

Chapter 16

Asset prices

Monika Piazzesi and Martin Schneider

16.1 Introduction

This chapter is about economies with multiple assets. The benchmark neoclassical growth model assumes that there is only one asset that agents can save in, capital, and hence only one rate of return. The model thus abstracts from differences in returns across types of capital, say business capital and housing. It further assumes that all capital is held directly: it is not “packaged” into other assets by capital owners or intermediaries. We do not distinguish, for example, firms’ equity and debt, household mortgages and housing equity, mutual fund shares, pension accounts or bank deposits.

Allowing for multiple assets explicitly is interesting for two reasons. First, we can study more issues than in the benchmark model. For example, we would like to understand why rates of return differ on average or why some asset prices are more volatile than others. We can also ask why asset positions are different across agents. With multiple assets, agents decide not only on savings—the total value of assets held—but also make a portfolio choice—they pick the share of savings in each available asset. We can then ask how asset price fluctuations affect agents’ welfare differently.

Second, allowing for multiple assets is typically important for conclusions we draw about allocations, or macroeconomic aggregates, even if we are not interested in asset prices or portfolios per se. One reason is that asset prices provide important moments to discipline assumptions on agents’ objectives. For example, when we introduce equity and debt explicitly into the neoclassical growth model, we can infer what preferences of the representative agent are needed to match the average return on equity vs debt.¹ When markets are incomplete, there is another reason to study multiple assets: market structure matters directly for allocations. For example, agents bear less risk if there are instruments through which they can mutually insure each other.

¹Asset prices are therefore relevant for macroeconomics even when markets are complete and allocations are efficient. Since complete markets allow us to find allocations from the planner problem and determine asset prices only in a second step, it is tempting to think that asset prices “don’t matter.” However, they are still important to test the model.

16.2 Background on household portfolio and asset prices

To motivate the study of multiple assets, we start with some background on household portfolios and asset price properties. Two common assets in household portfolios are housing and stocks. In the United States, most households (currently 65%) own a house, while half of households own stocks (currently 52%). These participation rates vary by country. For example, some countries like Germany have lower homeownership rates (47%), while other countries like Italy have higher homeownership rates (75%).

An important factor in determining the cross-country variation in stock market participation are pension systems. In the U.S., for example, workers choose how much of their income to save in pension plans and how to allocate these savings to various assets. Many households, therefore, participate in the stock market through these pension plans. In contrast, households in European countries often receive pension payments from the government, which are financed through taxes. In these countries, households often save less and do not actively face an asset allocation decision.

Within countries, there is much household heterogeneity in portfolio shares, the fraction of wealth invested in any given asset. Young households often do not have much wealth. Therefore, they borrow money from a bank to buy a house and thus hold a leveraged position in a risky asset. Older households accumulate wealth, slowly pay down their mortgage balance, and invest in both housing and stocks. There is also a substantial share of households with little savings at all.

Investing in long-lived assets can be risky. For example, stocks make future dividend payments that are random from today's point of view. Houses provide housing services that are random. For example, housing services can suddenly deteriorate in a natural disaster such as an earthquake. In addition to the randomness of assets' cash flows, the resale value of the asset is uncertain. The *payoff* of an asset in the next period, which consists of the cash flow and the resale value, can thus be either unexpectedly high or low.

In the data, the values of long-lived assets are highly volatile over time relative to their cash flows. Since both asset value and cash flows grow over time, it is helpful to study the ratio of asset price to cash flows. For stocks, we can measure the stock price at the end of the year and divide it by the dividends the stock paid during the year. Similarly, we can measure the house price at the end of the year and divide it by the value of housing services derived from the house during the year. For a currently rented house, the value of its housing services is simply the annual rent payment. For an owner-occupied house, the value of housing services is estimated based on data from comparable rentals.

Figure 16.1 shows the two ratios of asset price to cash flow in the United States for the years 1929-2023. The black line represents the ratio of stock values to the dividends paid per year by these stocks. The gray line represents an analogous ratio for housing, which divides the value of residential housing by the housing services it provides. These series show dramatic movements. Sometimes, these movements go together (such as in the Great Depression of the early 1930s) and sometimes in opposite directions (as in the Great Inflation of the late 1970s). Judging from these series, investing in stocks in 1999 was a bad idea, as they became significantly cheaper seven years later, while investing in housing was a good idea. How can we explain these movements, and how are they related to the rest of the economy?

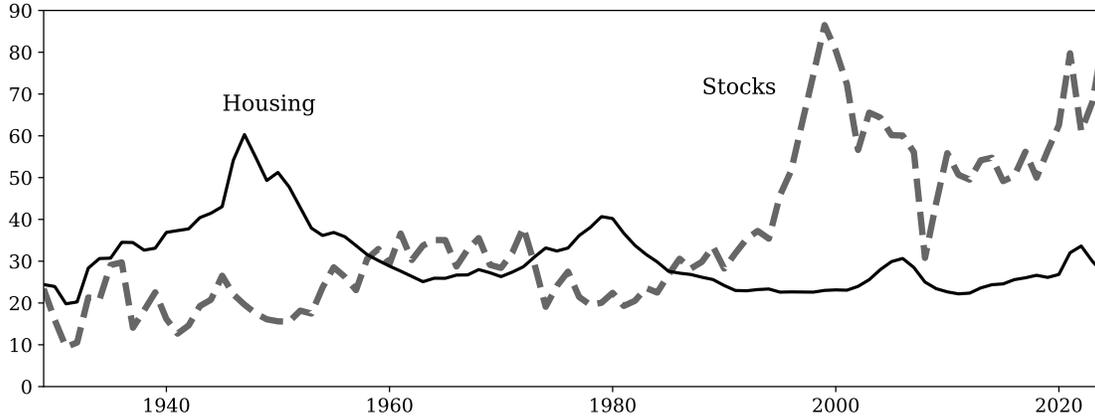


Figure 16.1: The ratio of asset value to cash flow for housing and stocks in the United States, 1930-2023.

16.3 Dynamic stochastic endowment economy

To understand the valuation of risky assets, we consider an endowment economy with an infinite horizon. At each date t , a state $\omega_t \in \Omega$ gets realized. There are S possible states. The history up to date t is $\omega^t = \{\omega_0, \omega_1, \dots, \omega_t\}$ and has probability $\pi_t(\omega^t)$. We can connect the history at date t to its unique predecessor history ω^{t-1} by writing $\omega^t = \{\omega^{t-1}, \omega_t\}$. The conditional probability of the state ω_t is $\pi_t(\omega^t | \omega^{t-1})$ or, equivalently, $\pi_t(\omega_t | \omega^{t-1})$. There is a set \mathcal{I} of agents so that the number of agents can be finite or infinite. These agents differ in their endowments of units of goods. Agent i has preferences described by a utility function $U_i(c_i)$ over consumption plans $c_i = \{c_{i,t}(\omega^t)\}_{t, \omega^t}$ and receives an endowment $y_i = \{y_{i,t}(\omega^t)\}_{t, \omega^t}$. The aggregate endowment is the sum of individual endowments $y_t(\omega^t) = \sum_{i \in \mathcal{I}} y_{i,t}(\omega^t)$.

A standard example for a utility function is expected utility

$$U_i(c_i) = \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega} \pi_t(\omega^t) \beta^t u(c_{i,t}(\omega^t)).$$

For understanding the dynamics of macro quantities, it is often sufficient to stay within this class of preferences. The most common example for the felicity function $u(x) = x^{1-\gamma}/(1-\gamma)$ is power utility, in which case the coefficient of relative risk aversion $\gamma = 1/\sigma$ is the inverse of the elasticity of intertemporal substitution σ (EIS). This specification thus tightly connects the agent's behavior towards gambles in static decisions (where γ matters) to the agent's behavior towards intertemporal choice in deterministic settings (where σ matters.) For understanding the properties of asset prices, however, it will be important to derive results for more general utility functions. The leading example is Epstein-Zin utility, a utility function that separates risk aversion from the EIS, $\gamma \neq 1/\sigma$. This chapter will therefore derive results based on more general utility functions $U_i(c_i)$.

Planner problem A feasible allocation $\{c_{i,t}(\omega^t)\}_{t,\omega^t,i}$ of consumption across agents in this economy satisfies

$$\sum_{i \in \mathcal{I}} c_{i,t}(\omega^t) \leq \sum_{i \in \mathcal{I}} y_{i,t}(\omega^t) = y_t(\omega^t) \quad \text{for all } t, \omega^t.$$

The planner assigns a weight λ_i to each agent and selects a feasible allocation that maximizes the welfare function

$$U_\lambda(y) := \max_{\{c_i\}_{i \in \mathcal{I}}} \sum_{i \in \mathcal{I}} \lambda_i U_i(c_i) \quad \text{subject to} \quad \sum_{i \in \mathcal{I}} c_{i,t}(\omega^t) \leq y_t(\omega^t).$$

The Lagrangian function of this problem is

$$\mathcal{L} = \sum_{i \in \mathcal{I}} \left\{ \lambda_i U_i(c_i) - \sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} \mu_t(\omega^t) \left(\sum_{i \in \mathcal{I}} c_{i,t}(\omega^t) - y_t(\omega^t) \right) \right\}.$$

The first-order condition (FOC) with respect to consumption of agent i at date t after history ω^t is

$$\lambda_i \partial U_i(c_i) / \partial c_{i,t}(\omega^t) = \mu_t(\omega^t). \quad (16.1)$$

An important property of this FOC is that every agent's weighted marginal utility is equated to a Lagrange multiplier that is independent of agent i . The envelope theorem says that the derivative of the Planner's Lagrangian function with respect to the aggregate endowment at date t after history ω^t is

$$\partial U_\lambda(y) / \partial y_t(\omega^t) = \mu_t(\omega^t). \quad (16.2)$$

If we combine the FOCs for agent i 's consumption at different dates, we see that the planner equalizes the marginal rates of substitution across all agents i with the planner's MRS:

$$\frac{\partial U_i(c_i) / \partial c_{i,t+1}(\omega^{t+1})}{\partial U_i(c_i) / \partial c_{i,t}(\omega^t)} = \frac{\partial U_\lambda(y) / \partial y_{t+1}(\omega^{t+1})}{\partial U_\lambda(y) / \partial y_t(\omega^t)}. \quad (16.3)$$

