How Do Finance Specialists Think?*

Ronald W. Anderson†

January 17, 2008

1 Introduction

In this talk I will try to explain what I understand to be the key ideas that run through academic and practitioner work on financial questions and which define the intellectual make-up of finance specialists. Finance cannot claim to be a core discipline of the social sciences. However, its stature has risen with the growing importance of the financial sector and also because of some genuinely new and important ideas that have come from scientific study of finance problems. Finance borrows from sister disciplines. Indeed, finance as it is taught and researched at LSE can be viewed as an off-shoot of economics. But it builds on other fields such as statistics, mathematics, accounting and more recently psychology as well.

2 The problem of finance

The way finance specialists think is of course shaped by the problems they study. What is the subject studied in finance? Historically, it has been developed by economists who have studied the problem of *allocating resources over time*. Here the key questions are:

- How much should be saved and invested in projects that will pay off in the future?
- Among the alternatives available what is the best form for the investments to take?

It was recognized early on that in answering questions of this sort we need to formulate views on what the future is likely to bring.

- If in the future our income will be relatively low because of adverse economic conditions it would be prudent to increase our savings now.

*Talk prepared for the LSE Series, "How Social Scientists Think"

†London School of Economics and CEPR. r.w.anderson@lse.ac.uk.

- The best of the alternative investment projects will be the one that produces the product that will be in highest demand, i.e., will produce the highest return.

However, at the present we do not know for sure whether future economic conditions will be good or poor, and we do not know for sure which project will realize the highest returns. That is, allocation of resources over time inevitably hits up against the uncertainty of future events. The proper formulation of thinking about these problems awaited tools for the study of choice making under uncertainty.

- The development of *expected utility theory* by von Neumann and Morgenstern, Friedman and Savage and others was key.
- The state contingent claims approach made the key insight that goods we consume can be distinguished by both the date when we consume and the *state of nature* prevailing when we consume them.

With the development of these ideas, it became increasingly clear that the scope of finance is to study the problem of *allocating resources over time and over states of nature.* Moving resources over time and states of nature is what is accomplished in financial contracts or securities which *shift risks* among contracting parties.

These questions were given an elegant and complete analysis by Arrow, Debreu and others economists who explored the properties of *general equilibrium* in an abstract economic system. The key results turn on whether or not *markets are complete*, that is, whether exchanges between all states and time can be effected. When markets are complete, a competitive equilibrium will be efficient in the sense of Pareto. Furthermore, *any* efficient allocation can be achieved in a competitive equilibrium through a system of taxes and transfers on endowments.

Now if we consider that there are a large number of possible states of natures the assumption of complete markets would seem to be implausible. For example if we postulate an economy with 1000 distinct types of goods and 1000 states of nature, we would need one million (1000 x 1000) markets. Arrow, showed that securities markets play a particular role in the economy that allow us to overcome this problem of excessive dimensionality. He showed that if we have a complete set of markets for state contingent claims on a numeraire good (i.e., on 1000 markets in the example above) and if agents can predict the relative prices that would prevail in the alternative states of nature, then markets are effectively complete and the optimality properties of complete markets will prevail.

This discussion may leave the impression that the notion of general equilibrium in a system of effectively complete markets is an elegant abstraction that could hardly be of practical relevance in understanding the workings of the real financial sector. In fact, the tools of choice making under uncertainty were given very practical application in finance with the development of modern portfolio theory by Tobin and Markowitz in the early

1950's. Building on this about ten years later, the Capital Asset Pricing Model (CAPM) was developed by Sharpe, Lintner, Mossin and Treyner thus becoming the most prominent early example of general equilibrium analysis to be put to a very practical purpose. The CAPM is an example of a *pricing model* that has proved very useful in finance. It says that the risk premium that prevails for a security, i.e., the return on that security in excess of the risk-free rate of return, should be a linear function of that security's *systematic risk* (reflecting the correlation of that market with an index of returns for the market overall).

3 The emergence of modern finance

The ideas that I have sketched so far are part of the general intellectual bagage of all economists. What is special about finance? What are the key concepts which define a distinctive finance approach to issues? It think that the answer is given in a number of developments between the late 1950's and the late 1970's building on the economics of uncertainty which made us aware of the power of the logic of *arbitrage*. I fix the time period here with reference to the publication of the paper by Modigliani and Miller in 1958 and to the publication in 1979 of the paper by Harrison and Kreps. Of course, arbitrage is an old notion in economics and underlies the venerable *law of one price*. However, starting with Modigliani and Miller arbitrage arguments have been used to strip away complexity to focus on essentials in financial analysis. It was shown that often you can go very far in understanding a problem without postulating all the structure about preferences, production technology and so forth as typically required in a fully formulated economic model.

