# Chapter 1

# Introduction

Experience, how much and of what, is a valuable commodity. It is a major difference between an airline pilot and a New York Cab driver, a surgeon and a butcher, a succesful financeer and a cashier at your local grocers. Experience with data, with its analysis, experience constructing portfolios, trading, and even experience losing money (one experience we all think we could do without) are all part of the education of the financially literate. Of course, few of us have the courage to approach the manager of our local bank and ask for a few million so we can acquire this experience, and fewer still managers have the courage to acceed to our request. The "joy of simulation" is that you do not need to have a Boeing 767 to fly one, and that you don't need millions of dollars to acquire a considerable experience valuing financial products, constructing portfolios and testing trading rules. Of course if your trading rule is to buy condos in Florida because you expect boomers to all wish to retire there, a computer simulation will do little to help you since the ingredients to your decision are largely psychological (yours and theirs), but if it is that you should hedge your current investment in condos using financial derivatives real estate companies, then the methods of computer simulation become relevant.

This book concerns the simulation and analysis of models for financial markets, particularly traded assets like stocks, bonds. We pay particular attention to financial derivatives such as options and futures. These are financial instruments which derive their value from some associated asset. For example a call option is written on a particular stock, and its value depends on the price of the stock at expiry. But there are many other types of financial derivatives, traded on assets such as bonds, currency markets or foreign exchange markets, and commodities. Indeed there is a growing interest in so-called "real options", those written on some real-world physical process such as the temperature or the amount of rainfall.

In general, an option gives the holder a right, not an obligation, to sell or buy a prescribed asset (the underlying asset) at a price determined by the contract (the exercise or strike price). For example if you own a call option on shares of IBM with expiry date Oct. 20, 2000 and exercise price $120, then on October 20, 2000 you have the right to purchase a fixed number , say 100 shares of IBM at the price $120. If IBM is selling for $130 on that date, then your option is worth $10 per share on expiry. If IBM is selling for $120 or less, then your option is worthless. We need to know what a fair value would be for this option when it is sold, say on February 1, 2000. Determining this fair value relies on sophisticated models both for the movements in the underlying asset and the relationship of this asset with the derivative, and is the subject of a large part of this book. You may have bought an IBM option for two possible reasons, either because you are speculating on an increase in the stock price, or to hedge a promise that you have made to deliver IBM stocks to someone in the future against possible increases in the stock price. The second use of derivatives is similar to the use of an insurance policy against movements in an asset price that could damage or bankrupt the holder of a portfolio. It is this second use of derivatives that has fueled most of the phenomenal growth in their trading. With the globalization of economies, industries are subject to

more and more economic forces that they are unable to control but nevertheless wish some form of insurance against. This requires hedges against a whole litany of disadvantageous moves of the market such as increases in the cost of borrowing, decreases in the value of assets held, changes in a foreign currency exchange rates, etc.

The advanced theory of finance, like many areas where advanced mathematics plays an important part, is undergoing a revolution aided and abetted by the computer and the proliferation of powerful simulation and symbolic mathematical tools. This is the mathematical equivalent of the invention of the printing press. The numerical and computational power once reserved for the most highly trained mathematicians, scientists or engineers is now available to any competent programmer.

One of the first hurdles faced before adopting stochastic or random models in finance is the recognition that for all practical purposes, the prices of equities in an efficient market are *random variables*, that is while they may show some dependence on fiscal and economic processes and policies, they have a component of randomness that makes them unpredictable. This appears on the surface to be contrary to the training we all receive that every effect has a cause, and every change in the price of a stock must be driven by some factor in the company or the economy. But we should remember that random models are often applied to systems that are essentially causal when measuring and analyzing the various factors influencing the process and their effects is too monumental a task. Even in the simple toss of a fair coin, the result is predetermined by the forces applied to the coin during and after it is tossed. In spite of this, we model it as a random variable because we have insufficient information on these forces to make a more accurate prediction of the outcome. Most financial processes in an advanced economy are of a similar nature. Exchange rates, interest rates and equity prices are subject to the pressures of a large number of traders, government agencies, speculators, as well as the forces applied by international

trade and the flow of information. In the aggregate there is an extraordinary number of forces and information that influence the process. While we might hope to predict some features of the process such as the average change in price or the volatility, a precise estimate of the price of an asset one year from today is clearly impossible. This is the basic argument necessitating stochastic models in finance. Adoption of a stochastic model does neither implies that the process is pure noise nor that we are unable to forecast. Such a model is adopted whenever we acknowledge that a process is not *perfectly* predictable and the *non-predictable* component of the process is of sufficient importance to warrant modeling.

