1
Introduction

1.1 What finance theory is about

How much should you save? and *How much risk should you bear?* When we think about these questions, it becomes clear pretty quickly that they are of great importance to our overall material well-being. Saving is essential because most of us will retire at some point. From that point onwards, although we will still be consuming, we will receive no more labor income. Moreover, we will all face significant economic risks during our lives—the risk of losing our job, for instance, or—much worse—of becoming unable to work because of illness or other misfortunes. Clearly, the risks we are exposed to can have a huge effect on our future life, and it is therefore essential to make rational decisions about how much risk to bear.

Important as these questions are for each one of us, individuals' decisions about saving and risk-taking also matter for society as a whole. Total saving determines the amount of investment that the economy as a whole can realize and thus affects future production possibilities. The amount of risk that people are willing to bear determines whether risky projects will be undertaken. Individual decisions in the face of future retirement and risk and the capital requirements of more or less risky investment projects are coordinated through financial markets. If markets work well, risk is allocated to those people who are least hurt by it, impatient people get to consume before they earn (by taking out a loan), and capital is allocated to those projects that generate the most attractive risk-return profile. Finance is concerned with the determination of those prices that equalize demand and supply on these markets and with their effect on the allocation of capital and risk across agents in the economy.

Finance theory is also useful in interpreting financial market prices in ways that are of interest for public policy and social welfare issues. Robert Lucas

(1987), for instance, has examined the social costs of business cycles. This is obviously important for economic policy making, but it is also important for macroeconomic theory. To learn the answer to this question, we need to know how much people dislike risk, that is, variations in income. More specifically, in order to judge how expensive business cycles are, we need to determine a price that people would be prepared to pay to avoid the income variations caused by business cycles. Modern asset pricing theory allows us—at least in principle—to do just that.

1.2 Some history of thought

General equilibrium theory, macroeconomics, and asset pricing theory are three fields in economics that have converged more and more over the last thirty years or so. In this section we consider how this convergence came about.

1.2.1 General equilibrium theory

General equilibrium theory is an approach to describing the behavior of an economy as a whole by working out the optimal behavior of each member of the set of *agents* that make up the economy, and looking for a point of mutual compatibility or consistency. The theory assumes that individuals do not interact with each other directly. Interaction occurs only indirectly, through *anonymous markets* on which prices (exchange rates for different commodities) are posted. A second assumption that goes hand in hand with anonymity is that each individual is small in relation to the market, so that everyone neglects his own influence on market prices. This assumption is called *perfect competition*. Models that make these two assumptions are called *Walrasian*, in honor of Léon Walras (1874) who was the first to formulate such a model. We say that the economy is in *equilibrium* if, at a certain price, each individual buys or sells the optimal quantities (given his tastes and possibilities) of all commodities and the total supply of each commodity equals the total demand for it.[1]

Modern general equilibrium theory in the tradition of Arrow & Debreu (1954) accommodates a large number of different goods and very diverse preferences of individuals. This research has established conditions that guarantee the existence of an equilibrium. It has also developed properties of equilibrium allocations, such as the welfare theorems. The two welfare

[1]Note that it is left unspecified in this model who posts the prices, since everyone takes them as given.

theorems demonstrate that market equilibrium allocations and socially efficient allocations are equivalent, under some conditions. This equivalence will be extremely useful for our purposes.

This theory was significantly advanced by Hirshleifer (1965, 1966) and Radner (1972). These authors built financial markets into the model and thus provided the first crucial ingredient for making general equilibrium theory applicable to finance. Moreover, their work opened up the possibility of analyzing financial markets that are incomplete in the sense that the available financial instruments may not be sufficient to trade all individual risks efficiently. This incompleteness opens the door to various sorts of coordination failures in a market economy.

By the late 1950s, general equilibrium theory had become the cornerstone of microeconomics, and remained so until it was slowly pushed aside by advances in game theory and information economics in the 1970s. General equilibrium theory has, however, received a new lease of life through its applications to the theory of macroeconomic fluctuations and the theory of asset pricing.

