Macroeconomics, Heterogeneity, and Inequality

An Introductory Course

[Preliminary and Incomplete]

Dylan Gowans

Contents

Introduction 3

Chapter 1. Simple Keynesian Models with Inequality 7

Chapter 2. The Permanent Income Hypothesis 17

Chapter 3. Empirical Evidence on MPCs 28

Chapter 4. Credit Constraints 42

Chapter 5. Models with Uncertainty 58

Chapter 6. The Distribution of Assets and Wealth Inequality 67

Chapter 7: An Introduction to Two-Asset Models 73

Introduction

Macroeconomics, Heterogeneity and Inequality

The goal of this course is to introduce students to the study of heterogeneity in macroeconomic models. In essence, it answers the question: when is it important to allow agents in models to differ, whether in income or in wealth?

This might seem like a strange question. In the real world there are household which differ in income and wealth and therefore, if we want our model to better approximate reality, surely it makes sense for our models to reflect this fact. However, until recently, it was common for macroeconomic models to be based on a single household — the representative consumer — whose consumption and savings decisions were to stand in for the many such decisions by heterogeneous households throughout the economy.

Perhaps, then, the decision about heterogeneity among agents is best reflected using the economist's favourite tool: the comparison of costs and benefits.

There is a cost to including heterogeneity in macroeconomic models. The more heterogeneity we have, the more variables we need to keep track of. Often it is already difficult enough to solve a model with a single representative household. The many non-linearities in the resulting system of equations means we need a computer to solve it. The more different agents we have, the more compute we need to solve the model.

If the cost is complexity, what are the benefits? Arguably there are two. First, an economy with many different kinds of households may behave differently than an economy populated by one "average" household. If, for example, an economy has very rich and very poor individuals, these individuals may behave differently in aggregate than one middle income consumer.

Second, we may be interested in heterogeneity itself. Perhaps we want to know economic fluctuations will impact poor households rather than households "on average". Or else we may want to know how a given policy impacts economic or wealth inequality. These kinds of questions will necessitate more complex models that include differences across households.

Arguably, over the past few decades, the balance of costs and benefits has shifted more towards benefits.

On the cost side, computing power continues to increase and become more accessible. Whereas some older economists will remember booking time in the university's "computer room" where they would have to bring their punch cards of computer code, today the same code can run much more quickly on a personal computer.

Meanwhile, both of the listed benefits have become more salient. In the first place, both empirical and theoretical research has suggested that taking heterogeneity into account is important when considering the impacts of macroeconomic policy — for example, the effects of government stimulus or of monetary policy. In general, we will see that poorer and low-wealth individuals will respond more strongly to these policies than others. Therefore, knowing the fractions of such consumers and how their behaviour differs can help us to design effective policies.

Second, especially after the global financial crisis of 2008-09, there has been increasing interest in the study of inequality. Many had the popular understanding after the crisis that the

burdens fell disproportionately on the poor and middle classes, while the upper classes — who some suggest are at fault for the crisis — were generally protected. The Occupy Wall Street movement made famous the distinction between the top 1% and the rest. And Thomas Piketty's 2013 *Capital in the 21st Century* showed that the top 1%'s shares of income and wealth in major rich-world countries had increased substantially since the 1970s.

These forces have converged to create a new class of heterogeneous agent models which have two-way causality between inequality and the macroeconomy. In other words, the macroeconomy impacts inequality, and the shape of inequality impacts how the macroeconomy behaves. We can represent this "mathematically" in the following way:

$$\text{Macroeconomy} \leftrightarrow \text{Inequality}$$

Structure of the Course

This course is designed for upper level undergraduate students, or as a primer for graduate students before undertaking more detailed study. As mentioned earlier, macroeconomic models with heterogeneity can be very complex, and often don't yield analytical solutions. Instead, they usually require a computer to solve. This has been a barrier to bringing undergraduate students up to the frontier of knowledge. The result is that an undergraduate student may get their degree with knowledge that is "out of date".[1] They risk entering the workforce with an incomplete understanding of current economic research, and arguably an incomplete understanding of the functioning of the economy. It will be harder for them to engage with modern papers or to understand some modern models.

This course is designed to alleviate this problem. The goal is to give students the intuition they need to understand frontier research, without going into the nitty-gritty details of computational economics needed to write and solve these models. This also sets up students who wish to continue into graduate studies. The hope is that, having already been introduced to the intuition of the models, students will find it easier to focus on the mathematics and computation necessary to further their studies of frontier models in graduate school.

Indeed, this method of education should not seem unfamiliar to students and teachers of economics. For example, the goal of first-year economics courses is to introduce students to the intuition of economic models often without going into the details mathematical details of optimization. The same types of models are then revisited in later years, but in greater mathematical detail. In this way, intuition precedes method.

This is not to say that this course lacks mathematical rigour. Rather, the course meets students where they are in terms of mathematical tools. Students should already be familiar with optimization theory (solving consumer and firm optimization problems) and perhaps have some understanding of representative agent macroeconomics — although this latter is not necessary, and it will be covered in Chapter 2. The course takes these tools and pushes them as far as they can go to intuitively understand models with heterogeneity. As we hope to show, this is not an unrealistic enterprise. The complexity of modern models often comes from many

[1] I put "out of date" in quotation marks because I know it is deliberately provocative. Not all models must have heterogeneity. As we will see, for some questions, a representative household may be sufficient. Moreover, these earlier models are still routinely used in governments, in central banks, and among academics. However, models with heterogeneity are becoming more common as they become more familiar to newly trained economists. It's not clear whether models with heterogeneity will become the dominant style of macroeconomic modelling, but it certainly has already become *a* dominant strand.

elements interacting at once: keeping track of distributions, uncertainty, non-linear constraints, many state variables, etc. Looking at these complications one at a time often helps us to understand how the models are working and to develop the intuition necessary to understand what happens when we include them all.

The course itself is structured in a historical way. That is, we start with macroeconomics as it exists in the post-WWII era, and then show the developments which got us to where we are now. We will explore the development of macroeconomics through three eras of macroeconomic modelling.

1. Post-WWII Keynesiansm
2. Representative Agent Models (i.e. “Modern Macro”)
3. Heterogeneous Agent Models

This serves several purposes.

First, it is often useful for students to understand how we got where we are. Often, this helps to explain the abstractions and assumptions that found our current models. For example, the movement from post-WWII Keynesianism to modern macroeconomics explains the importance of microfoundations. Moving from modern macro to models of heterogeneity explains the benefits of modelling heterogeneity as opposed to relying on a representative consumer. In each case, we can take stock of the empirical and theoretical developments which led to these movements.

Second, early models tend to be simpler, while later models tend to add further levels of complexity. Thus, starting with the simpler base often primes us to understand the more complex superstructure. For example, post-WWII Keynesian models make ad-hoc assumptions about consumption behaviour, while modern macro seeks to micro found these. Meanwhile, models of heterogeneity build on these earlier microfounded models by introducing non-linear constraints.



Figure 0.1 - The teleological development of macroeconomics

Third, although the last point makes it sound like the development of macroeconomics has been teleological — one development building on the last — there is some sense in which it has been circular. The conclusions of post-WWII Keynesianism were that marginal propensities to consume (MPCs) were important determinants of macroeconomic outcomes. From a policy perspective, this meant that government expenditure could be an effective means of boosting aggregate demand in the economy. The movement towards modern macro and the representative agent framework effectively removed these conclusions. In these models, MPCs are so low that they don’t matter, and government expenditures are essentially irrelevant to aggregate demand.

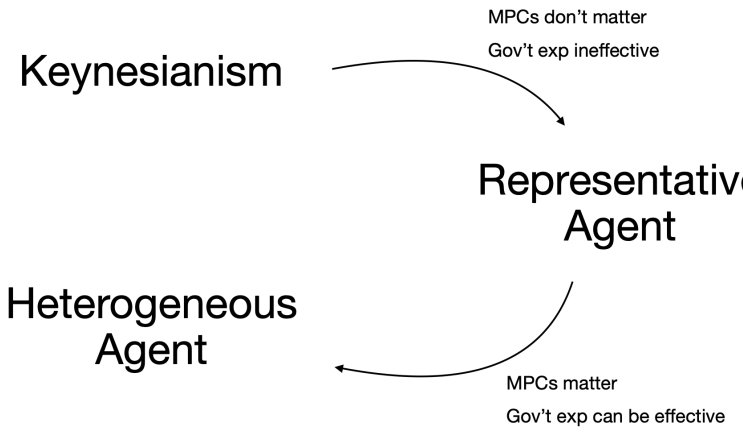

Figure 0.2 - The circular development of macroeconomics

The movement to heterogenous agent models brings us back towards the Keynesian world where MPCs are relevant variables again and a world in which government expenditures can be an effective means of dealing with economic downturns.

Taken together, we can see why it makes sense to start with early Keynesian models. This will give us a sense of where we're going, but at a lower level of complexity. Thus, it can provide intuition for the workings of heterogeneous agent models without students getting lost in the mathematics yet. It's also nice to have a sense of the policy relevance so we don't feel like we're doing math for math's sake!

Chapter 1. Simple Keynesian Models with Inequality

We start our discussion of heterogeneous agent models by going back to macroeconomics as it developed post-WWII. As mentioned in the introduction, this serves a dual purpose. First, it sets us up to understand the models which came in the 1970s and 80s which developed in reaction to these early Keynesian models. Second, it gives us a sense of where we are going in this course. A large part of this course is dedicated to understanding how marginal propensities to consume (MPCs) may differ across agents and how this will impact the macroeconomy. Seeing how this works first in a simplified model gives us an idea of what we're working towards. It helps us to understand the kinds of policy questions we are trying to answer with the more rigorous math developed later on.

A Basic Model

Keynes's *The General Theory of Employment, Interest, and Money* is considered to have marked the birth of macroeconomics. Its focus on the interactions of macroeconomic aggregates like money, employment and the interest rate defined the agenda for macroeconomists in the early post-war years.

Keynes was writing during the Great Depression. However, even prior to that, in the 1920s, his home country of Britain had been experiencing high unemployment. Thus, his research and writings were centred on the reasons and possible cures for low output and high unemployment.



Keynes (right) with his friend, the artist Duncan Grant

Keynes argued that the Great Depression was a function of insufficient aggregate demand. Employment was low because there was not enough demand to sustain it. The solution, therefore, was to find ways to increase demand. If, for instance, businesses chose to invest more, or the government spent more on infrastructure, this higher spending would increase the demand for goods and services. Businesses would thereby be incentivized to produce more and would hire more workers accordingly.

What's more, this first round of spending — say from the government building new bridges — leads to even more spending, as the builders of the bridge receive income which they in turn spend, which in turn creates more jobs and more spending, etc. This effect was referred to as the "multiplier effect".

The amount of additional expenditures created will depend on how much of their newly received income a person will spend. Imagine for example that workers employed in building a new bridge spent none of their new income. Then there would be no multiplier effect; the only impact on income would be the first round of spending coming from the government purchasing the bridge. Similarly, if workers spent 50% of their new income this would generate a smaller multiplier than if they spent 75%.

The math is relatively simple to work out. Imagine a simple closed economy. Since there is no trade, total income (or GDP), Y, is the sum of consumption, investment and government expenditures:

$$Y = C + G + I$$

Moreover, assume that a representative household receives all income from production, and consumes a fraction α of it and saves the rest. This parameter, α, is referred to as the marginal propensity to consume (MPC). Government expenditure is exogenous, depending only on the decisions of bureaucrats and politicians. Meanwhile, we assume that investment is a function only of the interest rate. This latter assumption allows for the central bank to have an impact on GDP through its effect on the interest rate. If the central bank lowers the interest rate, the cost of investment will go down, thus leading firms to invest.

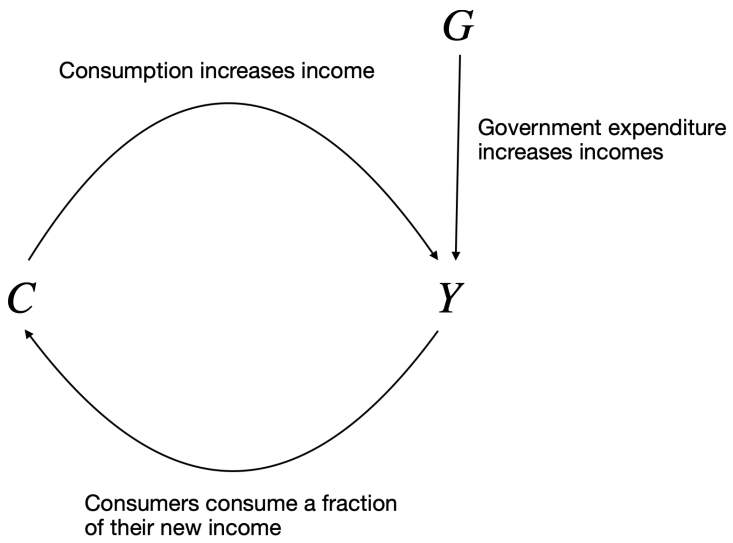

Figure 1.1 - The Multiplier Effect

Writing this all down mathematically:

$$C = \alpha Y$$

And plugging this into our GDP equation we have:

$$Y = \alpha Y + G + I(r)$$
$$Y = \frac{G + I(r)}{1 - \alpha}$$

Therefore we can see that, *ceteris paribus*, income will be higher the higher is α.[2] The intuition is simple. The higher is the MPC, the more expenditure consumption results from an increase in either G or I, and therefore the higher will be income.

Moreover, imagine that the government is fighting a recession and wants to expand income and employment. In this model, one way to do this is to expand its expenditures. The new government spending will lead to higher incomes, which in turn lead to higher consumption, etc. Note, therefore, that the effectiveness of this government expenditure program will depend on the marginal propensity to consume.

$$\frac{dY}{dG} = \frac{1}{1-\alpha}$$

Likewise, the effectiveness of the central bank's interest rate setting will depend on marginal propensity to consume.

$$\frac{dY}{dr} = \frac{I'(r)}{1-\alpha}$$

Therefore, a great deal depends on the marginal propensity to consume in this model!

Adding Heterogeneity - Income Inequality and MPCs

With so much riding on the marginal propensity to consume, some Keynesians began to consider what determines this important variable. One important factor seemed to be income. Keynes himself suggested that the rich tended to save a larger fraction of their income than the poor. And in the final chapter of *The General Theory*, he suggested that a more equal distribution of income would be conducive to economic output and economic growth:

> *Moreover, experience suggests that… measures for redistribution of incomes in a way likely to raise the propensity to consume may prove positively favourable to the growth of capital.*

Indeed, that the distribution of income was important for determining aggregate spending was not an uncommon view during the Great Depression. None other than the head of the Federal Reserve, Marriner Eccles, espoused similar views. Giving his views for the cause of the Great Depression in a talk in 1936, he attributes low spending power to the unequal distribution of income:

> *There was a maldistribution of the national income which resulted in getting consumer buying power out of relation to our productive capacity.*

Later, a whole class of Keynesian models developed among Keynes' students at Cambridge which tied economic output to the functional distribution of income. In these models, income received as wages was consumed with MPC=1, while income received as profits was

---

[2] α is subtracted from the denominator. Therefore, the higher is α, the lower will be the denominator, and therefore the higher will be Y.

completely saved, with MPC=0. The idea was that workers spent all, or nearly all, of their income, while capitalists saved all, or nearly all, of their income. We won't go into detail on these models here, but interested students can learn about them in the appendix to this chapter.

Interestingly, following the Great Recession of 2008-09, a similar idea has reemerged about the impact of the maldistribution of income. Economists Mian, Straub, and Sufi have a series of papers in which they suggest that, because the rich save more of their income, increasing income inequality has led to a decrease in aggregate demand. From the abstract of one of their papers:

> *Saving rates are significantly higher for high income households within a given birth cohort relative to other households in the same birth cohort, and there has been a large rise in income shares for high income households since the 1980s. The result has been a large rise in saving by high income earners since the 1980s, which is the exact same time period during which r* has fallen.*[3]

To explore these ideas, we develop a small extension to our model to include income inequality. Imagine now that households differ in their income. Here we do away with the representative household and introduce multiple households. In particular, assume now that income is split between a rich household and a poor household.

$$Y_i = \beta_i Y \quad \text{for } i \in \{P, R\}$$

Where

$$\beta_R + \beta_P = 1$$

Now, assume that the marginal propensity to consume of rich households is *lower* than for poor households. Poor households are concerned with meeting their basic needs, and less concerned with saving, while rich households already have much in the way of consumption goods, and are less pressed to consume more. Income now becomes:

$$\begin{aligned} Y &= C_R + C_P + I + G \\ Y &= \alpha_R Y_R + \alpha_P Y_P + I + G \\ Y &= \alpha_R (1 - \beta_P) Y + \alpha_P \beta_P Y + I + G \end{aligned}$$

Solving for Y:

$$Y = \frac{G + I}{1 - \left((1 - \beta_P)\alpha_R + \beta_P \alpha_P\right)}$$

$$Y = \frac{G + I}{1 - \bar{\alpha}(\beta_P)}$$

[3] r* is a measure of the interest rate that the central bank must set to maintain the economy at full employment. A decline in r* is a symptom of lower aggregate demand, all else equal.

Where $\bar{\alpha}$ is the average marginal propensity to consume across individuals, which obviously depends on the distribution of income. The greater is the share of income going to poor households, the higher will be the marginal propensity to consume, and therefore, all else equal, the higher will be income, and the greater will be the effectiveness of government policy.

Adding an Income Tax System

Up until now, we've assumed that government expenditures are deficit financed. Now let's add a tax system. Note that the different MPCs across the rich and the poor mean that the government could craft budget-neutral tax policy which is able to stimulate the economy. For example, imagine now that the government runs a progressive income tax system, and levies $t_R > t_p$ on the incomes of the rich and poor respectively. Consumption decisions are now made based on after-tax income. That is, consumers receive income, pay their taxes, and then allocate what is leftover between savings and consumption. First, let's define disposable income (i.e. after-tax income) for rich and poor households.

$$Y_R^d = (1 - t_R)Y_R$$
$$Y_P^d = (1 - t_P)Y_P$$

Now consumption becomes a function of disposable (after-tax) income:

$$C_R = \alpha_R Y_R^d$$
$$C_P = \alpha_P Y_P^d$$

Adding this to the GDP equation gives:

$$Y = \alpha_R(1 - t_R)(1 - \beta_P)Y + \alpha_P(1 - t_P)\beta_P Y + I + G$$

And solving gives:

$$Y = \frac{G + I}{1 - \Big(\alpha_R(1 - t_r)(1 - \beta_P) + \alpha_P(1 - t_P)\beta_P\Big)}$$

Note that the denominator is similar to the case without an income tax system. It is still the weighted average of the rich and poor marginal propensities to consume. The only difference is that it now takes into account the tax system.

Now imagine that the government wants to run a budget-neutral change in the income tax system. It wants to increase taxes on the rich while reducing them on the poor while keeping the overall budget the same. Since total tax revenues are:

$$T = t_R(1 - \beta_P)Y + t_P\beta_P Y$$

Then it is contemplating a change in tax rates which leave total tax revenue unchanged, or:

$$0 = \Delta T$$

$$0 = \Delta t_R(1-\beta_P)Y + \Delta t_P \beta_P Y$$

$$\Delta t_p = -\Delta t_R \frac{1-\beta_P}{\beta_P}$$

This last equation simply gives us the change in the tax rate on the poor relative to the change in the tax rate of the rich so that total tax revenues are unchanged.

For simplicity, call the term in brackets in the denominator of the GDP equation γ. Call this the effective marginal propensity to consume:

$$\gamma = \alpha_R(1-t_R)(1-\beta_P) + \alpha_P(1-t_P)\beta_P$$

How does the effective marginal propensity to consume change if we change income tax rates? The change in the marginal propensity to consume is given by:

$$\Delta\gamma = -\Delta t_R \alpha_R(1-\beta_P) - \Delta t_P \alpha_P \beta_P$$

Or plugging in our equation for Δt_P:

$$\begin{aligned}\Delta\gamma = & \quad -\Delta t_R \alpha_R(1-\beta_P) + \Delta t_R \frac{1-\beta_P}{\beta_P}\alpha_P\beta_P \\ = & \quad \Delta t_R(1-\beta_P)(\alpha_P - \alpha_R)\end{aligned}$$

Note that since the marginal propensity to consume of poor households is higher than rich households, we have that:

$$\alpha_P > \alpha_R \implies \Delta\gamma > 0$$

Again, this should be intuitive. The government is effectively redistributing income from households with a low marginal propensity to consume to those with a high marginal propensity to consume. The result will be a larger effective MPC, and therefore an increase in income.