Arrow-Debreu economy Suppose agents trade Arrow-Debreu securities at date 0, as in Section 7.4 of Chapter 7. These securities are claims to a unit of consumption at a specific date t after a specific history ω^t . The time-0 price of this security is the Arrow-Debreu price of consumption at date t after history ω^t denoted $p_t^0(\omega^t)$. Given consumption prices $p_t^0(\omega^t)$ and the endowment y_i of agent i , the budget set of agent i is

$$B^{AD}(p^0, y_i) = \left\{ \{c_{i,t}(\omega^t)\}_{t,\omega^t} : \sum_{t,\omega^t} p_t^0(\omega^t) c_{i,t}(\omega^t) \leq \sum_{t,\omega^t} p_t^0(\omega^t) y_{i,t}(\omega^t) \right\}.$$

Agent i selects the consumption plan in the budget set to maximize utility. The Lagrangian function of the agent's problem is

$$\mathcal{L} = U_i(c_i) - \alpha_i \left(\sum_{t=0}^{\infty} \sum_{\omega^t \in \Omega^t} p_t^0(\omega^t) c_{i,t}(\omega^t) - p_t^0(\omega^t) y_{i,t}(\omega^t) \right).$$

The first-order condition with respect to consumption at date t after history ω^t is

$$\partial U_i(c_i) / \partial c_{i,t}(\omega^t) = \alpha_i p_t^0(\omega^t). \quad (16.4)$$

As long as the budget constraint binds, the Lagrange multiplier α_i is strictly positive, and we can divide by $\alpha_i > 0$. We again obtain the key property that every agent's weighted marginal utility is equated to a variable, here the price of consumption, that is independent of agent i . By combining FOCs of agent i at different dates, we can see that trading at date 0 equalizes the marginal rates of substitution across agents:

$$\frac{\partial U_i(c_i) / \partial c_{i,t+1}(\omega^{t+1})}{\partial U_i(c_i) / \partial c_{i,t}(\omega^t)} = \frac{p_{t+1}^0(\omega^{t+1})}{p_t^0(\omega^t)}. \quad (16.5)$$

Arrow-Debreu equilibrium An equilibrium in the Arrow-Debreu economy consists of a consumption allocation $\{c_i^*\}_{i \in \mathcal{I}}$ and consumption prices p^{0*} such that (i) the consumption plan c_i^* is in the budget set $B^{AD}(p^{0*}, y_i)$ for given consumption prices p^{0*} and maximizes the utility of agent i , (ii) markets clear $\sum_{i \in \mathcal{I}} c_{i,t}(\omega^t) = y_t(\omega^t)$ for all dates t and histories ω^t .

First Welfare Theorem The optimality conditions in the Arrow-Debreu economy (16.4) and the planner problem (16.1) look similar. The planner wants marginal utilities of agents i and j to be collinear

$$\lambda_i \partial U_i(c_i) / \partial c_{i,t}(\omega^t) = \lambda_j \partial U_j(c_j) / \partial c_{j,t}(\omega^t) = \mu_t(\omega^t). \quad (16.6)$$

The optimization by individual agents given Arrow-Debreu prices $p_t^0(\omega^t)$ also leads to collinear marginal utilities

$$\frac{1}{\alpha_i} \partial U_i(c_i) / \partial c_{i,t}(\omega^t) = \frac{1}{\alpha_j} \partial U_j(c_j) / \partial c_{j,t}(\omega^t) = p_t^0(\omega^t). \quad (16.7)$$

If we select agent weights $\lambda_i = 1/\alpha_i$ in equation (16.6), we obtain the same conditions as in equation (16.7). The consumption allocations in the Arrow-Debreu equilibrium and the planner problem are thus the same. This is an implication of the First Welfare Theorem, which says that the consumption allocation in the Arrow-Debreu equilibrium is Pareto optimal. This property will ensure that marginal utilities of different agents are collinear and marginal rates of substitution across agents, as well as across planner and agents, are identical.

Sequential markets We now introduce repeated trading in consumption and asset markets, which involves assets like those we see being traded in actual asset markets instead of Arrow-Debreu securities. Suppose that agents can trade N assets and buy consumption in spot markets at every date t . Each asset n pays a dividend $d_t^n(\omega^t)$ in units of consumption at date t after history ω^t . The N -dimensional vector $d_t(\omega^t)$ contains the dividends of all N assets. The dividend stream $\{d_t(\omega^t)\}_{t, \omega^t}$ of the assets is exogenously given.

The N -dimensional vector of asset prices $p_t(\omega^t)$ describes the units of consumption needed to purchase the assets at date t after history ω^t . Our convention is that assets are traded

ex dividend: the buyer of the assets at date t will receive the dividends $d_{t+1}(\omega^{t+1})$ at date $t + 1$ after history ω^{t+1} . The asset prices $p_t(\omega^t)$ are thus quoted *ex dividend*, which means without the dividend: the asset prices do not contain the dividends paid at date t .

Agents can buy or sell assets, starting from some initial holdings. Unless stated explicitly otherwise, we assume that the agents have no assets initially, $\theta_{-1}^n = 0$ for all n . If agents buy asset n , they hold a positive number $\theta_t^n(\omega^t) > 0$ of the asset. If they sell asset n , they short $\theta_t^n(\omega^t) < 0$ many assets. Alternatively, they can choose to not have any number of asset n , $\theta_t^n(\omega^t) = 0$. The N -dimensional vector $\theta_t(\omega^t)$ collects all asset holdings. We call $\theta = \{\theta_t(\omega^t)\}_{t,\omega^t}$ a *trading strategy*.

With sequential trading, the budget set of agent i is

$$B(p, y_i) = \left\{ \{c_{i,t}(\omega^t)\}_{t,\omega^t} : \text{there is a } \theta_i \text{ such that } c_{i,t}(\omega^t) + p_t(\omega^t)^\top \theta_{i,t}(\omega^t) \right. \\ \left. \leq y_{i,t}(\omega^t) + (d_t(\omega^t) + p_t(\omega^t))^\top \theta_{i,t-1}(\omega^{t-1}) \text{ for every } t, \omega^t \in \Omega^t \right\}. \quad (16.8)$$

Spending on consumption and current asset holdings has to be less than the endowment and payoffs from previous asset holdings. These payoffs consists of dividends plus the current value of the assets. Agent i chooses a consumption plan $c_i \in B(p, y_i)$ given asset prices p and the endowment y_i to maximize utility $U_i(c_i)$.

Sequential market equilibrium A competitive equilibrium is a collection of consumption plans $c_i^* = \{c_{i,t}^*(\omega^t)\}_{t,\omega^t}$ and trading strategies $\theta_i^* = \{\theta_{i,t}^*(\omega^t)\}_{t,\omega^t}$ for each agent i and asset prices $p^* = \{p_t^*(\omega^t)\}_{t,\omega^t}$ such that (i) c_i^* is in the budget set $B(p^*, y_i)$ for given asset prices p^* and the agent's endowment y_i , and maximizes utility $U^i(c_i)$, (ii) good markets clear at every date after every history $\sum_{i \in \mathcal{I}} c_{i,t}^*(\omega^t) = \sum_{i \in \mathcal{I}} y_{i,t}^*(\omega^t)$, and (iii) asset markets clear $\sum_{i \in \mathcal{I}} \theta_{i,t}^*(\omega^t) = 0_{N \times 1}$ at every date after every history.

Here, we are assuming that assets are in *zero net supply*, the asset holdings across agents sum to zero. In some applications, we will want to assume that there is some exogenous supply $\bar{\theta}^j$ of asset j . For example, we may want to allow for a positive supply of government bonds. In this case, government bond market clearing becomes $\sum_{i \in \mathcal{I}} \theta_{i,t}^j = \bar{\theta}^j$.

Preview If agents have a rich set of assets to trade at every date t , markets are *complete*. Intuitively, agents with access to a rich set of assets can easily shift consumption units across time and states of the world. Since agents take asset prices as given, they trade assets at date t until they equate their valuation of the assets' payoffs at date $t + 1$. Since payoffs are defined as baskets of consumption units delivered in the various states of the world ω_{t+1} next period, equating asset prices is equivalent to agreeing on the value of each future consumption unit, as long as agents have enough assets to trade. Formally, markets are complete if for any sequence $\{z_t(\omega^t)\}$, there exists a trading strategy $\theta = \{\theta_t(\omega^t)\}$ with payoffs $z_t(\omega^t) = (d_t(\omega^t) + p_t(\omega^t))^\top \theta_{t-1}(\omega^{t-1}) - p_t(\omega^t)^\top \theta_t(\omega^t)$ for all histories and times $t > 0$.

With complete markets, we obtain three important results. First, the set of equilibrium allocations in the Arrow-Debreu economy and the sequential markets economy are the same. Second, the First Welfare Theorem holds, so that the equilibrium allocation is Pareto optimal. Third, there exists a representative agent. Without complete markets, all three results are lost.

The three results are key for understanding the degree to which agents' heterogeneity can matter for asset valuation. With complete markets, a single-agent model is sufficient for understanding asset prices because there is agreement among all agents on how to value consumption in each history. For heterogeneity to matter, markets have to be incomplete. With fewer assets to trade, agents will agree on the value of the assets they trade, at least in the absence of financial frictions (such as short-sale constraints). However, agents may disagree about the value of consumption in the various histories.

We will first gain intuition about these results in a version of this environment with only two dates. In the two-period economy, assets are traded at date 0 and not re-traded again at date 1. We will then study an infinite horizon economy in which consumption and assets are re-traded at every date.

16.4 Asset trading with two periods

Suppose there are S different states of the world that can happen tomorrow (at date 1). Each state ω has some probability $\pi(\omega) > 0$ with $\sum_{\omega \in \Omega} \pi(\omega) = 1$. Agent i receives an endowment $y_{i,0}$ units of consumption at date 0 and an endowment of $y_{i,1}(\omega)$ units of consumption in state ω at date 1. The agent has utility function $U_i(c_i)$ over consumption plans $c_i = \{c_{i,0}, \{c_{i,1}(\omega)\}_{\omega \in \Omega}\}$.

The agent can buy a portfolio θ of the N assets at date 0. We collect the N -dimensional payoffs $d(\omega)$ at date 1 in state $\omega \in \Omega$ in an $S \times N$ payoff matrix D for every state and asset tomorrow, as in Chapter 7 (Section 7.4). If there are Arrow securities for each state at date 1, then the payoff matrix is equal to the identity matrix, $D = I$. The N -dimensional vector p contains the asset prices at date 0, so p^n is the price of asset n in units of consumption at date 0. A portfolio θ involves spending $\sum_{n=1}^N p^n \theta^n = p^\top \theta$ units of consumption at date 0. At date 1, the agent receives portfolio payoffs $d(\omega)^\top \theta$ units of consumption in state ω at date 1. We collect the portfolio payoffs in an S -dimensional vector $D\theta$.

