

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Valuing Risky Income Streams in Incomplete Markets

Clare Johnson

Supervised by Dr Peter Ouwehand
Department of Mathematics and Applied Mathematics,
University of Cape Town

3 September 2001

Abstract

Empirical evidence suggesting that world financial markets are incomplete leads to the question of how best to price and hedge contingent claims and derivative securities in incomplete markets. The focus of this dissertation is on a model proposed by Carr, Geman and Madan [7], which combines elements of arbitrage pricing theory with expected utility maximisation to decide whether a risky investment opportunity is worth undertaking or not.

An account of the state of the art of pricing and hedging in incomplete markets is followed by a detailed exposition of the new model. A chapter which details the issues which arise when the model is extended treats multiple time periods, continuous time, and an infinite state space. It is not entirely obvious in each case how the model may be extended, and current work is considered along with some new suggestions to address these issues. A small battery of computer simulations based on the proposed multiple period model is performed using a trinomial tree structure. A justification for using the new model rather than finite difference or classical multinomial tree methods is provided in the form of an argument which establishes the validity of a new approach in cases when the Black-Scholes formulation cannot be applied, chiefly when the market is incomplete.

Complete listing of the MATLAB code written to test the model, and plots of a representative cross-section of results obtained after running these scripts are included as appendices.

University of Cape Town

Contents

1	Introduction	1
2	Review of Pricing Techniques	3
3	An Incomplete Market Model	9
3.1	Introducing the model	9
3.1.1	Notation and terminology	11
3.2	The First Fundamental Theorem	13
3.2.1	Proof of the theorem	14
3.3	The Second Fundamental Theorem	16
3.3.1	Proof of the theorem	17
3.4	Pricing derivatives	18
3.4.1	Acceptably complete markets	18
3.4.2	Persistently incomplete markets	18
3.4.3	A worked example	20
3.5	Discussion	21
4	Extensions and Modifications	23
4.1	An alternative approach	23
4.2	Infinite state space	28
4.3	Multiple time periods	30
4.3.1	Traditional extension to multiple time periods	31
4.3.2	Other approaches in acceptably complete markets	33
4.3.3	Persistently incomplete markets	34
4.4	Continuous-time as a limit of discrete-time	34
4.4.1	Asset price processes	35
4.4.2	Pricing contingent claims	36
4.4.3	Discussion	37
4.5	Continuous time by direct extension	38
5	Implementation and Testing	39
5.1	Implementation	39
5.1.1	Generating the measures	40
5.1.2	Classifying the measures	40

5.1.3	Calibrating the floors	41
5.1.4	Solving for the representative state pricing function	43
5.1.5	Scaling the programs	44
5.2	Testing	44
5.2.1	The trinomial model	44
5.2.2	Calibration of parameters	46
5.2.3	Outline of the testing scheme	46
5.3	Results	47
5.3.1	Acceptably complete market	47
5.3.2	Persistently incomplete market	49
6	Conclusions and further work	51
A	Some representative test results	53
A.1	Acceptably complete market	54
A.2	Persistently incomplete market	56
B	Program Listing	57
B.1	Driver program: generate.m	57
B.2	Complete market: complete.m	58
B.3	Incomplete market: incomplete.m	58
B.4	Initialisation: setuptrinomial.m	61
B.5	Trinomial tree: trinomial.m	61
B.6	Discounting function: disc.m	65
B.7	Singular value decomposition solver: solve.m	65
B.8	Backsubstitution: svdbksb.m	65
B.9	Bid price: bid.m	66
B.10	Asking price: ask.m	66
B.11	PCx interface: pcx.m	67
B.12	Call option boundary condition: call.m	68
B.13	Put option boundary condition: put.m	68

Chapter 1

Introduction

A pervasive problem faced by market makers and investors in world financial markets is that of deciding whether to accept or reject a given investment opportunity. The question of whether the risks associated with accepting an opportunity outweigh its potential rewards determines the attractiveness of the opportunity to the market participant.

The classical Black-Scholes model [6] assumes that contingent claims may be perfectly hedged to eliminate all risk. Arbitrage arguments then dictate that the return of the resulting portfolio must equal the risk-free interest rate. Using this idea and several assumptions about the structure and behaviour of the market, the Black-Scholes equation for a derivative security based on some underlying asset is derived. Solutions to this partial differential equation may be obtained by assuming that all investors are indifferent to risk, an assumption which is permitted because the equation is independent of the expected return of the underlying. So-called risk-neutral valuation is a powerful tool when markets satisfy the assumptions underlying the Black-Scholes model, but in the real world this is almost never the case. As a result, theories analogous to no-arbitrage pricing have developed, which take into account the inherent riskiness attached to being unable to perfectly hedge or replicate a contingent claim – a scenario which arises when the market is incomplete.

The focus of this exposition will thus be on the central question of valuing contingent claims in incomplete markets. The pricing methodology proposed to address this problem in Carr, Geman and Madan [7] is examined in detail.

Chapter 2 presents a review of the major research on pricing and hedging in incomplete markets to date. In particular, the link between the classical theory of no-arbitrage pricing in a complete market and much of the work on extending that theory to incomplete securities market models is examined. Chapter 3 presents an introduction to the pricing system which will be discussed in the remaining chapters, and which is based on Carr, Geman and Madan [7]. This initial outline develops the model in a single-period

market with a finite number of assets and states of the world.

An assessment of the effects of and problems associated with changing the underlying assumptions made in initially building the model follows in Chapter 4. Relaxing these assumptions leads to market models which admit multiple time periods, continuous time, and infinite state spaces. Moreover, an alternative derivation of the major results of Chapter 3 is provided, which uses the fundamental results on arbitrage pricing in a complete market to establish the analogous results in an incomplete market. Chapter 5 details computer simulations of an incomplete market using a trinomial tree structure, the most concise incomplete market model possible. Contingent claims are priced in this market using the results and methodology of the model described in Chapters 3 and 4.

Further discussion of the results obtained in Chapter 5, conclusions and suggestions for further work follow in Chapter 6. Graphs showing some results representative of those obtained by from the MATLAB [22] test scripts appear in Appendix A. A full listing of all the MATLAB code written to perform the testing of the model appears in Appendix B.

Chapter 2

Review of Incomplete Market Pricing and Hedging Techniques

From the seminal work of Harrison and Kreps [17] on derivative pricing and hedging in arbitrage-free markets, two results stand out. First, if there are no free lunches, there exists a state pricing function, which is a positive linear functional from which a risk-neutral probability measure may be constructed. This measure is equivalent to the reference or objective probability measure on the state space, and under it the discounted asset prices in the securities market model are martingales. For this reason the measure is known as an equivalent martingale measure. Second, if the market is complete – that is, there are at least as many traded assets as there are states of the world (or sources of uncertainty, in a continuous-time setting) – then the state pricing function, and hence the equivalent martingale measure, is unique.

These results, which are referred to as the First and Second Fundamental Theorems of Asset Pricing respectively, have been the foundation for the study of contingent claim pricing and hedging in complete securities markets. Empirical evidence and common sense, however, suggest that financial markets in the real world are incomplete. Examples of studies which make a case for incompleteness in world markets are Figlewski and Green [11] and Aiyagari [1].

Using simulations performed on historical data (as opposed to a Monte Carlo simulation which produces its own “data”), Figlewski and Green [11] shows that an institution attempting to execute a delta hedge of options using the Black-Scholes model, even using optimal volatility forecasts, will be exposed to considerable risks. This is symptomatic of the real-world failure of the assumptions made in the derivation of the Black-Scholes model [6]. A market model that does not include transaction costs, stochastic

volatility and stochastic interest rates, for example, has a smaller state space than one that does. The assumptions of continuous trading and continuous price processes increase the size of the asset space. The combined effect of shrinking the state space and increasing the size of the asset space pushes the market towards a complete state. In a discrete-time complete market the asset space is of comparable size to the state space: the inflated state space and diminished asset space which result from the relaxation of the Black-Scholes world-view are thus likely to lead to an incomplete market.

In a continuous-time environment a market may only be complete when the number of sources of randomness does not exceed the number of risky assets. Stochastic volatility and interest rates add extra sources of uncertainty, even though continuous trading may facilitate portfolio replication. An incomplete market may also be the result of modifying the Black-Scholes assumptions in a continuous-time setting.

Aiyagari [1] suggests that the historically high premium for investing in equities and the comparatively low risk-free rate can be explained by incorporating transaction costs and market incompleteness into the economists' standard market model, which assumes complete, frictionless markets as its starting point. The historical average return on equities (1949-1978) is 7%, against an average return on 90-day US Treasury Bills of less than 1%. Introducing the assumption of market incompleteness lowers the predicted risk-free rate, whereas transaction costs enlarge the equity premium predicted by the model to be more in line with observed trends.

The main obstacles present in an incomplete market are that the equivalent martingale measure is not unique, and not every contingent claim is attainable. Claims which cannot be replicated are intrinsically risky, and even those claims for which a replicating portfolio does exist may not have a unique price. Given a contingent claim X with expiry $t = T$, and a bank account or numeraire process B_t , traditional pricing theory provides the upper bound

$$V_+(X) = \sup\{\mathbb{E}_{\mathbb{Q}}(B_T^{-1}X) : \mathbb{Q} \text{ an equivalent martingale measure}\} \quad (2.1)$$

and the lower bound

$$V_-(X) = \inf\{\mathbb{E}_{\mathbb{Q}}(B_T^{-1}X) : \mathbb{Q} \text{ an equivalent martingale measure}\} \quad (2.2)$$

for the time $t = 0$ value of the claim (Pliska [27]).

The collection of equivalent martingale measures on the state space thus provides a range of prices for a claim, and it is incumbent upon the market-maker to choose one of these measures to value the claim so that he may decide whether to buy or sell it, and how to hedge it. Karatzas [24] describes a way of choosing one of the martingale measures based on the principle of no-arbitrage and the personal risk preferences (in the form of a utility function) of the investor, arguing that it is not possible, based only on the

principle that the market should not admit arbitrage, to choose a unique price for a contingent claim. For this reason the investor's personal attitude towards risk is considered.

The hedging problem in incomplete markets is addressed in El-Karoui and Quenez [10], and, since the price of a contingent claim is intimately related to the cost of portfolios which hedge it, their results touch on the pricing issue as well. The *bid-ask spread* is defined to be the size of the interval of prices determined by all the available risk neutral measures. To achieve this, the set of all possible prices for a contingent claim is considered, each price obtained with a different probability measure. (If every contingent claim is attainable, this set is a singleton containing the unique risk-neutral measure determined by the system: by arbitrage arguments, the unique price of any claim is given by the discounted value of its replicating portfolio under this measure.) If a unique price for a contingent claim is not available, the infimum and supremum of the set of possible prices is examined. By arbitrage arguments, the price of the claim must lie between these two bounds. The infimum of the set of all possible prices is shown to be the bid price. This is the largest price for which a self-financing hedging portfolio with the same value as the claim at time zero may be sold. The difference between the value of the claim and the price of the hedging portfolio is a process increasing from 0 at time zero, meaning that should the buyer wish to close out his position at any time, he can sell the claim he holds and buy back the hedge portfolio to make a profit. The selling or asking price is shown to be the supremum of the set of all possible prices of the contingent claim, which is the smallest price for which the seller may hedge the contingent claim (by purchasing some hedge portfolio), and guarantee himself a profit. The seller may buy back the claim he sold and get rid of the hedge portfolio he holds to lock in a profit at any time. This latter method has become known in the literature as *superhedging*.

A drawback of the superhedging technique is that it can be prohibitively expensive. Föllmer and Leukert [13] suggests an alternative hedging technique called quantile hedging, designed to work in both incomplete and complete markets. In an incomplete market, a suitable hedging portfolio for a contingent claim is not immediately obvious, as a perfect hedge may be impossible. Moreover, in a complete market, an investor may not wish to hedge a claim completely, because although a perfect hedge eliminates risk, it also eliminates the chance of making an unexpected profit. Two methods are considered. The first method maximises the probability of a successful hedge (under the reference probability measure), given a fixed initial cost which represents the maximum amount of money the investor is willing to pay to hedge a claim. The second method bears some similarities to Value at Risk (VaR), in that a shortfall probability ϵ (the chance of the hedge failing) of, say, 1%, is specified, and the minimal cost of a hedge which will succeed with probability $1 - \epsilon$ is calculated. Unlike the static VaR approach, this

strategy is dynamic, and the investor is able to react to price movements as they occur. For this reason, a smaller amount of capital is required. A drawback that this method shares with VaR, however, is that the size of the shortfall at expiry (when the hedge has failed to cover the liabilities ensuing from having written or purchased the claim) is not taken into account.

This problem is addressed by the same authors in Föllmer and Leukert [14]. Shortfall risk is defined as the expected value of the shortfall of the hedge, weighted by some convex loss function $l(x)$ on $[0, \infty)$. An example of such a function, which indicates the investor's risk aversion, is $l(x) = x^p$ for $p \geq 1$. Hedging strategies which minimise this shortfall risk, given some restriction on the amount of capital the investor is prepared to put forward, are considered.

Other approaches have focused on trying to strengthen the notion of an arbitrage-free market in order to place bounds on the price of a claim which cannot be replicated. By making additional assumptions about the market, Cochrane and Saá-Requejo [9], for example, attempts to obtain narrow bounds (so-called *good-deal bounds*) on a reasonable price for a security. The suggestion is made that investors will not only unquestioningly exploit arbitrage opportunities until they disappear, but also opportunities to purchase assets with an extraordinarily high Sharpe ratio (the expected excess return of the asset over and above the risk-free rate, divided by the asset's standard deviation). These opportunities are called *good deals*. The effect of investors immediately purchasing these assets whenever they became available would drive down the excess return on the high Sharpe ratio assets, lowering their Sharpe ratios and returning the market to equilibrium. Empirical tests on S&P500 index option prices using upper and lower bounds calculated with this method show that good-deal bounds are a good approximation of market behaviour. A drawback of this family of methods is that subjective probabilities still need to be specified, and very little of the price information that does exist in the market is actively used in the pricing procedure.

Černý and Hodges [8] also explores the theory of good-deal pricing, and shows that all pricing structures derived by generalising the theory of no-arbitrage pricing are essentially the same. Adopting the terminology of Cochrane and Saá-Requejo [9], a good deal is defined to be the opportunity to buy a desirable claim at no cost, and it is demonstrated that the more attractive the investments that are excluded from the market (examples are arbitrage opportunities and high Sharpe ratio assets), the narrower the band of fair prices obtained for a security. Just as the Extension Theorem of economics (Kreps [25], and also Pliska [27]) shows that any arbitrage-free incomplete market can be extended to a complete market by adding some "fictitious" assets, Černý and Hodges [8] shows that any incomplete market with no good deals (however those are defined) may be extended in the same way to a complete market with no good deals. Portfolio optimisation

problems may be solved using the unique risk-neutral measure yielded by the extended market, with additional constraints to prevent positions from being taken in the fictitious assets. This approach does not carry through usefully to the problem of pricing claims uniquely, though the no-good-deal bounds will provide an indication of the range in which the claim's price lies. In the extended, complete market, it *will* be possible to replicate every contingent claim, but the replications may involve securities that do not exist, leading to contingent claim prices that are not useful.

Central to most of the approaches to pricing and hedging in incomplete markets is the need to measure the riskiness of a position, and to compare it – ideally, numerically – to the rewards that may be expected. Techniques such as Value at Risk attempt to measure risk, and to express it in useful language, but the fact that the very definition of risk is sometimes disputed suggests the difficulties inherent in characterising it with a number or collection of numbers.

Artzner, Delbaen, Eber and Heath [3] defines risk as the variability of a future position, as opposed to a definition which encompasses the change in wealth between two dates. Risks are random variables on a state space at a future date, representing the possible future values of portfolios which are currently held. The assumption is made that future states of the world are known, but that the probability of an event occurring is not generally agreed upon. Risks are classified as acceptable or unacceptable using a risk measure, which yields a real number indicating how risky a position is. The authors point out that, while representing a risk as a single number seems to indicate a great loss of information, it reflects the fact that the decision whether to take on a risk or not is essentially binary, requiring a yes or no.

Holding an unacceptably risky position would require that the investor modify the position, perhaps by adding some “reference instrument” (typically bonds or money in the bank) to the position to make it acceptable. The cost of adding these instruments is the measure of the riskiness of the position. If X is a risk, and ρ is some risk measure, then $\rho(X) > 0$ if X is unacceptable (a positive amount of money needs to be deposited into the bank), and $\rho(X) \leq 0$ if X is acceptable (the amount $\rho(X)$ may be safely withdrawn from the bank). An elegant set of axioms which an economically sensible risk measure might be expected to satisfy is proposed. A risk measure which satisfies axioms of translation invariance, subadditivity, positive homogeneity, and monotonicity is called a *coherent risk measure*. Moreover, Artzner, Delbaen, Eber and Heath [3] shows that the widely used Value at Risk measure is not coherent, and suggests other ways of constructing such a measure. In a proposition closely related to the models of Carr, Geman and Madan [7] and Heath and Ku [19] which will be considered in subsequent chapters, they show that every coherent risk measure may be obtained by considering the supremum of the risks of each of a collection of “scenarios”, or possible future outcomes. Further discussion of this topic is found in

Chapter 4.

Intuitively, it makes sense to attempt to compare risks using a partial ordering to express preferences. Several studies (Černý and Hodges [8] and Jaschke and Küchler [23], for example) define such preference relations on the set of all risks, and Jaschke and Küchler [23] shows that these partial orderings are in one-to-one correspondence with the set of all risk measures. The theories of risk measures and of good-deal pricing are drawn together and a one-to-one correspondence between good-deal bounds and risk measures is also established. Each risk measure has an associated *acceptance set*, which is the set of all risks which are acceptable under that measure. In this framework, sets of acceptable risks are cones¹, a requirement that Artzner, Delbaen, Eber and Heath [3] describes as “less natural” than the other axioms on acceptance sets. While cones have useful mathematical properties, a conic acceptance set implies that arbitrarily large positions may be taken in acceptably risky opportunities. There is thus no risk associated with a dangerously large position in what may have only been a moderately risky portfolio to begin with. (The pricing framework of Carr, Geman and Madan [7] which subsequent chapters will consider does not have this cone property, for this very reason, which is explained further in Section 3.1.1.) Conversely, sets of acceptable risks may be used to define a risk measure. This correspondence is also one-to-one.

