

Lecture notes on Macroeconomics and Finance

Ec 745 - Fall 2010

Francois Gourio
Boston University

Warning: these notes are incomplete and may contain some mistakes or typos.

January 26, 2011

Contents

1	Organization; stylized facts; the standard model	2
1.1	Organization	2
1.2	Stylized facts about asset markets	2
1.3	Asset pricing basics	4
1.4	The standard model = CRRA utility and lognormal iid consumption process	5
1.5	The present value model with constant discount rates	9
2	General results on stochastic discount factor: existence, uniqueness. Hansen-Jagannathan bounds.	11
3	Euler equation estimation and GMM; Campbell-Shiller decomposition	21
3.1	Euler equation tests	21
3.2	Euler equation tests	22
3.2.1	GMM general formulas	23
3.3	Campbell and Shiller decomposition	24
3.4	Hansen and Sargent (1980) Prediction Formulas	27
4	Additional material: Disasters	29
5	Habits models	31
6	Basic intuition	31
7	Campbell and Cochrane Model	32
7.1	Key mechanism	34
7.2	Solving the model	35

8	Recursive Utility	36
8.1	The SDF and the market return	37
8.2	Log-normal iid results	40
8.3	Long-run risk model (Bansal and Yaron JF 2004)	43
8.4	Tallarini/Hansen computations when IES=1	45
8.5	Log-linearization (Campbell 1993 AER/1996 JPE)	47
8.6	Savings and Portfolio Choice with Recursive Utility	50
8.7	Other “exotic” preferences	51
9	Yield curve	51
10	Corporate bonds	60
11	Heterogeneity and asset pricing (Rough notes)	67
12	Production Economies	76

1 Organization; stylized facts; the standard model

1.1 Organization

- Check the syllabus (esp. requirements).
- Check your email regularly (once every day) and the blackboard website.
- These notes are rough, and may contain typos. Please report typos that you find.
- Work through these notes with a pen and paper, check that you can derive all the key results.
- Speak up in class! Any question or comment is good. If you don’t understand, ask me to repeat or explain again.

1.2 Stylized facts about asset markets

Before describing models it is useful to have some ideas about the key facts. (Read Cochrane “New Facts in Finance”.) Of course to some extent the facts you want to look at depend on the theory, but some general background is useful.

Here’s a summary: I will cite the numbers for the US 1947-2008 sample but they hold more generally (across different time periods and across different countries):

- The average real return on holding the SP500 or a similarly broad index is about 8% per year. This number comes with a large standard error, though: $\sigma(E(R)) = \frac{\sigma(R)}{\sqrt{T}}$ hence the confidence interval is quite large.
- Stock returns are very volatile: $\sigma(R)$ is about 17% per year.
- Stock returns show little serial correlation ($\rho = .08$ in q data, $-.04$ in annual data).
- The average risk free rate is about 1% per year (as proxied by US Tbill, after inflation).

(e) The risk-free rate is not very volatile ($\sigma(R)$ around 2% for annual data), but it is persistent ($\rho = .6$ in annual data), leading to some medium-run variation.

(f) Combining the previous facts, the equity premium - the difference between stock and Tbill average returns - is large, about 7%.

(g) Stock returns are predictable (time-series): there is some “mean-reversion” at medium-run horizon. A typical regression is

$$R_{t \rightarrow t+k}^e = \alpha + \beta \frac{D_t}{P_t} + \varepsilon_{t+k},$$

or

$$R_{t \rightarrow t+k}^e - R_{t \rightarrow t+k}^f = \alpha + \beta \frac{D_t}{P_t} + \varepsilon_{t+k},$$

where $t \rightarrow t+k$ means the cumulated return from time t to time $t+k$. The table below reports the slope and the R^2 for both regressions.

Data	1y	2y	3y	4y	1y	2y	3y	4y
coeff	3.83	7.42	11.57	15.81	3.39	6.44	9.99	13.54
tstat	2.47	3.13	4.04	4.35	2.18	2.74	3.58	3.83
R2	0.07	0.11	0.18	0.20	0.06	0.09	0.15	0.17

The slope coefficient is positive: when the price is high compared to dividends, the future return is low on average. There are also other variables which predict returns, such as the term premium (the difference between long rates and short rates) and the short-term interest rate. These regressions were initially perceived as a challenge to “efficient markets” since returns are forecastable. As we will see during the class, the fact that returns are predictable need not indicate a violation of frictionless markets, rational expectations, or rationality. Big debate on the statistical/economic significance of this. (1) Statistical: the persistence of P/D makes these regressions somewhat spurious; in particular the t-stats need to be adjusted and the long-horizon forecasts are not more significant than the 1 year forecast despite much larger R^2 . These regressions are highly unstable and have little out-of-sample forecasting power. (2) Economic: do we really believe that there are large changes over time in expected returns: why don’t investors time the market (invest in stocks when returns are expected to be high and go to bonds when expected returns are low)?

(h) There is a large heterogeneity in average stock returns across firms. Small firms have higher returns on average (the size anomaly). Firms with low Tobin’s q (or low D/P ratio, or low Earnings/Price ratio, or low Book/Market) also have higher returns on average (the value premium). Last, firms which have had positive returns recently tend to keep having positive returns for a (short) while (The momentum anomaly.) [Check Ken French’s webpage.] In principle these differences in average returns could be due to risk but often the firms with higher mean returns do not appear more risky according to the standard risk measures, such as the CAPM (we’ll discuss this in more detail later).

We will discuss later in more details facts about bond returns and currency returns, but here are a few reference facts on these:

(i) Long-term bond returns are not much higher than short-term bond returns (about 2.5% vs. 1%), but they are quite volatile ($\sigma = 7.5\%$ per year), where the short-term bond return is not volatile as explained above.

(j) Bond returns are predictable too - buy long term bonds when their prices are low (i.e. yields high) compared to short-term bonds - i.e. when the yield curve is upward sloping - this generates a positive return on average as bond prices mean-revert.

(k) Same thing with currencies - investing in currencies with high interest rates tends to generate a positive excess return (the forward premium puzzle, aka uncovered interest rate parity (UIP) puzzle, aka currency carry trade).

Of course, there are a lot of different asset classes in the world, such as real estate, corporate bonds, commodities, venture capital, private equity, etc. The general ideas we cover apply to all asset classes, but there are of course specific interesting aspects to each asset class. Moreover, the institutional environment in which you invest is different (e.g. the transaction costs, the taxes, the cost of shorting the asset, etc.).

1.3 Asset pricing basics

Deriving the Euler Equation

Consider a consumer who has expected utility over the payoffs today and tomorrow:

$$u(C_t) + \beta E_t(u(C_{t+1})),$$

where u is increasing and concave (and satisfy a Inada condition, $\lim_{c \rightarrow 0} u'(c) = +\infty$). Suppose this consumer has some starting allocation, e.g. his labor income Q_t in period t and Q_{t+1} in period $t + 1$. Suppose he can trade an asset with a payoff X_{t+1} . For instance, it could be a stock that he can buy or sell at price P_t , in this case the payoff next period is the stock price P_{t+1} plus dividend D_{t+1} , $X_{t+1} = P_{t+1} + D_{t+1}$. The payoff X_{t+1} is a random variable; at date t , an investor does not know exactly how much he will get from his investment at date $t + 1$.

Let ξ be the number of stock shares that the investor chooses to buy.

$$\text{Max}_{\xi} u(C_t) + E_t[\beta u(C_{t+1})]$$

subject to:

$$\begin{aligned} C_t &= Q_t - P_t \xi, \\ C_{t+1} &= Q_{t+1} + X_{t+1} \xi. \end{aligned}$$

Substituting the constraints into the objective, and setting the derivative with respect to ξ to zero, yields:

$$P_t u'(C_t) = E_t[\beta u'(C_{t+1}) X_{t+1}],$$

where $P_t u'(C_t)$ is the loss in utility if the investor buys another unit of the asset, and $E_t[\beta u'(C_{t+1}) X_{t+1}]$ is the expected and discounted increase in utility he obtains from the extra payoff X_{t+1} . The investor continues to buy or sell the asset until the marginal loss equals the marginal gain.

Note: this assumes an interior solution – in particular, you can buy or sell the asset no short sales, borrowing constraints etc.

The basic pricing formula is thus:

$$P_t = E_t \left[\beta \frac{u'(C_{t+1})}{u'(C_t)} X_{t+1} \right].$$

This formula is the central formula of the class (and indeed, of modern finance). *This is joint restriction on asset returns and consumption.*

We can express this in a more general way using a “stochastic discount factor” (SDF) M_{t+1} as

$$P_t = E_t[M_{t+1} X_{t+1}]. \tag{1}$$

In this example, the SDF is:

$$M_{t+1} \equiv \beta \frac{u'(C_{t+1})}{u'(C_t)}.$$

In the next class, we will consider in more general settings the existence and uniqueness of a SDF satisfying the equation (1).

The formula above applies to many cases:

- For stocks, the payoff X_{t+1} is the price next period P_{t+1} and the dividend D_{t+1} .
- For a one-period real bond, the payoff is 1: you buy it at price P_t and you get 1 unit of good next period.
- Can apply to options too.

Alternatively, we can think of a return on an asset in the following way: you pay 1 unit today, and you receive R_{t+1} units (goods) tomorrow. The Euler equation is thus:

$$E_t[M_{t+1}R_{t+1}] = 1. \quad (2)$$

A risk-free interest rate corresponds to the following case: you pay 1 unit today, and you receive R_t^f units tomorrow, where R^f is known at date t (which is why it is risk-free).

Real or nominal stochastic discount factors and returns: so far everything is in real units. You can do this (e.g. measure real returns in the data), or you can measure returns in nominal terms (in dollars). To do so, Assume that prices P_t^n , P_{t+1}^n and payoffs X_{t+1}^n are nominal. Let Π_t be the price index (the CPI) at date t . Then the Euler equation is:

$$E_t \left(M_{t+1} \frac{X_{t+1}^n}{P_{t+1}^n} \frac{\Pi_t}{\Pi_{t+1}} \right) = 1.$$

While we derived these formula by looking at only choices made at t for period $t+1$, they hold for much more general models of course, since the preceding (“variational”) argument also applies there. The beauty of these equations is that they are *partial information* method: you only need to specify a subset of the full model, in order to obtain interesting, testable implications.

1.4 The standard model = CRRA utility and lognormal iid consumption process

By “standard model” I mean a representative consumer with CRRA utility. The Euler Equation for any asset return R_{t+1} (return from t to $t+1$) is:

$$E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} \right) = 1,$$

where E_t is the conditional expectation, β the discount factor, and C_t is aggregate consumption per capita.

Let’s first consider the case with no uncertainty: then the interest rate between two periods is

$$R_{t,t+k}^f = \frac{1}{\beta^k} \left(\frac{C_{t+k}}{C_t} \right)^\gamma.$$

Hence: real interest rates are high when people are impatient, i.e β is low. If everyone wants to consume now, it takes a high interest rate to convince them to save.

Real interest rates are high when consumption growth is high. In times of high interest rates, it pays for investors to delay consumption, leading to high growth of consumption over time.

Real interest rates are more sensitive to consumption growth if γ is large. If the utility function is highly curved, the investor cares a lot about smoothing consumption. Thus it takes a larger interest rate change to induce him to a given consumption growth.

Uncertainty

Let's start with some simple computation in the case of *iid* lognormal consumption growth:

$$\begin{aligned}\Delta \log C_t &= \mu + \sigma \varepsilon_t, \\ \varepsilon_t &\text{ iid } N(0, 1).\end{aligned}$$

Recall that if X is normal with mean μ and variance σ^2 , then $Ee^X = e^{\mu + \frac{\sigma^2}{2}}$.

The risk-free rate is known in advance by definition, so we can 'pull it out' of the expectation:

$$\begin{aligned}R_{f,t+1} E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right) &= 1. \\ R_{f,t+1} &= E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right)^{-1} \\ R_{f,t+1} &= \frac{1}{\beta} e^{\gamma\mu - \frac{\gamma^2\sigma^2}{2}} \\ \log R_{f,t+1} &\simeq r_{f,t+1} \simeq -\log \beta + \gamma\mu - \frac{\gamma^2\sigma^2}{2}\end{aligned}$$

Comments:

(1) The risk-free rate is constant.

(2) Risk free rate determined by (a) impatience β , (b) intertemporal substitution γ and growth rate μ , [as in the case with certainty] and (c) "precautionary savings". More generally, if $\Delta \log C_{t+1}$ is *conditionally* lognormally distributed with mean $E_t \Delta \log C_{t+1}$ and $Var_t \Delta \log C_{t+1}$, the risk-free rate is:

$$r_{f,t+1} \simeq -\log \beta + \gamma E_t \Delta \log C_{t+1} - \frac{\gamma^2 Var_t \Delta \log C_{t+1}}{2},$$

hence a higher growth or lower uncertainty increases the risk-free rate. Note that for low risk aversion the precautionary savings effect appears to be small: since σ is about 2% (if period = one year), $\frac{\gamma^2\sigma^2}{2} = \frac{0.02^2}{2} = 0.0002 = 0.02\% = 2$ basis points if $\gamma = 1$ (log utility).

When the risk-free rate is constant, the term structure is completely flat, at a constant level: this is because the price of a N -period bond (a security which promises one unit of consumption in N periods) is pinned down by arbitrage. Hence, bonds of *all maturities* (not just one period bonds) are risk-free. (We'll go back to this when we discuss bonds, but you can prove this by brute force by calculating the price of an asset which pays off 1 at $t + 2$ instead of $t + 1$).

Stock Price

Define a stock as an asset which pays a dividend $D_t = C_t$ in all states of the world. (We'll discuss later why this may be a reasonable proxy.)

$$P_t = E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right)$$

$$\frac{P_t}{D_t} = \frac{P_t}{C_t} = E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \left(\frac{P_{t+1}}{D_{t+1}} + 1 \right) \right)$$

Iterating forward:

$$\begin{aligned} \frac{P_t}{C_t} &= \sum_{k \geq 1} E_t \left(\beta^k \left(\frac{C_{t+k}}{C_t} \right)^{1-\gamma} \right) \\ &= \sum_{k \geq 1} \beta^k E_t \left(\left(\frac{C_{t+k}}{C_{t+k-1}} \right)^{1-\gamma} \times \dots \times \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \right) \\ &= \sum_{k \geq 1} \beta^k \left(e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}} \right)^k \\ \frac{P_t}{C_t} &= \frac{\beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}}{1 - \beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}}. \end{aligned}$$

(Note how I used independence and then the lognormal formula.) Conclusion: *the P/D ratio is constant*. That's because the expectations of future consumption growth is independent of today's consumption. More generally, when expectations of future consumption and dividend growth are independent of the state today, and utility is CRRA, the P/D ratio is constant.

The expected return on equity is:

$$\begin{aligned} E_t \left(\frac{P_{t+1} + D_{t+1}}{P_t} \right) &= E_t \left(\frac{\frac{P_{t+1}}{D_{t+1}} + 1}{\frac{P_t}{D_t}} \frac{D_{t+1}}{D_t} \right) \\ &= \frac{\frac{\beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}}{1 - \beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}} + 1}{\frac{\beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}}{1 - \beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}}} E_t \left(\frac{D_{t+1}}{D_t} \right) \\ &= \frac{e^{\mu + \frac{\sigma^2}{2}}}{\beta e^{(1-\gamma)\mu + (1-\gamma)^2 \frac{\sigma^2}{2}}} \\ &= \frac{1}{\beta} e^{\gamma\mu + (2\gamma - \gamma^2) \frac{\sigma^2}{2}}. \end{aligned}$$

The geometric risk premium is:

$$\frac{ER^e}{R^f} = \frac{\frac{1}{\beta} e^{\gamma\mu + (2\gamma - \gamma^2) \frac{\sigma^2}{2}}}{\frac{1}{\beta} e^{\gamma\mu - \gamma^2 \frac{\sigma^2}{2}}} = e^{\gamma\sigma^2}$$

For the US, $\sigma = 2\%$ so this number is small, about $\gamma \times .0004$ or 0.4 basis points for log utility. That's the equity premium puzzle of Mehra and Prescott.

Realized returns: the return itself is

$$R_{t+1} = \frac{\frac{P_{t+1}}{D_{t+1}} + 1}{\frac{P_t}{D_t}} \frac{D_{t+1}}{D_t} = \frac{1}{\beta} e^{(\gamma-1)\mu + (2\gamma - \gamma^2 - 1) \frac{\sigma^2}{2}} e^{\mu + \sigma\varepsilon_{t+1}}.$$

so returns are iid lognormal.

Volatility Puzzle

In the standard model, the ratio of P_t to $D_t (= C_t)$ is constant. The data counterpart (eg the price-dividend ratio) is very volatile. More generally, since the P/D ratio is constant, the volatility of the return R_{t+1} is approximately equal to the volatility of dividend growth:

$$R_{t+1} = \frac{\frac{P_{t+1}}{D_{t+1}} + 1}{\frac{P_t}{D_t}} \frac{D_{t+1}}{D_t}$$

and if P_t/D_t is constant and large, then $\sigma(R_{t+1}) \simeq \sigma(D_{t+1}/D_t)$. This is also very counterfactual.

Cross-Sectional Tests

The model implies that the return on an asset depends on the covariance with consumption growth:

$$\begin{aligned} E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} (R_{t+1} - R_{t+1}^f) \right) &= 0, \\ E_t (R_{t+1} - R_{t+1}^f) E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right) &= -Cov \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}, R_{t+1} - R_{t+1}^f \right) \\ E_t (R_{t+1} - R_{t+1}^f) &= -R_{t+1}^f Cov \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}, R_{t+1} - R_{t+1}^f \right) \end{aligned}$$

Intuition: insurance: want assets which pay off in bad times i.e. when consumption growth is low. If the covariance is positive, the asset provides insurance, everybody wants to buy it and as a result it has a low expected return (or a high price, which is the same thing given that we're holding the payoff fixed: $R_t = E_t(X_{t+1})/P_t$). This implication can be tested in the data and does not work well. See the pb set 1 which will ask you to work out in detail an example of this.

A slightly more general formula

Start from:

$$E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{i,t+1} \right) = 1$$

Assume $(\log R_{t,t+1}, \Delta \log c_{t+1})$ is jointly lognormal, then:

$$E_t e^{\log \beta - \gamma \Delta \log c_{t+1} + \log R_{t+1}} = 1$$

$$e^{\log \beta - \gamma E_t \Delta \log c_{t+1} + E_t \log R_{t+1} + \frac{\gamma^2}{2} Var_t \Delta \log c_{t+1} + \frac{1}{2} V_t \log R_{t+1} - \gamma Cov(\Delta \log c_{t+1}, \log R_{t+1})} = 1$$

$$0 = \log \beta - \gamma E_t \Delta \log c_{t+1} + E_t \log R_{t+1} + \frac{\gamma^2}{2} Var_t \Delta \log c_{t+1} + \frac{1}{2} Var_t \log R_{t+1} - \gamma Cov(\Delta \log c_{t+1}, \log R_{t+1})$$

Write the same equation for the risk-free rate:

$$\log R_{t+1}^f = -\log \beta + \gamma E_t \Delta \log c_{t+1} - \frac{\gamma^2}{2} Var_t \Delta \log c_{t+1}$$

Subtract from the previous equation, this yields the CCAPM (Consumption Capital Asset Pricing Model).

$$E_t \log \left(\frac{R_{t+1}^i}{R_{t+1}^f} \right) + \frac{1}{2} Var_t \log R_{t+1}^i = \log E_t \left(\frac{R_{t+1}^i}{R_{t+1}^f} \right) = \gamma Cov(\Delta \log c_{t+1}, \log R_{t+1}^i),$$

i.e. differences in average returns relative to the risk-free rate should be explained solely by the covariance with consumption growth. (Note the assumption of joint lognormality is not innocuous: you wouldn't want to apply this to a very nonlinear asset such as an option.)

In our example above we had

$$Cov(\Delta \log c_{t+1}, \log R_{t+1}^i) = Var(\Delta \log c_{t+1})$$

since R_{t+1}^i is proportional to consumption growth. We can consider a more general model, e.g. if the “equity asset” pays a flow of dividends such that

$$\Delta \log D_{t+1} = \mu_D + \sigma_D \varepsilon_{t+1},$$

while

$$\Delta \log C_{t+1} = \mu_C + \sigma_C \varepsilon_{t+1},$$

with e.g. $\sigma_D > \sigma_C$, can show that the P-D ratio is constant and the return is proportional to dividend growth, leading to

$$Cov(\Delta \log c_{t+1}, \log R_{t+1}^i) = \sigma_C \sigma_D.$$

1.5 The present value model with constant discount rates

This model is useful because many people use it implicitly, and it has a historical role when discussing the “efficient market hypothesis”. It’s part of the basic language and allows us to think about some basic experiments. It’s also often used as synonymous with “efficient markets” (but this is not quite correct). Consider an asset which pays a dividend flow $\{D_t\}$, and assume that this asset has a constant expected rate of return r : for all $t \geq 0$,

$$1 + r = E_t \left(\frac{P_{t+1} + D_{t+1}}{P_t} \right).$$

Rewrite this as

$$P_t = \frac{E_t(D_{t+1})}{1+r} + \frac{E_t(P_{t+1})}{1+r}.$$

We can iterate this equation forward, assuming that $\lim_{s \rightarrow \infty} E_t(P_{t+s})/(1+r)^s = 0$, to get:

$$P_t = \sum_{k=1}^{\infty} \frac{E_t(D_{t+k})}{(1+r)^k}.$$

i.e. price = expected present discounted value of dividends. The discount rate r that we must apply is

the expected return on the asset. (Note: you can generalize this formula to the case where r varies over time.)

Use this formula to discuss the effect of various kind of news about future dividends on prices and returns:

- (i) an increase in D_{t+1} by 1\$, announced at time t ;
- (ii) an increase in D_{t+k} by 1\$, announced at time $t + 2$;
- (iii) an increase in D_{t+1}, D_{t+2}, \dots , and so on by 1\$, at all future dates, announced at time t .
- (iv) an increase in D_{t+k} of 1\$ combined with a decrease of D_{t+k+1} of $(1+r)$ \$, both announced at time $t + 2$.

ADD PLOTS IN CLASS

Plot the time path of prices, dividends, and returns in each case.

What happens if the discount rate r suddenly changes?

[Note: the “standard model” discussed above can be used to find r .]

Question: suppose that D_{t+1} is an AR(1), what is the form of P_{t+1} and $\frac{P_{t+1}}{D_{t+1}}$? What if $\frac{D_{t+1}}{D_t}$ is an AR(1)?

Gordon formula

A special case is the Gordon formula: assume that dividends grow at an average rate g , i.e. $E_t\left(\frac{D_{t+1}}{D_t}\right) = 1 + g$, then

$$P_t = \sum_{k=1}^{\infty} \frac{E_t(D_{t+k})}{(1+r)^k} = \sum_{k=1}^{\infty} \frac{(1+g)^{k+1}D_t}{(1+r)^k} = D_t \frac{1+g}{1 - \frac{1+g}{1+r}} = D_t \frac{(1+g)(1+r)}{r-g}.$$

For small r and g (or in continuous time), we get

$$\frac{P_t}{D_t} = \frac{1}{r-g}.$$

Note that we need $r > g$ for this sum to be finite! Note that if $r - g$ is small, a small change in either r or g leads to large price fluctuations.

Example: suppose that $\frac{D_{t+1}}{D_t} = 1 + g + \varepsilon_{t+1}$, with ε_{t+1} iid $N(0, 1)$. What is the effect of a good realization of ε_{t+1} on the dividend, on the price, on the P-D ratio, and on the return?

2 General results on stochastic discount factor: existence, uniqueness. Hansen-Jagannathan bounds.

So far we have looked at the Euler equation, derived for any agent from its FOC given available returns. Here we take a step back and look at this in more generality. This will turn out to provide some useful results before we go into building models. In this lecture I introduce the *stochastic discount factor* approach, which main benefit is to have results under weak assumptions.

A stochastic discount factor (SDF) is a stochastic process (i.e., sequence of random variables) $\{M_{t,t+1}\}$ such that for any security with stochastic payoff x_{t+1} at time $t+1$, the price of that security at time t is:

$$p_t = E_t(M_{t,t+1}x_{t+1}). \quad (3)$$

Or equivalently, for any asset return $R_{t,t+1}$, the following condition holds:

$$E_t(M_{t,t+1}R_{t,t+1}) = 1.$$

(This is equivalent by definition of the return = payoff / price: $R_{t,t+1} = x_{t+1}/p_t$.)

Notation: often we use only one index, the end of period, i.e. we write $E_t(M_{t+1}R_{t+1}) = 1$.

Example: in a representative agent economy with preferences $E \sum_{t=0}^{\infty} \beta^t u(c_t)$, the SDF is:

$$M_{t,t+1} = \frac{\beta u'(c_{t+1})}{u'(c_t)}. \quad (4)$$

What if there is more than one agent? It depends on the completeness of markets. If markets are complete, we show that the SDF exists and is unique. (Indeed, when markets are complete, you can define an “aggregate utility function” and the formula 4 applies for this utility function.) If markets are incomplete, but there is “no arbitrage opportunity”, there always exists a positive SDF, but there can be several SDFs which all satisfy equation 3.

Further, we derive some results which tell us what a “good” SDF should “look like”, i.e. what properties the process $\{M_{t+1}\}$ needs to possess to make the model fit the data. This can serve as a “diagnostic test” that can be used to see the properties that *any* SDF must satisfy, given the observed properties of asset returns. The main result is the Hansen-Jagannathan (JPE 1991) volatility bound: the SDF must be quite volatile. For instance, given the volatility of consumption growth, this tells us that the utility function must be very curved to generate a highly volatile SDF. See Cochrane and Hansen (1992 NBER macro) or Alvarez and Jermann (Econometrica 2004, “the persistence..”) for additional general implications. In particular, Alvarez and Jermann establish (roughly) that the SDF must have permanent innovations.

Basic properties of the SDF

The gross risk-free rate is the inverse of the conditional expectation of a SDF: just pull out the risk-free return (it is known at time t , by definition):

$$R_{t+1}^f E_t(M_{t+1}) = 1 \implies R_{t+1}^f = \frac{1}{E_t(M_{t+1})}.$$

The CAPM can be stated as saying that the log SDF is a linear function of the market return: $\log M_{t,t+1} = a - bR_{t+1}^m$. The Consumption CAPM (aka CCAPM) of course implies $\log M_{t,t+1} = \log \beta - \gamma \Delta \log C_{t+1}$ using equation 4.

Note that we always have the relations, for any returns i and the risk-free rate:

$$E_t(M_{t+1}R_{t+1}^i) = 1$$

$$E_t \left(M_{t+1} R_{t+1}^f \right) = 1$$

hence:

$$E_t(R_{t+1}^e - R_{t+1}^f) = \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})} \text{Corr}_t(M_{t+1}, R_{t+1}^e) \sigma_t(R_{t+1}^e)$$

the equity risk premium depends on the *market price of risk* $\frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})}$, the conditional volatility of the return $\sigma_t(R_{t+1}^e)$, and the conditional correlation of the SDF and return. Assets which have high average excess returns (i.e. $E_t(R_{t+1}^e - R_{t+1}^f) > 0$) must have volatile returns ($\sigma_t(R_{t+1}^e) > 0$) and/or their returns must be strongly correlated with the SDF ($\text{Corr}_t(M_{t+1}, R_{t+1}^e) > 0$). This is an equilibrium relation: if someone is holding these assets, it must be that their return compensates for their risk. Here risk means covariance with the SDF. The SDF measures “bad economic times”, i.e. the SDF is high in “bad” states of the world (In the simple representative agent, the SDF is marginal utility growth, and marginal utility is high in bad times.)

