CHAPTER 7

The Stock Market, Rational Expectations, and the Efficient Market Hypothesis

Ritsu Kitagawa

Columbia Business School

Created: June 9, 2026 ▪ Last modified: June 9, 2026

Outline

Reading and objectives

- Mishkin, *The Economics of Money, Banking, and Financial Markets*, 13th ed.
- This deck: **Chapter 7** — The Stock Market, the Theory of Rational Expectations, and the Efficient Market Hypothesis.
- We move from *bonds* (Ch. 4–6) to *equity*, then build the general theory of how forward-looking markets price assets.

Learning objectives

7.1 Calculate the price of common stock.

7.2 Recognize the impact of new information on stock prices.

7.3 Compare and contrast adaptive and rational expectations.

7.4 Explain why arbitrage opportunities imply that the efficient market hypothesis holds.

7.5 Identify the implications of the efficient market hypothesis for investors.

7.6 Summarize why behavioral finance questions market efficiency.

PART 1

Valuing Common Stock

What is a share of common stock?

- A claim on a firm's residual cash flows (dividends) and a residual claim on assets.
- Two cash flows to a shareholder:
  - **Dividends** paid while the share is held.
  - The **sale price** when the share is sold.
- As with any asset, the price today is the *present value* of expected future cash flows, discounted at the required return on equity k_e.

k_e is the return investors demand to hold the stock — it embeds the risk-free rate plus a risk premium for the stock's risk.

Mishkin Ch. 7, Section 7.1.

The one-period valuation model

Hold the stock for one year, collect a dividend, then sell.

$$P_0 = \frac{\text{Div}_1}{1+k_e} + \frac{P_1}{1+k_e}.$$

- P_0 — price today.
- Div_1 — dividend expected at the end of year 1.
- k_e — required return on equity.
- P_1 — expected sale price at the end of year 1.

Example

$\text{Div}_1 = \$0.16$, $P_1 = \$60$, $k_e = 12\%$. $P_0 = \frac{0.16}{1.12} + \frac{60}{1.12} = \$53.71$.

The generalized dividend valuation model

But P_1 itself is the present value of dividends from year 2 onward. Substitute repeatedly:

$$P_0 = \sum_{t=1}^{\infty} \frac{\text{Div}_t}{(1+k_e)^t}.$$

- The value of a stock today is the present value of *all* expected future dividends.
- The final sale price drops out: a sale price far in the future, discounted, contributes negligibly.
- The price is determined *only* by the expected future dividend stream.

Mishkin Ch. 7, Section 7.1.

The Gordon growth model

Assume dividends grow at a constant rate g forever, with $g < k_e$. The infinite sum collapses:

$$P_0 = \frac{\text{Div}_0(1+g)}{k_e - g} = \frac{\text{Div}_1}{k_e - g}.$$

- Div_0 — the most recent dividend paid.
- g — expected constant growth rate of dividends.
- k_e — required return on equity.

Why $g < k_e$?

If dividends grew at least as fast as the discount rate, the present value would be infinite. The condition keeps the price finite.

Gordon growth — worked example

A firm just paid $\mathrm{Div}_0 = \$3.00$. Investors expect $g = 5\%$ and require $k_e = 8\%$.

$$P_0 \;=\; \frac{3.00 \times 1.05}{0.08 - 0.05} \;=\; \frac{3.15}{0.03} \;=\; \$105.$$

- The price is highly sensitive to the spread $k_e - g$.
- If k_e falls to 7%: $P_0 = 3.15/0.02 = \$157.50$ — a 1-pp drop in k_e raises the price 50%.
- Small changes in expectations $\Rightarrow$ large price swings. Keep this in mind for the applications.

PART 2

How the Market Sets Prices

How the market sets stock prices

- The price is set by the buyer willing to pay the *highest* price.
- That buyer is typically the one who can make best use of the asset — e.g. has superior information that lowers its perceived risk (lower k_e) or raises expected dividends (higher g).
- **Information matters.** When news arrives about a firm, market participants revise Div, g, or k_e, and the price changes immediately.
- Because new information arrives constantly, prices change frequently.

Mishkin Ch. 7, Section 7.2.

