Chapter 16: Asset Prices

Chapter authors: Monika Piazzesi and Martin Schneider

August 2026

Outline

Introduction and Background

Why Multiple Assets?

The benchmark neoclassical growth model has one asset (capital) and one return. Allowing multiple assets matters for two reasons:

1. **New questions**: Why do average returns differ across assets? Why are some prices more volatile than others? Why do portfolios differ across agents? How do asset price movements affect welfare differently?
2. **Discipline for macro models**: Asset prices pin down preference parameters (e.g., what preferences match equity vs. debt returns?). Under incomplete markets, the asset structure matters directly for allocations and risk sharing.

Even with complete markets and efficient allocations, asset prices remain important to **test** the model.

Household Portfolios and Price-Cash Flow Ratios

Participation: 65% of US households own a house; 52% own stocks. Cross-country variation (homeownership: Germany 47%, Italy 75%) partly reflects pension systems.

Heterogeneity: young households hold leveraged housing positions; older households accumulate wealth, pay down mortgages, hold housing and stocks; many hold little wealth at all.

Price–Cash Flow Ratios

The ratio of asset value to cash flow for housing and stocks, US 1930–2023

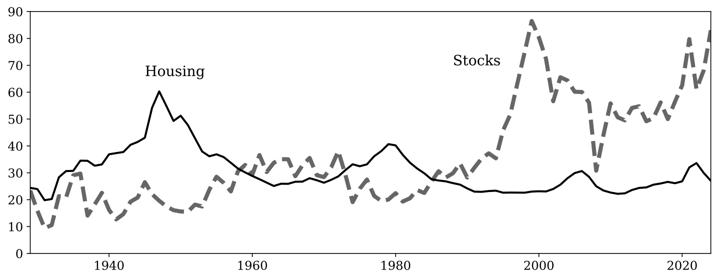

- Price–dividend ratio for stocks; value of residential housing over housing services for housing
- Both ratios are highly volatile; they sometimes comove (Great Depression) and sometimes diverge (Great Inflation)

Dynamic Stochastic Endowment Economy

Environment

States $\omega_t \in \Omega$ (S possible), histories $\omega^t = \{\omega_0, \dots, \omega_t\}$ with probability $\pi_t(\omega^t)$. Agents $i \in I$ with utility $U_i(c_i)$ over plans $c_i = \{c_{i,t}(\omega^t)\}$ and endowments y_i; aggregate endowment $y_t(\omega^t) = \sum_{i \in I} y_{i,t}(\omega^t)$.

Expected utility example:

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

Power utility $u(x) = x^{1-\gamma}/(1-\gamma)$ ties risk aversion γ to the inverse of the EIS: $\gamma = 1/\sigma$.

Environment: Preferences

For asset pricing, more general utility matters — the leading example is **Epstein-Zin utility**, which separates risk aversion from the EIS ($\gamma \neq 1/\sigma$). The chapter derives results for general U_i.

Planner Problem and Arrow-Debreu Equilibrium

Planner: maximize $\sum_{i \in I} \lambda_i U_i(c_i)$ subject to $\sum_{i \in I} c_{i,t}(\omega^t) \leq y_t(\omega^t)$. FOC:

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

Weighted marginal utilities equated to a multiplier **independent of** i; all MRS equalized.

Arrow-Debreu Equilibrium

Arrow-Debreu: trade date-0 claims at prices $p_t^0(\omega^t)$. Agent FOC:

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

First Welfare Theorem: choosing $\lambda_i = 1/\alpha_i$, the planner's and equilibrium conditions coincide. The AD allocation is Pareto optimal; marginal utilities of all agents are collinear.

Sequential Markets

N assets with exogenous dividends $d_t(\omega^t)$, traded **ex dividend** at prices $p_t(\omega^t)$. Trading strategies $\theta = \{\theta_t(\omega^t)\}$ (long, short, or zero positions). Budget set:

$$\begin{aligned} & c_{i,t}(\omega^t) + p_t(\omega^t)^\top \theta_{i,t}(\omega^t) \\ \leq \; & y_{i,t}(\omega^t) + \left(d_t(\omega^t) + p_t(\omega^t)\right)^\top \theta_{i,t-1}(\omega^{t-1}) \end{aligned}$$

Equilibrium: plans and strategies maximize utility in the budget set; goods markets clear; asset markets clear ($\sum_i \theta^*_{i,t}(\omega^t) = 0$, or $\bar{\theta}^j$ if asset j in positive supply).