1.2.2 Macroeconomics

When John Maynard Keynes (1936) developed his *General Theory*, the world was in disarray. Mass unemployment and mass bankruptcy had erupted—first in the U.S.A. but then quickly spreading throughout the capitalist world. Possibly because of these events, Keynes chose a style of model that broke with the tradition of classical economics. His model did not feature individual agents explicitly, nor did it feature dynamics of any sort. Instead, he focused on the interdependence between different aggregate variables. In that sense, Keynes's model is a general equilibrium model, yet one in which the aggregate demand and supply functions are not developed from an individual optimization perspective. This became most clear in Hicks's (1937) version of Keynes's model, which came to dominate macroeconomic thinking.

This lack of microfoundation led to problems associated with the endogenous determination of expectations. Clearly, expectations should affect an individual's decisions. We would expect that rational decision makers will try to collect information if faulty decisions are costly. Hicks's version of Keynes's model really lacked a convincing theory of expectations. This omission led to increasing dissatisfaction with the model on purely theoretical grounds and ultimately to a dramatic empirical failure with the stagflation of the 1970s, which was an impossible event in the Keynes–Hicks orthodoxy.

These developments gave impetus to a new, or rather renewed old, approach,[2] namely to construct dynamic models of aggregate economic fluctuations based on individual decisions together with shocks of some sort (most prominently to technology). The *rational expectations revolution* in macroeconomics is nothing but a simplified version of Radner's (1972) idea of an "equilibrium of plans, prices, and price expectations." The early macro versions of this idea were simplified, in the sense that agents were assumed to have an unbiased expectation of the mean of stochastic variables only, whereas in Radner's model agents have correct state-contingent expectations.

In essence, this is what the *New Classical* and later *Real Business Cycle* theory consist of: computable dynamic stochastic general equilibrium models. Compared with traditional general equilibrium theory, the macroeconomic variants are typically simpler because they feature only one good and one agent, and give scant attention to the conditions for aggregation. They also make much stronger assumptions concerning preferences and technology in order to get easily computable equilibria. Modern *New Keynesian* theory and the *New Neoclassical Synthesis* models deviate from the Walrasian orthodoxy by introducing various frictions into the model. But they, too, work within the general equilibrium framework and assume representative goods and agents.[3,4]

1.2.3 Finance

Finance theory started out as a field of business administration. Sensible decision making about how to finance operations is obviously vital for any firm, and the placement of free reserves into financial assets can have a substantial impact on the profitability of the enterprise. Markowitz's (1952) *mean–variance mechanics* was a breakthrough, offering a much more sophisticated decision rule than was common at the time, but one that was still simple to apply.

Markowitz's contribution serves as a tool for decision making; accordingly, his research is silent about the determination of asset prices. Their stochastic properties are taken as given. Subsequently, emphasis shifted away from

[2]Ramsey (1928) is a precursor to real business cycle theory and therefore an early contributor to a theory that only much later became part of "modern macroeconomics."

[3]Despite the fact that frictions render the first welfare theorem inoperative and thus remove the basis for aggregation!

[4]See Woodford (forthcoming) for an excellent survey of the development of macroeconomics.

using properties of asset prices to guide decisions, towards *explaining* asset prices. The *capital asset pricing model* (CAPM) of Sharpe (1964) and others assumes that the economy is populated by Markowitzian mean–variance decision makers. With the help of some additional assumptions, Sharpe concludes that the market portfolio must be mean–variance efficient, and that every agent must hold a mixture of the risk-free asset and the market portfolio (*two-fund separation theorem*). Most significantly, this theory implies that only that part of the risk of an asset that is correlated with the whole market carries a premium in equilibrium.

Today finance is largely concerned with the implications of *no arbitrage conditions* for asset prices. Absence of arbitrage opportunities is a weak form of rationality or equilibrium requirement.[5] An arbitrage portfolio is a portfolio that guarantees positive payoffs but whose price is zero or negative. If such a portfolio exists, it is possible to generate infinite payoffs without taking any risk. An absence of arbitrage opportunities means that asset prices must be such that no arbitrage portfolio exists. This is certainly reasonable, but the no arbitrage assumption alone does not allow one to incorporate all the economic fundamentals of preferences and endowments which arguably drive the decisions about intertemporal allocation and risk exposure; nor does it fully exploit the implications of market equilibrium. Equilibrium requires more: namely that the total supply of each asset equates the total demand for it. Relating asset prices to the extensive data of the economy in this way not only makes for a more complete (and hopefully more precise) theory, but also allows for interpretations of these prices that are beyond the range of possibilities when using just the assumption of arbitrage conditions alone.

