

Individual Consumption and Aggregate Implications

by

Jonathan A. Parker

B.A., Yale University (1988)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1996

© Jonathan A. Parker, MCMXCVI. All rights reserved.

The author hereby grants to Massachusetts Institute of Technology
permission to reproduce and to distribute publicly paper and
electronic copies of this thesis document in whole or in part.

Signature of Author ..

.....
Department of Economics
May 16, 1996

Certified by

.....
Ricardo J. Caballero
Professor of Economics
Thesis Supervisor

Certified by

.....
Olivier Jean Blanchard
Professor of Economics
Thesis Supervisor

Accepted by

.....
Richard S. Eckaus
Chairperson, Departmental Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 10 1996

ARCHIVES

Individual Consumption and Aggregate Implications

by

Jonathan A. Parker

Submitted to the Department of Economics on May 16, 1996,
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

This thesis consists of three essays on consumption.

Chapter 2 exploits a natural experiment provided by the pattern of Social Security tax withholding to test whether household consumption responds to expected changes in take-home pay. According to the basic Permanent Income Hypothesis, consumption should not respond to expected changes in income. In the U.S., Social Security taxes cause predictable swings in after-tax income due to both legislated increases and annual contribution limits. I find large and significant violations of consumption smoothing: consumption increases one-half percent for every expected, one percent increase in after-tax income.

Chapter 3 employs a synthetic cohort technique and Consumer Expenditure Survey data to construct average age-profiles of consumption and income over the working lives of typical households across different education and occupation groups. Even after controlling for family and cohort effects, consumption profiles are not flat, and seem to track income at young ages. Using these profiles, we estimate a structural model of optimal life-cycle consumption expenditures in the presence of realistic income uncertainty. The model fits the data quite well. Consumer behavior changes strikingly over the life-cycle due mostly to the expected profile of income. Young consumers behave as “buffer-stock” agents. Around age 43, the typical household starts accumulating liquid assets for retirement and its behavior mimics more closely that of a certainty equivalent consumer. Finally, we decompose saving over the life cycle into retirement and precautionary saving.

In Chapter 4, I drop the standard simplifying assumption that goods are purchased continuously in time and by all buyers. I analyze a market in which buyers optimally time a single purchase and in which sellers have market power. Market dynamics have many features of business cycles. Expected fluctuations in demand are inherently fluctuations in the elasticity of demand, and lead to smaller markups on the up-side of booms. Buyer intertemporal optimization opposes this force – generating real price stickiness and smoothing prices over time. These mechanisms produce prices which are less variable than quantities, countercyclical markups, and persistence of demand shocks. In industry

data, consumer goods for which timing is likely to be important exhibit less real price response to demand-driven movements in sales.

Thesis Supervisor: Ricardo J. Caballero
Title: Professor of Economics

Thesis Supervisor: Olivier Jean Blanchard
Title: Professor of Economics

Contents

1	Introduction	13
2	The Reaction of Household Consumption to Predictable Changes in Payroll Tax Rates	20
2.1	Introduction	20
2.2	Consumption Smoothing	24
2.2.1	Consumption Theory	24
2.2.2	The Basic LCH/PIH and the Linear Euler Equation	26
2.2.3	Three Alternatives Hypotheses to Consumption Smoothing	27
2.3	A Brief History of Excess Sensitivity Tests	30
2.4	The Test of the Basic LCH/PIH	33
2.5	The Consumer Expenditure Survey	35
2.6	Estimation and Results	38
2.7	Evaluating Alternative Hypotheses	44
2.8	Conclusion	48
A	Appendix: The Consumer Expenditure Survey	50
3	Consumption Over the Lifecycle (with Pierre-Olivier Gourinchas)	53
3.1	Introduction	53
3.2	Consumption Behavior with Stochastic Income	60
3.2.1	The Canonical Model with Labor Income Uncertainty	60

3.2.2	Solving for Optimal Consumer Behavior	65
3.2.3	Characterization of Individual Consumption Behavior	67
3.3	Data and Consumption and Income Profiles	70
3.3.1	Profile Construction Methodology	70
3.3.2	The Consumer Expenditure Survey and Our Use of it	73
3.3.3	Life Cycle Profiles	76
3.4	Estimation Strategy	78
3.4.1	Method of Simulated Moments (MSM) Estimation	78
3.4.2	Remaining Calibration and Data	82
3.5	Estimation Results	85
3.5.1	Results for the Entire Sample	85
3.5.2	Disaggregated Evidence	90
3.5.3	Robustness Checks and Extensions	92
3.6	Conclusion	96
B	Appendix: Solving the Consumer Optimization Problem	118
B.1	Euler equation	118
B.2	Gauss-Hermite quadrature	118
B.3	Consumption rules	119
C	Appendix: Data	121
D	Appendix: Method of Simulated Moments	123
4	The Timing of Purchases, Market Power, and Economic Fluctuations	129
4.1	Introduction	129
4.2	The Market Structure	133
4.2.1	Consumers	134
4.2.2	Sellers	137
4.2.3	Market Equilibrium	138
4.3	Properties of Dynamic Equilibria	141

4.3.1	The Steady State	142
4.3.2	Equilibrium and the Coase Conjecture	142
4.3.3	Demand Fluctuations with Constant Marginal Cost	143
4.3.4	Demand Fluctuations with Increasing Marginal Cost	144
4.4	Properties of Dynamic Collusive Equilibria	146
4.4.1	Sellers	146
4.4.2	Market Equilibrium	147
4.4.3	The Steady State	148
4.4.4	Demand Fluctuations with Increasing Marginal Cost	148
4.5	Related Models of Markups and Durable Goods	149
4.6	Empirical Evidence	151
4.6.1	The Data	151
4.6.2	Testing Price Smoothness and Dynamics	153
4.6.3	Constructing Markups	156
4.6.4	Testing Markup and Dynamics	158
4.7	Conclusion	162
E	Appendix: Proofs of Lemmas	168
F	Appendix: Industry Data	171
5	References	175

List of Tables

2.1	The Social Security Tax Structure, 1980-1993	22
2.2	Sample Statistics	37
2.3	The Reaction of Nondurable Consumption	38
2.4	Alternative Measures of the Social Security Tax Variable	40
2.5	High Consumption Subsample and Treatment vs. Control Group	42
2.6	Treatment, Control and Neither Group Statistics	43
2.7	Results for Asset and Age Splits	45
2.8	The Reaction Across Categories of Consumption	48
3.1	Cell Sizes for Consumption Profiles	75
3.2	F-Tests for Flatness of Consumption Profiles	77
3.3	Variance of Income Shocks	83
3.4	Data for Retirement Consumption Rule	84
3.5	Results of Structural Estimation	86
3.6	Benchmark Certainty Case, $\rho = 0.4897$	91
3.7	Estimates from the Stochastic Model, $\rho = 0.4897$	92
3.8	Robustness Checks, $\rho = 0.4897$	93
4.1	Real Price Regressions	154
4.2	Markup Regressions	159
4.3	Markup Regressions on Subsample	161
F.1	Durability and Infrequency of Purchase, Part I	172

F.2 Durability and Infrequency of Purchase, Part II 174

List of Figures

- 3-1 Consumption Rules without Retirement 98
- 3-2 Consumption Rules with Retirement 98
- 3-3 Sample Individual Profile 99
- 3-4 Sample Individual Profile with Zero Income Shock 99
- 3-5 Expected Cash on Hand 100
- 3-6 Household Consumption and Income over the Lifecycle 101
- 3-7 Family Size over the Lifecycle 101
- 3-8 Consumption and Income with and without Cohort Adjustments 102
- 3-9 Consumption and Income with and without Family Adjustments 102
- 3-10 Total Liquid Asset to Income Ratio 103
- 3-11 Consumption(O) and Income(+) over the Lifecycle by Education 104
- 3-12 Consumption(O) and Income(+) over the Lifecycle by Occupation 105
- 3-13 Rescaled Smoothed Consumption over the Lifecycle 106
- 3-14 Income, Consumption and Consumption Predicted by the CEQ LCH . . . 107
- 3-15 Income, Consumption and CEQ Consumption by Education 108
- 3-16 Income, Consumption and CEQ Consumption by Occupation 109
- 3-17 Simulated and Actual Consumption Profiles 110
- 3-18 Simulated Consumption Profiles for Different Discount Rates 111
- 3-19 The Predicted Savings Rate Profile 112
- 3-20 Target Cash on Hand 113
- 3-21 Life-Cycle and Buffer-Stock Savings 114
- 3-22 Simulated and Actual Consumption Profiles by Education 115

3-23	Simulated and Actual Consumption Profiles by Occupation	116
3-24	Robustness Checks: Simulated and Actual Consumption Profiles	117
4-1	Example Market Dynamics	164
4-2	Fluctuations with Constant Marginal Costs	165
4-3	Fluctuations with Increasing Marginal Costs	166
4-4	Fluctuations with Colluding Firms and Increasing Marginal Costs	167

Acknowledgments

I am forever and deeply indebted to the many talented people who, during the course of my life, have taught me not only knowledge, but more importantly the pleasures of thought and the imagination to create understanding. I cannot possibly thank here by name the full complement of those to whom thanks are owed.

First and foremost, I am extremely grateful to my advisers, Ricardo Caballero and Olivier Blanchard, who have provided inspiration, insight, and guidance without which this thesis would never have been completed. I will spend my career aspiring to match the quality of their research, teaching, and advising.

I have benefitted from the extremely high quality of formal and informal teaching at M.I.T. I thank those professors who have forged my thinking in classrooms, offices and hallways. My fellow students have shown boundless patience, generosity and wisdom, and my thesis owes much to their clarity of thought and insights. I am especially grateful to those who worked alongside me in study groups, helped me prepare for generals, and eased my transition into second-year classes when I first arrived. Finally, I thank Pierre-Olivier Gourinchas for sharing his friendship and his brilliance.

This thesis also owes much to the friends I made during my years in Ann Arbor and to the hospitality of the Michigan Economics Department. In particular, I thank Charles Fleischman for sparking my interest in macroeconomics and for listening to, reading, and helping to shape many rough ideas.

I deeply appreciate my superb teachers at Yale, Roxbury Latin, and Meriwether Lewis.

I thank my parents and brother for the motivation, imagination, love, and many meals which I needed to complete this thesis. I cannot thank Anne enough for being there and for providing dissertation support, brilliant insight, and spelling assistance.

I thank the National Science Foundation, the Sloan Foundation, and the World Economy Laboratory for providing the financial support for my dissertation.

Finally, I thank everyone, many already mentioned, who has helped me to laugh along the way.

Chapter 1

Introduction

The allocation of national output between consumption and savings is arguably the most important determinant of economic growth and business cycles. Consumers purchase roughly two thirds of national output annually, and recessions often are amplified and sometimes are caused by shortfalls in household demand for consumption goods. The difference between household consumption and income – savings – is the major determinant of the resources available for new investment and the price of new investment, that is, the interest rate of the economy. Finally, the stock of wealth held by households determines the wealth of the nation.

To understand business cycles and growth requires understanding the forces that drive national saving and consumption. This thesis investigates these driving forces by studying the savings and consumption decisions of individual households. I seek to shed light on the ways in which household demand for consumption goods deviates from the canonical theories of consumption. Further I seek to qualify and quantify how these deviations impact the evolution of aggregate consumption and savings and thus the evolution of the economy.

The main theories of consumption behavior, the Life-Cycle Hypothesis and Permanent-Income Hypothesis, predict that households should use savings to stabilize their level of consumption. That is, if income is expected to decline, a household should initially save more so that consumption does not fall when income declines. This implication is called

consumption smoothing. While optimizing models of consumer behavior do not always predict perfect smoothing of consumption, consumption smoothing is widely considered to be a reasonable characterization of household behavior, and, in turn, of aggregate behavior. To the extent that consumption smoothing fails, aggregate and individual models should incorporate alternative formulations of consumer decisions. These alternative formulations have the potential to increase our understanding of macroeconomic events and clarify the role of policy in maintaining economic growth. For example, and of special relevance given the test conducted in Chapter 2, if consumption smoothing fails, then tax policy may have significant effects on household expenditures and can be used, if desired, to stimulate aggregate demand in recessions. As another example, relevant to the research in Chapter 3, individual retirement accounts may help to increase national saving because such accounts cannot be used costlessly to “buffer” consumption in bad times.

In Chapter 2, I conduct a novel test of the basic Permanent Income Hypothesis, which demonstrates that consumption smoothing is a poor approximation to consumer behavior over *short horizons*, such as a few months. Using household-level consumption data, I examine whether families increase their expenditures on consumption contemporaneously with expected changes in after-tax income that are driven by changes in Social Security tax rates. If consumption smoothing were correct, households would not increase their consumption at all. In the United States, individuals with wage and salary income are subject to Social Security tax withholding of around seven percent of their gross pay up to an annual maximum income level. Social Security tax rates provide two sources of expected variation in income which I use to test consumption smoothing. First, there have been a series of pre-announced tax rate increases in the 1980's which have affected all wage and salary earners. Due to different shares of labor income across households and to the fact that some earners are not subject to Social Security taxes, these changes produce different changes in income for different households. Second, when income earned in a calendar year reaches the maximum taxable amount, an individual's take-home pay increases because Social Security taxes are no longer withheld from the individual's

paycheck. In January of the following year, when withholding begins again, take-home pay falls again.

I find that individual consumers change their consumption in response to the expected fluctuations in income induced by the Social Security tax system. I estimate that an expected, one percent change in after-tax income decreases nondurable consumption by a half of a percent. I also attempt to distinguish among alternative hypotheses which might explain this failure. Here the results of the chapter are less clear cut. There is little evidence for or against the hypothesis that consumption tracks income more for low-asset households, as would be the case if consumption smoothing failed either due to the presence of liquidity constraints or due to individual uncertainty generating significant precautionary saving. There is however some evidence – consistent with that found in chapter 3 – that young households have a greater response of consumption to expected changes in income. Since young households are more likely to have lower assets, this finding constitutes some evidence in favor of theories of precautionary saving or liquidity constraints.

I also evaluate a model in which consumers are near-rational or face adjustment costs so that consumers allow their expenditures to track income to some extent. I test whether the response of consumption to expected changes in income is strongest in subcategories of consumption in which households 1) can substitute purchases across time with little utility loss and 2) in which purchases generally do not employ credit. I indeed find evidence consistent with the near-rationality story: semi-durable goods seem to exhibit the largest violations of consumption smoothing.

Chapter 3, which is joint work with Pierre-Olivier Gourinchas, turns to testing whether consumption smoothing is a reasonable characterization of household behavior at *long horizons*, that is, over the life-cycle. We begin this task in a nonstructural manner by demonstrating that the age profile of consumption for a typical household is not flat and is related to the household's expected profile in income. Importantly, this result holds after controlling for family composition and cohort effects. We confirm this finding across subsets of the population constructed by grouping households into different ed-

educational attainments and occupational groupings, which have different typical profiles of income. We conclude that profiles of consumption are significantly hump-shaped and that consumption tracks income more closely early in life.

Taking a more structural approach, we next embed realistic levels of income uncertainty within the canonical model of life-cycle consumer behavior, as in Carroll (1993a) and Zeldes (1989b). The solution to this model would yield consumption smoothing if uncertainty were ignored. However, under uncertainty, consumption depends on the path of expected income. Thus consumer behavior varies systematically over the life cycle. Specifically, early in life, when expected income growth is low, consumers may behave as “buffer-stock” agents, consuming roughly their income and saving only small amounts to buffer against bad income draws. This buffer-stock behavior requires that consumers be impatient relative to the expected growth rate of income. As households age, expected income growth declines. Consequently, older households, whether impatient or not, behave more like certainty-equivalent consumers, smoothing consumption and saving for retirement.

We use a Simulated Method of Moments technique to estimate a partially-calibrated version of our structural model. That is, we pick starting values for the discount rate and the intertemporal elasticity of substitution of consumers, and solve for the consumption functions of such consumers at every age. By simulating the lives of many consumers using these consumption functions, we create predicted profiles of average consumption. We then iterate to match the predicted profiles to those from the data, and estimate an intertemporal elasticity of substitution of 2.04 and a discount rate of 3.9%. Thus reasonable parameters are consistent with our interpretation of consumer behavior.

Estimated consumption functions for households change strikingly across ages, and imply that consumers behave like “buffer-stock” consumers early in their working lives and only do a reasonable job of consumption smoothing as retirement nears. Thus, a significant fraction of consumers consists of target savers, who do not smooth consumption at high frequencies. Around age 43, households make the transition from buffer-stock consumers to life-cycle consumers who smooth consumption.

Chapter 3 also contributes to the debate on the determinants of wealth accumulation. In the fitted model, saving and consumption at each age are determined by the interactions between the precautionary and retirement savings motives. We find that early in life, life-cycle savings are negative because households would like to borrow against expected future labor income. However, uncertainty causes households to build a buffer stock of savings, implying that precautionary saving is positive. Late in life, labor-income uncertainty is mostly resolved, and consumers run down their buffer stocks, while retirement saving becomes positive and large. Both the fact that most households hold few, if any, assets and the fact that most households do not start saving for retirement until late in life often have been interpreted as evidence that household behavior is irrational. We find that these facts are instead consistent with rational, Life-Cycle behavior when one explicitly accounts for labor income uncertainty.

While Chapters 2 and 3 seek to test the relevance of consumption smoothing at the household level, Chapter 4 focuses instead on a market in which consumers must choose when to buy a single item of a single good. This structure, which I consider for a single market, exhibits many of the characteristics of the aggregate economy during business cycles. Thus, I argue that economic models should focus not only on *how much* consumers choose to buy in every period, but also upon *when* consumers choose to transact.

More specifically, I analyze a market with three crucial characteristics. Sellers have market power, so that changes in the effective elasticity of demand can lead to changes in the markup between price and marginal cost. Second, buyers choose when to make discrete, lumpy purchases. Finally, the number of consumers who would like to purchase in each period fluctuates exogenously through time.

I demonstrate that two main forces propagate and alter cycles in this market. First, fluctuations in the distribution of buyers waiting to buy represent changes in the share of consumers who are just about to buy, that is, changes in the elasticity of demand. Firms respond to these changes by changing prices in a way so as to increase sales and decrease markups at the beginnings of booms, while decreasing sales and increasing markups at the starts of slumps. Thus, in this market, the ends of recessions and booms are amplified.

However, a second characteristic of dynamic equilibria fights this amplification: the ability of buyers to delay or accelerate their purchases smooths prices. Sellers cannot raise prices relative to surrounding periods without losing sales. The ability of consumers to time purchases constrains price movements through time by increasing the effective intertemporal elasticity of substitution faced by sellers. With increasing marginal costs of production, the lack of movement in real price over cycles leads naturally to a countercyclical markup.

Together, these mechanisms create market dynamics which mimic quite closely those of aggregate business cycles. First, the stickiness of the real price implies that prices will vary less than quantities over cycles. This then provides an answer to the puzzle of why prices do not decline to “clear markets” in recessions. That is, prices are clearing markets, but intertemporal substitution of purchases keeps prices from moving and quantities adjust. Second, as noted, the markup of price over marginal cost moves countercyclically. During booms more (nominal) aggregate output is produced without increasing the cost of inputs to production proportionally. Finally, also under the assumption of increasing marginal costs, timing generates propagation of demand shocks. Because time periods are not separate markets, each clearing independently, high prices in one period cause buyers to move transactions to nearby periods. Thus a temporary increase in demand increases sales over several periods.

I use industry data to demonstrate that, as predicted by the model, consumer goods for which timing is likely to be important do exhibit less real price response to demand-driven movements in sales. Evidence on the behavior of markups is murky and neither supports nor rejects the model of timing.

The findings of this dissertation have two main themes. First, consumption smoothing does not accurately capture the behavior of many consumers in the economy. At high frequencies, that is within a three-month period, household consumption changes fifty cents for each dollar that income changes expectedly. At lower frequencies, that is over the span of the life cycle, households consume roughly their income when young and only smooth consumption relative to their expected profile of income after age 45. Second,

explicit consideration of both the timing of consumer purchases and of the individual uncertainty facing consumers can generate more empirically realistic market dynamics and consumer behavior. The former generates partial-equilibrium cycles with many of the main features of business cycles, while the latter reproduces life-cycle consumption patterns. The combination can match the observed high-frequency consumption-income parallels.

Chapter 2

The Reaction of Household Consumption to Predictable Changes in Payroll Tax Rates

2.1 Introduction

This chapter evaluates the key implication of the basic Life Cycle/Permanent Income Hypothesis (LCH/PIH): that predictable changes in income should have no effect on the growth rate of consumption, and only new information should lead consumers to revise consumption behavior.¹ This implication— often called Hall’s Martingale hypothesis — is important for understanding the effectiveness and optimal timing of fiscal policy, the causes and propagation of business cycles, and the effects of income fluctuations on the growth of the economy. As is well known, to the extent that the basic LCH/PIH fails to hold, predictable changes in tax rates can have real effects on consumption and investment, and government debt policy may affect national savings and possibly growth rates.

Using microeconomic consumption data from the Consumer Expenditure Survey

¹As I discuss in the subsequent section, the theory predicts smoothing in marginal utility which does not always imply smoothing of consumption.

(CEX), I examine whether households increase their expenditures on consumption contemporaneously with expected changes in Social Security tax rates.² Individuals with wage and salary income earned in the United States are subject to Social Security tax withholding of around seven percent of their gross pay up to an annual maximum income level. Social Security tax rates provide two sources of variation. First, there have been a series of pre-announced tax rate increases in the 1980's which have affected all wage and salary earners. Due to different shares of labor income across households and to the fact that some earners are not subject to Social Security taxes, these changes produce different changes in income for different households. Second, when income earned in a calendar year reaches the maximum taxable amount, an individual's take-home pay increases because Social Security taxes are no longer withheld from the individual's paycheck. In January of the following year, when withholding begins again, take-home pay falls again. Table 2.1 shows the rates and caps for the sample period of 1980 to 1993.

Under the null hypothesis, household consumption should not respond to these changes in tax rates since they are expected. The small changes in tax rates are legislated and announced well ahead of time. High-income individuals who hit the tax cap generally see this fluctuation in their after-tax wages annually, and, after 1982, the changes in the tax cap attempt to adjust the cap for average wage growth.³ In the CEX, the average annual wage and salary income among heads of households who hit the tax cap is \$53,169. In the middle year of my sample, 1986, a person with this wage income would have a \$317 temporary increase in monthly after-tax income from mid-October until the end of December. The method of Social Security taxation thus presents a nice natural experiment to evaluate Hall's Martingale hypothesis.

Tests of the basic LCH/PIH in this chapter avoid several of the weaknesses of earlier

²I am indebted to Joel Slemrod for originally pointing me to the Social Security tax cap as a nice test of the basic LCH/PIH.

³Beginning in 1982, the maximum contribution was adjusted upward automatically based on the average annual percent wage change and then rounded to the nearest figure divisible by \$300. From 1990 to 1992 this transition rule was accelerated because previous adjustments had ignored non-wage and deferred compensation which had been growing more rapidly than wage compensation.

Table 2.1: THE SOCIAL SECURITY TAX STRUCTURE, 1980-1993

YEAR	TAX RATE (PERCENT)	MAXIMUM ANNUAL CONTRIBUTION PER EARNER	MAXIMUM ANNUAL TAXABLE EARNINGS
1980	6.13	1,588	25,900
1981	6.65	1,975	29,700
1982	6.70	2,171	32,400
1983	6.70	2,392	35,700
1984*	6.70	2,533	37,800
1985	7.05	2,792	39,600
1986	7.15	3,003	42,000
1987	7.15	3,132	43,800
1988	7.51	3,380	45,000
1989	7.51	3,605	48,000
1990	7.65	3,924	51,300
1991			
OASDI:	6.20	3,311	53,400
HI:	1.45	1,812	125,000
1992			
OASDI:	6.20	3,441	55,500
HI:	1.45	1,888	130,200
1993			
OASDI:	6.20	3,571	57,600
HI:	1.45	1,958	135,000

* The tax rate in 1984 includes the tax credit.

Data refers to OASDI and HI except where noted. Source: Social Security Administration (1990) Tables 2.A3, 2.A4, and 2.A5 and The Social Security Bulletin, (January, 1993).

analyses.⁴ First, the tax change is exogenous to the consumer and observable to the researcher from labor income information; thus expected income changes can be measured with reasonable accuracy and questionable methods for obtaining expected income changes are avoided. Second, the Consumer Expenditure Survey provides data on many different types of consumption expenditures at the household level. Previous microeconomic studies often have had to make do with change in food consumption as a measure of nondurable consumption. Third, the dataset is representative of the entire U.S. population, although to some extent, the Social Security experiment relies more heavily on the wealthy. Many previous studies have focussed on non-representative groups such as

⁴Section 2.3 discusses the previous literature in more detail.

union workers or the elderly. Finally, if one is concerned with estimating the effect of fiscal policy on consumption, this test yields exactly the parameter of interest. I estimate the extent to which current spending changes in response to expected, tax-driven changes in income. The distinction among tax-driven changes in income and other changes can be important once the basic LCH/PIH is rejected, since some alternative theories of consumption behavior suggest that the source of the expected increase in part determines the magnitude of the consumption response.

To cut to the chase, I find that individual consumers change their consumption in response to the expected fluctuations in income induced by the Social Security tax system. I estimate that an expected, one percent change in after-tax income decreases nondurable consumption by a half of a percent. While not statistically significant, similar point estimates are obtained when identification is derived from the differences in behavior between a treatment group of earners who pay Social Security taxes and a control group of earners who do not. Due to large standard errors, I uncover little to no evidence for the hypothesis that families with less cash-on-hand react more to expected changes in income. Finally, the strongest violations of consumption smoothing are found in subcategories of consumption in which households can substitute purchases across time with little utility loss and in which purchases generally do not employ credit.

In the next section, I present the basic LCH/PIH and alternative hypotheses. Section 3 of this chapter contains an overview of previous empirical research on this topic. Section 4 describes the methodology used to test the basic LCH/PIH. In Section 5, I describe the data, and in Section 6 present the results and considers their robustness. Section 7 pushes the test further to evaluate the alternative hypotheses, while the final section concludes the chapter. In Appendix A, I provide the details of the data and variable construction.

2.2 Consumption Smoothing

2.2.1 Consumption Theory

The starting point for most modern studies of consumption is the following canonical model of consumer optimization. Households choose consumption expenditures to maximize utility subject to a lifetime budget constraint:

$$Max \quad E_s \left[\sum_{t=s}^T \beta^{t-s} u(c_t) \nu(z_t) + \beta^{T+1} V(A_{T+1}, z_{T+1}) \right] \quad (2.1)$$

$$S.T. \quad A_{t+1} = R_t(A_t + (1 - \tau_t)Y_t - c_t) \quad (2.2)$$

$$A_s \quad \text{given} \quad (2.3)$$

$$A_{T+1} \geq 0 \quad (2.4)$$

where E_s is the expectations operator conditional on all information available at time s ; β is the discount factor; c_t is consumption expenditures and consumption flow in period t ; $u(\cdot)$ is an intertemporally-separable felicity function which is stable through time, increasing and concave; z_t contains variables such as family size which change over the life cycle and change marginal utility (through the function $\nu(\cdot)$); A_t is household assets; R_t is the gross real interest rate; Y_t is stochastic labor income;⁵ τ_t is the Social Security tax rate on labor income; and finally $V(\cdot, \cdot)$ captures the value of assets left at death in the form of bequests. Households are not allowed to die with negative assets. When $V \equiv 0$, equations (2.1)-(2.4) represent the life-cycle hypothesis. When individuals live S years and $V(A_{T+1}, z_{T+1}) \equiv \sum_{t=T+1}^{T+S} \beta^{t-T-1} u(c_t) \nu(z_t) + \beta^{T+S+1} V(A_{T+S+1}, z_{T+S+1})$, these equations represent the permanent income hypothesis (with occasional positive asset restrictions).⁶

⁵Following convention, time t variables are assumed known when the household chooses consumption in t .

⁶Friedman (1957) and Modigliani and Brumberg (1956).

Prior to the last period of life, the Euler equation for the problem is:

$$u'(c_t)\nu(z_t) = \beta R_t E_t [u'(c_{t+1})\nu(z_{t+1})]. \quad (2.5)$$

This equation implies that households seek to equate marginal utility across time, with possible fluctuations due both to changes in the relative price of consumption across periods, R_t , and to changes in variables such as family size, z_t .

Under rational expectations, if z_t is constant and the interest rate equals the discount rate, equation (2.5) implies that marginal utility follows a Martingale process. Since households plan to equate marginal utility today and tomorrow, any *ex post* difference between marginal utility today and tomorrow must be due to information not available today. That is, $u'(c_{t+1}) = u'(c_t) + \varepsilon_{t+1}$, where the innovation at time $t + 1$ is orthogonal to all information available to the households at time t . In his seminal paper, Hall (1978) derives this result and shows that, if in addition the utility function is quadratic, then equation (2.5) implies that consumption follows a Martingale process and $c_{t+1} = c_t + \varepsilon_{t+1}$. Under the assumptions made by Hall (1978), one can proceed to test this theory by testing whether information available to households at time t can help to predict consumption growth between t and $t + 1$.⁷

Instead of making the assumptions used by Hall (1978), the more recent literature takes Hall's basic insights but assumes that the utility function is of the constant relative risk aversion (CRRA) form and consumption is log-normally distributed. In this case, if z_{t+1} is known at time t , the Euler equation is:

$$\Delta \ln(c_{t+1}) = \sigma \ln(\beta R_t) + \sigma \ln \left(\frac{\nu(z_{t+1})}{\nu(z_t)} \right) + \frac{\rho}{2} E_t \text{Var} (\Delta \ln(c_{t+1})) + \varepsilon_{t+1} \quad (2.6)$$

where σ is the intertemporal elasticity of substitution and ρ the coefficient of relative risk aversion.⁸

⁷Hall (1978) tests and rejects the Martingale hypothesis on aggregate consumption data for the United States.

⁸Under these assumptions $\rho = \frac{1}{\sigma}$. The Euler equation is often also motivated as a Taylor expansion. Whichever motivation is used, $\rho = \frac{1}{\sigma}$ is generally tested rather than imposed.

2.2.2 The Basic LCH/PIH and the Linear Euler Equation

A large body of literature has proceeded under the assumption that the expected variance of the log of consumption is orthogonal to information available at time t . In this case, after conditioning on the real interest rate and variables which plausibly alter the marginal utility of consumption over the life cycle, equation (2.6) implies that the percent change in consumption should be uncorrelated with any variables dated t or earlier. I refer to this as the key implication of the **basic LCH/PIH** and to the following equation as the **linear consumption Euler equation**:

$$\Delta \ln(c_{t+1}) = \sigma \ln(\beta R_t) + \sigma \ln \left(\frac{\nu(z_{t+1})}{\nu(z_t)} \right) + \epsilon'_{t+1}, \quad (2.7)$$

where I have simply added the variance term to the error term. However, as discussed in the next subsection, the variance term is unlikely to be orthogonal to all information available at time t . That the basic LCH/PIH is not generally theoretically true has not stopped and should not stop its wide use. Parsimony is the essence of modelling, and many insights can be drawn, especially at the aggregate level, from the assumption that the linear Euler equation is true at the household level.⁹ Further, to date the empirical evidence rejecting the basic LCH/PIH as a reasonable approximation to reality is not overwhelming. Given that the basic Euler equation is both highly useful and widely used, whether it is approximately true or not is perhaps the most important unresolved issue in the study of consumption.

If one estimates equation (2.7) and variables known as of time t do not enter significantly, then the coefficient on the real interest rate provides an estimate of the intertemporal elasticity of consumption, and this and the constant yield an estimate of the discount rate.

Most previous studies choose to test whether expected changes in household income have predictive power for changes in consumption in equation (2.7). One reason to

⁹Nearly all of general equilibrium macroeconomics proceeds under the assumption that the basic LCH/PIH is roughly true.

test expected income is historical: prior to Hall (1978), economists generally modelled consumption as a function of current and past income and made no formal distinction between expected and unexpected changes in income. In addition, however, testing whether expected income growth is correlated with consumption growth allows one to evaluate the economic importance of the rejection and to assess what alternative hypothesis might be causing the failure. For this second reason, I remain in this tradition and test equation (2.7) by including a measure of the change in after-tax income due to individual-specific changes in Social Security tax rates.

2.2.3 Three Alternatives Hypotheses to Consumption Smoothing

In this subsection I briefly describe three reasons why the linear Euler equation might be rejected. The first two imply that this failure should be concentrated among low-asset households. The last predicts that more durable goods that are not generally purchased using credit should track expected changes in income more closely.

The most obvious reason for the linear Euler equation not to hold is that the expected variance term in equation (2.5) is correlated with the included time t variable.¹⁰ Consider as an example, the realistic specification in which income is lognormally distributed and, as above, the utility function is CRRA. First, note that in this case households are *unwilling* to borrow against future income. Why? If a household were to arrive in its last period of life with negative assets and it were to receive an income of less than this amount, then it would have to consume nothing. Since this yields infinite marginal utility, the household would never allow this to be a possibility. By iterating backward through life, it follows that households will never borrow. Thus a household with few assets and low current income consumes roughly its income – rather than some fraction of its

¹⁰The role of uncertainty in linear Euler equation violations was first explored by Zeldes (1989b) and then Deaton (1991). Carroll (1993a) demonstrates that within this framework, assets are closely correlated with the variance term. See also Chapter 3. Dynan (1993), and Caballero (1990a) for studies that evaluate the importance of this mechanism.

permanent or life-cycle income. If income is expected to grow, as it does so, consumption will rise. Consequently, expected income growth and asset levels can be used to predict consumption growth. The mechanism for this relation is exactly the expected variance of consumption: because a household with few assets is unable to smooth consumption in response to negative shocks to income, its expected variance of consumption is high and its expected consumption growth is high.¹¹

This point is more general than the example above. Carroll and Kimball (forthcoming) demonstrate that a wide range of utility functions (the HARA class) and income processes generate consumption functions which depend upon the share of the expected present discounted value of total wealth that is comprised of currently available assets. To summarize, this alternative hypothesis predicts that the variance term in the consumption Euler equation is correlated with household assets and thus expected income growth should have predictive power for consumption growth. The relationship should be strongest for households with few assets.

A second reason for failure of the linear Euler equation is that some households might face liquidity constraints.¹² When a household is liquidity constrained, its Euler equation ceases to hold and it is *unable* to smooth consumption. With liquidity constraints, predictable rises in income, as would occur following a temporary layoff, for example, can lead to corresponding rises in consumption. For concreteness, consider the following additional two constraints on the household optimization problem (equations (2.1) to (2.4)):

$$A_t \geq 0 \quad \forall t \tag{2.8}$$

$$c_t = \hat{c} \quad \text{if } A_{t+1} = 0 \text{ and } c_t \leq \hat{c} \tag{2.9}$$

where \hat{c} is a government-provided strictly positive floor on consumption. Equation (2.8) implies that due to liquidity constraints, households are unable to borrow against future

¹¹This requires prudence of the utility function, a feature which CRRA exhibits. See Kimball (1990).

¹²This setup is motivated by that considered by Hubbard, Skinner and Zeldes (1994).

income, while equation (2.9) implies that if the household consumes all its income and assets, the government provides a consumption floor of \hat{c} . In this scenario, a household with no assets that has high expected future income is unable to consume more than \hat{c} due to liquidity constraints. When the income of such a household rises, consumption rises also. Thus, as in the first alternative considered, expected income growth should have predictive power for consumption growth for low-asset households.¹³

Third, and last, consider a model of near-rational consumers who allow consumption to track income provided that this strategy does not take them too far from the level of consumption which a fully rational agent would choose.¹⁴ Thus, households that get a few hundred dollars extra in take-home pay for a few months simply spend this money when they get it, rather than completely smoothing consumption expenditures. Since small deviations from optimal behavior have small costs, only small deviations from full rationality or small costs of shifting income across time are required for this story to be plausible.

The testable implications of this alternative theory are twofold. First, expenditures should track income more closely for goods that are purchased using current income—that is, that are not purchased using credit. Second, expenditures should track income more closely for goods with high intertemporal elasticities of substitution of expenditures—that is, for which swings in consumption provide little utility loss. If an expenditure on a good provides lasting utility, then the effective intertemporal substitutability of such a purchase is much higher.¹⁵ For example, relative to the swing in expenditures, there is a small gain in utility for a household which takes its monthly trip to the movies a week early. Buying lunch a few hours later may cause a large utility loss. Thus expenditures should respond to expected income most strongly for goods that provide some lasting

¹³Note that in both of the scenarios discussed, the possibility of having low assets (or being liquidity constrained) in the future changes the consumption decision even for households with many assets. That is, the basic LCH/PIH is not strictly true even for households with large wealth holdings. However, for these households, the linear Euler equation should provide a close approximation to the non-linear Euler equation.

¹⁴This hypothesis is similar in spirit to that proposed by Hall and Mishkin (1982).

¹⁵Heaton (1993) and Chapter 4 both explore the implications of higher intertemporal substitution due to lasting utility/durability.

utility flow or durability are that are generally financed out of income or cash on hand.¹⁶

2.3 A Brief History of Excess Sensitivity Tests

In keeping with its importance, the basic life-cycle/permanent-income model of individual consumption has been subject to much empirical scrutiny. Tests of the linear Euler equation on aggregate data generally reject the implications of the basic theory.¹⁷ For the representative consumer, expected changes in income seem to be correlated with expected changes in consumption. This does not imply that the linear Euler equation fails at the household level however. First, in aggregate data it is difficult to measure expected income growth. This is due both to data problems and the instrumental variables methodology used to capture expectations.¹⁸ It is worth noting however that by conducting tests at the individual level, one can get around these issues. If the household's linear Euler equation is rejected, then the aggregate version surely must not hold. Secondly, the basic LCH/PIH may be true at the household level and yet the aggregate version may fail if there is a failure of any one of the conditions necessary to derive a linear aggregate Euler equation from linear household Euler equations.¹⁹

¹⁶If there are adjustment costs associated with upgrading durable goods, then shocks to income feed through to purchases with delays. Caballero (1993) finds that expenditures on durable goods do not follow the theoretically predicted first-order moving-average process and posits that this is due to non-convex costs of adjusting the stocks of these goods. If this is true, then the timing of these purchases is influenced by the cumulation of innovations to household wealth since the previous adjustment (which took all information known at that point into account). This can invalidate some tests because shocks to income cause delayed increases in income and consumption expenditures. The current test is robust to this specification. If the announcement of Social Security taxes contains news to the household about its permanent income, there is no reason for the expenditure which adjusts the household's stock of durables to occur at the same time as the expected income fluctuation; it is just as likely to occur in the months before or after.

¹⁷See, for example Hall (1978), Flavin (1981), Poterba (1988), Campbell and Mankiw (1989), Wilcox (1989), Caballero (1990b), Carroll (1992), Levenson (1993).

¹⁸See Wilcox (1989) and Nelson and Startz (1990) respectively. Expectations are measures as the projection of change in income on lagged variables. The estimates are sensitive to which lagged variables are used, the fit of these expectations equations are poor, and reverse regressions – expected consumption on income – suggest that the equation is misspecified. Poor fit in the expectations equation implies that the estimated coefficient on expected income growth is biased towards the coefficient which would be found on unexpected income growth which should be positive.

¹⁹For example, the presence of overlapping generation, heterogeneous information sets, different intertemporal substitution across individuals of different wealth levels, and the construction of aggregate

Empirical studies of microeconomic consumption behavior yield more mixed results. The main methodology for testing the basic LCH/PIH at the household level was pioneered by Hall and Mishkin (1982). The strategy consists of using the Panel Study of Income Dynamics (PSID) to examine whether consumption responds to expected changes in income.²⁰ There are two main weaknesses with this line of attack. First, the only measure of consumption in the PSID is “usual weekly” food consumption, which is likely to have a lower elasticity than many other types of consumption since a certain amount of food consumption is of necessities. The specific question used in the PSID survey does not clearly delineate whether the “usual week” is usual for the time of the interview or for the previous calendar year, to which many questions (and the previous survey question) refer. The methodology has been extended by Altonji and Siow (1987), Zeldes (1989a), Runkle (1991) and Japelli, Pischke and Souleles (1995), and there is no consensus among these papers as to whether the linear Euler equation is rejected by the data.