Now if we accept that the price of a stock is a random variable, what are the constants in our model? Is a dollar of constant value, and if so, the dollar of which nation? Or should we accept one unit of a index what in some sense represents a share of the global economy as the constant? This question concerns our choice of what is called the "numeraire" in deference to the French influence on the theory of probability, or the process against which the value of our assets will be measured. We will see that there is not a unique answer to this question, nor does that matter for most purposes. We can use a bond denominated in Canadian dollars as the numeraire or one in US dollars. Provided we account for the variability in the exchange rate, the price of an asset will be the same. So to some extent our choice of numeraire is arbitrary- we may pick whatever is most convenient for the problem at hand.

One of the most important modern tools for analyzing a stochastic system is simulation. Simulation is the imitation of a real-world process or system. It is essentially a model, often a mathematical model of a process. In finance, a basic model for the evolution of stock prices, interest rates, exchange rates etc. would be necessary to determine a fair price of a derivative security. Simulations, like purely mathematical models, usually make assumptions about the behaviour of the system being modelled. This model requires inputs, often

called the parameters of the model and outputs a result which might measure the performance of a system, the price of a given financial instrument, or the weights on a portfolio chosen to have some desirable property. We usually construct the model in such a way that inputs are easily changed over a given set of values, as this allows for a more complete picture of the possible outcomes.

Why use simulation? The simple answer is that is that it transfers work to the computer. Models can be handled which have greater complexity, and fewer assumptions, and a more faithful representation of the real-world than those that can be handled tractable by pure mathematical analysis are possible. By changing parameters we can examine interactions, and sensitivities of the system to various factors. Experimenters may either use a simulation to provide a numerical answer to a question, assign a price to a given asset, identify optimal settings for controllable parameters, examine the effect of exogenous variables or identify which of several schemes is more efficient or more profitable. The variables that have the greatest effect on a system can be isolated. We can also use simulation to verify the results obtained from an analytic solution. For example many of the tractable models used in finance to select portfolios and price derivatives are wrong. They put too little weight on the extreme observations, the large positive and negative movements (crashes), which have the most dramatic effect on the results. Is this lack of fit of major concern when we use a standard model such as the Black-Scholes model to price a derivative? Questions such as this one can be answered in part by examining simulations which accord more closely with the real world, but which are intractable to mathematical analysis.

Simulation is also used to answer questions starting with "what if". For example, What would be the result if interest rates rose 3 percentage points over the next 12 months? In engineering, determining what would happen under more extreme circumstances is often referred to as stress testing and simulation is a particularly valuable tool here since the scenarios we are concerned about are

those that we observe too rarely to have a substantial experience of. Simulations are used, for example, to determine the effect of an aircraft of flying under extreme conditions and is used to analyse the flight data information in the event of an accident. Simulation often provides experience at a lower cost than the alternatives.

But these advantages are not without some sacrifice. Two individuals may choose to model the same phenomenon in different ways, and as a result, may have quite different simulation results. Because the output from a simulation is random, it is sometimes harder to analyze- some statistical experience and tools are a valuable asset. Building models and writing simulation code is not always easy. Time is required both to construct the simulation, validate it, and to analyze the results. And simulation does not render mathematical analysis unnecessary. If a reasonably simple analytic expression for a solution exists, it is always preferable to a simulation. While a simulation may provide an approximate numerical answer at one or more possible parameter values, only an expression for the solution provides insight to the way in which it responds to the individual parameters, the sensitivities of the solution.

In constructing a simulation, you should be conscious of a number of distinct steps;

1. Formulate the problem at hand. Why do we need to use simulation?

2. Set the objectives as specifically as possible. This should include what measures on the process are of most interest.

3. Suggest candidate models. Which of these are closest to the real-world? Which are fairly easy to write computer code for? What parameter values are of interest?

4. If possible, collect real data and identify which of the above models is most appropriate. Which does the best job of generating the general

characteristics of the real data?

