FIN 395: Asset Pricing Theory

III. Arbitrage and the Stochastic Discount Factor

Daniel Neuhann

daniel.neuhann@mccombs.utexas.edu

Spring 2026

Introduction

- We explore the classic notion of no arbitrage and its implications.
- This is the most powerful preference-free concept in asset pricing.
- This will lead us to characterize objects generally useful in asset pricing
  - state prices, stochastic discount factor (SDF), martingale equivalent measures,...

- We can then relate these objects to utility maximization and Euler equations.
- Ask what can be recovered from asset prices and ways to specify SDF.
- We discuss the sharp contrast between complete and incomplete markets.

- **No Arbitrage and the Stochastic Discount Factor**
- Bounds on SDFs as a Diagnostic Tool
- Applications of SDF and Risk-neutral Measures

Basic Framework: Static Version

- Two dates $t = \{0, 1\}$.
- S states of the world at $t = 1$.
- Complete, symmetric information among market participants.
- Competitive markets: agents take prices as given.
- N securities, security i is a payoff vector $d_i = \begin{bmatrix} d_i^1 & d_i^2 & \cdots & d_i^S \end{bmatrix}$.
- Payoff matrix D with $D_{is} = d_i^s$, $i \in \{1, ..., N\}$ and $s \in \{1, ..., S\}$.
- Portfolio: vector $\omega \in \mathbb{R}^N$.

Asset Span and Market Completeness

- Payoff of portfolio in state s is $\sum_i \omega_i D_{is}$ and payoff vector $= D'\omega$.
- The **asset span** associated with the payoff matrix is defined to be

$$\mathcal{M} = \left\{ \mathbf{z} \in \mathbb{R}^S : \mathbf{z} = D'\omega \text{ for some } \omega \in \mathbb{R}^N \right\}$$

  - $\mathcal{M}$ is a linear subspace of $\mathbb{R}^S$.
  - Complete markets implies $\mathcal{M} = \mathbb{R}^S$.
  - Markets are complete if and only if $rank\,(D) = S$

- Canonical form of complete markets is set of **Arrow-Debreu** securities.
  - S securities, security s satisfies $d_s^s = 1$ and 0 otherwise.

- Security i is **redundant** if $d_i = D'\omega$ with $\omega_i = 0$.
- That is, asset is redundant if payoff can be replicated with other assets.

No Arbitrage

- An **arbitrage** is the possibility of positive payoffs at some date and state of the world, with no possibility of a negative cash flow at any date or state.

- We are interested in price systems that satisfy **no arbitrage**.
- That is, price systems such that no arbitrages exist.

- A weak requirement on asset prices.
- Also a **necessary condition for equilibrium** in financial markets.
- Preference-free in that it relies only on "more $\succ$ less."

Arbitrage Portfolios and State Prices

- Let the vector of asset prices be $\mathbf{q} \in \mathbb{R}^N$.
- The cost of a portfolio ω is $\mathbf{q}' \cdot \omega = \sum_i q_i \omega_i$.
- An **arbitrage portfolio** $\omega \in \mathbb{R}^N$ is a portfolio such that

$$\mathbf{q}' \cdot \omega \leq 0 \text{ and } D' \cdot \omega > \mathbf{0},$$

or a portfolio such that

$$\mathbf{q}' \cdot \omega < 0 \text{ and } D' \cdot \omega \geq \mathbf{0},$$

State Prices

Definition (State Prices)
Given a vector of asset prices $\mathbf{q}$, a vector of **state prices** ψ is such that

$$\mathbf{q} = D \cdot \psi.$$

- State prices can be interpreted as the marginal cost of a state-contingent payout.
- The key restriction is that it is a *linear* object.

Theorem (Fundamental Theorem of Finance)
Let D be a $N \times S$ matrix, and $\mathbf{q} \in \mathbb{R}^N$. There is no ω in $\mathbb{R}^N$ satisfying

$$\mathbf{q}' \cdot \omega \leq 0 \text{ and } D' \cdot \omega \geq \mathbf{0},$$

with at least one strict inequality **iff** *there exists $\psi >> \mathbf{0} \in \mathbb{R}^S$ s.t. $\mathbf{q} = D \cdot \psi$.*

In words: arbitrage holds **iff** there exists a strictly positive vector of state prices.

Proof. (IF).
Suppose there exists $\psi >> \mathbf{0} \in \mathbb{R}^S$ such that $\mathbf{q} = D \cdot \psi$, and that there exists ω satisfying $\mathbf{q}' \cdot \omega \leq 0$ and $D' \cdot \omega \geq \mathbf{0}$. Then

$$0 \leq \psi' D' \cdot \omega = \mathbf{q}' \cdot \omega \leq 0.$$

This implies $\mathbf{q}' \cdot \omega = 0$. Since $\psi >> \mathbf{0} \in \mathbb{R}^S$, then $D' \cdot \omega = 0$. □

Two Useful Mathematical Results (Referenced from Duffie)

Theorem (Separating Hyperplane Theorem)
Suppose that A and B are convex disjoint subjects of $\mathbb{R}^N$. There is some nonzero linear functional F such that $F(x) \leq F(y)$ for each x in A and y in B. Moreover, if x is in the interior of A or y is in the interior of B, then $F(x) < F(y)$. Furthermore, if A is closed and B is compact, then F can be chosen such that $F(x) < F(y)$ for all x in A and y in B.

Theorem (Linear Separation of Cones)
Suppose M and k are closed convex cones in R^n that intersect precisely at zero. If K does not contain a linear subspace other than $\{0\}$, then there is a non-zero linear functional F such that $F(x) < F(y)$ for each in x in M and each nonzero y in K.

Proof. (ONLY IF first part).

