Pricing contingent claims on illiquid securities in continuous time

Master Thesis Econometrics and Management Science – Quantitative Finance

Author
Sebastiaan H.L.C.G. Vermeulen
Student ID: 388049

Supervisor
dr. R-J. Lange

Second assessor
dr. H.J.W.G. Kole

Advisors
bc. P. T. Kloek – Ortec Finance
ir. S. N. Singor – Ortec Finance

ERASMUS UNIVERSITEIT ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS

ORTEC FINANCE

Abstract

We complement arbitrage restrictions with rationality arguments to obtain tight 'good deal' bounds on the set of arbitrage-free valuations in incomplete markets with a complete submarket. We show how to select the stochastic discount factors which induce the lowest and highest option prices under risk-neutral no-good-deal valuation. We find that rationality measured by coherent risk measures restrict the volatility of the stochastic discount factor independent of the model parameters and available products, making coherent restrictions robust against model misspecification and giving them an interpretation that is consistent over different calibrations. Good deal bounds are derived in a general pricing model with stochastic volatility and a stochastic interest rate, and the implications of using good deal bounds in a simulation pricing framework are discussed: Most notably, without the completeness assumption the classical Heston stochastic volatility leads to the possibility of exploding discount processes. To mitigate this issue we propose a modified stochastic volatility model. If risk preferences are incoherent, we find that good deal bounds do not necessarily lead to sensible solutions without imposing additional constraints on the market.

Keywords: Good deal pricing, incomplete market, risk measure, continuous time, arbitrage

Contents

Chapter 1

Introduction

Insurance companies face the issue of pricing long-maturity pension- and insurance products with embedded path-dependent optionality such as interest rate floors or contingent claims on the insurer's profit. These options are written on the insurer's assets, which likely includes illiquid products such as private equity and unlisted real estate. Because these products appear on the liabilities side of an insurer's balance sheet, the Solvency II Directive issued by the European Union dictates that the pricing of these options is done in a 'market-consistent' manner, and as a result, derivatives need to be priced in a risk-neutral setting when continuous hedging is out of the question and prices of the underlying or derivatives thereof are infrequently observed.

To provide a flexible pricing instrument, Ortec Finance generates a set of risk-neutral scenarios nested within a set of real world scenarios. Hence for each point in the future and each real-world scenario, a set of risk-neutral scenarios is available to determine a risk-neutral price for derived products. To generate these risk-neutral scenarios, sample paths are generated via Monte Carlo simulation from a risk-neutral probability measure, which dictates a single price of risk for each financial product. If the market were complete this measure would be uniquely defined, but when the underlying is not actively traded, continuous hedging is out of the question, the completeness assumption fails and there is an infinite set of measures to choose from. For the purposes of a simulation framework, the actual set of risk-neutral measures is of prime importance, and less so a closed-form pricing identity for specific derivatives, since these do not generally exist for stochastic volatility and stochastic interest rate models.

This thesis researches how bounds can be derived on the set of risk-neutral measures in such a manner that the restricted set can be used to simulate scenarios for pricing purposes. These bounds are derived in a general framework such that practitioners can choose the restricting factor which best suits their market views. The implications of modeling choices on price bounds are discussed and it is shown that price bounds are well behaved in the presence of estimation errors, in the sense that small estimation errors lead to small deviations in the price bounds.

Since even in liquid markets transaction costs prohibit frictionless trading, the market completeness assumption has been a topic of discussion since the iconic paper Black and Scholes (1973) introduced a derivative pricing identity which relies on market completeness to have a unique solution. One of the first papers to accept market incompleteness as fact and instead complement the no-arbitrage restriction with other rationality-founded restrictions is Cochrane and Saa-Requejo (2000), who argue that trading opportunities which seem to good to be true are improbable and hence should be disregarded. They quantify 'too good to be true' as a limit on the Sharpe ratio of the payoff process and derive a stochastic differential equation for the lower and upper bounds of the risk-neutral price process under constant volatility and risk-free interest rate. They use the argument of Hansen and Jagannathan (1991) to show that a bound on the Sharpe ratio is equivalent to a bound on the variance of the stochastic discount factor.

Similarly, Bernardo and Ledoit (2000) define good deals in terms of the gain-loss ratio of a payoff process. Černý (2003) demonstrates how bounds on the Sharpe ratio can be translated to bounds on a truncated quadratic utility function, and generalizes the concept of good-deal bounds to convex utility functions. Other notable generalizations include convex risk measures (Arai, 2017), neighborhoods of benchmark risk-neutral measures (Bondarenko and Longarela, 2009) and the extension from diffusion processes to jump-diffusion processes (Björk and Slinko, 2005), or to Lévy processes (Klöppel and Schweizer, 2007). Xu (2006) discusses price bounds based on coherent risk measures from a theoretical perspective, proving the existence and some other properties of coherent risk price bounds for very general single period models. The set of measures that present arbitrage opportunities and the set of measures that present good deals are generally not nested, complicating the derivation of bounds. It will be shown that one solution is to express the bounds on the payoff process as bounds on the stochastic discount factor.

In discrete time, finding upper and lower bounds is done by solving a constrained optimization problem, while for diffusion processes bounds can be simply expressed in terms of the discount factor diffusion differential (Černý, 2003). Moreover, in discrete time, Černý (2003) shows that the bound on expected utility translates to a bound on the stochastic discount factor expectation in a different manner for each utility function, while if the payoff process follows an Itô diffusion, a single upper bound on the discount factor diffusion differential can be derived for the logarithmic, negative exponential, power and truncated quadratic utility function. However, this bound is not necessarily tight on longer horizons. This is formalized by Klöppel and Schweizer (2007), who show that local restrictions on the more general Lévy process result in well-behaved bound processes while global restrictions do not.

Often, illiquid securities can be observed or at least argued to correlate with liquid products. This fact can be exploited to separate hedgeable risk and unhedgeable risk within the illiquid security. In this framework only the unhedgeable risk must be bounded, since in the subspace spanned by the hedgeable risk the risk-neutral measure is uniquely identified, resulting in tighter bounds. (see for example Cochrane and Saa-Requejo (2000); Davis (2006); Floroiu and Pelsser (2013)).

Our contribution to existing literature is the treatise of good deal bounds in the presence of stochastic interest and stochastic volatility, which we show to be a nontrivial extension of models with deterministic interest rate and volatility. In addition, we show how coherent risk measures lead to bound specifications that do not depend on model parameters, giving them a single interpretation for all products that can be valued. To the author's knowledge, this is the first work to give explicit price bound formulae for a large class of good deal restrictions. Other works typically either discuss nonconstructive bound formulations for classes of restrictions or treat specific restrictions. The derivation of coherent risk measure price bounds crucially depends on the following simple yet elegant insight: Because (1) coherent risk measures are positively homogeneous and translation invariant and (2) normally distributed random variables are affine functions (with positive slope) of standard normal variables, the risk of a normally distributed variable can be expressed as an affine function of the risk of a standard normal variable.

Chapter 2 starts with notational conventions, after which the preliminaries for this thesis and some technical features of the tools we will require are discussed, including the relationship between risk measures and good deals. The market model and the valuation of products in the model is discussed in chapter 3. In chapter 4 the derivation of price bounds from risk measures is discussed. Chapter 5 discusses the relation between parameters and price bounds, and compares bounds to the complete market price. An example case demonstrates how price bounds can be used to quantify the pricing error between a complete market model and observed market prices. Lastly chapter 6 provides some concluding remarks and suggestions for future research.