In response to the problems with food consumption, recent research has turned to using the Consumer Expenditure Survey (CEX) which contains much better information on household-level consumption. Using a similar technique, Attanasio and Weber (1995) find evidence that the linear Euler equation is violated and evidence that it is not.²¹ Lusardi (forthcoming) uses a two-sample instrumental variables procedure to match the higher-quality income data from the PSID with consumption data from the CEX. This yields a more powerful test than use of either dataset alone, and she rejects the Martingale hypothesis.

However, Hall-Mishkin type tests still suffer from a second set of weaknesses: it is difficult to construct quality measures of expected income growth. Expected changes in

price indexes and aggregate consumption all stand as reasonable candidates. For some of these arguments see Wilcox (1992), Deaton (1992), Pischke (1995), and Attanasio and Weber (1995).

²⁰Hall and Mishkin (1982) test more than just this, but this correlation was the main violation they uncovered.

²¹Attanasio and Weber (1995) in Table 4 show that expected income growth predicts food consumption growth. The paper shows that when expected change in other nondurable consumption is included as an independent variable, it has a positive coefficient and the coefficient on expected labor income remains positive but becomes statistically insignificant. In contrast to the authors, I interpret this as demonstrating that predictable movements in consumption are correlated with predictable movements in income.

income are measured by predicting the change in income using predetermined variables such as education, occupation, industry, and age.²² Thus, in part, these studies are examining whether the age-profile of consumption tracks the age-profile of income. If there are permanent unexpected shocks to the age profile of earnings, then consumption should track these changes. The outlined procedure would include these changes in its measure of expected income change, and thus incorrectly reject the linear Euler equation. Further, there is the possibility that preference parameters in part determine which occupation, education and industry groups individuals select into, and these same parameters also determine the growth rate of consumption. This again could lead to spurious rejection. For example, high education groups tend to have steeper income growth. If patient people get higher levels of education, then because patient people have higher consumption growth, consumption will seem to track expected changes in income.

Finally, most individual income fluctuation comes from idiosyncratic sources which are not predictable from information about an individual's age, occupation, education, etc. While the instruments employed in these studies do a reasonable job of predicting labor income, the studies employ a large set of instruments.²³ Instrumental variables estimation with weak instruments can have large finite-sample biases even in large finite samples.²⁴

In response to this second set of weaknesses, several studies have used natural experiments to identify households which experience expected fluctuations in income. These studies generally find economically significant consumption responses which are only borderline statistically significant. Shea (1995) matches members of the PSID to publicly observable union contracts. Contracts cover several years and include publicly-known provisions for wage growth. The paper finds that the elasticity of food consumption

²²Attanasio and Weber (1995) do not use industry, occupation, or education, but instead use age and lagged changes in consumption, family size, income and aggregate variables.

²³Lusardi (forthcoming) reports an R^2 of around 0.01 when predicting income in the PSID. Attanasio and Weber (1995) have an R^2 of 0.24 in their first stage prediction of labor income using 21 instruments and with 288 observations on cohort data. Since the cohort techniques already effectively average the data, the true R^2 of the first stage regression is much smaller. The CEX is likely to have a slightly worse fit in the first stage due to the lower quality of the income data.

²⁴See Bound, Jaeger and Baker (1995).

with respect to these expected increases in wages is around one.²⁵ One question this study cannot address is whether the linear Euler equation fails only for this particular population. Souleles (1995) uses the CEX to see whether consumption increases when households receive their Federal income tax refunds, relative to when they first sent in their tax returns. The results suggest that consumers spend roughly 40% of their refunds in the quarter in which they receive them. Finally, Shapiro and Slemrod (1995) use survey methods to ask whether people saved or consumed the extra income from a pre-announced change in Federal income tax withholding in 1992. They find that consumption increases by 40% of the amount of additional take-home pay.

To conclude, the evidence that expected income growth predicts consumption growth is building, but is not conclusive. Deaton (1992), in his survey book on consumption, summarizes the micro evidence against the basic PIH as, “weaker and less transparent than in the aggregate data.”²⁶

2.4 The Test of the Basic LCH/PIH

As discussed in the introduction to this chapter, I test the linear Euler equation by examining whether household consumption rises when any individual earner in a household undergoes a change in his or her Social Security tax rate. I employ the following empirical specification of the linear Euler equation:

$$\Delta \ln(c_{ht+1}) = \alpha_1 z_h + \alpha_2 \Delta \tau_{ht}^{ss} + \alpha_3 m_t + \alpha_4 y_t + \epsilon_{ht+1} \quad (2.10)$$

where h indexes households; m is a complete set of month dummies; y is a complete set of year dummies less one; z contains a second-order polynomial in family size and a fourth-order polynomial in age that capture the fact that household consumption is

²⁵This finding is statistically significant at the 10% level for the entire sample, and at the 5% level for some subsamples.

²⁶Deaton (1992) p. 160.

generally not flat over the life cycle.²⁷ Finally, $\Delta\tau^{ss}$ is the change in the fraction of household after-tax income which is paid in Social Security taxes. According to the basic LCH/PIH α_2 should equal zero.

Three points about equation (2.10) are worth emphasizing. First, estimating the equation in first differences removes any household-specific effects. If these effects were not removed, a negative relationship between Social Security taxes and consumption might simply be due to the fact that only high-income households hit the cap and high-income households are more likely to have high levels of consumption.

Second, the regression includes a complete set of month effects.²⁸ Without the month dummies, it would be possible that the rise in consumption which occurs at the end of the calendar year would be attributed to the tax variable, which, on average, falls at the end of the calendar year. The month dummies however should completely absorb this 'Christmas' effect.

Third, in addition to the month dummies, time dummies for years are included in the regression. This allows the possibility that interest rates and other macro shocks affect families' consumption growth rates.

Equation (2.10) then has three sources of variation which identify the effect of the changing tax rates on consumption. First, high-income households hit the Social Security tax cap in different months, and τ^{ss} falls and after-tax income rises.²⁹ There is no variation along this dimension between December and January when all households who hit the cap begin to pay Social Security payroll taxes again. Second, for different households, Social Security taxes represent different amounts of after-tax income. For example, a household with two earners in which one earner hits the tax cap undergoes a smaller change in its after-tax income than a household with only one earner who hits

²⁷The age polynomial and the time dummies will also pick up cohort effects. This presents no problem for inference about α_2 but does imply that the year and time dummies should not be interpreted structurally.

²⁸Since, as is discussed in the data section, individual observations are actually overlapping three-month periods so that the 12 month dummies actually each represent a three-month period.

²⁹After 1992, the tax caps for OASDI and HI differ and so individuals can hit both caps at different times in a year.

the cap.³⁰ Finally, there are small changes in the Social Security tax rate across calendar years, which, like the tax caps, are public knowledge well in advance of their becoming effective. Thus all households sometimes experience small changes in after-tax income between December and January.³¹

2.5 The Consumer Expenditure Survey

I use data from the Family, Member, and Detailed Expenditure files of the Consumer Expenditure Surveys (CEX) for the years 1980 to 1993. The CEX is structured as a rotating panel of households. Each household is interviewed 5 times. In the first interview the CEX procedures are explained to the members of the household and they are asked to keep track of their expenditures for future interviews. Demographic information is collected and a population weight assigned to the household. This weight is the only data released from the first interview. Each household is subsequently interviewed four more times, once every three months. In each of these interviews detailed consumption expenditure information is collected on the past 3 months' expenditures. In each family's second and fifth interviews, a more detailed set of demographic and income information is collected. In these interviews, the family reports its pre-tax and post-tax income over the previous 12 months. In the fifth interview, each household is asked about its holdings of several categories of liquid assets and how much these holdings have changed over the past 12 months (i.e. since immediately prior to the start of their consumption reporting).

Families rotate into and out of the Survey, so that new households are being added every month. The Survey interviews about 1,500 households each month, and only about half of households contribute a complete one-year panel of four consumption observations each covering three months. I extract and merge data on families and individual members to make an unbalanced, overlapping panel of households covering January, 1980 to

³⁰As discussed in more detail in the next section, the payroll tax rate is adjusted from the percent of gross labor income to the percent of after-tax family income using information reported in the first interview about the retrospective 12 months.

³¹In Section 6 of this chapter, I explore the consequences of eliminating this last source of variation.

February, 1993. I drop any family which is missing the second interview reports of family size, age of reference person, or age of spouse.³² Households are also dropped if before-tax household income or after-tax household income in the second interview is topcoded, incomplete or missing. Differenced observations are dropped if nondurable consumption changes by more than 100%.³³

I construct a measure of nondurable consumption by summing expenditures on the following categories of goods: food, excluding food as pay and school meals; alcohol; house-furnishings and equipment excluding furniture, major appliances, and floor coverings; apparel; services; transportation excluding new and used vehicle spending and financing; entertainment; personal care; reading; and tobacco and smoking.

I calculate the share of after-tax income paid in Social Security taxes, $\Delta\tau^{ss}$, as the first difference of

$$\tau^{ss} = \sum_{i=1}^2 (1 - \tau_h) \sum_{s=1}^3 \left(\frac{Y_{is}}{Y_h} \right) \tau_{ss} D_{is}^{cap} \quad (2.11)$$

where i indexes individuals and s months in an interview period; τ_h is the average household tax rate; Y_i and Y_h are individual and household pre-tax income respectively; and τ_{ss} is the statutory Social Security tax rate in that month. D^{cap} is a variable which equals 1 if the individual has not hit the tax cap; a fraction representing the fraction of the month the individual pays Social Security if the individual hits the tax cap during the month; and equals 0 otherwise. All variables except D^{cap} are calculated from information given in the second interview pertaining to the 12 months prior to the second interview. To construct the best possible measure of when an individual hits the tax cap, D^{cap} is calculated from the retrospective labor income variable from the final interview. Consequently, D^{cap} uses information not available to the household when consumption decisions are made. Although it is unlikely that the month in which an individual hits the cap is systematically correlated with income shocks during the survey, it is technically

³²The last is conditional on having a spouse. The reference person in the CEX is defined as the person who is most responsible for paying the bills. This person is then interviewed, as he or she is probably the most knowledgeable about household expenditures.

³³I use other cuts when using other dependent variables. The results do not change much if outliers are not excluded. I am more precise on this point in the next section.

Table 2.2: SAMPLE STATISTICS

	MEAN	STANDARD DEVIATION	PERCENT OF TOTAL CONS.
AVERAGE MONTHLY LEVELS			
TOTAL CONSUMPTION	1,440	1272	100
NODURABLE CONSUMPTION	821	666	57
NON-GIFT CONSUMPTION			
TOTAL	1405	1243	98
NONDURABLE	796	641	55
FOOD+ALCOHOL	324	233	23
APPAREL+SERVICES	115	166	8
ENTERTAIN.+PERSONAL CARE	143	272	10
FAMILY SIZE	2.64	1.54	
AGE	46.5	17.5	
AFTER-TAX INCOME	1,972	1,626	137
BEFORE-TAX INCOME	2,167	1,811	181
AVERAGE FIRST DIFFERENCES			
TOTAL CONSUMPTION	0.008	0.501	
NONDURABLE CONSUMPTION	-0.003	0.381	
NON-GIFT NONDUR. CONS.	-0.004	0.380	
$\Delta\tau^{ss}$	0.000	0.005	
SUBSAMPLE: N=2778			
HEAD HITS SOC.SEC. CAP	MEAN	S.D.	MAX
MONTHS SOC. SEC. COVERS	9.62	1.76	
HEAD LABOR INCOME	53,169	16,089	
$\Delta\tau^{ss}$	0.022	0.015	0.10

Based on sample for non-gift nondurable consumption regressions. Total consumption excludes expenditures on mortgages, health care, pensions, education, and cash contributions. See Appendix A for details.

possible. Thus I check the robustness of the results to calculating D^{cap} based on the individual salary report from the second interview instead.

Finally, τ^{ss} is set to zero for any individual who might not be paying Social Security taxes. I do this based on an individual's employment history, occupation, industry, reported Social Security contributions, and retirement plan payments. More details are provided in the next section and in Appendix A.

Table 2.2 presents some summary statistics on the sample. There are 148,831 total observations (NT) on 63,527 total households ($\equiv N$) in the sample. 36,108 households contribute a full 3 differenced observations; 13,088 contribute only 2 differenced obser-

Table 2.3: THE REACTION OF NONDURABLE CONSUMPTION

DEPENDENT VARIABLE:	NT	OLS	FGLS
% Δ NONDURABLE CONS.	148,831	-0.648 (0.199)	-0.572 (0.186)
% Δ NONDURABLE NON-GIFT CONS.	148,921	-0.527 (0.210)	-0.467 (0.184)

variations.

2.6 Estimation and Results

I begin by estimating equation (2.10) by ordinary least squares (*OLS*). Due largely to measurement error in the level of consumption, observations are serially correlated within households. The first-order serial correlation is -0.4 while higher orders are insignificantly different from zero.³⁴ I construct standard errors which are consistent in the face of arbitrary heteroskedasticity and within-household serial correlation.³⁵

The first row of Table 2.3 reports the results of *OLS* estimation of equation (2.10). The point estimate implies that when after-tax income falls by one percent due to an increase in the Social Security tax rate, nondurable consumption falls by 65%. The coefficient is statistically significant at the 99% level. The last column reports the results of using a Feasible Generalized Least Squares estimator, which is consistent and has the advantage of being efficient if the covariance matrix is correctly specified. I estimate

³⁴With no noise in the consumption data, time aggregation implies that change in consumption should follow an *MA*(1) with a positive coefficient of 0.25. Under the assumptions that consumption is truly a random walk and the measurement error is classical, observed first-order serial correlation of -0.4 implies that 86 percent of the variance in observed change in consumption is due to measurement error. Given that these are the best consumption data available, it is not surprising that conclusive tests of the linear Euler equation are rare.

³⁵Standard errors are

$$(X'X)^{-1} \left(\sum_{n=1}^N X_n' e_n e_n' X_n \right) (X'X)^{-1}$$

where X is the full $NT \times K$ matrix of data, X_n is the $T_{N_n} \times K$ matrix of data for each household, and e_n is the vector of *OLS* residuals for that household. These errors are also robust to household-specific random effects in equation (2.10).

an unparameterized covariance matrix of error terms which applies to all households: $\hat{\Omega}_{ij} \equiv N_{ij}^{-1} \sum_{n=1}^N e_{ni}e_{nj}$ where N_{ij} is the number of individuals for whom one observes both the i^{th} and j^{th} difference ($i, j \in (1, 2, 3)$), and e_{ni} is household n 's i^{th} OLS residual. The point estimate is slightly lower but also highly statistically significant.

Despite the month dummies, one might be concerned that the seasonal in consumption, caused in part by Christmas, differs across households in the same way that the Social Security tax variable does. In the CEX expenditure files, every expenditure which is a gift for someone not living in the household is flagged. The second row of Table 2.3 reports the coefficient on the Social Security tax variable when all such gift purchases are excluded from the definition of consumption. The point estimates decline slightly and remain highly significant, implying a half percent decline in nondurable consumption for every percent decline in after-tax income. The remainder of the chapter excludes gift purchases from all measures of consumption.³⁶

Another possible weakness of the results so far is that measurement error in the Social Security tax variable could be *attenuating* the estimated coefficients. The income data are not of the quality of the income data in the PSID, for example.³⁷ However, in addition to the income data employed so far, which is the individual's reported earnings over the past 12 months, each individual is also asked the amount of their last paycheck and the length of the pay period.³⁸ From this information I construct another measure of annual income and then $\Delta\tau^{**}$. The first column of Table 2.4 shows the results of FGLS estimation when the average of these two measures is used. The estimated coefficient does indeed rise slightly, consistent with there being measurement error in both income reports.³⁹

³⁶An alternative specification one might consider is to perform the analysis in levels. The results of this experiment lead to even greater statistical significance, but much smaller economic effects. Another question one might ask is whether the results in the logarithmic specification change when outliers are not dropped. Including outliers led to standard errors of around 0.22 for both estimation techniques and a slight decline in coefficients so that the four averaged -0.399 .

³⁷See Lusardi (forthcoming).

³⁸This income measure is similar to that in the PSID and is likely to be superior to the main variable which is most commonly used.

³⁹I do not instrument one measure with the other since classical measurement error in income generates non-classical measurement error in the Social Security tax variable which in turn causes two-stage least

Table 2.4: ALTERNATIVE MEASURES OF THE SOCIAL SECURITY TAX VARIABLE

INDEPENDENT VARIABLE: <i>NT</i> = 148,921	AVERAGE OF MEASURES	INTERVIEW 2 MEASURE	AVERAGE OF INT. 2 MEASURES
COEFFICIENT	-0.558	-0.473	-0.540
STANDARD ERROR	(0.201)	(0.194)	(0.206)

FGLS estimation on non-gift, nondurable consumption.

As mentioned in the previous section, the construction of $\Delta\tau^{ss}$ uses information from the end of each household's tenure in the survey. This technically violates the information structure of the test, although it seems unlikely to have generated the results so far. To confirm this, I calculate another measure of the Social Security tax rate from the retrospective annual income report in the second interview.⁴⁰ The results are reported in the middle column in Table 2.4. Finally, I average both Social Security tax measures from the second interview information and estimate equation (2.10) again. The final column of Table 2.4 is again consistent with the presence of measurement error and provides further confirmation that consumption falls by half a percent for every percent increase in the effective net Social Security tax rate.

While the previous results would seem a strong rejection of the linear Euler equation, I take steps to eliminate possible alternative interpretations. Why might one be concerned with the results so far? First, note that the fit of the regressions is small: the R^2 are just less than 1%. This is due, in part, to the fact that the Social Security tax changes are small relative to the swings in consumption (true and due to measurement error).⁴¹ Second, the sample size is very large. Under these conditions, only a small correlation between the error term and the independent variable could perhaps cause the significant results. Note, however, that such a story is likely to be at odds with the fact that a better measure of $\Delta\tau^{ss}$ increases the estimated coefficient. Further, one would have to argue

squares estimates to be inconsistent.

⁴⁰This is the *first* interview in which consumption data are collected.

⁴¹There is no problem of the sort encountered in weak instrumental variables regressions. The point estimates are unbiased and the standard errors are correct for making statistical inference.

Table 2.5: HIGH CONSUMPTION SUBSAMPLE AND TREATMENT VS. CONTROL GROUP

DEPENDENT VARIABLE:	NT	MAIN MEASURE	AVERAGE OF MEASURES	AVERAGE OF INT. 2 MEAS.
HIGH CONS. SUBSAMPLE	18,108	-0.539 (0.293)	-0.506 (0.327)	-0.259 (0.331)
TREATMENT GROUP VS. CONTROL GROUP	148,921	-0.310 (0.552)	-0.575 (0.664)	-0.680 (0.666)

FGLS estimation on non-gift, nondurable consumption.

that it is pure coincidence that the coefficient estimate is negative and of a reasonable magnitude.

Nevertheless, two scenarios seem possible. First, suppose that households for which the small increases in Social Security tax rates represent the largest fraction of after-tax income are also those households with the largest consumption declines between December and January. In this unlikely event, the change in the Social Security tax variable will be spuriously correlated with the change in consumption. Second, perhaps there are macroeconomic effects contemporaneous with the January increases in the Social Security tax rate which cause consumption to grow less in the six years of the increases than in the seven years without. I take two tacks to eliminate these possibilities.

First, I restrict the sample to households in which average consumption is over \$30,000 (1987) dollars a year. This makes the sample more homogeneous in terms of income and consumption levels. It also decreases the sample size by an order of magnitude, and increases the percent of households who hit the Social Security tax cap. The standard deviation of $\Delta\tau^{ss}$ in this subsample is 1.0%, double that in the entire sample. If the skeptical view is correct, the estimated parameter should decline in this sample. The first row of Table 2.5 shows that this is generally not the case. The standard errors increase, but, apart from statistical significance, the story is the same.

The second tack I take is more elaborate. I employ a control group of identical households who do not pay Social Security taxes to further eliminate the alternative explanations raised above. That is, I identify the coefficient of interest using only the

variation across otherwise identical households that do and do not pay Social Security taxes.

I assign each earner to one of three groups: a treatment group comprised of those individuals who are almost certainly subject to Social Security tax withholding; a control group comprised of those individuals who almost certainly are not; and a “neither” group comprised of the remaining individuals.⁴² Assignment is made on the basis of five sets of information from individuals’ fifth interviews.⁴³ First, the CEX contains information on both occupation and industry of the reference person and spouse. Individuals who are self-employed or, prior to 1984, Federal government employees, are assigned to the control group.⁴⁴ Second, individuals are asked to estimate their contributions to Social Security over the past 12 months. Starting in 1982, individuals also are asked whether Social Security is normally deducted from their paycheck, and, starting in 1986 whether this deduction covers Medicare only. These variables are used to assign a second set of individuals to groups. Third, individuals who change jobs or work more than one job may overpay Social Security taxes and thus I cannot calculate when they would hit the tax cap. The CEX asks whether households overpaid Social Security in the last 12 months, and any individual in such a household is assigned to the neither group. Fourth, if an individual spends time unemployed I cannot be sure they are employed during any tax change nor can I accurately predict when they will hit the tax cap. Thus any individual reporting fewer than 50 weeks worked in the past year is put into the neither group. Finally, I use family reported Railroad Retirement contributions, which were merged with the Social Security system in 1985, to make a further set of assignments. Appendix A give a complete description of the assignment of individuals to treatment, control, and neither groups. In sum, I have a high degree of confidence that anyone in the treatment group is employed and paying Social Security taxes and that my calculation of when they

⁴²All previous results only employ the treatment group.

⁴³If the final interview is missing, I use second interview information.

⁴⁴Self-employed individuals pay Social Security taxes, but they do so as part of their paying quarterly estimated tax payments. Thus there is no withholding pattern which directly affects after-tax income. After 1984, some Federal government employees were covered by Social Security instead of Federal government retirement plans.

Table 2.6: TREATMENT, CONTROL AND NEITHER GROUP STATISTICS

TOTAL SAMPLE	TOTAL NUMBER OF INDIVIDUALS	NUMBER IN TREATMENT GROUP	NUMBER IN CONTROL GROUP	NUMBER IN NEITHER GROUP
HEADS	63,527	29,038	4,172	30,317
SPOUSES	63,527	40,077	1,412	22,038
SUBSAMPLE:				
HOUSEHOLDS				
HITTING SS CAP				
WHO HITS CAP:				
HEAD	3,622	2,778	274	570
SPOUSE	148	96	16	36
BOTH	35			
HEAD		22	9	5
SPOUSE		25	5	5

might hit the tax cap is correct. Table 2.6 presents the breakdown of individuals among the three groups.

I generate a hypothetical Social Security tax variable for every individual and construct two additional variables: $\Delta\tau^{ssE}$, the change in the Social Security tax rate that every household would have experienced had it been paying Social Security Taxes; and $\Delta\tau^{ssN}$, the hypothetical change in the Social Security tax rate for households that are in the “neither” group. I then estimate

$$\Delta \ln(c_{it+1}) = \alpha_1 z_i + \alpha_2 \Delta\tau_i^{ss} + \alpha_3 \Delta\tau_i^{ssN} + \alpha_4 \Delta\tau_i^{ssE} + \alpha_5 m_t + \alpha_6 y_t + \epsilon_{t+1}. \quad (2.12)$$

The significance of α_2 again provides a test of the basic LCH/PIH; however the identification comes only from the difference between the consumption response of the control group and that of the treatment group.

The bottom row of Table 2.5 displays the results of estimating equation (2.12), that is, including the Social Security tax variable for each group. The point estimates are in

line with the earlier results although they are now statistically insignificant. The large standard errors are consistent with the fact that the control group is quite small. If it were simply the case that, due to aggregate factors, consumption happened to decrease in Januaries in which the Social Security tax was raised, then one would expect these point estimates to return to zero. In fact consumption falls for those households who are paying Social Security taxes and not for those that are not.

The stability of both sets of point estimates in Table 2.5 to these alternative identification strategies suggests that the significant coefficients found for the previous regressions are not driven by odd cross-household seasonal patterns in consumption. Nor are the results coming solely from a correlation between the December to January change in consumption and increases in the Social Security tax rate.

2.7 Evaluating Alternative Hypotheses

The previous section demonstrates that household consumption responds to pre-announced changes in Social Security tax rates. This constitutes a rejection of the basic LCH/PIH and the linear Euler equation. In this section, I investigate the causes of this rejection. I do this by investigating both whether the violations of the linear Euler equation are strongest among low-asset households and which subcategories of consumption show the largest violations.

If the linear Euler equation fails due to liquidity constraints or a correlation between the expected variance of consumption and the growth rate of consumption, then the relationship between expected income growth and consumption growth should be strongest for those households with few liquid assets. Unfortunately, asset information in the CEX is quite limited and is often topcoded or missing. Further, the CEX only asks asset information in the final interview. However, this interview contains information about 4 categories of liquid assets and about how much the wealth in each category has changed over the past 12 months. From these sets of information, I construct a measure of the value of checking accounts, savings accounts, stock and mutual fund holdings and bonds

Table 2.7: RESULTS FOR ASSET AND AGE SPLITS

INTERACTION VARIABLE	NT	$\Delta\tau^{ss}$	$\Delta\tau^{ss}$ INTERACTION
HIGH ASSET RATIO	74,909	-0.246 (0.271)	-0.670 (0.396)
YOUNG AGE	131,073	-0.392 (0.288)	-0.185 (0.393)
RESULTS FOR TREATMENT VS. CONTROL			
HIGH ASSET RATIO	74,909	-2.441 (0.932)	3.525 (1.329)
YOUNG AGE	131,073	-0.439 (0.962)	-0.617 (1.322)

FGLS estimation on non-gift, nondurable consumption.

as of immediately before each household's first consumption observation.⁴⁵ The ratio of this variable to average annual household consumption yields a measure of how likely a household is to be liquidity constrained or, nearly equivalently, whether consumption is likely to track income due to the changes in the expected variance of consumption. I re-estimate the linear Euler equation adding an interaction between the Social Security tax variable and a dummy variable which is set to one if the household has more than enough assets to finance a half a year of consumption or if any asset category is topcoded.

I find little evidence that the Euler equation failure is concentrated among households with the fewest assets. Table 2.7 reports the results of FGLS estimation of nondurable consumption on the average measure of $\Delta\tau^{ss}$ and its interaction with the high-assets dummy variable. The first row of results suggests that households with the ability to finance more than half a year of consumption from assets seem to react more to the changes in Social Security tax rates than do low-asset individuals, although the result is not quite statistically significant at the 95% level. Row 3 of Table 2.7 shows the same regression but identifying the coefficients using the difference between treatment and control group reactions; the results are reversed. The significant point estimates suggest

⁴⁵Details are provided in Appendix A.

that low-asset households increase nondurable consumption in excess of two percent for every percent change in take-home pay, while high-asset households decrease consumption by around one percent. The size of these estimates and the reversal of sign makes it difficult to take much from this experiment.⁴⁶ It may well be that since the asset data is retrospective, the significant amount of noise in the data is correlated in some way with the growth rate of consumption— or with households' increasing or decreasing fortunes— and it is this aspect of the data which the coefficients reflect.⁴⁷

As another tack, I use age to proxy for the probability of being liquidity constrained or having few assets.⁴⁸ Young households have larger expected income growth and fewer assets than older households. Due to either liquidity constraints or the optimal choice not to borrow (as discussed in subsection 2.2.3), young households may be more likely to violate the linear Euler equation. Rows 2 and 4 report a similar pair of regressions to those for the asset classification but based on whether a household is younger than 43.⁴⁹ In Chapter 3, I estimate that the typical household moves from “buffer-stock” type behavior to behavior more consistent with that of the basic LCH/PIH around age 43. Rows 2 and 4 of Table 2.7 reports the results of these regressions. While the coefficients are not statistically significant, they are consistent with the hypothesis that younger households react more strongly to changes in expected income.⁵⁰

The second hypothesis I seek to evaluate is that households allow consumption to track income provided that this leads to only small deviations in marginal utility from the fully-optimal plan. As discussed in subsection 2.2.3, this hypothesis implies that expenditures should respond to expected income for those goods which are somewhat durable and yet purchased out of cash-on-hand and not separately financed. Thus this

⁴⁶Interacting the ratio itself leads to similar conclusions.

⁴⁷Since I discard individuals who are not fully employed during the previous 12 months, the sample I examine has fewer candidates to be liquidity constrained than the population which is typically studied.

⁴⁸For arguments on why this might be see for example Chapter 3. Also, Japelli et al. (1995) find that age is a significant predictor of whether a household reports that it is liquidity constrained.

⁴⁹I also drop households older than 70.

⁵⁰It should be noted that, for everyone except those who hit the tax cap, the changes in Social Security taxes are always increases. In the face of liquidity constraints or buffer-stock behavior, consumers do a much better job of smoothing consumption across expected declines in income.

Table 2.8: THE REACTION ACROSS CATEGORIES OF CONSUMPTION

DEPENDENT VARIABLE CONSUMPTION CATEGORY	NT	AVERAGE OF MEASURES	AVERAGE OF INT. 2 MEASURES
TOTAL	142,924	-0.250 (0.206)	-0.286 (0.210)
FOOD AND ALCOHOL	145,985	0.534 (0.203)	-0.084 (0.209)
ENTERTAINMENT AND PERSONAL CARE	143,238	-0.896 (0.404)	-0.850 (0.413)
APPAREL AND SERVICES	123,278	-3.004 (0.527)	-2.571 (0.540)

FGLS estimation on non-gift consumption.

theory predicts that households simply time the purchase of items such as clothing, nice dinners out, or small electronic items to coincide with months in which they have more cash on hand. Table 2.8 shows the results of estimating equation (2.10) across subcategories of goods.

The elasticity of total expenditures to expected changes in take-home pay is estimated as 0.25. While the standard error is similar to those found in the nondurables regressions, the lower point estimate implies that the response is statistically insignificant. Nondurable consumption is reacting more than durable consumption. One interpretation of this finding is that there are adjustment costs in the sale and purchase of durables and therefore they do not react as quickly or as much as nondurables. An alternative explanation is that durable goods are often purchased with credit, as in an automotive loan, so that there is no reason for expenditures on durable goods to be linked to fluctuations in income as expenditure on nondurables are.

Rows 2 through 4 of Table 2.8 show the results for estimation involving three subcategories of consumption. First, note that food consumption does not seem to respond at all to expected changes in income. This is not consistent with the alternative the-

ory at hand, since dinners out and alcohol purchases can both be easily timed to track income.⁵¹ However, food consumption also consists of purchases of necessities which are relatively inelastic. Apparel and services consumption reacts the most of all categories, with an elasticity of between 2.5 and 3. Since apparel and services make up 8% of total consumption, these estimates suggest that for every extra expected dollar of income which a Social Security tax change induces, 24 cents are spent on apparel and services. Finally Entertainment and Personal Care, which also includes expenditures on reading materials and tobacco and smoking supplies, shows a just slightly stronger reaction than total nondurable consumption. The pattern and magnitude of coefficients is similar in treatment and control group regressions, although the standard errors increase on a scale almost exactly as that for nondurable consumption.

In sum, there is some evidence in favor of the hypothesis that subcategories of consumption which respond most are those in which expenditures can be most easily substituted across short intervals of time.

2.8 Conclusion

This chapter finds that consumers do not perfectly smooth consumption at high frequencies. Contrary to the basic LCH/PIH, consumption reacts to predictable changes in tax rates. The findings are economically significant: consumption falls roughly half as much as the expected decline in income.

Three alternative theories to the linear Euler equation are entertained, each of which can explain why the basic PIH/LCH fails. First, when one models individual uncertainty explicitly, the optimal path of consumption generally will not follow a Martingale. Consumption would track expected income if, due to impatience and expected income growth, consumers were often in situations in which they would like to consume out of future income but do not, for fear of bad shocks to their future incomes causing extremely

⁵¹It is interesting to note that this is the only subcategory of consumption which can be used by PSID-based studies.

low levels of utility – the buffer-stock theory of consumption. A second theory with similar implications is that consumers face liquidity constraints. Thus, some fraction of households will have consumption increases track expected increases in income. The common prediction of these theories is that low-asset households should account for Euler equation violations. Tests to evaluate this prediction yield inconclusive results, although some scant evidence is found that young households, which generally have lower assets and higher consumption growth, are more likely to have consumption track income.

The second theory which this chapter considers posits that at high frequencies, many goods are somewhat “durable” and purchases can be substituted across time with little loss in utility. Near-rational households may simply choose to make more of these slightly durable purchases when income is higher. In addition to being able to explain the rejection of the basic LCH/PIH, this theory is consistent with the evidence from subcategories of consumption expenditures. I find the strongest violations of consumption smoothing in subcategories of consumption in which households can easily substitute purchases across time and in which purchases are generally made with cash or very short-term credit.

Whichever interpretation one chooses to take, the mounting evidence against consumption smoothing has far-reaching implications. First, if expected changes in tax rates influence contemporaneous consumption behavior, then fiscal stabilization, such as that undertaken by President Bush in 1993 or that provided by automatic fiscal stabilizers, may have large and important effects on consumption. Second, while my results do not imply that consumption smoothing is a poor approximation over long horizons, they do imply that when studying shorter horizons, such as those addressed in business cycle models, the linear Euler equation is not the correct structural equation to employ.

Appendix A

The Consumer Expenditure Survey

I use the *CEX* family, member and detailed expenditure files for years 1980 to 1993, as kindly provided by the National Bureau of Economic Research. Most information about the *CEX* is obtained from Bureau Of Labor Statistics (1980-1993) and conversations with Bureau of Labor Statistics (*BLS*) statisticians. Households should not be matched across 1985 to 1986, and are not. Care is taken to assure consistency in the data despite variable classification changes through time, and across reference person and spouse. Information was kindly provided by the Division of the *CEX* in the *BLS* about various issues including the matching of occupation codes from 1980 – 81 to later years.

As discussed in the text, households are discarded if they are missing any of the information necessary for the regressions, or if they are classified as incomplete income reporters or if family income is topcoded or missing in the second interview. Age is the average of both head and spouse if there is a spouse, otherwise it is the head's. Due to some extreme reports, I reset reported tax rates above 66% to 25%, and similarly for below 10%. The results are reasonably insensitive to this correction.

I assign individuals to treatment, control and neither groups as described in the following paragraph. At the end of the procedure, I move any individual assigned to both the treatment and the control group into the neither group.¹

¹The procedure never assigns an individual into the neither group and one of the other groups.

An individual is assigned to the control group if he or she reports his or her industry as Federal government employee prior to 1984 or reports his or her occupation as self-employed. I then create a government category which contains all remaining individuals who report their industry as any level of government employee. Further, all individuals in any family that reports paying into a government retirement account and that has no members already assigned to the government groups are moved into the government group. Next any individual in the government group which reports paying Social Security in their normal paycheck or during the last 12 months is assigned to the control group. Those who report not paying Social Security in either of these questions are put into the treatment group.

Next, any individual who is missing industry or occupation data is put into the treatment group unless they report not paying Social Security taxes in either of these variables. Households which report overpaying Social Security in the past 12 months are put into neither. Individuals who report working less than 50 weeks in the past 12 months are put into neither. Finally, if a household reports paying into a Railroad retirement account, if there is only one earner, that earner is put into the control group prior to 1985 and into the treatment group after 1985. If the family has multiple earners, all earners are assigned to the neither group.

Consumption data is compiled from the detailed expenditure files. I first calculate monthly expenditures and then average them to get consumption at a monthly rate. The denominator for the average is the number of nonzero consumption months for nondurable consumption. Thus for example, if no apparel expenditures are reported but nondurable consumption is positive, this is considered a valid month of data on apparel purchases. If nondurable purchases are zero, then there is assumed to be no data for this month. Some consumption observations are reported to have occurred prior to the the three months recall period for an interview and some in a later month. *BLS* statisticians recommend treating these expenditures as if they occurred in the reported month. Households with only 1 or 2 months of consumption data (one percent of the sample) were dropped. Households with more than 12 months of data have the last few observations dropped unless the first observation is less than the 13th in which case the first observation was dropped and the next 12 used.

Total consumption is defined as total expenditures less outlays for mortgage payments,

education, health care, pensions, and cash contributions. Food expenditures are all expenditures on food and alcohol less food as pay and school meals. Apparel and Services is the sum of these two *CEX* categories. Entertainment and Personal Care is the sum of Entertainment, Personal Care, Reading, and Tobacco and Smoking expenditures. Nondurable consumption is defined in the main text.

Observations with large changes in consumption between adjacent 3-month periods are dropped. For total, food, and nondurable consumption, observations which had changes in excess of 100% were dropped. For both Apparel and Services and for Entertainment and Personal Care observations were dropped only if the change in consumption exceeded 250%. All cuts are just beyond two standard deviations in the ex ante data.

The main measure of income is constructed by extrapolating using the second and fifth interview reports to make a best guess about income for any month. If income changes by more than 25% between the second and fifth reports, or if the second interview measure is missing, I simply use the fifth interview measure.

Chapter 3

Consumption Over the Lifecycle

with Pierre-Olivier Gourinchas

3.1 Introduction

This chapter analyzes consumption and savings behavior of households over their working lives, focusing on estimation of age-specific consumption functions and their implications. We are motivated by four observations.

First, household consumption and savings decisions are arguably among the most important determinants of economic growth and business cycles. Consumer expenditures account for two thirds of national output and a large percent of output fluctuations. The difference between consumption and income—savings—determines the stock of wealth, which in turn determines the interest rate, the level and perhaps the growth rate of output. To understand business cycles and growth, we must thus first understand household consumption behavior.

Second, better methodology, data, and creative use of natural experiments have led to more frequent and more convincing rejections of the most widely-used model of consumption behavior, the certainty equivalent life-cycle hypothesis (henceforth, CEQ LCH). At the individual level, empirical studies typically test the central implication of the CEQ LCH - the Martingale hypothesis - by testing whether consumption responds to *expected*

changes in income. Despite generally poor-quality individual data on consumption, the CEQ LCH is often rejected.¹ On aggregate data, the CEQ LCH is even more convincingly rejected.²

Third, consumption smoothing does not seem to be a good characterization of low frequency consumption movements and savings behavior. According to the 1992 Survey of Consumer Finances (SCF), only 15% of the respondents reported retirement as their primary motive for saving while 42% cited liquidity needs.³ Life-cycle savings seem to occur late in the working lives of consumers. Median holdings of very liquid assets for households under age 50 are \$3,900 while median holdings of non-housing non-business wealth are just under \$13,000.⁴

Lastly, recent theoretical work (Zeldes (1989b), Deaton (1991), and Carroll (1992)) demonstrates that the Martingale hypothesis and consumption smoothing can be bad approximations to consumer behavior when agents face large amounts of individual uncertainty. Carroll (1993a) and Hubbard et al. (1994) among others have shown that income uncertainty can generate a positive covariance between expected income changes and consumption at low and high frequencies through precautionary savings.

In this chapter, we examine individual consumption data, estimate a structural model with realistic levels of income uncertainty, and finally reinterpret life-cycle consumption and asset accumulation behavior within the context of the model. We measure, exploit and analyze the systematic age-pattern of consumption profiles. Age-heterogeneity in consumption behavior results from (a) the interaction between, and relative strengths of, retirement and precautionary motives for saving at different ages and (b) the changing

¹It should also be noted that many of these papers technically test the permanent income hypothesis rather than the LCH. The differences can be pronounced at the aggregate level, but both theories predict no response of consumption to expected changes in income at the household level.

²Good recent work includes Lusardi (forthcoming), Shea (1995), Souleles (1995), Levenson (1993), Carroll and Summers (1991), Campbell and Mankiw (1989) and Chapter 2. See Deaton (1992) for an excellent survey of the field or Dynan (1993) for a recent approach to estimating an Euler equation with a precautionary term.

³See Carroll (1993a) and Carroll (1992).

⁴As reported in Carroll and Samwick (1994). This does not include Social Security and pension wealth, which constitute a large fraction of wealth at retirement. Note that distribution of wealth and its constituents is always strongly skewed, with a mean far exceeding the median.

slope of the income profile. We are successful along several dimensions.

First, we provide new evidence on the failure of the Certainty Equivalent (CEQ) LCH at the microeconomic level. We do this by demonstrating that consumption age-profiles averaged across time and households are not flat and are related to expected profiles in income. Importantly, this result still holds after controlling for family composition and cohort effects — two potential reasons for the observed hump-shape of consumption. Our approach involves using the best available data on consumption expenditures in the US, the Consumer Expenditure Survey (CEX) from 1980 to 1993, giving us data on around 40,000 households. Using weak identifying assumptions, we construct consumption and income profiles across the working lives of “typical” men of five different educational attainments and five different occupational groupings. Consumption and income profiles are both hump-shaped, and consumption tracks income reasonably well early in life.

