Understanding Economics · Chapter 5

Money, Banking & Interest Rates — narration script

1 · The puzzle

Welcome to Chapter 5 of Understanding Economics. In this chapter, we'll see what money actually is, how banks create it, and what the central bank does. Then we'll follow how interest rates travel from a meeting in Washington to the loan on your car or home.

[SLIDE: A calendar page with "Committee meeting, Washington" circled. Beside it, a phone showing a loan-rate quote ticking up. Caption: "Why does this move that?"]

Here's a puzzle. A committee of about a dozen people meets in a room in Washington. They don't lend anyone a mortgage. They don't run your bank. They don't set the price of a car loan. And yet, around the time they meet, the rate on a new mortgage can jump, the interest on your savings account changes, and the rate on your credit card follows a few weeks later.

How does a meeting in one room reach a loan you take out at a bank across the country? And here's the stranger part: sometimes mortgage rates move a lot before the committee even meets, and then barely move on the day. By the end of this chapter you'll see the main reason why, and you'll be able to read a headline about the central bank and know what it likely means for your own borrowing and saving.

2 · What money is

[SLIDE: Three icons in a row with labels. A handshake: "Medium of exchange". A ruler: "Unit of account". A piggy bank: "Store of value".]

Let's start with something that seems too obvious to explain: money. What is it?

Imagine a world without it. You're a baker, and you want a haircut. You'd have to find a barber who happens to want bread, right now, and agree on how many loaves a haircut is worth. That's barter, and it's painfully slow. Economists call the problem the need for a "double coincidence of wants": each side has to want exactly what the other has.

Money fixes that. It does three jobs.

First, it's a medium of exchange. Everyone accepts it, so you can sell bread for money and use the money for a haircut. The barber doesn't need to want bread.

Second, it's a unit of account. It's the measuring stick we use for prices. Instead of knowing that a haircut is worth eight loaves, or half a pair of shoes, everything is priced in one unit: dollars.

Third, it's a store of value. You can earn it today and spend it next month, and it still buys roughly the same amount, as long as inflation stays low. That "as long as" is important, and it's a big part of this chapter.

[SLIDE: A gold coin, then a paper banknote, then a payment card and a phone. Caption: "Money is whatever people trust to do the three jobs."]

Over history, lots of things have done these jobs: shells, salt, cattle, gold and silver coins. Today, in most countries, money is fiat money. It isn't backed by gold or anything else. A dollar bill is just paper, and most dollars don't exist as paper at all. They're numbers in bank accounts. So why does it work? Because the government says it must be accepted to settle debts and pay taxes, and, above all, because everybody trusts that everybody else will take it. Money is, in a real sense, a shared agreement.

So how much money is there? It depends where you draw the line. Economists use measures with names like M1 and M2. The narrow one, M1, is basically cash plus money in checking and savings accounts: money you can spend right away or move with a tap. In the U.S., savings accounts have been counted in M1 since 2020. The broader one, M2, adds things that take a little longer to turn into spending money, like small certificates of deposit and retail money-market funds. You don't need to memorize these. The point is that most money is bank deposits, not cash. And that leads to a surprising question: where do bank deposits come from?

What about crypto? It's a fair question. Something like Bitcoin can work as a medium of exchange in some places, but its price swings so much that, so far, it's done the store-of-value and unit-of-account jobs poorly. Very few prices are set in it. Whether that changes is a debate we'll leave to the future.

3 · How banks create money

[SLIDE: A bank in the middle. On the left, savers depositing money. On the right, borrowers taking out loans: a car, a small business, a college degree.]

Most people think of a bank as a vault: you put money in, it sits there, you take it out. That's not how it works. A bank sits between savers and borrowers. It makes loans, and it pays interest on deposits and its other funding. It earns money on the gap between the interest it charges borrowers and what it pays for that funding.

Say, with illustrative numbers, a bank pays savers one percent and charges borrowers six. On every hundred dollars it lends, it collects six dollars of interest and pays one to the saver. The five-dollar

gap has to cover the bank's staff, its branches, the loans that never get paid back, and, if anything's left, a profit for its owners. That gap is why banks want to lend. Money sitting in a vault earns them nothing.