Construction of a SDF with complete markets, and risk-neutral probabilities. For simplicity consider a one-period world: there are S possible states, indexed by s , each with probability π_s , and we trade before the resolution of uncertainty. Let q_s be the state-contingent price (the price of an Arrow-Debreu security which pays one unit of good in state s). Consider a security with payoffs $\{d_s\}_{s=1}^S$. With complete markets, we can find the price of this security by “unpacking” it:

$$p(\{d\}) = \sum_{s=1}^S q_s d_s.$$

To go from this expression to the SDF notation requires just a bit of manipulation:

$$\begin{aligned} p(\{d\}) &= \sum_{s=1}^S \pi_s \frac{q_s}{\pi_s} d_s \\ &= E_\pi(md), \end{aligned}$$

i.e. the price is the expectation, under the true probability distribution π , of the payoffs times the SDF. Note we defined the SDF state by state by $m_s = q_s/\pi_s$.

Define $\hat{\pi}_s = \frac{\pi_s m_s}{\sum_k \pi_k m_k} = \pi_s m_s / E_\pi(m)$. By definition $\hat{\pi}_s$ lies between 0 and 1 and $\sum_{s=1}^S \hat{\pi}_s = 1$ so it is a probability distribution. Then

$$\begin{aligned} p(\{d\}) &= \sum_{s=1}^S \pi_s m_s d_s \\ &= E_\pi(m) \sum_{s=1}^S \hat{\pi}_s d_s \\ &= \frac{1}{R^f} E_{\hat{\pi}}(d). \end{aligned}$$

i.e. we get that the price is simply the discounted expected value of payoffs, under the probability distribution $\hat{\pi}$. This probability distribution is called the *risk-neutral measure*. Hence we have the usual “Price = PDV of dividends” condition which works after adjusting the probability measure. One interpretation of these risk neutral probabilities is that these are the subjective probabilities that the market assigns to states of nature, assuming investors care only about mean return (i.e. are risk-neutral). But this is really just vocabulary (often you will hear people in finance talk about “risk-neutral measure”).

Incomplete Markets: an example

Consider a one-period economy with N agents (indexed by i); there is uncertainty about the state s and we trade before that uncertainty is resolved. Each agent i is maximizing a general utility function

$$\max U_i(c_{i1}, \dots, c_{iS})$$

where c_s for $s = 1 \dots S$ is consumption in state s tomorrow. There are L securities available, with prices q_l and payoff d_{ls} in state s . We normalize the supply of each security to 1. Consumption is obtained as an endowment plus the payoffs of assets, i.e.

$$c_{is} = y_{is} + \sum_{l=1}^L \theta_{il} d_{ls},$$

where y_{is} is the endowment in state s and θ_{il} is the number of securities l bought by agent i . The budget constraint is

$$\sum_{l=1}^L q_l \theta_{il} \leq w_i,$$

where w_i is some initial endowment of wealth.

An equilibrium is a price vector $q = \{q_l\}$ and shares $\{\theta_{il}\}_{i,l}$ such that each agent maximizes and there is market-clearing for each security:

$$\forall l = 1 \dots L : \sum_{i=1}^I \theta_{il} = 1.$$

To characterize the equilibrium, we can consider the FOC of agent i wrt θ_{il} :

$$\sum_{s=1}^S \frac{\partial u_i}{\partial c_s} d_{ls} = q_l \lambda_i,$$

where λ_i is the multiplier on the wealth constraint. (This assumes an interior solution.)

Thus for any investor we have the Euler equation

$$\sum_{s=1}^S \frac{\partial u_i}{\partial c_s} \frac{d_{ls}}{\lambda_i q_l} = 1,$$

and so we can use as SDF the marginal rate of substitution of **any consumer who is unconstrained**. But it need not be true that any agent is unconstrained in all states. We now prove that in general, there exists at least one SDF, provided that there are no arbitrage opportunities.

Incomplete markets: No Arbitrage Pricing

Definition: the system of securities characterized by $\{d_{ls}\}$ and $\{q_l\}$ is arbitrage-free (or: there is no arbitrage opportunity, NOAO) if there is no vector of portfolio choices $\{\theta_l\}$ such that both

$$(1) \sum_{l=1}^L q_l \theta_l \leq 0 \text{ (i.e. } \theta \text{ is free, or has a negative price),}$$

and (2): $\forall s = 1 \dots S, \sum_{l=1}^L d_{ls} \theta_l \geq 0$, and > 0 for some state s (i.e. θ gives positive payoffs in all states).

Clearly, if the prices $\{q_l\}$ result from a competitive equilibrium as in the previous section, it must be that there is NOAO (otherwise consumers would demand an infinite amount of the combination which is free and gives a positive payoff.) But NOAO is more general and as a result requires very few assumptions.

Consequence of no arbitrage: assume you have two assets x and y and that in all possible states, x will pay you more dividends than y . Then the price of x should be (weakly) greater than the price of y . Hence, this basic principle imposes restrictions across prices of securities. In some cases, NOAO allows you to pin down exactly the price of some other securities. Even w/o complete markets, assets that are *redundant*, i.e. which can be replicated by existing assets, can be priced if we assume no arbitrage.

Example: static economy with 7 possible states, 3 assets are traded, payoff matrix listing in each row the payoff of each asset in each of the states:

$$D = \begin{pmatrix} 1 & 2 & 1 & 0 & 0 & 0 & 0 \\ 1 & 3 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Here, the third asset = the second asset minus the first one. Thus its price is the difference of the prices of the first two assets. An important practical application of NOAO is option pricing: Black and Scholes noted that the payoff of an option is (roughly, for small time intervals) a linear combination of the payoff of a stock and a bond, hence its value can be deduced from these two prices.

General Results on SDF

There are three important general results about the SDF:

- (1) The law of one price holds iff there is (at least) one stochastic discount factor.
- (2) There are no arbitrage opportunities iff there exists (at least) one strictly positive SDF.
- (3) The stochastic discount factor is unique iff markets are complete.

Here are sketches of the proofs. To avoid technicalities, I will again consider only one-period economies. (See Cochrane or Duffie, "Dynamic Asset Pricing Theory" chapter 1 for more details).

(1) The law of one price (LOOP) says that the price of a bundle of two securities is the sum of the prices of each security, i.e. $P(x + y) = P(x) + P(y)$. (We also assume $P(\lambda x) = \lambda P(x)$ where λ is a real number - we can divide securities.) This means that the function which assigns a price to a payoff is linear. A general theorem states that linear functions can be represented as an inner product, i.e. $P(x) = (x, w) = \sum_s w_s x_s$ from where we get a SDF by defining $m_s = w_s / \pi_s : P(x) = \sum_s \pi_s m_s x_s$.

(2) No arbitrage opportunities means that there is no trade which gives a sure profit, i.e. if $x \geq 0$ then $p(x) \geq 0$ and if $x_s > 0$ for some state, then $p(x) > 0$. The proof follows from a hyperplane separating theorem. (See Duffie.)

(3) The price of an Arrow-Debreu security which pays one unit of good in state s is $p_s = m_s \pi_s$. If markets are complete, such a security exists and there cannot be two SDF since both imply $p_s = m_s \pi_s = m_s^* \pi_s \rightarrow m_s = m_s^*$ for all states s . Conversely, if there is a unique SDF, markets must be complete. To see this, suppose that markets are incomplete, so there is a state j which is not tradable. Then, take a SDF m_s . We can construct another SDF m_s^* , which shows that there is no uniqueness. We simply set $m_s^* = m_s$ for $s \neq j$, and $m_j^* = m_j + 1$. By definition, m_s^* still satisfies the pricing equation (it prices any asset spanned by all the states but s).

Hansen and Jagannathan (JPE 1991)

HJ show that to be consistent with observed market returns, the SDF must be quite volatile. The basic result is actually simple. Consider any return R and the risk-free return R^f , we have, using that

$$-\sigma(x)\sigma(y) \leq Cov(x, y) \leq \sigma(x)\sigma(y) :$$

$$\begin{aligned} E(mR) &= 1 \\ E(m)R^f &= 1 \\ E(mR) &= E(m)R^f \\ Cov(m, R) &= E(m)R^f - E(m)E(R) \\ -\sigma(m)\sigma(R) &\leq E(m)(R^f - E(R)) \leq \sigma(m)\sigma(R) \\ -\frac{\sigma(m)}{E(m)} &\leq \frac{E(R) - R^f}{\sigma(R)} \leq \frac{\sigma(m)}{E(m)}. \end{aligned}$$

The number $\sigma(m)/E(m)$ is a measure of the volatility of the SDF called the *market price of risk*. The quantity $\frac{E(R)-R^f}{\sigma(R)}$ is called the *Sharpe Ratio* of a security and measures its return relative to its volatility. Hence given observed asset returns, we can find out how volatile the SDF must be. Recall that $E(m)$ is the inverse of the risk-free rate so we also have a good idea of what it is.

For the whole US stock market, the mean excess return is $E(R) - R^f$ approximately 6 to 8%. The standard deviation is roughly 16%, so overall the Sharpe ratio is 0.4 to 0.5. If you look at subsets of the stock market, you can relatively find Sharpe ratios of 1. For instance, the return on a portfolio long value stocks and short growth stocks is roughly 6% per year, and its volatility is about 8%.

These ‘‘Hansen-Jagannathan bounds’’ are popular because any model that does not satisfy them will be at odds with the data. This is a very weak test which you can easily do, for instance pick a utility function, plug the consumption data (or other data) in it, and see if the resulting m is volatile enough. If not, you can reject the model right away!

One can use these bounds for any asset. However a more refined version is to use data on returns *jointly*, i.e. given that you know that for $i = 1 \dots N$:

$$E(mR^i) = 1,$$

what can you say about the volatility of m ? HJ 1991 work out the implications of this system of equalities, which are of course stronger than the inequality you get from any single equation. For $N = 2$, i.e. 2 assets, you obtain a parabola in the $E(m)/\sigma(m)$ space. (e.g. see the figure in Tallarini 2000.) [ADD PICTURE IN CLASS.]

Review of General Equilibrium, Aggregation, Risk Sharing

These notes give some background on risk-sharing and the conditions under which a representative agent exists. I will not go over these notes in class, but they are useful background.

Review of General Equilibrium

We are applying the general equilibrium (GE) model of Arrow and Debreu to economies with time and uncertainty. This will allow us to justify under some conditions the use of a representative agent. In GE theory, an economy is defined by:

(1) a commodity space, i.e. the set of goods in the economy; mathematically a subset X of some vector space.

(2) a number of consumers $i = 1 \dots I$, each of which has some preferences U_i over consumption bundles in X ; each consumer also has an endowment $e_i \in X$ and shares in the technology j , θ_{ij} .

(3) a number of firms $j = 1 \dots J$, each of which has a technology defined by a production set $Y_j \subset X$.

The consumer i 's problem is:

$$\begin{aligned} & \max_{x \in X} U_i(x) \\ \text{s.t.} \quad & p \cdot x \leq p \cdot e_i + \sum_{j=1}^J \theta_{ij} \pi_j(p), \end{aligned}$$

yielding a net demand $x_i(p)$. (Note that $x_i(p)$ is a vector, the dimension of which is the dimension of the commodity space.)

The firm j 's problem is:

$$\pi_j(p) = \max_{y \in Y_j} p \cdot y,$$

yielding a net supply $y_j(p)$. ($\pi_j(p)$ is a scalar.)

A *competitive equilibrium* is a price vector p such that:

$$\sum_{i=1}^I x_i(p) = \sum_{i=1}^I e_i + \sum_{j=1}^J y_j(p).$$

i.e.: (i) each consumer solves his program, (ii) each firm solves his program, (iii) all markets clear.

Implicit in this definition is the fact that there are complete markets, i.e. one market for each good, and each consumer can buy or sell whatever amounts he wants in each market, subject only to his resource constraint.

A *feasible allocation* is a list of consumption vectors for each consumer, and of firm plans for each firm, that satisfies the aggregate resource constraint; mathematically $\{x_i\}, \{y_j\}$ such that

$$\sum_{i=1}^I x_i = \sum_{i=1}^I e_i + \sum_{j=1}^J y_j.$$

A *Pareto-optimal* allocation is a feasible allocation such that there is no feasible allocation that makes all consumers as well off, and at least one strictly better off.

An allocation $\{x_i\}, \{y_j\}$ is Pareto-optimal allocation iff there exists a set of Pareto weights $\{\lambda_i\}$ with $\lambda_i \geq 0$ for all i and $\lambda_i > 0$ for some i , such that $\{x_i\}, \{y_j\}$ solves

$$\begin{aligned} & \max_{\{x_i\}, \{y_j\}} \sum_{i=1}^I \lambda_i u(x_i) \\ \text{s.t.} \quad & \sum_{i=1}^I x_i \leq \sum_{i=1}^I e_i + \sum_{j=1}^J y_j. \end{aligned} \tag{5}$$

I now state the three key results. All rely on complete markets, i.e. there is a market for each good.

(1) First Welfare Theorem: if utility functions are strictly monotonic, any competitive equilibrium is a Pareto optimum.

(2) Second Welfare Theorem: under convexity assumptions¹, any Pareto optimum is a competitive equilibrium for some initial redistribution of wealth.

¹i.e. if utility functions are quasi-concave and production sets are convex.

(3) Existence of equilibrium: under convexity assumptions, there exists at least one competitive equilibrium.

Application to Economies with Time and Uncertainty

To apply this theory to dynamic models, we only need to reinterpret the commodity space. Now a good is defined not only by its physical characteristics, but also by the date and state of nature in which it is available.

Example 1: a two-period endowment economy. One consumer, with utility $u(c_1) + \beta u(c_2)$. The commodity space is $[0, +\infty)^2$. The price vector is (p_1, p_2) . There is no production. The income (endowment) is $(y_1, y_2) \in X$. The budget constraint is $p_1 c_1 + p_2 c_2 \leq p_1 y_1 + p_2 y_2$.

Example 2: an infinite-horizon economy without uncertainty. One consumer, with utility $\sum_{t=0}^{\infty} \beta^t u(c_t)$. X is the set of all sequences $\{c_t\}$ with nonnegative elements, and the price vector is $p = (p_0, p_1, \dots)$ also in the set of sequences of nonnegative numbers. The budget constraint is $\sum_{t=0}^{\infty} p_t c_t \leq \sum_{t=0}^{\infty} p_t y_t$.

Example 3: infinite horizon with uncertainty, the commodity space would include all state-contingent paths: $c_t(s^t)$, where s^t is the history (s_0, \dots, s_t) . See below.

Application of GE theory to dynamic stochastic models: Complete Markets and Consumption Allocation

Histories

We denote by $s \in S$ the “state of the economy” which keeps track of all the relevant information, e.g. income shocks, news about future income, etc.

Example 1: one agent, his income can be either high or low, then $S = \{high, low\}$ and $s = \{high\}$ or $s = \{low\}$.

Example 2: two agents A and B , each of which has an income which can be either high or low, then $S = \{hh, hl, lh, ll\}$.

s_0 is the initial known condition, and s_t is the realized state at time t . We denote a history by $s^t = (s_1, \dots, s_t) \in S^t$. The probability of history s^t occurring is denoted $\pi_t(s^t)$. Of course $\sum_{s^t \in S^t} \pi_t(s^t) = 1$ for all t , and for all histories s^{t-1} .

Example 3: Assume s_t is Markov with transition $P_{s, s'}$. Then $\pi_t(s^t) = P_{s_0 s_1} \times P_{s_1 s_2} \times \dots \times P_{s_{t-1} s_t}$.

A basic GE economy to study consumption insurance

(a) commodity space: the set of all possible stochastic processes: $\{c_t(s^t)\}$ all t, s^t .

(b) preferences: we will restrict preferences considerably, instead of defining them as functions of the entire stochastic processes, i.e. $U(\{c_t(s^t)\})$, I will assume that preferences satisfy (1) expected utility [i.e. separability across states of nature], and (2) time-separability:

$$\begin{aligned} U &= E \left[\sum_{t \geq 0} \beta^t u^i(c_t(s^t) \mid s_0) \right], \\ &= \sum_{t \geq 0} \sum_{s^t \in S^t} \beta^t \pi_t(s^t \mid s_0) u^i(c_t(s^t)). \end{aligned} \quad (6)$$

note that the utility functions can be different across agents, but I assume the beliefs are the same. There are I agents, indexed by i , with utilities $u^i(\cdot)$, endowments $e_{i,t}(s^t)$. Let $E_t(s^t) = \sum_{i=1}^I e_{i,t}(s^t)$ be the aggregate endowment at time t after history s^t .

(c) technology: this is an endowment economy. i.e. the only technology is free disposal.

Representative consumer

Assume markets are complete. Then any CE is PO. Any PO can be obtained from (10). Defining

$$V(E) = \max_{\{x_i\}} \sum_{i=1}^I \lambda_i u_i(x_i)$$

$$s.t. \quad : \quad \sum_{i=1}^I x_i \leq E,$$

we see that the economy is equivalent to one where there is only one consumer, with utility function V . This justifies using a representative consumer (RC).

Note: in general the utility function of the RC depends on the weights λ_i , which correspond to the initial wealth distribution. Only if all utility functions are homothetic (i.e., all income elasticities are unity) is the distribution of wealth always irrelevant to find the aggregate allocation and prices. (Of course the wealth distribution is always relevant to compute the individual allocation, but we may not care so much about it.) We will see an example below of representative consumer.

RESULT: if $\{x_i, y_j, p\}$ is an equilibrium for the economy $\{e_i, u_i, Y_j\}$, then $\{\sum x_i, y_j, p\}$ is an eq for the representative consumer economy.

Pareto Optima: Social Planning Problem

I solve for a competitive equilibrium with complete markets, so I can use a planner problem.

$$\max_{\{c_{i,t}(s^t)\}_{i,t,s^t}} \sum_{i=1}^I \lambda_i \sum_{t \geq 0} \sum_{s^t \in S^t} \beta^t \pi_t(s^t | s_0) u^i(c_{i,t}(s^t))$$

$$s.t. \quad : \quad \forall t, s^t : \sum_{i=1}^I c_{i,t}(s^t) = \sum_{i=1}^I e_{i,t}(s^t) = E_t(s^t).$$

The FOC w.r.t. $c_{i,t}(s^t)$ is (denoting $\mu_t(s^t)$ the Lagrange multiplier on the resource constraint):

$$\lambda_i \beta^t \pi_t(s^t) u'_i(c_{i,t}(s^t)) = \mu_t(s^t),$$

thus for any agent $i = 2 \dots I$

$$\frac{\lambda_i}{\lambda_1} \frac{u'_i(c_{it}(s^t))}{u'_1(c_{1t}(s^t))} = 1$$

$$\rightarrow c_{it}(s^t) = u_i'^{-1} \left(\frac{\lambda_1}{\lambda_i} u'_1(c_{1t}(s^t)) \right)$$

Plugging in the RC yields $c_{1t}(s^t)$. Clearly $c_{1t}(s^t)$ depends only on $E_t(s^t)$, not on the distribution of this endowment across agents $\{e_{i,t}\}$. This is the key result, stated in:

Perfect Risk Sharing and History-Independence

Under complete markets and expected discounted utility, there is perfect risk-sharing: the consumption of each agent depends only on the aggregate consumption, and not on his individual income. Moreover

allocations are independent of history, i.e. of the sequence of realization of shocks s^t , and depends only on the current realization of the endowment $E_t(s^t)$.

Decentralization: Competitive Equilibrium

We can decentralize this planner problem. We need one price for each good, i.e. each time-state combination. Let $q_t(s^t)$ = price at time 0 of unit of consumption at time t in state s^t . A competitive equilibrium is an allocation such that

(1) Taking as given the price vector $\{q_t(s^t)\}$, each agent $i = 1 \dots I$ maximizes his utility function (6) subject to the (time-zero) budget constraint:

$$\sum_{t \geq 0} \sum_{s^t \in S^t} q_t(s^t) c_{i,t}(s^t) \leq \sum_{t \geq 0} \sum_{s^t \in S^t} q_t(s^t) e_{i,t}(s^t).$$

(2) Markets clear: $\forall t \geq 0, \forall s^t \in S^t$:

$$\sum_{i=1}^I c_{i,t}(s^t) = \sum_{i=1}^I e_{i,t}(s^t) = E_t(s^t).$$

In class: check that this gives the same allocation.

Representative Consumer: two special cases

Assume that all agents have the utility $u(c) = c^{1-\gamma}/(1-\gamma)$. Agent i has endowment $\{e_{i,t}(s^t)\}$. Let the equilibrium allocation be $\{c_{i,t}(s^t)\}$ and the prices $\{q_t(s^t)\}$. Then $\{q_t(s^t)\}$ are equilibrium prices for the representative agent economy, with preferences $u(c) = c^{1-\gamma}/(1-\gamma)$ and endowment $\{\sum_i e_{i,t}(s^t)\}$.

Proof: to be done in class.

Note: you can similarly find the RC if utility are CARA, and agents differ in risk aversion.

Empirical Tests of Full Consumption Insurance

Cochrane (1991) and Mace (1991) were the first to test this implication with US data. They use the PSID to test for perfect consumption insurance by running regressions of the type

$$\Delta \log C_{i,t} = \alpha + \beta \Delta \log C_t + \gamma Z_{i,t} + \varepsilon_{i,t},$$

and testing for $\gamma = 0$. Z could be individual income growth, an unemployment indicator, a sickness indicator... This literature concludes that individual histories do matter, i.e. $\gamma \neq 0$, so there does not appear to be full insurance in the US.

Townsend (1994) test similarly for perfect insurance in Indian villages. Does an individual's consumption depends on his income, once you control for the village's total consumption. He finds that the complete model is rejected too, but by a narrow margin. Using data from Thailand the rejection is "larger".

Possible Extensions

If utility depends in a separable way of leisure (i.e. $u(c) + v(l)$), or something else (e.g. illness), the consumption predictions are unaffected. On the other hand, because the complete markets model makes marginal utilities of consumptions across agents all proportional across time and states, if *marginal utility of consumption* depends on something else besides consumption, then the complete markets model does not predict that consumption moves in lockstep between agents (e.g. $u(c, l)$).

Testing Complete Markets using Asset Prices

Under complete markets, we have a RC whose utility depends only on aggregate consumption. As we'll see in the next class, this imply we can use his utility function to price assets. It has proved very hard to find a utility function that rationalizes facts about asset prices using aggregate consumption only. This is an “indirect” rejection of complete markets.

3 Euler equation estimation and GMM; Campbell-Shiller decomposition

3.1 Euler equation tests

This lecture considers some further implications and tests of the “standard” asset pricing model: for any gross return R_{t+1} , we have the equation:

$$E_t \left(\frac{\beta u'(C_{t+1})}{u'(C_t)} R_{t+1} \right) = 1,$$

hence for any gross return R_{t+1} , we have

$$E_t \left(\frac{\beta u'(C_{t+1})}{u'(C_t)} (R_{t+1} - R_{t+1}^f) \right) = 0.$$

These equations hold for any asset which you can buy or sell (at the margin): stocks, bonds, options, commodities, etc.

Key intuition: assets are priced according to their covariance with consumption growth. Assets which pay off when consumption growth is high (in good times) are risky, hence they will need to have high expected returns for investors to want to hold them. Another way to say the same thing: these assets will trade at a discount: the price will be low to compensate for the riskiness. (Given some expected dividends, a lower price today implies a higher expected return.)

Today’s lecture = discuss how to test for the Euler equation using GMM. This leads to another rejection of the model. So in the rest of the course, we will study extensions of the model that try to improve its fit:

- either by changing preferences and going away from CRRA, expected utility (e.g. habits, Epstein-Zin, multiple goods);
- or by changing the endowment process (e.g. disasters, long-run risk);
- or by changing the market structure (short sales constraints, heterogeneous agents with incomplete markets).

From a macro point of view, these changes are going to affect the analysis of many standard issues:

- different preferences will affect business cycles, welfare costs, public finance (tax cut) estimates;
- endowments processes in the end mean different shocks and/or technologies;
- incomplete markets could have an even more substantial effect, e.g. on welfare.

The logic of Euler equation applies of course even to models others than the CRRA representative agent model. Many models can be tested by GMM, e.g. if you make up a utility function of consumption c_t and other observable stuff x_t , say $U(c_t, x_t)$, then we have

$$M_{t+1} = \frac{\beta U_1(c_{t+1}, x_{t+1})}{U_1(c_t, x_t)},$$

and $E_t(M_{t+1}R_{t+1}) = 1$ can similarly be tested.

3.2 Euler equation tests

Lecture 1 studied the implication of the Euler equation, plus an assumption that consumption growth is iid and dividends = consumption. This allows for a clean solution which leads to some intuition, and these assumptions are not too bad empirically², but they are not really necessary.

In contrast, the EE approach makes very little assumptions - no need to specify the shocks, the technology (as you would in a GE model), or the endowment process. Start from the fact that, for any return R_{t+1} ,

$$E_t \left(\frac{\beta u'(C_{t+1})}{u'(C_t)} R_{t+1} - 1 \right) = 0,$$

specify $u(C) = \frac{C^{1-\gamma}}{1-\gamma}$. For any variable Z_t known at time t , we have:

$$Z_t E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) = 0,$$

$$E_t \left(Z_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) \right) = 0,$$

$$E \left(Z_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) \right) = 0.$$

This is a set of restrictions that the theory implies on data on consumption, returns and Z_t . No need for any distributional assumption, except that returns and consumption growth must be stationary. Notice that if Z_t is uncorrelated with future returns and consumption, then it does not add any restriction since

$$E \left(Z_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) \right) = E(Z_t) E \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right).$$

So interesting Z_t are variables which forecast future returns or future consumption growth.

This equation holds for any returns R_{t+1}, Z_t , so we have NK equations, where $N = \#$ of returns and $K = \#$ of “instruments” Z_t , and we have 2 parameters (β, γ) . We can estimate using the GMM framework, see Cochrane’s book for a very nice discussion.

Intuitively, we simply pick β and γ to set the sample version of these equations equal to zero, i.e.

$$\frac{1}{T} \sum_{t=1}^T Z_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) = 0.$$

Different cases:

- No Z_t , and apply this to excess aggregate stock returns: pick γ such that

$$\frac{1}{T} \sum_{t=1}^T \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(R_{t+1}^e - R_{t+1}^f \right) = 0.$$

Note: this may not be possible! i.e. with the data there may not be a γ s.t. this holds.

²Consumption are not equal to dividends, but two effects offset: dividends are more volatile than consumption, but they are not perfectly correlated.

- No Z_t , find β and γ to match both stock returns and the risk-free rate:

$$\frac{1}{T} \sum_{t=1}^T \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1}^e - 1 \right) = 0,$$

$$\frac{1}{T} \sum_{t=1}^T \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1}^f - 1 \right) = 0.$$

Two eqns in two unknowns. Again, this may not be possible to find both β and γ in sample.

- If more than equations, then of course (generically) cannot set all the equations equal to zero. In practice (as in the other cases above if you can't find a zero), you minimize the weighted errors:

$$\min_{\beta, \gamma} g' W g,$$

where $g = \frac{1}{T} \sum_{t=1}^T Z_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right)$ and W is a weighting matrix.

- Can pick $W = I$, or W to have the moments have the same scale, or otherwise pick W in a statistically optimal way. Formula for W is the asymptotic variance,

$$W = \text{Var} \left(\frac{1}{T} \sum_{t=1}^T \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1} - 1 \right) \right).$$

See Cochrane's discussion of Hansen and Singleton in his survey paper, "Financial markets and the economy".

- Overidentification: if $NK > \#$ of parameters, test that $\frac{1}{T} \sum_{t=1}^T \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R_{t+1}^f - 1 \right)$ is close enough to 0 for all returns.
- Of course can apply this approach to any utility function.
- Hansen and Singleton (1983) take a different tack on this estimation: they assume joint log-normality of consumption growth and returns to derive a maximum likelihood estimator.

3.2.1 GMM general formulas

These are taken from Cochrane's book - this is just a summary.

- Express a model as $E[f(x_t, b)] = 0$.
- The general GMM estimate \hat{b} is given by:

$$a_T g_T(\hat{b}) = 0,$$

where

$$g_T(b) \equiv \frac{1}{T} \sum_{t=1}^T f(x_t, b),$$

and a_T is a matrix that defines which linear combination of $g_T(b)$ will be set to zero.