Second, embedding realistic income uncertainty into the canonical model of life-cycle consumer behavior substantially improves the fit of the predicted life-cycle consumption profile. To demonstrate this, we write down a consumer maximization problem with a retirement period and explicit income uncertainty. The solution to this model will be the standard CEQ LCH consumption rules when uncertainty is ignored. Consumption will depend on the interest rate, the intertemporal elasticity of substitution, the discount factor and the present discounted value of income. However, under uncertainty, consumption will also depend on the *path* of expected income. Thus consumer behavior will vary systematically over the life cycle. When expected income growth and the discount rate are low relative to the interest rate, consumers’ behavior will remain similar to that of standard life-cycle consumers. If, on the other hand, expected income growth or the discount rate are large relative to the interest rate, consumers will behave as “buffer-stock” agents, consuming roughly their income and saving only small amounts to buffer against bad income draws. As households age, income growth declines. Consequently, the retirement savings motive will enter consumers’ horizons. They will save more and behave more like certainty-equivalent consumers. The model can potentially deliver average consumption profiles which are more concave than income profiles. It is important to

emphasize that we do not assume our results here. With a sufficiently low discount rate, the average consumption profile would be very similar to that of the certainty equivalent case.

Positing that the average income profile for a given group corresponds to the expected income profile and incorporating calibrated individual-specific income shocks, we estimate consumption functions for consumers in each occupation and education group. By simulating the lives of many consumers using these consumption functions we create predicted average consumption profiles. We then estimate the discount rate and the intertemporal elasticity of substitution by a Method of Simulated Moments procedure. The average household has an intertemporal elasticity of substitution of 2.04 and a discount rate of 3.9%.⁵ The discount factor is tightly estimated, and the estimated discount rates decline weakly with educational attainment. It is worth stressing that the estimated coefficients are within a “reasonable” range. In particular, buffer-stock behavior arises early in life *due to the steepness of the income profiles at young ages*.

The model fits the data quite well and does an excellent job of capturing the main features of the consumption profiles. To the best of our knowledge, this represents the first structural estimation of consumption functions over the life cycle which incorporates precautionary savings.

Third, we find strikingly different consumption functions for households at different ages: consumers behave like “buffer-stock” consumers early in their working lives and more like CEQ LCH consumers as retirement nears. We show that *households make the transition from buffer-stock to LCH behavior just after age 42*. This confirms our initial intuition: relative movements of the consumption and income profiles reveal a great deal of information about the relative strength of the two savings motives. We conclude that a large fraction of consumers consists of *target savers*, for whom the Euler equation, as typically tested, should be expected to fail. This is, in part, a confirmation of Carroll (1993a) and (1993b) which argue, based on asset data, that buffer-stock models apply

⁵The latter estimate is very sensitive to the assumption about the real after tax interest rate, 3% in our benchmark analysis. We discuss robustness after presenting our results.

only to households before ages 45 to 50.

Fourth, this chapter contributes to the debate on the determinants of wealth accumulation. In our model, saving and consumption at each age are determined by the interactions between precautionary and retirement savings motives. Defining all wealth accumulation at retirement as life-cycle savings, we can decompose saving into precautionary and life-cycle saving.⁶ Early in life, households would like to borrow against expected future labor income. Consequently, life-cycle savings are negative. However, uncertainty causes households to build a buffer stock of savings, implying that precautionary saving is positive. Late in life, labor-income uncertainty is mostly resolved, and consumers run down their buffer stocks, while retirement saving becomes positive and large. Thus, the calibrated model matches the basic features of asset data. Both the fact that most households hold few, if any, assets and the fact that most households do not start saving for retirement until late in life have often been interpreted either as evidence against the LCH or evidence against forward-looking consumers. Our fitted model suggests that these facts are instead consistent with the LCH augmented to include income uncertainty and consistent with forward-looking optimizing behavior.

Our research builds on many previous studies of life-cycle consumption behavior.

Several papers have used micro-consumption and income data to construct life-cycle profiles of consumption and income. Kotlikoff and Summers (1981) construct synthetic life-cycle profiles of consumption and income and present some evidence that consumption falls significantly below income only after age 50. Carroll and Summers (1991) report that consumption tracks income across countries, education and occupation groups, providing additional evidence that life-cycle savings do not seem to occur until late in life. Both studies find that consumption tracks income over the lives of households until around age 50. However, these studies are using *cross-sectional* data to infer *time-series* behavior. Thus, the close correlation between consumption and income may come from cohort-effects: on average, young families have larger lifetime resources and hence consume

⁶We do not enter the debate on the relative importance of retirement versus bequest savings. Implicitly, our retirement savings measure will include both.

more.⁷ Further, changes in family size and consumer-needs over the life cycle may impart a hump-shape to consumption which does not come from a failure of CEQ consumption smoothing. We address these issues by adjusting both consumption and income for life-cycle changes in family size and for cohort effects.

More recently, Attanasio and Browning (1995) and Attanasio, Bank, Meghir and Weber (1995), using data from the UK Family Expenditure Survey (FES) and the US Consumer Expenditure Survey (CEX) respectively, have examined life-cycle behavior adjusted for cohort and family size effects. Attanasio et al. (1995) shows that the residuals from a regression of consumption on family composition and labor supply variables are uncorrelated with age. However, if the true consumption profile is hump shaped over the life cycle, this regression suffers from an omitted variable bias, which will incorrectly assign the hump to changes in demographics. Both papers also draw life-cycle implications from certainty equivalent Euler equation estimation under flexible representations of preferences. They do not reject the certainty equivalent life-cycle model.⁸ Here again, however, the instruments used in the Euler equation estimation are likely to be correlated with the omitted precautionary term. This overestimates the share of the consumption hump attributed to labor supply and family size variables.⁹

We are also building on previous studies which parameterize and simulate life-cycle consumption models with uncertainty. Hubbard et al. (1994) and Carroll (1993a) show that the optimal consumption choices of consumers lead to profiles which are hump-shaped and track income over the early part of life for some parameterizations. Hubbard, Skinner and Zeldes (1995) go further and choose simulated profiles so as to try to

⁷See for example Attanasio and Weber (1995).

⁸Since adjusted profiles are not reported, it is not clear whether their finding is merely the result of adding covariates - labor supply and family size - so that the instrumental variable technique they employ no longer has the power to detect the relationship. We note that one of the best papers that uses occupation and education to predict expected income growth, Lusardi (forthcoming), finds that consumption does track income.

⁹Both of these papers separately deflate components of consumption, which can eliminate a consumption-income parallel. If liquidity constraints are binding or consumers are buffer-stock then *nominal* consumption will track *nominal* income. Finally, both papers look only at nondurable consumption, which, as we demonstrate and discuss, is not simply a scaled down version of total consumption.

reproduce constructed profiles of assets over the life-cycle.¹⁰ They re-interpret low-asset holding by most households as driven by means-tested government programs.¹¹ Our approach goes beyond those studies by estimating a structural model of consumption.

Further, Palumbo (1994) uses individual consumption, income and asset data to estimate individual consumption functions for retirees. We choose to rely on average profiles precisely because we do not believe that the individual-level data are of sufficient quality to support the employed technique in general.¹²

Finally, nearly all previous Euler equation estimations of these parameters ignore the precautionary term in the Euler equation, a potentially serious flaw. Two recent papers immune to this problem are Carroll and Samwick (1994) and Barsky, Juster, Kimball and Shapiro (1995). The former, using asset data and a theoretical framework similar to ours, finds that the discount rate is poorly identified. The latter, using survey questions about preferences over lotteries and income paths, estimates an intertemporal elasticity of substitution and a discount rate both lower than what we estimate. We are exploiting lower frequency movements in the data than typical Euler-equation tests. High-frequency Euler equation tests might reject the CEQ PIH, while the CEQ PIH could still be a reasonably accurate model for low frequency analysis.

The structure of this chapter is as follows. In the next section, we lay out a model of consumer maximization and its implications for the construction of consumption and income profiles. We describe the numerical dynamic programming techniques used to solve the model and present characterizations of optimal behavior. Section 3.3 describes the data, discusses empirical issues involved in constructing our life-cycle profiles, and presents graphs of the profiles. Section 3.4 introduces the method of Simulated Moments methodology for estimating the model. Finally, the last two sections of this chapter

¹⁰These authors do not correct for family-size or cohort effects.

¹¹We do not address this alternate interpretation, but simply note that we believe that the heterogeneity in skills, abilities, and wealth across people starting their working lives makes the low-asset trap of their model very relevant for a small subset of the population and much less relevant for the typical household. See also Carroll and Samwick (1994) for a critique of the implications of their approach.

¹²For instance, Palumbo (1994) must use a scaled-up measure of food consumption as his measure of household consumption and must make various assumptions about each individual's expected health dynamics.

present the results of the estimation and conclude. Appendices contain more detailed descriptions of the numerical optimization, the CEX data, and econometric technique.

3.2 Consumption Behavior with Stochastic Income

We begin by setting up a model of consumer behavior incorporating two saving motives: retirement and precautionary. The life-cycle saving motive results from the finite lifetime of individuals and from their retirement period. Income uncertainty at the individual level provides incentives for precautionary savings.

By nature, we are dealing with a non-stationary problem, as expected income follows a deterministic path and the permanence of the shocks depends upon the consumer's age. The consumer's program must be solved recursively, keeping track of consumption rules at each age. But this results precisely from our initial intuition: the systematic age-pattern of consumption functions will reflect the interaction of the two saving motives and will translate into some definite life-cycle profile. Conversely, the age-pattern of the profiles we construct will allow us to identify life-cycle consumption functions.

3.2.1 The Canonical Model with Labor Income Uncertainty

Our starting point is the basic discrete-time, life-cycle model of consumption behavior. Consumers live for N periods and work for $T < N$, where both T and N are exogenous and fixed. In every period $t < T$, the consumer receives a stochastic income Y_t and consumes C_t . There is only one asset in the economy, totally liquid and yielding a constant gross, after-tax, real interest rate R . Our unit of analysis is the household. We assume that preferences take the standard additively separable expected utility form, with a discount factor β :

$$U = E \left[\sum_{t=1}^N \beta^t u(C_t, Z_t) + \beta^{N+1} V_{N+1}(W_{N+1}, Z_{N+1}) \right], \quad (3.1)$$

where W_t represents total financial wealth and Z_t is a vector of household characteristics (e.g. family size). V_{N+1} represents the value to the consumer of any assets left at the time of death, capturing any bequest motive. The consumer maximizes (3.1) given an initial wealth level W_1 , and the constraint that terminal wealth is non-negative $W_{N+1} \geq 0$. The dynamic budget constraint is:

$$W_{t+1} = R (W_t + Y_t - C_t). \quad (3.2)$$

We further assume that the felicity function $u(\cdot, \cdot)$ is of the Constant Relative Risk Aversion (*CRRA*) form, with intertemporal elasticity of substitution $1/\rho$, and multiplicatively separable in Z :¹³

$$u(C, Z) = v(Z) \frac{C^{1-\rho}}{1-\rho}.$$

If income were certain, the solution to this program would be standard: the consumer would choose a consumption path such that:

$$\frac{C_{t+1}}{C_t} = \left(\beta R \frac{v(Z_{t+1})}{v(Z_t)} \right)^{\frac{1}{\rho}}. \quad (3.3)$$

With constant individual characteristics, (3.3) implies a constant growth rate of consumption. Consumption increases (respectively decreases) over time when the interest rate is larger (respectively smaller) than the discount rate. The growth rate of consumption (as opposed to its level) is independent of the income profile. The consumption level is then determined by the lifetime budget constraint and the terminal value function. The desire to smooth consumption over the entire lifetime will induce households to save for retirement and for bequest during their working lives.

¹³Equivalently, ρ is the coefficient of relative risk aversion. This is a well known feature of additively separable expected utilities. This chapter will not explore alternatives such as non-expected utilities or habit formation. (See Kreps and Porteus (1978) and Heaton (1993) for these issues.)

When individual characteristics vary over the life cycle, the growth rate of consumption may change accordingly. For instance, if the marginal utility of consumption increases with family size, consumption will grow as family size increases, and decrease as children leave the household. These variations in individual characteristics may induce a positive correlation between consumption and income over the life cycle.

The certainty (or certainty-equivalent) LCH provides extremely valuable insights on the determinants of consumption and savings. However, by deliberately assuming away individual uncertainty, it may be missing an important part of consumer behavior.

With individual income uncertainty and prudence, households will hold precautionary savings to insure themselves against future contingencies. The variation in this precautionary motive has far-reaching and striking implications. The main consequence of income uncertainty is to increase the slope of the consumption profile (provided that consumers are prudent). Hubbard et al. (1994) demonstrate that this uncertainty can lead to hump-shaped consumption profiles as households save for precautionary reasons early in life and run down these assets during retirement due to lower levels of uncertainty and an increased probability of death. Carroll (1992) and Deaton (1991) analyze the case in which consumers are also *impatient*: in the absence of uncertainty, households would like to borrow in order to finance a high level of current consumption. In addition, Deaton (1991) imposes liquidity constraints while Carroll (1992) sets up a model in which consumers choose never to borrow. In either rendition, assets play the role of a *buffer stock* against bad income shocks. Consumers have a target level of liquid assets, above which impatience dominates and assets are decumulated, and below which the precautionary motive dominates and assets are accumulated. Thus the theory predicts a correlation of expected income growth and consumption growth at both high and low frequency. Over the life-cycle, consumption will appear to track income.

In the rest of the chapter, we explicitly incorporate uninsurable idiosyncratic income uncertainty in addition to our retirement saving motive. We adopt Carroll's (1992) formulation, and decompose the labor income process into a permanent, P_{jt} , and a transitory,

U_{jt} , components (where j indexes occupation and education groups):

$$\begin{aligned} Y_{jt} &= P_{jt}U_{jt} \\ P_{jt} &= G_{jt}P_{jt-1}N_{jt} \end{aligned} \tag{3.4}$$

The transitory shocks, U_{jt} , are independent and identically distributed, take the value 0 with probability $p \geq 0$, and are otherwise log-normally distributed so that $\ln U_{jt}$ has mean zero and variance σ_{uj}^2 . The log of the permanent component of income, $\ln P_{jt}$, evolves as a random walk with drift. G_{jt} is a deterministic growth factor (specific to age t and group j) while $\ln N_{jt}$, the shock to the permanent component of income, is independently and identically normally distributed with mean zero and variance σ_{nj}^2 . Thus income evolves as a *nonstationary*, serially correlated process, with both permanent and transitory shocks, and a positive probability of zero income in every period.¹⁴

Two points are worth noting. First, permanent shocks are only as permanent as the length of the working life: all shocks are ultimately transitory, as consumers retire and die. In the last working period, transitory and permanent shocks are equivalent. As a consequence, the propensity to consume out of “permanent” shocks to income will decrease with age, a point emphasized by Clarida (1991). This property holds true for the CEQ LCH also. Second, in this setup consumers will choose never to borrow against future labor income. This follows from (a) there being a strictly positive probability that labor income will be arbitrarily close to zero for the rest of the working life and (b) the Inada condition $\lim_{c \rightarrow 0} u'(c) = \infty$.¹⁵ It is important to note that this holds true, even when p , the probability of strictly zero income, is set to zero. Suppose the household were to borrow in the next to last working period. Then, with strictly positive probability it would be left without any wealth in the last working period. The household would then have an infinite expected marginal utility. Simple backward induction implies that it will never be optimal to borrow. Thus, in this setup, the precautionary motive acts as

¹⁴While Abowd and Card (1989) found that change in labor income was best characterized by an $MA(2)$ process, they also found little gain in moving from an $MA(1)$ to an $MA(2)$.

¹⁵A condition always satisfied with isoelastic utility and positive risk aversion.

a *self-imposed* liquidity constraint.¹⁶

Going from the model to the data, we need to make three assumptions. First, note that in order to solve the consumer's problem as stated, we need to specify both the nature of uncertainty during retirement and a bequest function. While there have been good attempts at modelling consumer behavior after retirement,¹⁷ we feel that we know too little about the form that uncertainty takes after retirement to use our methodology and draw inferences from post-retirement behavior. Uncertainty arises from different sources— medical expenses, the timing of death and asset returns. *Inter-vivos* bequests are important. Although these sources of uncertainty are also present to some extent in the last working years, labor income uncertainty seems to be the most important form of uncertainty. Further, high quality information on household asset holdings together with consumption and income are not available. Given that investment income, Social Security, and pensions represent the main sources of income during retirement, it is currently difficult to establish consumption patterns as a function of total wealth.

Even with a proper treatment of retirement issues, one would have to make a guess about the bequest function. Therefore we decided instead to make use of Bellman's optimality principle, and truncate our problem at retirement.¹⁸ Defining the value function at time t , V_t , our problem becomes:

$$\begin{aligned}
 V_\tau(X_\tau, P_\tau, Z_\tau) &= \max_{c_\tau, \dots, c_T} E_\tau \left[\sum_{t=\tau}^T \beta^{t-\tau} v(Z_t) \frac{C_t^{1-\rho}}{1-\rho} + \beta^{T-\tau} V_T(X_T, P_T, Z_T) \right] \\
 \text{s.t. } X_{t+1} &= R(X_t - C_t) + Y_{t+1},
 \end{aligned} \tag{3.5}$$

¹⁶If instead we had assumed a strictly positive lower bound on income, the consumer could borrow up to the present discounted value of certain future income.

¹⁷See Hubbard et al. (1994) and Palumbo (1994).

¹⁸Our approach does not emphasize the relative importance of retirement versus bequest wealth. The salvage value function accommodates a bequest motive. See Modigliani and Brumberg (1956) and the ensuing debate with Kotlikoff and Summers (1981) about the relative importance of life-cycle and bequest savings.

where we define cash on hand X_{t+1} as total available financial resources at time $t + 1$:

$$X_{t+1} = W_{t+1} + Y_{t+1} = R(X_t - C_t) + Y_{t+1}.$$

Second, the model imposes a single vehicle for precautionary and retirement savings, since there is only one asset. In practice, much of retirement savings is accumulated in the form of illiquid assets, only available after retirement.¹⁹ This suggests that the relevant model of consumption behavior should incorporate an additional asset which is illiquid and accessible only after retirement. However, this would substantially complicate the problem by introducing another control variable (how much to save in liquid versus illiquid assets) and state variable (illiquid assets). In order to keep our estimation procedure feasible, we instead assume that consumers will receive at the age of retirement some accumulated illiquid wealth, proportional to their permanent income. This illiquid wealth accumulates exogenously and cannot be borrowed against. Effectively, this imposes a borrowing constraint $W_t \geq 0$. We denote accumulated illiquid wealth as H_t and total financial wealth after retirement as $A_t = H_t + W_t$.

Lastly, we need to postulate a salvage value function which summarizes the consumer's problem at retirement time. We choose a functional form which maintains the tractability of the problem and is flexible enough to allow robustness checks:

$$V_T(X_T, H_T, Z_T) = k v(Z_T) (X_T + H_T)^{1-\rho} \quad (3.6)$$

This functional form is exactly correct if the only source of uncertainty after retirement is the time of death. When we move to estimation, we will calibrate the parameters of the associated consumption rule at retirement using information on consumer wealth, income, and consumption.

¹⁹Social Security wealth is definitely illiquid and is only available as annuities after retirement. Early withdrawal of pension and savings vehicles targeted for retirement purposes, such as IRA's, 401k plans and Keogh, is often penalized, if allowed at all. One might also consider housing wealth as part of retirement wealth. Empirical evidence suggests that households run down their housing wealth only extremely late in life.

3.2.2 Solving for Optimal Consumer Behavior

The setup of the problem combined with our particular choice of retirement value function makes the problem homogeneous of degree $1 - \rho$ in the permanent component of income. Since a second state variable would render our estimation procedure unfeasible, we assume here that individual characteristics are constant throughout the life cycle: $Z_t = Z = 1$. Furthermore, since our data does not track households over time, we cannot calibrate the family-size process for each household over its life cycle. We will instead directly control for family effects when constructing our profiles. This allows us to write the optimal consumption rule as a function of a single state variable, x_t , the ratio of cash on hand to permanent income:

$$x_{t+1} \equiv X_{t+1}/P_{t+1} = (x_t - c_t) \frac{R}{G_{t+1}N_{t+1}} + U_{t+1}. \quad (3.7)$$

We can then derive the Euler equation in any period prior to retirement:

$$\begin{aligned} u'(c_t(x_t)) &= \beta R E[u'(c_{t+1}(x_{t+1})G_{t+1}N_{t+1})] \\ &= \beta R \{p E[u'(c_{t+1}(x_{t+1})G_{t+1}N_{t+1})|U_{t+1} = 0] + \\ &\quad (1 - p) E[u'(c_{t+1}(x_{t+1})G_{t+1}N_{t+1})|U_{t+1} > 0]\}. \end{aligned} \quad (3.8)$$

where lowercase letters are normalized by the permanent component of income, and $c_t(x_t)$ represents the optimal consumption rule at time t as a function of normalized cash on hand x_t . Next period expected marginal utility can be decomposed according to Bayes formula into expected marginal utility conditional on zero future income, and expected marginal utility conditional on a strictly positive income.

The solution to the consumer problem consists of a set of consumption rules $\{c_t(x_t)\}_{1 \leq t \leq T}$. In the last working period, under our previous assumptions, consumption will be linear in cash on hand:

$$c_T(x_T) = \gamma_0 + \gamma_1 x_T. \quad (3.9)$$

where $\gamma_0 = \gamma_1 h_T$.²⁰ Consumption in the next to last period is then found as the solution to (3.8) for all values of cash on hand, where we replace c_T using (3.9). Solving recursively generates c_{T-1}, \dots, c_1 .²¹ A complete description of the solution method is provided in Appendix B.

3.2.3 Characterization of Individual Consumption Behavior

Figure 3-1 shows the consumption rules at various ages, when the permanent income profile is flat ($G_t = G = 1$), there is no retirement period ($\gamma_0 = 0, \gamma_1 = 1$), working life starts at age 25 and ends at age 65, and consumers are impatient.²² When permanent income growth is constant, the finite horizon problem converges in the limit to the infinite horizon one, as we move further away from retirement.²³ Consumption is always positive, increasing and concave in cash on hand. One can also show that cash on hand can only increase if the income draw is sufficiently large.²⁴ Early in life, households will exhibit the standard *buffer stock behavior*: for low level of assets, typically less than the permanent component of their income ($x \leq 1$), households will consume most, but never all, of their financial wealth, and move to the next period with a very low level of assets. At high levels of cash on hand, households will consume a smaller fraction of cash on hand, but always enough so that they expect to run down their assets.

As death nears, the consumer faces less and less uncertainty from labor income shocks.

²⁰In the case of full certainty after retirement, it is straightforward to show that: $\gamma_1 = \frac{1-\kappa}{1-\kappa N^{-T+1}}$ where $\kappa = \beta^{1/\rho} R^{1/\rho-1}$.

²¹The consumption rules have to be found numerically, as no closed form solution exists for this problem. We did this using a discretization method. See Judd (1993). We solve the Euler equation (3.8) recursively on a grid of normalized cash-on-hand values. The consumption function is then interpolated between the points on the grid. In order to capture the curvature of the consumption rules at low values of cash on hand, the discretization grid is finer for $x \in [0, 2]$. See Caballero (1990a) for a closed form solution under exponential utility.

²²Other relevant parameters are $\beta = 0.931$, $\rho = 1.13$ and $R = 1.05$. These parameters generate buffer-stock behavior. Income uncertainty is the average amount, discussed later, and presented in Table 3.3.

²³In other words, the solution to the supremum problem in the infinite-horizon case is obtained as the fixed point of the associated functional equation. Theorem 9.12 in Stokey, Lucas and Prescott (1989) applies even though returns are unbounded, as long as the discount factor is strictly less than 1.

²⁴This condition is analyzed in more details in Deaton (1991) and Ayagari (1993). In the infinite horizon case, this guarantees that cash on hand has an ergodic distribution.

It is then rational for an impatient consumer to start *running down* their buffer of assets: the consumption rules converge progressively towards the 45° line. In the last period, obviously, the household consumes everything. Consumers save only in order to buffer income shocks. We note also that from age 25 to 55, our consumer has roughly the same consumption rule. Buffer stock savings are run down quite late in life.

Figure 3-2 shows consumption rules, for a household facing the same expected income growth, now assuming that the retirement consumption rule, at age 65 is more realistically characterized by:²⁵

$$\gamma_0 = 0.384, \gamma_1 = 0.049.$$

The two households have the same consumption rules early in life—governed by the common solution to the infinite horizon problem. In other words, both households will behave as standard buffer stock households in their youth. Now however, the agent will have to *accumulate* enough wealth for retirement purposes. As retirement nears, savings must *increase*. Note also that with impatient consumers, the retirement savings motive matters only late in life.

Figures 3-3 and 3-4 display randomly drawn profiles of consumption for households facing typical paths of income, retirement rules, and income uncertainty.²⁶ In both profiles, consumption tracks income early in life, and diverges later in life. Notice further that unexpected transitory shocks are better smoothed later in life, despite the fact that they contain greater information about total resources for the remainder of the life. Smoothing is easier later in life since retirement saving also acts as a large buffer.

What are the necessary conditions to generate buffer-stock behavior early in life? Previous characterizations have addressed the problem in a stationary environment. Specifically, Deaton (1991) defines buffer-stock consumers as consumers who would borrow against future income, were it not for uncertainty. In the CEQ LCH, the slope of nor-

²⁵We discuss calibration of γ in section 3.4.2.

²⁶What is typical will be discussed in detail in the next two sections of this chapter.

malized consumption is given by (see (3.3)):

$$\ln\left(\frac{c_{t+1}}{c_t}\right) = \frac{1}{\rho} \ln(R\beta) - \ln(G) \simeq \frac{1}{\rho}(r - \delta) - g, \quad (3.10)$$

where G is the gross rate of income growth. The ratio of consumption to income increases whenever the right hand side is positive. A higher preference for the present or a lower interest rate will make buffer stock behavior more likely. However, households with a low discount rate and facing a high interest rate may still decide to behave as buffer-stock consumers, saving only a small fraction of their resources *if their income profile is steep enough*.

With i.i.d. income growth G , uncertainty, and an infinite horizon, Deaton (1991) shows that, agents are buffer-stock if and only if:

$$R\beta E[G^{-\rho}] < 1. \quad (3.11)$$

As Carroll (1992) emphasizes, buffer-stock agents have a *desired level of cash on hand relative to permanent income*. The existence of such target is both necessary and sufficient for buffer stock behavior in the infinite horizon framework. At low levels of assets, agents on average save to build the buffer, and cash on hand is expected to increase. For large levels of normalized cash on hand, the precautionary motive vanishes and agents increase consumption. This target level of cash on hand is defined as the fixed point of the mapping from current to expected future cash on hand:²⁷

$$\bar{x}_t = E_t[x_{t+1} | x_t = \bar{x}_t]. \quad (3.12)$$

²⁷In our finite-horizon framework, this characterization will be necessary but not sufficient in the following sense. When the interest rate is lower than the discount rate, agents will have a target level of cash on hand throughout their working lives. For large levels of cash on hand, impatience will cause consumers to run down these assets. At low levels of cash on hand, consumers must save either for liquidity needs or retirement. However, near retirement, this target level will reflect life-cycle considerations and is likely to be extremely large. The behavior of such agents is unlike buffer stock behavior as previously described.

Figure 3-5 presents the mapping $\bar{x}_t \rightarrow E_t [x_{t+1} | x_t = \bar{x}_t]$ at various ages. It is convex, increasing and initially above the 45⁰ line, so that in general it need not have a fixed point.

3.3 Data and Consumption and Income Profiles

3.3.1 Profile Construction Methodology

The construction of our profiles is motivated by the model presented in the previous section. We want the income profiles to be usable as inputs to the consumer optimization problem and the consumption profiles to be comparable to average consumption paths from the optimal program. That is, we are interested in constructing the income growth profile $\{G_t\}_{t=25}^{65}$ and the average consumption profile $\{\bar{C}_t\}_{t=25}^{65}$. Consumption at a certain age t will be the log-average across the distribution of cash-on-hand, permanent component of income, and consumer characteristics for that age:

$$\bar{C}_t \equiv \exp \left[\int \int \ln c_t(x_t) P_t dF(x_t) dF(P_t) \right]$$

In order to construct such a profile, we must address three issues. First, due to its excessive noisiness, we do not exploit the limited panel nature of the Consumer Expenditure Survey, but instead rely on data from repeated cross-sections. Consequently, in our sample, birthyear and age will be correlated. Households observed at age sixty, say, will have been born long before those we observe at young ages and will have on average lower lifetime resources, and lower levels of income and consumption at each age. Ignoring birthyear effects would lead to a negative bias in our estimate of the slope of income and consumption growth, especially, late in life.²⁸ Second, the model refers to household consumption, adjusted for family size ($Z = 1$). Since household size is hump-shaped over the life cycle, the correlation found in previous studies between consumption and income

²⁸This point has been emphasized recently by Attanasio and Weber (1995).

over the life cycle may disappear after correcting for family size.²⁹ Finally, we are interested in exploiting some variation in expected income profiles across households. Thus, while we do conduct our exercise for the average household, we also focus our analysis on income and corresponding consumption profiles for subgroups of the US population defined by occupation and education groups. We assign households to these groups on the basis of the male head. Male labor force participation over the life cycle is high and stable, giving us more data and robust profiles.

We posit the following effects model for the natural logarithm of consumption for individual i , in education/occupation group j , of age a , in year t :

$$\ln C_{ji}^{at} = f_i^a + \pi_j^a + b_i + y^t + \varepsilon_{ji}^{at} \quad (3.13)$$

That is, consumption of individual i is determined by a family size effect, f_i^a ; an effect, π_j^a , specific to their age, a , and education/occupation group, j ; a cohort effect, b_i ; a year effect, y^t ; and an idiosyncratic, individual effect ε_{ji}^{at} . We are interested in recovering π_j^a . By doing so, we create a profile which has a constant family size over the life cycle, and also correct for the fact that we do not actually follow the same individuals over their entire lives. As discussed in Deaton (1985), it is not possible to remove the linear component of the time and cohort effects without also removing the average (across education/occupation groups) age profile of consumption.³⁰ We make the identifying restriction that time effects are related to business cycles and thus are well captured by the partial correlation of consumption with the regional unemployment rate.

Our procedure can then be summarized as follows. First we put all the data into real

²⁹Attanasio et al. (1995) make this argument, and find that, after correcting for family size, consumption is no longer significantly dependent on expected changes in income at high frequency. These authors do not demonstrate that the profiles are flat after making this correction. It is worth noting that both this paper and our research assume that family size is exogenous. If the buffer-stock model is correct and having children is a form of consumption, then the decision to have children is affected by the expected path of income. By correcting for family size composition over the lifecycle, one is removing that portion of changes in consumption driven by expected income changes. As we will see, however, the profiles presented subsequently suggest that correcting for family size attenuates but does not eliminate the consumption-income parallel.

³⁰This follows from the annoying identity that interview year less age equals birthyear.

terms using the Gross Domestic Product implicit price deflator for personal consumption expenditures.³¹ Second we generate family size adjustments— f_i^a in equation (3.13)— and apply them to all the consumption and income data so that households have a constant effective size over the life cycle.³² Then, to construct unsmoothed profiles, we estimate the following model, over households with male heads aged 25 to 65, by weighted least squares with weights based on the CEX population weights:

$$\ln X_i - \hat{f}_i = \pi a_i + \eta b_i + \nu \mathcal{U}_i + \tau Ret_i + \varepsilon_i, \quad (3.14)$$

where X is either consumption or income, a_i is a complete set of age dummies crossed with education or occupation group dummies, b is a complete set of cohort dummies (less one), \mathcal{U} is the Census region unemployment rate in year t , and Ret is a dummy for each group which is equal to 1 when the respondent is retired. Profiles are constructed by predicting $\ln X_i$ for each age and grouping, setting the cohort and unemployment rates are at their average values and the retirement dummies to zero. Smooth profiles are estimated by replacing the age and cohort dummies by fifth order polynomials, and extending the highest age to 70 to avoid some of the endpoint problems commonly encountered with polynomial smoothing.

Income profiles are used to construct estimates of $\{G_t\}$ for the consumer problem. Recalling that, $\ln Y_t^i = \ln P_t^i + \ln U_t^i = \ln G_t + \ln P_{t-1}^i + \ln N_t^i + \ln U_t^i$, after removing the cohort, family, and time effects, our procedure is in effect taking a sample average over a large number of individuals, M , with the same characteristics:

$$\frac{1}{M} \sum_{i=1}^M \ln Y_t^i = g_t + \frac{1}{M} \sum_{i=1}^M \ln Y_{t-1}^i + \frac{1}{M} \sum_{i=1}^M \ln N_t^i + \frac{1}{M} \sum_{i=1}^M \ln U_t^i - \frac{1}{M} \sum_{i=1}^M \ln U_{t-1}^i$$

³¹Again, it is important not to use different deflators for different items within consumption or for income and consumption. This could break the relationship between cash on hand and consumption in nominal terms which is the relationship predicted by the buffer-stock theory.

³²We construct f_a^i by running equation (3.14) without separate age effects by subgroup and with family size dummies on the right-hand side. The effective family size is then 2.8 (the sample average). We also experimented with exogenous family size adjustments—assuming f_a^i is simply family size raised to the power -0.7 . This led to profiles which were noisier and flatter early in life.

Applying the Law of Large Numbers, the probability limits of the last three term are all zero. Hence, we get

$$plim \left(\frac{1}{M} \sum_{i=1}^M \ln Y_t^i - \frac{1}{M} \sum_{i=1}^M \ln Y_{t-1}^i \right) = \ln G_t.$$

Thus simply first differencing our log-average income levels gives the income growth rates which are input into the simulated model. $\bar{Y}_t = \exp[\frac{1}{M} \sum_{i=1}^M \ln Y_t^i]$ provides an estimate of average income by age.

3.3.2 The Consumer Expenditure Survey and Our Use of it

The main data source for the consumption and income profiles is the Consumer Expenditure Survey (*CEX*). The *CEX* contains information about consumption expenditures, demographics, income and assets, for a large sample of the US population. The Survey is conducted by the Bureau of Labor Statistics in order to construct baskets of goods for use in the bases for the Consumer Price Index, and has been run continuously since 1980. We use data from 1980 to 1993 from the family, member, and detailed expenditure files. The survey is known to have excellent coverage of consumption expenditures, to have reasonable data on liquid assets, and to have income information of moderate quality.³³ The survey interviews about 5500 households each quarter. In a household's first interview, the CEX procedures are explained to them and information is collected so that they can be assigned a population weight. They are then interviewed four more times (once every three months) about detailed consumption expenditures over the previous three months. In interviews two and five, income information is collected, and in the final interview asset information is collected. Families rotate through the process, so that about 25% of households leave and are replaced in each quarter. About half of all households make it through all the interviews.

Each household contributes one datapoint to our sample. For each household we construct a measure of household income and consumption, and assign it to an occupation

³³See Lusardi (forthcoming), Attanasio (1994), and Branch (1994).

group, an education group, a birth cohort, an interview year, and a Census region. In order to obtain a high quality sample which tracks men and has the required information, we drop a significant portion of the data and make a series of adjustments. A detailed description of the data preparation is contained in Appendix *C*, however, we will make note of three major points here, and then turn to our definitions of consumption and income.

First, we dropped households which are classified as incomplete income reporters, which had any of the crucial variables missing, or which reported changes in age over the course of the survey greater than one year or negative. Households dropped when constructing profiles by occupation remain in the education profiles. We did not use the occupational classifications of Armed forces, Service workers, and Farming, forestry, or fishing due to small cell sizes. Similarly we do not analyze the group of households with male heads holding less than 9 years of schooling due to very few younger households. Second, we dropped all households with male heads younger than 25 or older than 70, since as discussed above we are choosing to focus on the working life. Cell sizes are reported in Table 3.1. Third, while topcoding is very infrequent in consumption information, the household annual income variable reflects summation over a topcoded item for roughly half a percent of our households. Since, in most years, topcoding occurs at \$100,000 in income subcategories, reported individual annual labor income is the source of almost all income topcoding problems. However, households are also asked the gross amount of their paycheck and what length of time-period this paycheck covers. By multiplying these two variables together, we construct a second measure of annual labor income. Topcoding on this variable occurs only for a few cases. We correct our measure of after-tax family income by replacing the reported annual labor income in family income with our constructed measure whenever the family income variable is topcoded. We are able to correct almost all topcoding.

Finally, we consider how best to construct measures of income and consumption which match the concepts in the theoretical model. First, we choose to define consumption as total household expenditures less those on education, medical care, and mortgage

Table 3.1: CELL SIZES FOR CONSUMPTION PROFILES

	HOUSEHOLDS	PERCENT
TOTAL	43031	
EDUCATION GROUP		
Some Highschool	4409	10.2
Highschool Degree	12906	30.0
Some College	10027	23.3
College Degree	6570	15.3
Grad/Prof School	6087	14.1
OCCUPATION GROUP		
Managerial and Professional Specialty	13215	34.3
Technical, Sales, and Admin. Support	6976	18.1
Precision Production, Craft and Repair	5061	13.1
Operators, Fabricators, and Laborers	7116	18.5
Self-Employed	3479	9.0

Note: this table shows counts used in constructing unsmoothed consumption profiles. More observations are used for smoothed profiles and fewer are available for income profiles.

interest. These categories of expenditure do not provide current utility but rather are either investments or negative income shocks.³⁴ They are also excluded from our income definition.

It should be noted that our model refers to nondurable consumption at annual frequencies. Since we are averaging expenditures across a large number of individuals and looking across one-year horizons, the distinction between durables and nondurables is less likely to matter. Since the buffer stock model gives strong predictions about consumption tracking income, it is important not to break the consumption-income link when studying consumption.³⁵

Our measure of income is comprised of after-tax family income less Social Security tax payments, mortgage interest, expenditures on medical care, spending on education, pension contributions and after-tax asset and interest income. For the first five adjustments, the related expenditures do not provide current utility but are either non-liquid

³⁴We are arguing here that user cost of housing - repairs, maintenance, utilities, and housing services - captures the expenditures made for consumption on housing.

³⁵When computing power has increased, this issue is likely to be surmountable by adding a state variable for the stock of durable goods and taking a stand on the size and relevance of adjustment costs.

investments or, in the case of medical care, simply losses of income. Further these expenditures involve a large amount of commitment and are hard to substitute intertemporally. Were we to include pension contributions in income, for example, our measure of liquid assets which could be used to buffer bad income shocks would include pension wealth. We remove asset income since the input to our theoretical model is a profile of income net of liquid asset returns. Attanasio (1994) looks at savings behavior and finds large savings by the typical household over the entire life. If we consider a larger measure of income and take the relative levels of consumption and income seriously, we also find significant savings over the life cycle. Our preferred interpretation of pension and housing wealth however is that it is illiquid, and thus that the typical household has little in liquid assets with which to buffer income shocks early in life.

3.3.3 Life Cycle Profiles

Figure 3-6 presents the estimated consumption and income profiles for our entire sample.³⁶ Even after correcting for cohort, time, and family effects, both profiles are still hump shaped and still track each other early in life. Consumption lies above income from age 25 to 28. Reflecting on our own experiences, we may interpret this as underreporting the assistance which is provided by intergenerational transfers early in life. After these first few years, consumption rises with income from age 30 to age 45, when consumption drops significantly below income. This tracking is however a lot less than is observed in profiles constructed by simply averaging cross-sections. As stated above, the two main reasons for this are the changes in family size over the life cycle and the different wealth and incomes of different cohorts. Figure 3-7 displays the profile of average family size over the life cycle. Figures 3-8 and 3-9 present the consumption and income profiles without the cohort adjustment and without the family size adjustment, respectively. In each case the unadjusted profiles are more hump shaped, and seem to track each other

³⁶We get reasonable relative levels of consumption and income as does Attanasio (1994), who uses relative levels in his analysis of saving in the U.S. However, the levels may not be tightly identified by the data. In our results section, we will present evidence on estimation which instead uses only information from the changes in income and consumption.

Table 3.2: F-TESTS FOR FLATNESS OF CONSUMPTION PROFILES

	F-STAT.		F-STAT.
WHOLE SAMPLE	1121		
EDUCATION GROUP		OCCUPATION GROUP	
Some Highschool	637	Managerial and Professional Specialty	934
Highschool degree	1015	Technical, Sales, and Admin. Support	803
Some College	1023	Precision Production, Craft and Repair	676
College Degree	804	Operators, Fabricators, and Laborers	790
Grad/Prof School	533	Self-Employed	604

Note: the F-statistic is distributed as a $F(39, 43000)$. The critical value is 1.40.

more closely. These profiles are more directly comparable to those shown in Kotlikoff and Summers (1981) and Carroll and Summers (1991) which do not correct for family size over the life cycle or for cohort effects.