The agent maximizes utility

$$\max_{c_i \in B(p, y_i)} U_i(c_i) \quad (16.9)$$

where the budget set is

$$B(p, y_i) = \left\{ (c_{i,0}, \{c_{i,1}(\omega)\}_{\omega \in \Omega}) : \begin{array}{l} c_{i,0} + p^\top \theta_i \leq y_{i,0} \\ c_{i,1}(\omega) \leq y_{i,1}(\omega) + d(\omega)^\top \theta_i \end{array} \text{ for some } \theta_i \in \mathbb{R}^N \right\}. \quad (16.10)$$

Since the utility function is strictly increasing, the budget equations will bind, so they hold with equality. The budget equations for consumption in the various states ω at date 1 involve the portfolio payoff in each state.

Arbitrage We define a “free lunch,” or arbitrage opportunity, as a portfolio that is either free today, $p^\top \theta \leq 0$ and has non-negative payoffs tomorrow $D\theta \geq 0$ with one strictly positive payoff in at least one of the states, or a portfolio with a strictly negative value $p^\top \theta < 0$ that has nonnegative payoffs tomorrow $D\theta \geq 0$. An arbitrage is a strategy that would be infinitely

attractive for any agent who prefers more consumption over less. It makes sense to assume that such a strategy does not exist.

The compact way of defining an arbitrage opportunity is with the $(S + 1) \times N$ matrix

$$W = \begin{bmatrix} -p^\top \\ D \end{bmatrix}, \quad (16.11)$$

which contains the payoff at date 0 and for every state at date 1. An arbitrage is a portfolio θ such that its payoff $W\theta \geq 0$ is nonnegative and strictly positive either today or tomorrow in at least one state of the world. Markets are arbitrage-free if $D\theta \geq 0$ implies $p^\top\theta \geq 0$ and $D\theta \geq 0$ with one strict inequality implies $p^\top\theta > 0$.

Set of attainable payoffs The set of attainable payoffs is

$$\mathcal{M} = \{x \in \mathbb{R}^{S+1} : \text{there is a portfolio } \theta \in \mathbb{R}^N \text{ such that } x = W\theta\},$$

where W is the matrix in equation (16.11). If the system of asset prices p allows for arbitrage opportunities, then the set \mathcal{M} includes elements x that are non-negative and have at least one strictly positive entry.

State prices The price at date 0 of one unit of consumption in state ω at date 1 is called a state price. The contingent claim that will deliver one unit of consumption in state ω and zero in all other states is called an Arrow security.² A state price is thus an asset price, the price of an Arrow security. The absence of arbitrage implies that each state price $q(\omega)$ must be strictly positive. If the price of an Arrow security were zero or negative, one could construct a costless strategy—or even one that yields a gain today—that delivers a strictly positive payoff in state ω and zero otherwise, which constitutes an arbitrage opportunity.

Law of one price The law of one price says that two assets with identical payoffs in every state must trade at the same price. Again, if this was not true, then a strategy that buys the cheaper asset and sells the more expensive asset has zero payoffs at date 1 but makes a strictly positive payoff at date 0 — an arbitrage. For a set of state prices q , the law of one price therefore implies that the price of the n th asset is its future payoffs in the various states valued by the state prices.

$$p^n = \sum_{\omega \in \Omega} d^n(\omega) q(\omega). \quad (16.12)$$

By writing the key property of state prices in matrix form $p = D^\top q$, we can see that the vector of consumption prices

$$\hat{q} = \begin{pmatrix} 1 \\ q \end{pmatrix}, \quad (16.13)$$

²The Arrow-Debreu security that we introduced earlier is a contingent claim trading at date 0 for consumption at any date t , including $t > 1$, after history ω^t . An Arrow security is a contingent claim for *next period* consumption in one of the states of the world. Of course, the two concepts coincide in a model with only two periods. However, the statement that state prices are Arrow security prices will always be true.

which assigns a normalized price of 1 to consumption at date 0 and states prices q to consumption in the various states tomorrow, is orthogonal to the set of attainable payoffs \mathcal{M} . In other words, for any $x \in \mathcal{M}$, we have that $x^\top \hat{q} = 0$. This follows from (16.12) because for asset n we have $-p^n + \sum_{\omega \in \Omega} d^n(\omega) q(\omega) = 0$ and the attainable payoffs in \mathcal{M} are simply linear combinations of the payoffs from trading the individual assets.

Fundamental theorem of asset pricing For given asset prices p , the absence of arbitrage is equivalent with the existence of state prices. In general, these state prices may not be unique, so we may have $p = D^\top q = D^\top q'$, where both q and q' are state prices. If markets are complete ($\text{rank}(D) = S$), state prices are unique. If the number of assets equals the number of states, $N = S$, the state prices can be computed by inverting the payoff matrix, $q = (D^\top)^{-1} p$. If there are more assets than states, $N > S$, there are redundant assets, and we can eliminate some of them to obtain a $S \times S$ payoff matrix D with rank S . The state prices can then again be computed by inverting the payoff matrix, $q = (D^\top)^{-1} p$. If $N < S$, the rank of the payoff matrix is not S and markets are incomplete. If this is the case, we may have many state prices.

Below we sketch the proof. For the equivalence of the two statements, we need to show that the existence of state prices implies no arbitrage. Suppose we have some state prices $q \in \mathbb{R}_{++}^S$ for which $p = D^\top q$, then the value of any portfolio $\theta \in \mathbb{R}^N$ satisfies $p^\top \theta = q^\top D\theta$. If $D\theta \geq 0$, then $p^\top \theta = q^\top D\theta \geq 0$ since state prices are strictly positive. If $D\theta > 0$, then $p^\top \theta = q^\top D\theta > 0$. Thus, there are no arbitrage opportunities given asset prices q .

The other direction of the proof—showing that no arbitrage implies the existence of state prices—is more work, but is instructive as to what state prices do. The key step in the proof separates the set of attainable payoffs \mathcal{M} from the set of arbitrage opportunities $\mathcal{K} = \mathbb{R}_+^{S+1}$ as depicted in Figure 16.2 for the case of $S = 1$ tomorrow. The horizontal axis represents payoffs at date 0, while the vertical axis measures payoffs at date 1. To show that \mathcal{M} and \mathcal{K} only intersect at the origin, $\mathcal{M} \cap \mathcal{K} = \{0\}$, the proof uses a separating hyperplane theorem. The theorem implies that there is a nonzero linear function $F : \mathbb{R}^{S+1} \rightarrow \mathbb{R}$ such that $F(z) > 0$ for all nonzero $z \in \mathcal{K}$ and $F(x) = 0$ for all $x \in \mathcal{M}$.³ This means there are some strictly positive coefficients $\tilde{q} = (\tilde{q}_0, \tilde{q}_1) \in \mathbb{R}_{++}^{1+S}$, that represent the nonzero linear function: $F(x_0, x_1) = \tilde{q}_0 x_0 + \tilde{q}_1 x_1$, where $x_0 \in \mathbb{R}$ and $x_1 \in \mathbb{R}^S$. The vector \tilde{q} , depicted as an arrow in Figure 16.2, is orthogonal to the set \mathcal{M} , because if we evaluate the function at an element in \mathcal{M} , we get $-\tilde{q}_0 p^\top \theta + \tilde{q}_1^\top D\theta = 0$ for all $\theta \in \mathbb{R}^N$. The vector \tilde{q} shown as an arrow in Figure 16.2 thus contains consumption prices (16.13), but does not necessarily assign a price of one to consumption at date 0. However, since $\tilde{q}_0 > 0$, we can normalize the price of consumption at date 0 on the horizontal axis to one, and find the state prices $q = \tilde{q}_1 / \tilde{q}_0$ on the vertical axis.

The essence of this proof is the existence of a linear function $F : \mathbb{R}^{S+1} \rightarrow \mathbb{R}$ with strictly positive coefficients \tilde{q} that values all assets and portfolios of assets. When we evaluate this

³The separating hyperplane theorem is for cones. A cone is a set X for which if $x \in X$ implies that $\lambda x \in X$ for any $\lambda > 0$. Both \mathcal{M} and \mathcal{K} are closed convex cones in \mathbb{R}^{S+1} that intersect at the origin. The separating hyperplane theorem implies that there exists a nonzero linear function $F : \mathbb{R}^{S+1} \rightarrow \mathbb{R}$ such that $F(z) < F(x)$ for all z in \mathcal{M} and $F(x) > 0$ for all nonzero $x \in \mathcal{K}$. However, since \mathcal{M} is a linear space, $-z$ is also in \mathcal{M} . The inequality thus implies that $F(z) = 0$ for all z in \mathcal{M} since F is nonzero. Otherwise we could always start with $F(z) < F(x)$ and get $F(-z) > F(x)$.

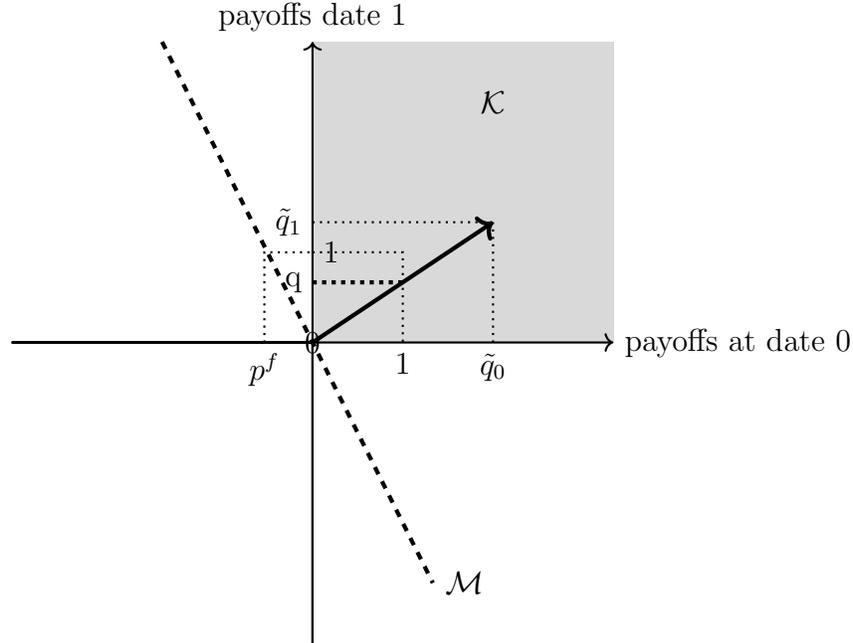


Figure 16.2: The set of attainable payoffs \mathcal{M} and the set of arbitrages \mathcal{K} separated by a linear function $F(x) = \tilde{q}_0 x_0 + \tilde{q}_1^\top x_1 = 0$ for $x \in \mathcal{M}$ with coefficients \tilde{q} that are orthogonal to the elements of \mathcal{M} .

linear function for any attainable payoff, the function ties the date-0 value of any traded asset to the weighted sum of its date-1 payoffs, where the weights are state prices $q = \tilde{q}_1/\tilde{q}_0$. It is important to note that while the price p of any traded asset is given, the state prices q may not be unique. The theorem thus guarantees the existence of some state prices that justify the given prices p , but does not imply that the state prices themselves are unique, unless markets are complete. The following two examples illustrate this principle.

