Financial Economics Notes

Jordan Pandolfo

Contents

Some Info

These notes are pulled pretty much from two sources: 1) directly from Leroy and Werner's *Principles of Financial Economics, Second Edition* and 2) from the lectures notes and papers that correspond to Anmol Bhandari's section of the course. While the main notes and proofs focus on the two-date version of an asset market economy, most proofs and theorems generalize almost perfectly to the multi-date or infinite-time economies. The reason for this focus is that the intuition and math is easier to follow/instill. These notes (like the book) do not include more recent material from the course on ambiguity, endogenous debt constraints, speculative trade and rational expectations equilibrium, but this should be a sufficient introduction to foundational material for financial economies in discrete time and discrete state space.

1 Equilibirum in Security Markets

- Markets are *complete* if and only if the payoff matrix X is full rank.
- Security prices p are marginal utility-weighted dividends.
- Securities market equilibrium exists under mild conditions.

A security $x_j \in \mathbb{R}^S$ has payoffs in S states of nature. For an economy with J securities, define the $J \times S$ matrix X as the *payoff matrix*. Further, define a *portfolio* h as the J-vector, describing the amount of each security held by the agent. The portfolio has a payoff $hX = \sum_{j=1}^{J} h_j x_j \in \mathbb{R}^S$. We can therefore define the *asset span* $\mathcal{M}$ as

$$\mathcal{M} = \{z \in \mathbb{R}^S : z = hX \text{ for some } h \in \mathbb{R}^J\}.$$

Thus, the asset span consists of all possible payoffs an agent could achieve in the S states of nature. If $\mathcal{M} = \mathbb{R}^S$, we say markets are *complete*; otherwise, they are *incomplete*.

Theorem. *Markets are complete iff the payoff matrix X has a rank S.*

Proof. If $J < S$, then there are an insufficient number of assets to span the space $\mathbb{R}^S$. If $J > S$, then there exists a basis of S linearly independent rows that can replicate the remaining $J - S$; thus, WLOG consider a square matrix. Given the columns are linearly independent, the matrix X has an inverse and $h = X^{-1}z$ maps back to a solution for h for any point z. □

Prices of the securities can be represented by the J-vector p and the price of a portfolio h is simply the inner product ph. In a similar sense, the returns on a security j is simply $\frac{x_j}{p_j}$.

Consider an economy with a finite number I of agents with utility functions $u^i : \mathbb{R}_+^{S+1} \to \mathbb{R}$ who have endowments ω_i. A *securities market economy* is one in which the agents' endowments lie within the asset span. Consider the agents problem

$$\max_{c_0, c_1, h} \quad u(c_0, c_1(1), ..., c_1(s)) \tag{1}$$
$$s.t. c_0 \leq \omega_0 - ph$$
$$s.t. c_1 \leq \omega_1 + hX$$

where I make explicit the utility function over states. Further, notice the second constraint is a system of inequalities for all states $s = 1, 2, ..., S$. After applying first-order conditions, we observe

$$p_j = \sum_{s=1}^{S} x_{js} \frac{u_{1,s}}{u_0}; \tag{2}$$

that is, the the price of an asset is equal to the weighted sum of its payoffs, where the weight is the marginal rate of substitution in its corresponding state.

Definition. *A security markets equilibrium consists of a price vector* p*, portfolio choice* h_i*, consumption allocation* $\{c_0^i, c_1^i\}$ *for each agent such that agents maximize their utility and markets clear:*

1. $\sum_i h_i = 0$
2. $\sum_i c_j^i = \sum_i \omega_j^i \quad j = 0, 1.$

Theorem. *If each agent's consumption plan is restricted to be positive, initial endowments strictly positive, a portfolio with a positive, nonzero payoff exists, and utility functions are strictly increasing and quasi-concave, then an equilibrium in security markets exists.*

2 Linear Pricing

- Define the payoff pricing functional $q(z)$.
- q is linear (on the asset span $\mathcal{M}$) if and only if the law of one price holds.
- Define state prices $q \in \mathbb{R}^S$, using Arrow securities. Observe relation between security prices p and q.

Definition. *The law of one price states that if*

$$hX = h^*X \Rightarrow ph = ph^*.$$

That is to say, if two portfolios attain the same point in the payoff space, then they must be the same price.

Proposition. *The law of one price holds iff every portfolio with zero payoff has zero price.*

Proof. ($\Rightarrow$): Suppose the Law of One Price holds. Further, suppose that $hX = \tilde{h}X = 0$ but $ph = p\tilde{h} > 0$. Consider a $\lambda > 1$. Then $\lambda hX = 0$ still and $ph \neq p(\lambda h)$ which contradicts the law of one price.

($\Leftarrow$): Suppose that for all portfolios h such that $hX = 0$, we have $ph = 0$. Consider portfolios $h, \tilde{h}$ such that $hX = \tilde{h}X$ but $ph > p\tilde{h}$. Then, we also observe

$$(h - \tilde{h})X = \bar{h}X = 0$$

but $p\bar{h} > 0$ which contradicts our assumption. □

Definition. *Given security prices p, define the payoff pricing functional $q : \mathcal{M} \to \mathbb{R}$ as a function that assigns a price to each payoff in $\mathcal{M} \subseteq \mathbb{R}^S$; that is,*

$$q(z) = \{w : w = ph \text{ for some } h \text{ such that } z = hX\} \quad \forall z \in \mathcal{M}.$$

Obviously, if the law of one price holds, the payuoff pricing functional q must be single-valued.

Theorem. *The law of one price holds iff q is a linear functional on the asset span $\mathcal{M}$.*

Proof. ($\Rightarrow$) : Assume that the Law of One Price holds. Note that this implies that q must be single-valued. Now, consider two payoffs $z = hX, \tilde{z} = \tilde{h}X \in \mathcal{M}$. For arbitrary $\lambda, \mu \in \mathbb{R}$ we can generate $\lambda z + \mu\tilde{z}$ with portfolio $\lambda h + \mu\tilde{h}$. The pricing functional is

$$\begin{aligned} q(\lambda z + \mu \tilde{z}) =& \lambda ph + \mu p\tilde{h} && \text{(By Law One Price)} \\ =& \lambda q(z) + \mu q(\tilde{z}) \end{aligned}$$

which implies linearity.

($\Leftarrow$): Suppose that q is a linear functional on $\mathcal{M}$. Then, for any portfolio $h, \tilde{h}$ such that $hX = \tilde{h}X$ we must have $q(hX) = q(\tilde{h}X)$. □

Theorem. *If agents have strictly increasing utility functions, the law of one price will hold in equilibrium.*

Proof. Given equilibrium prices p^*, suppose that the Law of One Price does not hold. Then, there exists a zero payoff portfolio with negative price (i.e. an h such that $hX = 0$ and $ph < 0$). Thus, for any equilibrium plan h^*, the agent can choose $\tilde{h} = h^* + h$ along with consumption $(c_0^* - p^*h, c_1^*)$ which is budget feasible and increases consumption. Since utility is strictly increasing in consumption, this strategy is preferable and thus h^* cannot be an equilibrium. A contradiction. □

Now, if we have i) complete markets, we know the payoff pricing functional is unique and if ii) the law of one price holds, the payoff pricing functional will be linear. Further, by assuming strictly increasing preferences, we guarantee the law of one price to hold, in equilibrium. For *Arrow securities* which have a unit vector payoff e_s for all states $s = 1, 2, ..., S$, we can recover the unique payoff price $q_s = q(e_s)$. Create the vector of Arrow-Debreu or *state* prices $q \in \mathbb{R}^S$. Further, given linearity of the pricing functional, we can price every payoff with the vector q; that is, for any payoff $z \in \mathbb{R}^S$, its price is

$$q(z) = qz = \sum_{s=1}^{S} q_s z_s$$

is simply the inner product between the states prices and the payoff values. Further, the price of each security is simply $p_j = qx_j$ or in matrix notation,

$$p = Xq. \tag{1}$$

Once again, we know q is unique because markets are complete, which means that the matrix X has rank S and can be inverted (assuming no redundant securities exist). Thus, solving for state prices is solving the system of equations in (1). In Section 1 equation (1), we setup and solved the agents utility maximizatino problem. We did so by choosing consumption and a portfolio of assets. Now, we can recast that optimization problem as one of choosing payoffs z instead of portfolios h:

$$\begin{aligned} \underset{c_0, c_1, z}{max} \quad & u(c_0, c_1) \qquad (2) \\ s.t. c_0 &\leq \omega_0 - qz \\ s.t. c_1 &\leq \omega_1 + z. \end{aligned}$$

Given strictly increasing utility, we have to maximize $u(\omega_0 - qz, \omega_1(1) + z(1), ..., \omega_1(S) + z(S))$ with respect to z_s for each possible payoff $s = 1, 2, ..., S$. The first-order conditions of this problem lead to $q_s = \frac{u_1(s)}{u_0(s)}$ for each state of the world. Thus, coupled with Section 1 equation (2), we get

$$p_j = \sum_{s=1}^{S} x_{j,s} \frac{u_{1,s}}{u_0} = \sum_{s=1}^{S} x_{j,s} q_s = x_j q. \tag{3}$$

Thus, in a consumption-based model, the price of a security is its marginal utility-weighted dividends/payoffs <u>and</u> this marginal utility ratio serves as the unique vector of state prices (given market completeness and law of one price).

3 Arbitrage and Positive Pricing

- q is linear and strictly positive if and only if there is no arbitrage.
- If agent preferences are strictly increasing, there cannot exist arbitrage in equilibrium; hence, q will be linear and strictly positive, in equilibrium.

Definition. *An arbitrage is a portfolio h such that $ph \leq 0$ and $hX \geq 0$ with one (or both) strict inequalities.*

Proposition. *The law of one price holds iff there does not exist an arbitrage.*

Proof. Rely on zero payoff property of Law of One Price which states that

$$\begin{aligned} \text{Law of One Price holds } &\iff \left\{ hX = 0 \Rightarrow ph = 0 \right\} \\ &\iff \nexists \quad h \quad \text{such that } hX = 0, ph < 0 \end{aligned}$$

which is the definition of an arbitrage. □

Theorem. *The payoff pricing functional q is linear and strictly positive iff there is no arbitrage.*

Proof. $(\Rightarrow)$: Suppose the payoff pricing functional were strictly positive and linear. Further, suppose arbitrage exists; that is, we have $z = hX = 0$ and $ph < 0$ which means

$$\begin{aligned} q(z) < 0 \Rightarrow & q \cdot z < 0 \\ \Rightarrow & q(hX) < 0 \end{aligned}$$

which is a contradiction, given $hX = 0$. Remember, the zero payoff must have zero price under linearity/law of one price.

$(\Leftarrow)$: Assume there exists no arbitrage. Now, consider the spaces

$$\begin{aligned} L =& \mathbb{R} \times \mathbb{R}^S \\ M =& \{-q(z), z\} \subset L \\ K =& \mathbb{R}_+ \times \mathbb{R}_+^S \subset L \end{aligned}$$

for any payoff $z \in \mathbb{R}^S$. No arbitrage implies $K \cap M = \{0\}$. Given that both K and M are closed, convex spaces, the Separating Hyperplane Theorem implies there exists a linear, nonzero functional $F : \mathbb{R}^{S+1} \to \mathbb{R}$ such that $F(k) = F'k > F(m) = F'm$ for all interior points $k \in K$ and $m \in M$.

It must be the case that $F(m) = 0$. To see this, suppose point $m \in M$ were attained with portfolio h (i.e. $F(m) = F'm = -F_1 q(hX) + F'_{-1}(hX)$ which is smaller than $F'k$ for all $k \in K$ by the hyperplane theorem). The set M includes all prices and payoffs attainable from any portfolio, so it must also include $-h$. Thus, we would have $F_1 q(hX) - F'_{-1}(hX) > F'k$ for all $k \in K$ which would contradict the hyperplane theorem.

Next, note that the linear functional used as the hyperplane is represented simply by the vector $F \in \mathbb{R} \times \mathbb{R}^S$. Given that $F(k) > F(m) = 0$, and k is a strictly positive vector, it must be the case that $F >> 0$. Let's define

$F = (\alpha, \psi)$ where α is a scalar. Choose a payoff m with corresponding $q(m)$ in M; then,

$$F(m) = -\alpha q(m) + \psi \cdot m = 0$$

which leads to $q(m) = \frac{\psi}{\alpha} \cdot m$ and thus we have the existence of a strictly positive payoff pricing functional. $\square$

Now, consider this within the context of a security market economy.

Theorem. *If at any given price p, an agent's optimal portofolio exists, and if the utility functions of the agents are strictly increasing, then there does not exist an arbitrage.*

Further, if at given security prices there is no arbitrage and agent consumption is restricted to be positive, then there exists an optimal portofolio.

Theorem. *If agent utility functions are strictly increasing, then there is no arbitrage at equilibrium security prices and the equilibrium payoff pricing functional is linear and strictly positive.*

4 Valuation

- Define the valuation functional $Q(z)$ and construct it from $q(z)$.
- Security prices exclude arbitrage if and only if there exists a strictly positive Q.
- Given no arbitrage, security markets are complete if and only if there exists a unique, strictly positive Q.

Consider cases in which markets are not complete. The payoff pricing functional is defined with a domain in the asset span. We can define the valuation functional $Q : \mathbb{R}^S \to \mathbb{R}$ which is a linear functional on the entire space $\mathbb{R}^S$ with the restriction that

$$Q(z) = q(z)$$

for all payoffs within the asset span. So, the valuation functional coincides with the payoff functional on the asset span and can also price payoffs on the larger contingent claim space $\mathbb{R}^S$.

Now, in the case of incomplete markets, let's show how the valuation functional can be constructed/extended from the underlying payoff pricing functional. Essentially, we will start with a payoff pricing functional q which only spans $S - J$ dimensions, find a payoff $\hat{z}$ not in the span and then define a new operator to include this new payoff. First, define bounds on the payoff pricing functional

$$\begin{aligned} q_u(z) =& \min_h \{ph : hX \geq z\} \\ q_l(z) =& \max_h \{ph : hx \leq z\} \end{aligned}$$

for all $z \in \mathbb{R}^S$. We have the following result for these functionals within the asset span.

Proposition. *If security prices exclude arbitrage,*

$$q_u(z) = q_l(z) = q(z).$$

This just states that we have uniqueness of the price of a payoff when we consider payoffs in the asset span. Further, it markets are complete, then the asset span is the space $\mathbb{R}^S$ and we once again have uniqueness such that $Q(z) = q(z)$ for all $z \in \mathbb{R}^S$. Now, consider the following proposition for payoffs outside the asset span.

Proposition. *If security prices exclude arbitrage, then*

$$q_u(z) > q_l(z) \quad \forall z \notin \mathcal{M}.$$

So, (given no arbitrage) for any given set of security prices, we know that there exists a unique payoff pricing functional within the asset span; but, when we consider payoffs outside the span, we allow for the possibility of a range of values. To illustrate this result, consider the example on the next page.

Example. *Consider two securities*

$$x_1 = (1, 1, 1)$$
$$x_2 = (1, 2, 4)$$

such that we have $J < S$. Suppose market prices are $(p_1, p_2) = (\frac{1}{2}, 1)$. A call option with strike price 3 on asset x_2 has the payoff $c(k = 3) = (0, 0, 1)$. We can show that $c(k = 3)$ is not in the asset span. So, let's compute an interval of permissible payoff prices for the portfolio, using q_l and q_u, which correspond to the call option payoff, labelled c_3.

$$q_l(c_3) = \max_{h_1, h_2} \quad \frac{1}{2}h_1 + h_2$$
$$s.t. h_1 + h_2 \leq 0$$
$$s.t. h_1 + 2h_2 \leq 0$$
$$s.t. h_1 + 4h_2 \leq 1$$

where the constraints are simply the requirement $h_1 x_1 + h_2 x_2 \leq c_3$, written in form of system of equations. Solve this and we get $q_l(c_3) = 0$.

Now, consider the other problem

$$q_u(c_3) = \min_{h_1, h_2} \quad \frac{1}{2}h_1 + h_2$$
$$s.t. h_1 + h_2 \geq 0$$
$$s.t. h_1 + 2h_2 \geq 0$$
$$s.t. h_1 + 4h_2 \geq 1$$

which leads to $q_u(c_3) = \frac{1}{6}$. Thus, we obtain all possible permissible prices $q(c_3) \in [0, \frac{1}{6}]$ to price the call option with strike price 3.

Now, construct Q in the following way: consider the asset span $\mathcal{M} \in \mathbb{R}^{S-J} \subset \mathbb{R}^S$. Choose a payoff $\hat{z} \notin \mathcal{M}$. Define the space

$$\mathcal{N} = \{z + \lambda \hat{z} : z \in \mathcal{M} \text{ and } \lambda \in \mathbb{R}\}$$

which is in the space $\mathbb{R}^{S-J+1}$ and is thus the $span\{x_1, ..., x_J, \hat{z}\}$. Choose any scalar π such that $q_l(\hat{z}) \leq \pi \leq q_u(\hat{z})$ (where these functionals have domain $\mathcal{M}$). Now, define the valuation functional $Q : \mathcal{N} \to \mathbb{R}$ as

$$Q(z + \lambda \hat{z}) = q(z) + \lambda \pi.$$

Keep repeating this process until we have a space $\mathcal{N}$ of dimension S. Then you're done.