Despite the fact that the asset data in the CEX is of lesser quality than data from sources like the Survey of Consumer Finances (SCF), we can see an interesting age pattern in the profile of asset accumulation, which mimics that found in more accurate surveys. Figure 3-10 shows a life-cycle profile of the ratio of total liquid assets to income. Liquid assets in the CEX are the value of holdings of stocks and bonds, and the cash held in savings and checking accounts. We observe that the typical household accelerates its building of a stock of liquid assets around age 45. This age pattern also shows up in the profiles of income and consumption and is the crucial feature that will help us to pin down the time-heterogeneity in consumer behavior: until around age 45, people consume roughly their income, saving very little in the form of liquid assets.

Figures 3-11 and 3-12 give some evidence that consumption and income track each other across subgroups of the population defined by education and occupation levels. These graphs are unfortunately noisy. However, despite the noise in the data, one can see that the occupation and education groups with the most pronounced humps in income present the most pronounced humps in consumption. Further, we can formally reject the null hypothesis that the consumption profiles are flat. Table 3.2 presents F-tests of the equality of 40 age dummies in a profile regression which also includes age (when

appropriate, interacted with subgroup). We get very strong rejection of a constant slope in all our profiles.³⁷ We could also proceed to test whether the income profiles are significant in predicting consumption profiles across occupation and education groups. But this is essentially a Hall and Mishkin (1982)-style test of the CEQ consumption Euler equation. Lusardi (forthcoming) performs such a test, merging Panel Study of Income Dynamics (PSID) data with the CEX consumption to get the best possible combination of data, and she rejects the null hypothesis that the consumption profiles are unrelated to the income profiles.

Finally, Figure 3-13 displays the profile for total, nondurable, and food consumption, all rescaled to the same mean. This figure demonstrates that total consumption over the life cycle is not simply a scaled up version of nondurable consumption. Expenditures on durable goods involve spending income when the good is purchased, and receiving a utility flow over time. To the extent that consumers would like to borrow against future income, purchases of durable goods tighten this constraint by moving expenditures forward in time relative to utility flow. If consumers are buffer stock when young and if nondurable and durable consumption are substitutes, then one would expect to see total consumption rising slower early in life and peaking later than nondurable consumption. The relative profiles of these consumption categories are at least consistent with what we shall conclude subsequently: that households are behaving as buffer stock consumers early in their lives.

³⁷Of course it is possible that we have omitted a key preference shifter which varies over the lifecycle. As Attanasio et al. (1995) note, labor supply is an obvious candidate. Thus, we also test for flatness conditioning on labor income variables. We include in the regressions, annual hours worked by the male head and annual hours worked by the female head. We are still able to reject flatness in the total and in all subgroups. We do not use profiles constructed after removing this correlation in our later analysis, since our theoretical model implicitly assumes that utility is additively separable in leisure and consumption expenditures.

3.4 Estimation Strategy

3.4.1 Method of Simulated Moments (MSM) Estimation

According to section 3.2, consumption at age t for individual i depends on cash on hand x_t^i , the realization of permanent component of income P_t^i , the entire path of expected permanent income growth $\mathfrak{G}_T = \{G_t\}_{t=1}^T$, and the parameters of the consumption problem $\tilde{\theta}' = (\beta, \rho, R, \gamma_0, \gamma_1, p, \sigma_u^2, \sigma_n^2)$, an 8×1 vector.³⁸ In practice, it will not be possible to estimate directly all the elements of $\tilde{\theta}$. Instead, we will calibrate most of the parameters using existing micro data and will focus on the estimation of the structural parameters of the utility function β and ρ . In other words, we rewrite $\tilde{\theta} = (\theta', \chi')'$ where $\theta = (\beta, \rho)'$ and $\chi = (R, \gamma_0, \gamma_1, p, \sigma_u^2, \sigma_n^2)'$. The elements of χ will be calibrated from micro data and we will estimate only the elements of θ . Defining the vector of state variable $z_t^i = (x_t^i, P_t^i)$, for individual i , we postulate the following data-generating process:

$$\ln C_t^i = \ln c_t(z_t^i, \theta; \chi, \mathfrak{G}_T) + \epsilon_t^i = \ln(c_t(x_t^i, \theta; \chi, \mathfrak{G}_T) P_t^i) + \epsilon_t^i, \quad (3.15)$$

where ϵ_t^i is an idiosyncratic shock that represents measurement error in consumption levels and satisfies $E[\epsilon_t^i | z_t^i] = 0$.³⁹ We are interested in estimating θ . Direct estimation of (3.15) is not possible since we have only poor information on individual assets x_t^i .⁴⁰ We do however observe χ and average consumption at each age, $\ln \bar{C}_t$, as defined in the previous section. This suggests that we can look directly at *the unconditional expectation of log-consumption at each age*:

³⁸For ease of notation and consistency with our theoretical model, we assume in this subsection that age runs from 1 to T .

³⁹ ϵ_t^i may also encompass missing variables such as class. Our approach remains correct as long as ϵ_t^i and the state variables are independent.

⁴⁰This is the approach taken by Palumbo (1994) who uses maximum likelihood estimation and PSID data to estimate structural parameters during the retirement period. One could also estimate consumption functions non-parametrically, much like Gross (1994), using Kernel methods, estimates investment functions of firms facing liquidity constraints.

$$\ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T) \equiv E[\ln C_t(z, \theta; \chi, \mathfrak{G}_T)] = \int \ln C_t(z, \theta; \chi, \mathfrak{G}_T) dF_t(z). \quad (3.16)$$

This says that average consumption of households equals the average level of consumption over both values of cash-on-hand and levels of permanent income. Our approach consists in matching the T moment conditions:

$$E[\ln C_t - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T)] = 0.$$

Defining the sample moment at age t , $g_t(\theta; \chi, \mathfrak{G}_T) = \ln \bar{C}_t - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T)$, and $g(\theta; \chi, \mathfrak{G}_T)$ the vector of moments, the estimation procedure minimizes

$$g(\theta; \chi, \mathfrak{G}_T)' W_T g(\theta; \chi, \mathfrak{G}_T)$$

where W_T is a weighting matrix. With a weighting matrix equal to the identity, this is equivalent to minimizing the distance between the average and the estimated profile. That is, considering the life-cycle profile, we minimize:

$$S(\theta; \chi, \mathfrak{G}_T) = \sum_{t=1}^T (\ln \bar{C}_t - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T))^2. \quad (3.17)$$

Unfortunately we do not directly observe $\ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T)$, since we do not observe the distributions of permanent income or of cash-on-hand. Further, we do not have analytic solutions for the consumption functions or how they change with alteration of our key parameters. Thus, instead of computing the actual expectation with respect to the true distribution of z , we use a Monte-Carlo integration method and perform a *Method of Simulated Moments (MSM)* estimation. We draw a random sample of shocks to income $\{U_t^i, N_t^i\}_{i=1}^L$, as defined in (3.4), calculate the associated paths of consumption and cash

on hand, and compute the simulated sample average:

$$\ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T) = \frac{1}{L} \sum_{j=1}^L \ln C_t(z_t^j, \theta; \chi, \mathfrak{G}_T).$$

We then choose parameters to minimize:

$$\hat{S}(\theta; \chi, \mathfrak{G}_T) = \sum_{t=1}^T (\ln \bar{C}_t - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T))^2.$$

In practice, we simulate $\ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T)$ by running 20,000 independent income processes (temporary and permanent) for 40 years, and computing in each year the associated consumptions. That is, we average over 20,000 profiles as in Figures 3-3 and 3-4.

Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the *MSM* estimator $\hat{\theta}$ is both consistent and asymptotically normally distributed. Denoting the number of observations at age t as $I(t)$, $I = \frac{1}{T} \sum_{t=1}^T I(t)$, $\tilde{g}_t = \ln \bar{C}_t - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T)$, and $\tilde{g} = (\tilde{g}_1, \dots, \tilde{g}_T)'$:

$$\sqrt{I}(\hat{\theta} - \theta_0) \rightsquigarrow \mathfrak{N}(0, V),$$

where V is estimated by:

$$\hat{V} = (\hat{D}'\hat{D})^{-1} \hat{D}'\hat{\Omega}\hat{D} (\hat{D}'\hat{D})^{-1} \quad (3.18)$$

$$\hat{D} = \left. \frac{\partial \tilde{g}}{\partial \theta'} \right|_{\theta=\hat{\theta}} \quad (3.19)$$

$$\hat{\Omega} = \text{avar}(\tilde{g}) \quad (3.20)$$

Efficient estimation is obtained when $W_T = \Omega^{-1}$. In practice, we use $\hat{W}_T = \hat{\Omega}^{-1}$. This methodology also provides a useful overidentifying restriction test. If the model is correctly specified, the statistic

$$\chi_{T-2} = I \bar{g}(\theta; \chi, \mathfrak{G}_T)' \hat{\Omega}^{-1} \bar{g}(\theta; \chi, \mathfrak{G}_T)$$

is distributed asymptotically as Chi-squared with $T - 2$ degrees of freedom.

As discussed previously, the levels of our profiles might be misestimated. To test the robustness of our results, we also estimate without using information on the level of consumption. Assuming that the bias is constant through time, the moment conditions become:

$$E \left[\ln C_t - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T) - \alpha \right] = 0 \quad ; \quad \forall 1 \leq t \leq T,$$

for some unknown constant α . Since we are not interested in estimating this parameter, we instead rewrite the moment condition in first difference:

$$E \left[\Delta \ln C_t - \Delta \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T) \right] = 0 \quad ; \quad \forall 2 \leq t \leq T,$$

where Δ is the time-difference operator. This amounts to performing the estimation in first differences, that is, minimizing:⁴¹

$$\tilde{S}(\theta; \chi, \mathfrak{G}_T) = \sum_{t=2}^T \left(\Delta \ln \bar{C}_t - \Delta \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T) \right)^2.$$

The rest of the procedure is identical, with the exception that we now have only $T - 1$ moments.

3.4.2 Remaining Calibration and Data

In order to find the consumption rules for the consumers in our dynamic program, we must still specify the elements of χ . Computing power currently limits us to searching over only two parameters, and to checking robustness and relative explanatory power as

⁴¹with an identity weighting matrix.

we change other parameters. We now present our calibration of the remaining parameters in our model.

1. The after-tax real interest rate is set as $r = R - 1 = 3\%$ per annum. This is roughly the average real interest rate on high grade municipal bonds over the sample period of our data.
2. The variance of the permanent and transitory components of shocks to income, σ_v^2 and σ_n^2 , are taken from Carroll and Samwick (1994). Carroll and Samwick (1994) estimate these parameters from the Panel Study of Income Dynamics (PSID), which provides repeated high-quality measures of household income. The estimation procedure is based on income differences of different lengths and correctly assigns the relative importance of transitory and permanent income shocks even in the presence of significant moving average correlation of transitory shocks, up to an $MA(2)$. The procedure and data employed are designed to estimate the parameters for the income process in Carroll (1992)— that is, exactly the income process we have specified.⁴² Table 3.3 displays the variances of the permanent and transitory shocks across education and education groups.
3. We follow Carroll (1992) and set the probability of zero income to $p = 0.005$. This calibration comes from the PSID and again is estimated with the goal of calibrating exactly the income process which we are considering.
4. Given our assumption on the retirement value function, the optimal consumption function in the last working period is linear in the permanent component of income and the level of cash on hand, or in normalized terms: $c_T = \gamma_0 + \gamma_1 x_T$. This results, as discussed earlier, from the assumption that post-retirement income— pension and Social Security income— is illiquid, cannot be borrowed against before retirement, and that the only source of uncertainty after retirement is the time of death. In

⁴²The definitions of occupation in the PSID and CEX do not exactly overlap, so that we are required to make rather crude adjustments to one cell.

Table 3.3: VARIANCE OF INCOME SHOCKS

GROUP	VARIANCE OF PERMANENT SHOCK	VARIANCE OF TRANSITORY SHOCK
TOTAL	0.0217	0.0440
OCCUPATION		
Managerial and Prof. Speciality	0.0180	0.0357
Tech., Sales, and Admin. Support	0.0235	0.0361
Precision Prod., Craft, and Repair	0.0175	0.0432
Operators and Laborers	0.0299	0.0458
Self Employed	0.0165	0.0926
EDUCATION		
Some Highschool	0.0214	0.0658
Highschool Degree	0.0277	0.0431
Some College	0.0238	0.0342
College Graduate	0.0146	0.0385
Graduate School	0.0115	0.05

Source: Carroll and Samwick (1994).

order to calibrate the parameters of this consumption rule, we first construct an income profile which adds Social Security and pension contributions to the income measure we consider. We do not use housing wealth to calibrate this parameter since most elderly do not run down the asset value of their housing. The difference between this profile and the consumption profile is accumulated at the assumed interest rate and gives an estimate of total resources at retirement, $W_T + H_T$. A similar calculation using our main income measure gives a measure of liquid assets at retirement, W_T . Finally we calculate the required parameters using data from smoothed profiles at retirement as:

$$\begin{aligned} \gamma_1 &= \frac{C_T}{W_T + H_T + Y_T} \\ \gamma_0 &= \gamma_1 \frac{H_T}{Y_T}. \end{aligned} \tag{3.21}$$

The results are summarized in the Table 3.4.⁴³

⁴³We find significantly larger assets on both counts than appears in wealth data, as for example in Venti and Wise (1993).

Table 3.4: DATA FOR RETIREMENT CONSUMPTION RULE

	CONS.	INCOME	LIQUID	TOTAL	γ_0	γ_1
			ASSETS	RESOURCES		
TOTAL AVERAGE	18349	23839	164616	352490	0.384	0.049
EDUCATION GROUP						
Some Highschool	13631	16112	48296	156788	0.531	0.079
Highschool Graduate	14556	17791	146852	311519	0.409	0.044
Some College	18768	24865	215380	427021	0.354	0.042
College Graduate	19628	27375	344287	607399	0.297	0.031
Graduate School	24410	32690	345442	665461	0.342	0.035
OCCUPATION GROUP						
Managerial and Prof.	22461	37732	406053	706256	0.240	0.030
Tech., Sales, Admin.	21394	30809	288376	520212	0.292	0.039
Precision Prod., Craft	17010	26726	252064	436494	0.253	0.037
Operators, Laborers	15557	23438	229031	370511	0.238	0.039
Self Employed	19654	25350	9097	201430	0.658	0.087

Source: CEX 1980-1993.

5. Since households generally begin life with some assets, we capture this by assuming that initial cash on hand, x_0 , is equal to 0.3 times permanent income at age 25, a number consistent with the CEX.⁴⁴

3.5 Estimation Results

We first estimate both the discount factor and the intertemporal elasticity of substitution for the average household. We then turn to disaggregated results, by education and occupation groups. Lastly, we discuss the robustness of our results to the calibrated parameters.

3.5.1 Results for the Entire Sample

We start by asking what the standard Life-Cycle theory would predict, assuming away all uncertainty. Although this constitutes a crude attempt at matching the data, it serves

⁴⁴As it turns out, the results are mostly insensitive to this assumption. See subsection 3.5.3.

as a useful benchmark against which to evaluate the rest of our results. To give the best chance to the CEQ LCH, we perform first difference estimation of the CEQ LCH, not asking it to fit the mean of the consumption profile, as discussed at the end of subsection 3.4.1. Under certainty, equation (3.3) holds, implying, after controlling for individual characteristics, a constant growth rate of consumption over the working period:

$$\Delta \ln \bar{C}_t = \frac{1}{\rho} \ln(\beta R) \equiv \xi. \quad (3.22)$$

We estimate ξ from the coefficient on age in a least-squares regression on individual data. This procedure seems trivial only because of our earlier efforts to remove changing family-size and cohort effects. It is precisely this simplicity which gives the CEQ LCH its power. From our estimate of ξ , we use the delta method to recover the discount factor and its standard error, postulating a real interest rate of three percent and a coefficient of relative risk aversion of 0.49.⁴⁵ The latter choice matters little since consumption is estimated to be nearly flat. The former matters a lot, and changes our estimates one for one.⁴⁶ We estimate a discount rate of 2.56% with a standard error of 0.05.⁴⁷ However, and not surprisingly given Figure 3-6, the certainty model performs poorly when it comes to explaining the dynamics of consumption across the life cycle. The estimated profile does not capture the hump shape in consumption, as Figure 3-14 demonstrates. Were we to present figures which showed both our fitted values and the data unadjusted for family size and cohort effects, the CEQ LCH would look better and our procedure less naive.

Next we use our structural model to estimate both the discount rate and the coefficient of relative risk aversion. The resulting estimates are reported in Table 3.5.⁴⁸

⁴⁵This is the estimate from our MSM procedure with uncertainty. While in our stochastic model we are able to estimate the coefficient of relative risk aversion and the discount factor, here they are not separately identified.

⁴⁶This estimation procedure only captures the substitution effect, as is clear from (3.22). Income and wealth effects change the *level* of the consumption path, not its *slope*.

⁴⁷The standard error is not robust to serial correlation and probably is underestimated by an order of magnitude.

⁴⁸MSM estimation is performed in 3 steps. First, we do a broad grid search over the parameter space,

Table 3.5: RESULTS OF STRUCTURAL ESTIMATION

METHOD	LEVEL	LEVEL	DIFFERENCES
OPTIMAL WEIGHTING	No	Yes	No
β	0.9615	0.9625	0.9603
S.E.	(0.0043)	(0.0042)	(0.0048)
δ (%)	4.0017	3.8891	4.1313
S.E.	(0.472)	(0.4527)	(0.537)
ρ	0.5437	0.4897	0.558
S.E.	(0.2092)	(0.2068)	(0.2404)
χ^2	288.22	288.05	514.66

Note: *MSM* estimation for entire group in levels and first differences. Cell size is 43031. The last row reports a test of the overidentifying restrictions distributed as a Chi-squared with 38 and 37 degrees of freedom respectively. The critical values at 5% are 53.10 and 51.91. Efficient estimates are reported in the second column with a weighting matrix $\hat{\Omega}$ computed from the first step.

Estimation in levels yields tight estimates of both parameters of interest.⁴⁹ The structure which our model imposes on the data gives us strong predictions about preferences. Although first-difference estimates are slightly less precise, they do not appear to contradict the level estimates. Moreover, the overidentifying restrictions are more strongly rejected for the difference method. Therefore, in what follows, we will refer only to level estimates. Efficient estimation has a minor impact on the point estimates or their standard deviations. In both level and first difference estimation, we can reject strongly the overidentifying restrictions. The 5% critical value for a χ^2 (38) is 53.10, and 51.91 for a χ^2 (37). This is not entirely surprising, given the number of moments we use. Moreover, the initial and last moments (average consumption at young ages and just before retirement) are each possibly misspecified if we have misspecified the initial distribution of cash on hand or the retirement rule.

then a finer grid search and lastly use an optimization routine. This ensures that the program identifies the global minimum. The programs are CPU-intensive: one evaluation of the objective function requires approximately 45 minutes on an experimental P6 (or "686") chip running at 133 Mhz (courtesy of Intel Corp.). A 25x25 grid search therefore requires 12 days. For details, see Appendixes *B* and *D*.

⁴⁹An earlier working paper version of this chapter presented pictures of our objective function which demonstrated that the optima lay in a valley in (β, ρ) space. Many interpreted this as implying that the separate identification of β and ρ was tenuous— that we could really only accurately estimate a linear relation between parameters. We want to emphasize that the valley picture is correct, but that the valley has a clear minimum which gives us estimates of both parameters.

Our estimates of the discount rate are close to the interest rate. In the efficient estimation case, we cannot reject the hypothesis that the two are equal at standard levels of confidence. It is worth noting that the discount factor which we estimate is within a reasonable range. Using information on the elasticity of assets with respect to uncertainty, Carroll and Samwick (1994) estimate that discount rate is in the vicinity of 10-15% and argue that even higher discount rates are needed to rationalize the findings of Hubbard et al. (1994). Our lower discount rate, however, does not imply that households are not impatient enough to generate buffer stock behavior. Lower levels of impatience generate buffer stock behavior when combined with steep income profiles.

With our estimates in hand, we can now address how well the stochastic model fits the life-cycle consumption profile. Figure 3-17 plots the simulated and actual consumption data along with the income profile, and a 95% pointwise confidence interval for the simulated profile. The stochastic life-cycle model does a much better job at fitting the consumption profile than the certainty line.⁵⁰ Consumption tracks income until around age 40 – 42 and then falls sharply, as the household starts building up its retirement savings. Simulated consumption never exceeds income, except in the first periods of life.⁵¹ The tight structure imposed by the model is able to deliver good predictions in terms of consumption dynamics, despite having only two free parameters to work with.

Why are we able, within the context of our model to obtain such tight estimates of the discount rate? Figure 3-18 plots various simulated profiles for different β between 0.91 and 0.95, corresponding to a discount rate between 4.71 and 9.98 percent. It is immediately obvious that the profiles are quite sensitive to this parameter. With a higher discount factor, the agent is willing to save more and earlier for retirement purposes. The consumption path exhibits less of a hump shape, and may even be increasing over the entire working life. On the other hand, for more impatient consumers, consumption parallels income until much later in life and then falls more precipitously to build assets for retirement. This implies a stronger concavity of the consumption profile. Thus, our

⁵⁰Unfortunately, the two cases we analyzed are not nested, so that hypothesis testing is not possible.

⁵¹This result is simply an artifact due to our assumptions regarding the initial level of cash on hand.

method will yield tight estimates of the discount factor precisely because the discount factor drives the hump shape in consumption.

We now turn to the question of how household behavior changes over the life cycle. Figure 3-19 displays the average household saving rate when there are no initial assets ($x_0 = 0$). Two distinct phases in the consumer's life are visible. Until about age 40, households built their buffer and then consume roughly their income. Around age 40, retirement considerations induce an increase in saving. From then until the age of retirement, households consume much less than their current income.

We can put some additional structure on these two phases by looking at the target level of cash on hand at each age, as defined in section 3.2. Figure 3-5, already reported, plots next period expected cash on hand as a function of current cash on hand for various ages. Figure 3-20 directly computes the target level of cash on hand for consumers aged 25 – 48. One can see from the graph that the target level of cash on hand remains small early in life, around 1.3 times permanent income. Shortly after 40, the target increases substantially, as consumers try to build their retirement nest-egg. This figure shows a dramatic change in behavior. When the target level of liquid wealth is small, agents are “buffer stock”. Their consumption closely follows their income. Around age 41, agents desire to accumulate assets for retirement. With a large stock of wealth relative to permanent income, consumers can smooth high frequency movements in income. Their behavior more closely mimics that of certainty-equivalent consumers.

None of these results are assumed in our model. If we had found either consumers to be more patient or flatter adjusted income profiles, households could have behaved as life-cycle consumers for their entire lives. Similarly, very impatient consumers could have been buffer stock consumers all of their working lives, relying on illiquid wealth to finance consumption during retirement.

We can also decompose total saving at each age into life-cycle and buffer-stock saving. Our previous discussion might lead the reader to think that agents have no concern for retirement when they are young and no concern for labor income uncertainty later in

life. This is incorrect since consumers are rational and perfectly foresee their retirement needs. To proceed, we define total saving as the discounted variation in financial wealth from one period to the next, using our simulated profile.⁵²

$$S_t = (W_{t+1} - W_t) / R = (R - 1) / R W_t + Y_t - C_t.$$

Saving is equal to investment plus labor income minus consumption, i.e. to disposable income minus consumption. Next, for the estimated parameters, we compute the consumption path $\{C_t^{LC}\}$, that would occur under certainty.⁵³ We then define life-cycle saving as the difference between total income and life-cycle consumption

$$S_t^{LC} = (W_{t+1}^{LC} - W_t^{LC}) / R = (R - 1) / R W_t^{LC} + Y_t - C_t^{LC}$$

and buffer stock saving is defined as the complement. Figure 3-21 plots the precautionary saving, liquid and total life-cycle saving. The latter is defined by adding back to income pension and social security contributions. Given the estimated discount rate, CEQ life-cycle consumers would like to borrow early in life. However, precautionary saving motives cause them to hold a positive buffer stock of wealth. Around age 40, in accordance with our previous characterization, life-cycle savings becomes larger than precautionary savings. The need to build retirement savings sets in. As asset levels increase, the expected variance of consumption declines, decreasing the precautionary saving motive. This latter effect, which our previous decomposition masked, induces the agent to *run down* the buffer. As a result, the total saving rate later in life is smaller than under the certainty equivalent framework.

⁵²The discount comes from our assumption that income is received and consumption occurs at the beginning of the period. See equation (3.2).

⁵³In order to do this, we calibrate the certainty case, so as to yield the same consumption rule at retirement. Consumers effectively have strong bequest motives under these assumptions.

Table 3.6: BENCHMARK CERTAINTY CASE, $\rho = 0.4897$

GROUP	δ (percent)	S.E.
TOTAL	2.56	0.045
EDUCATION		
Some Highschool	3.12	0.125
Highschool Graduate	3.12	0.074
Some College	3.02	0.092
College Graduate	2.94	0.111
Graduate School	2.78	0.128
OCCUPATION		
Managerial and Prof.	2.71	0.083
Tech., Sales, Admin.	2.89	0.109
Precision Prod., Craft	2.95	0.119
Operators, Laborers	2.95	0.102
Self-Employed	2.99	0.169

Note: estimation based on first differences of profiles from Figures 3-11 and 3-12

3.5.2 Disaggregated Evidence

We next fit the model to each occupation and education cell separately. To begin with, we present the results for the CEQ LCH case across education and occupation groups in Table 3.6. Note that here and subsequently, we do not incorporate possible cross-cells correlation.⁵⁴

Again, the results indicate that the discount rate is accurately estimated and is roughly equal to the interest rate. Higher education levels tend to have a higher discount factor. Figures 3-15 and 3-16 present these fitted profiles. One can see the increase in the slope for higher educational groups.

To estimate across cells using our stochastic model, we simply follow the procedure described above. However, due to computing power constraints, we are unable to search over the entire (β, ρ) space for all cells. Since the results are more sensitive to the discount factor, we fix the intertemporal elasticity of substitution at its aggregate value, $1/0.49 = 2.04$, and search across values of the discount rate. Each cell's optimization is

⁵⁴Note also that the same individuals are allocated both to an education and an occupation cell. Therefore, results across education/occupation are not independent.

Table 3.7: ESTIMATES FROM THE STOCHASTIC MODEL, $\rho = 0.4897$

GROUP	β	S.E.	δ	S.E.	χ^2
EDUCATION					
Some Highschool	0.9607	(1.33 10^{-5})	4.09	(0.001)	256.73
Highschool Graduate	0.9592	(2.84 10^{-5})	4.25	(0.003)	107.41
Some College	0.9629	(2.37 10^{-4})	3.84	(0.025)	85.25
College Graduate	0.9679	(1.62 10^{-5})	3.31	(0.002)	115.36
Graduate School	0.9656	(2.00 10^{-5})	3.56	(0.002)	165.96
OCCUPATION					
Managerial and Prof.	0.9621	(1.75 10^{-4})	3.94	(0.018)	118.11
Tech., Sales, Admin.	0.9669	(2.49 10^{-5})	3.42	(0.003)	74.92
Precision Prod., Craft	0.9641	(4.60 10^{-5})	3.72	(0.005)	123.43
Operators, Laborers	0.9655	(4.29 10^{-5})	3.57	(0.005)	36.53
Self Employed	0.9555	(6.57 10^{-3})	4.66	(0.719)	67.09

Note: MSM estimation in levels over β . Cell size given in Table 3.1. The last column reports a test of the overidentifying restrictions distributed as a Chi-squared with 38 degrees of freedom. The critical value at 5% is 53.10.

run with a different income profile (Figures ?? and ??), income uncertainty (Table 3.3), and retirement consumption rules (Table 3.4), while we impose a constant probability of zero income ($p = 0.5\%$) and real interest rate ($R = 1.03$). The results are summarized in Table 3.7. Except for the last occupation cell (Self-Employed), the parameters are remarkably close to those estimated using the aggregate profile. As we observed with the benchmark case, the discount rate decreases weakly with education, and ranges from 3.31% to 4.25%. The associated fitted profiles are displayed in Figures 3-22 and 3-23. The fit is quite good for most cells, except Self-Employed and College Graduates.⁵⁵ For all groups except the first educational group, the consumption profile is humped. As in the aggregate case, we reject the overidentifying restrictions in all cases. We conclude, given the similarity in the estimates that our results are quite robust to heterogenous shocks and processes.

Further, as in the aggregate case, the estimated discount factors are consistently

⁵⁵Note that for all cells the standard errors are smaller than for the aggregate estimation. This simply reflects the fact that ρ is fixed. The correlation between ρ and β in the aggregate estimation increases the standard errors. We hope in the future to be able to estimate both parameters for each cell.

Table 3.8: ROBUSTNESS CHECKS, $\rho = 0.4897$

		β	S.E.	δ	S.E.	χ^2
1	$R = 1.05, \gamma_0 = 0.493; \gamma_1 = 0.035$	0.9441	$(3.28 \cdot 10^{-4})$	5.92	(0.036)	266.25
2	$p = 0.05;$	0.9631	$(8.56 \cdot 10^{-5})$	3.83	(0.009)	275.07
3	$\sigma_u^2 = 0; \sigma_n^2 = 4.79 \cdot 10^{-4}$	0.9632	$(6.26 \cdot 10^{-5})$	3.81	(0.007)	435.67
4	$\gamma_0 = 0.535; \gamma_1 = 0.074$	0.9612	$(5.4 \cdot 10^{-5})$	4.03	(0.006)	471.92
5	$x_0 = 0.0$	0.9623	$(1.9 \cdot 10^{-4})$	3.91	(0.020)	489.96
6	Large Income; $\gamma_0 = 0.241; \gamma_1 = 0.035$	0.9899	(0.1676)	1.01	(17.10)	30596

Notes: *MSM* estimation in levels over β . Cell size is 43031. The last column reports a test of the overidentifying restrictions distributed as a Chi-squared with 38 degrees of freedom. The critical value at 5% is 53.10. (1) assumes an interest rate of 5% and recomputes accordingly the last working period consumption rule; (4) computes the last working period consumption rule using asset data from Venti and Wise (1993). Large income in (6) includes mortgage payments and pension contributions.

lower in our estimation of the stochastic model, than in the CEQ LCH baseline case. We can reject that the interest rate and the discount rate are equal in all cells, except Self-Employed.

3.5.3 Robustness Checks and Extensions

Our estimation procedure depends on the calibrated parameters, χ . In this section we investigate the robustness of our results to these parameters. Due to computing constraints, we have only been able to check the robustness with respect to the discount factor. In what follows, we again maintain a constant intertemporal elasticity of substitution equal to 2.04. The results are reported in Table 3.8.

The estimate of the discount rate is, as mentioned, most sensitive to our choice of interest rate. This is not surprising, since, early in life the difference between the interest rate and the discount rate is a key determinant of whether consumers exhibit buffer stock behavior. Late in life, consumers will behave in a manner more consistent with the CEQ LCH, in which the change in consumption is driven by the product βR .

So far we assumed a real interest rate of 3% a year. Although this is close to the long run real rate on liquid assets, some savings are held in longer term assets, yielding

on average higher returns. Therefore, we re-estimate the discount rate assuming a real interest rate of 5%.⁵⁶ The estimate of the discount factor is now $\beta = 0.9441$, with a standard deviation of $3.28 \cdot 10^{-4}$. This implies a discount factor of 5.92%, almost exactly two hundred basis points higher than our previous estimate. Thus it is important to note that our results do not provide a tight estimate of the discount factor per se. However, *our model gives precise estimates of the difference between the interest rate and the discount rate*. This indicates that we capture mostly the substitution effect. The simulated profile reported in Figure 3-24.A is roughly similar to our benchmark case, although consumption is smaller later in life to reflect the smaller γ .

We next check the robustness of our results to the probability of zero income, p . A higher p , by increasing uncertainty, should lead to a larger buffer stock, implying less life-cycle savings late in life. Thus, an increase in p is likely to yield a higher discount rate, to counteract the increase in precautionary savings. Our results indicate that this effect is quite weak. When $p = 5\%$, a tenfold increase, the discount rate actually decrease marginally to 3.83%. On the other hand, given unemployment benefits, government assistance programs, one might argue that households can never experience zero income. Decreasing p has even smaller effects on the estimated discount rate. Looking at the simulated profile (Figure 3-24.B), we see that the change in p affects mostly consumption early in life. After age 40, asset level can buffer transitory fluctuations in income.

Next, we investigate the sensitivity of our results to the variance of permanent and transitory components, σ_u^2 , and σ_n^2 . We decreased the uncertainty faced by the agent from both sources of shocks. Intuitively, this should lead to a smaller buffer and should give a lower estimate of the discount rate. Our test eliminates altogether transitory shocks and reduces the variance of the permanent shocks to $\sigma_n^2 = 4.8 \cdot 10^{-3}$. Our estimated discount rate, 3.81%, confirms our intuition and reemphasizes that our results are reliant upon the underlying individual uncertainty. Figure 3-24.C displays the simulated profile. Consumption is higher early in life and lower later, as the lower buffer translates into

⁵⁶With a 5% interest rate, we recompute the last period consumption rule as $\gamma_0 = 0.493$; $\gamma_1 = 0.035$. Thus illiquid assets are more important and consumption is less sensitive to current cash on hand.

smaller accumulated assets.

An important assumption of our model concerns the retirement rule. As described in section 3.4.2, the consumption rule at retirement is calibrated from CEX data, using both income and asset reports. To the extent that asset data are not accurate, our consumption rule is likely to be mismeasured. This, in turn, will affect the life-cycle profile as we near retirement. In order to test robustness to our hypothesis, we calibrate the retirement consumption rule using asset data reported in Venti and Wise (1993). The resulting values are $\gamma_0 = 0.535$; $\gamma_1 = 0.074$. Looking at equation (3.21), this indicates both a substantially smaller asset accumulation (as γ_1 increases) and a larger share of illiquid assets (as γ_0 increases). With a larger γ_0 , the agent can rely on illiquid saving at retirement time. This should lead to a smaller liquid asset accumulation and to a lower discount rate. A larger γ_1 has ambiguous effects. It increases the level of consumption out of cash on hand in the last working period. This is compatible both with a higher discount rate (as the household will consume more of a given wealth) and a lower discount rate (if the households accumulates more wealth). Thus, the actual consumption profile may indicate more or less preference for the present. We estimate a discount rate of 4.03%, slightly higher than the benchmark estimate. This suggests that the presence of illiquid assets play an important role in the precision of our estimates and that the increase in γ_1 dominates. As Figure 3-24.D shows, assumptions about the last working period consumption rule affect the simulated profile substantially.

We then check the validity of the assumption $x_0 = 0.3$. As for the retirement rule, this is likely to affect our estimates by shifting the consumption profile at young ages. We reestimate our aggregate problem assuming that $x_0 = 0$. The estimated β , 0.9623, is extremely close to our original one. With a lower initial level of cash on hand, households build their buffers early in life. This results in slightly lower consumption for the first few years. However, from then on, the consumption profile is similar. Thus, this assumption does not have a large impact on our estimates.

Our estimation procedure is also extremely sensitive to our assumed income profile. Permanent income growth is a key variable and determines to a large extent buffer-

stock behavior, as we have demonstrated. Our definition of income subtracts pension contributions and mortgage payments, as they are likely to reflect illiquid saving, for which our theory is ill-equipped. However, it is possible in some circumstances to draw on voluntary pension contribution. Similarly, housing wealth is not entirely illiquid. Thus we reestimate our model assuming that these components of income contribute to liquid savings instead of illiquid savings. The associated values of the consumption rule at retirement are $\gamma_0 = 0.241$; $\gamma_1 = 0.035$. Since pension contributions are now part of liquid savings, the accumulated illiquid savings and γ_0 are lower. In effect, this amounts to shifting upwards the income profile and increases substantially measured savings. Not surprisingly, this can only be matched by a lower discount rate. The estimated discount rate is 1.01%, well below the interest rate. The parameters are poorly estimated and the overidentifying restriction is enormous. This increase in savings can only be matched by assuming that consumers are extremely patient. However, this fails to capture the life-cycle pattern of consumption: the simulated profile grows exponentially.

We conclude that, except with respect to our definition of liquid income, our estimates and inference are reasonably robust to the calibration of our model.

3.6 Conclusion

Macroeconomic models generally represent the consumer as an infinitely-lived, rational, representative agent, who behaves in accordance with the Permanent Income Hypothesis. Some analyses provide explicit microeconomic justification for this assumption, by deriving an insurance system which protects individuals from any idiosyncratic consumption risk (e.g. Rogerson (1988)) or by assuming that individual budget constraints never bind, so that aggregate behavior mimics individual behavior (e.g. Barro (1974)). Other models incorporate certain forms of individual heterogeneity when aggregating (e.g. Blanchard (1985)). The resulting representative agent will not, in general, have the same characteristics as individual consumers.

However, nearly all currently employed macroeconomic models include some form of a

representative agent facing representative shocks. Such a representative agent framework constitutes an extremely powerful tool, mainly because the restrictions that optimizing behavior place on individuals will carry over to the aggregate economy. In this chapter, while we work in an explicitly partial equilibrium environment, we find important and substantial deviations from the canonical representative consumer model.

Using individual-level data to construct average profiles of income and consumption over the working lives of households, we demonstrated that consumption remains hump-shaped, even after controlling for family and cohort effects. We then developed a model of consumption behavior embedding realistic levels of income uncertainty and estimated *individual* consumption functions using the Method of Simulated Moments. The model fits well and yields tight estimates of the discount rate and intertemporal elasticity of substitution. To the best of our knowledge, this is the first study which uses explicit individual uncertainty and the life-cycle profile of consumption to identify structural parameters of the utility function. The results indicate that consumers hold only a buffer stock of liquid assets in order to offset labor income fluctuations, until around age 42. Then, they start saving actively for retirement purposes. These two phases in consumer behavior are quite distinct and are at the heart of our identification procedure.

This age-heterogeneity of consumer behavior has important implications for aggregate consumption. In particular, it may rationalize Campbell and Mankiw (1989) finding that roughly 40% of all agents are “hand to mouth”. In our interpretation, this may simply reflect the consumption of young households. In later work, we plan to investigate in more detail the aggregate implications of the age-heterogeneity of consumers.