What is arbitrage? It is the purchase of one collection of goods or securities and the simultaneous sale of another collection of goods and securities which produces a gain in at least one state of nature without incurring a loss in any state of nature.

An arbitrage is a private money machine. As such, people can be relied upon to pursue arbitrage opportunities without limit. Therefore, prices will be forced to adjust to reflect *no arbitrage conditions*. Thus the logic of arbitrage leads us to rules for relative pricing of securities.

The first and still most important example of how this logic shapes finance thinking is the Modigliani-Miller result that in the absence of market frictions the value of the equity, debt and other liabilities of the firm must have a total value equal to the value of the assets of the firm. Merton Miller was fond of explaining this notion by reference to a pizza pie. It does not matter whether you cut the pie into many slices or few, it is still has the same total number of calories. This deceptively simple idea turns out to have considerable power in simplifying arguments and in imposing a higher degree of consistency in financial analysis than was prevalent before their contribution.

For example, what is the relation of return on equity and return on the assets of the firm? For a firm with just equity and debt M-M says the value of the firm's assets equals the sum of it debt value and its equity, or in symbols,

$$A = D + E$$

Now defining return on assets as $ROA = dA/A$, return on equity as $ROE = dE/E$, differentiating our equation and doing a little algebra shows that if the firm issues only risk-free debt we have,

$$ROE - r = \frac{A}{E}(ROA - r)$$

where r is the risk-free rate of return. This gives the useful insight that firm managers may seek higher returns to shareholders by increasing returns on assets, but they may also pursue them by simply increasing the amount leverage and thus exposing the firm to higher risk.

Now, hopefully you will immediately notice that in my last sentence I have been somewhat inconsistent. Why? Because I had premised my statement on the assumption that the firm only issue risk-free debt, but if the firm increases leverage to the point that debt is no longer risk-free then my equation needs to be modified. This is an illustration of how M-M leads us to impose greater clarity in the analysis of financial problems.

The logic of M-M can be extended to any security might wish to consider. Interestingly, this leads us to an insight as to how values of securities can be related to their expected payoffs in the future. Suppose that there is a complete market for securities that payoff one unit of numraire in a given state of nature s and nothing otherwise, and let that security's price be given as π_s. Now consider a security that pays off 1 unit in every state. This security is equivalent to a portfolio consisting of one unit of each state contingent claims. It is also equivalent to investing $1/(1+r)$ today at the risk-free rate r. So its no-arbitrage value is,

$$V_0 = \sum_s \pi_s = \frac{1}{(1+r)}$$

Now let us use the prices π_s and the interest rate r to define a set of variables $\pi_s^* = \pi_s(1+r)$. Notice that these must be positive because otherwise we could buy something for the future and be paid for doing so today which would be a clear arbitrage. Furthermore, by construction they sum to unity $\sum \pi_s^* = 1$. So this set of modified prices defines a probability distribution over future states of the world. This distribution is called the "risk-neutral" probability distribution in finance. The reason for this terminology can be understood by noting that the no-arbitrage value of an arbitrary security paying a variable amount y_s in each state s is,

$$V = \sum_s \pi_s y_s = \sum_s \frac{\pi_s^*}{(1+r)} y_s = \frac{1}{(1+r)} E^* y_s$$

That is, the value of any security can be expressed as the present discounted value of its expected payoff tomorrow where expectations are taken with respect to the *risk neutral probability distribution* and discounting is done at the risk-free rate. We stress that if the prices π_s are the market prices of the elementary state claims then this expression gives the fair value in the market of the security. Even though agents may be averse to risk, securities are priced in the market as though agents were risk-neutral but calculated expectations using the probability distribution π_s^*.

Now the market value of the asset V computed in this way should be distinguished from the actuarially fair value of the asset, V^a which would be based on the *statistical* probability distribution over states. If the latter probabilities are denoted π_s^a then we have,

$$V^a = \frac{1}{(1+r)} \sum_s \pi_s^a y_s = \frac{1}{(1+r)} E y_s$$

Note that in Ey_s expectation is taken with respect to the statistical distribution of states. The difference $V^a - V$ reflects the discount that the market imposes in order the bear the risk involved in holding the payoffs $\{y_s\}$. Another way in finance the market value of the security is expressed is using the *risk adjusted* rate of return on the security r^* defined by, $V = \frac{1}{(1+r^*)} E y_s$.