- Hansen (1982) tells us that the asymptotic distribution of the GMM estimate is:

$$\sqrt{T}(\hat{b} - b) \rightarrow N[0, (ad)^{-1}aSa'^{-1}],$$

where

$$\begin{aligned} d &\equiv E \left[\frac{\partial f}{\partial b'}(x_t, b) \right] = \frac{\partial g_T(b)}{\partial b'}, \\ a &\equiv a_T, \\ S &\equiv \sum_{j=-\infty}^{\infty} E[f(x_t, b)f(x_{t-j}, b)']. \end{aligned}$$

- Hansen (1982) gives us the sampling distribution of the moments $g_T(b)$:

$$\sqrt{T}g_T(\hat{b}) \rightarrow N[0, (I - d(ad)^{-1}a)S(I - d(ad)^{-1}a)'].$$

- The efficient estimate is obtained by setting:

$$a = d'^{-1}$$

In this case,

$$\begin{aligned} \sqrt{T}(\hat{b} - b) &\rightarrow N[0, (d'^{-1}d)^{-1}], \\ TJ_T = Tg_T(\hat{b})'^{-1}g_T(\hat{b}) &= \chi^2(\#moments - \#parameters), \\ TJ_T(restricted) - TJ_T(unrestricted) &= \chi^2(\#restrictions). \end{aligned}$$

- A familiar example: OLS.

$$\text{Min}_{\beta} E_T[(y_t - \beta'x_t)^2]$$

In this case:

$$\begin{aligned} f(x_t, \beta) &= x_t(y_t - \beta'x_t) = x_t\varepsilon_t, \\ g_T(\beta) &= E_T[x_t(y_t - \beta'x_t)], \\ a_T &= I, \\ d &= -E_T(x_t x_t'), \\ \hat{\beta} &= E_T[x_t x_t'^{-1} E_T(x_t y_t)] \\ \text{var}(\hat{\beta}) &= \frac{1}{T} E_T[x_t x_t'^{-1} \left[\sum_{j=-\infty}^{\infty} E(\varepsilon_t x_t x_{t-j}' \varepsilon_{t-j}') \right] E_T[x_t x_t'^{-1}] \end{aligned}$$

3.3 Campbell and Shiller decomposition

Campbell and Shiller (1988, Review of Financial Studies) introduced a log-linear approximation to the present-value identity (prices = Present Discounted Value [PDV] of dividends). They used this approximation to discuss the **sources of stock price volatility**. This formula has proven extremely useful in applied work, because it is easy to use it in conjunction with linear time series models (e.g. VARs). Also known as “Discount Rate - Cash Flow Decompositions”

Present Discounted Values

Start from the definition of return:

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}.$$

You can rewrite this as “price equal PDV of dividends” by rewriting and iterating forward:

$$\begin{aligned}
 P_t &= \frac{D_{t+1} + P_{t+1}}{R_{t+1}} \\
 &= \frac{D_{t+1}}{R_{t+1}} + \frac{D_{t+2}}{R_{t+1}R_{t+2}} + \dots \\
 &= \sum_{k=1}^{\infty} \frac{D_{t+k}}{R_{t,t+k}},
 \end{aligned}$$

where $R_{t,t+k}$ is the **realized** return on the asset (not the risk-free rate!!!) from t to $t+k$:

$$R_{t,t+k} = R_{t+1} \times \dots \times R_{t+k}.$$

Of course this equation, price = PDV of dividends, is for now an accounting identity - there’s no assumption about behavior. (Just a “no-bubble” condition to take out the limit: $\lim_{k \rightarrow \infty} \frac{D_{t+k}}{R_{t,t+k}} = 0$.) This equation is useful, but it is nonlinear. C-S find a way to linearize it.

The approximation. Use lowercase letters for logs: $r_t = \log(1 + R_t)$, $p_t = \log(P_t)$, $d_t = \log(D_t)$:

$$\begin{aligned}
 r_{t+1} &= \log(P_{t+1} + D_{t+1}) - \log(P_t) \\
 &= p_{t+1} + \log\left(1 + \frac{D_{t+1}}{P_{t+1}}\right) - p_t \\
 &= p_{t+1} - p_t + \log(1 + \exp(d_{t+1} - p_{t+1}))
 \end{aligned}$$

Now the key step: do a first-order Taylor approximation to the function $f(x) = \log(1 + \exp(x))$ around $x = \bar{d} - \bar{p}$ where \bar{a} denotes the sample average value:

$$\log(1 + \exp(d_{t+1} - p_{t+1})) \simeq k + (1 - \rho)(d_{t+1} - p_{t+1}),$$

with:

$$\begin{aligned}
 \rho &= \frac{1}{1 + \exp(\bar{d} - \bar{p})} \\
 k &= -\log \rho + (1 - \rho) \log(1/\rho - 1).
 \end{aligned}$$

This yields:

$$\begin{aligned}
 r_{t+1} &= \rho p_{t+1} - p_t + k + (1 - \rho)d_{t+1}, \\
 p_t &= \rho p_{t+1} + k + (1 - \rho)d_{t+1} - r_{t+1}.
 \end{aligned}$$

Iterating forward yields:

$$\begin{aligned}
 p_t &= \frac{k}{1 - \rho} + \sum_{j \geq 1} \rho^{j-1} ((1 - \rho) d_{t+j} - r_{t+j}), \\
 &= \frac{k}{1 - \rho} + \sum_{j \geq 0} \rho^j ((1 - \rho) d_{t+1+j} - r_{t+1+j}).
 \end{aligned}$$

This identity holds for any returns, prices and dividends ex-post. (And thus ex-ante, if you take conditional expectations of this equation.) It is an accounting identity without any behavioral assumption. Can rewrite it as the sources of volatility of P/D:

$$p_t - d_t = \frac{k}{1 - \rho} + \sum_{j \geq 0} \rho^j (\Delta d_{t+1+j} - r_{t+1+j})$$

This formula can also be expressed in terms of “return shocks”. Using the formula $r_{t+1} = \rho p_{t+1} - p_t + k + (1 - \rho)d_{t+1}$, and applying $E_{t+1} - E_t$ to both sides yields:

$$\begin{aligned}
r_{t+1} - E_t r_{t+1} &= \rho(p_{t+1} - E_t p_{t+1}) + (1 - \rho)(d_{t+1} - E_t d_{t+1}) \\
&= \rho(E_{t+1} - E_t) \sum_{j \geq 0} \rho^j ((1 - \rho) d_{t+2+j} - r_{t+2+j}) + (1 - \rho)(d_{t+1} - E_t d_{t+1}) \\
&= (E_{t+1} - E_t) \sum_{s \geq 1} \rho^s ((1 - \rho) d_{t+1+s} - r_{t+1+s}) + (1 - \rho)(d_{t+1} - E_t d_{t+1}) \\
&= (1 - \rho)(E_{t+1} - E_t) \sum_{s \geq 0} \rho^s d_{t+1+s} - (E_{t+1} - E_t) \sum_{s \geq 1} \rho^s r_{t+1+s} \\
&= (E_{t+1} - E_t) \sum_{s \geq 0} \rho^s \Delta d_{t+1+s} - (E_{t+1} - E_t) \sum_{s \geq 1} \rho^s r_{t+1+s}.
\end{aligned}$$

An unexpectedly good stock return must occur because either the current dividend went up, or expectations of future dividends go up, or because expectations of future returns go down. The first two terms are a standard “cash flow effect” and the second is an expected return or risk premium effect: the price goes up if the risk premium or risk-free interest rate go down.

Hence, stock price volatility can come from either volatility of future dividends or volatility of expected future returns. Which of these terms contribute more to volatility empirically? One way to go to the data is to fit a VAR to dividend growth and returns and use the VAR to compute the implied decomposition:

$$\begin{pmatrix} \Delta d_{t+1} \\ r_{t+1} \\ x_{t+1} \end{pmatrix} = A(L) \begin{pmatrix} \Delta d_t \\ r_t \\ x_t \end{pmatrix} + \begin{pmatrix} \varepsilon_{t+1}^d \\ \varepsilon_{t+1}^r \\ \varepsilon_{t+1}^x \end{pmatrix},$$

and iterating on this VAR one can compute the forecast $E_t \Delta d_{t+1+j}$ and then the revision of the forecast. By stacking the vector $\begin{pmatrix} \Delta d_{t+1} \\ r_{t+1} \\ x_{t+1} \end{pmatrix}$ and its lags, we can assume the VAR has an order 1; so $z_{t+1} = Az_t + \varepsilon_{t+1}$ with $E_t \varepsilon_t' = \Sigma$. As a result, $E_t \sum_{j \geq 0} \beta^j z_{t+j} = \sum_{j \geq 0} \beta^j A^j z_t = (I - \beta A)^{-1} z_t$ and by premultiplying by the correct row vector, you obtain the forecast of the PV of the component of z that you want. Moreover,

$$\begin{aligned}
(E_{t+1} - E_t) \sum_{j \geq 0} \beta^j z_{t+j} &= z_t + \beta(I - \beta A)^{-1} z_{t+1} - (I - \beta A)^{-1} z_t \\
&= z_t + \beta(I - \beta A)^{-1} (Az_t + \varepsilon_{t+1}) - (I - \beta A)^{-1} z_t \\
&= (I - \beta A)^{-1} (I - \beta A + \beta A - I) z_t + \beta(I - \beta A)^{-1} \varepsilon_{t+1} \\
&= \beta(I - \beta A)^{-1} \varepsilon_{t+1}.
\end{aligned}$$

(See the related Hansen-Sargent formula I attach below - no need to get all the details right now.)

It turns out the second term dominates by a very large margin. Roughly, in the data, it is hard to forecast dividends. Returns are predictable (a bit!) as we saw before. That’s why the second term dominates. Since the risk-free interest rate does not move much in the data, it means the changes in expected returns are mainly changes in risk premia. This is terribly important. It suggests we want to think of models where these risk premia change over time. One model is Campbell-Cochrane, which has time-varying risk aversion. Another model is to have stochastic volatility in consumption growth. Sometimes consumption growth is volatile (“uncertain times”), so the risk premium is high, and sometime it is low.

One can get more results by making additional assumptions about the asset pricing model and using this decomposition. See Campbell 1993 AER, 1996 JPE. I will summarize these papers later.

Historical note: the early literature on “excess stock price volatility” was assuming that the second term in the Campbell-Shiller decomposition was nil, i.e. assuming that discount rates are constant. For instance, the seminal paper by Shiller (1980 AER) was testing if prices are too volatile, assuming that the correct price should be $p_t = E_t \sum_{k \geq 0} \rho^k d_{t+k}$. Then, given $Var(E_t(X)) \leq Var(X)$ for any random variable X , Shiller tested if

$$Var(p_t) < Var\left(\sum_{k \geq 0} \rho^k d_{t+k}\right),$$

where the RHS can be computed given *realized* dividends. Shiller found that in the data, $Var\left(\sum_{k \geq 0} \rho^k d_{t+k}\right) < Var(p_t)$, and interpreted this as “excess stock price volatility”. However, it is unclear if his rejection is a rejection of “rationality” or of the constant discount rate model (see Cochrane’s 1991 JME review for a nice discussion). As Cochrane notes, there is now no reason to examine this particular implication rather than the Euler equation, which is more informative. (A side issue in this controversy was the stationarity of prices and the impact on statistical test; the reasonable way to do this test is to write this as a test of volatility of the p/d ratio.)

Cochrane’s decomposition: Cochrane (1992, RFS)

3.4 Hansen and Sargent (1980) Prediction Formulas

This is a useful formula which gives the expectation of a present value of a stochastic process, or revisions to this present value. Suppose we want to compute $E_t \sum_{j \geq 0} \beta^j x_{t+j}$. Write the Wold decomposition as $x_t = A(L)w_t$, where $A(L)$ is a lag operator. (For instance, $A(L) = 1 - \theta L$ and $x_t = w_t - \theta w_{t-1}$, or $A(L) = \frac{1}{1-\rho L}$ and $x_t - \rho x_{t-1} = w_t$.) Define the $[\cdot]_+$ operator as deletion of terms with negative powers, e.g. $[2L + 3 + 4L^{-1} - 3L^{-2}]_+ = 2L + 3$. Note that for any lag operator $B(L)$, we have $E_t(B(L)w_t) = [B(L)]_+ w_t$.

$$\begin{aligned} E_t \sum_{j \geq 0} \beta^j x_{t+j} &= E_t \sum_{j \geq 0} \beta^j A(L) \varepsilon_{t+j} \\ &= E_t \sum_{j \geq 0} \beta^j A(L) L^{-j} \varepsilon_t \\ &= \left[\sum_{j \geq 0} \beta^j A(L) L^{-j} \right]_+ \varepsilon_t \\ &= \left[\frac{LA(L)}{L - \beta} \right]_+ \varepsilon_t \\ &= \frac{LA(L) - \beta A(\beta)}{L - \beta} \varepsilon_t. \end{aligned}$$

The line before last computes the geometric series. The last line is more subtle: note that $\frac{LA(L) - \beta A(\beta)}{L - \beta} = D(L)$ is only a polynomial (i.e. positive lags), since we can factor a $L - \beta$ in the numerator. Since $\left[\frac{LA(L) - \beta A(\beta)}{L - \beta} \right]_+ = \left[\frac{LA(L)}{L - \beta} \right]_+ - \beta \left[\frac{1}{L - \beta} \right]_+$ and $\left[\frac{1}{L - \beta} \right]_+ = 0$, we have $\left[\frac{LA(L) - \beta A(\beta)}{L - \beta} \right]_+ = \frac{LA(L) - \beta A(\beta)}{L - \beta}$. (This is because $\frac{1}{L - \beta} = -\frac{1}{\beta} \frac{1}{1 - \frac{L}{\beta}}$ and since $1/\beta > 1$, this equation must be solved forward: $\frac{1}{1 - \frac{L}{\beta}} = \frac{L^{-1}}{L^{-1} - \frac{1}{\beta}} = \frac{L^{-1}\beta}{L^{-1}\beta - 1} = \frac{-L^{-1}\beta}{1 - L^{-1}\beta} = -L^{-1}\beta \sum_{k \geq 0} \beta^k L^{-k}$.)

Apply this formula to compute the expectation as of $t - 1$:

$$\begin{aligned}
E_{t-1} \sum_{j \geq 0} \beta^j A(L) \varepsilon_{t+j} &= E_{t-1} E_t \sum_{j \geq 0} \beta^j A(L) \varepsilon_{t+j} = E_{t-1} \frac{LA(L) - \beta A(\beta)}{L - \beta} \varepsilon_t \\
&= \frac{\frac{LA(L) - \beta A(\beta)}{L - \beta} - \frac{-\beta A(\beta)}{-\beta}}{L} \varepsilon_{t-1} = \frac{LA(L) - \beta A(\beta) - A(\beta)(L - \beta)}{L(L - \beta)} \varepsilon_{t-1} \\
&= \frac{A(L) - A(\beta)}{L - \beta} \varepsilon_{t-1}.
\end{aligned}$$

Thus we obtain a formula for the revision of expectations:

$$\begin{aligned}
(E_t - E_{t-1}) \sum_{j \geq 0} \beta^j A(L) \varepsilon_{t+j} &= \left(\frac{LA(L) - \beta A(\beta)}{L - \beta} - \left(\frac{A(L) - A(\beta)}{L - \beta} \right) L \right) \varepsilon_t, \\
&= A(\beta) \varepsilon_t.
\end{aligned}$$

This formula is much used, eg in the PIH-rational expectation literature (eg Quah (1989 JPE)). Note how in the case of Campbell-Shiller, $A(L) = (I - AL)^{-1}$ and $A(\beta) = (I - A\beta)^{-1}$.

4 Additional material: Disasters

Motivation from an empirical standpoint: lots of crashes in the stock market, as we saw recently. They tend to happen during big recessions / depressions. Are we capturing this well with our standard model?

Motivation from a logical/theoretical standpoint:

Given the failure of the standard model, we need to modify the standard model (CRRA representative agent model with iid log-normal consumption growth). Rather than changing preferences, we may change instead the endowment process. For instance, one may consider an AR(1):

$$\Delta \log C_{t+1} = \rho \Delta \log C_t + (1 - \rho)\mu + \sigma \varepsilon_{t+1}.$$

This is what Mehra-Prescott did in their 1985 JME paper, and you will study this in PS2. Another possibility is that there is a small risk of a very large decline in stocks, e.g.

$$\begin{aligned} \Delta \log C_{t+1} &= \mu + \sigma \varepsilon_{t+1}, \text{ with probability } 1 - p, \\ &= \mu + \sigma \varepsilon_{t+1} + \log(1 - b), \text{ with probability } p. \end{aligned}$$

where ε_{t+1} is iid $N(0, 1)$ and $0 < b < 1$ is the size of the disaster. Hence, in period $t + 1$, with probability p , consumption drops by a factor b . For instance if $b = .2$, then consumption drops by 20%. With this formulation, the drop is permanent. Overall, we still maintain the assumption of iid consumption growth, but the distribution is non-normal - specifically, it has negative skewness and positive kurtosis.

Given this process we can compute asset prices. The risk-free rate is

$$\log R^f = -\log \beta + \gamma\mu - \frac{\gamma^2 \sigma^2}{2} - \log(1 - p + p(1 - b)^{-\gamma}).$$

When $p = 0$, this formula collapses to the standard result of the iid lognormal model. When $p > 0$, we see that the risk-free rate is lower (since $0 < b < 1$), and the higher the probability of disaster, the lower the risk-free rate. This reflects that a higher probability of disaster reduces expected growth and increases risk, and thus increases the demand for savings, both for intertemporal substitution and for precautionary reasons. This drives the risk-free rate down. (Quantitatively, the risk effect is bigger than the substitution effect.)

For stocks, we can obtain the P-D ratio through the usual equation:

$$\frac{P_t}{D_t} = E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \left(\frac{P_{t+1}}{D_{t+1}} + 1 \right) \right).$$

Given that consumption growth is iid, the P-D ratio is constant, hence $P_t/D_t = q$ with

$$q = \beta e^{(1-\gamma)\mu + (1-\gamma)\frac{\sigma^2}{2}} (1 - p + p(1 - b)^{1-\gamma}) (1 + q),$$

and we have the log expected return on the equity:

$$\log(ER^e) = \gamma\mu - \frac{\gamma^2 \sigma^2}{2} + \gamma\sigma^2 - \log \beta + \log \frac{(1 - p + p(1 - b))}{(1 - p + p(1 - b)^{1-\gamma})}.$$

Last, the log equity premium is obtained as:

$$\log \left(\frac{ER^e}{R^f} \right) = \gamma\sigma^2 + \log \frac{(1 - p + p(1 - b)) (1 - p + p(1 - b)^{-\gamma})}{(1 - p + p(1 - b)^{1-\gamma})}.$$

The first term is, again, the standard term from the log-normal model. The second term is the risk premium driven by disaster risk: it is 0 if $p = 0$, and taking derivatives in this expression shows that this

is an increasing function of p when p is small enough.³ This term can be substantial, with CRRA utility a small chance of a very large consumption drop matters a lot more than frequent small shocks.

The economic intuition is the same as the standard model, but it works off large shocks: stocks do very poorly in disasters, i.e. in states when consumption growth is very low.

Notes on the literature

- Role of sample selection: no *data* in bad times/bad countries. So probably we overestimate the equity premium. Do we really believe that the premium is 7% for shocks with std dev of consumption 2% going forward? See Cochrane for a good discussion.
- But: disaster model is not sample selection! (There is a risk premium even if we observe the full sample. Of course, if we observe a sample where no disaster has occurred, we estimate a higher “risk premium”. This is what is called the Peso problem.)
- Before Barro (2006), people thought that the disaster size and probability had to be very large, and hence the explanation was not credible. Barro (2006) measures large output declines in over 20 countries during the XXth century, and finds that they are often large, and rather frequent. Subsequent work by Barro and Ursua looks at a larger set of countries, longer periods, and uses consumption data.
- However, there is still some debate on these numbers. Barro measured them from peak to trough, which is difficult to map into a model. More recent work by Barro, Nakamura, Steinsson and Ursua estimates a full model.
- Also: the disaster model can be tested using the Euler equation – more precisely, when we test the rep. agent model with CRRA utility we do not make any distributional assumption on consumption growth. (Julliard and Gosh.) The only problem may be that the sample does not include “enough” disasters.
- Can also test more precisely the disaster model using option data: should make deep out of the money put options very expensive to less deep out of the money options. Doesn’t seem to work very well if you assume constant prob of disaster and CRRA utility (Backus, Chernov and Martin), but may work better if one is a bit less
- Recoveries following disasters? Empirical debate (e.g. Reinhart and Rogoff, “the myth of economic recovery”, Barro, Nakamura et al., my AER P&P paper). Not even clear how this affects the results.
- Implications of the disaster model for the cross-section of asset returns?
- Note, however, that this model generates a constant P-D ratio as cons. growth is iid. One interesting extension is to make the probability (or size) of disaster time varying. This would generate a variation in risk premia over time. One challenge is to keep interest rate volatility limited, however. With CRRA, counterintuitive implications. An increase in prob of disaster raises stock prices, if $CRRA > 1$ ($IES < 1$). Similar to the effect of an increase in volatility of consumption in the standard lognormal model. Hence, need Epstein-Zin preferences with both IES and risk aversion > 1 . Wachter 2009, Gabaix 2008, Gourio 2008 FRL. Leverage helps also (dividends more volatile than consumption). Extension with production: Gourio 2010.

³Because disasters are a binomial variable, the uncertainty is highest for intermediate values of p , and hence the risk premium is not increasing over the entire range of values: if p is large enough, a further increase reduces the uncertainty and thus the risk premium. This remark is not important in practice because disasters are always calibrated as rare events.

5 Habits models

One reaction to the rejection of the CRRA representative agent model is to consider more general utility functions. A popular choice is models with habits: agents enjoy not the flow of consumption today, but its excess over a varying habit level, e.g.

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

with u given, for instance, by the formula

$$u(c, h) = \frac{(c - \theta h)^{1-\sigma}}{1 - \sigma},$$

where $\theta > 0$ is a parameter and h_t is the habit level.

The habit level h_t satisfies itself a law of motion, e.g. $h_t = (1 - \delta)h_{t-1} + c_{t-1} = \sum_{j=1}^{\infty} (1 - \delta)^{j-1} c_{t-j}$, i.e. it is a function of past consumption choices. There are different types of habits models, which can be classified depending on (1) the specific functional form for u , (2) the law of motion of h , (3) whether habits are internal or external. Internal means that the individual's own consumption enters the habit, and he takes this into account when making consumption decisions, whereas external means that h_t is exogenous to the consumer, and the c_t in the law of motion for h_t is the aggregate consumption; this last case is often referred to as “catching up with the Joneses” i.e. you care about your consumption relative to the average of other's people consumption (“conspicuous consumption”). This implies an externality - people consume “too much”.⁴

Habits can be defended as a “psychological law” (“repetition of a stimulus diminishes the response to it”; or potentially justified by natural evolution and selection⁵). Alternatively, habits can capture a time-varying subsistence level (as in Stone-Geary preferences), or again alternatively they can be seen as a “reduced form” for consumption commitments: agents cannot change their consumption easily, many of their expenses are fixed, and at the margin they must cut on the part of consumption which is flexible. (Some papers try to show that a model with consumption commitments leads to the same reduced form as habits; see e.g. Chetty and Saez, QJE 2007).

Microevidence for habits/catching up preferences is relatively sparse. Obviously this is a difficult topic but it would seem important. Note that the habit model implies large welfare cost of business cycles, since agents dislike small reductions in consumption. Finally, note that habits are also used in macro DSGE literature, here the goal in general is not to match asset prices but to generate a slower response of consumption to shocks.

6 Basic intuition

Unfortunately most habits model do not admit “useful” closed-form solutions. The basic intuition why they can give (a) a large equity premium and (b) a time-varying equity premium is however simple:

(a) If agent's utility is $v(c, h) = u(c - h)$ instead of $u(c)$, and h grows over time so that its distance to c is always rather small, then for a given volatility of c there is more volatility of $c - h$:

$$\frac{\Delta(c - h)}{c - h} = \frac{\Delta c}{c} \frac{c}{c - h} > \frac{\Delta c}{c}.$$

⁴Ljungqvist and Uhlig (AER 2001) show that, if labor supply is elastic, people work too much, hence income taxes can be Pareto-improving. Moreover, under some conditions, these taxes should be countercyclical.

⁵See for instance Szentes and Robson, for an example of preferences derived from evolutionary theory.

This is just a “leverage” effect coming from the “subsistence level” h . Hence for a given volatility of c , we certainly get more volatility of marginal utility of consumption $\frac{\partial u}{\partial c}$. This will allow us to become closer to the Hansen-Jagannathan bounds: marginal utility of consumption is volatile, which is what you need for the equity premium puzzle.

(b) When agents’ consumption becomes closer to the habit level h , they fear further negative shocks since their utility is concave. Technically, the curvature of the utility function, i.e. the “index of relative risk aversion”, is

$$I = \frac{-cv''(c)}{v'(c)} = \frac{-cu''(c-h)}{u'(c-h)},$$

and if u is CRRA, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, then

$$I = \gamma \frac{c}{c-h},$$

so that the relative risk aversion will vary with the consumption. As $c \rightarrow h$, $I \rightarrow \infty$. Hence “time-varying risk aversion”, and hence time-varying risk premia.⁶

Of course a technical issue with habits model is that you must make sure that consumption never falls below the habit, otherwise utility is not well defined! We now turn to a particular habit model, by Campbell and Cochrane. CC. They impose a specific functional form, which allows them to match a number of asset pricing facts.

7 Campbell and Cochrane Model

I use their notation.

- Utility:

$$E \sum_{t=0}^{\infty} \beta^t \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma},$$

where γ denotes the risk-aversion coefficient, X_t the external habit level and C_t consumption.

- Define the surplus consumption ratio $S_t \equiv (C_t - X_t)/C_t$. Note that $0 < S_t < 1$.
- Assume that $s_t = \log(S_t)$ is related to consumption through the following *heteroskedastic* AR(1) process:

$$s_{t+1} = (1 - \phi)\bar{s} + \phi s_t + \lambda(s_t)(\Delta c_{t+1} - g). \quad (7)$$

Lowercase letters correspond to logs, $\lambda(s_t)$ is the sensitivity function, and g is the average growth rate of the log-normal consumption process.

- This is a generalization of a standard AR(1), i.e. $X_t = (1 - \delta)X_{t-1} + C_{t-1}$. The function $\lambda(\cdot)$ introduces a nonlinearity, which will prove important.
- Assume that consumption growth is *i.i.d* and log-normal:

$$\Delta c_{t+1} = g + u_{t+1}, \text{ where } u_{t+1} \sim \text{i.i.d. } N(0, \sigma^2).$$

- External habits: The habit is assumed here to depend only on aggregate, not on individual, consumption. Thus, the inter-temporal marginal rate of substitution is here:

$$M_{t+1} = \beta \frac{U_c(C_{t+1}, X_{t+1})}{U_c(C_t, X_t)} = \beta \left(\frac{S_{t+1} C_{t+1}}{S_t C_t} \right)^{-\gamma} = \beta e^{-\gamma[g + (\phi-1)(s_t - \bar{s}) + (1 + \lambda(s_t))(\Delta c_{t+1} - g)]}. \quad (8)$$

⁶Note that some habits model use a ratio formulation, $u(c, h) = (c/h)^{1-\gamma}/(1-\gamma)$. This formulation does not generate (b).