As an example, suppose there are two states, $S = 2$, and two assets: a stock and a bond that trade at given prices $p = (p^s, p^b)^\top$ and have a payoff matrix

$$D = \begin{pmatrix} \delta_1 & 1 \\ \delta_2 & 1 \end{pmatrix}, \quad \delta_1 > 1 > \delta_2 > 0.$$

We impose two arbitrage restrictions. First, asset prices p have to be strictly positive, since the payoffs are strictly positive in both states. Second, we impose that the return on the stock cannot dominate that on the bond in all states (and vice versa): $\delta_1/p^s > 1/p^b > \delta_2/p^s$. The return is defined as the asset's payoff at date 1 divided by its price at date 0.

To determine state prices, we note that markets are complete markets, because $\text{rank}(D) = S = 2$. Moreover, we have as many states as assets $S = N$. In this case, we can recover unique state prices from asset prices by simply inverting the transpose of the payoff matrix:

$$q = (D^\top)^{-1} p = \frac{1}{\delta_1 - \delta_2} \begin{pmatrix} 1 & -\delta_2 \\ -1 & \delta_1 \end{pmatrix} \begin{pmatrix} p^s \\ p^b \end{pmatrix} = \frac{1}{\delta_1 - \delta_2} \begin{pmatrix} p^s - \delta_2 p^b \\ \delta_1 p^b - p^s \end{pmatrix}.$$

The state prices must be strictly positive, so we need the arbitrage conditions $\delta_1 p^b > p^s > \delta_2 p^b > 0$ as before.

The set of attainable payoffs tomorrow $D\theta$ is the \mathbb{R}^2 . For every payoff $x \in \mathbb{R}^2$, there is a portfolio $\theta = D^{-1}x$ that will generate the payoff x . The value of the payoff x is the cost of that portfolio

$$p^\top \theta = p^\top D^{-1}x = \frac{1}{\delta_1 - \delta_2} (p^s - p^f \delta_2, p^f \delta_1 - p^s)^\top x.$$

Under the arbitrage conditions we imposed, the value of any nonnegative vector x with at least one strictly positive entry is strictly positive.

As another example, suppose again $S = 2$. However, we now only have one asset, a stock, with the same payoffs as before. D consists of a single column:

$$D = \begin{pmatrix} \delta_1 \\ \delta_2 \end{pmatrix}.$$

We have incomplete markets, because $\text{rank}(D) = 1 < S$. We are now restricted to payoffs $[\delta_1 \ \delta_2]^\top \theta$ along a line in \mathbb{R}^2 . Suppose $\delta_1 > \delta_2 > 0$. No arbitrage then requires that the stock price is positive. There are many $q \in \mathbb{R}_{++}^2$ so that $p^s = \sum_{\omega \in \Omega} d^s(\omega) q(\omega)$, where $d^s(\omega)$ is the payoff of the stock in state ω . Given our matrix D , we have $p^s = \delta_1 q(1) + \delta_2 q(2)$. For example, suppose $\delta_1 = 2$, $\delta_2 = 0.5$, and $p^s = 1$. Any $q(1) \in (0, 0.5)$ and $q(2) = 2 - 4q(1)$ are valid state prices, because they solve the single equation $p^s = 2q(1) + 0.5q(2)$.

Connection between no arbitrage and individual agent optimality If there is a solution to the individual agent problem, that implies that there is no arbitrage opportunities. Suppose there were arbitrage opportunities available. In this case, agents can add a free lunch to their budget sets. Therefore, there cannot be an optimal solution since any candidate optimal portfolio θ^* can always be improved upon with another free lunch on top of it.

The converse statement also holds: If the utility function is continuous and there are no arbitrage opportunities, then there is a solution to the individual agent problem. The key here is to apply the fundamental theorem that implies that no arbitrage implies the existence of state prices to rewrite the budget set (16.10) using the state prices. This rewriting of the budget set helps to show that it is compact, so that the maximum theorem applies and guarantees the existence of a solution to the optimization problem.

Rewriting the budget constraint If there are no arbitrage opportunities, we have state prices q that satisfy $p = D^\top q$. We use these state prices to rewrite the budget equation for date 0 and obtain the net trade away from the endowment of agent i

$$c_{i,0} - y_{i,0} \leq -p^\top \theta_i = -q^\top D\theta_i = -\sum_{\omega \in \Omega} q(\omega) d(\omega)^\top \theta_i.$$

For the budget equation at date 1, agent i 's net trade in state ω is

$$c_{i,1}(\omega) - y_{i,1}(\omega) \leq d(\omega)^\top \theta_i.$$

We can multiply the net trade in state ω by the state price $q(\omega)$ and add up all net trades at date 0 and the different states ω at date 1. We then obtain a single budget constraint

$$c_{i,0} - y_{i,0} + \sum_{\omega \in \Omega} q(\omega) (c_{i,1}(\omega) - y_{i,1}(\omega)) \leq 0.$$

If markets are arbitrage-free, the budget set $B(p, y_i)$ in equation (16.10) can thus be rewritten as

$$B(p, y_i) = \left\{ (c_{i,0}, \{c_{i,1}(\omega)\}_{\omega \in \Omega}) : c_{i,0} - y_{i,0} + \sum_{\omega \in \Omega} q(\omega) (c_{i,1}(\omega) - y_{i,1}(\omega)) \leq 0 \quad (16.14) \right. \\ \left. \text{and } \{c_{i,1}(\omega) - y_{i,1}(\omega)\}_{\omega \in \Omega} \in \text{span}(D) \right\}.$$

The span restriction ensures that the agent's net trades can be achieved with the given assets in the economy. This restriction can reduce the budget set compared to the budget set in an Arrow-Debreu economy with consumption at date 0 as the numeraire and Arrow-Debreu prices q for consumption at date 1. If the span of the payoff matrix is a strict subspace of \mathbb{R}^S , the agent cannot freely obtain all consumption bundles even if they are affordable given the state prices q . The reason is that given the asset payoff matrix D , there may not be enough assets to achieve the net trades $c_{i,1}(\omega) - y_{i,1}(\omega)$ away from the endowment $y_{i,1}(\omega)$ that are required to get the desired bundle. Moreover, state prices are not unique.

As an example, suppose the agent does not care about consumption at date 0 and also has no endowment at date 0. The budget constraint for date 0 is $p^\top \theta \leq 0$ and for state ω at date 1 is $c_{i,1}(\omega) - y_{i,1}(\omega) \leq d(\omega)^\top \theta_i \leq 0$. With two states, $S = 1$, we can draw the budget set with complete markets as the light gray shaded area in Figure 16.3. The agent can reach any point below or on the dashed budget line going through the endowment $y_{i,1}$. With incomplete markets, the agent may be more constrained. For example, with only a riskless bond, the agent can add or subtract one unit of consumption in both states, and thus can only reach points on the thick black line with unit slope starting at the endowment point.

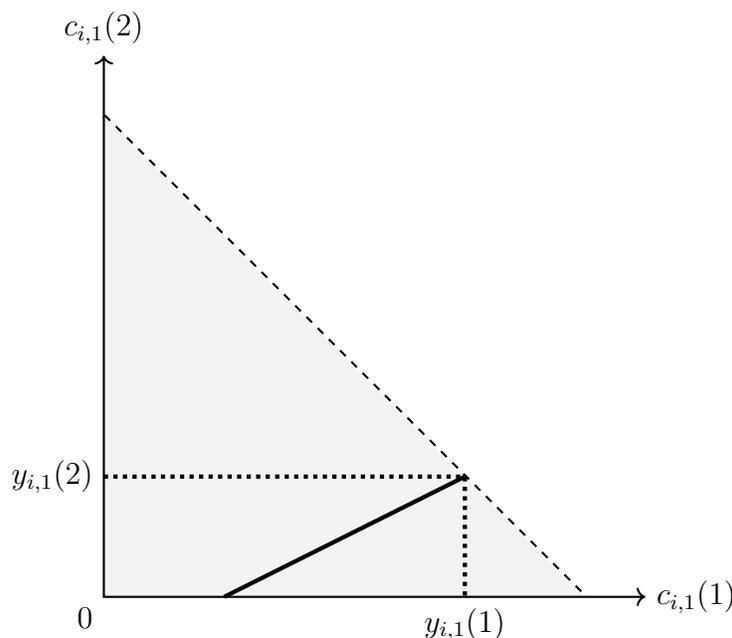


Figure 16.3: Budget sets with complete markets as dashed line and with only a riskless bond as solid line

Optimality conditions for asset holdings The Lagrangian for the single agent problem is

$$\mathcal{L} = U_i(c_i) - \psi_i(c_{i,0} + p^\top \theta_i - y_{i,0}) - \sum_{\omega \in \Omega} \eta_i(\omega) (c_{i,1}(\omega) - y_{i,1}(\omega) - d(\omega)^\top \theta_i),$$

where $\eta_i(\omega)$ is the Lagrange multiplier on the budget equation in state ω . The first-order condition for consumption at date 0 and in state ω at date 1 are

$$\frac{\partial U_i(c_i)}{\partial c_{i,0}} = \psi_i > 0 \text{ and } \frac{\partial U_i(c_i)}{\partial c_{i,1}(\omega)} = \eta_i(\omega) > 0.$$

The first-order condition for holdings θ_i^n of asset n is

$$-\psi_i p^n + \sum_{\omega \in \Omega} \eta_i(\omega) d^n(\omega) = 0.$$

Subjective state prices Combining, we obtain

$$p^n = \sum_{\omega \in \Omega} d^n(\omega) q_i^n(\omega), \text{ where } q_i^n(\omega) = \frac{\partial U_i(c_i) / \partial c_{i,1}(\omega)}{\partial U_i(c_i) / \partial c_{i,0}} \quad (16.15)$$

are state prices that satisfy the law of one price (16.12).

