Notes for Econ202A: Consumption

Pierre-Olivier Gourinchas
UC Berkeley

Fall 2014

Disclaimer: These notes are riddled with inconsistencies, typos and omissions. Use at your own peril.

1 Introduction

Where the second part of econ202A fits?

- Change in focus: the first part of the course focused on the big picture: long run growth, what drives improvements in standards of living.
- This part of the course looks more closely at pieces of models. We will focus on four pieces:
  - consumption-saving. Large part of national output.
  - investment. Most volatile part of national output.
  - open economy. Difference between S and I is the current account.
  - financial markets (and crises). Because we learned the hard way that it matters a lot!

2 Consumption under Certainty

2.1 A Canonical Model

A Canonical Model of Consumption under Certainty

- A household (of size 1!) lives T periods (from $t = 0$ to $t = T - 1$). Lifetime preferences defined over consumption sequences $\{c_t\}_{t=1}^{T}$:

$$U = \sum_{t=0}^{T-1} \beta^t u(c_t) \tag{1}$$

  where $0 < \beta < 1$ is the discount factor, c_t is the household's consumption in period t and $u(c)$ measures the utility the household derives from consuming c_t in period t. $u(c)$ satisfies the 'usual' conditions:

  - $u'(c) > 0$,
  - $u''(c) < 0$,
  - $\lim_{c \to 0} u'(c) = \infty$
  - $\lim_{c \to \infty} u'(c) = 0$

- Seems like a reasonable problem to analyze.

2.2 Questioning the Assumptions

Yet, this representation of preferences embeds a number of assumptions. Some of these assumptions have some micro-foundations, but to be honest, the main advantage of this representation is its convenience and tractability. So let's start by reviewing the assumptions:

- **Uncertainty**. In particular, there is uncertainty about what T is. Whose T are we talking about anyway? What about children? This is probably not a fundamental assumption. We will introduce uncertainty later. This is not essential for now.

- **Aggregation**. Aggregate consumption expenditures represent expenditures on many different goods: $c_t = \sum_i p_{i,t} c_{i,t}$ over commodities i (where I am assuming that aggregate consumption is the numeraire). If preferences are homothetic over individual commodities, then it is possible to 'aggregate' preferences of the form $u(\mathbf{c}, \mathbf{p})$ into an expression of the form $u(c)$ where $c = \mathbf{p}.\mathbf{c}$

- **Separation**. Other arguments enter utility: labor supply etc... The implicit assumption here is that preferences are separable over these different arguments: $u(c) + v(z)$.

- **Time additivity**. The marginal utility of consumption at time t only depends on consumption expenditures at that time.
  - What about durable goods, i.e. goods that provide utility over many periods? Distinction between consumption expenditures (what we pay when we purchase the goods) and consumption services (the usage flow of the good in a given period). The preferences are defined over consumption services but the budget constraint records consumption expenditures. Stock-flow distinction.
  - What if utility depends on previous consumption decisions, e.g. $u(c_t, H_t)$ where H_t is a habit level acquired through past consumption decisions? Habit formation would correspond to a situation where $\partial H_t / \partial c_s > 0$ for $s < t$ and $\partial^2 u / (\partial c \partial H) > 0$. In words: past consumption increases my habit, and a higher habit increases my marginal utility of consumption today. Internal habit.
  - What if utility depends on the consumption of others, e.g. $u(c_t, \bar{C}_t)$ where $\bar{C}_t$ is the aggregate consumption of 'others' (catching up with the Joneses). External habit. As the name suggests, external habit implies an externality of my consumption on other people's utility that may require corrective taxation).

- **Intertemporal Marginal Rate of Substitution**. Consider two consecutive periods t and $t+1$. The IMRS between t and $t+1$ seen from period 1 is $\beta^{t+1} u'(c_{t+1}) / \beta^t u'(c_t)$. The same IMRS seen from time t is $\beta u'(c_{t+1}) / u'(c_t)$. The two are equal! Key property that arises from exponential discounting (Strotz (1957)). Example: 1 apple now, vs 2 apples in two weeks. Answer should not change with the time at which we consider the choice (period 1 or period t). Substantial body of experimental evidence suggests that the present is more salient then exponential discounting.

Suppose instead that $U = u(c_0) + \theta \sum_{t=1}^{T-1} \beta^t u(c_t)$ with $0 < \theta < 1$ represent the lifetime preferences of the household in period 1. Notice that θ only applies to future utility (salience of the present). quasi-hyperbolic discounting (see Laibson (1996)).

The problem is that preferences become time-inconsistent: next period, the household would like to re-optimize if given a chance. Not the case with exponential discounting (check this):

$$max_{c_t, c_{t+1}, \ldots c_{T-1}} \sum_{s=1}^{t-1} \beta^s u(c_s^*) + \sum_{s=t}^{T-1} \beta^s u(c_s)$$

2.3 The Intertemporal Budget Constraint

Since there is no uncertainty, all financial assets should pay the same return (can you explain why?). Let's denote $R = 1 + r$ the gross real interest rate between any two periods, assumed constant. The budget constraint of the agent is:

$$a_{t+1} = R(a_t + y_t - c_t)$$

a_t denotes the financial assets held at the beginning of the period, and y_t is the non-financial income of the household during period t. [Note that this way of writing the budget constraint assumes that interest is earned 'overnight' i.e. as we transition from period t to $t + 1$.]
We can derive the intertemporal budget constraint of the household by solving forward for a_t and substituting repeatedly to get:

$$a_0 = R^{-1} a_1 - y_0 + c_0 = \ldots = \sum_{t=0}^{T-1} R^{-t}(c_t - y_t) + R^{-T} a_T$$

Since the household cannot die in debt T, we know that $a_T \geq 0$ and the intertemporal budget constraint becomes:

$$\sum_{t=0}^{T-1} R^{-t} c_t \leq a_0 + \sum_{t=0}^{T-1} R^{-t} y_t$$

$$\sum_{t=0}^{T-1} R^{-t} c_t \leq a_0 + \sum_{t=0}^{T-1} R^{-t} y_t \qquad (2)$$

Interpretation:

- the present value of consumption equals initial financial wealth (a_1) + present value of human wealth ($\sum_{t=0}^{T-1} R^{-t} y_t$).
- the term on the right hand side is the economically relevant measure of total wealth: financial + non-financial.
- the combination of time-additive preferences and an additive intertemporal budget constraint is what makes the problem so tractable (Ghez & Becker (1975))

2.4 Optimal Consumption-Saving under Certainty

Optimal Consumption-Saving under Certainty.

The problem of the household is to maximize (1) subject to (2):

$$\max_{\{c_t\}_{t=0}^{T-1}} \sum_{t=0}^{T-1} \beta^t u(c_t)$$

$$s.t.$$

$$\sum_{t=0}^{T-1} R^{-t} c_t \leq a_0 + \sum_{t=0}^{T-1} R^{-t} y_t$$

We can solve this problem by setting-up the Lagrangian (where $\lambda > 0$ is the Lagrange multiplier on the intertemporal constraint):

$$\mathcal{L} = \sum_{t=0}^{T-1} \beta^t u(c_t) + \lambda \left(a_0 + \sum_{t=0}^{T-1} R^{-t} y_t - \sum_{t=0}^{T-1} R^{-t} c_t \right)$$

The first order condition for c_t is:

$$u'(c_t) = (\beta R)^{-t} \lambda$$

[Note: from this you should be able to infer that the IBC will hold with equality). Can you see why?]

Interpretation:

$$u'(c_t) = (\beta R)^{-t} \lambda$$

- β captures impatience, i.e. the preference for the present. Makes us want to consume now.
- R determines the return on saving. A higher R makes us want to consume later (is that really the case? More later....)
- Marginal utility will be decreasing over time if $\beta R > 1$ and increasing otherwise.
- Since marginal utility decreases with consumption, this implies that consumption will be increasing over time when $\beta R > 1$ and decrease otherwise.
- When $\beta R = 1$ the two forces balance each other out and consumption becomes flat.
- Note that this gives us some key information on the slope of the consumption profile over time, but not on the consumption level.

2.5 A Special Case: $\beta R = 1$

$$u'(c_t) = \lambda$$

This implies that consumption is constant over time: $c_t = \bar{c}$. Substitute this into the intertemporal budget constraint to obtain:

$$\sum_{t=0}^{T-1} R^{-t}\bar{c} = \frac{1-R^{-T}}{1-R^{-1}}\bar{c} = \frac{1-\beta^T}{1-\beta}\bar{c} = a_0 + \sum_{t=0}^{T-1} R^{-t} y_t$$

from which we obtain:

$$\bar{c} = \frac{1-\beta}{1-\beta^T}\left(a_0 + \sum_{t=0}^{T-1} R^{-t} y_t\right)$$

Observe:

- consumption is a function of total wealth.
- the marginal propensity to consume is $(1-\beta)/(1-\beta^T)$ and converges to $1-\beta$ when the horizon extends ($T \to \infty$).
- if $\beta = 0.96$ (a reasonable estimate), this gives $R = 1/0.96 = 1.0416$. Then we should consume about 4% of total wealth every period.

2.6 The Permanent Income Hypothesis

Friedman's (1957) Permanent Income.

$$\bar{c} = \frac{1-\beta}{1-\beta^T}\left(a_0 + \sum_{t=0}^{T-1} R^{-t} y_t\right)$$

- This is Friedman's permanent income hypothesis. Individual consumption is not determined by income in that period, but by lifetime resources, unlike Keynesian consumption functions of the form $c_t = a + by_t$.
- Friedman actually defines permanent income as the right hand side of this equation. This is the annuity value of total resources.
- This implies that consumption should not respond much to transitory changes in income, since these will not affect much permanent income, but should respond if there are changes in your permanent income.
  - you earn an extra $200 today
  - you just got tenured and learn that starting next year, your income will double.
  - you learn that you won $10m at the state lottery

2.7 Understanding Estimated Consumption Functions

Keynes (1936) argues that 'aggregate consumption mainly depends on the amount of aggregate income,' 'is a stable function,' and 'increases less than proportionately with income.'

In other words, Keynes argues for a consumption function of the type $c_t = a + by_t$.

Empirically, it matters whether we look (a) in the cross section or (b) in the time series.
This looks quite different from Friedman's permanent income which we can write as $c_t = y_t^p$ where y_t^P is simply permanent income.

Yet, Friedman's PIH can account for the empirical observations. Observe that we can write $y_t = y_t^P + y_t^T$ where y_t^T is the transitory component of income. An OLS regression of consumption on income yields:

$$\begin{aligned}\hat{b} &= \frac{\text{cov}(c_t, y_t)}{\text{var}(y_t)} = \frac{\text{cov}(y_t^P, y_t)}{\text{var}(y_t)} = \frac{\text{var}(y_t^P)}{\text{var}(y_t)} < 1 \\ \hat{a} &= E(c_t) - \hat{b}E(y_t) = (1-\hat{b})E(y_t^P)\end{aligned}$$

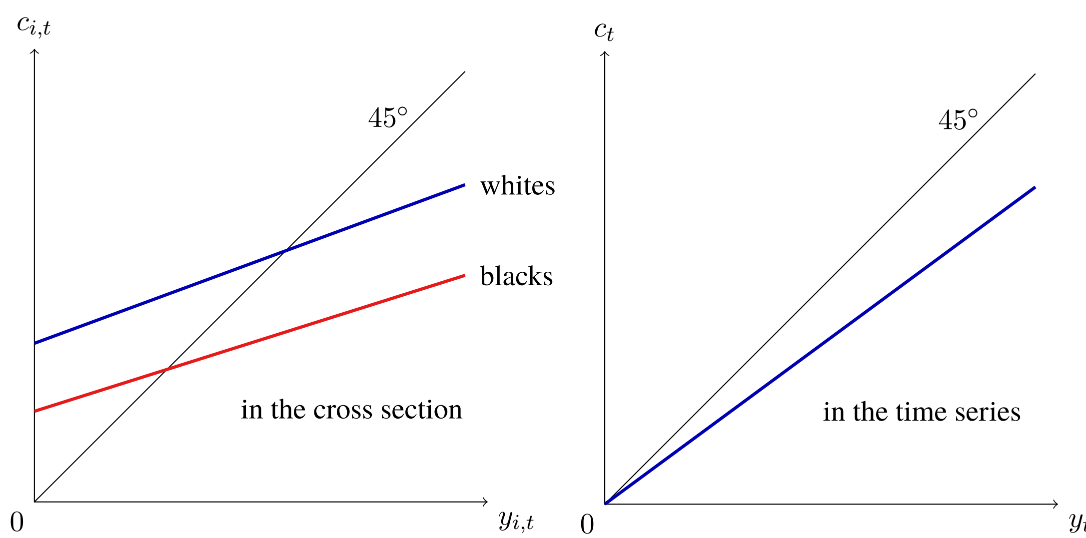

Figure 1: Keynesian consumption functions and the PIH

- In the cross section: more variations from $y_{i,t}^T$: lower $\hat{b}$.
- in the time series, more variation from y_t^P: higher $\hat{b}$
- in the cross section: lower intercept $\hat{a}$ if lower $E(y_t^P)$

2.8 The LifeCycle Model under certainty

Modigliani and Brumberg (1954) consider the lifecycle implications of the previous model. Suppose that people live T periods (from 1 to T) and that $\beta = R = 1$. (Note: this is a stronger assumption than $\beta R = 1$)

The PIH model tells us that consumption is constant and equal to the permanent income of the agent: $c = \bar{c} = y^P$.

This is irrespective of the income profile $\{y_t\}$ over time. Suppose now that the agent works for $N < T$ periods, earning income y, then retires.

The household saves $y - \bar{c}$ when working, then dissaves $-\bar{c}$ when retired.

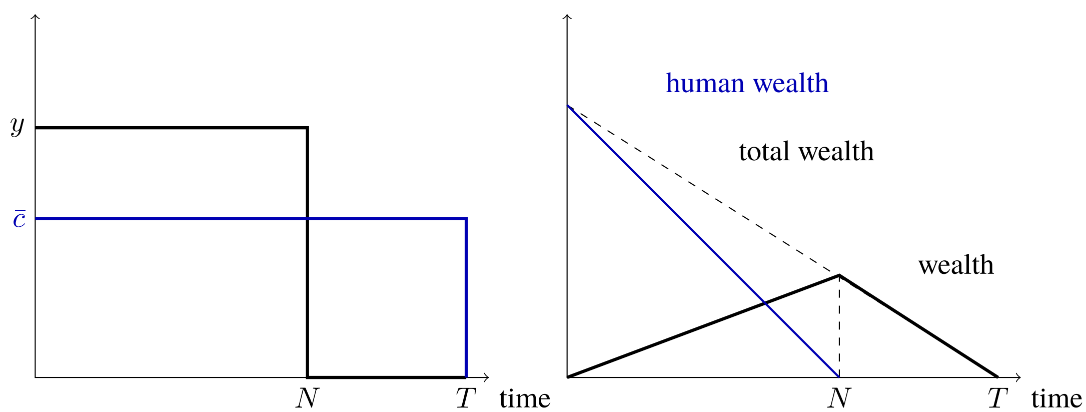

Figure 2: The Baseline Life-Cycle Model

Equations for the simple Modigliani-Brumberg (1954) lifecycle model:

$$
\begin{aligned}
\bar{c} &= \frac{N}{T} y \\
a_t &= t \frac{T-N}{T} y \quad \text{for } t \leq N \\
a_t &= \frac{N}{T}(T-t) y \quad \text{for } N \leq t \leq T \\
h_t &= (N-t) y \quad \text{human wealth} \\
w_t &= a_t + h_t = \frac{N}{T}(T-t) y \quad \text{for all } t \leq T
\end{aligned}
$$

Assume there is no growth. Then, we have the following (aggregating across cohorts):

$$\bar{w} = \frac{N}{2T}y(T+1)$$
$$\bar{h} = \frac{N}{2T}y(N+1) \quad ; \quad \bar{a} = \frac{N}{2T}y(T-N)$$

Note:

- The household runs total wealth to 0
- Human wealth runs out at $t = N$. It is supplemented by financial wealth
- We can have positive financial wealth even if there is no bequest motive.
- The ratio of human to financial wealth $\bar{h}/\bar{a}$ does not depend on income y (it is equal to $(N+1)/(T-N)$.
- The details of the social security system matter. This describes a fully funded system (or even more precisely, what should happen if there is NO social security system and no bequest motive). What if we have a society where the young take care of the old (China); or an unfunded system where the government taxes the young to support the old? What happens to
  - consumption profiles?
  - income?
  - private saving?