Proposition. *If $q : \mathcal{M} \to \mathbb{R}$ is strictly positive, then $Q : \mathcal{N} \to \mathbb{R}$ is strictly positive.*

Now, we restate a couple theorems from the previous section in terms of the valuation functional.

Theorem. *The First Fundamental Theorem of Finance states that security prices exclude arbitrage iff there exists a strictly positive valuation functional.*

Proof. $(\Rightarrow)$: We know from previous theorem that no arbitrage implies a strictly positive payoff pricing functional q on the asset span $\mathcal{M}$. From the last proposition, a strictly positive q with no arbitrage implies a strictly positive Q.

$(\Leftarrow)$: We know that a strictly positive valuation functional implies a strictly positive payoff functional, by definition. Thus, from our previous theorem, there cannot be arbitrage with security prices p. □

Theorem. *Suppose security prices exclude arbitrage. Then security markets are complete iff there exists a unique strictly positive valuation functional.*

Proof. $(\Rightarrow)$: If security markets are complete, we know that the payoff pricing functional is unique on the asset span $\mathcal{M}$ which is $\mathbb{R}^S$. Thus, $Q(z) = q(z) \quad \forall z \in \mathbb{R}^S$.

$(\Leftarrow)$: Assume we have a unique strictly positive valuation functional. Suppose markets are not complete. Then $\mathcal{M} \subset \mathbb{R}^S$ and for any payoff $\hat{z} \notin \mathcal{M}$, we have $q_l(z) < q_u(z)$. Recall, we construct the valuation functional from q by choosing an element $\hat{z} \notin \mathcal{M}$with values define by $Q(z + \lambda\hat{z}) = q(z) + \lambda\pi$ where we choose π. Given that $\pi \in [q_l(\hat{z}), q_u(\hat{z})]$ comes from a non-degenerate interval, we can have multiple values of the valuation functional. A contradiction. □

5 State Prices and Risk-Neutral Probabilities

- The risk free return $\bar{r}$ is the inverse of the sum of state prices q_s.
- Define a *risk-neutral* probability measure π^*.
- The price of a portfolio, attaining payoff z, is its discounted, expected value under the π^*-measure.
- For any asset $E_*[r] = \bar{r}$ under the π^*-measure.

Just like we did for state prices in complete markets with $q(e_s) = q_s$, we can identify the valuation functional by $Q(e_s) = q_s$ for the basis vectors e_s. Given linearity of the valuation functional, we have

$$Q(z) = qz$$

and this is called the state-price representation of the valuation functional Q.[1] Further, we know we can derive the system of equations

$$p = Xq$$

such that state prices are a solution to the system of J equations with S unknowns.

Theorem. *There exists a strictly positive valuation functional iff there exists a strictly positive solution to* $p = Xq$.

Proof. ($\Rightarrow$):If we have a strictly positive valuation functional, we can use the basis functions to construct a strictly positive q that satisfies $p = Xq$, as was shown above.

($\Leftarrow$): If there exists a strictly positive solution q to the matrix equation, then define a strictly linear operator $Q(z) = qz$. Then, for any payoff $z \in \mathbb{R}^S$, there exists a portfolio h such that $z = hX$ and we write

$$Q(z) = q(hX) = ph$$

which prices the the payoff and corresponding portfolio. Thus, $Q(z)$ is a well-defined valuation functional and is strictly positive. □

Consider a payoff that does not depend on the state of the world s. Such a payoff is *risk free.* Consider a non-zero payoff k that is risk-free (i.e. it's payoff vector is $(k, k, ..., k)'$). Then, it has the risk-free return $\bar{r}$ of

$$\bar{r} = \frac{k}{\sum_s q_s k} = \frac{1}{\sum_s q_s}. \tag{1}$$

Thus, the inverse of the sum of state (Arrow-Debreu) prices gives the risk-free return. Now, define a *risk-neutral probability* $\pi_s^* = \bar{r} q_s$ such that

$$\pi_s^* = \bar{r} q_s = \frac{q_s}{\sum_s q_s} \tag{2}$$

is just a re-scaling of the state prices. Further, π^* is a probability measure. Given this, we can consider the expected value of a payoff z. We could multiply $\frac{\bar{r}}{\bar{r}}$ to the valuation functional of a payoff z:

$$\frac{\bar{r}}{\bar{r}} Q(z) = \frac{1}{\bar{r}} \sum_s \bar{r} q_s z_s = \frac{1}{\bar{r}} \sum_s \pi_s^* z_s = \frac{1}{\bar{r}} E_*[z]. \tag{3}$$

[1] I'm using the same vector notation q instead of Q because often times we will be assuming complete markets, in which case $Q(z) = q(z)$ for all contingent payoffs z so we just stick to the standard state price vector q.

Thus, (3) states that the price of a portfolio attaining a payoff $z \in \mathbb{R}^S$ is the discounted expected value of the payoff, with respect to the risk-neutral probabilities. Further, given we know that $p_j = qx_j$, we get

$$p_j = \sum_s x_{j,s} q_s = \frac{1}{\bar{r}} E_*[x_j] \tag{4}$$

which also implies $\bar{r} = E_*[r_j]$ for all assets $j = 1, 2, ..., J$.[2] That is, the expected return (under the risk-neutral measure) of any asset is equal to the risk-free return. This is the sense in which we use the world *neutral*.

[2]There is a key decomposition to take away from equation (4). While the consumption-based model we have thus far considered doesn't explicitly include randomness, it easily could. We would then have a *natural* probability measure π over possible states of the world tomorrow. Recall from Section 1 equation (1), equilibrium security prices are determined via $p_j = \sum_s \frac{u_{1,s}}{u_0} x_{j,s}$ or $p_j = \sum_s \frac{\pi_{1,s}}{\pi_0} \frac{u_{1,s}}{u_0} x_{j,s}$ if randomness were modeled. Thus, the pricing of an asset depends upon its natural probability distribution of dividends/payoffs <u>and also</u> the consumer's marginal rate of substitution with respect to each state s. If we used power utility functions, then the parameter σ would capture the consumer's aversion to risk, depending on how large it was set. So, prices are the product of both i) agent beliefs about π and ii) agent attitude towards risk, via σ. When we switch to a risk-netrual measure π^*, we embed the agent's risk attitude inside the probability measure so that prices are simply calculated as a discounted expectation. If we observe dividends and prices in the data, we can calculate π^*; what is π and σ? That is the tougher question.

6 Optimal Portfolios with Multiple Risky Securities

- The agent problem of choosing consumption and a portfolio (c, h) is recast as one of choosing amounts of wealth a_j invested in asset j

Consider the portfolio choice problem

$$\begin{aligned} \underset{c_1,h}{max} \quad & E[v(c_1)] \\ s.t. \quad & ph = \omega_0 \\ s.t. \quad & c_1 = \omega_1 + hX \end{aligned}$$

This is the same problem as Section 1, but here the agent postpones consumption at date 0 and only considers consumption at date 1. Define $\hat{h}$ as a portfolio such that $\omega_1 = \hat{h}X$. Then the second constraint is written $c_1 = (\hat{h} + h)X$. Thus, define a_j as the amount of wealth invested in security j, written

$$a_j = p_j(\hat{h}_j + h_j)$$

and re-write the date 1 budget constraint as

$$c_1 = \sum_j \frac{a_j x_j}{p_j} = \sum_j a_j r_j.$$

This says that your consumption is equal to the returns of all the assets, scaled by the amount you invested in each asset. Now, the maximization problem can be re-stated as

$$\begin{aligned} \underset{\{a_j\}}{max} \quad & E[v(\sum_j a_j r_j)] \\ s.t. \sum_j a_j = \omega \end{aligned}$$

where ω is defined as agent's total wealth $\omega_0 + p\hat{h}$. Say we have a risk-free security (without loss of generality, security 1), then we can substitute for a_1 and write the problem as

$$\underset{a_2,\ldots,a_J}{max} \quad E[v(\omega\bar{r} + \sum_{j=2}^{J} a_j(r_j - \bar{r})].$$

In the case of just a risk-free security and one risky security, we have the first-order condition

$$E[v'\big(\omega\bar{r} + a^*(r - \bar{r})\big)(r - \bar{r})] = 0.$$

7 Consumption-Based Security Pricing

- Equilibrium conditions provide $\bar{r} = E[\frac{v_0}{v_1}]$.
- The *consumption-based security pricing* formula is derived.
- The *Hansen-Jagannathan* bounds are derived.

In general for a securities market economy with randomness, we have the equilibrium condition

$$p_j v_0 = E[v_1 x_j] \tag{1}$$

for a security j, which can be written $v_0 = E[v_1 r_j]$; thus, if there is a risk-free security, we observe

$$\bar{r} = E[\frac{v_0}{v_1}] = E_0[\frac{\frac{\partial v(c)}{\partial c_1}}{\frac{\partial v(c)}{\partial c_0}}]. \tag{2}$$

Recall that the marginal rate of substitution v_1/v_0 served as the payoff pricing functional q for an asset. Thus, the inverse of the risk-free rate $\bar{r}$ serves as a vector of state prices for equilibrium prices p^* (i.e. $\frac{1}{\bar{r}} = E[\frac{v_1}{v_0}] = E[q]$). This follows from market completeness whereby the agent will equalize his expected marginal rate of substitution across all states.

Use the covariance formula decomposition to write

$$\begin{aligned} v_0 =& E[v_1 r] \\ =& cov(v_1, r) + E[v_1]E[r] \\ \Rightarrow& E[r] = \frac{v_0}{E[v_1]} - \frac{cov(v_1, r)}{E[v_1]} \\ \Rightarrow& E[r] - \bar{r} = -\bar{r}\frac{cov(v_1, r)}{v_0} \end{aligned} \tag{3}$$

where the last line follows from (2). Equation (3) is called the *consumption-based security pricing* formula. It states that the excess return of a risky asset is proportional to the covariance of its returns with marginal utility at date 1. For example, consider a strictly increasing and concave utility function. If state s consumption c_s is high and return r_s is high, then we have low $v_1(c_s)$ and a therefore negative covariance, implying a positive risk premium. Agents have high demand for an asset if it pays out when times are bad; this asset, pays out large when consumption is large and pays out low when consumption is low which is not ideal. Therefore, to hold it, the investor must be compensated for the risk with a positive premium.

Definition. *Two contingent claims y and z are co-monotone if $(y_s - y_t)(z_s - z_t) \geq 0$ for all states s and t.*

Proposition. *If random variables y and z are co-monotone, then cov(y,z) ≥ 0.*

Theorem. *If an agent is risk averse, then expected return $E[r]$ is greater than the risk-free return $\bar{r}$ for every return r that is co-monotone with optimal consumption.*

Proof. Risk aversion implies a strictly increasing and strictly concave utility function; that is,

$$u'(c) > 0 \quad \text{and} \quad u''(c) < 0 \quad \forall c.$$

Further, given that c^* is co-monotnone with return r, we know that $cov(r, c^*) > 0$. High returns r correspond to high consumption c^* which corresponds to low marginal utility $\frac{\partial v}{\partial c^*}$; therefore,

$$-\bar{r}\frac{cov(\frac{\partial v}{\partial c_1^*}, r)}{\frac{\partial v}{\partial c_0^*}} > 0$$

and the corresponding consumption-based security pricing formula shows positive excess returns. □

Using the fact that $v_0 = E[v_1 r]$ and $\bar{r} = \frac{v_0}{E[v_1]}$, we have

$$E[v_1(r - \bar{r})] = 0. \tag{4}$$

Note that the correlation ρ between v_1 and $r - \bar{r}$ can be written

$$\begin{aligned}
\rho =& \frac{E[v_1(r-\bar{r})] - E[v_1]E[r-\bar{r}]}{\sigma(v_1)\sigma(r)} \\
=& -\frac{E[v_1]E[r-\bar{r}]}{\sigma(v_1)\sigma(r)} && \text{(From 4)} \\
\Rightarrow& |\sigma(v_1)\rho| = \frac{|E[v_1]E[r-\bar{r}]|}{\sigma(r)} \\
\Rightarrow& \frac{\sigma(v_1)}{E(v_1)} \geq \frac{|E[r]-\bar{r}|}{\sigma(r)} && \text{(From } |\rho| \leq 1) \\
\Rightarrow& \frac{\sigma(\frac{v_1}{v_0})}{E(\frac{v_1}{v_0})} \geq \frac{|E[r]-\bar{r}|}{\sigma(r)}
\end{aligned}$$

because v_0 is occurs at the present and therefore has no variance or uncertainty. Let's simplify notation and define state price vector $q = \frac{v_1}{v_0}$ on the RHS above. Given that this holds for all risky securities, we have the result:

$$\frac{\sigma(q)}{E[q]} \geq \sup_r \frac{|E[r]-\bar{r}|}{\sigma(r)} \tag{5}$$

The RHS of equation (5) is called a Sharpe-ratio and is the excess expected return of a security, normalized by its volatility. Equation (5) is called the *Hansen-Jagannathan* bounds for the state prices (or stochastic discount factor). Given data on prices and dividends, we can found permissible values for parameters used in the agent utility function.

8 The Expectations and Pricing Kernels

- Introduce the pricing kernel k_q as another representation of the payoff pricing function $q(z) = E[k_q z]$, with respect to some measure π.
- The Riesz Representation theorem implies k_q is unique on the asset span $\mathcal{M}$.
- Define a stochastic discount factor (SDF) $m \in \mathbb{R}^S$ and show k_q is the unique projection of all m onto $\mathcal{M}$.

Recall that the inner product is a function that maps from a $\mathcal{H} \times \mathcal{H}$ into the reals, where $\mathcal{H}$ is some vector space. Inner products are symmetric, linear and strictly positive when applied to the same element of $\mathcal{H}$. Represent the inner product between x and y as $x \cdot y$. Given an inner product, we can define the norm of a vector in $\mathcal{H}$ as

$$||x|| = \sqrt{x \cdot x}$$

which is an operator that satisfies i) the triangle inequality and ii) the Cauchy-Schwartz inequality.

Definition. *A Hilbert space is a vector space H which is equipped with an inner product and complete with respect to the norm induced by its inner product.*

Completeness of a space means that if there is a Cauchy sequence in the space which converges, then the point of convergence is also in the space.

Definition. *Two vectors $x, y \in \mathcal{H}$ are orthogonal (denoted $x \perp y$) if and only if their inner product is zero; that is,*

$$x \perp y \quad iff \quad x \cdot y = 0.$$

If we have a collection of vectors $\{z_i\}_{i=1}^n$ in a Hilbert space that are all orthogonal to one another, then we call this an *orthonormal system*. This orthonormal system makes up an *orthonormal basis* for its linear span. Any orthonormal system (of nonzero vectors) is linearly independent; that is, no one vector z_i can be replicated by linear combinations of vectors z_j with $j \neq i$. Now, on to orhogonality with respect to subspaces of the Hilbert space.

Definition. *A vector $x \in \mathcal{H}$ is orthogonal to a linear subspace $\mathcal{Z} \subset \mathcal{H}$ if and only if x is orthogonal to every vector $z \in \mathcal{Z}$.*

Further, the set of vectors x which are orthogonal to the subspace $\mathcal{Z}$ is called the *orthogonal complement* of $\mathcal{Z}$ and denoted as $\mathcal{Z}^{\perp}$. It, too, is a linear subspace of $\mathcal{H}$.

Theorem. *For any subspace $\mathcal{Z}$ of Hilbert space $\mathcal{H}$ and a vector $x \in \mathcal{H}$, $\exists$ a unique vector $x^{\mathcal{Z}} \in \mathcal{Z}$ and $x^{\mathcal{Z}^\perp} \in \mathcal{Z}^\perp$ such that $x = x^{\mathcal{Z}} + x^{\mathcal{Z}^\perp}$.*

Proof. Let $\{z_1, ..., z_n\}$ be an orthogonal basis for $\mathcal{Z}$ where $z_i \in \mathbb{R}^{dim(\mathcal{Z})}$ for all $i = 1, 2, ..., n$. Define vector

$$x^{\mathcal{Z}} = \sum_{i=1}^{n} \frac{x \cdot z_i}{z_i \cdot z_i} z_i \tag{1}$$

and the vector $y = x - x^{\mathcal{Z}}$. Because $x^{\mathcal{Z}}$ is a linear combination of the orthogonal basis vectors, it is itself in $\mathcal{Z}$. Further, we can compute

$$\begin{aligned} y \cdot z_j =& (x - \sum_{i=1}^{n} \frac{x \cdot z_i}{z_i \cdot z_i} z_i) \cdot z_j \\ =& x \cdot z_j - \sum_{i \neq j} \frac{x \cdot z_i}{z_i \cdot z_i} z_i \cdot z_j - x \cdot z_j \\ =& 0 \quad \forall j \end{aligned}$$

because $z_i \cdot z_j = 0$ for all $i \neq j$. Therefore, y is orthogonal to $\mathcal{Z}$ and we write $y \in \mathcal{Z}^\perp$.

