Cambridge University Press
978-1-107-02412-0 - Principles of Financial Economics: Second Edition
Stephen F. Leroy and Jan Werner
Excerpt
More information

Part One

Equilibrium and Arbitrage

1

Equilibrium in Security Markets

1.1 Introduction

The analytical framework in the classical finance models discussed in this book is largely the same as in general equilibrium theory: agents, acting as price-takers, exchange claims on consumption to maximize their respective utilities. Because the focus in financial economics is somewhat different from that in mainstream economics, we will ask for greater generality in some directions while sacrificing generality in favor of simplification in other directions.

As an example of greater generality, it is assumed that uncertainty will always be explicitly incorporated in the analysis. We do not assert that there is any special merit in doing so; the point is simply that the area of economics that deals with the same concerns as finance but concentrates on production rather than uncertainty has a different name (capital theory). Another example is that markets are incomplete: the Arrow–Debreu assumption of complete markets is an important special case, but in general it will not be assumed that agents can purchase any imaginable payoff pattern on securities markets.

As an example of simplification, it is assumed that only one good is consumed and that there is no production. Again, the specialization to a single-good exchange economy is adopted only to focus attention on the concerns that are distinctive to finance rather than microeconomics, in which it is assumed that there are many goods (some produced), or capital theory, in which production economies are analyzed in an intertemporal setting.

In addition to those simplifications motivated by the distinctive concerns of finance, classical finance shares many of the same restrictions as Walrasian equilibrium analysis: agents treat the market structure as given, implying that no one tries to create new trading opportunities, and the abstract Walrasian auctioneer must be introduced to establish prices. Markets are competitive and free of transaction costs (except possibly costs implied by trading restrictions, as analyzed in

Chapter 6), and they clear instantaneously. Finally, it is assumed that all agents have the same information. This last assumption largely defines the term "classical"; most of the best work now being done in finance assumes asymmetric information and therefore lies outside the framework of this book.

Even students whose primary interest is in the economics of asymmetric information are well advised to devote some effort to understanding how financial markets work under symmetric information before passing to the much more difficult general case.

1.2 Security Markets

Securities are traded at date 0, and their payoffs are realized at date 1. Date 0, the present, is certain, whereas any of S states can occur at date 1, representing the uncertain future.

Security j is identified by its payoff x_j, an element of $\mathcal{R}^S$, where x_{js} denotes the payoff that the holder of one share of security j receives in state s at date 1. Payoffs are in units of the consumption good. They may be positive, zero, or negative. There exists a finite number J of securities with payoffs $x_1, \ldots, x_J$, $x_j \in \mathcal{R}^S$, taken as given.

The $J \times S$ matrix X of payoffs of all securities

$$X = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_J \end{bmatrix} \tag{1.1}$$

is the *payoff matrix*. Here vectors x_j are understood to be row vectors. In general, vectors are understood to be either row vectors or column vectors, as the context requires.

A *portfolio* is composed of holdings of the J securities. These holdings may be positive, zero, or negative. A positive holding of a security means a long position in that security, whereas a negative holding means a short position (short sale). Thus, short sales are allowed (except in Chapters 6 and 7).

A portfolio is denoted by a J-dimensional vector h, where h_j denotes the holding of security j. The *portfolio payoff* is $\sum_j h_j x_j$ and can be represented as hX.

The set of payoffs available via trades in security markets is the *asset span* and is denoted by $\mathcal{M}$:

$$\mathcal{M} = \{z \in \mathcal{R}^S : z = hX \text{ for some } h \in \mathcal{R}^J\}. \tag{1.2}$$

Thus $\mathcal{M}$ is the subspace of $\mathcal{R}^S$ spanned by the security payoffs, that is, the row span of the payoff matrix X. If $\mathcal{M} = \mathcal{R}^S$, then markets are *complete*. If $\mathcal{M}$ is a proper subspace of $\mathcal{R}^S$, then markets are *incomplete*. When markets are complete, any date-1 consumption plan (that is, any element of $\mathcal{R}^S$) can be obtained as a portfolio payoff but perhaps not uniquely.

Theorem 1.2.1 *Markets are complete iff the payoff matrix X has rank S.*[1]

Proof: Asset span $\mathcal{M}$ equals the whole space $\mathcal{R}^S$ iff the equation $z = hX$, with J unknowns h_j, has a solution for every $z \in \mathcal{R}^S$. A necessary and sufficient condition for this is that X has rank S. □

A security is *redundant* if its payoff can be generated as the payoff of a portfolio of other securities. There are no redundant securities iff the payoff matrix X has rank J.

The prices of securities at date 0 are denoted by a J-dimensional vector $p = (p_1, \ldots, p_J)$. The price of portfolio h at security prices p is $ph = \sum_j p_j h_j$.