What's "Missing" From These Models?

Keynesian models were developed during the Great Depression when aggregate demand was seen to be insufficient. Thus, the simplest models see the economy as demand-determined.[4] To increase output, we simply needed to increase demand. The supply-side was irrelevant, because demand was far below what the economy could reasonably provide.

[4] Note that this isn't true of all Keynesian models, or even all early Keynesian models, which quickly incorporated a supply side. Moreover, modern "New Keynesian" models incorporate all of the elements which we identify as "missing" below.

Thus, there are several things that we might think of as being "missing" from these models, and which we might want to add.

- **No "Production Function"**: The production side of the economy wasn't modelled. For example, we have not modelled capital inputs, nor does the production function have the usual decreasing returns to factor inputs.

- **No Labour or Capital Supply:** These models are of demand constrained economies. The implicit assumption is that there is an effectively infinite amount of factors of production willing to work at the prevailing rates of pay. The only constraint on output is demand. Therefore, if demand goes up, so does output.

- **Savings and Consumption Decisions are Based on Rules of Thumb:** The decision to save and consume is based on the MPC. But what determines the MPC? Could there be reasons that MPC will change over time? Could the MPC change because of policy?

- **Investment is Exogenous:** Similarly to the above, we have left investment exogenous. Note that this differs from models where investment is the result of optimizing behaviour of the firm based on prices and the productivity of capital.

- **What's the result of saving?** We don't model the asset acquisition of households or firms. This is related to the point about savings and investment being exogenous. As households and firms accumulate assets by savings and investing, shouldn't this change their future behaviour? For this we would require a dynamic model.

What Should We Take Away from These Models?

These are models of aggregate demand. Aggregate demand still plays a role in other models, but often it is tempered by other factors — for example, less than totally elastic supply.

However, the logic of these models will remain important. First, in Keynesian models such as these, MPCs are important determinants of the system. They tell us, for example, how effective a given government expenditure plan will be at fighting a recession. Take for example the case of "stimulus cheques" which were sent to U.S. citizens following the Great Recession. These will be effective at stimulating demand only if consumers spend a large fraction of them when they receive them. If, on the other hand, they have very low MPCs, the stimulus effort will be ineffective. We will return to this exact question starting in Chapter 3.

Second, we have seen how differences in MPCs across individuals can be important. If certain individuals (e.g. poorer households) have higher MPCs, then redistributing income to them can have positive aggregate demand impacts. Or, thinking again about government stimulus, efforts by government which increase the incomes of high MPC individuals will tend to me more effective at stimulating aggregate demand.

Next we will see how macroeconomics developed starting in the 1970s, and the ways this changed the assumptions and results of macroeconomic models. To anticipate a bit, we will see that economists like Milton Friedman began to suggest that MPCs were very low, and therefore were not an important determinant of macroeconomic activity.

Practice Problems

Appendix: Cambridge, Kalecki, and Worker-Capitalist Models

As mentioned above, although Keynes made passing reference to inequality and the macroeconomy, he never fleshed out the idea much. However, Keynes had several followers at Cambridge who, in subsequent years, would further work out models of aggregate demand which included inequality.[5]

We present here a simplified version of a model developed by Michal Kalecki (pronounced Kah-LET-skee). In this model we explicitly include two factors of production: labour and capital.[6] Workers receive wages while capitalists receive profits. Therefore, total income is given by:

$$Y = wN + \pi$$

Where N is the amount of labour, while π is nominal profits.

These worker-capitalist models employ what is perhaps an extreme assumption about savings: workers have a MPC of 1 — meaning they spend all their marginal income — while capitalists have a MPC of 0. Capitalists still consume, but their consumption is not based on (current) income. We assume no government sector. Therefore, from the expenditure side, nominal GDP is also given by:

$$Y = C_W + C_C + I$$

Since workers consume all of their income, $C_W = wN$. Substituting into the above:

$$Y = wN + C_C + I$$

Finally, capitalist consumption and investment are both exogenous in the short period. For simplicity, call both of these $a = C_C + I$.

$$Y = \underbrace{wN}_{\text{Workers Consumption}} + \underbrace{a}_{\text{Autonomous Expenditure}}$$

[5] Kalecki, whose model we will see shortly, may have worked out much of it at the same time but independently to Keynes! It's said that when Kalecki received a copy of *The General Theory* in 1936, he realized that it covered many of the same issues he had been working out on his own for the past few years. He fell ill as a result, staying in bed for several days.

[6] Note that in the previous model we didn't consider what factors of production were being used to earn income. In a sense it didn't matter since all that mattered was aggregate income, not its split between capital and labour.

We've considered GDP from the income and expenditure side. Finally let's consider it from the production side. Assume that workers produce z units of output per unit of labour. Therefore, nominal production is given by:

$$Y = pzN$$

Quickly, note that this allows us to figure out the profit and labour share of income. In particular, the labour share of income is given by:

$$\begin{aligned} LS &= \frac{wN}{pzN} \\ &= \frac{w}{pz} \end{aligned}$$

Note that each worker produces pz dollars of output, but receives w. Similarly, the capital share is given by:

$$\begin{aligned} CS &= \frac{pzN - wN}{pzN} \\ &= 1 - \frac{w}{pz} \end{aligned}$$

Substituting into the above:

$$pzN = wN + a$$

$$N = \frac{a}{pz - w}$$

Plugging back into the production equation:

$$\begin{aligned} Y &= pzN \\ Y &= pz\frac{a}{pz - w} \\ Y &= \frac{a}{1 - \frac{w}{pz}} \\ Y &= \frac{a}{1 - LS} \\ Y &= \frac{a}{CS} \end{aligned}$$

Note that this equation is very similar to the ones we've seen in past sections. The numerator is exogenous expenditures: in previous sections we had the exogenous expenditures $G + I$ in the numerator. Meanwhile, the denominator is related to the marginal propensity to consume. Because workers consume all of their marginal income while capitalists save all of their

marginal income, the effective marginal propensity to consume is just equal to the labour share of income. For each dollar that is earned, fraction LS is consumed because this is the fraction that goes to workers, while fraction CS is saved, because this is the fraction going to capitalists.

One implication of this is that output is responsive to the factor division of income. Thus, policies which increase the real wage rate will tend to increase output. Take for example an increase in the minimum wage. Some other models predict that output will decrease following an increase in the minimum wage. In these models, the higher wage makes labour unaffordable given its marginal productivity and therefore fewer workers are hired and output falls.

In this model, output and labour demand are only a function of aggregate demand. Firms hire workers if there is demand for them. So long as $w < pz$, the marginal product of labour is greater than its cost, and so it will be profitable to hire workers if there exists demand for their output. In these conditions, increases in the real wage — for example, brought about by an increase in the minimum wage — can increase output.

What You Should Know After this Chapter

Intuition

- **Marginal propensities to consume** are an important determinant of aggregate demand.
- **Heterogeneity in MPCs** means that the distribution of income can be an important determinant of macroeconomic outcomes.
- The value of MPCs and how they are distributed can help to determine the effectiveness of **government stimulus** policy.

Technical Skills

- Know how to derive aggregate demand given the MPC.
- Know how output and changes in output depend on the MPC.
- Know how to derive an “effective MPC” given the distribution of after-tax income.

Chapter 2. The Permanent Income Hypothesis

An interesting dichotomy developed in economic theory after the second world war. Macroeconomics developed on its own, along the lines of the Keynesian model. Meanwhile, what we now think of as microeconomics, having originated largely in the 1800s, continued to be developed along separate lines.

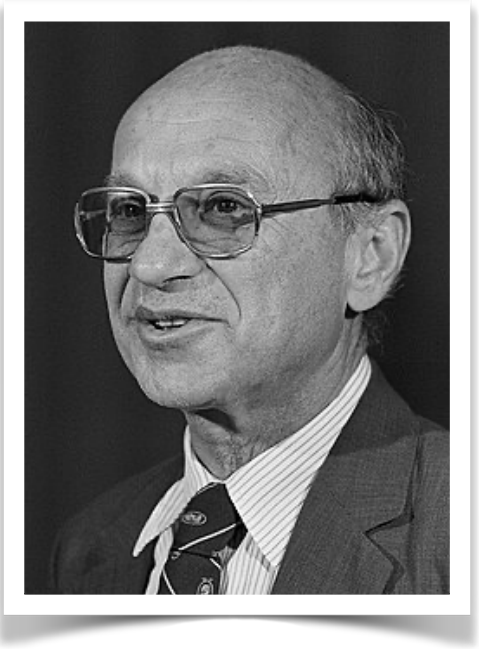

Milton Friedman

Microeconomics focused on models of consumer and firm optimization as the basis for behaviour. This seemed in contrast to some of the simpler Keynesian models, like the ones presented above, in which many decisions were exogenous or based on simple rules of thumb.

Thus, not surprisingly, some economists began to apply the lessons of microeconomics to macroeconomics. One such instance, which we will explore here, is the theory of consumption and savings, which has its origins in the works of Irving Fisher in the 1930s, but was further developed by Milton Friedman and Franco Modigliani in the 1950s.

What will be important to note is how the conclusions of this model differ substantially from the Keynesian model, in particular what they imply for MPCs and for heterogeneity. We will see that the basic Friedman model — called the Permanent Income Hypothesis — finds that MPCs should be very low and finds little reason they should differ across consumers.

Friedman's conclusions that MPCs were low also suggested that government stimulus — the favoured Keynesian policy for fighting the economic cycle — was relatively powerless and thus ill-advised.

Friedman's conclusions were matched his politics. Classical liberals like Friedman were uncomfortable with Keynesian economics which they saw as encouraging a large state with significant intervention in the economy. Meanwhile, Friedman supported limited government intervention. For example, in his classic 1962 book *Capitalism and Freedom*, he writes:

> *Government is necessary to preserve our freedom, it is an instrument through which we can exercise our freedom; yet by concentrating power in political hands, it is also a threat to freedom. Even though the men who wield this power initially be of good will and even though they be not corrupted by the power they exercise, the power will both attract and form men of a different stamp.*

Thus, Friedman's economic conclusions went in the same direction as his political views: governments should be smaller, and one power in particular that should be removed was the responsibility to fight recessions through government spending.[7]

[7] None of this is to suggest that Friedman was somehow intellectually dishonest. It is often the case that our economic views align with our political views. Both reflect common sets of beliefs about people, society and values.

Related to Friedman's work on the consumption function was the movement to introduce microfoundations in macroeconomic models. In a sense, this was a way to relieve the tension between a microeconomics that was developing separately to macroeconomics. As we saw, Keynesian macroeconomic models tended assume simple parameter relationships between variables. A great example is the Keynesian consumption function, which saw consumption as being linearly related to income.

In the 1970s, Robert Lucas formulated what is known as the Lucas Critique. His point was that it was foolish to assume that these "structural parameters" would be unchanging in response to policy changes. For example, there was no reason to assume that the MPC parameter, α, in the Keynesian consumption function should stay the same across time or in response to other changes in the economy. A better modelling methodology would be to make these parameters a function of consumer (or firm) optimization. In this way, the parameter will not necessarily be a constant, but will respond in different ways depending on what else is happening in the economy.

To give an example, and to anticipate a little what comes next, imagine how a consumer might respond differently to being given $100 today with no strings attached to a case where they are given $100 today, but know that they will have to pay an equivalent amount of taxes back in the future. In the first case, we might expect their consumption response to be higher than in the second. In the second, it's possible the consumer will want to save some of their windfall earnings to be able to pay back the necessary taxes in the future. Note how this differs from the Keynesian model, in which receiving $100 today should have similar effects on consumption, regardless of the policy context.

Thus, Friedman's model is an example of early microfoundations. Instead of taking the consumption function as being given by certain unchanging parameters, it models these as the result of consumer optimization.

The Consumer's Problem

In this section, we develop the consumer's problem along the lines formulated by Friedman in his 1957 book *A Theory of the Consumption Function*.

Unlike the Keynesian consumption function, we will here develop a model of consumption and savings which is based on consumer optimization.[8]

Imagine a household who wants to maximize utility from now into the future. For simplicity, utility will be a function of consumption only. The household receives a known stream of income from now into the future. This last assumption can be easily modified to make income uncertain, but with a predictable distribution. We will return to the case with uncertainty later on.

[8] There are some that would argue that, although it is not based on consumer optimization, the Keynesian consumption function is a better description of reality. These economists say that consumers are constrained in their ability to fully 'optimize' their consumption-savings decisions as the problem is too complex and uncertain. Instead, consumers rely on simple rules of thumb to make their consumption and savings decisions. In fact, once you start to save for retirement, you will find advice replete with such rules of thumb, e.g. "Save 20% of your net income for retirement".

The household can save and borrow at known rate of interest r.[9]

Thus the household's problem is:

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

$$\text{s.t.} \quad c_t = (1+r)s_{t-1} + y_t - s_t$$

Admittedly we don't want to solve for two choice variables (c_t, s_t). Luckily we can simplify the problem.

Start with period 0, in which, we assume that $s_{-1} = 0$. Then the budget constraint can be written as:

$$s_0 = y_0 - c_0$$

Plugging this into the next period budget constraint gives:

$$s_1 = y_1 - c_1 + (y_0 - c_0)(1+r)$$

Iterating again…

$$s_2 = y_2 - c_2 + (y_1 - c_1)(1+r) + (y_0 - c_0)(1+r)^2$$

In general as we iterate we have:

$$s_T = \sum_{t=0}^{T} (y_t - c_t)(1+r)^{T-t}$$

Or dividing through by $(1+r)^T$:

$$\frac{s_T}{(1+r)^T} = \sum_{t=0}^{T} \frac{(y_t - c_t)}{(1+r)^t}$$

We can keep iterating again and again. Eventually, in the "final" period, the consumer has no reason to save, and therefore:

$$s_\infty = 0$$

So the budget constraint will become:

---

[9] Again, this could be modified to make r dependent on time period or to be uncertain.

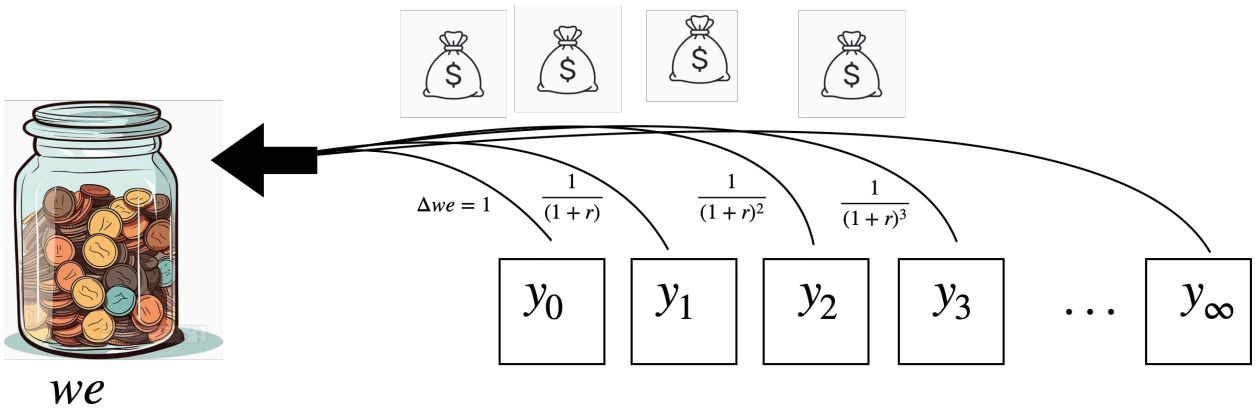

Lifetime wealth is as if the consumer has borrowed against all of their income…

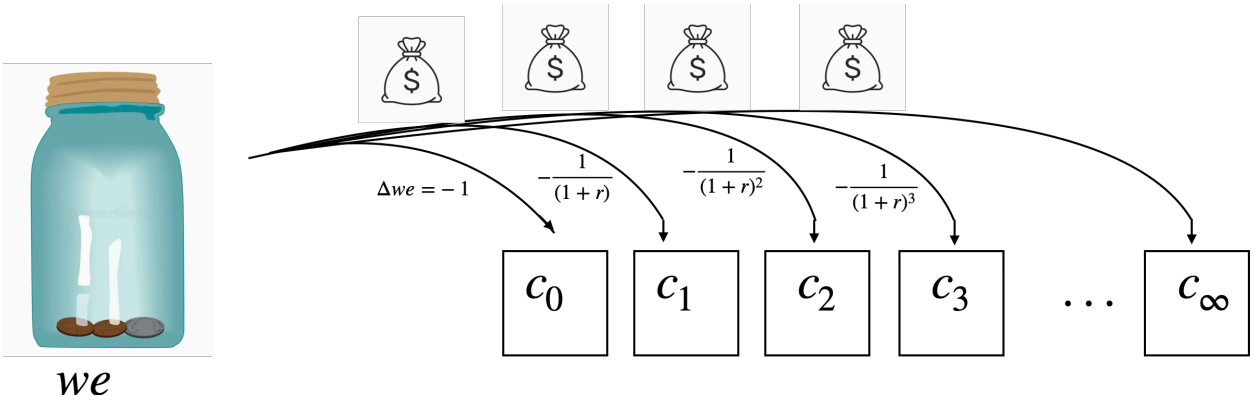

…which is then allocated to consumption across all periods based on the cost of consumption in terms of we

$$0 = \sum_{t=0}^{\infty} \frac{y_t - c_t}{(1+r)^t}$$

Or rearranging:

$$\sum_{t=0}^{\infty} \frac{c_t}{(1+r)^t} = \sum_{t=0}^{\infty} \frac{y_t}{(1+r)^t}$$

First, note that this is useful to use because it allows us to eliminate a choice variable from the problem, therefore making it easier to solve.

Secondly, inspecting this equation, you will see that it has a reasonable interpretation. The right-hand side of the equation is the present value of lifetime income. Since the household is able to borrow against future income, this is the amount of money they would have today if they borrowed the maximum possible. The household could then take that money and split it into consumption across all lifetime periods

The left-hand side of the equation then is the consumer's consumption choice. Note that the price of consumption in period t is $1/(1+r)^t$. Why? The price is how much income the household must sacrifice today to get 1 unit in the future. In other words. To get 1 more unit t periods from now means that the consumer is saving for t periods and accumulating interest. Therefore:

$$\Delta y_0(1+r)^t = \Delta c_t$$

$$\frac{\Delta y_0}{\Delta c_t} = \frac{1}{(1+r)^t}$$

Which shows the price of future consumption in terms of current income. This is represented in the next two diagrams.

Therefore, the consumer's problem can be written more succinctly as:

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

$$\text{s.t.} \quad \sum_{t=0}^{\infty} \frac{c_t}{(1+r)^t} = \sum_{t=0}^{\infty} \frac{y_t}{(1+r)^t} = we$$

Consumer Optimization Solution

For simplicity, let's choose a particular functional form for per-period utility:

$$u(c_t) = \ln c_t$$

Which makes the consumer's problem:

$$\max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \ln c_t$$

$$\text{s.t.} \quad \sum_{t=0}^{\infty} \frac{c_t}{(1+r)^t} = \sum_{t=0}^{\infty} \frac{y_t}{(1+r)^t} = we$$

This is an example of intertemporal optimization. Intertemporal means across time since the consumer is maximizing their consumption across periods. The first order condition for the intertemporal optimization problem is known as an Euler equation, which shows how consumption changes over time.