As an example, suppose agents have the same beliefs π and the same discount factor β . The utility function (16.9) takes the form

$$U_i(c_i) = \frac{c_{i,0}^{1-\gamma}}{1-\gamma} + \beta \sum_{\omega \in \Omega} \pi(\omega) \frac{c_{i,1}(\omega)^{1-\gamma}}{1-\gamma},$$

where γ is the coefficient of relative risk aversion, and $\pi(\omega)$ is the probability of state s . The resulting state prices are $q_i^n(\omega) = \beta \pi(\omega) (c_{i,1}(\omega) / c_{i,0})^{-\gamma}$, where $c_{i,1}(\omega) / c_{i,0}$ is agent i 's consumption growth rate from date 0 to state ω at date 1. A risk-averse agent is thus willing to pay a high price for consumption in state s when the consumption growth rate in that state is low. If the agent is risk neutral, $\gamma = 0$, the state prices $\beta \pi(\omega)$ just reflect the discounted probabilities of the states at date 1.

Disagreement about subjective state prices The agent index i is an important reminder that state prices are subjective. When some markets are missing, agents cannot freely trade consumption claims at date 0. The missing markets allow agents to disagree about the price $q_i(\omega)$ of a unit of consumption good in state ω at date 1: their marginal rates of substitution of consumption at date 0 for consumption in state ω at date 1 may be different. The subjective state price $q_i(\omega)$ of agent i may be different from the subjective state price $q_j(\omega)$ of agent j . In the example with expected power utility, the consumption growth rates may differ across agents. Moreover, agents can differ in their coefficients of relative risk aversion γ_i , discount factors β_i or beliefs π_i .

The optimality conditions for asset holdings lead to agreement among agents *about asset prices*. The first-order conditions (16.15) make sure that agents agree about the valuation of

the linear combination of payoffs. This highlights that assets promise bundles of consumption goods at date 1. By optimally trading these assets, agents will agree on the value of the bundle, but may disagree about the individual prices of consumption in the various states of the world. Going back to Example 2, agents i and j will agree about the stock price $p^s = 2q_i i(1) + 0.5q_i(2) = 2q_j(1) + 0.5q_j(2) = 1$, for example because they value consumption in each state according to the individual-specific state prices $q_i = (1/4, 1)^\top$ and $q_j = (1/3, 2/3)^\top$. With complete markets, there are enough assets with prices p^n that agents trade, that the first-order conditions (16.15) will lead them to agree on state prices. Put differently, at date 0 agents will assign the same value to extra units of consumption in the S states at date 1, $q_i = q_j$.

Agreement with complete markets If the span of the payoff matrix D is the entire \mathbb{R}^S , markets are complete. In this case, the span restriction in the budget set does not bind. The budget set with sequential trading is then identical to the budget set in an Arrow-Debreu economy with consumption at date 0 as the numeraire and state prices q for consumption at date 1. Moreover, the state prices q that determine the asset prices $p = D^\top q$ in the sequential trading economy are unique. The unique state prices solve $q = (D^\top)^{-1} p$ if $S = N$. If $N > S$, there are redundant assets, and we can compute unique state prices based on a subset of these assets that span the entire \mathbb{R}^S . Agents thus agree on the value of consumption in different states.

Stochastic discount factor If no arbitrage holds, there are state prices q , which we can use to define a stochastic discount factor $M(\omega) = q(\omega) / \pi(\omega)$. The price of asset n satisfies $p^n = \sum_{\omega \in \Omega} d^n(\omega) q(\omega) = \mathbb{E}[d^n(\omega) M(\omega)]$, while the return on the asset $R^n(\omega) = d^n(\omega) / p^n$ satisfies $1 = \mathbb{E}[R^n(\omega) M(\omega)]$.

The stochastic discount factor is a measure of how hungry the agent is for consumption. When $M(\omega)$ is high, the agent is hungry and desperate for more consumption. Any payoffs in those hungry “high M ” states contribute a lot to the value of the asset, while payoffs in “low M ” states do not contribute much. In the case of expected power utility, the stochastic discount factor is $M(\omega) = \beta (c_1(\omega) / c_0)^{-\gamma}$, which illustrates that the agent is hungry in states of the world ω in which the growth rate of consumption is low. If the agent is risk neutral, the stochastic discount factor is a constant, $M(\omega) = \beta$ in all states ω ; the agent thus assigns equal value to consumption in the various states of the world.

Risk-neutral pricing A risk-free bond pays one unit of consumption in every state of the world at date 1. If agents are risk neutral, the bond price satisfies $p^b = 1/R^f = \mathbb{E}[M(\omega)] = \beta$. The price of any other asset n is given by its expected discounted payoffs

$$p^n = \mathbb{E}[d^n(\omega) M(\omega)] = \frac{\mathbb{E}[d^n(\omega)]}{R^f}, \quad (16.16)$$

where the discount rate is the risk-free interest rate.

Risk premium More generally, without assuming risk neutrality, the price of asset n satisfies

$$p^n = \frac{\mathbb{E}[d^n(\omega)]}{R^f} - \text{risk adjustment},$$

where the risk adjustment reduces the price of the asset. What determines the risk adjustment? Using the definition of covariance,

$$\begin{aligned} p^n &= \mathbb{E}[d^n(\omega) M(\omega)] \\ &= \mathbb{E}[d^n(\omega)] \mathbb{E}[M(\omega)] + \text{Cov}(d^n(\omega), M(\omega)) \\ &= \frac{\mathbb{E}[d^n(\omega)]}{R^f} - (-\text{Cov}(d^n(\omega), M(\omega))), \end{aligned}$$

we can see that the risk adjustment is minus the covariance of the asset's payoffs with the stochastic discount factor. The formula highlights that an insurance asset is particularly valuable: its payoffs are high precisely when the agent is hungry, the stochastic discount factor is high. An asset like a stock market index has payoffs that negatively covary with the stochastic discount factor: the stock market does well in economic booms, when consumption tends to be high, and badly in economic recessions, when consumption tends to be low. Compared to risk neutral pricing, this negative covariance is a positive risk adjustment which lowers the stock price.

An important insight is that not all risk matters for asset valuation. Only systematic risk, defined as risk that covaries with the stochastic discount factor, matters for asset values. Unsystematic risk does not affect asset values. Put differently, variance is not the appropriate measure of risk in financial markets, *covariance* is what matters.

Alternatively, we can rewrite the equation $1 = \mathbb{E}[R^n(\omega) M(\omega)]$ using the definition of covariance as

$$1 = \mathbb{E}[R^n(\omega) M(\omega)] + \text{Cov}(R^n(\omega), M(\omega)) \iff \mathbb{E}[R^n(\omega)] - R^f = -\frac{\text{cov}(R^n(\omega), M)}{\mathbb{E}[M]}. \quad (16.17)$$

The risk premium, defined as the expected excess return of the asset over the risk-free rate, is determined by minus the covariance of the asset's return with the pricing kernel (divided by the expected stochastic discount factor, which is a number close to one.) An insurance asset has positive covariance with the pricing kernel and therefore has a negative premium: the agent is willing to pay an insurance premium to get insurance. An asset like a stock has negative covariance with the pricing kernel and therefore has a positive equity premium: the agent has to be compensated with a high mean return to hold the asset.

Risk neutral probabilities We can define $\pi^*(\omega) = q(\omega) / \sum_{\omega' \in \Omega} q(\omega')$ and obtain probabilities that are strictly positive and sum up to one. The sum of state prices in the denominator of the probabilities, $\sum_{\omega' \in \Omega} q(\omega') = p^b = 1/R^f$, represent the value of one unit of consumption in every state next period, which is simply the price of a riskless bond. Therefore, asset prices are equal to the value of their payoff based on the state prices and, inserting $q(\omega) = \pi^*(\omega) / R^f$, we have

$$p^n = \sum_{\omega \in \Omega} d^n(\omega) q(\omega) = \frac{\sum_{\omega \in \Omega} d^n(\omega) \pi^*(\omega)}{R^f} = \frac{\mathbb{E}^*[d^n(\omega)]}{R^f}, \quad (16.18)$$

where the expectation \mathbb{E}^* is computed using probabilities π^* . The formula (16.18) looks like equation (16.16) which we obtained by assuming that agents are risk neutral. However, the probability measure π^* is not the actual probability measure π . Instead, it is a distorted probability measure under which asset prices are determined *as if* agents were risk-neutral. The probability measure π^* is therefore called risk-neutral probability. On Wall Street, financial economists often do not specify utility functions and determine equilibrium consumption streams to compute state prices or the pricing kernel. Instead, they back out the risk neutral probability π^* from observed asset prices and then use it, for example, to determine the price of derivatives. This chapter shows that state prices, pricing kernels, and risk-neutral probabilities are equivalent labels to talk about the determination of asset prices. Moreover, the risk-neutral probability π^* is unique if markets are complete, otherwise there may be many such probabilities, one for each possible state price.

Representative agent With complete markets, the equilibrium asset prices of the heterogeneous-agent economy are identical to those in a representative-agent economy. More precisely, suppose a consumption allocation $\{c_i^*\}_{i \in \mathcal{I}}$ and portfolio holdings $\{\theta^{*i}\}_{i \in \mathcal{I}}$ together with asset prices p^* are a competitive equilibrium for the heterogeneous-agent economy with complete markets, then the consumption allocation $c = y$ and zero portfolio holdings together with the same asset prices p^* are a competitive equilibrium for the representative-agent economy with endowment $y = \sum_{i \in \mathcal{I}} y_i$ and preferences $U_\lambda \equiv \sum_{i \in \mathcal{I}} \lambda_i U_i$ where $\lambda_i = 1/\alpha_i$ and α_i is the Lagrange multiplier on the Arrow-Debreu budget constraint in the competitive equilibrium. The key to the proof of this result is to exploit the connection between the planner's problem and the decentralized equilibrium discussed in section 16.3. The decentralized equilibrium coincides with the solution to the planner's problem with the particular weights λ_i on the households. We can then interpret the planner's objective U_λ as the utility function of a representative agent.

With complete markets, we can read off the (unique) equilibrium state prices from the marginal rate of substitution of the representative agent. When we choose agent weights $\lambda_i = 1/\alpha_i$, combine the first-order conditions of individual agents (16.4) with the planner's first-order conditions (16.1) and envelope condition (16.2), we can write

$$q(\omega) = \frac{\partial U_\lambda(y) / \partial y_1(\omega)}{\partial U_\lambda(y) / \partial y_0}.$$

16.5 Dynamic asset trading

To define an arbitrage and to write the set of attainable payoffs, it will be useful to have a notation for the payoffs that can be generated by a particular trading strategy.