Chapter 2

Preliminaries

In this chapter the concepts we require to make our case are briefly discussed. We start with some notational conventions. Following econometric conventions, primes are used to denote the transpose of a vector or matrix. All vectors are column vectors; row vectors are denoted as transposed column vectors. Throughout this thesis, capital letters will denote stochastic processes, and subscripted capital letters denote the stochastic variable which the process attains at the index given by the subscript. The space of stochastic processes on the space U will be denoted as $\{U_t\}_{t\in[0,\infty)}$. M is reserved for pricing kernels (discussed in more detail in section 2.2) and B is reserved for the bank account, also known as risk-free price process or numeraire. $\rho_{x,y}$ is the correlation coefficient between x and y and $\bar{\rho}_{x,y} = \sqrt{1-\rho^2_{x,y}}$. W invariably refers to a (possibly multivariate) Wiener process. We will be working on a single filtered probability space $\{\Omega, \mathcal{F}, \mathbb{P}\}$ with filtration $\{\mathcal{F}_t\}$. The filtration is assumed to be chosen such that all processes are adapted to this filtration. The projection of x on y is denoted by $\text{proj}_y(x)$, and $\text{orth}_y(x)$ refers to the orthogonal remainder $x - \text{proj}_y(x)$. The only norm used in this thesis is the 2-norm, which is the norm induced by the inner product: $\|x\| = \sqrt{x'x}$. The next section establishes that all the relevant spaces have an inner product defined. Wiener processes, stochastic calculus and Itô's results on stochastic processes are used without introduction. For the construction of Wiener processes (also known as Brownian motions) we refer to Schilling (2017). Stochastic calculus and Itô's results are treated in for example Baxter and Rennie (1996); Cochrane (2009); Back (2010).

2.1 The incomplete market

In this paper we follow the reasoning of Cochrane and Saa-Requejo (2000); Davis (2006); Floroiu and Pelsser (2013): We assume that some, but not all, financial products are traded frequently enough to allow for a continuous hedging strategy such that the market for these products and their derivatives is at least dynamically complete. These efficient assets collectively form the efficient market $\mathcal{M}^c$, a subset of the market $\mathcal{M}$.

Products in $\mathcal{M} \setminus \mathcal{M}^c$ are less frequently traded, thus cannot be continuously hedged and therefore introduce incompleteness to the market. These products correlate imperfectly with hedgeable products and they are the focal point of this thesis.

Without loss of generality, all equity price processes are assumed to accumulate any associated cash flows, such as dividend, in their price. We also expect all price processes to have finite variance. This empirically reasonable restriction yields that

$$\mathcal{M} \subset \mathcal{L}^2 := \left\{ X_t \mid \int_\Omega X_t^2 \mathrm{d}\mathbb{P} < \infty, t \geq 0 \right\}.$$

The technically redundant subscript t is included here to stress the fact that products in the market are sequences; the individual elements of $\mathcal{M}$ do not have a practical interpretation beyond being a process observed at a specific time. We take the usual implicit quotient[1] over almost-everywhere equality in $\mathcal{M}$ such that $\mathcal{M}$ is a subset of a Hilbert space[2]. The prime reason to work in a Hilbert space is because in such spaces the Riesz representation theorem applies: Roughly put, every linear function $\phi(X)$ can be written as an inner product $\langle Y_\phi, X\rangle$ for some Y_ϕ and all X, which in a probabilistic context equals the expectation $\mathbb{E}[Y_\phi X]$. For more details on the Riesz representation theorem, see for example Rynne and Youngson (2008). Although not strictly necessary, the bijection between linear functions and expectation makes the notion of pricing kernels (see section 2.2) much more natural.

Much of the results presented in this thesis can be extended to other $\mathcal{L}^p$-spaces and even to Orlicz spaces, but these generalizations do not offer insights into our problem that cannot be derived in the simpler and more practical $\mathcal{L}^2$ case. With practicality in mind we therefore refrain from these generalizations. See Arai and Fukasawa (2014); Arai (2017) for an overview of what is still possible regarding good deal bounds under more general conditions.

2.2 Pricing kernels and probability measures

To prevent trivial arbitrage strategies, the price of a sum of payoffs must equal the sum of prices for the individual payoffs; this economic concept is known as the Law of one Price. If the Law of one Price does not hold, an investor could either buy the sum of payoffs and sell short the individual payoffs or vice versa, and collect the positive price difference while having a perfectly hedged position. By extension, the price of a payoff process must be a linear function of the payoffs.

Next, note that the payoff associated with products that do not generate a cashflow must be the price for which it can be sold at any time, given that the product exchanges hands at that time. This should be a time-consistent relation: The current price as derived from the expected future sales price should not depend on the time in the future where the product is expected to be sold. If this were the case, an investor could buy the same product at the same time for different prices, depending on when the investor expects to sell the product again, which is an unrealistic scenario for obvious reasons.

[1] Technically, $\mathcal{L}^2$ is not a Hilbert space because if $X \in \mathcal{L}^2$, then X and $Y : Y(\omega) = X(\omega) + 1_{\omega=\omega_0}$ are two different $\mathcal{L}^2$ functions, yet $\|X - Y\|^2 = \int_\Omega 1_{\omega=\omega_0} \mathrm{d}\mathbb{P}(\omega) = 0$. By taking the quotient over almost-everywhere equality, X and Y are considered equal and $\mathcal{L}^2$ becomes a proper Hilbert space. For more details on quotient spaces, see for example Schilling (2017).

[2] Simply put, a Hilbert spaces is a vector space with an inner product and without 'holes'; every Cauchy sequence converges to a point in the space. For some more background, see for example Rynne and Youngson (2008).

We can therefore define a linear valuation operator $\phi_t : \mathcal{M} \to \mathbb{R}$ for each $t \geq 0$ with the property $\phi_t(V_{t+s}) = \phi_t(V_t), \forall s, t \geq 0$. Then by the Riesz-Fréchet theorem there exists a process $M \subset \mathcal{L}^2$ such that $\phi_t(V_{t+s}) = \mathbb{E}[M_{t+s}V_{t+s} \mid \mathcal{F}_t]$ for all $V \subset \mathcal{M}$ and all $t, s \geq 0$. This notion leads to the definition of the pricing kernel or stochastic discount factor (SDF).

Definition 1. A Stochastic Discount Factor or pricing kernel M is a positive random process that correctly prices each process in $\mathcal{M}$. M correctly prices the price process X with associated payoff process C if

$$M_t X_t = \mathbb{E}_t \left[\int_t^\infty M_s C_s \mathrm{d}s \right].$$

In particular, if X does not generate an intermediate payoff, then MX is a martingale; the only payoff is generated by selling X_{t+s} for payoff X_{t+s}, i.e.

$$M_t X_t = \mathbb{E}_t[M_{t+s} X_{t+s}] \quad \forall s, t \geq 0.$$

Here we defined $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_t]$, the expectation operator conditional on the natural filtration of $\cdot$ up to and including t.

Without loss of generality we fix $M_0 = 1$.

Remark 1. The restriction $M > 0$ makes economical sense because each positive future payoff, that occurs with positive probability, should currently represent a positive value. This restriction is also a sufficient condition for the absence of arbitrage opportunities on each process correctly priced by M. In literature where M is defined from a consumption or utility framework, the restriction $M > 0$ doesn't follow by definition but rather by assumption. Since we constructed M directly from economically desirable properties, it was most suitable to include $M > 0$ in the definition.