The solution takes the usual form in which the "bang-for-the-buck" is equalized across the consumer's choice variables. In other words, marginal utility per dollar spent must be the same across all choice variables, otherwise the consumers could be better off by reallocating. In our case, the choice variables are consumption in each period, and the "prices" represent the

amount of lifetime wealth a consumer must give up to increase consumption in a given period. In general, the solution is of the form:

$$MU(c_t)/\tilde{p}_t = MU(c_{t+1})/\tilde{p}_{t+1}$$

Where we use $\tilde{p}_t$ to represent the price of consumption in period t in terms of lifetime wealth, we. As we've seen, this price is $1/(1+r)^t$, meaning that it costs less to "buy" consumption in a later period, since we only need to set aside a portion of what is needed today and allow it to grow with the rate of interest. Substituting this:

$$\beta^t u'(c_t)(1+r)^t = \beta^{t+1} u'(c_{t+1})(1+r)^{t+1}$$

$$\frac{c_{t+1}}{c_t} = \beta(1+r)$$

This is the consumer's Euler equation, which dictates how consumption grows over time. Note that we can use it to express the relationship of consumption for any period with respect to c_0.

$$\begin{aligned} c_1 &= \beta(1+r)c_0 \\ c_2 &= \beta^2(1+r)^2 c_0 \\ c_3 &= \beta^3(1+r)^3 c_0 \\ &\dots \\ c_t &= \beta^t(1+r)^t c_0 \end{aligned}$$

Plugging this into the budget constraint gives:

$$\sum_{t=0}^{\infty} \beta^t c_0 = we$$

$$c_0 \sum_{t=0}^{\infty} \beta^t = we$$

$$c_0 \frac{1}{1-\beta} = we$$

Which allows us to solve for c_0 in terms of exogenous parameters:

$$c_0 = we(1-\beta)$$

Low MPCs

What is the consumer's MPC in this model? Recall that the MPC is defined as the change in current consumption resulting from a change in current income. Thus:

$$\begin{aligned} \frac{dc_0}{dy_0} &= (1-\beta)\,\frac{dwe}{dy_0} \\ &= (1-\beta)\,(1) \\ &= (1-\beta) \end{aligned}$$

The MPC will depend on the value of β, which in macroeconomic models usually takes an annual value of around .97. Note that this implies an MPC of only 3%!

Why is this the case? Milton Friedman formulated this as the "permanent income hypothesis". Friedman's point was that consumption choices are based on lifetime income, not on current income. The consumer will optimally smooth consumption across periods. Thus any increase in income today will be shared among all future periods. Given that the lifespan of a consumer is long, this means saving most of it to be able to spread it out into the future. The result is that consumers save most windfall gains, and have a very low MPC.

With such a low MPC, Keynesian policies — say increases in government expenditures — will have little additional impacts on overall activity. Put another way, recall that the multiplier is equal to:

$$\begin{aligned} multiplier &= \frac{1}{1-\alpha} \\ &= \frac{1}{1-0.03} \\ &= 1.03 \end{aligned}$$

In other words, this says that, *ceteris paribus*, an increase in government expenditure will only increase GDP by slightly more than the initial impact. Moreover, low MPCs mean that other forms of government stimulus, for example sending consumers government stimulus cheques, will be very ineffective at increasing aggregate demand.

Unimportance of Heterogeneity

Note also that MPCs depend only on β. There is no sense in which they depend on income. Imagine two consumers, one rich and the other poor. Total consumption is:

$$C_R + C_P = we_R(1-\beta) + we_P(1-\beta)$$

$$C_R + C_P = (1-\beta)we$$

$$C = (1-\beta)we$$

Where we is aggregate lifetime income, and C is aggregate consumption.

One implication of this is that it is not important to model heterogeneity. It doesn't matter how income (and lifetime income) is distributed across households, all that matters is the aggregate! Thus, it will be far simpler to model the economy as being populated by one representative household.

One important implication of this is that the distribution of income has no impact on aggregate consumption. For example, what happens if we redistribute income from the rich to the poor household? If total lifetime income is unaffected, this should have no impact on current consumption. Mathematically:

$$\Delta C = (1-\beta)\Delta we$$

$$\Delta C = (1-\beta)(\Delta we_P + \Delta we_R)$$

$$\Delta C = (1-\beta)(\Delta y_{0,P} + \Delta y_{0,R})$$

And if we are simply redistributing income, we have that any income gained by the poor is lost by the rich, therefore:

$$\Delta y_{0,P} + \Delta y_{0,R} = 0$$

And thus,

$$\Delta C = 0$$

Changes in Future Income are Important Too

Recall that in the basic Keynesian model, current consumption depends only on current income. Note that the permanent income hypothesis, on the other hand, says that current consumption should depend on current income *and* future income. This is evident mathematically since consumption depends on lifetime income, we, and this is a function of all future income.

This also has important implications for government policy. Consider the case of a government running a deficit to give stimulus cheques to its citizens. Recall that in the basic Keynesian model, this would mean that consumption would increase by the size of the MPC:

$$\Delta C_0 = \alpha M$$

Where M is the size of the stimulus cheque. In the basic Keynesian case, all that matter is the current change in income. Future government policy, for example how it will pay for its current deficit, doesn't factor in at all to the consumer's decision.

Now imagine how a permanent income hypothesis consumer would behave. The change in consumption will instead be a function of the consumer's lifetime change in income. In other words:

$$\Delta c_0 = (1-\beta)\Delta we$$

$$\Delta c_0 = (1-\beta)\sum_{t=0}^{\infty}\frac{\Delta y_t}{(1+r)^t}$$

Since present consumption depends on lifetime income, it matters how the deficit will be financed. For example, imagine that the government will borrow today to finance the stimulus cheques. But to pay back the deficit, it will raise taxes in the following period.

In other words:

$$T_1 = M(1+r)$$

Note that taxes in period 1, T_1, are larger than the stimulus package given that they also must account for the interest incurred by borrowing.

For the consumer, this means we must keep track of their change in consumption in period 0 and period 1. In other words:

$$\Delta c_0 = (1-\beta)\left(\Delta y_0 + \frac{\Delta y_1}{(1+r)}\right)$$

The change in (after-tax) income in period 0 is the stimulus payment, M, while in period 1 is the taxes the consumer must pay, T_1. Therefore:

$$\Delta c_0 = (1-\beta)\left(M - \frac{T}{1+r}\right)$$

$$\Delta c_0 = (1-\beta)\left(M - \frac{M(1+r)}{1+r}\right)$$

$$\Delta c_0 = (1-\beta)(0)$$

$$\Delta c_0 = 0$$

Thus, the consumer will not increase consumption at all in response to the stimulus cheques! The reason is that they know that their future taxes will increase in such a way that lifetime income stays unchanged. Since consumption is based on lifetime income, and not current income, they will not change their current consumption. This result is known as Ricardian Equivalence. Although named after 19th century economist David Ricardo, it was developed in its current form by Robert J. Barro in the 1970s.

Practice Problems:

1. Imagine a rich consumer and a poor consumer. Both are infinitely lived with per period utility given by $u(c_t) = \ln c_t$. Both have a known stream of income, y_t. The only difference between them is that the poor consumer has a lower discount factor. In particular:

$\beta_R = 0.97$
$\beta_P = 0.93$

(a) In words, what does a lower β mean about a consumer's preferences?
(b) Derive the MPCs for both consumers. Is their relationship to one another what the Keynesian model in chapter 1 would predict?
(c) Imagine now that the lifetime wealths of the two consumers are $we_R = 500, \ we_P = 100$. Imagine that the government taxes the rich consumer $10 and redistributes it to the poor consumer. By how much will consumption demand change in period 0? By how much will it change in period 1?

2. Imagine consumers who lives for two periods. Utility is given by:

$$U = c_0^\alpha + \beta c_1^\alpha, \qquad \alpha \in (0,1)$$

Income is known by the consumer in both periods. The consumer can borrow and lend at interest rate r. Income grows over the consumer's life. In particular, we have that

$$y_0 = y_L < y_1 = y_H$$

For simplicity, assume that $\beta(1 + r) = 1$.

(a) Show that in equilibrium $c_0 = c_1$.
(b) What is the marginal propensity to consume out of current income?
(c) Now imagine that the government begins to tax the rich to redistribute to the poor. In particular, it withdraws T dollars in taxes from a rich consumer, and transfers the same T dollars to a poor consumer. What will be the poor consumer's marginal propensity to consume out of the transfer of T dollars?
(d) Explain how and why the results in (b) and (c) differ.
(e) What does this suggest about redistributive policy as a means of increasing aggregate demand? How does it differ from what we saw in the Keynesian model?

3. Imagine an infinitely lived consumer with per-period utility given by:

$$u(c_t) = \ln c_t$$

The consumer discounts future utility at rate β and can borrow and save at rate $(1 + r)$.

(a) Derive the consumer's Euler equation.
(b) Derive an expression for current period consumption in terms of lifetime wealth.

Now imagine that the government issues stimulus cheques in period 0 worth M. To finance this, the government will levy taxes in period 1 and period 2, i.e. $T_1, \ T_2$.

(c) What is the size of the deficit in period 0?
(d) What is the size of the deficit in period 1, after the government levies T_1?
(e) Therefore, given T_1, what is the value of T_2 that will eliminate the deficit in period 2?

(f) What is Ricardian Equivalence? Use the values of T_1 and T_2 that you derived along with the consumer's budget constraint to show that it holds in this case.

What You Should Know After this Chapter

Intuition

- For consumers in the Permanent Income Hypothesis:
    - **MPCs are low** because of consumption smoothing
    - **MPCs don't vary much across individuals** since they depend only on smoothing parameters and permanent income
    - **Consumption depends on current and future income**

- Friedman's hypothesis suggests that government expenditure should have little impact on aggregate demand. This is in contrast to early Keynesian models.

Technical Skills

- Know how to derive and solve the consumer's problem
- Be able to show the MPC given the solution to the consumer's problem

Chapter 3. Empirical Evidence on MPCs

Friedman's Permanent Income Hypothesis makes very different hypotheses about MPCs than do the early Keynesian models. In this section we look first at the empirical evidence that Friedman used to support his hypothesis, and then at more recent evidence which seems to contradict it.

Early Evidence on MPCs

In the early post-war era, a typical method for estimating MPCs was to look at consumption across households and compare this to their income. Things were more complicated than this — for example, how do you define income? How do you define consumption? What time period should you look at? — but you can think of this as the basis of empirical studies. The results might look something like in Figure 3.1:

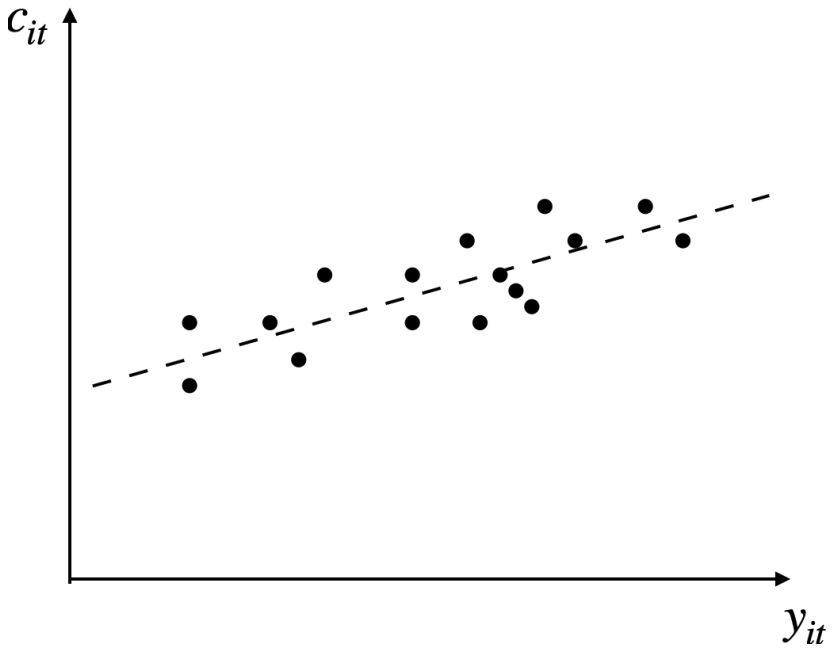

Figure 3.1 - A typical post-war estimation of the MPC

The x-axis shows incomes for many households, possibly over multiple periods, while the y-axis shows consumption. The slope of the regression line gives the estimated MPC. Studies tended to find shapes like this with a positive y-intercept and an MPC<1, suggesting that average propensities to consume (i.e. c/y) were falling as income increases. In other words, on average, the rich saved a greater share of their income than the poor.[10]

[10] For example, in our sample figure the poorest consumers have positive consumption but near 0 income, suggesting that their average propensity to consume is greater than 1 — i.e. they are consuming more than their incomes. Conversely, because the slope of the regression line is below one, the very rich are likely consuming less than their income, suggesting an average propensity to consume that is less than 1.

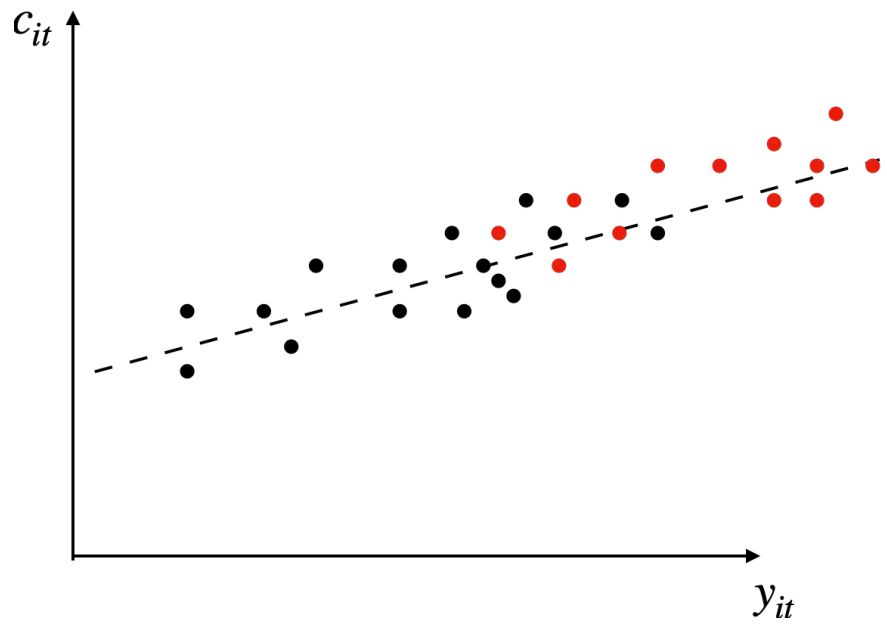

Figure 3.2 - What we might expect to see as income grows

Given that the rich save more on average, does this mean that total savings rates will increase as the economy grows? If we interpret y as real income, this seems to be the implication of the figure above. Imagine for example that we do the same exercise again, but after 10 years during which average incomes have grown. We might expect to see something like Figure 3.2. As the country has grown richer, we've simply moved along the MPC curve. In this case, note that average average consumption to income ratios have decreased, and average rates of saving have increased.

However, this is not the case. Famously, Simon Kuznets found that, in the United States, income growth was *not* associated with an increasing savings rate. Rather, the savings rate was remarkably constant over time! This became known as the Kuznets paradox given that it conflicted with the simple Keynesian story.

Effectively, a consumer making $1,000 in 1910 was different than a consumer making $1,000 in 1920, even when that income is adjusted for inflation. Such a consumer was just as rich in both periods in their ability to buy goods and services, but in 1910 they might have been relatively rich compared to their fellow citizens, while in 1920 they might have been relatively poor. Importantly, the consumer in 1910 would have had a higher savings rate than the consumer in 1920. This is shown in Figure 3.3.

Friedman's Critique Applied

Several hypotheses were advanced to explain the Kuznets paradox. For example, the Relative Income Hypothesis, advanced by James Duesenberry, suggested that consumption rates were determined by relative income rather than absolute income. The poor felt pressure to "keep up with the Joneses" — i.e. to consume similar baskets to their richer counterparts — thus leading them to have a relatively low savings rate. Meanwhile, no such pressure existed on the rich, who therefore had a high savings rate. As total income grows, the same pressures exist,

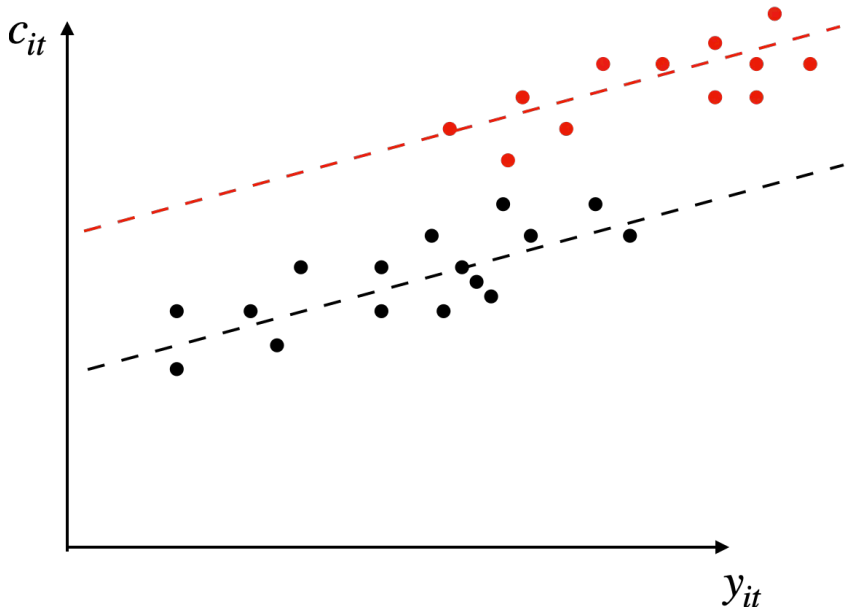

Figure 3.3 - What we actually saw as income grew

maintaining the low savings rate of the relatively poor and the high savings rate of the relatively rich.

Friedman framed his own explanation in terms of the theory we saw in the last chapter. First, consider how Figure 3.1 seems to conflict with Friedman's hypothesis. Imagine for example a case where $\beta = (1 + r)^{-1}$, which implies perfect consumption smoothing across time (i.e. $c_t = c_{t+1} = c_{t+2} = \ldots$)[11] According to the model of the consumer:

1. **Changes in consumption won't respond at all to expected changes in income.** Expected changes in income don't change lifetime wealth, and therefore won't change consumption over time. Instead, consumption is perfectly smoothed and doesn't change from period to period.
2. **Changes in consumption will respond very little to unexpected temporary changes in income.** This is the result we saw in Chapter 2. A windfall increase in today's income will be shared across today and all future periods, meaning the increase in today's consumption will be very small.
3. **Changes in consumption will respond 1-to-1 to unexpected permanent changes in income.** Imagine the consumer has an unexpected income shock that increases y_t by 20% in every period. Then lifetime wealth will increase by 20%. Since consumption is proportional to lifetime wealth, consumption today (and in all future periods) will also increase by 20%.

Points (1) and (2) suggest we should find an MPC of around 0, while point (3) suggests an MPC of 1 out of permanent increases in income. How does this translate into what we see in Figure 3.1?

[11] The conclusions that follow don't depend on $\beta = (1 + r)^{-1}$. We are simply using it here for simplicity.

Imagine two kinds of individuals. One kind has a high lifetime wealth while the other has a low lifetime wealth. Both have income which fluctuates over time, however. The Permanent Income Hypothesis would predict that their consumption and income paths will look something like Figure 3.4.

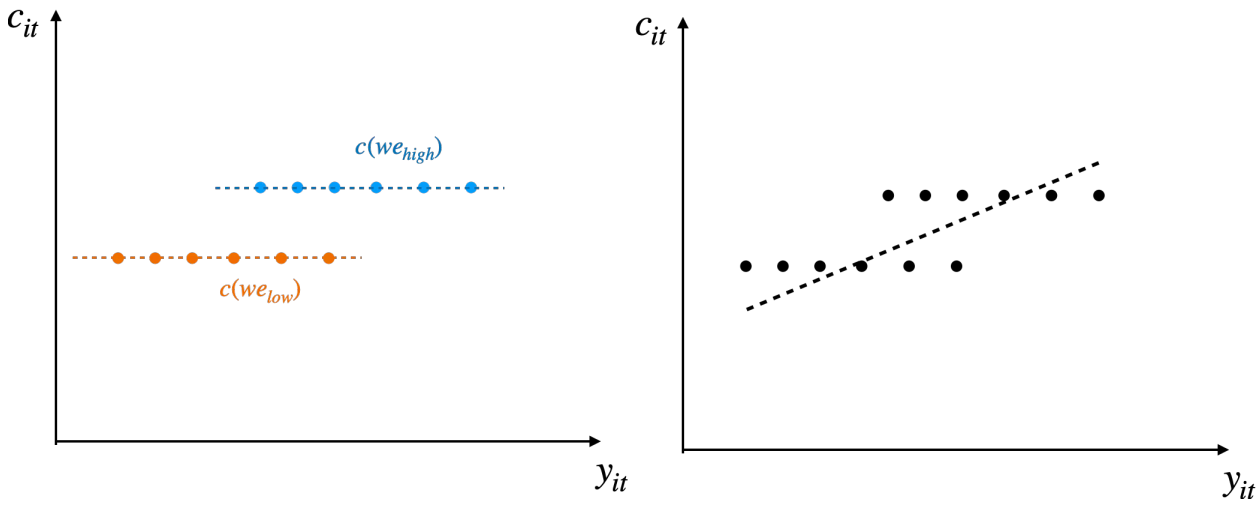

Figure 3.4 - Friedmanian consumers on their own and in aggregate

On the left hand side of the figure we have the income and consumption paths of the high lifetime wealth and low lifetime wealth consumers. Both have income which fluctuates over time, however consumption does not change with these predictable changes in income as consumption is based on lifetime wealth, not present income. The consumption of the high lifetime wealth individual is higher since, on average, their income is higher and therefore their lifetime wealth is higher.