Application — monetary policy and stock prices

The Gordon model gives two channels through which a Fed *rate cut* raises stock prices:

1. **Required-return channel.** Lower interest rates make bonds (the alternative asset) less attractive, so investors accept a lower k_e. The denominator $k_e - g$ falls $\Rightarrow$ P_0 rises.
2. **Growth channel.** Lower rates stimulate the economy, raising expected dividend growth g. Again $k_e - g$ falls $\Rightarrow$ P_0 rises.

Both channels push in the same direction — which is why stock markets often rally on news of monetary easing.

Mishkin Ch. 7, Application.

Application — the 2020 coronavirus crash

- The Dow fell from a peak of **29,551** on Feb 12, 2020 to **18,561** on Mar 20 — a **37%** decline in five weeks.
- Read it through the Gordon model:
  - Expected dividend growth g *collapsed* as lockdowns hit earnings.
  - The required return k_e *rose* as risk and uncertainty spiked.
- Both moves widen $k_e - g \Rightarrow$ a sharp fall in P_0 — one of the steepest declines in U.S. history.

The same arithmetic that gives gentle re-pricings in normal times produces crashes when expectations move violently.

Mishkin Ch. 7, Application.

PART 3

Rational Expectations

Adaptive expectations

How do people form the expectations (Div, g, future prices) that drive asset prices?

Adaptive expectations

Expectations are formed *only* from past values of the variable.

- Expectations change *slowly*, as new data accumulates.
- **Problem:** people clearly use more than past data — they react to news, policy announcements, and forecasts, and can change their minds fast.
- Backward-looking expectations are too rigid to describe real financial markets.

Mishkin Ch. 7, Section 7.3.

Rational expectations — the idea

Rational expectations

Expectations are identical to the *optimal forecast* (the best guess) using *all available information*.

$$X^e = X^{of}, \qquad X^{of} = \text{optimal forecast using all available information.}$$

- Forward-looking: uses *everything* relevant, not just the past.
- Does *not* require forecasts to be perfectly accurate — only that errors are not systematic.

Mishkin Ch. 7, Section 7.3.

When can a rational forecast still be wrong?

A rational expectation equals the optimal forecast, yet a given prediction may be inaccurate when:

1. It takes *too much effort* to use all information — the gain is not worth the cost.
2. The forecaster is *unaware* of some relevant information.

Two implications

- If the way a variable moves changes, the way expectations are formed changes too (e.g. a new Fed policy rule).
- Forecast errors are, *on average, zero* and *cannot be predicted ahead of time*.

Why expect rationality in financial markets?

The incentive to use information optimally is *especially strong* in financial markets:

- People with better forecasts of the future **get rich**.
- Poor forecasters lose money and are driven out.
- Competition pushes surviving market prices toward optimal forecasts.

Applied to financial markets, rational expectations becomes the **efficient market hypothesis** (the theory of efficient capital markets) — the topic of the rest of the chapter.

Mishkin Ch. 7, Section 7.3.

PART 4

The Efficient Market Hypothesis

Setup — the return on a security

Recall the one-period return: capital gain plus cash payment, over the purchase price.

$$R = \frac{P_{t+1} - P_t + C}{P_t},$$

where P_t is the price now, P_{t+1} the price at the end of the period, and C the coupon or dividend.

- At time t we know P_t and C; the future price P_{t+1} is *unknown*.
- We must form an expectation P^e_{t+1}, giving an *expected return* R^e.

Mishkin Ch. 7, Section 7.4.

Deriving the hypothesis

Step 1. Rational expectations: the expected price equals the optimal forecast, $P^e_{t+1} = P^{of}_{t+1}$, so $R^e = R^{of}$.

Step 2. Supply and demand: in equilibrium the expected return equals the equilibrium return R^* required by the market, $R^e = R^*$.

Combine:

$$R^{of} \;=\; R^*.$$

Efficient Market Hypothesis

Current prices are set so that the *optimal forecast of a security's return using all available information equals the security's equilibrium return*. Equivalently: **a security's price fully reflects all available information.**

Rationale — arbitrage eliminates easy profits

Why must $R^{of} = R^*$ hold? Because unexploited profit opportunities get arbitraged away.

- If $R^{of} > R^*$: the security is a bargain. Investors buy $\Rightarrow P_t \uparrow \Rightarrow R^{of} \downarrow$, until $R^{of} = R^*$.
- If $R^{of} < R^*$: the security is overpriced. Investors sell $\Rightarrow P_t \downarrow \Rightarrow R^{of} \uparrow$, until $R^{of} = R^*$.

