



Economics and the Equity Market: A Microeconomics Course Application

1

Economics encompasses two broad subjects: macroeconomics and microeconomics. Macroeconomics deals with an economy in aggregate and addresses issues such as inflation, unemployment, interest rates, and economic growth. We present a macroeconomic perspective in Chap. 3. Microeconomics, the focus of this chapter, operates, as its name indicates, on the micro level, addressing household consumption decisions and the production decisions of firms. In this chapter, we focus on the parallels (and a few differences) between a standard microeconomics formulation (a household's selection of an optimal consumption bundle) and a standard finance model (an investor's selection of a portfolio that optimally combines a riskless asset – cash – and a risky equity portfolio). The finance formulation is the *Capital Asset Pricing Model (CAPM)*. CAPM is a keystone of what is known as modern portfolio theory, the originator of which is Harry Markowitz who was awarded a Nobel Memorial Prize in Economic Sciences in 1990 for having developed the theory of portfolio choice.

In both formulations, price plays a central role as it guides the decisions of both households and investors. Along with the decisions of households and firms, determination of an *equilibrium price* is of paramount importance. The price variable is so important that microeconomics courses can carry and have carried the name “price theory.”

It is one thing to analyze price equilibrium in a theoretical model, and something else for an equilibrium price to actually be attained in a real-world market, especially one where prices are changing with great frequency, as is the case in an equity market. A primary function of a financial marketplace such as the New York Stock Exchange or Nasdaq is to facilitate attainment of equilibrium prices, an objective referred to as *price discovery*. Effective price discovery, however, is not easily achieved. We discuss this in considerably more detail in Chap. 2 (Finance) and in Chap. 3 (Macroeconomics).

For most of our discussion in this chapter, we assume, as is standard in much microeconomics, that the marketplace is a totally frictionless environment. By *frictionless*, we mean that there are no fees, taxes, or other impediments to buying and

selling, which, therefore, are costless activities. Only toward the end of the chapter do we relax this assumption and consider the impact that friction can have on price determination and the operations of a real-world equity market.

While a theoretical, frictionless market equilibrium might not be fully achieved in a real-world marketplace, an unobservable equilibrium price nevertheless exerts a force that improves the quality of market outcomes. This force merits being understood. By way of analogy, one might think of the power of the Gulf Stream, a strong, deep sea ocean current that brings warm water into the Atlantic Ocean from the Gulf of Mexico, moves up the Atlantic coast, and branches out to Europe. A ship crossing the Atlantic should take account of the Gulf Stream, but the vessel also has to contend with the winds, waves, and storms on the surface of the sea. One might equate the power of the Gulf Stream with the force exerted by an unobservable frictionless market equilibrium price and equate the wind, waves, and storms with frictions that buffet real-world, non-frictionless markets.

1.1 Microeconomics in a Nutshell

The terms *optimum*, *maximum*, and *equilibrium* play a key role in microeconomic analysis. Households are assumed to make “optimal” decisions when confronted by something that lies at the heart of a microeconomic problem: resolving a trade-off between alternative possibilities (e.g., get a little more of this and a little less of that, or vice versa). Optimality is achieved with regard to the decision maker’s single, ultimate goal – maximize his/her personal utility. In a two-good environment (X and Y, for simplicity), a household determines the optimal amount of X to buy relative to Y when, because of a resource constraint (income or wealth), more of X can be obtained if and only if less of Y is obtained, and vice versa (more Y and less X). Having allocated its scarce resources optimally and, in so doing, having maximized utility, a household is in equilibrium.

A firm makes two optimal decisions in order to achieve a single goal – the maximization of profits. In a two-input environment (again, we are keeping it simple), a firm maximizes profits by (1) optimally combining L (let us call it labor) and C (let us call it physical capital) and (2) producing an optimal quantity of its product (let us stay with X). When a firm has done this, it too is in equilibrium.

A household’s utility maximizing decisions are made with respect to tastes, income, and the prices of X and Y. An X-producing firm’s profit maximizing decisions are made with respect to technology, the price of the product it is producing (the price of X), and the prices of its factors of production (L and C). When numerous households are consuming X, when many firms are producing X, and when all households and all firms are in equilibrium, the market for X is in equilibrium. We can obtain this equilibrium with the use of a downward sloping market demand curve to consume X and an upward sloping market supply curve to produce X.¹ The

¹ Intuitively stated, a price decrease attracts buyers to buy more, and a price increase attracts sellers to supply more.

intersection between these two curves identifies the equilibrium price of X and the quantity of X produced and consumed.