The right hand side shows what would happen if we aggregated the data points from both types of individuals and ran a regression of consumption on present income. We get the usual Keynesian form: a positive y-intercept and a positive MPC with slope of less than 1.

Note, however, that the right hand panel is misleading as a measure of MPCs. If we were to give each type of consumer a transfer in this period, their consumption would not increase by the measured MPC.[12]

How does Friedman's theory hold up to the Kuznets paradox? Imagine we perform the same regression 10 years apart after incomes have grown. Now both the low and high lifetime wealth consumers have higher average incomes over their lifetime and therefore their mass of points has shifted up. The higher lifetime wealth also means higher consumption, and therefore the points have also shifted to the right. Drawing our regression lines leads to what we see empirically: it appears as though the the MPC is more or less constant, but our y-intercept has shifted up, suggesting that average savings are more or less the same. This is all shown in Figure 3.5.

Note that this is also intuitive. Average savings rates won't increase with real income growth because savings are determined by the need to smooth income over time, not in any way on

[12] To put this in econometric terms, what we really need is fixed effects for consumers. However, this will then require reliable panel data with sufficient years to get a good estimate of the fixed effects.

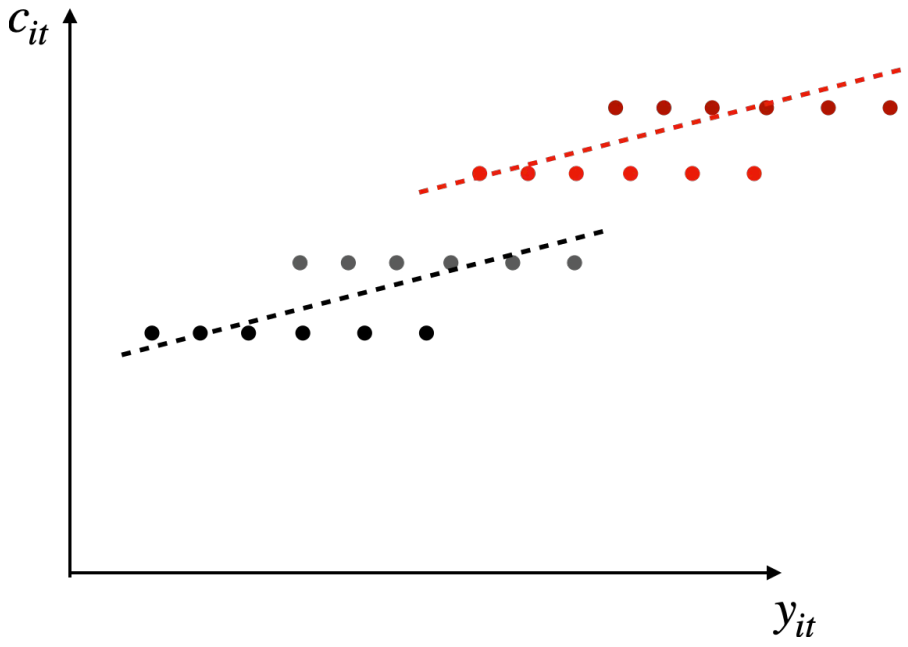

Figure 3.5 - Income growth over time with Friedmanian consumers

the absolute value of income. In a sense, the intensity of the need to save doesn't fluctuate with average real income. As intuition, consider the following mathematical example. Imagine an individual who lives for only two periods, with predictable income in both periods, with log per-period utility, and for whom $\beta = (1+r)^{-1}$. Then the individual's Euler equation is:

$$\frac{c_2}{c_1} = 1 \quad \Longrightarrow \quad c_2 = c_1 = c$$

Plugging into their budget constraint and solving we get:

$$c = \frac{1+r}{2+r}\left(y_1 + \frac{y_2}{1+r}\right)$$

We can therefore calculate savings in each period.

$$s_1 = y_1 - c_1$$
$$s_1 = y_1 - \left(\frac{1+r}{2+r}y_1 + \frac{1}{2+r}y_2\right)$$
$$s_1 = \frac{y_1 - y_2}{2+r}$$

And therefore the savings rate is:

$$\frac{s_1}{y_1} = \frac{1}{2+r}(1 - g_y)$$

Where $g_y = y_2/y_1$ is the gross growth rate of the individual's income across the two periods.

We can do the same for second period savings and we get:

$$\frac{s_2}{y_2} = \frac{1+r}{2+r}\left(1 - \frac{1}{g_y}\right)$$

Note that in neither case does the savings rate depend on the *absolute* value of income. Rather, the savings rate is a function of the dispersion of income across time. As long as the relative dispersion doesn't change with increases in real income, then savings rates will stay more or less the same.

Difficulties of Testing the Friedman Hypothesis

Friedman showed that his model fit with the facts of the day: it was able to generate apparent aggregate MPCs between 0 and 1, and could explain the Kuznets paradox. However, he did not subject many of the theories implications to rigorous testing.

Several empirical studies in the 1980s began to test Friedman's hypothesis seriously. Note that this is not an easy task. Recall the theory's main implications:

1. Changes in consumption won't respond to expected changes in income.
2. Changes in consumption will be small following unexpected transitory changes in income.
3. Changes in consumption will be large following unexpected permanent changes in income.

Therefore there are four possible kinds of change in income: expected-transitory, expected-permanent, unexpected-transitory, and unexpected-permanent. Depending on which kind of change, the model predicts different kinds of changes in consumption.

Table 1 - Changes in Consumption Resulting from Different Changes in Income

	Transitory	**Permanent**
Expected	No effect	No effect
Unexpected	Small	Large

Thus, it becomes important to be able to distinguish between different kinds of income changes. To disprove the model, it's not good enough to simply show, for example, that a change in income leads to a large impact on consumption — if that change in income was an unexpected permanent change, then that fits with the model. However, if it was an expected or transitory change, that suggests that the model is not a good description of real consumer behaviour.

Tests Based on the Euler Equation

Early tests of the permanent income hypothesis attacked the first row of Table 1. Friedman suggested that expected changes in income should have no effect on consumption. Therefore, if we can find cases where expected changes in income *do* seem to impact consumption, then this would be a violation of the hypothesis.

Starting with Hall (1978), the tests relied on the the Euler equation. Because the theory takes into account both expected and unexpected changes in income, we have to reformulate the model to take into account uncertainty.

The Consumer's Problem With Uncertainty

$$\max_{\{c_t\}_{t=0}^{\infty}} u(c_0) + \sum_{t=1}^{\infty} \beta^t \mathbb{E}\Big(u(c_t)\Big)$$

$$\text{s.t.} \quad c_0 + \sum_{t=1}^{\infty} \frac{\mathbb{E}(c_t)}{(1+r)^t} = y_0 + \sum_{t=1}^{\infty} \frac{\mathbb{E}(y_t)}{(1+r)^t}$$

The equation looks a bit more complex, but the logic is still the same. Now given their *expected* lifetime wealth, the consumer makes plans for consumption from now into the future. We've separated out current period values from the sums and expectations to be clear that y_0 and c_0 are known in the current period. However, because future income is uncertain, so is future consumption.

Despite this, the problem has a similar solution. The first order condition — the Euler equation — is given by:

$$\frac{\mathbb{E}\Big(u'(c_1)\Big)}{u'(c_0)} = \beta(1+r)$$

The Consumer Equates Expected Consumption (Given Assumptions)

If we make the usual assumption that $\beta = (1+r)^{-1}$, this becomes:

$$u'(c_0) = \mathbb{E}\Big(u'(c_1)\Big)$$

Which says that the consumer will allocate consumption between today and tomorrow so that the *expected* marginal utilities will be equal.

It also simplifies the problem if we assume that marginal utility is linear. For example, if we assume that

$$u(c_t) = Ac_t - c_t^2/2$$

Then the Euler equation simplifies even further:

$$\mathbb{E}\Big(A - c_1\Big) = A - c_0$$

$$\mathbb{E}(c_1) = c_0$$

If marginal utility is linear, then the Euler equation says that the consumer will equate expected consumption across periods.

Consumption Should Behave like a Random Walk

Substituting the Euler equation into the budget constraint, we get that consumption will be some fraction of expected lifetime wealth:

$$c_0 = \mathbb{E}_0(c_1) = \gamma \mathbb{E}_0(we)$$

Each period, the consumer's expectation of lifetime wealth will change based on unexpected shocks to income and to information about expected future income. Therefore, expected wealth will evolve like a random walk:

$$\mathbb{E}_1(we) = \mathbb{E}_0(we) + \varepsilon_1$$

Since consumption in any period is a fraction of lifetime wealth, this means that consumption will also evolve like a random walk:

$$c_1 = c_0 + \epsilon_1$$

In any period, the consumer will receive new information about lifetime wealth, and therefore will update their expectation for current and future consumption. If this information is completely random and uncorrelated, then consumption will follow a random walk.

Or put differently:

$$\Delta c_t = \epsilon_t$$

The Growth of Consumption Should not Depend on Current Variables

Given our assumptions, in the no-uncertainty case, consumption wouldn't grow at all. The assumptions we've made lead the consumer to perfectly smooth across time. When there is uncertainty, each period new information is revealed that leads the consumer to update their expectation of lifetime wealth, and therefore to update current consumption.

Note that this means that information available at t should not influence the change in consumption at $t + 1$. The reason is that the change in consumption between t and $t + 1$ is the result of new information at time $t + 1$ only. All information from t is already factored into the consumer's plans for consumption in the next period. If there was no new information in the next period, consumption would not change at all. It's only when new information becomes

available in $t+1$ that the consumer changes their consumption plans, and therefore c_{t+1} differs from c_t. This is shown in Figure 3.6.

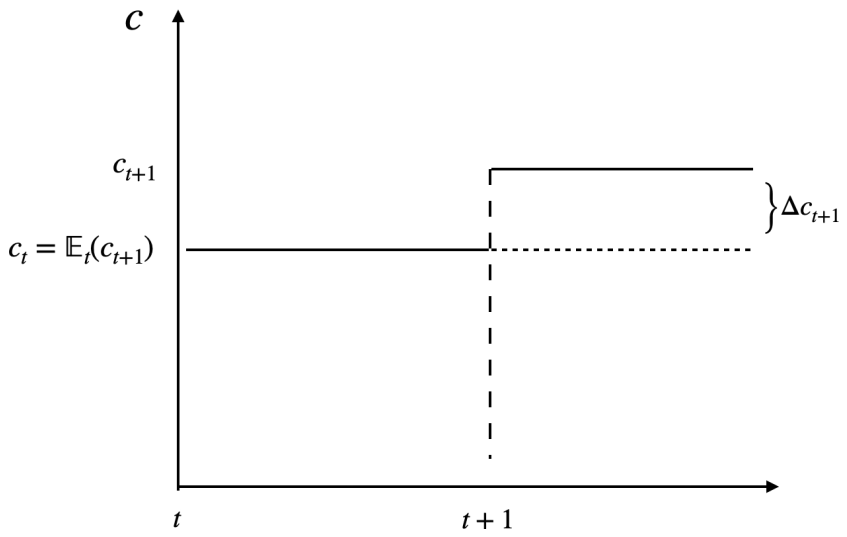

Figure 3.6 - The consumer receives new information at t+1 which leads them to consume more than expected

Therefore, for example, knowing the consumer's income in period t should not matter for the change in consumption in the next period. The consumer has already factored in all information from knowing their income today into estimating their lifetime wealth. It's only when new information arrives in the next period — say that income in $t+1$ is higher than expected — that consumption may change.

Tests of the Euler Equation

We mention two such tests of the Euler equation. First, Zeldes (1989) uses this approach and finds that current income is a predictor of the change in consumption for the approximately 20% of his sample which is liquidity constrained — that is, who may have problems borrowing to be able to effectively smooth their consumption. We will deal more with liquidity constraints in the next section. What you should keep in mind here is that, for these consumers, Zeldes finds a negative correlation between current income and consumption growth. He interprets this to suggest that these consumers are violating the permanent income hypothesis by increasing current consumption in response to current income, without also increasing future consumption.

This is shown in Figure 3.7. The left hand side shows what the permanent income hypothesis would predict if income increased at time t. Because the consumer is smoothing consumption, they would share the income increases across all periods. Therefore, consumption in time t would increase, but so would consumption in time $t+1$. Overall then, the change in consumption should be unaffected. Instead, Zeldes finds something more like in the right hand side. Here, the increase in income leads to an increase in current consumption but not future consumption — the consumer is not smoothing! The result is that the increase in

current income leads to a fall in the expected rate of change of consumption. In our figure, consumption falls between the two periods.

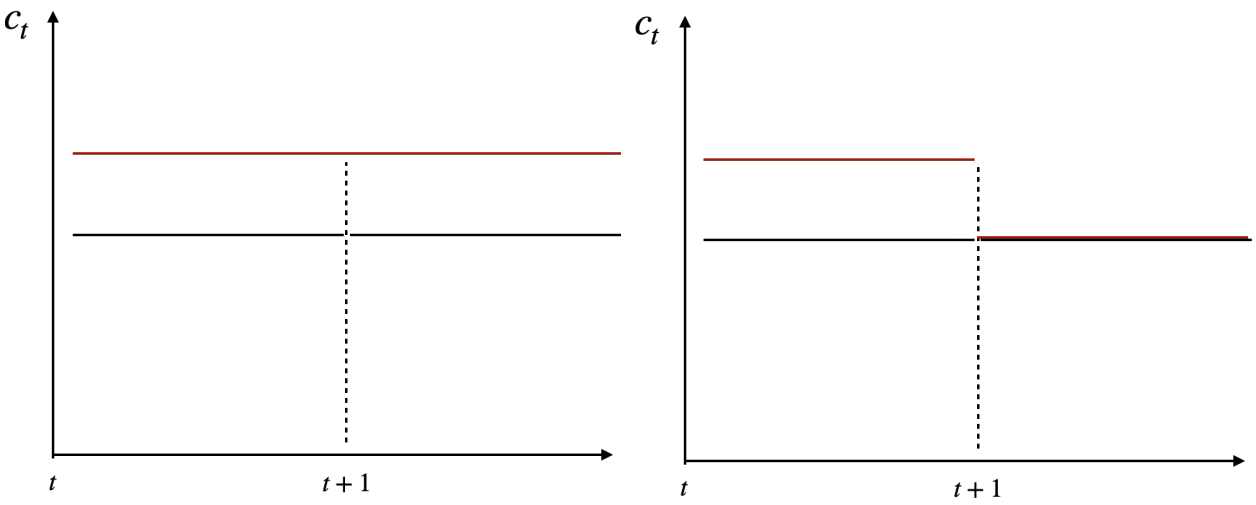

Figure 3.7 - A permanent income consumer (left) and what Zeldes finds (right) in response to a change in current income. The black line is the counterfactual with no income increase in period t. The red line is the effect of the income increase.

Similarly, Campbell and Mankiw (1989) find several lagged variables that are able to predict changes in consumption, thus violating the permanent income hypothesis. Importantly, they use this to estimate that about 50% of consumers are “hand-to-mouth”, suggesting that their consumption changes 1-for-1 with changes in income.

What you should take from both of these studies is that early on there was evidence that some consumers may have been behaving like the permanent income hypothesis, but that a significant fraction were behaving more like Keynesian consumers, responding significantly to changes in current income.

Modern Tests: Natural Experiments

Recently, tests of the permanent income hypothesis have been based on attacking the second row of Table 1 by looking at the implications of unexpected changes in income. To do so, they employ natural experiments. Researchers look for cases of unexpected transitory changes in income — e.g. winning the lottery, or receipt of stimulus cheques — and observe the changes in consumption. The permanent income hypothesis would suggest that these should be small, however researchers often find significant and large changes in consumption. There are actually many such papers, but we survey only a few here.

Parker et al. (2013)

One way Keynesians believed that the government could respond to a fall in aggregate demand — like a recession — is by providing cheques to households. The logic is that households, receiving this new income, would spend it thus increasing aggregate demand.

The permanent income hypothesis suggests that such stimulus measures would be futile. Stimulus cheques are (unexpected) transitory increases in income. The fact that they are transitory means that the impact on lifetime wealth should be small, and therefore the impact on consumption should also be small. Instead, consumers will save a majority of the new stimulus cheques to smooth their income over time.

Therefore, observing how consumers behave when they receive stimulus cheques is a test of the permanent income hypothesis. Thus, Parker et al. (2013) estimate the consumption response to the stimulus cheques sent to American households in 2008 in response to the Great Recession.

Stimulus cheques were sent out over several weeks, so the authors could compare households who had received their stimulus cheques already to those who had not to estimate the consumption response. When consumers received their cheque was effectively random, determined by the last digits of their social security numbers.

The authors find that, within the three months that payments were disbursed, consumers spent 12-30% of their cheques on non-durable consumption — a share much higher than what the permanent income hypothesis would predict.[13] Moreover, expenditures on durable goods like cars also increased, bringing the total expenditure to between 50-90% of the value of the cheques. The authors find that these effects are higher for low-income and older consumers, and that they are higher for homeowners rather than renters.

Fagereng et al. (2021)

Unless one is clairvoyant, or else one is very good at cheating, winning the lottery is about as close as we can get to an unexpected, transitory income shock. The difficulty in using lottery winnings to estimate MPCs comes from data limitations. Researchers require a large sample of lottery winners observed over many years as well as a sample of non-lottery winners to whom to compare. Ideally they would like to know the amount lottery winners win, their income and their consumption.

Fagereng et al. (2021) therefore were given a great opportunity when they cleverly discovered they could use Norwegian tax data to do exactly this. Norwegian lotteries are run by the state (just like OLG in Ontario). Although winnings are tax free, winners of above 10,000 NOK (~1,100 USD) must declare these winnings on their taxes. Tax returns allow us to observe other income and to estimate consumption based on income and the tax payer's change in wealth.[14]

The authors find that MPCs are high, with winners spending about 50% of their winnings in the first year, and about 90% within 5 years.

Interestingly, the MPC seems to decline with the amount of winnings. Winners of small prizes (< 2,000 USD) tend to spend almost all of it within the year, whereas for large prize winners (>8,000 USD), the estimated MPC after a year is less than 50%.

---

[13] In a practice problem below, you are asked to work out what the permanent income hypothesis would predict.

[14] In the simplest case, the change in wealth is equal to savings: $\Delta w_t = s_t$. Moreover, since all income is either consumed or saved, we can use this measure to back out consumption:

$$y_t = c_t + s_t$$
$$y_t = c_t + \Delta w_t$$
$$c_t = y_t - \Delta w_t$$

The authors also find that estimates vary by age — with older individuals saving a greater share of their winnings on average[15] — and with liquid wealth, with less liquid individuals having greater MPCs.[16]

---

[15] A practice problem asks you to compare this to the predictions of the permanent income hypothesis.

[16] We return to the liquid wealth issue in the next chapter.

Practice Problems

1. **Implications of the PIH for MPCs**. Thus far we have talked vaguely about the implication of the permanent income hypothesis for MPCs, suggesting that they will be "small" for transitory increases in consumption. We worked out once instance with an infinitely lived consumer. But no consumer lives forever. So what are realistic MPCs?