Remark 2. Pricing via SDFs has the implication that when X and Y have the same payoff structure almost everywhere, then both must have the same price. This Law of one Price property makes pricing kernels an attractive instrument to value derivatives since a derivative's payoff can at some point in the future be expressed as a function of the underlying, such that the value of the derivative can be determined from the underlying.

Within the complete market $\mathcal{M}^c$, only a single SDF can exist; this result is known as the Second Fundamental Theorem of Asset Pricing. Conversely, the uniqueness of an SDF implies completeness of the market. Without the assumption of a complete market, stochastic discount factors are not unique and investors must decide on a specific pricing kernel. An investor's choice of SDF reflects how the investor feels about possible future developments. In fact, an SDF can by its very definition be interpreted as the marginal rate of intertemporal substitution, and as such, the investor's utility function implies a specific choice of SDF.

In addition to SDFs, there are several equivalent approaches to risk-neutral pricing. One of the most well-known approaches and the de facto approach to simulation-based risk-neutral pricing is a change of measure to a so-called risk-neutral measure. The relation between these two concepts is formalized in Definition 2.

Definition 2. If B is the (positive) risk-free price process and M is an SDF for all price processes $X \subset \mathcal{M}$, then there is a probability measure $\mathbb{Q}$ such that $\mathbb{P}(M_t X_t \in A) = \mathbb{Q}(X_t/B_t \in A)$ for each event $A \in \mathcal{F}$, each $t \geq 0$ and each price process $X \subset \mathcal{M}$. Measure $\mathbb{Q}$ is the risk-neutral measure defined by M.

The existence of $\mathbb{Q}$ is evident since Definition 2 gives a recipe that can be used to pointwise define $\mathbb{Q}(X \in A)$ for all $X \in \mathcal{M}, A \subset \Omega$. Note that $\mathbb{P}$ and $\mathbb{Q}$ are equivalent in the sense that both are absolutely continuous[3] with respect to the other. We can thus write $M = \frac{1}{B}\frac{d\mathbb{Q}}{d\mathbb{P}}$, where $\frac{d\mathbb{Q}}{d\mathbb{P}}$ is the Radon-Nikodym derivative, from which we can conclude that for given M and B, $\mathbb{Q}$ is unique.

Moreover, the discounted price process X/B is a martingale under $\mathbb{Q}$ since $\mathbb{E}_{t,\mathbb{Q}}[X/B] = \mathbb{E}_{t,\mathbb{P}}[MX]$. For this reason, $\mathbb{Q}$ is known as an Equivalent Martingale Measure (EMM) of $\mathbb{P}$. Each EMM of $\mathbb{P}$ also induces an SDF, which is almost-everywhere unique in the sense that if M and M' are SDFs induced by the same EMM, then $M - M' = 0$ almost everywhere, which can again be deduced from the existence of $\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$.

2.3 Risk-neutral pricing through Monte Carlo simulation

By definition the expectation of any payoff process X discounted with an SDF M should equal the current price of X. It is therefore possible to deduce the value of a derivative as the M-discounted payoff as described by the derivative's dependence on the underlying: For example, a European call option C with maturity T can be valued as the M-discounted maximum of either 0 or the value of the underlying S minus the strike price K: $M_t C_t = \mathbb{E}_t[M_T(S_T - K)^+]$.

With the exception of some specific cases, the M-discounted expectation of derivatives can be rather involved. Notably so when the derivative depends on intermediate values of the underlying or when M or the underlying follow general processes (such as jump diffusions or processes with covarying unobserved variables like conditional variance), situations that often arise in insurance valuation. In such cases the expected value of a derivative can be approximated by sampling price paths from the specification of X as well as pricing kernel paths from the specification of M, and then average the value of the derivative for each of the individual paths of X discounted with M.

If M and X are drift-diffusion processes, it is sufficient to know the stochastic differential equations describing X and M. Itô's lemma gives the identity

$$\mathrm{d}(MX) = \mathrm{d}MX + M\mathrm{d}X + \mathrm{d}M\mathrm{d}X,$$

which shows that the diffusion process driving MX is given by the diffusion processes of X and M. MX has no drift since this would violate the martingale property.

A practical issue arises in Monte Carlo simulation: It is not possible to sample continuous-time paths since such a path will not have a finite-dimensional representation. To overcome this problem a number of discretization methods have been developed that sample a finite number of observations at times $0 < \tau_0 < \tau_1 < \cdots < T$ from the approximate distribution of the process X observed at times $\tau_0, \tau_1 \ldots$ respectively. In some corner cases the mismatch between the approximate discrete distribution and the true continuous distribution leads to new problems. One case is particularly relevant to us; the Euler discretization of Cox-Ingersoll-Ross processes introduces a positive probability for negative values, something that is impossible under the true (continuous-time) distribution. In addition to widely used heuristics such as reflecting all sampled negative values to their absolute values, Andersen (2007) describes two discretization schemes that more accurately proxy the true distribution, thereby alleviating the problem.

[3] A measure ν is absolutely continuous with respect to measure μ if and only if $\nu(A) > 0 \Rightarrow \mu(A) > 0$.

2.4 Risk measures and good deals

Risk measures are functions that allow one to qualitatively compare the probability distributions of (possibly risky) payoffs. Every risk measure corresponds to a preference profile, since a risk measure can be used to order payoffs according to the risk perceived by the measure. Well-known examples of risk measures include Value at Risk (VaR) and Expected Shortfall (ES)[4]. For a preference ordering to be sensible, a number of constraints can be formulated, which lead to the definition of convex and coherent risk measures.

Definition 3 (Convex risk measure)**.** A convex risk measure is a function $\psi : \mathcal{L}^2 \to \mathbb{R}$ which satisfies the following constraints:

convexity $\psi(\lambda X + (1-\lambda)Y) \leq \lambda\psi(X) + (1-\lambda)\psi(Y)$ for all $\lambda \in [0, 1]$,

translation invariance (or cash invariance) $\psi(X+a) = \psi(X)-a$ for all (deterministic) $a \in \mathbb{R}$,

monotonicity if $X \leq Y$ almost everywhere, then $\psi(X) \geq \psi(Y)$.

Definition 4 (Coherent risk measure)**.** Let ψ be a convex risk measure. If, additionaly, ψ satisfies the following constraint:

positive homogeneity $\psi(\lambda X) = \lambda\psi(X)$ for all $\lambda \geq 0$,

then ψ is a coherent risk measure.

These constraints can be motivated as follows. Convexity implies both that risk cannot be increased through a merger, and that two risk factors cannot negatively influence each other simply because an investor has both items in a single portfolio. Translation invariance implies that the value (in terms of risk) of a single unit of cash is independent of the payoff X, as well as the size of the cash deposit. Monotonicity guarantees consistency: if a payoff Y will be at least as high as X regardless of the state of the world at the cash-out moment, then Y cannot be considered more risky than X. Lastly, positive homogeneity implies that an investor is neither risk averse nor risk loving. The investor values risk according to its dynamics and scales the value of risk appropriately with with the scale of the risk itself. It is a well-known result that VaR is neither convex or coherent, while Expected Shortfall is a coherent risk measure.

Then we impose a technical restriction: from here on forward risk measures are expected to be lower semi-continuous such that, by the Fenchel-Moreau theorem (see for example Rockafellar (1997) for some background and a proof) all risk measures are equal to their biconjugate. Although convex or coherent functions are not by default lower semi-continuous, this hardly poses a restriction in practical applications because lower semi-continuity is often a desirable property for risk measures: if payoff Y is marginally larger than X, i.e. $X \leq Y \leq X + \delta$ for a small $\delta > 0$, then one would expect that the risk of Y is only marginally smaller than X; that is, $\psi(Y) > \psi(X) - \epsilon$ for a small ϵ.