- Let $Q = \left\{(-\mathbf{q}' \cdot \omega,\ D' \cdot \omega) \mid \omega \in \mathbb{R}^N\right\}$, and $K = \mathbb{R}_+ \times \mathbb{R}_+^S$.
- For no arbitrage to hold, we must have $Q \cap K = \{0\}$. (Else there exists at least one portfolio such that $\mathbf{q}' \cdot \omega \leq 0$, $D' \cdot \omega \geq 0$ and one equality strict).
- This implies there exists an open cone $K' \supset K - \{0\}$ with $Q \cap K' = \varnothing$.

□

Proof. (ONLY IF second part).

- By the Separating Hyperplane Theorem, there exists $\phi \neq \mathbf{0}$ with $\phi' \cdot \mathbf{z} \leq \phi' \cdot \mathbf{x}$ for each $\mathbf{z} \in Q$ and $\mathbf{x} \in K$.
- Since Q is linear and K is a cone, $\phi' \cdot \mathbf{z} = 0$ for each $\mathbf{z} \in Q$, and $\phi' \cdot \mathbf{x} > 0$ for each $\mathbf{x} \neq \mathbf{0}$ and $\mathbf{x} \in K$ with $\phi' >> \mathbf{0}$.
- Let us express $\phi = \left(\alpha, \tilde{\phi}\right)$ with $\alpha \in \mathbb{R}$, and $\psi = \frac{\tilde{\phi}}{\alpha}$, and $\mathbf{z} = (-\mathbf{q}' \cdot \omega, D' \cdot \omega)$ Then

$$\phi' \cdot \mathbf{z} = 0 \Longrightarrow \tilde{\phi}' \cdot \left(D' \cdot \omega\right) - \alpha \mathbf{q}' \cdot \omega = (D \cdot \psi)' \cdot \omega - \mathbf{q}' \cdot \omega = 0.$$

Since this must hold for each $\omega \in \mathbb{R}^N$, it follows

$$\mathbf{q} = D \cdot \psi.$$

□

Definition and Use of “No Arbitrage” Terminology

- We have considered two forms of pure arbitrage:
    1. a zero-cost portfolio that can generate positive profit with some probability.
    2. a positive payoff today that generates no losses tomorrow.

- This is the appropriate formal view of no-arbitrage.
- It is not necessarily the same as the colloquial use.
- For example, there is also “risk arbitrage”:
    - There is “mispricing” but arbitrageurs face risk in exploiting it for profit.
    - We will return to this issue when we discuss "limits of arbitrage"

Law of One Price (LOOP)

Definition (Law of one Price)
Any two assets with the same cash flows must have the same price.

- LOOP is implied by no arbitrage. This follows from existence of state prices.
- Under no arbitrage, a portfolio ω with payoff schedule $D' \cdot \omega$ costs

$$\omega' \cdot \mathbb{E}^p [D \cdot \pi] = \omega' \cdot \mathbf{q},$$

- The weighted sum of prices of underlying securities that replicate payoffs.
- Frequent application: A risk-free portfolio must earn the risk-free rate.
- This observation be used to price derivatives.

Some Useful Properties of State Prices

- State prices ψ, allow us to price any asset i given payoffs D_i with

$$\mathbf{q} = D_i \cdot \psi = \sum_s \psi_s D_{is}$$

- A one-period riskfree bond with a payoff of 1 in every state has price

$$q_0 = \mathbf{1}'_{S\times 1} \cdot \psi = \sum_s \psi_s$$

- The price of this bond is inverse of the riskfree rate, $q_0 = \frac{1}{R^f}$. Hence

$$R^f = \left(\mathbf{1}'_{S\times 1} \cdot \psi\right)^{-1} = \left(\sum_s \psi_s\right)^{-1}.$$

Risk-neutral Measure

- We can always rewrite the state price vector as

$$\psi = \left(\mathbf{1}'_{S\times 1}\cdot\psi\right)\cdot\frac{\psi}{\mathbf{1}'_{S\times 1}\cdot\psi} = \frac{1}{R^f}\cdot\mu,$$

- Observe that

$$\mu = \frac{\psi}{\mathbf{1}'_{S\times 1}\cdot\psi}$$

  is a **probability measure** because $1 \geq \mu_i > 0$ for all i and $\sum_i \mu_i = 1$.

- That is, we have an alternative probability distribution over states.
  - We have constructed it by discounting by the risk-free rate.
  - So it is called the **risk-neutral measure**.

Risk-neutral Measure

- Given our construction, we can always write prices as

$$\mathbf{q} = D \cdot \psi = \frac{1}{R^f} D \cdot \mu.$$

- Since μ is a measure, this is just an **expectation under** μ.
- That is, for some vector $\mathbf{0} \leq \nu \in \mathbb{R}^S$ with $\mathbf{1}'_{S\times 1} \cdot \nu = 1$, we can always write

$$\begin{aligned} \mathbb{E}^{\nu}[\mathbf{x}] &= \mathbf{x}' \cdot \nu = \sum_s x_s v_s \\ \mathbb{E}^{\nu}[D] &= D \cdot \nu, \end{aligned}$$

  for vector $\mathbf{x}$ and $N \times S$ matrix D.

- Hence we can express prices as

$$\mathbf{q} = \frac{1}{R^f} \mathbb{E}^{\mu}[D].$$

Risk-neutral Measure

- Let $\mathbf{p}$ denote the "physical" (true) probability measure.
- A risk-neutral agent would price assets according to

$$\mathbf{q} = \frac{1}{R^f} \mathbb{E}^p [D] .$$

- We have just shown that a risk-averse agent prices assets according to

$$\mathbf{q} = \frac{1}{R^f} \mathbb{E}^\mu [D] .$$

- μ is the *risk-neutral measure* because risk-averse agents price assets **as if** they were risk-neutral and the probability distribution is μ.

- Risk attitudes are reflected in the tilting of probability.
- Valuable states receive higher probability-weighting under μ than $\mathbf{p}$.