With this demonstration of the close connections between various theories on the problem of comparing and pricing risky income streams, economic and financial ideas which originate from a wide range of viewpoints are converging in the literature. Another property of these ideas, and of the theory which will be discussed in Chapter 3, is that the frameworks lie between the two schools of thought on arbitrage pricing and expected utility maximisation. Both models have their disadvantages and advantages, but by employing a system which requires less specification of agents’ personal preferences than is demanded by expected utility maximisation, and more than just the asset and state space required by arbitrage pricing theories, a happy compromise may be reached. Of course, the penalty for providing weaker inputs than are normally required by the utility maximisation theories is that less powerful results are obtained. On the other hand a slightly more comprehensive collection of inputs than is usually required to perform arbitrage pricing theory returns handsome rewards in an incomplete market setting. Whereas a wide range of prices for a derivative security was often forthcoming (see equations (2.1) and (2.2)), new pricing theories are able to tighten the bounds on the price of a security significantly, or even to select a single, “best” price from among the possibilities.

¹Recall that C is a cone if $C + C \subseteq C$ and $\lambda C \subseteq C$ for all $\lambda \geq 0$.

Chapter 3

An Incomplete Market Model

3.1 Introducing the model

A model to deal with the problems associated with pricing contingent claims in incomplete financial markets is suggested in Carr, Geman and Madan [7]. The central problem faced by an investor in an incomplete market is considered to be whether or not to accept a given (risky) investment opportunity, and mechanisms are developed which are intended to assist the investor in making a wise decision.

The existing tools for this branch of decision-making include arbitrage pricing theory (APT) and expected utility maximisation. The theory of expected utility maximisation is based on the assumption that an investor will accept an opportunity if and only if it increases his expected utility. While this is a powerful mechanism for financial decision-making, the main shortcoming of the theory is the scope of the inputs it requires. These include the investor's utility function, the stochastic processes followed by all assets, and the investor's initial wealth or endowment. In practice it is difficult to specify and calibrate these preferences and beliefs so that they are consistent and meaningful. The optimisation process is also very sensitive to the inputs, making the model unstable and in danger of delivering inconsistent results.

Arbitrage pricing theory produces results consistent with expected utility maximisation, and requires only that the state space and asset space be specified as inputs to the model. This means that in certain situations APT can be used instead of expected utility maximisation, simplifying the decision-making but providing equally useful results. Moreover, owing to the simplicity of the inputs required, the results obtained are independent of any beliefs or preferences the investor may hold. APT attempts to replicate an investment opportunity's payoff using the available assets. If this replication is possible, arbitrage arguments are used to determine a course of action,

which often involves synthesising arbitrages when investment opportunities are mispriced.

When the market is complete, APT is guaranteed to be able to replicate the payoff to any opportunity and a decision may be made based on arbitrage arguments. In an incomplete market, it is not always possible to execute a synthetic arbitrage strategy based on the value of a replicating portfolio, and using APT to evaluate investment opportunities becomes considerably more complex. An investor who rejects all investment opportunities which *cannot* be synthesised using the existing liquid assets will almost certainly be rejecting opportunities with finite, reasonable risks which a wide range of investors would consider to increase their expected utility.

These “controlled risks”, which are outweighed by their associated rewards, are the motivation for the theory developed in Carr, Geman and Madan [7]. In the tradition of the good-deal pricing theory introduced in Cochrane and Saá-Requejo [9], the principle of market efficiency is redefined to state that there does not exist a portfolio which would be viewed as beneficial by all investors, since such a portfolio would be traded until its price changed and made it less desirable. *Strictly acceptable opportunities* – those opportunities which are seen as having negligible associated risk compared with their potential benefits – are classified as opportunities which would be accepted by all rational investors. Strictly acceptable opportunities are the “good deals” of this model.

The only inputs required by the model are the state and asset spaces, and a family of probability measures on the state space called *test measures*, intended to represent the range of views of the world held by the market-makers. Like Artzner, Delbaen, Eber and Heath [3], the model assumes that the probability of each future state of the world occurring is not generally agreed upon. Thus each hypothetical investor can be thought of as being equipped with one of the test measures, with which he evaluates investment opportunities. An investor considers an opportunity’s rewards to outweigh its risks if the expected return of the opportunity exceeds a fixed, non-positive scalar value called a *floor*, which is specific to that investor (that is, to that investor’s test measure).

The results obtained using the new model reflect the fact that its inputs are more comprehensive than those required by arbitrage pricing theory but less so than those for expected utility maximisation. APT and expected utility maximisation are in fact limiting cases of this theory, a characteristic of all good-deal pricing theories (Černý and Hodges [8]). For example, the policy of accepting only arbitrage opportunities can be put into practice by having one valuation measure per state, which assigns probability 1 to that state and 0 to all the others.

The model is unable to answer some of the issues which utility theory can address, such as which individual investment decisions are optimal. Unlike APT, however, the model is able to decide which risky investment

opportunities are worth accepting, making it of use in an incomplete market setting.

3.1.1 Notation and terminology

This section establishes some notation which will be used in subsequent chapters, and which will have the same meanings, where possible, in the MATLAB [22] code in Appendix B.

The state space is assumed to be finite, equipped with a filtration $\mathcal{F} = \{\mathcal{F}_0, \mathcal{F}_T\}$ and a reference probability measure \mathbb{P} . Let S_t^i denote the i th traded security at time t , $1 \leq i \leq N$, and assume that the benchmark or numeraire process is B_t , with $B_0 = 1$. This process is assumed to be almost surely strictly positive. There are thus $N + 1$ traded assets. Let $\mathbf{S}_t = (B_t, S_t^1, \dots, S_t^N)$ be the $(N + 1)$ -dimensional process describing the time t values of the traded assets.

When a security's value has been discounted using this process, it will be written as \bar{S}_t^i . The time interval under consideration will range from 0 to the finite horizon T , and trading is possible only at $t = 0$ and $t = T$.

Denote by \mathbf{X} the set of all possible payoffs resulting from trading strategies or investment opportunities. Payoffs elements of \mathbb{R}^K , where K is the number of states of the world, and are \mathcal{F}_T -measurable.

Definition 3.1.1. An \mathcal{F}_T -measurable random variable (payoff) $X \in \mathbf{X}$ which is offered to the trader at zero initial cost is called an *opportunity*. (Opportunities with non-zero initial cost may be offset by borrowing at the risk-free rate in order to satisfy this definition.) Opportunities will also be referred to as *contingent claims*.

The expected return of an opportunity is evaluated using each of a collection of probability measures on the state space, known as *stress test measures*. Each stress test measure is assumed to be absolutely continuous with respect to the reference measure \mathbb{P} . Every such measure \mathbb{P}_m has an associated floor $f_m < 0$, a real number representing the maximum expected trading loss which is admissible under that measure. Denote by \mathcal{S}^s the set containing all the stress test measures.

A test measure with a floor $f_m = 0$ is called a *valuation measure*. The set of all valuation measures on Ω is denoted \mathcal{S}^v . The valuation test measures are also assumed to be absolutely continuous with respect to the reference measure \mathbb{P} . Let $\mathcal{S} = \mathcal{S}^v \cup \mathcal{S}^s$, and let M be the cardinality of \mathcal{S} .

The crux of the model outlined in Carr, Geman and Madan [7] is the extension of the collection of opportunities which any "reasonable" person would accept. The classical theory considers only arbitrage opportunities in this category, but here the collection is extended to include the idea of an *acceptable opportunity*.

Definition 3.1.2. An opportunity $X \in \mathbf{X}$ is called an *acceptable opportunity* if, for each $\mathbb{P}_m \in \mathcal{S}$,

$$\mathbb{E}_{\mathbb{P}_m}(X) \geq f_m$$

Let \mathcal{A} denote the set of all acceptable opportunities.

An acceptable opportunity is one which is considered to offer reasonable expected returns from a variety of perspectives (represented by the test measures). Arbitrage opportunities are acceptable opportunities, and it is thus clear that when there are no acceptable opportunities there can be no arbitrage opportunities.

Arbitrage opportunities may be scaled arbitrarily and still guarantee a non-negative return in every state of the world. Acceptable opportunities, however, cannot necessarily be scaled upwards (by increasing the position by a constant multiple) and remain acceptable. It is for this reason that the set \mathcal{A} is not a cone, whereas Artzner, Delbaen, Eber and Heath [3] and Jaschke and Küchler [23] develop cones of acceptable risks. This restriction is a result of the negative floors associated with each stress test measure, and is economically sensible, particularly in an illiquid economy where an excessively large position in any instrument places the holder at the risk of being unable to liquidate the position should he need to.

The concept of an acceptable opportunity is further refined by that of a strictly acceptable opportunity, this economic model's good deal. Such opportunities offer a positive expected payoff under at least one of the valuation measures. This refinement is analogous to the observation that a zero payoff at a zero cost is not an arbitrage opportunity in any useful sense.

Definition 3.1.3. An acceptable opportunity $X \in \mathbf{X}$ is called *strictly acceptable* if, for some $m \in \mathcal{S}^v$,

$$\mathbb{E}_{\mathbb{P}_m}(X) > 0$$

Let \mathcal{A}^+ denote the set of all strictly acceptable opportunities.

Let the elements of the matrix A be defined by

$$A_{i+1,j} = S_T^i(\omega_j)$$

and set the first row of A to be $(B_T(\omega_1), B_T(\omega_2), \dots, B_T(\omega_K))$, the final values of the numeraire process. A is thus the $(N+1) \times K$ matrix containing the time $t = T$ values of the traded assets in each state of the world. Let B be the $K \times M$ matrix whose columns B_m comprise the valuation and stress test measures. Let \vec{f} be a M -vector containing the floors associated with each measure, in the same order as the columns of B are arranged.

Let \vec{p} be the $(N+1)$ -vector of initial asset prices,

$$\vec{p} = \begin{pmatrix} B_0 \\ S_0^1 \\ \dots \\ S_0^N \end{pmatrix}$$

Then, using this notation, if there are no strictly acceptable opportunities in the economy, there does *not* exist an opportunity $\vec{\alpha}$ such that:

1. $\vec{\alpha}^T \vec{p} = 0$ (zero initial cost), and
2. $\vec{\alpha}^T AB \geq \vec{f}^T$ (its expected value under each test measure weakly exceeds that measure's floor), and
3. $\vec{\alpha}^T AB_m > 0$ for some column B_m of B corresponding to a valuation measure (its expected value under at least one of the valuation measures is strictly positive).

3.2 The First Fundamental Theorem of Asset Pricing

Before attempting to generalise the first fundamental result of Harrison and Kreps [17], the *representative state pricing function* is defined. This function allows us to distinguish between the (possibly many) classical state pricing functions which may exist in an incomplete market.

Definition 3.2.1. Recall that a *state pricing function* is a K -vector $\pi > 0$ such that, for each asset S_t^n in the asset space,

$$S_0^n = \sum_{\omega \in \Omega} \pi(\omega) S_T^n(\omega) \quad (3.5)$$

Asset pricing theory defines a martingale measure \mathbb{Q} using the state pricing function, setting

$$\mathbb{Q}(\omega) = \pi(\omega) B_T \quad (3.6)$$

Note that, according to Definition 3.2.1 (also Pliska [27]), the components of the state pricing function do not sum to 1 unless the risk-free rate is zero or the discounted asset prices are used on the right-hand side of equation (3.5). The state pricing function therefore needs to be normalised as in equation (3.6) to create a probability measure. If the risk-free rate r is discretely compounded and constant over the period $[0, T]$,

$$\sum_{\omega \in \Omega} \pi(\omega) = \frac{1}{1+r}$$

In an incomplete market, a more general object than a state pricing function is required. A *representative state pricing function* is a strict convex combination of the valuation measures \mathbb{P}_m . Its name is derived from the fact that every valuation test measure which assigns a positive probability to a

state $\omega \in \Omega$ contributes to the value of the representative state pricing function at ω .

Definition 3.2.2. A state pricing function π is a *representative state pricing function* if there exists a set of strictly positive weights w_m with each element corresponding to a $\mathbb{P}_m \in \mathcal{S}^v$, such that

$$\pi(\omega) = \sum_{\mathbb{P}_m \in \mathcal{S}^v} w_m \mathbb{P}_m(\omega) \quad (3.8)$$

and $w_m < 1$ for each m .

It is desirable that the weights defining the representative state pricing function sum to 1 so that it may be used completely analogously to the way in which a risk-neutral measure is used in a complete market. This can be achieved by working with the discounted $t = T$ asset prices when calculating the representative state pricing function, or by normalising after the fact (which is essentially the same approach, since one would be dividing the weights by B_T).

The existence of a representative state pricing function can be recast as there existing a $\vec{w} \geq 0$, with $w_m > 0$ for valuation measures \mathbb{P}_m and $w_k = 0$ for stress test measures \mathbb{P}_k , such that $\vec{p} = AB\vec{w}$. Because of the construction of \vec{w} and the definition of the vector \vec{f} , it follows that $\vec{f}^T \vec{w} = 0$.

3.2.1 Proof of the theorem

The following corollary to the Separating Hyperplane Theorem (see, for example, Pliska [27]) is required in the proof of Theorem 3.2.1:

Corollary 3.2.1. *Suppose that K is a convex, compact subset of \mathbb{R}^n and L is a subspace of \mathbb{R}^n with $L \cap K = \emptyset$. Then there is a vector $\vec{\pi}$ such that*

1. $\vec{\pi} \cdot \vec{l} = 0$ for each $\vec{l} \in L$
2. $\vec{\pi} \cdot \vec{k} > 0$ for each $\vec{k} \in K$

An analogous result to the first Fundamental Theorem of Asset Pricing of Harrison and Kreps [17] may now be proved, set in an incomplete market.

Theorem 3.2.1 (First Fundamental Theorem of Asset Pricing). *A single-period market has no strictly acceptable opportunities if and only if there exists a representative state pricing function.*

Proof. (\Leftarrow) Let a representative state pricing function be defined by its vector of weights $\vec{w} > 0$. Suppose that a strictly acceptable opportunity $\vec{\alpha}$, exists.

Then, since $\vec{p} = AB\vec{w}$, if $\vec{\alpha}^T \vec{p} = 0$, it follows that $\vec{\alpha}^T AB\vec{w} = 0$.

Since $\vec{\alpha}$ is a strictly acceptable opportunity, for some valuation measure \mathbb{P}_j , given by the column B_j of B , it must hold that

$$\vec{\alpha}^T AB_j > 0 \quad (3.9)$$

Now, because $\vec{\alpha}$ must be acceptable, it also holds that

$$\vec{\alpha}^T AB_m \geq 0 \quad (3.10)$$

for each valuation measure \mathbb{P}_m in the set of test measures \mathcal{S} (since valuation measures have floors of zero, and acceptable opportunities have expected payoffs weakly exceeding those floors).

Now $w_k = 0$ for each stress test measure \mathbb{P}_k , so that those columns of B corresponding to valuation measures make no contribution to the matrix equation

$$\vec{\alpha}^T AB\vec{w} = 0$$

Equations (3.9) and (3.10) thus contradict the fact that $\vec{\alpha}^T AB\vec{w} = 0$. It follows that there can be no strictly acceptable opportunities.

(\Rightarrow) Suppose now that there are no strictly acceptable opportunities.

Denote by B^v the submatrix of B created by deleting those columns which correspond to stress test measures, and let k be the column dimension of B^v .

Let $C_1 = \{\vec{\alpha}^T AB^v : \vec{\alpha}^T \vec{p} = 0\}$ be the linear space of expected payoffs from all zero-cost portfolios generated by the valuation measures.

Let

$$C_2 = \left\{ \vec{y} : \vec{y} \geq 0, \sum_{i=1}^k y_i = 1 \right\}$$

Then C_1 and C_2 are closed, convex sets. C_2 is also bounded, so it is compact. Since there are no strictly acceptable opportunities, $C_1 \cap C_2 = \emptyset$. For if $\vec{x} \in C_1 \cap C_2$, it follows that $\vec{x} \geq 0$ and since $\sum x_i = 1$, \vec{x} has at least one non-zero component. Moreover, $\vec{x} = \vec{\alpha}^T AB^v$ for some $\vec{\alpha}$, and $\vec{\alpha}^T \vec{p} = 0$. This means that \vec{x} is the vector of expected returns under each valuation measure on a zero-cost portfolio. Since it has at least one strictly positive component, \vec{x} represents a strictly acceptable opportunity, contradicting our assumption.

Then, by a corollary to the separating hyperplane theorem (Corollary 3.2.1), there exists a k -vector \vec{w} such that if $\vec{\alpha}^T AB \in C_1$ and $\vec{y} \in C_2$, both of the following are true:

$$\vec{\alpha}^T AB^v \vec{w} = 0 \quad (3.13a)$$

$$\vec{w}^T \vec{y} > 0 \quad (3.13b)$$

Equation (3.13b) implies that $\vec{w} > 0$.

Additionally,

$$\alpha^T AB^v \vec{w} = 0 = \vec{\alpha}^T \vec{p}$$

so the vectors $AB^v \vec{w}$ and \vec{p} are parallel in k -space. By scaling \vec{w} appropriately, its properties (equations (3.13a) and (3.13b) above) may be retained and the equality

$$AB^v \vec{w} = \vec{p} \tag{3.15}$$

will be satisfied.

The representative state pricing function is then defined to be $B^v \vec{w}$, or, setting $w_i = 0$ where the column B_i of B is a stress test measure, by $B\vec{w}$. \square

It is possible that there will be several solutions to the matrix equation (3.15). If there is only one solution, the representative state pricing function is obviously unique.

If the weights of the representative state pricing function \vec{w} are normalised to sum to unity, security prices calculated using the representative state pricing function will be forward prices. That is, they will need to be discounted, using the ruling risk-free rate, to the present day.

3.3 The Second Fundamental Theorem of Asset Pricing

The classical Second Fundamental Theorem of asset pricing (Harrison and Kreps [17]) shows the equivalence of market completeness with the uniqueness of the state pricing function. The state pricing function exists when there is no arbitrage.