Complete Markets: Three Results

Preview — with complete markets:

1. AD and sequential equilibrium allocations coincide
2. First Welfare Theorem holds (Pareto optimal)
3. A representative agent exists

All three are lost under incomplete markets — heterogeneity then matters for asset valuation.

Asset Trading with Two Periods

Setup and Arbitrage

Two dates; S states tomorrow with probabilities $\pi(\omega) > 0$. N assets, prices $p \in \mathbb{R}^N$, payoffs collected in the $S \times N$ matrix D. Portfolio θ costs $p^\top\theta$ today, pays $D\theta$ tomorrow.

Arbitrage: a portfolio with payoff matrix

$$W = \begin{pmatrix} -p^\top \\ D \end{pmatrix}$$

such that $W\theta \geq 0$ with at least one strict inequality — a free lunch.

Attainable Payoffs and the Law of One Price

Set of attainable payoffs:
$M = \{x \in \mathbb{R}^{S+1} : x = W\theta \text{ for some } \theta \in \mathbb{R}^N\}$

Law of one price: identical payoffs $\Rightarrow$ identical prices. With state prices $q(\omega) > 0$ (prices of Arrow securities):

$$p^n = \sum_{\omega \in \Omega} d^n(\omega) q(\omega), \qquad \text{or} \qquad p = D^\top q$$

Fundamental Theorem of Asset Pricing

Theorem: For given asset prices p, the absence of arbitrage is equivalent to the existence of strictly positive state prices q with $p = D^\top q$.

- Complete markets (rank$(D) = S$): state prices **unique**; if $N = S$, $q = (D^\top)^{-1} p$
- Incomplete markets (rank$(D) < S$): **many** state-price vectors support the same p

Proof idea: ($\Leftarrow$) $Dq > 0$ values any nonnegative payoff positively. ($\Rightarrow$) Separate the cones M and $K = \mathbb{R}_+^{S+1}$ (intersecting only at 0) with a hyperplane; its strictly positive normal $\tilde{q}$, normalized by $\tilde{q}_0$, yields the state prices $q = \tilde{q}_1 / \tilde{q}_0$.

Fundamental Theorem: Example

Example (stock and bond, $S = N = 2$, $\delta_1 > 1 > \delta_2 > 0$): no arbitrage requires $\delta_1 p^b > p^s > \delta_2 p^b$; inverting $D^\top$ gives unique q. With only the stock: a continuum of valid q.

Separating $\mathcal{M}$ from the Arbitrage Cone $\mathcal{K}$

Attainable payoffs $\mathcal{M}$ and arbitrages $\mathcal{K}$, for $S = 1$

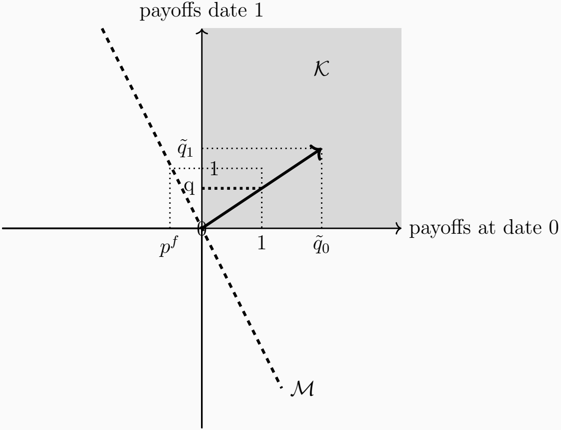

- $\tilde{q} \perp \mathcal{M}$; normalizing by $\tilde{q}_0 > 0$ gives state prices $q = \tilde{q}_1/\tilde{q}_0$

Rewriting the Budget Set; Optimality

With state prices, the period budget constraints consolidate into a single constraint plus a **span restriction**:

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

$$\{c_{i,1}(\omega) - y_{i,1}(\omega)\}_\omega \in \mathsf{span}(D)$$

Budget Sets: Complete Markets vs. a Riskless Bond

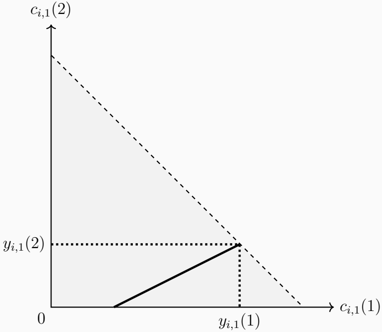

- Two states; no date-0 consumption or endowment
- Complete markets (dashed): any point below or on the line through $y_{i,1}$
- Riskless bond (solid): one unit added or subtracted in *both* states

Optimality and Subjective State Prices

Agent optimality: FOC for asset n combines into

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

Subjective state prices: with power utility and common beliefs, $q_i(\omega) = \beta\pi(\omega)(c_{i,1}(\omega)/c_{i,0})^{-\gamma}$. Agents **agree on asset prices** (bundles) but may **disagree on state prices** when markets are incomplete; with complete markets, $q_i = q_j = q$ (unique).