   To make some headway on this problem, imagine a consumer with known income stream between today and period T. At period T, the consumer dies. The consumer maximizes lifetime utility given by:

   $$U = \sum_{t=0}^{T} \beta^t \ln(c_t)$$

   Note here that T represents how many years the consumer believes they have left to live. Assume that consumers all live for 80 years. Therefore, $T = 80 - age$.

   Finally, for simplicity, imagine that $\beta = (1 + r)^{-1}$

   A) Write down the consumer's problem.

   B) Solve for the Euler equation. What does this imply about the growth of consumption over time? (Would this be different if we had not made assumptions about β?)

   C) Solve for consumption in period 0.

   D) What is the MPC out of an increase in current income? How does this depend on T?

   E) Given that T is the number of years the consumer has left to live, what does this mean about the effect of age on MPC? (Assume that all individuals live for 80 years, and therefore $T = 80 - age$, such that a lower T is associated with a higher age.)

   F) How do your conclusions in (E) compare with the the findings of Fagereng et al. regarding age?

   G) Finally, assume that the annual interest rate is $r = 0.04$. What is the MPC for someone who is 40 years old (and therefore has 40 years left to live)? (You may need to use a computer for this one, apologies!) How does this MPC compare to the estimates from Parker et al. (2013)?

What You Should Know After this Chapter

Intuition

- The Permanent Income Hypothesis can explain some features of the consumption-income data:
    - Aggregate income and consumption seem to have a positive MPC and a positive y-intercept
    - When the economy grows, the savings rate does not increase
- The predicted change in consumption following a change in income depends on whether the change in income is predicted and whether it is permanent or transitory.
- However, modern tests of the PIH find several inconsistencies with the data:
    - For some fraction of consumers, income growth is predicted by current income
    - Consumers who receive stimulus cheques or lottery winnings tend to spend a sizeable fraction of it within the year its received.

Technical Skills

- Be able to show how a PIH consumer should behave following an income shock if it is expected/unexpected or transitory/permanent.

Chapter 4. Credit Constraints

Perhaps it seems strange that many of the empirical papers above focused on how MPCs differed by liquid wealth. However, there is a reason.

Several authors have pointed out that the permanent income hypothesis won't hold in cases where consumers faced constraints on borrowing. Some important references are Flemming (1972), Schechtman (1976), and Bewley (1977).

There are good reasons for including such borrowing limits in our models. Obtaining consumer credit isn't particularly easy. For some durable goods — like homes or cars — we can often borrow against the collateral of the good. However, it's much harder to go to a bank and receive an unsecured loan on the promise that, in expectation, your income will be higher in 10 years and you'll be able to pay it back. Notice also that the bank faces serious informational problems — they don't know how much income you'll make in 10 years, nor whether you'll actually put in the effort necessary to earn that money.

It's true that many consumers have access to credit cards, which are an example of easily accessible unsecured consumer credit. However, these tend to have very high interest rates, certainly much higher than the rates at which consumers can lend, and usually have borrowing limits.

Thus, it's true that consumers usually *can* borrow. However, the ability to do so is limited and often must take place at very high interest rates. A simple assumption which approximates the real life situation, then, will be to simply allow no consumer borrowing.

Credit Constraints Can Increase MPCs

The intuition and implications for MPCs can be made clear with a simple example. Imagine a consumer who lives for two periods. In the first period they are poor with no savings or income, while in the second period they know they will receive a great inheritance from a rich relative.

How would the permanent income hypothesis suggest this individual would behave? Because their income is very low in the first period, the marginal utility of consumption will be very high, and vice versa for the second period when they are rich. Consequently, it would be optimal for the consumer to borrow in the first period to be able to smooth out their consumption.

However, imagine that the consumer cannot borrow — perhaps banks find it too risky to lend against a future inheritance. Then the consumer cannot optimally smooth their consumption: the marginal utility in the first period will be very high, and in the second period very low! If only they could trade consumption across periods by borrowing!

So what would happen in this case if the consumer won the lottery in the first period? In a permanent income hypothesis world, they would save part of their winnings to smooth consumption across time. However, in this case, there's no point in saving as the marginal utility of consumption tomorrow will be very low! In fact, they can better smooth their consumption by consuming most —or all — of their lottery winnings today.

Numerical Example

Let us illustrate this example with a small numerical and graphical example. Assume that the consumer's problem is given by:

$$\begin{aligned} \max_{c_0, c_1, s_0} \quad & U = u(c_0) + \beta u(c_1) \\ \text{s.t.} \quad & c_1 = y_1 + s_0(1+r) \\ & c_0 = y_0 - s_0 \\ & s_0 \geq 0 \end{aligned}$$

Note that this is our usual consumer's problem, but we've included a borrowing constraint which says that savings in the first first period cannot be negative, i.e. the consumer cannot borrow.

Let's given some particular values to our example.

$u(c) = 10c - c^2/2$
$y_0 = 1$
$y_1 = 9$
$\beta^{-1} = (1+r) = 1$

The Unconstrained Case

For the sake of exposition, let's first consider the unconstrained case. We know how to solve this. First, we find the consumer's Euler equation:[17]

$$u'(c_0) = \beta(1+r)u'(c_1)$$

Given our assumptions, this gives:

$$\begin{aligned} 10 - c_0 &= 10 - c_1 \\ c_0 &= c_1 \end{aligned}$$

In other words, it will be optimal for the consumer to perfectly smooth their consumption across periods.

Note that, as in the past, we can combine our per-period budget constraints into a single constraint:

$$c_0 + \frac{c_1}{1+r} = y_0 + \frac{y_1}{1+r}$$

Solving for c_0:

---

[17] Note that it is (relatively!) easy to find the Euler equation in this simple two-period case. Simply combine the two budget constraints and plug them into the utility function:

$$\max_{c_0} U = u(c_0) + \beta u\big(y_1 + (y_0 - c_0)(1+r)\big)$$

This allows us to convert our optimization problem into a problem with one variable and no constraints. Taking the derivative with respect to c_0 gives the Euler equation.

$$
\begin{aligned}
c_0 + \frac{c_0}{1+r} &= y_0 + \frac{y_1}{1+r} \\
c_0\left(\frac{2+r}{1+r}\right) &= y_0 + \frac{y_1}{1+r} \\
c_0 &= \frac{(1+r)y_0 + y_1}{2+r} \\
c_0 &= 5
\end{aligned}
$$

Since the consumer wants to perfectly smooth consumption, this implies that:

$$c_0^* = c_1^* = 5$$

Note that to make the math easy, we have assumed that the interest rate is 0, which leads the consumption choice to be the average of income across both periods. Keep in mind that with a non zero interest rate, this will not be the case!

What does this imply for saving in period 0? The consumer's income is very low in the first period ($y_0 = 1$), and so this will entail borrowing in the first period against future income. That is:

$$
\begin{aligned}
s_0 &= y_0 - c_0 \\
&= 1 - 5 \\
&= -4
\end{aligned}
$$

Graphically, we can think about the consumer choosing consumption points on their period utility functions:

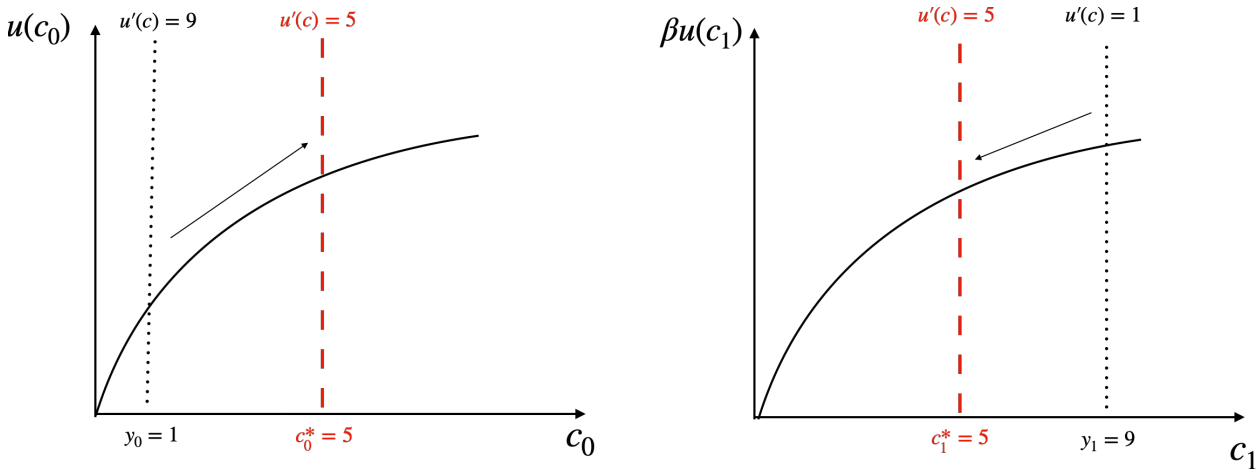

Figure 4.1 - The Consumer's Solution with No Borrowing Constraint

The black dotted lines show the case where the consumer consumes only their income in each period without borrowing or lending. Note that this means that the consumer's marginal utility of consumption in period 0 is 9, while it is 1 in period 1. I.e.

$$u'(y_0) = 9 > \beta u'(y_1) = 1$$

Clearly then it is worthwhile for the consume to borrow today so that they can consume more now and less later. Consequently, they will borrow until the marginal utility of consumption is equalized across periods, which happens at $c_0^* = c_1^* = 5$.

Now what happens if the consumer wins the lottery in period 0? Put another way, imagine now that their income is $y_0' = 7$. Plugging this in to the consumer's problem gives:

$$c_0'^* = c_1'^* = 8$$

In other words, the consumer received $6 more in the first period, and shares the income gain equally across periods (8=5+6/2).

Returning to our Figure 4.1, it is clear why this is the case. Imagine the consumer allocating their income gain dollar-by-dollar across periods. If they first increase their consumption in the first period, we move to the right on the period 0 diagram. This leads to a decrease in marginal utility in period 0 (i.e. $u'(c_0 = 6) = 4$). When looking at how to allocate their next dollar, they compare the marginal utilities in both periods: $u'(c_0) = 4 < u'(c_1) = 5$. Since marginal utility in period 1 is now higher than in period 0, the next dollar should be allocated there. Etc. This leads to an equal sharing of the income gain across periods.

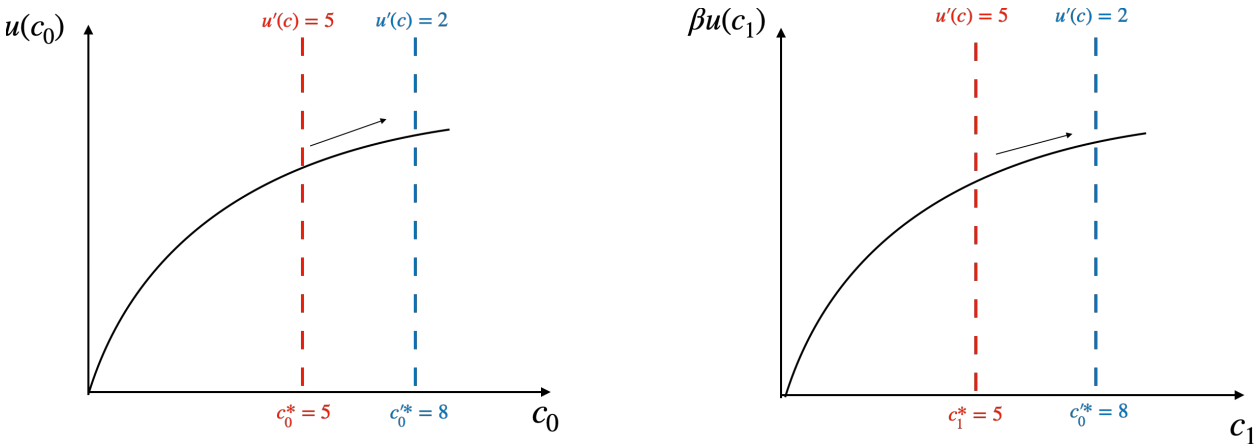

Figure 4.2 - An increase in present income is shared across periods in the unconstrained case

What is the MPC then? The consumer's income increased by 6, and their consumption increased by 3. Therefore we have an MPC of 0.5.

Constrained Case

Now what happens if the consumer is borrowing constrained?

For the optimal solution, the consumer is wishes to bring consumption forward in time, i.e. wishes to borrow. What happens if the consumer cannot borrow? Then in the first period they must consume their income. Because they did not borrow, this means they also must consume their income in the second period.

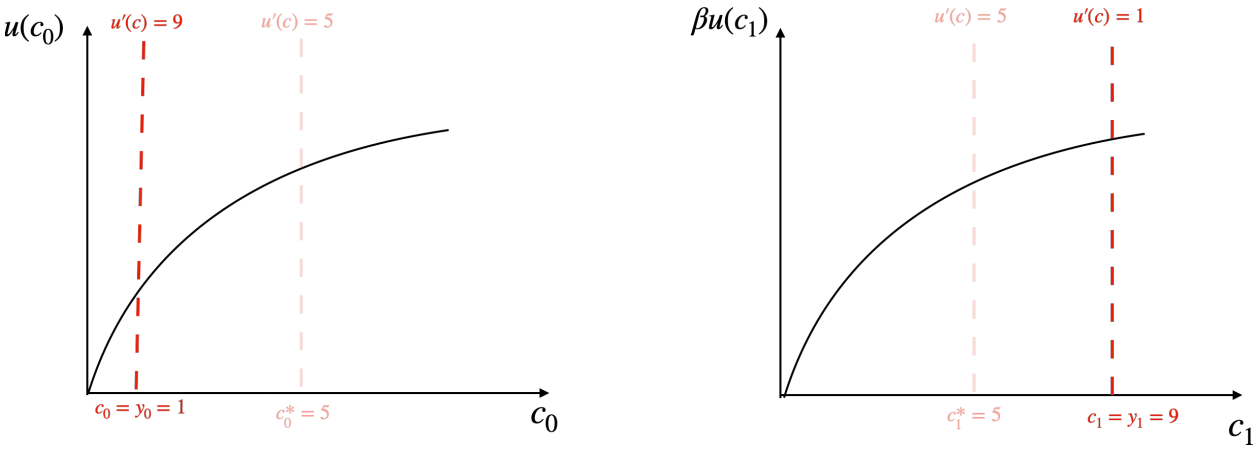

Figure 4.3 - The Consumer cannot borrow and must consume their income

Now what happens when the consumer receives an income increase in period 0? Again, assume that $y_1' = 7$.

Consider once again the consumer deciding how to allocate this additional lifetime income between the two periods. Imagine that they allocate one dollar at a time. To do so, they compare the marginal utility of consumption in each period. They will allocate the dollar of income to the period where marginal utility is highest. These decisions are represented in Table 4.1:

Income Left to Allocate	c_0	c_1	$u'(c_0)$	$u'(c_1)$	Where \$1 better spent
6	1	9	9	1	c_0
5	2	9	8	1	c_0
4	3	9	7	1	c_0
3	4	9	6	1	c_0
2	5	9	5	1	c_0
1	6	9	4	1	c_0

0	7	9	3	1	-

Table 4.1 - The dollar-by-dollar allocation decision of a constrained consumer

Even if the consumer allocates all of their income to period 0, the marginal utility of consumption is still higher in this period than in period 1! If anything, they would like to allocate even more income to period 0, but are constrained. This is shown graphically in Figure 4.4.

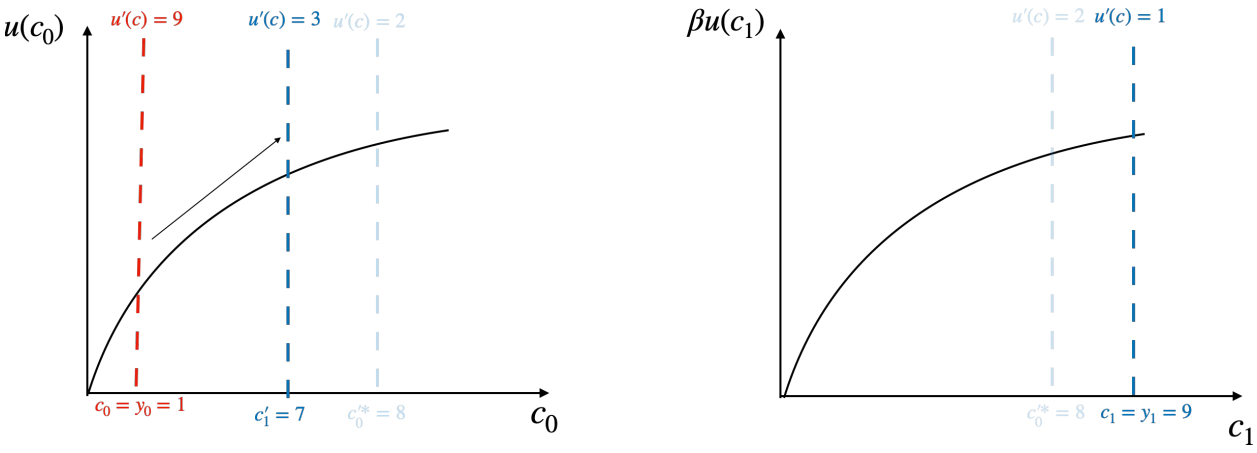

Figure 4.4 - Only first period consumption increases in response to a period 0 income shock

The unconstrained optimum is shown in faded lines. We see that the increase in income in the first period allows the consumer to approach the unconstrained optimum. Note that this would not be the case if the income shock happened in period 1! The consumer can't borrow, so consumption in period 0 would stay as it is. The advantage of the consumer receiving the income shock in period 0 is that it helps to reduce the borrowing constraint — the consumer wants to consume more in the first period, but can't!

What is the MPC in this case? The consumer consumes the entirety of the income shock. Therefore MPC=1.

Credit Constraints Don't Bind with High Liquid Wealth

Recall that in the empirical results from Chapter 3, a great deal was made about differences between consumers in "liquid wealth". Liquid wealth is any asset that a consumer can easily draw on to fund consumption. Money in a savings account or a chequing account are good examples of liquid wealth. However, a house — most people's most valuable asset — is not an example of liquid wealth since it is hard to use this asset to buy consumption goods.[18]

[18] A consumer would, in general, need to sell their house first to get the funds necessary to purchase consumption goods. Note that there are some financial instruments, such as home equity lines of credit, which allow consumers to borrow against the value of their home to fund consumption. This makes housing somewhat more "liquid" for our purposes.

In our simple model, there is only one savings asset — a savings account — which is perfectly liquid, as consumers can easily use their savings to purchase consumption goods. Moreover, a consumer's income is paid in money, another form of liquid wealth.

Consequently, in our model, define liquid wealth as:

$$l_t = s_{t-1}(1 + r) + y_t$$

In other words, liquid wealth is equal to the consumer's accumulated savings from the past period plus their current period income.

For our purposes, note that a consumer with high amounts of liquid wealth will not be constrained and will therefore have a lower MPC.

Return to our model from the previous section, but now assume that income receipts are reversed. That is, the consumer receives a high income in the first period and a low income in the second period:

$$y_0 = 9$$
$$y_1 = 1$$

What is the consumer's liquid wealth in the first period? By our definition:

$$l_0 = y_0 = 9$$

Liquid wealth is simply equal to income since the consumer has no accumulated savings.

The consumer's preferences are the same, and, given our assumptions about interest rates, the consumer's lifetime wealth is also the same.[19] Thus, the unconstrained solution will be identical:

$$c_0^* = c_1^* = 5$$

This means that now the consumer will be savings in the first period

$$s_0 = y_0 - c_0 = 9 - 5 = 4$$

---

[19] Recall that lifetime wealth is

$$we = y_0 + y_1/(1 + r).$$

Because we assumed that $(1 + r) = 1$, this simplifies to:

$$we = y_0 + y_1$$

and therefore it doesn't matter in what order the consumer receives income (i.e. whether $\{y_0, y_1\} = \{9,1\}$ or $\{y_0, y_1\} = \{1,9\}$). However, note that if the interest rate was different, then the order the consumer receives income in would matter.

And then consumes their savings in the second period. This is shown in Figure 4.5. However, note that this will also be the result in the constrained case! The consumer is not able to borrow, but this is not an issue since they only want to save.