Acceptable opportunities are investment opportunities which carry a non-zero probability of a loss in some state(s) of the world, but are not risk-free in the same way as arbitrage opportunities are required to be. In Carr, Geman and Madan [7], a generalisation of the idea of market completeness is suggested, in order to accommodate this property of acceptable opportunities.

Definition 3.3.1. A single-period market is said to be *acceptably complete* if, for every payoff $X \in \mathbf{X}$, there exists a hedge portfolio $\vec{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_N)$, such that for each $\mathbb{P}_m \in \mathcal{S}$,

$$\mathbb{E}_{\mathbb{P}_m}(\alpha_0 B_T + \sum_{n=1}^N \alpha_n S_T^n) - \mathbb{E}_{\mathbb{P}_m}(X) = f_m \tag{3.16}$$

This represents a weakening of the requirement that assets can be perfectly hedged (with zero residual), as in a complete market. In this setting, it is required that the hedged investment is just acceptable under each test measure. Negative expected payoffs are permitted, but they must exceed the

floors of the test measures. An incomplete market which is not acceptably complete will be referred to as a *persistently incomplete market*.

In an incomplete market, the many choices for the state pricing function can lead to widely different results when deciding on the price of a derivative. Before proving a theorem which allows us to avoid this problem in an acceptably complete market, a further condition on the asset space and space of measures is required.

Uniqueness of the representative state pricing function is only guaranteed when the matrix AB is of full rank. It is insufficient to demand only that B be of rank M , since this does not guarantee that the product AB has sufficient rank to yield a unique solution to the representative state pricing function.

Definition 3.3.2. The test measures are said to satisfy the *uniqueness condition* if the matrix AB is of rank M .¹

It is thus required that for each set of potential test results, a portfolio of the traded assets can be found to produce those results. The acceptable completeness criterion can thus be re-stated as requiring that for all payoffs \vec{x} there exists an $\vec{\alpha}$ such that

$$\vec{\alpha}^T AB = \vec{x}^T B + \vec{f}^T \quad (3.17)$$

where \vec{f} is the vector of floors associated with the test measures.

3.3.1 Proof of the theorem

In analogy to the classical theorem proved in Harrison and Kreps [17], which shows that in a complete and arbitrage-free market the risk-neutral measure is unique, the theorem proved here guarantees the uniqueness of the representative state pricing function when the market is acceptably complete.

Theorem 3.3.1 (The Second Fundamental Theorem of Asset Pricing). *The single-period market is acceptably complete if and only if the representative state pricing function is unique, provided that the uniqueness condition holds.*

Proof. (\Leftarrow) Suppose that the representative state pricing function defined by the weights \vec{w} is unique, and that the uniqueness condition holds.

This immediately implies that $\text{ran}(AB) = \mathbb{R}^M$ where M is the number of test measures, and so the market is acceptably complete.

(\Rightarrow) If the market is acceptably complete, for each payoff $\vec{x} \in \mathbf{X}$ there exists an $\vec{\alpha}$ such that

$$\vec{\alpha}^T AB = \vec{x}^T B + \vec{f}^T \quad (3.18)$$

¹The condition laid down in the original paper [7] (the *under-testing* condition) is that the matrix B be of rank M . This, however, does not guarantee that the product AB has rank M , hence this modified condition.

The uniqueness condition guarantees that the rank of AB is M . Thus the representative state pricing function, which is defined by the solution to the equation

$$p = ABw \quad (3.19)$$

with $w_m = 0$ for stress-test measures \mathbb{P}_m , is unique. \square

3.4 Pricing derivatives

3.4.1 Acceptably complete markets

When the number of assets in the marketplace weakly exceeds the number of valuation test measures (probability measures with floors equal to zero), a representative state pricing function can be uniquely determined. This allows one to price derivative securities uniquely.

If the representative state pricing function can be uniquely determined, it can be used analogously to the risk-neutral measure of classical asset-pricing theory. In the single-period model, the price of an asset is simply the discounted expected value (under the normalised representative state pricing function) of its payoff at expiry.

3.4.2 Persistently incomplete markets

In the event that the representative state pricing function cannot be uniquely determined, it is still possible to decide whether to purchase a derivative security, or to find a reasonable value for which to sell it. The possible solutions to the matrix equation

$$\vec{p} = AB\vec{w} \quad (3.20)$$

will provide a number of “state pricing functions”, yielding a range of fair values for the derivative, analogous to the bounds on the price of a contingent claim given by equations (2.1) and (2.2).

Given a derivative security with a payoff $\vec{x} \in \mathbf{X}$, an investor may wish to find the smallest price for which the derivative should be sold so that his hedge of the derivative will be acceptable (that is, will have an expected payoff under each test measure which exceeds its floor).

Let $\vec{\alpha}$ denote the proposed hedge portfolio for the derivative. An investor who has issued the derivative will have to pay the amount $\vec{x}^T B$ to the buyer at expiry. He expects to receive $\vec{\alpha}^T AB$ from the hedge, and would like to pay as little as possible to set up this hedge. The linear programming problem

$$\min_{\vec{\alpha}} \vec{\alpha}^T \vec{p} \quad \text{subject to} \quad \vec{\alpha}^T AB - \vec{x}^T B \geq \vec{f}^T \quad (3.21)$$

expresses this desire, including the constraint that the total expected return (the expected return on the hedge, less the derivative payoff) accruing to the investor at expiry should be acceptable.

If a solution $a(\bar{x})$ to equation (3.21) exists, it is the offer or *asking price* for the derivative with payoff \bar{x} .

The *bid price* for the derivative can be found in a similar manner. The market participant who wishes to buy the derivative would wish to maximise the amount of cash he receives at $t = 0$ from setting up a hedge portfolio $\bar{\alpha}$. The hedge would be constructed by issuing liabilities (including borrowing at the risk-free rate) – effectively selling a portfolio, which would result in an initial cash inflow of $\bar{\alpha}^T \bar{p}$. This inflow must be sufficiently large to ensure that the payoff from the derivative, less the cost of closing down the hedge position, is acceptable.

The problem of finding the bid price for the derivative can be formulated as the linear program

$$\max_{\bar{\alpha}} \bar{\alpha}^T \bar{p} \quad \text{subject to} \quad \bar{x}^T B - \bar{\alpha}^T AB \geq \bar{f}^T \quad (3.22)$$

which, if a solution exists, yields the bid price $b(\bar{x})$ for the opportunity \bar{x} .

The bid-ask spread $s(\bar{x}) = a(\bar{x}) - b(\bar{x})$ as formulated in equations (3.22) and (3.21) is positive or zero (Theorem 3, Carr, Geman and Madan [7]). Moreover, if a derivative is valued using a representative state pricing function, the resulting price will lie between the bid and the offer price for the security. This can be seen at once by considering the dual to the bid equation (3.22), which is

$$\min_{\bar{w}} \bar{x}^T B \bar{w} - \bar{f}^T \bar{w} \quad \text{subject to} \quad AB \bar{w} = \bar{p} \quad \text{and} \quad \bar{w} \geq 0 \quad (3.23)$$

and the dual to (3.21), the ask problem:

$$\max_{\bar{w}} \bar{x}^T B \bar{w} - \bar{f}^T \bar{w} \quad \text{subject to} \quad AB \bar{w} = \bar{p} \quad \text{and} \quad \bar{w} \geq 0 \quad (3.24)$$

It is clear from equation (3.24) that the asking price is found by maximising over representative state pricing functions, and that the bid price is determined by minimising over the same set. What also emerges from this observation is the already established fact that, in an acceptably complete market, since there is only one representative state pricing function, the bid and the offer price for a derivative must be equal.

When the market is not acceptably complete it may still be possible to hedge a derivative so that the total investment is exactly acceptable (in the sense that the expected returns are precisely the floors of each measure). If this is the case, it holds that the bid-ask spread is still zero, implying that derivative securities may be priced uniquely.

A further result (Theorem 5, Carr, Geman and Madan [7]) shows that the bid-ask spread, $s(\bar{x})$, increases as the scale of the trade is increased. That is, if $\lambda \in (0, 1)$, $s(\lambda \bar{x}) \leq s(\bar{x})$. This shows that the bid-ask spread for an opportunity may be narrowed by decreasing the scale of the investment.

3.4.3 A worked example

This section describes an example to illustrate the case of an acceptably complete market, set in a single period economy with dates 0 and 1, and three states of the world. The model has two assets, a unit bond and a stock. The continuously compounded risk-free rate is 10%, and the bond has an initial price of 1. The time 0 price of the stock is 100. With two assets and three states of the world, this economy is incomplete.

The time $t = 1$ matrix of the values of the two assets across the three possible states of the world is:

	ω_1	ω_2	ω_3
Bond	1.1	1.1	1.1
Stock	90	100	105

This is the matrix A of the theory developed in previous sections. There are two valuation measures, \mathbb{P}_1 and \mathbb{P}_2 . The probabilities they assign to the three future states are:

	\mathbb{P}_1	\mathbb{P}_2
ω_1	4/5	1/7
ω_2	1/10	1/7
ω_3	1/10	5/7

This is the matrix B of the preceding theorems and definitions. The vector of initial prices, \vec{p} , is given by

$$\vec{p} = \begin{pmatrix} 1 \\ 100 \end{pmatrix}$$

This market is acceptably complete and has no strictly acceptable opportunities. For this reason the matrix equation $\vec{p} = AB\vec{w}$ may be solved to obtain the weights \vec{w} which comprise the representative state pricing function. Using the method of singular value decomposition (described in Flannery *et al* [12] and in Section 5.1.4) because the matrix AB is close to singular, the (normalised) weights vector

$$\vec{w} = \begin{pmatrix} 0.4680 \\ 0.5320 \end{pmatrix}$$

is obtained.

These are the weights to be placed on the two valuation measures. The resulting linear combination of measures may be used as a “risk-neutral” measure in this environment. Prices of contingent claims may be calculated by calculating the expected value of the payoff of the claim, using the representative state pricing function. It is necessary to discount the resulting expected values to reflect the fact that the weights have been normalised.

To illustrate how this framework is used to price contingent claims, the price is calculated for a European call option on the stock, with strike 98, expiring at time $t = 1$. The option has payoff $(0, 2, 7)$ in the three states of the world.

Under the measure \mathbb{P}_1 , the option payoff has an expected value of 0.9, and under \mathbb{P}_2 it has an expected value of $\frac{37}{7}$ or 5.2857. These expected values must be weighted by the weights \bar{w} , added together, and discounted to get the call value. The call option's value is

$$0.9048(0.4680 \times 0.9 + 0.5320 \times 5.2857) = 2.9292$$

The value of the call is thus 2.9292. The value of the corresponding European put option (with payoff $(8, 0, 0)$ at expiry) may be calculated as 3.2610.

A slightly larger market structure with more than two traded assets is required before meaningful calculations may be performed to find the optimal hedge portfolios which solve equations (3.22) and (3.21). If there are only two or three risky securities, then the likelihood of being able to perform any sort of hedge for a derivative, other than simply purchasing or shorting the underlying, is unlikely.

3.5 Discussion

The theory developed in Carr, Geman and Madan [7] is closely related to the theories of no-arbitrage and risk-neutral pricing. Section 4.1 shows that, through a transformation of the market model, all the no-arbitrage pricing machinery developed in Harrison and Kreps [17] and subsequent papers can be brought to bear on the problem of pricing contingent claims in a single-period incomplete market setting. Attempts to extend the finite-state, single-period model described in this chapter to encompass an infinite state space, multiple periods, and continuous time are described in Chapter 4. Chapter 5 shows that a quasi-Monte Carlo method may be used to pick test measures, making the model eminently suitable for computational simulations.

The model of Carr, Geman and Madan [7] encapsulates elements of no-arbitrage pricing as well as utility theory. In this latter respect it shifts the focus of expected utility maximisation from the individual to a representative group of market participants, without losing an awareness of the fact that differing views of market participants impact securities prices.

In addition to these attractive properties, real-world use of the structure is feasible. Artzner, Delbaen, Eber and Heath [3] and the Matif Exchange [26] describe a related system called the SPAN[®] system, used by, among others, the Matif exchange in Paris and the Chicago Mercantile Exchange, to calculate the margin payments required on portfolios of futures and other

instruments. The SPAN system is used to assess risk, or to determine which positions are “acceptable”. Pricing of instruments is not necessary for an exchange, as the market has determined the price of the instrument and all the exchange has to do is to protect itself against possible losses stemming from investor default.

The risk of a futures position is assessed using each of what is essentially a collection of sixteen probability measures, reflecting sixteen different scenarios. Examples of these scenarios are up or down moves of volatility combined with various movements of the futures price. Two of the scenarios reflect possible extreme up and down movements of the futures price. The maximum of the expected losses incurred under the first fourteen measures and 35% of the loss incurred using the two extreme scenarios gives the measure of the riskiness of the position. The futures position is acceptable if the account of the investor holding the position has a balance greater than or equal to the maximum expected loss. If necessary the margin account is topped up as a result of a margin call. This corresponds closely to the definition of a risk measure formulated by Artzner, Delbaen, Eber and Heath [3], and outlined in Chapter 2.

If \mathcal{Q} denotes the family of probability measures, then the risk measure ρ used by the SPAN system may be generally specified as

$$\rho(X) = \sup\{\mathbb{E}_{\mathbb{P}}(-X/B_T) : \mathbb{P} \in \mathcal{Q}\}$$

where X is the random variable representing the future value of a position or portfolio. This is a coherent risk measure as defined in Artzner, Delbaen, Eber and Heath [3].

When this measurement of the riskiness of a position is compared to the balance in the investor’s margin account, the methodology becomes closely analogous to the test for acceptability of an opportunity of Definition 3.1.2. Similar systems are used by nearly twenty other exchanges around the world to assess the risk of positions held by market participants.

Chapter 4

Extensions and Modifications

Carr, Geman and Madan [7] formulates the valuation model in a single-period economy with a finite number of states of the world, and references are made in the paper to extensions of the model to an infinite state space. Heath and Ku [19] provides the beginning of a continuous-time treatment of the model, formulating the notion of a risk-measure based on a number of probability measures. The problem of pricing is not addressed, however, and it proves to be a nontrivial matter to define a representative state pricing function for multiple time periods and continuous time.

This chapter describes an alternative approach to the single-period model, which makes available much of the machinery developed for dealing with complete securities market models. A discussion of the difficulties inherent in defining multiple-period and continuous-time representative state pricing functions is followed by a section dealing with the results of Heath and Ku [19].

4.1 An alternative approach

The framework which has been detailed in Chapter 3 is, in many ways, directly analogous to the familiar risk-neutral environment of the complete market. This section demonstrates that this analogy carries through much further than may at first be apparent¹. By transforming an incomplete market, equipped with the valuation and stress test measures of the model detailed in Carr, Geman and Madan [7], it is possible to obtain a complete market in which the probability measures play the roles of states of the world, and the representative state pricing function is the unique risk-neutral measure.

It should be stressed that this transformation does not represent anything new mathematically, but that it provides a different perspective from

¹Peter Ouwehand is responsible for this insight.

which the same calculations as before are performed, but with a different meaning. Moreover, by considering the market as if it is a complete, multiple-period market, powerful and well-known results from classical complete-market asset pricing theory are available in a framework which was originally unfamiliar.

Recall that the securities market model described in Chapter 3 has the following elements:

- Assets B_t, S_t^1, \dots, S_t^N ,
- States $\omega_1, \omega_2, \dots, \omega_K$,
- Test measures $\mathbb{P}_1, \mathbb{P}_2, \dots, \mathbb{P}_M$ of which $\mathbb{P}_1, \dots, \mathbb{P}_E$ are valuation measures, and $\mathbb{P}_{E+1}, \dots, \mathbb{P}_M$ are stress test measures.

Let \mathcal{E} denote this economy. Let $\hat{\mathcal{E}}$ be the economy transformed to have the following structure:

- Assets B_t, S_t^1, \dots, S_t^N ,
- “States” $\mathbb{P}_1, \dots, \mathbb{P}_E$

An asset with payoff S_T^n across states, has value $S_T^n(\mathbb{P}_m) = \mathbb{E}_{\mathbb{P}_m}(S_T^n)$ in state \mathbb{P}_m . If B^v denotes the matrix B of earlier sections, but with only the columns corresponding to valuation measures, then

$$S_T^n(\mathbb{P}_m) = \mathbb{E}_{\mathbb{P}_m}(S_T^n) = (AB^v)_{nm} \quad (4.1)$$

This economy is complete if $E \leq N$.

The economy $\hat{\mathcal{E}}$ has the form of the usual single period securities market model. Denote by D the matrix of final asset values which usually characterises such an economy, where because of equation (4.1),

$$D = \begin{pmatrix} \mathbb{E}_{\mathbb{P}_1}(B_T) & \mathbb{E}_{\mathbb{P}_2}(B_T) & \dots & \mathbb{E}_{\mathbb{P}_E}(B_T) \\ \mathbb{E}_{\mathbb{P}_1}(S_T^1) & \mathbb{E}_{\mathbb{P}_2}(S_T^1) & \dots & \mathbb{E}_{\mathbb{P}_E}(S_T^1) \\ \dots & \dots & \dots & \dots \\ \mathbb{E}_{\mathbb{P}_1}(S_T^N) & \mathbb{E}_{\mathbb{P}_2}(S_T^N) & \dots & \mathbb{E}_{\mathbb{P}_E}(S_T^N) \end{pmatrix}$$

and denote by \vec{p} the vector of initial asset prices B_0, S_0^1, \dots, S_0^N . A state pricing function $\vec{\pi}$ (Definition 3.2.1) for this economy may be found by solving the equation $\vec{p} = D\vec{\pi}$. Moreover, there is a neat correspondence between \mathcal{E} and $\hat{\mathcal{E}}$:

Proposition 4.1.1. *There are no strictly acceptable opportunities in \mathcal{E} if and only if there are no arbitrage opportunities in $\hat{\mathcal{E}}$.*

Proof. (\Rightarrow) Suppose that there is an arbitrage opportunity, $\vec{\alpha}$, in $\hat{\mathcal{E}}$. That is, $\vec{\alpha}^T \vec{p} = 0$ ($\vec{\alpha}$ has zero cost) and $\vec{\alpha}^T D > 0$ ($\vec{\alpha}$ has non-negative payoff over all states and a strictly positive value in some state).