Payoffs generated by trading strategy The payoffs $\{x_t(\omega^t)\}_{t, \omega^t}$ generated by a particular trading strategy $\theta = \{\theta_t(\omega^t)\}$ are defined as

$$x_t(\omega^t) = \sum_{n=1}^N (p_t^n(\omega^t) + d_t^n(\omega^t)) \theta_{t-1}^n(\omega^{t-1}) - p_t^n(\omega^t) \theta_t^n(\omega^t).$$

With no initial asset holdings ($\theta_{-1}(\omega^0) = 0_{N \times 1}$), the initial payoff is minus the spending on the first portfolio, $x_0(\omega^0) = -p_0(\omega^0)^\top \theta_0(\omega^0)$. At date 1 after history ω^1 , the initial portfolio $\theta_0(\omega^0)$ has payoff $(p_1(\omega^1) + d_1(\omega^1))^\top \theta_0(\omega^0)$ and the agent spends $p_1(\omega^1)^\top \theta_1(\omega^1)$ on the new portfolio, and so on.

Arbitrage Since agents trade at given asset prices p , an important question is whether assets offer a free lunch. An arbitrage is a trading strategy that delivers such a free lunch; these are units of consumption at some date and after some history for which agents don't have to pay anything. Formally, we define an arbitrage as a trading strategy θ that generates nonnegative payoffs $x_t(\omega^t) \geq 0$ for all dates t and histories $\omega^t \in \Omega^t$, with one strict inequality.

Set of attainable payoffs By forming a portfolio θ of assets, agents can obtain payoffs in the set

$$\mathcal{M} = \left\{ \{x_t(\omega^t)\}_{t, \omega^t} : \text{there is a trading strategy } \theta \text{ that generates } x \right\}.$$

It will depend on the set of assets whether we can generate arbitrary payoffs x . We will say that *markets are complete* if any payoff stream $\{x_t(\omega^t)\}_{t \geq 0, \omega^t}$ starting at date 1 can be generated at some cost at date 0. Otherwise markets are *incomplete*.

State prices The price at date t after history ω^t of consumption at date $t+1$ in state ω_{t+1} is denoted $q_{t+1}(\omega_{t+1}|\omega^t)$. This price is the value of a contingent claim that pays out one unit of consumption tomorrow, which is called Arrow security.

Law of one price, fundamental theorem, and individual optimality The properties we established in the two-period settings carry over to the dynamic trading environment. The law of one price implies that the price of asset n can be written as its payoffs valued by the state prices

$$p_t^n(\omega^t) = \sum_{\omega_{t+1}} (d_{t+1}^n(\omega_{t+1}) + p_{t+1}^n(\omega_{t+1})) q_{t+1}(\omega_{t+1}|\omega^t).$$

The fundamental theorem of asset pricing states that no arbitrage is equivalent to the existence of state prices $q_{t+1}(\omega_{t+1}|\omega^t)$ that satisfy the previous equation. The proof of the theorem involves a separating hyperplane theorem that separates the set of attainable payoffs \mathcal{M} from the set of arbitrage opportunities $\mathcal{K} = \{z_t(\omega^t) \text{ with } z_t(\omega^t) \geq 0\}$. Moreover, if there is a solution to the individual agent problem, that implies there are no arbitrages. Conversely, if the utility function is continuous and there are no arbitrages, there exists a solution to the individual agent problem.

Optimality conditions The Lagrangian for the single agent problem is

$$\begin{aligned} \mathcal{L} = & U_i(c_i) - \sum_{t, \omega^t} \eta_{i,t}(\omega^t) \left(c_{i,t}(\omega^t) + p_t(\omega^t)^\top \theta_{i,t}(\omega^t) \right. \\ & \left. - y_{i,t}(\omega^t) - (d_t(\omega^t) + p_t(\omega^t))^\top \theta_{i,t-1}(\omega^{t-1}) \right). \end{aligned}$$

The first-order condition for consumption and holdings of asset n at date t after history ω^t are

$$\partial U_i(c_i) / \partial c_{i,t}(\omega^t) = \eta_{i,t}(\omega^t)$$

and

$$\eta_{i,t}(\omega^t) p_t^n(\omega^t) - \sum_{\omega^{t+1}} \eta_{i,t+1}(\omega^{t+1}) (d_{t+1}^n(\omega^{t+1}) + p_{t+1}^n(\omega^{t+1})) = 0.$$

Subjective state prices Combining, we generalize (16.15) to the multiperiod case

$$p_t^n(\omega^t) = \sum_{\omega_{t+1}} (d_{t+1}^n(\omega^{t+1}) + p_{t+1}^n(\omega^{t+1})) q_{i,t+1}(\omega_{t+1}|\omega^t), \quad (16.19)$$

where

$$q_{i,t+1}(\omega_{t+1}|\omega^t) = \frac{\partial U_i(c_i) / \partial c_{i,t+1}(\omega^{t+1})}{\partial U_i(c_i) / \partial c_{i,t}(\omega^t)}$$

and

$$M_{i,t+1}(\omega_{t+1}|\omega^t) = \frac{q_{i,t+1}(\omega_{t+1}|\omega^t)}{\pi_{t+1}(\omega_{t+1}|\omega^t)}.$$

Expected power utility Suppose the utility function takes the form

$$U_i(c_i) = \sum_{t,\omega^t} \pi_t(\omega^t) \beta^t \frac{c_{i,t}(\omega^t)^{1-\gamma}}{1-\gamma}. \quad (16.20)$$

The state price is

$$q_{i,t+1}(\omega_{t+1}|\omega^t) = \frac{\pi_{t+1}(\omega^{t+1}) \beta^{t+1}(\omega^{t+1}) c_{i,t+1}(\omega^t)^{-\gamma}}{\pi_t(\omega^t) \beta^t(\omega^t) c_{i,t}(\omega^t)^{-\gamma}} = \pi_{t+1}(\omega_{t+1}|\omega^t) \frac{\beta c_{i,t+1}(\omega^t)^{-\gamma}}{c_{i,t}(\omega^t)^{-\gamma}}. \quad (16.21)$$

Agreement or disagreement about state prices With complete markets, any payoff stream can be generated by trading the assets in the economy. When markets are complete, there are many equations (16.19) that determine asset prices as payoffs times the state prices. In particular, we can choose a full set of Arrow securities to complete markets and see that state prices, as well as the stochastic discount factor, will be unique. When markets are incomplete, there are fewer equations (16.19) than the number of state prices, and so state prices will not be unique.

Returns The return on asset n is defined as

$$R_{t+1}^n(\omega^{t+1}) = \frac{d_{t+1}^n(\omega^{t+1}) + p_{t+1}^n(\omega^{t+1})}{p_t^n(\omega^t)}. \quad (16.22)$$

When M is a stochastic discount factor, the pricing equation (16.19) implies that the return satisfies the Euler equation

$$\mathbb{E} [M_{t+1}(\omega_{t+1}|\omega^t) R_{t+1}^n(\omega^{t+1}) | \omega^t] = 1. \quad (16.23)$$

Using the definition of conditional covariance, writing the expectation conditional on history ω^t as $\mathbb{E}_t[\cdot]$, and suppressing the dependence of returns and stochastic discount factor on history, we obtain

$$\mathbb{E}_t [M_{t+1}] \mathbb{E}_t [R_{t+1}^n] + \text{Cov}_t (M_{t+1}, R_{t+1}^n) = 1,$$

implying

$$\mathbb{E}_t [R_{t+1}^n] - R_{t+1}^f = -\frac{\text{Cov}_t (M_{t+1}, R_{t+1}^n)}{\mathbb{E}_t [M_{t+1}]}.$$

The last equation determines the conditional expected excess return of asset n by the conditional covariance of its return with the stochastic discount factor. The difference to equation (16.17) is that the conditional covariance matters, which may vary over time. As a consequence, there will be times in which the conditional expected excess return is high, while it will be low at other times.

16.6 The equity premium puzzle and risk-free rate puzzle

We now want to understand the quantitative properties of equilibrium returns in a model with sequential markets trading. For this purpose, we make specific assumptions on preferences and endowments. In particular, we assume that there is a single agent with expected power utility (16.20) who receives an endowment stream y . We now add assets in zero net supply to this economy and solve for equilibrium asset prices such that the single agent chooses not to buy or sell these assets in equilibrium. Put differently, the agent chooses not to trade at given prices and just consumes the endowment.

More formally, a competitive equilibrium in this model consists of consumption c^* , a trading strategy θ^* and asset prices p^* such that (i) c^* solves the agent's optimization problem given asset prices p^* , (ii) goods markets clear $c^* = y$, and (iii) asset markets clear $\theta^* = 0$. As long as markets are complete, the aggregation result says that the implications of this specification will be identical to those of a heterogeneous agent model in which agents have utility (16.20) and the aggregate endowment is y . An important lesson is that the quantitative implications of the single-agent model thus apply in much more general environments.

An equivalent formulation of this environment is the *Lucas tree model*. This formulation views stocks as trees that produce fruits every period. The fruits are the dividends that the owner of the trees receives. This is an interesting example, where some of the initial holdings of assets θ_{-1} are nonzero: in period 0, the single agent already owns the trees and can then buy or sell shares of the trees, which are in non-zero net supply. We normalize the supply of trees to one. There is no endowment. Given equilibrium prices p^* , the single agent chooses to hold the single tree $\theta^{s,*} = 1$. Other assets, such as a one-period riskless bond, are in zero net supply and the single agent chooses to hold zero given p^* . In this formulation, equilibrium consumption equals the dividend $c^* = d$.

In the endowment economy, we assume that the endowment y has a growth rate that follows a two-state Markov chain with transition matrix Π . The equivalent Lucas tree model

formulation assumes that dividends d follow this Markov chain:

$$y_{t+1} = y_t G_{t+1}, \text{ where } G_{t+1} \in \{G_L, G_H\} \text{ with transition matrix } \Pi = \begin{pmatrix} \phi & 1 - \phi \\ 1 - \phi & \phi \end{pmatrix}.$$

The growth rate can either be low, G_L , or high, $G_H > G_L$. The transition matrix of the Markov chain is symmetric, so that the stationary probability π , which satisfies $\pi = \Pi^\top \pi$, spends half the time in each state, $\pi_1 = \pi_2 = 1/2$.