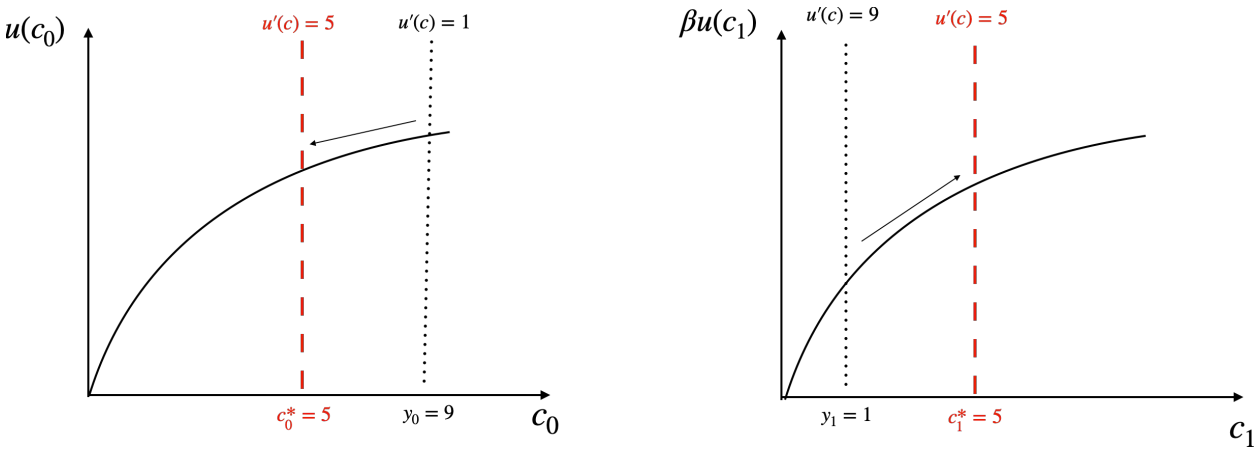

Figure 4.5 - Constrained and Unconstrained cases are the same with high liquid wealth

Note that this means that if the consumer receives an income shock in period 0, they will share it across periods, and so their MPC will be lower than if they had low liquid wealth.

Credit Constraints and Precautionary Savings

Obviously the consumer does not want to be constrained. There is a real utility loss to not being able to smooth consumption. The constraint means that they must consume a lot in cases where they have low marginal utility and very little in cases where they have high marginal utility — exactly the opposite of what they would want!

The result is that, in unconstrained periods, consumers will have an incentive to save more to tide themselves over in constrained periods. Another way to think about this is that, because consumers can't smooth as well by bringing consumption forward in time (i.e. borrowing), they compensate by saving more to be able to smooth better between constrained and unconstrained situations. (We will see more what this means in the example that follows.) We refer to this as precautionary savings.

Numerical Example

To see this precautionary savings motive in action, we need to consider now a three-period model. In the first period the consumer has a high income, in the second period they fall on hard times, and in the third period they regain their original level of income. The middle period can be thought of as a period of unemployment. In an ideal unconstrained world, the consumer would smooth their consumption by saving in the first period, but also by borrowing in the second period. In this way, they are taking consumption from both the first and last period to compensate for their low income in the period of unemployment. However, in the face of credit constraints, they will be unable to borrow against the last period, and thus will compensate by saving more in the first period.

Once more, consider a consumer like in the last section, with per-period utility given by:

$$u(c) = 10c - c^2/2$$

But now the consumer live for three periods, lifetime utility is given by:

$$U = u(c_0) + \beta u(c_1) + \beta^2 u(c_2)$$

Once again, let's make the simplifying assumption that saving yields no interest, and that the interest rate is equal to the inverse of the discount factor. I.e.

$$\beta^{-1} = (1 + r) = 1$$

Finally, assume that the consumer goes through a spell of unemployment in period 1. Therefore, income is given by:

$$y_0 = 6$$
$$y_1 = 0$$
$$y_2 = 6$$

Unconstrained Case

Once again, let's set up the consumer's problem for the unconstrained case:

$$\max_{c_0, c_1, c_2} U = u(c_0) + \beta u(c_1) + \beta^2 u(c_2)$$

$$\text{s.t.} \quad c_0 + \frac{c_1}{1+r} + \frac{c_2}{(1+r)^2} = y_0 + \frac{y_1}{1+r} + \frac{y_2}{(1+r)^2}$$

Given our simplifying assumptions, the Euler equation gives:

$$c_0 = c_1 = c_2$$

I.e. the consumer is perfectly smoothing when there are no borrowing constraints.

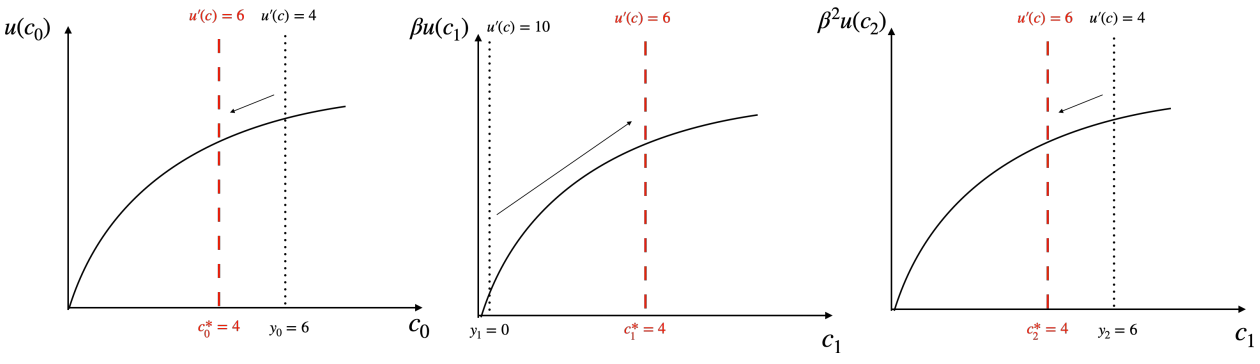

Figure 4.5 - In the unconstrained case, the consumer uses income from the first and last periods to smooth consumption in the second period

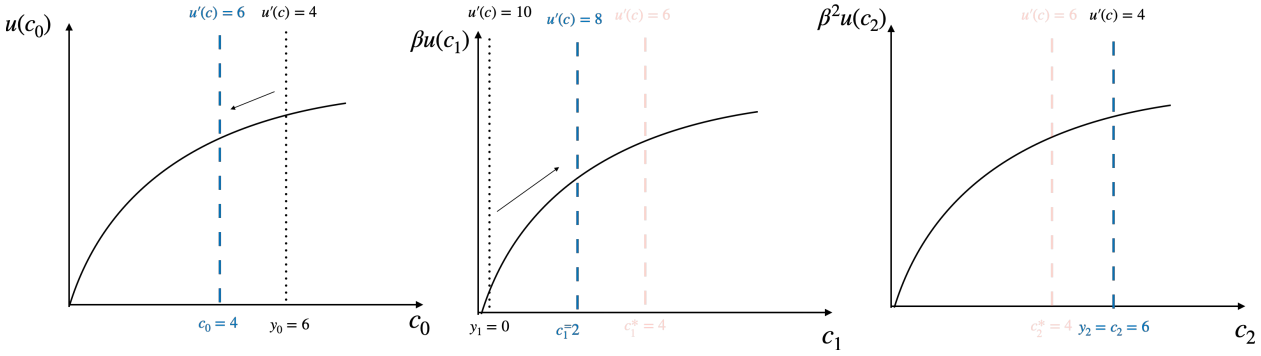

Figure 4.6 - Saving the same as in the unconstrained case is suboptimal

Plugging into the budget constraint and using the fact that $(1 + r) = 1$:

$$
\begin{aligned}
c_0\left(1 + \frac{1}{1+r} + \frac{1}{(1+r)^2}\right) &= y_0 + \frac{y_1}{1+r} + \frac{y_2}{(1+r)^2} \\
3c_0 &= 6 + 0 + 6 \\
c_0 &= 4
\end{aligned}
$$

And since the consumer is perfectly smoothing:

$$c_0^* = c_1^* = c_2^* = 4$$

Again, the fact that the discount factor and the gross interest rate are equal to 1 means that the consumer consumes the average of their total lifetime income in each period.

Note that this means that the consumer saves in period 0 and borrows in period 1. Therefore, they are smoothing their consumption by taking income from both period 0 and period 2 to compensate for their very low income in period 1.

Graphically:

Constrained Case

Now imagine that the consumer faces a borrowing constraint. Note that this means that they can no longer borrow in period 1, and thus they are unable to move some income from the period 2 to period 1. How should they react? We will see that this will lead them to save *more* in the first period!

To see why, consider the following. The consumer cannot borrow in period 1, and therefore they cannot shift consumption from period 2 to period 1. What if they don't borrow in period 1, but still save in period 0 like they do in the unconstrained case. This is shown in the Figure 4.6

Starting from the right of the figure, note that because they cannot borrow, consumption in period 2 is equal to income in period 2. However, the consumer still saves $2 in period 0, as they would in the unconstrained case. This means that consumption in period 0 is 6-2=4, and in period 1 is 0+2=2.

However, note that the consumer could do better!

Look at the marginal utility of consumption across periods. Currently, marginal utility of consumption is lowest in period 2 (4) and is highest in period 1 (8). Therefore, ideally the consumer would like to shift consumption from period 2 to 1. However, because this would entail borrowing in period 1, this is not possible.

However, consider that marginal utility in period 0 is also lower than in period 1 (6 vs. 8). Therefore, shifting consumption from 0 to 1 should be beneficial to the consumer. To shift consumption from period 0 to 1, the consumer simply must save, which is allowed!

Note that if the consumer saves one more unit, they equalize marginal utility across periods 0 and 1. Note that the only remaining inequalities in marginal utility ($u'(c_2) > u'(c_1) = u'(c_2)$) cannot be addressed since this would entail borrowing. Therefore, the consumer is at the constrained optimum.

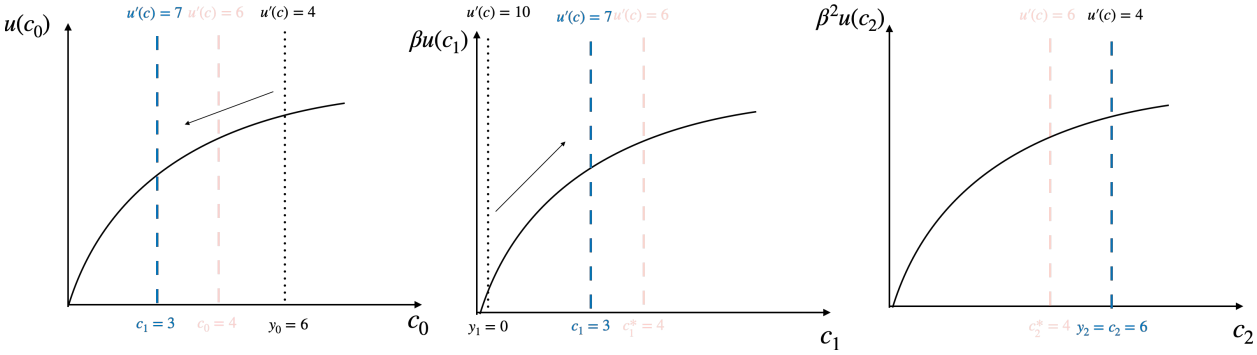

Figure 4.7 - The consumer will save more in period 0 in the constrained case

Saving Your Way Out of the Constraint

The last two subsections can be combined to form an interesting result. First, we saw that consumers with with high liquid wealth will behave like permanent income consumers. These consumers have enough wealth to effectively smooth against drops in current or future income. The result is that income shocks in the current period will have a low MPC: the consumer is not constrained in this period, and therefore they are happy to share the gain in income across all future periods by saving most of it.

Second, we saw that the presence of future potential constraints will lead the consumer to save more today. If the consumer accumulates enough liquid wealth, then they will not be constrained in the future.

Taking these two together, it's possible that very few consumers will actually face constraints over time. When times are good, consumers will save to ensure against future negative income shocks.

Numerical Example

Consider again a situation similar to that in the previous section, in which a consumer's income alternates each period. In recessionary periods, they receive a low income and in boom periods they receive a high income.

Assume the that we alternate each period between boom and recession, starting with a recession. Therefore, the consumer's income stream is:

$$\{y_0, y_1, y_2, y_3, \dots\} = \{0,6,0,6,...\}$$

Or more specifically:

$$y = \begin{cases} 6 & \text{if } t \text{ odd} \\ 0 & \text{otherwise} \end{cases}$$

Let's assume that the consumer lives from some arbitrary period of time, T, where T is an odd number, meaning that the consumer ends their life on a high note in a boom period! Once again, make all the same assumptions from the last section:

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

Where

$$u(c) = 10c - c^2/2$$

And once again, we make the simplifying assumption:

$$\beta^{-1} = (1 + r) = 1$$

Unconstrained Case

In the unconstrained case, the consumer solves:

$$\max_{\{c_t\}} \sum_{t=0}^{T} \beta^t u(c_t)$$
$$\text{s.t.} \sum_{t=0}^{T} \frac{c_t}{(1+r)^t} = \sum_{t=0}^{T} \frac{y_t}{(1+r)^t}$$

As we've seen, the Euler equation is:

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

implying perfect smoothing. Subbing this into the budget constraint and using the fact that $(1 + r) = 1$:

$$Tc = \sum_{t=0}^{T} y_t$$

$$c = \bar{y}$$

Where $\bar{y}$ is the average income across all periods. Since income alternates between 0 and 6 (and there are an equal number of each), this implies that:

$$\begin{aligned} c^* &= (0+6)/2 \\ &= 3 \end{aligned}$$

The perfect smoothing means that the consumer consumes their average income each period. Since income fluctuates between 0 and 6, they consume 3 units in each period.

Constrained Case

Here we will show that, aside from the first and last period, the consumer will replicate the unconstrained case.

The consumer's solution in the first period is obvious. They receive income $y_0 = 0$, but cannot borrow to compensate. Therefore they must consume their income, and c_0.

What about subsequent periods? Note that in subsequent periods the consumer can mimic the unconstrained case. Consider the diagram in Figure 4.8.

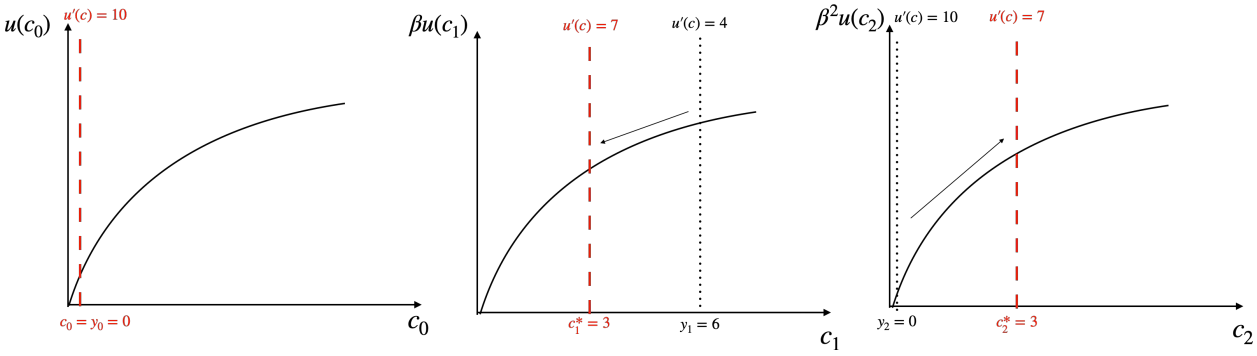

Figure 4.8 - The constrained consumer at the beginning of their life can equalize marginal utility in periods 1 and 2, but not period 0.

After period 0, the consumer will try to equalize marginal utilities of consumption across periods as best as they can. Note that they can accomplish this by saving $3 in good periods, and consuming their savings in bad periods. So, in period 1 the consumer receives income $y_1 = 6$ but finds it is optimal to save 3 to consume next period when they have no income. Note that this will equalize marginal utilities across these two periods.

What about in period T? Consider this case represented in Figure 4.9:

In the two preceding periods ($T-2,\ T-1$), the consumer maintains their strategy of perfectly smoothing income across alternating good and bad periods. This means that marginal utility is equalized in these periods. However, what about the final period? Here, marginal utility is low given the consumer's relatively high income. In an unconstrained world, the consumer would borrow against this income to increase their consumption in other periods. However, because of the borrowing constraint, the consumer has no way of equalizing marginal utility in period T with marginal utility in subsequent periods!

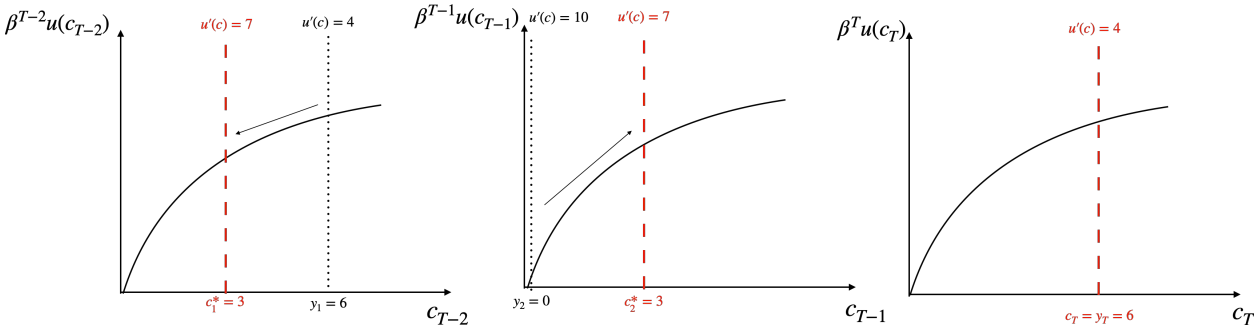

Figure 4.9 - The constrained consumer cannot borrow against their final period high income

Note that we know this is the consumer's constrained optimal solution because there is no way to make them better off. Marginal utility is equalized across all periods except for the first and last. However, the consumer has no way to equalize marginal utility in these periods: in period 0, they would like to consume more by borrowing, while in period T, they have too much income against which they would like to borrow.

So what would happen if the consumer received a small income shock in period 1? They would share it across all subsequent periods except for T. The reason we can tell this is to again think about the consumer allocating the income shock across periods in an optimal way. Since marginal utilities are already equalized, except for period T, in which the marginal utility is already low, there is no gain to allocate all of the income shock to consumption in any given period. If the consumer allocates a small fraction to consumption in period 1, this will lead the marginal utility of consumption in this period to be lower than in period 2, 3, 4,… At which point it will be optimal to allocate the next bit of the income shock to one of these periods. Etc.[20]

What can we conclude from this? Savings can be a very good means of smoothing consumption, especially when income fluctuates around some average. If there will always be low income periods in the future, the consumer can predictably save in high income periods to smooth consumption in the low income periods. The borrowing constraint only holds when the consumer holds low liquid wealth (like in period 0) because they have no been able to save enough already.

[20] A practice problem will have you work out exactly what the MPC will be in this case.

Practice Problems

1. **The Consumer's Problem with and without credit constraints.** Imagine a consumer with period period utility given by $u(c) = 50c - 3c^2/2$, who maximizes utility given by:

   $U = u(c_0) + \beta u(c_1)$
   Assume that:

   $\beta^{-1} = (1 + r) = 2$

   Income in each period is:

   $y_0 = 5$
   $y_1 = 11$

   Unconstrained Case
   A) Find the Euler equation.
   B) Use the Euler equation and the lifetime budget constraint to solve for consumption in each period.
   C) Does the consumer borrow or lend?
   D) What is the MPC from an increase in current period income (y_0)?

   Constrained Case
   E) Imagine now the consumer is borrowing constrained. What will be consumption in each period?
   F) What is the marginal utility of consumption in each period?
   G) Imagine that the consumer wins $3 in the lottery in period 0. What happens to consumption in both periods? What is the MPC out of new income?
   H) Now imagine that the consumer wins $9. What happens to consumption in both periods? What is the MPC? Why is it different than in the case in (F)?

   (Hint for G,H: What would be the solution in the unconstrained case? Is the consumer still constrained?)