Martingale Equivalent Measure (MEM)

- Risk-neutral measure is special case of a martingale equivalent measure.
  - Equivalence: $\mu'(s) > 0$ if and only if $p(s) > 0$.
  - Martingale: price process is a martingale under μ'.
- Each MEM is defined with respect to a "numeraire" asset.
- The risk-neutral measure's numeraire is a short-term risk-free asset.

- Another common numeraire is a T-period bond.
- This leads to a T-forward measure.

Definition (Utility Maximization Problem)
An agent has endowment $e \geq 0$ and strictly increasing utility $U : \mathbb{R}^s_+ \to R$. Dividends and prices of N securities are given by the pair (D, q). Given portfolio θ, the budget set is

$$B(q, e) = \{e + D'\theta \in \mathbb{R}^s_+ : \theta \in \mathbb{R}^N, q \cdot \theta \leq 0\}.$$

The utility maximization problem is:

$$\max_{c \in B(q,e)} \quad U(c) \tag{UMP}$$

State Prices and Utility Maximization

Assume there exists a portfolio θ^0 with $D'\theta^0 > 0$. This implies that the wealth constraint on portfolio choices is binding, $q\theta^* = 0$.

Theorem
If there is a solution to UMP, there is no arbitrage. If U is continuous and there is no arbitrage, there is a solution to UMP.

Theorem
Suppose U is strictly concave and differentiable at some $c^* = e + D'\theta^* >> 0$ *with* $q \cdot \theta^* = 0$*. Then* c^* *is a solution to UMP* **if and only if** $\lambda \partial U(c^*)$ *is a state-price vector for a scalar* $\lambda > 0$*.*

Stochastic Discount Factor (SDF)

- Another way of rewriting our pricing equation is

$$\mathbf{q} = D \cdot \psi = \left(D \cdot \frac{\psi}{\mathbf{p}} \right) \cdot \mathbf{p} = (\pi \cdot D) \cdot \mathbf{p}$$
$$= \mathbb{E}^p \left[\pi \cdot D \right]$$

  where $\frac{\psi}{\mathbf{p}}$ is interpreted as $\frac{\psi_i}{p_i}$ element-by-element

- We call π a **stochastic discount factor** (or state-price density.)
  - It is a random variable that discounts all payoffs.
  - Another common notation is M, as in $\mathbf{q} = \mathbb{E}^p[M \cdot D]$.
- An SDF prices *all assets*. It all comes down to finding the "right" one.

Stochastic Discount Factor (SDF)

- Recall that asset i's return is just $R_i^s = \frac{D_{is}}{q_i}$. Hence we have

$$1 = \mathbb{E}^p \left[\pi \cdot \mathbf{R}_i\right].$$

- For any portfolio ω with $\mathbf{q} \cdot \omega \neq 0$, if $R_\omega^s = \frac{\left(D' \cdot \omega\right)_s}{\mathbf{q} \cdot \omega}$, then

$$\mathbb{E}^p \left[\pi \cdot \mathbf{R}_\omega\right] = 1.$$

- For any risk-free portfolio with return R_f, we have $1 = \mathbb{E}^p \left[\pi\right] R_f$. Hence

$$R_f = \frac{1}{\mathbb{E}^p \left[\pi\right]}.$$

Utility Maximization and the SDF: An Example

- Two dates, one of S states realized at date 2. Probability of state s is $p(s)$.
- A full set of Arrow-Debreu securities. Claim on state s has price $q(s)$.
- Initial wealth is W_0, and $C(s)$ denotes purchases of the s claim.

$$\max_{\{C(s)\}_{s=1}^{S}} \quad u(C_0) + \sum_{s=1}^{S} \beta p(s) u(C(s))$$

$$\text{s.t} \quad C_0 + \sum_{s=1}^{S} q(s) C(s) = W_0.$$

- First-order condition for any $C(s)$ is

$$\beta p(s) u'(C(s)) = q(s) u'(C_0).$$

Utility Maximization and the SDF: An Example

- D is a diagonal matrix of ones. Hence A-D security prices **are** state prices.
- Rearranging the first-order conditions gives

$$q(s) = \frac{\beta p(s) u'(C(s))}{u'(C_0)}.$$

- Hence the stochastic discount factor is the ratio of marginal utilities:

$$\pi(s) = \frac{q(s)}{p(s)} = \frac{\beta u'(C(s))}{u'(C_0)}.$$

- With CRRA utility, it is a function of consumption growth:

$$\pi(s) = \frac{q(s)}{p(s)} = \beta \left(\frac{C(s)}{C_0} \right)^{-\gamma}$$

- This is the starting point of consumption-based asset pricing.

Uniqueness of State Prices

- If markets are complete, there exists a unique state price vector $\psi >> 0$.
  - Hence there is a unique SDF and unique MEM.
- If markets are incomplete, there are infinitely many state prices vectors.

Simple argument:

1. Suppose that m^* is an SDF. Then $\mathbf{q} = \mathbb{E}[m^* \cdot D]$.
2. Now pick some $\epsilon \in \mathbb{R}^S$ orthogonal to D, i.e. $\mathbb{E}[\epsilon D] = 0$.
3. Then $m' = m^* + \epsilon$ is also an SDF because

$$\mathbf{q} = \mathbb{E}[(m^* + \epsilon) \cdot D] = \mathbb{E}[m^* \cdot D] + \mathbb{E}[\epsilon \cdot D] = \mathbb{E}[m^* \cdot D].$$

4. Fails in complete markets because there does not exist an orthogonal ϵ.

Some Additional Implications of this Construction

- There *does* exist a unique SDF in the asset span (or payoff space) $\mathcal{M}$.
- This is because we can always project any SDF onto $\mathcal{M}$, i.e.

$$m = m^* + \epsilon \qquad \text{where} \qquad m = proj(m|\mathcal{M}) \quad \text{and} \quad \mathbb{E}[\epsilon D] = 0.$$

- Easy to verify that m^* is a valid SDF, with identical pricing implications.
    - The projected SDF is sometimes called the "mimicking portfolio" for m.
    - It best summarizes the pricing information in *all* SDFs.