And here's the surprising part. When a bank lends, it creates money. Let's slow down and see how a modern bank does it, with illustrative numbers.

[SLIDE: A simple balance sheet for "First Bank", drawn as a T. Left side, titled "What the bank owns": "New loan to you $20,000". Right side, titled "What the bank owes": "New deposit in your account $20,000". Caption: "A loan creates a matching deposit." Illustrative.]

Say First Bank approves you for a twenty-thousand-dollar car loan. Here's the bank's balance sheet. Economists draw it as a big T. On the left, what the bank owns: your loan, a promise to pay twenty thousand dollars back with interest. On the right, what the bank owes: a brand-new deposit of twenty thousand dollars in your account. The bank didn't take that money out of another customer's account. It simply credited yours. The loan and the deposit appear at the same moment, and twenty thousand dollars of money now exists that didn't a minute ago. Nobody printed anything. The loan created it. And when you pay the loan back, the reverse happens: the deposit and the loan shrink together, and that money disappears.

So can a bank lend without limit? No. Once the loan is made, the bank has to manage three things.

[SLIDE: Three tiles. "Funding & liquidity: can it settle payments and pay depositors?" "Capital: can it absorb losses?" "A good bet: will the borrower repay?"]

First, funding and liquidity. You'll spend that twenty thousand at a car dealer who probably banks somewhere else, so First Bank has to send reserves, the money banks use to pay each other, over to the dealer's bank. It needs enough reserves, or a way to borrow them, to settle payments like that, plus enough cash for depositors who want theirs. What it pays to get that funding feeds into the rate it charges you. Second, capital: the owners' own money in the bank, which absorbs losses if loans go bad. Rules require banks to hold enough of it relative to what they lend. Third, the plain question of whether the loan is a good bet, and whether people want to borrow at the rate on offer. Those, not a fixed formula, are what limit lending in practice.

You'll often see a simpler story in textbooks, and it's worth knowing as exactly that: a simplified model.

[SLIDE: "Simplified textbook model". A stack of shrinking bars: $1,000, $900, $810, $729… Caption: "If every bank kept 10% and lent the rest: total deposits ≤ $1,000 ÷ 0.10 = $10,000." Illustrative.]

Picture a thousand dollars of cash deposited at a bank that keeps ten percent and lends out the other nine hundred. That gets spent and deposited at a second bank, which keeps ninety and lends eight

hundred and ten, and so on, each round a bit smaller. Add up the rounds and deposits approach ten thousand dollars: a thousand divided by ten percent. That shortcut is called the money multiplier.

Treat it as a picture, not a rulebook. It gets the big idea right: bank lending expands the money supply, and it can expand it many times over. But real banks don't wait for a deposit and then lend out a fixed share of it, and in the U.S. required reserve ratios were set to zero in 2020. The real limits are the ones we just saw: funding, capital and willing borrowers.

One more piece, because it sets up the next section. Remember reserves, the money banks use to settle payments with each other? Most of them aren't cash in a vault. They're money the bank keeps in its own account at the central bank. And only one institution can create brand-new reserves: the central bank itself. Ordinary banks create deposits when they lend. The central bank creates the reserves underneath the whole system. That's a big part of where its power comes from.

Here's your first check.

[SLIDE: "Pause and predict #1". Text: "People get nervous about banks, pull cash out of their accounts, and keep it at home. What happens to the total amount of money in the economy, and to lending?"]

Pause and predict. Suppose people get nervous about banks. They pull cash out of their accounts and keep it at home. What happens to the total amount of money in the economy? And what happens to lending? Pause and think, then come back.

Here's the answer, and it surprises a lot of people. At first, the total amount of money doesn't change. When you withdraw a hundred dollars, your bank deposit goes down a hundred and the cash in your pocket goes up a hundred. Both count as money, so you've only changed its form. The trouble is a separate step. Cash leaving the banks drains their reserves, and if lots of people do it at once, banks get squeezed. They turn cautious, hold more cash, and make fewer new loans. And since new loans create new deposits, less lending can shrink the money supply over time, and spending slows with it.