But $D = AB^v$, so $\vec{\alpha}^T AB^v > 0$. Since the matrix B^v is comprised of only the valuation measures in the economy \mathcal{E} , it follows that, for every \mathbb{P}_m , $1 \leq m \leq E$, the expected value of $\vec{\alpha}$, or $\mathbb{E}_{\mathbb{P}_m}(\vec{\alpha}^T A)$, is non-negative, and that for some n , with $1 \leq n \leq E$, $\mathbb{E}_{\mathbb{P}_n}(\vec{\alpha}^T A) > 0$.

To fulfil the criterion of strict acceptability completely, an opportunity must also be acceptable under not just the valuation measures, but also all the stress test measures, which remains to be shown.

A property of arbitrage opportunities is that they may be scaled arbitrarily and still remain arbitrages. Thus there exists a $\lambda \in \mathbb{R}^+$ such that

$$\mathbb{E}_{\mathbb{P}_m}(\lambda \vec{\alpha}^T A) \geq f_m$$

for $E + 1 \leq m \leq M$.

It follows that if $\vec{\alpha}$ is an arbitrage opportunity in $\hat{\mathcal{E}}$, a positive scalar multiple of $\vec{\alpha}$ is a strictly acceptable opportunity in \mathcal{E} .

(\Leftarrow) Now suppose that $\vec{\alpha}$ is a strictly acceptable opportunity in \mathcal{E} . Then

$$\mathbb{E}_{\mathbb{P}_m}(\vec{\alpha}^T A) \geq 0 \quad \text{for every } m, 1 \leq m \leq E \quad (4.4a)$$

$$\mathbb{E}_{\mathbb{P}_n}(\vec{\alpha}^T A) > 0 \quad \text{for some } n, 1 \leq n \leq E \quad (4.4b)$$

Expanding the expectations in equations (4.4a) and (4.4b), and writing S_T^0 for B_T yields

$$\sum_{j=0}^N \alpha_j \sum_{\omega \in \Omega} \mathbb{P}_m(\omega) S_T^j(\omega) \geq 0 \quad \text{for every } m, 1 \leq m \leq E \quad (4.5a)$$

$$\sum_{j=0}^N \alpha_j \sum_{\omega \in \Omega} \mathbb{P}_n(\omega) S_T^j(\omega) > 0 \quad \text{for some } n, 1 \leq n \leq E \quad (4.5b)$$

It follows that $\vec{\alpha}^T AB^v \geq 0$, with at least one nonzero element. Now $AB^v = D$ in $\hat{\mathcal{E}}$, so $\vec{\alpha}$ is an arbitrage opportunity in $\hat{\mathcal{E}}$. It follows that the absence of arbitrage opportunities in $\hat{\mathcal{E}}$ implies the absence of strictly acceptable opportunities in \mathcal{E} . \square

The incomplete economy \mathcal{E} (with no strictly acceptable opportunities) may thus be transformed into a complete market. Then the First Fundamental Theorem of Asset Pricing of Harrison and Kreps [17] may be used to find an equivalent martingale measure, and contingent claims may be priced in the complete market using this measure. It turns out that an equivalent martingale measure in the economy $\hat{\mathcal{E}}$ is a representative state pricing function in \mathcal{E} , making these prices useful in the original incomplete market.

This usefulness stems from the following argument. Suppose that there is the right number E of valuation measures to make the transformed market $\hat{\mathcal{E}}$ complete, $E \leq N$, and that the no-arbitrage condition holds. Then there is a unique risk-neutral measure \mathbb{Q} in $\hat{\mathcal{E}}$, denoting $\mathbb{Q}(\omega_k)$ by w_k . Then for every asset S^n the following equalities hold:

$$\begin{aligned}
 E_{\mathbb{Q}}(S_T^n) &= \sum_{m=1}^E \mathbb{Q}(\mathbb{P}_m) S^n(\mathbb{P}_m) \\
 &= \sum_{m=1}^E w_m \mathbb{E}_{\mathbb{P}_m}(S^n) \\
 &= \sum_{m=1}^E w_m \sum_{\omega \in \Omega} \mathbb{P}_m(\omega) S^n(\omega) \\
 &= \sum_{\omega \in \Omega} \left(\sum_{m=1}^E w_m \mathbb{P}_m(\omega) \right) S^n(\omega) \\
 &= \sum_{\omega \in \Omega} \pi(\omega) S^n(\omega)
 \end{aligned}$$

using equation (3.8), the definition of the representative state pricing function $\pi(\omega)$, to move from the second last line to the last line.

Discounting both sides of the equalities above by multiplying by $p_0 = B_0$ yields

$$\begin{aligned}
 B_0 E_{\mathbb{Q}}(S_T^n) &= B_0 \sum_{\omega \in \Omega} \pi(\omega) S^n(\omega) \\
 &= S_0^n
 \end{aligned}$$

which follows from the definition of the representative state pricing function, and from the comments on normalising the weights which appear at the end of Section 3.2.

The equivalence of the risk-neutral measure and the representative state pricing function stems from the fact that $\vec{p} = D\vec{\pi}$ is solved for the state pricing function $\vec{\pi}$ in the standard economy, but since $D = AB^v$, the equation in question is actually $\vec{p} = AB^v\vec{\pi}$. The solution to this equation is $\vec{\pi}$, the weights which define the representative state pricing function, and which have been called \vec{w} previously.

Intuitively this correspondence between the pricing functions is appealing. Probability measures play the role of states of the world in the transformed economy, and the representative state pricing function places weights on the valuation measures, indicating their importance in pricing. In an economy like $\hat{\mathcal{E}}$, the state pricing function $\vec{\pi}$ (from which the equivalent martingale measure is determined) derives its name from the “value” it assigns to each state in that π_j is the cost of a security that pays 1 in state ω_j and 0 in all other states.

All the risk-neutral pricing theory which has been derived for complete markets is now at our disposal in an incomplete market, provided there are sufficient valuation measures. Some of the theory may be circumvented to provide a quick, alternative proof of the First Fundamental Theorem, Theorem 3.2.1.

Theorem 4.1.1 (First Fundamental Theorem of Asset Pricing). *The market has no strictly acceptable opportunities if and only if there exists a representative state pricing function.*

Proof. Consider the securities market models \mathcal{E} and $\hat{\mathcal{E}}$ as before.

There are no strictly acceptable opportunities in \mathcal{E}

\Leftrightarrow There is no arbitrage in $\hat{\mathcal{E}}$ (by Proposition 4.1.1)

\Leftrightarrow There exists a risk-neutral measure for $\hat{\mathcal{E}}$, by classical results

\Leftrightarrow There exists a representative state pricing function for \mathcal{E} . \square

The relationship between \mathcal{E} and $\hat{\mathcal{E}}$ is deeper when the property of (acceptable) completeness holds.

Proposition 4.1.2. *$\hat{\mathcal{E}}$ is complete if and only if \mathcal{E} is acceptably complete.*

Proof. (\Rightarrow) Suppose that \mathcal{E} is not acceptably complete. Then $\text{rank}(AB) < M$. To obtain the matrix B^v some columns of B are removed, and hence some columns of AB . It follows that $\text{rank}(AB^v) < E$, so $\text{rank}(D) < E$. This implies that $\hat{\mathcal{E}}$ is not complete.

(\Leftarrow) Suppose that \mathcal{E} is acceptably complete. Then $\text{rank}(AB) = M$. Now AB is an $(N + 1) \times M$ matrix, so all its columns must be linearly independent. It follows that the rank of AB^v is E , since $M - E$ of these columns are removed to form the $N \times E$ matrix AB^v .

But in $\hat{\mathcal{E}}$, $D = AB^v$, so $\text{rank}(D) = E$. It follows that $\hat{\mathcal{E}}$ is complete. \square

This, in turn, allows a short-circuit proof of the Second Fundamental Theorem, Theorem 3.3.1.

Theorem 4.1.2 (The Second Fundamental Theorem of Asset Pricing). *The market is acceptably complete if and only if the representative state pricing function is unique.*

Proof. \mathcal{E} is acceptably complete and has no strictly acceptable opportunities

\Leftrightarrow $\hat{\mathcal{E}}$ is complete and arbitrage-free (by Proposition 4.1.2)

\Leftrightarrow There exists a unique risk-neutral measure for $\hat{\mathcal{E}}$, by classical results

\Leftrightarrow There exists a unique representative state pricing function for \mathcal{E} . \square

These results highlight the close relationship between classical arbitrage pricing theory, as developed in Harrison and Kreps [17], for example, and

the model proposed in Carr, Geman and Madan [7]. Since the representative state pricing function and the risk-neutral equivalent martingale measure in the original and transformed economies are exactly the same, prices may be calculated in the complete economy and will be applicable in the incomplete, original market.

4.2 Infinite state space

As the securities market model is refined to become a closer and closer approximation of a real-world financial market, it is natural to assume that both the number of states of the world and the number of time periods would increase. The focus of this section is on single-period market models with infinitely large state spaces. A continuum of eventual states of nature renders any discrete-time financial market with finitely many assets grossly incomplete.

Carr, Geman and Madan [7] mentions in passing how the two fundamental theorems of asset pricing would generalise to an infinite (continuous) state space setting. What follows is a brief treatment of the topic, working in the infinite probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a finite number of assets and valuation measures. Definitions are re-stated to conform to the new setting where necessary.

The first fundamental theorem generalises easily. Here, the matrix C plays the role that the matrix product AB does in the finite state space setting of Section 3.2.

Let each valuation measure \mathbb{P}_m on the space $(\Omega, \mathcal{F}, \mathbb{P})$ be defined by $d\mathbb{P}_m/d\mathbb{P}$, its Radon-Nikodym derivative with respect to the reference measure \mathbb{P} . These derivatives are \mathcal{F}_T -measurable and \mathbb{P} -integrable over Ω . Moreover, assume that the final asset prices $S_T^n(\omega)$ and B_T are \mathcal{F}_T -measurable and integrable with respect to each of the valuation test measures.

Define a matrix C , of the asset valuations with respect to each test measure, the analog of the quantity AB in a finite state space:

$$C_{n+1,m} = \mathbb{E}_{\mathbb{P}_m}(S_T^n) = \int S_T^n(\omega) \frac{d\mathbb{P}_m}{d\mathbb{P}} d\mathbb{P} \quad (4.6)$$

with the first row of C containing the expected value of the numeraire process under each test measure.

Let $\vec{p} = (B_0, S_0^1, \dots, S_0^N)^T$ be the vector of initial asset prices, as before.

Theorem 4.2.1 (First Fundamental Theorem of Asset Pricing). *The market has no strictly acceptable opportunities if and only if there exists a representative state pricing function.*

Proof. (\Rightarrow) Suppose that there are no strictly acceptable opportunities. Let $\vec{\alpha}$ be a portfolio of the traded assets such that $\vec{\alpha}^T \vec{p} = 0$. Then, since $\vec{\alpha}$

cannot be a strictly acceptable opportunity, it *cannot* be the case that

$$\bar{\alpha}^T C \geq 0$$

and that

$$\bar{\alpha}^T C \neq 0$$

By the classical First Theorem of Asset Pricing, since no strictly acceptable opportunities implies that there are no arbitrage opportunities, there must exist a $\bar{w} > 0$ such that

$$\bar{p} = C\bar{w}$$

or, equivalently, using the definition of the matrix C , such that

$$p_n = S_0^n = \int S_T^n(\omega) \left(\sum_{\mathbb{P}_m \in \mathcal{S}^v} w_m \frac{d\mathbb{P}_m}{d\mathbb{P}} \right) d\mathbb{P}$$

It follows that \bar{w} defines a representative state pricing function, in view of Equations (3.5) and (3.8).

(\Leftarrow) Suppose that there exists a representative state pricing function, so there exists a $\bar{w} > 0$ such that $\bar{p} = C\bar{w}$. Suppose that for some trading strategy $\bar{\alpha}$, $\bar{\alpha}^T C \geq 0$ and $\bar{\alpha}^T C \neq 0$. That is, the expected final return of the portfolio $\bar{\alpha}$ is non-zero and is positive under one of the test measures.

Then $\bar{\alpha}^T \bar{p} = \bar{\alpha}^T C\bar{w}$, the initial price of $\bar{\alpha}$, is positive and non-zero, since $\bar{w} > 0$ and $\bar{p} \neq 0$. So there are no strictly acceptable opportunities. \square

Recall from earlier chapters that when the number of assets in the marketplace weakly exceeds the number of valuation test measures (probability measures with an associated floor of 0), the representative state pricing function can be uniquely determined. Classical discrete-time asset pricing theory requires for a unique state pricing function that the number of *assets* must weakly exceed the number of states of the world. In this setting, that requirement is impossible to satisfy, but, as before, the generalisation of the Second Fundamental Theorem of Asset Pricing guarantees the uniqueness of the representative state pricing function if a less stringent condition is met, namely that of no strictly acceptable opportunities.

The uniqueness condition of Definition (3.3.2) is reformulated in the setting of a continuum of states of the world.

Definition 4.2.1. The test measures are said to satisfy the condition of *under-testing* if, for all test results $\vec{c} = (c_m : \mathbb{P}_m \in \mathcal{S}, 1 \leq m \leq M)$ there exists an \mathcal{F} -measurable random variable X such that

$$\int X(\omega) \frac{d\mathbb{Q}_M}{d\mathbb{P}} d\mathbb{P} = c_m \quad (4.10)$$

Recall that a set of test results is an M -vector comprising the expected values of a particular security under each test measure, as defined in Section 3.3. In continuous state space, the expected value of a security under a particular measure becomes the integral in equation (4.10). Moreover, the purely mathematical characterisation of acceptable completeness given by equation (3.17) can be reformulated to require the existence of $\vec{\alpha}$ such that

$$\vec{\alpha}^T C = \vec{c} + \vec{f}^T \quad (4.11)$$

in the continuous state space setting.

Theorem 4.2.2 (The Second Fundamental Theorem of Asset Pricing). *The market is acceptably complete if and only if the representative state pricing function is unique, provided that the under testing condition holds.*

Proof. (\Rightarrow) Suppose that the representative state pricing function is unique. Then there is a unique solution to the matrix equation $\vec{p} = C\vec{w}$. Hence from linear algebra it follows that $\ker(C) = \{0\}$, where C is the matrix of test results defined in equation (4.6). This means that $\text{ran}(C) = \mathbb{R}^M$, where M is the number of test measures.

It follows from this property of C that there is a unique solution $\vec{\alpha}$ to the matrix equation

$$\vec{\alpha}^T C = \vec{c} + \vec{f}^T \quad (4.12)$$

where C , \vec{c} and \vec{f} are as before. The solution $\vec{\alpha}$ represents a portfolio of the traded assets. But equation (4.12) is just the alternative characterisation of acceptable completeness as in equation (4.11). So if the representative state pricing function is unique, the market is acceptably complete.

(\Leftarrow) Now suppose that the market is acceptably complete. Then, since the asset space is under-tested, $\text{ran}(C) = \mathbb{R}^M$. Hence $\ker(C) = \{0\}$, and this in turn implies that the representative state pricing function is unique. \square

4.3 Multiple time periods

Given a set of trading dates $\tau = \{0 = t_0 < \dots < t_n = T\}$ and traded securities $\mathbf{S} = (B, S^1, \dots, S^N)$ with $B_0 = 1$, a traditional multiple-period securities market model may be constructed by concatenating a number of single-period market models to obtain a multinomial tree structure. The risk-free rate is assumed to be a constant r , so $B_{t_k} = e^{rt_k}$. A filtration $\mathcal{F} = \{\mathcal{F}_{t_k}\}_{t_k \in \tau}$ for the tree is implied in a natural way by the refinement of information about the state of the world as time passes.

The test measures may be completely specified on the multiple-period model. The time t_k probability of a state ω occurring at time t_{k+1} may be expressed using conditional probabilities based on the information contained in \mathcal{F}_k . In this way the test measures may be seen to operate on the underlying single-period models as well.

4.3.1 Traditional extension to multiple time periods

Owing to its similarity to the risk-neutral measure, it makes sense to define the representative state pricing function over multiple time periods in the same way as the single-period risk-neutral measure is extended from one to multiple periods. This may be done by calculating a representative state pricing function for each of the single-period models comprising the multiple-period market model, and (if they all exist) multiplying along the paths through the tree to find the multiple-period representative state pricing function.

Bearing in mind that a single-period market without strictly acceptable opportunities implies the existence of a representative state pricing function (Theorem 3.2.1), the following proposition provides reassurance that attempts to find a representative state pricing function for a multiple-period model may not be in vain.

Proposition 4.3.1. *There are no strictly acceptable opportunities in a multiple-period securities market model if and only if there are no strictly acceptable opportunities in each of its component single-period models.*

Proof. (\Rightarrow) Suppose that there exists a strictly acceptable opportunity X in one of the single period models which spans the time period $[t_k, t_{k+1})$ and which emanates from the state ω_0 . Then for every $\mathbb{P}_m \in \mathcal{S}$,

$$\mathbb{E}_{\mathbb{P}_m}(X|\mathcal{F}_{t_k})(\omega_0) \geq f_m \quad (4.13)$$

and

$$\mathbb{E}_{\mathbb{P}_j}(X|\mathcal{F}_{t_k})(\omega_0) > 0 \quad (4.14)$$

for some $1 \leq j \leq M$.

This opportunity can be turned into a global strictly acceptable opportunity Y . The opportunity Y can be exercised by doing nothing until time t_k , and then purchasing $e^{-r(T-t_{k+1})}$ units of the opportunity X . At t_{k+1} in state ω , take the payoff $X(\omega)e^{-r(T-t_{k+1})}$ from X and invest it in the riskless security if it is positive, or borrow the amount $X(\omega)e^{-r(T-t_{k+1})}$ if it is negative. At time T , the payoff from Y in all the states emanating from each time t_{k+1} -state ω will be $X(\omega)$.