The *return* r_j on security j is its payoff x_j divided by its price p_j (assumed to be nonzero; the return on a payoff with zero price is undefined):

$$r_j = \frac{x_j}{p_j}. \tag{1.3}$$

Thus, "return" means gross return ("net return" equals gross return minus one). Throughout we will be working with gross returns.

Frequently the practice in the finance literature is to specify the asset span using the returns on the securities rather than their payoffs. The asset span is the subspace of $\mathcal{R}^S$ spanned by the returns on the securities.

The following example illustrates the concepts introduced earlier.

Example 1.2.1 Let there be three states and two securities. Security 1 is risk free and has payoff $x_1 = (1, 1, 1)$. Security 2 is risky with $x_2 = (1, 2, 2)$. The payoff matrix is

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 2 \end{bmatrix}.$$

The asset span is $\mathcal{M} = \{(z_1, z_2, z_3) : z_1 = h_1 + h_2, z_2 = h_1 + 2h_2, z_3 = h_1 + 2h_2, \text{ for some } (h_1, h_2)\}$ – a two-dimensional subspace of R^3. By inspection, $\mathcal{M} = \{(z_1, z_2, z_3) : z_2 = z_3\}$. At prices $p_1 = 0.8$ and $p_2 = 1.25$, security returns are $r_1 = (1.25, 1.25, 1.25)$ and $r_2 = (0.8, 1.6, 1.6)$. □

[1] Here and throughout this book, "A iff B," an abbreviation for "A if and only if B," has the same meaning as "A is equivalent to B" and as "A is a necessary and sufficient condition for B." Therefore, proving necessity in "A iff B" means proving "B implies A," whereas proving sufficiency means proving "A implies B."

1.3 Agents

In the most general case (pending discussion of the multidate model), agents consume at both dates 0 and 1. Consumption at date 0 is represented by the scalar c_0, whereas consumption at date 1 is represented by the S-dimensional vector $c_1 = (c_{11}, \ldots, c_{1S})$, where c_{1s} represents consumption conditional on state s. Consumption c_{1s} will be denoted by c_s when no confusion can result.

At times we will restrict the set of admissible consumption plans. The most common restriction will be that c_0 and c_1 be positive.[2] However, when using particular utility functions it is generally necessary to impose restrictions other than, or in addition to, positivity. For example, the logarithmic utility function presumes that consumption is strictly positive, whereas the quadratic utility function $u(c) = -\sum_{s=1}^{S}(c_s - \alpha)^2$ has acceptable properties only when $c_s \leq \alpha$. However, under the quadratic utility function, unlike the logarithmic function, zero or negative consumption poses no difficulties.

There is a finite number I of agents. Agent i's preferences are indicated by a continuous utility function $u^i : \mathcal{R}_+^{S+1} \to \mathcal{R}$, in the case in which admissible consumptions are restricted to be positive and $u^i(c_0, c_1)$ is the utility of consumption plan (c_0, c_1). Agent i's endowment is w_0^i at date 0 and w_1^i at date 1.

A *securities market economy* is an economy in which all agents' endowments lie in the asset span. In that case one can think of agents as endowed with initial portfolios of securities (see Section 1.7).

Utility function u is *increasing at date* 0 if $u(c_0', c_1) \geq u(c_0, c_1)$ whenever $c_0' \geq c_0$ for every c_1; it is *increasing at date* 1 if $u(c_0, c_1') \geq u(c_0, c_1)$ whenever $c_1' \geq c_1$ for every c_0. It is *strictly increasing at date* 0 if $u(c_0', c_1) > u(c_0, c_1)$ whenever $c_0' > c_0$ for every c_1 and *strictly increasing at date* 1 if $u(c_0, c_1') > u(c_0, c_1)$ whenever $c_1' > c_1$ for every c_0. If u is (strictly) increasing at date 0 and at date 1, then u is (strictly) increasing.

Utility functions and endowments typically differ across agents; nevertheless, the superscript i will frequently be deleted when no confusion can result.