Stochastic Discount Factor and Risk Premia

Define $M(\omega) = q(\omega)/\pi(\omega)$. Then:

$$p^n = \mathbb{E}[d^n(\omega)M(\omega)], \qquad 1 = \mathbb{E}[R^n(\omega)M(\omega)]$$

M measures hunger for consumption: power utility gives $M(\omega) = \beta(c_1(\omega)/c_0)^{-\gamma}$ — high when consumption growth is low. Risk neutrality: $M = \beta$, $p^n = \mathbb{E}[d^n]/R^f$.

Risk Adjustment and Risk Premia

Risk adjustment via the covariance decomposition:

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

$$\mathbb{E}[R^n(\omega)] - R^f = -\frac{\mathsf{Cov}(R^n(\omega), M(\omega))}{\mathbb{E}[M(\omega)]}$$

Insurance assets (pay in hungry states): positive covariance $\Rightarrow$ **negative** premium. Stocks (pay in booms): negative covariance $\Rightarrow$ **positive** premium. Only **systematic** risk (covariance with M) matters — not variance.

Risk-Neutral Probabilities and the Representative Agent

Risk-neutral probabilities: $\pi^*(\omega) = q(\omega)/\sum_{\omega'} q(\omega')$, with $\sum_{\omega'} q(\omega') = p^b = 1/R^f$:

$$p^n = \frac{\mathbb{E}^*[d^n(\omega)]}{R^f}$$

Prices look risk-neutral under the distorted measure π^*. State prices, SDFs, and risk-neutral probabilities are equivalent languages; π^* is unique iff markets are complete. Practitioners back out π^* from observed prices to value derivatives.

Representative agent: with complete markets, equilibrium prices of the heterogeneous-agent economy equal those of a representative-agent economy with endowment $y = \sum_i y_i$ and utility $U_\lambda = \sum_i \lambda_i U_i$, $\lambda_i = 1/\alpha_i$. State prices read off the representative agent's MRS:

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

Dynamic Asset Trading

Infinite Horizon: Payoffs, Arbitrage, State Prices

Payoffs generated by strategy θ:

$$x_t(\omega^t) = \left(p_t(\omega^t) + d_t(\omega^t)\right)^\top \theta_{t-1}(\omega^{t-1}) - p_t(\omega^t)^\top \theta_t(\omega^t)$$

Arbitrage: a strategy with $x_t(\omega^t) \geq 0$ for all (t, ω^t), one strict. Markets **complete** if any payoff stream from date 1 onward can be generated at some date-0 cost.

Law of one price (state prices $q_{t+1}(\omega_{t+1}|\omega^t)$):

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

The fundamental theorem, the separating-hyperplane proof, and the equivalence between no arbitrage and existence of a solution to the agent problem all carry over from the two-period case.

Optimality and Subjective State Prices

Agent FOCs give subjective state prices and the SDF:

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

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

With expected power utility:

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

The Euler Equation

Returns $R^n_{t+1} = (d^n_{t+1} + p^n_{t+1})/p^n_t$ satisfy the **Euler equation**:

$$\mathbb{E}_t[M_{t+1} R^n_{t+1}] = 1$$

$$\mathbb{E}_t[R^n_{t+1}] - R^f_{t+1} = -\frac{\mathsf{Cov}_t(M_{t+1}, R^n_{t+1})}{\mathbb{E}_t[M_{t+1}]}$$

Conditional covariances can vary over time $\Rightarrow$ time-varying expected excess returns.