And in the extreme, you get a bank run. Remember, a bank only keeps a slice of deposits on hand. If everyone asks for their money on the same day, even a healthy bank can't pay. So the fear of a bank failing can make it fail. That's why most countries insure bank deposits. In the U.S., deposit insurance covers up to two hundred and fifty thousand dollars per depositor, per bank, for each type of account ownership. If you know your money is safe, you have no reason to run. Rules on how much capital banks must hold, and regular check-ups on their finances, do the same job from the other side.

It helps to know the two different ways a bank can get into trouble. One is running out of cash. The bank's loans are fine, but they're tied up for years, and depositors want their money today. That's a liquidity problem, and a loan from the central bank can solve it, because the bank is good for the

money eventually. The other is worse. The bank's loans have gone bad, so what it owns is worth less than what it owes. That's called being insolvent, and no short-term loan fixes it. The bank's owners lose their stake, and the bank gets closed, sold or rescued. Capital is the cushion between those two: the owners' own money in the bank, which absorbs losses before depositors ever feel them. That's why regulators care so much about how much capital a bank holds.

4 · The Fed and its tools

[SLIDE: "The Federal Reserve: the central bank of the United States". Two goals underneath: "Maximum employment" and "Stable prices (about 2% inflation a year)".]

Every big economy has a central bank: a bank for the banks and for the government. In the U.S. it's the Federal Reserve, usually just called the Fed. It was set up by Congress in 1913. It came after a long run of banking panics, when there was no one to step in and lend to sound banks during a run.

The Fed has a job description from Congress, usually called its dual mandate: maximum employment and stable prices. It also aims for moderate long-term interest rates. In practice, "stable prices" means the Fed aims for inflation of about two percent a year over time. Not zero, for the reasons we saw in Chapter 4: a little inflation gives some room, and falling prices tend to be bad news.

And here's the committee from the puzzle. The decisions about interest rates are made by the Federal Open Market Committee, the FOMC. It has twelve voting members: the seven governors on the Fed's board in Washington, the president of the New York Fed, and four of the other eleven regional Fed presidents, who take turns. It meets eight times a year.

[GRAPH: A horizontal band labeled "Target range for the federal funds rate". A line wiggling inside the band labeled "Actual overnight rate between banks". Above the band, a line labeled "Rate the Fed pays banks on their reserves (near the top)"; below, a floor line.]

What the committee actually decides is a target for one particular interest rate: the federal funds rate. That's the rate banks charge each other for overnight loans. Banks keep accounts at the Fed, and at the end of each day some have more reserves than they need and some have less. The ones with extra lend to the ones short, overnight. The Fed sets a target range for that rate, like a band a quarter of a percentage point wide.

[SLIDE: "Real example: November 2026. Here's what these numbers looked like when this course was made."]

Let's look at a real example. Here's what these numbers looked like in November 2026, when this course was made. They're here to show you how to read the numbers, so focus on what each one means.

[SLIDE: "Federal funds target range: 3.75%–4.00%. Held at the October 28, 2026 meeting (10–2 vote). Real example, November 2026. Source: Federal Reserve."]

The committee held that range at its October meeting, by a vote of ten to two.

How does the committee decide? Mostly by reading the scoreboard from Chapter 4. Before each meeting, members pore over the latest numbers on jobs, wages, spending, GDP and, above all, inflation, plus reports from businesses in each region. If inflation is running well above two percent, they lean toward higher rates. If unemployment is rising and inflation is tame, they lean toward lower rates. When the two goals point in opposite directions, which happens, that's when the arguments get serious, and when you see members vote against the decision.

But how does the Fed make a market rate land inside its band? It doesn't order banks to charge it. It uses tools.

Tool one, today the main one: the interest rate it pays on the reserves banks keep at the Fed. Think about it. If the Fed pays banks, say, four percent, risk-free, on money parked at the Fed, no bank will lend its reserves to another bank overnight for much less than that. Why would it? So raising the rate it pays on reserves pulls the whole overnight rate up with it. Lowering it pulls the rate down. The Fed has a second, similar tool that reaches some other big financial firms, which helps set a floor under the band.