- All differences in SDFs are about **non-marketed payoffs**.
- Another way of saying this: non-traded payoffs may have different prices.

Some Basic Implications

- SDF is a random variable. Use $cov(\pi, D) = \mathbb{E}[\pi \cdot D] - \mathbb{E}(\pi)\mathbb{E}(D)$ to give

$$\begin{aligned} \mathbf{q} &= \mathbb{E}[D]\mathbb{E}[\pi] + cov(\pi, D) \\ &= \frac{\mathbb{E}[D]}{R_f} + cov(\pi, D) \end{aligned}$$

- Prices = risk-neutral discounted value + *risk adjustment*.
- Using the utility maximization example, we have

$$\mathbf{q} = \frac{\mathbb{E}[D]}{R_f} + \frac{cov(\beta u'(C(s), D(s))}{u'(C_0)}.$$

- *Only assets that co-vary with marginal utility get risk adjustment.*

Some Basic Implications

- We can rewrite the basic equation to reflect returns. For a particular asset,

$$1 = \mathbb{E}[\pi \cdot R_i]$$

- Hence the decomposition is $1 = \mathbb{E}[\pi]\mathbb{E}[R_i] + cov(\pi, R_i)$. Since $R_f = \mathbb{E}[\pi]^{-1}$,

$$\begin{aligned}\mathbb{E}[R_i] - R_f &= -R_f cov(\pi, R_i) \\ &= -\frac{cov(u'(c(s)), R_i)}{\mathbb{E}[u'(c(s))}.\end{aligned}$$

- This is just saying that excess returns depend only on covariance with π.
- Another way of saying that idionsyncratic risk has no risk premium.

Beta Representations

- The expected return equation can be written as

$$\mathbb{E}[R_i] = R_f + \left(\frac{cov(R_i, \pi)}{var(\pi)}\right)\left(-\frac{var(\pi)}{\mathbb{E}[\pi]}\right)$$

  or equivalently,

$$\mathbb{E}[R_i] = R_f + \beta_{i,\pi}\lambda_\pi.$$

- The β is the **regression coefficient** of R_i on π.
- The λ is akin to a **price of risk** as determined by the SDF.
- This provides a simple link to empirical analysis.

Multi-factor Structure of SDF

- Assume SDF is a linear combination of K common factors $f_{k,t+1}$, $k = 1, ..., K$, that are mean-zero and are orthogonal to each other

$$\pi_{t+1} = a_t - \sum_{k=1}^{K} b_{kt} f_{k,t+1}.$$

- Then it follows that

$$\begin{aligned} -Cov_t^p\left[\pi_{t+1}, R_{i,t+1} - R_t^f\right] &= \sum_{k=1}^{K} b_{kt}\sigma_{ikt} = \sum_{k=1}^{K}\left(b_{kt}\sigma_{kt}^2\right)\cdot\left(\frac{\sigma_{ikt}}{\sigma_{kt}^2}\right) \\ &= \sum_{k=1}^{K}\lambda_{kt}\beta_{ikt}. \end{aligned}$$

  - sometimes referred to as a "beta-lambda" decomposition in empirical tests

- This form of SDF requires strong assumptions: e.g. static, or M-V utility.

No Arbitrage in Multi-period Models

- We now move to a dynamic setting with $T+1$ dates, $t = 0, 1, \dots, T$.
- The key difference is that payoffs are now partly endogenous.

- Information is represented by a filtration $\mathcal{F}$ (as defined in Introduction.)
- An adapted process is a sequence $X = \{X_0, \dots, X_T\}$ such that X_t is a random variable with respect $\{\Omega, \mathcal{F}_t\}$ for each t.

- A *security* (or asset) is a claim to an adapted dividend process δ, where δ_t denotes the dividend paid at time t. Each security is associated with an adapted *price process* S, where S_t denotes the ex-dividend price at time t.

No Arbitrage in Multi-period Models

- Assume there are n securities defined by the $\mathbf{R}^N$-valued dividend process $\delta = (\delta^{(1)}, \ldots, \delta^{(N)})$. The associated price process is $S = (S^{(1)}, \ldots, S^{(N)})$.

- A *trading strategy* is an adapted process θ in $\mathbf{R}^N$, where $\theta_t = \{\theta_t^{(1)}, \ldots, \theta_t^{(N)}\}$ represents the portfolio (absolute, not in shares) held at time t.
- The dividend process δ^θ generated by a trading strategy is defined by

$$\delta_t^\theta = \underbrace{\theta_{t-1}(S_t + \delta_t)}_{\text{Payoff from previous period}} - \underbrace{\theta_t \cdot S_t}_{\text{Current Expenditures}} .$$

- Trading strategy θ is an **arbitrage** if $\delta^\theta > 0$. ("Something for nothing").

No Arbitrage in Multi-period Models

- We will use Θ to denote the space of trading strategies.
- For any $\theta, \psi \in \Theta$ and scalars a, b, **assume** $a\delta^\theta + b\delta^\psi = \delta^{a\theta + b\psi}$.
- The marketed subspace (or asset span) $\mathcal{M} = \{\delta^\theta : \theta \in \Theta\}$ is a linear space.

- A **deflator** is any strictly postive adapted process.
- Definition: Deflator π is a **state price deflator** if, for all t,

$$S_t = \frac{1}{\pi_t} E_t \left(\sum_{j=t+1}^{T} \pi_j \delta_j \right)$$

- Observation: π is a state-price deflator if and only if

$$\theta_t \cdot S_t = \frac{1}{\pi_t} E_t \left(\sum_{j=t+1}^{T} \pi_j \delta_j^\theta \right) \quad \text{for all trading strategies } \theta$$

- This is the natural extension of a state-price deflator from the static case.