These basic valuation expressions are relevant to one of the most famous pricing relations in finance, namely, the Black-Scholes formula for the pricing of call options, that is, for securities giving the right but not the obligation to purchase an underlying asset at some date in the future at an exercise price fixed today. Suppose that the value of the underlying asset, e.g., a share of common stock, has a price S_t at time t, the exercise price in the call option is X and the exercise date is T, then the logic of no-arbitrage means that the value of the call at maturity must be, $Max(S_T - X, 0)$. The problem of pricing such an option was studied in the 1900 doctoral thesis of Louis Bachelier and subsequently by others who found various expressions for the actuarially fair value which in our notation can be expressed as,

$$C^a = \frac{1}{(1+r)} E Max(S_T - X, 0)$$

Now this expression was unsatisfactory for two reasons. First, it does not tell us the market value of the call option, i.e., the risk premium. Second the calculation of the expectation requires us to take a view on the future course of the stock market, something which in efficient markets is very hard to do. For example, the famous economist Paul Samuelson tried his hand at this problem under the assumption that the stock price followed a geometric Brownian motion which can be denoted,

$$dS = \mu S dt + \sigma S dz \tag{1}$$

where dS is the change of the stock price over a very small time interval dt and dz is a Brownian motion, i.e., random variable following a normal distribution over dt. The resulting expected value $EMax(S_T - X, 0)$ involves an integral expression where the drift of the stock process μ appears.

Now Black and Scholes studied this problem and building on a comment from Samuelson's student Bob Merton, they found that they could construct a sequence of arbitrage portfolios involving the underlying stock and short term lending such that the portfolio was riskless over short time periods dt. Applying the logic of arbitrage which implies that the value of the portfolio must grow at the risk-free rate r they were able to derive a partial differential equation that must be obeyed by the price of the call option in the absence of arbitrage. They were able to solve that equation and found an expression which can be written in our notation as,

$$C = \frac{1}{(1+r)} E^* Max(S_T - X, 0)$$

where the expectation E^* is taken with respect to the risk-neutralized process,

$$dS = rSdt + \sigma S dz \tag{2}$$

Let us write this solution as $C(S_t, X, r, T, \sigma)$. This is remarkable because it gives us an expression for the *fair market value* of the option and because its calculation does not require us to take a view on the direction of the stock market, i.e., the parameter μ.

The importance of this work went beyond the applications to the buying and selling of options or the equivalent. In fact, it gave us an expression for pricing any security dependent upon a risk following a geometric Brownian motion as in equation (1). If the payoff of the security at maturity satisfies a known function $y(S_T)$ at some future date T, then in the absence of arbitrage its market value today is $V^+ = \frac{1}{(1+r)} E^* y(S_T)$ in which expectations are taken with respect to the risk-neutral process (2). Furthermore, this unique probability distribution for pricing over all the possible realizations S_T can be written as $f(s)$ and can be inferred from the prices of a complete set of options on the stock using the equation,

$$f(s) = \gamma C_X(S_t, s, r, T, \sigma)$$

where C_X is the partial derivative with respect to the second argument, the exercise price X, and γ is a negative constant chosen to assure $\int_{\underline{S}}^{\overline{S}} f(s)ds = 1$ over the support $[\underline{S}, \overline{S}]$.

Again the risk neutral probability distribution $f(s)$ will differ from the statistical distribution denoted $f^a(s)$, say, by an amount that will reflect the price in equilibrium of bearing the risk. Exactly how the two distributions differ is given by an object called the pricing kernal, denoted $k(s)$, and defined implicitly by the valuation relations.

$$V^+ = \frac{1}{(1+r)} \int_{\underline{S}}^{\overline{S}} y(s)f(s)ds = \frac{1}{(1+r)} \int_{\underline{S}}^{\overline{S}} k(s)y(s)f^a(s)ds$$

Further explorations into asset pricing in equilibrium with complete markets has established a formula for the pricing kernal as $k(s) = \frac{U'(C_T(s))}{U'(C_t)}$, where $U'(C_T(s))$ is the marginal utility of consumption at time T in state s. That is, in equilibrium the pricing kernal is given by the marginal rate of substitution between state-time (s, T) and today, t. Finally, it was established in a variety of general settings, not just for geometric Brownian motions, that so long as the system of markets is effectively complete, the risk-neutral density and therefore the pricing kernal are unique.