- In contrast, with internal habits, the consumer is forward-looking and realizes that increasing C today will result in a higher habit in the future. In this case the SDF is more complicated:

$$M_{t+1} = \beta \frac{U_c(t+1) + \sum_{j=1}^{\infty} \beta^j U_x(t+1+j) \frac{\partial X_{t+1+j}}{\partial C_{t+1}}}{U_c(t) + \sum_{j=1}^{\infty} \beta^j U_x(t+j) \frac{\partial X_{t+j}}{\partial C_t}}.$$

- Campbell and Cochrane use the following sensitivity function:

$$\lambda(s_t) = \frac{1}{\bar{S}} \sqrt{1 - 2(s_t - \bar{s})} - 1, \text{ when } s_t \leq s_{\max}, 0 \text{ elsewhere,}$$

where \bar{S} and s_{\max} are respectively the steady-state and upper bound of the surplus-consumption ratio, which we set as:

$$\bar{S} = \sigma \sqrt{\frac{\gamma}{1 - \phi}},$$

and

$$s_{\max} = \bar{s} + \frac{1 - \bar{S}^2}{2}.$$

- These two values for \bar{S} and s_{\max} are one possible choice, which they justify to make the habit locally predetermined.
- This sensitivity function allows them to have a constant risk-free interest rate. To see this, note that the risk-free rate is

$$r_{t+1}^f = -\log \beta + \gamma g - \gamma(1 - \phi)(s_t - \bar{s}) - \frac{\gamma^2 \sigma^2}{2} (1 + \lambda(s_t))^2.$$

Two effects where s_t appears: intertemporal substitution and precautionary savings. CC offset these two effects by picking λ such that

$$\gamma(1 - \phi)(s_t - \bar{s}) + \frac{\gamma^2 \sigma^2}{2} (1 + \lambda(s_t))^2 = \text{constant.}$$

(Intertemporal substitution is here relative to the habit level – while consumption growth is iid, the habit moves over time predictably, generating a desire to smooth predictable variations in marginal utility.)

- Note that by picking a different λ , you can have a risk-free rate which depends, say, linearly on the state variable s_t , i.e.

$$r_{t+1}^f = A - B s_t.$$

Extensions of the CC model to study the yield curve (Wachter, JFE) or the forward premium puzzle (Verdelhan, JF) consider the case $B < 0$ and $B > 0$, modifying slightly the CC model. Depending on the value of the structural parameters, the model implies constant, pro- or counter-cyclical interest rates.

- Interpretation of B ? Consumption smoothing and precautionary savings affect the real interest rate, and the parameter B here summarizes these two different effects.
 - In good times, after a series of positive consumption shocks that result in a high surplus consumption ratio s , the agent wants to save more in order to smooth marginal utility. This leads to a decrease in the interest rate through an intertemporal substitution effect.
 - But, in good times, the representative agent is less risk-averse (the local curvature of his utility function is γ/S_t). He is less interested in saving, leading to an increase in the real interest rate through a precautionary saving effect. The case $B < 0$ is thus the one in which the precautionary effect overcomes the substitution effect. As a result, interest rates are low in bad times and high in good times.

7.1 Key mechanism

- Time-varying local risk-aversion coefficient:

$$\gamma_t = -\frac{CU_{CC}}{U_C} = \frac{\gamma}{S_t}.$$

- Counter-cyclical market price of risk. Start from:

$$SR_t = \left| \frac{E_t(R_{t+1}^e)}{\sigma_t(R_{t+1}^e)} \right| \leq \frac{\sigma_t(M_{t+1})}{E_t(M_{t+1})} = MPR_t,$$

with an equality for assets that perfectly correlated with the SDF.

- In this model, the SDF is conditionally log-normally distributed, hence we can apply the general formula:

$$SR_t = \sqrt{e^{Var_t(\log M_{t+1})} - 1} \simeq \sigma_t(\log M_{t+1})$$

Proof: Juste use the log-normal formula $E(\exp(X)) = \exp(E(X) + \frac{1}{2}Var(X))$ and compute

$$E_t(M_{t+1}) = e^{E_t(\log M_{t+1}) + \frac{1}{2}Var_t(\log M_{t+1})},$$

and

$$\begin{aligned} Var_t(M_{t+1}) &= E_t(M_{t+1}^2) - [E_t(M_{t+1})]^2, \\ &= e^{2E_t(\log M_{t+1}) + 2Var_t(\log M_{t+1})} - e^{2E_t(\log M_{t+1}) + Var_t(\log M_{t+1})}, \end{aligned}$$

hence the result. In our specific case,

$$MPR_t = \gamma\sigma(1 + \lambda(s_t)) = \frac{\gamma\sigma}{\bar{S}} \sqrt{1 - 2(s_t - \bar{s})}.$$

At the steady-state, $\bar{SR} = \gamma\sigma/\bar{S}$, but the market price of risk is countercyclical, and hence so is the sha

- The model matches the level of the riskless rate and the equity return, and their volatilities.
- Look carefully at figures 2, 3, 4, 5: how do prices move with the state variable s_t ?
- Intuition: no variation in expected cash flows, no variation in the risk-free rate, but a lot of variation in risk premia. Risk premia are higher after a sequence of some bad shocks, which push C close to X , hence a very low s . People are very afraid of further decreases in consumption, hence they become more risk averse and are willing to hold risky assets only if these assets give a very nice return on average.
- Note how consumption growth is not the key source of risk, but covariance with the s_t variable is.
- The model matches the time-series predictability evidence: dividend growth is not predictable, but returns are predictable, and the volatility of the P-D ratio is accounted for by this later term (“discount rate news”).
- Read all of section 3D especially.
- Figure 9: feed consumption data in the model to plot the implied P-D ratio - it tracks the data well.
- Long-run equity premium: because of mean-reversion in stock prices, excess returns on stocks at long horizons are even more puzzling than the standard one-period ahead puzzle. CC note that if the state variable is stationary, the long-run std dev of the SDF will not depend on the current state. Key point: in their model, $S^{-\gamma}$ is not stationary!

- Conditional vs. unconditional correlations: conditional correlation is (nearly) perfect between consumption growth and stock returns, but the unconditional correlation is quite messy.
- Role of time-aggregation in some moments.
- Patterns in volatility: volatility higher in bad times.
- Section 4 not so useful.

7.2 Solving the model

Solving numerically this model is somewhat complicated, because of the important nonlinearities. However, the general method is standard. The aggregate market is represented as a claim to the future consumption stream. Let P_t denote the ex-dividend price of this claim. Then, $E_t[M_{t+1}R_{t+1}] = 1$ implies that in equilibrium P_t satisfies:

$$E_t \left(M_{t+1} \frac{P_{t+1} + C_{t+1}}{P_t} \right) = 1$$

$$\begin{aligned} \frac{P_t}{C_t} &= E_t \left(M_{t+1} \left(\frac{P_{t+1}}{C_{t+1}} + 1 \right) \frac{C_{t+1}}{C_t} \right) \\ &= E_t \left(\beta \left(\frac{S_{t+1}}{S_t} \right)^{-\gamma} \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \left(\frac{P_{t+1}}{C_{t+1}} + 1 \right) \right) \end{aligned}$$

The state variable is s_t . We solve for a fixed point, i.e. $\frac{P_t}{C_t} = h(s_t)$ with

$$h(s_t) = E_{u_{t+1}} \left(\beta \left(\frac{S_{t+1}}{S_t} \right)^{-\gamma} \left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} (h(s_{t+1}) + 1) \right),$$

with

$$\begin{aligned} \Delta c_{t+1} &= g + u_{t+1}, \\ s_{t+1} &= (1 - \phi)\bar{s} + \phi s_t + \lambda(s_t)u_{t+1}, \\ u_{t+1} &\sim i.i.d.N(0, \sigma^2). \end{aligned}$$

i.e.

$$h(s) = \beta \int_{-\infty}^{\infty} \exp(-\gamma((1 - \phi)(\bar{s} - s) + \lambda(s)u) + (1 - \gamma)g + (1 - \gamma)u) (h((1 - \phi)\bar{s} + \phi s + \lambda(s)u) + 1) d\Phi(u). \quad (9)$$

One way to do it:

(1) set a grid for s , $\{s_1, s_2, \dots, s_N\}$; (2) take a guess for h , i.e. $\{h(s_1), h(s_2), \dots, h(s_N)\}$; (3) for each value of s in the grid, compute the RHS of equation (9) numerically. you need to interpolate h to find its value at the point $(1 - \phi)\bar{s} + \phi s + \lambda(s)u$ since it lies outside the grid. To compute the integral numerically, one way is to use a quadrature, e.g. $\int k(u)du = \sum_i \omega_i k(u_i)$ for some points u_i . (4) This gives you a new guess for h using equation (9), and you keep iterating on this equation until convergence.

Once we know h , we have the P-D ratio and the rest can be computed simply, as in Mehra-Prescott. (See Wachter, Finance research letters, for a note detailing the computation of that model.)

8 Recursive Utility

Basics

Epstein and Zin (1989 JPE, 1991 Ecta) following work by Kreps and Porteus introduced a class of preferences which allow to break the link between risk aversion and intertemporal substitution. These preferences have proved very useful in applied work in asset pricing, portfolio choice, and are becoming more prevalent in macroeconomics.⁷

To understand the formulation, recall the standard expected utility time-separable preferences are defined as

$$V_t = E_t \sum_{s=0}^{\infty} \beta^{s-t} u(c_{t+s}),$$

but we can also define them recursively as

$$V_t = u(c_t) + \beta E_t V_{t+1},$$

or equivalently (this is just a scaling):

$$V_t = (1 - \beta)u(c_t) + \beta E_t (V_{t+1}).$$

EZ preferences generalize this: they are defined recursively over current (known) consumption and a certainty equivalent $R_t(V_{t+1})$ of tomorrow's utility V_{t+1} :

$$V_t = F(c_t, R_t(V_{t+1})), \quad (10)$$

where

$$R_t(V_{t+1}) = G^{-1}(E_t G(V_{t+1})), \quad (11)$$

with F and G increasing and concave. Note that $R_t(V_{t+1}) = V_{t+1}$ if there is no uncertainty on V_{t+1} . The more concave G is, and the more uncertain V_{t+1} is, the lower is $R_t(V_{t+1})$.

Most of the literature considers simple functional forms for F and G :

$$\begin{aligned} \rho > 0 : F(c, z) &= ((1 - \beta)c^{1-\rho} + \beta z^{1-\rho})^{\frac{1}{1-\rho}}, \\ \alpha > 0 : G(x) &= \frac{x^{1-\alpha}}{1-\alpha}. \end{aligned}$$

Note I will use the following limits:⁸

$$\begin{aligned} \rho = 1 : F(c, z) &= c^{1-\beta} z^\beta. \\ \alpha = 1 : G(x) &= \log x. \end{aligned}$$

⁷Potentially there are alternative preferences that also break down this link between IES and risk aversion.

⁸Proofs

$$\begin{aligned} ((1 - \beta)c^{1-\rho} + \beta z^{1-\rho})^{\frac{1}{1-\rho}} &= \exp\left(\frac{1}{1-\rho} \log\left(1 + (1-\rho)\left((1-\beta)\frac{c^{1-\rho}-1}{1-\rho} + \beta\frac{z^{1-\rho}-1}{1-\rho}\right)\right)\right) \\ &\simeq \exp\left((1-\beta)\frac{c^{1-\rho}-1}{1-\rho} + \beta\frac{z^{1-\rho}-1}{1-\rho}\right) \\ &\simeq \exp((1-\beta)\ln c + \beta\ln z) \\ &\simeq c^{1-\beta} z^\beta. \end{aligned}$$

Hence

$$\begin{aligned}\alpha > 0 &: R_t(V_{t+1}) = E_t(V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}}, \\ \alpha = 1 &: R_t(V_{t+1}) = \exp(E_t \log(V_{t+1})).\end{aligned}$$

- Special case: if $\alpha = \rho$ or if consumption is deterministic: we have the usual standard time-separable expected discounted utility with discount factor β and IES = $\frac{1}{\rho}$, risk aversion $\alpha = \rho$.
- In general α is the relative risk aversion coefficient for static gambles and ρ is the inverse of the intertemporal elasticity of substitution for deterministic variations.
- To see this, suppose you have consumption c today and that there is some uncertainty tomorrow about which consumption path you will have: either $\{c_L, \bar{c}, \bar{c}, \dots\}$ or $\{c_H, \bar{c}, \bar{c}, \dots\}$, each has prob $\frac{1}{2}$. Then the utility today is $V = F(c, G^{-1}(\frac{1}{2}G(V_L) + \frac{1}{2}G(V_H)))$, where $V_L = F(c_L, \bar{c})$ and $V_H = F(c_H, \bar{c})$. Clearly the curvature of G determines how adverse you are to the uncertainty, e.g. if G is linear you only care about the expected value.
- Discuss simple example with two lotteries:

- lottery A pays in each period $t = 1, 2, \dots$ c_h or c_l , the probability is $\frac{1}{2}$ and the outcome is iid across period;

- lottery B pays starting at $t = 1$ either c_h at all future dates for sure, or c_l at all future date for sure; there is a single draw at time $t = 1$.

With expected utility, you are indifferent between these lotteries, but with EZ lottery B is preferred iff $\alpha > \rho$.

- In general, early resolution of uncertainty is preferred if and only if $\alpha > \rho$ i.e. risk aversion $> \frac{1}{\text{IES}}$. This is another way to motivate these preferences, since early resolution seems intuitively preferable.
- Technically, EZ is an extension of EU which relaxes the independence axiom. Recall the independence axiom is: if $x \succeq y$, then for any $z, \alpha : \alpha x + (1-\alpha)z \succeq \alpha y + (1-\alpha)z$. With EZ, “Intertemporal composition of risk matters”: we cannot reduce compound lotteries.

8.1 The SDF and the market return

For any asset pricing model, the empirical implications are all

As we have discussed before, the SDF summarizes all the we need the SDF implied by these preferences to obtain empirical predictions. The first result gives the SDF, by simply computing the intertemporal MRS.

Result: The stochastic discount factor is

$$M_{t,t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{V_{t+1}}{R_t(V_{t+1})} \right)^{\rho-\alpha}.$$

Proof:

Start with the definition of the IMRS:

$$M_{t,t+1}(s^t, s_{t+1}) = \frac{\partial V_t / \partial c_{t+1}(s^t, s_{t+1})}{\partial V_t / \partial c_t(s^t)}.$$

Differentiation of (10) w.r.t. c_t :

$$\frac{\partial V_t}{\partial c_t} = F_1(c_t, R_t(V_{t+1})).$$

Differentiation of (10) w.r.t. $c_{t+1}(s^t, s_{t+1})$:

$$\begin{aligned} \frac{\partial V_t}{\partial c_{t+1}(s^t, s_{t+1})} &= F_2(c_t, R_t(V_{t+1})) \times \frac{\partial R_t(V_{t+1})}{\partial V_{t+1}(s^t, s_{t+1})} \times \frac{\partial V_{t+1}(s^t, s_{t+1})}{\partial c_{t+1}(s^t, s_{t+1})} \\ &= F_2(c_t, R_t(V_{t+1})) \times (E_t V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}-1} \times V_{t+1}^{-\alpha} \times F_1(c_{t+1}(s^t, s_{t+1}), R_{t+1}(V_{t+2})). \end{aligned}$$

since

$$\frac{\partial R_t(V_{t+1})}{\partial V_{t+1}(s^t, s_{t+1})} = (E_t V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}-1} V_{t+1}^{-\alpha}.$$

Note now:

$$\begin{aligned} F_1(c_t, R_t(V_{t+1})) &= (1-\beta)c_t^{-\rho} \left((1-\beta)c_t^{1-\rho} + \beta R_t(V_{t+1})^{1-\rho} \right)^{\frac{1}{1-\rho}-1} \\ &= (1-\beta)c_t^{-\rho} F(c_t, R_t(V_{t+1}))^\rho. \end{aligned}$$

$$\begin{aligned} F_2(c_t, R_t(V_{t+1})) &= \beta R_t(V_{t+1})^{-\rho} F(c_t, R_t(V_{t+1}))^\rho, \\ &= \beta R_t(V_{t+1})^{-\rho} V_t^\rho. \end{aligned}$$

I get finally

$$\begin{aligned} M_{t,t+1}(s^t, s_{t+1}) &= \frac{F_2(c_t, R_t(V_{t+1})) \times (E_t V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}-1} \times V_{t+1}^{-\alpha} \times F_1(c_{t+1}(s^t, s_{t+1}), R_{t+1}(V_{t+2}))}{F_1(c_t, R_t(V_{t+1}))} \\ &= \frac{\beta R_t(V_{t+1})^{-\rho} F(c_t, R_t(V_{t+1}))^\rho (E_t V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}-1} V_{t+1}^{-\alpha} c_{t+1}^{-\rho} F(c_{t+1}, R_{t+1}(V_{t+2}))^\rho}{c_t^{-\rho} F(c_t, R_t(V_{t+1}))^\rho} \\ &= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left[R_t(V_{t+1})^{-\rho} \times (E_t V_{t+1}^{1-\alpha})^{\frac{1}{1-\alpha}-1} \times V_{t+1}^{-\alpha} F(c_{t+1}, R_{t+1}(V_{t+2}))^\rho \right] \\ &= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left[R_t(V_{t+1})^{-\rho} \times R_t(V_{t+1})^\alpha \times V_{t+1}^{-\alpha} \times V_{t+1}^\rho \right] \\ &= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{V_{t+1}}{R_t(V_{t+1})} \right)^{\rho-\alpha}. \end{aligned}$$

Empirical implementation (Epstein-Zin, JPE 1989)

Our formula for the SDF has an obvious problem: it depends on future utilities $\frac{V_{t+1}}{R_t(V_{t+1})}$, which are not observed - only consumption is. Epstein and Zin produced an empirical implementation of these preferences using the fact that the *return on wealth* can be substituted in instead of future utility (the second term of the discount factor). This allows to do (i) GMM estimation as Hansen-Singleton (1983), or (ii) to use log-linear-log-normal approximations.

The derivation is as follows. Define wealth as the PDV of consumption, or recursively as:

$$W_t = c_t + E_t (M_{t,t+1} W_{t+1}).$$

Note that this includes the current consumption – this is the cum-dividend price, not the ex-dividend price that we often define.

Result: $W_t = V_t / F_1 (c_t, R_t (V_{t+1}))$.

Proof: Guess and verify that satisfies this recursion.

$$\begin{aligned} W_t &= c_t + E_t (m_{t,t+1} W_{t+1}) \\ &= c_t + E_t \left(\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{V_{t+1}}{R_t(V_{t+1})} \right)^{\rho-\alpha} \frac{V_{t+1}}{(1-\beta)c_{t+1}^{-\rho} F(c_{t+1}, R_{t+1}(V_{t+2}))^\rho} \right) \\ &= c_t + E_t \left(\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} V_{t+1}^{1-\alpha} R_t(V_{t+1})^{\alpha-\rho} \frac{1}{(1-\beta)c_{t+1}^{-\rho}} \right) \\ &\quad \frac{V_t}{F_1(c_t, R_t(V_{t+1}))} (1-\beta)c_t^{-\rho} R_t(V_{t+1})^{\rho-\alpha} \stackrel{?}{=} (1-\beta)c_t^{1-\rho} R_t(V_{t+1})^{\rho-\alpha} + \beta E_t (V_{t+1}^{1-\alpha}) \\ &\quad \frac{V_t}{(1-\beta)c_t^{-\rho} F(c_t, R_t(V_{t+1}))^\rho} (1-\beta)c_t^{-\rho} R_t(V_{t+1})^{\rho-\alpha} \stackrel{?}{=} (1-\beta)c_t^{1-\rho} R_t(V_{t+1})^{\rho-\alpha} + \beta E_t (V_{t+1}^{1-\alpha}) \\ &\quad V_t^{1-\rho} R_t(V_{t+1})^{\rho-\alpha} \stackrel{?}{=} (1-\beta)c_t^{1-\rho} R_t(V_{t+1})^{\rho-\alpha} + \beta E_t (V_{t+1}^{1-\alpha}) \\ &\quad V_t^{1-\rho} \stackrel{?}{=} (1-\beta)c_t^{1-\rho} + \beta R_t(V_{t+1})^{1-\rho}, \end{aligned}$$

which is true. This confirms our guess.

Note that I can now define the return on the *wealth portfolio*:

$$R_{t,t+1} = \frac{W_{t+1}}{W_t - C_t}.$$

(This is the return definition when the price is cum-dividend.) Simple algebra shows that $R_{t,t+1} = \left\{ \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{R_t(V_{t+1})}{V_{t+1}} \right)^{1-\rho} \right\}^{-1}$.

Proof:

$$\begin{aligned} R_{t,t+1} &= \frac{V_{t+1}}{F_1(c_{t+1}, R_{t+1}(V_{t+2}))} \frac{1}{\left(\frac{V_t}{F_1(c_t, R_t(V_{t+1}))} - C_t \right)} \\ &= \frac{V_{t+1}}{F_1(c_{t+1}, R_{t+1}(V_{t+2}))} \frac{F_1(c_t, R_t(V_{t+1}))}{(V_t - C_t F_1(c_t, R_t(V_{t+1})))} \\ &= \frac{V_{t+1}}{(1-\beta)c_{t+1}^{-\rho} V_{t+1}^\rho} \frac{(1-\beta)c_t^{-\rho} F(c_t, R_t(V_{t+1}))^\rho}{\beta R_t(V_{t+1})^{1-\rho} F(c_t, R_t(V_{t+1}))^\rho} \\ R_{t,t+1} &= \left\{ \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{R_t(V_{t+1})}{V_{t+1}} \right)^{1-\rho} \right\}^{-1}. \end{aligned}$$

Hence the SDF can be expressed as a function of the return:

$$\begin{aligned}
M_{t,t+1} &= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left(\frac{V_{t+1}}{R_t(V_{t+1})} \right)^{\rho-\alpha} \\
&= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho} \left\{ \frac{R_{t,t+1}^{-1}}{\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho}} \right\}^{\frac{\rho-\alpha}{\rho-1}} \\
&= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho \left(1 - \frac{\rho-\alpha}{\rho-1} \right)} R_{t,t+1}^{\frac{\rho-\alpha}{1-\rho}} \\
&= \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\rho \frac{\alpha-1}{\rho-1}} R_{t,t+1}^{\frac{\rho-\alpha}{1-\rho}}.
\end{aligned}$$

This allows to implement this model empirically since **the SDF can now be measured**. Epstein and Zin proxy the return on wealth by the return on a broad stock market return. An obvious criticism is that the a lot of wealth is not traded on the stock market (private firms, human capital, housing...); however these returns may be correlated with the stock market return. Future work (e.g. Campbell 1996 JPE) try to add human capital (see below).

Note that assets' risk are measured as the covariance with the SDF, so the Epstein-Zin utility rationalizes a formula which is a mix of the CAPM and the CCAPM:

$$\log \left(\frac{E_t R_{t+1}^i}{R_{t+1}^f} \right) = \rho \text{Cov}_t(R_{t+1}^i, \Delta \log C_{t+1}) + \frac{\rho - \alpha}{1 - \alpha} \text{Cov}_t(R_{t+1}^i, R_{t+1}^m).$$

Note that if $\rho = \alpha$, the second term disappears - we are back to expected utility.

Empirically, the extra free parameter of EZ allows to improve on the standard CRRA model. However the solutions of asset pricing puzzles with EZ utility require high risk aversion (except, to some extent, the Bansal-Yaron paper we will mention later). This formula is, however, limited by the

A remark on the consumption-wealth ratio: given that

$$\begin{aligned}
W_t &= \frac{V_t}{F_1(c_t, R_t(V_{t+1}))} \\
&= \frac{V_t}{(1 - \beta)c_t^{-\rho} F(c_t, R_t(V_{t+1}))^\rho} \\
&= \frac{V_t^{1-\rho}}{(1 - \beta)c_t^{-\rho}}. \\
\frac{c_t}{W_t} &= (1 - \beta) \frac{c_t^{1-\rho}}{V_t^{1-\rho}},
\end{aligned}$$

so that if $\rho = 1$ (log utility in IES), the consumption-wealth ratio is constant, but in general the consumption-wealth ratio encodes the state variable V_t/c_t . In the data, Lettau and Ludvigson (2001 JPE) show that a high consumption-wealth ratio forecasts low returns. They suggest that the consumption-wealth ratio is a measure of “time-varying risk aversion” which can help price assets in the cross-section.

8.2 Log-normal iid results

We revisit the standard log-normal asset pricing computations with E-Z utility. Suppose that

$$\Delta \log C_{t+1} = \mu_c - \frac{\sigma_c^2}{2} + \sigma_c \varepsilon_{t+1},$$

and we seek the price of an equity paying out $\{D_t\}$, with

$$\Delta \log D_{t+1} = \mu_d - \frac{\sigma_d^2}{2} + \sigma_d \varepsilon_{t+1}.$$

We have:

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\alpha} \left(\frac{V_{t+1}}{E_t(V_{t+1}^{1-\theta})^{\frac{1}{1-\theta}}} \right)^{\alpha-\theta},$$

with $\alpha = 1/\text{IES}$ and $\theta = \text{risk aversion}$.

To compute asset prices, first rewrite the SDF as

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\theta} \frac{\left(\frac{V_{t+1}}{C_{t+1}} \right)^{\alpha-\theta}}{E_t \left\{ \left(\frac{V_{t+1}}{C_{t+1}} \right)^{1-\theta} \left(\frac{C_{t+1}}{C_t} \right)^{1-\theta} \right\}^{\frac{\alpha-\theta}{1-\theta}}}.$$

Next, note that the utility recursion implies

$$\frac{V_t}{C_t} = \left(1 - \beta + \beta E_t \left(\left(\frac{V_{t+1}}{C_{t+1}} \right)^{1-\theta} \left(\frac{C_{t+1}}{C_t} \right)^{1-\theta} \right)^{\frac{1-\alpha}{1-\theta}} \right)^{\frac{1}{1-\alpha}},$$

and given the iid assumption, $\frac{V_t}{C_t}$ is constant, and satisfies

$$v^{1-\alpha} = 1 - \beta + \beta v^{1-\alpha} E \left(\left(\frac{C_{t+1}}{C_t} \right)^{1-\theta} \right)^{\frac{1-\alpha}{1-\theta}}.$$

Hence, the SDF is

$$\begin{aligned} M_{t+1} &= \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\theta} \frac{1}{E_t \left\{ \left(\frac{C_{t+1}}{C_t} \right)^{1-\theta} \right\}^{\frac{\alpha-\theta}{1-\theta}}}, \\ &= \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\theta} \frac{1}{\left\{ e^{(1-\theta)\mu_c - \theta(1-\theta)\frac{\sigma_c^2}{2}} \right\}^{\frac{\alpha-\theta}{1-\theta}}}, \\ &= \beta e^{(\theta-\alpha)\mu_c + \theta(\alpha-\theta)\frac{\sigma_c^2}{2}} \left(\frac{C_{t+1}}{C_t} \right)^{-\theta}. \end{aligned}$$

Hence, the risk-free rate is constant and equal to

$$\begin{aligned} R^f &= \frac{1}{E_t(M_{t+1})}, \\ &= \frac{1}{\beta e^{(\theta-\alpha)\mu_c + \theta(\alpha-\theta)\frac{\sigma_c^2}{2}} E_t \left(\frac{C_{t+1}}{C_t} \right)^{-\theta}}, \\ &= \frac{1}{\beta e^{(\theta-\alpha)\mu_c + \theta(\alpha-\theta)\frac{\sigma_c^2}{2}} e^{-\theta\mu_c + \theta\frac{\sigma_c^2}{2} + \frac{\theta^2\sigma_c^2}{2}}}, \\ &= \frac{e^{\alpha\mu_c - \theta(\alpha+1)\frac{\sigma_c^2}{2}}}{\beta}. \end{aligned}$$

First result:

$$\log(R^f) = -\log\beta + \alpha\mu_c - \theta(1 + \alpha)\frac{\sigma_c^2}{2}$$

Key points:

- (1) This is the same formula as before if $\theta = \alpha$.
- (2) The effect of the growth rate μ_c on R^f is governed by α , the inverse of the IES.
- (3) The effect of σ_c on R_f is zero if risk aversion $\theta = 0$.
- (4) But in general, the strength of the effect depends on the IES. Intuition (?): when α is big, people are reluctant to change their consumption path, so the higher uncertainty does not lead to much change in savings, and interest rates have to move a lot.