1.4 Consumption and Portfolio Choice

At date 0 agents consume their date-0 endowments less the value of their security purchases. At date 1 they consume their date-1 endowments plus their security

[2] Our convention on inequalities is as follows: for two vectors $x, y \in \mathcal{R}^n$,

$$x \geq y \text{ means that } x_i \geq y_i \ \forall i; \quad x \text{ is greater than } y,$$
$$x > y \text{ means that } x \geq y \text{ and } x \neq y; \quad x \text{ is greater than but not equal to } y,$$
$$x \gg y \text{ means that } x_i > y_i \ \forall i; \quad x \text{ is strictly greater than } y.$$

For a vector x, *positive* means $x \geq 0$, *positive and nonzero* means $x > 0$, and *strictly positive* means $x \gg 0$. These definitions apply to scalars as well. For scalars, "positive and nonzero" is equivalent to "strictly positive."

payoffs. The agent's consumption-portfolio choice problem is

$$\max_{c_0, c_1, h} u(c_0, c_1), \tag{1.4}$$

subject to

$$c_0 \leq w_0 - ph \tag{1.5}$$

$$c_1 \leq w_1 + hX, \tag{1.6}$$

and a restriction that consumption be positive, $c_0 \geq 0,\ c_1 \geq 0$, if that restriction is imposed.

When, as in Chapters 11 and 13, we want to analyze an agent's optimal portfolio abstracting from the effects of intertemporal consumption choice, we will consider a simplified model in which date-0 consumption does not enter the utility function. The agent's choice problem is then

$$\max_{c_1, h} u(c_1), \tag{1.7}$$

subject to

$$ph \leq w_0 \tag{1.8}$$

and

$$c_1 \leq w_1 + hX. \tag{1.9}$$

1.5 First-Order Conditions

If the utility function u is differentiable, the first-order conditions for a solution to the consumption-portfolio choice problem (1.4)–(1.6) (if the constraint $c_0 \geq 0$, $c_1 \geq 0$ is imposed) are

$$\partial_0 u(c_0, c_1) - \lambda \leq 0, \qquad [\partial_0 u(c_0, c_1) - \lambda]c_0 = 0 \tag{1.10}$$

$$\partial_s u(c_0, c_1) - \mu_s \leq 0, \qquad [\partial_s u(c_0, c_1) - \mu_s]c_s = 0\,, \qquad \forall s \tag{1.11}$$

$$\lambda p = X\mu, \tag{1.12}$$

where λ and $\mu = (\mu_1, \ldots, \mu_S)$ are positive Lagrange multipliers.[3]

[3] If f is a function of a single variable, its first derivative is indicated $f'(x)$ or, when no confusion can result, f'. Similarly, the second derivative is indicated $f''(x)$ or f''. The partial derivative of a function f of two variables x and y with respect to the first variable is indicated $partial_x f(x, y)$ or $\partial_x f$.

Frequently the function in question is a utility function u, and the argument is (c_0, c_1), where, as noted earlier, c_0 is a scalar and c_1 is an S-vector. In that case the partial derivative of the function u with respect to c_0 is denoted $\partial_0 u(c_0, c_1)$ or $\partial_0 u$, and the partial derivative with respect to c_s is denoted $\partial_s u(c_0, c_1)$ or $\partial_s u$. The vector of S partial derivatives with respect to c_s, for all s is denoted $\partial_1 u(c_0, c_1)$ or $\partial_1 u$.

Note that there exists the possibility of confusion: the subscript "1" can indicate either the vector of date-1 partial derivatives or the (scalar) partial derivative with respect to consumption in state 1. The context will always make the intended meaning clear.

If u is quasi-concave, then these conditions are sufficient as well as necessary. If it is assumed that the solution is interior and that $\partial_0 u > 0$, inequalities (1.10) and (1.11) are satisfied with equality. Then Eq. (1.12) becomes

$$p = X\frac{\partial_1 u}{\partial_0 u} \tag{1.13}$$

with typical equation

$$p_j = \sum_s x_{js}\frac{\partial_s u}{\partial_0 u}, \tag{1.14}$$

where we now – and henceforth – delete the argument of u in the first-order conditions. Equation (1.14) says that the price of security j (which is the cost in units of date-0 consumption of a unit increase in the holding of the jth security) is equal to the sum over states of its payoff in each state multiplied by the marginal rate of substitution between consumption in that state and consumption at date 0.

The first-order conditions for problem (1.7) with no consumption at date 0 are as follows:

$$\partial_s u - \mu_s \leq 0, \qquad (\partial_s u - \mu_s)c_s = 0, \quad \forall s \tag{1.15}$$

$$\lambda p = X\mu. \tag{1.16}$$

At an interior solution, Eq. (1.16) becomes

$$\lambda p = X\partial_1 u \tag{1.17}$$

with typical element

$$\lambda p_j = \sum_s x_{js}\partial_s u. \tag{1.18}$$

Because security prices are denominated in units of an abstract numeraire, all we can say about marginal-utility-weighted payoffs is that their sums over states are proportional to security prices.