Tool two: open market operations. The Fed buys and sells government bonds. When it buys bonds from banks, it pays by adding reserves to their accounts. More reserves floating around tends to push rates down. Selling bonds does the opposite. For most of the twentieth century, this was the Fed's main tool.

Tool three: the discount window. The Fed lends directly to banks that need cash, usually at a rate at or above the top of its target range. It's a backstop. A bank that can always borrow from the Fed in a pinch shouldn't have to sell everything in a panic. That's the "lender of last resort" role, and it's one way the Fed fights bank runs.

And one more, for emergencies: large-scale bond buying, often called quantitative easing, or QE. When the Fed has already cut its short-term rate to near zero and the economy still needs help, it can buy huge amounts of longer-term bonds to push longer-term interest rates down directly. The Fed used this in and after the 2008 financial crisis and again in 2020. Running it in reverse, letting those bonds roll off, is often called quantitative tightening.

Finally, a word on independence. The Fed's leaders are appointed by the president and confirmed by the Senate, but the Fed doesn't need approval for each decision. The idea is that interest-rate

decisions should look at the long run, not the next election, because the temptation to keep rates low to give the economy a quick boost is always there, and the inflation bill tends to arrive later.

5 · Interest rates and why they move

[GRAPH: "The market for loans". Vertical axis "Interest rate (%)". Horizontal axis "Amount lent ($ billions)". An upward-sloping curve labeled "Supply of loans: savers". A downward-sloping curve labeled "Demand for loans: borrowers". They cross at an illustrative 5%.]

Now let's step back. What is an interest rate, really? It's a price. The price of borrowing money for a period of time. And like any price in this course, it's set by supply and demand.

Here's the market for loans. Economists call it the loanable funds market. On one side are savers, who supply money: people with savings accounts, retirement funds, businesses with spare cash. Higher interest rates reward saving, so the supply curve slopes up. On the other side are borrowers: families buying cars or homes, businesses building factories, students, and the government. Higher rates make borrowing more expensive, so fewer projects are worth it, and demand slopes down. Where they cross, with these illustrative curves, you get the market interest rate. Say five percent.

So what moves rates? Anything that shifts these curves.

Why does borrowing have a price at all? Because of opportunity cost, from Chapter 1. If you lend someone a thousand dollars for a year, you can't spend it, and you can't invest it somewhere else, for that year. You're also taking the risk that you won't get it back, and that prices will rise in the meantime. Interest is what pays you for all of that. And most people would rather have something now than the same thing later, so to get them to wait, you have to offer them a bit more later. Interest is the reward for waiting.

If people get nervous and save more, the supply of loans shifts right, and rates tend to fall. If businesses get excited about new projects and want to borrow more, demand shifts right, and rates tend to rise. And here's an idea we promised in Chapter 4: if the government borrows a lot to cover big deficits, that's more demand for loans, which pushes rates up. Higher rates then make some private borrowers, like a business deciding whether to build a new store, give up.

[GRAPH: The loans market again. The demand curve shifts right, labeled "Government borrows more". The crossing point moves up from an illustrative 5% to 5.5%. A bracket along the horizontal axis labeled "Private borrowing squeezed out".]

Here it is on the graph, with illustrative numbers. The government's extra borrowing shifts the demand for loans to the right. The crossing point moves up, from five percent to five and a half. At the new, higher rate, the government gets its money. But some families and businesses who would have borrowed at five percent decide it's not worth it at five and a half. Economists call that

crowding out: government borrowing can push some private borrowing out of the market. How big the effect is, in real life, is debated, and it depends on how much savings are available from at home and abroad.

But one rate is never the whole story. There are lots of interest rates at the same time: on a credit card, a car loan, a savings account, a government bond. Why do they differ? Two main reasons, plus one ingredient that's inside all of them.

[SLIDE: Three tiles. "Risk: will I get paid back?" "Time: how long is my money tied up?" "Inflation: built into every nominal rate. What will the money be worth when it comes back?"]

One: risk. Lenders charge more when they're less sure of being paid back. A loan to the U.S. government is considered about as safe as it gets. A car loan is secured by the car, so if you don't pay, the lender can take it. A credit card balance is backed by nothing but your promise. So, generally, credit cards charge more than car loans, and car loans more than government bonds. That extra is a risk premium.