2.9 Saving and Growth in the LifeCycle Model

How does growth affect saving in the lifecycle model?

- Start with zero growth: the age-profile = cross-section. Aggregate wealth is constant and aggregate saving equals 0. The young save, and the old dissave
- population growth: more young saving, so saving increases with population growth.
- productivity growth: more complex and depends on how productivity growth affects each cohort's income:
  - If productivity growth is across cohorts (i.e. each cohort's income is constant but younger cohorts have a higher income profile) then productivity growth increases saving. (Why? b/c the young save more than the old dissave)
  - but instead if productivity growth increases income over a worker's lifetime, then young workers may decide to borrow against higher future income in middle age. In that case, faster growth can reduce savings.

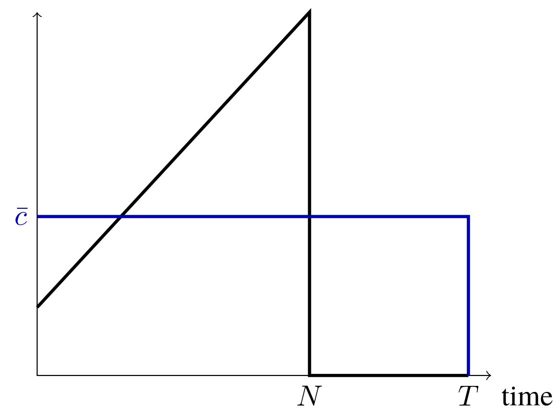

Figure 3: The Life-Cycle Model when Income grows over time

2.10 Interest Rate Elasticity of Saving

The response of consumption and savings to changes in interest rates is an important question. Think about:

- the transmission of monetary policy (changing the real interest rate)
- changes to the tax code that affects rates of returns on savings. you have seen in the first part of this course how changes in savings can affect growth rates temporarily (if growth is exogenous) and potentially permanently (if growth is endogenous)

Consider the first-order condition again:

$$u'(c_t) = (\beta R)^{-t} \lambda$$

Rewrite it in two consecutive periods and eliminate λ. This is the Euler equation under certainty:

$$u'(c_t) = (\beta R)\, u'(c_{t+1})$$

Consider CRRA preferences: $u(c) = c^{1-\theta}/(1-\theta)$. [We have already seen these preferences when solving the Ramsey-Cass-Koopmans problem: θ represents both the CRRA coefficient and the inverse of the (IES).]

Substitute to get:

$$c_{t+1}/c_t = (\beta R)^{1/\theta}$$

- if $1/\theta = 0$ (Leontieff) then c_t is flat regardless of the interest rate. No substitutability
- for $1/\theta < 1$: weak substitution effects
- for $\theta = 1$ income and substitution effects cancel out (log preferences)

- for $1/\theta > 1$: strong substitution effects
- if $1/\theta \to \infty$ then c_t becomes very responsive to the interest rate. In the limit, consumption growth becomes so responsive that the interest rate R will have to stay 'close' to $1/\beta$ to ensure that consumption growth does not become too extreme.

$$c_{t+1}/c_t = (\beta R)^{1/\theta}$$

In general, consumption growth should be responsive to changes in the interest rate. On can rewrite:

$$d\ln(c_{t+1}/c_t) = (1/\theta)d\ln R = (1/\theta)d\ln(1+r) \approx (1/\theta)dr$$

An increase in the real interest rate by 100bp should increase consumption growth by $1/\theta\%$.

But the analysis is incomplete: we need to figure out by how much consumption itself changes.

What about the overall effect?

To simplify things, let's consider first a two-period problem (with $a_0 = 0$)

$$\begin{aligned} c_1/c_0 &= (\beta R)^{1/\theta} \\ c_0 + \frac{c_1}{R} &= y_0 + \frac{y_1}{R} \end{aligned}$$

Let's start with a simple case where $y_1 = 0$:

We have:

- a substitution effect. Keeping the utility level constant, the change in interest rate leads us to substitute consumption today for consumption tomorrow: c_0 falls, c_1 increases.
- an income effect: the budget line rotates around $(y_0, 0)$. This means more consumption can be afforded in each period. This increases c_0 and c_1.

The effect on c_1 is unambiguous. The effect on c_0 is ambiguous. When $1/\theta > 1$ the substitution effect dominates so that c_0 falls.

Mathematically:

$$c_0(1 + (\beta R)^{1/\theta}/R) = y_0$$

2.10.1 The 2-period case with $y_1 = 0$

Now, let's consider what happens when $y_1 \neq 0$

- the budget curve rotates around (y_0, y_1).
- in addition to the income and substitution effects, there is a wealth effect: future income is worth less to the household. This reduces c_0 and c_1.

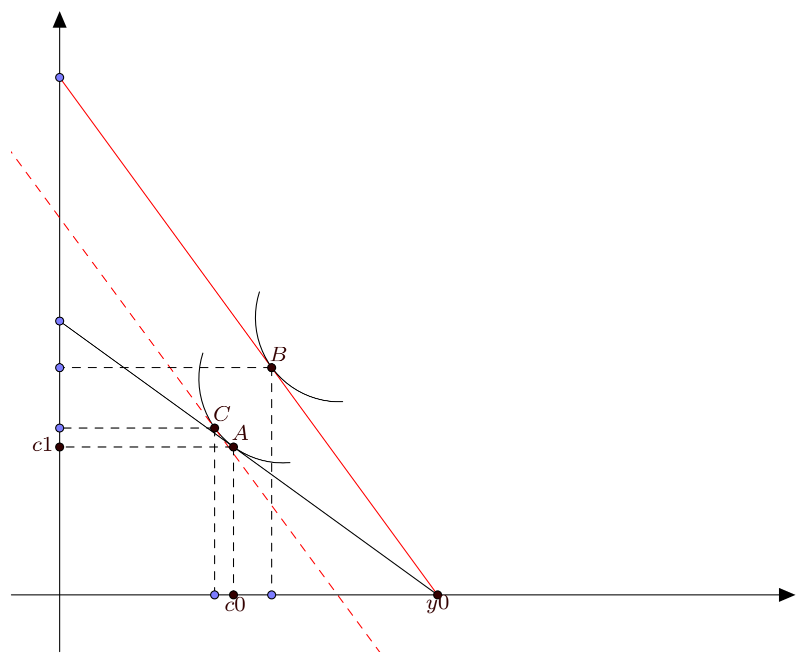

Figure 4: An increase in interest rates when $y_1 = 0$

- the net effect often depends on whether the household is a net borrower or lender.
  - if $c_0 = y_0$ and $c_1 = y_1$ then there is no income and wealth effect (why? because the initial consumption bundle remains on the new budget line). So there is only a substitution effect and c_0 falls.
  - if the household is initially a saver (i.e. $c_0 < y_0$) then the income effect dominates the wealth effect and the overall effect on c_0 is indeterminate
  - if the household is initially a borrower ($c_0 > y_0$) then the income effect is weaker than the wealth effect. Therefore c_0 falls.

Mathematically:

$$c_0(1 + \beta^{1/\theta} R^{1/\theta - 1}) = (y_0 + y_1/R)$$

2.10.2 The 2-period case with $y_1 \neq 0$ and $c_0 = y_0$

The substitution effect is the only effect. c_0 falls.

2.10.3 The 2-period case when $c_0 < y_0$ (lender)

Ambiguous. Income effect stronger than wealth effect.

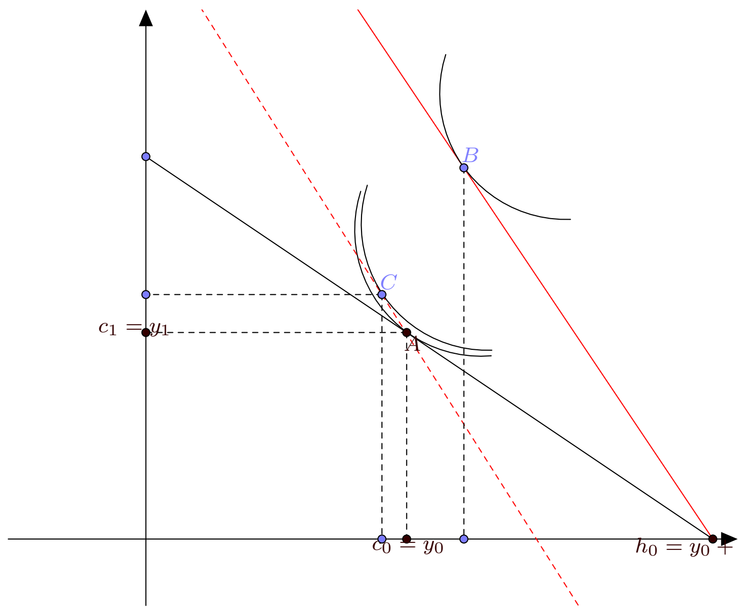

Figure 5: An increase in interest rates when $c_0 = y_0$ and $c_1 = y_1$

2.10.4 2-period case when $c_0 > y_0$ (borrower)

wealth effect stronger than income effect. c_0 falls.

Savings and Interest Rates: Recap:

The literature often considers the case $\theta = 1$ as a benchmark, where income and substitution effects cancel out, leaving the saving rate independent of the interest rate;

Furthermore, empirical estimates of the elasticity of intertemporal substitution suggest relatively low numbers for $1/\theta$, especially since there are also income and wealth effects.

It is tempting to conclude from this that (a) the slope of consumption growth and (b) the level of consumption will be largely unaffected by changes in the interest rate. However, this answer can be misleading for a number of reasons:

- lifetime horizon. But even if the IES is small, it can have a large impact over a lifetime (Summers 1981).
- This omits the wealth effect. Even if income and substitution effects cancel out, a change in interest rates affects human wealth, and this leads to a change in consumption in the PIH-LC model.
- Finally, the nature of the change in interest rates matters. For instance, a change in interest rates due to tax changes may be offset somewhere else to leave government

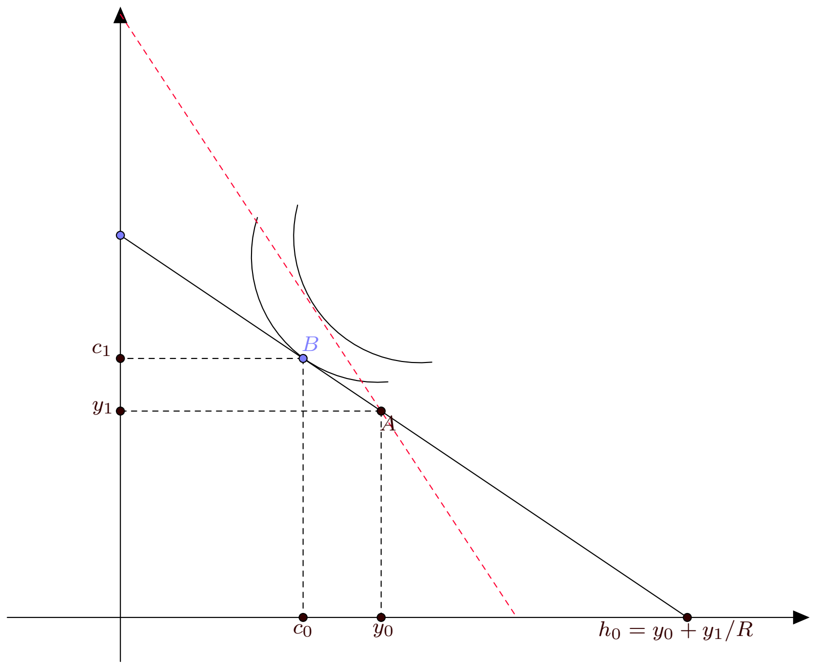

Figure 6: An increase in interest rates for a saver ($c_0 < y_0$)

revenues unchanged. In that case, there is no income effect and only the substitution effect. This might not be very helpful if the IES is small anyway.

2.11 The LifeCycle Model under Certainty Again

Consider now the case where R and β differ from 1. In addition, suppose that $a_0 = 0$ and that y is constant as before.

The Euler equation with CRRA preferences implies:

$$c_t = (\beta R)^{t/\theta} c_0$$

Substituting into the budget constraint, we obtain:

$$\begin{aligned} c_0 &= \frac{1-(\beta R^{1-\theta})^{1/\theta}}{1-(\beta R^{1-\theta})^{T/\theta}} \sum_{t=0}^{T-1} R^{-t} y_t \\ c_0 &= \frac{1-(\beta R^{1-\theta})^{1/\theta}}{1-(\beta R^{1-\theta})^{T/\theta}} \frac{1-R^{-N}}{1-R^{-1}} y_0 \end{aligned}$$

Suppose that $\beta R > 1$ so that consumption grows over time, even it θ is low. If the horizon T is long enough relative to the working period, consumption must be much higher at the end

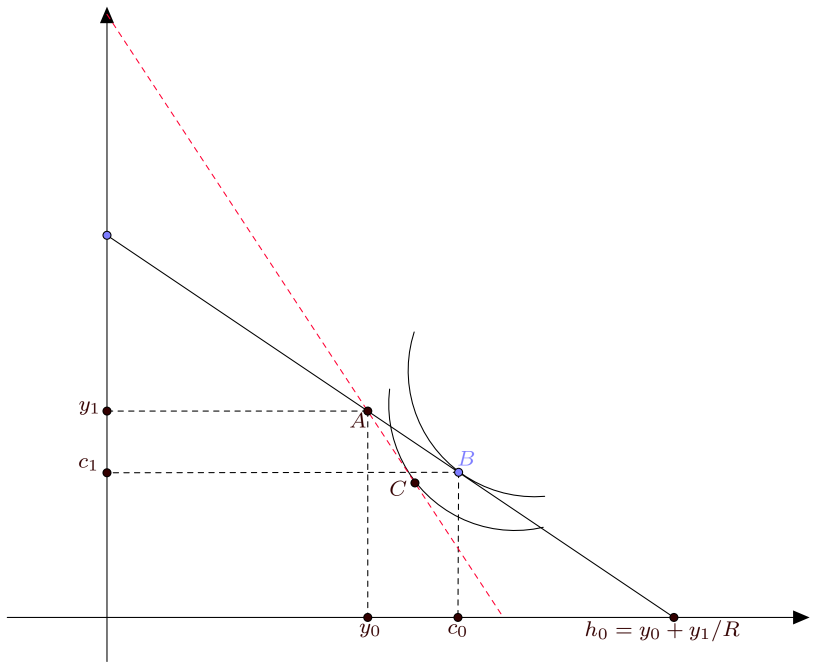

Figure 7: An increase in interest rates for a borrower ($c_0 > y_0$)

of life than at the beginning: the agent must accumulate a large stock of wealth. Aggregate wealth and saving may be highly responsive to changes in interest rates. See Summers 1981.

3 Consumption under Uncertainty

Last class we looked at the consumption model under certainty. The model provides important insights:

- consumption is a function of total wealth (permanent income)
- the slope of the consumption profile is controlled by the discount rate, the interest rate and the intertemporal elasticity of substitution
- in a lifecycle environment, there is a substantial amount of life-cycle wealth accumulation. In the simple model, the amount of wealth is $\bar{a}/y = N/(2T)(T-N) = 40/(120)(20) = 800/120 = 6.66$
- the elasticity of aggregate saving to the interest rate is complex.

The model needs to be extended to allow for uncertainty. Precautionary saving is another reason why households decide to save. We start with the canonical model, augmented for uncertainty.

TABLE 1—THE INTEREST ELASTICITY OF AGGREGATE SAVINGS

	Value of r		
	.04	.06	.08
$\gamma=.5$			
η_r	3.71	2.26	2.44
S/WL	.121	.274	.451
$\gamma=0$			
η_r	3.36	1.89	1.87
S/WL	.068	.142	.21[illegible]
$\gamma=-.5$			
η_r	3.09	1.71	1.54
S/WL	.049	.096	.13[illegible]
$\gamma=-1$			
η_r	2.87	1.59	1.37
S/WL	.038	.073	.09[illegible]
$\gamma=-2$			
η_r	2.38	1.45	.12[illegible]
S/WL	.028	.048	.06[illegible]
$\gamma=-5$			
η_r	.741	1.09	1.18
S/WL	.014	.019	.02[illegible]

Note: The calculation assumes $n=.015$, $g=.02$, T' $T=40$, and $\delta=.03$. The savings rate is measured fraction of labor income.