Note the critical role played by the prices of X, Y, L, and C in enabling households and firms to achieve equilibrium. As we have already noted, price is a key variable in microeconomics. The price of X, for instance, is a guiding light since it signals what has to be given up in return for more X. Because we have assumed that many individuals consume X and that many firms produce X, we take the market for X to be *perfectly competitive*. Accordingly, no participant, on either the demand or the supply side, is big enough to individually affect the price of X. Thus, all participants, both households and firms, are *price takers*, and the price that each one of them faces is determined by all participants as they meet collectively in the marketplace for X.


1.2 Microeconomic Analysis Goes to an Equity Market

Now let X be a share of stock, and take a microeconomics journey to an equity market to see how optimal investment decisions are made and how a stock's share price is determined. Following standard microeconomic methodology, we start by making assumptions that let us get answers that we are looking for while keeping the analysis as simple as possible. Here is our first assumption: the equity market is frictionless. Our second assumption is that, as with the consumer choice model, all investors are small, retail customers, and they are sizable enough in number so that no one of them has the power to individually affect the price of shares. Thus, the market is perfectly competitive, and all participants are price takers.

Like the highly simplified two-good consumer choice analysis of households, our treatment of the equity market deals with the optimal allocation of scarce financial resources between, not two goods, but two assets: a risky stock portfolio and cash (the riskless asset). Should the investor hold more stock and less cash or more cash and less stock? Like the consumer, the investor makes this allocation decision with reference to one goal: maximize expected personal utility. Why "expected utility," not just "utility"? For a simple reason. With the introduction of a risky asset, we are operating in an uncertain environment where the outcome is unknown.

As in the simplified two-good consumer model, allocation in the two-asset investment model is made with reference to the decision maker's tastes. But it is not taste or a preference for stock versus cash per se that matters. It is taste for two attributes of a financial asset: risk and return. For cash, risk is zero and the return is low; for a stock portfolio, there is risk and the return is higher. By substituting stock for cash, the investor increases his/her expected return and accepts more risk; by substituting cash for stock, the investor decreases his/her expected return and incurs less risk. The allocation decision is made with regard to this risk, return trade-off.

Exhibit 1.1 Indifference curve: a good versus good trade-off




In the two-good consumer model, X and Y both deliver utility, and thus we have a “good versus good” trade-off. In the two-asset model, positive returns deliver increased utility, but investors are risk averse and so risk delivers “disutility” (i.e., a decrease in utility). Accordingly, with regard to the financial assets, we are dealing with a “good versus bad” trade-off.

1.3 Risk, Return Indifference Curves

Exhibit 1.1 displays a key microeconomics tool, an *indifference curve*, to depict a consumer’s tastes for a good versus good trade-off. Because an indifference curve is the locus of points that deliver the same utility and because X and Y both yield positive utility, the indifference curve is downward sloping (it is also convex to the origin). All consumption combinations above and to the right of the indifference curve shown in Exhibit 1.1 yield more utility than combinations on the indifference curve, and all combinations below and to the left yield less. Accordingly, any X, Y combination above and to the right of the curve shown in Exhibit 1.1 lies on a higher indifference curve, and any X, Y combination below and to the left lies on a lower indifference curve.

Exhibit 1.2 presents an indifference curve that depicts an investor’s tastes for a good versus bad trade-off. Here, the axes are labeled, respectively, risk and return. The properties of the curve shown in Exhibit 1.2 are the same as those shown in Exhibit 1.1 with one exception: the curve in Exhibit 1.2 is upward sloping. Why? Because investors are risk averse, they view risk as a bad. So if you get more risk, you also must get more return to stay on the same indifference curve.

Exhibit 1.2 Indifference curve: a good versus bad trade-off



1.4 The Constraint

More structure is required for an optimal solution to be obtained. Specifically, a constraint is needed and, for the consumer, there is a constraint because income and wealth are scarce resources (for the investor, the situation is a bit different and we get to it shortly). For the consumer, the resource constraint is called a *budget constraint*. Let us take a look. Denote the consumer's income by M , and let M be allocated entirely to X and Y . Letting Y be the quantity of Y and X be the quantity of X , we have

$$M = P_Y Y + P_X X \quad (1.1)$$

Rearranging gives

$$Y = \frac{M}{P_Y} - \left(\frac{P_X}{P_Y} \right) X \quad (1.2)$$

Equation 1.2 is the budget constraint. We show it along with the indifference curve in Exhibit 1.3.

In Exhibit 1.3, a dot marks the point where the indifference curve is tangent to the budget constraint. The point of tangency marks the spot where the highest indifference curve (and thus the highest level of utility) can be reached (any other point along the budget constraint lies on a lower indifference curve). Accordingly, the point of tangency identifies the optimal (utility maximizing) combination of X and Y to consume.