This theory of arbitrage pricing was essentially complete by the early 1980's. Since then academic finance has been busy pursuing implications of the variety of ways the assumptions of this elegant theory might be violated in the real world. We will hint briefly at some of these directions of development in the next section. However, it is important to stress that Keynes' dictum that the practical business person of today is hostage to the thinking of some defunct academic of the recent past was never truer than in financial markets of today. Indeed, the whole business of secruitisation, structured finance and all the other aspects of what has become known as the "slicing and dicing" of risks are nothing other than elaborate exercises in the application of the complete markets tools we have outlined here.

4 Finance in a world of market imperfections

The theory of complete financial markets we have described strips away some of the details of real world financial operations and arrives at a coherent and rather complete vision of how savers, investors and firms are knit together to channel funds into good productive uses and to shift risks onto agents best placed to bear them. The assumptions made in this analysis, e.g., that agents can buy and sell a security on the market at the same price and that the quantity they sell will not impact price, are not strictly true. But in some markets, for example, the market for foreign exchange transactions among well-established banks the assumptions seem fairly close to the being fulfilled. However, in other cases, market imperfections are hard to ignore, and finance specialists have explored what happens when frictions are built into the models. Unfortunately, the answer to this investigation appears to be "many different things", and the result is not a unitary "theory of imperfect capital markets" comparable in its completeness and coherence to the theory of complete markets.

4.1 Transactions costs

The reason that market imperfections tend to blur predictions can be considered by looking at a basic arbitrage trade such as trilateral currency arbitrage. For example if the US-UK exchange rate is $2/£ and the US-EU exchange rate is $1.5/€ then in frictionless markets the EU-UK exchange rate must be € 1.3333/ £ or else there is a profitable arbitrage. Now if buying and selling currency involves a transactions cost of about 1% as represented by US-UK ask of $2.01/£ and an bid of $1.99/£ and similar spread of $1.49/€ bid and $1.51/€ ask then we will have an implied no-arbitrage range of EU-UK exchange rates given by € 1.31$\leq$ *bid* $\leq$ *ask* $\leq$€ 1.3490 . Notice the that the arbitrage range has widened to more than 2% reflecting the fact that an arbitrage trade would be on the unfavorable side of the trade each step along the way. Thus transactions costs can easily render complicated arbitrages uneconomic.

This idea that a range of prices may be compatible with no-arbitrage carries over to arbitrages in general. Recall that in frictionless markets that are effectively complete there is a *unique* risk-neutral probability distribution that is would not permit arbitrage. Now when markets are incomplete there is a multiplicity of risk-neutral probability distributions that do not permit arbitrage. Equivalently there are a multiplicity of pricing kernals $k_t(s)$ consistent with the statistical distribution $f^a(s)$ of the underlying risk. In this case the valuation of securities cannot be determined unambiguously by arbitrage considerations alone. How do finance specialists resolve this problem? There are several approaches that can be found in the literature. *Financial economists* tend to rely on the idea of *equlibrium pricing*, i.e., that a security's price will reflect a balance of supply and demand. Thus some analysts will determine the pricing kernal $k_t(s) = \frac{U'(C_T(s))}{U'(C_t)}$ by reference to an equilibrium model where the preferences $U(.)$, technology etc are explicitly spelled out. Others will take a more reduced-form approach and posit a convenient form for the pricing kernal as a function of conditioning variables, either observable (e.g., GDP, employment etc.) or unobservable (latent), and determine the estimate the kernal statistically. This contrasts with the approach taken by mathematicians working in finance who have attempted to select among alternative risk-neutral pricing distributions on the basis of additional properties (e.g., minimal entropy) which are thought to be plausible.

4.2 Corporate finance

Taking market imperfections seriously has transformed our thinking about the major areas of financial application: corporate finance and asset pricing. Modern corporate finance looks at various frictions that lead to the violation of the M-M result so that financial policy may have an impact on the value of the firm. The main frictions that were first considered, starting with Miller and Modigliani themselves, were corporate taxation and bankruptcy costs. This led to a first generation of models, loosely known as "trade-off" theories which held that firms would choose leverage and other aspects of financial policy to balance the tax advantages of debt versus potential costs of financial distress. A richer vein of research sprang from the observation that in the modern corporation there is an important asymmetry between the position of corporate *insiders* such as senior managers or controlling share holders and corporate *outsiders* such as small share holders or creditors. This tends to create a wedge between the external cost of capital and the internal cost of capital. In such a world, the way financial operations are organized can have considerable impact on the value of the firm, the efficiency, and indeed on the prosperity of the economy generally.