Figure 8: Table 1 from Summers (1981)

3.1 The Canonical Model

3.1.1 the set-up

The household has the following preferences over consumption sequences:

$$U = E\left[\sum_{t=0}^{\infty} \beta^t u(c_t)|\Omega_0\right] \tag{3}$$

Notice two differences with the model under certainty:

1. First I assume that the horizon is infinite. This is mostly to show you how to use solve the model in that case. Nothing substantial rests on that hypothesis and I will point out as we go where things might be different if we have a finite horizon. Formally, you may think that households care about their offsprings and apply the same discount rate.

2. The term $E[.|\Omega_0]$ captures expectations conditional on information available at time $t = 0$. This information set is denoted Ω_0. This is also an important assumption. It implies that preferences are separable over states and over time.. To see this, suppose that there are S_t possible states of the world in period t and that each of them has probability (as of time 0) given by $\pi_{s,t}$. Then we can write the utility as:

$$U = \sum_{t=0}^{\infty} \sum_{s \in S_t} \beta^t \pi_{s,t} u(c_{s,t})$$

This double separation imposes strong structure on preferences, but it simplifies tremendously the analysis.

3. To lighten the notation, I will indifferently write $E_t[.]$ or $E[.|\Omega_t]$ to indicate conditional expectation as of time t.

The household budget constraint takes the same form as before, except that now, I will suppose that households face some uncertain interest rate $\tilde{R}_{t+1}$ and an uncertain future income $\tilde{y}_t$. In this notation, the $\tilde{}$ indicates that a variable is stochastic (as seen from previous periods). The budget constraint then takes the form:

$$a_{t+1} = \tilde{R}_{t+1}(a_t + \tilde{y}_t - c_t) \tag{4}$$

3.1.2 Recursive Representation

The problem is to maximize (3) subject to (4), and any other restriction on consumption and asset choices, for a given initial level of wealth a_0. For instance, we know that we only consider positive consumption: $c_t \geq 0$.

We have also already discussed the fact that the household will not be allowed to run Ponzi-like schemes:

$$\lim_{T \to \infty} \beta^T a_T \geq 0$$

This constraint holds in the uncertain case, along all possible consumption sequences (technically, it holds almost surely).[1] But there might be other constraints on assets holdings. For instance, the household may be prevented from borrowing beyond a certain limit:

$$a_t \geq \underline{a}$$

At time t, a_t is a state variable of the household consumption problem, in the sense that it is pre-determined by the previous actions of the households and is beyond its control.

[1] It is not sufficient that the No-Ponzi condition holds in expectation, that is $E_0[\lim_{T\to\infty} \beta^T a_T] \geq 0$. If this were the case, then there would be possible paths with non-zero measure where the No-Ponzi condition would be violated. Along these paths, lenders would have to agree to provide an infinite level of consumption to the household. Note also that if the NPC holds a.s., then it holds in expectation, while the reverse is obviously not true.

We are going to assume that income and return realizations are iid, so that $\tilde{y}_t$ and $\tilde{R}_t$ are not state variables of the household problem. This is mostly to keep notations simple. It would be quite straightforward to extend the set-up to a case where $\tilde{y}$ and $\tilde{R}$ have a Markov structure.

Remark 1 *In some situations, it is easier to use cash-on-hand x_t as the state variable, where x_t is defined as : $x_t = a_t + \tilde{y}_t$. x_t represents the resources available for consumption and saving to the household, after the realization of current income. The budget constraint becomes:*

$$x_{t+1} = \tilde{R}_{t+1}(x_t - c_t) + \tilde{y}_{t+1}$$

Since financial assets a_t are the sole state variable, we can write the value function that maximizes the utility of the agent as a function of the state variable a:

$$v(a_0) = max_{\{c_t\}_{t=0}^{\infty}} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right]$$

Given the nicely recursive structure of the problem, we write the Bellman equation as follows. Suppose that the level of assets is a in a given period. Consumption that period must satisfy:

$$\begin{aligned} v(a_t) &= max_{c_t \in \mathcal{C}_t} \quad u(c_t) + \beta E_t[v(a_{t+1})] \\ \text{s.t.} & \\ a_{t+1} &= \tilde{R}_{t+1}(a_{t+1} + \tilde{y}_t - c_t) \end{aligned}$$

where $\mathcal{C}_t$ denotes the set of permissible consumption choices at time t. Notice that it is the same value function that enters on both sides of this equation. So, one way to think about the household problem is that the Bellman equation defines the value function as fixed point of a functional equation. There are various theorems that establish existence and uniqueness of this fixed point, when the Bellman equation is well-behaved -as is the case here.

Remark 2 *If the functional equation is contraction mapping, then the Bellman equation has a unique solution AND this solution can be found by iterating on the value function. This provides a convenient (if not especially rapid) way to characterize numerically the value function (value function iteration).*

3.1.3 Optimal Consumption and Euler Equation

We start by assuming that the solution is interior to the set $\mathcal{C}_t$. The first-order condition of the above problem yields:

$$u'(c_t) = \beta E_t[v'(a_{t+1})\tilde{R}_{t+1}]$$

Let's now consider what happens when there is a small change in a_t on the household value function $v(a_t)$. To calculate $v'(a_t)$, let's take a full derivative of the Bellman equation.[2] The total variation is:

$$v'(a_t)da = u'(c_t)dc_t + \beta E_t[v'(a_{t+1})\tilde{R}_{t+1}(da - dc)]$$

where dc denotes the change in optimal consumption for a given small change in a. Regrouping terms, we obtain:

$$v'(a_t)da = (u'(c_t) - \beta E[v'(a_{t+1})\tilde{R}_{t+1}])dc + \beta E[v'(a_{t+1})\tilde{R}_{t+1}]da$$

The first term on the right hand side is zero from the first-order condition of the problem. So we are left with:

$$v'(a_t) = \beta E_t[v'(a_{t+1})\tilde{R}_{t+1}]$$

This is a straightforward application of the Envelope Theorem.

Combining the first order condition and the Envelope theorem, we conclude that:

$$u'(c_t) = v'(a_t)$$

Substituting back into the first order condition, we obtain the well-known Euler Equation under uncertainty:

$$u'(c_t) = \beta E_t[\tilde{R}_{t+1}u'(c_{t+1})] \tag{5}$$

What is the intuition for the Euler equation? A variational argument might help. Suppose that we reduce consumption from the optimal path in period t by ϵ, and increase consumption by $\tilde{R}_{t+1}\epsilon$ next period (so that we are back on the optimal consumption path after period $t+1$). The marginal disutility (as of time t) of reducing consumption in t is $u'(c_t)\epsilon$. The marginal increase utility from higher consumption in $t+1$ (as of time t) is $\beta E_t[\tilde{R}_{t+1}u'(c_{t+1})$. For a small ϵ the two should be equal (otherwise the proposed consumption is not optimal to start with).

Note that the discount rate β and the interest rate $\tilde{R}_{t+1}$ still play opposing force on consumption growth, so the insights from the certainty case do carry over to the uncertain case. But we now also have to take into account uncertainty over future returns and future marginal utility.

Remark 3 *The derivation above assume that consumption is interior. What would happen if consumption is at the boundary. For instance, suppose that we impose the conditions that $0 \leq c \leq a + y$ (how should we interpret this condition?). What form does the Euler equation take?*

[2] We are assuming that the value function is differentiable, which is not always the case. See Stokey, Lucas and Prescott (1983) for more details on this.

3.2 The Certainty Equivalent (CEQ)

The Euler equation provides some important insights into consumption behavior, but in its general form, it is not very tractable. We now make a number of simplifying assumptions, following Hall (1978).

First, we assume that there is no uncertainty in interest rates, so $\tilde{R}_t = R$. Moreover, we assume that there is no tilt in consumption profiles, that is $\beta R = 1$.

Second, we will consider a very particular form of preferences:

$$u(c) = \alpha c - \gamma c^2/2; \quad \gamma > 0, \alpha > 0$$

In other words, preferences are quadratic over consumption. These preferences are very weird from a number of points of view:

- even if $\alpha > 0$, utility turns negative for sufficiently large consumption
- these preferences admit negative consumption (they definitely violate Inada's conditions)

So why would we want to make these crazy assumptions? Two possible justifications are:

- we could think of these preferences as a second order approximation of utility for more general utility functions. If we think about it this way, then it would suggest that this may not be such a bad approximation for relatively small changes in consumption over time.
- these preferences have the important property that marginal utility is linear, or equivalently that the second derivative is constant: $u'(c) = \alpha - \gamma c$ and $u''(c) = -\gamma$.

Let's make these two assumptions and substitute into the Euler equation (5) to obtain:

$$c_t = E_t[c_{t+1}] \tag{6}$$

The stunning result here is that consumption follows a Random Walk. This means that changes in consumption are unpredictable. To see how stunning it is, recall that if we had no uncertainty (and $\beta R = 1$), then we would get

$$u'(c_t) = u'(c_{t+1)}$$

and so consumption would be constant over time, and therefore entirely predictable. Instead, once we introduce uncertainty, consumption becomes entirely unpredictable!

To see what is going on, it helps to solve for the level of consumption in the CEQ case. To do this, let's first derive the Intertemporal Budget Constraint. First recall that the dynamic budget constraint is:

$$a_{t+1} = R(a_t + \tilde{y}_t - c_t)$$

Let's solve this sequence forward *for a given sequence of consumption and income realization*:

$$a_0 = R^-1(a_1 + c_0 - \tilde{y}_0 = ... = \sum_0^\infty R^{-t}(c_t - \tilde{y}_t) + \lim_{T\to\infty} R^{-T} a_T$$

With the No-Ponzi condition, the last term has to be positive, so the intertemporal budget constraint takes the form:

$$\sum_0^\infty R^{-t}c_t \leq a_0 + \sum_0^\infty R^{-t}\tilde{y}_t$$

Notice that this intertemporal budget constraint does not have an expectation term: it has to hold along any possible realization of income and consumption: it holds almost surely. But if it holds almost surely, then we are allowed to take expectations and the following also holds:

$$E_0[\sum_0^\infty R^{-t}c_t] \leq a_0 + E_0[\sum_0^\infty R^{-t}\tilde{y}_t]$$

The next step is to observe that we can move the expectation inside the summation, and use the fact that under the random walk hypothesis, the following holds:

$$E_0 c_t = E_0 E_{t-1} c_t = E_0 c_{t-1} = ... = c_0$$

where the second term follows from the Law of Iterated Expectations, the third one from the fact that consumption follows a random walk at t and the last one from iterating the argument.

It follows that consumption at time $t = 0$ must satisfy:

$$\begin{aligned} c_0 \frac{1}{1 - R^{-1}} &= a_0 + E_0[\sum_0^\infty R^{-t}\tilde{y}_t] \\ c_0 &= (1-\beta)\left(a_0 + E_0\left[\sum_0^\infty R^{-t}\tilde{y}_t\right]\right) \end{aligned}$$

where we used the assumption that $R^{-1} = \beta$ and the definition of x_t in the second line. What this tells us is that consumption follows the PIH in expectation. The term in parenthesis on the right hand side is expected total wealth, where the expectation is over future labor income.

This is why the model is called the 'certainty equivalent' model: as far as consumption decisions are concerned, the household behaves as if future income was certain and equal to its expected value. The source of this behavior can be traced back to the assumption of quadratic utility. Note that the Euler equation in the CEQ model is:

$$c_t = E_t c_{t+1}$$

what this tells us is that the household is smoothing consumption, but taking future consumption as if it were certain and equal to its expected value. But if you retrace your steps, you will see that this result arises from the Euler equation in general form:

$$u'(c_t) = E_t u'(c_{t+1})$$

and the fact that marginal utility is linear when utility is quadratic: $u'(c) = \alpha - \gamma c$. Anticipating on the next lecture, this tells you that this result will not hold in the more general case where marginal utility is not linear.

Why are changes in consumption unpredictable, while the consumption level itself seems to follow a minor modification of the PIH? To see what is going on, consider consumption in two consecutive periods, t and $t+1$:

$$\begin{aligned} c_t &= (1-\beta)\left(a_t + E_t\left[\sum_{s=t}^{\infty} R^{-(s-t)}\tilde{y}_s\right]\right) \\ c_{t+1} &= (1-\beta)\left(a_{t+1} + E_{t+1}\left[\sum_{s=t+1}^{\infty} R^{-(s-(t+1))}\tilde{y}_s\right]\right) \end{aligned}$$

Take the difference and substitute $a_{t+1} = R(a_t + \tilde{y}_t - c_t)$ to obtain:

$$\begin{aligned} c_{t+1} - c_t &= (1-\beta)\left(a_{t+1} - a_t + \sum_{s=t+1}^{\infty} R^{-(s-(t+1))}E_{t+1}\tilde{y}_s - \sum_{s=t}^{\infty} R^{-(s-(t))}E_t\tilde{y}_s\right) \\ &= (1-\beta)\left(R(a_t + \tilde{y}_t - c_t) - a_t + \sum_{s=t+1}^{\infty} R^{-(s-(t+1))}E_{t+1}\tilde{y}_s - \sum_{s=t}^{\infty} R^{-(s-(t))}E_t\tilde{y}_s\right) \\ &= (1-\beta)\left((R-1)a_t + R\tilde{y}_t - Rc_t + R\sum_{s=t+1}^{\infty} R^{-(s-t)}E_{t+1}\tilde{y}_s - \sum_{s=t}^{\infty} R^{-(s-(t))}E_t\tilde{y}_s\right) \\ &= (1-\beta)\left(-R\sum_{s=t+1}^{\infty} R^{-(s-t)}E_t\tilde{y}_s + R\sum_{s=t+1}^{\infty} R^{-(s-t)}E_{t+1}\tilde{y}_s\right) \end{aligned}$$

using the expression for c_t and the fact that $\beta R = 1$, we obtain finally:

$$c_{t+1} - c_t = (R-1)\sum_{s=t+1}^{\infty} R^{-(s-t)}(E_{t+1}\tilde{y}_s - E_t\tilde{y}_s) \tag{7}$$

Notice that the term in the summation on the right hand side of (7) is $E_{t+1}\tilde{y}_s - E_t\tilde{y}_s$, that is, the revision in expectations about future income. Of course, this revision is unpredictable as of period t, otherwise it would already have been incorporated in the current expectation $E_t\tilde{y}_s$!

This gives us a very nice result: the change in consumption is related to the news the household receives about future income. We will see also that it provides us with a way to test the certainty equivalent model.

Remark 4 *You can check that if you take expectations as of time t on both side of this equation, you recover $E_t c_{t+1} - c_t = 0$.*

Example 1 *Consider the case where income follows an AR(1) process:*

$$\tilde{y}_{t+1} = \rho \tilde{y}_t + \eta_{t+1}; \quad 0 < \rho \leq 1$$

Then we can easily check that $E_t \tilde{y}_s = \rho^{s-t} \tilde{y}_t$. Substituting back into (7), we obtain after some easy manipulations:

$$c_{t+1} - c_t = \frac{1-\beta}{1-\rho\beta} \eta_{t+1}$$

Since $\rho \leq 1$, $c_{t+1} - c_t \leq \eta_{t+1}$, that is, consumption in general responds less than 1 for 1 to a change in income. The case where consumption moves 1 for 1 is when $\rho = 1$, i.e. income itself is a random walk.

3.3 Tests of the Certainty Equivalent Model

3.3.1 Testing the Euler Equation

The literature up until Hall (1978) used to attempt to derive closed form solution for consumption (i.e. a consumption function) and estimate it. But a closed form solution for the consumption function is often not available. So instead, the literature would try to identify the determinants of consumption and estimate empirically the relationship between consumption and its determinants. This would not allow for a rigorous test of the theory. In addition, the regression typically faced a serious problem of identification since income (the most common right hand side variable) is not exogenous.