The Equity Premium and Risk-Free Rate Puzzles

The Mehra-Prescott Environment

Single agent with power utility, endowment y (equivalently, the **Lucas tree** model with $c^* = d$). Assets in zero net supply priced so the agent chooses not to trade. Growth follows a two-state Markov chain:

$$y_{t+1} = y_t G_{t+1}, \qquad G_{t+1} \in \{G_L, G_H\}, \qquad \Pi = \begin{pmatrix} \phi & 1-\phi \\ 1-\phi & \phi \end{pmatrix}$$

Mehra-Prescott: Prices and Returns

Stock $=$ claim to the endowment. Price-dividend ratio $v_t = p_t^s / y_t$:

$$v_i = \sum_{j=1}^{2} \Pi_{ij} \beta G_j^{1-\gamma} (1 + v_j)$$
$$\Rightarrow \quad v = (I - \beta A)^{-1} \beta A \mathbf{1}, \quad A_{ij} = \Pi_{ij} G_j^{1-\gamma}$$

Returns: $r_{ij}^s = G_j(1 + v_j)/v_i - 1$; riskless bond price $p_i^f = \sum_{j=1}^{2} \Pi_{ij} \beta G_j^{-\gamma}$. Unconditional moments use the stationary distribution ($\pi_1 = \pi_2 = 1/2$).

Data and the Puzzle

Sample moments of aggregate consumption growth and real returns, US 1929–2024

	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		

Equity premium $\approx 7\%$; stock volatility 18%; risk-free rate $\approx 1\%$.

The Equity Premium Puzzle

Equity premium for equilibria with a risk-free rate below 4%, $\beta \in [0,1]$, $\gamma \in [0,100]$

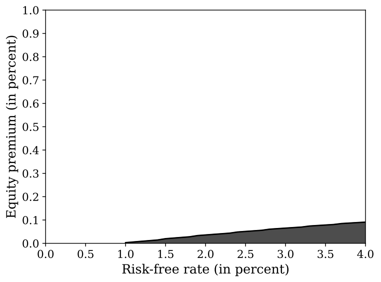

- The model's equity premium stays **below 0.1%** while the risk-free rate quickly exceeds 1%
- Qualitative success, massive quantitative failure: the **equity premium** and **risk-free rate** puzzles

The Lognormal Model

Closed Forms with Lognormal Growth

Assume $g_{t+1} = \log G_{t+1} \sim N(\mu_g, \sigma_g^2)$; with power utility $\log M_{t+1} = \log \beta - \gamma g_{t+1}$.

Risk-free rate:

$$r_t^f = -\log \beta + \gamma \mu_g - \frac{1}{2}\gamma^2 \sigma_g^2$$

- **Intertemporal smoothing**: higher $\mu_g \Rightarrow$ agents want to borrow $\Rightarrow$ higher r^f
- **Precautionary saving**: higher $\sigma_g \Rightarrow$ agents want to save $\Rightarrow$ lower r^f
- Power utility forces **both** motives to be governed by the same γ; with the data's $\mu_g = 1.75\%$, $\sigma_g^2 = 0.054\%$, the smoothing motive dominates: higher γ raises r^f

Lognormal Model: Premia and Rates

Equity premium (joint lognormality of R^s and M):

$$\mathbb{E}_t[r^s_{t+1}] - r^f_t + \frac{1}{2}\sigma^2_r = \gamma \mathsf{Cov}_t(g_{t+1}, r^s_{t+1})$$

This is the **consumption CAPM**: exposure to consumption risk (the covariance) is the quantity of risk; γ is its price.

The Joint Puzzle and Epstein-Zin

From Table 16.1: $\sigma_c \sigma_r = 0.0233 \times 0.1836 = 0.0043$. Even with **perfect correlation** between consumption growth and returns, the premium is at most $\gamma \times 0.0043$.

Matching the 7% premium requires $\gamma \approx 16$ — but then the risk-free rate equation pushes r^f far above its 1% historical mean. This is the **joint equity premium and risk-free rate puzzle**.

Epstein-Zin Preferences

Epstein-Zin resolution: separate the two roles.

- The smoothing motive in the risk-free rate is governed by $1/\sigma$ (the inverse EIS)
- The equity premium is still governed by risk aversion γ

High γ delivers a high premium (high compensation for consumption risk) while a moderate EIS keeps the risk-free rate low.

Excess Volatility and Return Predictability

The Excess Volatility Puzzle

With i.i.d. lognormal growth and power utility, the price-dividend ratio is **constant**:

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

Excess Volatility: Return Decomposition

Log stock returns decompose as:

$$r^s_{t+1} = g_{t+1} + \log \frac{1 + v_{t+1}}{v_t}$$

Constant $v \Rightarrow$ return volatility = consumption growth volatility $\approx$ **2%**, vs. **18%** in the data. This is the **excess volatility puzzle** (Shiller, Nobel 2013).