Since equations (4.13) and (4.14) hold for X , and for every $\mathbb{P}_m \in \mathcal{S}$

$$\mathbb{E}_{\mathbb{P}_m}(Y) = \mathbb{E}_{\mathbb{P}_m}(X|\mathcal{F}_{t_k})$$

it follows that the payoff from Y satisfies the criteria for Y to be strictly acceptable. It follows that if there are no strictly acceptable opportunities in the multi-period model, there are no strictly acceptable opportunities in any of the single-period models.

(\Leftarrow) Now suppose that there are no strictly acceptable opportunities in any of the single-period models, but that Y is a strictly acceptable opportunity available over the period $[0, T]$. Then for every $\mathbb{P}_m \in \mathcal{S}$,

$$\mathbb{E}_{\mathbb{P}_m}(Y) \geq f_m$$

and

$$\mathbb{E}_{\mathbb{P}_j}(Y) > 0$$

for some $1 \leq j \leq M$.

But Y cannot be a strictly acceptable opportunity in any of the single-period models, so for every $\mathbb{P}_m \in \mathcal{S}$ and for every $t_k \in \tau$,

$$\mathbb{E}_{\mathbb{P}_m}(Y|\mathcal{F}_{t_k}) < f_m$$

Using the properties of conditional expectation (Williams [29]), it holds that

$$\mathbb{E}_{\mathbb{P}_m}(Y) = \mathbb{E}_{\mathbb{P}_m}(\mathbb{E}_{\mathbb{P}_m}(Y|\mathcal{F}_{t_k}))$$

for every $\mathbb{P}_m \in \mathcal{S}$ and for every $t_k \in \tau$. But $\mathbb{E}_{\mathbb{P}_m}(Y|\mathcal{F}_{t_k}) < f_m$, so for every $\mathbb{P}_m \in \mathcal{S}$, it must be true that $\mathbb{E}_{\mathbb{P}_m}(Y) < f_m$, a contradiction. So if there are no strictly acceptable opportunities in any of the single-period models, there can be no strictly acceptable opportunities in the multiple-period model. \square

The multiplicative solution for extending the representative state function to apply to markets with multiple trading dates is an intuitively natural one, but a difficulty arises when the nature of the representative state pricing function, which in the single-period model is a linear combination of test measures, is considered.

It is impossible to express the multiple-period representative state pricing function as it is defined above as a convex combination of the test measures, with a single constant weights vector \vec{w} to define the representative state pricing function for all time. Multiplication across time periods introduces nonzero cross terms, and if internal states are considered the equations become inconsistent.

If there are n time periods, solving for the weights would involve a set of linear equations in \vec{w} for each state ω in the tree. The time at which ω occurs will determine how many equations in \vec{w} there are at that time-step.

$$\left(\sum_m w_m^1 \mathbb{P}_m^1(\omega)\right) \dots \left(\sum_m w_m^{n,j} \mathbb{P}_m^n(\omega)\right) = \sum_m w_m \mathbb{P}_m(\omega) \quad (4.20a)$$

$$\left(\sum_m w_m^1 \mathbb{P}_m^1(\omega)\right) \dots \left(\sum_m w_m^{n-1,j} \mathbb{P}_m^{n-1,j}(\omega)\right) = \sum_m w_m \mathbb{P}_m(\omega) \quad (4.20b)$$

... = ...

$$\left(\sum_m w_m^k \mathbb{P}_m^1(\omega)\right) = \sum_m w_m \mathbb{P}_m(\omega) \quad (4.20c)$$

For example, there will be as many of equation (4.20a) (each in a different ω) as there are terminal states of the world. Similarly, there will be as many of equation (4.20c) as there are states of the world emanating from the initial state at $t = 0$.

Equations (4.20a) and (4.20c) will contradict each other, as equation (4.20c) implies that the weights vector \bar{w} is simply the weights \bar{w}^1 for the root single-period model situated at $t = 0$, whereas equation (4.20a) indicates otherwise. The system above is thus likely to be not only over-determined, but blatantly inconsistent. It is clear that the multiple-period representative state pricing function cannot, in general, be expressed as a convex linear combination of the test measures like its single-period counterpart.

4.3.2 Other approaches in acceptably complete markets

A variety of other approaches to defining the multiple-period representative state pricing function were considered. Solutions which focused on a new way to define the weights \bar{w} in a way consistent with the single-period model, such as multiplying or averaging the weights from each of the single-period models, were discarded. In all cases the “representative state pricing functions” defined in this way and then restricted to an underlying single-period model differed from the already-existing representative state pricing functions on the same single period model. This difference introduced strictly acceptable opportunities – and even arbitrage opportunities – to the single-period models, as it causes two different prices to exist for each opportunity.

Three choices thus exist with regard to the multiple-period representative state pricing function:

1. The naive definition which exactly parallels that of the multiple-period risk-neutral measure may be used to obtain prices which are consistent with the underlying single-period models. This representative state pricing function, however, cannot be expressed as a constant convex combination of the test measures on the multiple-period model, although its restriction on each single-period model admits such a representation.
2. Persisting in attempts to define the multiple-period representative state pricing function as a convex combination of the test measures leads to the possibility of allowing the weights to vary with time. The multiple-period representative state pricing function could then be defined as the product of the single-period representative state pricing functions as above, using the weights for the current single-period model as the weights at each time step.

An economic argument may be produced both for and against this suggestion: on one hand, the importance of a particular test measure

may change as time passes and the progression of states of the world becomes apparent. On the other, since the test measures are fully specified at $t = 0$, a decision about how important each one is in pricing should be possible at the start.

Allowing the weights to vary with time changes the representative state pricing function from a “probability measure” to a function of time and state: a random variable. It would also indicate that the continuous-time representative state pricing function, if it exists, is a function of time. Further research could focus on whether this is a feasible approach.

3. A completely different, more sophisticated approach to the multiple-period representative state pricing function may be possible, which somehow encapsulates the notion of weighting the test measures but remains consistent with the existing single-period structure.

4.3.3 Persistently incomplete markets

Surprisingly, the problem seems to be much easier when some or all of the underlying single-period markets are persistently incomplete. The bid-ask price may be calculated at time T , and at all other times by working backwards through the tree, replacing the payoff \bar{x} in equations (3.22) and (3.21) with the bid or ask price at the previous time-step.

4.4 Continuous-time as a limit of discrete-time

The assumption that there are a fixed, finite number of trading dates in the economy makes the set Ω of possible states of the world seem impoverished in comparison to the real world. Relaxing the assumption of a finite number of trading dates leads, in the limit, to the case of continuous time. While this is not a natural assumption either (exchanges do close for weekends, overnight and on holidays), it is perhaps a closer approximation of reality than the extreme of finitely many trading times.

This section suggests that a discrete-time extension of the model of Carr, Geman and Madan [7] may be used to form a sequence of successive approximations to the continuous-time model with progressively smaller time intervals. The suggested approach uses the theory of He [18] and Bingham and Kiesel [5], which shows that if the sequence of discrete-time approximations is properly constructed, contingent claim prices converge to their correct continuous-time values. Unanswered questions remain, however, and should be the subject of future research.

The complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is equipped with a filtration $\mathcal{F} = (\mathcal{F}_t)_{t \geq 0}$, which is the natural right-continuous filtration generated by the d -dimensional Brownian motion $\{\mathbf{B}_t\}_{t \geq 0}$. In a continuous-time securities

market model \mathcal{M} , the price processes of both the risky and riskless assets are specified by stochastic differential equations.

Keeping the notation conventions established in Section 3.1.1, the risky assets at time t are denoted by $S_t^1, S_t^2, \dots, S_t^N$. The numeraire process, denoted B_t , is described by the stochastic differential equation

$$\begin{aligned} dB_t &= B_t r(\mathbf{S}_t, t) dt \\ B_0 &= 1 \end{aligned}$$

whereas the lognormal price process of the risky assets at time t is modelled by the stochastic differential equation

$$d\mathbf{S}_t = \mu(\mathbf{S}_t, t) dt + \sigma(\mathbf{S}_t, t) dB_t$$

with $\mathbf{S}_t = (S_t^1, S_t^2, \dots, S_t^N)^T$.

The d independent Brownian motions model the systematic risks present in the market. The element $\sigma_{ij}(t)$ of the volatility matrix σ may be interpreted as the degree to which the j^{th} source of randomness, that is, $\{B_\tau^j\}_{\tau \geq 0}$, affects the price of the i^{th} risky asset at the instant t in time.

This market model is complete when the number of risky assets equals the number of sources of uncertainty in the model, which is true when $N = d$. We shall assume that the number of sources of uncertainty exceeds the number of risky assets, that is, that $N < d$, so the market is incomplete.

The superscript (n) will be used globally to indicate that the object in question belongs to $\mathcal{M}^{(n)}$, the n th approximation – with n time steps – to the continuous-time model \mathcal{M} .

4.4.1 Asset price processes

He [18] describes how the diffusion processes of the stock prices and numeraire process may be approximated by discrete-time processes which converge weakly to the original processes. The increments of the d independent Brownian motions may be approximated at the k th time step by d uncorrelated random variables $(\varepsilon_1^k, \varepsilon_2^k, \dots, \varepsilon_d^k)$ which need *not* be independent.

The sample paths of the discrete-time processes $\mathbf{S}^{(n)}$ and $B^{(n)}$ constructed by this method are piecewise constant and right-continuous with left limits. The process $\mathbf{S}^{(n)}$ is a $(d+1)$ -nomial one, meaning that $d+1$ states of the world emanate from each non-terminal state. Jumps in the value of $\mathbf{S}^{(n)}$ and $B_t^{(n)}$ occur at times given by a mesh $\tau^{(n)} = \{0 = t_0^{(n)} < t_1^{(n)} < \dots < t_n^{(n)} = T\}$, where

$$\lim_{n \rightarrow \infty} \sup_k |t_k^{(n)} - t_{k-1}^{(n)}| = 0$$

and $\tau = [0, T]$. The process $\mathbf{S}^{(n)} = \{\mathbf{S}_t^{(n)} : t \in \tau\}$ is constructed to be a finite Markov chain. This reflects the efficiency of the market: all information about the past values of \mathbf{S}_t is incorporated into the current stock prices.

The n th discretised process approximating the price process of the N risky assets at time t_{k+1} is a solution to the stochastic difference equation

$$\mathbf{S}_{t_{k+1}}^{(n)} = \mathbf{S}_{t_k}^{(n)} + \frac{\boldsymbol{\mu}(\mathbf{S}_{t_k}, t_k)}{n} + \frac{\boldsymbol{\sigma}_{ij}(\mathbf{S}_{t_k}, t_k)}{\sqrt{n}} \cdot \bar{\boldsymbol{\varepsilon}}^k$$

and the discretised numeraire process B_t solves the difference equation

$$B_{t_{k+1}}^{(n)} = B_{t_k}^{(n)} \left(1 + \frac{r(\mathbf{S}_{t_k}, t_k)}{n} \right) \quad (4.24)$$

at time t_{k+1} . From these difference equations it is evident that the random vector, $\bar{\boldsymbol{\varepsilon}}^k$, approximates the random increments of the Brownian motion from the k th time step to the $(k+1)$ st time step.

The construction of the processes $\mathbf{S}^{(n)}$ and $B^{(n)}$ guarantees that they converge weakly (Williams [29]) to the original diffusion processes \mathbf{S} and B .

Definition 4.4.1. A sequence of random variables $\{X_n\}_{n \in \mathbb{N}}$ on the space (Ω, \mathcal{F}) is said to *converge weakly* to the random variable X , defined on the probability space (Ω, \mathcal{F}) if, for all bounded, continuous functions f on Ω ,

$$\mathbb{E}(h(X_n)) \rightarrow \mathbb{E}(h(X))$$

4.4.2 Pricing contingent claims

The usefulness of the discretised asset price processes depends upon whether the test measures may be viewed as the weak limits of a sequence of probability measures which operate on the multiple-period market models, and on the convergence properties of the resulting representative state pricing functions. Both of these requirements are nontrivial, and it is not immediately obvious whether they can be satisfied. Weak convergence of probability measures is defined similarly (Shiryayev [28]) to weak convergence of random variables.

Definition 4.4.2. A sequence of probability measures $\{\mu_n\}_{n \in \mathbb{N}}$ is said to *converge weakly* to the probability measure μ , defined on the probability space (Ω, \mathcal{F}) if, for all bounded, continuous functions f on Ω ,

$$\mathbb{E}_{\mu_n}(f) \rightarrow \mathbb{E}_{\mu}(f)$$

The difficulty of restricting the test measures to the tree structure of the multiple trading date model stems from the fact that the events in the approximating multiple-period models correspond to sets of measure zero (for example, the event " $S^1 = 100$ ") under \mathbb{P} in the original space $(\Omega, \mathcal{F}, \mathbb{P})$. Since the test measures are absolutely continuous with respect to \mathbb{P} , these sets have measure zero under each test measure as well. It is thus insufficient to restrict the test measures to the discrete-time models by applying them

directly to the events in each discrete-time model, without augmenting these sets in some way. A similar problem is faced with respect to defining the reference measure \mathbb{P} on the multinomial trees $\{\mathcal{M}^{(n)}\}_{n \in \mathbb{N}}$.

Once the test measures are defined for each element $\mathcal{M}^{(n)}$ of the sequence discrete-time models, and assuming that each $\mathcal{M}^{(n)}$ is acceptably complete, a representative state pricing function $\mathbb{Q}^{(n)}$ may be calculated for each $\mathcal{M}^{(n)}$, for each $n \in \mathbb{N}$. Ideally the Radon-Nikodym derivatives $L^{(n)} = d\mathbb{Q}^{(n)}/d\mathbb{P}^{(n)}$ of the representative state pricing functions with respect to the discrete-time reference measures would converge weakly as the number of time steps is increased. This does not seem unreasonable, if the restricted test and reference measures are also converging weakly to the original measures on (Ω, \mathcal{F}) .

If we suppose that the sequence $\mathcal{M}^{(n)} = (\Omega^{(n)}, \mathcal{F}^{(n)}, \mathbb{P}^{(n)}, \mathbf{S}^{(n)}, B^{(n)}, L^{(n)})$ converges weakly to $\mathcal{M} = (\Omega, \mathcal{F}, \mathbb{P}, \mathbf{S}, B, L)$, where L is an \mathcal{F}_T -measurable random variable, then the risk-neutral valuation formula (using the representative state pricing function) which is used to value contingent claims guarantees the following result, proved (for a complete market with a weakly converging sequence of risk-neutral measures) in Bingham and Kiesel [5].

Proposition 4.4.1. *Suppose that the sequence*

$$M^{(n)} = (\Omega^{(n)}, \mathcal{F}^{(n)}, \mathbb{P}^{(n)}, \mathbf{S}^{(n)}, B^{(n)}, L^{(n)})$$

of acceptably complete discrete-time securities market models converges weakly to $\mathcal{M} = (\Omega, \mathcal{F}, \mathbb{P}, \mathbf{S}, B, L)$. Suppose also that the densities $L^{(n)}(T)$ are uniformly integrable.

Then if the bounded functions $X = f(\mathbf{S})$ and $X^{(n)} = f(\mathbf{S}^{(n)})$ represent the payoffs of contingent claims in the continuous-time and discrete-time markets respectively, with time 0 prices Y_0 and $Y_0^{(n)}$, it follows that

$$Y_0 = \lim_{n \rightarrow \infty} Y_0^{(n)}$$

4.4.3 Discussion

The advantage of an approach which views continuous-time as a limiting case of the discrete-time treatment of Carr, Geman and Madan [7] model is that it is ideally suited to computational simulation. The continuous-time model may be approximated by a discrete-time model with as many time-steps as desired: if the accuracy is insufficient, the number of time-steps can be increased to provide a better approximation to the continuous-time model.

Although it is potentially computationally expensive, the method of successive approximations suggested in Section 4.4 could use the pricing framework for discrete-time models. Further work could concentrate on establishing whether this is indeed a viable approach, as it is contingent on a tractable

definition for the multiple-period representative state pricing function, and convergence properties of the weights \vec{w} and the test and reference measures.

4.5 Continuous time by direct extension

Current work by Heath and Ku [19] focuses on a representation of a risk measure in terms of a family of probability (test) measures, each with an associated scalar value providing an indication of the risk tolerance of the test measure. This risk measure is related to the representative state pricing function of Carr, Geman and Madan [7].

Heath and Ku [19] extends the ideas of Artzner, Delbaen, Eber and Heath [3] to define *coherent dynamic risk measures*. These are coherent risk measures as described in Chapter 2, with associated *acceptance sets* of risks (investment opportunities) which are acceptable. The axioms satisfied by these risk measures coincide with some of the axioms on risk measures, but there is an additional continuity requirement and no stipulation that acceptance sets form a positively homogeneous cone. This is in line with the properties of acceptable risks set forth in Carr, Geman and Madan [7].

While *dynamic* risk measures need not satisfy a subadditivity condition in general, Heath and Ku [19] requires that they are subadditive when a convex combination of opportunities is considered. That is, if ρ is a dynamic risk measure, X and Y are risks, and $0 \leq \lambda \leq 1$, then $\rho(\lambda X + (1 - \lambda)Y) \leq \lambda\rho(X) + (1 - \lambda)\rho(Y)$.

Heath and Ku [19] demonstrates that a coherent dynamic risk measure ρ may be represented using a family $\mathcal{Q} = \{\mathbb{P}_m : 1 \leq m \leq M\}$ of probability measures on Ω , each with an associated non-positive constant f_m , where $1 \leq m \leq M$, such that

$$\rho(X) = - \inf_{\mathbb{P}_m \in \mathcal{Q}} \{\mathbb{E}_{\mathbb{P}_m}(X/B_T) + f_m\}$$

and where B_T is the final value of the numeraire process. This is an extension of Proposition 4.1 of Artzner, Delbaen, Eber and Heath [3], which is a representation theorem for risk measures like the Chicago Mercantile Exchange's SPAN system, described in Section 3.5.

The theory is developed in a continuous time setting, and shows that expectation under each of the family of probability measures may be represented by expectation under another probability measure at a randomised stopping time. Combining these results a representation is obtained for coherent dynamic risk measures in terms of a family of probability measures, each with an associated stopping time and floor.

gramming problem must be solved to obtain the bid-ask spread for a derivative, if possible. The software used to perform the linear optimisation was *PCx* [2], a freely available interior-point predictor-corrector linear programming package which can interface with MATLAB scripts.