5. Implement the model. Write computer code to run simulations.

6. Verify (debug) the model. Using simple special cases, insure that the code is doing what you think it is doing.

7. Validate the model. Ensure that it generates data with the characteristics of the real data.

8. Determine simulation design parameters. How many simulations are to be run and what alternatives are to be simulated?

9. Run the simulation. Collect and analyse the output.

10. Are there surprises? Do we need to change the model or the parameters? Do we need more runs?

11. Finally we document the results and conclusions in the light of the simulation results. Tables of numbers are to be avoided. Well-chosen graphs are often better ways of gleaning qualitative information from a simulation.

In this book, we will not always follow our own advice, leaving some of the above steps for the reader to fill in. Nevertheless, the importance of model validation, for example, cannot be overstated. Particularly in finance where data is often plentiful, highly complex mathematical models are too often applied without any evidence that they fit the observed data adequately. The reader is advised to consult and address the points in each of the steps above with each new simulation (and many of the examples in this text).

## Example

Let us consider the following example illustrating a simple use for a simulation model. We are considering a buy-out bid for the shares of a company. Although the company's stock is presently valued at around $11.50 per share, a careful analysis has determined that it fits sufficiently well with our current

assets that if the buy-out were successful, it would be worth approximately \$14.00 per share in our hands. We are considering only three alternatives, an immediate cash offer of \$12.00, \$13.00 or \$14.00 per share for outstanding shares of the company. Naturally we would like to bid as little as possible, but we expect a competitor to virtually simultaneously make a bid for the company and the competitor values the shares differently. The competitor has three bidding strategies that we will simply identify as I, II, and III. There are costs associated with any pair of strategies (our bid-competitor's bidding strategy) including costs associated with losing a given bid to the competitor or paying too much for the company. In other words, the payoff to our firm depends on the amount bid by the competitor and the possible scenarios are as given in the following table.

| | | | Competitor's | Strategy | |
|---|---|---|---|---|---|
| | Bid | I | II | III | |
| Your | 12 | 3 | 2 | -2 | |
| Bid | 13 | 1 | -4 | 4 | |
| | 14 | 0 | -5 | 5 | |

The payoffs to the competitor are somewhat different and given below

| | | Competitor's | Strategy | |
|---|---|---|---|---|
| | | I | II | III |
| Your | 12 | -1 | -2 | 3 |
| Bid | 13 | 0 | 4 | -6 |
| | 14 | 0 | 5 | -5 |

For example, the combination of your bid=\$13 per share and your competitor's strategy $II$ results in a loss of 4 units (for example four dollars per share) to you and a gain of 4 units to your competitor. However it is not always the case that the your loss is the same as your competitor's gain. A game with this property is called a *zero-sum game* and these are much easier to analyze analytically. Define the $3 \times 3$ matrix of payoffs to your company by $A$ and the

payoff matrix to your competitor by $B$,

$$A = \begin{pmatrix} 3 & 2 & -2 \\ 1 & -4 & 4 \\ 0 & -5 & 5 \end{pmatrix}, \quad B = \begin{pmatrix} -1 & -2 & 3 \\ 0 & 4 & -6 \\ 0 & 5 & -5 \end{pmatrix}.$$

Provided that you play strategy $i = 1, 2, 3$ (i.e. bid \$12,\$13,\$14 with probabilities $p_1, p_2, p_3$ respectively and the probabilities of the competitor's strategies are $q_1, q_2, q_3$. Then if we denote

$$p = \begin{pmatrix} p_1 \\ p_2 \\ p_3 \end{pmatrix}, \text{ and } q = \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix},$$

we can write the expected payoff to you in the form $\sum_{i=1}^{3} \sum_{j=1}^{3} p_i A_{ij} q_j$. When written as a vector-matrix product, this takes the form $p^T A q$. This might be thought of as the average return to your firm in the long run if this game were repeated many times, although in the real world, the game is played only once. *If the vector $q$ were known to you,* you would clearly choose $p_i = 1$ for the row $i$ corresponding to the maximum component of $Aq$ since this maximizes your payoff. Similarly if your competitor knew $p$, they would choose $q_j = 1$ for the column $j$ corresponding to the maximum component of $p^T B$. Over the long haul, if this game were indeed repeated may times, you would likely keep track of your opponent's frequencies and replace the unknown probabilities by the frequencies. However, we assume that both the actual move made by your opponent and the probabilities that they use in selecting their move are unknown to you at the time you commit to your strategy. However, if the game is repeated many times, each player obtains information about their opponent's taste in moves, and this would seem to be a reasonable approach to building a simulation model for this game. Suppose the game is played repeatedly, with each of the two players updating their estimated probabilities using information gathered about their opponent's historical use of their available strategies. We