1.6 Left and Right Inverses of the Payoff Matrix

The payoff matrix X has an inverse iff it is a square matrix ($J = S$) and is of full rank. Neither of these properties is assumed to be true in general. However, even if X is not square, it may have a *left inverse*, defined as a matrix L that satisfies $LX = I_S$, where I_S is the $S \times S$ identity matrix. A left inverse exists iff X is of rank S, which occurs if $J \geq S$ and the columns of X are linearly independent. A left inverse, if it exists, is unique iff there are no redundant securities. Iff a left inverse of X exists, the asset span $\mathcal{M}$ coincides with the date-1 consumption space $\mathcal{R}^S$, and thus markets are complete.

If markets are complete, the vectors of marginal rates of substitution of all agents (whose optimal consumption is interior) are the same and can be inferred uniquely from security prices. To see this, premultiply Eq. (1.13) by a left inverse L to obtain

$$Lp = \frac{\partial_1 u}{\partial_0 u}. \tag{1.19}$$

If markets are incomplete, the vectors of marginal rates of substitution may differ across agents.

Similarly, X may have a *right inverse*, which is defined as a matrix R that satisfies $XR = I_J$. The right inverse exists if X is of rank J, which occurs if $J \leq S$ and the rows of X are linearly independent. Then no security is redundant. Right inverses are unique iff markets are complete. Any date-1 consumption plan c_1 such that $c_1 - w_1$ belongs to the asset span is associated with a unique portfolio

$$h = (c_1 - w_1)R, \tag{1.20}$$

which is derived by postmultiplying Eq. (1.6) by R.

L and R, as defined by

$$L = (X'X)^{-1}X' \tag{1.21}$$

$$R = X'(XX')^{-1}, \tag{1.22}$$

where the prime indicates transposition, are a left inverse and a right inverse, respectively, of X. As these expressions make clear, L exists iff $X'X$ is invertible, whereas R exists iff XX' is invertible.

The payoff matrix X is invertible iff both left and right inverses exist. In that case both L and R are unique. Under the assumptions thus far, none of the following four possibilities is ruled out:

1. Both left and right inverses exist. In that case both are unique.
2. The left inverse exists, but the right inverse does not exist. In that case the left inverse is not unique.
3. The right inverse exists, but the left inverse does not exist. In that case the right inverse is not unique.
4. Neither directional inverse exists.

1.7 General Equilibrium

An *equilibrium* in security markets consists of a vector of security prices p, a portfolio allocation $\{h^i\}$, and a consumption allocation $\{(c_0^i, c_1^i)\}$ such that (1) portfolio h^i and consumption plan (c_0^i, c_1^i) are a solution to agent i's choice problem (1.4) at prices p, and (2) markets clear; that is

$$\sum_i h^i = 0, \tag{1.23}$$

and

$$\sum_i c_0^i \leq \bar{w}_0 \equiv \sum_i w_0^i, \qquad \sum_i c_1^i \leq \bar{w}_1 \equiv \sum_i w_1^i. \tag{1.24}$$

The portfolio market-clearing condition (1.23) implies, by summing agents' budget constraints, the consumption market-clearing condition (1.24). If agents' utility functions are strictly increasing so that all budget constraints hold with equality, and if there are no redundant securities (X has a right inverse), then the converse is also true. If, in contrast, there are redundant securities, then there are many portfolio allocations associated with a market-clearing consumption allocation. At least one of these portfolio allocations is market clearing.

In the simplified model in which date-0 consumption does not enter utility functions, each agent's equilibrium portfolio and date-1 consumption plan are a solution to the choice problem (1.7). Agents' endowments at date 0 are equal to zero, and thus there is zero demand and zero supply of date-0 consumption. Security prices are denominated in units of an abstract numeraire and are determined up to a strictly positive scale factor.

Example 1.7.1 There are two states at date 1 and two agents who consume only at date 1 and have the same utility function

$$u\left(c_1^i, c_2^i\right) = \frac{1}{2}\ln\left(c_1^i\right) + \frac{1}{2}\ln\left(c_2^i\right), \tag{1.25}$$

for $i = 1, 2$. Their date-0 endowments are zero. Date-1 endowments are $w_1^1 = (3, 0)$ and $w_1^2 = (0, 3)$. There are two securities with payoffs

$$x_1 = (1, 1) \qquad \text{and } x_2 = (1, 0). \tag{1.26}$$

We do not present a complete derivation of an equilibrium. Instead, we start from a conjecture that the consumption allocation consisting of risk-free date-1 consumption plans $c_1^1 = c_1^2 = (3/2, 3/2)$ is an equilibrium allocation. We need to verify that, indeed, this is an equilibrium allocation and then find equilibrium security prices and portfolios.

Agent 1's marginal utilities of consumption at $c^1 = (3/2, 3/2)$ are $1/3$ for both states 1 and 2. They are the same for agent 2 at $c^1 = (3/2, 3/2)$. We can check that the first-order conditions (1.18) hold for both agents with security prices set