What You Should Know After this Chapter

Intuition

- Credit constrained consumers will tend to have higher MPCs.
    - Because they are unable to smooth by borrowing, if their current income is relatively low it may be efficient to consume most of a positive income shock.
    - This means that **MPCs** will tend to be high.
- Because they do not like being constrained, consumers will try to save more.
- In this way they can **save their way out of the constraint.**
- **C**onsumers with lots of liquid wealth will behave like permanent income consumers.

Technical Skills

- Be able to demonstrate the above intuition with the use of diagrams.
- Be able to show mathematically how non-constrained consumers will behave in cases of fluctuating income.

Chapter 5. Models with Uncertainty

Chapter 4 established that credit constraints mean that consumers don't always operate like permanent income consumers. In particular, consumers who are constrained will tend to have high MPCs. In theory, this might bring us back towards a Keynesian type world like the one explored in Chapter 1. Moreover, it suggests that, if we were to build a model of the economy, it would be important to keep track of which consumers are constrained as this might be important for the behaviour of the economy.

However, note that at the end of Chapter 4, we also discussed how consumers can "save their way out of the constraint". Meaning that, given time, consumers can effectively smooth consumption across periods through savings alone. These consumers will tend to have higher liquid wealth and will behave like permanent income consumers and with low MPCs.

So we are left with two conflicting stories: (1) credit constraints matter for consumption behaviour, and (2) maybe they don't matter that much as sufficient savings eliminate the credit constraint.

To disentangle these stories, we can try to use a simple model to get at the general equilibrium impacts of consumption and saving. Such a model would track consumers over time. Consumers who save accumulate liquid assets which they can then use in periods where their incomes are lower.

However, first we must introduce an important element: uncertainty. Thus far, we have generally been dealing with cases with certain streams of income. That is, the consumer knows exactly what their income will be in all future periods. This is probably a strong assumption — you and I might have a good idea of what our income will be in the next year or so, but whether we face unemployment or get a raise or find a new job is uncertain, and becomes increasingly uncertain the further into the future we look. Thus, most macro models assume that consumers know only the possible distribution of their future stream of income, but not the actual values.

In general, we will see that the conclusions of chapter 4 — using the model with certainty — carry over to the model with uncertainty. In particular, consumers will have high MPCs close to the constraint. However, consumers with high levels of assets will behave like an unconstrained consumer. Finally, consumers will attempt to save their way out of the constraint, in that they will increase their savings in such a way that they are unlikely to be constrained in the future.

Constraints and Uncertainty

Adding uncertainty changes the consumer's problem only slightly. For an infinitely-lived consumer, this changes the consumer's problem to be:

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

Note that in the utility function, utility takes an expectation operator. Because consumption will depend on the consumer's income, and income is uncertain, consumption is also uncertain. For example, in the face of a high income shock in period t, the consumer will make a different consumption choice than if they receive a low income.

Brief Numerical Example

As an example, consider a consumer with per-period utility given by:

$$u(c_t) = 100c_t + c_t^2/2$$

The consumer faces an equal probability each period of being in recession, in which case they receive income $y_l = 1$ or of being in a boom, in which case they receive $y_h = 4$. In other words, the income process is given by:

$$y_t = \begin{cases} 4 & \text{with } p = 0.5 \\ 1 & \text{with } p = 0.5 \end{cases}$$

However, the consumer does not know in advance what will be their income in any given period. The consumer is infinitely-lived (meaning they live from $t = 0,...,\infty$). The consumer maximizes their lifetime expected utility.

We once again assume that $\beta(1 + r) = 1$, although we allow for positive interest, setting $\beta = 0.95$.

Unconstrained Solution

We begin by finding the unconstrained solution. This is important in its own right because it demonstrates how to solve the consumer's problem with uncertainty. It will also be important when we look at the constrained solution, because we will be able to compare the two.

In the unconstrained case, we have that the consumer solves:

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

$$\text{s.t.} \quad \sum_{t=0}^{\infty} \frac{\mathbb{E}(c_t)}{(1+r)^t} = s_{-1}(1+r) + \sum_{t=0}^{\infty} \frac{\mathbb{E}(y_t)}{(1+r)^t}$$

We've modified the budget constraint to allow for the fact that at period 0 the consumer might have some stock of assets already accumulated. This explains the first term on the right hand side of the budget constraint.

The Euler equation is thus similar to what we've seen before, but with an expectation operator:

$$\mathbb{E}u'(c_t) = \beta(1+r)\mathbb{E}u'(c_{t+1})$$

First, since we're assuming that $\beta(1 + r) = 1$, this simplifies to:

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

Which says that the consumer wants to equalize expected marginal utility across periods. We can think of the consumer at period 0 making plans for their consumption in every future period conditional on the realizations of income. In other words, at period 0 they are making a detailed plan of what to consume for any history of realizations of income. Thus, in period 2, they might have a different consumption choice if their history of income realizations was $\{y_l, y_l, y_h\}$ as opposed to $\{y_l, y_h, y_h\}$. Each such realization of incomes has a given probability, and therefore the consumer can calculate the expected marginal utility in each period.

Next, since in period 0 the consumer knows their consumption, we can rewrite this as:

$$u'(c_0) = \mathbb{E}u'(c_t) \quad \forall t$$

Finally, using the definition of per-period utility:

$$100 - c_0 = \mathbb{E}(100 - c_t)$$

$$100 - c_0 = 100 - \mathbb{E}(c_t)$$

$$c_0 = \mathbb{E}(c_t) \quad \forall t$$

Which says that the consumer wants to smooth expected consumption across periods.

Plug this into the budget constraint:

$$\sum_{t=0}^{\infty} \frac{\mathbb{E}(c_t)}{(1+r)^t} = s_{-1}(1+r) + \sum_{t=0}^{\infty} \frac{\mathbb{E}(y_t)}{(1+r)^t}$$

$$c_0 \sum_{t=0}^{\infty} \frac{1}{(1+r)^t} = s_{-1}(1+r) + y_0 + \bar{y} \sum_{t=1}^{\infty} \frac{1}{(1+r)^t}$$

$$c_0 \frac{1+r}{r} = s_{-1}(1+r) + y_0 + \bar{y} \frac{1}{r}$$

$$c_0 = s_{-1} r + \frac{r y_0 + \bar{y}}{1+r}$$

In the second line, we separate out period 0 income, which is known to the consumer and we call expected income $\mathbb{E}(y) = \bar{y}$, which we can think of as the average income the consumer will receive over their (infinite) lifetime. The third line uses the properties of infinite series to solve for the sums.

The fourth line says that consumption at period 0 will be equal to any interest earned on current assets, plus the weighted average of current income and expected future income.

Finally, note that we can use this to solve for a consumer's savings in period 0. Start with the fact that savings in period 0 are everything left over after the consumer receives income, interest income and makes some consumption choice.

$$s_0 = y_0 + s_{-1}(1+r) - c_0$$

$$s_0 = y_0 + s_{-1}(1+r) - \left(s_{-1}r + \frac{r y_0 + \bar{y}}{1+r}\right)$$

$$s_0 = y_0\left(1 - \frac{r}{1+r}\right) + s_{-1} - \frac{\bar{y}}{1+r}$$

$$s_0 = s_{-1} + \frac{y_0 - \bar{y}}{1+r}$$

$$s_0 = s_{-1} + \beta(y_0 - \bar{y})$$

Where the final line makes use of the fact that $\beta^{-1} = (1+r)$.

This says that the consumer will increase or decrease their savings depending on the realization of y_0. For example, to use our numerical example, we have that $y_h = 4$, $y_l = 1$, and since each happen with equal probability, $\bar{y} = 2.5$.

Therefore, in good times, the consumer increases their stock of savings:

$$\Delta s = \beta(y_h - \bar{y}) = \beta(1.5)$$

And in bad times, they will decrease their stock of savings (or borrow if savings are less than or equal to zero):

$$\Delta s = \beta(y_l - \bar{y}) = -\beta(1.5)$$

In the limit, when $(1+r) = \beta = 1$, the consumer will always save any extra income above what is expected and dissave (or borrow) when income falls below the expectation.

What does this mean about changes in a consumer's asset holdings on average? In good times they increase their asset holdings, and in bad times they decrease their asset holdings by an equivalent amount. Since good and bad times happen with equal likelihood, on average their change in assets will be 0.

Constrained Case

The consumer's problem with borrowing constraints is given by:

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

$$\text{s.t. } c_t = y_t + s_{t-1}(1+r) - s_t$$

$$s_t \geq 0$$

We will not solve the consumer's constrained problem analytically here. Solving the problem in a world with certainty and constraints is hard enough; adding uncertainty makes things very difficult and usual means there is no analytical solution. Instead, economists solve these models with a computer using techniques not explored in this chapter.

However, we can discuss the shape of the solution. First, consider a consumer solving this problem. In any period, they observe their income and they observe any savings brought over from the previous period. Based on this information, the consumer has to decide how much to consume and save. In other words, the solution takes the form of functions:

$$c_t = c(s_{t-1}, y_t)$$
$$s_t = s(s_{t-1}, y_t)$$

These functions are known as the consumer's policy functions. The variables in the brackets are the consumer's state variables. They are the values the consumer knows at the beginning of the period that inform their optimal behaviour.

Using a computer, it is relatively easy to solve for the consumer's policy functions. Figure 5.1 shows the savings functions for a consumer in a recession ("low income") and in a boom ("high income"). The figure also contains a 45 degree line. If a point were to lie on the 45 degree line, that would mean that savings today are the same as savings tomorrow, or, put another way, the consumer is maintaining the same level of assets across time periods, and is neither saving nor dissaving.

A few things are important to notice. First, consider the red line, which is the consumer's savings function in a boom. This line relates the consumer's savings decision to the stock of assets they begin the period with. For example, it says that a consumer in a boom with assets of at the beginning of the period $10 will choose next period savings to be equal to around $11.5.

Note first that this consumer behaves in the way we would expect: in boom times they save, while in recessionary times they dissave. Interestingly, both lines are fairly straight with a slope of approximately 1. The major exception is the recessionary state line close to the origin. Here, the consumer has low income and would like to borrow but can't because of the borrowing constraint. Consequently, the consumer is constrained, and sets their savings to 0.

Next we turn to the consumption functions in Figure 5.2. These show consumption as a function of previous year's savings. Again, note two things. First, consumption is higher in the high state than the low state. Recall, however, that this is also the case in the unconstrained case when $(1 + r) > 1$. Second, notice that consumption is very low in the low state when previous year's assets are 0. The reason is that this consumer is completely constrained and end up consuming all of their income. Hence we can see that $c(y = y_l, s_{-1} = 0) = y_l = 1$.

Next, notice that we can also get a good idea of MPCs. The MPC is how much more a consumer would consume today if we gave them one more unit of income. Note that this is equivalent to giving the consumer one more unit of assets, adjusted for the interest rate. In other words, consider the consumer's budget constraint:

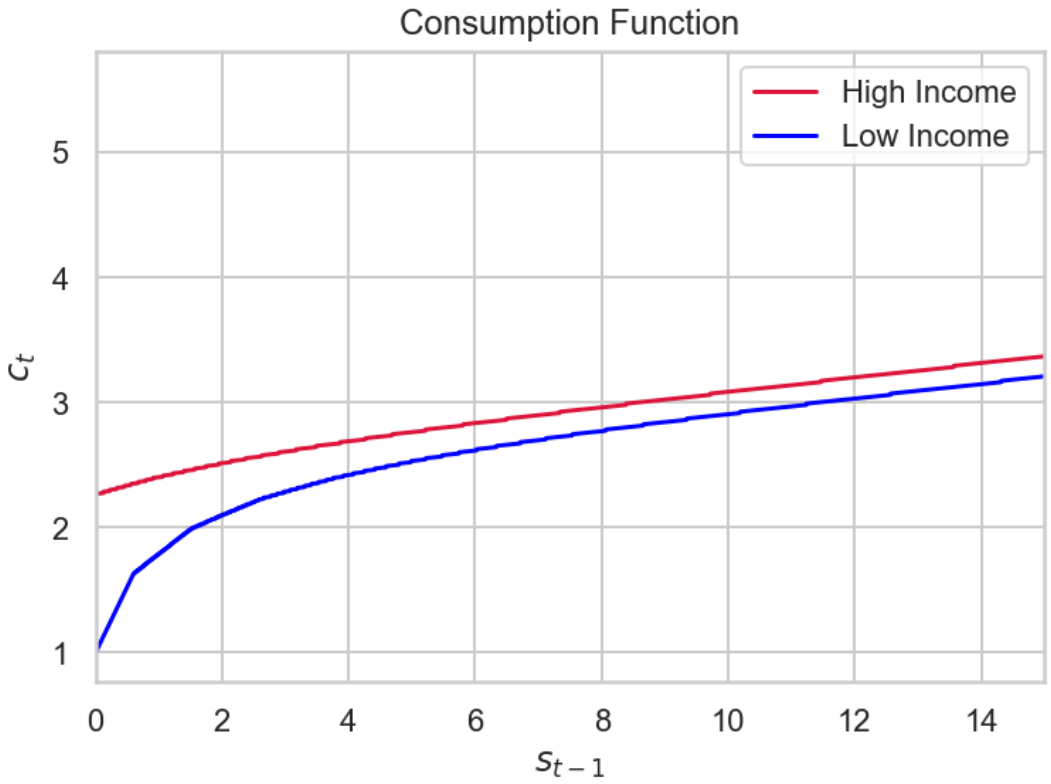

Figure 5.1 - Savings Functions

$$c_t = y_t + s_{t-1}(1 + r) - s_t$$

The consumer would be indifferent between receiving one more unit of income, or $1/(1 + r)$ units of assets. Or in other words:

$$MPC = \frac{dc_t}{dy_t} = \frac{dc_t}{ds_{t-1}/(1 + r)}$$

Note that the slope of the blue line is very high close to the y-intercept, suggesting that the consumer's MPC is very high for low rates of assets in the low state. However, both lines flatten as the consumer accumulates assets, suggesting that consumers with lots of liquid wealth will have very low MPCs.

In fact, as the consumer's assets increase, they become more like permanent income consumers. Consider Figure 5.3 which shows the consumer's consumption function as compared to the consumption function in the unconstrained case. We can see that as assets increase, the constrained consumer behaves more and more like the unconstrained consumer.

This makes sense given that the more assets a consumer holds, the less likely they are to be constrained anytime in the future, and therefore the more they can behave like in the permanent income hypothesis. Put another way, for high enough savings, a consumer will likely never have to borrow in the future. If they don't have to borrow, they will never face a

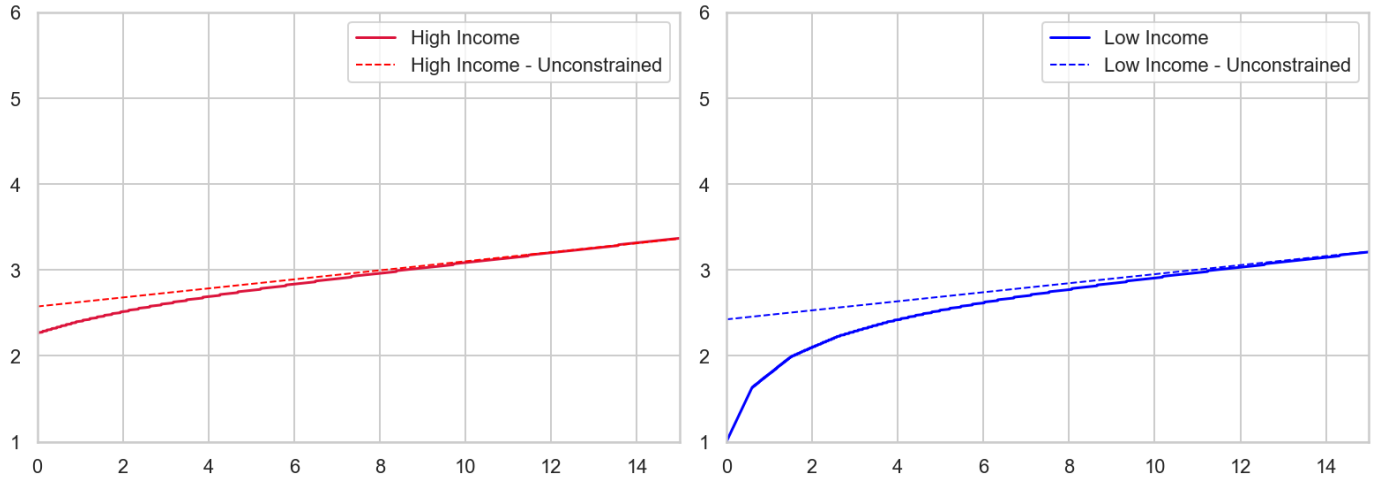

Figure 5.3 - Consumption Functions, Constrained and Unconstrained

constraint, and therefore they are in the exact same situation as a permanent income consumer. Notice that we saw something very similar in the case with certainty in Chapter 4, in which consumers with enough liquid assets behave just like a permanent income consumer.

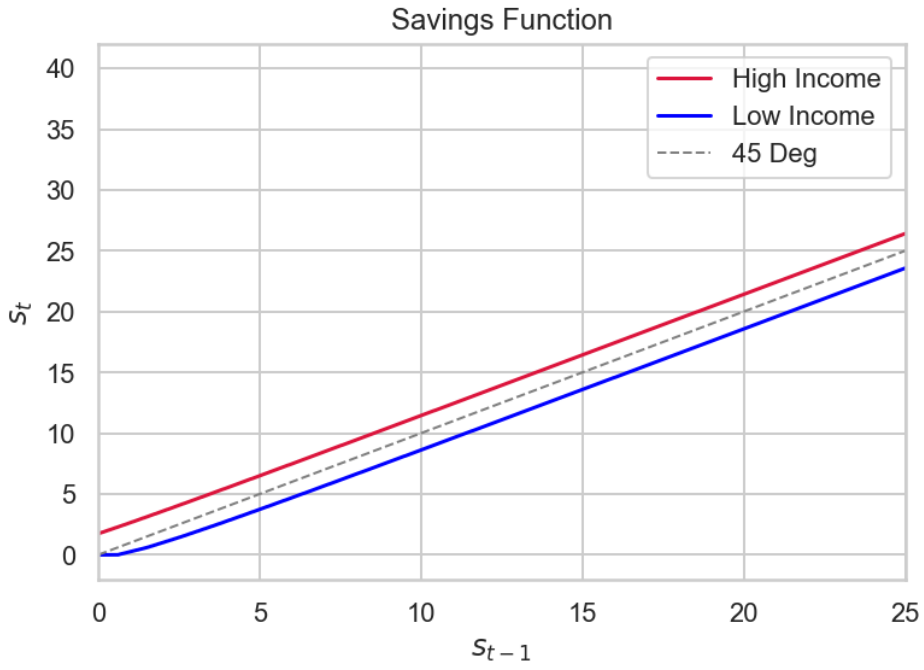

Figure 5.4 - Savings Functions of the Constrained Consumer

Also interesting is that consumers will try to save their way out of the constraint. Thus, when times are good, but liquid assets are low, the consumer will save a relatively large part of their good income shock because there is a high value in escaping the constraint. Consider the

consumer's savings functions, shown in Figure 5.4. This function relates asset holdings in period t to asset holdings in period $t-1$.

A point on the 45 degree line would mean that a consumer keeps their asset holdings the same. Note that in good times the savings function is everywhere above the 45 degree line, meaning that a consumer will increase their asset holdings when they receive a positive income shock. On the other hand, in bad times, the consumer will generally decrease their asset holdings, i.e. dissave and consume some of their asset stock.

Away from the constraint, the consumer behaves like a permanent income consumer increasing and decreasing their assets by equal amounts in good and bad times respectively. Thus, away from the constraint, the savings functions are equally spaced above and below the 45 degree line.

However, notice that close to the constraint, two things are true. First, consumers save a larger part of their income when they are in good times (represented by the larger vertical distance between the red line and the 45 degree line). But also, because of the constraint, their savings cannot fall below 0. Both of these mean that the average change in assets is high close to the constraint, while away from it, it approaches 0. This is shown in Figure 5.5. In other words, the consumer tries to save their way out of the constraint.