The constraint is not as readily obtained for the investor as it is for the consumer. Instead of being determined by income and relative prices, the constraint depends on the risk-free rate for cash and on the risk and return combinations that are available. Regarding the risky component, are we dealing with one stock or with a portfolio of stocks? The answer is both. We first focus on the set of risky stocks and see how they can all be brought together to form one risky portfolio that includes all

Exhibit 1.3 Indifference curve and budget constraint

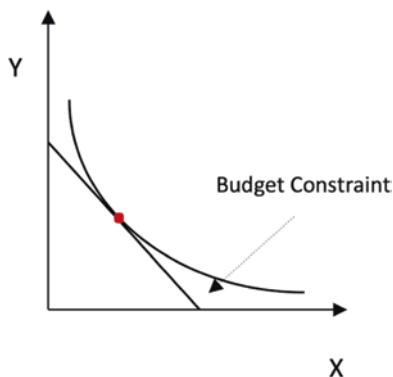



Exhibit 1.4 The efficient frontier



stocks. This all-inclusive portfolio is called the *market portfolio*.² Each stock's expected return, its standard deviation, and its covariance with the returns of other risky stocks are what we need to do this.³ As we proceed, for the moment, let us keep cash to the side.

We can assess risk and return for all single-stock and multi-stock portfolios, with risk measured by the standard deviation of returns and expected return defined as the expected percentage price change for an individual stock and for a stock portfolio. Exhibit 1.4 shows a set of arbitrarily selected dots that represents a risk, return mapping for single-stock and multi-stock portfolios. The curve to the left of the dots is an outer envelope that shows the stock or stock portfolio for which the standard deviation is the lowest for any given return. The upward sloping, concave portion of the envelope, which also shows the maximum return that can be obtained for any given level of risk, is the *efficient frontier*. Portfolios that are inside the efficient frontier are inefficient because a higher expected return is available for each level of

²As we discuss in more detail below, this market portfolio should include all risky assets. In practice, the market portfolio is usually represented by a broad-based index such as the SP500.

³In finance, we typically use standard deviation as a measure of risk because it measures the dispersion around an average, or expected, rate of return on an asset.

Covariance (or its statistical cousin, correlation) measures how closely one stock's returns move in tandem with another stock's returns.

risk and, equivalently, a lower risk is available for each level of expected return. Only portfolios on the efficient frontier maximize return for a risk and minimize risk for a return, which explains why the curve is called the efficient frontier.

There is more; we are not there yet. We have not taken account of cash, the risk-free asset. This is how to proceed. The expected return on the investor’s portfolio, $E(R_p)$, is

$$E(R_p) = (1 - k) \cdot R_f + k \cdot E(R_m) \tag{1.3}$$

where k (a weight) is the percentage invested in the risky portfolio, R_f is the risk-free rate, and $E(R_m)$ is the expected return on the market portfolio. SD_p , the standard deviation of returns for a portfolio that contains both a risky stock portfolio and cash, is⁴

$$SD_p = k \cdot SD_m \tag{1.4}$$

where SD_m is the standard deviation of returns for the market portfolio. From Eq. 1.4, we have $k = SD_p / SD_m$. Substituting for k into 1.3 and rearranging gives

$$E(R_p) = R_f + \left[\frac{E(R_m) - R_f}{SD_m} \right] SD_p \tag{1.5}$$

What specific portfolio on the efficient frontier should be held in combination with cash? As before, the objective is to identify a portfolio that offers the highest expected return for any level of risk (measured by the standard deviation, SD) or, equivalently, that minimizes SD for any level of expected return. Two upward sloping lines are displayed in Exhibit 1.5. Both lines have the same intercept, R_f , and both pass through a portfolio on the efficient frontier. But one line crosses through




Exhibit 1.5 The capital market line

⁴With regard to Eq. 1.4, keep in mind that for cash (the risk-free asset), $SD = 0$. Thus, from $R_p = kR_m + (1 - k)R_{\text{cash}}$, taking variances, we have $\text{Var}(R_p) = k^2 \text{Var}(R_m)$, from which Eq. 1.4 follows.