Instead, Hall argued that we can test the theory by directly testing the first-order condition of the model, i.e. the Euler equation. Under rational expectation, any variable can be expressed as the sum of its conditional expectation and an innovation term, orthogonal to any information available at time t:

$$c_{t+1} = E_t c_{t+1} + \epsilon_{t+1} = c_t + \epsilon_{t+1}$$

where $E_t \epsilon_{t+1} = 0$ and the second equality uses (6). So the theory implies that c_t contains all the relevant information necessary to predict c_{t+1}.

Under the null hypothesis that the theory is correct, a regression of the form:

$$c_{t+1} = a + b \ c_t + c \ x_t + \varepsilon_{t+1} \tag{8}$$

where x_t is any variable available at time t to the household should yield:

$$\hat{a} = 0; \quad \hat{b} = 1; \quad \hat{c} = 0$$

What is important in that regression is that it does not matter whether y_t is exogenous or not (the key problem with the consumption function estimation approach). The key test is whether $\hat{c} = 0$ or not. If we find some variables, known as of time t that can help predict next period's consumption after controlling for current consumption, then the theory has to be incorrect.

Remark 5 *Notice that the theory does not say that consumption should not react to current income. In other words, if we run the regression*

$$c_{t+1} = a + b \ \ c_t + c \ \ x_t + d \ \ y_{t+1} + \varepsilon_{t+1}$$

there is no presumption that $\hat{d}$ should be equal to 0.

Remark 6 *Notice that given (7), we know that*

$$\epsilon_{t+1} = (R-1) \sum_{s=t+1}^{\infty} R^{-(s-t)} (E_{t+1}\tilde{y}_s - E_t\tilde{y}_s)$$

One could think that this would provide another way to test the theory. For instance, when income follows an AR(1) process as above, we know that the consumption innovation is given by:

$$\epsilon_{t+1} = \frac{1-\beta}{1-\rho\beta}\eta_{t+1}$$

so the innovation to consumption ϵ_{t+1} and the innovation to income η_{t+1} are linked in a very precise way. However, this could be exploited only if the household learns about the change in its income as it happens. If instead, the household learns about a change in its income before it is realized, this is when consumption will change, and not when the actual change in income occurs. Unless the econometrician has information on when the information becomes available to the household (more on this below), then the relationship above will not be terribly useful. Testing the first order condition remains valid, however, since any information known at time t to the household should not help predict future consumption.

Hall (1978) tests the CEQ model using aggregate quarterly data on non-durable real consumption per capita and real disposable income per capita. The results (see attached table) suggest that indeed lagged income is not helpful in predicting future consumption (on top of lagged consumption).

Equations Relating Consumption to Lagged Consumption and Past Levels of Real Disposable Income

	Equation No. and Equation	R^2	s	D-W	F	F^*
3.1	$c_t = -16 + 1.024\, c_{t-1} - .010\, y_{t-1}$ (11) (.044) (.032)	.9988	14.7	1.71	.1	3.9
3.2	$c_t = -23 + 1.076\, c_{t-1} + .049\, y_{t-1} - .051\, y_{t-2}$ (11) (.047) (.043) (.052)					
	$-.023\, y_{t-3} - .024\, y_{t-4}$ (.051) (.037)	.9989	14.4	2.02	2.0	2.4
3.3	$c_t = -25 + 1.113\, c_{t\ \ 1} + \sum_{i=1}^{12} \beta_i y_{t-1}$ $\Sigma\beta_i = .077$ (11) (.054) (.040)	.9988	14.6	1.92	2.0	2.7

Figure 9: Table 3 in Hall (1978).

3.3.2 Allowing for time-variation in interest rate: the log-linearized Euler equation

The regression (8) imposes that the gross real interest is constant and equal to the inverse of the discount factor. It also imposes that preferences are quadratic. We can relax both assumptions yet obtain a result very similar to the CEQ, as long as we are looking at small deviations around the equilibrium.[3]

To see how this is done, consider the Euler equation of the general model:

$$u'(c_t) = \beta E_t[\tilde{R}_{t+1} u'(c_{t+1})]$$

Assume that $\tilde{R}_{t+1}$ is known as of time t. This would be the case if $\tilde{R}_{t+1}$ is the return on a one-period risk free bond between t and $t+1$. Assume further that preferences are CRRA so that $u'(c) = c^{-\theta}$, with $\theta > 0$.

The Euler equation takes the form:

$$c_t^{-\theta} = \beta R_{t+1} E_t[c_{t+1}^{-\theta}]$$

We can rewrite this as follows:

$$\begin{aligned} 1 &= \beta R_{t+1} E_t[c_{t+1}^{-\theta} c_t^{\theta}] \\ 1 &= \exp(-\rho + r_{t+1}) E_t[\exp(-\theta \Delta \ln c_{t+1})] \\ 0 &= -\rho + r_{t+1} + \ln E_t[\exp(-\theta \ln(c_{t+1}/c_t)] \end{aligned}$$

where we define $\rho = -\ln \beta$ and $r_t = \ln R_{t+1}$ and where the third line takes logs. Assume now that $\Delta \ln c_{t+1}$ is conditionally normally distributed. Then, the Euler equation takes the form:

$$0 = -\rho + r_{t+1} - \theta E_t \Delta \ln c_{t+1} + \frac{1}{2}\theta^2 V_t \Delta \ln c_{t+1}$$

[3]Recall that we motivated the CEQ as a second order approximation of preferences around the equilibrium. Instead of taking a second order approximation of preferences then solving for optimal smoothing, we can take a first order condition of the first order condition of the general consumption-saving problem.

where $V_t \Delta \ln c_{t+1}$ is the conditional variance of consumption growth.[4] If consumption growth is not conditionally normally distributed, this expression is a second-order approximation.

Re-arranging, we obtain:

$$E_t \Delta \ln c_{t+1} \approx \frac{1}{\theta}(r_{t+1} - \rho) + \frac{1}{2}\theta V_t \Delta \ln c_{t+1} \tag{9}$$

If we ignore the conditional variance term, and assume that the interest rate is equal to the discount rate ($r_t = \rho$), then we obtain an expression similar to the CEQ:

$$E_t \ln c_{t+1} \approx \ln c_t.$$

log-consumption follows a random-walk.[5]

If the interest rate is not constant, but we still ignore the variance term, we recover that expected consumption growth depends on the difference between the interest rate and the growth rate, scaled by the IES $1/\theta$: $E_t \Delta \ln c_{t+1} = 1/\theta(r_{t+1} - \rho)$.

As we will see a bit later, the variance term captures the precautionary savings component of consumption growth. It is always positive, increasing the growth rate of consumption.

For now, let's assume that the variance term is either zero, or constant. The log-linearized Euler equation leads to the following empirical specification:

$$\ln c_{t+1} = a + b \;\; \ln c_t + c \;\; x_t + d \;\; r_{t+1} + \varepsilon_{t+1}$$

and, if the equation is correctly specified, the point estimate $\hat{d}$ should be the Intertemporal Elasticity of Substitution $1/\theta$.[6]

Equations of that form have been estimated in literally hundred of papers. The goals of these regressions are usually two-fold:

1. estimate $1/\theta$ from $\hat{d}$, the IES from the coefficient on r_{t+1}
2. test the orthogonality restriction that information available at time t does not predict consumption growth: $\hat{c} = 0$. For example, expected income growth $E_t \Delta \ln \tilde{y}_{t+1}$ should not help predict consumption growth.

The literature typically finds:

[4] This results from the fact that if x is distributed $\mathcal{N}(\mu, \sigma)$ then $E[\exp(x)] = \exp(\mu + 1/2\sigma^2)$.

[5] The original CEQ model states that consumption in levels follows a random walk. The log-linearized result states that it is log-consumption that follows a random walk. The two are not very different for small deviations. Moreover, a random walk in logs is probably a better empirical specification given that consumption (and its innovations) grow over time. Campbell and Mankiw (1989) test the CEQ in logs.

[6] The constant a captures the sum of the impatience terms ρ/θ and the –constant– precautionary term $\theta/2 V_t \Delta \ln c_{t+1}$ and so does not provide useful information.

1. estimates of $1/\theta$ between 0 and 0.2, i.e. very low estimates of the sensitivity of consumption growth to the interest rate. Recall however, that aggregate saving can still be quite sensitive to the interest rate because people have long lifetimes (so the effects build up over long periods)

2. $\hat{c}$ is positive and significant when using $E_t \Delta \ln \tilde{y}_{t+1}$ as a regressor. This means that we reject the strict CEQ restriction. Expected income growth predicts expected consumption growth. This is sometimes referred to as the 'Excess Sensitivity of Consumption'. The conclusion is we reject the joint assumptions that (1) the Euler equation is true; (2) the utility function is CRRA; or (3) the linearization is accurate;

Remark 7 *The excess sensitivity is in response to variables that consumption should not respond too. In addition to excess sensitivity, there is an excess smoothness puzzle, whereby if income changes are very persistent, then innovations to consumption should be more volatile than innovations to income. One can see this is we assume the following process for income:* $\Delta \tilde{y}_{t+1} = \mu_{t+1} + \gamma \mu_t$, *with* $\gamma > 0$, *i.e. a MA(1) for income change. Then one can show that*

$$\epsilon_{t+1} = (1+\gamma)\mu_{t+1}$$

Deaton (1987) observed that despite $\gamma > 0$, *consumption innovations appear less volatile than income innovations.*

3.3.3 Campbell and Mankiw (1989)

This is an example of a paper estimating an equation similar to (9). They start from the baseline CEQ model (consumption is a random walk). Rather then simply test the null that all consumers are CEQ consumers, so that aggregate consumption also follows a random walk, they specify an alternative where a fraction λ of consumers are 'hand-to-mouth' with $c_t^h = y_t^h$ while the remaining consumers are CEQ with $\Delta c_{t+1}^r = \epsilon_{t+1}^r$. Aggregate consumption change is then the sum of the consumption change of the hand-to-mouth consumers and of the CEQ consumers:

$$\Delta c_{t+1} = \lambda \Delta y_{t+1} + (1-\lambda)\epsilon_{t+1}$$

Since ϵ_{t+1} is not observed, it is treated as the residual of the regression.

However, ϵ_{t+1} is likely to be correlated with income changes, so the equation above cannot be estimated directly. Instead, income changes need to be instrumented, using as instruments any lagged variables (hence orthogonal to ϵ_{t+1}) and good predictors of income growth. Equivalently, this consists in constructing a measure of predicted income change $E_t \Delta y_{t+1}$ and regressing realized consumption changes on predicted income changes.

This approach provides an alternative to the null that all consumers are CEQ. It even allows to quantify the share of consumers that are hand-to-mouth, as measured by λ. The estimates indicate that λ is quite high, on the order of 0.48 in the US (see Table 10).[7]

In addition, they run the regression including the real interest rate (see Table 11). They use as real interest rate the 3-month T-bill rate over the quarter minus the rate of change in the PCE deflator. They instrument the real interest rate using lagged real interest rates (since the regression should use ex-ante real interest rates).

The specification becomes:

$$\Delta c_{t+1} = \lambda \Delta y_{t+1} + (1-\lambda)/\theta r_{t+1} + (1-\lambda)\epsilon_{t+1}$$

Notice that the coefficient on the real interest rate is not the IES, but the IES multiplied by the share of CEQ consumers. Hence, $\lambda = 0$, the coefficient is going to be biased downwards.

Table 2 EVIDENCE FROM ABROAD
$\Delta c_t = \mu + \lambda \Delta y_t$

	Country (sample period)	First-stage regressions: Δc equation	First-stage regressions: Δy equation	λ estimate (s.e.)	Test of restrictions
1	Canada	0.047	0.090	0.616	0.007
	(1963–1986)	(0.127)	(0.030)	(0.215)	(0.263)
2	France	0.083	0.166	1.095	−0.055
	(1970–1986)	(0.091)	(0.015)	(0.341)	(0.714)
3	Germany	0.028	0.086	0.646	−0.030
	(1962–1986)	(0.211)	(0.031)	(0.182)	(0.639)
4	Italy	0.195	0.356	0.400	−0.034
	(1973–1986)	(0.013)	(0.000)	(0.094)	(0.488)
5	Japan	0.087	0.205	0.553	0.018
	(1959–1986)	(0.020)	(0.000)	(0.096)	(0.178)
6	United Kingdom	0.092	0.127	0.221	0.086
	(1957–1986)	(0.012)	(0.002)	(0.153)	(0.010)
7	United States	0.040	0.079	0.478	0.004
	(1953–1986)	(0.092)	(0.014)	(0.158)	(0.269)

Note: For all countries, the consumption data are total spending. The set of instruments is: $\Delta y_{t-2}, \ldots, \Delta y_{t-4}, \ldots, \Delta c_{t-2}, \ldots, \Delta c_{t-4}. c_{t-2} - y_{t-2}$. Also see note, Table 1.

Figure 10: Campbell and Mankiw (1989), Table 2.

[7] Note that Campbell and Mankiw in fact log consumption when they estimate their regression, so in fact they are working with a specification close to the log-linearized CEQ.

Table 5 UNITED STATES, 1953–1986
$\Delta c_t = \mu + \lambda \Delta y_t + \theta r_t$

Row	*Instruments*	*First-stage regressions* Δc	Δy	r	λ (s.e.)	θ (s.e.)	*Test of restrictions*
1	None (OLS)	—	—	—	0.294 (0.041)	0.150 (0.070)	—
2	$\Delta y_{t-2}, \ldots, \Delta y_{t-4}$ $r_{t-2}, \ldots, r_{t-4}$	0.045 (0.061)	0.030 (0.125)	0.471 (0.000)	0.438 (0.189)	0.080 (0.123)	−0.010 (0.441)
3	$\Delta c_{t-2}, \ldots, \Delta c_{t-4}$ $r_{t-2}, \ldots, r_{t-4}$	0.062 (0.026)	0.046 (0.060)	0.455 (0.000)	0.467 (0.152)	0.089 (0.110)	−0.006 (0.391)
4	$\Delta i_{t-2}, \ldots, \Delta i_{t-4}$ $r_{t-2}, \ldots, r_{t-4}$	0.092 (0.005)	0.034 (0.106)	0.431 (0.000)	0.657 (0.212)	0.016 (0.146)	−0.022 (0.665)

Note: See Table 1

Figure 11: Campbell and Mankiw (1989), Table 5.

3.3.4 Household level data: Shea (1995), Parker (1999), Souleles (1999) and Hsieh (2003)

While early papers (like Hall (1978) or Campbell and Mankiw (1989)) tested the CEQ using aggregate data, the literature quickly moved to testing the CEQ using household level data. There are a number of reasons why this is more satisfying:

- aggregate consumption does not fluctuate much. By contrast, household level consumption can fluctuate a lot more
- we can hope to identify more precisely expected future income for some categories of workers (Shea (1995)) and so we get cross sectional variation that adds to the power of the tests
- we can use 'natural experiments' where households learn about future income in a measurable way, and then measure the response of consumption (Parker (1999), Souleles (1999), Hsieh (2003)).

But household level data can also be problematic: accurate time series data on consumption and income for a given household is hard to obtain. One needs to worry about sampling weights, sample attrition...

- Shea (1995, AER): uses the PSID. PSID is a panel with reasonable income data, but only food consumption. It also oversample poor households. Shea matches head of households to union contracts and uses publicly available information on union wage growth to construct a measure of expected wage growth. He assigns respondents to unions with national or regional bargaining: trucking, postal service, railroads... or lumber in the Pacific Northwest, shipping on the East Coast.... He matches other respondents to dominant employers in some areas (e.g. worker in the automobile industry living near flint will be assigned to G.M.). Ends up with 647 observations from 285 households. Finds that expected wage growth predicts expected consumption growth.

- Parker (1999) uses the Consumption Expenditure Survey, used to construct the consumption weights for the Consumer Price index. Not a panel, but repeated cross sections. Each household stays in the CEX for 4 quarters, which yields 5 interviews, with income and demographic data collected in interviews 2 and 5. The CEX has excellent coverage of consumption, but relatively poor measures of income and also has a very limited time dimension for each household. Parker uses the fact that different households will hit the cap at different times during the calendar year and therefore will see a drop in payroll taxes in different quarters. Crucially, the household knows in advance whether he/she will hit the cap in a given quarter, so consumption should not respond. Yet, he finds a predictable response of consumption. The elasticity of consumption to predictable income is around 0.5.