5.1.1 Generating the measures

An open issue in Carr, Geman and Madan [7] was that of choosing the stress test and valuation measures required by the model. Acting on a suggestion by one of the authors¹ of [7], measures were generated as random, uniformly distributed points in the unit hypercube with a vertex at the origin and another at $(1, 1, \dots, 1)$ in \mathbb{R}^K , where K is the number of possible states of the world. Each point is normalised to sum to unity, producing a probability measure.

The rationale behind this Monte Carlo-type strategy is that the test measures should “fill up” the space of all possible measures, representing as varied a collection of views of the world as possible.

A further suggestion put forward in Carr, Geman and Madan [7] is that measures be chosen using utility functions. The probability of a state occurring under a specific measure would be proportional to the marginal utility of that state under a particular utility function, given its time T wealth level. Since the quasi-Monte Carlo method described above lends itself to computer simulation it was favoured for the tests described here.

5.1.2 Classifying the measures

A further requirement of the model is that each measure be either a valuation or a stress test measure, and that it have an associated scalar value, or floor, which reflects its designation.

In deciding whether a point represents a valuation or a stress test measure, a geometric argument is used. Since valuation measures have associated floors of zero, in order for an opportunity to be acceptable it should not produce a negative expected return when evaluated using one of these measures.

Intuitively, one would hope that an opportunity evaluated using a measure which considers one state of the world, ω_0 say, far more likely than any of the others would not produce a negative expected return. This could indicate that the opportunity leads to a loss in the state of the world which is considered most probable, and should thus be rejected by a rational investor. (Of course it is also possible that the opportunity leads to significant losses in the other less probable states, and a gain in state ω_0 , but an argument can be made for rejecting such an opportunity as well.)

¹Email correspondence with Peter Carr

Chapter 5

Implementation and Testing

In order to begin to gauge the feasibility of computational use of the model described in Carr, Geman and Madan [7], computer programs were written implementing the model for a small-scale securities market model. A set of tests was performed, pricing vanilla European call and put options for a range of strikes and volatility estimates.

The scale of the tests was limited by the fact that the model becomes computationally expensive as the number of assets, test measures and states of the world is increased. The linear optimisation, in particular, which is necessary in a persistently incomplete market to find the bid-ask spread, requires considerable numerical accuracy to provide reasonable results. The trinomial model was chosen as the basis for testing, as it accommodates the Carr, Geman and Madan [7] pricing methodology without too much modification. The issue of how to define the multiple-period representative state pricing function is also neatly circumvented.

5.1 Implementation

MATLAB [22] scripts were written to simulate single- and multiple-period securities markets, with a set of assets described by their initial prices and final values in each one of a finite number of possible states of the world. The representative state pricing functions of the multiple-period model were created by multiplying the single-period representative state pricing functions over each time-step, as is usually done with multiple time-period risk-neutral measures. Justification for this choice of multiple-period representative state pricing function is found in Section 5.2.1.

For the sake of generality, code was also produced to calculate the risk-neutral measure which exists if the market is complete. The incomplete case, however, is more complex, and several possible sub-states of the market must be catered for so that derivatives may be priced.

In the event that the market is not acceptably complete, a linear pro-

Bearing in mind the preceding argument, if a point (before normalisation) is found to lie outside of the K -sphere with center $(0.5, 0.5, \dots, 0.5)$ and radius 0.5, and in one of K critical regions of the hypercube, its floor is set to zero and it is deemed to be a valuation measure. Otherwise its floor is set to

$$f_m = \frac{D - \frac{K-1}{\sqrt{K}}}{\frac{K-1}{\sqrt{K}}}$$

where D is the perpendicular distance from the point in question to the main diagonal of the hypercube, joining the origin to the point $(1, \dots, 1)$. The value $\frac{K-1}{\sqrt{K}}$ is the maximum distance between the diagonal of a K -dimensional hypercube and any of the points inside the unit hypercube.

The floor assigned to the measure thus provides an indication of how “far away” the measure is from one which assigns an equal probability to every state of the world. (The whole diagonal of the hypercube will collapse to a single measure which has this property.)

The critical region in which a point must lie in order to be a valuation measure comprises the K corners of the hypercube, bounded by the sides of the hypercube and the face of the sphere, in which one co-ordinate of the point (before normalisation) is greater than 0.5, and all the others are less than 0.5. In three dimensions, this corresponds to the three regions of the unit cube bounded by any two of the $x-y$, the $y-z$ and the $x-z$ planes, and the surface of the sphere with radius 0.5, centered at $(0.5, 0.5, 0.5)$. Figure 5.1 illustrates this.

5.1.3 Calibrating the floors

Note that the floors f_m chosen by the MATLAB script have the property that $0 \leq f_m < -1$. The model requires simply that floors be non-positive real numbers, and does not specify how large or small their values should be. The intention in Carr, Geman and Madan [7], however, is that the floor associated with a measure represent an absolute lower bound for expected loss under that measure, and that it should also curtail the scale of the opportunities an investor is prepared to accept.

An investor using a stress test measure with a floor of -50 would consider an opportunity with expected gains of -49 acceptable, but would reject one with expected gains of -51. The floor would also prevent the investor from taking advantage of a scalar multiple of an acceptable opportunity (such as investing twice over in one with an expected return of -26) if the expected losses are too great.

The probability measure provides an indication of how likely a hypothetical investor considers each state of the world, but in practice it proved to be infeasible to assign a floor of arbitrary size to a stress test measure, as

Figure 5.1: The dark shading indicates the regions from which test measures are drawn when working in three dimensions (i.e. there are three future states of the world)

it requires speculation about the individual investor's preferences, level of wealth, and risk-aversion.

Since the floors chosen by the MATLAB implementation of the model lay in the interval $(-1, 0]$, some of the role of the floors in determining what level of investment is acceptable under a particular stress test measure was lost. It was thus necessary to find a way of evaluating an investment, which may have an expected loss greater than the lowest possible floor and still be acceptable.

The solution was to "normalise" the opportunity's payoff across states to a number close in magnitude to 1, and to compare the expected value of this normalised payoff to the floor. The normalisation was done in proportion to the expected value of an opportunity under the measure. The larger the

absolute value of the opportunity's unmodified expected value, the larger one would require the normalised payoffs to be, so that the measure could still play a role in limiting the scale of investments.

If \vec{x} is the payoff of an investment opportunity or contingent claim, $\vec{x} = (x_1, \dots, x_K)$, and its price at time $t = 0$ is x_0 , the algorithm used to determine whether \vec{x} is acceptable under a measure \mathbb{P}_m is as follows:

- Calculate $q = \mathbb{E}_{\mathbb{P}_m}(\vec{x}) - x_0 e^{rT}$, the expected payoff of the opportunity, assuming that it was financed by borrowing from the bank at a risk-free rate r .
- “Normalise” q by dividing the each component by the q -norm of \vec{x} , $\|\vec{x}\|_q$, where q is as above and

$$\|x\|_q = \left(\sum_{i=1}^K |x_i|^q \right)^{1/q}$$

- Compare the quantity $q/\|x\|_q$ to the floor, f_m , associated with the measure \mathbb{P}_m .
- If the floor lies below this normalised expected return, the opportunity \vec{x} is acceptable, otherwise it is not.

5.1.4 Solving for the representative state pricing function

The representative state pricing function can be found by solving the matrix equation (3.20) for the weights \vec{w} . It is likely that the matrix AB will not be invertible, in which case the equation may be solved in a least-squares sense. Then the representative state pricing function would be an exact solution to the equation

$$B^T A^T \vec{p} = B^T A^T AB \vec{w} \quad (5.3)$$

In practice it turns out that the matrix $B^T A^T AB$ is also ill-conditioned. The method of singular value decomposition, as detailed in Flannery *et al* [12], is used to solve the system of equations prescribed by equation (5.3) to obtain a numerically useful solution. The matrix $B^T A^T AB$ is decomposed into a diagonal matrix W , and two matrices U and V , such that $U^T U = V^T V = \mathbb{I}$, the $K \times K$ identity matrix. These matrices have the property that

$$B^T A^T AB = UWV^T$$

Inverting the square matrix $B^T A^T AB$ becomes a fairly straightforward matter when it is noted that, since U and V are orthogonal, their inverses are their transposes. The matrix W can be inverted simply by inverting each non-zero element on its diagonal. The method of singular value decomposition involves replacing those elements on the diagonal of W which

are close to zero (and which would thus blow up when inverted) with zeros in the matrix W^{-1} . This amounts to throwing away those equations which would be very susceptible to round-off error, and prevents the solution from being corrupted by numerical error.

While the solution obtained by using this method is by no means exact, it is a more consistently stable approximation to the solution than that produced by using MATLAB's matrix inversion methods.

5.1.5 Scaling the programs

It should be noted that the ratio of the volume of the hypersphere to the volume of the hypercube in \mathbb{R}^K approaches 0 as $K \rightarrow \infty$ (albeit fairly slowly). This means that if measures are simply generated at random, the number of valuation measures identified using the geometric argument will dwindle as the number of possible states of the world is increased.

Since valuation measures are used in pricing and stress test measures are not, common sense requires that a certain number of the test measures must be valuation measures. For this reason it was necessary to fix the number of valuation and stress test measures generated by the program so that at least a minimum number of valuation measures could be produced.

It should also be noted that linear optimisation problems can be computationally intensive, and it was established that, while useful with a few assets, states of the world and test measures, the MATLAB scripts written to implement the model did not admit scaling of the state and asset spaces. Numerical error led to nearly random results for large systems, compounded by multiple approximations to the solutions to under-determined systems of linear equations.

5.2 Testing

5.2.1 The trinomial model

The binomial model for contingent claim pricing assumes that the price of a stock can move either up or down with a fixed probability, at certain predefined points in the future. The resulting price process can be modelled by a binomial tree. Since the securities market for the single period binomial model is assumed to have a stock – the risky asset, and a bond or bank account – the riskless asset, and two possible states of the world, the single period binomial model is always complete. It is possible to show that an arbitrage-free multiple period binomial tree model is complete if and only if its component single period models are all complete. The binomial model is thus always complete.

Another member of the family of tree pricing models is the trinomial tree method (described, for example, in Bingham and Kiesel [5]), with one

risky and one riskless asset and three future states of the world. It is perhaps the simplest example of an incomplete securities market model. It is clear that, for the market to be arbitrage-free, the stock should be able to offer returns both better and worse than the risk-free rate. While the binomial model provides meaningful prices, in that the prices represent the cost of replicating portfolios, trinomial trees find favour among practitioners because they provide an extra degree of freedom in parameter choice. An example of their usefulness is in modelling processes with parameters which depend on an underlying index.

Hull and White [21] shows that the trinomial tree approach is equivalent to the explicit finite difference method. A drawback of the explicit finite difference method is that the quantities which are analogues of the probabilities of an up, down or across movement of the stock price in the trinomial tree may be negative, depending on the choice of parameters for the model. This leads to the calculation of possibly incorrect option values. In addition to this difficulty, the derivation of the finite difference method relies upon a Black-Scholes-type equation which the derivative is assumed to satisfy.

Chapter 2 indicated that the assumptions made in deriving the Black-Scholes model (such as continuous trading, constant volatility, continuous price processes and perfectly divisible assets) have the effect of inflating the asset space and shrinking the state space, leading to a complete market model. The derivation of this equation may also rely on synthesising the payoff of the contingent claim by using a replicating portfolio. In an incomplete market, replication is likely to be impossible. Therefore it does not necessarily hold that the derivative satisfies a Black-Scholes type equation, and the finite difference method cannot, strictly speaking, be used.

Using a version of the model of Carr, Geman and Madan [7], however, allows one to guarantee that the probabilities describing the movement of the stock price will be positive: these probabilities will be given by the representative state pricing function $B\bar{w}$ if the market is acceptably complete. Contingent claims may then be priced as if the representative state pricing function is just the unique risk-neutral measure of the binomial tree model, by calculating the discounted expected values of the claim's payoff backwards through the tree. If the market is not acceptably complete, linear optimisation may yield the claim's bid-ask spread at each node of the tree.

This adaptation of the trinomial model does *not* require a definition of a multiple-period pricing function (risk-neutral or representative). The three probabilities required to govern the movement of stock prices may be specified at the start, and are the same over each single-period model. These probabilities are selected by generating valuation measures and calculating a representative state pricing function for the first single-period model. Since the stock price movements are a fixed percentage of their previous levels, the representative state pricing function is the same over all periods.

5.2.2 Calibration of parameters

The binomial model is defined by the multiplicative factors describing the movement of the stock price, from S_0 at time $t = 0$ to S_0u and S_0d , for example, and by the probabilities of these movements. Because there are three future states of the world emanating from any single node of a trinomial tree, a third movement, S_0m , such that $S_0d < S_0m < S_0u$, is also possible. Some models choose m so that both m and u represent an increase in the stock price; others choose m to be a second downward move. Because the no-arbitrage condition for a trinomial lattice requires that $d < e^{r\Delta t} < u$, it is natural to choose m to be growth at the risk-free rate. In the tree of *discounted* stock prices, then, the price at the middle node emanating from a state would be unchanged from the source state.

The values for u , m and d which were used are:

$$u = e^{\sigma\sqrt{3\Delta t}} \quad (5.5a)$$

$$m = e^{r\Delta t} \quad (5.5b)$$

$$d = e^{-\sigma\sqrt{3\Delta t}} = \frac{1}{u} \quad (5.5c)$$

Note from equations (5.5a) and (5.5c) that an estimate of the volatility of the asset price is required. There is a single risky asset S in the economy. The risk-free rate is assumed to be a constant 10%.

Traditional trinomial tree models choose the probabilities of the various price movements to depend on the risk-free rate and the volatility of the asset. In this case, these probabilities are provided by the Carr model (in the form of a representative state pricing function) described in Chapter 3, where possible. If a unique representative state pricing function cannot be found, the bid-ask spread for contracts may be calculated using the linear programming methods described in Section 3.4.2.

5.2.3 Outline of the testing scheme

In the context of the trinomial model described above, two valuation measures give an acceptably complete economy and with three valuation measures the economy is not acceptably complete. No stress test measures were used. The tests were performed in an acceptably complete market model and in an incomplete market which was not acceptably complete. The contingent claims upon which the tests were performed were one-year European call and put options, with varying strike prices.

Acceptably complete market Pairs of valuation measures were generated. For each pair of valuation measures, attempts were made to value European call and put options with strike prices between 95 and 105. For each strike price, volatilities $5\% \leq \sigma \leq 15\%$ were used as

inputs to the trinomial model (the volatility of the stock affects the size of the up and down movements of its price in the trinomial tree). A three period trinomial tree was used for these tests.

In order to compare results with a varying number of time steps, another pair of valuation measures was generated. The two European options were priced using this set of measures, in trees with between 1 and 16 periods.

Persistently incomplete market The same set of tests was performed, but three valuation measures instead of two were generated at each step. The resulting market is incomplete and not acceptably complete. With only one risky asset, however, most of the time it is not possible to construct optimal hedge portfolios from which to calculate the bid and offer prices. The linear optimisation to calculate the bid-ask spread is thus likely to fail in this instance, depending on the weights assigned by the test measures to the future states of the world.

5.3 Results

The tests performed were not exhaustive, but some general conclusions may be drawn from the results obtained. Although different sets of test measures lead to different representative state pricing functions, prices did not differ widely between each run of tests. If a batch of test measures are in agreement that a particular state is very unlikely, derivatives which pay a significant amount in that state will have a lower than "expected" value.

In isolation it is impossible to state categorically that the values obtained for the European call and put options would be useful to a potential market participant. Future work using real market data and enough measures and risky assets to simulate a real financial market would be able to reach conclusions of this kind. Care should thus be taken when interpreting the test results, and in this section some general inferences will be drawn from the test results obtained in market models which are acceptably complete and persistently incomplete.

5.3.1 Acceptably complete market

The general characteristics of the results drawn from repeated runs of the MATLAB scripts can be summarised as follows:

1. The test results obtained with two test measures in the trinomial model indicate that the general properties of call and put option prices with respect to changing strike and volatility remain the same as in complete markets. Increasing the volatility of the underlying leads to higher option prices – this is expected, as the volatility estimate is used in

the calibration of the size of the up and down movements of the stock price. Similarly, call and put option prices decreased and increased as expected when the strike price was increased. These are basic “sanity checks” which indicate that both the model and the test programs are not fundamentally flawed.

2. A naive comparison² of the test results with the Black-Scholes prices (obtained using the closed-form solutions to the Black-Scholes equation for European call and put options) showed that the programs produced prices lying in a band around the Black-Scholes price. Most often, however, call options were “overvalued” and put options had a far smaller value than the Black-Scholes value predicted.

This trend is possibly an indication of the difficulties associated with successful hedging in an incomplete market. The trinomial model under consideration is seriously constrained, with a single risky asset and the numeraire providing the only raw materials for hedging.

3. Put option prices decreased and call option prices increased, or vice versa, as the number of trading dates was increased and all other parameters were held constant. This is because the acceptably complete market model used only two test measures and allowed three future states of the world. The small dimensions of the model led to the representative state pricing function placing a particular emphasis on one or two of the three future states. If, for example, a downward-movement of the stock price is considered far more likely than its counterparts, opportunities which have nonzero payoffs in the bottom regions of the tree (here, the put option) will increase in value as the number of terminal states (and hence opportunities for a positive payoff) are increased. This leads to option prices which either decay or grow as the number of trading dates is increased. Figure A.3 illustrates this behaviour.

The same set of tests was performed choosing measures as uniformly distributed random points in the unit hypercube, in case the geometric argument which distinguished between stress test and valuation measures described in Section 5.1.2 was to blame for this phenomenon. This did not solve the problem, however, and prices remained prone to gradual change as the number of trading dates was increased.

It is possible that this trend will not prevail in a multinomial tree structure with many more than three future states of the world, as a larger collection of test measures could be used to obtain a more balanced representative state pricing function.

²The reasons why this comparison is economically meaningless, and why the Black-Scholes equation does not hold in an incomplete market environment are detailed in Chapter 2 and Section 5.2.1.