No Arbitrage in Multi-period Models

Theorem (Fundamental Theorem of Finance: Multi-period Version)
The dividend-price pair (δ, S) *admits no arbitrage* **if and only if** *there is a state-price deflator.*

- Proof is similar to the static version using the Riesz Representation Theorem.
- Can be extended to infinite states and time. See Duffie Chapter 2 and 5.

Deflated Gains Process

- The gains process for (δ, S) is defined as

$$G_t = S_t + \sum_{j=1}^{t} \delta_j,$$

which is the history of dividends and the current price

- For any deflator $\gamma >> 0$, the **deflated gains process** is

$$G_t^{\gamma} = \gamma_t S_t + \sum_{j=1}^{t} \gamma_t \delta_t,$$

which is the gains process discounted by γ. (Similar to change in numeraire.)

Theorem
π is a state-price deflator if and only if $S_T = 0$ and G^{π} is a martingale.

- As in the static case, we can define equivalent measures to gain insight.
- This requires a numeraire. Assume that there is risk-free short-term asset. If you invest 1 in the risk-free asset at date t, you receive $1 + r_t$ at date $t + 1$ (the *short rate*).
- For any t and $\tau \leq T$, define

$$R_{t,\tau} = (1 + r_t)(1 + r_{t+1}) \cdots (1 + r_{\tau-1}).$$

Definition
An alternative probability measure Q is a martingale equivalent measure if

$$S_t = \mathbb{E}_t^Q \left(\sum_{j=t+1}^{T} \frac{\delta_j}{R_{t,j}} \right).$$

- This is risk-neutral discounting of dividends under Q.

Martingale Equivalent Measure

- We have similar results as before:
  1. There is no arbitrage iff there is a martingale equivalent measure.
  2. Each state-price deflator is associated with a unique MEM.
  3. There is a unique MEM if and only if markets are complete.
- See Duffie Chapter 2 for the full derivation.

Relationship to Utility Maximization

- Consider the canonical dynamic utility maximization framework.
- Additively time-separable with strictly concave flow utility u.
- Then no arbitrage implies that the following stochastic Euler equation holds:

$$S_t = \frac{1}{u_t'(c_t^*)} \mathbb{E}_t \left[S_\tau u_\tau'(c_\tau^*) + \sum_{[j=t+1]}^{\tau} \delta_j u_j'(c_j^*) \right] \quad \text{for all } t \text{ and } \tau \geq t.$$

Concrete Implications in a Canonical Framework

- We now study the implications of the general model in a canonical set-up.
- Specifically, an infinite-horizon endowment economy with risky assets.

- An agent has wealth A_t at time t and wants to maximize lifetime utility

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j}), \quad 0 < \beta < 1.$$

(i) $\mathbb{E}_t$ is the expectation given information known at t

(ii) β is the subjective discount factor

(iii) c_{t+j} is the agent's consumption in period $t + j$

(iv) u is concave, strictly increasing, twice continously differentiable

Concrete Implications in a Canonical Framework

- There are two assets: a one-period bond and risky equity.

- Bonds earn the risk-free rate R_t. We call L_t the gross payout on the agent's bond holdings between periods t and $t+1$, with present value $L_t R_t^{-1}$.
  - $L_t < 0$ indicates borrowing. We impose the borrowing constraint $L_t \geq -b_b$.

- An equity position s_t entitles the agent to a risky share of dividends y_t. The ex-dividend price of equity at date t is p_t.
  - $s_t < 0$ indicates a short position, and we impose the constraint $s_t \geq -b_s$.
  - The equity payout y_t is the only source of risk. We can make various assumptions about the driving stochastic process.

Intertemporal Euler Equations

- The agent's budget constraint in period t is

$$c_t + R_t^{-1} L_t + p_t s_t \leq A_t.$$

- The evolution of wealth satisfies

$$A_{t+1} = L_t + (p_{t+1} + y_{t+1} s_t).$$

- This is a dynamic programming problem.
  - The state variables are A_t and the history of y.
  - The controls are L_t and s_t.

Intertemporal Euler Equations

- Suppose that the borrowing constraints do not bind.
- First-order conditions associated with L_t and s_t are

$$u'(c_t)R_t^{-1} = \mathbb{E}_t \beta u'(c_{t+1}).$$

$$u'(c_t)p_t = \mathbb{E}_t \beta (y_{t+1} + p_{t+1}) u'(c_{t+1}).$$

- These are so-called **Euler equations** pinning down intertemporal optimality.
- They impose joint restrictions on consumption, income, and asset prices.

- Any solution must also satisfy **transversality conditions**:

$$\lim_{k\to\infty} \mathbb{E}_t \beta^k u'(c_{t+k}) R_{t+k}^{-1} L_{t+k} = 0.$$

$$\lim_{k\to\infty} \mathbb{E}_t \beta^k u'(c_{t+k}) p_{t+k} s_{t+k} = 0.$$

A Famous Question: Are Prices Martingales?

- An intuitive definition of *market efficiency* is that prices are martingales: all information is incorporated today, so we can't "predict" price movements.
- The Euler equation for equity shows that this will generically fail. We write

$$p_t = \mathbb{E}_t \beta (y_{t+1} + p_{t+1}) \frac{u'(c_{t+1})}{u'(c_t)}$$

- Using the covariance formula $\mathbb{E}xy = \mathbb{E}x\mathbb{E}y + cov(x, y)$ implies that

$$p_t = \beta E_t (y_{t+1} + p_{t+1}) \mathbb{E}_t \frac{u'(c_{t+1})}{u'(c_t)} + \beta cov_t \left[(y_{t+1} + p_{t+1}), \frac{u'(c_{t+1}}{u'(c_t)} \right].$$

- For prices to be a martingale, we require $\mathbb{E}_t \frac{u'(c_{t+1})}{u'(c_t)} = const$ and $cov[\cdot] = 0$.
- Both restrictions require very strong assumptions.