Figure 3-1: Consumption Rules without Retirement

$$\beta=0.963, \rho=0.490, \gamma_1=1, G=1$$




Figure 3-2: Consumption Rules with Retirement

$$\beta=0.963, \rho=0.490, \gamma_0=0.384, \gamma_1=0.048, G=1$$




Figure 3-3: Sample Individual Profile

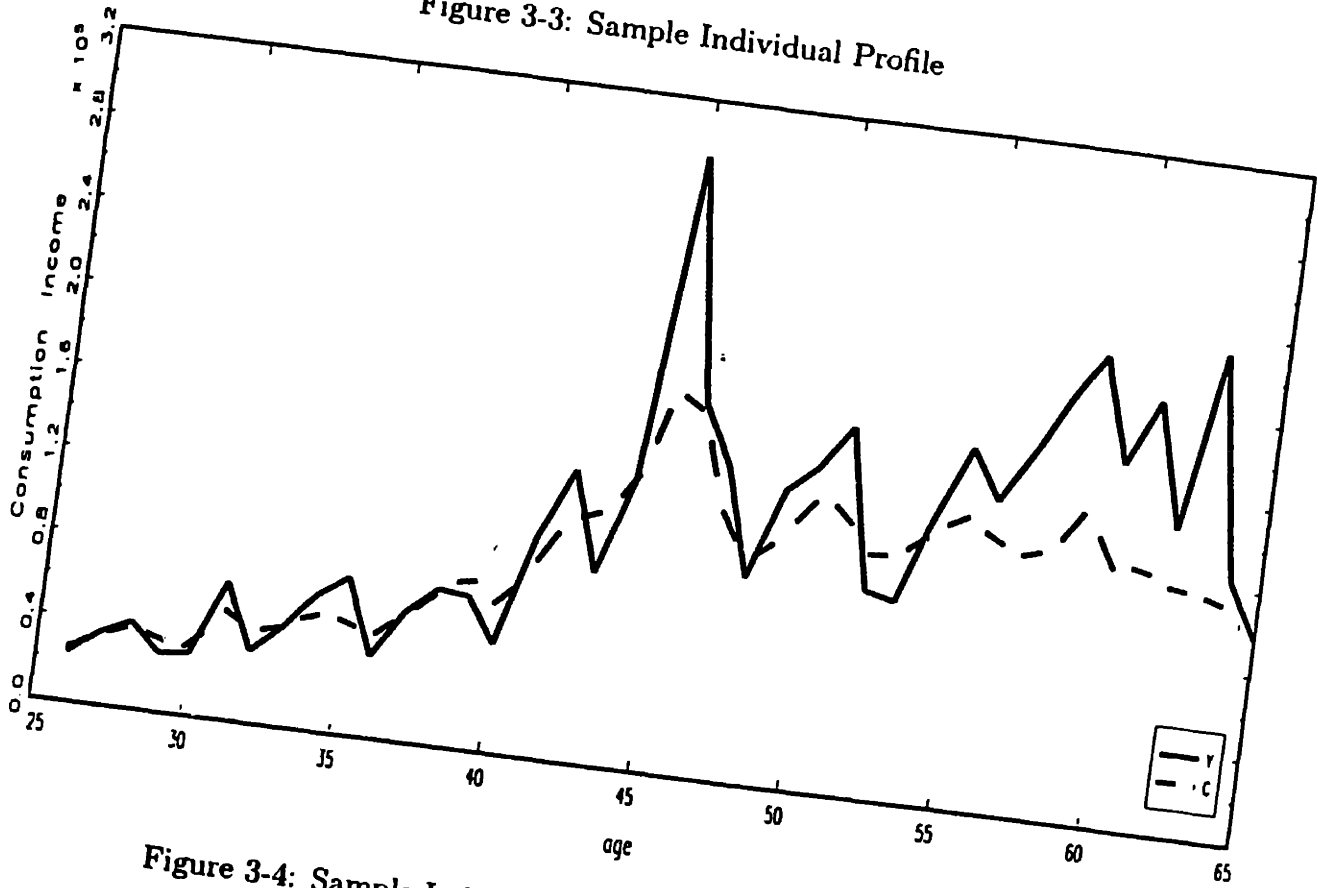


Figure 3-4: Sample Individual Profile with Zero Income Shock

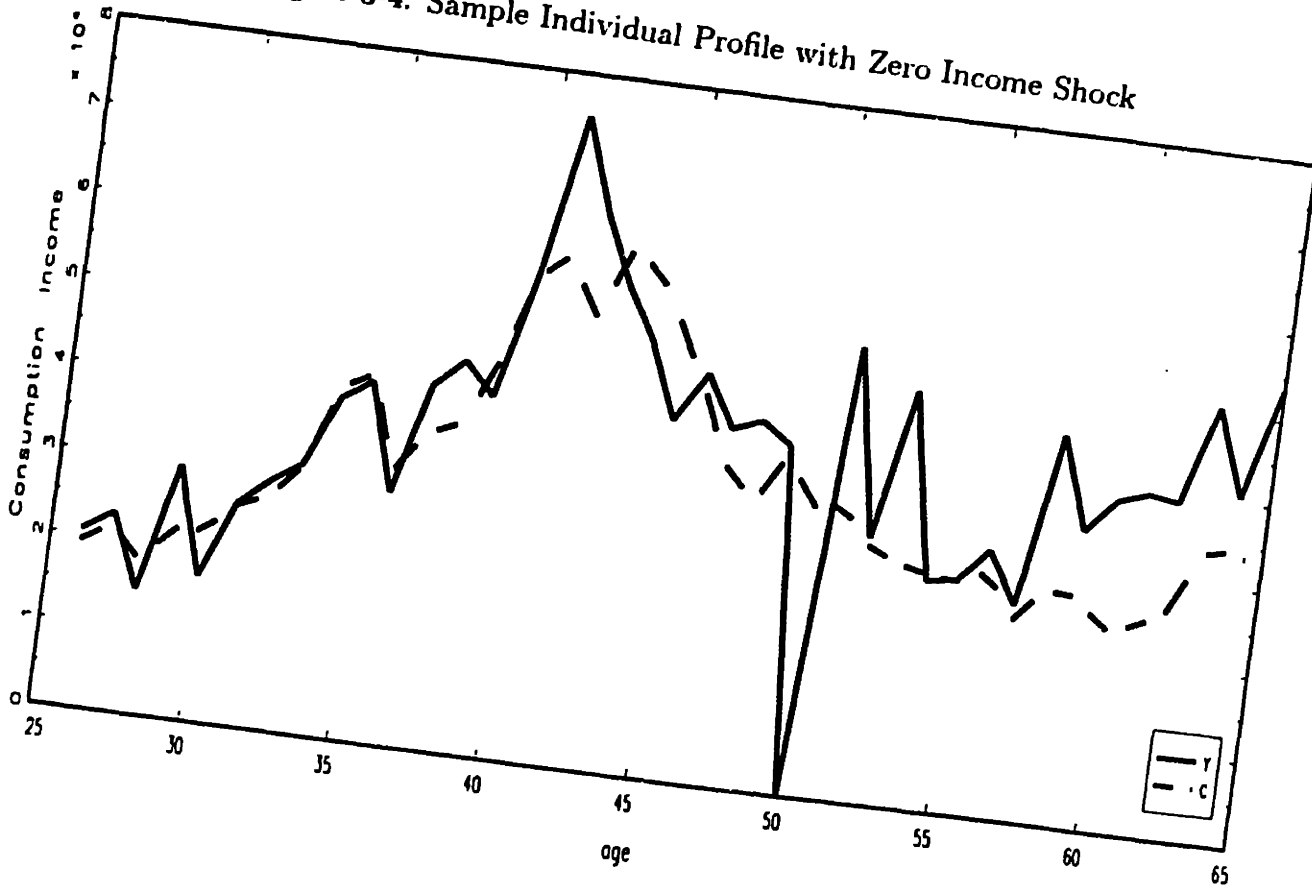


Figure 3-5: Expected Cash on Hand

$$\beta=0.963, \rho=0.490$$




Figure 3-6: Household Consumption and Income over the Lifecycle




Figure 3-7: Family Size over the Lifecycle




Figure 3-8: Consumption and Income with and without Cohort Adjustments




Figure 3-9: Consumption and Income with and without Family Adjustments

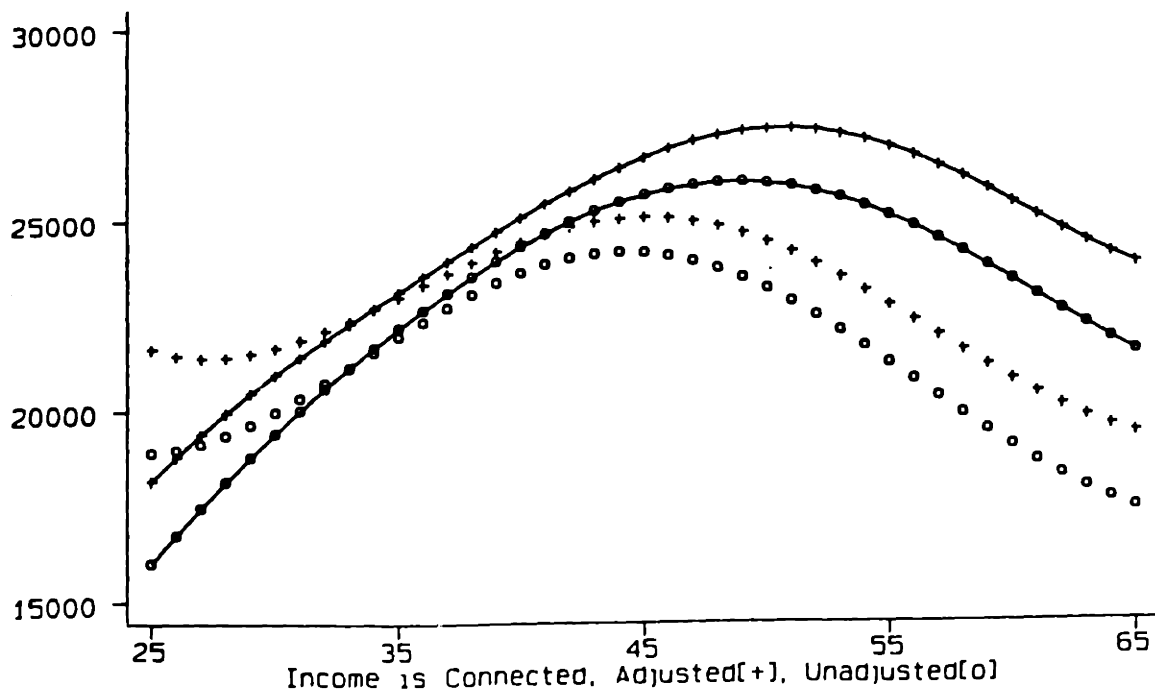


Figure 3-10: Total Liquid Asset to Income Ratio



Figure 3-11: Consumption(O) and Income(+) over the Lifecycle by Education

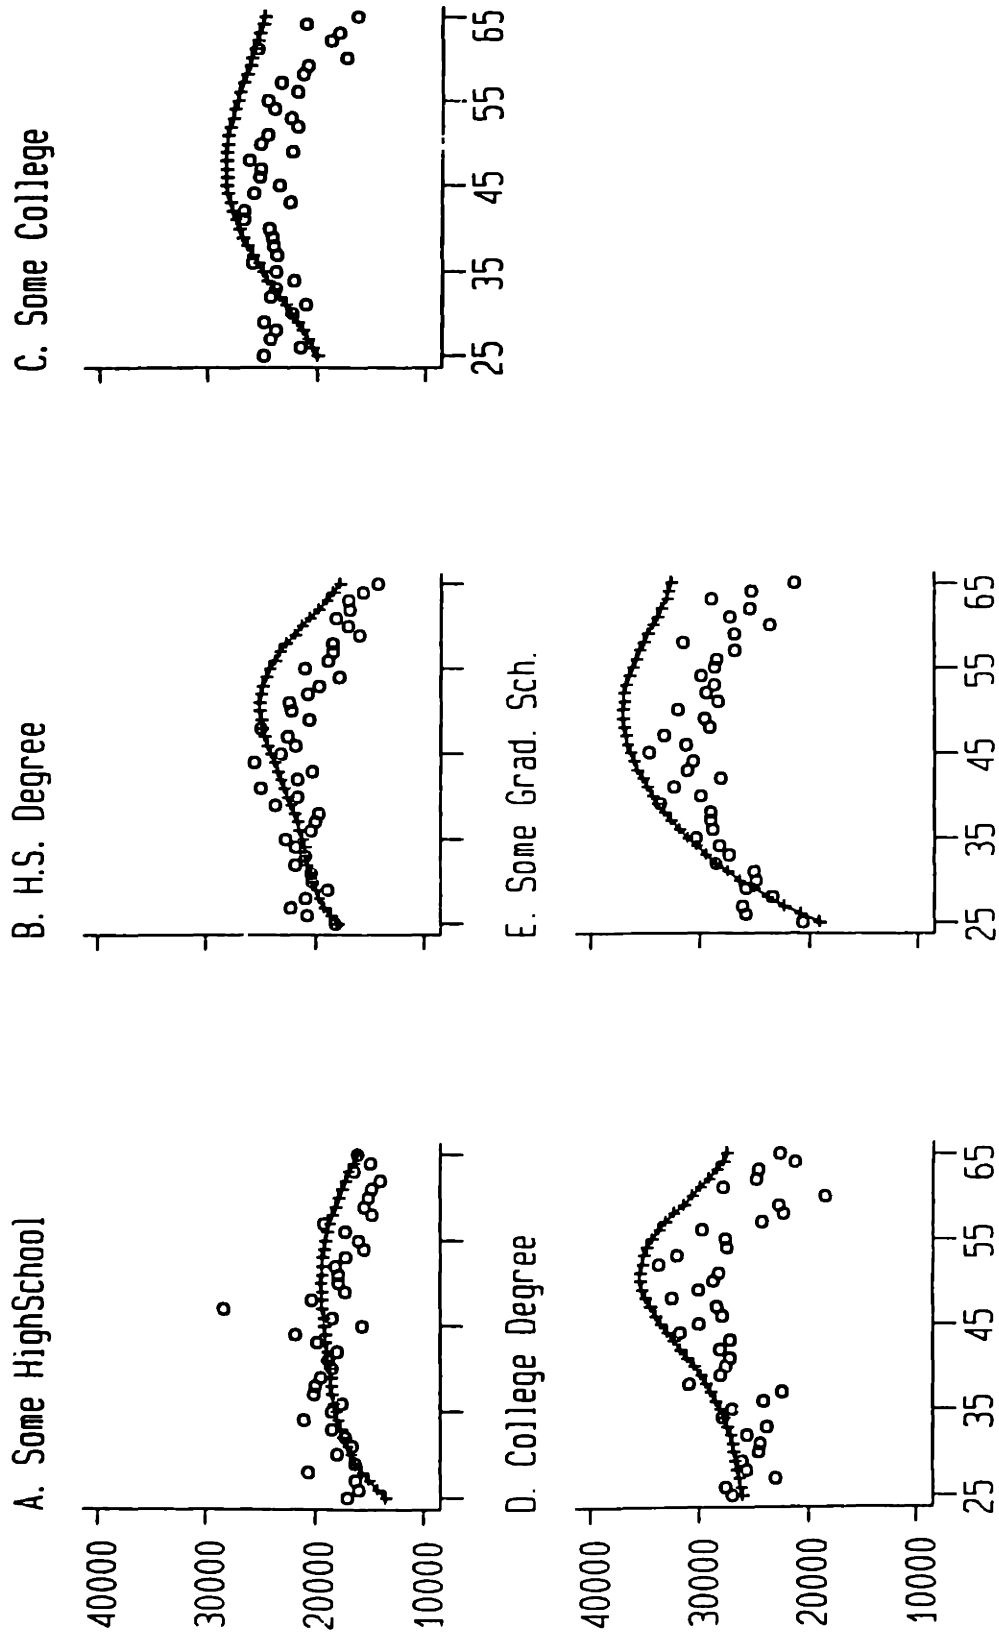


Figure 3-12: Consumption(O) and Income(+) over the Lifecycle by Occupation

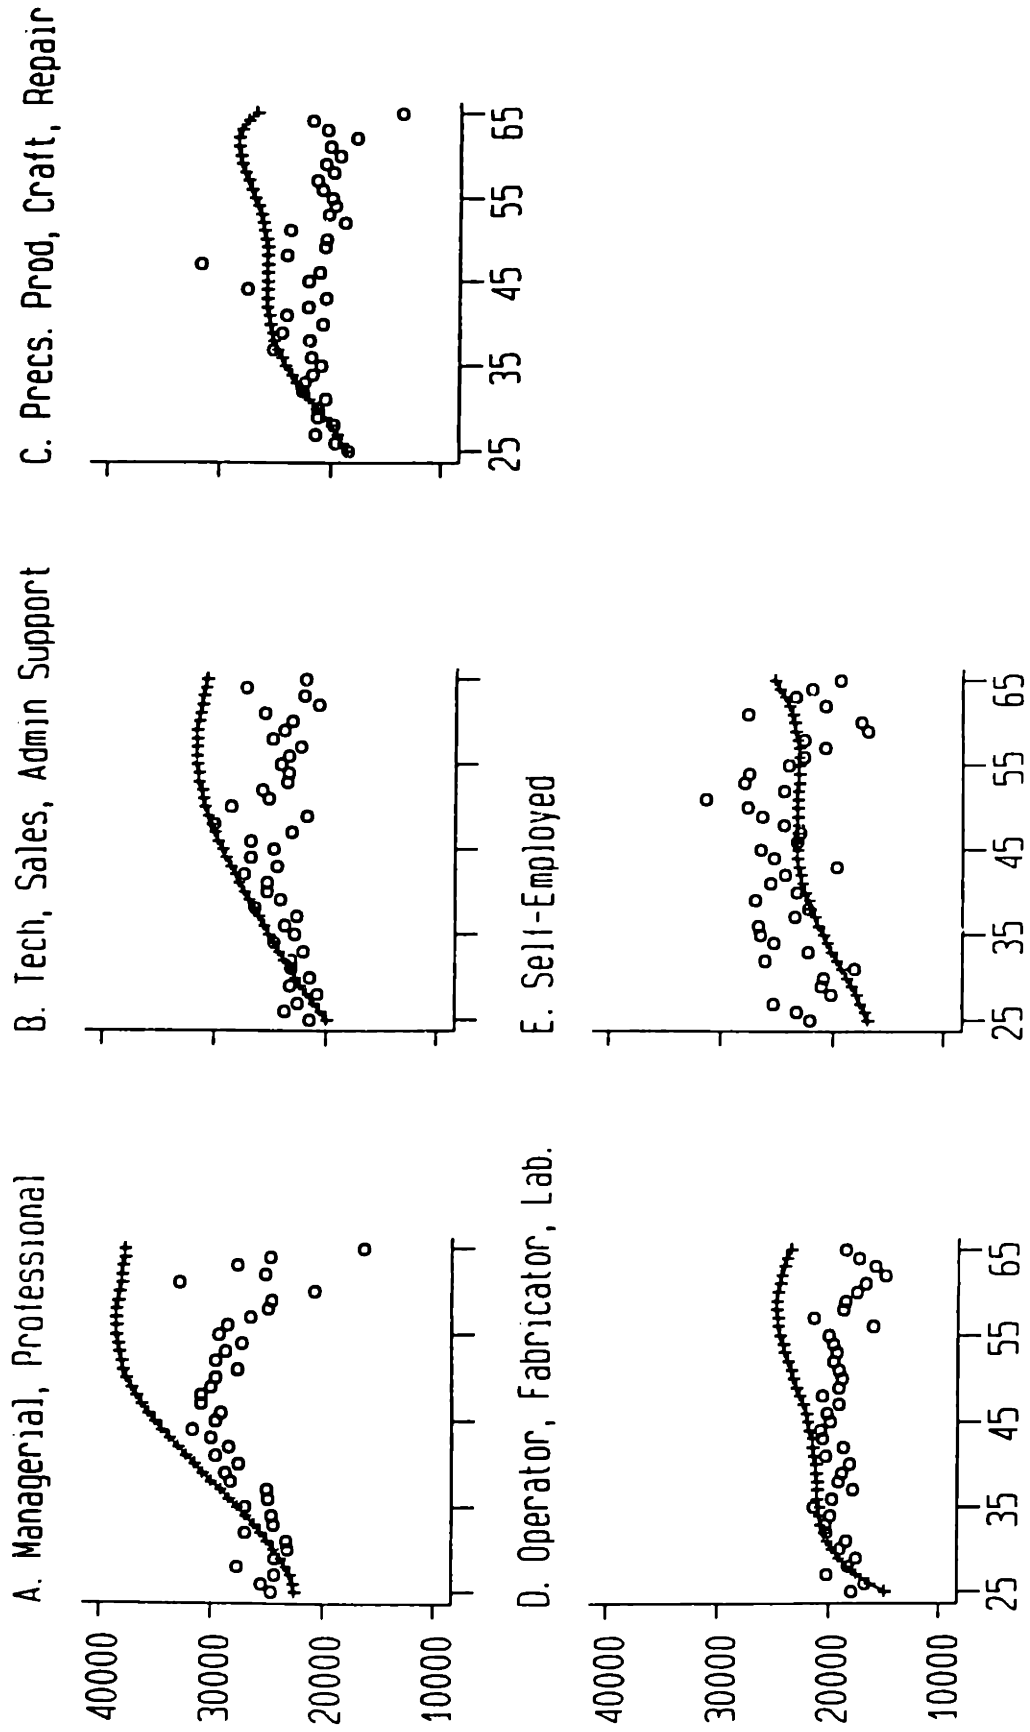


Figure 3-13: Rescaled Smoothed Consumption over the Lifecycle




Figure 3-14: Income, Consumption and Consumption Predicted by the CEQ LCH

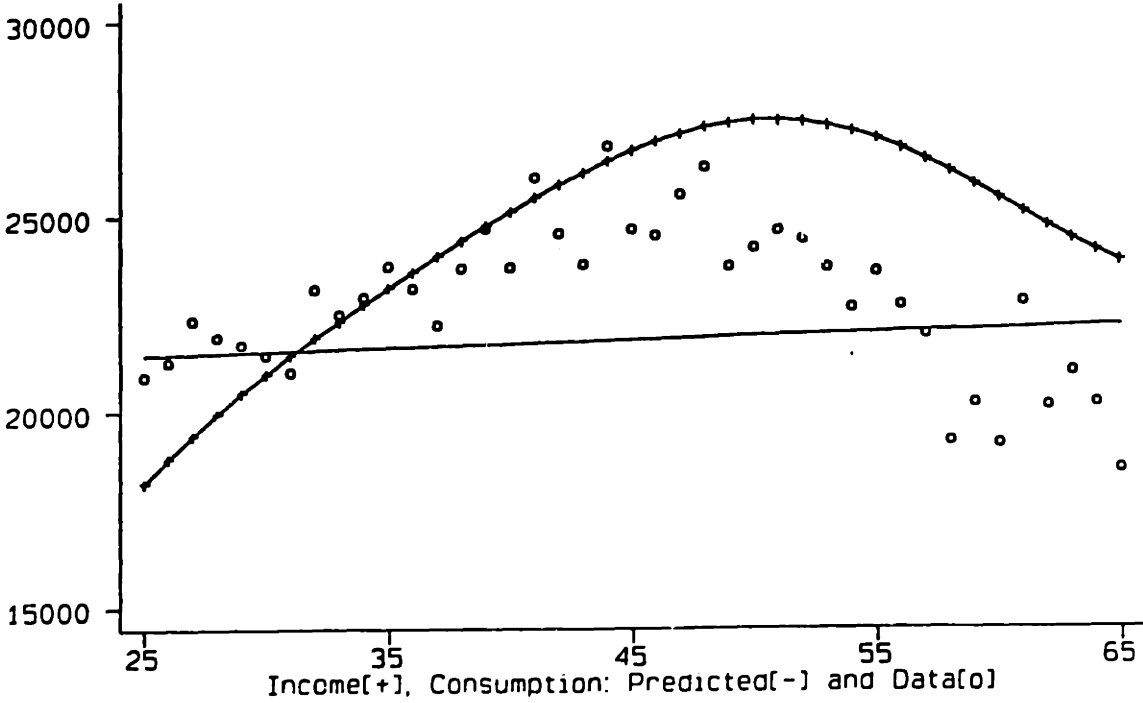




Figure 3-15: Income, Consumption and CEQ Consumption by Education


A. Some HighSchool




B. H.S. Degree



C. Some College



D. College Degree



E. Some Grad. Sch.




Figure 3-16: Income, Consumption and CEQ Consumption by Occupation

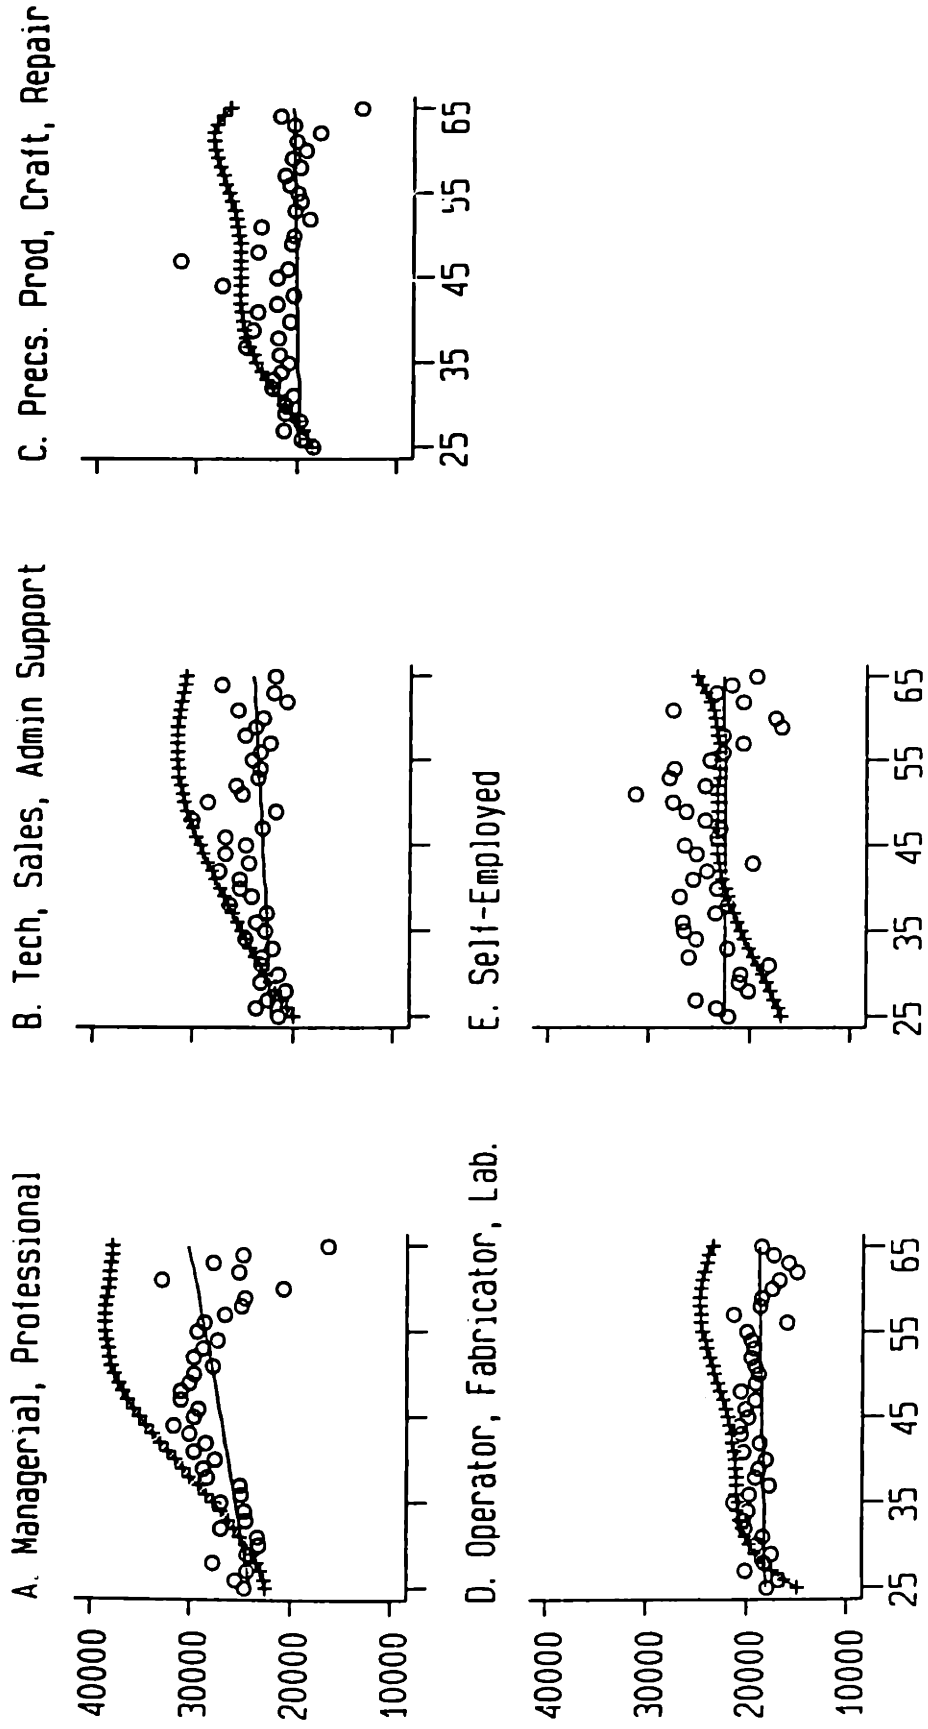


Figure 3-17: Simulated and Actual Consumption Profiles

(confidence bands), $\beta=0.963$, $\rho=0.490$




Figure 3-18: Simulated Consumption Profiles for Different Discount Rates

$\bar{R} = 3\%, \rho = 0.490$




Figure 3-19: The Predicted Savings Rate Profile




Figure 3-20: Target Cash on Hand

$\beta=0.932, \mu=1.573, R=37\%$

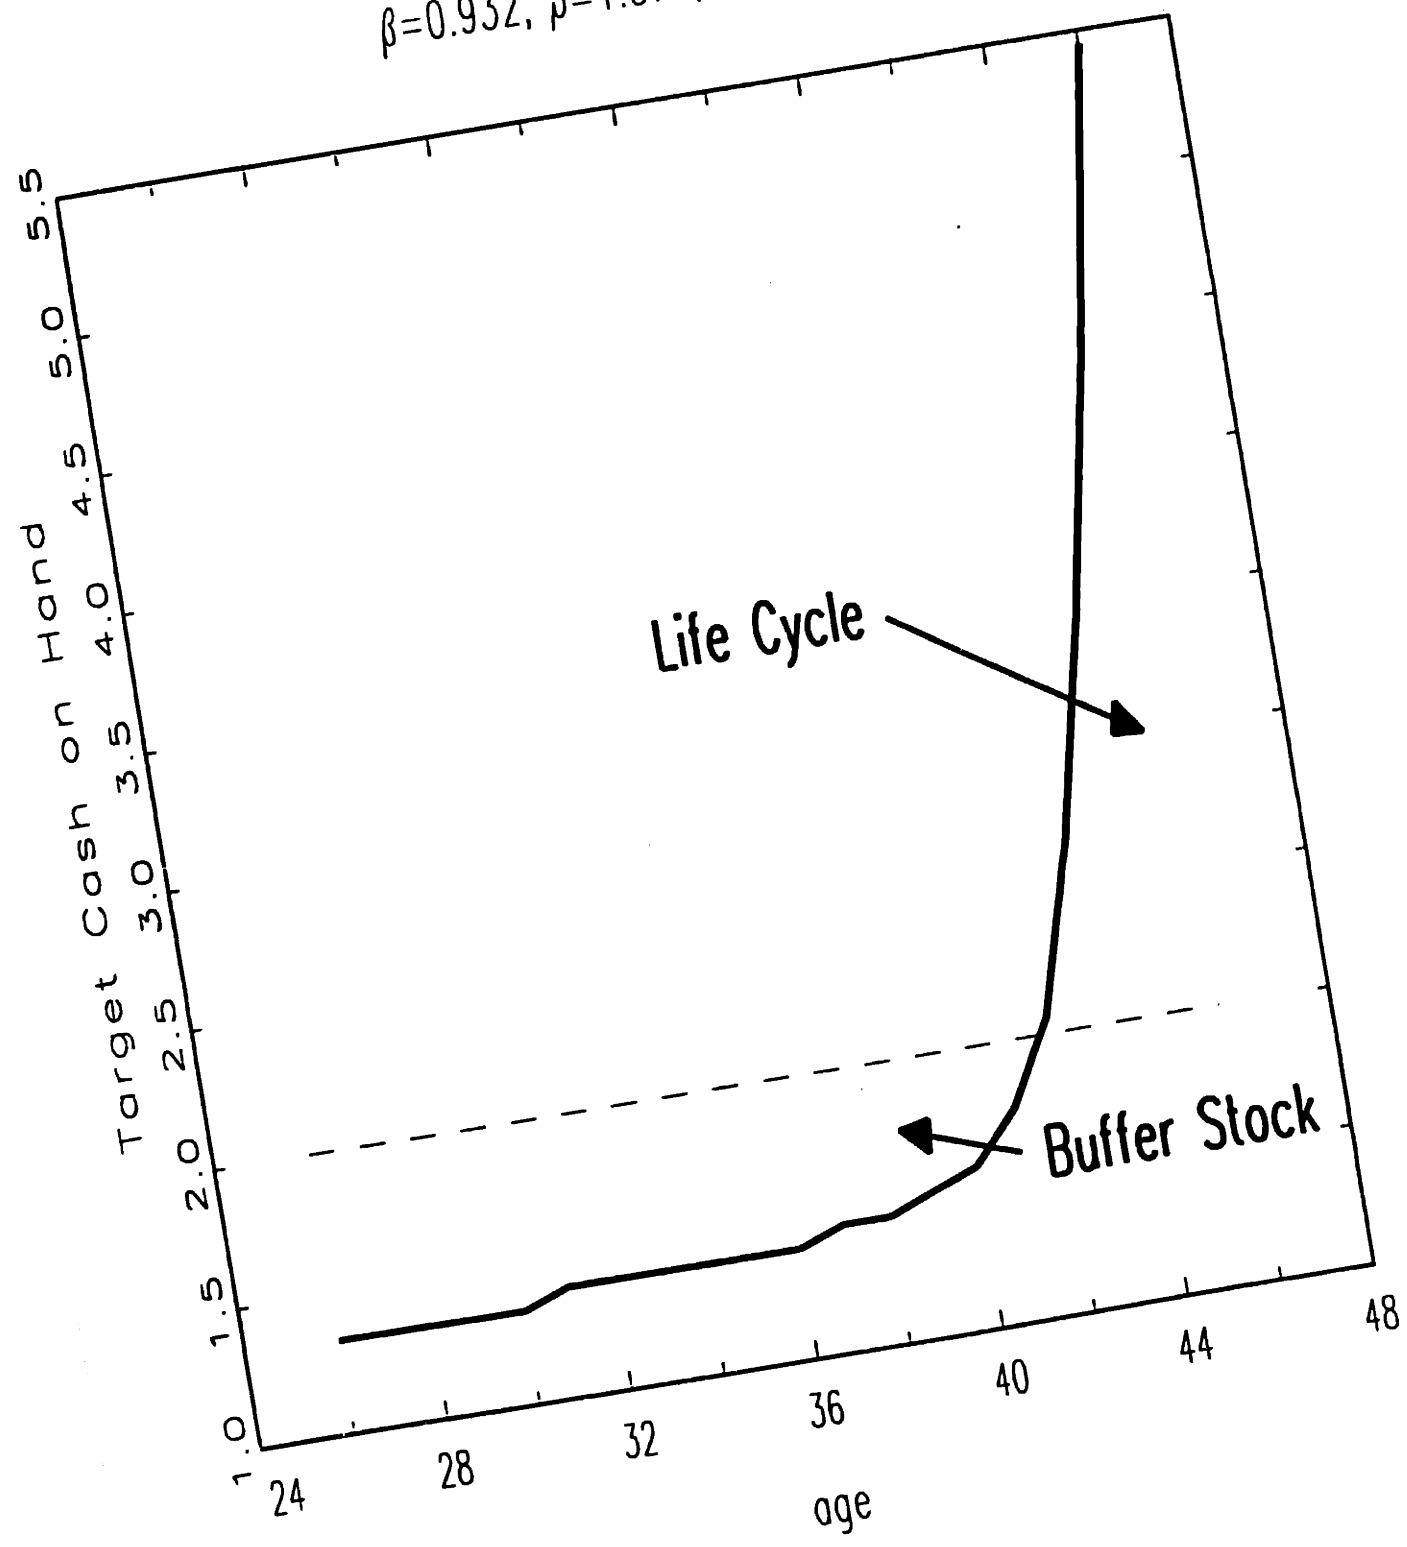


Figure 3-21: Life-Cycle and Buffer-Stock Savings

$$\beta = 0.963, \rho = 0.490; x_0 = 0$$




Figure 3-22: Simulated and Actual Consumption Profiles by Education

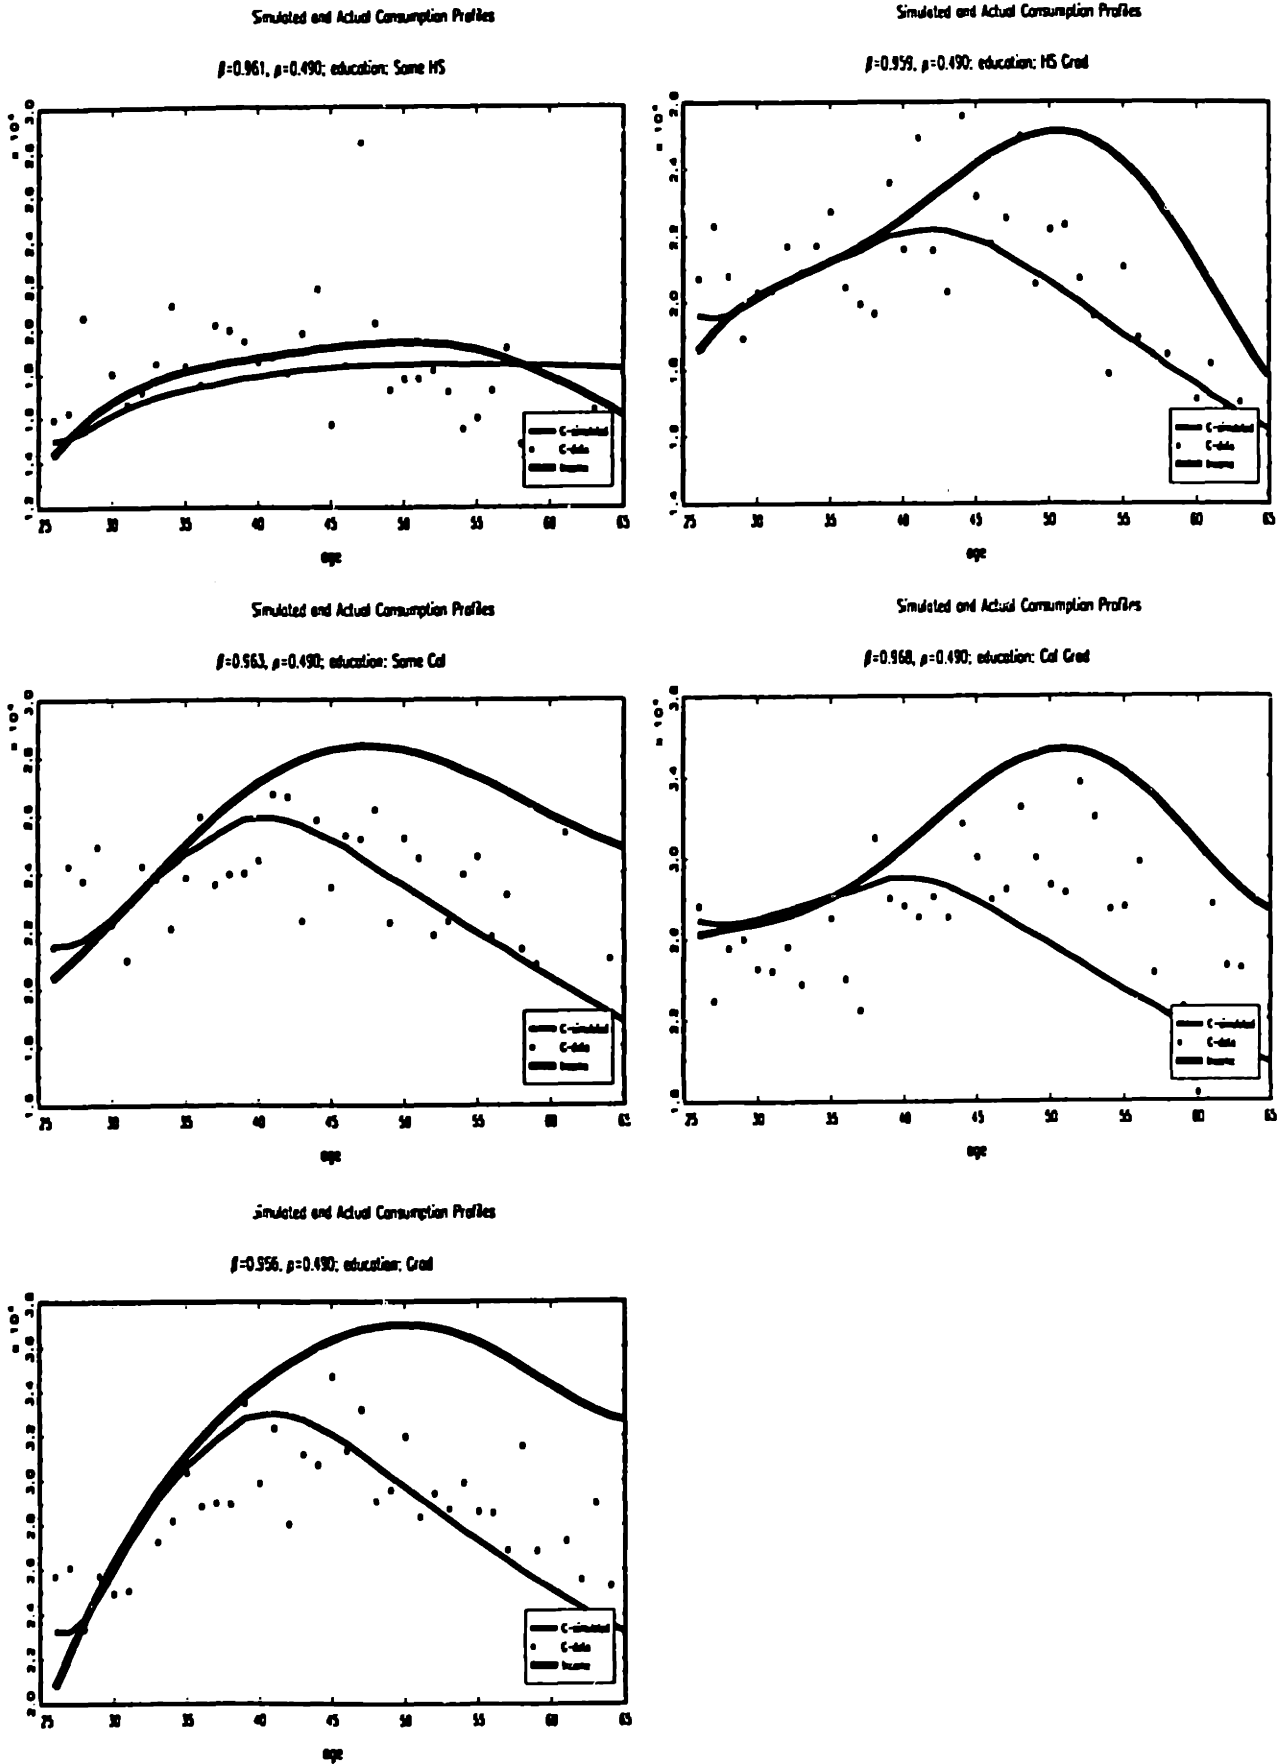


Figure 3-23: Simulated and Actual Consumption Profiles by Occupation

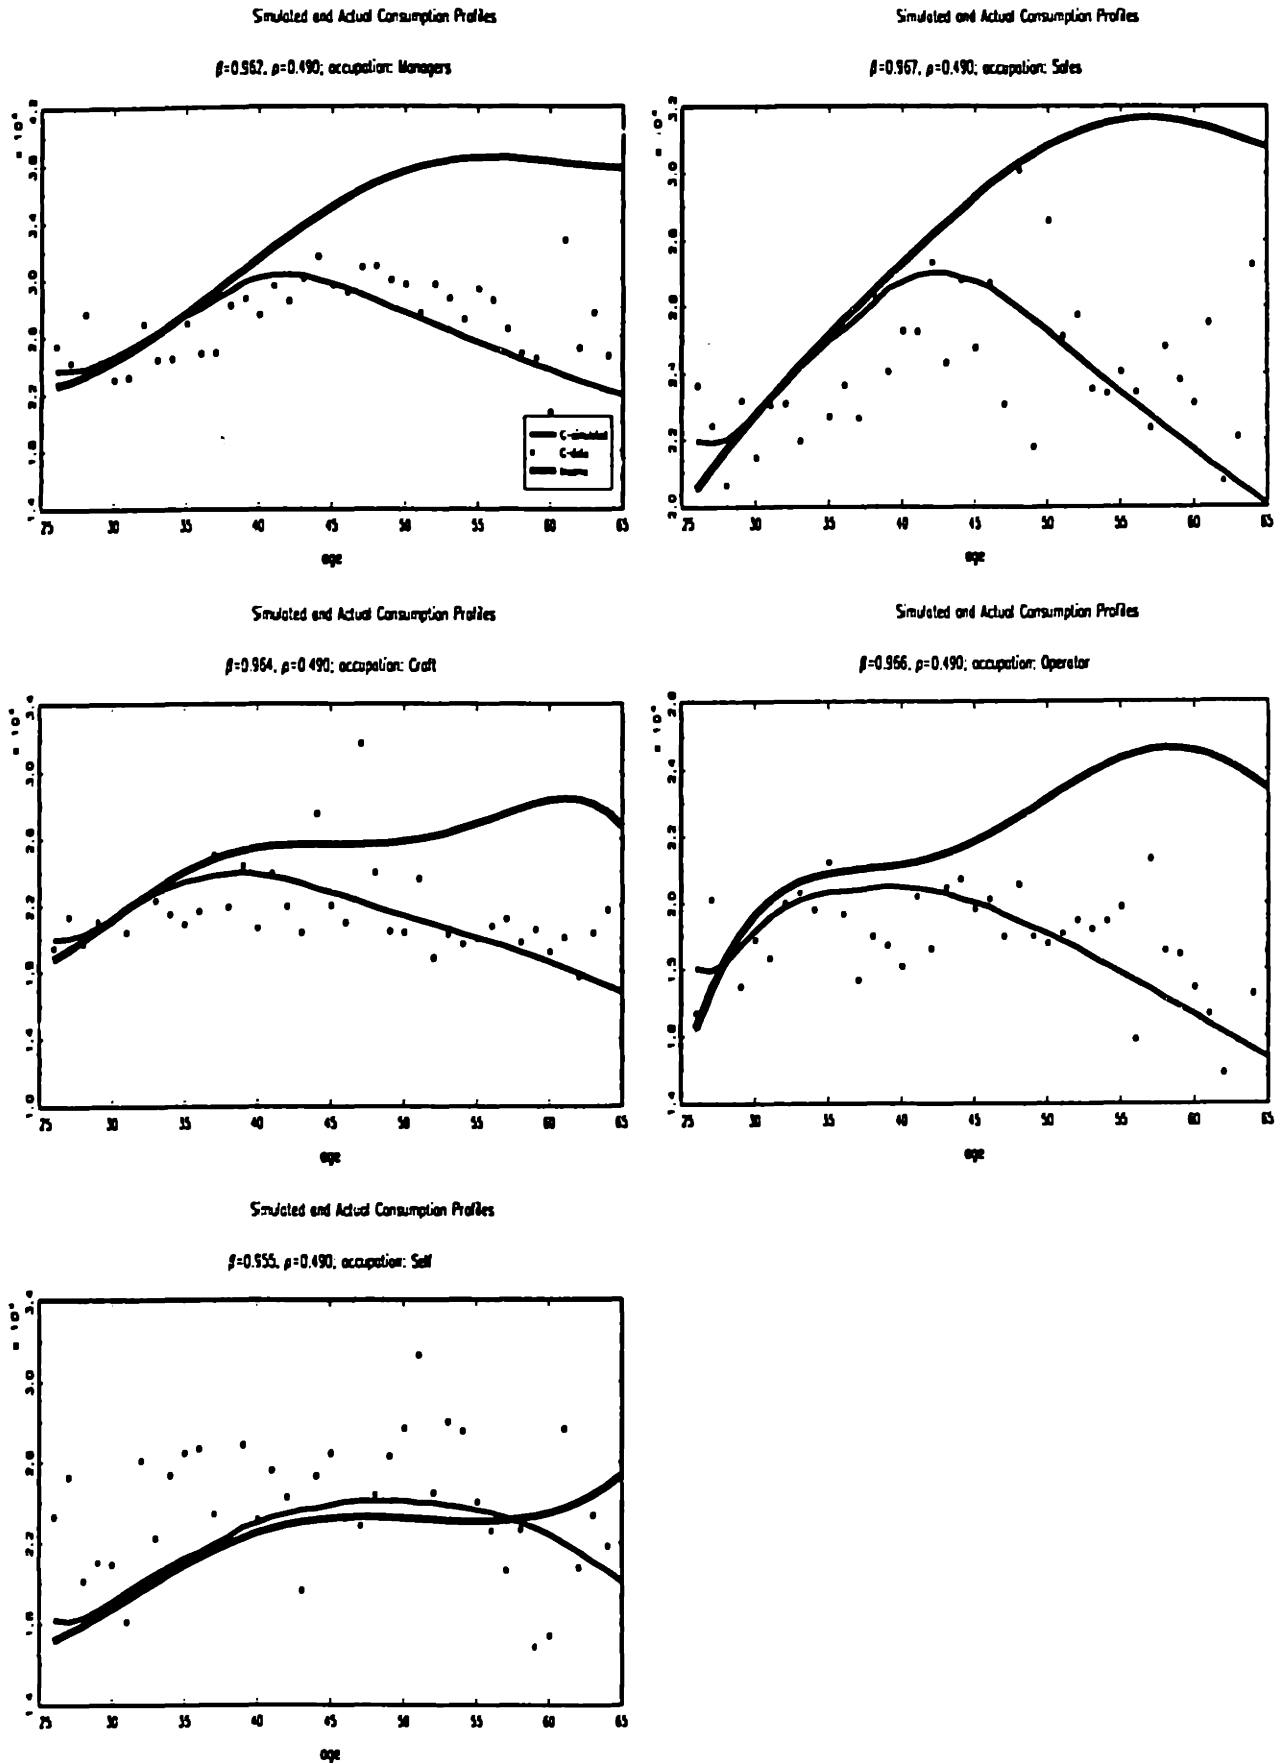
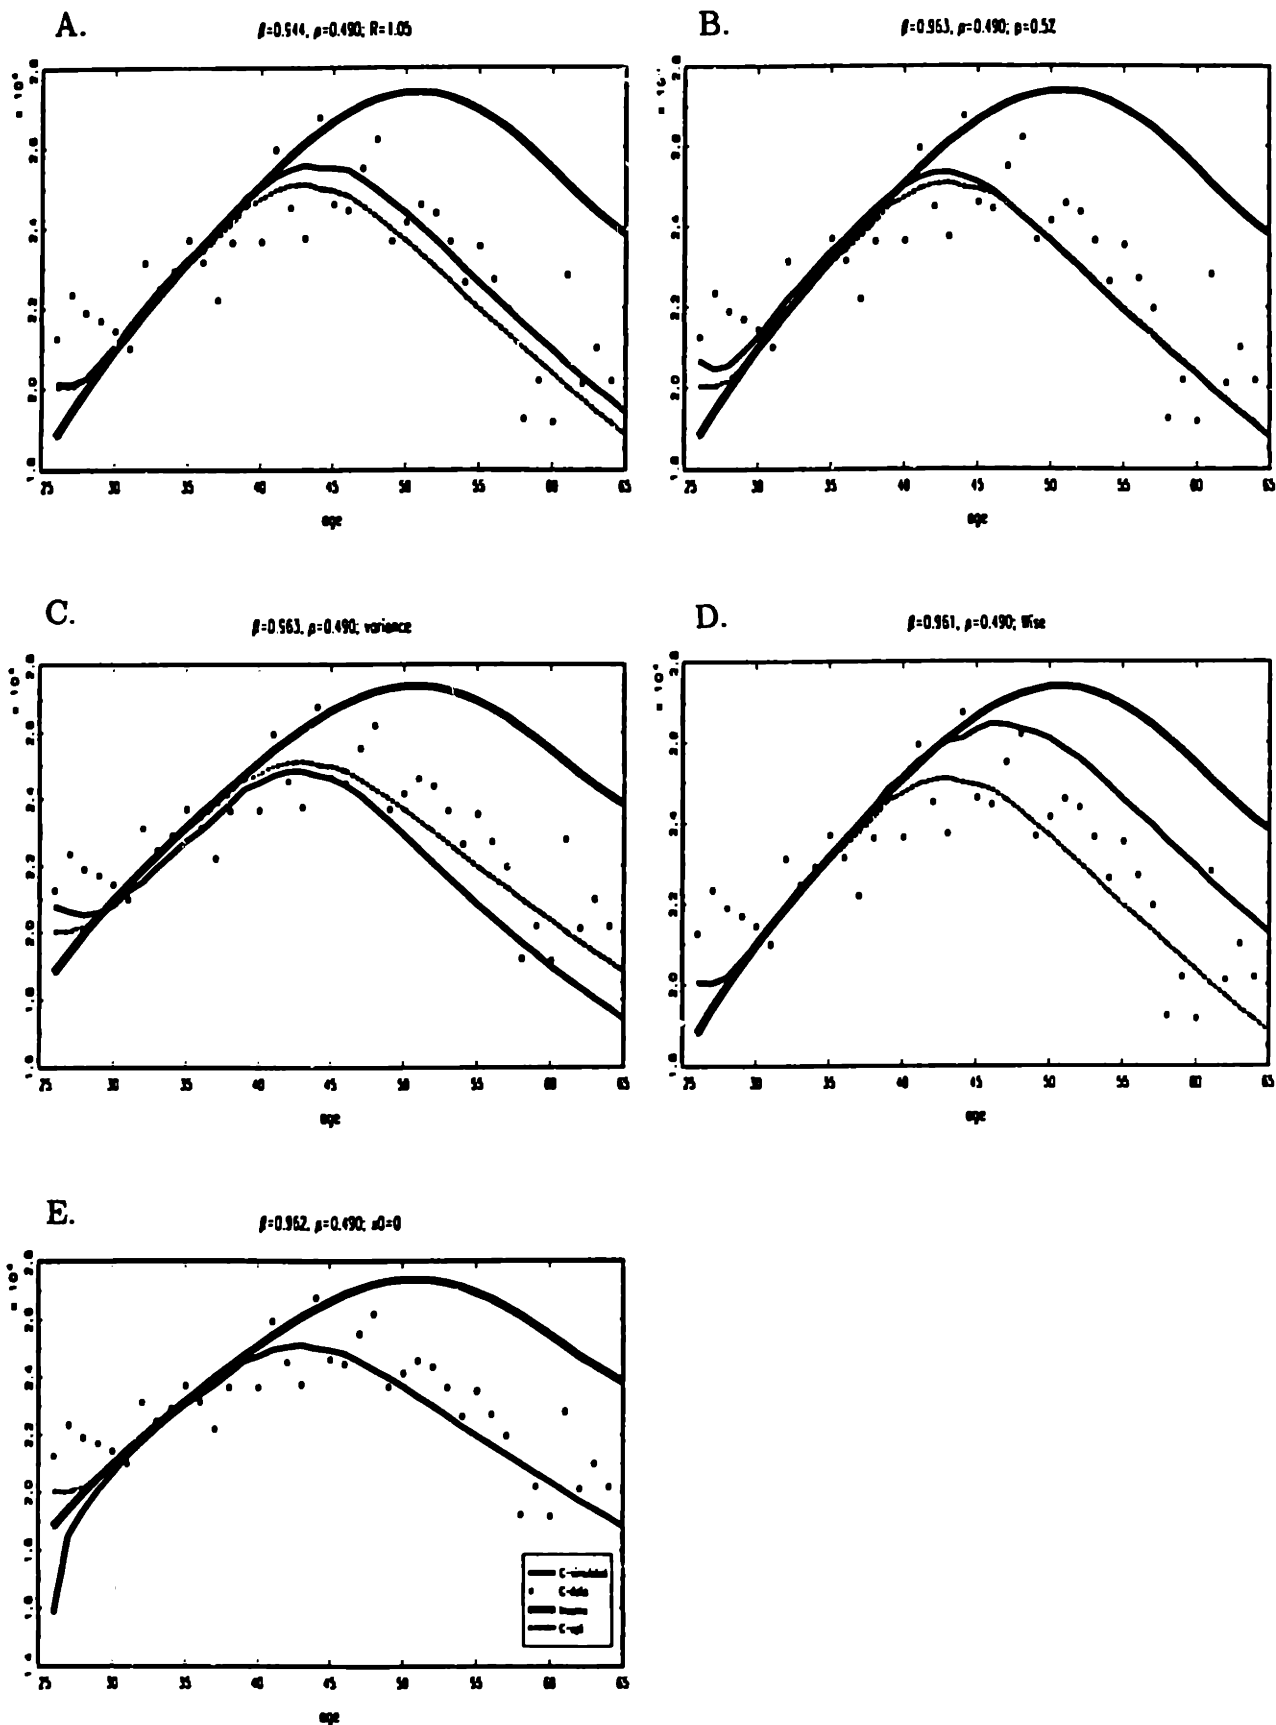


Figure 3-24: Robustness Checks: Simulated and Actual Consumption Profiles



Appendix B

Solving the Consumer Optimization Problem

In this appendix we describe our approach to solving numerically the consumer problem.

B.1 Euler equation

The algorithm exploits the recursive structure of the consumer problem by solving the Euler equation. Given a consumption rule at age $t + 1$, $c_{t+1}(\cdot)$, the algorithm solves for the consumption rule $c_t(x_t)$ that satisfies for any x_t :

$$\begin{aligned} u'(c_t(x_t)) &= \beta R E[u'(c_{t+1}(x_{t+1}))] && \text{(B.1)} \\ &= \beta R (p E[u'(c_{t+1}(x_{t+1}))|U_{t+1} = 0] + (1 - p) E[u'(c_{t+1}(x_{t+1}))|U_{t+1} > 0]). \end{aligned}$$

B.2 Gauss-Hermite quadrature

Assume for the time being that we know how to compute $c_{t+1}(\cdot)$ for all values of cash on hand. Our first problem consist in evaluating the expectation in (B.1). One can rewrite the Euler equation using the Intertemporal budget constraint,

$$x_{t+1} = X_{t+1}/P_{t+1} = (x_t - c_t) \frac{R}{G_{t+1}N_{t+1}} + U_{t+1}. \quad (\text{B.2})$$

as:

$$u'(c_t(x)) = \beta R \left(p E \left[u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}N} \right) G_{t+1}N \right) \right] + (1-p) E \left[u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}N} + U \right) G_{t+1}N \right) \middle| U > 0 \right] \right).$$

Since N and U are log normally distributed, the natural way to evaluate these integrals is to perform a two dimensional Gauss-Hermite quadrature:

$$\begin{aligned} E [u'(c_{t+1}(x_{t+1})G_{t+1}N)] &= \int u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}N} + U \right) G_{t+1}N \right) dF(N) dF(U) \\ &= \int_{-\infty}^{\infty} f_t(n, u) e^{-n^2} e^{-u^2} du dn \\ &= \sum_{i,j} f_t(n_i, u_j) \omega_{ij}, \end{aligned}$$

where $f_t(n, u) = \frac{1}{\pi} u' \left(c_{t+1} \left((x - c_t) \frac{R}{G_{t+1}} e^{-\sqrt{2}\sigma_n n} + e^{-\sqrt{2}\sigma_u u} \right) G_{t+1} e^{-\sqrt{2}\sigma_n n} \right)$. The weights ω_{ij} and nodes n_i, u_j are tabulated in Judd (1993). In practice, we performed a quadrature of order 12.

One can then find the root of the Euler equation at any point x using a standard Newton method. In practice, we constrain the root to be positive and less than x , the current cash on hand. As discussed in the text, this restriction is always satisfied when there are no illiquid assets. Since illiquid assets cannot be borrowed against, it is also satisfied in their presence.

B.3 Consumption rules

We initialize the algorithm with the consumption rule at retirement: $c_T(x_T) = \gamma_0 + \gamma_1 x_T$. One can show that the consumption rules for this problem are continuously differentiable as

long as there are no liquidity constraints. However, in the presence of liquidity constraints, the consumption rules may exhibit a kink. See Deaton (1991) and Ayagari (1993). We effectively impose a liquidity constraint by not allowing the household to borrow against illiquid assets. This indicates that smooth approximation methods, as advocated by Judd (1992) are inappropriate in that case. Instead, we will use a standard discretization method: we specify an exogenous grid for cash-on-hand: $\{x^j\}_{j=1}^J \subset [0, x^{\max}]$. In order to capture the curvature of the consumption rule at low values of cash on hand, the grid will be finer for $x \in [0, 2]$. In practice, for each value of cash on hand on the grid, x^j , we find the associated consumption c^j that satisfies (B.1). In choosing the size and coarseness of the grid, we face the usual trade-off between precision and computing time. Adding points on the grid gives a finer approximation of the consumption rules. Since the consumption rule at age $t + 1$ is the input necessary to get the consumption rule at age t , imprecisions could compound over time. On the other hand, the Euler equation is the innermost loop of the entire algorithm. With 100 points on the grid and 40 time period, we must solve 4000 solutions to (B.1). This takes approximately 45 minutes on a P6 chip running at 133Mhz, courtesy of Intel Corporation. We also face a difficult decision regarding the range of cash on hand, x^{\max} . For small values, cash on hand in sample is likely to move out of the grid. Consumption will then be evaluated using extrapolation methods, always much less precise than interpolation. On the other hand, increasing the range with a fixed number of points implies less precise estimation of the curvature. One solution consists in endogenizing the grid so that, for instance, cash on hand remains within the grid with probability 0.95. We adopted the simpler approach consisting in checking that cash on hand, in the simulations, remains *strictly* inside the grid. In practice, we took $x^{\max} = 40$ and $J = 100$, with 50 points between 0 and 2. We checked the quality of the approximation by solving the stationary infinite horizon problem and checking the rate of convergence to the fixed point of the functional Bellman equation.

Appendix C

Data

We use the CEX family, member and detailed expenditure files for years 1980 to 1993, as kindly provided by the NBER. Most of our information about the CEX is obtained from Bureau of Labor Statistics (1980-1993) and conversations with BLS statisticians. Households are discarded if they are missing any of the information necessary for the regressions, if they report changes in age from the second to fifth interview of more than a year or less than zero years, if they are classified as incomplete income reporters, or if their reporting implies less than \$1000 in annual income or consumption.

We use information about the reference person to assign the household to cells, unless the reference person is female. In this case we use the spouses information. If there is no spouse, or his information is missing, the household is discarded. When this cut was made it eliminated 20% of the sample. All information besides individual labor income and consumption is taken from the family files. Values are assigned to a household based on information gathered in the fifth interview, otherwise information is used from the second interview, or, if it is not available, the household is discarded. Households should not be matched across 1985 to 1986, and are not. Care is taken to assure consistency in our data despite variable classification changes through time, and across reference person and spouse. Information was kindly provided by the Division of the CEX in the Bureau of Labor Statistics about various issues including the matching of occupation codes from 1980-81 to later years.

Pension contributions, income, Social Security contributions, and all asset income all refer

to the past twelve months. Our definition of pension contributions is the sum over the CEX subcategories and thus includes private pensions, public pensions, Railroad Retirement pensions, and self-employed, IRA, and Keogh plans. If the after-tax family income variables is topcoded, reference person and spouse labor incomes are subtracted and we add, for each, the variable created by multiplying the earnings in last paycheck by the appropriate pay period. These labor income variables are the sole variables from the member files used. Assets and asset income refers to the sum over savings accounts, checking accounts, bonds, and stocks, as of the time of interview. Each household is assigned to a year based on the midpoint between the first and fifth interview if both data are available; otherwise simply the single interview date is used. Age is the average of both interviews if both are available, otherwise it is the single one available. Due to some extreme reports, we reset reported tax rates above 50% back to 50%, and below zero percent to zero. We perform a similar exercise for Social Security contribution rates and pension contribution rates, using 25% as the upper bound.

Consumption data is compiled from the detailed expenditure files as all expenditures by a household except for those for health care, mortgage interest, and education. The consumption level is then the average monthly expenditure times twelve. Five percent of households have consumption data for 4, 7, 10, 13, or 14 months and these households' consumption are treated as if they were over 3, 6, 9, and 12 months. That is the recall interview period extended beyond the basic three months and some expenditures are recorded in a later month. BLS statisticians recommend treating these expenditures as if they occurred in the preceding month. Those covering 1 or 2 months (one percent of the sample) were dropped.

The unemployment rates merged to the CEX are the regional unemployment rates for civilian population from the Household survey conducted and published by the Bureau of Labor Statistics in "Employment and Earnings." The GNP IPD PCE is from Council Of Economic Advisors (1995).

Appendix D

Method of Simulated Moments

According to section 3.2, consumption at age t for individual i depends on cash on hand x_t^i , the realization of permanent component of income P_t^i , the entire path of expected permanent income growth $\mathfrak{G}_T = \{G_t\}_{t=1}^T$, and the parameters of the consumption problem $\bar{\theta}^i = (\beta, \rho, R, \gamma_0, \gamma_1, p, \sigma_u^2, \sigma_n^2)$, an 8×1 vector. In practice, it will not be possible to estimate directly all the elements of $\bar{\theta}$. Instead, we rewrite $\bar{\theta} = (\theta', \chi')'$ where $\theta = (\beta, \rho)'$ and $\chi = (R, \gamma_0, \gamma_1, p, \sigma_u^2, \sigma_n^2)'$. We assume that θ belongs to some compact set $\Theta \subset \mathfrak{R}^2$. The elements of χ will be inputs into the estimation procedure. Defining the vector of state variable $z_t^i = (x_t^i, P_t^i)$, we postulate the following data-generating process:

$$\ln C_t^i = \ln c_t(z_t^i, \theta; \chi, \mathfrak{G}_T) + \epsilon_t^i = \ln(c_t(z_t^i, \theta; \chi, \mathfrak{G}_T) P_t^i) + \epsilon_t^i, \quad (\text{D.1})$$

where ϵ_t^i is an idiosyncratic shock that represents measurement error in consumption levels. We are interested in estimating θ . If we were able to observe simultaneously the level of cash on hand of consumers and the level of their permanent component of income (as well as their consumption), θ could be estimated using Hansen's *GMM* (1982) on individual level data. More precisely, one would write the following moment conditions:

$$E[h(w^i, \theta_0; \chi, \mathfrak{G}_T)] = 0, \quad (\text{D.2})$$

where θ_0 is the true parameter vector, and

$$\begin{aligned} w_t^i &= (\ln C_t^i, z_t^i)' \\ h_t(w_t^i, \theta; \chi, \mathfrak{G}_T) &= (\ln C_t^i - \ln C_t(z_t^i, \theta; \chi, \mathfrak{G}_T)) \frac{\partial \ln C_t(z_t^i, \theta; \chi, \mathfrak{G}_T)}{\partial \theta'}. \end{aligned}$$

w is a $T \times 3$ vector and $h(w, \theta, \mathfrak{G}_T)$ is a $T \times 2$ vector.

The estimation procedure would then minimize:

$$g(\theta; \chi, \mathfrak{G}_T)' W g(\theta; \chi, \mathfrak{G}_T), \quad (\text{D.3})$$

where $g(\theta; \chi, \mathfrak{G}_T) = \text{vec} \left(\frac{1}{T} \sum_{i=1}^I h(w^i, \theta; \chi, \mathfrak{G}_T) \right)$ is a $2T \times 1$ vector and W is a weighting matrix. In practice, the number of cross-section observations for each age varies in the sample, so that $I = I_t$. We do not write explicitly this extra time-dependence in order to keep the notations simpler. The first difficulty with (D.2) is that quality panel data on consumption, asset and income information for individual households are not available in any US dataset. Therefore direct estimation using (D.3) is not possible. We do however observe χ and average consumption at each age $\ln \bar{C}_t \equiv \frac{1}{I(t)} \sum_{i=1}^{I(t)} \ln C_t^i$. This suggests that we can circumvent the problem by looking directly at *the unconditional expectation of consumption at each age*:

$$\ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T) \equiv E[\ln C_t(z_t, \theta; \chi, \mathfrak{G}_T)] = \int \ln C_t(z, \theta; \chi, \mathfrak{G}_T) dF_t(z), \quad (\text{D.4})$$

where the unconditional distributions of normalized cash on hand and permanent income depend on age t . We write the T moment conditions as

$$E \left[h(\ln C^i, \theta_0; \chi, \mathfrak{G}_T) \right] = 0, \quad (\text{D.5})$$