Now for the equity premium: write the equation for the P-D ratio, denoted q , and which is constant given the iid assumption, as

$$\begin{aligned} q &= E_t \left(M_{t+1} \frac{D_{t+1}}{D_t} (1 + q) \right), \\ \frac{q}{1 + q} &= E_t \left(\beta e^{(\theta - \alpha)\mu_c + \theta(\alpha - \theta)\frac{\sigma_c^2}{2}} \left(\frac{C_{t+1}}{C_t} \right)^{-\theta} \frac{D_{t+1}}{D_t} \right), \\ &= \beta e^{(\theta - \alpha)\mu_c + \theta(\alpha - \theta)\frac{\sigma_c^2}{2}} e^{\mu_d - \frac{\sigma_d^2}{2} - \theta \left(\mu_c - \frac{\sigma_c^2}{2} \right) + \frac{(\sigma_d - \theta\sigma_c)^2}{2}}, \\ &= \beta e^{\mu_d - \alpha\mu_c + \theta(\alpha + 1)\frac{\sigma_c^2}{2}} e^{-\theta\sigma_d\sigma_c}, \\ &= \frac{1}{R_f} e^{\mu_d - \theta\sigma_d\sigma_c}. \end{aligned}$$

The expected return on equity is thus

$$\begin{aligned} E_t(R_{t+1}^e) &= E_t \left(\frac{q + 1}{q} \frac{D_{t+1}}{D_t} \right), \\ &= R_f e^{\theta\sigma_d\sigma_c}, \end{aligned}$$

and the (log) equity premium is:

$$\log \left(\frac{E_t R_{t+1}^e}{R_{t+1}^f} \right) = \theta\sigma_d\sigma_c.$$

Second set of results:

- (1) Same formula as before if $\theta = \alpha$.
- (2) The risk premium is determined by risk aversion θ .

Third set of result: the P-D ratio is

$$q = \frac{H}{1 - H}$$

with

$$H = \frac{1}{R_f} e^{\mu_d - \theta\sigma_d\sigma_c}.$$

hence:

- (1) Suppose that $\mu_d = \mu_c = \mu$, then an increase in μ will increase asset prices provided that $\alpha < 1$ i.e. the IES is large enough, so that interest rates do not increase too much.

(2) Suppose that $\sigma_d = \sigma_c = \sigma$, then an increase in σ will decrease asset prices provided that $\theta(1 - \alpha) > 0$, i.e. $\alpha < 1$ again. The intuition is that an increase in σ leads people to save more in risk-free assets, lowering the RF rate, and to be reluctant to hold risky assets, increasing the equity premium. When the IES is high, the interest rate does not change much, and the risk premium effect dominates.

These two comparative statics motivate the Bansal-Yaron (JF, 2004) model and calibration.

Note: we can allow for “leverage”, i.e. dividends more volatile than consumption, i.e. $\sigma_d = \lambda\sigma_c$ with $\lambda > 1$. How does this affect the results above?

8.3 Long-run risk model (Bansal and Yaron JF 2004)

Background: in CRRA model an increase in the growth rate of the economy reduces asset prices if $\gamma > 1$. That's because of the low IES: interest rates go up by more than the growth rate of cash flows. (Think Gordon model: $\frac{P}{D} = \frac{1}{r-g}$, with $r = -\log(\beta) + \gamma g - \frac{\gamma^2 \sigma^2}{2}$.) Intuitively, this sound strange, since we may expect asset prices to increase when expected growth goes up. If the IES was large, i.e. $\gamma < 1$, we would get the more intuitive sign. This would be consistent with the data, in that interest rates do not move much. (This is related to the framework of PS2 - time-varying growth rate.) However, $\gamma < 1$ with CRRA utility implies low risk aversion, very much at odds with the equity premium. This motivates introducing Epstein-Zin preferences - we can have high risk aversion to have a shot at the equity premium, and still high IES to get that shocks to expected growth drive asset prices up.

Bansal and Yaron (JF 2004) study the implications of recursive utility with both IES and risk aversion greater than unity, when consumption growth has a highly persistent component, e.g.:

$$\begin{aligned}\Delta \log C_t &= \mu + x_t + u_t, \\ x_t &= \rho x_{t-1} + v_t.\end{aligned}$$

A persistent negative shock to consumption growth (i.e. a negative v) decreases asset prices strongly since the effect on dividends/consumption is very persistent. Their full model also incorporates (a) different processes for dividends and consumption, and (b) stochastic volatility:

$$\begin{aligned}\Delta \log C_t &= \mu + x_t + \sigma_{t-1} u_t, \\ x_t &= \rho x_{t-1} + \sigma_{t-1} v_t, \\ \Delta \log D_t &= \mu_d + \phi x_t + \phi_d \sigma_{t-1} u_t, \\ \sigma_{t+1}^2 &= \sigma^2 + v_1 (\sigma_t^2 - \sigma^2) + \sigma_w^2 w_{t+1}.\end{aligned}$$

Two state variables: x_t = conditional mean of consumption growth rate, and σ_t^2 = time-varying variance of consumption growth rate.

Log-linear approximations a la Campbell-Shiller lead to the following expression for the return on a consumption claim:

$$r_{c,t+1} = \kappa_0 + \kappa_1 z_{t+1} - \kappa_2 z_t + \Delta \log C_{t+1},$$

where $z_t = \log \frac{P_t}{C_t}$ satisfies:

$$z_t = A_0 + A_1 x_t + A_2 \sigma_t^2,$$

and

$$A_1 = \frac{1 - \frac{1}{\psi}}{1 - \kappa_1 \rho}.$$

Here, we find again the same result as in the IID model: an increase in x on the P-D ratio raises asset prices only if $\text{IES} > 1$ ($\psi < 1$). Also, $A_2 < 0$ if IES and $\text{RA} > 1$, so that higher uncertainty then lowers asset prices. These comparative statics are intuitive and appealing.

Similarly, for a dividend claim,

$$A_{1,m} = \frac{\phi - \frac{1}{\psi}}{1 - \kappa_{1m}\rho},$$

and $A_{2m} < 0$ if $\text{IES} > 1$. Now $\phi > 1$ because of “leverage”, hence the threshold value for ψ is lower than 1.

SDF innovation (in log):

$$m_{t+1} - E_t m_{t+1} = \lambda_{m\eta} \sigma_t \eta_{t+1} - \lambda_{me} \sigma_t e_{t+1} - \lambda_{mw} \sigma_w w_{t+1},$$

the SDF loads on the three shocks: vol shock, iid shock, and long-run shock. As a result,

$$E_t(r_{m,t+1} - r_{f,t+1}) + \frac{1}{2} \text{Var}_t(r_{m,t+1}) = \beta_{me} \lambda_{me} \sigma_t^2 + \beta_{mw} \lambda_{mw} \sigma_w^2.$$

The iid shock to dividends (u_{t+1}) is not priced, because it is assumed to be uncorrelated with the iid shock to consumption (e_{t+1}). Risk premia rise when σ_t rises.

Without the shock to σ_t , the model implies large risk premia and volatile stock returns, since stock prices move a lot when x_t moves. However, the model also then implies (nearly) constant risk premia. Hence, expected returns are not forecastable, but cash flows are. This is very much at odds with the standard empirical results coming from the Campbell-Shiller decomposition.

The stochastic volatility shock helps the model match the data in these dimensions. An increase in σ_t lowers the P-D ratio and increases expected future returns. In contrast to the Campbell-Cochrane model, which works off *time-varying risk aversion* through the surplus-consumption ratio s_t , the Bansal-Yaron model works off time-varying *risk* through the conditional volatility σ_t .

A few words on the typical calibration:

- very high persistence for x : $\rho = .98$,
- high persistence for volatility shock too,
- these two shocks have significant variances;
- risk aversion 10,
- IES 1.5.

Evidence and discussion:

(a) The shocks to x in the model imply that a cash-flow discount rate decomposition would produce strong effects of cash flow variations on prices. The volatility shock reduce this, because they act as “discount rate shocks”. Still, the model may produce too much cash flow news compared to the data. More recent calibrations of the model give more importance to the volatility shock.

(b) Direct evidence for the time series process for consumption? Very difficult to establish empirically that there is long-run risk. BY show that this process would be very difficult to distinguish empirically from a standard iid in growth rate model, if the vol of v is much smaller than the vol of u . Of course, the two models have very different implications for asset prices. Re time-varying volatility: there is some in returns and perhaps in consumption, but it is rather short-lived. So if you fit directly the time series

process to the data, you don't get much direct evidence, certainly the persistence that they assume is fairly high.

(c) On the other hand, indirect estimation by asking, for which parameters does the model fit best asset prices as well as the consumption process suggest an important role for long-run risk.

(d) Hansen and Sargent note that the paper relies on the assumption that agents are able to observe these long-run shocks perfectly in real time, even though in reality we as economists/econometricians have difficulty observing them. It would be interesting to examine the model with learning (or with robustness).

(e) IES: the paper was controversial because it assumed a large IES, while many researchers have estimated rather low IES (e.g. Hall 1988 JPE) by estimating the Euler equation with power utility for T-bills. However, some recent studies find somewhat larger IES if you look at stockholders' consumption (Vissing-Jorgensen, 2002). Some results suggest an aggregation bias (Guvenen, 2005). And BY show that in their model, the regression would suggest an $IES > 1$ (though not as small as the typical estimates).

(f) Discussion: the Bansal-Yaron model has become a very popular framework for asset pricing. Like the Campbell Cochrane, it has the simplicity and tractability of a representative agent, and it matches well the standard facts of large, volatile risk premia, and volatile stock returns. It has the advantage of being more testable, and less "reverse-engineered". But, the presence of long-run risk in macro time series remains controversial. Also, time-varying volatility is not an obvious feature of the data. Still, many extensions of the model are studied now international context, to yield curve, to option pricing, etc.

(g) Instead of the lognormal time-varying risk, one can have a time-varying risk of disaster, as in Gourio (2008), Wachter (2009), or Gabaix (2009). The model works similarly, and does quite well. The time-varying risk of disaster plays a role similar to that of the time-varying volatility. The calibration is more difficult, because disasters are rarely observed.

8.4 Tallarini/Hansen computations when $IES=1$

Another approach to implement Epstein-Zin utility empirically is to solve exactly for the value V_t , given a tractable consumption process. This implies we do not need to introduce the market return in the SDF. This section does this computation in the special case $\rho = 1$, i.e. an IES equal to unity, and for arbitrary autocorrelations of consumption growth. There is a fair bit of algebra but I give all of this to you for free! First, write the recursion for $\sigma = 1$:

$$V_t = C_t^{1-\beta} (E_t V_{t+1}^{1-\alpha})^{\frac{\beta}{1-\alpha}}$$

$$\log V_t = (1-\beta) \log C_t + \frac{\beta}{1-\alpha} \log E_t (V_{t+1}^{1-\alpha})$$

$$\log V_t = (1-\beta) \log C_t + \frac{\beta}{1-\alpha} \log E_t \exp((1-\alpha) \log V_{t+1})$$

Assume that consumption growth follows some $MA(\infty)$ process,

$$\Delta \log C_t = \gamma(L)w_t + \mu_c$$

$$w_t \text{ iid } N(0, I)$$

where w_t can be a vector.

From our previous computations, the stochastic discount factor is:

$$m_{t,t+1} = \beta \frac{c_t}{c_{t+1}} \frac{\exp((1-\alpha) \ln V_{t+1})}{E_t \exp((1-\alpha) \ln V_{t+1})}$$

We now guess and verify that $v_t = \xi c_t + \tau(L)w_t + \mu_v$ (where lowercase denotes log). Once we verify this guess and find the coefficients ξ , $\tau(L)$ and μ_v we have the value v_t and hence the stochastic discount factor $m_{t,t+1}$. OK, here is now the algebra:

Verification:

$$\begin{aligned} & E_t \exp((1-\alpha)v_{t+1}) \\ = & E_t \exp((1-\alpha)(\xi c_{t+1} - \xi c_t + \xi c_t + \tau(L)w_{t+1} + \mu_v)), \\ = & \exp((1-\alpha)(\xi c_t + \mu_v + \xi \mu_c)) E_t \exp\left(\frac{\sigma}{2}((\xi \gamma(L) + \tau(L))w_{t+1})\right), \\ = & \exp((1-\alpha)(\xi c_t + \mu_v + \xi \mu_c)) \exp\left((1-\alpha)E_t(\xi \gamma(L) + \tau(L))w_{t+1} + \frac{(1-\alpha)^2}{2}(\xi \gamma(0) + \tau(0))'(\xi \gamma(0) + \tau(0))\right), \\ = & \exp((1-\alpha)(\xi c_t + \mu_v + \xi \mu_c)) \exp\left((1-\alpha)\left[\frac{\xi \gamma(L) + \tau(L)}{L}\right]_+ w_t + \frac{(1-\alpha)^2}{2}(\xi \gamma(0) + \tau(0))'(\xi \gamma(0) + \tau(0))\right). \end{aligned}$$

Hence,

$$\begin{aligned} v_t &= (1-\beta)c_t + \frac{\beta}{1-\alpha} \log E_t \exp((1-\alpha)v_{t+1}), \\ &= (1-\beta)c_t + \beta(\xi c_t + \mu_v + \xi \mu_c) + \beta \left[\frac{\xi \gamma(L) + \tau(L)}{L}\right]_+ w_t + \frac{\beta(1-\alpha)}{2}(\xi \gamma(0) + \tau(0))'(\xi \gamma(0) + \tau(0)), \\ &= (1-\beta + \beta\xi)c_t + \beta(\mu_v + \xi \mu_c) + \beta \left[\frac{\xi \gamma(L) + \tau(L)}{L}\right]_+ w_t + \frac{\beta(1-\alpha)}{2}(\xi \gamma(0) + \tau(0))'(\xi \gamma(0) + \tau(0)), \end{aligned}$$

where $[A(L)]_+$ = terms with nonnegative degrees of this polynomial fraction. (See the section on ‘‘Hansen-Sargent prediction formulas’’ for some explanations.)

Identification with our guess: $v_t = \xi c_t + \tau(L)w_t + \mu_v$:

$$\begin{aligned} \xi &= (1-\beta + \beta\xi) \Rightarrow \xi = 1, \\ \mu_v &= \beta(\mu_v + \xi \mu_c) + \frac{\beta(1-\alpha)}{2}(\xi \gamma(0) + \tau(0))'(\xi \gamma(0) + \tau(0)) \\ &\Rightarrow \mu_v(1-\beta) = \beta \mu_c + \frac{\beta\sigma}{4}(\gamma(0) + \tau(0))'(\gamma(0) + \tau(0)), \end{aligned}$$

and

$$\tau(L) = \beta \left[\frac{\gamma(L) + \tau(L)}{L}\right]_+$$

Clearly ξ is uniquely determined and so is μ_v given $\tau(L)$. To see how the last equation fully determines $\tau(L)$, it is useful to use z -transforms:

$$\tau(z) = \beta \frac{\gamma(z) + \tau(z) - \tau(0) - \gamma(0)}{z},$$

Hence

$$\begin{aligned} \tau(\beta) &= \gamma(\beta) + \tau(\beta) - \tau(0) - \gamma(0), \\ \tau(0) &= \gamma(\beta) - \gamma(0), \\ \tau(z)z &= \beta(\tau(z) + \gamma(z) - \gamma(\beta)), \\ \tau(z)(z-\beta) &= \gamma(z) - \gamma(\beta), \\ \tau(z) &= \frac{\gamma(z) - \gamma(\beta)}{z-\beta}. \end{aligned}$$

which has a pole for $z = \beta$, so this is a polynomial with only nonnegative degrees. Done!

Computing asset prices

Given the expression for m_{t+1} , it is not too hard to compute the risk free rate and the risk premium. The key result is the following. Since dividends and consumption are jointly log-normal, the risk premium is given by:

$$\begin{aligned} \log E_t \left(\frac{R_{i,t+1}}{R_{f,t+1}} \right) &= E_t \log \left(\frac{R_{i,t+1}}{R_{f,t+1}} \right) + \frac{1}{2} \text{Var}_t \log (R_{i,t+1}), \\ &= -\text{Cov}_t(\log R_{i,t+1}, \log m_{t,t+1}), \\ &= -\text{Cov}_t(\log R_{i,t+1}, \log m_{t,t+1} - E_t \log m_{t,t+1}), \end{aligned}$$

i.e. the covariance with the innovation to the SDF. In our case, this innovation is:

$$\begin{aligned} \log m_{t+1} - E_t \log m_{t+1} &= -(\Delta \log C_{t+1} - E_t \Delta \log C_{t+1}) + (1 - \alpha)(v_{t+1} - E_t v_{t+1}) \\ &= -\gamma(0)w_{t+1} + (1 - \alpha)(\gamma(0)w_{t+1} + \tau(0)w_{t+1}) \\ &= [-\gamma(0) + (1 - \alpha)\gamma(\beta)]w_{t+1}. \end{aligned}$$

Where the last line follows from $\tau(0) = \gamma(\beta) - \gamma(0)$. The first term $\gamma(0)$ is the standard one (risk premium associated with log utility). The second one is new and depends on the future path of consumption as measured by $\gamma(\beta) = \sum_{j \geq 0} \gamma_j \beta^j$. This is the *long-run risk* which now plays a role with EZ utility. It is directly related to the idea that the “intertemporal composition of risk matters”. Note that if $\beta \rightarrow 1$, this term can be substantial and high for high risk aversion. This all comes from the form of the SDF with Epstein-Zin utility – future utility matter. So even if consumption does not change much today, if it changes in the future, the utility reacts. (More generally, what matters is that the future value changes – it could change either because of expected changes in future consumption, or because of higher risk.)

8.5 Log-linearization (Campbell 1993 AER/1996 JPE)

This is the third main approach to Epstein-Zin utility. Campbell uses some log-linear approximations to derive implications even if the IES is not one. (Some of these log-linear approximations can be useful more generally, which is why I go over them.) The first use is to show analytically some implications of EZ. The second use is to estimate the model. Just like Campbell and Shiller did a log-linear approximation of the return, Campbell writes a log-linear approximation to the budget constraint. He writes the budget constraint as:

$$W_{t+1} = R_{t+1}^m (W_t - C_t),$$

where W_t is wealth and R_{t+1}^m is the return on wealth (the “market return”, if by market we include all savings instruments). Now do all the steps similar to Campbell-Shiller:

$$\begin{aligned} \frac{W_{t+1}}{W_t} &= R_{t+1}^m \left(1 - \frac{C_t}{W_t} \right), \\ \Delta \log W_{t+1} = w_{t+1} - w_t &= r_{t+1}^m + \log \left(1 - \frac{C_t}{W_t} \right), \end{aligned}$$

$$\begin{aligned} \log \left(1 - \frac{C_t}{W_t} \right) &= \log(1 - \exp(c_t - w_t)), \\ &\simeq k + \left(1 - \frac{1}{\rho} \right) (c_t - w_t). \end{aligned}$$

with $\rho = \frac{W-C}{W} < 1$. This yields:

$$\Delta w_{t+1} = r_{t+1}^m + k + \left(1 - \frac{1}{\rho}\right)(c_t - w_t). \quad (12)$$

We can rewrite this equation as:

$$\Delta c_{t+1} + (c_t - w_t) - (c_{t+1} - w_{t+1}) = r_{t+1}^m + k + \left(1 - \frac{1}{\rho}\right)(c_t - w_t), \quad (13)$$

$$\frac{1}{\rho}(c_t - w_t) - (c_{t+1} - w_{t+1}) = r_{t+1}^m + k - \Delta c_{t+1},$$

$$(c_t - w_t) - \rho(c_{t+1} - w_{t+1}) = \rho(r_{t+1}^m + k - \Delta c_{t+1}),$$

Iterating forward yields:

$$\begin{aligned} c_t - w_t &= \sum_{j \geq 1} \rho^j (r_{t+j}^m + k - \Delta c_{t+j}), \\ &= \frac{k\rho}{1-\rho} + \sum_{j \geq 1} \rho^j (r_{t+j}^m - \Delta c_{t+j}). \end{aligned}$$

Just like in the Campbell-Shiller decomposition, this is just an *accounting identity*, which holds ex-post as well as ex-ante. This holds also in expectation:

$$c_t - w_t = \frac{k\rho}{1-\rho} + E_t \sum_{j \geq 1} \rho^j (r_{t+j}^m - \Delta c_{t+j}).$$

When the consumption-wealth ratio is high, it means that either future returns will be high or future consumption growth will be low (so that the C/W ratio returns to its average).

Another way to state this equality is to apply the operator $E_{t+1} - E_t$ to equation 13 (this operator cancels all terms known at time t):

$$(E_{t+1} - E_t) \Delta c_{t+1} = (E_{t+1} - E_t)(c_{t+1} - w_{t+1}) + (E_{t+1} - E_t) r_{t+1}^m$$

$$\begin{aligned} c_{t+1} - E_t c_{t+1} &= (E_{t+1} - E_t) \sum_{j \geq 1} \rho^j (r_{t+1+j}^m - \Delta c_{t+j}) + r_{t+1}^m - E_t r_{t+1}^m \\ &= (E_{t+1} - E_t) \sum_{j \geq 0} \rho^j r_{t+1+j}^m - (E_{t+1} - E_t) \sum_{j \geq 1} \rho^j \Delta c_{t+1+j}, \end{aligned}$$

so that a good news for consumption must come from either a good current return, a good news about future returns, or a downward revision in consumption growth in the future.

Euler equations

Up to now we have only log-linearized the budget constraint. We can use this empirically to measure what explains changes in consumption (just like we did with returns), but the interesting part is to note that consumption, wealth and returns are also tied by the optimality of consumer choice. *This will allow us to substitute either consumption or the market returns out of the SDF.* Campbell thinks the consumption data is low quality, so he substitutes out consumption (hence the title of the paper). Other people prefer to substitute out returns.

To do this, we assume that all second moments (variances and covariances) are constant. [There is a bit of a consistency issue here since expected returns are time-varying while second moments are constant – a bit tricky, but it seems the approximation is OK – see the sections of the paper with heteroskedasticity – this is in any case a useful source of intuition.]

Recall the SDF formula with EZ utility when we substitute out the return on wealth:

$$m_{t,t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\rho(\alpha-1)}{\rho-1}} R_{m,t+1}^{\frac{\rho-\alpha}{1-\rho}},$$

When we write the Euler equation for the market return, we obtain:

$$E_t \left(\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\rho}{\rho-1}(\alpha-1)} R_{m,t+1}^{\frac{\rho-\alpha}{1-\rho}+1} \right) = 1.$$

This leads to:

$$\log \beta - \rho \frac{\alpha-1}{\rho-1} E_t (\Delta \log C_{t+1}) + \frac{1-\alpha}{1-\rho} E_t \log (R_{m,t+1}) + \zeta = 0,$$

where ζ is a constant regrouping the (constant) conditional variances and covariances. Hence a first result:

$$E_t \Delta \log C_{t+1} = k + \frac{1}{\rho} E_t (r_{m,t+1}),$$

where k is a constant. Expected consumption growth moves proportionally to the expected stock return, with the IES $1/\rho$ governing the proportionality.

I now derive the equation for the excess returns. Start from the EE for any return and for the risk-free rate:

$$E_t (m_{t+1} R_{i,t+1}) = 1$$

$$E_t \log m_{t+1} + \frac{1}{2} Var_t \log m_{t+1} + E_t \log R_{i,t+1} + \frac{1}{2} Var_t \log R_{i,t+1} + Cov_t (\log m_{t+1}, \log R_{i,t+1}) = 0$$

$$E_t \log m_{t+1} + \frac{1}{2} Var_t \log m_{t+1} + E_t \log R_{f,t+1} = 0$$

Subtracting these equations yields

$$\begin{aligned} \log E_t \frac{R_{i,t+1}}{R_{f,t+1}} &= E_t \log \frac{R_{i,t+1}}{R_{f,t+1}} + \frac{1}{2} Var_t \log R_{i,t+1}, \\ &= -Cov_t (\log m_{t+1}, \log R_{i,t+1}), \\ &= \frac{\rho(\alpha-1)}{\rho-1} Cov_t (\Delta \log c_{t+1}, \log R_{i,t+1}) + \frac{\rho-\alpha}{1-\rho} Cov_t (\log R_{m,t+1}, \log R_{i,t+1}). \end{aligned} \quad (14)$$

You can see how both consumption growth and the market return are risk factors in this equation.⁹

Campbell's last step: use the consumption EE $E_t \Delta c_{t+1} = k + \frac{1}{\rho} E_t r_{m,t+1}$ in the PVBC found above:

$$c_{t+1} - E_t c_{t+1} = (E_{t+1} - E_t) \sum_{j \geq 0} \rho^j r_{t+1+j}^m - (E_{t+1} - E_t) \sum_{j \geq 1} \rho^j \Delta c_{t+1+j}.$$

Thus

$$\begin{aligned} c_{t+1} - E_t c_{t+1} &= (E_{t+1} - E_t) \sum_{j \geq 0} \rho^j r_{t+1+j}^m - \frac{1}{\rho} (E_{t+1} - E_t) \sum_{j \geq 1} \rho^j \Delta r_{m,t+1+j} \\ &= r_{m,t+1} - E_t r_{m,t+1} + \left(1 - \frac{1}{\rho} \right) (E_{t+1} - E_t) \sum_{j \geq 1} \rho^j r_{m,t+1+j}. \end{aligned} \quad (15)$$

⁹Campbell uses a different notation, but he has the same results: $\sigma = \frac{1}{\rho}$, $\alpha = \gamma$, and $\theta = \frac{1-\gamma}{1-\frac{1}{\rho}} = \frac{1-\alpha}{1-\rho}$. So $1-\theta = \frac{\rho-\alpha}{1-\rho}$ and $\frac{\theta}{\sigma} = \theta\rho = \frac{(1-\alpha)\rho}{1-\rho}$.

We can use this equation in 14 to substitute out consumption. Hence we get what he calls the ‘‘CAPM+’’ formula:

$$\begin{aligned} \log E_t \frac{R_{i,t+1}}{R_{f,t+1}} &= \left(\frac{\rho(\alpha - 1)}{\rho - 1} + \frac{\rho - \alpha}{\rho - 1} \right) Cov_t(r_{m,t+1}, r_{i,t+1}) \\ &\quad + \frac{\rho(\alpha - 1)}{\rho - 1} \left(1 - \frac{1}{\rho} \right) Cov_t \left((E_{t+1} - E_t) \sum_{j \geq 1} \rho^j r_{m,t+1+j}, r_{i,t+1} \right) \\ &= \alpha Cov_t(r_{m,t+1}, r_{i,t+1}) + (\alpha - 1) Cov_t \left((E_{t+1} - E_t) \sum_{j \geq 1} \rho^j r_{m,t+1+j}, r_{i,t+1} \right) \end{aligned}$$

The novelty is that expectations of future returns now matter. Investors dislike assets that do badly when the market does badly (the first term, which is just the usual CAPM effect), but they also like/dislike assets which do badly *when expected future returns are bad*. Whether this is a ‘‘like’’ or ‘‘dislike’’ depends on whether α (risk aversion) is greater or smaller than 1. Intuition: if future returns are good, this would lead you to decrease consumption today to invest more (substitution effect), but the future returns also would make you invest less through the wealth effect.

This formula can be implemented empirically if you use a VAR to measure the news to future market returns, $(E_{t+1} - E_t) \sum_{j \geq 1} \rho^j r_{m,t+1+j}$. Campbell (1996) does this. He measures the market return as a weighted average of the SP500 return and a return on human capital. Because human capital price is unobserved, one needs to make an assumption - Campbell assumes that the returns are the same than on the SP500. Alternative assumptions are possible (e.g. constant return, or a more general model, see Baxter and Jermann (1997), Lustig, Verdelhan and Van Nieuwerburgh (2010)). Campbell then tests this equation in the cross-section: do assets which have high average returns also have high covariance with one of these two terms? He finds some support for the mechanism.

Note: by substituting out returns rather than consumption, you can obtain a similar formula where the risk premium on an asset is determined by its covariance not only with current consumption growth, but also with news about future consumption growth. The ‘‘Tallarini-Hansen’’ computations above are one example of these (when $IES = 1$). Malloy, Moskowitz and Vissing-Jorgensen (2009) use such an approach to estimate a model that uses stockholders’ consumption data.