We've shown that we can decompose a point x into a $\mathcal{Z}$ component plus a $\mathcal{Z}^\perp$ component. Let's now prove uniqueness. Suppose $x = x_1^{\mathcal{Z}} + y_1 = x_2^{\mathcal{Z}} + y_2$ for $z_1^{\mathcal{Z}}, z_2^{\mathcal{Z}} \in \mathcal{Z}$ and y_1, y_2 in the complement space. Applying Pythageron's theorem, we observe

$$\begin{aligned} ||y_1||^2 =& ||x_2^{\mathcal{Z}} - x_1^{\mathcal{Z}}||^2 + ||y_2||^2 \\ ||y_2||^2 =& ||x_1^{\mathcal{Z}} - x_2^{\mathcal{Z}}||^2 + ||y_1||^2 \end{aligned}$$

From these equations, and given that norms are non-negative and symmetric, it must be that $||y_1|| = ||y_2||$ which implies $||x_1^Z - x_2^Z||^2 = 0$; thus, the points must be the same and we have uniqueness. □

So, any vector can be decomposed into two components (with respect to a subspace $\mathcal{Z}$) and these components are orthogonal to one another. Further, the Hilbert space can be written as $\mathcal{H} = \mathcal{Z} + \mathcal{Z}^\perp$ where $\mathcal{Z} \cap \mathcal{Z}^\perp = \{0\}$. Thus, for a particular x, $x^{\mathcal{Z}}$ is an orthogonal projection of x onto the subspace $\mathcal{Z}$. Along with the orthogonal basis z for the space $\mathcal{Z}$, it can be represented as

$$\begin{aligned} x^{\mathcal{Z}} =& \sum_{i=1}^{n} \frac{x \cdot z_i}{z_i \cdot z_i} z_i \\ =& \sum_{i=1}^{n} \frac{E[xz_i]}{E[z_i^2]} z_i \end{aligned}$$

where the last line follows if we define the expectation operator $E[xz]$ as the inner product for the space $\mathcal{H}$. Thus, this projection is the same as the linear regression of x on the basis z_i's. This is the "predictable" part of x when being restricted to the space $\mathcal{Z}$.

Taking Stock: In Section 2, we showed a representation of the payoff pricing function $q(z)$ via state prices (i.e. $q(z) = qz$ where vector q was found from pricing Arrow securities). In Section 5 equation (3), we showed a representation of the payoff pricing function as a discounted expectation with respect to risk-neutral probabilities. The following Riesz Representation theorem provides yet another representation of $q(z)$ but also proves uniqueness and restricts the associated pricing kernel k_q to the asset span $\mathcal{M}$.

Theorem. *The Riesz Representation theorem states that if $F : \mathcal{H} \to \mathbb{R}$ is a continuous, linear functional on the Hilbert space, then there exists a unique vector $k_f \in \mathcal{H}$ such that*

$$F(x) = k_f \cdot x \quad \forall x \in \mathcal{H}. \tag{2}$$

For a space $\mathbb{R}^S$, we can find the unique kernel k_f by $k_{f,s} = F(e_s)$ for basis vectors $e_s \in \mathbb{R}^S$ and for all $s = 1, 2, ..., S$. Then, for any point x, we have $F(x) = k_f x$. Further, we have $F(x) = \sum_s \pi_s \frac{k_{f,s}}{\pi_s} x_s = E[k_f x]$ for an expectations representation and some probability measure π.

Now, within the context of asset pricing, we will consider the Hilbert space $\mathbb{R}^S$ and the asset span $\mathcal{M}$ which is a linear subspace of $\mathbb{R}^S$. The two functionals of interest in this space are the *expectations functional* and the payoff pricing functional $q(z)$. When we use probability measure π, consider this the objective probability distribution for states of the world or agent's subjective beliefs.

Definition. *The expectations functional E maps a payoff $z \in \mathcal{M}$ into its expected values $E[z]$.*[3]

Thus, the Riesz kernel associated with the expectations functional is the unique vector k_e such that $E(z) = E[k_e z]$ for all payoffs z in the asset span $\mathcal{M}$ where k_e is also in the asset span.[4]

Example. *Consider two securities with payoffs in $\mathbb{R}^3$:*

$$x_1 = (1, 1, 0)$$
$$x_2 = (0, 1, 1)$$

where we have a uniform probability of $1/3$ for each state. Thus, an expectation kernel must satisfy

$$\frac{2}{3} = E[k_e x_1]$$
$$\frac{2}{3} = E[k_e x_2].$$

The kernel lies in the asset span, so we have $k_e = h_1 x_1 + h_2 x_2$ must hold. Thus, viewing k_e as an asset, it has payoffs $(h_1, h_1 + h_2, h_2)$. This leads to a unique solution $k_e = (\frac{2}{3}, \frac{4}{3}, \frac{2}{3})$.

[3] We are assuming some underlying probability measure dictates the distribution of payoffs/states.

[4] Notice that if the risk-free payoff (of say 1) is in the asset span, then $k_e = 1$.

Definition. *For the payoff pricing functional q on the asset span, the associated Riesz kernel is k_q which is a unique payoff in $\mathcal{M}$ satisfying*

$$q(z) = E[k_q z] \quad \forall z \in \mathcal{M}. \tag{3}$$

Call k_q the pricing kernel in the asset span. Recall, if there is no arbitrage, we can derive strictly positive state prices q such that the payoff pricing function is

$$\begin{aligned} q(z) =& \sum_s q_s z_s \\ =& E[\frac{q}{\pi} z] \end{aligned}$$

for some measure π and any payoff z in the asset span. This implies (along with equation 3),

$$E[(\frac{q}{\pi} - k_q)z = 0$$

and shows that k_q is the unique orthogonal projection of $\frac{q}{\pi}$ onto the asset span $\mathcal{M}$.

Definition. *Any contingent claim $m \in \mathbb{R}^S$ that satisfies $q(z) = E[mz]$ for all $z \in \mathcal{M}$ is called a stochastic discount factor.*

Non-exhaustive examples of stochastic discount factors include the pricing kernel k_q, agent marginal rate of substitution $\frac{v_1}{v_0}$ and re-scaled state prices $\frac{q}{\pi}$. The pricing kernel differs from these in that it is unique with respect to lying in the asset span $\mathcal{M}$. Thus, when markets are incomplete, there can exist many state price vectors q and stochastic discount factors $\tilde{m}$ which will <u>all</u> map (via an orthogonal projection) to unique k_q. If markets are complete, then $k_q = \frac{q}{\pi} = \frac{v_1}{v_0}$.

Lastly, if there is a risk-free payoff in the asset span, then

$$E[k_q] = E[k_q k_e] = \frac{1}{\bar{r}}. \tag{4}$$

The first equality follows from the expectations functional being applied to k_q. The last equality is a general result from the payoff pricing functional: For risk-free payoff a, we have $q(a) = E[k_q a] \Rightarrow \frac{q(a)}{a} = E[k_q]$ where $\bar{r} = \frac{a}{q(a)}$ is a risk-free return by definition.

9 The Mean-Variance Frontier Payoffs

- A mean-variance frontier payoff lies in the span of $\{k_e, k_q\}$.
- The standard deviation of agent MRS must be greater than or equal to $\sigma(k_q)$.
- Derive the Beta pricing model.

Definition. *A payoff z is a mean-variance frontier payoff if for all other payoffs x where i) $q(z) = q(x)$ and ii) $E[z] = E[x]$ we observe $\sigma(x) > \sigma(z)$.*

For any price and mean combination, a mean-variance frontier payoff is the one with least variance. Let's define $\mathcal{E}$ as the span between the kernels k_q and k_e. It is therefore a subspace of the asset span $\mathcal{M}$, given that the kernels both lie in the asset span.

Theorem. *A payoff is a mean-variance frontier payoff iff it lies in the span of the expectations kernel and the pricing kernel.*

Proof. $(\Leftarrow)$: Consider any payoff $z \in \mathcal{M}$ in the asset span. It's orthogonal projection onto the kernel span $\mathcal{E}$ is decomposed as

$$z = z^{\mathcal{E}} + \epsilon$$

where $z^{\mathcal{E}} \in \mathcal{E}$ and $\epsilon \in \mathcal{E}^{\perp}$. Given this, we know $\epsilon \perp k_q, k_e$; that is, the residual component is orthogonal to both the pricing and expectations kernel. This further implies

$$\begin{aligned} E[k_e\epsilon] =& E[\epsilon] = 0 \\ q(\epsilon) =& E[k_q\epsilon] = 0 \end{aligned}$$

so that from the decomposition of z, we observe

$$\begin{aligned} E[z] =& E[z^{\mathcal{E}}] \\ q(z) =& q(z^{\mathcal{E}}). \end{aligned}$$

We know $cov(\epsilon, z^{\mathcal{E}}) = 0$ so that

$$\begin{aligned} var(z) =& var(z^{\mathcal{E}}) + var(\epsilon) \\ \Rightarrow& var(z) \geq var(z^{\mathcal{E}}). \end{aligned}$$

Thus, take any $z^{\mathcal{E}}$ in the kernel span $\mathcal{E}$. All other payoffs with the same expectation and price have a higher variance, given the decomposition above.

(⇒) : Do proof by contradiction. Suppose we have a payoff z such that there does not exist a z' with

$$\begin{aligned} q(z) =& q(z') \\ E[z] =& E[z'] \\ var(z) >& var(z') \end{aligned}$$

but $z \notin \mathcal{E}$. Then, z can be decomposed into $z = z^{\mathcal{E}} + \epsilon$. Through the same procedure as was done in the (⇐) direction, we can show z is not a mean-variance payoff; hence, a contradiction. □

Define *frontier returns* as frontier payoffs divided by their price. Another way of thinking about this is that frontier returns are frontier payoffs with a price of 1. The returns r_e and r_q corresponding to the return on k_e and k_q are by construction frontier returns. They are

$$r_e = \frac{k_e}{q(k_e)} = \frac{k_e}{E[k_e k_q]} = \frac{k_e}{E[k_q]} \tag{1}$$

$$r_q = \frac{k_q}{q(k_q)} = \frac{k_q}{E[k_q^2]} \tag{2}$$

where (1) comes from (4) in Section 8. We know that if a payoff is on the frontier, it is in the span of the kernels. Equivalently, if a return is on the frontier, it is a line passing through the returns (r_e, r_q) which cab be indexed by scalar λ and written

$$r_\lambda = \lambda r_e + (1-\lambda) r_q \quad \text{for} \lambda \in (-\infty, +\infty). \tag{3}$$

Re-write this as $r_\lambda = r_e + \lambda(r_q - r_e)$ and note that

$$\begin{aligned} E[r_\lambda] =& E[r_e] + \lambda(E[r_q] - E[r_e]) \\ var(r_\lambda) =& var(r_e) + \lambda^2 var(r_q - r_e) + 2\lambda cov(r_e, r_q - r_e). \end{aligned}$$

Claim. *If the risk-free rate is in the asset span, then $r_e = \bar{r}$.*

Proof. If a risk-free payoff (say WLOG 1) is in the asset span, then we know from Section 8 equation (4), $E[k_q] = \frac{1}{\bar{r}}$. Further, using the expectations functional on the risk-free payoff 1, we have

$$E[1] = E[k_e 1] \Rightarrow E[k_e] = 1.$$

$k_{e,s} = 1 \quad \forall s$ satisfies this equation, and given that the expectations kernel is unique on $\mathcal{M}$, $k_e = 1$ is the only solution. Now, use equation (1) to arrive at

$$\begin{aligned} r_e =& \frac{k_e}{E[k_e k_q]} \\ =& \frac{1}{E[k_q]} && \text{(From last result)} \\ =& \bar{r}. \end{aligned}$$

□

Thus, under the assumption of a risk-free rate $\bar{r}$ in the asset span, re-write equation (3) as

$$r_\lambda = \bar{r} + \lambda(r_q - \bar{r}) \quad \forall z \in (-\infty, \infty). \tag{4}$$

Claim. *Given the risk-free payoff is in the asset span, $\bar{r} > E[r_q]$.*

Proof. Using the variance formula, we have

$$\begin{aligned} E[k_q^2] =& E[k_q]^2 + var(k_q) \\ <& E[k_q]^2 \end{aligned} \tag{5}$$

because the variance is non-zero for the pricing kernel. Next, take expectations of equation (2),

$$\begin{aligned} E[r_q] =& E\left[\frac{k_q}{E[k_q^2]}\right] \\ <& \frac{E[k_q]}{E[k_q]^2} && \text{(From (5))} \\ =& \bar{r}. && \text{(Equation (4) Section 8)} \end{aligned}$$

□

What to take away from this? Using frontier returns with (4), and in the $\big(E[r_\lambda], \lambda\big)$ plane, the intercept is at $E[r_e]$ and the function is decreasing in λ.

Observation. *Given the risk-free payoff is in the asset span, frontier returns r_λ as in equation (4) have variance $var(r_\lambda) = \lambda^2 var(r_q)$ with standard deviation $\sigma(r_\lambda) = |\lambda|\sigma(r_q)$.*

What to take away from this? Using frontier return formula (4), and in the $\big(\lambda, \sigma(r)\big)$ plane, At $\lambda = 0$, we have $r_0 = \bar{r}$ and this is the minimum variance point with $\sigma(r_0) = 0$. As we increase λ (in negative and positive directions), we get symmetric, increasing lines from the origin. The positive direction is the "bad" direction because it not only increases variance but lowers expected returns. Thus, when we move to the classic $\big(E[r_\lambda], \sigma(r_\lambda)\big)$ representation of the mean-variance frontier, there are two frontier points with the same standard deviation, but the higher one is the *efficient*/"negative λ" one!

Thus, when looking at the relation

$$\sigma(\frac{v_1}{v_0}) \geq \sup_r \frac{|E[r] - \bar{r}|}{\bar{r}\sigma(r)}$$

from Section 7, we can show that this supremum must be a frontier return

Proposition. *If r is a frontier return, it must achieve the supremum Sharpe ratio.*

Making use of the properties that $r_q = k_q/E[k_q^2]$ and $\bar{r} = 1/E[k_q]$, we have from the HJ bounds of section 7,

$$\begin{aligned} \sigma(\frac{v_1}{v_0}) \geq& \frac{|E[r_q] - \bar{r}|}{\bar{r}\sigma(r_q)} \\ =& \frac{|\frac{E[k_q]}{E[k_q^2]} - \frac{1}{E[k_1]}|}{\sigma(k_q)} E[k_q]E[k_q^2] \\ =& \sigma(k_q). \end{aligned}$$

This result states that an agent's standard deviation of marginal rate of substitution must be larger than the standard deviation of the pricing kernel. After specifying a utility function and using consumption data, this allows us to test restrictions on the values of preference parameters.

9.1 Beta Pricing

When the risk-free return is in the asset span, it obviously has zero covariance with a return $\bar{r}$. It can be shown that for any r_γ on the frontier there exists an r_μ on the frontier that has zero-covariance with r_γ. You can use these two returns as the span for $\mathcal{E}$ instead of (r_e, r_q). Thus, if you have some return r_j, it can be projected onto the plane of frontier payoffs $\mathcal{E}$ and is decomposed as

$$r_j = r_j^{\mathcal{E}} + \epsilon_j \tag{6}$$

where $r_j^{\mathcal{E}} \in \mathcal{E}$ and $\epsilon_j \in \mathcal{E}^\perp$. Because it is a projection, ϵ_j is orthogonal to both kernels and therefore has zero price. Thus, using returns r_γ and r_μ to describe the frontier line, we can write the return in the form of (4) as

$$\begin{aligned} r_j =& r_\mu + \beta_j(r_\gamma - r_\mu) + \epsilon_j \\ \Rightarrow & E[r_j] = E[r_\mu] + \beta_j(E[r_\gamma] - E[r_\mu]). \end{aligned} \tag{7}$$

Of course, if we use the risk-free return $\bar{r} = r_\mu$, this equation is written

$$E[r_j] = \bar{r} + \beta_j(E[r_\gamma] - \bar{r}) \tag{8}$$

and taking the covariance of (7) (subbing $\bar{r} = r_\mu$) with respect to r_γ, we get

$$\begin{aligned} cov(r_j, r_\gamma) =& cov(\bar{r}, r_\gamma) + \beta_j cov(r_\gamma - \bar{r}, r_\gamma) + cov(\epsilon_j, r_\gamma) \\ =& \beta_j cov(r_\gamma - \bar{r}, r_\gamma) \\ \Rightarrow & \beta_j = \frac{cov(r_j, r_\gamma)}{var(r_\gamma)} \end{aligned} \tag{9}$$

because r_γ is uncorrelated with ϵ_j and $\bar{r}$. Therefore, we get

$$E[r_j] = \bar{r} + \beta_j(E[r_\gamma] - \bar{r}) \tag{10}$$

where the β coefficient is from the linear regression of return r_j on return r_γ. This states that the risk-premium on a security j is proportional to premium of a frontier return r_γ and is scaled by the assets covariance with the frontier return. In many circumstances, we can think of the frontier return r_γ as the return of the market r_m or some index for the overall performance of the stock market.

10 Equilibrium in a Multidate Security Markets

- Define a multidate securities economy and its competitive equilibrium.