We want to think of the aggregate stock market as a claim to aggregate endowment, $d_t = y_t$, which is in zero net supply. The value of the stock market is therefore

$$p_t^s = \mathbb{E}_t [M_{t+1} (y_{t+1} + p_{t+1}^s)], \text{ with } M_{t+1} = \beta (c_{t+1}/c_t)^{-\gamma}. \quad (16.24)$$

In equilibrium, consumption equals the endowment, $c = y$, which means the growth rate of consumption equals $c_{t+1}/c_t = G_{t+1}$. Since the economy is growing, we solve for the stationary price-dividend ratio $v_t := p_t^s/y_t$

$$\frac{p_t^s}{y_t} = \mathbb{E}_t \left[M_{t+1} \frac{y_{t+1} (1 + p_{t+1}^s/y_{t+1})}{y_t} \right] \quad \text{or} \quad v_t = \mathbb{E}_t [\beta G_{t+1}^{1-\gamma} (1 + v_{t+1})].$$

Since we can substitute recursively for the price-dividend ratio v_{t+1} on the right-hand side, and the growth rate is a Markov process, the solution for v will be a function of only the current growth rate G_t .⁴ Using the symbols of the Markov chain, and writing v_i for the price-dividend ratio in state i , we get

$$v_i = \sum_{j=1}^2 \Pi_{ij} \beta G_j^{1-\gamma} (1 + v_j), \text{ or in matrix form } \mathbf{v} = \beta A (\mathbf{1} + \mathbf{v}),$$

where the matrix A is given by

$$A = \begin{bmatrix} \Pi_{11} G_1^{1-\gamma} & \Pi_{12} G_2^{1-\gamma} \\ \Pi_{21} G_1^{1-\gamma} & \Pi_{22} G_2^{1-\gamma} \end{bmatrix}.$$

We can therefore solve for the vector $\mathbf{v} = (I_{2 \times 2} - \beta A)^{-1} \beta A \mathbf{1}$ by matrix inversion and obtain the price-dividend ratios in every state of the Markov chain.

We also want to solve for the (net) return on stocks

$$r_{t+1}^s = \frac{y_{t+1} + p_{t+1}^s - p_t^s}{p_t^s} = \frac{y_{t+1} (1 + v_{t+1})}{p_t^s} - 1 = \frac{(y_{t+1}/y_t) (1 + v_{t+1})}{v_t} - 1.$$

With our Markov chain, we want to solve

$$r_{ij}^s = g_j \frac{(1 + v_j)}{v_i} - 1 \quad (16.25)$$

⁴With recursive substitution, we obtain

$$v_t = \mathbb{E}_t [\beta G_{t+1}^{1-\gamma} + \beta^2 G_{t+1}^{1-\gamma} G_{t+2}^{1-\gamma} + \beta^3 G_{t+1}^{1-\gamma} G_{t+2}^{1-\gamma} G_{t+3}^{1-\gamma} \dots].$$

The conditional expected value of the return uses the transition probability Π_{ij} from state i to state j , while the unconditional expected value of the return uses the stationary probability π_i of state i

$$\mathbb{E}_i [r_{ij}^s] = \sum_{j=1}^2 \Pi_{ij} r_{ij}^s \text{ and } \mathbb{E} [r_{ij}^s] = \sum_{i=1}^2 \pi_i \mathbb{E}_i [r_{ij}^s]. \quad (16.26)$$

The price of a one-period riskless bond satisfies

$$p_t^f = \mathbb{E}_t [M_{t+1}] \text{ or } p_i^f = \sum_{j=1}^2 \Pi_{ij} \beta G_j^{-\gamma}.$$

The real rate and its unconditional expected value are

$$r_i^f = \frac{1}{p_i^f} - 1 \text{ and } \mathbb{E} [r_i^f] = \sum_{i=1}^2 \pi_i r_i^f. \quad (16.27)$$

The equity premium is defined as the difference between the expected stock return (16.26) and the mean risk-free rate (16.27). This unconditional moment can be compared to the average return difference in the data. One can also study a conditional version of the equity premium, defined as the difference between the conditional expected value of the stock return and the risk-free rate. This moment, which conditions on the available information at time t , can be compared to predictability regressions of future excess returns on stocks on regressors known at time t .

Our model has two preference parameters $\{\beta, \gamma\}$ and three parameters that describe the dynamics of the endowment process $\{G_L, G_H, \phi\}$. When we specify $G_L = 1 + \mu - \delta$ and $G_H = 1 + \mu + \delta$, the parameter μ is the mean consumption growth rate and δ generates variance in the consumption growth rate. The parameter ϕ is the probability of remaining in the current state of the Markov chain; it governs the persistence of consumption growth. We use the generalized method of moments to select the three parameters μ , δ and ϕ that describe the dynamics of the consumption growth rate. Table 16.1 shows three empirical moments that we want to match: average per-capita real consumption growth, its standard deviation (or volatility), and its autocorrelation. In the United States, consumption growth has a mean and volatility of roughly 2 percent, and a low autocorrelation of 10 percent.⁵

Table 16.1 shows the high average real return on the S&P 500, an index of 500 large companies which captures the behavior of the overall U.S. stock market. By returning 8 percent per year to investors, U.S. stocks have had a strong performance on average, especially compared to the 1 percent per year that investors would have achieved by holding a safe asset, the 3-month T-bill. The resulting (unconditional) equity premium of 7 percent

⁵We start from moments of the gross growth rate

$$\mathbb{E} [G_t] = \pi_1 G_L + \pi_2 G_H, \mathbb{E} [G_t^2] = \pi_1 G_L^2 + \pi_2 G_H^2, \mathbb{E} [G_t G_{t-1}] = \pi_1 \Pi_{11} G_L^2 + \pi_1 \Pi_{12} G_L G_H + \pi_2 \Pi_{21} G_H G_L + \pi_2 \Pi_{22} G_H^2.$$

The mean, variance and autocorrelation of the net growth rate are: $\mathbb{E} [G_t] - 1$, $\sigma^2(G_t - 1) = \sigma^2(G_t) = \mathbb{E} [G_t^2] - \mathbb{E} [G_t]^2$, and $\rho(G_t - 1, G_{t-1} - 1) = \rho(G_t, G_{t-1}) = \text{Cov}(G_t, G_{t-1}) / \sigma^2(G_t) = (\mathbb{E} (G_t G_{t-1}) - \mathbb{E} (G_t)^2) / \sigma^2(G_t)$.

Table 16.1: Sample moments of aggregate consumption growth and real returns

	Consumption growth	Real return on S&P	risk-free rate
Mean	0.0175	0.0828	0.0074
Standard deviation	0.0233	0.1836	0.0380
Autocorrelation	0.1069		

Notes: Real per-capita data on nondurables and services consumption constructed as Fisher index based on annual sample 1929-2024 of NIPA Tables 2.3.5 in billions of Dollars, chain-type quantity indexes, 2.3.3., price indices 2.3.4., and population numbers in NIPA Table 2.1. Annual data from 1930 to 2024 on S&P returns are from Bob Shiller’s "ie-data.xls" file at Yale. Annual data on the risk-free rate is from Bob Shiller’s "chapt26.xlsx" file combined with data on the 3-month T-bill rate from 1954 to 2024 from FRED. To compute real returns and rates, we subtract the annual inflation rate based on nondurable and service consumption.

represents compensation for taking risk: the return on the S&P 500 has a volatility of 18 percent.

Once we estimate the dynamics of the endowment process, the only two free parameters are the discount factor β and risk aversion γ . Figure 16.4 computes the equity premium and risk-free rate for a discount factor β between 0 and 1, and a risk aversion γ between 0 and 100. The plot only shows equilibria in which the risk-free rate is below 4. The model is a qualitative success: the equity premium and the risk-free rate are both positive. However, the figure shows the massive quantitative failure of the model: the equity premium is less than 0.1 percent, while the risk-free rate quickly rises above its historical average of 1 percent. The tension is known as the equity premium puzzle and the risk-free rate puzzle.

16.7 Lognormal model

Computational results make it hard to understand the intuition behind the quantitative failure of the complete-markets, rational expectations model with power utility. To obtain analytical expressions, let us now assume that endowment growth is normally distributed:

$$g_{t+1} = \log(G_{t+1}) \sim N(\mu_g, \sigma_g^2) \quad (16.28)$$

With power utility, the log pricing kernel is log-normal as well, $\log M_{t+1} = \log \beta - \gamma g_{t+1}$. In this case, we can derive a closed-form expression for the risk-free rate. First, we determine the price of a one-period risk-free bond

$$p_t^f = \mathbb{E}_t[M_{t+1}] = \beta \mathbb{E}_t[e^{-\gamma g_{t+1}}] = \beta e^{-\gamma \mu_g + \frac{1}{2} \gamma^2 \sigma_g^2},$$

which implies that the risk-free rate is

$$r_t^f = \log(R_t^f) = \log\left(\frac{1}{p_t^f}\right) = -\log \beta + \gamma \mu_g - \frac{1}{2} \gamma^2 \sigma_g^2. \quad (16.29)$$

There are two key motives that determine the equilibrium risk-free rate. The *intertemporal smoothing* motive says that a higher expected consumption growth μ_g leads to a higher risk-free rate. The reason is that agents would like to smooth their consumption, try to borrow,

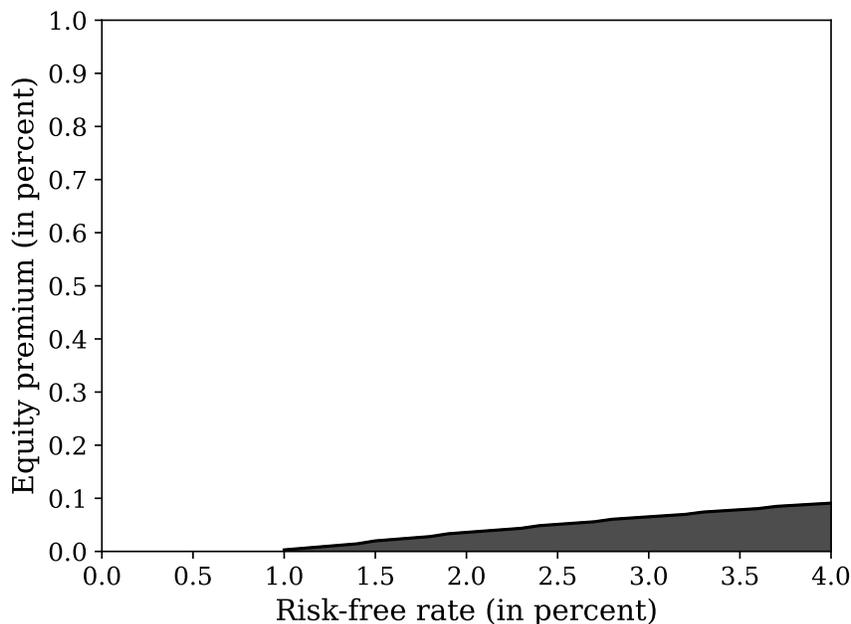


Figure 16.4: The shaded area shows the equity premium for equilibria in which the risk-free rate is below 4 percent for preference parameters $\beta \in [0, 1]$ and $\gamma \in [0, 100]$.

and the risk-free rate increases to clear the bond market. The *precautionary savings* motive says that a higher volatility of consumption growth σ_g lowers the risk-free rate. To smooth consumption across states, agents try to save, and the risk-free rate decreases the clear the bond market.