5.3.2 Persistently incomplete market

In an incomplete market which is not acceptably complete, the values of optimal hedge portfolios provide the bid and ask prices for derivatives. Even though the trinomial model as it is considered here offers very limited opportunities for hedging, with its paucity of traded assets, these bid-ask spreads sometimes yielded very narrow bounds for the European option prices. Moreover, even in a market where there is no unique representative state pricing function, it is sometimes possible for the buyer and seller of an investment opportunity to hedge the opportunity so that it is just acceptable, resulting in a bid-ask spread of zero and a unique price for the opportunity.

Few other generalisations can be made about the test results obtained in persistently incomplete market models, save that the call and put option bid and ask prices displayed the same general characteristics (such as increasing as the volatility increases) as the unique option prices calculated in the acceptably complete market model.

University of Cape Town

Chapter 6

Conclusions and further work

Further work on the model proposed in Carr, Geman and Madan [7] could fall into the following broad areas:

1. An unambiguous definition of the multiple-period representative state pricing function is necessary before further work on the multiple trading date or continuous-time models is possible. This definition should be consistent with the single-period model as described in Carr, Geman and Madan [7], and retain the “representative” character of the representative state pricing function.
2. Further work on the development of the continuous-time model is necessary to flesh out the skeleton structure developed in Heath and Ku [19]. A comparison between the results obtained by modelling a continuous-time market as a sequence of discrete-time models, and those obtained using the coherent dynamic risk measures of Heath and Ku [19] and Artzner, Delbaen, Eber and Heath [3], would be informative. Particular attention could be paid to the computational viability of each approach, as well as a general result indicating how quickly the sequence of discrete-time models converges to provide useful pricing estimates.
3. Large-scale tests of the model, with a large number of test measures, should be performed to ascertain the economic viability of the prices obtained using a representative state pricing function. Considerable computing power is required to obtain meaningful results using this method in a large securities market model. Both numerical accuracy and speed suffer as the number of assets, time periods and states of the world are increased, particularly in markets which are not acceptably complete.

4. Research into alternative methods of choosing valuation and stress test measures may be used to determine the most suitable methods. The method described in Section 5.1.2 is only one of many possible routes to take with quasi-Monte Carlo methods, and none of the tests described here used utility theory to choose test measures, as suggested in Carr, Geman and Madan [7].
5. A comparison of this and other incomplete market pricing methodologies, such as those described in Cochrane and Saá-Requejo [9], Föllmer and Leukert [13] and [14], and El-Karoui and Quenez [10], would provide an indication of the comparative tightness of the bounds on contingent claim prices provided by this model in persistently incomplete markets. A study of the method in the context of the work done by Černy and Hodges [8] would establish how “attractive” the strictly acceptable opportunities of the model are, and how narrow its bid-ask spreads are in comparison to other no good-deal models.

Scrutiny of the model of Carr, Geman and Madan [7] leads to the conclusion that it is practical to use as a basis for computational studies of contingent claim values in incomplete markets. Its usefulness for determining real-world prices in incomplete markets depends upon the method whereby the representative state pricing function is defined for multiple periods and continuous time. The fact that related models are already used in practice, however (see, for example, the Matif SA handbook [26]), recommends it as an economically viable risk measure.

Appendix A

Some representative test results

This appendix contains graphs illustrating some representative results of the various test runs performed. Since each test run generated a new set of test measures reflecting investors' differing views of the future, prices differed slightly from run to run even though other market conditions (such as stock prices and volatility) remained the same. The graphs shown here are therefore a representative cross-section of the test results obtained.

A.1 Acceptably complete market

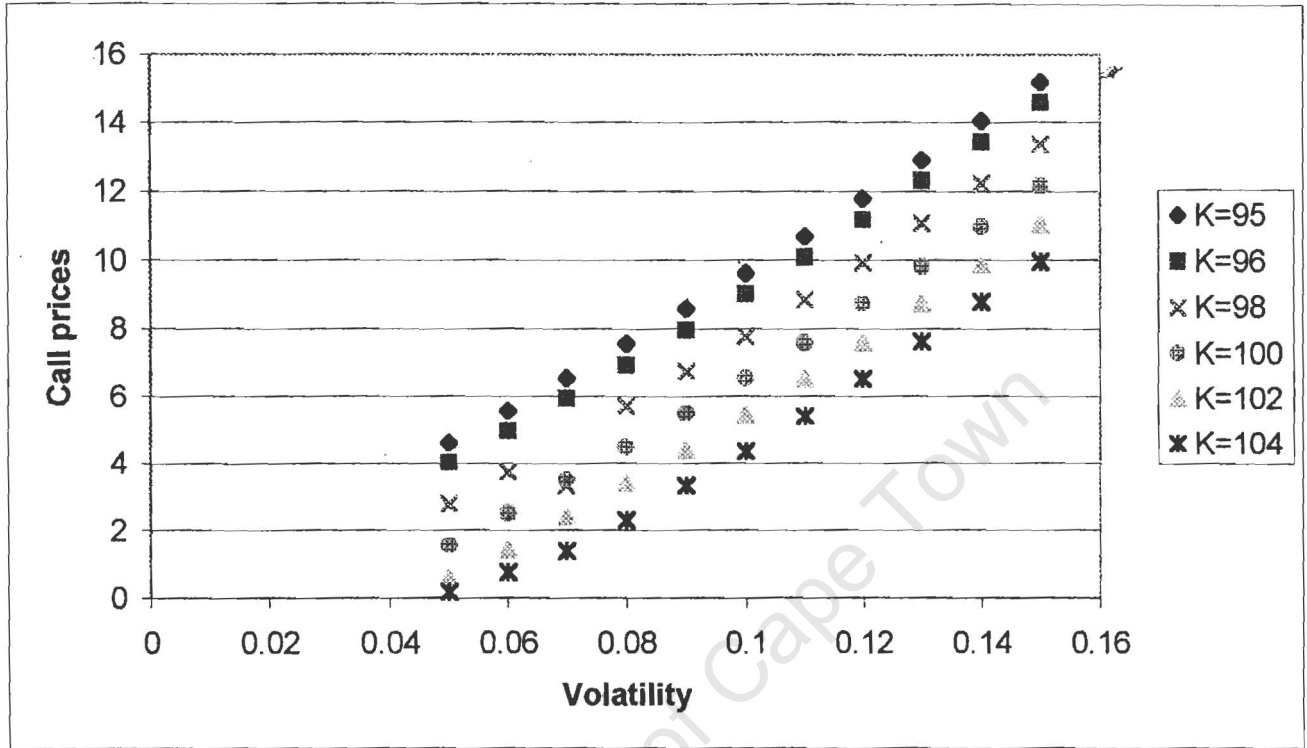


Figure A.1: Trinomial model prices for a European call option with varying strike on a stock with varying volatility

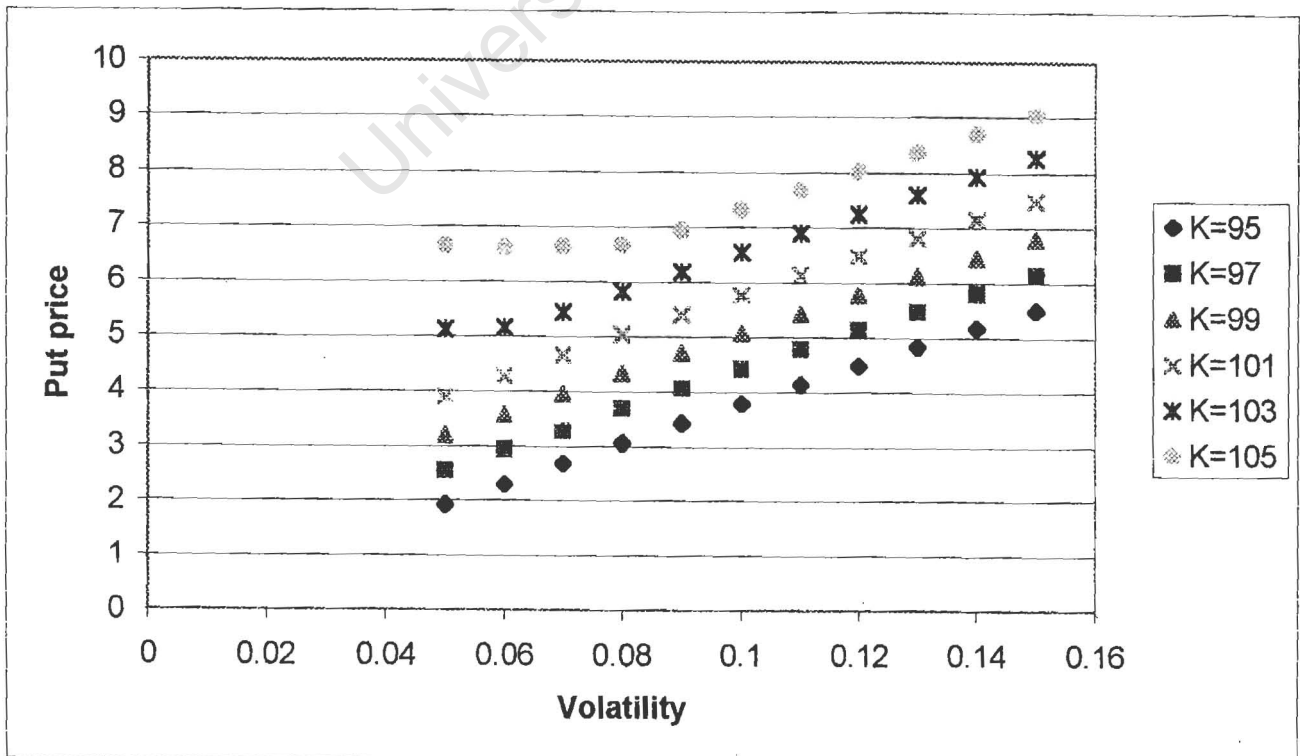
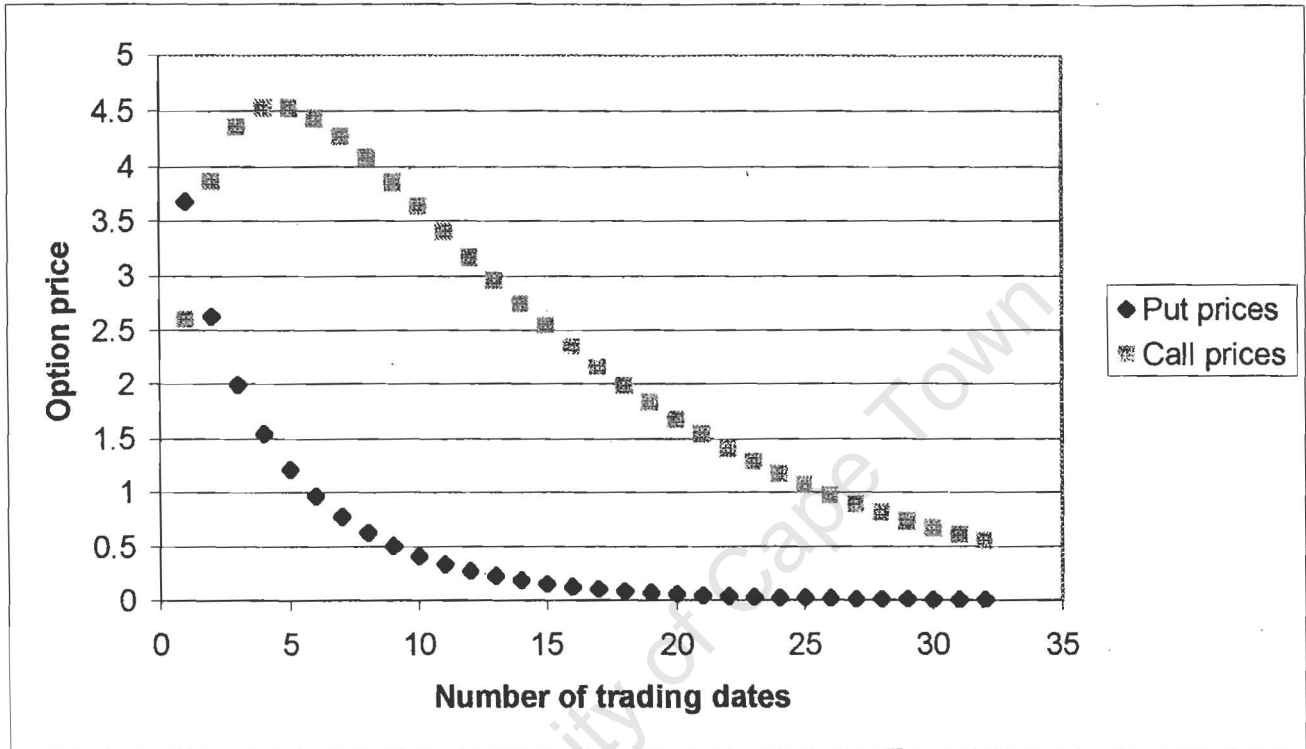


Figure A.2: Trinomial model prices for a European put option with varying strike on a stock with varying volatility

Figure A.3: European call and put option prices with strike 100, expiry 1, on a stock with initial price 100 and volatility 10%, priced using the trinomial tree method with increasing numbers of trading dates.



A.2 Persistently incomplete market




Figure A.4: The bid and ask prices for a European call option with strike 100 on a stock with volatility 10%

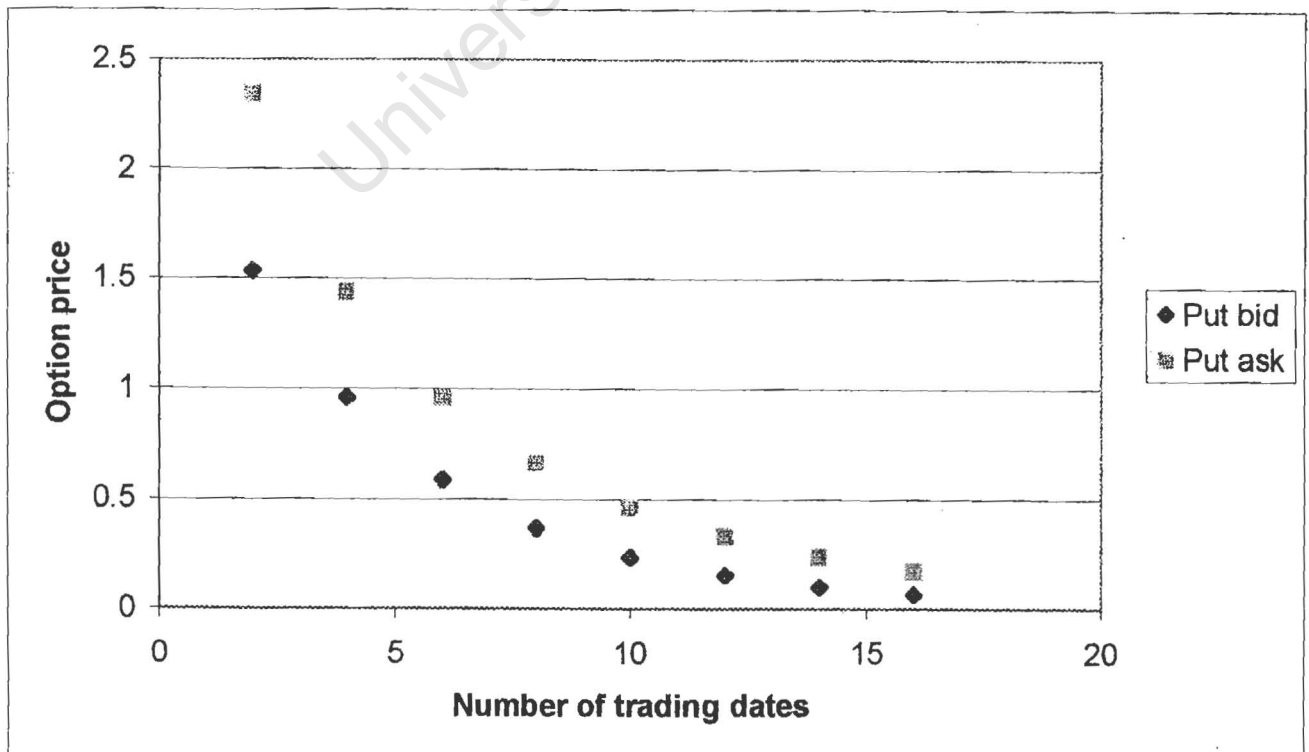


Figure A.5: The bid and ask prices for a European put option with strike 100 on a stock with volatility 10%

Appendix B

Program Listing

A full program listing of the MATLAB [22] code written to simulate the workings of the model follows.