Now, we expand our analysis to the case of an economy taking place at $T \geq 2$ dates (but still a finite number). Let there be multiple dates $t = 0, 1, ..., T$ and a securities economy with J securities. These securities pay out non-negative dividends $x_j(s_t)$ and carry the price $p_j(s_t)$ in the event s_t at date t. An agent can hold a quantity $h_j(s_t)$ of security j at date t. We call $h = (h_0, h_1, ..., h_T)$ a *portfolio strategy* where h_t is a $S \times J$ dimensional object because it specifies the holdings of each of the J securities in any of the S possible events at date t. The *gross payoff* of a portfolio at date t is $\big(p(s_t) + x(s_t)\big)h(s_{t-1})$. Further, net a *payoff* is written

$$z(h,p)(s_t) = (p(s_t) + x(s_t))h(s_{t-1}) - p(s_t)h(s_t) \quad \forall s_t, t \tag{1}$$

and is the gross payoff, less the cost of buying a new portfolio.[5] Define the *asset span* as

$$\mathcal{M}(p) = \{(z_1, ..., z_T) \in \mathbb{R}^{T\times S} : z_t = z_t(h,p) \text{ for some } h, \text{ and all } t \geq 1\}. \tag{2}$$

A couple things to note. The asset span is a function of the price vector across all future dates. Previously, the asset span only depended upon the exogenous process for payoffs in the matrix X; now, the span depends upon the vector of securities prices. Further, like the one-date model, at each date t, there are (WLOG) a set of S possible states of the world. Thus, the asset span includes all the possible states of the world at time t and for all times $t = 0, 1, ..., T$ and is thus a $S \times T$ dimensional object.

Definition. *Security markets are dynamically complete at the prices p if any consumption plan over future dates can be obtained as the payoff of a portfolio strategy (i.e. $\mathcal{M}(p) = \mathbb{R}^k$ where $k = S \times T$). If instead $\mathcal{M}(p)$ is a subset of $\mathbb{R}^K$, then we say markets are incomplete.*

An agent problem in this setup can be written

$$\begin{aligned} &\underset{c,h}{max} \quad u(c) \\ &s.t. c(s_0) = \omega(s_0) - p(s_0)h(s_0) \\ &s.t. c(s_t) = \omega(s_t) + z(h,p)(s_t) \quad \forall s_t, \forall t. \end{aligned} \tag{3}$$

If we attach a multiplier λ to the budget constraint, we have the first-order condition

$$\lambda_t p_t = \sum_{s_{t+1}|s_t} [p_{t+1} + x_{t+1}]\lambda_{t+1} \quad \forall s_t \tag{4}$$

where I've suppresed the s notation. Further, we observe $\lambda_t = u_c(t)$; thus, combine conditions to arrive at

$$p_t = \sum_{s_{t+1}|s_t} [p_{t+1} + x_{t+1}] \frac{u_c(t+1)}{u_c(t)}. \tag{5}$$

[5] When we write $z_t(h,p)$, we are referring to the S-vector of net payoffs for any realization s_t at date t. Further, $z(h,p)$ is an $S \times T$ object. Also, keep track of the dimensionality of the price vector p which has the price of J assets in T time periods over S possible states of the world for each t.

Definition. *An equilibrium in multidate security markets consists of security prices* p*, a set of portfolio strategies* $\{h^i\}$*, a consumption plan* $\{c^i\}$ *for agents* $i = 1, 2, ..., I$ *such that, given prices,*

1. (c^i, h^i) *solve the agent's problem*

2. *Markets clear*

   a. $\sum_i h^i(s_t) = 0 \quad \forall s_t, t$

   b. $\sum_i c^i(s_T) = \sum_i \omega^i(s_t) \quad \forall s_t, t$

11 Multidate Arbitrage and Positivity

- Define a payoff pricing functional $q(z)$ and notion of arbitrage in multidate security markets.
- Given the Law of One Price, $q(z)$ is single-valued and linear.
- $q(z)$ is strictly positive if and only if there doesn't exist arbitrage.

Definition. *The Law of One Price states that*

$$\text{if } z(h,p) = z(h',p) \Rightarrow p_0 h_0 = p_0 h'_0.$$

This states that, at given prices p, if two portfolios h and $\tilde{h}$ guarantee the same payoff $z \in \mathbb{R}^S$, both portfolios must have the same date 0 price ($p_0 h_0$) where $p_0, h_0 \in \mathbb{R}^J$.

Proposition. *The Law of One Price holds iff for every portfolio h with $z(h,p) = 0$, we observe $p_0 h_0 = 0$.*

Proof. This is an analogue theorem and proof for the two-date model in which the law of one price holds if and only if a strategy h with payoff vector $z = 0$ has price $ph = 0$. □

Define the *payoff pricing functional* $q : \mathcal{M} \to \mathbb{R}$ as

$$q(z) = \{w : w = p_0 h_0 \text{ for all } h \text{ with } z = z(h,p)\} \quad \forall z \in \mathcal{M}. \tag{1}$$

So, given a payoff $z \in \mathbb{R}^{S \times T}$, $q(\cdot)$ will provide the prices of all portfolios that possibly could generate this payoff. "Prices" because q could potentially be a correspondence. Given the Law of One Price, the payoff pricing functional is single-valued and hence linear on the asset span $\mathcal{M}$. Take note: if you buy a security at date 0 and hold it for all time, it generates the payoffs $x_j(s_t)$; therefore, the dividend stream is in the asset span and is therefore priced by q. In particular, for $x_j \in \mathcal{M}$, $q(x_j) = p_{j0}$.

Definition. *An arbitrage is a portfolio h that has $z(h,p) \geq 0$ and $p_0 h_0 \leq 0$ with one or both inequalities strict.*

Theorem. *The payoff pricing functional q is strictly positive iff there does not exist arbitrage.*

Proof. No arbitrage here means $p_0 h_0 > 0$ whenever $z(h,p) > 0$. Further, note that pricing the payoff reveals $q(z(h,p)) = p_0 h_0$. Thus, no arbitrage precisely corresponds to $q's$ being a stritly positive vector with respect to the asset span $\mathcal{M}(p)$. □

Definition. *We can define a one-period arbitrage at event s_t as a date-t portfolio $h(s_t)$ that has $p(s_t)h(s_t) \leq 0$ and*

$$[p(s_{t+1}) + x(s_{t+1})]h(s_t) \geq 0 \quad \forall s_{t+1}.$$

with one or both inequalities strict.

When the payoff pricing functional we are talking about is associated with equilibrium security prices p^*, we call it an *equilibrium payoff pricing functional.*

Theorem. *If agents' utility functions are strictly increasing, then there is no arbitrage at equilibrium security prices. Further, the equilibrium payoff pricing functional q is strictly positive.*

12 Multidate Dynamically Complete Markets

- Markets are dynamically complete if and only if the number of (non-redundant) securities is equal to the number of future states s_{t+1}, for any preceding s_t.
- Event prices are strictly positive if and only if there does not exist arbitrage.
- Derive an equation relating security prices p with event prices q.
- Define and construct the valuation functional $Q(z)$ for payoffs $z \in \mathbb{R}^k$.

Even though there may not be as many secruties as there are events and dates, the ability to trade at future dates allows the agents an ability to rebalance their portfolio and complete markets. This will be shown below. Another way in which markets can be completed would be if agents are able to trade Arrow securities for all future events s_t. These securities pay out a dividend of 1 only in their corresponding event s_t. The intuition is simple, you can achieve any payoff you want by scaling up or down the quantities of Arrow dividends you purchase; therefore, you can span the space $\mathbb{R}^k$ completely.[6]

Definition. *Define the one-period payoff matrix at event s_t as a $J \times \kappa(s_t)$ with entries $p_j(s_{t+1}) + x_j(s_{t+1})$ for all rows $j = 1, 2, ..., J$. $\kappa(s_t)$ is the number of successors s_{t+1} that can occur, given that s_t has been realized.*

So, each row of the payoff matrix represents a security, and each column contains the gross payoff of holding one unit of that security.

Theorem. *Markets are dynamically complete iff the one-period payoff matrix in each nonterminal event s_t is of rank $\kappa(s_t)$.*

Proof. For any event s_t, if $J < \kappa(s_t)$, then there are an insufficient number of assets to span the space $\mathbb{R}^{\kappa(s_t)}$. If $J > \kappa(s_t)$, then there exists a basis of $\kappa(s_t)$ linearly independent rows that can replicate the remaining $J - \kappa(s_t)$; thus, WLOG consider a square matrix. Given the columns are linearly independent, the matrix X has an inverse and $h(s_t) = X^{-1}(s_t)z$ maps back to a solution for the portfolio h needed to reach a payoff point z. □

If markets are dynamically complete and the Law of One price holds, then we have a unique linear payoff pricing functional q on the asset span $\mathcal{M}(p)$ which is simply the space $\mathbb{R}^k$. To define the values of the pricing functional, we use Arrow securities. In the context of the multidate model, an Arrow security with dividend in period s_t is represented by a S vector with a 1 in the s_T position and zero everywhere else. Then, define $q(s_t) = q(e(s_t))$ as the *event price* associated with an Arrow security paying out 1 in state s_t. Because every payoff $z \in \mathbb{R}^k$ is simply scaled Arrow securities (scaled by $z * 1$), we can price it using event prices:

$$q(z) = q(\sum_{s \in S} z(s)e(s)) = \sum_{s \in S} q(s)z(s) \tag{1}$$

where S is all the events and times that can occur. We can equivalently write (with an inner product notation)

[6] In the multidate market setup, to avoid continually using notation $S \times T$ for the asset span, I just use k. So does Jan the Man.

this as $q(z) = qz$ where q is a vector of event prices.[7] Thus, event prices are strictly positive iff the payoff pricing functional is strictly positive.

Theorem. *Event prices in a multidate security economy are strictly positive iff there does not exist an arbitrage.*

Proof. This proof method follows from Section 3's two-date proof and can be extended to the multidate economy. □

Proposition. *Event prices satisfy*

$$q(s_t)p_j(s_t) = \sum_{s_{t+1}|s_t} q(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] \tag{2}$$

for every event s_t *and all* $t \geq 0$.

Proof. Consider buying security j at event s_t and selling it at time $t+1$ for all possible s_{t+1}. Call this strategy $\hat{h}$. Thus, we have

$$\begin{aligned} z(\hat{h}, p)(s_t) &= -\, p_j(s_t) \\ z(\hat{h}, p)(s_{t+1}) &= p_j(s_{t+1}) + x_j(s_{t+1}) \quad \forall s_{t+1} \end{aligned}$$

and lastly $z(\hat{h}, p)(\xi_t) = 0$ for all other possible events and times.

For this payoff $z(\hat{h}, p)$, the payoff pricing functional gives us $q(z(\hat{h}, p)) = p_0\hat{h}_0 = 0$. Then, apply event prices in form $q(z) = qz = 0$ to obtain

$$0 = \sum_{s_{t+1}|s_t} q(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] - q(s_t)p_j(s_t).$$

and we obtain the identity. □

Now, consider the agent problem again

$$\begin{aligned} \max_{c,h} \quad & u(c) \\ s.t. \quad & c_0 = \omega_0 - p_0 h_0 \\ s.t. \quad & c_t = \omega_t + z_t(h, p) \quad \forall t \geq 1. \end{aligned} \tag{3}$$

Note: the date 0 value of the portfolio h must be $p_0 h_0 = q(c_1^+ - \omega_1^+)$ where the + notation represents all future dates $t \geq 1$; this states that the strategy which generates all future consumption less future endowments must be the price of the date 0 portfolio. The second constraint can be written as $c_{1+} - \omega_{1+} \in \mathcal{M}(p)$, which just says that the stream of consumption less endowments must be an allocation/payoff in the asset span.

[7] In the two-date model, we had a single-valued, unique payoff pricing functional and used Arrow securities to recover the vector q of *state prices*. Here, we do the same thing in a larger state-space, but now we are calling the corresponding vector q as containing all the *event prices*. I think this is just to distinguish between the two models. They are conceptually and mathematically the same thing.

If markets are complete, this constraint is trivially satisfied. Thus, (in complete markets) we can re-cast this problem with just the first constraint as

$$\begin{aligned} \max_{c} \quad & u(c) \\ s.t. \quad & c_0 + qc_1^+ = \omega_0 + q\omega_1^+ \end{aligned} \tag{4}$$

because the payoff pricing functional is linear. Optimization of (4) yields

$$q(s_t) = \frac{u_c(s_t)}{u_c(s_0)} \quad \forall s_t, t, \tag{5}$$

given we normalize $q(s_0) = 1$. Under market completeness, we have an equivalence between solving problem (3) by choosing $\{c, h\}$ with respect to equilibrium prices p and solving problem (4) by choosing $\{c\}$ with respect to event prices q, which are derived from the system in (2). The problem in (4) is a date-0, or contingent claims problem, whereas the problem in (3) is for a security market equilibrium. With no arbitrage and completeness, both problems solve for the same $\{c^*\}$; lacking a complete market, there can exist multiple event prices satisfying (2), leading to a divergence between $\{c^*\}$ in the security market and $\{c^q\}$ in problem (4).

Theorem. *If security markets are dynamically complete under equilibrium security prices and agents' utility functions are strictly increasing, then every equilibrium consumption allocation is Pareto optimal.*

Proof. Utilize the same proof method as you would for the 2-date model. Posit an equilibrium allocation (c^*, h^*) and suppose there exists another allocation $(\tilde{c}, \tilde{h})$ which Pareto dominates it. Through strictly increasing utility, such a allocation can be proven to be infeasible. □

Often, it may be the case that the payoff pricing functional q is a proper subset of the contingent commodity space $\mathbb{R}^k$. In this case, as before, we define the *valuation functional* $Q : \mathbb{R}^k \to \mathbb{R}$ where

$$Q(z) = q(z) \quad \text{for every } z \in \mathcal{M}(p).$$

Theorem. *Fundamental Theorem of Finance (for Multidate markets). Security prices exlcude arbitrage iff there exists a strictly positive valuation functional.*

Proof. Use same proof method from Section 4. In this version, we have

$$\begin{aligned} q_u(z) =& \min_{h}\{p_0 h_0 : z(h, p) \geq z\} \\ q_l(z) =& \max_{h}\{p_0 h_0 : z(h, p) \leq z\} \end{aligned}$$

for anypayoff z. Then construct $Q(\cdot)$ by identifying a payoff z^* not in $\mathcal{M}$ and so forth. □

Theorem. *Suppose that security prices exclude arbitrage. Then security markets are dynamically complete iff there exists a unique strictly positive valuation functional.*

Absence of arbitrage guarantees a strictly positive payoff and valuation pricing functional. Further, when we assume dynamically complete markets, which implies that $q(\cdot)$ has $\mathbb{R}^k$ as a domain. Therefore, given that q is unique on its domain, the payoff pricing functional is equivalent to the valuation functional on the contingent commodity space.

13 Event Prices, Risk-Neutral Probabilities and the Pricing Kernel in Multidate Markets

- Define the risk-neutral measure π^* and show its properties.
- Define the pricing kernel k_q and show its relation to event prices q.

When markets are dynamically complete, the payoff pricing functional q is defined by its pricing of Arrow securities $e(s) \in \mathbb{R}^k$. When markets are incomplete, some Arrow securities cannot be priced by q. Despite this, given the absence of arbitrage, the payoff pricing functional can be extended to a strictly positive valuation functional Q which can price all Arrow securities. Thus,

$$q(s) = Q(e(s)) \tag{1}$$

for every event $e(s)$ is a unit vector in $\mathbb{R}^k$ with unit payoff in only the event s. We call $q(s)$ the *event price*. Thus, for any payoff z, we have $Q(z) = \sum_{s_t \in z} q(s_t)z(s_t) = \sum_{t=1}^{T} \sum_{s_t \in F_t} q(s_t)z(s_t)$ where F_t is the set of events which relate to the payoff z. It was shown in the previous section that this implies that the event prices are the solution to the linear system of equations

$$q(s_t)p_j(s_t) = \sum_{s_{t+1}|s_t} q(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] \tag{2}$$

for all t and all s_t. There exists a strictly positive valuation functional iff there exists a strictly positive solution to the above system. Further, if markets are incomplete, there are many valuation functionals and therefore many solutions to the above system.

Using the event price representation of the valuation functional, if we were to buy a security j at event s_t and hold until the end of the economy, we would observe

$$p_j(s_t) = \frac{1}{q(s_t)} \sum_{\tau=t+1}^{T} \sum_{s_{t+1}|s_t} q(s_\tau x_j(s_\tau) \tag{3}$$

The *one-period return* on a security j in event s_{t+1} is

$$r_j(s_{t+1}) = \frac{p_j(s_{t+1}) + x_j(s_{t+1})}{p_j(s_t)} \tag{4}$$

which we write as $r_{j,t+1}$ as shorthand. Further, if a return does not depend upon a future realization, it is considered *risk free* and written as $\bar{r}(s_{t+1})$ in all events s_{t+1}, where its value is known at time t.

Definition. *If at all dates and in every possible event there exists a a strictly positive risk-free, one-period return, then define the discount factor $\rho(s_t)$ as*

$$\rho(s_t) = \prod_{\tau=1}^{t} [\bar{r}(s_\tau)]^{-1} \quad t = 1, 2, ..., T. \tag{5}$$

The discount factor is simply the inverse of the product of risk-free returns from time zero to now (date t). This is a $\mathcal{F}_{t-1}$-measurable function (i.e. we know the value of the risk-free return at date t and therefore know the value of the discount factor at date t). This is just to say that we know what the discount factor will be next period.