- Souleles (1999) looks at income tax refunds, using the CEX. Since income tax refunds are known in advance (taxpayers know their income of the previous year and calculate the size of the refund when they file), consumption should not change when the refund is received. The CEX survey asks the household about tax refunds in the second and final interviews. Souleles finds that consumption responds to the tax refund. In a more recent paper, Parker, Souleles, Johnson and McClellan (2013, AER) add a module to the CEX questionnaire during the period of the 2008 stimulus payment (a tax cut). Cross sectional variation arises from the random timing of receipts, based on the last two digits of the recipient SSN. The amount of the tax cut was also well publicized and known in advance ($300-$600 for individuals, $600-$1200 for couples....). They find that households spent 12-30% of the tax cut.

- Hsieh (2003) instead exploits the annual payments of the Alaska Permanent fund to all Alaskan residents. The amount disbursed is large and known in advance. Using the CEX again, Hsieh finds that there is little evidence of a change in consumption during the quarter of the receipt of the fund. Variation comes from changes over time in the size of the transfers as well as cross sectional variation in family size. Importantly, Hsieh finds that –while consumption does not respond to the payment of the Alaskan Permanent Fund– it does respond to tax refunds, suggesting that there is an issue of 'saliency' of the change in income.

Why does expected income growth predict consumption growth?

- Consumption and leisure are substitutes (Aguiar and Hurst 2005 JPE). Home production and work-related expenditures. Relevant when thinking about drop in consumption at retirement. Food expenditures drops because home production (not recorded in food expenditures) rises dramatically.

TABLE 2—RESPONSE OF CONSUMPTION TO ALASKA PFD

	dlog(*Nondurable consumption*)			dlog(*Durable consumption*)		
	(1)	(2)	(3)	(4)	(5)	(6)
$\frac{PFD_t \times Family\ Size_h}{Family\ Income_h}$	0.0002 (0.0324)	−0.0167 (0.0336)	−0.0034 (0.0328)	−0.1659 (0.0878)	−0.1741 (0.0916)	−0.1488 (0.0890)
Controls for:						
Family size	No	No	Yes	No	No	Yes
Year dummies	No	Yes	No	No	Yes	No
Number of observations	806	806	806	806	806	806

Notes: Dependent variable is $\log(C_{IV}/C_{III})$. Standard errors are in parentheses. All regressions are ordinary least squares (OLS) and include a quadratic in age and changes in the number of children and adults in the household.

Figure 12: Hsieh (2003): Table 2

- Households support lots of dependent in mid-life (children, college...). Hence they have no choice but to have high consumption when income is high
- households are liquidity constrained and impatients (more on this later)
- some consumers are 'rule of thumb' consumers (Campbell and Mankiw 1989)
- welfare costs of not optimizing constantly are second order

3.3.5 A Detour: GMM Estimation

With CRRA preferences, the general equation to be estimated is

$$c_t^{-\theta} = \beta E_t[\tilde{R}_{t+1} c_{t+1}^{-\theta}]$$

This is a non-linear equation with at least two parameters: θ and β. The general idea of the Hansen and Singleton (1982) Generalized Method of Moments (GMM) estimation method is to define $z_{t+1} = \beta \tilde{R}_{t+1} c_{t+1}^{-\theta} - c_t^{-\theta}$. Then, under the null that the model is correct, $E_t z_{t+1} = 0$. In other words, z_{t+1} is unpredictable based on any variable dated t or earlier. Suppose we consider such a variable $w_{i,t}$. Then we must have $E[z_{t+1} w_{i,t}] = 0$. The sample analog of this moment condition is: $v_i = \sum_{t=0}^{T-1} w_{i,t} z_{t+1} = 0$. If we have J such variables, then we can estimate the parameters of the model by minimizing $v' \Omega v$ where v 'stacks' all the moment conditions: $v = (v_1, v_2, ... v_J)'$ and where Ω is some weighting matrix. This method is very general (and is one of the reasons Hansen received the Nobel prize in 2013).

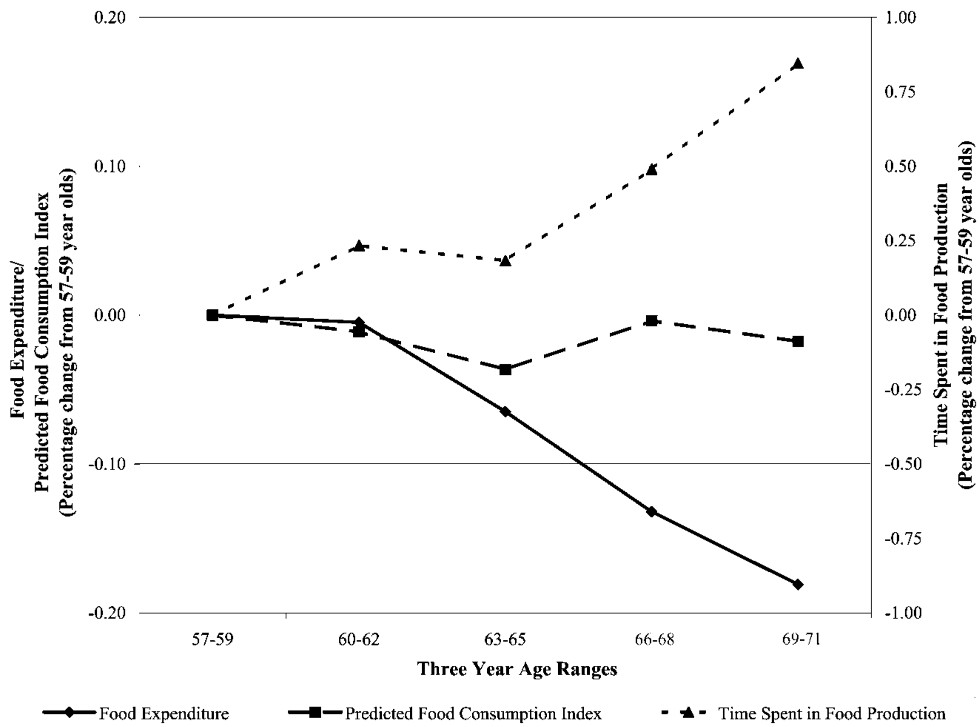

Figure 13: Aguiar & Hurst (2005), figure 1.

4 Moving beyond the CEQ

The CEQ model provides some interesting insights, mainly consumption smoothing (i.e. the fact that consumption is going to respond to permanent changes more than to transitory ones). Yet the model is soundly rejected.

In particular, consumption growth responds to income growth. The model also predicts that consumption profile should be dictated by the interest rate and the rate of time preference, not by the actual timing of income. This is also rejected: Carroll and Summers show that consumption profiles track income profiles for different occupations or in different countries. See figures 14 and 15.

Most households have relatively little wealth, so their consumption will mostly track their income. Yet, the use their modest saving to ‘buffer’ income shocks. We are interested in understanding these ‘buffer stock’ households.

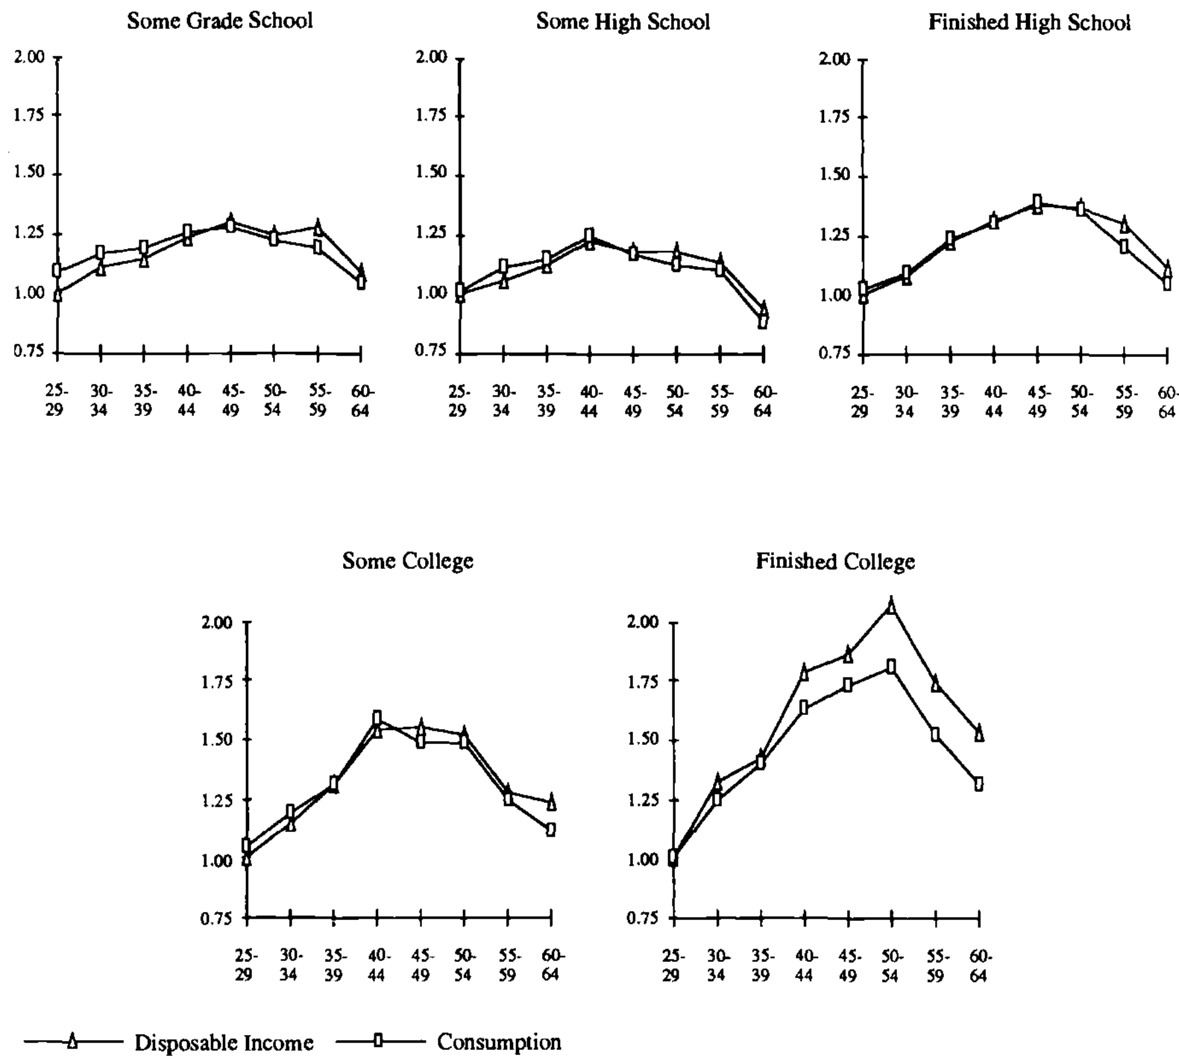

Fig. 10.7*a* Income and consumption profiles by educational group, 1960–61 CES
Source: Calculations by authors using CES tapes.

Figure 14: Carroll & Summers (1991), figure 10.7a

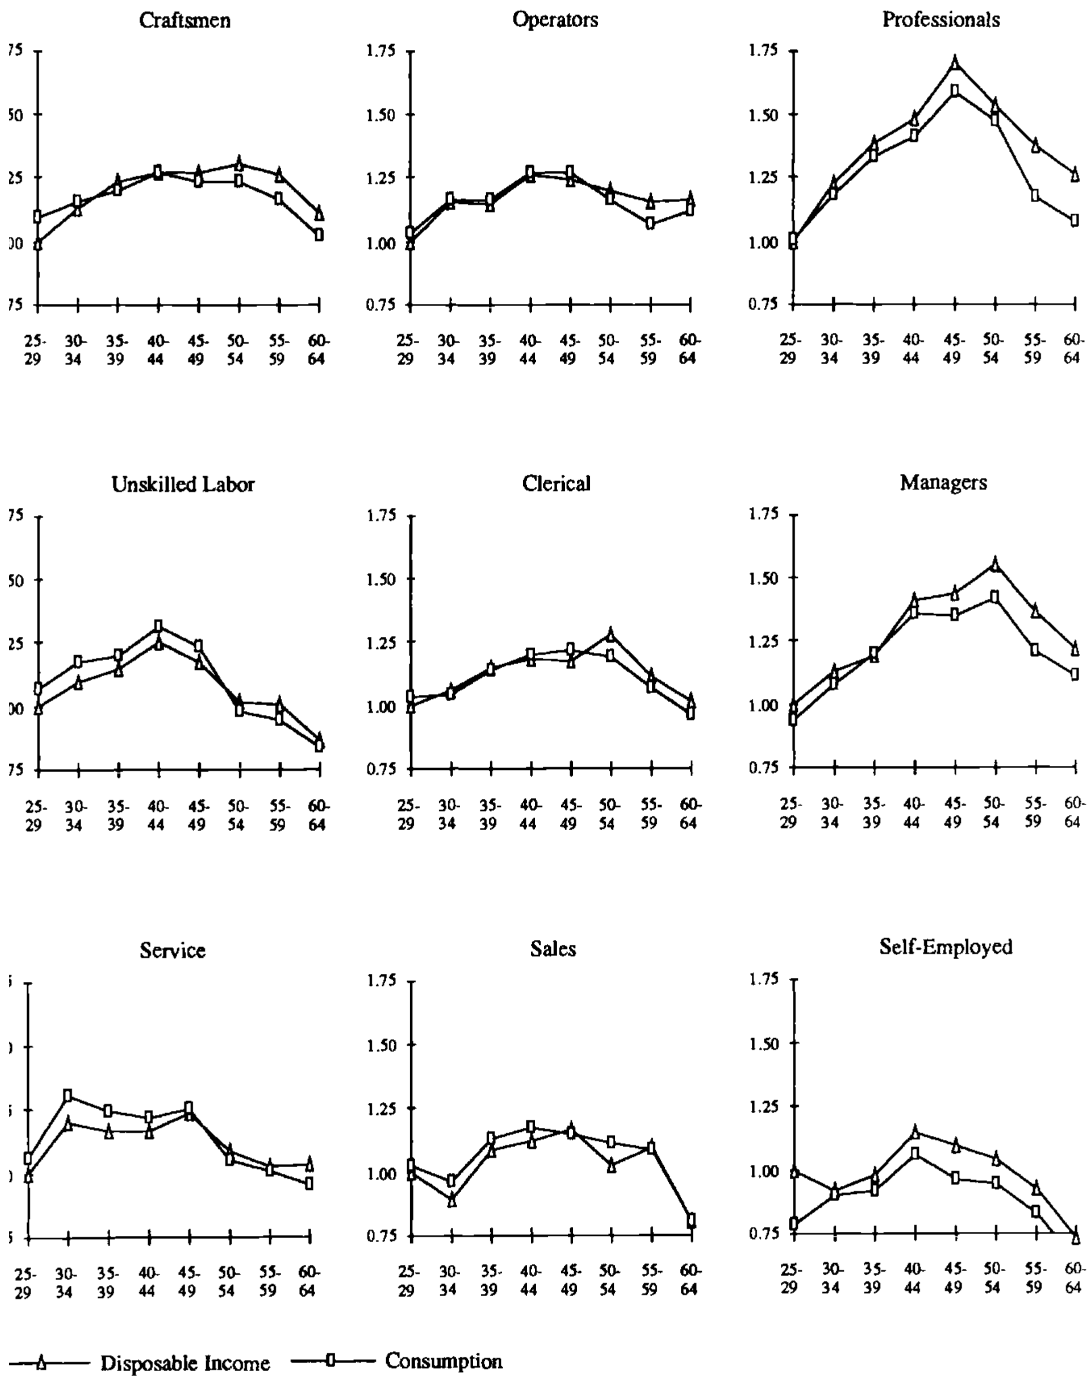

Fig. 10.7*b* Income and consumption profiles by occupational group, 1960–61 CES

Source: Calculations by authors using CES tapes.