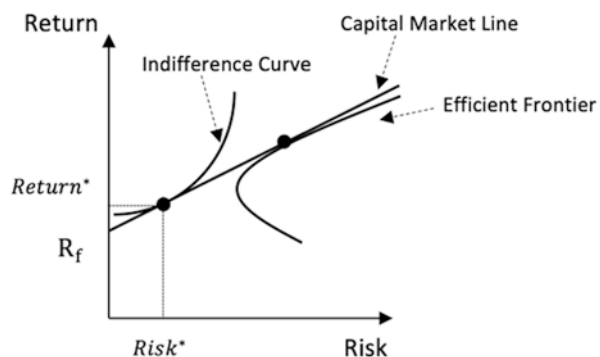
the efficient frontier, while the other is tangent to the efficient frontier. Which one of the two do you think is preferable? Answer: the line that is tangent dominates because, being steeper, it maximizes the expected return for each level of risk and, equivalently, it minimizes risk for each expected return.

Now go back to the flatter line in Exhibit 1.5 and picture rotating it counterclockwise with its intercept fixed at R_f . As it is rotated counterclockwise, it offers increasingly better risk, return combinations until it becomes tangent to the efficient frontier, at which point its slope is maximized. Thus, the line that is tangent is the one that we want. It is equivalent to the budget constraint shown in Exhibit 1.3. We call this line the *Capital Market Line*. The line forms the basis of a widely used financial model that is commonly referred to as the Capital Asset Pricing Model (CAPM). The portfolio that is on the efficient frontier at the point of tangency with the Capital Market Line is the *market portfolio*. We can now understand Eq. 1.5 as depicting the Capital Market Line. As noted earlier, the market portfolio should include all risky assets in the world but is usually represented by a large set of US stocks such as the S&P 500 stock index.

Unlike the budget constraint in the consumer choice model that varies from individual to individual according to his/her income and wealth, the Capital Market Line is the same for every investor, and so too is the market portfolio. But investors' tastes for risk vs. return differ from person to person, and each, with reference to his/her own indifference curves, selects the utility maximizing combination of the market portfolio and the risk-free asset. Where would the utility maximizing combination lie? It lies at the point where the investor's indifference curve is tangent to the Capital Market Line, as shown in Exhibit 1.6. Exhibit 1.7 contrasts the optimality solutions for two individuals with different tastes for risk and return.

Notice in Exhibit 1.6 that the indifference curve's point of tangency is to the left and below the Capital Market Line's point of tangency with the efficient frontier. This means that a relatively risk-averse individual will hold long positions in both cash and the market portfolio (i.e., the weight k will be less than 1). What if the indifference curve's point of tangency is to the right and above the Capital Market Line's point of tangency with the efficient frontier? In this case, the *less* risk-averse investor would borrow cash and invest a larger amount in the market portfolio (i.e.,

Exhibit 1.6 Investor optimality






Exhibit 1.7 Optimality solution for two different investors

the weight, k , will be greater than 1). We discuss the weight later in this chapter with respect to Eq. 1.9.

There we have it! This is how an investor's risk, return trade-off is resolved. To illustrate, consider two investors with indifference curves that are different. Sam will accept additional risk but only for a substantial increase in expected return. Rachel, on the other hand, will tolerate added risk for smaller improvements in expected return. Rachel's optimal portfolio is riskier than Sam's.

1.5 Demand Curve to Hold Shares of the Market Portfolio

Thus far, we see that the investment model closely parallels the standard consumer choice model. Now for the next, more advanced step. From indifference curves and budget constraints, we can obtain demand curves for the goods and services that households consume. From our risk, return analysis, are we able to obtain a downward sloping demand curve to hold shares of the market portfolio and of individual stocks?

For an individual stock, the answer in a frictionless, perfectly liquid environment is simple. Investors do not have tastes for the individual stocks themselves – all that matters to them is risk and return. In the Capital Asset Pricing Model, risk is measured by the beta coefficient, and two stocks or portfolio of stocks that have the same beta coefficient are perfect substitutes for each other.⁵ As such, they should be

⁵A beta coefficient (β) is an estimate of a stock's riskiness relative to the riskiness of the market portfolio (R_m).

Beta is a measure of systematic risk. Risk that refers to that component of a stock's return that is not related to the market return is referred to as non-systematic risk. Non-systematic risk is not priced because, in the frictionless environment, it can be eliminated by portfolio diversification (whereas systematic risk cannot be eliminated by diversification).

priced to yield the same expected return. And all stocks do have perfect substitutes because the risk of any stock can be replicated by an appropriately weighted combination of two other stocks. Because all stocks have perfect substitutes, the demand curve for each stock is horizontal (and thus infinitely elastic) at a price $P_{m0} = a$ that we explain below with respect to Eq. 1.6.

The demand curve for the market portfolio, however, is downward sloping. To obtain it requires analysis, and to this end, we first set forth our assumptions.