Define a *risk-neutral probability* at date T as

$$\pi^*(s_T) = \frac{q(s_T)}{\rho(s_T)}$$

and for all other dates as

$$\pi^*(s_t) = \sum_{\text{All possible } S_T} \pi^*(s_T).$$

This notation is confusing. Basically, the probability of event s_t is the sum of all the probabilities s_T that can be reached, given that s_t has occurred. Given this, the risk-neutral probability of event s_t satisfies

$$\pi^*(s_t) = \frac{q(s_t)}{\rho(s_t)} \tag{6}$$

so it is the event prices, scaled by the corresponding discount factor. Knowing this, start with the event price representation (2) to get an expression in terms of the risk-neutral expectation:

$$\begin{aligned} p_j(s_t) =& \frac{1}{q(s_t)} \sum_{s_{t+1}|s_t} q(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] \\ =& \frac{1}{q(s_t)} \rho(s_{t+1}) \sum_{s_{t+1}|s_t} \frac{q(s_{t+1})}{\rho(s_{t+1})} [p_j(s_{t+1}) + x_j(s_{t+1})] \\ =& \frac{1}{\pi^*(s_t)} \frac{1}{\bar{r}(s_{t+1})} \sum_{s_{t+1}|s_t} \pi^*(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] \\ =& \frac{1}{\bar{r}(s_{t+1})} \sum_{s_{t+1}|s_t} \pi^*(s_{t+1}|s_t)[p_j(s_{t+1}) + x_j(s_{t+1})] \\ =& \frac{1}{\bar{r}(s_{t+1})} E_{t,*}[p_j(s_{t+1}) + x_j(s_{t+1})]. \end{aligned} \tag{7}$$

where $\pi^*(s_{t+1}|s_t) = \frac{\pi^*(s_{t+1})}{\pi^*(s_t)}$. Thus, the price of a security j is its discounted, expected price plus dividend payout tomorrow, with respect to the risk-neutral probability measure. From flipping the position of $p_j(s_t)$ and $\bar{r}(s_{t+1})$, we get the expression

$$E_{t,*}[r_{j,t+1}] = \bar{r}(s_{t+1}) \tag{8}$$

such that the risk-netrual expected return of a risky security is equal to the risk-free return available at that date. If we substitute the risk-neutral probabilities π^* into the valuation equation $Q(z) = \sum_t \sum_{s_t} q(s_t)z(s_t)$, we observe

$$Q(z) = \sum_{t=1}^{T} E_*[\rho_t z_t] \quad \forall z_t \in \mathbb{R}^k \tag{9}$$

and further can derive

$$q(z) = \sum_{t=1}^{T} E_*[\rho_t z_t] \quad \forall z_t \in \mathcal{M}(p) \tag{10}$$

$$p_{j,0} = \sum_{t=1}^{T} E_*[\rho_t z_{jt}]. \tag{11}$$

where this last equality illustrates that the date-0 price of a security is equal to the expected, discounted sum of future dividend streams, with respect to the risk-neutral measure.

Now, we can also talk about the payoff pricing kernel in multidate markets which serves an analogous role to the two-period model (as in Section 8).

Definition. *The pricing kernel k_q is a payoff in $\mathcal{M}(p)$ such that*

$$q(z) = \sum_{t=1}^{T} E[k_{qt} z_t] \tag{12}$$

where expectation E is taken with respect to the natural probability measure π.

This value is unique and in the asset span. The multidate extension here is that the kernel is a $T \times S$ dimensional object. Thus, we can think of several ways in which to represent the price of a portfolio attaining some payoff z: one with natural expectation, using the pricing kernel, and another way with risk-netrual expectations. Observe

$$\begin{aligned} q(z) &= \sum_{t=1}^{T} E_*[\rho_t z_t] \\ &= \sum_{t=1}^{T} E[k_{qt} z_t] \\ &= \sum_{t=1}^{T} \sum_{s_t} \pi(s_t) k_q(s_t) z(s_t) \end{aligned}$$

where in a two-period model, this would just be written as $q(z) = \frac{1}{\bar{r}} E_*[z] = E[k_q z]$.

Proposition. *The pricing kernel can be used (by definition) in the following relations*

$$k_{qt} = E[k_{q,t+1} r_{j,t+1}]$$

and if the security in question is risk-free, we get

$$k_{qt} = \bar{r}_{t+1} E_t[k_{q,t+1}].$$

Proof. Just as we derived the event prices system of equations, consider a strategy of only buying security j in event s_t and selling it at date $t+1$ for any event s_{t+1}. We will get the following equality

$$k_{qt} p_j(s_t) = \sum_{s_{t+1}|s_t} \pi(s_{t+1}|s_t) k_q(s_{t+1}) [p_j(s_{t+1}) + x_j(s_{t+1})]$$

which can then be written as

$$k_{qt} = E[k_{q,t+1} r_{j,t+1} | s_t].$$

If the security in question is risk-free, then it can be pulled out of the operator for the last equality of the proposition. □

Proposition. *In a dynamically complete market, the pricing kernel is given by*

$$k_q(s_t) = \frac{q(s_t)}{\pi(s_t)}.$$

Same as section 8. In complete markets, q is unique on $\mathbb{R}^k$ and $\sum_s \pi_s \frac{q_s}{\pi_s} z_s = E[k_q z]$ is the unique solution for the kernel on the span $\mathcal{M} = R^k$.

14 Martingale Property of Gains

- Define the *gain* G_t and *discounted gain* $\hat{G}_t$.
- Discounted gains are martingales with respect to the risk-neutral measure.

Definition. *A sequence of random variables* $\{y_t\}_{t=0}^{T}$*, which are* $\mathcal{F}_t$*-measurable with respect to the filtration* $\mathcal{F}_t$*, are called a martingale with respect to the probability measure* π *if*

$$E_t[y_\tau] = y_t \quad \forall \tau \geq t$$

Definition. *The gain at date* $t \geq 1$ *on a strategy* h *is*

$$G(s_t) = p(s_t)h(s_t) + \frac{1}{\rho_t}\sum_{\tau=1}^{t} \rho_\tau z_\tau(h,p)(s_\tau). \tag{1}$$

What does this measure? The gain gives you the date-t value of a portfolio strategy h plus the sum of net payoffs of that strategy from all previous dates, where those net payoffs are re-invested at the risk-free rate. If we multiplied the definition through by ρ_t, then we would have the *discounted gain*

$$\hat{G}(s_t) = \rho_t p(s_t)h(s_t) + \sum_{\tau=1}^{t} \rho_\tau z_\tau(h,p)(s_\tau) \tag{2}$$

which is now a date-0 measure of a portfolio strategy. If we looked at the gain on holding a single security, it would be simply $\rho_t p_j(s_t) + \sum_{\tau=1}^{t} \rho_\tau x_j(s_\tau)$.

Theorem. *The discounted gain on any portfolio strategy is a martingale under the risk-neutral probability measure. That is*

$$E_{*,t}[\hat{G}_\tau] = \hat{G}_t \quad \forall \tau \geq t.$$

Further, given pricing kernel k_{qt}*, we have*

$$E_t[k_{q\tau}G_\tau] = k_{q\tau}G_t$$

with respect to the natural probability measure π *and note this is for the gain, not discounted gain.*

Proof. I just prove the first part. We can show

$$\hat{G}_{t+1} - \hat{G}_t = \rho_{t+1}[(p_{t+1} + x_{t+1})h_t] - \rho_t p_t h_t.$$

Take risk-neutral conditional expectations on both sides of this equation, and recall from $p_{jt} = \frac{1}{\bar{r}_{t+1}} E_{*,t}[p_{j,t+1} + x_{j,t+1}]$ we get

$$\rho_t p_t h_t = \rho_{t+1} E_{*,t}[(p_{t+1} + x_{t+1})h_t]$$

which is the desired result. □

15 Conditional Consumption-Based Security Pricing in Multidate Market

We have expression

$$\bar{r}_{t+1} = \frac{1}{\delta} \frac{v'(c_t)}{E_t[v'(c_{t+1})]}.$$

We eventually get to security-pricing equation

$$E_t[r_{j,t+1}] = \bar{r}_{t+1} - \delta \bar{r}_{t+1} \frac{cov_t(v'(c_{t+1}), r_{j,t+1})}{v'(c_t)}$$

using the same methods as in Section 7.

16 Equilibrium in Infinite-Time Security Markets

- Define the agent's infinite-time problem under debt constraints and equilibrium condition.
- Define an infinite-time equilibrium under debt constraints.

Many securities (such as publicly-traded stocks) are thought of as, or priced as, an infinite stream of potential dividends. Thus, in moving towards an infinite time economy, we potentially step toward a closer approximation to reality. This new setup is not innocuous: new problems, such as the infinite rolling-over of debt, have to be addressed. Now, formalize an agent's preferences as

$$u(c) = \sum_{t=0}^{\infty} \delta^t E[v(s_t)] \tag{1}$$

for positive levels of consumption, a discount factor $\delta \in (0,1)$ and a strictly increasing, continuous utility function v. Further, each agent has a consumption endowment process, specified as $\omega^i = (\omega_0^i, \omega_1^i, ...) \in \mathbb{R}_+^\infty$ for each agent i. These sum to aggregate consumption endowments $\bar{\omega} = \sum_i \omega_i$. In addition, agent's are endowed with an initial portfolio endowment $\hat{h}_0^i \in \mathbb{R}^J$ and there are no more portfolio endowments at future dates. These initial endowments aggregate to $\bar{h}_0^i = \sum_i \hat{h}_0^i > 0$ so there is a positive supply. Lastly, we specify the *effective consumption endowment* of an agent as $\hat{\omega}^i(s_t) = \omega^i(s_t) + \hat{h}_0^i x(s_t)$ as the sum of the agent's consumption endowment and dividends on portfolio endowment, in event s_t.

Agent budget constraints are

$$\begin{aligned} c(s_0) + p(s_0)h(s_0) =& \omega^i(s_0) + p(s_0)\hat{h}_0^i \\ c(s_t) + p(s_t)h(s_t) =& \omega^i(s_t) + [p(s_t) + x(s_t)]h(s_{t-1}) \quad \forall s_t \end{aligned} \tag{2}$$

Definition. *A debt constraint, with respect to threshold $D(s_{t+1})$ is*

$$[p(s_{t+1}) + x(s_{t+1})]h(s)t) \geq -D(s_{t+1} \quad \forall s_{t+1} \tag{3}$$

and a borrowing constraint, with respect to threshold $B(s_t)$, is

$$p(s_t)h(s_t) \geq -B(s_t) \tag{4}$$

Note that there can potentially be a vector of debt constraints for each state of the world tomorrow at s_{t+1}. Both constraints are with respect to date t portfolio decisions, but differ in that debt constraints restrict the agent according to outcomes tomorrow whereas borrowing constraints are a limit on quantities today.

The first-order conditions of the problem with an attached debt constraint leads to the pricing equation

$$p(s_t) = \sum_{s_{t+1}} [p(s_{t+1}) + x(s_{t+1})][\frac{u_c(t+1)}{u_c(t)} + \frac{\mu(s_{t+1})}{u_c(t)}] \tag{5}$$

where μ is the Lagrange multiplier associated with the debt constraint. When the debt constraint is not binding, we get the normal condition. When the debt constraint is binding, it creates a wedge in the pricing equation.

Definition. *An equilibrium under debt constraints is a price process p, consumption/portfolio allocation* $\{c^i, h^i\}_{i=1}^I$ *such that the allocations are solutions to the agent problem, subject to budget and debt constraints and the transversality condition. Further, markets clear:*

1. $\sum_i h^i(s_t) = \bar{h}_0$

2. $\sum_i c^i(s_t) = \bar{\omega}(s_t) + x(s_t)\bar{h}_0 \quad \forall s_t, t.$

17 Arbitrage, Valuation and Price Bubbles in Infinite Time

- Define arbitrage under debt and borrowing constraints.
- These forms of arbitrage are equivalent to one-period arbitrage.
- No arbitrage under debt/borrowing restrictions is equivalent to strictly positive event prices.
- Define a price bubble σ. Under mild(?) conditions, price bubbles are zero in equilibrium.

Definition. *An arbitrage under debt constraints is a portfolio strategy h such that*

$$p_0 h_0 \leq 0 \quad \textit{and} \quad z(h,p)(s_t) \geq 0 \quad \forall s_t, t$$

with one or two strict inequalities, and

$$[p(s_{t+1}) + x(s_{t+1})]h(s_t) \geq 0 \quad \forall s_{t+1}.$$

This is a common definition for arbitrage with the added restriction that the agent cannot have positive debt tomorrow (s_{t+1}). Thus, an arbitrage under debt constraints is a trading strategy that i) costs nothing today, ii) guarantees a positive future in future states and iii) does not require a negative gross payoff in any event. Another way to think about the third condition is that an arbitrage is a strategy h that an agent can technically add to an existing strategy $\hat{h}$ without violating constraints. So, an arbitrage under debt constraints must be able to be added to any strategy without violating any possible debt constraints; the one that satisfies this is the one that does not incur positive levels of debt for all s_t.

Definition. *A one-period arbitrage is a event s_t strategy $h(s_t)$ such that*

$$p(s_t)h(s_t) \leq 0 \quad \textit{and} \quad [p(s_{t+1}) + x(s_{t+1})]h(s_t) \geq 0 \quad \forall s_{t+1}$$

with at least one strict inequality.

Proposition. *Security prices exclude arbitrage under debt constraints iff they exclude one-period arbitrage in every event.*

Proof. ($\Rightarrow$): Assume $\nexists h$ for an arbitrage under debt constraints. Looking at the definitions above, we see that a one-period arbitrage is an arbitrage under debt constraints.

($\Leftarrow$): Assume $\nexists h(s_t)$, a one-period arbitrage for all events s_t. Suppose, for contradiction $\exists$ an arbitrage under debt constraints h. This implies

$$p_0 h_0 \leq 0 \quad \text{and} \quad [p(s_1) + x(s_1)]h_0 \geq 0 \quad \forall s_1$$

by definition of an arbitrage under debt constraints. Given that there is no one-period arbitrage, this implies $p_0 h_0 = [p(s_1) + x(s_1)]h_0 = 0$ for all s_1. Because an arbitrage under debt constraints requires

$$z(h,p)(s_1) \geq 0 \Rightarrow p(s_1)h(s_1) \leq 0.$$

Moving forward in time, we also have (by definition) that $[p(s_2) + x(s_2)]h(s_1) \geq 0$ for all s_2 and the assumption of no one-period arbitrage again imposes $p(s_1)h(s_1) = [p(s_2) + x(s_2)]h(s_1) = 0$. All future time periods follow, inductively. In summary, we have a strategy h such that $p_0 h_0 = 0$ and $z(h,p)(s_t) = 0$ for all s_t, which is not an arbitrage under debt constraint. A contradiction. □

Definition. *Event prices in infinite time security markets are defined as a sequence $q \in \mathbb{R}^\infty$ satisfying the equations*

$$q(s_t)p_j(s_t) = \sum_{s_{t+1}|s_t} q(s_{t+1})[p_j(s_{t+1}) + x_j(s_{t+1})] \tag{1}$$

for all events s_t and securities $j = 1, 2, ..., J$.

Theorem. *Security prices exclude arbitrage under debt constraints iff there exist strictly positive event prices.*

Proof. Recall Stiemke's Lemma.[8] Stiemke's Lemma can be applied to the two date model and shows that if there does not exist a one-period arbitrage strategy $h(s_t)$, then equation (1) has a strictly positive solution $q \in \mathbb{R}^S$.

From our previous proposition, we know there is an equivalence between one-period arbitrage and arbitrage under debt constraints. Therefore, because the event s_t is arbitrary and holds for all t, we have a strictly positive solution in (1) for $q \in \mathbb{R}^{S \times T}$. □

Definition. *An arbitrage under borrowing constraints is a strategy h such that*

$$p_0 h_0 \leq 0 \quad \textit{and} \quad z(h,p)(s_t) \geq 0 \quad \forall s_t$$

with at least one strict inequality and

$$p(s_t)h(s_t) \geq 0 \quad \forall s_t.$$

Thus, this definition of arbitrage differs from the debt constraint one in the sense that it restricts the values of a portfolio today (what you borrow) versus the values of its realized return, tomorrow (what you owe).

[8] Stiemke's Lemma states that $\nexists a \in \mathbb{R}^m$ such that

$$aY \geq 0 \quad \text{and} \quad ay \leq 0$$

with at least one strict inequality <u>if and only if</u> $\exists b \in \mathbb{R}^n$ such that $y = Yb$ and $b >> 0$. For the two-date model, this can be written as $\nexists h \in \mathbb{R}^J$ such that

$$hX \geq 0 \quad \text{and} \quad hp \leq 0$$

with at least one strict inequality <u>if and only if</u> $\exists q \in \mathbb{R}^S$ such that $p = Xq$ and $q >> 0$.