Figure 15: Carroll & Summers (1991), figure 10.7b

4.1 Precautionary Saving

If we move away from linear marginal utility (i.e. quadratic utility), we open the door to precautionary saving. In that case, the household will care about higher moments of future consumption. Suppose that $u'(.)$ is convex, i.e. that $u'''(.)$ is positive. In that case, we know from the convexity of u' that:

$$E_t[u'(c_{t+1})] \geq u'(E_t[c_{t+1}])$$

and the inequality is strict when there is uncertainty about future consumption. This implies that we cannot have $c_t = E_t[c_{t+1}]$. If this were the case, then :

$$u'(c_t) = u'(E_t[c_{t+1}]) < E_t[u'(c_{t+1})]$$

which would violate the Euler equation. Consumption c_t has to be lower than in the CEQ case , that is

$$c_t < E_t[c_{t+1}]$$

Uncertainty depresses current consumption and thus raises saving. This is known as precautionary saving (or saving for a rainy day).

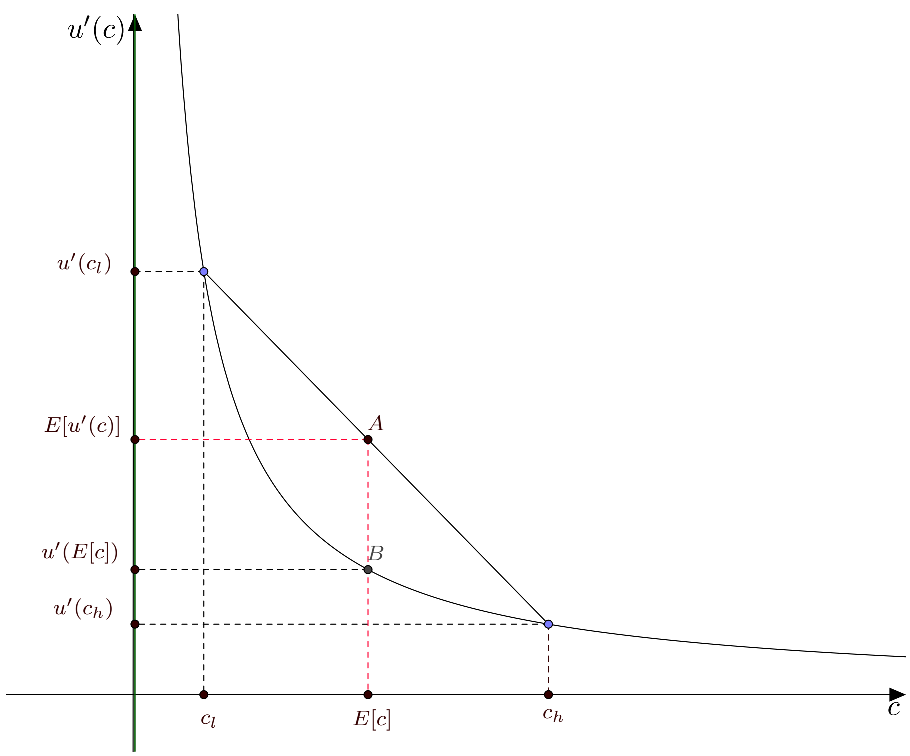

Figure 16: Convex marginal utility

For general preferences, Kimball (1990) showed that what matters for the precautionary motive is the concavity of $-u'(.)$ (or equivalently the convexity of $u'(.)$). He defined a coefficient of relative prudence as:[8]

$$CRP = -\frac{u'''(c)c}{u''(c)}$$

Remark 8 *For CRRA preferences $u'(c) = c^{-\theta}$ and the coefficient of relative prudence is constant and equal to:*

$$CRP = -\frac{\theta(1+\theta)c^{-\theta-1}}{-\theta c^{-\theta-1}} = 1+\theta$$

To see the role that prudence plays, consider the case where $\beta R = 1$ (so that a CEQ household would want to keep consumption flat over time) and let's perform a second-order Taylor expansion of the Euler equation around c_t:

$$\begin{aligned} E[u'(c_{t+1})] &\approx E[u'(c_t) + u''(c_t)(c_{t+1} - c_t) + 1/2u'''(c_t)(c_{t+1} - c_t)^2] \\ &\approx u'(c_t) + u''(c_t)E[(c_{t+1} - c_t)] + 1/2u'''(c_t)E_t[(c_{t+1} - c_t)^2] \end{aligned}$$

Now, substitute into the Euler equation (with $\beta R = 1$) to obtain:

$$u'(c_t) = u'(c_t) + u''(c_t)E[(c_{t+1} - c_t)] + 1/2u'''(c_t)E_t[(c_{t+1} - c_t)^2]$$

Rearrange to solve for expected consumption growth:

$$E[\frac{c_{t+1} - c_t}{c_t}] = -\frac{u'''(c_t)c_t}{u''(c_t)}E_t[(\frac{c_{t+1} - c_t}{c_t})^2] = CRP \;\; E_t[(\frac{c_{t+1} - c_t}{c_t})^2]$$

The slope of consumption growth is controlled by the coefficient of relative prudence and something that looks like the conditional variance of consumption growth (it's not quite the variance since $c_t \neq E_t[c_{t+1}]$).

Everything else equal, precautionary saving tends to increase expected consumption growth. How does this modify the Euler equation estimation?

To answer this question, let's consider again the log-linearized Euler equation (9) rewritten below for convenience:[9]

$$E_t \Delta \ln c_{t+1} \approx \frac{1}{\theta}(r_{t+1} - \rho) + \frac{1}{2}\theta V_t \Delta \ln c_{t+1}$$

[8] There is also a coefficient of absolute prudence, defined as $-u'''(c)/u''(c)$.

[9] Technical observation: in the log-linear Euler equation the conditional variance is scaled by θ while in the derivation above, it is scaled by the CRP $1+\theta$. The reason for the difference is that in the log-linear Euler equation there is an adjustment for the concavity of the log (Jensen's inequality).

This gives the following empirical specification:

$$\Delta \ln c_{t+1} = \frac{1}{\theta}(r_{t+1} - \rho) + \frac{1}{2}\theta V_t \Delta \ln c_{t+1} + \varepsilon_{t+1}$$

Suppose for the time being that we ignore the precautionary term in this expression, i.e. that we lump it with the error term of the regression ε_{t+1} and estimate the first-order log-linear Euler equation. This would be valid if we satisfy the orthogonality condition that the precautionary term is orthogonal to the interest rate. Carroll (1997) [death to the log-linearized Euler equation!] argues that this is unlikely to be the case. The reason is that both the interest rate r_{t+1} and the conditional variance of consumption growth are endogenous objects and are likely to interact in the general equilibrium of the economy. This implies that the precautionary saving term is an omitted variable that is likely to be correlated with the equilibrium interest rate.

To see how this works most simply, suppose that the economy is on a balanced growth path where households face some level of idiosyncratic risk, but no aggregate risk.[10] Along such a balanced growth path, the growth rate of aggregate consumption must equal the growth rate of aggregate income which is certain, since there is no aggregate risk, and which we denote $g_y = \Delta \ln y_{t+1}$. If households are ex-ante identical, facing the same amount of risk etc... they will all choose the same expected consumption growth, therefore also equal to g_y. It follows from the second order log-linearized Euler equation that:

$$g_y \approx \frac{1}{\theta}(r_{t+1} - \rho) + \frac{1}{2}\theta V_t \Delta \ln c_{t+1}$$

This equation tells us that in the aggregate equilibrium, the variance of idiosyncratic consumption growth and the interest rate will be related by:

$$r_{t+1} = \theta g_y + \rho - \frac{\theta^2}{2} V_t \Delta \ln c_{t+1}$$

A higher amount of uninsurable idiosyncratic risk (as measured by $V_t \Delta \ln c_{t+1}$) will be associated with a lower real riskfree interest rate r_{t+1}. Carroll concludes that estimating the first-order log-linearized Euler equation (i.e. without controlling for precautionary saving) is likely to be seriously misspecified. This can explain in particular why the estimated IES is very low.

A way to address this critique would be to incorporate directly in the regression a term that controls for the importance of precautionary saving, i.e. for the term $\frac{\theta}{2} V_t \Delta \ln c_{t+1}$ in the regression. This is what Dynan (1993) does by adding proxies for income uncertainty. But it is difficult to obtain such estimates in the first place, and if we try to instrument for the precautionary saving motive, we have to be careful to find instruments that for precautionary saving that are independent from the interest rate, not an easy task.

[10] Implicitly this requires that some risk sharing opportunities are not exploited. Otherwise, households would like to diversify their idiosyncratic risk away. Technically, such models are called Bewley models, after Bewley (1977). For a seminal Bewley model, see Aiyagari (1994).

4.2 The Buffer Stock Model

Intuitively, precautionary saving tilts-up consumption profiles and therefore leads to more saving and wealth accumulation. Consider a household that faces income uncertainty. If that household has a high wealth level, then heuristically income uncertainty should not matter much and therefore consumption should not be too different from the certainty equivalent framework (CEQ). We know that in that case, what controls the slope of the expected consumption profile (and therefore of subsequent wealth) is whether βR is smaller of greater than 1.

- If $\beta R > 1$, the household is patient and would like to save. In that case, the precautionary and smoothing motive push in the same direction: eventually, the household will manage to accumulate enough assets to insure against income fluctuations. In fact, if $\beta R > 1$ an infinitely lived household would accumulate an unbounded level of assets.
- If $\beta R = 1$, the argument is a bit more subtle, but the result is the same. Here, the household would like to smooth marginal utility. It will be able to do this by accumulating an unbounded amount of wealth.[11] The upshot is that if $\beta R \geq 1$ the model is not terribly interesting: the household would just accumulate vast amounts of wealth, enough to be indifferent to the impact of income fluctuations on marginal utility. This is neither interesting nor realistic!
- The last case is when $\beta R < 1$. In that case, the household is impatient. A CEQ household would run choose to consume more today and run down assets. But by running down assets, it increases the strength of the precautionary saving motive since income fluctuations are more likely to impact marginal utility. So this case present an interesting tension: on the one hand, the household would like to save to smooth fluctuations in marginal utility. On the other hand, it wants to consume now and prefers not to accumulate wealth. The result from this tension is that the household will aim to achieve a certain target level of liquid wealth, but not more. Once households have accumulated this target level of wealth, consumption will tend to track income at high frequency (even in response to predictable income change), thus potentially explaining the excess sensitivity puzzle. It can also explain why consumption tracks income at low frequency (explaining the Carroll-Summers (1991) empirical patterns in figures 14 and 15). This is the buffer-stock model.

Let's flesh the details of that model out. Consider a household with standard preferences:

$$U = \max_{\{c_t\}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

[11] This result is formally established by Schechtman (1975) and Bewley (1977). See Deaton (1991) for a discussion.

and with a budget constraint:

$$a_{t+1} = R(a_t + \tilde{y}_t - c_t)$$

The household faces a constant interest rate R but a stochastic income stream $\{\tilde{y}_t\}$, where we assume for simplicity that $\tilde{y}_t$ is independently identically distributed every period. We assume that $\beta R < 1$ so that, if there was no uncertainty, the household would prefer to consume now and would run down assets over time, and even borrow against future income.

How much would the household borrow? If $y_{\min} \geq 0$ is the lowest possible realization of income every period, then it is immediate to show that the household would not be able to run its asset levels below $a_{\min} = -y_{\min}/(R-1)$.[12] If the household borrowed a larger amount at any point in time, there would be a strictly positive probability that it would not be able to repay. In other words, $a_{\min}$ is the natural borrowing limit faced by the household. It is the present value of the lowest possible income the household would receive from now on, and $a_t \geq a_{\min}$. Of course, it is possible that the household faces a stronger liquidity constraint than the natural borrowing limit, if access to credit markets is limited. This is a relevant feature of the world since many people face limited access to credit markets.

In order to fix ideas, we are going to consider an extreme case where the household cannot borrow at all. That is, we impose the restriction that:[13]

$$a_t \geq 0$$

If there was no uncertainty, the solution to the household consumption-saving problem would be quite straightforward: it would run down initial assets a_0, then set consumption equal to income. With uncertainty, this is not going to be optimal for the reasons discussed above: it would leave the household exposed to too much fluctuations in marginal utility.

Therefore, there should be some target level of liquid wealth that the household would like to revert to.

We can write the income fluctuations problem as (see Deaton (1991)):

$$U = \max_{\{c_t\}} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to:

$$\begin{aligned} a_{t+1} &= R(a_t + \tilde{y}_t - c_t) \\ a_t &\geq 0 \\ c_t &\geq 0 \end{aligned}$$

[12]This condition derives from the intertemporal budget constraint and the requirement that consumption remain positive.

[13]This corresponds to the natural borrowing limit if $y_{\min} = 0$.

It is useful to express the problem in terms of 'cash on hand' x_t, defined as the amount of liquid resources the household has access to at the beginning of the period:

$$x_t = a_t + y_t$$

The constraints of the problem become:

$$\begin{aligned} x_{t+1} &= R(x_t - c_t) + \tilde{y}_{t+1} \\ 0 &\leq c_t \leq x_t \end{aligned}$$

Let's define $v(x_t)$ the value function of this problem. We can write the associated Bellman equation:

$$\begin{aligned} v(x_t) &= max_{c_t} u(c_t) + \beta E_t[v(x_{t+1})] \\ \text{s.t.} & \\ x_{t+1} &= R(x_t - c_t) + \tilde{y}_{t+1} \\ 0 &\leq c_t \leq x_t \end{aligned}$$

The first order condition associated with this Bellman equation is:

$$u'(c_t) = \beta R E_t[v'(c_{t+1})] + \lambda_t$$

where λ_t is the Lagrange multiplier associated with the constraint $c_t \leq x_t$.[14] The complementary slackness condition is:

$$\lambda_t(x_t - c_t) = 0$$

For the usual envelope reasons, the marginal value of cash on hand satisfies:

$$v'(x_t) = \beta R E_t[v'(x_{t+1})] + \lambda_t = u'(c_t)$$

It follows that:

- when the credit constraint does **not** bind, the usual Euler equation holds:

$$u'(c_t) = \beta R E_t[u'(c_{t+1})]$$

- when the credit constraint binds, $\lambda_t > 0$ and $c_t = x_t$ and

$$u'(x_t) > \beta R E_t[u'(c_{t+1})]$$

We can summarize both cases as follows:

$$u'(c_t) = \max \left\langle \beta R E_t[u'(c_{t+1})], u'(x_t) \right\rangle \tag{10}$$

The credit constraint $c_t \leq x_t$ operates in two ways:

[14] Technically there is another Lagrange multiplier associated with the constraint $c_t \geq 0$, but this one will never bind as long as the Inada conditions are satisfied, so we ignore it here.

1. If the household is constrained at time t, it is forced to consume less than desired.

2. The credit constraint also matters, even in periods where it does not bind directly, because of the likelihood that it will bind in the future. Technically, this is encoded in $E_t[u'(c_{t+1})]$. The curvature of marginal utility leads the household to save more to reduce the likelihood of being constrained in the future.

This model cannot be solved in closed form. Instead, we have to resort to numerical techniques to characterize optimal consumption behavior. Denote $c_t = f(x_t)$ the optimal consumption rule followed by the household. It is not a function of time because the problem is recursive and stationary. We can then rewrite the Euler equation as:

$$
\begin{aligned}
u'(f(x_t)) &= \max\left\langle \beta R E_t[u'(f(x_{t+1}))], u'(x_t)\right\rangle && (11)\\
x_{t+1} &= R(x_t - f(x_t)) + \tilde{y}_{t+1} && (12)
\end{aligned}
$$

The problem becomes one of solving for the function $f(.)$. The right hand side of equation (14) defines a functional equation:

$$
\begin{aligned}
T(f)(x) &= u'^{-1}\left(\max\left\langle \beta R E[u'(f(x_{+1}))], u'(x)\right\rangle\right)\\
x_{+1} &= R(x - f(x)) + \tilde{y}_{+1}
\end{aligned}
$$

where u'^{-1} is the inverse of the marginal utility (assumed well defined). The optimal consumption rule is then a fixed point of the operator $T(f)$:

$$f(x) = T(f)(x)$$

Not surprisingly, the regularity condition that ensures that this operator has a unique fixed point is $\beta R < 1$, i.e. precisely the requirement that the household is impatient.[15]

Moreover, this fixed point can be obtained by iteration. Suppose that we have a candidate consumption function $c(x) = f^n(x)$. Then we can construct $f^{n+1}(x)$ as

$$f^{n+1}(x) = T(f^n)(x)$$

i.e. as the solution of:

$$
\begin{aligned}
u'(f^{n+1}(x)) &= \max\left\langle \beta R E_t[u'(f^n(x_{+1}))], u'(x)\right\rangle && (13)\\
x_{+1} &= R(x - f^n(x)) + \tilde{y}_{+1} && (14)
\end{aligned}
$$

The sequence $f^n(x)$ converges uniformly to $f(x)$, i.e. $\lim_{n\to\infty} ||f^n(x) - f(x)|| = 0$ where $||.||$ is some Euclidean distance.