[4]The 100α% value-at-risk of X is the value $-k$ of X for which $\mathbb{P}[X < k] = \alpha$. The Expected Shortfall equals $\mathbb{E}[-X \mid X < k]$. Expected Shortfall is also known as conditional VaR (CVaR) and Average VaR (AVaR).

2.4.1 Good deals

It stands to reason that an investor whose risk profile is adequately described by risk measure ψ will consider a transaction to be a good deal if that transaction exposes him or her to negative risk. Every transaction involves a payoff C_X and an initial price P_X, and the investor's risk is $\psi(C_X - P_X) = \psi(C_X) + P_X$[5]. As one would expect, any transaction can thus be made into a good deal by lowering the initial price. As such, a risk measure induces the acceptance set containing all good deals

$$\mathcal{A}_\psi := \{X \in \mathcal{L}^2 \mid \psi(X) \leq 0\},$$

and a payoff X does not present a good deal if and only if the price of X lies in the interval $[-\psi(X), \psi(-X)]$. To see this, note that an investor buying X at price P_X will value the risk in this transaction as $\psi(X - P_X) = \psi(X) + P_X$, such that this buyer considers the transaction a good deal if $P_X < -\psi(X)$. Similar reasoning from the seller's perspective leads to the upper bound. Note that these are bounds from the perspective of agents with risk preferences matching ψ. An individual will only ever be interested in one of the two limits during a single transaction, and in any one transaction the buyer and seller may also be using different risk measures, with the possibility that they cannot find common ground.

Each acceptance set also describes a risk measure. It can be verified that $\mathcal{A}_\psi$ is convex if and only if ψ is convex, and that $\mathcal{A}_\psi$ is a cone if and only if ψ is positively homogeneous. For a proof, see Lemma 1 in the appendix. Coherent risk measures ψ therefore impose a rationality restriction on the market in the form of the convex cone $\mathcal{A}_\psi$.

Good deal valuation can be seen as a generalization of arbitrage pricing. A deal to buy future payoff X for price P_X is considered an arbitrage opportunity when X is risky and $X - P_X \geq 0$ regardless of the state of the world at the pay-out moment. Viewing arbitrage opportunities as good deals, the risk measure corresponding to arbitrage pricing is

$$\psi_0(X) := \inf\{a \in \mathbb{R} \mid X + a \geq 0\},$$

or in words, $\psi_0(X)$ is the worst-case loss of X. Note that ψ_0 is a coherent risk measure: homogeneity, monotonicity and translation invariance can directly be verified from their definitions. Given that ψ_0 is homogeneous, to prove convexity it is sufficient to prove that $\rho_0(X+Y) \leq \psi_0(X) + \rho_0(Y)$. Since $X + \psi_0(X) + Y + \psi_0(Y) \geq 0$, $\psi_0(X + Y + \psi_0(X) + \psi_0(Y)) = \psi_0(X+Y) - \psi_0(X) - \psi_0(Y) \leq 0$, hence the first inequality holds and ψ_0 is coherent.

2.4.2 Good deals for coherent risk measures

The question remains how the acceptance set induced by ψ restricts the choice for arbitrage-free pricing kernels, or equivalently, EMMs. The representation theorem (initially proven by Artzner et al. (1999), later refined by Delbaen (2002); Ruszczyński and Shapiro (2006)) shows that a risk measure is coherent if and only if a set $L \subset \mathcal{L}^2$ and a set of probability measures $\mathcal{Q}$, the risk envelope, exist, such that for each $X \in \mathcal{M}$,

$$\psi(X) = \sup_{Y \in L} \left\{ \int_\Omega XY \mathrm{d}\mathbb{P} \right\} = \sup_{\mathbb{Q} \in \mathcal{Q}} \{\mathbb{E}_\mathbb{Q}[X/B]\}.$$

[5]Note that the perspective of the investor, i.e. whether he or she is a buyer or seller, can simply be modeled through the sign of C_X and P_X.

Note that the dual representation implicitly defines $\mathcal{Q}$ as the set of measures

$$\{\mathbb{Q} \mid \mathbb{E}_{\mathbb{Q}}[X/B] \leq \psi(X) \quad \text{for all } X \in \mathcal{L}^2\}.$$

This relation is formalized in Huber (1981, Chapter 10). The measures in $\mathcal{Q}$ are naturally absolutely continuous with respect to $\mathbb{P}$ due to the correspondence between L and $\mathcal{Q}$. On the other hand, $\mathbb{P}$ is absolutely continuous with respect to $\mathbb{Q} \in \mathcal{Q}$ if and only if ψ is relevant – a condition introduced in Delbaen (2002): if X is not positive and nondegenerate (i.e. truly nondeterministic) then X must have positive risk. It is shown in Arai and Fukasawa (2014) that the $\mathcal{Q}$ contains only EMMs if additionally $\psi(0) = 0$.

An attractive property of coherent risk measures is the conicity of the acceptance set. If a specific investment position is considered acceptable by an investor, then that investor presumably is also willing to take a leveraged position, which implies a conical acceptance set. Investors with coherent risk tastes are therefore indifferent with regard to the leverage in their portfolio and rely on practical arguments to determine their leverage, such as the limited willingness of lenders to provide credit.

2.4.3 Good deals for convex risk measures

The coherent case can be generalized as follows. Convex (not necessarily coherent) lower semi-continuous risk measures can be represented by the Legendre-Fenchel transform

$$\psi(X) = \sup_{Y \in \mathcal{L}^2} \{\mathbb{E}[XY] - \psi^*(Y)\}, \quad \text{where} \quad \psi^*(Y) := \sup_{X \in \mathcal{L}^2} \{\mathbb{E}[XY] - \psi(X)\}.$$

The connection to the coherent case is that if ψ is coherent, then ψ^* acts as an indicator function taking on the values 0 or ∞, thereby restricting the choices of Y to elements from L in the supremum formulation of ψ. The details are treated in Lemma 2.

If $\mathcal{A}_\psi$ is a cone, then

$$\psi^*(Y) \geq \sup_{X \in \mathcal{A}_\psi} \{\mathbb{E}[XY] - \psi(X)\}, \text{ such that } \psi^*(Y) = \begin{cases} 0 & \text{if } Y \in \mathcal{A}_\psi \\ \infty & \text{if } Y \notin \mathcal{A}_\psi \end{cases}$$

since $\psi(X) \leq 0 \leq \psi^*(Y)$ for $X, Y \in \mathcal{A}_\psi$. Consequently, ψ^* acts as the indicator function, restricting $\mathcal{Q}$ to $\mathcal{A}_\psi$ if ψ is coherent.

Convex risk measures are relevant to model risk-averse investor preferences. In particular, ψ is a convex function of the scale of X since $\psi(\lambda aX + (1-\lambda)bX) \leq \lambda\psi(aX) + (1-\lambda)\psi(bX)$ for any a, b and $\lambda \in [0, 1]$. This shows that an investor with a convex risk profile considers leverage as a means to change the exposure to risk. Given a cashflow C_X with price P_X and $\lambda \geq 0$, one can expect to also find the cashflow λC_X for price λP_X in the market, but an investor with convex risk preferences will generally not consider the investment attractive for all values of λ.