1.2.4 Macrofinance: a unified general–equilibrium–asset–pricing–business–cycle–theory

Traditional general equilibrium theory as well as macroeconomics focuses on the description of properties of equilibrium allocations. The scope is different, with macroeconomics concentrating on the time series (dynamical) aspects of aggregate measures of economic activity and traditional general equilibrium theory targeting questions of existence and efficiency of equilibria. But general equilibrium theory is also able to make statements about equilibrium prices. General equilibrium theory that focuses on the explanation of prices of financial assets may be called *equilibrium asset pricing theory*.

[5] Bachelier (1900), Fisher (1907), and Bronzin (1908) were very early contributors to this theory. I thank Heinz Zimmermann for pointing out Bronzin's book to me.

Equilibrium asset pricing theory in this generality (heterogeneous goods and agents and general preferences) has not enough structure to yield interesting results, though, and it has proved very fruitful to impose more structure on preferences.

It was (one of) Stiglitz's (1970) contribution(s) to connect finance more closely with economic theory. By explaining the demand for financial assets with a utility maximization problem whose ultimate goal is the optimal choice of consumption, he paved the way to Lucas's (1978) *tree model* and Breeden's (1979) *consumption capital asset pricing model* (CCAPM). Here we will call this model the *finance economy*, because it is a general equilibrium model that is simplified and specialized in exactly the way that financial economists have found useful. This model is the result of the combination of Arrow–Debreu–Radner general equilibrium theory and von Neumann–Morgenstern expected utility theory. The first welfare theorem of general equilibrium theory allows us to transform the general version of the model into a much simpler one-good, one-agent economy. We will discuss in detail how this aggregation can be performed. Expected utility theory gives much more structure to the behavior of people with respect to risk taking, and thus allows for a theory with more concrete predictions about equilibrium asset prices.

What exactly does the finance economy look like? First of all, there is only one agent, hence the equilibrium allocation is trivial: the single agent eats the output. Moreover, we are dealing with an exchange economy; that is to say, there is no production,[6] so the endowment is exogenous (but still stochastic). Thus, the equilibrium allocation is also exogenous: the single agent eats his own endowment.

What's the point of this? Why should a model like this be of interest? Since quantities are not the focus of the analysis, it is natural to choose a model in which the equilibrium allocation is given beforehand. Thus, moving from the traditional general equilibrium model to the finance economy simplifies the model in exactly the way we want: we lose information on items that we do not aim to explain, such as equilibrium trades, distribution, and allocations, but we keep all the information on equilibrium prices. Such a model is perfectly geared for investigating how changes of stochastic properties of endowments affect equilibrium prices of different kinds of securities.

The field that studies these relationships could be called *macrofinance*, because the objective is to explain financial market data with aggregate or macroeconomic shocks. By building on a complete (if simplified) general

[6] In fact, this is an exchange economy in which there is not even any exchange going on, because there is only one agent and the poor chap has no one to trade with.

equilibrium model, macrofinance provides microeconomic foundations for more conventional theories of finance. This microeconomic foundation helps us to gain deeper insights, because it relates asset price data to individual preferences over risk and time and aggregate consumption fluctuations, and thus allows us to interpret asset prices in terms of structural data of the economy.

1.3 The importance of the puzzles

This general model can be applied to a quite diverse set of objects. For instance, we can make predictions about the return rates of bonds of different maturities, thereby generating an equilibrium model of the term structure of interest rates. The model also makes predictions about the equilibrium return rate of risk-free bonds and of risky shares. The difference between the two is called the equity premium. Unfortunately, this application of the model fails miserably: it predicts a much higher risk-free return rate than what we observe in the data, and it predicts a very low (almost zero) equity premium, which is not at all what we typically observe in the data.

It is puzzling why anyone would invest in bonds rather than shares, given the large premium that equities offer. Standard theory could justify this if agents were subject to very large risks, or were very averse to being exposed to risk. National income accounting data tell us, however, that aggregate risk is small, and the assumption of very strong risk aversion contradicts experimental evidence. This empirical failure is called the *equity premium puzzle*. Similarly, consider that an impatient person's optimal consumption path is decreasing through time: he would rather consume early than late. The optimal consumption path of an infinitely patient person (having a discount factor of one) would be flat, because he would prefer to consume the same amount every year. But it is a fact that on average income grows every year. Thus, there is a tendency for people to dissave in order to transfer consumption from the future to the present. Of course, not everyone can dissave at the same time, and in equilibrium a high enough interest rate must provide an incentive for agents to postpone consumption. The trouble is that standard assumptions (a discount factor not exceeding one, moderate risk aversion) imply that the equilibrium interest rate should be much larger than the current market rate. This empirical failure has been named the *risk–free rate puzzle*.