8.6 Savings and Portfolio Choice with Recursive Utility

Weil (QJE 1990, REStud 1993) shows how to solve portfolio problems with recursive utility. (Basically you can do guess and verify.) This allows to generalize naturally the results with CRRA utility, except now we see the two coefficients (RA and IES) showing up in different terms. Weil studies for instance the effect of interest-rate uncertainty on savings. The Bellman equation is:

$$\begin{aligned} V(W) &= \max_c \left((1 - \beta)c^{1-\rho} + \beta (EV(W')^{1-\alpha})^{\frac{1-\rho}{1-\alpha}} \right)^{\frac{1}{1-\rho}}, \\ W' &= R(W - c), \\ \log R &\text{ iid } N(\mu, \sigma^2). \end{aligned}$$

Weil shows that interest rate uncertainty leads to more savings iff the $IES < 1$. In the same spirit, Gourio (2010) studies a RBC model with time-varying risk of disaster - the uncertainty about the rate of return leads to less savings, investment, and output.

8.7 Other “exotic” preferences

There are many more “exotic” preferences that people have used in asset pricing, for instance ambiguity (aversion to uncertainty as opposed to risk). To know more, a nice paper is “Exotic preferences for macroeconomics”, by Backus, Routledge and Zin (NBER macro annual, 2005).

9 Yield curve

Here we discuss the pricing of bonds (without default risk), aka fixed income. We start with some definitions, then use the standard consumption-based asset pricing to establish some general intuition. Finally, while in principle the same methods and results can be applied to bonds and to stocks, the methodologies used to formulate and estimate a model are somewhat different (e.g. a large share of the bond pricing literature uses latent factor models, see below). There are some signs of convergence, though.

Bond basics A zero-coupon n period bond is a claim to a sure payoff of 1 at time $t + n$. The price is denoted $P_t^{(n)}$ and it satisfies the recursion:

$$\begin{aligned} P_t^{(n)} &= E_t \left(M_{t+1} P_{t+1}^{(n-1)} \right), \\ P_t^{(0)} &= 1. \end{aligned}$$

We define the yield of a bond with maturity n at time t through the equation

$$P_t^{(n)} = \frac{1}{\left(1 + Y_t^{(n)}\right)^n},$$

i.e. it is the per period (e.g. per year) return that you get if you buy a bond today *and hold it until it matures*. The yield curve is a plot of $Y_t^{(n)}$ as a function of n , for fixed t . But of course, if you sell the bond before it matures, you may make a capital gain or loss and your realized return will not be equal to $1 + Y_t^{(n)}$. There is one exception: the one-period rate, which will mature: $Y_t^{(1)}$ is the *sure* return in this case.

The holding period return is the return if you buy a bond of maturity n at time t and sell it back at time $t + 1$:

$$R_{t+1}^{(n)} = \frac{P_{t+1}^{(n-1)}}{P_t^{(n)}}.$$

With $p = \log(P)$, $y = \log(1 + Y)$, and $hpr = \log R$, we have

$$\begin{aligned} p_t^{(n)} &= -ny_t^{(n)}, \\ hpr_{t+1}^{(n)} &= p_{t+1}^{(n-1)} - p_t^{(n)}. \end{aligned}$$

The forward rate is the rate at which you commit today to borrow/lend for one period, N periods from now:¹⁰

$$1 + F_t^{(N \rightarrow N+1)} = \frac{P_t^{(N)}}{P_t^{(N+1)}}.$$

¹⁰Suppose you buy one n -period zero-coupon bond and simultaneously sell x units of a $n + 1$ -period zero-coupon bond. Today's cash-flow is $xP_t^{(n+1)} - P_t^{(n)}$. At time $t + n$, you receive 1 and at time $t + n + 1$ you have to give back x . Choose x such that today's cash flow is zero. As a result, $F_t^{n \rightarrow n+1} = P_t^{(n)} / P_t^{(n+1)}$.

Facts about the yield curve are alternatively stated in terms of prices, yields, holding period returns, or forward rates. Due to the relations above, these are all related.

Yield curve with iid stochastic discount factor

Suppose that $\{M_{t+1}\}$ is iid. Then the yield curve is constant and flat, i.e. $y_t^{(n)} = y_t = \bar{y}$ for all t, n .

Proof: (math) let $p = E(M)$. We prove that $P_t^{(n)} = p^n$ for all t . This is true for $n = 0$ and $n = 1$ obviously. By induction, $P_t^{(n)} = E_t \left(M_{t+1} P_{t+1}^{(n-1)} \right) = E_t \left(M_{t+1} p^{n-1} \right) = p \cdot p^{n-1} = p^n$. If $P_t^{(n)} = p^n$, then $y_t^{(n)} = \bar{y}$ with $1 + \bar{y} = \frac{1}{p}$. End. Note: as you can see from the proof, we only need that $E_t(M_{t+1})$ is constant over time.

Proof: (econ) If $E_t(M_{t+1})$ is constant, then the short-rate $y_t^{(1)}$ is constant: $y_t^{(1)} = y_{t+1}^{(1)} = \dots = y_{t+k}^{(1)} = \dots$. We can then perfectly replicate a long-term bond using short-term bonds. Suppose you want a payoff of 1\$ at time $t + T$. You can invest in the one-period asset, then “roll over” i.e. reinvest in the same one-period asset. If the short rate is constant, you know at which rate you will be able to reinvest in the future. Your total return will be $(1 + y_t^{(1)})^T$. By no arbitrage, this must be $1 + y_t^{(T)}$. Note that it is critical for this argument that you know with certainty the path of future short-term rates.

Conclusion: if the short-rate is constant, all yields are constant, and hence bonds of all maturities have risk-free returns! (Whereas in general, only the one-period bond has a return that is risk-free.)

Nominal discount factors

Usually we use a real discount factor M_{t+1} (e.g., $M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$ depends on real consumption growth), and discount real payoffs (e.g. dividends need to be adjusted for inflation using the CPI). However, one can also use a nominal stochastic discount factor and discount nominal payoffs. That’s basically the difference between nominal and real interest rates.

Technically, if you consider the FOC for buying a nominal asset with price P_t and payoff D_t^{nom} , you find

$$\frac{P_t}{q_t} = E_t \left(M_{t+1} \frac{D_t^{nom}}{q_{t+1}} \right),$$

where q_t = price index (CPI). Hence, inflation is $\pi_{t+1} = q_{t+1}/q_t$. Why this FOC? I consider the real cost of buying the asset, and the real payoff, and apply the real SDF. Alternatively, we can write

$$\begin{aligned} P_t &= E_t \left(\frac{M_{t+1} q_t}{q_{t+1}} D_t^{nom} \right), \\ &= E_t \left(\frac{M_{t+1}}{\pi_{t+1}} D_t^{nom} \right), \end{aligned}$$

and we can define a nominal SDF

$$M_{t+1}^{nom} = \frac{M_{t+1}}{\pi_{t+1}},$$

so that for any nominal return R_{t+1}^{nom} ,

$$1 = E_t \left(M_{t+1}^{nom} R_{t+1}^{nom} \right).$$

In particular, the bond pricing recursion holds for nominal bonds using this nominal SDF: $P_t^{(n)} = E_t \left(M_{t+1}^{nom} P_{t+1}^{(n-1)} \right)$.

The key facts on nominal bonds and the nominal yield curve

See table 1, copied from Cochrane: New Facts in Finance, and table 2 below. These facts are obtained for the nominal US federal gov’t debt market. Assuming that the US gov’t does not default, these payoffs

are riskless in nominal terms but not in real terms (i.e. there is inflation risk). While the US federal gov't debt market is one of the largest and most liquid in the world, there are other very large markets of debt, including local government debt (state and municipalities) in the US, and mortgage debt (including agency debt, i.e. debt issued or guaranteed by Fanny Mae and Freddy Mac and essentially guaranteed by the US gov't).

- On average the yield curve is somewhat upward sloping; i.e. the yield on long-term bonds is larger than on short-term bonds.

- The slope of the yield curve, i.e. the difference between a long rate (≥ 5 years) and a short rate (≤ 1 year) is correlated with the business cycle: an inverted yield curve predicts a recession, and at the trough of the recession, the yield curve is steeply upward sloping.

- Long-term bond prices are fairly volatile; the std dev of the 10y return is about 8% per year, i.e. almost half that of stocks. In terms of yields, the std dev of yields as a function of maturity is hump-shaped. (But of course a given change in yield has a much bigger effect on the bond price for long maturity bonds.¹¹)

n (years)	$E(hpr)$	$s.e.$	$\sigma(hpr)$
1	5.83	.42	2.83
2	6.15	.54	3.65
3	6.40	.69	4.66
4	6.40	.85	5.71
5	6.36	.98	6.58

Table 1

- All yields (long and short) are highly correlated - they tend to move up and down together a lot;

- more precisely, one can do principal components to find the factors which move yields. The first, by far most important factor is the "level": all yields move up and down together; second, there is a "slope" effect i.e. long term yields and short term yields move in opposite direction; last, there is a "curvature" effect i.e. the concavity of the yield curve changes somewhat.

- violation of the expectation hypothesis & predictability of bond returns. See below.

Expectation hypothesis (EH)

The EH states that the expected log (holding period) returns on all bonds is the same:

$$E_t hpr_{t+1}^{(n)} = E_t hpr_{t+1}^{(1)} = y_t^{(1)}, \text{ all } n \geq 0.$$

This can be shown to be equivalent to: the N-period (log) yield is the average of expected future one-period (log) yields:

$$y_t^{(N)} = \frac{1}{N} E_t \left(y_t^{(1)} + y_{t+1}^{(1)} + \dots + y_{t+N-1}^{(1)} \right).$$

Another equivalent statement is, the forward rate equals the expected future spot rate (in logs):

$$f_t^{N \rightarrow N+1} = E_t \left(y_{t+N}^{(1)} \right).$$

(Exercise: prove the equivalence between these statements!)

¹¹If you have not done finance before, this is a duration effect. To illustrate it, compute the price of a 10y bond a 1y bond when both yields are 5%, then when both yields are 3%, and compare the % change in bond prices.

A weaker form of the EH is that these relations hold “up to a constant”, i.e.

$$\begin{aligned} E_t hpr_{t+1}^{(n)} &= y_{t+1}^{(1)} + \text{constant}, \\ y_t^{(N)} &= \frac{1}{N} E_t \left(y_t^{(1)} + y_{t+1}^{(1)} + \dots + y_{t+N-1}^{(1)} \right) + \text{constant}, \\ f_t^{N \rightarrow N+1} &= E_t \left(y_{t+N}^{(1)} \right) + \text{constant}, \end{aligned}$$

where the constant may depend on maturity n but not on time t .

Key point: the EH is almost assuming risk-neutrality - expected returns should be the same on all assets. (It is not quite that, because it is in logs instead of levels.)

Key intuition: under the EH, if the long-term yield is high today relative to the short yield, it must be that the short yield will rise in the future, so that if you invest in short rate only every period you will end up getting the same return at the end. A two period example:

$$2y_t^{(2)} = y_t^{(1)} + E_t y_{t+1}^{(1)}$$

hence if the yield curve is upward sloping, $y_t^{(1)} < y_t^{(2)}$, it must be that $y_t^{(2)} < E_t y_{t+1}^{(1)}$ i.e. the short rate will rise, so rolling-over one-period investments will bring the same return at the end as the long-term investment.

In the data, the expectation hypothesis does not work very well (though it is a decent start). One way to say this is to say that the expected return on bonds is forecastable, $E_t hpr_{t+1}^{(n)} - y_{t+1}^{(1)} = \alpha + \beta X_t$ i.e. there are times when investing in long-term bonds brings excess returns. One way to summarize the results is to go back to the example explaining the EH - in the data, on average, when short < long, the short yield does not increase enough in the future, so there is a positive excess return to borrowing short term and buying long-term bonds.

More precisely, the variable X_t that researchers use to predict returns on long-term bonds is usually based on current yields or forward rates. Fama and Bliss use the difference between the forward rate at time $t + n$ and the current short rate, to forecast the maturity n bond excess return. Cochrane and Piazzesi find that a particular combination of forward rates forecast all maturities of excess bonds returns. (Refs: Fama and Bliss (1988), Campbell and Shiller (1991), Cochrane and Piazzesi (2005)).

The expectation hypothesis is often tested through to the following equation:

$$y_{t+1}^{n-1} - y_t^n = \alpha + \beta_n \left(\frac{y_t^n - y_t^1}{n-1} \right) + \varepsilon_{t+1}.$$

The expectation hypothesis implies that $\beta_n = 1$. In the data, $\beta_n < 1$, often negative, and decreasing with the horizon n .

Table 1: Expectation Hypothesis Tests

$n = 2$	$n = 3$	$n = 4$	$n = 5$
<i>Slope Coefficients - 1961-1979</i>			
-1.03	-1.52	-1.55	-1.43
[0.65]	[0.71]	[0.83]	[0.96]
<i>Slope Coefficients - 1988-2006</i>			
0.61	-0.13	-0.19	-0.21
[0.89]	[0.90]	[0.97]	[1.02]

Summing up

Just like for stocks, we need a model which explains both (i) the mean return on long-term bonds, relative to short-term bonds, (ii) the volatility of long-term bond returns, and (iii) the variation over time in the expected returns. Things are somewhat simpler since cash flows are fixed in nominal terms.

The yield curve in a standard consumption-based model

This section discusses how the consumption-based allows to think about the yield curve. We start with the real yield curve and then turn to the nominal yield curve.

Real yield curve, iid lognormal case: in this case the risk-free rate (short-rate) is constant, hence all yields are constant and equal to the short rate, that is given by $y_t^{(1)} = \log R_t^f = -\log(\beta) + \gamma\mu_c - \gamma\frac{\sigma_c^2}{2}$.

Real yield curve, serial correlation in growth rates:

Key intuition: suppose that consumption growth is positively serially correlated. Hence, an increase in consumption today increases expected future consumption growth. Hence, it increases interest rates, since interest rates depend on expected future consumption growth. This reduces the value of long-term bonds. (Bond prices move inversely to interest rates: when the interest rate goes up, the value goes down since we discount a fixed payoff at a higher rate.) Hence, long-term bond prices fall when consumption goes up. This implies that long-term bonds allow you to hedge against consumption shocks, i.e. they are insurance rather than being risky. Hence, they will have a negative excess return. This means that the average return on the long-term bond is less than the average return on the short-term bond. This implies that the yield curve is downward sloping on average.¹²

Math proof: two ways to do this. *Method 1:* an example where we can compute everything fairly easily. Find the price of a one-period and a two-period bond, assuming expected utility/CRRA preferences and

$$\Delta \log C_{t+1} = (1 - \delta)\mu + \delta\Delta \log C_t + \sigma\varepsilon_{t+1},$$

with ε_t iid $N(0,1)$. Note that the state of the economy is last period's consumption growth rate: it is a sufficient statistic to compute all conditional expectations of the SDF and payoffs. Find the expected return on the two-period bond and the one-period bond. When is the yield curve upward sloping? Under which conditions is the yield curve upward sloping on average? This is a bit algebra-intensive, but it is a good exercise.

Method 2: a bit more general and insightful. first, recall the Campbell-Shiller log-linear approximation,

$$r_{i,t+1} - E_t r_{i,t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \Delta d_{i,t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} r_{i,t+1+j},$$

and consider a bond with an infinite maturity (a consol) that repays 1 unit every period, forever. The first term is zero, since expected cash flows do not change over time. Hence,

$$r_{c,t+1} - E_t r_{c,t+1} = - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{c,t+1+j}.$$

Moreover, we have

$$E_t r_{t+1} = \text{constant} + \gamma E_t \Delta \log c_{t+1},$$

i.e. expected returns are high when expected consumption growth is high, where γ is the inverse of the IES.

This equation also implies:

$$E_t r_{t+1+j} = \text{constant} + \gamma E_t \Delta \log c_{t+1+j},$$

and if $\Delta \log c_{t+1} = (1 - \delta)\mu + \delta\Delta \log c_t + \sigma\varepsilon_{t+1}$, then we have,

$$\begin{aligned} (E_{t+1} - E_t) r_{t+1+j} &= \gamma (E_{t+1} - E_t) \Delta \log c_{t+1+j} \\ &= \gamma \delta^j \varepsilon_{t+1}, \end{aligned}$$

¹²You can prove by induction on maturity n that, if $E(\log R^{(n)}) < E(\log R^{(1)})$ for all n , then $Ey^{(n)} < Ey^{(n-1)} < \dots < Ey^{(1)}$. Intuitively, if average returns are lower on long-term bonds, they will have on average lower yields since the yield is the return that you get if you buy and hold to maturity.

so that a good consumption growth surprise increases expected future consumption growth is $\delta > 0$, and hence leads to higher interest rates and higher expected returns in the future. Using the C-S decomposition, we find

$$r_{c,t+1} - E_t r_{c,t+1} = -\gamma \sum_{j=1}^{\infty} \rho^j \delta^j \varepsilon_{t+1} = \frac{-\gamma \rho \delta}{1 - \rho \delta} \varepsilon_{t+1}.$$

A good news today ($\varepsilon_{t+1} > 0$) leads to higher future returns later since cons. growth will be higher in the future. Hence, price fall today since discount rates increase. The consol risk premium premium is then

$$\begin{aligned} E_t \left(\log \frac{R_{t+1}^c}{R_{t+1}^f} \right) &= \gamma \text{Cov}_t(\Delta \log c_{t+1}, r_{c,t+1}) \\ &= \frac{-\gamma^2 \rho \delta}{1 - \rho \delta} \sigma_{\varepsilon}^2. \end{aligned}$$

Hence, this premium is negative if $\delta > 0$, and is positive if $\delta < 0$, as in the intuition above.

Extension to other shocks

Just like before, we can think of having shocks to uncertainty or shocks to expected future growth rates. The analysis for persistent consumption growth above directly gives the response of bond prices to changes in expected future growth rates. On the other hand, higher uncertainty drives the short rate down (that's our basic result from PS2), hence long-term bond price go up. Hence, the long term bond is a hedge against this shock. If this shock is "bad" for the agent (i.e. marginal utility goes up), long-term bonds will have a negative risk premium. Hence the yield curve will slope down on average. (This, and the previous analysis for $\delta > 0$, imply that in the long-run risk model the real yield curve is downward sloping.)

Nominal yield curve

Most facts are for the nominal yield curve. A model of the nominal yield curve requires a specification of inflation and its impact on the SDF. – e.g., in the basic CRRA model, the covariance of consumption growth with inflation. For instance, suppose you buy a 1-period nominal bond, but inflation is positively correlated with consumption growth. Then, the bond payoff will be higher in real terms when inflation is low i.e. when consumption growth is low, so the bond is a hedge against consumption growth. So this would give a negative risk premium over a "pure" risk-free asset. Inversely, if inflation is negatively correlated with consumption growth, you would get a positive risk premium. (PS4 asks you to do such simple computations.) Note that what matters is the inflation surprise – expected inflation is already built into the price of course. At short horizon inflation is easy to forecast (since it is persistent) so the inflation risk premia are more likely to play a role at long horizon.

As in the analysis of the real yield curve, we can use the C-S decomposition for a consol,

$$r_{c,t+1} - E_t r_{c,t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \Delta d_{c,t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} r_{c,t+1+j},$$

and note that $d_{c,t+1+j} = \log \left(\frac{1}{P_{t+1+j}} \right)$ is the log real payoff of a consol. Hence,

$$r_{c,t+1} - E_t r_{c,t+1} = - (E_{t+1} - E_t) \sum_{j=0}^{\infty} \Delta p_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} r_{c,t+1+j},$$

and assuming that both consumption growth and inflation follows AR(1), i.e.

$$\Delta \log c_{t+1} = (1 - \delta_c) \mu_c + \delta_c \Delta \log c_t + \sigma_c \varepsilon_{c,t+1},$$

$$\Delta p_{t+1} = (1 - \delta_p)\mu_p + \delta_p \Delta p_t + \sigma_p \varepsilon_{p,t+1},$$

we have

$$r_{c,t+1} - E_t r_{c,t+1} = \frac{-\delta_p}{1 - \delta_p} \varepsilon_{p,t+1} + \frac{-\gamma \rho_c \delta_c}{1 - \rho_c \delta_c} \varepsilon_{c,t+1},$$

and hence the risk premium is the risk premium for the second part (the same as above, i.e. the real risk premium) and the first part, the inflation risk premium:

$$\begin{aligned} E_t \left(\log \frac{R_{t+1}^c}{R_{t+1}^f} \right) &= \gamma \text{Cov}_t (\Delta \log c_{t+1}, r_{c,t+1}), \\ &= -\frac{\gamma^2 \rho \delta}{1 - \rho \delta} \sigma_\varepsilon^2 - \frac{\gamma \delta_p}{1 - \delta_p} \text{Cov}(\varepsilon_{c,t+1}, \varepsilon_{p,t+1}), \end{aligned}$$

e.g. if $\text{Cov}(\varepsilon_{c,t+1}, \varepsilon_{p,t+1}) > 0$, inflation is procyclical, and the risk premium is negative, while if inflation is countercyclical, the risk premium is positive. This is consistent with intuition above.

Macro intuition: supply vs. demand shocks

What's the correlation of inflation and consumption growth in the data? In the data set I gave you (1929-2009), it is somewhat positive (0.19), however, it is negative in the post WWII sample (1950-2009), i.e. the Great depression (deflation and low growth) matters a lot for this correlation. From a macro point of view, this is really about supply vs. demand shocks – supply shocks such as oil price increases lead to low growth and high inflation, as in the 1970s, while demand shocks like 2008 or 1929 lead to low inflation and low growth. See "Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds" (Campbell, Sunderam and Viceira), who argue that this correlation is time-varying and can account for some interesting dynamics.

News about inflation

More generally, if you do not hold a 1-period asset, the *realized* return on a long-term bond will be lower when there is news of higher future inflation. It will also be lower when there are news about future economic growth (since the real interest rate will go up then driving long-term bond prices down). Key question is, are these news good or bad today (i.e. for the SDF today).

Piazzesi and Schneider (NBER macro 2006) use Epstein-Zin preferences, to analyze this question. (They also incorporate learning, but that's something on top of the basic point above.) With Epstein-Zin, good states are states with high current or future consumption growth. Hence the key question is whether future inflation is correlated with current or future cons. growth. In post WWII data, they find a strongly negative correlation. (A simple suggestive correlation is between the 5-year moving average of inflation and cons. growth: it is -.60.)

More consumption-based asset pricing

So far we have looked at the basic consumption-based model. Of course you can use a different model to fit the yield curve. Some of the finer facts about the yield curve require time-varying risk premia, so the model needs to be rich enough.

Wachter (JFE 2007?) extends the Campbell-Cochrane model. Recall CC has constant interest rates, so the yield curve was flat. She changes the preference specification to obtain procyclical interest rates, which allows her to match some features of the yield curve. Adrien Verdelhan's paper considers a different preference specification, which implies countercyclical interest rates and helps him fit the forward premium puzzle (exchange rate premium).

Bansal and coauthors (see also Hasseltoft) have used the long-run risk model to fit the yield curve. This relies on correlations between long-run inflation and long-run consumption, and on stochastic volatility of inflation. There is some evidence for these correlations in the data.

TIPS

TIPS are inflation-indexed bonds, which have been traded in the United States since 1997. In the UK inflation indexed-bonds have been traded for a longer period of time. In principle TIPS are the analog to real bonds in our model. Because so far there has been relatively little data, there has not been much work attempting to fit models to TIPS, but this is becoming more and more doable now. In principle they allow for a neat test of the models: we should fit both real and nominal yields. In PS4, you will study some particularities of the pricing of TIPS (they are adjusted up but not down with inflation).

Historical data

Something that seems fascinating is to analyze the nominal and real yield curve under different monetary regimes, e.g. during the gold standard. Inflation dynamics were very different (mean-reversion in the price level!) which should lead to very different yield curves.

Link with macro models

There is of course a huge literature on DSGE models and monetary policy. These models have, however, constant or nearly constant risk premia, so in these models the EH works very well. Monetary policy sets the short term rate which then influences the long-term rate through the EH. As argued by Atkeson and Kehoe (NBER macro 2008) and others, these models are thus inconsistent with the basic facts about asset prices. [See, for instance, Palomino and coauthors, and Swanson and coauthors, for related work.]

Latent factor models

A large literature on the yield curve does not use microfounded (consumption-based) models but rather uses “affine models”, e.g.

$$\log M_{t+1} = -\frac{\lambda'_t \lambda_t}{2} - \lambda'_t \varepsilon_{t+1} - r_t^{(1)},$$

where $r_t^{(1)}$ is the short rate (so that $E_t M_{t+1} = 1/(1 + r_t^{(1)})$), and

$$\begin{aligned}\lambda_t &= \lambda_0 + \lambda_1 X_t, \\ r_t^{(1)} &= \delta_0 + \delta_1 X_t\end{aligned}$$

where X_t are observables or unobservables (aka latent) factors, which are assumed to follow a VAR(1):

$$X_{t+1} = \mu + \Phi X_t + \Sigma \varepsilon_{t+1}.$$

Note that λ_t is a vector above.

These models are relatively tractable (computing bond prices is easy thanks to the “affine” i.e. linear structure), and they allow for a time-varying price of risk λ_t . Because they allow for several factors, some of which may be unobservables, and some of which may be interest rates, they typically fit the yield curve much better than consumption-based model. The latent factor(s), if any, must be recovered as part of the estimation, through the Kalman filter. Estimation can be conducted by MLE (but there is some work to make the maximization converge!) or through other methods, e.g. simulation-based methods like SMM. The main limitation of these models is that they are not as nicely tied to economic fundamentals, and may be harder to interpret. They may be useful for pricing, though. (Industry practice is, I believe, to use continuous time models with a few factors, which are observed interest rates.)

10 Corporate bonds

Firms which need to raise funds for investment or other expenditures can use the equity market (IPO or SEO), or they can get bank loans, or other private financing (through private equity or hedge funds), but probably the single most important source of external finance is the corporate bond market (and associated commercial paper market for short-term borrowing). There is > 6,000Bn\$ of corporate bonds outstanding in the US.¹³

The interest rate at which firms can borrow depends on the firm's credit rating, which is based on the estimated probability of default by rating agencies (most famous: Standard & Poors, Moody's, Fitch). This interest rate is higher than the rate at which the US Federal government borrows, because of credit risk and because of the lower liquidity of corporate bonds.

See table: investment grade bonds vs. high yield.

Rating	10-year default probability	Default rate (annualized)
<i>Investment-grade Bonds (IG)</i>		
AAA	0.19%	0.02%
AA+	0.57%	0.06%
AA	0.89%	0.09%
AA-	1.15%	0.12%
A+	1.65%	0.17%
A	1.85%	0.19%
A-	2.44%	0.25%
BBB+	3.13%	0.32%
BBB	3.74%	0.38%
BBB-	7.26%	0.75%
<i>Speculative-grade Bonds (aka High yield or Junk bonds)</i>		
BB+	10.18%	1.07%
BB	13.53%	1.45%
BB-	18.46%	2.04%
B+	22.84%	2.59%
B	27.67%	3.24%
B-	34.98%	4.30%
CCC+	43.36%	5.68%
CCC	48.52%	6.64%
CC	77.00%	14.70%
C	95.00%	29.96%

Table 1: Historical default experience of Bonds Rated by Fitch

As seen in the attached figures, the spread between safe bonds and less safe ones is countercyclical. A striking example is the current recession, where the interest rate on government bonds has declined significantly, while the yield on BAA bonds (which are not so risky) has increased substantially.¹⁴ Of course this is *qualitatively* not surprising: the probability of default has increased since profits fell substantially.¹⁵ However, it is not clear that this is *quantitatively* enough to explain the magnitude of the increase - while default probabilities have gone up, they are not huge. Simple risk-neutral calculation: if spread = 3%, it means you expect 3% of default with zero recovery *per year*. If recovery rate = 50%, a 3% spread means a 6% probability of default!

¹³Flow of Funds data from the Fed.

¹⁴BAA are Moody's ratings. BAA corresponds to BBB of Fitch.