where θ_0 is the true parameter vector, $\ln C^i = \{\ln C_t^i\}_{t=1}^T$ and $h(\ln C^i, \theta, \mathfrak{G}_T)$ is a $T \times 1$ vector with t^{th} element:

$$h_t(\ln C^i, \theta; \chi, \mathfrak{G}_T) = \ln C_t^i - \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T).$$

However, at this point we encounter a second difficulty. The unconditional distribution

for the state variables at age t , $dF_t(z)$, is extremely cumbersome to evaluate, as well as the unconditional expectation (D.4).

The *Method of Simulated Moments*, as developed by Pakes and Pollard (1989) and Duffie and Singleton (1993) allows us to circumvent this difficulty. We can define a measurable transition function $\mathfrak{T} : \mathfrak{R}^2 \times \mathfrak{R}^2 \times \mathfrak{R}^T \times \Theta \rightarrow \mathfrak{R}^2$ that describes the dynamics of the state variables $z_{t+1} = (x_{t+1}, P_{t+1}) = \mathfrak{T}(z_t, \nu_{t+1}, \mathfrak{G}_T, \theta)$, where $\nu_t = (U_t, N_t)$ according to:

$$\begin{aligned} x_{t+1} &= \frac{R}{G_{t+1}N_{t+1}} \left(x_t - c_t(x_t^i, \theta; \mathfrak{G}_T) \right) + U_{t+1} \\ P_{t+1} &= G_t P_t N_{t+1}. \end{aligned} \quad (\text{D.6})$$

We omit the dependence in the calibrated parameters χ . U_{t+1} and N_{t+1} are respectively the permanent and transitory shocks to income. The first line of (D.6) is the normalized budget equation while the second line follows from our assumptions about the income process. This transition function can then be used to rewrite the unconditional expectation (D.4):

$$\begin{aligned} \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T) &= \int \ln C_t(z, \theta; \chi, \mathfrak{G}_T) dF_t(z) \\ &= \iint \ln C_t(\mathfrak{T}(z, \nu, \mathfrak{G}_T, \theta), \theta; \chi, \mathfrak{G}_T) dF_{t-1}(z) dF(\nu). \end{aligned} \quad (\text{D.7})$$

Note that the transition function depends on θ , through the consumption rule. (D.7) provides a convenient way to calculate the unconditional expectation is to use a Monte-Carlo integration. Assume that we have an $\mathfrak{R}^2 \times \mathfrak{R}^T$ -valued sequence of random variables $\{\hat{\nu}^i\}_{i=1}^{i=L}$ where $\hat{\nu}^i = (\hat{\nu}_1, \dots, \hat{\nu}_T)'$, identically distributed and independent of $\{\nu^i\}_{i=1}^{i=L}$. From any initial distribution $F(z_0)$ and candidate θ , we can generate the path of state variables according to (D.6):

$$\hat{z}_{t+1}^i = \mathfrak{T}(\hat{z}_t^i, \hat{\nu}_{t+1}^i, \mathfrak{G}_T, \theta); \quad \forall 1 \leq t \leq T \text{ and } 1 \leq i \leq L.$$

For large enough L , the unconditional expectation is then simulated by:

$$\ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T) \equiv \frac{1}{L} \sum_{i=1}^L \ln C_t(\hat{z}_t^i, \theta; \chi, \mathfrak{G}_T) \rightsquigarrow \ln \hat{C}_t(\theta; \chi, \mathfrak{G}_T).$$

For any parameter vector $\theta \in \Theta$ we can construct the T moments:

$$\begin{aligned} \bar{g}_t(\theta; \chi, \mathfrak{G}_T) &= \frac{1}{I(t)} \sum_{i=1}^{I(t)} \bar{h}_t(\ln C_t^i, \theta; \chi, \mathfrak{G}_T) \\ &= \frac{1}{I(t)} \sum_{i=1}^{I(t)} \ln C_t^i - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T) \\ &= \ln \bar{C}_t - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T). \end{aligned}$$

The estimation procedure minimizes:

$$\bar{g}(\theta; \chi, \mathfrak{G}_T)' W \bar{g}(\theta; \chi, \mathfrak{G}_T), \quad (\text{D.8})$$

where $\bar{g}(\theta; \chi, \mathfrak{G}_T) = (\bar{g}_1, \dots, \bar{g}_T)'$ is a $T \times 1$ vector and W is a weighting matrix. Note that in the case where $W = I_T$, the identity matrix, the estimation procedure is equivalent to minimizing the sum of square residuals:

$$S(\theta; \chi, \mathfrak{G}_T) = \sum_{t=1}^T \left(\ln \bar{C}_t - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T) \right)^2.$$

However, even though we are minimizing the sum of squared residuals, asymptotic results still apply as long as $I(t) \rightarrow \infty$ where $I(t)$ is the number of observations at age t . Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the *MSM* estimator $\hat{\theta}$ is both consistent and asymptotically normally distributed. Denoting $I = \frac{1}{T} \sum_{t=1}^T I(t)$,

$$\sqrt{I}(\hat{\theta} - \theta_0) \rightsquigarrow \mathfrak{N}(0, V),$$

with

$$\begin{aligned}
 V &= (D'WD)^{-1} D'W\Omega WD (D'WD)^{-1} \\
 D &= E [\partial\tilde{g}/\partial\theta'] \\
 \Omega &= \text{avar}(\tilde{g}) \\
 W &= \text{plim } W_I.
 \end{aligned}$$

In practice, the asymptotic variance-covariance matrix is estimated by:

$$\hat{V} = \frac{1}{I} (\hat{D}'W\hat{D})^{-1} \hat{D}'W\hat{\Omega}W\hat{D} (\hat{D}'W\hat{D})^{-1} \quad (\text{D.9})$$

$$\hat{D} = \left. \frac{\partial\tilde{g}}{\partial\theta'} \right|_{\theta=\hat{\theta}}. \quad (\text{D.10})$$

The t^{th} diagonal element of $\hat{\Omega}$ is given by

$$\hat{\Omega}_t = \frac{1}{I(t)} \sum_{i=1}^{I(t)} (\ln C_i^i - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T))^2.$$

Under the assumption of no serial correlation, the off-diagonal elements of $\hat{\Omega}$ are 0. However, a robust estimator can be constructed if $\hat{\Omega}_{t,t'}$ is defined as:

$$\hat{\Omega}_{t,t'} = (\ln \bar{C}_t - \ln \bar{C}_t(\theta; \chi, \mathfrak{G}_T)) (\ln \bar{C}_{t'} - \ln \bar{C}_{t'}(\theta; \chi, \mathfrak{G}_T)).$$

This methodology also provides a useful overidentifying restriction test. If the model is correct, the statistic

$$\chi_{T-2} = I \tilde{g}(\theta; \chi, \mathfrak{G}_T)' \hat{\Omega}^{-1} \tilde{g}(\theta; \chi, \mathfrak{G}_T)$$

is distributed asymptotically as Chi-squared with $T - 2$ degrees of freedom.

The optimal weighting matrix is $W = \Omega^{-1}$. The optimal weighting is implemented by first running the regression with an arbitrary weighting matrix, computing the associated $\hat{\Omega}^{-1}$, and then using this estimate in a second round of estimation.

In practice, we simulate $\ln \tilde{C}_t(\theta; \chi, \mathfrak{G}_T)$ by running $L = 20,000$ independent income processes for 40 years, and computing in each year the associated consumption and cash on hand. The first step used $W = I_T$. We also assume that the initial distribution of cash on hand is 0.3 times current income. This assumption captures the fact that most households do not start with no assets. Finally, for all households, we set the initial value of the permanent component of income to the estimated income level at age 25 from our profiles. We performed first a 25×25 grid search over the parameter space Θ . Then, we performed a second 25×25 grid search around the optimum. This guarantees that the procedure converges to the global minimum. Then, we used a standard minimization algorithm. Each grid search takes approximately 12 days of CPU time on a P6 or on a RSC6000. Once the optimum has been found, the gradient of the moment vector is evaluated numerically and the variance-covariance matrix estimated. For the disaggregated results and the robustness checks, a Brent algorithm was used.

Chapter 4

The Timing of Purchases, Market Power, and Economic Fluctuations

4.1 Introduction

All transactions are discrete; they occur at certain instants in time. Yet in most economic models, agents transact in every period or continuously, and do not choose the timing of their transactions.¹ This simplification leads to powerful insights, but it also obscures characteristics of markets in which the timing of transactions is important. Business cycles can be viewed as resulting from the temporal clustering of largely discrete decisions such as layoffs, purchases of durables, new product introductions, and plant retooling. Many recent models that explicitly allow agents to time transactions give new insights into business cycle phenomena.²

In this chapter, I analyze a market with three crucial characteristics. Sellers have market power, so that changes in the effective elasticity of demand can lead to changes in markups. Second, buyers choose when to make discrete, lumpy purchases. The com-

¹While many macroeconomic models, such as q -theory or consumption theory, focus on the optimal intertemporal allocation of purchases, every agent in such models transacts in every period or at every instant.

²See, for example, Caballero and Hammour (1994), Caballero, Engel and Haltiwanger (1995), Mortensen (1994), and Diamond (1994).

ination of these two features separates the current model from both previous models of industry markup dynamics and models with fixed adjustment costs. Finally, the number of consumers who would like to purchase in each period fluctuates exogenously though time.³

This structure provides a realistic baseline for examining the macroeconomy and many specific industries. Most vendors have some discretion in setting price and demand does change in predictable ways. Every expenditure occurs at an instant in time, and all expenditures yield some lasting utility or service flow – consumers do not continuously purchase airline tickets, stereo components, clothes, haircuts, or even caffeinated beverages.

This market structure also has interesting implications for market dynamics and business cycles. First, fluctuations in the distribution of buyers waiting to buy represent changes in the elasticity of demand and thus lead to price variation. These price movements increase sales and decrease markups at the beginnings of booms, while decreasing sales and increasing markups at the starts of slumps. However, a second characteristic of dynamic equilibria fights this amplification: the ability of buyers to delay or accelerate their purchases smooths prices. Sellers cannot raise prices relative to surrounding periods without losing sales. I call this feature of the market “real price stickiness.”⁴ With increasing marginal costs, real price stickiness leads naturally to a countercyclical markup.

More generally, the focus on the timing of purchases has important implications for macroeconomic fluctuations. This model offers an explanation for why quantities rather than prices seem to adjust over the business cycle. Time periods are not separate markets, each clearing independently. High prices in one period cause buyers to move transactions to nearby periods, so that the elasticity of demand is highly dependent on local price variation. On the other side of the market, sellers may choose not to decrease their prices

³In reality, these fluctuations are likely to be driven partly by the history of the market.

⁴As will be discussed later this has very different implications from either real price rigidity, which occurs when firms' objective functions are flat-topped, or price stickiness, which occurs when nominal prices cannot adjust.

much when demand is low because to do so would steal from their own future sales. Both buyer and seller behavior contribute to the smoothing of real prices.

The price fluctuations that do occur reduce the volatility of sales relative to the fluctuations which a constant price would induce.⁵ The timing of transactions thus creates a natural propagation mechanism. In a transitory boom, prices increase slightly and many buyers postpone purchases, increasing demand in later periods. Thus a shock in a single period is spread out over several periods.

A simple example of the dynamics of the model is illuminating.

First, consider a market in which there are a large number of identical consumers who choose to buy a new car when their old car has depreciated (deterministically) to a value s . Consumers are initially uniformly distributed over (s, ∞) . Firms are monopolistically competitive and produce using a constant marginal cost technology. When any firm considers lowering its price, it weighs the trade-offs among 1) lower profits from its own, current customers, 2) increased sales, stolen from the current customers of its competitors, and 3) increased sales stolen from itself and its competitors in the future.⁶ Let p^* denote the equilibrium price which balances these factors. Assume that p^* exceeds marginal cost. Further, assume that the market is in a stable equilibrium, as pictured in Figure 4-1 A: in every period each consumer's car depreciates by some small amount, δ , and those people who are at or below s buy a new car and exit the market, leaving the distribution of consumers unchanged.

Now consider perturbing the above market slightly by adding a small amount of consumers at one point. When these consumers are about to purchase because their cars have depreciated sufficiently, the importance of the third effect, "future demand stealing," increases for all firms. The elasticity of demand is higher because a change

⁵As noted by Caballero (1993) and Bils and Klenow (1995) durable goods indeed seem to fluctuate less than would be predicted based on changes in nondurable consumption; however both of these papers interpret this as solely due to adjustment costs.

⁶In this example I assume away concerns over lower future sales and decreased sales in the previous period. In Section 4 of this chapter, I examine a collusive equilibrium in which firms care about lower sales in the future. As in Bils (1989) and analyzed in Bulow (1982), dynamic consistency considerations imply that firms do not take the last effect into account.

in price generates an unusually large change in sales, and so firms cut prices. The endogenous reaction of firms amplifies the blip in demand by inducing some of this group of consumers to upgrade at the same time as those who would have purchased anyway. The time-pattern of sales is changed by this endogenous firm response, as shown in Figures 4-1 *B* and 4-1 *C*. Finally, note that the market displays a countercyclical markup in response to a change in demand— the ratio of price to marginal cost declines as sales increase.⁷

The previous discussion highlights the importance of firm's reactions. Price movements, however, are tightly constrained by consumer optimization. Consumers have the ability to substitute into low (out of high) price periods, so that today's price is constrained by expected future prices. Moreover, some buyers may have expected the price decline and thus delayed their purchases, adding further to sales at the beginning of the boom, and creating a decline in sales immediately before the increase. The endogenous decline in sales causes firms to cut prices immediately before the boom, thus creating a ripple effect on sales (and therefore optimal price) in the previous period.