A Famous Special Case: Risk Neutrality

- Suppose now that agents are risk neutral so that $u'(c_t) = const$. Then

$$p_t = \beta \mathbb{E}_t (y_{t+1} + p_{t+1})$$

- This has the general class of solutions

$$p_t = \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j y_{t+j} + \xi_t \left(\frac{1}{\beta} \right)^t .$$

  where ξ_t is any random process that satisfies $\mathbb{E}_t \xi_{t+1} = \xi_t$ (a martingale).
- This is the discounted sum of expected dividends plus a "bubble term."
  - The bubble term is typically zero in general equilibrium.

Constructing the Risk-Neutral Measure

- Denote the state of the economy by s_t and assume that it follows a Markov process with transition probabilites $\mu(s_{t+1}|s_t)$.

- Assume that an asset pays a dividend stream $\{d(s_t)\}_{t\geq 0}$.
- Using the Euler equation, the cum-dividend price follows the recursion

$$a(s_t) = d(s_t) + \beta \sum_{s_{t+1}} \mu(s_{t+1}|s_t) \frac{u'_{t+1}[c(s_{t+1})]}{u'_t[c(s_t)]} a(s_{t+1}).$$

Constructing the Risk-Neutral Measure

- We can rewrite this equation as

$$a(s_t) = d(s_t) + R_t^{-1} \sum_{s_{t+1}} \tilde{\mu}(s_{t+1}|s_t) a(s_{t+1}) = d(s_t) + R_t^{-1} \tilde{\mathbb{E}}_t a(s_{t+1}).$$

- To do so, we define the risk-free rate as the inverse sum of state prices,

$$R_t^{-1} = \beta \sum_{s_{t+1}} \mu(s_{t+1}|s_t) \frac{u'_{t+1}[c(s_{t+1})]}{u'_t[c(s_t)]}$$

  and the risk-neutral **transition measure** as

$$\tilde{\mu}(s_{t+1}|s_t) = R_t \beta \frac{u'_{t+1}[c(s_{t+1})]}{u'_t[c(s_t)]} \mu(s_{t+1}|s_t)$$

  where multiplying by R_t ensures that the “twisted” measure is in $(0, 1)$.

- We can then construct an “overall” risk-neutral measure using

Verifying the Martingale Part

- Consider an asset that pays $d_T = d(s_T)$ at date T and 0 otherwise.
- Prices satisfiy $a_T(s_T) = d(s_T)$ and

$$a_t(s_t) = \mathbb{E}_{s_t}\beta^{T-t}\left[\frac{u'(c_T(s_T))}{u'(c_t(s_t))}\right]a_T(s_T).$$

- Now consider some $t < T$ and define the deflated process

$$\tilde{a}_{t+j} = \frac{a_{t+j}}{R_t R_{t+1}\cdots R_{t+j-1}} \quad \text{for } j = 1\ldots T-1.$$

- Then we can verify that $\tilde{\mathbb{E}}_t\tilde{a}_{t+j} = \tilde{a}_t(s_t)$ where $\tilde{a}_t(s_t) = a(s_t) - d(s_t)$.

- An equivalent statement of the same result is that

$$\tilde{\mathbb{E}}[a(s_{t+1}|s_t)] = R_t[a(s_t) - d(s_t)].$$

Adjusting for interest rates and dividends, the asset prices is a martingale **with respect to the risk-neutral measure**!

Asset Prices in (General) Equilibrium

- Where does the consumption process and the risk-free rate come from?
- To do so, we need to specify a production possibility frontier.
- We now specialize our economy to the Lucas (1978) economy.

- There is a large number of identical agents. The only durable good in the economy is set of identical "trees," one for each person in the economy.
- At the beginning of period t, each tree yields a dividend of "fruit" y_t. The fruit is not storable, but the tree is perfectly durable. The **state** is $s_t = y_t$.
- There is a time-invariant transition p.d.f $Prob(s_{t+1}) \leq s' | s_t = s) = F(s', s)$.

- Number of shares in a tree is normalized to 1. Bonds are in zero net supply.

Asset Prices in (General) Equilibrium

- Since all agents are identical, we consider a single "representative agent."
  (More on when and why we can do this later).

- Since this an endowment economy, the market clearing condition is $c_t = y_t$.
  (This is **indifference pricing**: no trade on the equilibrium path).
- Using this restriction yields the following Euler equations:

$$u'(y_t)R_t^{-1} = \mathbb{E}_t \beta u'(y_{t+1}).$$

$$u'(y_t)p_t = \mathbb{E}_t \beta (y_{t+1} + p_{t+1}) u'(y_{t+1}).$$

- Now use the law of iterated expectations $\mathbb{E}_t \mathbb{E}_{t+1}(\cdot) = \mathbb{E}_t(\cdot)$ and iterate:

$$u'(y_t)p_t = \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j u'(y_{t+j}) y_{t+j} + \mathbb{E}_t \lim_{k \to \infty} \beta^k u'(y_{t+k}) p_{t+k}.$$

Asset Prices in (General) Equilibrium

- The limiting term on RHS must be zero (Why?). Hence the asset price is

$$p_t = \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j \frac{u'(y_{t+j})}{u'(y_t)} y_{t+j}.$$

- This a nice generalization of our previous risk-neutral discounting formula.
- Risk aversion and aggregate risk $\Rightarrow$ time-varying stochastic discount rates.
- Under a Markov transition matrix, we also know that $p_t = p(s_t)$.

- A simple example with a particular clean solution is log utility.
- If $u(c) = \log(c)$ then $u'(c) = c^{-1}$. Hence our solution becomes

$$p_t = \mathbb{E}_t \sum_{j=1}^{\infty} \beta^j y_t = \frac{\beta}{1-\beta} y_t.$$

What does this say about price cyclicality and sensitivity to shocks?