B.1 Driver program: generate.m

```
clear; global A prices r T strike K sigma

r = 0.1;           % The risk-free rate
T = 1;            % The time period
strike = 100;
% Number of states of the world
K = 3; sigma = 0.1;

% These are the prices of the respective assets
load prices.dat;
% These are the opportunity values in different states
load A.dat;
A = a;
clear a;

% COMPLETE MARKET
if rank(A) == K
    disp('The market is complete.');
```

University of Cape Town

```
    complete;
% INCOMPLETE MARKET
else
    fprintf('The market is incomplete: rank(A)=%g ', rank(A) );
    fprintf('and there are %g states of the world\n\n', K);
    ac = incomplete(0, 2, 2);
    trinomial(ac, 2);
```

```
end
```

B.2 Complete market: complete.m

```
% The market is complete
global A prices

% We can find a state price vector
p = solve(A, prices); p = p/sum(p);
disp('Risk-neutral measure:')
disp(p);
```

B.3 Incomplete market: incomplete.m

```
function acc_complete = incomplete( M1, M2 )
% In the case of an incomplete market
% M1 is number of stress test measures to use
% M2 is number of valuation measures
global A B prices r T strike K w floors

M = M1+M2;
% The matrix to hold the measures
B = zeros(K, M);
% For the weights (if they exist)
w = zeros(M, 1);
floors = zeros(M, 1);
strictly = 0;
acceptable = 0;
acc_complete = 0;
rand('state', sum(100*clock));

% To keep track of the number of each kind of measure
Nval = 0;
Nstress = 0;
i = 1;
% The maximum value a floor could assume
maxfloor = (K-1)/sqrt(K);
while i <= M
    % Using uniformly distributed random numbers
    new = rand(K, 1);
    while any(new) <= eps
        new = rand(K, 1);
    end
    % while
```

```

meas_norm = norm(new, 2);
if (meas_norm > 1) & (any(new) > 0.5) & (sum(new < 0.5) == K - 1)
    & (Nval < M2)
    % Are we outside the sphere? (valuation measure)
    floors(i) = 0;
    Nval = Nval + 1;
    % Normalise the measure (not essential) and save it
    new = new/sum(new);
    B(:, i) = new;
    i = i+1;
elseif (Nstress < M1)
    % Or not? (stress test measure)
    % Projection onto the diagonal
    diag = sum(new)/K;
    distance = norm(new - diag, 2);
    floors(i) = (distance - maxfloor)/maxfloor;
    Nstress = Nstress + 1;
    % Normalise the measure (not essential) and save it
    new = new/sum(new);
    B(:, i) = new;
    i = i+1;
end
end % if-else
end % while

disp('Measures:');
disp(B);
disp('Floors:');
disp(floors);

% Are there any acceptable opportunities?
acceptable = zeros(size(prices)); for j = 1 : size(A, 1)
    for k = 1 : M
        opp = A(j, :);
        measure = B(:, k);
        % Run through each measure
        expect = opp*measure - prices(j);
        % "Normalise" it according to its expected value
        % The larger the expected value, the less we want to adjust it
        % We divide by the p-norm where p is expected return
        factor = norm(opp, abs(expect));
        opp = opp/factor;
        price = prices(j)/factor;
        % The normalised expected return
        expect = opp*measure - price;
    end
end

```

```

% Test for acceptability
if expect > floors(k)
    acceptable(k) = 1;
end % if
% Test for strict acceptability - compare to t=0 cashflow
% We use zero-cost portfolios: take a position and borrow from bank
if floors(k) == 0 & acceptable(j)
    & (opp-prices(j)*exp(r*T)) * measure > 0
    strictly = 1;
    disp('This is strictly acceptable:');
    disp(opp);
    disp('under this measure:');
    disp(measure);
    break;
end % if
end % for k
end % for j

% Modify A, B so they don't contain rows/cols for stress measures
% This is because the weights on the stress test measures are zero
% Need to do this so we can solve for the weights easily
if ~strictly
% There are no strictly acceptable opportunities so we can find w
% ACCEPTABLE COMPLETENESS
if rank(A*B) == M
    disp('The market is acceptably complete.');
```

$$Q = B' * A' * A * B;$$

```

    acc_complete = 1;
    if rank(B) >= M
        % The range of B is  $R^M$ , so under-testing holds
        disp('Under-testing holds.');
```

$$\text{statep} = B' * A' * \text{prices};$$

```

    w = solve(Q, statep);
    % Set stress test measures' weights explicitly to exactly 0
    % Solving doesn't always do this - it's approximate
    w = w .* (floors == 0);
    % Normalise weights
    w = w./sum(w);
    disp('Weights:');
    disp(w);
else
    disp('There are too many measures and too few assets: over-testing');
end
    % if-else (rank B)
% NOT ACCEPTABLY COMPLETE BUT EXISTS AN RSPF BECAUSE NSAO
else
    fprintf('The market is not acceptably complete:
        rank(AB) = %g < %g\n\n', rank(A*B), M);
    acc_complete = 0;
end
    % if-else (rank AB)
end
    % if no strictly acceptable opportunities

```

B.4 Initialisation: setuptrinomial.m

```

function y = setuptrinomial(P)
% Prepares for the trinomial model to be run next
global A B prices r T strike K w floors sigma

% Final number of states - leaf nodes (recombining)
final_states = 2*P+1;
dt = T/P;

% Up and down moves of stock price:
stock_tree = zeros(final_states, P+1);
u = exp(sigma*sqrt(3*dt));
d = 1/u;
m = exp(r*dt);
midpt = ceil(final_states/2);

A(2, :) = [prices(2)*d prices(2)*m prices(2)*u];

```

B.5 Trinomial tree: trinomial.m

```

function y = trinomial(acc_complete, P)
% Does a trinomial tree of P periods on the call and put option
global A B prices r T strike K w floors sigma

```

```

% The final number of states of the world - leaf nodes (recombining)
final_states = 2*P+1;
dt = T/P;

% Up and down moves of stock price:
stock_tree = zeros(final_states, P+1);
u = exp(sigma*sqrt(3*dt));
d = 1/u;
m = exp(r*dt);
midpt = ceil(final_states/2);
% Generate the asset price tree (using the second row of A)
% Set up the middle row (stays the same)
stock_tree(midpt, :) = prices(2);
% Which column we are in:
for j = 2 : P+1
    for i = 1 : j-1
        % This is how many ups we need
        stock_tree(midpt+i, j) = stock_tree(midpt+i-1, j-1)*u;
        % And downs
        stock_tree(midpt-i, j) = stock_tree(midpt-i+1, j-1)*d;
    end
    % One middle
    stock_tree(midpt, j) = stock_tree(midpt, j-1)*m;
end

if ~acc_complete
    % For each security at each time we have a bid-ask spread
    put_ask_tree = zeros(final_states, P+1);
    call_ask_tree = zeros(final_states, P+1);
    put_bid_tree = zeros(final_states, P+1);
    call_bid_tree = zeros(final_states, P+1);

    % Final payoffs of the options
    call_bid_tree(:, P+1) = call(stock_tree(:, P+1)', strike)';
    put_bid_tree(:, P+1) = put(stock_tree(:, P+1)', strike)';
    call_ask_tree(:, P+1) = call(stock_tree(:, P+1)', strike)';
    put_ask_tree(:, P+1) = put(stock_tree(:, P+1)', strike)';

    % Now discount for as many periods as necessary
    for k = P : -1 : 1
        % How many values above middle in that column:
        for m = 1 : k-1
            % The top half

```

```

call_ask_tree(midpt+m, k) = ask(B, r, dt,
    call_ask_tree(midpt+m-1:midpt+m+1, k+1)');
put_ask_tree(midpt+m, k) = ask(B, r, dt,
    put_ask_tree(midpt+m-1:midpt+m+1, k+1)');
call_bid_tree(midpt+m, k) = bid(B, r, dt,
    call_bid_tree(midpt+m-1:midpt+m+1, k+1)');
put_bid_tree(midpt+m, k) = bid(B, r, dt,
    put_bid_tree(midpt+m-1:midpt+m+1, k+1)');
% The bottom half
call_ask_tree(midpt-m, k) = ask(B, r, dt,
    call_ask_tree(midpt-m-1:midpt-m+1, k+1)');
put_ask_tree(midpt-m, k) = ask(B, r, dt,
    put_ask_tree(midpt-m-1:midpt-m+1, k+1)');
call_bid_tree(midpt-m, k) = bid(B, r, dt,
    call_bid_tree(midpt-m-1:midpt-m+1, k+1)');
put_bid_tree(midpt-m, k) = bid(B, r, dt,
    put_bid_tree(midpt-m-1:midpt-m+1, k+1)');
if any(any(call_ask_tree + call_bid_tree +
    put_bid_tree + put_ask_tree == inf))
    y = 0;
    disp('LP problem was infeasible - sorry!');
    return;
elseif any(any(call_ask_tree + call_bid_tree +
    put_bid_tree + put_ask_tree == nan))
    y = 0;
    disp('LP problem was infeasible - sorry!');
    return;
end
end
% The middle
call_ask_tree(midpt, k) = ask(B, r, dt,
    call_ask_tree(midpt-1:midpt+1, k+1)');
put_ask_tree(midpt, k) = ask(B, r, dt,
    put_ask_tree(midpt-1:midpt+1, k+1)');
call_bid_tree(midpt, k) = bid(B, r, dt,
    call_bid_tree(midpt-1:midpt+1, k+1)');
put_bid_tree(midpt, k) = bid(B, r, dt,
    put_bid_tree(midpt-1:midpt+1, k+1)');
end

else
    % We expect a unique price for each security
    put_tree = zeros(final_states, P+1);
    call_tree = zeros(final_states, P+1);

```

```

% Work out the probabilities - down, middle, up:
probs = B*w;
% Final payoffs of the options
call_tree(:, P+1) = call(stock_tree(:, P+1)', strike)';
put_tree(:, P+1) = put(stock_tree(:, P+1)', strike)';
% Now discount for as many periods as necessary
for k = P : -1 : 1
    % How many values above middle in that column:
    for m = 1 : k-1
        % The top half
        call_tree(midpt+m, k) = disc(B, w, r, dt,
            call_tree(midpt+m-1:midpt+m+1, k+1)');
        put_tree(midpt+m, k) = disc(B, w, r, dt,
            put_tree(midpt+m-1:midpt+m+1, k+1)');
        % The bottom half
        call_tree(midpt-m, k) = disc(B, w, r, dt,
            call_tree(midpt-m-1:midpt-m+1, k+1)');
        put_tree(midpt-m, k) = disc(B, w, r, dt,
            put_tree(midpt-m-1:midpt-m+1, k+1)');
    end
    % The middle
    call_tree(midpt, k) = disc(B, w, r, dt,
        call_tree(midpt-1:midpt+1, k+1)');
    put_tree(midpt, k) = disc(B, w, r, dt,
        put_tree(midpt-1:midpt+1, k+1)');
end
end

disp('Stock prices:');
disp(stock_tree);

if acc_complete
    disp('Call prices:');
    disp(call_tree);
    disp('Put prices:');
    disp(put_tree);
else
    disp('Call bid prices:');
    disp(call_bid_tree);
    disp('Call ask prices:');
    disp(call_ask_tree);
    disp('Put bid prices:');
    disp(put_bid_tree);
    disp('Put ask prices:');
end

```

```

    disp(put_ask_tree);
end
y = 1;

```

B.6 Discounting function: disc.m

```

function y = disc(B, w, r, T, vals)
% This takes discounted expected values back through the tree
    rspf = B*w;
    y = exp(-r*T)*vals*rspf;

```

B.7 Singular value decomposition solver: solve.m

```

function x = solve(A, b)
% Solve the linear system Ay = b where y > 0
% using singular value decomposition - this is for
% when the system is badly scaled

[U, W, V] = svd(A);
% Remove rows/cols corresponding to zero rows of W
u = [];
w = [];
for k = 1 : size(W, 1)
    if any(W(k, :))
        w = [w; W(k, :)];
        u = [u U(:, k)];
    end
end

% Perform backsubstitution
x = svdbksb(u, w, V, b);

```

B.8 Backsubstitution: svdbksb.m

```

function x = svdbksb(u, w, v, b)

% Solves A.x = B for a vector x, where A is specified by arrays U, W, V
% as returned by svd.m m and n are the dimensions of A, and will be
% equal for square matrices. b is the input right-hand side.
m = size(w, 1); n = size(w, 2);

x = zeros(n, 1);
tmp = zeros(1,n);

```

```

for j = 1 : n
    % Calculate U T B.
    s=0;
    if w(j)
        % Nonzero result only if wj is nonzero.
        for i = 1 : m
            s = s + u(i, j)*b(i);
        end
        s = s/w(j);      % This is the divide by wj .
    end
    tmp(j)=s;
end

for j = 1 : n
    % Matrix multiply by V to get answer.
    s=0;
    for k = 1 : n
        s = s + v(j, k)*tmp(k);
    end
    x(j)=s;
end

```

B.9 Bid price: bid.m

```

function y = bid(B, r, T, vals)
% Calculates the bid price for a derivative
global A prices floors

c = vals*B - floors';
H = A*B;
b = prices;
[x, z, status] = pcx(H, b, c);
% For the bid price:
if strcmp(status, 'optimal')
    y = exp(-r*T)*prices'*z;
else
    % Sorry Dave
    y = inf;
end

```

B.10 Asking price: ask.m

```

function y = ask(B, r, T, vals)

```

```

% Calculates the asking price for a derivative
global A prices floors
% For the asking price:
b = -prices;
H = -(A*B);
c = -floors' - vals*B;
[z, x, status] = pcx(H, b, c);
if strcmp(status, 'optimal')
    y = exp(-r*T)*prices'*x;
else
    % Danger, Will Robinson, danger
    y = inf;
end          % if

```

B.11 *PCx* interface: pcx.m

```

function [x,z,status] = pcx(A,b,c)

% [x,z,status] = pcx(A,b,c)
%
% solves primal-dual pair of LPs:
%
% * primal problem:
%   minimize    c'*x
%   subject to  A*x = b
%               x >= 0
%
% * dual problem
%   maximize    b'*z
%   subject to  A'*z <= c
%
%
% INPUT ARGUMENTS:
%
% - A:    mxn-matrix
% - b:    m-vector
% - c:    n-vector
%
%
% OUTPUT ARGUMENTS
%
% x,z: the primal and dual solutions if status = 'optimal'
%      x=z=[] otherwise

```

Bibliography

- [1] S. Aiyagari. Explaining financial market facts: The importance of incomplete markets and transaction costs. *Federal Reserve Bank of Minneapolis Quarterly Review*, 17(1):17–31, Winter 1993.
- [2] Argonne National Laboratory and Northwestern University. Pcx linear programming software. Available at <http://www-unix.mcs.anl.gov/otc/Tools/PCx/index.html>, 1998.
- [3] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1998.
- [4] P. Billingsley. *Probability and Measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, New York, 1979.
- [5] N. Bingham and R. Kiesel. *Risk-Neutral Valuation: Pricing and Hedging Financial Derivatives*. Springer Finance. Springer, London, 1998.
- [6] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659, May 1973.
- [7] P. Carr, H. Geman, and D. Madan. Pricing and hedging in incomplete markets. *Journal of Financial Economics (forthcoming)*, 2001.
- [8] A. Černy and S. Hodges. The theory of good-deal pricing in financial markets. In H. Geman, D. Madan, S. Pliska, and T. Vorst, editors, *Selected Proceedings of the First Bachelier Conference (Held in Paris, 2000)*. Springer, forthcoming.
- [9] J. Cochrane and J. Saó-Requejo. Beyond arbitrage: Good-deal asset pricing in incomplete markets. *Journal of Political Economy*, 108(1):79–119, 2000.
- [10] N. El-Karoui and M. Quenez. Dynamic programming and pricing of contingent claims in an incomplete market. *SIAM Journal on Control and Optimisation*, 33(1):29–66, 1995.
- [11] S. Figlewski and T. Green. Market risk and model risk for a financial institution writing options. *Journal of Finance (forthcoming)*, 2001.

```
% status: string; possible values: optimal, infeasible, unknown
%
% pcx.m saves data in a .mat file and calls a routine based on
% PCx version beta-2.0

A = sparse(A);
b = full(b);
c = full(c);
if (size(b,2) > size(b,1))
    b=b';
end
if (size(c,2) > size(c,1))
    c=c';
end
save temp A b c

!pcxwin temp.mat

load temp;
delete temp.mat
```

B.12 Call option boundary condition: call.m

```
function y = payoff(S, K)
% The payoff for a European call option
y = max(S - K, 0);
```

B.13 Put option boundary condition: put.m

```
function y = payoff(S, K)
% The payoff for a European put option
y = max(K - S, 0);
```

- [12] B. Flannery, W. Press, S. Teukolsky, and W. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1992.
- [13] H. Föllmer and P. Leukert. Quantile hedging. *Finance and Stochastics*, 3(3):251–273, 1999.
- [14] H. Föllmer and P. Leukert. Efficient hedging: Cost versus shortfall risk. *Finance and Stochastics*, 4(2):117–146, 2000.
- [15] J. Fraleigh and R. Beauregard. *Linear Algebra*. Addison-Wesley, 2nd edition, 1991.
- [16] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford Science Publications. Clarendon Press, Oxford, 2nd edition, 1998.
- [17] J. Harrison and D. Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20:381–408, 1979.
- [18] H. He. Convergence from discrete- to continuous-time contingent claim prices. *Review of Financial Studies*, 3(4):523–546, 1990.
- [19] D. Heath and H. Ku. Coherent dynamic risk measures. Working paper, Carnegie Mellon University, 2001.
- [20] J. Hull. *Options, Futures and Other Derivatives*. Prentice-Hall, 4th edition, 2000.
- [21] J. Hull and A. White. Valuing derivative securities using the explicit finite difference method. *Journal of Financial and Quantitative Analysis*, 25:87–100, March 1990.
- [22] The Mathworks Inc. Matlab version 5. <http://www.mathworks.com>, 1996.
- [23] S. Jaschke and U. Küchler. Coherent risk measures and good-deal bounds. *Finance and Stochastics*, 5(2):181–200, 2001.
- [24] I. Karatzas. *Lectures on the Mathematics of Finance*. CRM Monograph Series. American Mathematical Society, Providence, 1991.
- [25] D. Kreps. Arbitrage and equilibrium in economies with infinitely many commodities. *Journal of Mathematical Economics*, 8:15–35, 1981.
- [26] Matif SA. SPAN[©] performance bond calculation method. Available at <http://www.matif.com>, 2001.
- [27] S. Pliska. *Introduction to Mathematical Finance: Discrete Time Models*. Blackwell Publishers, 1997.

- [28] A. Shiriyayev. *Probability*. Graduate Texts in Mathematics. Springer-Verlag, New York, 1984.
- [29] D. Williams. *Probability with Martingales*. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge, 1999.

University of Cape Town

University of Cape Town

Acknowledgements

Thank you Peter Ouwehand for supervising this thesis. I am grateful for your interest, patience and good ideas. I would also like to express my gratitude to the Department of Mathematics at the University of Cape Town for a mind-altering six years – it is going to be difficult to leave.

Email correspondence with Peter Carr, primary author of Carr, Geman and Madan [7], the foundation upon which much of this dissertation is built, was very helpful and encouraging in the early stages of the work. The idea of choosing the test measures using a Monte Carlo-type method comes from him and one of his co-authors, Dilip Madan. He also provided me with some reading material which would not otherwise have come to my attention. I appreciate too the generosity of David Heath, who sent me an outline of some of his work which is still in development, namely Heath and Ku [19].

Thank you also to my parents for their patience and care, to Sarah for her boundless and inspiring enthusiasm for mathematics, and to James, for always asking what (and how) I was doing.