026	1,130	2,089
1989	1,413	1,980	1,462	1,942
1990	1,705	2,095	1,429	2,236
1991	1,906	2,052	1,525	2,222
1992	1,103	2,091	1,183	2,198
1993	1,090	2,101	1,145	2,200
No children	1,719	2,000	1,478	2,024
One child	1,417	2,167	1,383	2,346
Two children	1,499	2,053	1,386	2,154
Three-Four children	1,316	2,147	1,272	2,384
Bank Law 1 = no	1,557	2,094	1,410	2,180
Bank Law 1 = yes	1,557	2,060	1,410	2,161
Bank Law 2 = no	1,557	2,094	1,410	2,253
Bank Law 2 = yes	1,557	2,051	1,410	2,124

Table 5.6: How much extra do households wish to borrow?

Group	level	log-level
All	1,642	3,975
Men	1,603	3,742
Women	1,767	4,730
Married	1,807	4,416
Unmarried	1,304	3,032
Age 25	1,778	4,813
Age 30	1,698	4,248
Age 35	1,656	3,946
Age 40	1,628	3,773
Age 45	1,588	3,625
Age 50	1,514	3,409
Age 55	1,386	3,052
Some High School	838	1,519
Completed High school	1,599	3,804
Some College	1,729	4,598
College degree	1,797	4,024
White	1,706	4,160
Black	1,168	2,308
1988	1,484	3,217
1989	1,611	4,150
1990	1,684	3,994
1991	1,756	4,258
1992	1,487	3,342
1993	1,481	3,239
No children	1,706	4,185
One child	1,580	3,869
Two children	1,625	3,903
Three-Four children	1,550	3,557

Chapter 6

Concluding Remarks:

The main body of work, in chapters 3-5 addressed some current issues in consumption by exploring differences in the institutional framework in the different US states. Chapter 3 investigated the effect of state bankruptcy laws on debt holdings among US households. It restricted attention to households who were borrowing to smooth consumption and specifically excluded households who were investing in a farm, or non-farm, business. The results suggest that restrictions on fully enforcing debts by allowing households to keep at least some assets when they default reduces the amount of debt that households are allowed to hold. This might suggests that optimality would imply arbitrarily large punishments for default, something that is not in fact observed. An important argument against such large punishments is that it might reduce the insurance element of default: that is defaulting when households realise a negative and permanent shock to income allows the idiosyncratic component of this risk to be shared among all households. Chapter 3 also found weak evidence that bankruptcy did allow agents to share risk in this way.

This idea of risk-sharing was pursued in chapter 4. If agents have both temporary and permanent shocks to income, then the full risk-sharing suggests that the idiosyncratic component of the income shock can be fully insured, and hence will not enter into *changes* in consumption. However, agents can not fully insure against all their idiosyncratic risk. Chapter 4 showed how the proportion of the risk that can be insured, say ϕ can be estimated and related to some policy instrument, such as the degree to which taxes are re-distributive. The regressions showed that making taxes more re-distributive unambiguously reduced the ability of agents to insure themselves against idiosyncratic shocks to their income process. The advantage of the methodology is that it is not necessary to state and describe the alternative insurance mechanisms that are available to the agent.

Chapter 5 started with the assumption that at least some agents are credit-constrained. It assumes a reduced form for the supply of credit and the demand for credit for a reasonably homogeneous sample of US households. By constructing estimates of these equations it is possible to estimate the proportion of US households who are credit-constrained (estimated at between 26% and 31%) and, perhaps more interestingly to see how credit constraints differ among households and how much more credit constrained households would like to borrow. Unfortunately, estimation has only been successful for the case in which the unexplained component of the supply and the demand equations are uncorrelated: probably an implausible assumption. Nevertheless, the results suggest that many groups who have low observed borrowings, such as black and poorly educated households, do so because their demand for debt is low. This has important policy implications: for instance, encouraging banks to lend to these households may make little difference to these households: instead it is important understand why these households do not use the credit markets, are they for instance more risk averse? The chapter also saw some differences with age, but the differences were not as large as might be expected if credit constraints explained why income and consumption followed each other at life-cycle frequencies.

Bibliography

[1] Amel, D., (1993) "State laws affecting the geographic expansion of commercial banks", *Working Paper, Board of Governors of the Federal Reserve System.*

[2] Asdrubali, Sørensen, B., and Yosha, O., (1996), "Channels of Interstate Risksharing: United States 1963-1990", *Quarterly Journal of Economics*, 104, pp. 1081-1110.

[3] Attanasio, O., and Davis, S., (1996), "Relative wage movements and the distribution of consumption", *Journal of Political Economy*, 104, pp. 1227-1262.

[4] Attanasio, O., and Jappelli, T., (2001), "Intertemporal choice and the cross-sectional variance of marginal utility", *The Review of Economics and Statistics*, 83(1), pp. 13-27.