may record number of times each strategy is used by each player and hope that the relative frequencies approach a sensible limit. This is carried out by the following Matlab function;

```
function [p,q]=nonzerosum(A,B,nsim)
% A and B are payoff matrices to the two participants in a game.
Outputs
%mixed strategies p and q determined by simulation conducted nsim
times
n=size(A);                        % A and B have  the same size
p=ones(1,n(1)); q=ones(n(2),1); % initialize with positive weights
on all strategies
for i=1:nsim                    % runs the simulation nsim times
[m,s]=max(A*q);                 % s=index of optimal strategy for
us
[m,t]=max(p*B);                 % =index of optimal strategy for
competitor
p(s)=p(s)+1;                    % augment counts for us
q(t)=q(t)+1;                    % augment counts for competitor
end
p=p-ones(1,n(1)); p=p/sum(p);  %remove initial weights from counts
and then
q=q-ones(n(2),1); q=q/sum(q);  % convert counts to relative frequencies
```

The following output results from running this function for 50,000 simulations.

```
[p,q]=nonzerosum(A,B,50000)
```

This results in approximately $p' = [\frac{2}{3}\ 0\ \frac{1}{3}]$ and $q' = [0 \quad \frac{1}{2} \quad \frac{1}{2}]$ with an average payoff to us of 0 and to the competitor 1/3. This seems to indicate that the strategies should be "mixed" or random. You should choose a bid of \$12.00 with probability around 2/3, and \$14.00 with probability 1/3. It appears that

the competitor need only toss a fair coin and select between $B$ and $C$ based on its outcome. Why randomize your choice? The average value of the game to you is 0 if you use the probabilities above (in fact if your competitor chooses probabilities $q' = [0 \quad \frac{1}{2} \quad \frac{1}{2}]$ it doesn't matter what your frequencies are, your average is 0). If you were to believe a single fixed strategy is always your "best" then your competitor could presumably determine what your "best" strategy is and act to reduce your return (i.e. substantially less than 0) while increasing theirs. Only randomization provides the necessary insurance that neither player can guess the strategy to be employed by the other. This is a rather simple example of a two-person game with non-constant sum (in the sense that A+B is not a constant matrix). Mathematical analysis of such games can be quite complex. In such case, provided we can ensure cooperation, participants may cooperate for a greater total return.

There is no assurance that the solution above is optimal. In fact the above solution is worth an average of 0 per game to us and 1/3 to our competitor. If we revise our strategy to $p' = [\frac{2}{3} \ \frac{2}{9} \ \frac{1}{9}]$, for example, our average return is still 0 but we have succeeded in reducing that of our opponent to 1/9. The solution we arrived at in this case seems to be sensible solution, achieved with little effort. Evidently, in a game such as this, there is no clear definition of what an optimal strategy would be, since one might plan one's play based on the worst case, or the best case scenario, or something in between such as an average? Do you attempt to collaborate with your competitor for greater total return and then subsequently divide this in some fashion? This simulation has emulated a simple form of competitor behaviour and arrived at a reasonable solution, the best we can hope for without further assumptions.

There remains the question of how we actually select a bid with probabilities 2/3, 0 and 1/3 respectively. First let us assume that we are able to choose a "random number" $U$ in the interval [0,1] so that the probability that it falls in any given subinterval is proportional to the length of that subinterval. This

means that the random number has a uniform distribution on the interval [0,1]. Then we could determine our bid based on the value of this random number from the following table;

| | | | |
|---|---|---|---|
| If | $U < 2/3$ | | $2/3 \leq U < 1$ |
| Bid | 12 | 13 | 14 |

The way in which $U$ is generated on a computer will be discussed in more detail in chapter 2, but for the present note that each of the three alternative bids have the correct probabilities.