Risk-Sharing and The Existence of Representative Agents

- Consider a complete-markets model with different agents indexed by i.
- Assume that these agents agree on the probability distribution over states.
- The first-order condition for any Arrow-Debreu claim on state s_{t+1} is

$$\beta\mu(s_{t+1}|s_t)u_i'(c_{t+1}^i(s_{t+1})) = q(s_{t+1}|s_t)u_i'(c_t^i(s_t)).$$

- We can rearrange this to solve for the claims price:

$$\beta\mu(s_{t+1}|s_t)\frac{u_i'(c_{t+1}^i(s_{t+1}))}{u_i'(c_t^i(s_t))} = q(s_{t+1}|s_t)$$

  We can also take ratios of two states, s_{t+1} and s_{t+1}', to give

$$\frac{\mu(s_{t+1}|s_t)u_i'(c_{t+1}^i(s_{t+1}))}{\mu(s_{t+1}'|s_t)u_i'(c_{t+1}^i(s_{t+1}'))} = \frac{q(s_{t+1}|s_t)}{q(s_{t+1}'|s_t)}$$

Risk-Sharing and The Existence of Representative Agents

- The right-hand side of these equations is the same for all agents!
- Leads to perfect risk-sharing: marginal utilities are aligned state-by-state.

- Moreover, agents have the same ordering of marginal utility across states.
- We can construct a utility function that represents this ordering. This is the utility function of the representative agent. It need not be related directly to underlying agent preferences.

- This argument works only in complete markets. (Why?)

Another Convenient Feature of Complete Markets

- We can appeal to Welfare Theorems: market allocation is Pareto efficient.
- Hence we can solve a Social Planner's problem to obtain allocations.
- Social Planner's problem (appeal to First Welfare Theorem)

$$\max_{\{c_t^i, c_t^j\}} \eta_i \mathbb{E}\left[\sum_t \beta^t u_i\left(C_t^i\right)\right] + \eta_j \mathbb{E}\left[\sum_t \beta^t u_j\left(C_t^j\right)\right], \ s.t.\ c_t^i + c_t^j = C_t,$$

  with Pareto weights η_i and η_j, and FOC

$$\eta_i u_i'\left(C_t^i\right) = \eta_j u_j'\left(C_t^j\right).$$

- Condition reflects perfect risk-sharing, and holds ex-ante and ex-post.

Overview

- No Arbitrage and the Stochastic Discount Factor
- **Bounds on SDFs as a Diagnostic Tool**
- Applications of SDF and Risk-neutral Measures

Hansen-Jagannathan Bounds

- Recall that we have the basic asset pricing equation $1 = \mathbb{E}[m \cdot R_i]$.
- We showed that this can be decomposed as

$$1 = \mathbb{E}[m]\mathbb{E}[R_i] + cov(m, R_i) = \mathbb{E}[m]\mathbb{E}[R_i] + \rho_{m,R_i}\sigma(R_i)\sigma(m).$$

This implies the following, where $\rho'_{m,R_i} \in [-1, 1]$:

$$\mathbb{E}R_i = R_f - \rho_{m,R_i}\sigma(R_i)\frac{\sigma(m)}{\mathbb{E}m}.$$

Hence we can bound the volatility of the stochastic discount factor:

$$\frac{\sigma(m)}{\mathbb{E}m} \geq \frac{|\mathbb{E}[R_i^e]|}{\sigma(R_e)}.$$

Hansen-Jagannathan Bounds

- The right-hand side is the Sharpe Ratio of asset i.
- Hence we have a lower bounds on the volatility of the SDF.
- This bound is particularly tight if we know the risk-free rate $R_f = 1/\mathbb{E}m \approx 1$.
- It is also tightest for the asset with the highest Sharpe Ratio.

- We can find the tightest bound using mean-variance analysis.
- This produces a set of SDFs (also defined by mean-variance) that can price a given set of assets.

  This is a very useful tool for diagnosing the potential of a model.

- We will discuss this further. HJ provide full analysis without risk-free asset.

Entropy Bounds on SDF

- Alvarez and Jermann (2005) use entropy bounds to define a lower bound on the volatility of the "permanent component" of an SDF.
- That is, they decompose $M_t = M_t^P M_t^T$ where M_t^P is a martingale.
- This allows us to get information about the SDF from long-dated assets.

- Define entropy of a positive random variable X as

$$L(X) = \log E[X] - E[\log X] \geq 0.$$

  - for lognormal random variables, $L(X) = \frac{1}{2}\sigma_x^2$ $(x = \log X)$
  - $L(aX) = L(X)$ for constant a

Entropy Bounds on SDF

- In a finite-state model, we have

$$\pi = \frac{\psi}{\mathbf{p}} = \frac{1}{R^f} \cdot \frac{\mu}{\mathbf{p}},$$

  where again μ is risk-neutral measure

- It then follows if R^f is constant (returns i.i.d. in a dynamic setting) that

$$\begin{aligned} L(\pi) &= L\left(\frac{1}{R^f}\frac{\mu}{\mathbf{p}}\right) = L\left(\frac{\mu}{\mathbf{p}}\right) \\ &= \log E^p\left[\frac{\mu}{\mathbf{p}}\right] - E^p\left[\log\left(\frac{\mu}{\mathbf{p}}\right)\right] = -E^p\left[\log\left(\frac{\mu}{\mathbf{p}}\right)\right]. \end{aligned}$$

- Entropy of SDF is a measure of deviation of μ from $\mathbf{p}$. AJ (2005) show that

$$L(\pi) \geq E\left[r_j - r^f\right].$$

  - a high log risk premium implies high entropy of SDF.