Now consider the same market, but populated by firms with increasing marginal cost production technologies. During a boom, prices increase, but since consumers can wait to purchase until after the boom, the price increase is limited by expected future prices.⁸ As long as consumers are sufficiently patient, the markup will decline during the boom. Consumer optimization constrains real price movements, so that prices are smoother than marginal costs and the markup is countercyclical.

If the market structure is collusive or monopolistic, so that firms consider the effect of their price setting today on their sales in the future, then firms also have a price smoothing motive. In this case, firms do not cut margins much during slumps, since sales are not lost but merely postponed.⁹ As in the case in which firms were monopolistically competitive and had constant marginal cost, changes in the distribution increase sales and reduce

⁷I would argue that it is exactly this effect which leads retailers to offer pre-season sales on seasonal items, such as clothing.

⁸Consumers also may have been able to buy before the boom if it was expected.

⁹This point is also made in Domowitz, Hubbard and Petersen (1987).

markups at the start of booms.

After presenting a more formal dynamic model, I demonstrate that industry pricing behavior is consistent with the model's implications. For a subset of four-digit manufacturing industries – chosen because they correspond well to specific consumer goods – I obtain measures of the frequency at which households purchase the goods. I match these measures to the NBER productivity database. Consistent with the implications of the model, I find that demand-driven increases in sales cause smaller real price increases for those goods which are purchased less frequently – that is, for which the timing of purchases is more important. In addition to smaller price elasticities of demand, the model predicts that markups vary more countercyclically in industries where the timing of purchases is more important. Analysis of industry-level markups however is not informative due to large standard errors and unstable coefficients.

The balance of this chapter is laid out as follows. In the next section, I present a model in which buyers are price-takers and optimally time their purchases of a lumpy good. I consider optimal seller and buyer strategies when the market structure is monopolistically competitive and when it is monopolistic (actually under collusion). Sections 3 and 4 characterize the dynamics of prices, costs, and sales when there are fluctuations in the distribution of consumers. In section 5, I relate this theory of countercyclical markups to competing theories. Section 6 turns to industry-level data and tests the basic implications of the theory against other theories. In the final section of the chapter, I conclude. Appendix *E* contains proofs and Appendix *F* contains the details of data construction and sources.

4.2 The Market Structure

In this section, I present the buyer and seller optimization problems, and discuss the necessary conditions for equilibrium. In the model, a large number of consumers search over the prices offered by a large number of sellers. Search by consumers generates market power, in a similar manner to the steady state analysis in Diamond (1971).

4.2.1 Consumers

The demand side of the market consists of a large number of potential buyers or consumers.¹⁰ Each consumer has some amount of a good which depreciates deterministically through time. The good is lumpy, and consumers decide *when* to buy, not *how much* to buy. The good provides lasting utility or profit flow because it is a durable good or capital investment, because it must be bought “bundled” in a fixed quantity and is storable (like table salt), or because its consumption yields a stock which depreciates through time and gives a flow of utility (like expenditures on entertainment).

Consumers are price takers and decide when to purchase a single new good. After purchasing the good, consumers leave the market. There are large transactions costs: a buyer who upgrades to a new good gets nothing for its old good.¹¹ Consumers are differentiated by their individual stocks in any period, denoted k_{it} , that depreciate deterministically according to:¹²

$$k_{it} = (1 - \bar{\delta})k_{it-1}. \quad (4.1)$$

The money-metric utility (or profit) flow from the stock of the good is logarithmic,

$$u_{it} = \frac{1}{\lambda} \ln(k_{it}), \quad (4.2)$$

where λ is the marginal utility of wealth.¹³ An individual's utility flow evolves according to:

$$u_{it+1} = u_{it} - \delta \quad (4.3)$$

where $\delta = -\frac{1}{\lambda} \ln(1 - \bar{\delta}) \simeq \frac{\bar{\delta}}{\lambda}$.

In every period, each consumer, i , observes the price of one randomly drawn seller,

¹⁰Although I use the term consumers, the potential buyers may be firms considering purchasing new investment goods as well as households seeking to purchase consumer goods.

¹¹This assumption can be relaxed without changing the substance of the results.

¹²I denote an individual consumer's or seller's stock, price, or sales using lower case letters, and market prices and quantities using capitals.

¹³The model can be generalized by using a constant relative risk aversion utility function at significant cost in tractability. Similarly, deterministic decline in the marginal utility of wealth changes none of the main implications.

j , but does not automatically observe the price of any other seller. The consumer then chooses among the following three options. First, it can purchase the good at the posted price. In this case the consumer pays p_{jt} and receives v , which represents the dollar-denominated expected present discounted value of holding a new good.¹⁴ Second, the consumer can choose to do nothing, in which case it receives the benefits of its current stock, and, in the next period, faces the same decision with the price of a new randomly drawn seller.¹⁵ Finally, it can pay a search cost c , and observe the price of a new randomly drawn seller. It then faces the same decision again in the current period: purchase, wait, or continue searching.¹⁶

Each period thus consists of an infinite number of instants in which consumers can choose to search over prices/goods. If a buyer is intent on purchasing in a given period, it can visit as many sellers (and pay as much in search costs) as it wishes within a single period.

Prior to searching, consumers know the distribution of prices posted by sellers, and the payment of search costs yields only information about a specific seller's price. Thus, if a consumer ever prefers searching to delaying, it searches in the current period until it finds an acceptable price and purchases.¹⁷ Consequently, once any potential buyer chooses to search, the problem is a stationary search problem and consumers search across firms until they find a price below their reservation price. The value to a buyer of

¹⁴One might also motivate a constant v by arguing that buyers actually return to the market (at the top of their adjustment bands in an (s, S) type model) but that all buyers have an expected holding time and future purchase prices which are independent of the current state of the system. This assumption would be close to true with a significant amount of individual uncertainty and little predictable variation in price at long horizons.

In an early manuscript version of this chapter, I explicitly included a stochastic, match-specific value of a new good, v_{ij} , and demonstrated that the general implications are the same as those I present here.

¹⁵Search costs correspond to the costs a consumer incurs in test-driving an automobile or in going to a store and pricing its goods; similarly a firm incurs a cost when it evaluates how well a given seller's equipment meshes with its other capital goods and its employee's skills and needs.

¹⁶Since I model all heterogeneity in firms as occurring in their effective price, I am implicitly assuming that whatever differences there are in match quality are compensated through price adjustment. So, for example, if a firm chooses to produce tractors of low quality, consumers compensate by purchasing warranties, or paying for additional features or equipment which make these tractors equivalent to those of other firms.

¹⁷That is, if the value of searching exceeds the value of delaying, it does so after one round of search or any number, since search costs paid while searching are sunk.

having found price p_{jt} can be written recursively as:

$$V^s(p_{jt}) = \text{Max}\{v - p_{jt}, E_t[V^s(p_{kt})] - c\}. \quad (4.4)$$

The value to searching is the maximum of the value of buying at the current price, or paying c and going to a new seller and discovering its price.

The value of a consumer with a utility flow of u who sees a price offer of p_{jt} can now be written recursively as:

$$V_t(u, p_{jt}) = \text{Max}\{v - p_{jt}, u + \beta E_t[V_{t+1}(u - \delta, p_{kt+1})], E_t[V^s(p_{kt})] - c\} \quad (4.5)$$

where β is the consumer's discount factor. Equation (4.5) says that the value of having utility flow u in period t , and seeing the price of one good, p_{jt} , is equal to the maximum of 1) the value of purchasing at the observed price, 2) the expected value of delaying the decision by one period, and 3) the expected value of deciding to purchase in t , but doing some searching over prices first. Since consumers end search by purchasing, I henceforth refer to entering the search mode as purchasing.

Prices are bounded below so that the value function is bounded above for a positive depreciation rate and discount rate less than one.¹⁸ Further, the value function is decreasing in its first argument, which decreases through time; provided the first term is always positive, a finite purchase time is optimal. Note that the value function is a function of time— the buyer's optimal purchase date depends on the price path that it faces. The time-path of prices in turn depends on the true state of the system: the entire distribution of consumers over utility flows. In solving the model below, I look at cyclical equilibria, so that there are a finite number of value functions to solve.

In every period, each firm is randomly matched with a slice of consumers distributed according to $f_t(u)$. Let u_t^+ denote the utility flow of the buyer with the least amount of

¹⁸Marginal costs are assumed to be weakly increasing so that prices are bounded below by marginal cost when producing the first increment of output. The facts that utility flow depreciates linearly and firms discount future profits exponentially are sufficient but far from necessary to generate bounded returns in this setup. See Stokey et al. (1989) for the technical conditions.

the good in the market at the beginning of period t . The function $f_0(u)$ is assumed to be continuous, atomless, bounded above, and strictly positive over $[u_t^+, \bar{u}]$. The distribution evolves according to:

$$f_t(u) = \begin{cases} f_{t-1}(u + \delta t)\varepsilon_t(u) + g_t(u) & \bar{u} \geq u \geq u_t^+ \\ 0 & \text{otherwise} \end{cases} \quad (4.6)$$

where $\varepsilon_t(u)$ represents multiplicative, strictly positive shocks to the distribution and $g_t(u)$ is deterministic and strictly positive only over $[\bar{u} - \delta, \bar{u}]$.¹⁹ $\varepsilon_t(u)$ represents information revealed about the true and imperfectly known distribution at time t .

The evolution of u_t^+ follows from consumer optimization.

4.2.2 Sellers

There are an infinite number of identical sellers along the unit interval. Each seller posts a price in each period and sells to those consumers who arrive and have reservation prices above the seller's posted price. The search process gives each seller some market power. That there are an infinite number of sellers means that each one ignores the impact of its own price setting on its future sales. That is, there is a problem of the commons. Like competition for a natural resource, seller competition leads to rapid depletion of potential consumers. If firms colluded, they would choose to have much higher prices and sell to all consumers at later dates. Since the dynamic implications differ, I will present a characterization of the collusive game in Section 4 of this chapter.

In a symmetric Nash equilibrium in which all sellers post the same price and in which there is no collusion, no consumers actually search.

Sellers choose a sequence of prices, $\{p_{j_s}\}$, to maximize:

$$\sum_{s=t}^{\infty} R^{t-s} E_t \left[p_{j_s} q_s(p_{j_s}, P_s) - c_1 q_s(p_{j_s}, P_s) - \frac{1}{2} c_2 q_s(p_{j_s}, P_s)^2 \right] \quad (4.7)$$

taking the market price, P_s , as given. The function $q_s(p, P)$ is the amount of sales

¹⁹The function $g_t(u)$ provides a boundary condition for the distribution.

that a seller charging p makes in period s when the market price is P . This function is determined by the distribution of potential consumers and their optimal strategies. c_1 and c_2 represent quadratic costs of production, both of which are weakly positive.²⁰ Equation (4.7) reduces to a sequence of static problems in which sellers choose individual prices, taking present, past, and future market prices as given. The first order condition can be written as an inverse elasticity pricing rule:

$$\frac{p_{js} - c_1 - c_2 q_s(p_{js}, P_s)}{p_{js}} = \frac{1}{\epsilon_s^d(p_{js}, P_s)}, \quad (4.8)$$

where $\epsilon_s^d(p_{js}, P_s) \equiv -\frac{\partial q_s(p_{js}, P_s)}{\partial p_{js}} \frac{p_{js}}{q_s(p_{js}, P_s)}$, the elasticity of firm demand given the market price.

4.2.3 Market Equilibrium

In order to find the optimal price and quantity paths, I need to derive the function $q_t(p, P)$ which follows from buyer optimization given rational price expectations. Define u_t^* as the utility flow of the buyer with the smallest stock of the good after sales have been made in period t . Then by equation (4.3)

$$u_{t+1}^+ = u_t^* - \delta. \quad (4.9)$$

I begin with four lemmas which help to characterize the equilibrium. The proofs are contained in Appendix E.

Lemma 1 *No search. In a symmetric equilibrium, no consumers search and*

$$V_t(u, p_{jt}) = \text{Max}\{u + \beta E_t[V_{t+1}(u - \delta, P_{t+1})], v - p_{jt}\} \quad (4.10)$$

²⁰The problem could also contain fixed costs to entry and fixed costs in each period, both of which would determine market size and ensure that sellers not make net profits in excess of the usual rate of return. These costs are not relevant to the analysis at hand and are ignored.

Since in every period all firms charge the same price, any search has a total expected gain of minus the search cost. Define T_{tj} as the expected purchase date of consumer j conditional on information available at time t .

Lemma 2 *Skimming property.* $u_{ti} > u_{tj}$ implies $T_{ti} \leq T_{tj}$.

This Lemma follows from a revealed preference argument, which demonstrates that the value function is weakly increasing in its first argument. Lemma 2 implies that sales in any period can be calculated by finding the buyer who is indifferent between purchasing and waiting, and then summing over consumers with lower utility flows.

Lemma 3 u_t^* evolution. *Provided that sellers sell to some consumers in every period, u_t^* is defined by*

$$u_t^* = (1 - \beta)v + \beta E_t [P_{t+1}] - P_t \quad (4.11)$$

Equation (4.11) is derived from equation (4.10) by noting that the marginal consumers today must buy tomorrow, and, therefore, that their expected value to delaying is the expected utility of purchasing in the next period. Finally, I note that, under certainty, weak conditions are needed to imply positive sales in every period, and thus make equation (4.11) valid.

Lemma 4 *Positive sales.* If $\varepsilon_t(u) \equiv 0$, $q_0 > 0$, and $p_t \geq c_1 \forall t$, then $q_t > 0 \forall t$.

Lemmas 3 and 4 together imply that equation (4.11) gives u_t^* .

I now turn to deriving sales for any firm given the market price. First, if all sellers are charging P_t , then, for every seller the quantity of sales is the integral of the buyer distribution from u_t^+ to u_t^* :

$$Q_t(P_t) \equiv \int_{u_t^+}^{(1-\beta)v + \beta E_t [P_{t+1}] - P_t} f_t(u) du. \quad (4.12)$$

Consider a seller who deviates from the market price. If the seller cuts price, then it sells to a larger share of the consumers who see its price in the current period. However, it

gains no sales from other sellers since no consumer finds it worth searching across such a large number of sellers when it knows only one seller has cut its price.²¹ Thus for all price cuts:

$$q_{jt}(p_{jt}, P_t) \equiv \int_{u_t^+}^{(1-\beta)v + \beta E_t[P_{t+1}] - p_{jt}} f_t(u) du. \quad (4.13)$$

For price increases, the seller may lose some of its customers. However, as long as the price increase is smaller than c , no one leaves the seller to search, although some consumers may decide to delay purchase rather than buy or search. Thus, equation (4.13) determines demand for $p_{jt} \in [c_1, P_t + c)$; for higher prices, demand is zero. Note that $q_{jt}(P_t, P_t) = Q_t(P_t)$.

Given a set $(f_0(u), u_0^+, \{g_s(u), \varepsilon_s(u)\}_{t=0}^\infty)$, an equilibrium is a series $\{P_s, u_s^+\}_{t=0}^\infty$ which satisfies the following.

1. Consumers choose optimally whether to search and when to buy, so combining equations (4.9) and (4.11) yields:

$$u_t^+ = (1 - \beta)v + \beta E_t [P_{t+1}] - P_t - \delta. \quad (4.14)$$

2. The forcing term evolves according to equation (4.6).
3. P_t represents a global profit maximum for each firm.
4. Each seller optimizes, so that its first-order condition, which is the following nonlinear difference equation in prices, is satisfied:

$$P_t - c_1 - c_2 \int_{u_t^+}^{(1-\beta)v + \beta E_t[P_{t+1}] - P_t} f_t(u) du = \frac{\int_{u_t^+}^{(1-\beta)v + \beta E_t[P_{t+1}] - P_t} f_t(u) du}{f_t((1 - \beta)v + \beta E_t [P_{t+1}] - P_t)}. \quad (4.15)$$

²¹This follows from the assumption of an infinite number of sellers and random search. Modelling an equilibrium in which search occurs due to stochastic match-specific variation in v leads to the same main implications.

In equation (4.15), the integral terms are the market quantity sold in period t , and the right-hand-side is the negative inverse of the elasticity of demand. Note that, in choosing prices, firms take taking u_t^+ , P_t , and $E_t [P_{t+1}]$ as given, and that the first-order condition is evaluated at the market price.

Two main features of the equilibrium conditions drive interesting market dynamics. First, in equation (4.15), given a quantity of sales in t , the higher the density of consumers who are indifferent between purchasing in t and $t + 1$ (the denominator of the right-hand side), the lower is the markup. That is, when the distribution of consumers is increasing, the markup is likely to be lower, and consumers will purchase sooner than if markup were held constant. Second, the quantity of sales is determined relative to the expected market price in the next period. Sellers are unable to sell if today's price is set much above future prices. Nor is it feasible to plan to price well below future prices since, given such a planned future price, consumers who might choose to buy in $t + 1$ instead buy in t . In this case, sellers would have to cut future prices to satisfy their first-order conditions and to sell to any consumers in $t + 1$.

4.3 Properties of Dynamic Equilibria

Together conditions 1 – 4 are necessary and sufficient for a sequence of prices and lowest utility flows (u^+) to constitute an equilibrium. In general, however, the pair of difference equations (4.14) and (4.15) can be quite difficult to solve, even numerically. Moreover, the system may exhibit multiple solutions to these conditions.²² While there may be interesting economic implications of multiple equilibria in this dynamic system, I leave this to later research. For now, as I describe in the subsequent subsection, I simply choose a natural way to select an equilibrium path. I consider distributions in which

²²This possibility arises from the complexity of the nonlinear difference equations being solved (in addition to being nonlinear, equation (4.15) contains values of the state variable inside the forcing function) and from the lack of restriction on the demand curve in each period. The latter difficulty is endemic to market analysis, and most research eliminates the possibility by assuming a nice functional form for the demand curve. Since the demand curve in the current work depends crucially on past actions as well as the input specification, I choose the more flexible approach of examining each case.

certain important features of the solution can be explored and tractability and uniqueness maintained.

4.3.1 The Steady State

I define the steady state of the system as those constant quantities and prices that solve conditions 1 through 4 when the distribution of consumers over stocks is flat and evolves deterministically. Define α as the height of the distribution of consumers over utility flows, i.e. $f_t(u) = \alpha \forall t$.

Lemma 5. Steady state equilibrium. The steady state is uniquely determined by

$$\begin{aligned}
 P_{ss} &= c_1 + c_2\alpha\delta + \delta & (4.16) \\
 u_{ss}^+ &= (1 - \beta)(v - P_{ss}) - \delta \\
 Q_{ss} &= \alpha\delta
 \end{aligned}$$

Uniqueness of the steady state should not be surprising since the sources of possible multiple equilibria – nonlinearity in the difference equations and the unrestricted shape of the demand curve – are not present in the steady state because the distribution of consumers over utility flows is linear. The steady-state quantity sold depends solely on the number of consumers at each point in the distribution and the speed at which consumers' utility flows depreciate. Finally, note that the equilibrium markup rises with the depreciation rate.

4.3.2 Equilibrium and the Coase Conjecture

Before moving on to dynamics, I address a possible objection to the use of this model for analysis of markup dynamics. First, note that, as the length of a time period goes to zero, the markup goes to zero (in equation (4.16), this amounts to letting δ go to zero).²³

²³Despite the fact that the size of the markup goes to zero as the length of a period goes to zero, the percent fluctuation in the markup over the cycle remains constant.

This result follows from the same intuition as the well-known Coase (1972) conjecture.²⁴ Sellers cannot commit to keeping price high, so that as time periods become shorter, and follow one another more quickly, firms eventually sell to customers so rapidly that they run out of customers to sell to in any given interval of real time, and price falls to marginal cost. The faster price falls to marginal cost, the closer the initial price must be to marginal cost, so that in the limit price is always equal to marginal cost.²⁵ This implication is not a problem for the current model since it stems from assumptions made to preserve the tractability of the dynamic solutions. More specifically, either a model in continuous time with a matching function or one in which search takes a fixed amount of real time, would preserve the dynamics of interest and be immune to a Coasian critique.

4.3.3 Demand Fluctuations with Constant Marginal Cost

When the distribution of consumers over stocks fluctuates, the system of equations characterizing the equilibrium are not easily solved, and I approach the problem numerically. To characterize the solution, I find a solution to conditions 1 – 4 when the distribution of utility flows of consumers has infinitely repeating, deterministic cycles. That is, I solve the model under certainty, choose $f_0(u)$ to be periodic, and look for equilibria that have the same period, so that the price and quantity dynamics repeat exactly from cycle to cycle. I use a multiple shooting technique to find a series of prices and cutoff utility flows that satisfy the equilibrium conditions, taking my steady-state values as the starting points for this search process.²⁶ The distribution of consumers over stocks is taken to be a square wave. When interpreting the cycles, one must keep in mind the stylized nature of this example. In particular, since the simulations assume certainty, one should only

²⁴See Bulow (1982) and Stokey (1981) for proofs and discussion of the robustness of the Coase conjecture. Bond and Samuelson (1987) demonstrate that depreciation does not eliminate this result.

²⁵Another way of looking at this thought experiment is that as the length of a time period goes to zero, fewer and fewer buyers purchase in any given period, so that sellers no longer worry about losing profits on their inframarginal customers. It is the presence of inframarginal customers which generates markups.

²⁶See Judd (1993) for description of multiple shooting techniques. My program checks that all local optima for sellers are global optima, and also, for robustness, that quantity is weakly positive in every period.

consider the implications for expected changes in distributions.

Figure 4-2 displays the results of a typical simulation when firms have constant marginal costs.²⁷ The figures demonstrate that price is set in part according to the change in the height of the distribution. That is, when a large “lump” of consumers gets near to buying, price declines and quantity sold increases. This effect causes a reduced markup on the up-side of booms. Further, notice that this effect also drives down the price and quantity immediately before the boom. Consumers know that the price will fall in period 2, so they substitute towards that period; sellers respond by cutting price in period 1. Thus consumer optimization smooths price. Here, the sharp change in elasticity caused by the fluctuation in the distribution fights the price-smoothing effect of buyer intertemporal optimization. The upshot is that immediately before the boom, there is a “recession amplifier” as consumers delay purchasing to take advantage of the lower prices at the beginning of the boom.

A similar set of factors act at the end of the boom to amplify the end of the boom, as consumers substitute away from the high prices that occur at the beginning of the recession. After the recession begins, consumers have just been waiting to purchase and many have low stocks (Figure 4-2 C). Thus price declines only slowly due to the presence of fewer inframarginal consumers at the start of the low-demand time.

4.3.4 Demand Fluctuations with Increasing Marginal Cost

The second simulation employs the same cyclical distribution of consumers over utility flows, but now considers sellers with increasing marginal costs. Figure 4-3 displays the results of a typical simulation.

Three points can be taken from this experiment. First, the price-smoothing effect of consumer intertemporal optimization is much more powerful than the price responses of firms to changes in the height of the distribution. The price changes analyzed above are still present however, and are visible in the asymmetry of the figures. Thus, in this case,

²⁷Due to the stylized nature of the model, the model is not calibrated beyond setting the parameters so price exceeds marginal cost by 15 to 20 percent in all of the simulations.

sharp swings in the distribution are not necessary to get significant changes in markups and the timing of sales.

Second, the main effect of increasing marginal cost is to increase price in booms and therefore decrease quantity and smooth sales over the cycle. This effect generates propagation of increases in demand. All buyers postpone (or accelerate) purchase slightly so as to move towards lower-price periods.

Third, intertemporal substitution by consumers robs firms of their market power during booms. Buyers are less willing to purchase because past and future prices are lower. This substitution causes a reduction in the markup when sales are high.²⁸ Figure 4-3 C, shows u^+ , the utility flow for the buyer with the lowest utility flow. Immediately before the boom, u^+ rises as consumers buy earlier than they would have if prices had stayed constant. Thus when the boom begins, sellers find themselves facing consumers who have higher amounts of their old goods left. In competing to sell to these consumers, firms lower their markups. Sellers only slowly increase their prices once in the boom, as u^+ falls towards its equilibrium level. Inversely, at the end of the boom, consumers delay their purchases which decreases prices before the end of the boom and smooths the price and quantity fluctuation on the downside.

To summarize the implications of these two experiments, when marginal costs are flat, prices and markups should be low when quantity increases at the beginning of booms; when marginal costs are increasing, booms should be smoothed and markups should be countercyclical.

²⁸Because the costs of production are quadratic, the ratio of price to marginal cost declines in steady state with increases in the height of the distribution. Since this is not novel nor the interesting source of markup cyclicity in this chapter, I report price less marginal cost as the markup.

4.4 Properties of Dynamic Collusive Equilibria

4.4.1 Sellers

When sellers are able to collude, total profits are maximized by the sequence of prices which would be chosen by a monopolist who owned all the firms.²⁹ One can write this recursively as:

$$V_t(u_t^+) \equiv \max_{\{P_t\}} \left(P_t q_t(P_t, P_t) - c_1 q_t(P_t, P_t) - \frac{1}{2} c_2 q_t(P_t, P_t)^2 + R^{-1} E_t [V_{t+1}(u_{t+1}^+)] \right). \quad (4.17)$$

However, it may not be possible to maintain this optimal level of collusion.³⁰ If a firm deviates from the prescribed price sequence, it will cut price and try to steal consumers from the future demand of all firms. The usual punishment strategy involves all other firms setting prices to punish the defector in the subsequent period. Here however the optimal collusive arrangement can be maintained by subsequent pricing that makes the gross benefits of defecting zero. After seeing any price below the collusive price, all firms choose the largest price that makes any consumers who purchased from the defector wish that they had not. In equilibrium, these consumers therefore would not. Subsequently, the firms all return to collusion.³¹ As long as periods are not so far apart that these prices must be very different, collusion always can be maintained.³²

Using the first-order condition and the envelope condition yields the following in-

²⁹I continue to assume that the colluding firms or monopolist take ${}_t P_{t+1}$ as given when choosing P_t , as is standard. (See for example *Bils (1989)*). However, using a complicated self-punishment scheme, it may be possible for firms to internalize the effects of changes in today's price on consumer expectations of tomorrow's price.

³⁰Indeed, *Rotemberg and Saloner (1986)* argue that difficulty in maintaining collusion during booms causes prices to fall from monopolistic levels and leads to countercyclical markups.

³¹Note that this equilibrium is not renegotiation-proof.

³²I would argue that this is why collusion is possible in industries like the airline industry where the timing of purchases is important.

tertemporal Euler equation which is satisfied in the collusive equilibrium

$$q_t(P_t, P_t) + R^{-1} E_t \left[(P_{t+1} - c_1 - c_2 q_{t+1}(P_{t+1}, P_{t+1})) \frac{dq_{t+1}}{dP_t} \right] = (P_t - c_1 - c_2 q_t(P_t, P_t)) \frac{-dq_t}{dP_t}.$$

When the firms consider a price increase they balance the additional return on the inframarginal buyers, $q_t(P_t, P_t)$ and the gain in profits on increased sales in $t + 1$ against the decrease in profit from lost sales in the current period.

4.4.2 Market Equilibrium

Lemmas 1 – 3 still apply when firms are colluding, but firms may choose sales equal to zero.³³ Rather than keep track of these corner solutions, I will restrict attention to fluctuations in demand that generate positive sales in every period.

Equation (4.12) determines quantity, so that the firm's first-order conditions can be rewritten to replace condition 4 as:

$$\begin{aligned} -R_t^{-1} [P_{t+1} - c_1 - c_2 E_t(Q_{t+1}(P_{t+1}))] &+ (P_t - c_1 - c_2 Q_t(P_t)) & (4.18) \\ &= \frac{Q_t(P_t)}{f_t((1 - \beta)v + \beta E_t[P_{t+1}] - P_t)}. \end{aligned}$$

where I have used the fact that:

$$\frac{dQ_t(P_t)}{dP_t} = -f_t(u_t^*) = -f_{t+1}(u_t^* - \delta) = -f_{t+1}(u_{t+1}^*) = -\frac{dQ_{t+1}}{dP_t}. \quad (4.19)$$

The first equality comes from equation (4.12), the second from equation (4.6), the third from equation (4.9), the final one from equation (4.12), and I have assumed that there is no uncertainty about the height or evolution of the endpoint of the distribution.

³³While existence is still assured, the firm first-order condition does not hold in periods of zero sales.

4.4.3 The Steady State

When sellers are colluding the steady state price, quantity and cutoff utility flow are:³⁴

$$\begin{aligned}P_{ss} &= c_1 + c_2\alpha\delta + \frac{\delta}{1 - R^{-1}} \\u_{ss}^+ &= (1 - \beta)(v - P_{ss}) - \delta \\Q_{ss} &= \alpha\delta.\end{aligned}\tag{4.20}$$

Unlike in the noncollusive case, the steady-state price and markup would remain the same if time periods were to become infinitely fine.³⁵ The markup dynamics are also independent of the length of the time period; thus a Coasian criticism is not present here.

When firms collude and have constant marginal cost, they have a slight incentive to shift demand forward at the beginning of booms. Colluding firms seek to shift sales forward because sales today are worth more than future sales, and, if demand is increasing, because inframarginal sales are less important than marginal sales. However, when firms collude, they internalize the effect of sales today on sales tomorrow, so that the impact of fluctuations in demand on prices are an order of magnitude smaller than when the firms do not collude. That is, both firms and consumers seek to smooth prices, so that prices are effectively flat in this scenario. I now turn to the case of colluding firms with increasing marginal costs. In this case consumer optimization smooths price while firms seek to smooth markups.

4.4.4 Demand Fluctuations with Increasing Marginal Cost

Figure 4-4 displays the results of a typical simulation when colluding firms have increasing marginal costs.³⁶ Again, buyer intertemporal optimization acts to smooth price and quantity fluctuations. Sellers, on the other hand, try to keep markups roughly constant.

³⁴The steady state is unique. The proof follows exactly that of Lemma 5.

³⁵This follows from $\frac{\delta}{1-R^{-1}} = \frac{\delta}{r} + \delta$, which goes to $\frac{\delta}{r}$ as the length of each interval of time goes to zero.

³⁶Domowitz et al. (1987) also examine this case and argue that it implies a countercyclical markup.

They are willing to cut margins at the beginning of the boom, however, since getting buyers to buy early helps to keep costs lower. Firms also raise prices at the end of the boom because losing some buyers to the future helps them keep costs low. Thus, this market structure predicts that markups should be lowest at the start of booms, when demand is increasing, as in the first case analyzed. Price and quantity dynamics mimic the dynamics of the noncollusive case with increasing marginal costs: prices rise in booms and quantity fluctuations are smoothed and propagated.

4.5 Related Models of Markups and Durable Goods

I now discuss the relationship of the timing model to four related models of countercyclical markups. First, and most closely related are customer market models, as in Phelps and Winter (1970) and Bils (1989). In such models, firms have repeat customers, for whom interfirm competition is weak, and potential new customers, for whom interfirm competition is strong. When new customers are relatively more important than repeat customers— that is when demand is increasing— markups are low. This pricing behavior does not shift demand. Chevalier and Scharfstein (1994) analyze a variant of this model in which firms also face cash constraints and thus increase margins when cash flow is low. Markups are high when demand is low, and, again, the timing of demand fluctuations is unchanged. Thus, the basic predictions for the dynamics of sales and markups in the customer market are the same as for the timing model. However, the key difference is that the customer market story should apply to markets in which repeat purchases and switching costs are important, such as in the supermarket industry which Chevalier and Scharfstein (1994) analyze. The timing model predicts these dynamics in markets in which timing is key.

Second, Rotemberg and Saloner (1986) and Rotemberg and Woodford (1991) argue that implicit collusion among firms becomes more difficult when the size of profits in the present are greater than profits in the future. In order to maintain cooperation among firms, margins are therefore lowest when demand is decreasing. This model then has

the opposite implication from the timing model, and is more likely to apply to highly concentrated markets.

Third, Bilal (1991) models a market for durables in which high-income consumers are the marginal consumers deciding whether to purchase the good in recessions, while middle-income consumers are the marginal consumers in booms. Since there are far more middle-income consumers, the elasticity of demand is higher in booms, and therefore markups lower. While this model predicts that durable goods should have more countercyclical markups, it does not incorporate the effect of durability on the consumer side of the market. Further, there are many degrees of luxuriousness in durables and the relative importance of marginal purchasers may well be reversed for goods besides the most luxurious.

Finally, perhaps the most widespread theory of countercyclical markups comes from models of *nominal* price stickiness. In these models prices do not increase when demand increases because of costs of changing nominal prices. Such models generally require real rigidities – small profit losses to small deviations of price from its optimal level – and nominal rigidities – small costs of changing prices. Considering timing makes the *real* price sticky. That is, if firms were to charge prices much above expected future prices, buyers would delay purchases. Small deviations of prices from expected future prices can be very costly. Thus the real price is tightly constrained and will not be allowed to stray far from its optimal level. That is, real price *rigidity* is reduced by real price *stickiness*. With cyclical changes in the general price level, if the timing model applies to an industry, small costs of changing prices are less likely to have real effects and cause countercyclical markups.

Conlisk, Gerstner and Sobel (1984) present a model similar in spirit to the current model in that demand is shifted and markups are countercyclical. In their model, however, demand cannot be shifted forward and firms do not compete for sales. Rather two sorts of consumers enter the market every period. A monopolist sells to desperate consumers in every period, by never lowering price too fast. Then, every so often, price gets down to the reservation price of the low-demand consumers, who have been piling

up unsatiated in the market, and all consumers buy. In the next period, price jumps up to the reservation price of the desperate consumers and sales are only to the desperate consumers as price slowly declines again. This model thus has endogenous cycles, countercyclical markups, and the shifting of sales, but through a monopolist's optimal strategy rather than intertemporal competition for consumers.

4.6 Empirical Evidence

In order to quantify the importance of the price-smoothing effect of the timing of purchases I turn to industry data on prices, sales, and markups. I take three hypotheses from the previous sections. First, buyer intertemporal optimization smooths prices so that price reactions to demand fluctuations are small. Second, the first and third examples analyzed predict that when sales are increasing, prices and markups should be low. Finally, when marginal costs are increasing, markups should be strongly countercyclical. Thus I seek to test whether these hypotheses hold true for industries that sell goods for which the ability of buyers to time their purchases is important.

4.6.1 The Data

The main dataset employed is the National Bureau of Economic Research (NBER) productivity database, which contains annual data on industry inputs, sales, and prices at the level of the four-digit SIC code. The dataset covers 450 industries from 1958 to 1991. Its strength is careful attention to temporal consistency of industry and variable definitions. It includes measures of industry sales and inputs – including intermediate goods and raw materials – and price deflators for all inputs except the capital stock, where instead the database includes a deflator for new industry investment. Labor input is decomposed into production worker hours and nonproduction worker employment. Appendix *F* contains further details on the data.

The second set of data comes from Bils and Klenow (1995). They report durability measures taken from the Bureau of Economic Analysis (BEA) and insurance company

estimates which can be easily matched to the output of industries classified by four-digit SIC. I consider this measure of durability an imperfect measure of the concept of interest: how easily consumers can shift the purchase of a good through time. Bils and Klenow (1995) also use the Consumer Expenditure Survey (CEX) to construct Engel curves by good/industry for the same industries. However, they have difficulties with missing data because many households in the data do not purchase every good. They report these fractions, which for their work represent a nuisance. For the tests that I wish to conduct, these data are another imperfect measure of the concept of interest. From the reported statistics, I construct a variable that represents the percent of households which do not purchase a given good during a one year period. I set the number to 0 for all nondurable industries (and also use a 0 for the measure of durability). There is a high correlation between the two measures. For example, motorcycles are not purchased by 99 percent of households, refrigerators and freezers by 92 percent, and blinds and shades by 91 percent. Tables F.1 and F.2 in Appendix *F* gives the industry names, SIC codes, frequency of purchase and durability measures.

After eliminating industries for which frequency and durability data are not available, there are 109 industries.³⁷ Due to a large number of outliers in the first year, I use data from 1959 to 1991 on each industry, leaving a total of 3597 observations.

Third, I use measures of industry four-firm concentration ratios to capture differences in cyclicity which are due to market power. Rotemberg and Saloner (1986), as discussed, argue that collusion plays an important role in price smoothing. I use data from Rotemberg and Woodford (1991) on industries at the 2-digit SIC level that are based on 1967, which is roughly the middle of my sample. The data are reported in Appendix *F*.

Finally, in order to capture the demand-driven fluctuations in sales, I use the now standard Hall-Ramey instruments: a dummy for the political party of the U.S. President, real Federal government defense spending, and the price of oil deflated by the GDP deflator. As discussed subsequently, I use these instruments interacted with the cross-

³⁷These industries include all subindustries of SIC codes 20 and 21, which are nondurable industries, and those set of industries employed by Bils and Klenow (1995).

industry measure of durability.

4.6.2 Testing Price Smoothness and Dynamics

In order to be able test whether prices are lower when demand is high, all regressions are performed in levels after log-detrending each time-varying series separately for each industry.³⁸ This procedure also has the advantage of removing fixed industry effects which might be correlated with the dependent variables. Since there is substantial industry-level serial correlation, all standard errors are calculated so as to be consistent in the presence of arbitrary serial correlation as:

$$(X'X)^{-1} \left(\sum_{j=1}^J X_j' e_j e_j' X_j \right) (X'X)^{-1} \text{ and}$$

$$(\hat{X}'\hat{X})^{-1} \left(\sum_{j=1}^J \hat{X}_j' e_j e_j' \hat{X}_j \right) (\hat{X}'\hat{X})^{-1}$$

for ordinary least squares and two-stage least squares respectively.³⁹ j indexes industries and $e_j \equiv Y_j - X_j \hat{\beta}$, a $T = 33$ by 1 vector.

The first row of Table 4.1 shows the results of the regressing the real price of final sales⁴⁰ on real final sales, the percent of households who purchase the good, these two variables interacted, a constant, and a time trend.⁴¹ A one percent increase in sales for a typical nondurable or frequently purchased good is associated with a 0.13 percent decrease in the real price of that good. A good purchased by only half of households in

³⁸That is, I regress the logarithm of the variable in question on a constant and a time trend and then treat the residual as the datum. This procedure is done to make the regressions compatible with the markup regressions. One of the standard methods used to construct markups involves creating a log-detrended series. That industries may have a stochastic trend does not present a problem for the current estimation since asymptotic properties are derived from the number of industries going to infinity rather than the time dimension.

³⁹In practice this has a large impact, decreasing estimated standard errors on average by a factor of 3 to 4.

⁴⁰The price deflator for final shipments divided by the consumer price index (then log-detrended).

⁴¹The time trend and constant are included because the first observation has been dropped. No substantive results change when these two variables are omitted. If time dummies are included instead of a time trend similar conclusions concerning statistical significance of the interaction terms are reached, although magnitudes vary somewhat.

Table 4.1: REAL PRICE REGRESSIONS

	SALES	SALES* %NOTBUY	SALES* CR4	ΔSALES	ΔSALES* %NOTBUY	ΔSALES* CR4	N*T
OLS	-0.13 (0.03)	0.10 (0.03)					3597
IV	2.63 (1.26)	-3.19 (1.55)					3597
	2.53 (1.33)	-3.24 (1.72)	0.32 (0.41)				3465
	2.47 (0.73)	-2.99 (0.90)		-2.07 (1.77)	2.46 (2.16)		3597
	2.69 (0.99)	-3.63 (1.40)	0.44 (0.55)	-4.34 (3.68)	1.23 (2.49)	8.47 (9.65)	3465

All regressions also include a constant, a time trend, and the percent of households who do not purchase the good.

The instrument set for all regressions includes a time trend, the Hall-Ramey instruments, the durability of the industry's good, and the same interacted with the Hall-Ramey instruments. The instrument set for regressions with differenced right-hand-side variables also includes the Hall-Ramey instruments and interactions once lagged.

a year sees a typical decline in price of 0.08 percent. The negative relationship between sales and price represents the fact that some output increases are driven by supply-side factors such as increases in productivity and decreases in the cost of factor inputs. Note also that since I am using a subset of manufacturing industries and since industry output includes intermediate goods there is no reason for the average real price response to be zero.