Proposition. *Security prices exclude arbitrage under borrowing constraints if and only if they exclude one-period arbitrage.*

Proof. (⇐): Assume there does not exist one-period arbitrage. Suppose, for a contradiction, that there exists an arbitrage under borrowing constraints. This means there exists an h such that we observe

$$p_0 h_0, \leq 0 \quad p_0 h_0 \geq 0, \quad p(s_1)h(s_1) \geq 0, \quad [p(s_1) + x(s_1)]h_0 - p(s_1)h(s_1) \geq 0 \quad \forall s_1$$

with at least one strict inequality on date-0 price and date-1 payoffs. given the first two conditions, we have $p_0 h_0 = 0$, which implies there must exist some s_1' such that

$$\begin{aligned} [p(s_1')+&x(s_1')]h_0 - p(s_1')h(s_1') > 0 \\ \Rightarrow&[p(s_1') + x(s_1')]h(s_1') > 0 \qquad \text{(by } p(s_1')h(s_1') \geq 0\text{)} \end{aligned}$$

which along with $p_0 h_0 = 0$ constitutes a one-period arbitrage. A contradiction.

(⇒): Now, assume there does not exist arbitrage under borrowing constraints, but suppose, for a contradiction, that there exists one-period arbitrage. This implies there exists an $h(s_t)$ such that

$$p(s_t)h(s_t) \leq 0 \quad \text{and} \quad [p(s_{t+1} + x(s_{t+1})]h(s_t) \geq 0 \forall s_{t+1}$$

with at least one strict inequality. Now, consider a portfolio strategy $\tilde{h}$ defined as

$$\tilde{h}(s) = \begin{cases} h(s_t), & \text{when } s = s_t \\ 0, & \text{otherwise} \end{cases}$$

which implies $p_0 \tilde{h}_0 = 0$. Further, look at the payoff profile

$$z(\tilde{h}, p)(s) = \begin{cases} -p(s_t)h(s_t) \geq 0, & \text{when } s = s_t \\ [p(s_{t+1}) + x(s_{t+1})]h(s_t) \geq 0, & \text{for all } s = s_{t+1} \\ [p(s_{t+1}) + x(s_{t+1})]\tilde{h}(s_t) = 0, & \text{for all } s \neq s_t, s_{t+1} \end{cases}$$

where one of the inequalities will be strict depending on the one-period arbitrage. Notice: the first two payoffs are defined in terms of the one-period strategy h. Thus, we have a strategy $\tilde{h}$ with $p_0 \tilde{h}_0 = 0$ and $z(\tilde{h}, p)(s_t) \geq 0$ for all s_t with one strict inequality guaranteed, which is an arbitrage under borrowing constraints; thus, a contradiction. □

We can now express an analogous result for this form of arbitrage.

Theorem. *Security price exclude arbitrage under borrowing constraints iff there exist strictly positive event prices.*

Proof. Once again, given we know that an arbitrage under borrowing constraints is equivalent to a one-period arbitrage, we can use Stiemke's lemma and have necessity and sufficiency. □

Moving forward, many of the results in the multi-date security market setup more or less extend seamlessly to the infinite time security market. For markets to be dynamically complete in infinite time, we require that the number J of non-redundant securities is greater than or equal to the number of immediate successors of all events s_t. If markets are dynamically complete and security prices do not permit arbitrage under debt constraints, then there exists a unique solution of strictly positive event prices, given security prices p.

Just as before, using event prices q (which depend upon security prices p), the present value of a dividend stream for security j in event s_t is

$$\frac{1}{q(s_t)} \sum_{\tau=t+1}^{\infty} \sum_{s_\tau | s_{\tau-1}} q(s_\tau) x_j(s_\tau). \tag{2}$$

Note: if markets are not complete, there may exist multiple event prices, which means the present value of an asset depends upon which event price is used. It is a result that if a security has zero dividends after a date t, its present price is equal to its present value of dividend streams; that is,

$$p_j(s_t) = \frac{1}{q(s_t)} \sum_{\tau=t+1}^{T} \sum_{s_\tau | s_{\tau-1}} q(s_\tau) x_j(s_\tau). \tag{3}$$

Now, what happens to the present value calculation when a security admits positive, non-zero payoffs in perpetuity? Is there a divergence between the present value of dividends and the price of the security?

Definition. *Define a price bubble on security j after event s_t as*

$$\sigma_j(s_t) = p_j(s_t) - \frac{1}{q(s_t)} \sum_{\tau=t+1}^{\infty} \sum_{s_\tau | s_{\tau-1}} q(s_\tau) x_j(s_\tau). \tag{4}$$

Thus, a price bubble is exactly the difference between the current price of security j and its present value of dividends, as calculated by event prices.

Proposition. *It follows that*

$$\sigma_j(s_t) = \underset{T\to\infty}{lim} \sum_{S_t | s_t} q(s_T) p_j(s_T)$$

and

$$q(s_t)\sigma_j(s_t) = \sum_{s_{t+1}|s_t} q(s_{t+1})\sigma_j(s_{t+1}) \quad \forall j \forall s_t. \tag{5}$$

Definition. *Agents exhibit uniform impatience with respect to the effective aggregate endowment $\hat{\omega}$ if there exists $\gamma \in [0,1]$ such that*

$$u^i\big(c^i_-(s_t), c^i(s_t) + \hat{\omega}(s_t), \gamma c^i_+(s_t)\big) > u^i(c^i) \tag{6}$$

for every i, every s_t and every $c^i \in [0, \hat{\omega}]$. Further, $c_-(s_t)$ and $c_+(s_t)$ notation refers to all consumption before and after date s_t, respectively.

What does this definition mean? A uniform impatient agent has a fraction γ such that the utility from increasing consumption by $\hat{\omega}$ today and decreasing consumption by $(1-\gamma)\%$ at all future dates is better than staying with the current consumption plan. This definition (which will now be used as an assumption) seems a little odd but is not that restrictive; for instance, the discounted, time-separable utility functions exhibits uniform impatience.

Theorem. *Assume that agents' utility functions exhibit uniform impatience. Suppose that q is the sequence of event prices associated with a security markets equilibrium price vector p. If the present value of the aggregate endowment is finite, then the price bubble of every security is zero, <u>for all securities in strictly positive supply</u>.*

18 Arrow-Debreu Equilibrium in Inifnite Time

- Define an Arrow-Debreu (AD) equilibrium.
- Under sitrctly increasing utility and market completeness, if $(q, \{c^i\})$ is an AD equilibrium, then $\exists h$ such that $(p, \{c^i, h^i\})$ is a security market equilibrium (SM) under natural debt constraints.
- Under strictly increasing utility and market completeness and zero price bubbles, if $(p, \{c^i, h^i\})$ is a SM equilibrium under debt constraints, then $\exists! q$ such that q are event prices and $(q, \{c^i\})$ is the AD equilibrium.

An Arrow-Debreu market is date-0 market in the sense that all of the trade takes place at the beginning of time. Further, all contingent claims (i.e. all Arrow securities) can be traded, which implies a complete market. Given this type of market and its properties, we can make comparisons to the more realistic dynamic portfolio market in which agents trade/update their portfolio sequentially, after each date.

To begin, we postulate the existence of market at date 0 which trades over all future events s_t for all t. Prices are described by a positive, linear functional Q, which assigns a date 0 price to a payoff. This can often be represented by a sequence $q \in \mathbb{R}^\infty$ such that we write

$$Q(c) = \sum_{t=0}^{\infty} \sum_{s_t} q(s_t) c(s_t)$$

for any infinite stream of state-contingent consumption. This functional form assumption is known as *countably additive pricing*. Normalize the price to $q(s_0) = 1$. Agents solve the problem

$$\begin{aligned} \max_{c,h} \quad & u(c) \\ s.t. \sum_t \sum_{s_t} q(s_t) c(s_t) & \leq \sum_t \sum_{s_t} q(s_t) \hat{\omega}^i(s_t) \end{aligned} \tag{1}$$

plus the restriction that consumption be positive. The first-order conditions lead to

$$\frac{u_c(s_t)}{u_c(s_0)} = q(s_t) \tag{2}$$

which states that the price of a claim to unit consumption at s_t is equal to the MRS between consumption at that event and date 0 consumption.

Definition. *An Arrow-Debreu equilibrium is a price system q and consumption allocation $\{c^i\}_{i=1}^I$ such that the consumption is a solution to each agent's problem (1) and markets clear*

$$\sum_{i=1}^{I} c^i(s_t) = \hat{\omega}(s_t)$$

at all dates and all events.

Definition. *Natural debt bounds are*

$$N^i(s_t) = \frac{1}{q(s_t)} \sum_{\tau=t}^{\infty} \sum_{\tau} q(s_\tau)\omega^i(s_\tau). \tag{3}$$

If the negative of this quantity is used as a debt constraint, it says that an agent is not able to take out debt larger than the present value of all future endowment streams.

Theorem. *Suppose that $(q, \{c^i\})$ is an Arrow-Debreu equilibrium in contingent commodity markets and agents' utility functions are strictly increasing. If security markets are dynamically complete at security prices p, then there exists a portfolio strategy $\{h^i\}$ such that $(p, \{c^i, h^i\})$ is an equilibrium in security markets under natural debt constraints.*

Now, a bit of a converse.

Theorem. *Suppose that $(p, \{c^i, h^i\})$ is a security markets equilibrium under debt constraints and agents' utility functions are strictly increasing. If security markets are dynamically complete at prices p and price bubbles are zero, then $(q, \{c^i\})$ is an Arrow-Debreu equilibrium where q is the unique sequence of event prices.*

19 Speculative Trade

19.1 Harrison-Kreps (1978)

Consider a simple security economy for an infinite number of discrete dates $t = 0, 1, 2, ...$ with just a single security x_t which pays out a dividends as

$$x_t(s) = \begin{cases} 1, & \text{if s=1} \\ 0, & \text{if s=0} \end{cases} \quad (1)$$

so that there is a high state (1) and low state (0) in any given period. Markets are incomplete given debt limits $D_t = 0$, which is to say there is no short selling the asset. There is a positive aggregate supply of the security $\bar{h}_0 = 1$. Agents have preferences

$$u^i(c) = E^i[\sum_{t=0}^{\infty} \beta^t c_t \quad (2)$$

so that agents are risk-neutral and have their own beliefs about payoffs, embodied in different probability measures π^i. In particular, assume agents believe dividends follow a one-period Markov chain and have beliefs

$$Q^1 = \begin{bmatrix} 1/2 & 1/2 \\ 2/3 & 1/3 \end{bmatrix}$$

and

$$Q^2 = \begin{bmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{bmatrix}.$$

Consider the underlying *fundamental value* that each agent has towards holding the security forever. Denote this as $V^i(s)$ where s is the initial date-0 state. This is written

$$V^i(s) = \sum_{t=1}^{\infty} \beta^t E_0^i[x_t] \quad \forall s = 0, 1 \quad (3)$$

$$= \beta[\pi_s^i p_0 + (1 - \pi_s^i)(p_1 + 1)] \quad (4)$$

where (3) calculates fundamental value via taking expectation of the perpetual dividends steam and (4) utilizes the fact that date-0 value is the discounted, expected payoff of the security next period.[9] This ius a system of two equations and can also be represented as

$$P = \beta\big[I - \beta Q]^{-1} Q \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (5)$$

where P is the price vector and I is the identity matrix.

[9] In the low state, the payoff is just the low state price p_0 of the security, whereas in the high state the payoff is the high state price p_1 plus the dividend of 1.

Given a discount factor $\beta = 0.75$, we recover prices for the agents

$$(p_0^1, p_1^1) = (4/3, 11/9) \quad \text{and} \quad (p_0^2, p_1^2) = (16/11, 21/11) \tag{6}$$

where p_i^j is the price attached to the security by agent j given date-0 realization of state i. At this point, we obviously note that agent's have heterogeneous beliefs about the fundamental value of the security. Further, agent 2 is more optimistic in both starting scenarios!

These fundamental values exist in a security economy where the agent is alone with those beliefs. What happens when heterogenuous beliefs coexist in equilibrium? It must be true that

$$p_t = \max_i \quad \beta E_t^i[p_{t+1} + x_{t+1}] \tag{7}$$

where the price is set by the agent with the most optimistic view tomorrow. This equilibrium condition leads to the system of equations

$$p_0 = \beta[\frac{1}{2}p_0 + \frac{1}{2}(p_1 + 1)] \tag{8}$$

$$p_1 = \beta[\frac{1}{4}p_0 + \frac{3}{4}(p_1 + 1)] \tag{9}$$

where (8) is priced by agent 1's expectations and (9) by agent 2's: different beliefs but the <u>same</u> price system. This leads to solutions

$$(p_0^*, p_1^*) = (\frac{24}{13}, \frac{27}{13}) \tag{10}$$

and agent 1 holds all of the security ($\bar{h} = 1$) in state 0 and agent 2 holds all in high state 1. Notice that the equilibrium prices in (10) exceed the fundamental values of both agents in both states in (6)! Thus, there exists a speculative bubble for all events s_t in the model. How is this possible? From the equilibrium condition, the investor has to forecast both future dividends and prices; this means, the agent must implicitly take into account the beliefs of the other agent. Even if a state gives the agent a signal of low dividends to come, it may nonetheless forecast higher prices due to the other agent's optimistic beliefs. Here, agent's trade not only for dividends but also for resale value of the security.

19.2 Rational Expectations Equilibrium

Consider a two-period security market economy with I agents who each receive a signal $\sigma_i \in \Sigma_i$ which provides information about the state of the world tomorrow s. Agent i's signal is private to her so she does not observe others' signals, but she does observe the market price and she knows the joint probability distribution $\pi(s, \sigma_1, ..., \sigma_I)$ of a certain state and aggregate signal being observed. Here, we develop the notion of a price forecast function.

Definition. *A price forecast function Φ maps joint signals to security price vectors:*

$$\Phi : \Sigma_1 \times \cdots \times \Sigma_I \to \mathbb{R}^J_+ \tag{1}$$

where Σ_i is a finite set containing all the possible signals that agent i can receive, and we can represent the signal profile simply as Σ.

Therefore, if an agent observes a price p then she can maybe recover the signal via $\Phi^{-1}(p)$ or at least make some probability assessment based upon the function $\Phi(\cdot)$. Thus, the agent problem is

$$\begin{aligned} \underset{h, c_0^i, c_{1,s}^i}{max} \quad & \sum_{s=1}^{S} \pi(s|\sigma^i, \Phi = p) u^i(c_0^i, c_{1,s}^i) \\ s.t. \quad & c_{0p}^i h^i \leq \omega_0^i \\ s.t. \quad & c_{1,s}^i \leq w_{1,s}^i + x_s h^i \quad \forall s \end{aligned} \tag{2}$$

Definition. *A rational expectations equilibrium is a price forecast function Φ such that $\Phi(\sigma) = p$, where p is an equilibrium price vector, and each agent solves their problem (2), for any joint signal σ. Further, markets must clear:*

1. $\sum_i c_0 = \sum_i \omega_0^i$
2. $\sum_i h_i = \bar{h_i}$
3. $\sum_i c_{1,s}^i = \sum_i \omega_{1,s}^i + \bar{h} x_s \quad \forall s = 1, 2, ..., S.$

If the price forecast function is *revealing* (i.e. a one-to-one function) then the rational expectations equilibirum is not very interesting, as agents can perfectly recover the joing signal σ, conditioned on the observed price $\Phi(\sigma) = p$. A variation on this setup is to assume that all agents receive a *noisy signal* which is composed of the true value of the signal and an error term.

20 Recursive Preferences

Refer to Kreps and Porteus for seminal paper on recursive preferences with complete axiomization. For now, we consider Epstein-Zin preferences, which have a period t utility form of

$$U_t = [(1-\beta)c_t^{1-\rho} + \beta R_t(U_{t+1})^{1-\rho}]^{\frac{1}{1-\rho}}$$

where $R_t(U_{t+1}) = E_t[U_{t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}}$. This is model is characterized with three parameters (β, γ, ρ). When $\rho = \gamma$, the formula reduces to expected utility. Loosely, ρ is for inter-temporal elasticity and γ represents the agent's attitude towards risk.

To illustrate how this representation breaks from that of a standard expected utility representation, consider the following example. You have two identical consumption sequences with nearly identical information sets:

$$\{c_s\}_{s=0}^t \quad \text{with} \quad \{\mathcal{F}_t\}_{t=0}^\infty$$
$$\{c_s\}_{s=0}^t \quad \text{with} \quad \{\mathcal{F}_t^*\}_{t=0}^\infty$$

where $\mathcal{F}_t = \mathcal{F}_t^* \quad \forall t \neq \tau$ and $\mathcal{F}_\tau^* = \mathcal{F}_{\tau+1}$. The idea here is that the sequence of consumption is the same and the available information is the same for all dates except date τ, in which case $\mathcal{F}^*$ tells you tomorrow's consumption and $\mathcal{F}$ does not. Now, consider the value function transformations

$$\tilde{V}_t = \frac{U_t^{1-\rho}}{1-\rho}, \quad \tilde{V}_t^* = \frac{U_t^{*1-\rho}}{1-\rho}.$$

Now, re-write the Bellman as

$$\begin{aligned}
\tilde{V}_t =& \frac{1}{1-\rho}[(1-\beta)c_t^{1-\rho} + \beta R_t(U_{t+1})^{1-\rho}]^{\frac{1-\rho}{1-\rho}} \\
=& \frac{1-\beta}{1-\rho}c_t^{1-\rho} + \frac{\beta}{1-\rho}E_t[U_{t+1}^{1-\gamma}]^{\frac{1-\rho}{1-\gamma}} \\
=& \frac{1-\beta}{1-\rho}c_t^{1-\rho} + \frac{\beta}{1-\rho}E_t[((1-\rho)\tilde{V}_{t+1})^{\frac{1-\gamma}{1-\rho}}]^{\frac{1-\rho}{1-\gamma}}
\end{aligned}$$

and the same form follows through for $\tilde{V}_t^*$. Look at $\tilde{V}_\tau^*$:

$$\tilde{V}_\tau^* = \frac{1-\beta}{1-\rho}c_\tau^{1-\rho} + \beta\tilde{V}_{\tau+1}$$

where $E_{\tau^*}[((1-\rho)\tilde{V}_{\tau+1}^*)^{\frac{1-\gamma}{1-\rho}}]^{\frac{1-\rho}{1-\gamma}} = E_{\tau+1}[((1-\rho)\tilde{V}_{\tau+1}^*)^{\frac{1-\gamma}{1-\rho}}]^{\frac{1-\rho}{1-\gamma}} = (1-\rho)\tilde{V}_{\tau+1}$ because conditioning on $\mathcal{F}_{\tau+1}$ changes the value function (inside the expectation operator) from $\tilde{V}_{\tau+1}$ to $\tilde{V}_{\tau+1}$.