¹⁵Moreover, the recovery rate on the firms assets may have fallen because the market value of these assets, should they be seized and sold, is now lower.

Very simple example – one period debt

Consider a debt of 1\$ which may not be repaid next period. If it is not repaid, you will get 0\$ instead. The value of this debt today is

$$P_t = E_t(M_{t+1}x_{t+1}),$$

where $x_{t+1} = 1$ if repaid, and $x_{t+1} = 0$ if not. Let $p = \Pr(x_{t+1} = 0)$ the probability of default. If the discount factor is simply $M_{t+1} = \beta$, then

$$P_t = \beta(1 - p)$$

is the discounted expected payoff (where you compute the expected payoff by computing the prob of default). As long as M_{t+1} is independent of x_{t+1} , we have

$$P_t = E_t(M_{t+1}) E(x_{t+1}) = \frac{1 - p}{R_t^f}.$$

But in general, if default is correlated with the SDF M_{t+1} , we have

$$P_t = E_t(M_{t+1}) E(x_{t+1}) + Cov_t(M_{t+1}, x_{t+1}).$$

Typically, defaults occur in “bad times” when the marginal utility is high i.e. M_{t+1} is high, so $Cov_t(M_{t+1}, x_{t+1}) < 0$. This new term will hence reduce the price of the asset further down. Hence, there are two reasons why default affects prices: (i) it affects the expected cashflow, (ii) it makes the asset more risky (usually), i.e. lower payoffs in bad times.

Pricing defaultable debt - an example

Consider a firm, which has a random stream of profits $\{\pi_t\}_{t=0}^{\infty}$. The PDV of this profit stream is what we call the “asset value” of the firm A_t , and it satisfies

$$A_t = E_t(M_{t+1}(\pi_{t+1} + A_{t+1})) = \sum_{k=1}^{\infty} M_{t,t+k} \pi_{t+k},$$

where $M_{t,t+k}$ is the SDF between dates t and dates $t + k$, i.e. recursively

$$M_{t,t+k} = M_{t,t+k-1} \times M_{t+k}.$$

The standard financial structure is that firms issue debt and equity claims. Once the firm has issued some debt, it pays interest on this debt, however it may decide to default at any point in time if the value of the future profits is less than the cost of repaying the debt. For simplicity suppose that the debt is a consol with constant payout c . The ex-dividend *equity* value satisfies the recursion

$$V_t = E_t(M_{t+1} \max(\pi_{t+1} - c + V_{t+1}, 0)),$$

i.e. the firm can file for bankruptcy, if it is better for the equity holders to do so. Note that the firm will not shut down if π_{t+1} is low but future profits are expected to be high! In this case the dividend payout $\pi_{t+1} - c$ may be negative, i.e. there will be equity issuance, and existing equity holders will increase their participation in the firm (or new equity holders will join – more on this below).¹⁶

As a simple illustration, suppose that $M_{t+1} = \beta$ and that profits follow a simple AR(1): $\pi_{t+1} = \rho\pi_t + (1 - \rho)\bar{\pi} + \varepsilon_{t+1}$, then there will be a threshold π^* , such that the firm defaults if $\pi_t \leq \pi^*$. And, typically, $\pi^* - c < 0$ and not = 0.

The firm default realization is thus given by the characteristic function $1_{\pi_{t+1} - c + V_{t+1} \leq 0}$. We can rewrite this in terms of cum-dividend values, P_t and

$$P_t = \max(\pi_t - c + E_t(M_{t+1}P_{t+1}), 0),$$

¹⁶In a more realistic model, the firm could also issue new debt.

and the default event is $1_{P_t=0}$.

What is the value of the firm's debt, i.e. the value of the consol? It may default, hence its value is not that of the risk-free consol.¹⁷ Its value satisfies the recursion

$$B_t = E_t \left(M_{t+1} \left((c + B_{t+1}) 1_{P_{t+1}>0} + c 1_{P_{t+1}=0} \theta A_{t+1} \right) \right),$$

i.e.: either the firm does not default, then the coupon c is paid and the consol continues next period, or the firm does default, and then the bondholders recover a fraction $\theta \in (0, 1)$ of the asset value. (Alternative assumptions are possible here regarding what value bondholders recover.) The parameter $1 - \theta > 0$ captures “bankruptcy costs”, i.e. bondholders are not able to seize the entire value of assets (e.g. because they have to pay some legal fees to get the assets).

Formally, the equity is like a call option on the asset value, while the debt is like a put option.

Given a coupon c , a profit stream $\{\pi_t\}$, and a SDF M_{t+1} , it is straightforward to use the two recursions to find the price of equity and the debt: first, solve the asset value recursion; then, solve the equity value recursion; last, solve the debt value recursion (for which you need the equity value $P_t!$).

We can define the yield on the debt y_t as the discount rate which makes the PDV of cash flows (*assuming no default*) equal to the price:

$$B_t = \sum_{k=1}^{\infty} \frac{c}{(1 + y_t)^k},$$

and of course y_t will be greater than the yield on the risk-free consol y_t^{RF} , since there is a risk of default. The spread $y_t - y_t^{RF}$ will reflect both the likelihood of default, the recovery rate θ , and the risk premium - does the firm default in bad or good states. (Note that idiosyncratic risk matters for the spread since idiosyncratic risk can trigger default.)

The total firm value is the sum of the equity and debt value, i.e.

$$F_t = V_t + B_t.$$

It is easy to check that if $\theta = 1$, then $F_t = A_t$. In general however, there is a social loss because of bankruptcy costs ($\theta < 1$), hence $F_t < A_t$.

Credit spread puzzle

The “credit spread puzzle” is that simple models fail to replicate the magnitude of the credit spread in the data, given the low historical probabilities of default (and high recovery rates conditional on default). This can be thought of as the equivalent to the equity premium puzzle: spreads are “too high”. (It may be that some of the spread is due to liquidity, hence researchers sometimes concentrate on the spread between AAA and BAA issues not between BAA and Treasuries.)

Recently a literature has developed which tries to use similar models that have been used to solve the equity premium puzzle (e.g. Campbell-Cochrane, Bansal-Yaron, disasters) and applies them to the credit spread puzzle. This seems relevant since corporate bonds are very much exposed to “tail risk” – they pay out less in deep recessions.¹⁸

Optimal capital structure

¹⁷The value of the risk-free consol is given by

$$X_t = E_t (M_{t+1} (c + X_{t+1})) = c E_t \sum_{k=1}^{\infty} M_{t,t+k}.$$

¹⁸See papers by Strebulaev and Kuehn, Hui Chen, Colin-Dufresne, Goldstein and coauthors.

What we have done so far is take as given the choice of coupon c , i.e. the choice of how much debt to issue. This framework allows to consider the optimal choice of debt vs. equity. Consider the problem at time 0 - the firm has to decide how much debt vs. equity to issue. A cost of debt is that it can lead to default, which is (socially) costly because of $\theta < 1$. In the analysis above there was no benefit of debt, so firms would decide to use only equity.

In reality one incentive to issue debt is tax treatment - interest on debt can be deducted from corporate income. Using the notation of the previous setup, if the profit is π_t , and the coupon c , the payout to equity holders is

$$(1 - \tau)(\pi_t - c) = (1 - \tau)\pi_t - c + c\tau,$$

which encourages firms to increase the coupon payment ($c\tau$ is a subsidy).¹⁹

Incorporating this into the model, we can solve for the time 0 equity and debt value:

$$V_0(c; \theta, \tau)$$

and

$$B_0(c; \theta, \tau),$$

where I have written explicitly the dependence on parameters. The manager thus decides on the optimal coupon c , to maximize the total firm value:

$$\max_c \{V_0(c; \theta, \tau) + B_0(c; \theta, \tau)\}.$$

This coupon c implies a certain amount of debt issuance at time 0, i.e. a certain leverage, and a probability of default. Basic intuition: the higher θ , the more leverage. The higher τ , the more leverage. . The more volatile the profits, the less leverage. This is the basic trade-off theory of debt, but there are alternative “theories” (e.g. pecking order)

Note that once the debt has been issued, its price is going to evolve depending on the realization of shocks. Good shocks (i.e. increase in expected profits) push up both equity and bond value.

An increase in σ – the risk shifting problem

An interesting shock is an increase in volatility σ . Because of the option-like nature of equity and debt, an increase in volatility is good for equity holders, but is bad for bondholders (they have no upside, only a downside). This leads to the risk-shifting problem: equity holders, who run the firm, have an incentive to increase risk, because debt holders pay most of the price for this. (This may be a justification for covenants, i.e. strings attached to a corporate bond issue, e.g. rules which limit how much you can invest etc.)

Debt overhang problem

Inversely, there can be situations where debt holders would reap the benefits of equity holders’ decisions. One example is the “debt overhang problem”. Suppose a firm has issued some debt and then suffers bad shocks. Bond prices fall, and equity prices fall even more. Suppose the firm needs to raise funds e.g. because of temporarily low profits. It may be impossible for the firm to issue new equity. That is because an increase in the firm’s capital will benefit bondholders – bond prices will increase because the probability of default will fall. However, the investors who are buying new equity are not internalizing this benefit. Hence, it is possible for a firm with positive NPV projects to be unable to finance them! The problem is that the seniority (the order in which people receive payoffs: debt first, then equity) is fixed, so standard debt contracts are inefficient. If it was possible to have a most senior claim issued

¹⁹In reality the comparison of the tax treatment of equity and debt is not completely obvious (capital gains based on nominal, realized gains and dividend taxes vs. personal income tax for debt interest).

(something senior to debt), that would work, but it is legally impossible ex-post in general: debt holders have the right to oppose this.²⁰

Merton and Leland Model

Merton (1974) produced a formula for the value of defaultable debt, when the asset value follows a geometric brownian motion, based on the Black-Scholes formula. This is based on a replication argument – corporate debt + equity = asset value. Leland (1994) extended this model to have taxes and bankruptcy costs, hence making it a model of the capital structure. These models are cast in continuous time but the economics are very much the same as the above example. There are lots of possible extensions: finite maturity debt instead of consol as above, jumps instead of diffusions, different recovery values, etc.

Credit Default Swaps (CDS)

The CDS market has developed over the past 10 years, allowing investors to hedge their exposure to the credit risk of a firm. The CDS market is much more liquid than corporate bonds. The main consequence is that it has become much cheaper to short corporate bonds.

Liquidity issues

Corporate bonds are not traded much. As a result, there is a liquidity premium on corporate bonds, relative to Treasuries. For instance, AAA bonds have almost no default risk, yet the spread over Treasuries is quite large. See attached picture from a recent paper by Vissing-Jorgensen and Krishnamurthy where they show that debt/gdp is highly negatively correlated with the AAA-Treasury spread.

General equilibrium and models with credit frictions

While a large literature in macroeconomics has studied the effect of credit frictions on macroeconomic dynamics (e.g. Bernanke, Gertler and Gilchrist, Kiyotaki and Moore, Carlstrom and Fuerst), the vast majority of this research is cast in linearized DSGE models which fail to reproduce the level and the volatility of risk premia. In particular, these models talk about the “external finance premium”, which corresponds approximately to a corporate bond spread. In these linearized models, the external finance premium is mostly due to idiosyncratic default risk, and the spread reflects essentially only the probability of default (i.e. there is no adjustment for aggregate risk). Hence, I think one important question is to embed financial frictions in a model with large risk premia.

A two-period model with investment²¹

These notes give a very simple, two-period, partial equilibrium model where I add the choice of investment. So far we have been pricing exogenous cash flows, but of course the reason why we are interested in the price of corporate bonds (aka their yields) is that they affect the cost of capital. For simplicity I will assume that the discount factor is simply β - no adj for risk.

At time 1, the firm buys capital k , using equity issuance s and debt: the firm issues a debt with face value b , at unit price $q(k, b)$, hence the budget constraint:

$$\chi q(k, b)b + s = k \tag{16}$$

The parameter $\chi > 1$ reflects the tax shield effect - interest expenses are subsidized. To simplify I assume the deduction is done at issuance: for each dollar of debt issued, the firm receives a subsidy $\chi - 1$ \$.

²⁰During the crisis, banks may have been affected by the debt overhang problem. When the gov't essentially guaranteed the bank debt (i.e. they would not allow the banks to fail), this mitigated the issue. Of course, this amounted to a very large bailout to bondholders.

Some researchers have suggested that new state-dependent debt contracts, with explicit contingencies, be used in the future. (e.g. debt is converted to equity if a specified event happens.)

²¹See Miao 2005 JF for a dynamic model of investment and capital structure.

At time 2, the firm produces, and obtains a profit $\pi = zk^\alpha$, where z is an idiosyncratic shock, distributed according to a cumulative distribution function H , with mean 1 : $E(z) = \int_0^\infty sh(s)ds = 1$.

The firm will default if its profits are not large enough to repay its debt, i.e. if $z < z^*$, with

$$z^*k^\alpha = b. \quad (17)$$

If the firm does default, equity holders get nothing, while bondholders share the firm profits, net of bankruptcy costs. The parameter $0 < \theta < 1$ determines the share of profits which are recovered.

Debt value is given by a standard pricing equation, assuming risk neutrality and a discount factor β :

$$q(k, b) = \beta \left(\int_{z^*}^\infty dH(z) + \int_0^{z^*} \theta \frac{zk^\alpha}{b} dH(z) \right). \quad (18)$$

The firm equity value is

$$V = \beta \int_{z^*}^\infty (zk^\alpha - b) dH(z). \quad (19)$$

The firm picks k, b, s, z^* to maximize its present discounted value, $V - s$, subject to 16, 18 and 17. This can be rewritten as

$$\begin{aligned} & \max_{k, b, z^*} \left\{ \beta \int_{z^*}^\infty (zk^\alpha - b) dH(z) - k + \chi q(k, b)b \right\}, \\ \text{s.t.} \quad & : \\ q(k, b) &= \beta \left(\int_{z^*}^\infty dH(z) + \int_0^{z^*} \theta \frac{zk^\alpha}{b} dH(z) \right), \\ z^*k^\alpha &= b. \end{aligned}$$

or

$$\max_{k, z^*} \left\{ \beta \int_{z^*}^\infty k^\alpha (z - z^*) dH(z) - k + \chi \beta \left(z^*k^\alpha \int_{z^*}^\infty dH(z) + \int_0^{z^*} \theta zk^\alpha dH(z) \right) \right\}.$$

This program can be solved by writing the two first order conditions, with respect to k and z^* , which determine the optimal investment and financing. First, we have

$$1 = \beta \alpha k^{\alpha-1} \left\{ \int_{z^*}^\infty (z - z^*) dH(z) + \chi \left(z^* \int_{z^*}^\infty dH(z) + \theta \int_0^{z^*} z dH(z) \right) \right\},$$

which can be rearranged, given that $E(z) = 1$, as:

$$1 = \beta \alpha k^{\alpha-1} \left\{ 1 + (\chi - 1) z^* (1 - H(z^*)) + (\theta \chi - 1) \int_0^{z^*} z dH(z) \right\}. \quad (20)$$

Interpretation: note that if $\chi = \theta = 1$, we have the frictionless limit: $1 = \beta \alpha k^{\alpha-1} E(z)$, the usual MPK = user cost condition. When $\theta < 1, \chi > 1$, this is not true anymore, and the tax shield χ and the bankruptcy cost θ both tend to decrease the user cost.

The second first-order condition, with respect to z^* , yields, after rearrangement

$$((1 - H(z^*)) (\chi - 1) = \chi (1 - \theta) z^* h(z^*)). \quad (21)$$

Interpretation: this equation determines z^* , and hence the optimal leverage (increasing z^* means increasing the leverage and the probability of default). The benefit of higher leverage is the left-hand side: it is the tax shield conditional on no default. The cost of higher leverage is an increase in bankruptcy costs

through a higher probability of default - the change in probability of default is $h(z^*)$ at the margin, and the default costs are proportional to z . (Note: this expression appears to be similar to Bernanke, Gertler, Gilchrist 1999 handbook paper - they assume that the pdf h is such that $z \rightarrow \frac{zh(z)}{1-H(z)}$ is increasing, which guarantees a unique solution to this equation. This assumption seems fairly weak (most distributions satisfy it, e.g. the lognormal) and we can make it too).

Note that eqn 21 determines z^* , then eqn 20 determines the investment choice k .

From 21, z^* is determined by the parameters χ, θ , and the distribution H . The comparative statics w.r.t. χ and θ is clear. With regard to H , an interesting comparative static is a change in the payoff distribution H . If H becomes more risky, at least for most distributions, we will have less debt and less capital. Intuitively, the debt contract becomes more inefficient, making the cost of the debt higher. (Note: this is ex-ante higher risk, not ex-post: this is different from risk-shifting!)

More intro

$$P_t = E_t(M_{t+1}x_{t+1})$$

Define the yield as

$$P_t = \frac{1}{1 + rc_t}$$

and the risk-free yield as

$$\frac{1}{1 + rf_t} = E_t(M_{t+1})$$

Assume $x_{t+1} = 1$ w/p $1 - p$, and $x_{t+1} = 0$ w/p p . Then if risk-neutral, i.e. if $M_{t+1} = \beta$, then

$$rc_t - rf_t \approx p.$$

More generally, if recovery rate r , i.e. $x_{t+1} = r$ w/prob p , then

$$rc_t - rf_t \approx p(1 - r).$$

If not risk-neutral, and default correlated with business cycle, then

$$rc_t - rf_t \approx p(1 - r) + Cov(M_{t+1}, x_{t+1})$$

where the last term is a risk premium, that is positive only if default is more likely when the SDF is high i.e. in bad times.

Notes on CDOs

Consider a bunch of assets, which may or may not default, $x_{it} = 1$ or $x_{it} = 0$. These assets could be mortgages, loans, credit card balances, etc.

We can price a claim to the total payout of these assets, $Z_t = \sum_{i=1}^N x_{it}$. We can also tranche the asset, i.e. we split

First, suppose the pricing kernel is $M_{t+1} = \beta$ (risk-neutrality), and further suppose that the x_{it} are uncorrelated: each is 1 w/p $1 - p$ and 0 w/p p .

Second, suppose now that they are correlated. One question is how to model in a simple way their correlation.

These assets are subject to both idiosyncratic and aggregate risk.

UG notes + problem set.

11 Heterogeneity and asset pricing (Rough notes)

So far we have considered mostly models with representative agent. As we discussed in class, this can be justified if markets are complete. Even then, however, one question that arises is, what is the aggregate utility function, i.e. the utility function of the representative agent? Perhaps this function is not a simple CRRA utility, but has a more complicated functional form.

Of course, in reality, markets are not complete, and an interesting question is what kind of asset price dynamics are generated when agents differ in, say, their income/wealth, their risk aversion, their beliefs, or their access to markets. For instance, the more risk averse agents would want to buy safe assets, and the less risk-averse agent would then sell him these safe assets, i.e. borrow from him, and invest in risky assets. The less risk averse agent is hence leveraged. As the state of the economy changes, the relative wealth of these two agents will change, and this will affect asset prices.

The first section of this handout discusses a prototypical portfolio choice problem, which is useful as general background. Then I consider a simple example, that can be solved analytically. I then discuss briefly the related literature.

Background: a standard portfolio result in continuous time (Merton)

A standard finance question is, given the process for asset prices or asset returns, to pick a portfolio to maximize utility (i.e. maximize wealth, and smooth consumption). These models are typically partial equilibrium, i.e. they assume a stochastic process for returns, and potentially for labor income as well as other state variables.

Suppose there is no labor income, CRRA utility, and the investor has to decide at each point how to allocate his wealth between a stock and a short-term bond. The interest rate is constant, and the stock price follows a diffusion process:

$$\frac{dS}{S} = \mu dt + \sigma dZ_t,$$

where Z_t is a standard Brownian motion. This is the continuous time analog of a random walk,

$$\frac{\Delta S_t}{S_t} = \mu + \sigma \varepsilon_t,$$

with ε_t iid $N(0, 1)$.

The consumer problem is:

$$\begin{aligned} V(W_0) &= \max_{\{c_t, \alpha_t\}} \int_0^\infty e^{-\rho t} u(c_t) dt \\ \text{s.t.} &: dW_t = W_t(r dt + \alpha(\mu - r) dt - c_t dt) + W_t \alpha \sigma dZ_t. \end{aligned}$$

See Ec744 for how to solve this problem - you have to write the continuous time version of the Bellman equation, which in this case is

$$\rho V(W) = \max_{c, \alpha} \left\{ u(c) + V'(W) (r + \alpha(\mu - r) - c) W + V''(W) W^2 \frac{\alpha^2 \sigma^2}{2} \right\}. \quad (22)$$

Intuitive derivation: split between t and $t + dt$:

$$\begin{aligned} V(W_t) &= \max_{\{c_t, \alpha_t\}} \left\{ \int_0^{dt} e^{-\rho s} u(c_{t+s}) ds + \int_{dt}^\infty e^{-\rho s} u(c_{t+s}) ds \right\}, \\ &= \max_{c, \alpha} u(c_t) dt + e^{-\rho dt} EV(W_{t+dt}), \\ &= \max_{c, \alpha} u(c_t) dt + (1 - \rho dt)(V(W_t) + E(V'(W_t)dW_t) + \frac{1}{2}E(V''(W_t)dW_t^2)), \end{aligned}$$

use Ito's Lemma: $dW^2 = W^2\alpha^2\sigma^2dt$, and note that $E(dW_t) = W_t(rdt + \alpha(\mu - r)dt - c_tdt)$; simplify and you get equation (22)

FOCs w.r.t. c and α respectively,

$$u'(c) = V'(W),$$

and

$$V'(W)(\mu - r)W = -V''(W)W^2\alpha\sigma^2.$$

$$\alpha = \left(\frac{V'(W)}{-V''(W)W} \right) \left(\frac{\mu - r}{\sigma^2} \right).$$

Can guess and verify that $V(W) = \phi \frac{W^{1-\gamma}}{1-\gamma}$. This implies that consumption is a constant share of wealth, and so is savings. The key formula (see Cochrane's op-ed) is the portfolio share:

$$\alpha = \frac{1}{\gamma} \frac{\mu - r}{\sigma^2}$$

This formula gives reasonable outcomes - e.g. if $\mu - r = 7\%$ and $\sigma = 17\%$, then with $\gamma = 4.84$ we get a portfolio share of 50%.²²

Where does the equity premium puzzle show up then? Subtle: in consumption volatility. Here consumption is proportional to wealth, so $\sigma(\Delta \log c) = \sigma(d \log W)$ which is far too high. We implicitly assumed high consumption volatility.

Extensions there are plenty of (interesting) extensions to this basic setup:

- first, return process may be more complicated - the evidence is for some predictability, so there is a state variable x such that

$$dS = \mu(x)Sdt + \sigma(x)SdZ_t,$$

- second, labor income is risky, esp. at individual level (large idiosyncratic shocks),

- third, taxes,

- fourth, other assets, esp. housing,

- fifth, life cycle,

- sixth, borrowing constraints,

- seventh, more general utility functions.

etc.

The simple point is that all this is going to mean there are "state variables" x that shift the value function, so now the FOCs will be

$$u'(c) = V_W(W, x),$$

and

$$\alpha(x) = \left(\frac{V_W(W, x)}{-V_{WW}(W, x)W} \right) \left(\frac{\mu(x) - r(x)}{\sigma(x)^2} \right).$$

(1) Interpretation in terms of hedging - if high x means high MU of wealth, then you want to have more wealth in states with high x - you want to "hedge" the state variable x .

²²To be more precise: we would like to pick γ to match the observed portfolio share. In the data, people hold housing, hold gov't bonds (but have to pay taxes to the gov't too, so perhaps we net this out with ricardian equivalence), etc.

(2) Note that the SDF is $M_{t+1} = \frac{\beta V_W(W_{t+1}, x_{t+1})}{V_W(W_t, x_t)}$ - Merton (1973) makes it clear that you want to “hedge” changes in the state variables x_t . However, at the end of the day $u'(c) = V_W(W, x)$, so this is still a consumption model, which can be tested using the standard Euler equation approach.

(3) These problems can be complicated to solve if several assets, labor income, borrowing constraints, etc. – many state variables etc.

Aggregation with heterogeneity in risk aversion: a simple example

Question: what is the relation between the “aggregate” risk aversion, i.e. the risk aversion that you would estimate from a representative agent, and the “individual” risk aversion?

Consider the following simple one-period example. Consumers have to allocate their wealth between a real asset (“stock”) with exogenous random return R^E , normally distributed with mean μ and variance σ^2 , and a zero-net supply risk-free asset R^F . R^F needs to be determined in equilibrium (In contrast, the mean and variances of R^E are determined by technology – it is a real asset.) There is a large number of consumers, who all have a CARA utility function

$$u(c; \theta) = -\frac{\exp(-\theta c)}{\theta},$$

but they different risk aversion coefficients θ . Each consumer at time 0 decides how to allocate a wealth W between the stock and the bond. The problem is

$$\begin{aligned} \max_{\alpha} E(u(c; \theta)) &= E(u(WR; \theta)), \\ R &= (\alpha R^e + (1 - \alpha)R^f). \end{aligned}$$

with, given the normality:

$$E(u(WR; \theta)) = -\frac{1}{\theta} \exp\left(-\theta W E(R) + \frac{\theta^2 W^2}{2} \text{Var}(R)\right).$$

The first-order condition for a consumer yields the optimal portfolio share α :

$$\alpha^* = \alpha^*(\theta, \sigma^2, \mu, R^f, W) = \frac{\mu - R^f}{\theta \sigma^2 W}.$$

It is inversely proportional to wealth because of CARA (agent wants to invest a constant \$ amount in the risky asset).

Let $f(W, \theta)$ be the joint p.d.f. of W and θ in the population. Then total holdings of stocks are

$$\int \int \alpha^*(\theta, \sigma^2, \mu, R^f, W) W f(W, \theta) d\theta dW,$$

and total holdings of bonds are

$$\int \int (1 - \alpha^*(\theta, \sigma^2, \mu, R^f, W)) W f(W, \theta) dW d\theta.$$

Market-clearing requires that total holdings of bonds are 0, since they are in zero net supply:

$$\int \int (1 - \alpha^*(\gamma, \sigma^2, \mu, R^f)) W f(\gamma, W) d\gamma dW = 0,$$

or alternatively (i.e. equilibrium in the stock market, which is equivalent!)

$$\int \int \alpha^*(\theta, \sigma^2, \mu, R^f, W) W f(W, \theta) dW d\theta = \int \int W f(W, \theta) dW d\theta = E(W)$$

This yields

$$\int \frac{\mu - R^f}{\theta \sigma^2} f(W, \theta) dW d\theta = E(W).$$

Now let's consider three different cases.

Case 1: Everybody has the same θ . Then this equation is just:

$$\frac{\mu - R^f}{\theta \sigma^2} = E(W),$$

i.e. $\alpha^*(\gamma, \sigma^2, \mu, R^f) = 1$, since everybody must be holding stocks in equilibrium. This determines the equilibrium R^f given θ and the properties of the equity return μ, σ^2 as well as the initial wealth $E(W)$.

Case 2: everybody has the same W , but risk aversion is different in the population. Then the equilibrium R^f is determined from the condition

$$\frac{\mu - R^f}{\sigma^2} \int \frac{1}{\theta} f(\theta) d\theta = E(W).$$

Hence, the equilibrium R^f is the same as if there was only ONE agent, with risk aversion θ^{RA} (representative agent) such that

$$\frac{1}{\theta^{RA}} = \int \frac{1}{\theta} f(\theta) d\theta = E\left(\frac{1}{\theta}\right).$$

Since $x \rightarrow \frac{1}{x}$ is convex,

$$E\left(\frac{1}{\theta}\right) \geq \frac{1}{E(\theta)}$$

and

$$\theta^{RA} \leq E(\theta).$$

So the economy is equivalent (in terms of prices) to an economy with a representative agent, but with a risk aversion θ^{RA} , which is lower than the average risk aversion. This is because people sort according to risk aversion; low risk aversion people choose to hold more stocks, hence they matter more for the determination of prices. Of course this makes the equity premium harder to solve (i.e. the equity premium is that risk aversion needs to be higher in the aggregate than it seems to be in micro data, rather than the opposite)

Case 3: θ and W are correlated. It turns out not to matter in this case because of CARA preferences. (i.e. the argument of case 2 goes through exactly as before). The distribution of wealth is irrelevant.