Two: time. Lending for thirty years is usually riskier and less flexible than lending for three months, because a lot more can happen, including rates moving against you. So longer loans usually pay a higher rate. That extra for tying money up longer is called a term premium, and we'll use it in a minute. Longer doesn't always pay more, though, and when that flips, people pay attention. We'll see why in Chapter 6.

And then the ingredient inside every rate: inflation. This one is so important it needs its own moment.

[SLIDE: "Real interest rate ≈ nominal interest rate − expected inflation". Worked example, illustrative: "Savings account pays 5%. Inflation is 3%. Real return ≈ 2%." Second row: "Pays 2%, inflation 3% → real return ≈ −1%".]

Remember real versus nominal from Chapter 4. The interest rate you see, on a loan or a savings account, is the nominal rate. But what matters is what your money can actually buy. Say your savings account pays five percent, and prices rise three percent over the year. Your money grew five percent, but each dollar buys three percent less. So your real return is about two percent. That's the real interest rate: roughly the nominal rate minus inflation.

And if your account pays two percent while inflation is three, your real return is about minus one percent. You have more dollars and can buy less. That's why, when people expect higher inflation, lenders demand higher interest rates to protect themselves. Expected inflation gets built into interest rates. Hold on to that. It's a big part of the puzzle.

6 · From the Fed to your loan

[SLIDE: A chain of five boxes, left to right: "Fed's target rate" → "Overnight bank rates" → "Short-term rates: savings accounts, credit cards, adjustable loans" → "Long-term rates: 10-year Treasury" → "30-year mortgage rate". Arrows get thinner and dotted toward the right.]

Now we can connect the committee in the puzzle to the loan in your hand. Follow the chain.

The Fed sets a target for the overnight rate between banks. That rate is the foundation for other short-term rates almost right away. Banks change the rates they pay on savings accounts. Many credit cards and adjustable loans are tied to a benchmark that moves with the Fed's rate, so they change within a billing cycle or two. Those are the solid arrows on the left. Short-term rates follow the Fed closely.

But long-term rates are different. Look at the dotted arrows. The Fed doesn't set the rate on a ten-year government bond, or on a thirty-year mortgage. Those are set by markets: by millions of lenders and borrowers trading every day.

So how do long rates get decided? Think about it from a lender's point of view. Lending for ten years is a bit like lending for one year, then rolling it over, ten times in a row. So a ten-year rate is roughly two pieces. One: what lenders expect short-term rates to average over those ten years. Two: the term premium, the extra for tying money up that long and taking the risk that rates move against you.

[SLIDE: "Long-term rate ≈ expected average of future short-term rates + term premium". Small note underneath: "Expected inflation works through both pieces."]

So where does inflation fit? Inside both pieces. If people expect higher inflation, they expect the Fed to keep short rates higher, and they want a bigger cushion for the extra uncertainty. Inflation matters a lot, but it works through those two pieces, not as a third one stacked on top.

That word, "expect", does most of the work.

And how do you know what markets expect? You can actually see it. Traders buy and sell contracts whose value depends on where the Fed's rate will be in the coming months, so their prices reveal what the market thinks will happen. Financial news often reports it as something like "markets see a strong chance of a cut at the next meeting". Those odds change every day, as new numbers on jobs and inflation come out. A strong jobs report or a hot inflation reading can shift expectations, and long rates, before the committee has said a word.

[SLIDE: "Real example, November 2026": "Fed funds target: 3.75–4.00% (Federal Reserve, Oct 28 meeting)", "3-month Treasury: 4.19% (U.S. Treasury daily yield curve, Nov 6)", "10-year Treasury: 4.86% (U.S. Treasury daily yield curve, Nov 6)", "30-year fixed mortgage: 7.18% (Freddie Mac weekly survey, Nov 5)". A staircase

graphic with each rate as a step. Note on the 3-month step: "Above the Fed's top rate. One likely reason: markets expected another hike."]