1. There are two financial assets: cash and the market portfolio.
2. A single holding period with a starting point in time denoted by 0 and an ending point in time denoted by T.
3. At point in time 0, the investor is holding only cash.
4. The investor maximizes expected utility given his/her initial cash holdings, the risk-free rate (R_f), risk, and $E(P_{mT})$ which is the expected price of the market portfolio at point in time T.
5. $E(P_{mT})$ is independent of P_{m0} , the price of the market portfolio at point in time 0.
6. The environment is frictionless. There are no commissions, borrowing costs, short-sale restrictions, etc.

With these assumptions, we can obtain the investor's demand curve to hold shares of the market portfolio. The demand curve, presented with the price P_{m0} on the left-hand side, is the linear equation:⁶

$$P_{m0} = a - 2b \cdot N \quad (1.6)$$

where

- N is the number of shares of the market portfolio purchased and held
- $a = E(P_{mT})/(1 + R_f)$ is the intercept parameter
- $b = \pi/(1 + R_f)$ is the slope parameter
- π is a risk premium that reflects the investor's risk aversion⁷


This demand curve is shown in Exhibit 1.8. Several things are of interest regarding this demand curve:

1. The consumption of most goods is measured as a rate per a period of time, for instance, the amount of X consumed per day, per month, or per year. This is not the case for financial assets, for they are not "consumed" over time. Rather, they are "held," and the demand curve shows that the number of shares held depends on the price of shares.

⁶The derivation is in Ho, Schwartz, and Whitcomb (1985), and further discussion is provided in Francioni, Hazarika, Reck, and Schwartz (2010).

⁷The risk premium is the compensation required for an investor to hold a risky asset instead of a riskless asset (cash).

Exhibit 1.8 Demand to hold shares of the market portfolio



2. Assumption 5 indicates that $E(P_{mT})$, a term in the expression for the intercept parameter, is independent of P_{m0} , the price of the market portfolio at point in time 0. By this assumption, we ignore the unnecessary complexity of the price intercept being related to the price at which transactions are made, which would be the case if price conveyed an informational signal. With independence, as the current price of the market portfolio decreases while $E(P_{mT})$ is constant, the expected return over the holding period rises. Or let us state it this way: if the current price goes down and the expected future price stays constant, the expected return must go up.
3. The demand curve described by Eq. 1.6 and shown in Exhibit 1.8 is linear. Linearity is desirable because it is a good deal easier to work with. Note that price is on the left-hand side of Eq. 1.6, not its usual location on the right-hand side of a demand equation. We have switched the side simply to obtain linearity. To achieve linearity, one further assumption is required: that shares are priced such that the difference between the return on the market portfolio and the risk-free rate is small enough to ignore.⁸ In other words, we have to remain in the area of the demand curve that is reasonably close to the price intercept. This certainly makes sense. With linearity throughout, the demand curve would intersect the quantity axis, at which point a finite number of shares would be held at a price of zero. This would not make sense. If shares were free, they should be held in unlimited amounts, and thus, the demand curve must become curvilinear as the price approaches zero in order to capture this behavior.⁹
4. The slope of the negatively inclined demand curve reflects the investor's risk premium parameter, π . The more risk averse the investor, the greater is the slope. If the investor is risk neutral rather than risk averse, the slope, b , would be zero, and the demand curve would be horizontal (infinitely elastic) at the price intercept, a . Why? Because, with risk neutrality, the risky asset and the risk-free asset are perfect substitutes.

⁸This assumption is explained in further detail in Ho, Schwartz, and Whitcomb (1985).

⁹As the price becomes asymptotically close to zero, the number of shares demanded, N , would increase without bound. So a linear relationship is not possible for all price points.

5. The intercept parameter, a , contains the term, $E(P_{mT})$. As we discuss in some detail in Chapter 2, different investors can have different (divergent) expectations of this end-of-period price, and thus, the location of the demand curve will differ from investor to investor. In consumer choice theory, the demand curve for good X differs from consumer to consumer because of differences in their tastes for the product and their wealth. For the risky asset, demand differs among investors because of differences in their taste (really, distaste) for risk and/or their expectations of future returns.