This is called Euler equation iteration.[16]

Figure 17 shows the optimal consumption rule for this problem for the case where $y_{\min} > 0$. It has the following properties:

[15] Technically, this condition ensures that the operator $T(f)$ is a contraction mapping.

[16] Another approach, called value function iteration works with the value function $v(x_t)$ that solves the Bellman equation.

Certainty Equivalent	Buffer Stock Model
forward looking	much less forward looking
retirement saving	households will not save for retirement at age 20
consumption and income paths independent	Once you have your buffer, $g_c \approx g_y$
interest rate elasticity	small effect of interest rate
uncertainty does not matter	uncertainty matter

Table 1: Comparing CEQ and Buffer Stock Models

- consumption is a function of x, not y.
- below a certain threshold level x^*, the household prefers to consume all its assets: $c = x$. This is because the current marginal utility of consumption is very high.
- above x^*, the consumption rule is concave, and always below the certainty equivalent consumption
- we can represent expected consumption growth $E_t \Delta \ln c$ as a function of cash on hand x. It is a decreasing function:
  - for low levels of wealth, precautionary saving dominate, cash on hand will increase and consumption is expected to grow.
  - For high levels of cash on hand, consumption grows at rate $\beta R < 1$ so cash on hand decreases.
  - The target level of cash on hand can be defined as that level that remains constant (in expectations), i.e. the level x^{**} such that $E[x_{t+1}|x_t = x^{**}] = x^{**}$. Carroll (2012) shows that that expected consumption growth is below 1 and above βR at the target level of cash-on-handx^{**}.[17]
- Even if c is a function of x, once x is close to its target, c will move together with y: if y is expected to decline, then consumption will decline once x declines (not before): predictable movements in y will translate into movements in c.

Figure 18 reports the dynamics of the buffer stock model, as computed by Carroll (2012). We can summarize the two models as in table 1:

4.3 Consumption over the Life Cycle

See Gourinchas & Parker (2002) [GP]. Revisits the question of optimal consumption behavior:

- model with both lifecycle saving motive and precautionary saving motive
- structural estimation of the consumption function, i.e. not relying on Euler equation, or reduced form consumption functions

[17]For more on this, see Carroll (2012), "Theoretical foundations of buffer stock saving."

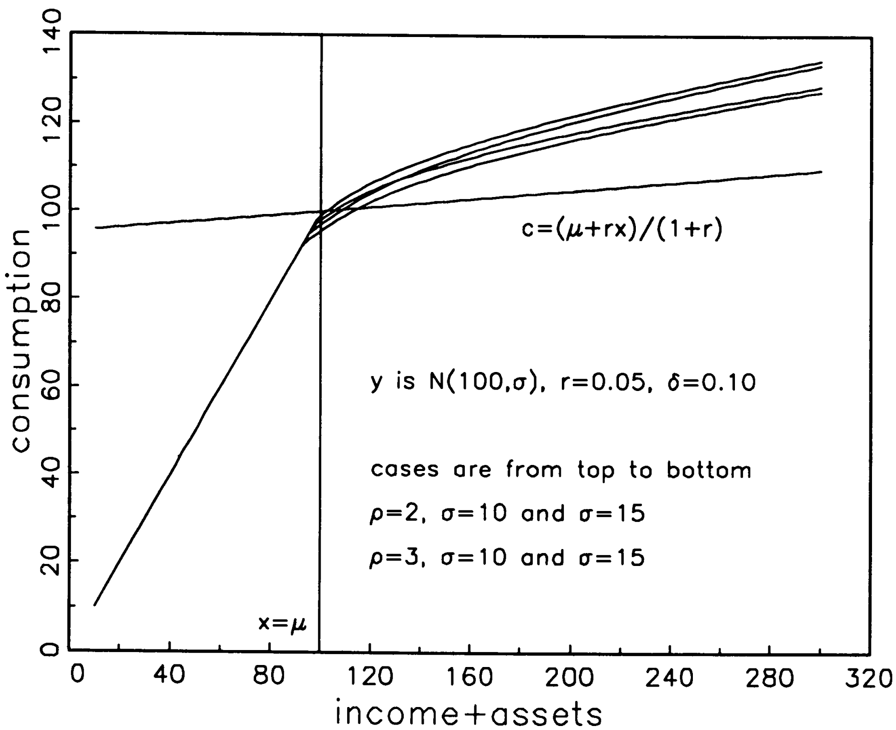

FIGURE 1.—Consumption functions for alternative utility functions and income dispersions.

Figure 17: Deaton (1991), figure 1

- estimation based on household level data using income and consumption expenditures

The estimation procedure consists in constructing age-profiles of consumption based on micro data and estimating the parameters of the consumption problems that best replicate these age profiles in the model.

4.3.1 The Model

Each household lives for T periods, works for N periods. GP truncate the problem at retirement by writing:

$$U = E_0[\sum_{t=0}^{N-1} \beta^t u(c_t) + \beta^N V_N(a_N)]$$

subject to:

$$a_{t+1} = R(a_t + y_t - c_t)$$

The function $V_N(.)$ summarizes preferences from retirement onwards, including any bequest motive. GP assume that preferences are CRRA: $u'(c) = c^{-\theta}$. Further, they assume

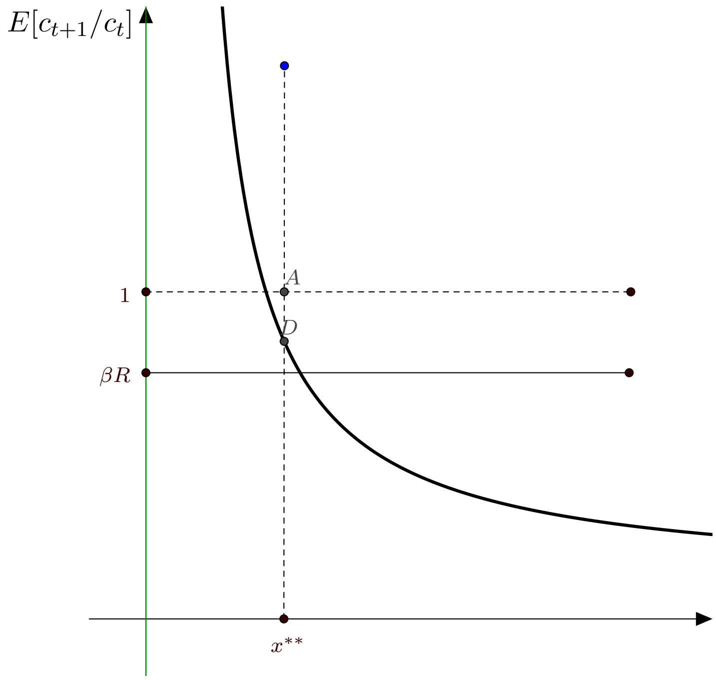

Figure 18: The Buffer Stock Model

that labor income y_t has a transitory and a permanent component:

$$\begin{aligned} y_t &= p_t \mu_t \\ p_t &= G_t p_{t-1} \eta_t \end{aligned}$$

where μ_t and η_t are iid. GP assume that $\mu_t = 0$ with some probability p. This is meant to capture unemployment risk. One implication is that $y_{\min} = 0$ so the natural borrowing limit is $a_{\min} = 0$. However, with preferences that satisfy the Inada conditions, the household will never choose to hit the borrowing limit.[18] Otherwise, $\ln \mu_t$ is $\mathcal{N}(0, \sigma_\mu^2)$. $\ln \eta_t$ is also normally distributed $\mathcal{N}(0, \sigma_\eta^2)$. The variance of these shocks and the unemployment risk are calibrated to household level data.

The problem features two state variables: cash on hand $x_t = a_t + y_t$ and the permanent level of income p_t (since the latter conditions how large future income will be). In general, the complexity of numerical problems grows exponentially with the number of state variables (curse of dimensionality). Even with modern computers and parallelization techniques, we cannot realistically solve problems with more than 1 or 2 state variables. Fortunately, the assumptions of the problem allow to implement a normalization that reduces the number of

[18]To see this, observe that if $a = 0$ in one period, then there is a strictly positive probability that the agent will have zero consumption next period if the unemployment shock is realized.

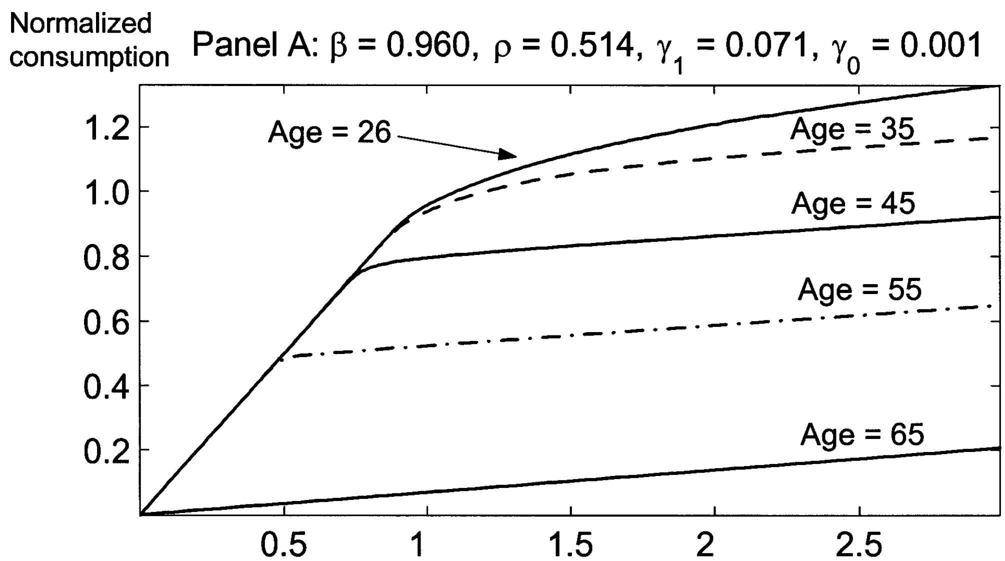

Figure 19: Gourinchas & Parker (2002), figure 1

state variables. Define $\hat{z}$ as the ratio of variable z to the permanent component of income: $\hat{z} = z/p$. Then, we can rewrite the problem's Euler equation as:

$$u'(\hat{c}_t(\hat{x}_t)) = \beta R E_t \left[u'(G_{t+1}\eta_{t+1}\hat{c}_{t+1}(\hat{x}_{t+1})) \right]$$

where

$$\hat{x}_{t+1} = R(\hat{x}_t - \hat{c}_t)/(G_{t+1}\eta_{t+1}) + \mu_{t+1}$$

Notice that the consumption rules are indexed by age, since households of different ages face different remaining horizons. So unlike the infinite horizon model, we don't need to look for a fixed point of the Euler equation, but iterate backwards, starting from the consumption rule at retirement, assumed to be linear in cash on hand:[19]

$$\hat{c}_N = \gamma_1 \hat{x}_N$$

For any value of the parameters, the resulting consumption functions can be evaluated numerically (see figure 19).

4.3.2 Estimating the Structural Model

Once the optimal consumption rules $\hat{c}_t(\hat{x})$ are evaluated, GP construct the age-consumption profiles $\bar{C}_t$ by aggregating over the distribution of possible realizations of the state variables

[19]This implies that in principle we don't need to ensure that the functional equation defines a contraction mapping. If the horizon were infinite, the condition for buffer stock behavior would be $\beta R E[(G\eta)^{-\theta}] < 1$

x and p:

$$\ln \bar{c}_t = E[\ln(\hat{c}_t(x_t/p_t)p_t)] = \int \ln(\hat{c}_t(x_t/p_t)p_t)dF_t(x_t, p_t; \psi)$$

where $dF_t(.,.)$ denotes the joint distribution, according to the model, of cash on hand and permanent income at age t. In practice, this joint distribution is a complex object to calculate. An important step is to construct the moments above by simulating a large number of households (a Monte Carlo simulation). The simulated moment converges to the true model moment as the size of the simulation increases.

The last step is in matching these simulated moments to the same moments constructed on the household level data. The algorithm chooses the vector of structural parameters to minimize the distance between these simulated moments and the data moments.

If we denote $g_t(\psi) = 1/I_t \sum_i \ln c_{i,t} - \ln \bar{c}_t(\psi)$ as the distance between simulated and data consumption at age t, then the estimator minimizes

$$g(\psi)'Wg(\psi)$$

where W is a weighting matrix and $g = (g_1, g_2, ..., g_{N-1})'$.

The results of the estimation indicate the following:

- Consumption tracks income over the lifecycle
- The estimated parameters (θ and β) are quite reasonable with $\beta = 0.96$ and $\theta = 0.51$, with $R - 1 = 3.44\%$. The model thus features "buffer stock" behavior in the sense that agents want to keep a constant target level of cash on hand around 1.2 times permanent income until around age 40.
- Around age 40, life cycle considerations kick in and saving increase markedly. At that point, precautionary saving become less relevant since the higher liquid wealth allows the household to smooth consumption. Hence in the latter phase of active life, households behave like CEQ consumers with consumption growth controlled by βR.
- in the early part of their lifecycle, income is growing but households are unable to borrow much. Hence precautionary saving dominates.
- The structural parameters are well identified precisely because the turning point at which savings increase is determined by the relative strength of the two saving motives.

The lifecycle model with uninsurable labor income risk has become a workhorse to evaluate quantitatively various policies. For instance, Scholz et al (2006) use it to evaluate whether Americans are saving optimally. Their model features uncertain lifetimes, uninsurable earnings, medical expenses, progressive taxation, government transfers and social security benefits. They use the model to compare, household by household, wealth predictions and find that the model accounts for more than 80% of the 1992 cross sectional variation in

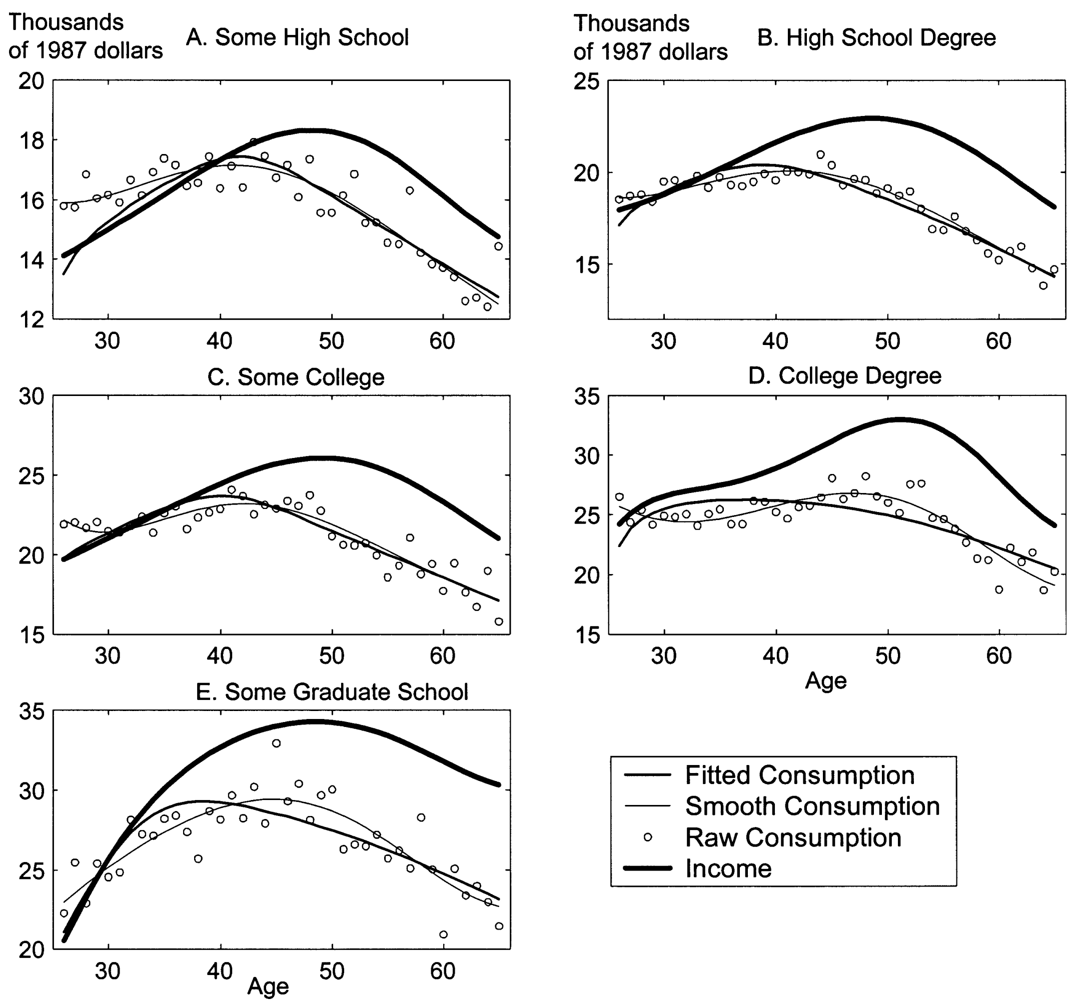

FIGURE 3.—Household consumption and income over the life cycle, by education group.