Here's the real chain from that November 2026 example. On the bottom step, the Fed's target range. Next, the three-month Treasury bill, a very short loan to the government, at about four point one nine percent. Then the ten-year Treasury, at about four point eight six. And the average thirty-year fixed mortgage, from Freddie Mac's weekly survey, at about seven point one eight.

Notice the staircase. Short rates usually sit close to the Fed's target. But look closely: here, the three-month rate is actually above the top of the Fed's range. How can a three-month loan pay more than the Fed's own rate? One likely reason is that markets expected another hike at the next meeting, in December. A three-month bill lives through that meeting, so investors want the rate they expect over those three months, not just today's. Other, more technical things can push bill rates around too, like how many bills the Treasury is selling, so read this as a likely explanation, not proof. Either way, it's a live example of this chapter's main idea: expectations can move rates before the committee says a word. The ten-year sits higher, reflecting where markets expected short rates to go over ten years, plus a term premium. And the mortgage sits higher again, because lending to a homeowner is riskier than lending to the government, and because homeowners can repay early when rates fall, which lenders charge for. That's why mortgage rates tend to move with the ten-year Treasury rate, more than with the Fed's rate directly.

[SLIDE: Worked example, illustrative: "$30,000 car loan over 5 years". "At 5%: about $566 a month, about $3,970 in total interest". "At 8%: about $608 a month, about $6,500 in total interest".]

And these differences add up. Here's a quick illustrative example with a car loan, to keep housing for Chapter 6. Borrow thirty thousand dollars for five years at five percent, and the monthly payment is about five hundred and sixty-six dollars. Over the whole loan, you pay roughly four thousand dollars in interest. Same loan at eight percent: about six hundred and eight dollars a month, and about six and a half thousand dollars in interest. Three points of interest rate means about forty-two dollars more every month, and more than sixty percent more interest over the life of the loan. For a borrower, that's what a rate move actually feels like.

Now here's your second check.

[SLIDE: "Pause and predict #2". Text: "For weeks, everyone has expected the Fed to raise its rate by a quarter point at its next meeting. On the day, it does exactly that. What do you think happens to mortgage rates that day?"]

Pause and predict. For weeks, everyone has expected the Fed to raise its rate by a quarter point at its next meeting. On the day of the meeting, it does exactly that, no surprises. What do you think happens to mortgage rates that day? Pause, take a guess, then come back.

The answer: probably not much. Maybe nothing at all. Because the move was already expected, lenders and bond traders had already built it into longer-term rates, weeks before. The mortgage rate rose back when the expectation formed, not on the day.

And that's the payoff for our opening puzzle.

[SLIDE: The opening picture again: the committee meeting and the loan quote. Three captions: "Markets move on expectations", "Surprises move rates on the day", "The words matter as much as the decision".]

Why does your mortgage rate jump when a committee in Washington meets? Here's the answer. The committee doesn't set your mortgage rate. It sets a short-term rate, the bottom step of the staircase. Your mortgage rate is set by markets, and it rides mostly on what lenders expect about where the Fed's rate is headed over the next several years, which depends heavily on where they think inflation is headed, plus a term premium.

So when does the rate jump on meeting day? When the committee surprises people. A bigger move than expected. Or, very often, just new words. Along with its decision, the Fed releases a statement, holds a press conference, and at some meetings publishes its members' own projections for where they think rates are going. If those signals suggest rates will stay higher for longer than people thought, longer-term rates can jump that afternoon, even with no change to today's rate at all. And if the meeting goes exactly as expected, the jump already happened weeks earlier, as expectations built up. Notice that this also explains why rates can move on days with no meeting at all: a surprising inflation report, a speech by a Fed official, a strong or weak jobs number. Each one changes what people expect the committee to do next. That's the stranger part of the puzzle: the mortgage rate moved before the meeting, because the market saw it coming.

There's even a twist. Sometimes the Fed cuts its rate and long-term rates go up. How? If markets decide the cut will lead to more inflation later, they expect the Fed to have to push rates back up down the road, and they want a bigger term premium for the extra uncertainty. Both pieces of the long rate rise. The Fed controls the bottom step. Expectations do most of the work on the rest of the staircase.