Chapter 3

The market model

For typical pension- and insurance questions there is an emphasis in market models on long-term forecasting accuracy and – arguably as important – model credibility. For this reason, a model should display some stylized facts regarding long-term effects: Volatility and interest rates should vary over time, but periods of high volatility or interest should be clustered. As such, both volatility and interest rate should be modeled as stochastic variables. In addition stochastic volatility helps to explain the volatility smile (Hull and White, 1987).

Following the Heath-Jarrow-Morton (HJM) framework of Heath et al. (1992), the forward rate $f_t(\tau)$ from time t to time τ is modeled as an Itô diffusion, and both the short rate r_t and the numeraire process are deterministically related to the forward rate:

$$\begin{aligned} \mathrm{d}f_t(\tau) =& \mu_{t,\tau}\mathrm{d}t + \sigma'_{f,t,\tau}\mathrm{d}W_t, \\ r_t :=& f_t(t), \\ \mathrm{d}B_t =& r_t B_t \mathrm{d}t. \end{aligned} \tag{3.1}$$

Both the scalar μ and the vector σ_f are allowed to vary with both the time t and maturity τ. Here σ_f is, like $\mathrm{d}W_t$, a vector-valued process because $f_t(\tau)$ is allowed to correlate with other market-driving processes, which manifests as the cosine between σ_f and the σ-s of other processes. Under these general specifications B has the undesired property that the expected value of the continuously compounded short rate is not necessarily equal to the forward price over the same timespan. To enforce equality the choice of μ requires that

$$\mathbb{E}_t[B_\tau] = \mathbb{E}_t\left[B_0 \exp\left(\int_t^\tau r_s \mathrm{d}s\right)\right] = \exp\left(\int_t^\tau f(t,s)\mathrm{d}s\right).$$

This restriction can be simplified to give the following choice of μ_f:

$$\mu_{t,\tau} = \sigma'_{f,t,\tau}\int_t^\tau \sigma_{f,t,s}\mathrm{d}s,$$

which is consistent with practical implementations of the HJM framework (Cheyette, 1992; Ritchken and Sankarasubramanian, 1995). The HJM interest rate model allows one to model the entire forward rate curve through the specification of $\sigma_{f,t,T}$ and the forward curve at $t = 0$. The HJM model also allows for negative interest rates, a property that has become a desirable feature for interest rate models in more recent years. It was also shown in Hoorens (2011) that the HJM model provides a better fit of the swaption market volatility skew than the stochastic interest model of Hull and White (1990).

For the volatility and price processes we consider a variant of the stochastic volatility model of Heston (1993), describing price processes S and V:

$$\begin{aligned} \mathrm{d}S_t/S_t &= (r_t + q_S)\mathrm{d}t + (1 + \sqrt{\nu_t})\sigma_S'\mathrm{d}W_t, \\ \mathrm{d}V_t/V_t &= (r_t + q_V)\mathrm{d}t + (1 + \sqrt{\nu_t})\sigma_V'\mathrm{d}W_t, \\ \mathrm{d}\nu_t &= \kappa(\xi - \nu_t)\mathrm{d}t + \sqrt{\nu_t}\sigma_\nu'\mathrm{d}W_t. \end{aligned} \tag{3.2}$$

Here S is a frequently traded asset and V is infrequently traded, and therefore introduces incompleteness in the market. The parameters in this model can be interpreted as follows: q_S and q_V are the risk premium of S and V, respectively. Each of the processes has its own exposure to the multi-dimensional state process $\mathrm{d}W_t$, which is modeled by the vector $\sigma_\bullet$. The process V is nonredundant if $\mathrm{orth}_{\sigma_S}(\sigma_V) \neq 0$. $\sqrt{\nu_t}$ acts as a measure of the degree of erratic changes in the world: unexpected price movements become greater as ν_t increases. As such, $\sqrt{\nu_t}$ can be interpreted as the market stress level. In this model, $\sqrt{\nu_t}$ is recurrent with asymptotic mean $\sqrt{\xi}$ and mean-reversion rate κ if $\kappa > 0$ and $\xi > 0$. ν_t reacts to the state of the world W via σ_ν.

The model 3.2 deviates from the classical Heston model in how the 'vol-of-vol' ν_t interacts with the volatility of S and V. In our model the processes S and V have minimal volatility $\|\sigma_S\|$ and $\|\sigma_V\|$ respectively, while in the original Heston model their volatility was allowed to reach 0. The reason for this change is that the constant excess returns q_S and q_V cannot be nonzero without some compensation. If the volatility of, say, S is allowed to become zero at some point, then a simple long position in S at that moment would be an arbitrage. The argument made here is a technicality in the sense that ν_t cannot stay 0 for a measurable amount of time, and as such these arbitrage opportunities are not feasibly exploitable. However, allowing any price process to haphazardly expose arbitrage opportunities is an inconsistency in the model that would prevent us from excluding arbitrage opportunities on a theoretical basis. As an example of what can go wrong, let for a moment X be given by $\mathrm{d}X_t/X_t = (r_t + q_X)\mathrm{d}t + \sqrt{\nu_t}\sigma_X'\mathrm{d}W_t$, and consider the Girsanov transformation of X to the risk-neutral measure. From analogy with the Black-Scholes model it can be deduced that the change of measure should be given by the Dóleans-Dade exponential of the process $\frac{-q_X}{\|\sigma_X\|\sqrt{\nu}}$. However, this process is discontinuous in $\{\nu = 0\}$ such that the Girsanov theorem does not hold. The Feller condition ($2\kappa\xi > \sigma_\nu'\sigma_\nu$) does not provide solace since it merely proves that $\{\nu = 0\}$ is a null set (Gikhman, 2011), and this is not enough to let $1/\sqrt{\nu}$ be continuous.

All parameters are expected to be different from zero to exclude degenerate or trivial cases. We further assume that $\kappa, \xi > 0$ to guarantee recurrence, otherwise the parameters $q_S, q_V, \rho_{V,S}, \kappa, \xi, \sigma_V, \sigma_\nu$ are unrestricted. The Feller condition is *not* assumed. The Feller condition restricts the permitted volatility skewness and is consistently violated in market observations (Andersen, 2007), while it fails to provide an advantage in our analysis. This model is the simplest non-trivial extension of the Heston model to include a new infrequently traded product.

Ignoring the process V for a moment, the stochastic volatility component ν induces market incompleteness if the volatility risk is not traded. To see this, it is sufficient to realize that if volatility risk is not traded, there is a set of EMMs that ascribe the same probabilities to random variates which are orthogonal to σ_S, but different probabilities to variates in the space spanned by σ_ν. EMMs from this set will value S the same, but induce different prices for contingent claims on S since contingent claim prices are sensitive to the variance of S. The extension theorem of Kreps (1981) yields that this market can be completed by adding appropriate securities, and it is shown in Bajeux-Besnainou and Rochet (1996) that a European call option on S completes the Heston-HJM-market. Instead of modeling this product explicitly, we assume that the part of σ_ν that is orthogonal to σ_S, is also orthogonal to σ_V. Economically, this restriction implies that V contains no information about ν that is not also in S. Practically, this means that the parameters describing ν can be identified via the volatility smile of options on S (Sircar and Papanicolaou, 1999).

Model (3.2) specifies risk premia for risky assets as a constant rate. In the presence of a stochastic short rate and stochastic volatility this is a plausible assumption; Investors cannot accurately determine the current value of the (unobserved) vol-of-vol ν. The best inference for future values of ν is therefore the (constant) unconditional mean ξ, which means that investors can only assess their exposure to risk as a constant, by which the risk premium should also be a constant. On the other hand, because the risk-free return is observed, changes in the short rate should reflect on the expected return of risky assets, hence the expected return on risky assets should consist of the short rate and a constant risk premium.