Equation (1.6) contains a number of parameters. Juggling them around, we can obtain the following equation:¹⁰

$$\frac{\pi'}{P_{m0}} = E(R_m) - R_f \quad (1.7)$$

where π' is the investor's marginal risk premium and, because of the division by P_{m0} , the left-hand side, like the right-hand side, has a percentage dimension. Equation 1.7 has an interesting interpretation, one that is equivalent to that of the consumer choice model. In the consumer choice model, at the point where the consumer's indifference curve is tangent to his/her budget constraint, we have

$$\frac{\text{Marginal Utility of X}}{\text{Marginal Utility of Y}} = \frac{\text{Price of X}}{\text{Price of Y}} \quad (1.8)$$

Equation 1.8 says that when the consumer has allocated his/her resources optimally, the rate at which he/she can substitute X for Y while keeping utility constant (the left-hand side of the equation) equals the rate at which X can be substituted for Y in the marketplace keeping total expenditures constant (the right-hand side of the equation). Thus, we see that the consumer harmonizes his/her tastes, on the margin, with the trade-off that is possible in the marketplace.

Equivalently, for the CAPM, at the point where the investor's indifference curve is tangent to the Capital Market Line (refer to Exhibit 1.6), the investor's marginal risk premium expressed as a percent of price (refer to Eq. 1.7) equals $(E(R_m) - R_f)$, which is the higher percentage return that the market offers the investor for accepting risk rather than holding cash. To understand this intuitively, think of the risk premium as a price: the price that the market will pay the investor for accepting risk. To repeat, we see that, in a frictionless world, the decision maker maximizes his/her expected utility by harmonizing his/her own tastes (on the margin) with the price of risk, which can be understood as the rate at which two inputs into utility (risk and return) can be substituted for each other in the market.

So this is what we have. In the consumer choice model, the rate at which X can be substituted for Y is the price of X relative to the price of Y. In CAPM, the rate at which risk can be substituted for certainty equals the risk premium that the market offers. In both the consumer choice model and CAPM, the decision maker

¹⁰Further discussion and its derivation are provided in Schwartz (1991) and in Francioni, Hazarika, Reck, and Schwartz (2010).

maximizes utility by harmonizing his/her utility trade-off (the left hand side of Eq. 1.7) with the trade-off that is possible in the market (the right hand side of Eq. 1.7).

By further juggling the parameters in Eq. (1.6), we can obtain¹¹

$$W = \frac{E(R_m) - r_f}{\text{Var}(r_m) R_R} \quad (1.9)$$

where W is an investor's weight in the market portfolio and R_R is a measure of his/her risk aversion. The equation makes intuitive sense: an investor's optimal weight in the risky market portfolio is greater (all else equal) the greater is the risk premium, $[E(R_m) - R_f]$, the less risky is the market portfolio $[\text{Var}(r_m)]$, and the less risk averse is the investor (the smaller is his/her R_R). Note that the only term in Eq. 1.9 that is specific to an investor is the risk aversion parameter, R_R .

Regarding the distribution of the weight W across investors, a participant with a relatively high R_R and $W < 1$ will have a long position in both cash and the risky market portfolio, while a participant with $W = 1$ will be neither a borrower nor a lender, and a participant with a relatively low R_R and $W > 1$ will borrow at the risk-free rate (i.e., acquire a short position in cash) so as to finance his/her leveraged long position in the risky market portfolio. Now, consider the dynamics that come into play when the long positions in cash of the $W < 1$ investors are the source of lending to the $W > 1$ investors. To this end, let us ask: how is equilibrium between the long and the short cash positions achieved when, to repeat, the long cash positions of some investors finance other investors' leveraged long positions in the market portfolio?

The total amount borrowed by $W > 1$ investors is brought into line with the total amount lent by $W < 1$ investors by an equilibrating variable which, consistent with standard microeconomics theory, is a price. The price is the risk premium, $E(R_m) - R_f$. In analyzing how the risk premium fulfills this role, let us for simplicity treat R_f as an exogenously determined constant so that the risk premium changes only with $E(R_m)$.¹² Note that the relative aggressiveness of the $W > 1$ investors exerts upward pressure on the share price of the market portfolio and, in so doing, decreases $E(R_m)$, while the relative nonaggressiveness of the $W < 1$ investors exerts downward pressure on the share price of the market portfolio and, in so doing, increases $E(R_m)$. When the upward and downward pressures are in balance, equilibrium has been achieved. In equilibrium, $E(R_m)$ and the risk premium would be at a level where the amount of cash borrowed by $W > 1$ investors just equals the amount that $W < 1$ investors desire to lend.

¹¹As noted above, further discussion and the equation's derivation are provided in Schwartz (1991) and in Francioni, Hazarika, Reck, and Schwartz (2010).

¹²In the financial markets, R_f is not determined by the balance between borrowing and lending in the equity markets alone. Other financial markets and the Federal Reserve Bank's macroeconomic policy are also major determinants of R_f . Regardless, we could obtain an equivalent result by holding $E(R_m)$ and allowing R_f to vary, rising when the demand for cash exceeds its supply, and falling when the supply of cash exceeds its demand.