Now, let us define the concepts of *early resolution* and *late resolution*. Define early resolution of uncertainty as $E[\tilde{V}_\tau^*|\mathcal{F}_{-\tau}^*]$ where $\mathcal{F}_{-\tau}^* = \mathcal{F}_\tau$ is the $*$ information set, in period τ, just before receiving the signal/information that gives $\mathcal{F}_{\tau+1}$. So, we derived the functional form of $\tilde{V}_\tau^*$ where the agent sees ahead to $\tau+1$ and now condition on that with date information set $\mathcal{F}_\tau$. Define late resolution as $E[\tilde{V}_\tau|\mathcal{F}_\tau]$. Thus, early resolution is preferred to late resolution when

$$\begin{aligned}
E[\tilde{V}_\tau^*|\mathcal{F}_{-\tau}^*] >& E[\tilde{V}_\tau|\mathcal{F}_\tau] \\
\Rightarrow \beta E_\tau[\tilde{V}_{\tau+1}] >& \frac{\beta}{1-\rho}E_\tau[((1-\rho)\tilde{V}_{\tau+1})^{\frac{1-\gamma}{1-\rho}}]^{\frac{1-\rho}{1-\gamma}} \\
\Rightarrow E_\tau[(1-\rho)\tilde{V}_{\tau+1}] >& E_\tau[((1-\rho)\tilde{V}_{\tau+1})^{\frac{1-\gamma}{1-\rho}}]^{\frac{1-\rho}{1-\gamma}}.
\end{aligned}$$

If we consider the function $\phi(x) = x^{\frac{1-\gamma}{1-\rho}}$ where $x = (1-\rho)\tilde{V}_{\tau+1}$ then the above inequality reduces to

$$E_\tau[\phi(x)^{\frac{1-\rho}{1-\gamma}}] > E_\tau[\phi(x)]^{\frac{1-\rho}{1-\gamma}}$$

where $\phi(\cdot)$ is a concave function if $\frac{1-\gamma}{1-\rho} \in [0,1] \Rightarrow \rho < \gamma$, which is a necessary and sufficient condition for the agent to prefer early resolution of uncertainty.

20.1 IID versus Uncertain Perpetuity

Proposition. *When* $\rho \to 1$,

$$U_t = c_t^{1-\beta}[R_t(U_{t+1})]^\beta.$$

Thus, taking logs,

$$\begin{aligned}\tilde{U}_t =&(1-\beta)\tilde{c}_t + \beta\tilde{R}_t(U_{t+1})\\ =&(1-\beta)logc_t + \frac{\beta}{1-\gamma}log(E_t[exp\{(1-\gamma)\tilde{U}_{t+1}\}])\\ =&(1-\beta)logc_t + \beta Q_t(\tilde{U}_{t+1})\end{aligned}$$

where $Q_t(\tilde{U}_{t+1}) = \frac{1}{1-\gamma}log(E_t[exp\{(1-\gamma)\tilde{U}_{t+1}\}])$.

Let's consider different cases of consumption streams for the agent with $\rho = 1$. For the first case (perpetual consumption case), assume that with probabilities $(p, 1-p)$ the change in log consumption will be $\Delta logc_{t+1} = g_{t+1} \in \{g_l, g_h\}$ forever after the first realization. That is, with probability p the agent will observe $\{g_l, g_l, g_l, ...\}$ and with probability $(1-p)$ the agent will observe $\{g_h, g_h, ...,\}$ forever. Since, we are talking about constrant growth rates of consumption, the value funciton is not stationary; therefore, consider the following de-meaning transformation (Note: Let U_t be the logged version)

$$\begin{aligned}U_t - c_t =& -\beta logc_t + \beta Q_t(U_{t+1})\\ =&\beta Q_t(U_{t+1} - c_t) && \text{(Check it)}\\ =&\beta Q_t(U_{t+1} - c_{t+1} + c_{t+1} - c_t)\end{aligned}$$

such that we have $v_t = U_t - c_t$ and write

$$v_t = \beta Q_t(v_{t+1} + g_{t+1})$$

where g_{t+1} is the log-growth rate of consumption in period $t+1$. This value function is stationary. Therefore, after uncertainty has been resolved, we simply have

$$\begin{aligned}&v_= \beta Q(v_+ g_i)\\ \Rightarrow &v_= \frac{\beta}{1-\gamma}log(exp\{(1-\gamma)(v_i + g_i)\})\\ \Rightarrow &v_i = \frac{\beta}{1-\beta}g_i \quad \forall i = l, h.\end{aligned}$$

That is to say the the value function, contingent on the realization of growth rates, is simply the present value of the constant stream g_i. On the other hand, at time 0, we have

$$\begin{aligned} v_0 =& \beta Q(v' + g') \\ =& \frac{\beta}{1-\gamma} log(E[exp\{\frac{1-\gamma}{1-\beta} g'\}]) \\ =& \frac{\beta}{1-\beta}\frac{1-\beta}{1-\gamma} log(E[exp\{\frac{1-\gamma}{1-\beta} g'\}]) \\ =& \frac{\beta}{1-\beta}\Omega_{\frac{1-\gamma}{1-\beta}}(g') \end{aligned} \tag{*}$$

where

$$\Omega_\theta(x) = \frac{1}{\theta} log(E[exp\{\theta x\}]) \quad \text{and} \quad \theta = \frac{1-\gamma}{1-\beta} \quad \text{and} \quad \theta > 0.$$

For the <u>second case</u> (the iid case), consider consumption growth rates as an iid process each period with $\{g_l, g_h\}$ and corresponding probabilities $(p, 1-p)$. Now, value function v is constant throughout all time periods. Label this one $\bar{v}$. We observe

$$\begin{aligned} \bar{v} =& \beta\bar{v} + \frac{\beta}{1-\gamma} log(E[exp\{(1-\gamma)g'\}]) \\ \Rightarrow& \bar{v}(1-\beta) = \beta\Omega_{1-\gamma}(g') \\ \Rightarrow& \bar{v} = \frac{\beta}{1-\beta}\Omega_{1-\gamma}(g'). \end{aligned} \tag{**}$$

Thus, for given $\rho = 1$, we can look at scenarios in which the agent would prefer the random perpetuity versus the iid case by comparing (*) with (**) as

$$\frac{\beta}{1-\beta}\Omega_{\frac{1-\gamma}{1-\beta}}(g') > \frac{\beta}{1-\beta}\Omega_{1-\gamma}(g') \tag{***}$$

$$\Rightarrow \Omega_{\frac{1-\gamma}{1-\beta}}(g') > \Omega_{1-\gamma}(g')$$

and you can see as $\gamma \to \rho = 1$ we head to the expected utility representation and the agent is indifferent. First, note that $\frac{1-\gamma}{1-\beta} > 1-\gamma$ if and only if $\gamma < 1$. Use the following proposition to establish whether or not Ω is an increasing or decreasing function.

Proposition. *Let m be a non-negative random variable with $E[m] = 1$ and define function $\phi(m) = mlogm$, then $E[\phi(m)] \geq 0$.*

A good way to compare (***) is to look at the function Ω and determine its derivative in the θ domain. Notice

$$\frac{\partial\Omega_\theta}{\partial\theta} = -\frac{log(E[exp\{\theta x\}])}{\theta^2} + \frac{1}{\theta}\frac{E[exp\{\theta x\}x]}{E[exp\{\theta x\}]}.$$

In the notation of the proposition, let $m = \frac{exp\{\theta x\}}{E[exp\{\theta x\}]}$ which is non-negative and $E[m] = 1$. From the proposition we know

$$\begin{aligned} E[\frac{e^{\theta x}}{E[e^{\theta x}]}]log(\frac{e^{\theta x}}{E[e^{\theta x}]}) =& -logE[e^{\theta x}] + E[\frac{e^{\theta x}\theta x}{E[e^{\theta x}]}] \\ =& -\frac{1}{\theta^2}logE[e^{\theta x}] + \frac{1}{\theta}E[\frac{e^{\theta x}x}{E[e^{\theta x}]}] \\ =& \frac{\partial\Omega_\theta}{\partial\theta} \geq 0. \end{aligned}$$

Therefore, Ω is an increasing function of θ. If $\gamma < 1$, we have $\frac{1-\gamma}{1-\beta}$ is greater than $1-\gamma$ and it must be that the perpetuity case (v_0) offers higher utility than the iid case ($\bar{v}$).

20.2 Stochastic Discount Factor

For a stochastic discount factor, we generally observe the form

$$m_{t,t+1} = \frac{\frac{\partial u_{t+1}}{\partial c_{t+1}} \frac{\partial u_t}{\partial u_{t+1}}}{\frac{u_t}{\partial c_t}}.$$

Thus, for the recursive preferences, calculate

$$\begin{aligned}\frac{\partial U_t}{\partial c_t} &= \ldots = (1-\beta) U_t^{\rho} c_t^{-\rho} \\ \frac{\partial U_t}{\partial U_{t+1}} &= \beta U_t^{\rho} R_t(U_{t+1})^{-\rho} u_{t+1}^{-\rho} (R_t(U_{t+1}))^{\gamma}\end{aligned}$$

which leads to stochastic discount factor

$$m_{t,t+1} = \beta (\frac{c_{t+1}}{c_t})^{-\rho} (\frac{R_t(U_{t+1})}{U_{t+1}})^{\gamma-\rho}$$

which looks like the classic SDF we get from the consumption-based asset pricing model, with the added component $(\frac{R_t(U_{t+1})}{U_{t+1}})^{\gamma-\rho}$.

21 Constantinides-Duffie

This model is based upon *Asset Pricing with Heterogenuous Consumers* (1996) by Constantinides and Duffie. Denote prices $\{\pi_t^j\}$ and dividends $\{d_t^j\}$ for all assets $j = 1, 2, ..., k$. We have a set A of agents with income processes $\{y_{it}\}_{i\in A}$ and an initial holding of assets $\{\theta_{i,-1}^j\} = \{\theta_{-1}^j\}_{i\in A}$, such that they have identical starting assets. Each agent solves

$$\max_{\{c_t^i,\{\theta_t^j\}_j\}_t} \quad E_0[\sum_t e^{-\rho t}\frac{c_{it}^{1-\sigma}}{1-\sigma}]$$
$$s.t. \quad c_{it} + \sum_j \pi_t^j \theta_{it}^j = y_{it} + \sum_j (\pi_t^i + d_t^j)\theta_{i,t-1}^j$$

and in equilibrium we must observe market clearance

$$\int_{i\in A} c_{it}di = \int_{i\in A} y_{it}di + \sum_j d_t^j$$
$$\int_{i\in A} \theta_{it}^j di = 1 \quad \forall j = 1, 2, ..., k$$

All agents maximizing, taking prices as given, along with market clearing conditions, defines a competitive equilibrium. Moving forward, the central idea of the paper is that if observe data on prices and returns, then we can construct A, $\{y_{it}\}$ and a stochastic discount factor to justify the the data as equilibrium outcomes.

Assumption. *No arbitrage.*

This implies there exists a sitrctly positive, time-zero discount factor

$$m_t^* = \prod_{j=0}^{t} m_{j,j+1}^*$$

such that prices satisfy

$$\pi_t^j = \frac{1}{m_t^*}E_t[\sum_{j=0}^{\infty} m_{t+j}^* d_{t+j}]$$
$$\Rightarrow 1 = E_t[m_{t,t+1}^* R_{t,t+1}^j]$$

which will hold for any asset j. A one-period ahead return is simply $R_{t,t+1}^j = \frac{\pi_{t+1}^j + d_{t+1}^j}{\pi_t^j}$.

Assumption.

$$\lim_{j\to\infty} E_t[m_{t+j}] = 0$$

Assumption.

$$\frac{m_{t+1}^*}{m_t^*} \geq e^{-\rho}(\frac{c_{t+1}}{c_t})^{-\sigma} \quad \forall t$$

where the former condition is a transversality-type condition on the SDF, whereas the latter is a little stranger/stronger (but holds up under statistical tests in the paper). For an agent i, we have the Euler equation

$$E_t[R_{t,t+1}^j e^{-\rho}(\frac{c_{i,t+1}}{c_{i,t}})^{-\sigma}] = 1.$$

22 Long-Run Risk

22.1 Baseline Model

This section is based upon the long-run risk model, due to Bansal-Yaron (2004). In particular, suppose we have an agent with Epstein-Zin preferences. Then for a return $R_{i,t+1}$, we have the standard asset pricing equation

$$E_t[\delta^\theta G_{t+1}^{-\frac{\theta}{\psi}} R_{a,t+1}^{-(1-\theta)} R_{i,t+1}] = 1 \tag{1}$$

where δ is the period discount fact, $R_{a,t+1}$ is the gross return on an asset which has aggregate consumption as its dividends, and G_{t+1} is the growth rate of aggregate consumption. Further, we have the parameterization $\theta = \frac{1-\psi}{1-1/\psi}$ where $\gamma \geq 0$ is the risk aversion parameter and $\psi \geq 0$ is the intertemporal parameter. Thus, we have

$$M_{t+1} = \delta^\theta G_{t+1}^{-\frac{\theta}{\psi}} R_{a,t+1}^{-(1-\theta)} \tag{2}$$

as the model stochastic discount factor. We assume that returns on asset i are observable but the R_a return is unobservable. From Campbell-Shiller, we utilize the approximation

$$r_{a,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + g_{t+1} \tag{3}$$

where lowercase variables are *logged* and $z_t = log(\frac{P_t}{C_t})$ is the log of the price-consumption ratio. Further, (κ_0, κ_1) are constants. From this we have

$$m_{t+1} = \theta log\delta - \frac{\theta}{\psi} g_{t+1} + (\theta - 1) r_{a,t+1}. \tag{4}$$

We now specify the following exogenous processes that drive the model:

$$\begin{aligned}
x_{t+1} =& \rho x_t + \phi_e \sigma_t e_{t+1} \\
g_{t+1} =& \mu + x_t + \sigma_t \eta_{t+1} \\
g_{d,t+1} =& \mu_d + \phi x_t + \phi_d \sigma_t + u_{t+1} \\
\sigma_{t+1}^2 =& \sigma^2 + \nu_1(\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1}
\end{aligned}$$

where g_d is the log of the growth rate for aggregate dividends. We assume $(e_{t+1}, \eta_{t+1}, u_{t+1}, w_{t+1})$ are iid standard normal random variables. Notice that for the growth rates $(g_t, g_{d,t})$, there is a persistent component x_t included along with time-varying variance (i.e. heteroskedasticity). For this model, it will be sufficient to work with the state variables (x_t, σ_t^2).