The puzzles have initiated an extraordinary research effort. We will review some of this large literature dedicated to resolving the empirical failure of the model. Correcting this failure, as well as the question *how* we correct it, is important in three dimensions. First, since modern macroeconomic

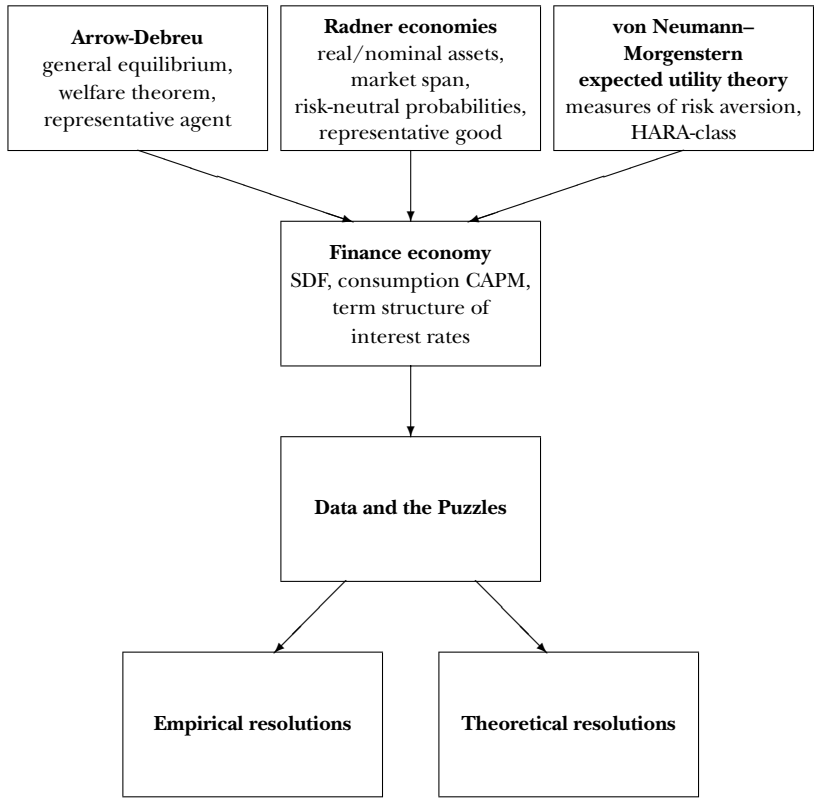

Figure 1.1. *Structure of the book.*

theory is an application of the same model as the one underlying asset pricing theory, the way we resolve the puzzles will ultimately also influence our thinking about the mechanics of growth and business cycles. Second, different approaches taken to resolving the puzzles will result in different predictions about the future return to capital we can expect. Whether this will be 10% or only 3% per year has an enormous effect on pension systems and, more generally, on the appropriate amount of saving that each of us should undertake. And third, the fashion in which the model has to be changed to match the empirical data will also affect our view about the average attitude toward risk and, accordingly, the social cost we attribute to aggregate risk.

1.4 Outline of the book

The book develops general equilibrium asset pricing theory from the bottom up. We start with Arrow–Debreu (chapter 2) and Radner general equilibrium theory (chapter 3). Combining this with von Neumann–Morgenstern utility theory (chapter 4), we derive a model that we call finance economy (chapter 5). This model is the workhorse of macrofinance. We consider several special cases, and also extend the model to cover many periods (chapter 6). Only at this point do we confront the model with the data and identify the asset pricing puzzles (chapter 7). We then explore some of the avenues researchers have taken to bring the theoretical predictions closer to empirical evidence (chapter 8). The last chapter concludes by offering three directions in which the asset pricing puzzles can be interpreted, and speculating on what these interpretations imply for the future development of the theory, and the likely future performance of financial assets. Figure 1.1 shows the structure of the book and relates the different parts to each other.