Figure 20: Gourinchas & Parker (2002), figure 2

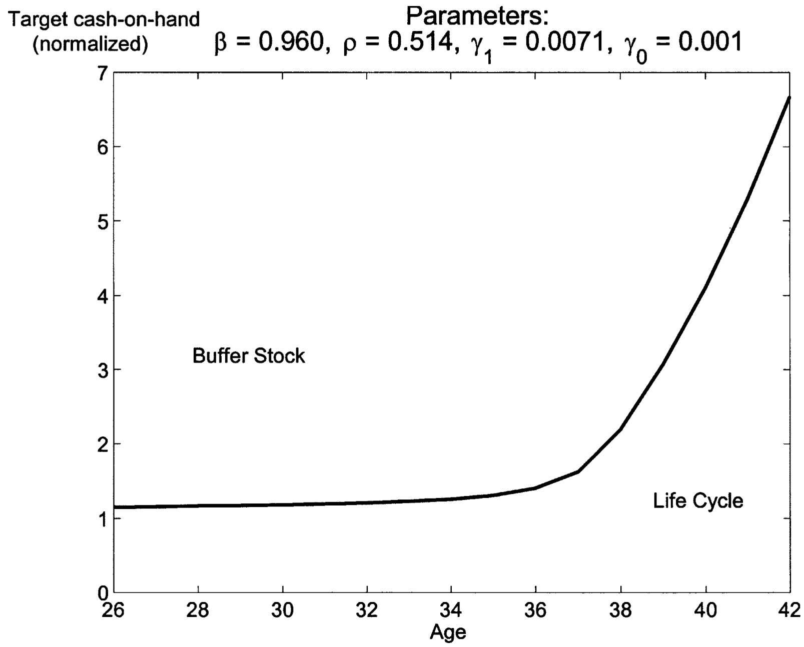

Figure 21: Gourinchas & Parker (2002), figure 6. Target cash on hand over age.

wealth. The paper uses data from the Health and Retirement Study (HRS) supplemented with restricted social security data on earning realizations throughout life. The results indicate that most households save enough for retirement, especially give the fact that they contribute to social security and employer retirement plans.

5 Asset Pricing

5.1 The Canonical Model Again with Multiple Assets

We now switch focus and use the canonical model to tell us about asset prices. Let's consider again the canonical model. Now assume that the household can invest in two assets:

- a riskless asset that pays a riskfree return R_{t+1}
- a risky asset that pays a risky return $\tilde{Z}_{t+1}$

The problem of the household is:

$$U = max_{\{c_t,\omega_t\}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to:

$$a_{t+1} = (R_{t+1}(1-\omega_t) + \tilde{Z}_{t+1}\omega_t)(a_t + \tilde{y}_t - c_t)$$

TABLE 1
Descriptive Statistics for the Health and Retirement Study (Dollar Amounts in 1992 Dollars)

Variable	Mean	Median	Standard Deviation
1991 earnings	\$35,958	\$28,976	\$39,368
Present discounted value of lifetime earnings	\$1,718,932	\$1,541,555	\$1,207,561
Defined-benefit pension wealth	\$106,041	\$17,327	\$191,407
Social security wealth	\$107,577	\$97,726	\$65,397
Net worth	\$225,928	\$102,600	\$464,314
Mean age (years)	55.7		4.7
Mean education (years)	12.7		3.4
Fraction male	.70		.46
Fraction black	.11		.31
Fraction Hispanic	.06		.25
Fraction couple	.66		.48
No high school diploma	.22		.41
High school diploma	.55		.50
College graduate	.12		.33
Postcollege education	.10		.30
Fraction self-employed	.15		.35
Fraction partially or fully retired	.29		.45

Source.—Authors' calculations from the 1992 HRS. The table is weighted by the 1992 HRS household analysis weights.

Figure 22: Scholz , Sheshadri & Khitatrakun (2006), Table 1

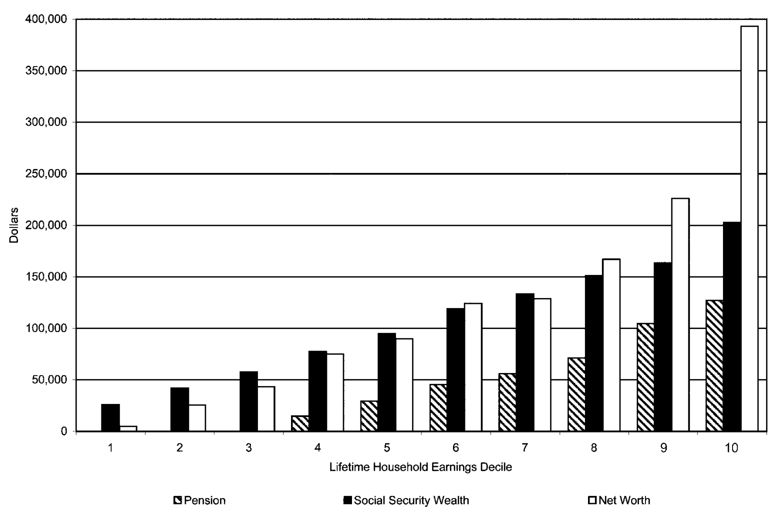

Fig. 1.—Median defined-benefit pension wealth, social security wealth, and net worth (excluding defined-benefit pensions) by lifetime earnings decile (1992 dollars).

Figure 23: Scholz , Sheshadri & Khitatrakun (2006), Figure 1

TABLE 2

Optimal Net Worth (Excluding Social Security and Defined-Benefit Pensions) and the Percentage of Population Failing to Meet Optimal Targets (Dollar Amounts in 1992 Dollars)

Group	Median Optimal Wealth Target (1)	Mean Optimal Wealth Target (2)	Percentage below Optimal Target (3)	Median Deficit (Conditional) (4)	Median Net Worth (5)	Median Social Security Wealth (6)	Median Defined-Benefit Pension Wealth (7)
All households	$63,116	$157,246	15.6%	$5,260	$102,600	$97,726	$17,327
No high school diploma	20,578	70,711	18.6	2,632	36,800	72,561	0
High school diploma	63,426	139,732	15.6	5,702	102,885	97,972	21,290
College degree	128,887	243,706	12.7	14,092	209,616	127,704	60,752
Postcollege education	158,926	333,713	13.2	23,234	253,000	129,320	152,781
Earnings decile:							
Lowest	2,050	48,445	30.4	2,481	5,000	26,202	0
2nd	13,781	55,898	28.7	3,328	25,500	42,159	0
3rd	26,698	84,582	21.8	5,948	43,485	57,844	0
4th	43,566	123,441	19.4	4,730	75,000	77,452	14,830
Middle	53,709	128,285	16.9	6,979	90,000	94,929	29,497
6th	76,462	131,565	10.8	10,000	124,348	119,011	45,613
7th	80,402	154,891	9.9	11,379	128,580	133,451	56,033
8th	101,034	180,643	5.5	21,036	167,000	151,397	71,373
9th	136,075	238,186	4.4	5,206	226,000	163,639	104,657
Highest	238,073	463,807	5.4	25,855	393,000	202,659	126,998

Source.—Authors' calculations as described in the text.

Figure 24: Scholz , Sheshadri & Khitatrakun (2006), Table 2

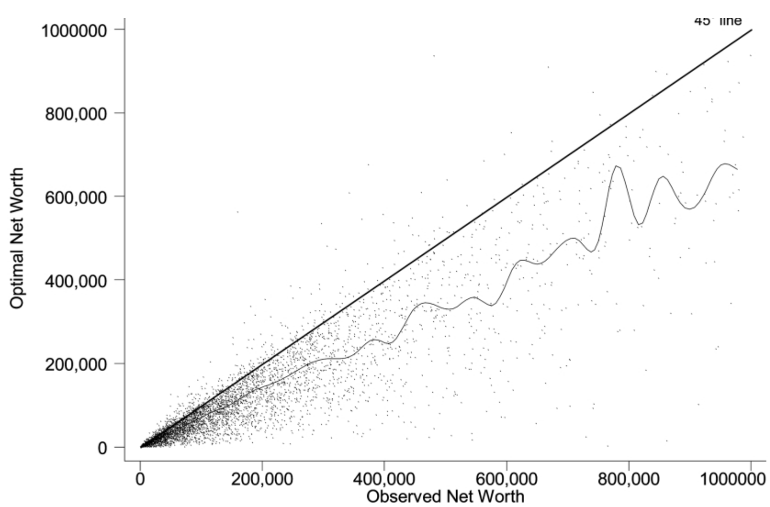

Fig. 2.—Scatter plot of optimal and actual wealth. Observed net wealth is constructed from the 1992 HRS. Optimal net worth comes from solving the baseline model described in the text.

Figure 25: Scholz , Sheshadri & Khitatrakun (2006), Figure 2

This problem features two control variables: how much to consume c_t, and how much to invest in the risky asset (ω_t). Following the same steps as before (Bellman equation, first order condition, envelope theorem), one can show that the following conditions hold:

$$\begin{aligned} u'(c_t) &= \beta E_t \left[u'(c_{t+1})(R_{t+1}(1-\omega_t) + \tilde{Z}_{t+1}\omega_t) \right] \\ R_{t+1} E_t \left[u'(c_{t+1}) \right] &= E_t \left[u'(c_{t+1}) \tilde{Z}_{t+1} \right] \end{aligned}$$

The first condition is simply the usual Euler equation. The second one is an asset pricing condition. In fact, combining the two equations, for any asset $\tilde{Z}_{i,t}$, optimal portfolio allocation requires that

$$1 = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \tilde{Z}_{i,t+1} \right]$$

How should we interpret this condition? Recall that we can write, for any two random variables: $E[XY] = E[X]E[Y] + cov(X, Y)$. Substituting, we obtain:

$$\begin{aligned} 1 &= E_t \left[\mathcal{M}_{t,t+1} \tilde{Z}_{i,t+1} \right] \\ &= E_t \left[\mathcal{M}_{t,t+1} \right] E_t \left[\tilde{Z}_{i,t+1} \right] + \text{cov}_t \left(\mathcal{M}_{t,t+1}, \tilde{Z}_{i,t+1} \right) \\ E_t \left[\tilde{Z}_{i,t+1} \right] &= R_{t+1} \left(1 - \text{cov}_t \left(\mathcal{M}_{t,t+1}, \tilde{Z}_{i,t+1} \right) \right) \end{aligned}$$

where we define $\mathcal{M}_{t,t+1} = \beta u'(c_{t+1})/u'(c_t)$, i.e. the intertemporal marginal rate of substitution and the last line uses the fact that, from the Euler equation,

$$E_t \left[\mathcal{M}_{t,t+1} \right] R_{t+1} = 1$$

The interpretation is now quite straightforward: an asset requires a high expected return if it covaries negatively with the pricing kernel $\mathcal{M}_{t,t+1}$, i.e. if the return on the asset is low when $\mathcal{M}$ is high. Now if we go back to the definition of the IMRS, we see that it is high when consumption is low. In other words, an asset requires a premium if it offers a poor return precisely at times when consumption is low.

The required excess return satisfies

$$E_t \left[\tilde{Z}_{i,t+1} \right] - R_{t+1} = -R_{t+1} \text{cov}_t \left(\mathcal{M}_{t,t+1}, \tilde{Z}_{i,t+1} \right)$$

5.2 Stock Prices: a Present Value Formula

We can use the previous pricing equation to evaluate the value of a stock. Suppose we have an asset with price P_t at time t, resale value P_{t+1} at time $t+1$ and a dividend d_{t+1} in period $t+1$. Then the return to the asset is:[20]

$$Z_{t+1} = \frac{P_{t+1} + d_{t+1}}{P_t}$$

[20] Here P_t is the price after the dividend in period t has been paid, i.e. an ex-dividend price

Substituting into the asset pricing equation, and solving forward, we obtain:

$$P_t = E_t \left[\sum_{s=1}^{\infty} \mathcal{M}_{t,t+s} d_{t+s} \right]$$

in other words, the stock price is the expected PDV of future dividends, evaluated using the pricing kernel defined recursively as $\mathcal{M}_{t,t+s} = \mathcal{M}_{t,t+s-1}\mathcal{M}_{t+s-1,t+s}$.

5.3 The Equity Premium

Consider the pricing equation derived earlier:

$$1 = E_t \left[\mathcal{M}_{t,t+1} \tilde{Z}_{i,t+1} \right]$$

Now, let's log-linearize in the case where utility is CRRA, assuming that consumption growth and asset returns are jointly lognormally distributed:

$$\begin{aligned}
1 &= \exp(-\rho) E_t \left[\exp(-\theta \Delta \ln c_{t+1}) \exp(\ln Z_{i,t+1}) \right] \\
1 &= \exp(-\rho) \exp\left(-\theta E_t \left[\Delta \ln c_{t+1}\right] + E_t \left[\ln Z_{i,t+1}\right] + 1/2 \mathrm{var}_t \left(-\theta \Delta \ln c_{t+1} + \ln Z_{i,t+1}\right)\right) \\
0 &= -\rho - \theta E_t \left[\Delta \ln c_{t+1}\right] + E_t \left[\ln Z_{i,t+1}\right] + 1/2 \mathrm{var}_t \left(-\theta \Delta \ln c_{t+1} + \ln Z_{i,t+1}\right)
\end{aligned}$$

where the last line takes logs. Now taking the difference between a risky and the riskess asset, we obtain:

$$\begin{aligned}
E_t \left[\ln Z_{i,t+1}\right] - r_{t+1} + 1/2 \mathrm{var}_t (Z_{i,t+1} - R_{t+1}) &= \theta \mathrm{cov}_t \left(\Delta \ln c_{t+1}, \ln Z_{i,t+1} - R_{t+1}\right) \\
E_t \left[\ln Z_{i,t+1}\right] + 1/2 \mathrm{var}_t (Z_{i,t+1} - R_{t+1}) - r_{t+1} &= \theta \mathrm{cov}_t \left(\Delta \ln c_{t+1}, \ln Z_{i,t+1} - R_{t+1}\right)
\end{aligned}$$

The last term on the left hand side in this equation is a Jensen's inequality term. The left hand side measures the equity premium. Mankiw and Zeldes estimate that it is about 6 percentage points. The standard deviation of consumption growth is 3.6 percentage points and the standard deviation of the excess return is 16.7 percentage points. The correlation of consumption growth and the excess return is 0.40. It follows that the right hand side is equal to $\theta \times 0.40 \times 3.6 \times 16.7$. To match the equity premium, the CRRA coefficient needs to be about 25.