1.6 What About the Supply Curve?

On to the next question. After we aggregate individual investor demand curves to get a market demand curve, how do we obtain the equilibrium price and number of shares of the market portfolio that each investor will hold? With regard to consumer choice analysis, we would do this by obtaining a supply curve to match against the demand curve to get an equilibrium solution. Can we now obtain a meaningful supply curve for shares?

One candidate for the supply curve that might come to mind is to take account of the number of shares of the market portfolio that are outstanding. To do this, we would aggregate the individual demand curves depicted in Exhibit 1.8 and match the aggregate with a vertical line located on the horizontal axis at the number of shares outstanding. Could this vertical line be the supply curve? We answer this question by considering what would happen if the number of shares outstanding is changed so that the location of the “supply curve” shifts. Say, the company splits its shares 2 for 1. Would this change the total value of shares outstanding? No, it would not. Following a 2 for 1 split, the price of shares would simply be cut in half. This being the case, what would the market’s demand to hold shares look like? It would be a downward sloping, convex curve that satisfies the condition that at all points along the curve, price times the number of shares outstanding is a constant. In other words, the demand curve would be a rectangular hyperbola. Such a curve would be valueless. Thus, the vertical line cannot be considered a supply curve.

In point of fact, there is no supply curve to match with the demand curve. Think about it. There are no separate suppliers, as there are for the consumption goods X and Y in the consumer choice model. In the financial model, any participant can be either a buyer (demander) or a seller (supplier) of shares, depending on the price of shares. This being the case, how do we proceed so as to identify an equilibrium price?

1.7 Buy and Sell Curves

There is another route to follow. As we take it, we relax the assumption that the investor’s starting position was 100% in cash and 0% in shares of the market portfolio. This assumption was not necessary for the analysis; we made it only to simplify the specification of the intercept and slope parameters, a and b .

Go back to the individual investor’s demand curve shown in Exhibit 1.8 and draw in a vertical line at the number of shares of the market portfolio that the investor is currently holding, N_{m0} . Such a line is shown in Exhibit 1.9. Go to the price where the vertical line intersects the demand curve. At all higher prices, the investor wants to hold fewer shares than are currently in his/her portfolio, and at all lower prices, the investor wants to hold more shares. In effect, prices *above* P_{m0} will turn the investor

into a *seller* of some of his/her shares, and prices *below* P_{m0} will cause the investor to be a *buyer* of more shares to add to his current holdings. We can use this insight to form buy and sell curves. To visualize this, direct your attention to Exhibit 1.9 and shift the price axis to the right, to where the vertical line is placed, so that part of the demand curve is in the negative quadrant. Flip that part of the demand curve over from the negative quadrant (where it is viewed as a negative buy) to the positive quadrant (where, without the minus sign, it is viewed as a positive sell). Relabel the upward sloping curve “Sell” and relabel the downward sloping portion “Buy,” as we have done in Exhibit 1.9, and relabel the horizontal axis Q which we do, not in Exhibit 1.9, but in Exhibit 1.10 From the demand curve, we now have a positively inclined sell curve and a negatively inclined buy curve, with the two curves branching off the vertical line at the price where the vertical line at N_{m0} intersects the investor’s demand curve.

Buy and sell curves of two investors are shown in Exhibit 1.10. This exhibit is related to Exhibit 1.9 except that the horizontal axis is labeled “Q” (for the number of shares bought or sold), not “N” (for the number of shares held). The buy/sell curves for the second market participant (which are labeled B2 and S2, respectively)

Exhibit 1.9 From the demand curve to buy and sell curves





Exhibit 1.10 Buy and sell curves of two participants




are higher than the curves for the first participant (which are labeled B1 and S1, respectively). Why do the buy/sell curves for the second participant have a higher price intercept than the buy/sell curves for the first participant? Either because initial share holdings are less for the second participant or because the second participant has a higher expectation of the end of period price, $E(P_{mT})$. Either way, his/her demand curve has a higher intercept on the price axis. Because the buy/sell curves are lower for the first participant than for the second, the upward sloping sell curve for the first participant intersects the downward sloping buy curve for the second participant. Accordingly, the first participant will sell shares at a price at which the second will buy, and thus, a trade can be made.

Now, let there be many buy/sell curves, all with different price intercepts. Aggregating these gives us the downward sloping market buy curve and the upward sloping market sell curve that are shown in Exhibit 1.11 (for simplicity, both curves are presented as linear). The upward sloping sell curve plays the role of a supply curve, but it is not what one might consider a traditional supply curve. In our setting, there is no traditional supply curve. There is a sell curve that can be matched with a buy curve, and this is all we need to obtain the equilibrium price of shares.