7 · Inflation vs interest rates

[SLIDE: The chain of effects when the Fed raises rates, as a row of dominoes: "Higher rates" → "Borrowing costs more, saving pays more" → "Less spending on cars, homes, equipment" → "Businesses sell and hire less" → "Price rises slow".]

So why does the Fed raise rates at all? Mostly to fight inflation. Let's see how, and what it costs.

When the Fed raises rates, borrowing gets more expensive. A car loan costs more each month, so some people wait. A business deciding whether to build a new warehouse finds the loan pricier, and some projects stop making sense. Credit card balances cost more, so people spend a bit less. At the same time, saving pays more, so some people save instead of spending. Put it all together and spending across the economy slows. Remember the speed limit from Chapter 4. When spending outruns what the economy can produce, prices rise. Slowing spending takes the pressure off, and price rises tend to slow.

When the economy is weak and unemployment is rising, the Fed does the opposite. It cuts rates. Borrowing gets cheaper, saving pays less, and spending picks up.

But it can only cut so far. Once its rate is close to zero, the Fed can't easily push it much lower, because people could just hold cash instead of accepting a negative rate on a deposit. Economists call that the zero lower bound. In 2008 and again in 2020, the Fed hit it, which is when it reached for the emergency tool from earlier: buying long-term bonds to push longer rates down directly.

And remember, the Fed isn't steering alone. The government's spending and taxes, from Chapter 4, are pushing on the same economy at the same time. If the government is spending heavily while the Fed is trying to cool things down, the two are pulling in opposite directions, and the Fed may have to raise rates further to get the same effect. When both push the same way, like cutting rates and boosting spending together in a deep recession, the effect is stronger. Economists call the combination the policy mix.

But there's a catch, and it's the hardest part of the Fed's job: timing. These effects take a long time, often a year or more, to work through the economy. So the Fed is always steering by looking ahead, like driving a big ship that turns slowly. Raise rates too little, too late, and inflation can take hold. Raise them too much, and you can tip the economy into a recession and throw people out of work.

That's the trade-off at the heart of central banking. In the short run, cooling inflation often means a weaker job market, at least for a while. Economists sometimes draw this as a curve called the Phillips curve: lower unemployment tends to come with higher inflation, and the other way round. But over longer periods, that trade-off tends to break down. If the Fed tries to keep unemployment below its natural level by keeping rates too low, people come to expect higher inflation, build it into their wages and prices, and you end up with more inflation and no lasting gain in jobs.

Which brings us back to expectations, from Chapter 4. If people believe the Fed will keep inflation near two percent, they set wages and prices that way, and inflation tends to stay near two percent. That makes the Fed's job much easier. This is why central banks talk so much, and why their words can move markets as much as their decisions. Credibility is one of their most valuable tools.

[SLIDE: "Real example, November 2026": "Consumer prices: +3.74% over 12 months (September data, BLS)", "Core prices: +2.56% (September data, BLS)", "Unemployment: 4.1% (October data, BLS)", "Fed funds target: 3.75–4.00% (Federal Reserve, Oct 28 meeting)". Small note: "Rough arithmetic, not an official measure: fed funds minus headline inflation ≈ 0; minus core ≈ +1.3 points."]

Let's put real numbers on this, with the November 2026 example. Consumer prices were up about three point seven percent from a year earlier, and core prices, without food and energy, about two point six percent. Unemployment was four point one percent. Here's a rough piece of arithmetic, not an official measure: take the middle of the Fed's range and subtract headline inflation, and the real short-term rate comes out around zero. Subtract core inflation instead, and it's a bit more than one percentage point above zero. Which inflation measure you use changes the picture, and that's the kind of judgment the committee argues about.

One last connection: printing money. Why can't a government just create more money to pay for things? Over short periods, more money and lower rates can boost spending. But if the money supply grows much faster than the economy's ability to make things, you get more dollars chasing the same goods, and prices rise. Over long periods, countries with very fast money growth have tended to have high inflation, and the most extreme cases of runaway inflation have usually involved governments creating money to pay their bills. That's another reason central banks are kept at arm's length from the people who do the spending.

8 · Recap, cheat sheet & self-check

[SLIDE: "Recap". Six short lines appear one at a time.]