To isolate the response of price to demand fluctuations, I instrument the measures of sales with the Hall-Ramey instruments.⁴² For two distinct reasons, the instruments also include the interaction of durability and the Hall-Ramey series. First, an aggregate demand shock does not increase demand equally across all industries. Economic theory suggests that demand increases much more for more durable goods, as expenditures

⁴²The regressions were also conducted using the real aggregate personal consumption expenditures series from the NIPA. All conclusions are robust to using this alternative instrument.

must move large amounts to adjust stocks. The interaction term increases the explanatory power of the instruments significantly. Second, the frequency of purchase is only an imperfect measure of the concept of interest. By using durability instead of the frequency of purchase interacted with the Hall-Ramey instruments, I correct for a possible attenuation bias due to the fact that I am using a noisy measure of the true concept of interest.

The second row of Table 4.1 shows the results of two-stage least squares estimation. The typical frequently purchased good now sees a price rise of two and a half percent for each percentage increase in sales due to demand. For a good purchased by only half of households, this number falls to one percent, and the difference is statistically significant.⁴³ Including the measure of industry concentration (*CR4*) does not alter this conclusion, nor does including the first difference of sales and its interaction with the frequency of purchase measure. Thus the first implication of the theory is confirmed in the data: prices rise less in response to increases in demand for goods which are purchased less frequently.⁴⁴

Rows 4 and 5 add the first difference of the sales variables to the regression in order to test whether, for infrequently purchased goods, prices are lower when quantities are increasing. Prices are lower when sales are increasing in general, and there is no statistically significant effect of frequency of purchase on this relationship.⁴⁵ Thus, the relationship predicted to hold for the subset of goods for which timing is important holds for all goods. It may well be that this additional general force which lowers prices when

⁴³The large change in the coefficient on sales in the theoretically predicted direction is evidence of good instruments. Further, the fits of the first stages are good. In rows 2 and 3, %NotBuy is predicted with an R^2 of 0.74; sales with an R^2 of 0.09; sales interacted with %NotBuy with an R^2 of 0.12. First differenced sales are predicted with an R^2 of 0.04, and, when interacted with %NotBuy, with an R^2 of 0.05.

⁴⁴Bils and Klenow (1995) do not include the interaction term and regress relative prices on relative labor-capital ratios in first differences and find insignificant and small relationship, even when they instrument. The main differences are that I am working in log-deviations from trend and more importantly that I use sales and an interaction term as my explanatory variables. A regression without the interaction term yields an insignificant coefficient on total sales.

⁴⁵Similar results are obtained if the change in sales at $t + 1$ is used instead of t (without a change in the timing of the instrument set).

sales are increasing causes the price-smoothing effect to become the dominant difference between goods for which timing is important and those for which it is not. Thus, prices are smoother for this subset of goods rather than lower when sales are increasing.

While prices are smoother for goods for which timing is important, there remains the possibility that marginal costs are heterogeneous across industries in just such a way as to generate smoother prices for less frequently purchased goods. That is, some combination of higher returns to scale and more elastic factor supplies implies that marginal costs are flatter or even decreasing for those goods which I find have smoother prices. To rule out this possibility, I perform a similar set of regressions using markups as the dependent variable. In doing so, I also seek to quantify the contribution of consumer intertemporal substitution of purchases to the cyclicity of the markup and thus the cyclical variability of production and sales.

4.6.3 Constructing Markups

Measuring markups is a difficult and controversial undertaking. Thus, I analyze three different constructed measures of markups, each based on a slightly different set of assumptions. The starting point for all of the measures is a standard production function in which real gross output is produced from labor input, capital, and intermediate goods.⁴⁶

$$Y = AF(L, K, M). \quad (4.21)$$

I omit time and industry subscripts for notational simplicity. Assuming that firms are price takers in factor markets and that factors are freely variable, cost minimization implies:

$$\frac{F_{JJ}}{Y} = \lambda \frac{P_{JJ}}{PY}, \quad (4.22)$$

⁴⁶I experimented with including production and nonproduction workers as separate inputs. Conclusions reached throughout this alternative analysis were similar if not slightly more favorable to the theory being tested. I chose to report this method since the only measure of compensation of production workers is wages, which is likely significantly more cyclical than total compensation of production workers. Thus, I use total payroll for all workers to measure the cost of labor input.

where P_J is the price of input J , λ is the Lagrange multiplier on the output constraint, and J is any factor for which the marginal product, F_J , is strictly positive and bounded for strictly positive and bounded levels of J . If this is true for all inputs, $\lambda = \mu$, the markup, defined as price divided by marginal cost. Then, the elasticity of output with respect to each factor input equals the markup times the ratio of the input's cost to total revenue.⁴⁷ Finally, define γ as the degree of returns to scale of the production function so that:

$$PY = \frac{\mu}{\gamma} \sum_J P_J J. \quad (4.23)$$

The first measure of markups which I consider is derived from three assumptions. First, γ is assumed constant across time for each industry. Second, on average there are no pure profits in each industry, so that revenues equal costs for each industry over the sample: $\frac{1}{T} \sum_t PY = \frac{1}{T} \sum_t (\sum_J P_J J)$. Finally, I assume that capital is quasi-fixed. Rearranging equation (4.23) and taking log-deviations from trend, the first markup measure is:

$$\widehat{\mu 1} = \frac{PY - \overline{PY} - (P_L L - \overline{P_L L}) - (P_M M - \overline{P_M M})}{\overline{PY}}, \quad (4.24)$$

where $\widehat{x} \equiv \frac{x - \bar{x}}{\bar{x}}$ and \bar{x} is the log-trend in x . While the assumption of capital fixity is rather crude, because the real capital stock does not move much over the cycle, the effects of possible fixity are small.⁴⁸

The second measure of markups is constructed by adding the additional assumption that the marginal product of labor is proportional to the ratio of labor input to real output. This is true, for example, of a Cobb-Douglas production function. Substituting

⁴⁷Hall (1988) originates the use of this methodology to estimate marginal costs (and thus markups). See Basu and Fernald (1995b) for a discussion of this general methodology and the importance of using gross output data.

⁴⁸Basu (1993) assumes freely variable capital and argue that the effects of fixity are likely small. The argument applies here in reverse. Further, an alternative approach is to assume that capital is freely variable and construct the nominal cost of capital. This can be done rather crudely under the assumption that capital is freely variable, so that $(P_K K)_t = P_{t-1}^I K_{t-1} + P_t^I I_t - P_t^I K_t$ where P^I is the price deflator for new investment, and I is new investment. When tried, the results are similar to those reported for $\mu 1$, but with slightly larger standard errors due most likely to the additional error that the noisy measure of the return to capital introduces.

into equation (4.22) and taking log-deviations from trend yields

$$\widehat{\mu 2} = \left(\frac{\widehat{PY}}{\widehat{P_L L}} \right), \quad (4.25)$$

Finally, I follow the method of Benabou (1992) that extends the procedure of Rotemberg and Woodford (1991) to include intermediate goods. First, one assumes that intermediate goods are used in strict proportion to output and that the production function exhibits constant returns to scale but there may be fixed costs. Therefore equation (4.22) applies only to capital and labor, $\lambda = \frac{\mu}{1 - \mu S_M}$ where S_M is the share of intermediate goods costs in total revenue, and equation (4.23) has $\gamma = 1$. Next, one assumes that free entry leads to the elimination of pure profits, so that the average cost shares of each input for each industry sum to one. Finally, one assumes that the cost share of labor in value added is equal to one minus the cost share of capital in value added. This assumption allows one to avoid having to calculate a cost of capital series, and it can be justified by assuming that capital and labor are combined using a Cobb-Douglas technology. Taking log-deviations and rearranging (see Benabou (1992)), I define my third markup series as

$$\widehat{\mu 3} = \frac{1}{1 + S_M \left(\frac{\mu}{1 - S_M} - 1 \right)} \left[-\widehat{S_H} - (1 - S_M - \mu) (S_H \widehat{H} + S_K \widehat{K} + S_M \widehat{S_M}) \right]. \quad (4.26)$$

As before, all hatted variables are log deviation from trend (for each industry), while variables without hats represent sample averages (again by industry). S_M , S_H , S_K represent the share of intermediate goods, labor, and capital in total revenue; and μ is the average markup. The only difference between this equation and equation (8) in Benabou (1992) is the term $\frac{\widehat{S_H}}{1 - S_M}$ and using $\frac{\widehat{S_M}}{1 - S_M}$ instead of the price series employed in Benabou (1992). The first is correcting a typo; the second substitution is taken because nominal shares are likely to be better measured than price deflators. Steady state and log-deviations can all be calculated from the NBER productivity database, except for the average markup which is set to 1.20 based on recent consensus.⁴⁹

⁴⁹See for example Basu (1993).

Table 4.2: MARKUP REGRESSIONS

MARKUP SERIES		SALES*		Δ SALES*	
		SALES	%NOTBUY	Δ SALES	%NOTBUY
μ_1	OLS	-0.12 (0.39)	0.12 (0.63)		
	IV	0.31 (0.63)	0.29 (0.93)		
	IV	-0.23 (0.73)	0.23 (1.00)	1.78 (2.32)	-2.48 (3.26)
μ_2	OLS	0.14 (0.02)	-0.11 (0.03)		
	IV	0.86 (0.43)	-1.12 (0.54)		
	IV	0.74 (0.42)	-1.00 (0.53)	-1.63 (1.05)	2.03 (1.27)
μ_3	OLS	-0.04 (0.01)	0.01 (0.02)		
	IV	-0.98 (0.40)	1.14 (0.49)		
	IV	-0.96 (0.25)	1.11 (0.31)	0.02 (0.47)	0.08 (0.56)

All regressions also include a constant, a time trend, and the percent of households who do not purchase the good.

The instrument set for all regressions includes a time trend, the Hall-Ramey instruments, the durability of the industry's good, and the same interacted with the Hall-Ramey instruments. The instrument set for regressions with differenced right-hand-side variables also includes the Hall-Ramey instruments and interactions once lagged.

4.6.4 Testing Markup and Dynamics

In Table 4.2, I present the results of regressions in which the dependent variable is the markup of price over marginal costs. The three pairs of rows each contain the results for one markup series. OLS regressions show markups to be roughly acyclical, falling between the findings of Domowitz, Hubbard and Petersen (1988) and Rotemberg and Woodford (1991).

Instrumental variables regressions capture the change in markups associated with demand-driven fluctuations. These regressions do not give clean answers about either markup hypothesis. First, markups seem slightly more countercyclical in response to demand fluctuations. Second, markups are more countercyclical for goods for which timing is more important only for the second markup measure. Evidence from the first markup series is inconclusive and evidence from the third shows less frequently purchased goods to have more procyclical markups. Third, the three series also give contradictory and weak evidence as to whether markups are lower when sales of infrequently purchased goods are increasing.⁵⁰

Why are the results so unstable and inconclusive? One possibility is that differences in the construction of the three markup series generate different answers. However, it is also possible that the pattern of industry-specific slopes of marginal cost is confounding inference. In industries with short-run increasing returns to scale, real price stickiness may increase markups in booms. That is, as the theoretical sections discuss, the impact of frequency of purchase depends on the slope of the marginal cost curve. To test this, I reestimate the markup regressions on two subsamples of industries.

First, I use only those industries that are in 2-digit industries which Basu and Fernald (1995a) find have decreasing returns to scale. The results for this subsample of industries are similar to the results reported in Table 4.2. This is not wholly surprising given that returns to scale is only one component of marginal cost, and differences in

⁵⁰When the regressions include industry concentration interacted with the quantity dependent variable, and/or its first difference, the interaction term is never significant. As in the price regressions, the addition of industry concentration variables does not alter the significance or magnitude of other coefficients.

Table 4.3: MARKUP REGRESSIONS ON SUBSAMPLE

MARKUP SERIES		SALES*		Δ SALES*	
		SALES	%NOTBUY	Δ SALES	%NOTBUY
$\mu 1$	IV	1.25 (0.73)	-1.41 (0.91)		
	IV	1.27 (0.60)	-1.42 (0.77)	0.09 (0.53)	-0.14 (0.85)
$\mu 2$	IV	1.58 (0.89)	-1.49 (0.77)		
	IV	-1.52 (0.73)	1.61 (0.85)	-0.25 (0.82)	0.40 (0.96)
$\mu 3$	IV	0.94 (0.54)	-1.06 (0.63)		
	IV	0.97 (0.41)	-1.10 (0.48)	0.23 (0.40)	-0.20 (0.47)

All regressions also include a constant, a time trend, and the percent of households who do not purchase the good.

The instrument set for all regressions includes a time trend, the Hall-Ramey instruments, the durability of the industry's good, and the same interacted with the Hall-Ramey instruments. The instrument set for regressions with differenced right-hand-side variables also includes the Hall-Ramey instruments and interactions once lagged.

factor elasticities may well be more important. Thus, as a second cut, I examine only the subsample of industries in which the instrumented correlation of sales and price is positive.⁵¹ This leaves 1815 observations. As is shown in Table 4.3, there is evidence that markups are less procyclical for infrequently purchased goods.⁵² There is little evidence, however, of a consistent relationship between markups and whether demand is increasing or decreasing.

In sum then, cross industry evidence suggests that prices are smoother in industries where the timing of purchases is important. However, the timing variable is potentially

⁵¹That is, for each industry separately and using the usual instruments, I run price on sales, a time trend, and the percent of households not buying the good. Then I use only those industries for which the coefficient on sales is positive.

⁵²There remains one puzzle however, which is that the relationship between price smoothness and the timing variable is reversed in this subsample.

correlated with the industry-specific slope of marginal cost. Unfortunately, evidence on markups which attempt to measure both marginal cost and price are less conclusive.

4.7 Conclusion

This chapter presents a model in which consumers' ability to time their purchases of goods amplifies their effective elasticity of intertemporal substitution. When firms have some market power, fluctuations in consumer demand are fluctuations in the elasticity of demand and lead to potentially important price dynamics. When marginal costs are increasing, the markup of price over marginal cost is countercyclical, and demand fluctuations are smoothed over time, or propagated.

Using industry data, I demonstrate that, as predicted by the model, the price responses to fluctuations in demand are smaller for those goods for which the timing of purchases is more important. But the evidence on the behavior of markups is less clear. Only a shred of evidence is found that markups are more countercyclical (or less procyclical) for those goods purchased most infrequently. In future empirical work, I plan to examine specific industries in which the timing of purchases is thought to be important and for which measures of marginal cost are simpler to come by.

Throughout this chapter, I assume that the distribution of consumers over the old stocks is exogenous. However, consider the example discussed in the introduction, except allow consumers to be distributed between two bands of adjustment. Now, when the distribution of consumers is perturbed, the endogenous response of sellers with a constant or decreasing marginal cost technology amplifies the shock. The distribution of consumers is moved further from a uniform distribution. The next time this larger group prepares to purchase new goods, the producers' pricing strategies may add still more consumers to this group, so that the demand fluctuation grows. It is possible that if the amount of individual uncertainty is not too great, the market may exhibit stable repeating fluctuations. Under these circumstances, it would be surprising if we did not see demand-driven business cycles.

Future research will embed these timing considerations in a general equilibrium stochastic growth model. The structure may be able to address two significant shortcomings in current general equilibrium business cycle theory. First, consideration of optimal timing of sales provides a propagation mechanism that may significantly improve the empirical fit of the model. Second, this extension may provide a theoretical alternative for the *ad hoc* assumption common in the Real Business Cycle literature that technology shocks are highly serially correlated. Instead of serially correlated technology, timing considerations could cause a single uncorrelated demand shock to generate changes in the markup lasting several periods. In the partial equilibrium model examined, a positive shock to demand is smoothed and a countercyclical markup arises. Thus, as the markup returns to normal and sales increase it might appear as if technology were improving in a highly serially correlated manner.

Figure 4-1: Example Market Dynamics

Figure 1A




Figure 1B




Figure 1C




Figure 4-2: Fluctuations with Constant Marginal Costs

Figure 2A
Price over a Cycle




Figure 2B
Actual Sales and Sales at Steady-State Price




Figure 2C
Cutoff Utility Flow at Start of Period




Figure 4-3: Fluctuations with Increasing Marginal Costs

Figure 3A
Price over a Cycle




Figure 3B
Actual Sales and Sales at Steady-State Price




Figure 3D
Cutoff Utility Flow at Start of Period




Figure 3C
Price Less Marginal Cost




Figure 4-4: Fluctuations with Colluding Firms and Increasing Marginal Costs

Figure 4A
Price over a Cycle




Figure 4B
Actual Sales and Sales at Steady-State Price




Figure 4C
Cutoff Utility Flow at Start of Period





Figure 4D
Price Less Marginal Cost



Appendix E

Proofs of Lemmas

In this appendix I restate and prove the lemmas in Chapter 4.

Lemma 1 *No search. In a symmetric equilibrium, no consumers search and*

$$V_t(u, p_{jt}) = \text{Max}\{u + \beta E_t[V_{t+1}(u - \delta, P_{t+1})], v - p_{jt}\} \quad (\text{E.1})$$

Proof. Since all firms charge the same price in every period, any search has a total expected gain of minus the search cost. ◊

Lemma 2 *Skimming property. $u_{it} > u_{jt}$ implies $T_{ti} \leq T_{tj}$.*

Proof. A buyer, i , with utility flow u_{it} can exactly imitate the strategy of a buyer, j , with a lower utility flow, in which case i receives the same return from purchasing but greater utility flow in every period before purchase. Thus $V_t(u, p_j)$ is weakly increasing in its first argument. Consider now the decision of each buyer as to whether to buy in t or wait, as captured by equation (E.1). Given that both u and $E_t[V_t(u - \delta, P_{t+1})]$ are greater for buyer i , the buyer with the lower utility flow, j , will always choose to purchase if buyer i does, and may choose to do so when buyer i does not. ◊

Lemma 3 *u_t^* evolution. Provided that sellers sell to some consumers in every period, u_t^* is defined by*

$$u_t^* = (1 - \beta)v + \beta_t P_{t+1} - P_t \quad (\text{E.2})$$

Proof. Consider consumers who are indifferent between purchasing in the current period and waiting. Allow the equilibrium to involve some of these indifferent consumers purchasing in the current period and some delaying their purchases.¹ Then it follows from Lemma (2), that all those with lower utility flows buy in the current period. Those with higher utility flows delay since $V_t(u, p_j)$ is increasing in u , $u + \beta E_t[V_{t+1}(u - \delta, P_{t+1})]$ is strictly increasing in u — the return to delaying is strictly increasing in u . If positive sales are made in every period, then those who are indifferent between purchasing and delaying in t will purchase in $t + 1$. u_t^* is then defined by indifference in equation (E.1) as

$$\begin{aligned} v - P_t &= u_t^* + \beta E_t[V_{t+1}(u_t^* - \delta)] \\ &= u_t^* + \beta(v - P_{t+1}) \end{aligned} \tag{E.3}$$

Rearranging yields equation (E.2).◊

Lemma 4 *Positive sales.* If $\varepsilon_t(u) \equiv 0$, $q_0 > 0$, and $p_t \geq c_1 \forall t$, then $q_t > 0 \forall t$.

Proof. First suppose market price were such that no sales were being made in period t . Profits to all sellers are zero. Then any individual seller could choose a price an arbitrarily small distance above the marginal cost at zero sales, c_1 , and, if it made positive sales, make a profit in t . Since there are an infinite number of sellers, selling to some consumers does not reduce expected future profits noticeably. Thus, any supposed market price greater than c_1 cannot coexist with zero sales. Suppose sales are made in period t and in period $T > t + 1$ and no sales are made between these dates. In period $T - 1$, the highest utility flow buyer weakly prefers buying in T :

$$v - P_{T-1} \leq u_{T-1}^+ + \beta(v - P_T). \tag{E.4}$$

At the end of period t , the highest utility flow consumer weakly prefers purchasing in t to purchasing in all other periods including $t + 1$:

$$v - P_t \geq u_t^* + \beta(v - P_{t+1}). \tag{E.5}$$

¹Since these consumers are measure zero to firms, whether they all purchase, wait, or mix is irrelevant for the equilibrium.

Since zero sales are made during the period between T and t , the highest utility flow individuals are the same and u evolves as: $u_s^* = u_s^+ = u_{s-1}^* - \delta$. Using this to eliminate the utility flows from equations E.4 and E.5 yields:

$$v - P_{T-1} \leq -(T - 1 - t)\delta + (1 - \beta)v + \beta P_{t+1} - P_t + \beta(v - P_T)$$

Note that in periods t and T sales are positive so that $P_t > c_1$, and $P_T > c_1$. while in periods $t + 1$ and $T - 1$, sales are zero so prices must be less than or equal to c_1 . Making these substitutions preserves the inequality and yields

$$0 \leq -(T - 1 - t)\delta \tag{E.6}$$

which can only be true if $T = t + 1$, that is if there is no intermediate period with no sales.◊

Lemma 5 *Steady state equilibrium. The steady state always exists is uniquely determined by*

$$\begin{aligned} P_{ss} &= c_1 + c_2\alpha\delta + \delta \\ u_{ss}^+ &= (1 - \beta)(v - P_{ss}) - \delta \\ Q_{ss} &= \alpha\delta \end{aligned}$$

Proof: In order for u_{ss}^+ to remain constant, Q_{ss} must equal $\alpha\delta$. Plugging this and the distribution function into the seller first-order condition (4.15) yields a unique P_{ss} . Equation (4.14) then gives a unique u_{ss}^+ . There is thus a unique candidate for a steady-state equilibrium. Existence then follows from the fact that the seller profit function is concave, a fact easily checked. The proof for the collusive/monopolist case is identical and omitted.◊

Appendix F

Industry Data

A complete description of the NBER productivity database can be found in Bartelsman and Gray (1994). The database and Bartelsman and Gray (1994) can be downloaded from the \pub\productivity directory on nber.harvard.edu by anonymous ftp. The main weakness of the database for my purposes is that the measure of sales covers all sales by firms within the SIC code. That is, my measure of sales includes sales of non-final goods. The buyers of the intermediate goods sold may have somewhat different abilities to time purchases than the buyers of the final consumer goods, which generate the measures of durability and frequency of purchase. All SIC codes are based on the 1972 categorization, as used by the NBER productivity database. Gross nominal production is calculated as total revenue less change in nominal inventories. Nominal inventories in $t - 1$ are multiplied by the current inventories price deflator and divided by the lagged inventories price deflator to make the nominal change consistent. These variables and nominal payments to labor and intermediate goods are included in the NBER database.

The measure of frequency of purchase represents the percent of households not reporting any consumption expenditures on items in this SIC code during a one-year period (1986). SIC codes defined as nondurable have these measures set to zero. These industries are all subindustries of 2-digit SIC code 20 and 21, food and kindred products and tobacco respectively. Tables F.1 and F.2 list the measures of infrequency of purchase and the measures of durability. Durability measures are based on the life expectancy tables of a major U.S. insur-

Table F.1: DURABILITY AND INFREQUENCY OF PURCHASE, PART I

SIC CODE	INDUSTRY	PERCENT NOT BUYING	DURABILITY (YEARS)
2251	WOMEN'S HOSIERY	33.9	1.0
2252	MEN'S HOSIERY	49.2	1.7
2271	WOVEN CARPETS AND RUGS	82.6	11.1
2272	TUFTED CARPETS AND RUGS	82.6	11.1
2279	CARPETS AND RUGS, NEC.	82.6	11.1
2311	MEN'S SUITS AND COATS	53.3	4.1
2321	MEN'S SHIRTS AND NIGHTWEAR	34.9	2.7
2322	MEN'S UNDERWEAR	56.1	2.2
2327	MEN'S TROUSERS	34.6	2.7
2328	MEN'S WORK CLOTHING	34.6	2.7
2331	WOMEN'S BLOUSES	36.6	2.3
2335	WOMEN'S DRESSES	49.9	4.0
2337	WOMEN'S COATS	55.5	4.3
2341	WOMEN'S UNDERWEAR	32.7	1.8
2342	BRASSIERS, GIRDLES, ETC.	32.7	1.8
2361	GIRL'S DRESSES AND BLOUSES	82.9	2.3
2391	CURTAINS AND DRAPES	82.9	4.2
2511	WOOD FURNITURE	61.7	8.1
2512	WOOD FURN. UPHOLSTERED	61.7	8.1
2514	METAL FURNITURE	61.7	8.1
2515	MATTRESSES AND BEDS	89.3	15.0
2591	BLINDS AND SHADES	90.2	10.9
2711	NEWSPAPERS*	0.0	0.0
2721	MAGAZINES*	0.0	0.0
2731	BOOKS PUBLISHING	44.5	11.0
2732	BOOKS PRINTING	44.5	11.0
2834	PRESCRIPTION DRUGS	0.0	0.0
2911	FUEL OIL AND GASOLINE	0.0	0.0
2992	MOTOR OIL	0.0	0.0
3011	TIRES	35.2	3.0
3143	MEN'S FOOTWARE	49.0	2.5
3144	WOMEN'S FOOTWARE	31.6	2.6

Source: Bils and Klenow (1995).

*These industries have their percent-not-buying measures set to zero since these goods are not purchased by everyone, yet they are nondurable in the sense that one chooses to buy the current issue or not at all. None of the results change significance or sign with this adjustment.

ance company. These life expectancies of goods are used by the company to adjust insurance claims for covered damages to these items and are weighted aggregates of slightly finer classifications. Durability measures for a subset of the industries (e.g. automobiles) are taken from Fixed Reproducible Tangible Wealth, 1925-89 by the Bureau of Economic Analysis. The reader is referred to Bils and Klenow (1995) for further details.

Industry concentration measures, taken from Rotemberg and Woodford (1991), estimate the share of total final sales accounted for by the four largest firms in 1967, roughly the midpoint of my sample. The concentration ratios are at the 2-digit level except motor vehicles and other transportation equipment which are split. They are as follows: SIC 20 : 0.345, SIC 21 : 0.736, SIC 22 : 0.341, SIC 23 : 0.197, SIC 25 : 0.216, SIC 27 : 0.189, SIC 28 : 0.499; SIC 29 : 0.329; SIC 30 : 0.691, SIC 31 : 0.245, SIC 32 : 0.374, SIC 35 : 0.363, SIC 36 : 0.450, SIC 371 : 0.808, SIC 372 – 9 : 0.501 ; SIC 38 : 0.478.

Table F.2: DURABILITY AND INFREQUENCY OF PURCHASE, PART II

SIC CODE	INDUSTRY	PERCENT NOT BUYING	DURABILITY (YEARS)
3161	LUGGAGE	88.5	17.5
3229	GLASSWARE	84.3	10.0
3262	CHINA	84.1	17.5
3263	COOKWARE	79.4	17.5
3524	LAWNMOWERS	85.7	7.5
3631	STOVES AND OVENS	88.3	14.1
3632	REFRIGERATORS AND FREEZERS	92.2	15.0
3633	WASHERS AND DRIERS	91.9	11.0
3634	PORTABLE HEATERS	87.7	11.3
3635	VACUUM CLEANERS	91.2	9.5
3645	LAMPS	83.9	16.7
3651	TV'S, VCR'S, AND STEREOS	53.3	11.9
3652	RECORDS AND TAPES	49.7	5.0
3661	TELEPHONES	80.7	7.1
3711	AUTOMOBILES	91.2	10.0
3713	LIGHT TRUCKS AND VANS	97.1	8.0
3732	BOATS	99.1	10.0
3751	MOTORCYCLES	90.2	8.6
3792	TRAILERS AND CAMPERS	98.6	8.0
3851	EYEGASSES AND CONTACTS	68.3	10.0
3861	FILM AND PHOTO EQUIP.	39.2	6.7
3873	CLOCKS AND WATCHES	64.1	15.5
3911	JEWELRY	55.8	5.5
3914	SILVERWARE	91.4	27.5
3931	MUSICAL INSTRUMENTS	91.5	13.0
3944	GAMES AND TOYS	41.9	5.0

Source: Bills and Klenow (1995).

References

- Abowd, John and David Card**, "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 1989, 57, 411–45.
- Altonji, Joseph and Aloysius Siow**, "Testing the Response of Consumption to Income Change with (Noisy) Panel Data," *Quarterly Journal of Economics*, 1987, 103, 293–328.
- Attanasio, Orazio**, "Personal Saving in the US," in James M. Poterba, ed., *International Comparisons of Household Saving*, University of Chicago Press, 1994.
- ___ and **Guglielmo Weber**, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, Dec 1995, 103 (6), 1121–57.
- ___, **James Bank, Costas Meghir, and Guglielmo Weber**, "Humps and Bumps in Lifetime Consumption," Nov 1995. mimeo.
- Attanazio, Orazio and Martin Browning**, "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, Dec 1995, 85, 1118–1137.
- Ayagari, Rao**, "Uninsured Idiosyncratic Risk and Aggregate Savings," *Quarterly Journal of Economics*, 1993, 109, 659–684.
- Barro, Robert**, "Are Government Bonds Net Wealth?," *Journal of Political Economy*,

1974, *82*, 1095–1117.

Barsky, Robert, Thomas Juster, Miles Kimball, and Matthew Shapiro,

“Preference Parameters and Behavioral Heterogeneity: an Experimental Approach in the Health and Retirement Survey,” 1995. manuscript, University of Michigan.

Bartelsman, Eric and Wayne Gray, “NBER Manufacturing Database,” December 1994. NBER Manuscript.

Basu, Susanto, “Constant Returns and Small Markups in U.S. Manufacturing,” April 1993. Center for Research on Economic and Social Theory and Department of Economics Working Paper Number 93-19, University of Michigan.

___ and **John G. Fernald,** “Aggregate Productivity and the Productivity of Aggregates,” December 1995. NBER Working Paper 5382.

___ and ___, “Are Apparent Productivity Spillovers a Figment of Measurement Error?,” *Journal of Monetary Economics*, August 1995.

Benabou, Roland, “Inflation and Markups,” *European Economic Review*, 1992, *36*, 566–74.

Bils, Mark, “Pricing in a Customer Market,” *Quarterly Journal of Economics*, November 1989, *104*, 699–718.

___, “Cyclical Pricing of Durable Goods,” September 1991. University of Chicago GSB mimeo.

___ and **Peter J. Klenow,** “Uncovering Curvature— Tests of Competing Business Cycle Models,” March 1995.

Blanchard, Olivier J., “Debt, Deficits, and Finite Horizons,” *Journal of Political Economy*, 1985, *93*, 223–47.

Bond, Eric W. and Larry Samuelson, “The Coase Conjecture Need not Hold for Durable Good Monopolies with Depreciation,” *Economic Letters*, 1987, *24*, 93–97.

Bound, John, David A. Jaeger, and Regina Baker, “Problems with Instrumental Variables Estimation when the Correlation between the Instruments and the Endogenous Explanatory Variable is Weak,” *Journal of the American Statistical Association*, Jun 1995, *90*, 443–50.

- Branch, E. Raphael**, "The Consumer Expenditure Survey: a comparative analysis," *Monthly Labor Review*, 1994, pp. 47–55.
- Bulow, Jeremy**, "Durable Goods Monopolists," *Journal of Political Economy*, 1982, 90, 314–32.
- Bureau Of Labor Statistics**, *Consumer Expenditure Survey, 1993: Interview Survey*, Washington, DC: US Department of Commerce, Bureau of the Census, 1980-1993. Ann Arbor, MI: Inter-university Consortium of Political and Social Research.
- Caballero, Ricardo J.**, "Consumption Puzzles and Precautionary Saving," *Journal of Monetary Economics*, 1990, 25, 113–36.
- ___, "Expenditure on Durable Goods: A Case for Slow Adjustment," *Quarterly Journal of Economics*, Aug 1990, 105, 727–43.
- ___, "Durable Goods: An Explanation for their Slow Adjustment," *Journal of Political Economy*, April 1993, 101, 351–84.
- ___ and **Mohamad L. Hammour**, "The Cleansing Effect of Recessions," *American Economic Review*, December 1994, 84, 1350–68.
- ___, **Eduardo M. R. A. Engel**, and **John C. Haltiwanger**, "Plant Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 1995, 1, 1–39.
- Campbell, John Y. and N. Gregory Mankiw**, "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in Oliver Jean Blanchard and Stanley Fischer, eds., *N.B.E.R. Macroeconomics Annual*, Cambridge: MIT Press 1989, pp. 185–215.
- Carroll, Christopher D.**, "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992, 2, 61–135.
- ___, "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," 1993. Federal Reserve Board mimeo.
- ___, "How Does Future Income Affect Current Consumption?," *Quarterly Journal of Economics*, 1993, 109, 111–147.

- ___ and **Andrew A. Samwick**, "The Nature of Precautionary Wealth," 1994.
manuscript.
- ___ and **Lawrence H. Summers**, "Consumption Growth Parallels Income Growth: Some New Evidence," in B.D. Bernheim and John B. Shoven, eds., *National Savings and Economic Performance*, Chicago University Press 1991.
- ___ and **Miles S. Kimball**, "On the Concavity of the Consumption Function," *Econometrica*, forthcoming.
- Chevalier, Judith and David Scharfstein**, "Capital Market Imperfections and Countercyclical Markups: Theory and Evidence," Jan 1994. NBER Working Paper No. 4614.
- Clarida, Richard**, "Aggregate Stochastic Implications of the Life-Cycle Hypothesis," *Quarterly Journal of Economics*, 1991, 106, 851-867.
- Coase, Ronald**, "Durability and Monopoly," *Journal of Law and Economics*, 1972, 15, 143-49.
- Conlisk, John, Eitan Gerstner, and Joel Sobel**, "Cyclic Pricing by a Durable Goods Monopolist," *Quarterly Journal of Economics*, August 1984, 99, 489-505.
- Council Of Economic Advisors**, *The Economic Report of the President*, Washington, DC: Government Printing Office, 1995.
- Deaton, Angus**, "Panel Data from Time Series and Cross Sections," *Journal of Econometrics*, 1985, 30, 109-26.
- ___, "Saving and Liquidity Constraints," *Econometrica*, 1991, 59, 1121-1142.
- ___, *Understanding Consumption*, Oxford: Clarendon Press, 1992.
- Diamond, Peter A.**, "A Model of Price Adjustment," *Journal of Economic Theory*, 1971, 3, 156-68.
- ___, *On Time*, Cambridge, UK: Cambridge University Press, 1994.
- Domowitz, Ian, R. Glenn Hubbard, and Bruce Petersen**, "Market Structure, Durable Goods, and Cyclical Fluctuations in Markups," June 1987. Mimeo, Northwestern University.
- ___, ___, and ___, "Market Structure and Cyclical Fluctuations in U.S.

- Manufacturing," *Review of Economics and Statistics*, February 1988, 70, 55–66.
- Duffie, Darrell and Kenneth Singleton**, "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, July 1993, 61, 929–952.
- Dynan, Karen**, "How Prudent Are Consumers?," *Journal of Political Economy*, Dec 1993, 101, 1104–1113.
- Flavin, Marjorie**, "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, 1981, 89, 974–1009.
- Friedman, Milton**, *A Theory of the Consumption Function*, Princeton: Princeton University Press, 1957.
- Gross, David B.**, "The Investment and Financing Decisions of Liquidity Constrained Firms," 1994. MIT mimeo.
- Hall, Robert E.**, "The Stochastic Implications of the Life Cycle Permanent Income Hypothesis," *Journal of Political Economy*, 1978, 86, 971–87.
- ___, "Intertemporal Substitution in Consumption," *Journal of Political Economy*, Apr 1988, 96.
- ___ and **Frederic S. Mishkin**, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 1982, 50, 461–481.
- Hansen, Lars**, "Large Sample Properties of Method of Moments Estimators," *Econometrica*, 1982, 50, 1029–1054.
- Heaton, John**, "The Interaction Between Time-nonseparable Preferences and Time Aggregation," Mar 1993.
- Hubbard, Glenn, Jonathan S. Skinner, and Stephen P. Zeldes**, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in Allan H. Meltzer and Charles I. Plosser, eds., *Carnegie-Rochester Conference Series on Public Policy*, Vol. 40 Jun 1994, pp. 59–125.
- ___, **Jonathan Skinner, and Stephen Zeldes**, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, Apr 1995, 103, 360–397.
- Japelli, Tulio, Jörn-Steffen Pischke, and Nicholas Souleles**, "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," 1995.

MIT mimeo.

Judd, Kenneth, "Projection Methods for Solving Aggregate Growth Models,"
Journal of Economic Theory, 1992, 58, 410–452.

—, "Numerical Methods in Economics," 1993. manuscript, Hoover.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large,"
Econometrica, 1990, 58, 53–73.

Kotlikoff, Laurence J. and Lawrence H. Summers, "The Contribution of
Intergenerational Transfers in Aggregate Capital Accumulation," *Journal of Political
Economy*, 1981, 86, 706–32.

Kreps, David and Evan Porteus, "Temporal Resolution of Uncertainty and
Dynamic Choice Theory," *Econometrica*, 1978, 46, 185–200.

Levenson, Alec R., "Do Consumers Respond to Future Income Shocks? Evidence
from Social Security Reform in Taiwan," 1993. mimeo, Milken Institute for Job and
Capital Formation.

Lusardi, Annamaria, "Euler Equations in Micro Data: Merging Data from Two
Samples," *Journal of Business and Economic Statistics*, forthcoming.

Modigliani, Franco and Richard Brumberg, "Utility Analysis and the
Consumption Function: an Interpretation of Cross-Section Data," in Kenneth K.
Kurihara, ed., *Post-Keynesian Economics*, New Brunswick, N.J. Rutgers University
Press 1956, pp. 388–436.

Mortensen, Dale T., "The Cyclical Behavior of Job and Worker Flows," *Journal of
Economic Dynamics and Control*, Nov 1994, 18 (6), 1121–42.

Nelson, C. E. and R. Startz, "Some Further Results on the Exact Small Sample
Properties of the IV Estimator," *Econometrica*, 1990, 58, 967–976.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization
Estimators," *Econometrica*, 1989, 57, 1027–1057.

Palumbo, Michael G., "Uncertain Medical Expenses and Precautionary Saving Near
the End of the Life Cycle," 1994. University of Houston Mimeo.

Phelps, Edmund S. and Sidney G. Winter, "Optimal Price Policy under

- Atomistic Competition,” in E.S. Phelps, ed., *Foundations of Inflation and Employment Theory*, New York, NY: Norton 1970.
- Pischke, Jörn-Steffen**, “Individual Income, Incomplete Information, and Aggregate Consumption,” *Econometrica*, Jul 1995, 63 (4), 805–840.
- Poterba, James M.**, “Are Consumers Forward Looking? Evidence from Fiscal Experiments,” *American Economic Review*, May 1988, 78, 413–18.
- Rogerson, Richard**, “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 1988, 21, 3–16.
- Rotemberg, Julio and Garth Saloner**, “A Supergame-Theoretic Model of Price Wars During Booms,” *American Economic Review*, June 1986, 76, 390–407.
- ___ and **Michael Woodford**, “Markups and the Business Cycle,” in Oliver Jean Blanchard and Stanley Fischer, eds., *N.B.E.R. Macroeconomics Annual, 1991*, Cambridge, MA: MIT Press 1991, pp. 63–129.
- Runkle, David**, “Liquidity Constraints and the Permanent Income Hypothesis,” *Journal of Monetary Economics*, 1991, 27.
- Shapiro, Matthew D. and Joel Slemrod**, “Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding,” *American Economic Review*, 1995, 85, 274–83.
- Shea, John**, “Union Contracts and the Life Cycle-Permanent Income Hypothesis,” *American Economic Review*, 1995, 85, 186–200.
- Social Security Administration**, *Social Security Bulletin Annual Statistical Supplement*, United States Department of Health and Human Services, Government Printing Office, 1990.
- Souleles, Nicholas**, “The Response of Household Consumption to Income Tax Refunds,” 1995. M.I.T. Mimeo.
- Stokey, Nancy**, “Rational Expectations and Durable Goods Pricing,” *Bell Journal of Economics*, 1981, 12, 112–28.
- Stokey, Nancy L., Robert E Lucas, and Edward Prescott**, *Recursive Methods in Economics Dynamics*, Harvard University Press, Cambridge, MA., 1989.

Venti, Steven and David Wise, "The Wealth of Cohorts: Retirement Saving and the Changing Assets of Older Americans," 1993. NBER Working Paper 4600.

Wilcox, David W., "Social Security Benefits, Consumption Expenditure, and the Life Cycle Hypothesis," *Journal of Political Economy*, Apr 1989, *97*, 288–304.

___ , "The Construction of U S. Consumption Data: Some Facts and Their Implications For Empirical Work," *American Economic Review*, Sep 1992, *82*, 922–941.

Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, Apr 1989, *97*, 305–346.

___ , "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 1989, *104*, 275–98.