Extension to CRRA: With CRRA, we could get an additional effect, because the shares would be independent of wealth and hence the amounts invested αW would depend on wealth. Using the formula for CRRA,

$$\alpha_{crra} = \alpha_{crra}(\theta, \sigma^2, \mu, R^f) = \frac{\mu - R^f}{\theta \sigma^2}.$$

we obtain the market clearing condition

$$\int \alpha^*(\theta, \sigma^2, \mu, R^f) W f(W, \theta) d\theta dW = \int W f(W, \theta) d\theta dW,$$

leading to

$$\frac{\mu - R^f}{\sigma^2} \int \frac{W}{\theta} f(W, \theta) dW d\theta = \int W f(W, \theta) d\theta dW,$$

which shows that the economy is equivalent to an economy with a representative agent with risk aversion

$$\theta^{RA} = \frac{\int W f(W, \theta) d\theta dW}{\int \frac{W}{\theta} f(W, \theta) dW d\theta}.$$

Note that now we overweight agents with larger wealth since their demand is larger. This formula implies that redistributing wealth from low risk aversion agents to high risk aversion agents would look like an increase in aggregate risk aversion θ^{RA} . Why is this experiment interesting? There are good reasons to suppose that the wealth of less risk-averse agents falls in recessions relative to the wealth of more risk-averse agents. By definition, less risk averse agents hold more risky portfolios, hence a bad shock will decrease their wealth more than more risk averse agents. Hence risk aversion would go up in downturns.²³ This suggests that this mechanism can account for the time-varying risk aversion postulated in Campbell and Cochrane (1999, JPE). Bottom line: the effective aggregate risk aversion depends on the distribution of risk aversion in the population. It will change with the distribution of wealth.

Leverage constraint

Note that in equilibrium, we must have $E(R^e) > R^f$ otherwise nobody would hold the equity asset. But some people may have levered positions, i.e. $\alpha > 1$, implying that they short the risk-free asset to buy equity. This strategy can be implemented in various ways (buying stocks on margin, using options, etc.), but it is typically costly. How do the implications of the model differ if agents cannot take highly levered positions? e.g. what if $\alpha \leq \bar{\alpha}$ is imposed as an additional constraint?

In this case, the agents who would optimally pick $\alpha^* \leq \bar{\alpha}$ are unaffected by the constraint; and the agents with $\alpha^* > \bar{\alpha}$ simply set $\alpha = \bar{\alpha}$. The market equilibrium condition is then,

$$\int_{\frac{\mu - R^f}{\theta \sigma^2 W} \geq \bar{\alpha}} \bar{\alpha} W dF(W, \theta) + \int_{\frac{\mu - R^f}{\theta \sigma^2 W} \leq \bar{\alpha}} \alpha^* W dF(W, \theta) = \int W dF(W, \theta).$$

This implies that we reduce the demand for equity of the less risk-averse agents. Hence, these agents will demand more risk-free assets (or supply less risk-free assets) to the other agents. As a result, the risk-free rate will have to fall. Given that the average equity return is fixed by technology, the expected equity premium rises. Hence, tighter constraints on portfolios are likely to raise the equity premium. Moreover, to the extent that more people become constrained in recessions, the risk premium would be countercyclical. (Why do people become more constrained? Not so clear.) Also, if for some reasons (e.g. breakdown of financial intermediation) the constraint is tightened, you would see some effects on asset prices. As a practical matter, starting in 2007, it became very difficult to take leveraged position in some assets such as MBS. This may have had an effect on their price.

Heterogeneity from realizations of idiosyncratic income shocks (very very rough summary)

There is a large literature on models with heterogeneous agents and incomplete markets (e.g. Heaton and Lucas, JPE 1996; Constantinides and Duffie, JPE 1996; Alvarez and Jermann, Econometrica 2000 and Review of Financial Studies 2001; Lustig 2005; Storesletten, Tellmer and Yaron, RED 2009, Guvenen 2010, Kruger and Lustig 2009). These models typically consider agents who have only access to a risk-free asset and a stock to save or borrow, they have risky labor income, and they may face some constraints on their trading. These constraints can be either, or a mix of, the following: (a) limited participation: some agents are not allowed to hold stocks directly (Guvenen - this is an approximation for the case where there is a participation cost); (b) borrowing constraints or collateral constraints; (c) costs of adjusting the portfolios; (d) limits on leverage. These models have attractive microfoundations but their asset pricing implications are not always very impressive, e.g. generating a large equity premium is already difficult, and matching the variation of risk premia is challenging. Incorporating additional assets or frictions in this model (e.g. housing) is on the research agenda. More precisely, while individual consumption is much more volatile than aggregate consumption, it is also less correlated with the market, and hence it is not clear that the equity premium will be higher. Generally speaking, agents in these models are able to smooth consumption well with only one or two assets, provided that shocks are not too persistent and trading constraints not too tight.

Mankiw (1986) intuition

²³This idea has been formalized, among others, by Kogan and Chang (JPE 2002).

If aggregate shocks are concentrated on a small set of people, the risk premium will be higher. (To be added – read the paper, very simple, JFE 1986, http://www.economics.harvard.edu/files/faculty/40_Equity_Premium.pdf)

Constantinides-Duffie (JPE 1996) (rough summary)

This important paper gives an example of economy where agents can trade, without constraints, a (potentially large) set of assets, and the equity premium is large. Indeed, *any* arbitrage-free prices can be generated (any pricing kernel M_{t+1}) through a particular process of labor income. The key idea is to assume a special type of income process for households, and to show that given this income process, no trade is an equilibrium.

All agents have CRRA utility (β, γ) . It is somewhat simpler to specify directly the consumption process

$$\log \left(\frac{c_{it+1}}{c_{it}} \right) = \eta_{it+1} y_{t+1} - \frac{y_{t+1}^2}{2},$$

with η_{it+1} iid $N(0, 1)$, and

$$y_{t+1} = \sqrt{\frac{2}{\gamma(\gamma+1)}} \sqrt{\log M_{t+1} + \beta + \gamma \log \left(\frac{C_{t+1}}{C_t} \right)}.$$

Because of the unit root, there is no incentive to save if you have a good shock (permanent income intuition). You can check directly that the restriction $E_t(M_{t+1} R_{i,t+1})$ hold for any asset – all consumers are on their first-order conditions. CD actually spell out an income process that generates this consumption process.

The general result of C-D is that any M_{t+1} can be generated as an equilibrium given through a certain process for y_{t+1} . Note that y_{t+1} is the cross-sectional variance of consumption growth. If this variance is constant, the model actually aggregates to a CRRA model, $\log M_{t+1} = -\gamma \log \left(\frac{C_{t+1}}{C_t} \right)$. Hence, it is not idiosyncratic risk *in itself* that leads to a failure of the Euler equation, but the fact that this idiosyncratic risk *changes* over time (Krueger and Lustig, 2009, show the necessity of this condition in a general setup.) An interesting simple case is when

$$y_{t+1}^2 = a + b \log \left(\frac{C_{t+1}}{C_t} \right),$$

and then aggregation holds, but the representative agent risk aversion is $\hat{\gamma} = \gamma - \gamma(\gamma+1)b/2$, so that if $b < 0$ (countercyclical idiosyncratic risk), the aggregate risk aversion is larger than the individual risk aversion. This would justify the large estimated risk aversion (but would not justify the fact that aggregate Euler equations are rejected). Some empirical work suggests that y_{t+1} (XS variance of consumption or income growth) is indeed countercyclical, but perhaps not enough for this mechanism to work.

One important difference with the discussion above (“heterogeneity from...”) is that income shocks are nonstationary (permanent) and the XS distribution widens in recessions. See also Storesletten, Telmer and Yaron (RED 2009) who show how life-cycle effects can also play a role.

Risk-sharing and the aggregate utility function: a case with positivity constraints

This is an alternative example where you can study risk-sharing and aggregation analytically. Assume we have complete markets. One agent is risk-neutral, so he insures the other agent, but he cannot have a negative consumption, which puts a limitation on the insurance. This is a simple way to capture the fact that some agents who are willing to take risks may play a large role in determining asset prices, since they fix the price of risk, but if their wealth is limited there are states where they cannot provide insurance.

There are S states of the world, which occur w/prob π_s , $s = 1 \dots S$. Agent N (for risk-neutral) has an endowment $y_{N,s}$ in state s . His utility function is

$$\sum_{s=1}^S \pi_s v(c_{N,s}),$$

with $v(x) = x$. Note that consumption has to be positive: $c_{N,s} \geq 0$.

Agent A (for risk-averse) has an endowment $y_{A,s}$ in state s . His utility function is

$$\sum_s \pi_s u(c_{A,s}),$$

with $u' > 0$ and $u'' < 0$, and $u'(0) = +\infty$.

Complete markets

The two agents can trade any state-contingent security. Market equilibrium requires that

$$\forall s = 1 \dots S : c_{A,s} + c_{N,s} = y_{A,s} + y_{N,s}.$$

Let q_s be the state-contingent price. Then the problem of agent N is

$$\begin{aligned} & \max \sum_s \pi_s v(c_{N,s}) \\ \text{s.t.} \quad & : \sum_s q_s c_{N,s} \leq \sum_s q_s y_{N,s} : [\lambda_N] \\ & c_{N,s} \geq 0 : [\mu_s] \end{aligned}$$

with FOCs

$$\begin{aligned} \pi_s + \mu_s &= \lambda_N q_s \\ \mu_s c_{N,s} &= 0 \end{aligned}$$

For agent A , we have:

$$\pi_s u'(c_{A,s}) = \lambda_A q_s.$$

The case where the constraint $c_{N,s} = 0$ does not bind.

Then we get perfect risk sharing, i.e. $c_{A,s} = \bar{c}$, and $c_{N,s} = Y_s - \bar{c}$. The prices are

$$q_s = \pi_s u'(\bar{c}) / \lambda_A,$$

i.e. they are simply proportional to the prob of the states. Expected returns are the same on all assets!

The general case

There are two possibilities, for each state - either $\mu_s > 0$, or $\mu_s = 0$. In the first case, $c_{N,s} = 0$ and hence $c_{A,s} = Y_s$. Thus, $q_s = \pi_s u'(Y_s) / \lambda_A$. Note that this requires:

$$\begin{aligned} \mu_s &= \lambda_N q_s - \pi_s > 0 \\ &= \pi_s \left(u'(c_{A,s}) \frac{\lambda_N}{\lambda_A} - 1 \right) > 0. \end{aligned}$$

In the second case, $\mu_s = 0$, then

$$\pi_s = \lambda_N q_s = \frac{\lambda_N}{\lambda_A} \pi_s u'(c_{A,s})$$

and

$$c_{A,s} = \bar{c},$$

where \bar{c} is defined through: $\lambda_N u'(\bar{c}) = \lambda_A$. Then $q_s = \frac{\pi_s}{\lambda_N} = \frac{\pi_s u'(\bar{c})}{\lambda_A}$. Note that state-contingent prices are always given by the risk-averse agent (bc the RN agent is not at an interior solution - the risk averse agent is the “marginal investor” here). Notice also that the state contingent price exhibits a *kink* for $c = \bar{c}$.

Summarizing, if $Y_s < \bar{c}$, then $c_{N,s} = 0$ and $c_{A,s} = Y_s$, and if $Y_s > \bar{c}$, then $c_{N,s} = Y_s - \bar{c}$ and $c_{A,s} = \bar{c}$. So this looks like a “debt contract”: agent A is perfectly insured as long as $Y_s > \bar{c}$, but if $Y_s < \bar{c}$ his payoff is just Y_s . (Plot payoffs and prices q_s as a function of Y_s .)

The level \bar{c} is determined through the budget constraint:

$$\sum_{s=1}^S q_s (c_{N,s} - y_{N,s}) = 0$$

$$\sum_{s:Y_s>\bar{c}} q_s (Y_s - \bar{c} - y_{N,s}) + \sum_{s:Y_s\leq\bar{c}} q_s (0 - y_{N,s}) = 0$$

$$\sum_{s:Y_s>\bar{c}} \frac{\pi_s u'(\bar{c})}{\lambda_A} (Y_s - \bar{c} - y_{N,s}) + \sum_{s:Y_s\leq\bar{c}} \frac{\pi_s u'(Y_s)}{\lambda_A} (0 - y_{N,s}) = 0$$

This is one eqn in \bar{c} . (Note that writing the budget constraint for the other agent of course is equivalent given Walras’ law.)

Aggregate Utility Function

Note that an alternative way to find the equilibrium is to define an artificial representative agent. Since markets are complete, the RA is defined through

$$V(C, \lambda) = \max_{c_A, c_N} \{ \lambda u(c_A) + c_N \}$$

$$s.t. \quad : \quad c_A + c_N = C$$

$$c_N \geq 0$$

for some Pareto weight λ .

Note that the function V has a kink for $C = \bar{c}$. More precisely,

$$V(C) = \lambda u(C) \text{ for } C < \bar{c}$$

$$= \lambda u(\bar{c}) + C - \bar{c} \text{ for } C > \bar{c}$$

and as a result, risk aversion is

$$RRA = \frac{-V''(C)C}{V'(C)} = -\frac{u''(C)C}{u'(C)} \text{ for } C < \bar{c}$$

and

$$RRA = 0 \text{ for } C > \bar{c}.$$

Bottom line: the aggregate utility function’s risk aversion is state-dependent. Low in good states, high in bad states, as in Campbell-Cochrane. There is perfect risk-sharing, only the risk-neutral agents cannot provide insurance against very bad states because their consumption must be positive.

Deleveraging

How is the consumption allocation discussed in the model above implemented in terms of portfolios? (e.g. suppose the model lasted more than one period) Loosely speaking, the risk-neutral agent buys risky assets, and sells risk-free assets to the risk-averse agent, i.e. the risk-neutral agent borrows from the risk-averse agent. When bad shocks happen, the risk-neutral agent has to sell the risky assets to pay

back the loans. Who can he sell the assets to in equilibrium? He has to sell them to the risk-averse agent! But this agent doesn't like the risky asset and he will be willing to hold these assets only if they sell for a very low price i.e. the expected return is high. This is the mechanism through which asset prices fall. The equilibrium market-clearing is critical, but it is not often measured or tested empirically.

Financial Intermediation

Perhaps one of the two agents above can be understood as “hedge fund” or “financial institution”. See Krishnamurthy and Zhe (“Intermediary Asset Pricing”) for such a model. They assume that households must hold securities through financial intermediaries. The wealth of financial intermediaries is then a state variable for the economy, and when it is low, risk premia are higher.

Differences in beliefs and differences in risk aversion

Geanakoplos and coauthors have written on the implications of belief heterogeneity when markets allow for default and the margin (or haircut) varies over time. See also the paper by Alp Simsek.

Guvenen (very rough)

He studies a RBC model with 2 agents, one (“stockholder”) who can invest in stock or bonds, and one (“worker”) who can only buy or sell bonds. He finds that stockholder insures the worker, against poor labor income realizations, and this insurance varies over the business cycle in such a way as to make the equity premium time-varying. Guvenen shows this model can be made consistent with the wealth distribution. The model implies that consumption of stockholders is quite volatile and correlated with asset prices. There is indeed evidence for this (Vissing-Jorgensen, Parker, Malloy, Moskowitz), but perhaps not quite “enough” to make the model work.

12 Production Economies

These notes discuss asset pricing in models with production. So far we have concentrated on endowment economies: the aggregate consumption and dividend are exogenous stochastic processes. In a sense, this is all what we need to determine asset prices, since the SDF depends on consumption only, and the dividends give us the payoffs of the assets we want to price. I call this the “demand” approach to asset prices, as it results from investors demand,

$$V_t = \mathbb{E}_t \sum_{j \geq 1} m_{t,t+j} D_{t+j}.$$

Of course in general equilibrium, consumption and dividends are endogenous, but, in order to get asset prices, we only need to know the equilibrium consumption process that is generated by the economy. We do not need to know the production technologies, financing constraints, shock processes, etc. that generate these equilibrium dividends and consumption processes.

There are still some reasons to be interested in production economies:

- First, a full explanation of the equity premium also requires to understand why the economy generates such a consumption and dividend processes, given that alternatives may be feasible. See section 1 below.

- Second, in reality one can produce assets (through investment) and they are not in fixed supply (trees). This leads to additional relations between quantities and prices, of the form

$$V_t = P_t^K K_t,$$

where P_t^K is the price of investment goods. This is what I call the “supply” approach to asset prices, as it results from the supply of capital (assets). Taking this into account can make it harder to solve the equity premium puzzle. Example of housing or dot-com bubbles: why do prices rise so fast? Standard endowment economy models assume the supply of these assets is fixed, and try to explain a change in demand for these assets, which then leads to a change in the price. But they forget another restriction: if it was easy to build new houses or new dot-coms, the price could not rise. (Usually in economics, we think of prices as being supply-determined, price = cost through competition.)

What I find interesting about the “supply” approach to asset prices is that you do not need to make an assumption regarding the stochastic discount factor, to obtain the value: you are using the fact that the value = PDV of dividends = cost of creating the asset. Of course, this approach requires some other assumption, e.g. on the cost of creating this asset. We cannot price anything: can only price the assets that are effectively produced. But in general equilibrium of course, the price of the asset must equal both the PDV of its dividends, and the cost of creating the asset. (In GE, price can be read from the supply or from the demand!)

- Last, from a macro point of view we are interested in the implications of our models for macro quantities and welfare. The asset pricing puzzles suggest something is terribly wrong with our basic model of preferences (repr. consumer with CRRA utility). Presumably if we fix our model to match asset prices, this will have important implications for quantities, and hence all the standard macro questions: business cycles, welfare cost of taxes and uncertainty, etc.

Adjustment cost model (Cochrane 1991 JF)

One simple model of how to create assets is the Q-theory, or adjustment cost model. Suppose a firm is just some physical capital. The firm can increase its assets by incurring some adjustment costs, $\phi \left(\frac{I_t}{K_t} \right)$. Under constant returns to scale, we have $V_t = P_t^K K_t$ and

$$P_t^K = 1 + \phi' \left(\frac{I_t}{K_t} \right),$$

and hence the stock return is

$$R_{t+1} = \frac{P_{t+1}^K (1 - \delta) + F_K(K_{t+1}, N_{t+1})}{P_t^K}.$$

Lots of testable predictions, which work reasonably well. However, the adjustment cost model also has some limitations: (1) need large curvature of adjustment costs to generate volatile returns, (2) correlation with data return is far from perfect, (3) no clear microfoundation given the observed lumpiness, (4) there are some additional problems if you embed it into a GE model as we will see. Common wisdom in macro is that the Q-theory model is rejected.

Aggregate time-series [Read the paper by Jermann, 1998 JME.]

Suppose we set up a standard RBC model:

$$E \sum_{t=0}^{\infty} \beta^t u(C_t, 1 - N_t),$$

$$K_{t+1} = (1 - \delta)K_t + I_t,$$

$$C_t + I_t \leq F(K_t, z_t N_t),$$

$$\log z_{t+1} = \mu + \rho \log z_t + \sigma \varepsilon_{t+1}, \text{ with } 0 < \rho \leq 1.$$

A shock to TFP ($\varepsilon_{t+1} > 0$) leads C, I, N, Y to rise. Permanent vs. transitory shock (plot IRF). Investment rises more than consumption on impact. This model is a common benchmark, with a risk aversion of 1 or 2. It replicates roughly the volatility of consumption growth. Consumption growth is autocorrelated but not very much (i.e. it is not far from iid). But of course with low risk aversion the risk premia are very small.

So one can be tempted to increase risk aversion so as to generate large risk premia.²⁴ However, this may not work! Consumption becomes smoother as the agent changes investment to avoid adjusting consumption. This is one central difficulty with production economies: the agent is not *forced* to accept a shock to consumption, but can vary investment, hours or other margins instead.

Jermann (1998) studied a model with fixed labor, and introduced habits to help with the equity premium. Habits similarly lead consumption to be very smooth, indeed too smooth compared to the data. (The agent responds to a positive technology shock by adjusting investment more and consumption less.) Because consumption is smooth, risk premia are not that high, despite the habit model.

Jermann's solution is to add adjustment costs. The capital accumulation is

$$K_{t+1} = (1 - \delta)K_t + \phi \left(\frac{I_t}{K_t} \right) K_t,$$

with $\phi' > 0$ and $\phi'' < 0$, so increasing the capital stock quickly is costly. The benchmark case is $\phi(x) = x$ linear. Jermann (1998) finds that adding these adjustment costs, consumption and investment now have both about the right volatility. Moreover, the equity premium is high, and the equity return is volatile. Hence, the conclusion: to get a large equity premium, consumers must care about the risk (high risk aversion or habits), and they must be *prevented* from doing something about it (fixed labor and large adjustment costs to the capital stock).

Limitations of Jermann's paper: the model has fixed labor, which is a serious limitation to study business cycles. The model has a very volatile risk-free rate, much more than in the data, because of the habits. (They do not have the Campbell-Cochrane form.) In terms of business cycle dynamics, the model

²⁴Because IES = 1 / risk aversion, we get a very low IES and hence very slow responses to shocks.

is similar to a RBC model with fixed labor. Shutting down the hours margin and adding the investment adjustment costs leads to less amplification, which is undesirable from a business cycle point of view.

Profit and dividend volatility

So far we have talked about how the consumption process generated from the model may change as you change risk aversion or add habits. What about the dividend process? It turns out to be at odds with the data. Defining dividends as profits net of investment, we have

$$D_t = Y_t - w_t N_t - I_t = (1 - \alpha)Y_t - I_t,$$

where α = capital share if production function is Cobb-Douglas. Two issues: (a) profits are not volatile enough: here profits are proportional to output, so they have the same volatility, but in the data profits are about 3 times more volatile than output. (There is surprisingly little research on this topic.) (b) Because in good times investment goes up, dividends are typically countercyclical in this model. Incorporating debt/equity choices and leverage can generate some additional volatility in dividends.

Return on capital volatility

Absent adjustment costs, the return on holding physical capital in the RBC model is

$$R_{t+1}^K = 1 - \delta + F_1(K_{t+1}, z_{t+1}N_{t+1}).$$

We have the FOC:

$$E_t \left(\frac{\beta u_1(C_{t+1}, 1 - N_{t+1})}{u_1(C_t, 1 - N_t)} R_{t+1}^K \right) = 1.$$

This return on capital is very smooth, since $1 - \delta$ is constant, and it is the largest share of the return; and $F_1(K_{t+1}, z_{t+1}N_{t+1})$ does not move much. Economic interpretation: you can produce (or destroy) assets at price one, at the margin, so the value of capital is always equal to one, except for the current dividend. Tobin $q = \frac{\text{market value}}{\text{production cost}} = 1$ which implies no volatility of asset prices.

Adjustment costs make it costly to build capital quickly, and hence market value of installed capital \neq production cost today. With CRS, average $Q = \text{marginal } Q = \frac{1}{\phi'(\frac{I_t}{K_t})}$, so a high investment rate means a high price of capital. The value of capital is then

$$V_t = q_t K_t,$$

and the return is

$$R_{t+1} = \frac{q_{t+1}(1 - \delta) + F_1(K_{t+1}, z_{t+1}N_{t+1})}{q_t},$$

so the return can be volatile because of “capital gains” i.e. change in the price, as in the data, if the investment moves around.

In practice these models tend to require large adjustment costs. These may have undesirable effects on business cycle dynamics. In reality I suspect that it's not just physical adjustment costs but also intangible capital (getting new capital is easy but getting employees, customers, and so on is hard, but this is rarely modeled or estimated; see Merz and Yashiv AER for an example with labor adjustment costs).

Intangible capital

More generally, intangible capital is useful to understand the level of q , e.g. the fact that $q > 1$ most of the time (the fact that $q < 1$ sometimes is even more of a puzzle). For instance, if a firm is formed by adding some “idea” to physical capital, then $q > 1$ since the value of the firm = value of idea + value of physical capital and hence value of firm / value of physical capital > 1 . Another important implication of intangible capital is that it is related to long-run growth. For instance, Schmid and Huang consider a Romer (1986) style model where firms invest in r&d, and consider the implications for firm values.

Similarly, Panageas, Garleanu and Yang consider a model in the Romer (1990) tradition. One interesting question is, what are the risk characteristics of intangible capital? May depend on complementarity of intangible and tangible capital, the costs of accumulating/decumulating it, etc.

Tallarini: does risk aversion affect business cycle dynamics?

Tallarini (2000 JME) used a RBC model with variable labor, and with Epstein-Zin utility, to study the effect of risk aversion on business cycle dynamics. He found that it was very small. A higher risk aversion changes the “steady-state” (i.e. the average level of capital: there is more capital since agents want to save more) but does not change the standard business cycle moments, such as $\sigma(\Delta \log C)$, $\sigma(\Delta \log Y)$, etc.

The intuition is that, in this model, the representative agent faces some constant risk. Increasing risk aversion changes his average behavior, but does not lead to different dynamics since risk is nearly constant. The IES matters, since it determines how fast you respond to a shock.

You can check my paper on “Disaster risk and business cycles”, for a production economy model which has an analytical result similar to Tallarini, and some more results when risk varies over time.

Cross-section [Read the paper by Lu Zhang, 2005 Journal of Finance.]

As we discussed in the class, there is heterogeneity of expected returns - in particular, firms with low Book/Market ratios have low average returns (and they do not seem to have a higher market or consumption beta).

People sometimes study these questions given exogenously specified (and estimated) time series processes for cash flows of these firms. But of course these cash flows must be generated, so there are some natural economic questions:

- why are these firms different? (i.e. sources of heterogeneity)
- why do they generate these different cash flows? (esp. given that the market gives very different values to these cash flow processes, so there are incentives to produce “safe” cash flows?)
- how do their investment, sales, etc. differ? does this matter somehow?

A natural source of heterogeneity is differences in productivity. Several paper consider the decision problem of a firm which maximizes the PDV of dividends, and picks investment, e.g.

$$V(K, z, x, s) = \max_{K'} \{ D(K, z, x, K') + E_{s'|s} M(s, s') V(K', z', x', s') \},$$

where

- x = idiosyncratic productivity, follows an AR(1)
- z = aggregate productivity, follows an AR(1)
- s = aggregate states, which determine the SDF $M(s, s')$

Dividend is

$$D(K, z, x, K') = \pi(K, z, x) - I = zxK^\alpha - f - I,$$

where f is a fixed cost, and

$$K' = (1 - \delta)K + \phi(I, K),$$

where ϕ is an adj. cost function, which may include irreversibility or asymmetric adj. costs.

Sometimes the model is GE, so $M(s, s')$ equals, say, $\beta C(s')^{-\gamma} / C(s)^{-\gamma}$, but often they just specify an exogenous process for the SDF which is empirically reasonable, and work out firm values and firms' decisions of investment from the above Bellman equation.

How can a model be consistent with the facts above?

(a) First, what are the firms with low B/M i.e. high Q ? Well, V is usually increasing in x and K , but $V(K, z, x, s)/K$ is decreasing in K if there are decreasing return to scale. A high x means you will be profitable for a while (since x is persistent), so your value goes up. Hence, firms with high x relative to their current K will have high Q .

(b) Why would these firms have low expected returns? Intuition in Zhang paper: asymmetric adj costs. A firm with a high x is investing, and making profits, so if there is a bad aggregate shock (low z) it can reduce investment easily and pay dividends. In contrast, a firm with a low x is currently not investing, if z falls it will have to disinvest, which is costly, and it will have very low (even negative) dividends, so it has low payouts in bad states. Hence the low x firms (which are the low Q firms) are more risky because they have less flexibility to smooth dividends in the face of aggregate shocks. See the Zhang paper for much more, some nice quantitative experiments (turning on/off the physical adj. costs and the SDF) and some interesting predictions for the pattern of investment, Q , and returns. Some of these implications have not yet been tested, however. There are extensions that allow for financial constraints, labor adj costs, etc., but this is still a fairly recent literature.