In the data the price-dividend ratio moves persistently between roughly 15 and 100: the second term generates the missing return volatility.

Dividend–price ratio ($\times 4$) with annualized excess stock returns over the next 10 years, US 1947–2024

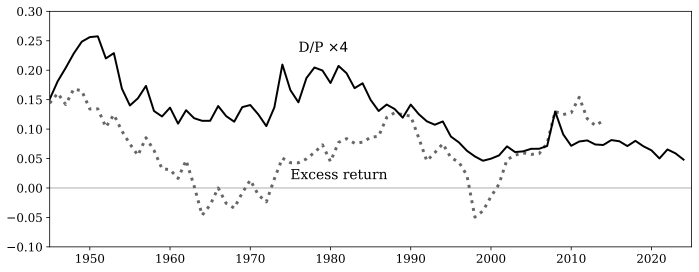

The two series move together: a higher dividend yield predicts higher future excess returns.

Return Forecasting Regressions

$$r^s_{t\to t+k} - r^f_{t\to t+k} = a + b\, d_t/p^s_t + \varepsilon_{t+k}$$

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

A high dividend yield predicts **high future excess returns**; predictability strengthens with the horizon (the Hansen-Hodrick correction applies for $k > 1$).

Three Approaches to Volatility and Predictability

1. **Time-varying growth moments under rational expectations**: time-varying $\mathbb{E}_t[G_{t+h}]$ moves the price-dividend ratio; time-varying $\mathsf{Cov}_t(g_{t+1}, r^s_{t+1})$ moves expected excess returns. Empirical support exists but is quantitatively weak.

Three Approaches to Volatility and Predictability (cont.)

2. **Subjective beliefs**: sentiment, ambiguity aversion, or other deviations from rational expectations make subjective conditional moments more variable than measured ones. Support comes from survey data on investors and CEOs.
3. **Preference shocks**: heteroskedastic preference shocks designed to match volatility and predictability. Since the shocks are unobservable, validation must come from other implications (household portfolio micro data, observed policy behavior).

Summary

Key Takeaways (I)

1. **Why multiple assets**: new questions (return differences, volatility, portfolios) and discipline for macro models; under incomplete markets, market structure matters for allocations.
2. **Complete markets trilogy**: AD and sequential equilibria coincide; First Welfare Theorem; a representative agent exists with $U_\lambda = \sum_i \lambda_i U_i$, $\lambda_i = 1/\alpha_i$. All three fail under incomplete markets.
3. **No arbitrage and the fundamental theorem**: no arbitrage $\Leftrightarrow$ strictly positive state prices with $p = D^\top q$ (separating hyperplane proof). State prices unique iff markets complete (rank$(D) = S$).

Key Takeaways (II)

4. **SDF pricing**: $M = q/\pi$; $p^n = \mathbb{E}[d^n M]$; $1 = \mathbb{E}_t[M_{t+1} R^n_{t+1}]$; risk premium $\mathbb{E}_t[R^n] - R^f = -\mathsf{Cov}_t(M, R^n)/\mathbb{E}_t[M]$. Only systematic risk (covariance, not variance) is priced. Risk-neutral probabilities $\pi^* = qR^f$ give equivalent pricing.
5. **Subjective state prices**: agents agree on traded asset values but can disagree on state prices when markets are incomplete.

Key Takeaways (III)

6. **Mehra-Prescott**: two-state Markov growth; $v = (I - \beta A)^{-1}\beta A\mathbf{1}$. Calibrated to consumption moments, the model delivers a premium below 0.1% — the **equity premium puzzle**.
7. **Lognormal model**: $r^f = -\log\beta + \gamma\mu_g - \frac{1}{2}\gamma^2\sigma_g^2$ (smoothing vs. precautionary motives); CCAPM premium $= \gamma\mathrm{Cov}_t(g, r^s)$. With $\sigma_c\sigma_r = 0.0043$, matching 7% needs $\gamma \approx 16$, which blows up r^f: the **joint puzzle**.
8. **Epstein-Zin**: separates risk aversion (premium) from the EIS (risk-free rate), reconciling a high premium with a low safe rate.

Key Takeaways (IV)

8. **Excess volatility**: i.i.d. growth $\Rightarrow$ constant price-dividend ratio $\Rightarrow$ return volatility 2% vs. 18% in the data (Shiller).
9. **Predictability**: high dividend yields forecast high future excess returns, increasingly so at long horizons (Table 16.2). Three modeling responses: time-varying moments, subjective beliefs, preference shocks.