With power utility, both motives are governed by the same parameter γ . With the sample moments from Table 16.1, $\mu_g = 1.75\%$ and $\sigma_g^2 = 0.054\%$, so that the intertemporal smoothing motive quantitatively dominates, except for extremely high values of risk aversion. In this case, higher risk aversion γ leads to a higher risk-free rate.

If the gross return on equity R^s and the pricing kernel are jointly log-normal, the Euler equation becomes

$$\begin{aligned} 1 &= \mathbb{E}_t [M_{t+1} R_{t+1}^s] = \mathbb{E}_t [\beta e^{-\gamma g_{t+1} + r_{t+1}^s}] \\ &= \beta \exp \left\{ -\gamma \mu_g + \mathbb{E}_t (r_{t+1}^s) + \frac{1}{2} (\gamma^2 \sigma_g^2 + \sigma_r^2 + 2 \text{Cov}_t (-\gamma g_{t+1}, r_{t+1}^s)) \right\} \end{aligned}$$

Taking logs and rearranging gives

$$\mathbb{E}_t (r_{t+1}^s) = -\log \beta + \gamma \mu_c - \frac{1}{2} \gamma^2 \sigma_c^2 - \frac{1}{2} \sigma_r^2 + \gamma \text{Cov}_t (g_{t+1}, r_{t+1}^s)$$

The first three terms in this expression are just the risk-free rate. Therefore, the equity premium is

$$\mathbb{E}_t (r_{t+1}^s) - r_t^f + \frac{1}{2} \sigma_r^2 = \gamma \text{Cov}_t (g_{t+1}, r_{t+1}^s) \quad (16.30)$$

The term $\frac{1}{2} \sigma_r^2$ is a Jensen's inequality term which is introduced by studying log returns, since $\log \mathbb{E}_t (R_{t+1}^s) = \mathbb{E}_t (r_{t+1}^s) + \frac{1}{2} \sigma_t^2 (r_{t+1}^s)$. The equity premium therefore includes the Jensen's term on the left-hand side.

The determination of the equity premium in equation (16.30) is the core of the *consumption-CAPM* or *CCAPM*. This name is derived from the CAPM (capital asset pricing model), which is derived with quadratic utility (or mean-variance preferences). According to the CAPM, assets that are highly exposed to overall stock-market risk earn higher returns on average. With power utility and log-normal consumption and (gross) returns, we obtain the CCAPM. According to the CCAPM, assets that are highly exposed to *consumption risk* earn higher returns. The equity premium is compensation for consumption risk exposures when holding stocks. The quantity of risk exposure is measured as the covariance of stock returns with consumption growth. The risk aversion coefficient γ represents the price for each unit of risk exposure. If investors are highly risk averse, they demand high compensation for being exposed to consumption risk.

We can write the covariance between consumption growth and stock returns as their correlation coefficient times the standard deviation of consumption growth σ_c and the standard deviation of stock returns σ_r . From Table 16.1, the product of these standard deviations is $\sigma_c\sigma_r = 0.0233 \times 0.1836 = 0.0043$. Even if consumption growth and stock returns are perfectly correlated, the right-hand side of equation (16.30) is, therefore, $\gamma \times 0.0043$. To obtain the 7% equity premium from Table 16.1, risk aversion would need to be high, $\gamma = 16$. With such a high risk aversion, the risk-free rate in equation (16.29) would be pushed to a high level, above its 1% mean in Table 16.1. This tension represents the joint equity premium and risk-free rate puzzle.

With the alternative assumption of Epstein-Zin utility (instead of power utility), we can reconcile a high equity premium with a low risk-free rate. The reason is that, with Epstein-Zin utility, the intertemporal smoothing motive in equation (16.29) is governed by the inverse of the elasticity of intertemporal substitution instead of risk aversion γ , while the parameter that enters the equity premium is still risk aversion. In this world, a high equity premium reflects that investors demand a high compensation for consumption risk, while they accept low compensation for holding risk-free bonds.

16.8 The excess volatility puzzle

Table 16.1 shows that consumption growth is only weakly autocorrelated. By assuming that consumption growth is i.i.d. log-normal, we obtain that the price-dividend ratio is constant:

$$v_t = \mathbb{E}_t \left[\beta G_{t+1}^{1-\gamma} (1 + v_{t+1}) \right] = \mathbb{E}_t \left[\beta G_{t+1}^{1-\gamma} + \beta^2 G_{t+1}^{1-\gamma} G_{t+2}^{1-\gamma} + \beta^3 G_{t+1}^{1-\gamma} G_{t+2}^{1-\gamma} G_{t+3}^{1-\gamma} \dots \right]. \quad (16.31)$$

The volatility of the price-dividend ratio determines the volatility of stock returns

$$r_{t+1}^s = \log \left(G_{t+1} \frac{(1 + v_{t+1})}{v_t} \right) = g_{t+1} + \log \frac{(1 + v_{t+1})}{v_t} \quad (16.32)$$

since the growth rate has a low volatility of 2 percent in Table 16.1. With a constant pd-ratio, the second term on the right-hand side of the equation (16.32) for stock returns is constant, implying that the volatility of the return is also 2 percent. This low volatility stands in stark contrast to the high historical volatility of 18 percent in Table 16.1. The resulting tension between the model and the data is called the excess volatility puzzle for which Robert Shiller won the Nobel Prize in 2013.

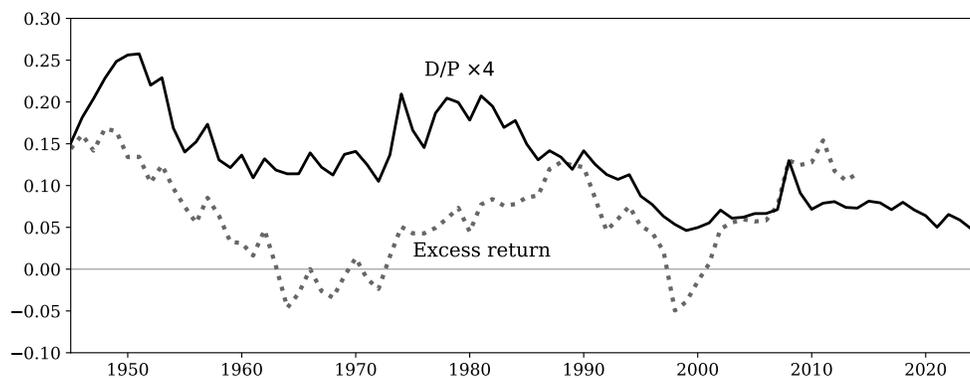


Figure 16.5: Dividend-price ratio together with excess returns on stocks over the next 10 years.

The preceding example assumes an iid growth rate g_{t+1} . Together with power utility, this assumption implies a constant price-dividend ratio and iid stock returns (16.32). Recent work on the quantitative asset-pricing implications of macro models has stressed that the price-dividend ratio varies dramatically over time. As a consequence, the second term in equation (16.32) is volatile and generates volatility in returns. Figure 16.5 shows annual postwar data on the inverse price-dividend ratio $1/v_t = d_t/p_t^s$, the so-called dividend yield, as a gray line together with (annualized) excess stock returns on the S&P 500 over the next decade in black. The dividend yield is multiplied by 4, so that the two series fit on the same vertical scale. The dividend yield moves around in a persistent way between 1% and 6.5%. This translates into persistent movements in the price-dividend ratio, ranging from 15 to 100. In particular, a high dividend yield tends to be associated with high subsequent excess returns, since the two series in Figure 16.5 move together.

Table 2 shows the results of predicting returns from buying the S&P 500 in year t and selling in year $t+k$. The prediction is based on the current dividend yield d_t/p_t^s . As suggested by Figure 16.5, the estimated slope coefficient in this regression is positive: a higher dividend yield predicts higher future excess returns. Table 16.2 shows that the slope-coefficient increases and becomes more significant over longer horizons k . The R^2 also increases with the horizon, meaning that excess returns can be forecasted more easily over longer holding periods, such as several years. Over shorter horizons, excess returns are harder to predict.

Table 16.2: Return forecasting regressions

Forecast horizon k	b	$t(b)$	R^2
1 year	2.41	1.78	0.04
5 year	16.85	2.03	0.17
10 year	51.49	3.20	0.28

Notes: The regression equation is $r_{t \rightarrow t+k}^s - r_{t \rightarrow t+k}^f = a + bd_t/p_t^s + \varepsilon_{t+k}$. The dependent variable is the excess return on the S&P 500 over the 3-month T-bill rate. Data are annual 1947-2024. The k -year regression t-statistic uses the Hansen and Hodrick (1980) correction when $k > 1$.

The literature has studied three main approaches to capture this predictability in macro

models. The approach deviates from the iid assumption on growth rates. This can be done by introducing time-varying conditional moments in growth rates, which are then measured from historical data in a rational expectations approach. From equation (16.31), we can see that time-varying first moments $\mathbb{E}_t[G_{t+h}]$ for $h > 0$ will generate movements in the price-dividend ratio v_t . From equation (16.30), we can see that time-variation in conditional second moments $\text{Cov}_t(g_{t+1}, r_{t+1}^s)$ will translate into time-variation in the conditional expectation of future excess returns. There is some empirical evidence supporting such time variation, but it is quantitatively not strong enough.

The second approach allows investors' subjective beliefs to feature more time variation in conditional moments of growth rates than what is measured in regressions. For example, these models assume sentiment, ambiguity-aversion or other deviations from rational expectations, which implies that the conditional expected value of growth rates varies over time. The empirical evidence in favor of these non-rational expectations comes from survey data on investors or CEOs.

The third approach takes the time-variation from preferences. For example, some models feature preference shocks that are heteroskedastic. These preference shocks are designed to match the evidence on volatility and predictability. Since these preference shocks are unobservable, the support for this modeling approach has to come from other implications of the model that match additional data (such as micro data on household portfolio choice) or that rationalize some observed policy actions by the government (such as the behavior of tax rates over the business cycle.)