3.1 Pricing kernels in the Heston-HJM model

Assuming that the risk-neutral discounted price process is also a diffusion, we can make an educated guess about the shape of the pricing kernel M. The process M may have both a drift and diffusion term, and without loss of generality the diffusion term can be split in terms that correlate with S, V, and ν. It is safe to assume that M does not correlate with any risk that is independent from S, V, and ν because, lacking financial compensation, this risk does not need to be discounted.

$$\frac{\mathrm{d}M_t}{M_t} = \alpha_t \mathrm{d}t + \sigma_M \mathrm{d}W_t := \alpha_t \mathrm{d}t + (\beta_t \sigma_S + \gamma_t \mathrm{orth}_{\sigma_S}(\sigma_V) + \delta_t \mathrm{orth}_{\sigma_S}(\sigma_\nu))' \mathrm{d}W_t. \tag{3.3}$$

Here $\beta_t \|\sigma_S\|$ is to be interpreted as the premium for the risk involved with investing a standardized unit in S, and $\gamma_t \|\mathrm{orth}_{\sigma_S}(\sigma_V)\|$ (resp. $\delta_t \|\mathrm{orth}_{\sigma_S}(\sigma_\nu)\|$) is the premium for the marginal risk pertaining V (resp. ν). The functions α and β can be identified via the martingale property of SM and BM. For nonnegative Itô diffusions to be a martingale it is necessary and sufficient to have zero drift, i.e.

$$\begin{aligned} 0 = \frac{\mathrm{d}(MB)_t}{M_t B_t} =& \mathrm{d}M_t/M_t + \mathrm{d}B_t/B_t + \frac{\mathrm{d}M_t \mathrm{d}B_t}{M_t B_t} \\ =& (r_t + \alpha_t)\mathrm{d}t + (\beta_t \sigma_S + \gamma_t \mathrm{orth}_{\sigma_S}(\sigma_V) + \delta_t \mathrm{orth}_{\sigma_S}(\sigma_\nu))' \mathrm{d}W_t, \end{aligned}$$

and likewise

$$\begin{aligned} 0 = \frac{\mathrm{d}(MS)_t}{M_t S_t} =& (r_t + q_S + \alpha_t + (1 + \sqrt{\nu_t})\beta_t \sigma_S' \sigma_S)\mathrm{d}t \\ &+ ((1 + \sqrt{\nu_t} + \beta_t)\sigma_S + \gamma_t \mathrm{orth}_{\sigma_S}(\sigma_V) + \delta_t \mathrm{orth}_{\sigma_S}(\sigma_\nu))' \mathrm{d}W_t, \end{aligned}$$

induce the restrictions

$$\alpha_t = -r_t, \qquad \beta_t = \frac{-q_S}{\sigma_S'\sigma_S(1+\sqrt{\nu_t})}. \tag{3.4}$$

Given that M correctly prices options on S, the price of a European call option on S with strike K and maturity T must match $\mathbb{E}[M_T(S_T - K)^+]$. While an analytical identity for δ_t in terms of the observed model parameters is not at hand because the distributions of M_t and S_t do not have closed form representations for all choices of σ_M and σ_S, the market price of volatility risk (and with it, δ_t) can be estimated by fitting the market-observed volatility surface of options on S to the theoretical volatility surface implied by the model. For the details pertaining this calibration, we refer to for example Heston (1993); Sircar and Papanicolaou (1999); Romano and Touzi (1997).

Given that V is infrequently traded we cannot assume that M uniquely prices V, and as such γ remains undetermined. Instead, we will introduce rationality bounds on γ to restrict the set of feasible pricing kernels. Note that by construction $M > 0$ for any $\alpha_t, \beta_t, \gamma_t, \delta_t$, such that M does not admit arbitrage opportunities, regardless of these parameters' values.

The shape of M signifies again why the variance of risky assets must be greater than or equal to some $\epsilon > 0$. If β_t were given by $\frac{-q_S}{\sigma_S'\sigma_S\sqrt{\nu_t}}$, then β_t would not be locally bounded and therefore M would not be an Itô process, while the results in this paper crucially depend on M having properties of Itô processes.

It can be verified that under restrictions (3.4) the discounted price processes of S and V do not depend on the short rate. The economic intuition behind this is the same as in complete market models: Money that is not invested in the market resides in a risk free account, and therefore the growth of a risky investment should be compared to the growth of the bank account. The fact that the short rate is stochastic is unrelated: Even when the change in interest is correlated with changes in a risky portfolio, the change in short rate can only be observed ex post and can therefore not be used ex ante to change the investor's market view – or by extension, the stochastic discount factor. The relevance of the interest rate model therefore largely stems from the valuation of derivatives, which may depend on the interest rate at maturity or the interest rate path through for example the strike price.

3.2 From option price bounds to variance restrictions

European, Bermudan and American options whose payoff is a convex function of the price of the underlying all share the property that their value monotonically varies with the measurable risk-neutral volatility process σ_t of the risky asset.

Janson and Tysk (2003); Ekström (2004) prove this for the case where σ_t is a general measurable process and where the risky asset resides in the complete market. Outside of the complete market, the problem arises that the discount factor volatility both influences the mean and variance of the discounted price process, such that their results cannot be extended. However, in a drift-diffusion framework the presence of a pricing kernel guarantees a monotonic relation between the volatility of M and the price of any priceable asset:

Theorem 1. *Let M be a pricing kernel as in* (3.3)*. For any asset X with drift process μ_X, diffusion process σ_X and cashflow process C_X, the value of X is a monotone function of γ.*

The principal implication of Theorem 1 is that options achieve their minimum and maximum values at the extreme values for γ_t, all other things left equal.

This implies that, for the purpose of determining good deal price bounds on convex options, we only need to calculate option prices using the pricing kernels with minimal and maximal variance, respectively, permitted under good deal bounds. For example, one could generate a single scenario set of real-world price-path samples, and calculate the minimum and maximum pricing kernel variance processes permitted under the good deal bound. The real-world scenario set can then be discounted using these two pricing kernels, resulting in two scenario sets of risk-neutral price-path samples. These sets could finally be used to price options, resulting in the approximate minimum and maximum option prices permitted under the good deal bound. As the number of samples in the scenario set approaches infinity, these option prices will converge almost surely to the minimum and maximum permitted option prices.

Proof of theorem 1. This proof is an adaptation of the proof of proposition 3 in Cochrane and Saa-Requejo (2000).