The equilibrium price, as shown in Exhibit 1.11, is set where the market buy and market sell curves intersect. Each participant buys or sells shares at this price and, in so doing, achieves an optimal cash-and-shares portfolio. Following the purchases and sales, each investor is holding the exact number of shares desired given the equilibrium price. At this point, there will be no follow-up trades. That is, there will be no desire on anyone's part to re-contract. The market is in equilibrium. This speaks to the efficiency of the perfectly efficient, perfectly liquid market. Let us label the frictionless market equilibrium price, P^* , as shown in Exhibit 1.11, and the equilibrium number of shares traded is labeled Q^* .

Exhibit 1.11 Market buy and sell curves



1.8 The Non-frictionless Market

What if the market is not frictionless? Will all trades still be made at the equilibrium price, P^* ? To answer this question, we introduce friction into the analysis.

There is one simple way to do this: do not allow participants to submit their full, continuous buy and sell curves. This constraint certainly is realistic. What would the response be if you were to contact your broker and, in placing an order, state the intercept and slope parameters of your buy/sell curves? The broker would certainly be confused and probably would think that you are a bit crazy! Of course, real-world investors like us do not do this. We simply state an order that stipulates a price and the number of shares to be bought or sold. This constraint is not at all equivalent to the imposition of a minimum or a maximum price limit. It is simply that the investor submits just one price and one number of shares to buy or to sell. Assume the selection is made optimally, given the participant's demand curve to hold shares and his/her expectation of what the clearing price will be.

While continuing to assume that the investor knows his/her complete buy and sell curves, we have introduced one basic reality of a non-frictionless market. Further, let us introduce a bit of market structure. This is something that is not germane to the frictionless world, but it is important when friction is introduced. Consider the order placement decisions of investors as they approach the 9:30 opening of an equity market such as the New York Stock Exchange or Nasdaq.

Before the market opens, no participant can know the equilibrium price of a stock, P^* , but each has an expectation of what it might be. Each participant's strategic order placement decision is based on two determinants: (1) his/her expectation concerning P^* and (2) the slope and intercept parameters of his/her own buy and sell curves.

Question: If all participants price and size their orders according to the above two determinants, will the market's opening price be P^* ?

Answer: Only by dumb luck! It would occur as a very special case that depends on the accuracy of investor expectations of P^* and on how their buy/sell curves are distributed around P^* .¹³

Conclusion: Our frictionless market gives a wonderfully efficient security market solution that one would expect from a microeconomic analysis of a perfectly competitive, perfectly liquid market. But the finding does not hold in a real-world financial market once friction is introduced.

¹³Ho, Schwartz, and Whitcomb (1985) provide further discussion and analysis.

1.9 Wrap Up: Microeconomics in a Non-frictionless Financial Market

This chapter is focused on the parallels between a household's selection of an optimal consumption bundle (product X and product Y) and an investor's selection of an optimal combination of the riskless asset (cash) and a risky asset (an equity portfolio). In our discussion, we first assume away all transaction costs and operate in a frictionless environment. Then, at the end of the chapter, we introduce friction by constraining the investor to submitting, not the complete, continuous buy and sell curves we used to identify an equilibrium price, but just one single-priced order that is optimally selected given the individual's complete buy and sell curves and his/her expectation of what the equilibrium price might be. Given this one real-world constraint, we have shown that market prices will generally deviate from perfectly competitive equilibrium values even in a market that comprises an unbounded number of order placers, each of whom is a price taker.

In the introduction to this chapter, we noted that an unobservable, frictionless market equilibrium price can exert its force on price determination in real-world markets even if realized transaction prices differ from equilibrium values. Think back to our analogy at the beginning of the chapter where we contrast the force of the deep-water Gulf Stream with the winds, waves, and storms that perturb the surface of the sea.

Like the winds, waves, and storms, a broad array of market frictions perturbs real-world financial market operations and price discovery. Consider the enormous difficulty of dealing with complex information sets that are huge, often imprecise, incomplete, inaccurate, and not equally shared by all. Recognize that investors can have divergent expectations based on publicly available information. Take a close look at trading costs such as bid-ask spreads, opportunity costs, and market impact costs. Understand why short-period (e.g., intraday) price volatility is sharply accentuated. Note the importance of the rules and regulations that define a market's structure. Appreciate the need for liquidity and the importance of achieving high-quality price discovery. We deal with all of this (and more) in the next chapter which turns to our finance course application.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