[5] Attanasio, O., and Rios-Rull, J., (2000), "Consumption smoothing in island economies: can public insurance reduce welfare?" *European Economic Review*, 44(7), pp. 1225-1258

[6] Attanasio, O. P., and Weber, G., (1993), "Consumption growth, the interest rate and aggregation", *Review of Economic Studies*, 60, pp631-649.

[7] Barsky R. B., and Miron J. A., (1989), "The seasonal cycle and the business cycle," *Journal of Political Economy*, 97, pp503-534

[8] Cambell, J. Y., (1987), "Does saving anticipate declining labour income? An alternative test for the permanent income hypothesis", *Econometrica*, 55, pp1249-1273.

[9] Carroll, Christopher D., and Summers, Lawrence H., (1991), "Consumption Growth Parallels Income Growth: Some New Evidence". In National Saving and Economic Performance, edited by B. Douglas Bernheim and John B. Shoven. Chicago Press (for NBER), pp305-343.

[10] Carroll, Christopher D., (1997), "Buffer-Stock saving and the lifecycle/permanent income hypothesis," *Quarterly Journal of Economics*, 112(1), pp1-56.

[11] Charles, (2000), "Consumer Bankruptcy Law and Consumer Borrowing", *CSEF working paper no. 40*, CSEF, University of Salerno.

[12] Cochrane, J., (1991), "A simple test of consumption insurance", *Journal of Political Economy*, 99, pp. 957-976

[13] Cox, D., and Jappelli, T., (1993) "The effect of borrowing constraints on consumer liabilities", *Journal of Money, Credit and Banking*, 25, pp197-213.

[14] Deaton, A., (1991), "Saving and liquidity constraints", *Econometrica*, 59, pp. 1221-1248

[15] Deaton, A., (1992), *Understanding Consumption*, Oxford: Clarendon.

[16] Deaton, A., (1993), "On risk, insurance and intra-village smoothing", preliminary draft.

[17] Deaton, A., and Paxson, C., (1994), "Intertemporal choice and inequality", *Journal of Political Economy*, 102, pp. 437-467

[18] Deaton, A., Gourinchas, P. O., and Paxson, C., (2000), "Social security and inequality over the lifecycle", NBER working paper no. 7570.

[19] Dercon, S. and Krishnan, P., (2000), "In sickness and in health: risk sharing within households in rural Ethiopia", *Journal of Political Economy*, 108, pp. 688-727.

[20] Duca, J. and Rosenthal, S., (1994), "Borrowing constraints and access to owner-occupied housing", *Regional Science and Urban Economics*, 24 (3), pp. 301-322.

[21] Eaton, J., and Gersovitz, M., (1981), "Debt with potential repudiation: theoretical and empirical analysis", *Review of Economic Studies*, 48, pp. 289-309.

[22] Flavin, M., (1981), "The adjustment of consumption to changing expectation about future income," *Journal of Political Economy*, 89, pp974-1009.

[23] Feenberg D., and Coutts E., 1993, "An Introduction to the TAXSIM Model", *Journal of Policy Analysis and Management*, 12(1).

[24] Garcia, R., Lusardi, A., and Ng, S., (1997), "Excess Sensitivity and Assymetries in Consumption: An Empirical Investigation", *Journal of Money, Credit and Banking*, 29(2), pp154-176.

[25] Gropp, R., Scholz, J. K., and White, M. J., (1997), "Personal Bankruptcy and Credit Supply and Demand", *Quarterly Journal of Economics*, 112(1), pp217-251.

[26] Gross, D., and Souleles, N., (2002), "Do liquidity constraints and interest rates matter for consumer behaviour? Evidence from credit card data", *Quarterly Journal of Economics*, 114(1), pp149-185.

[27] Hajivassiliou, V., (1987), "The external debt repayments problems of LDC's: an econometric model based on panel data", *Journal of Econometrics*, 36, pp205-230.

[28] Hajivassiliou, V., (1994), "A simulation estimation analysis of the external debt crisis of developing countries", *Journal of Applied Econometrics*, 9(2), pp109-132.

[29] Hajivassiliou, V., and Ioannides, Y., (1996), "Duality and Liquidity Constraints under uncertainty", *Journal of Economic Dynamics and Control*, 20, pp1177-1192.

[30] Hajivassiliou, V., and Ioannides, Y., (2001), "Dynamic Switching Regression Models of the Euler Equation with Imperfect Regime Classification", *LSE mimeo.*

[31] Hajivassiliou, V., and Ioannides, Y., (2002a), "Dynamic Switching Regression Models of the Euler Equation: Consumption and Liquidity Constraints", *LSE Sticerd discussion paper.*

[32] Hajivassiliou, V., and Ioannides, Y., (2002b), "Unemployment and Liquidity Constraints", *LSE Department of Economics working paper.*

[33] Hajivassiliou, V., and Ioannides, Y., (2002c), "An econometric explanation of discouragement effects in credit and labour markets: the role of expectations of binding quantity constraints", *LSE Department of Economics working paper.*

[34] Hall, R. E., (1978), "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence", *Journal of Political Economy*, 86, pp971-987.