The definition of M yields that

$$M_t X_t = \mathbb{E}_t \left[\int_t^\infty M_s C_{X_s} \mathrm{d}s \right]$$

which can be split by the law of iterated expectations as

$$M_t X_t = \mathbb{E}_t \left[\int_t^{t+\Delta t} M_s C_{X_s} \mathrm{d}s \right] + \mathbb{E}_t \left[\mathbb{E}_{t+\Delta t} \left[\int_{t+\Delta t}^\infty M_s C_{X_s} \mathrm{d}s \right] \right].$$

Taking the limit $\Delta t \downarrow 0$, we can switch to differentials finding

$$\begin{aligned} M_t X_t &= \mathbb{E}_t \left[M_t C_{X_t} \mathrm{d}t \right] + \mathbb{E}_t \left[M_{t+\mathrm{d}t} X_{t+\mathrm{d}t} \right] \\ &= M_t C_{X_t} \mathrm{d}t + \mathbb{E}_t \left[(M_t + \mathrm{d}M_t)(X_t + \mathrm{d}X_t) \right] \\ &= M_t C_{X_t} \mathrm{d}t + M_t X_t + \mathbb{E}_t \left[\mathrm{d}(M_t) X_t + M_t \mathrm{d}(X_t) + \mathrm{d}(M_t)\mathrm{d}(X_t) \right]. \end{aligned}$$

Reordering yields

$$M_t C_{X_t} \mathrm{d}t = (-\alpha_t M_t X_t - \mu_X M_t - \sigma_M' \sigma_X M_t)\mathrm{d}t.$$

The last step follows from the definition of M and an application of Itô's lemma. Since σ_M is both a linear function of γ_t and the only right-hand term involved with γ_t, the left-hand term is a linear function of γ_t whose sign depends on σ_X. Therefore the integral $I(\gamma) := \int_t^\infty M_t C_{X_t} \mathrm{d}t$ has the property that if $\tilde{\gamma} = \gamma + \hat{\gamma}$ where $\hat{\gamma}$ is a nonnegative process, then whether $I(\tilde{\gamma}) \geq I(\gamma)$ or $I(\tilde{\gamma}) \leq I(\gamma)$ only depends on whether $\sigma_M' \sigma_X$ is increasing or decreasing in γ, and not on the values of γ or $\tilde{\gamma}$. This relation is invariant to the expectation operator which concludes the proof. The inequality can be made strict if $\hat{\gamma}$ is a strictly positive process. □

Chapter 4

Rational restrictions

In chapter 3 we discussed a parametrization of the market and the set of pricing kernels. This set of pricing kernels is determined by a complete subset of the market and a single process-valued free parameter γ. This chapter formulates how rational expectations of the market lead to restrictions of the choice of γ. Combined with the monotonicity of option prices as a function of γ, proven in Theorem 1, rational restrictions can then be used to choose two processes for γ corresponding to the minimum and maximum option prices attainable under the rational expectations.

First, section 4.1 asserts necessary properties for rational restrictions. With these properties the idea of good deals is then used to distill rational restrictions on the market from risk measures. The specific case of coherent restrictions is treated in section 4.2. Finally section 4.3 discusses how the results for the coherent case can be generalized for incoherent risk measures.

4.1 Good deal bounds

This section formalizes properties that good deal bounds should adhere to in a continuous time setting. This is done is a somewhat axiomatic approach; it is natural to expect certain behavior from a price bound in the real world. Appropriate mathematical properties are therefore derived from these economical expectations.

4.1.1 Translation invariance and $\mathcal{F}_t$-adaptedness

The goal of good deal bounds is to exclude trading opportunities which are 'too good to be true', i.e. good deal bounds should depend only on the return of an investment, and not its absolute size. In particular should a price bound be taking the current value of a price process as given; This means that a price bound for X_{t_1}, calculated at t_0, should be expressible as a price bound on the change in X between t_0 and t_1; it makes no sense to constrain the (exogenous at time t_0) value of

X_{t_0}. To this end bounds need the following two properties.

- Bounds should be $\{\mathcal{F}_t\}_{t\in[0,\infty)}$-measurable: This means in particular that bounds are conditional on a filtration, and that $b(X_t \mid X_t) = X_t$ for bound b and any $\{\mathcal{F}_t\}_{t\in[0,\infty)}$-adapted process X.
- Bounds should be translation invariant: $b(X_t + Y \mid \mathcal{F}_s) = b(X_t \mid \mathcal{F}_s) + Y$ for bound b, for all $X_t, \mathcal{F}_s$ and all $\mathcal{F}_s$-measurable Y.

4.1.2 Time consistency

The idea of time consistency for price bounds presented here strongly relates to an equivalent concept for risk measures. The idea of a risk measure can be extended to a dynamic risk measure, which is a sequence of risk measures, each element of which values a stochastic variable at a different moment in time. A dynamic risk measure can be thought of as measuring the risk process of a stochastic variable. In discrete time, dynamic risk measures have been extensively studied (see Acciaio and Penner (2011) for an overview) but the continuous time case has some mathematical intricacies that prevent the derivation of equally general results. Time consistency for dynamic risk measures in discrete time is the requirement that if payoff X is deemed more risky than payoff Y at time t, then X should also be considered more risky than Y at times $s < t$. The rationale for time consistency is that if a risky event between times s and t can decrease the perceived risk of Y at time t below that of X, then this event should not weigh towards a higher observed risk of Y than of X at time s. Simply put, inconsistency can therefore be seen as misinterpreting important events during the assessment of risk.

The above characterization of risk is equivalent to a recursive relation: If $\{\psi_t\}_{t\in[0,\infty)}$ is a dynamic risk measure then $\{\psi_t\}_{t\in[0,\infty)}$ is time consistent if and only if $\psi_t(X) = \psi_t(\psi_{t+s}(X))$ for all nonnegative t, s. See lemma 3 for a proof. Since good deal bounds relate to an investor's preference through her risk measure, they should also be time-consistent; in particular should we have that a price bound for a product over a fixed period leads to the same price range as when one iteratively determines upper and lower bounds for said product over subintervals of the fixed period. Analogous to the dynamic risk measure case, if a price bound is inconsistent then one could argue that the price bound does not comprehensively capture the dynamics of the valued product, which is an obvious design flaw of the price bound. Time-consistency is graphically explained in Figure 4.1. The green (red) lines point connect the upper (lower) bounds b with their right tail as determined with the current price indicated by the left tail. The bounds determined at time 1 given time 0 should be the same whether determined over the entire interval at once, or determined conditional on intermediate bound values.

Mathematically, an (upper or lower) price bound is a continuous function that maps the possible development of a process between two points in time to an extreme value for that process at the end of the interval. That is, bound $b : [0,\infty) \times [0,\infty) \times \{\mathcal{L}^2_t\}_{t\in[0,\infty)} \to \mathbb{R}$ is a function that maps the distribution of the process X at time t_1 to an (upper or lower) price $b(X_{t_1} \mid X_{t_0})$ at t_0, for times $0 \le t_0 \le t_1$. Time consistency is then the requirement that $b(X_{t_1} \mid X_0) = b(X_{t_1} \mid b(X_{t_0} \mid X_0))$. Note that this formulation of time consistency has a strong analogy with the law of total expectation. Since this latter formulation of time-consistency is a recursive relation, price bounds for larger time intervals can simply be recursively *defined* from smaller interval price bounds. In a continuous-time

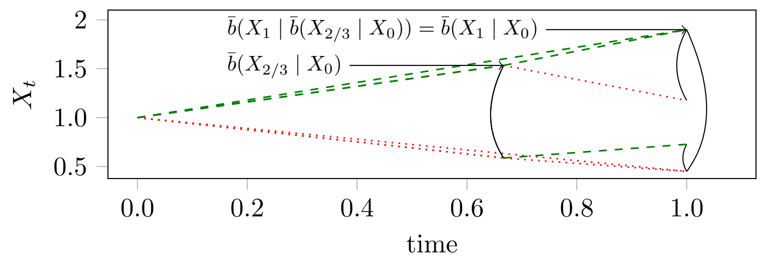

Figure 4.1: Green (red) lines connect the upper (lower) bound on the right to the current price on the left for the process X. The associated price bounds are shown as curved black lines. To be time-consistent, the price interval derived at time 0 for X_1 (right curve) should be equal to iterated price intervals (left curves) over period $t = 2/3$. The iterated price intervals are shown at the boundary values of the previous interval. $\bar{b}$ denotes the upper price bound.