Explorations of these problems have tended to take either of two distinct directions depending upon the way the asymmetry of insiders and outsiders is formulated. Some emphasize problems of moral hazard (hidden actions); whereas others focus on adverse selection (hidden types). From from these starting points early analyses consider *second best* financial structures and policies with an *exogenously given* array of possible financial instruments available, typically simple debt and equity. These studies consider when one or another of the financial contracts will be preferred and under certain assumptions arrive at clear hierarchies ("pecking-orders") among financing alternatives. Other studies explore the consequences of different organizational forms on the performance of the firm. Examples of these are studies of the organization of boards of directors and the structure of internal capital markets i.e., ways in which internal funds are allocated to different divisions of a multi-product firm. Subsequently, analysts have considered how the nature of the financial contracts themselves are determined endogenously in the interaction of insiders and outsiders. This has led to the study of *security design* where analyst have tended to use the tools of incomplete contracts theory or mechanism design.

The result of all these developments has been a rich body of theoretical corporate finance models which in few if any general conclusions appear robust to changes in the modeling assumptions. Naturally, many analysts have been strongly interested in going to the data to see which theories appear to be supported or at least to establish empirical regularities. This literature is so vast and the results are so varied as to defy any simple summary. However, I close this discussion of corporate finance by noting that recently empirical work has had a discernable impact on corporate finance theory in that it has been noted that relatively simple models taking into account basic frictions such as tax shields and bankruptcy costs do considerably better in accounting for patterns seen in the data if the analysis explicitly takes into account the inter-temporal nature of financing and investment decisions made by firms. Thus there has considerable recent interest in *dynamic trade-off* models of capital structure and financial policy.

4.3 Asset markets

In asset market research, much of the work over the last two decades has emerged from empirical studies of the implications of the *efficient markets hypothesis*. The theory of asset pricing under perfect markets sketched above has direct implications for how should asset values, V_t, should vary over time. To simplify notation assume the risk-free rate is zero, $r = 0$. Then by property of iterated expectations, $E_t V_{t+1} = E_t E_{t+1} y(S_T) = E_t y(S_T) = V_t$. Or $E_t(V_{t+1} - V_t) = 0$.

This says that the market value of an asset follows a martingale and asset prices are unpredictable at first order. This is a testable hypothesis, and financial markets produce in great abundance the security price data that are relevant for testing this. Initial studies of the random behavior of stock prices were generally supportive of this efficient markets hypothesis. However, further analysis uncovered a variety of ways these markets seem to violate the properties of efficient markets. Early examples of such pricing *anomalies* include the January effect, the small firm effect and the profitability of certain technical trading rules. Some of these apparent inefficiencies disappeared with closer scrutiny of the data or once transactions costs were taken into account. Others, such as the profitability of technical trading rules, could be explained by the fact that the predicted returns were not excessive once you took into account the greater riskiness of the the returns. Furthermore, some apparent predictability of risk-adjusted returns be accounted by time variations in the pricing kernal $k_t(s)$, or equivalently of the marginal rate of substitution function.

This pattern of empirical work uncovering apparent pricing anomalies and theorists coming up with more general theoretical explanations to account for them has continued unabated to the present. The participants tend to line up into either of two camps depending upon the class of explanations they tend to favor. In the last fifteen years or so there has been great interest in *behavioral* explanations which might suggest that some aspects of investor behavior which might be irrational in the sense that they are do not maximize a well-defined utility function or they do not process information in a correct manner (e.g., by doing Bayesian updating). Some of the thinking in this camp has built insights of psychology where concepts such as *over-confidence*, *envy* and *biased perceptions* have been widely used.

The other line of work tends to retain the assumption that agents are rational, and instead looks for explanations of pricing anomalies in more general representations of *agents' objectives* or in the *institutional environment*. Examples of the former line of research are those models that posit preferences exhibiting *habit formation* or *ambiguity aversion*. Examples of the latter are models that take into account *agency problems* that can emerge for example in *delegated investment management* or through *imperfect incentive schemes* for financial analysts.

5 Conclusion

If in this last section the discussion of recent developments in finance has left the impression that there is little consensus about the best approach to understanding financial imperfections, it is not just because we have been very brief. Had we gone into this in three times the depth and detail the result would have been a better awareness of the multiplicity of the ways finance specialists think about open issues but very little in the way of a settled, unitary theory. Nevertheless, I would argue that there is an implicit common thrust in most of current finance research. This is the shared goal of achieving a coherent body of theory that is as complete and internally consistent as the theory based on no-arbitrage in complete financial markets summarized in section 3. The paradigm of self-interested agents maximizing some objective subject to constraints imposed by the institutional environment has proved so rich and malleable that virtually no financial economist of my acquaintance makes any pretense of offering a revolutionary idea that would sweep this framework away.