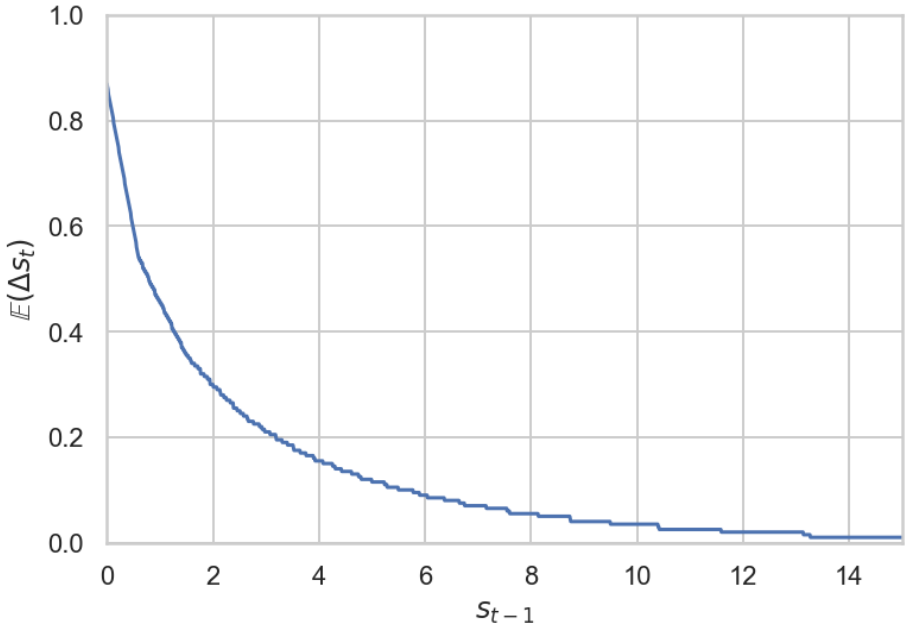

Figure 5.5 - Expected Change in Assets of the Constrained Consumer

A Comparison of to Models with No Uncertainty

Let's take a second to briefly recap what we've seen so far. Many of the conclusions of our model with uncertainty mirror the model with no uncertainty.

1. Consumers with low assets have a high marginal propensity to consume. This was evident by looking at the consumption function around the constraint.
2. When consumers have high enough assets, they behave like unconstrained (permanent income) consumers. We see this by comparing the consumption and savings functions to the unconstrained case.
3. Consumers will behave in such a way as to remove themselves from the constraint — i.e. saving more on average when assets are low.

How the model will behave in aggregate will depend on how these features interact. For example, how many consumers end up constrained depends on how effective consumers are at saving out of the constraint. The more effective, the more consumers will behave like unconstrained, permanent income consumers. And therefore, the better we are able to model the economy with a representative agent.

What You Should Know After this Chapter

Intuition

- Models with uncertainty behave very similarly to basic models:
    - Consumers with low assets have a high MPC
    - Consumers with high amounts of liquid wealth behave like unconstrained consumers
    - Consumers will try to save their way out of the constraint

Technical Skills

- Be able to solve the consumers problem with uncertainty without credit constraints and quadratic utility.
- Understand how the existence of credit constraints will change the result.

Chapter 6. The Distribution of Assets and Wealth Inequality

So far we have looked at the solutions to the consumer's problem which yields policy functions: a consumption function and a savings function. Both functions tell us how a given consumer will behave given their state variables: current income and accumulated assets.

How the economy will behave in aggregate will depend on the resulting distribution of these consumers. For example, if most consumers have high assets in equilibrium, this means that most consumers will behave like permanent income consumers and it might be easier to model the economy using a representative agent.

In this chapter we look at what determines this equilibrium distribution of assets. In general, these models are solved through simulation. In other words, once we have determined consumer's policy functions, we simulate random draws for income given the income process (or any other random variables in the model), and simulate how consumers react in aggregate.

The distribution of assets is important for the functioning of our model. But it is also important in its own right. If, for example, we are interested in the distribution of wealth across consumers, having a model which allows us to understand this distribution will be very useful. In particular, in this section we examine the determinants of the asset distribution, which is an important component in formulating a theory of wealth inequality.

Example Simulation of a Constrained Consumer

Take our model from Chapter 5. We can easily simulate the asset holdings of a given consumer given a set of random draws for income. This is shown in Figure 6.1. This consumer has several lucky draws early on in their lifetime, which leads them to accumulate a fair amount of savings. From period 20 to period 80, they have low income draws on average, leading them to consume most of their liquid wealth. They then replenish their assets.

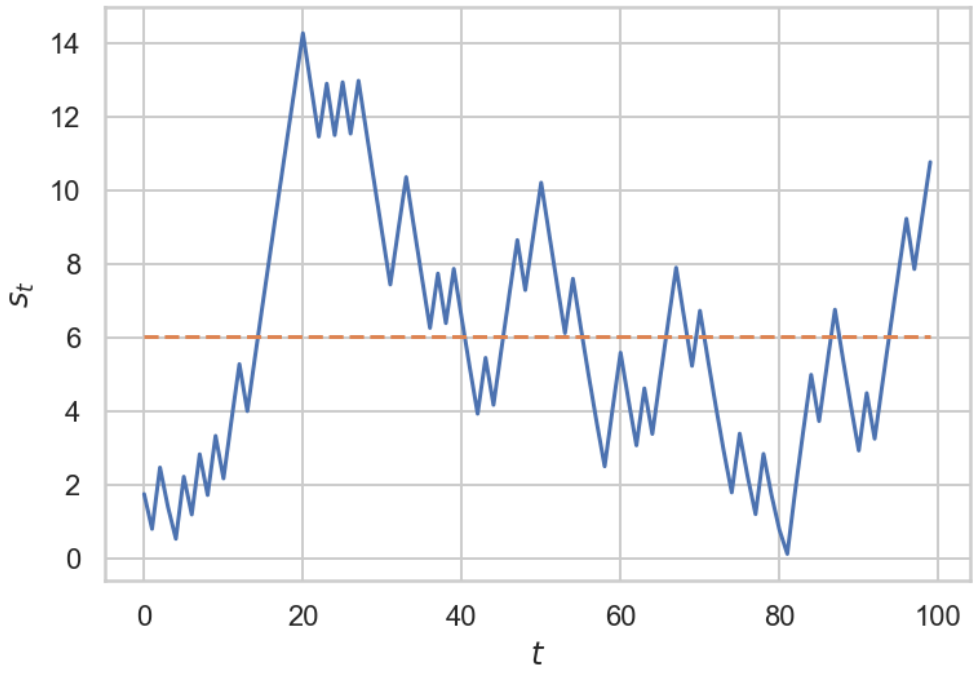

Figure 6.1 - The asset holdings over time of a simulated constrained consumer

We can easily simulate thousands of such consumers and look at the distribution of assets in each period. However, first let's think about what we should expect of the asset distribution.

The Case of Explosive Asset Distributions

Unconstrained Case

Let's start with our simple model from Chapter 5 with no constraints. Recall that in this case the savings function was given by:

$$s = s_{-1} + \beta(y - \bar{y})$$

Which says that savings will grow by amount $\beta(y - \bar{y})$. As we saw, this means that, on average, savings shouldn't change. This simple savings function can be reformulated as a simple algorithm:

$$\Delta s = \begin{cases} 1.5\beta & if \quad y = y_h \\ -1.5\beta & if \quad y = y_l \end{cases}$$

Or, to put it another way, savings can be represented as a random walk:

$$s = s_{-1} + \varepsilon$$

Where ε is a random variable with mean 0, i.e. $\mathbb{E}(\varepsilon) = 0$.

What happens to the mean of assets over time? As you may know, a random walk is stationary, and therefore it's mean will not change over time. Mathematically:

$$\mathbb{E}(s) = \mathbb{E}(s_{-1}) + \mathbb{E}(\varepsilon)$$

$$\mathbb{E}(s) = \mathbb{E}(s_{-1}) + 0$$

$$\mathbb{E}(s) = \mathbb{E}(s_{-1})$$

However, the variance is explosive. To see why, consider the asset choices of 100 unconstrained consumers over time shown in Figure 6.2. If all consumers start with 0 assets, if there are enough consumers, there will always be a consumer who receives all good income draws. This consumer always saves a fraction of their income draws, and therefore their assets grow continuously. Likewise, there will always be a consumer who receives all negative income draws. And of course, there will be consumers everywhere in between. The result is that the bounds of the distribution will increase with time, and so will the variance.

Constrained Case

The situation is even worse in our constrained case. With our simple model, high-asset consumers behave like permanent income consumers, and therefore may increase their assets without bound. *However*, low asset consumers are stuck at the asset constraint, and therefore

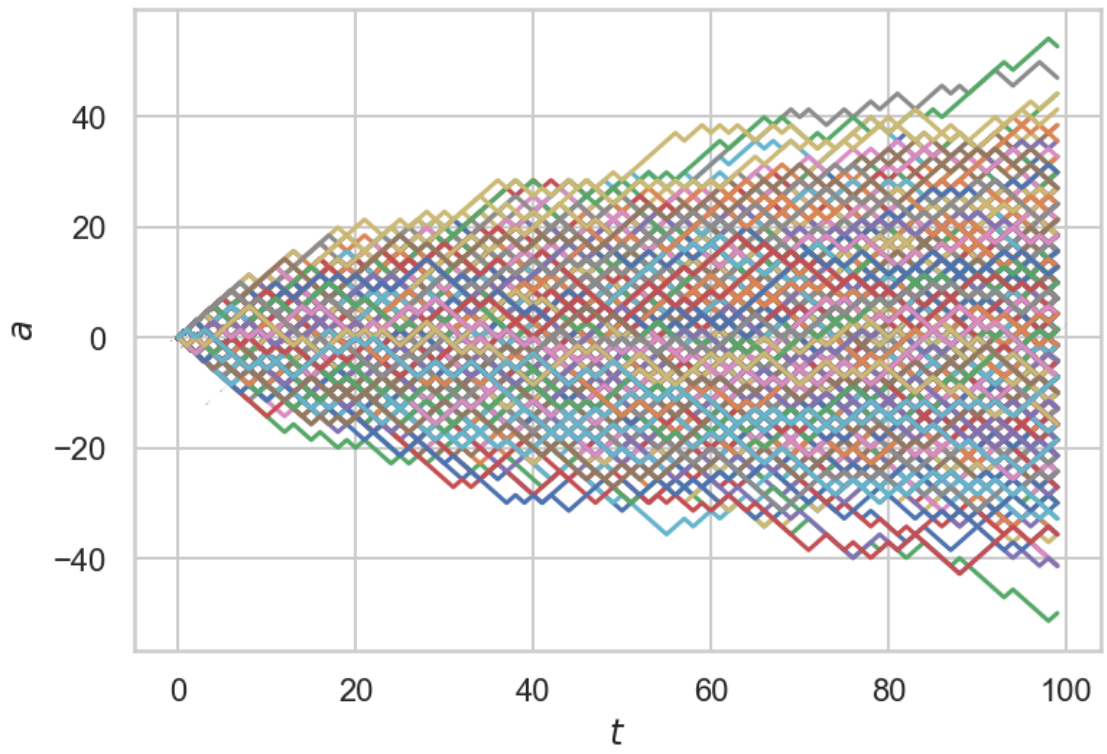

Figure 6.2 - Asset choices for unconstrained consumers over time

the distribution cannot expand symmetrically on either side of 0. The result is that average assets also increase over time! This is shown in Figure 6.3.

Fixing Explosive Distributions

It seems unlikely that the asset distribution will be explosive over time. It's true that wealth inequality has increased considerably over the last decades. But it's not clear that this is the result of explosive growth in the asset distribution, otherwise wealth inequality should be increasing all the time, rather than just over certain historical periods. So what can we do to our models to fix this? One option is to change the value of $\beta(1+r)$.

Recall that these parameters are important for the Euler equation. In general the Euler equation takes the form

$$u'(c_0) = \beta(1+r)\mathbb{E}u'(c_1)$$

Up until now, we frequently make the simplifying assumption that $\beta(1+r) = 1$. However, what if $\beta(1+r) < 1$? In general, this means that consumer cares less about the future. This is either because savings yields relatively little because the interest rate is low, or else it is because consumers discount future utility more. The result should be that consumers save less on average, and therefore assets grow less quickly.

CRRA Utility - A brief detour

Up until now we have been using quadratic utility functions because they allow for easy derivations. As we move to more complex models, we won't be doing as many derivations. Thus, let's switch to using the more common CRRA utility function:

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}$$

Note that this form has marginal utility:

$$u'(c_t) = c_t^{-\gamma}$$

For $\gamma = 1$, this will be equivalent to log utility.

Non-Explosive Distributions

With $\beta(1+r) = 1$, we will still have explosive distributions with CRRA preferences. The savings function and the distribution after 40 periods are shown in Figure 6.4. The story is the same as with quadratic utility. Consumers who receive positive income shocks always want to increase their assets. This can be seen by the fact that the high-income savings function is always above the 45 degree line. The result is that there will always be some consumer who receives all positive income shocks and whose assets increase without bound.

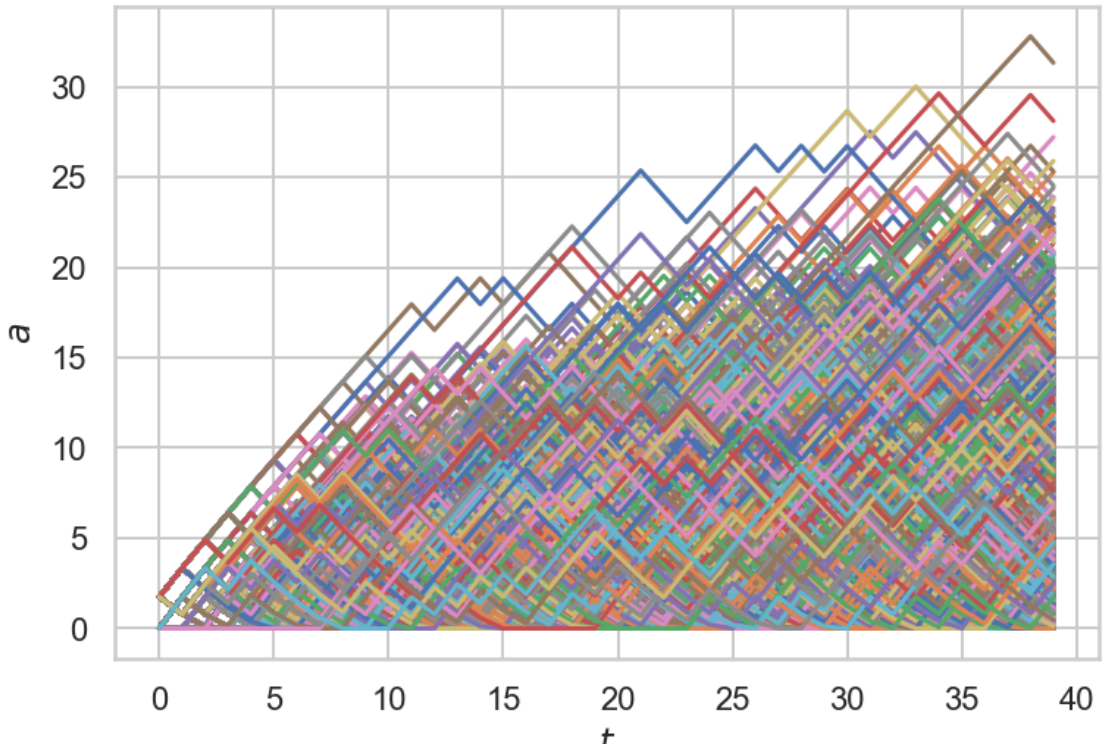

Figure 6.3 - Asset choices for constrained consumers over time

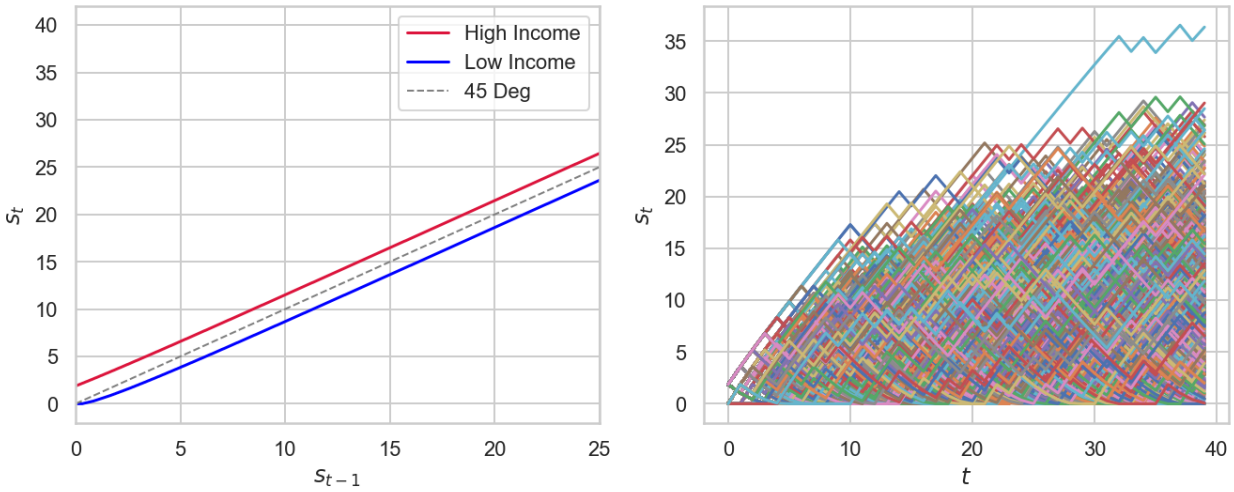

Figure 6.4 - Savings Function and Distribution of Assets after t periods, CRRA Preferences with $\beta(1+r)=1$

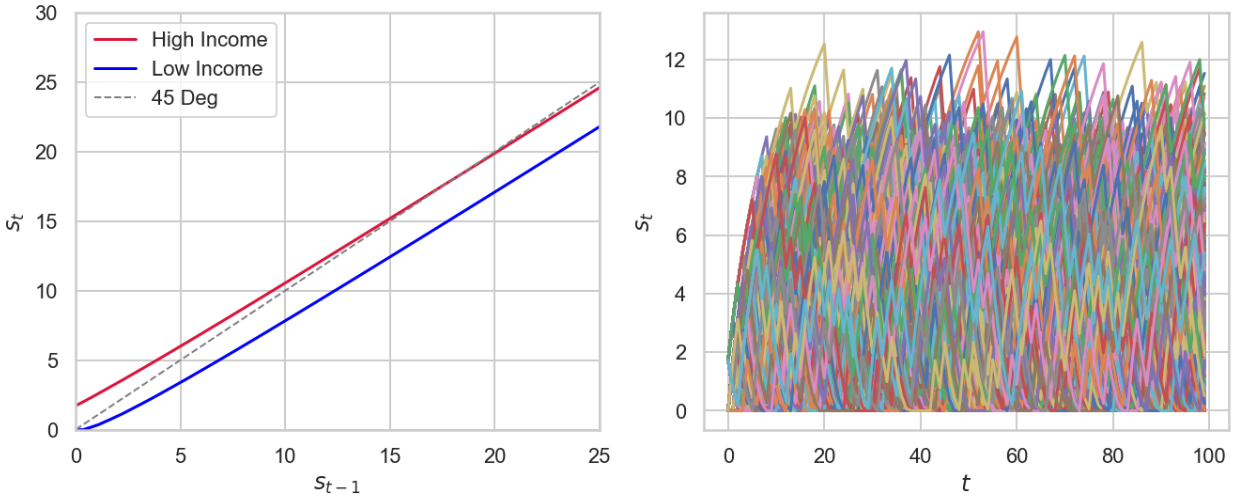

Figure 6.5 - Savings Function and Distribution of Assets after t periods, CRRA preferences, with $\beta(1+r)<1$

However, what happens if we set $\beta(1+r)<1$? Figure 6.5 shows the equivalent graphs for this case. First, notice the difference in the savings function. The slope of the savings function is now below 1 meaning that, in the high-income case, it eventually crosses the 45 degree line. This means that the consumer's assets will be bounded from above. Even if the consumer gets all positive income shocks, if their assets were ever to reach $20, they would want to keep them at exactly this amount. If their assets were to somehow grow above $20, they would want to dissave.

Why does the consumer behave in this way?

[TBC]

Chapter 7: An Introduction to Two-Asset Models

[TBC]