Entropy Bounds on SDF

- Proof: Since $E^p[\pi \cdot \mathbf{R}_i] = 1$, $E^p[\log \pi] + E^p[\log \mathbf{R}_i] \leq \log E^p[\pi \cdot \mathbf{R}_i] = 0$, from which follows that

$$E^p[\log \mathbf{R}_i] \leq -E^p[\log \pi].$$

- Allowing for time-variation in price of riskless asset: $1/R_t^f = E_t^p[\pi_{t+1}]$, entropy of riskless asset price is

$$L\left(R_t^{f-1}\right) = \log E_t^p\left[R_t^{f-1}\right] - E_t^p\left[\log R_t^{f-1}\right] = \log E_t^p[\pi] + E_t^p\left[\log R_t^f\right].$$

- It then follows from these two results that

$$\begin{aligned} L(\pi) &= \log E^p[\pi] - E^p[\log(\pi)] \\ &\geq L\left(R_t^{f-1}\right) + E^p\left[\log \mathbf{R}_i - \log R^f\right] \\ &\geq E^p\left[\log \mathbf{R}_i - \log R^f\right]. \end{aligned}$$

- Ultimately, Alvarez-Jermann provide a lower bound on the volatility of the permanent component of M_t relative to the overall volatility of M_t.
- The empirical implementation argues that prices of long-dated bonds reflect properties of the permanent component.
- To be consistent with the pricing of long-dated bounds, models must therefore have volatile *permanent* innovations to the SDF.

Overview

- No Arbitrage and the Stochastic Discount Factor
- Bounds on SDFs as a Diagnostic Tool
- **Applications of SDF and Risk-neutral Measures**

Physical vs Risk-neutral Measures

- Suppose we know the risk-neutral probabilities of a state. What good is this without knowing the true state prices?
- Black-Scholes show we can make progress on pricing redundant assets.
- The prototypical example is an option, which can be replicated using a stock and a bond.

- More generally, we can decompose state prices into

$$\psi_t = \pi_t p_t,$$

  which compose the risk-neutral measure through $\mu_{st+1} = \frac{\psi_{st+1}}{\sum_{s=1}^{S} \psi_{st+1}}$
- If we can observe μ_{st+1}, can we learn anything about π_t and p_t?

- In principle, one can recover μ_{t+1} from observing a cross-section of option prices that differ in their strike price K (Breeden and Litzenberger (1978)).
- For a call option for date T with price $C_0(K)$ for strike price K on underlying stock with price S_t, one has that

$$C_0(K) = E^Q\left[\frac{1}{R_{0,T}}(S_T - K)\cdot 1_{\{S_T \geq K\}}\right],$$

from which follows

$$\frac{\partial^2 C_0(K)}{\partial K^2} = \frac{\partial}{\partial K}\left(\frac{1}{R_{0,T}}(S_{sT} - K)\,\mu_{sT}\right) = -\frac{1}{R_{0,T}}\mu_{sT}.$$

(Here we are abstracting from issues of liquidity).
- We cannot decompose μ_{sT}, into π_t and p_t without more structure. High state price can reflect high marginal utility or high likelihood of the state

Ross Recovery Theorem (Ross (2015))

- With complete markets, FOC for representative agent is

$$U_i' \psi_{ij} = \beta U_j' p_{ij},$$

  where ψ_{ij} is the Arrow-Debreu state price for state j if current state is i, and we can interpret $U_i' = U_i'(c(\theta_i))$
- We can write this as SDF $\Lambda_{1,j} = \beta \frac{U_j'}{U_1'}$, letting the current state be 1.

Ross Recovery Theorem (Ross (2015))

- Stacking $\psi_t = \pi_t p_t$ into matrix form, one has the matrix equation

$$G\Psi = \beta P G,$$

where Ψ is the $S \times S$ matrix of state contingent Arrow-Debreu prices ψ_{ij}, P is the $S \times S$ matrix of natural probabilities p_{ij}, and G is the diagonal matrix with the undiscounted kernel

$$G = \frac{1}{U_1'} \begin{bmatrix} U_1' & 0 & 0 \\ 0 & U_i' & 0 \\ 0 & 0 & U_m' \end{bmatrix},$$

and no arbitrage is guaranteed with strictly positive state prices

Ross Recovery Theorem (Ross (2015))

- Manipulating the matrix equation

$$P = \frac{1}{\beta} G \Psi G^{-1}.$$

  We have S^2 equations with S^2 unknown probabilities, S unknown marginal utilities, and 1 discount rate

- Ross (2015) has insight that, since P is a stochastic matrix whose rows are transition probabilities, one also has the additional restriction

$$P \cdot \mathbf{1}_{S\times 1} = \mathbf{1}_{S\times 1}.$$

- He uses this to back out the physical measure. (Details in paper).
- Borovicka, Hansen, and Scheinkman (2016) argue Ross (2015) is able to recover p_t only using restrictions on the SDF that are easily rejected.

Nominal SDF

- Suppose nominal price of an asset is $P_t q_{it}$ that pays dividend $P_t \delta_{it}$, where P_t is the price level and q_{it} is the price in the numeraire.
- Then no arbitrage pricing requires

$$\hat{\pi}_t P_t q_{it} = E_t^p \left[\hat{\pi}_{t+1} P_{t+1} \left(\delta_{it+1} + q_{it+1} \right) \right],$$

which can be rewritten, with $i_{t+1} = \frac{P_{t+1}}{P_t}$ as inflation, as

$$q_{it} = E_t^p \left[\frac{\hat{\pi}_{t+1}}{\hat{\pi}_t} i_{t+1} \left(\delta_{it+1} + q_{it+1} \right) \right] = E_t^p \left[\frac{\pi_{t+1}}{\pi_t} \left(\delta_{it+1} + q_{it+1} \right) \right],$$

where second equality is from definition of pricing real assets

- It then follows the above holds state-by-state, and the nominal SDF $\hat{\Lambda}_{t,t+1} = \frac{\hat{\pi}_{t+1}}{\hat{\pi}_t}$ and the real SDF $\Lambda_{t,t+1}$ are related by

$$\Lambda_{t,t+1} = \hat{\Lambda}_{t,t+1} i_{t+1}.$$