We will conjecture the form

$$z_t = A_0 + A_1 x_t + A_2 \sigma_t^2 \tag{5}$$

and plug this into (3) and that into (1) to recover the z coefficients and ultimately the value of $r_{a,t+1}$. For the LHS of (1), we write

$$E_t[exp\{\theta log\delta - \frac{\theta}{\psi}g_{t+1} + \theta r_{a,t+1}\}] = 1 \qquad \text{(letting } r_{i,t+1} = r_{a,t+1}\text{)}$$

$$\Rightarrow E_t[exp\{\theta log\delta - \frac{\theta}{\psi}g_{t+1} + \theta(\kappa_0 + \kappa_1 z_{t+1} - z_t + g_{t+1})\}] = 1 \qquad \text{(subbing in (3))}$$

$$\Rightarrow E_t[exp\{\theta log\delta - \frac{\theta}{\psi}g_{t+1} + \theta(\kappa_0 + \kappa_1(A_0 + A_1 x_{t+1} + A_2\sigma_{t+1}^2) - (A_0 + A_1 x_t + A_2\sigma_t^2) + g_{t+1})\}] = 1$$

$$\text{(subbing in (5))}$$

$$\Rightarrow E_t[exp\{\theta log\delta - \frac{\theta}{\psi}(\mu + x_t + \sigma_t\eta_{t+1}) + \theta(\kappa_0 + \kappa_1(A_0 + A_1(\rho x_t + \phi_e\sigma_t e_{t+1}) + \qquad \text{(subbing in } g_{t+1} \text{ and } x_{t+1}\text{)}$$

$$+ A_2(\sigma^2 + \nu_1(\sigma_t^2 - \sigma^2) + \sigma_w w_{t+1}) - (A_0 + A_1 x_t + A_2\sigma_t^2) + \mu + x_t + \sigma_t\eta_{t+1}\}] = 1$$

where the last line takes expectations of lognormal random variables. Because the Euler equation must hold for all values of the stat variables (x_t, σ_t), we first collect all x_t terms and set to zero:

$$-\frac{\theta}{\psi}x_t + \theta(\kappa_1 A_1\rho x_t - A_1 x_t + x_t) = 0$$

$$\Rightarrow A_1 = \frac{1 - 1/\psi}{1 - \kappa_1\rho} \tag{6}$$

and then collect all σ_t^2 terms and set to zero yields

$$\theta[\kappa_1\nu_1 A_2\sigma_t^2 - A_2\sigma_t^2] + \frac{1}{2}[(\theta_{\frac{\theta}{\psi}})^2 + (\theta A_1\kappa_1\phi_e)^2]\sigma_t^2 = 0$$

$$\Rightarrow A_2 = \frac{\frac{1}{2}[(\theta - \theta/\psi)^2 + (\theta A_1\kappa_1\phi_e)^2]}{\theta(1 - \kappa_1\nu_1)}. \tag{7}$$

22.2 Cost of Uncertainty

We want to compare utility under two scenarios: one in which uncertainty is resolved gradually, and the other in which uncertainty is resolved at time $t = 1$. Assume the following dynamics:

$$log\Delta c_{t+1} = m + x_t + \sigma W_{c,t+1} \tag{1}$$

$$x_{t+1} = ax_t + \phi\sigma W_{x,t+1} \tag{2}$$

where W is iid with distribution $N(0, \Sigma)$. Further, assume agents have Epstein-Zin preferences with unitary elasticity of substitution, so we get

$$logU_t = (1-\beta)logc_t + \beta log(E_t[U_{t+1}^{\alpha}]^{\frac{1}{\alpha}}) \tag{3}$$

where $1 - \alpha$ is the relative risk aversion of the agent. Let's define the *timing premium* π^* as

$$\pi^* = 1 - \frac{U_0}{U_0^*} \tag{4}$$

where U_0 is the date-0 utility in the normal setting and U_0^* is utility in the setting in which risk is resolved at time $t = 1$. Let's first start with the normal setting and use guess and verify on the functional form. Guess

$$logU_t = logc_t + \phi_1 m + \phi_2 x_t + \phi_3. \tag{5}$$

Now, let's derive date-0 utility

$$\begin{aligned} logU_0 =& (1-\beta)logc_0 + \beta log(E_0[e^{\alpha logU_1}]^{\frac{1}{\alpha}}) \\ =& (1-\beta)logc_0 + \beta[\phi_1 m + \phi_3] + \beta log(E_0[e^{\alpha(logc_1 - logc_0 + logc_0 + \phi_2 x_1)}]^{\frac{1}{\alpha}}) \\ =& logc_0 + \beta[\phi_1 m + \phi_3] + \beta log(E_0[e^{\alpha log\Delta c_1 + \alpha\phi_2 x_1}]^{\frac{1}{\alpha}}) \\ =& logc_0 + \beta[\phi_1 m + \phi_3] + \beta log([e^{\alpha m + \alpha x_0 + \frac{\alpha^2\sigma^2}{2} + \alpha\phi_2 a x_0 + \frac{(\phi\alpha\phi_2\sigma)^2}{2}}]^{\frac{1}{\alpha}}) \\ =& logc_0 + m\beta(1+\phi_1) + x_0\beta(1 + a\phi_2) + \beta\left(\phi_3 + \frac{\sigma^2\alpha}{2}(1 + \phi^2\phi_2^2)\right). \end{aligned} \tag{4}$$

Then, matching coefficients we get

$$logU_0 = logc_0 + \frac{\beta}{1-\beta}m + \frac{\beta}{1-\beta a}x_0 + \frac{\alpha}{2}\frac{\beta\sigma^2}{1-\beta}[1 + \frac{\phi^2\beta^2}{(1-\beta a)^2}]. \tag{6}$$

Now, consider the case of total resolution of uncertainty at date $t = 1$. We would have,

$$\begin{aligned} logU_1^* =& (1-\beta)[logc_1 + \beta logc_2 + \beta^2 logc_3 + ...] \\ =& logc_0 + \beta_{t=1}^{\infty}\beta^{t-1}log\Delta c_t \end{aligned} \tag{7}$$

which, using the date-0 information set, is a normal random variable with mean

$$\begin{aligned} E_0[logU_1^*] =& logc_0 + \sum_{t=1}^{\infty}\beta^{t-1}E_0[log\Delta c_t] \\ =& logc_0 + \frac{m}{1-\beta} + \sum_{t=1}^{\infty}\beta^{t-1}E[x_{t-1}] \\ =& logc_0 + \frac{m}{1-\beta} + \frac{a}{1-\beta a}x_0 \end{aligned} \tag{8}$$

and a variance

$$var_0(logU_1^*) = \frac{\sigma^2}{1-\beta^2}\left(1 + \frac{\phi^2}{(1-\beta a)^2}\right). \tag{9}$$

We now have enough information to compute the date-0 utility in the case of uncertainty resolution at $t = 1$:

$$\begin{aligned} logU_0^* =& (1-\beta)logc_0 + \beta log\left(E_0[U_1^{*\alpha}]^{\frac{1}{\alpha}}\right) \\ =& logc_0 + \frac{\beta}{1-\beta}m + \frac{\beta}{1-\beta a}x_0 + \frac{\alpha}{2}\frac{\beta\sigma^2}{1-\beta^2}(1 + \frac{\phi^2\beta^2}{(1-\beta a)^2}). \end{aligned} \tag{10}$$

Therefore, compute

$$\begin{aligned} log\frac{U_0}{U_0^*} =& \frac{\alpha}{2}\beta\sigma^2\big[1 + \frac{\phi^2\beta^2}{(1-\beta a)^2}\big]\big(\frac{1}{1-\beta} - \frac{1}{1-\beta^2}\big) \\ =& \frac{\alpha}{2}\frac{\beta^2\sigma^2}{1-\beta^2}\big(1 + \frac{\phi^2\beta^2}{(1-\beta a)^2}\big). \end{aligned} \tag{11}$$

It follows, that

$$\pi^* = 1 - exp\{\frac{\alpha}{2}\frac{\beta^2\sigma^2}{1-\beta^2}\big(1 + \frac{\phi^2\beta^2}{(1-\beta a)^2}\big)\}. \tag{12}$$

Observation. *The timing premium π^* is positive if and only if $\alpha < 0$.*

Recall that we have a result (from the previous section) with Epstein-Zin preferences that early resolution of uncertainty is always preferred if *relative risk aversion* exceeds the inverse of elasticity of substitution; that is,

$$1 - \alpha > 1 - \rho.$$

In this example, we assume that $\rho = 0$ and we therefore get an equivalent result.

23 Habit Model

23.1 Baseline Model

The representative agent has preferences

$$E_t[\sum_t \delta^t \frac{(C_t - X_t)^{1-\gamma} - 1}{1-\gamma} \tag{1}$$

where δ is the discount factor and C_t is called the agent's *habit*. Note: pay attention to notation because capital letters represent untransformed variables, whereas lowercase stands for log-transformed ones. Define the surplus S_t as

$$S_t = \frac{C_t - X_t}{C_t}. \tag{2}$$

We specify a process for consumption growth and the surplus

$$\begin{aligned} log\Delta c_{t+1} =& g + \nu_{t+1} & (3) \\ s_{t+1} =& (1-\phi)\bar{s} + \phi s_t + \lambda(s_t)\nu_{t+1} & (4) \end{aligned}$$

where notice we are using the log transformations. For simplicity, we can take equation (3) as the specification for the endowment process. The agent has marginal utility $\delta^t(C_t - X_t)^{-\gamma} = \delta^t(C_t S_t)^{-\gamma}$, so we get a stochastic discount factor

$$\frac{M_{t+1}}{M_t} = \delta(\frac{S_{t+1}}{S_t})^{-\gamma}(\frac{C_{t+1}}{C_t})^{-\gamma} \tag{5}$$

written in log form as

$$\begin{aligned} log\Delta m_{t+1} =& log\delta - \gamma log\Delta c_{t+1} - \gamma log\Delta s_{t+1} \\ =& log\delta - \gamma\big[g + (1-\phi)\bar{s} + (\phi - 1)s_t + (1_\lambda(s_t))\nu_{t+1}\big] \\ =& log\delta - \gamma\big[g + (\phi-1)(s_t - \bar{s}) + (1+\lambda(s_t))\nu_{t+1}\big] & (6) \end{aligned}$$

such that $\frac{M_{t+1}}{M_t}$ is a lognormal random variable with mean

$$e^{log\delta - \gamma\big(g + (\phi-1)(s_t - \bar{s})\big) + \frac{\gamma^2\sigma_\nu^2}{2}(1+\lambda(s_t))^2}. \tag{7}$$

Then, consider a discount bond with price q as $q = E_t[\frac{M_{t+1}}{M_t}]$ which implies a gross risk-free interest rate $R^f = \frac{1}{q}$ and net rate $r^f \approx logR^f$ such that

$$r^f \approx -log\delta + \gamma g - \gamma(1-\phi)(s_t - \bar{s}) - \frac{\gamma^2\sigma_\nu^2}{2}(1+\lambda(s_t))^2. \tag{8}$$

So far, we do not have a functional form for $\lambda(\cdot)$ so we impose three restrictions/assumptions to pin down its form. We assume

$$r^f \text{ is constant.} \tag{9a}$$

$$\frac{\partial x_{t+1}}{\partial c_{t+1}}|_{S_{t+1}=\bar{S}} = 0 \tag{9b}$$

$$\frac{\partial[\frac{\partial x_{t+1}}{\partial c_{t+1}}]}{\partial s_{t+1}}|_{S_{t+1}=\bar{S}} = 0 \tag{9c}$$

Notice that these derivatives are with respect to log-transformed variables and evaluated at the untransformed steady state surplus value $\bar{S}$. From 9)a), we get that $(1+\lambda(s_t))^2$ is a linear function of s_t to balance the equation (8) at a constant value. Restrictions 9)b) and 9)c) require that habits are slow-changing with respect to consumption and the surplus when the surplus ratio is near its steady state, respectively.

To work on 9)b) first, we need an expression for x_{t+1} in terms of c_t. From (4) (with (3) plugged in on the error term), we get

$$\begin{aligned} X_{t+1} =& C_{t+1}[1-e^{s_{t+1}}] \\ &\Rightarrow x_{t+1} = c_{t+1} + log(1-e^{s_{t+1}}) \\ &\Rightarrow \frac{\partial x_{t+1}}{\partial c_{t+1}} = 1 - \frac{e^{s_{t+1}}\lambda(s_t)}{1-e^{s_{t+1}}} \end{aligned} \tag{10}$$

which can be rewritten as

$$\begin{aligned} \frac{\partial x_{t+1}}{\partial c_{t+1}} =& 1 - \frac{e^{s_t}\lambda(s_t)}{1-e^{s_t}} && \text{(assuming } s_{t+1} \approx s_t) \\ =& 1 - \frac{\lambda(s_t)}{e^{-s_t}-1}. && (11) \end{aligned}$$

The approximation made in the first line comes from the assumption (and empirical observation) that $\phi \approx 1$. Now, apply restriction (9b) to (10) to get

$$\begin{aligned} \lambda(\bar{s}) =& e^{-\bar{s}} - 1 \\ =& \frac{1}{\bar{S}} - 1 \end{aligned} \tag{12}$$

where we have subbed back in the untransformed steady state surplus $\bar{S}$. Now, moving towards restriction (9c), take the derivative of (10) with respect to surplus s_t:

$$\begin{aligned} \frac{\partial[\frac{\partial x_{t+1}}{\partial c_{t+1}}]}{\partial s_{t+1}}|_{S_{t+1}=\bar{S}} =& \frac{\partial[1-\frac{\lambda(s_t)}{e^{-s_t}-1}]}{\partial s_t} \\ =& -\frac{\lambda'(s_t)}{e^{-s_t}-1} - \frac{\lambda(s_t)e^{-s_t}}{(e^{-s_t}-1)^2} = 0 \end{aligned}$$

which evaluated at $\bar{S}$ gives

$$\lambda'(\bar{s}) = -\frac{1}{\bar{S}}. \tag{13}$$

Now, applying the restriction on the risk-free rate:

$$\begin{aligned} \frac{dr^f}{ds_t}|_{s_t=\bar{s}} =& -\gamma(1-\phi) - \gamma^2\sigma_\nu^2(1+\lambda(\bar{s}))\lambda'(\bar{s}) = 0 \\ \Rightarrow& \frac{\gamma\sigma_\nu^2}{\bar{S}^2} = 1-\phi \\ \Rightarrow& \bar{S} = \sqrt{\frac{\gamma\sigma_\nu^2}{1-\phi}}. \end{aligned} \tag{14}$$

Given this, we can recover a functional form for $\lambda(s_t)$; specifically,

$$\lambda(s_t) = \begin{cases} \frac{1}{\bar{S}}\sqrt{1-2(s_t-\bar{s})} - 1, & \text{if } s_t \leq s_{max} \\ 0, & \text{if } s_t > s_{max} \end{cases} \tag{15}$$

where

$$s_{max} = \bar{s} + \frac{1}{2}(1-\bar{S}^2). \tag{16}$$

23.2 Time-Varying Risk-free Rate

23.3 Value Function Iteration

We can write market return as $E_t[M_{t+1}(\frac{P_{t+1}+C_{t+1}}{P_t})] = 1$ where dividends are treated as equivalent to consumption claims at date t. Thus, we can write the Price-Consumption (Price-Dividend) ratio as

$$\frac{P_t}{C_t}(s_t) = E_t[M_{t+1}\frac{C_{t+1}}{C_t}(1 + \frac{P_{t+1}}{C_{t+1}}(s_{t+1}))] \tag{9}$$

such that returns can be written as

$$\frac{(P_{t+1}/C_{t+1}) + 1}{P_t/C_t}\frac{C_{t+1}}{C_t}. \tag{10}$$

To compute the numerical risk-free rate, we can use the first asset-pricing equation in (8) and do numerical integration on the first moment of the SDF. This is

$$\begin{aligned} E_t[M_{t+1}] =& \delta E_t[(\frac{C_{t+1}}{C_t}\frac{S_{t+1}}{S_t})^{-\gamma}] \\ =& \delta E_t[e^{-\gamma} e^{\Delta c_{t+1}} e^{s_{t+1}/s_t}] \\ =& \delta E_t[e^{-\gamma} e^{g+\nu_{t+1}} e^{(1-\phi)(\bar{s}-s_t)+\lambda(s_t)\nu_{t+1}}] \\ =& \delta e^{-\gamma(g+(1-\phi)(\bar{s}-s_t)} \int_{-\infty}^{\infty} e^{-\gamma(1+\lambda(s_t)v} f(v) d\nu \end{aligned} \tag{11}$$

where $f(\nu)$ is a normal probability distribution with mean zero and variance σ_v^2.

Next, we look to numerically solve for the ratio $\frac{P_t}{C_t}(s_t)$ over a grid of points for the logged surplus ratio s_t where the domain ranges from 0 to S_max (i.e. $S_{max} = e^{s_{max}}$). Write

$$\begin{aligned} \frac{P_t}{C_t}(s_t) =& E_t[M_{t+1}\frac{C_{t+1}}{C_t}(1 + \frac{P_{t+1}}{C_{t+1}}(s_{t+1}))] \\ =& \delta e^{-\gamma(g+(1-\phi)(\bar{s}-s_t)+g} E_t[e^{v_{t+1}(1-\gamma(1+\lambda(s_t)))}(1 + \frac{P_{t+1}}{C_{t+1}}((1-\phi)\bar{s} + \phi s_t + \lambda(s_t)\nu_{t+1})] \\ =& \delta e^{g(1-\gamma)-\gamma(1-\phi)(\bar{s}-s_t)} \int_{-\infty}^{\infty} e^{(1-\gamma(1+\lambda(s_t)))\nu}(1 + \frac{P_{t+1}}{C_{t+1}}((1-\phi)\bar{s} + \phi s_t + \lambda(s_t)\nu)) f(\nu) d\nu \end{aligned} \tag{12}$$

where $f(\nu)$ is the same pdf.

24 Appendix

24.1 Log Normal Distribution

A *log normal* random variable x is one such that its log-transform is normally distributed. Thus, if we define $x \sim N(\mu, \sigma^2)$, then

$$E[log(x)] = \mu \quad \text{and} \quad Var(log(x)) = \sigma^2.$$

In addition, we can derive results for the non-logged version of the random variable x. In particular, if $x \sim ln(\mu, \sigma^2)$ and we are looking at the linear combination ax for some scalar $a \in \mathbb{R}$, we observe

$$ax \sim N\left(e^{a\mu + \frac{a^2}{2}\sigma^2}, (e^{a^2\sigma^2} - 1)(e^{2a\mu + a^2\sigma^2})\right)$$

24.2 Log Utility as Limit of Power Utility

24.3 Projections