[35] Hall, R. E., and Mishkin, F. S., (1982), "The sensitivity of consumption to transitory income: estimates from panel data on households", *Econometrica*, 50, pp461-481.

[36] Hartley, M. J., and Mallela, P., (1977) "The asymptotic properties of a maximum likelihood estimator for a model of markets in disequilibrium", *Econometrica*, 45, pp1205-1220.

[37] Hayashi, F., (1987), "Tests for liquidity constraints: a critical survey", in T. Bewley, ed., *Advances in Econometrics, Fifth World Congress, vol. 2*, Cambridge University Press, pp91-120.

[38] Hwang, H. -S., (1980), "A test of a disequilibrium model", *Journal of Econometrics*, 12, pp319-334.

[39] Hynes, Richard, and Bekowitz, Jeremy, (1998), "Bankruptcy Exemptions and the Market for Mortgage Loans", *Wharton School Working Paper no. 98-17*, University of Pennsylvania.

[40] Jackson, T., (1986), *The Logic and Limits of Bankruptcy Law*, Cambridge MA: Harvard University Press.

[41] Jaffee, D., M., and Russel, T., (1976), "Imperfect information, uncertainty, and credit rationing", *Quarterly Journal of Economics*, 80, pp651-666

[42] Jappelli, T. and Pistaferri, L., (1999), "Intertemporal choice and consumption mobility", *CSEF working paper* no. 23

[43] Jappelli, T., Pischke, S., and Souleles, N., (1998), "Testing for liquidity constraints in Euler equations with complementary data sources" *Review of Economics and Statistics*, 80(2), pp. 251-262.

[44] Johnson, Norman, L., and Kotz, Samuel, (1972), "Distributions in Statistics: Continuous Multivariate Distributions" *Wiley Series in Probability and Mathematical Statistics*, John Wiley & Sons Inc., New York.

[45] Kehoe, Timothy J., and Levine, David K., (1993), "Debt-Constrained Asset Markets," *Review of Economic Studies*, 60, pp865-888.

[46] Kocherlakota, Narayana, (1996), "Implications of efficient risk sharing without commitment", *Review of Economic Studies*, 63, pp595-609.

[47] Kooiman, P., van Dijk, H. K., and Thurik, A. R., (1985) "Likelihood diagnostics and Bayesian analysis of a micro-disequilibrium model for retail services", *Journal of Econometrics*, 29, pp323-354.

[48] Kreugar and Perri, (1999), "Risk-sharing, private insurance markets or redistributive taxes", mimeo.

[49] Kroszner, R. S., and Strahan, P. E., (1999) "What drives deregulation? Economics and Politics of the relaxation of branching restrictions", *Quarterly Journal of Economics*, 114, pp1437-1467

[50] Mace, B., (1991), "Consumption volatility: borrowing constraints or full insurance", *Journal of Political Economy*, 99, pp928-956.

[51] MaCurdy, T. E., (1982), "The use of time series processes to model the error structure of earnings in a longitudinal data analysis", *Journal of Econometrics*, 18, pp82-114

[52] Miron J. A., (1986), "Financial Panics, the seasonality of the nominal interest rate, and the founding of the Fed" *American Economic Review*, 76, pp125-140

[53] Mirrlees, J. A., (1971), "Optimal tax theory: a synthesis", *Review of Economic Studies*, pp. 327-358.

[54] Perraudin W. R.M., and Sørensen, B., (2000), "The demand for risky assets: sample selection and household portfolios", *Journal of Econometrics*, 97, pp. 117-144.

[55] Quandt, R., (1988), *The Econometrics of Disequilibrium*, Basil Blackwell, London.

[56] Scholz, J. K., (1996), "In work benefits in the United States: the earned income tax credit", *Economic Journal*, 106(434), pp. 156-169.

[57] Stiglitz, Joseph E., and Weiss, Andrew, (1981), "Credit Rationing in Markets with incomplete information", *American Economic Review*, 71, pp393-410.

[58] Townsend, R.M., (1994), "Risk and insurance in village India", *Econometrica*, 62 (3), pp. 539-591.

[59] Varian, H., 1980, "Redistributive taxation as social insurance", *Journal of Public Economics*, 14, pp. 49-68.

[60] Westlaw, *Annotated State Codes*, various.

[61] Zeldes, S., (1989), "Consumption and liquidity constraints: an empirical investigation", *Journal of Political Economy*, 97, pp305-346.

[62] Zeldes, Stephen P., (1989), "Optimal consumption with stochastic income: deviations from certainty equivalence", *Quarterly Journal of Economics*, 104, pp275-298.