Key insight

In an efficient market, *all unexploited profit opportunities are eliminated*. You do *not* need every investor to be informed — enough informed arbitrageurs at the margin are sufficient to set the price.

Mishkin Ch. 7, Section 7.4.

Implication — stock prices follow a random walk

The sharpest implication of the EMH: **changes in stock prices are essentially unpredictable.**

Random walk

Future price changes are unpredictable from currently available information — the best forecast of tomorrow's price is (close to) today's price.

- Prices move only on *new, unexpected* information — and news, by definition, cannot be predicted.
- Any predictable component of returns would be arbitraged away immediately.
- **Corollary:** "technical analysis" — forecasting future prices from past price patterns — cannot systematically beat the market.

Mishkin Ch. 7, Section 7.5.

Global — do exchange rates follow a random walk?

- Though usually applied to stocks, the EMH also implies that **foreign exchange rates follow a random walk**.
- *Random walk:* future changes are essentially unpredictable from past changes — the best forecast of tomorrow's rate is today's rate.
- If exchange-rate changes were predictable, traders would exploit them, eliminating the predictability.
- The data largely confirm random-walk behavior in major exchange rates.

Mishkin Ch. 7, Global box.

PART 5

Evidence, Investing, and Behavioral Finance

How valuable are published investment reports?

- Information in newspapers and analysts' published reports is *readily available* to many participants — and is therefore *already in the price*.
- Acting on it will not, on average, earn abnormally high returns.
- The empirical evidence largely confirms: recommendations from investment advisers do *not* help us outperform the market on average.

This is the EMH's most testable and most uncomfortable prediction for the asset-management industry.

Mishkin Ch. 7, Section 7.5.

The EMH prescription for the investor

Practical advice that follows directly from market efficiency:

- Published recommendations cannot help you beat the market.
- A "hot tip" is probably *already* reflected in the price.
- Prices respond to announcements *only* when the information is *new and unexpected*.
- For the small investor, a low-cost **buy-and-hold** strategy (e.g. a broad index fund) is the most sensible approach.

Mishkin Ch. 7, Section 7.5.

Application — what do crashes tell us?

- Efficient-markets theory does *not* rule out large price changes: dramatic news can produce a large, justified re-pricing.
- But economists struggle to find *fundamental* news big enough to explain Black Monday (1987, -22% in a day) or the 2000–01 tech crash.
- **Lesson:** factors *other than* market fundamentals may move asset prices — a crack in the strong form of the hypothesis.

Mishkin Ch. 7, Application.

Efficiency vs. "efficient markets"

A subtle but important distinction:

- **EMH (the version we derived):** prices reflect all *available information*; you cannot systematically beat the market. This is well supported.
- **Stronger view:** prices always equal true *market fundamentals* (the correct present value of future income), so markets allocate capital perfectly.

The gap

Stock prices are more volatile and less predictable than the fundamentals-only view allows. The evidence casts doubt on the *stronger* claim, even where the basic EMH survives.

Mishkin Ch. 7, Section 7.6.

Behavioral finance

Behavioral finance applies psychology and sociology to explain departures from efficiency:

- **Loss aversion** + limits to short selling $\Rightarrow$ overpriced stocks can persist (pessimists can't easily bet against them).
- **Overconfidence** $\Rightarrow$ excessive trading volume relative to what rational models predict.
- **Overconfidence + social contagion** $\Rightarrow$ speculative *bubbles*.

Behavioral finance does not discard the EMH so much as bound it: arbitrage is powerful but *limited*, so prices can wander from fundamentals for a time.

Mishkin Ch. 7, Section 7.6.

PART 6

Wrap-up

Key terms from Chapter 7

- one-period valuation model
- generalized dividend valuation model
- Gordon growth model
- required return on equity k_e
- adaptive expectations
- rational expectations
- optimal forecast
- efficient market hypothesis
- unexploited profit opportunity
- arbitrage
- random walk
- market fundamentals
- buy-and-hold strategy
- behavioral finance

Takeaways

- Stock prices are the present value of expected future dividends; the Gordon growth model makes prices very sensitive to $k_e - g$.
- Rational expectations = optimal forecasts using all information; strong in financial markets where good forecasters profit.
- EMH: arbitrage drives prices to reflect all available information, so you cannot systematically beat the market — buy and hold.
- Crashes and behavioral evidence show the limits: prices can detach from fundamentals.