Let's lock it in, in six lines.

Money does three jobs: medium of exchange, unit of account, and store of value. Today's money is fiat money, and most of it is bank deposits, not cash.

Banks create money when they lend: a new loan creates a matching new deposit. Then the bank manages its funding, liquidity and capital. The money multiplier is a simplified textbook model of the same idea. Deposit insurance and the Fed as lender of last resort exist to stop bank runs.

The Fed's dual mandate is maximum employment and stable prices, which in practice means about two percent inflation. Its committee, the FOMC, sets a target range for the overnight rate between banks, mainly by setting the rate it pays on bank reserves.

An interest rate is the price of borrowing, set by savers and borrowers. Rates differ by risk and by time, and every nominal rate builds in expected inflation. The real rate is roughly the nominal rate minus inflation.

The Fed controls short rates closely. Long rates, like the ten-year Treasury and mortgages, are set by markets: roughly expected future short rates plus a term premium. So they move on surprises and signals, often before the decision day itself.

Higher rates slow spending and, with a lag, inflation, at the risk of slowing jobs too. Lower rates do the opposite.

[SLIDE: "Cheat sheet". Left: the bank T-account, "a loan creates a matching deposit", the three limits (funding & liquidity, capital, a good bet), and a small box "Textbook model: ceiling = first deposit ÷ share kept". Middle: the loanable funds graph, with arrows for what shifts each curve, and "real ≈ nominal − expected inflation". Right: the staircase from the Fed's target to the mortgage rate, with "long rate ≈ expected short rates + term premium", the November 2026 numbers labeled "Real example, November 2026", and the domino chain from higher rates to slower inflation.]

Here's your cheat sheet. On the left, how banks create money. In the middle, what sets interest rates and how to turn a nominal rate into a real one. On the right, the staircase from the Fed's rate up to the rate on a mortgage, with the November 2026 example, and the chain from higher rates to lower inflation. Any time you see a headline about the Fed, ask two questions. Was it a surprise? And what does it signal about where rates are going?

Now three questions to check yourself. Pause after each one, answer out loud, then let it play.

Question one. Your bank approves you for a ten-thousand-dollar loan and credits it to your checking account. Right after that, has the total amount of money changed? And what does the bank have to manage next?

Answer. Yes. There's ten thousand dollars of new money: a new deposit in your account, matched by your new loan on the bank's books. Nobody else's balance went down. Next, the bank has to make sure it can settle the payment when you spend the money and it moves to another bank, keep enough cash for depositors, and hold enough capital in case loans go bad. When you repay, that money disappears again.

Question two. Your savings account pays four percent. Prices are rising three percent a year. Next year, your bank raises your rate to five percent, but inflation rises to six. Are you better or worse off in real terms?

Answer. Worse off. This year your real return is about plus one percent, four minus three. Next year it's about minus one, five minus six. Your interest rate went up, but your money buys less. Always subtract inflation.

Question three. The Fed leaves its rate unchanged, but in its statement it signals that it expects to keep rates higher for longer than markets had thought. What probably happens to the ten-year

Treasury rate and mortgage rates, and why?

Answer. They probably rise. Long-term rates depend on where people expect short-term rates to be over the coming years. A signal of higher rates for longer raises that expected path, so long rates go up, even though today's rate didn't change. That's the puzzle again: the words can matter as much as the decision.

[SLIDE: "Where we're going". The 6 chapters listed, 1 to 5 checked: "1 Foundations & Supply and Demand", "2 How Markets Work", "3 When Markets Fail", "4 The Whole Economy", "5 Money, the Fed & Interest Rates", "6 Trade, Stocks & Bonds, and Real Estate". Chapter 6 highlighted: "Next up".]

That's Chapter 5. You now know what money is, how banks create it, what the Fed does, and how a decision in one room travels out to the loans and savings of millions of people.

In Chapter 6, the last one, we take all of it out into the world. Why countries trade, and what tariffs do. How stocks and bonds work, and why rising interest rates can push their prices down. And then we put the whole course to work on one of the biggest markets of all: real estate. See you in Chapter 6.