setting, time consistency therefore implies that price bounds on X_t can be expressed as restrictions of the differential of X_t. Time consistency in conjunction with translation invariance yields that

$$b(X_t \mid X_0) = b(X_{t-\Delta t}+[X_t - X_{t-\Delta t}] \mid b(X_{t-\Delta t} \mid X_0)) = b(X_{t-\Delta t} \mid X_0)+b(\int_{t-\Delta t}^{t} \mathrm{d}X_s \mid b(X_{t-\Delta t} \mid X_0)),$$

and by repeating the above process one can find

$$b(X_t \mid X_0) = \sum_{i=1}^{n} b(\int_{t_{i-1}}^{t_i} \mathrm{d}X_s \mid b(X_{t_{i-1}} \mid X_0)) \quad \text{with } 0 = t_0 < t_1 < \cdots < t_n = t, \tag{4.1}$$

which shows that it makes sense to define $b(X)$ as the sum of bounds on small increments in X. Given the continuity of X as a function of time, the limit

$$\lim_{\Delta t \downarrow 0} b(\int_{t-\Delta t}^{t} \mathrm{d}X_s \mid b(X_{t-\Delta t} \mid X_0)) = b(\mathrm{d}X_t \mid b(X_t \mid X_0))$$

is well defined. This in turn guarantees that the sum (4.1) is well defined as $n \uparrow \infty$. In particular we find that

$$b(X_t \mid X_0) = \int_0^t b(\mathrm{d}X_s \mid X_s),$$

such that if a process X satisfies the bound $b(\mathrm{d}X_t)$, then it also satisfies the bound $b(X_t)$. In the remainder of this thesis we will refer to restrictions of the form $b(\mathrm{d}X_t)$ as instantaneous restrictions.

4.2 Coherent risk measure restrictions

In the continuous-time drift-diffusion framework an interesting property arises for instantaneous coherent and lower semi-continuous risk measure restrictions: All risk restrictions can be expressed as restrictions of the discount factor volatility in terms of the risk of standard normal variables, and therefore independent of the model parameters. These restrictions are time-consistent price bounds on all products in the market.

Theorem 2. *Let ψ be a coherent lower semi-continuous risk measure in a market as described in chapter 3. Under the time-consistent pricing restriction*

$$-\psi(X_t/B_t) \le \mathbb{E}[M_t X_t] \le \psi(-X_t/B_t) \quad \textit{for all } t \ge 0, \tag{4.2}$$

where $\mathbb{E}[MX]$ is the price process of the generic payoff process X discounted with SDF M, ψ induces the restriction

$$\|\sigma_M\| \le |\psi(Z)|,$$

where Z is a standard normal random variable.

Remark 3. Note that X is explicitly time-discounted using the risk-free account, because M should value both risk and time, while ψ only values risk.

The intuition behind Theorem 2 is straightforward. When price processes are drift-diffusion processes, the risky part involved with small changes in a price process is approximately normally distributed. Since (1) normally distributed variables scale linearly with volatility and (2) coherent risk measures are positively homogeneous, the instantaneous risk also scales linearly with the volatility. As a result, price restrictions are insensitive to the diffusion of price processes. Since coherent risk measures are translation invariant, also the drift of price processes cannot influence price restrictions. Consequently, price restrictions can be expressed for a generic normally distributed variable. Since M is defined by its drift and diffusion, and since the drift is determined by the bank account, restrictions only need to address the volatility of M.

The inequality in Theorem 2 is only tight in a market-wide sense: In a market where the traded products are not strictly limited to some known set, one can reasonably expect to find products for which the bound is tight. On the other hand, we are free to select a subset of the complete market in which the bound is verifiably slack. The complete market however is of course of limited relevance in pricing contingent claims on illiquid securities.

The inequality $\|\sigma_M\| \le |\psi(Z)|$ works in two ways. First the efficient market yields a minimum value for $\|\sigma_M\|$, effectively restricting the risk measures investors may consider if they want to participate in the market. If an investor's risk preference yields $\psi(Z) < \|\sigma_M\|$ with $\gamma = 0$, that investor considers the market-wide risk/reward tradeoff unfavorable and therefore such investors will retain their wealth in a strict subset of the complete market with the risk/reward tradeoff of their liking. Second, given the efficient market and a feasible risk measure, the inequality in Theorem 2 limits the choice of γ since $\|\sigma_M\|$ is an increasing function of $\|\gamma\|$. This result can intuitively be interpreted as follows: Investors with coherent risk preferences can restrict their expectations about products outside of the complete market based on information retrieved from the complete market and their own beliefs.

Furthermore, the implications of Theorem 2 reach beyond the pricing of contingent claims. The restriction of the pricing kernel affects the valuation of all products in the market, and therefore is related to asset pricing problems where the assets fit our market model and assets are possibly illiquid.

Theorem 2 shows that in continuous time the coherent risk measure restriction can be written in our market model as

$$\|\sigma_M\|^2 = \beta_t^2 \sigma_S' \sigma_S + \gamma_t^2 \mathrm{orth}_{\sigma_S}(\sigma_V)' \mathrm{orth}_{\sigma_S}(\sigma_V) + \delta_t^2 \mathrm{orth}_{\sigma_S}(\sigma_\nu)' \mathrm{orth}_{\sigma_S}(\sigma_\nu) \le \bar{\sigma}_M^2$$

or equivalently

$$\gamma_t^2 \leq \bar{\gamma}_t^2 = \frac{\bar{\sigma}_M^2 - q_S^2(\nu_t\sigma_S'\sigma_S)^{-1} - \delta_t^2 \mathrm{orth}_{\sigma_S}(\sigma_\nu)'\sigma_\nu}{\mathrm{orth}_{\sigma_S}(\sigma_V)'\sigma_V}, \tag{4.3}$$

using the pairwise orthogonality of the risk components $\sigma_S, \mathrm{orth}_{\sigma_S}(\sigma_V)$, and $\mathrm{orth}_{\sigma_S}(\sigma_\nu)$. Here $\bar{\sigma}_M := \psi(Z)$ is independent of the model parameters. This inequality yields two extreme values for γ_t, namely $\bar{\gamma}_t$ and $-\bar{\gamma}_t$. By virtue of the results in section 3.2, these values must correspond with minimum and maximum prices for products attainable under restriction 4.3.

Proof of Theorem 2. Consider a generic scalar drift-diffusion process $\mathrm{d}X_t = \mu(X_t, t)\mathrm{d}t + \sigma(X_t, t)'\mathrm{d}W_t$ and a coherent risk measure ψ. By the definition of Itô processes as the limit of normally distributed increments, as the time increment Δt goes to zero, the increment $\Delta X/\Delta t := (X_{t+\Delta t} - X_t)/\Delta t$ approaches a normally distributed random variable with mean $\mu(X_t, t)$ and variance $\sigma(X_t, t)'\sigma(X_t, t)$. Therefore, using the lower semi-continuity of ψ and the continuity of X_t as a function of t, $\psi(\Delta X/\Delta t)$ will approach $\psi(\mu(X_t, t) + \|\sigma(X_t, t)\|Z)$, where Z is a standard normal random variable.

The instantaneous risk of X_t is thus given by

$$\psi(\mu(X_t, t)\mathrm{d}t + \|\sigma(X_t, t)\|Z\mathrm{d}t) = \|\sigma(X_t, t)\|\psi(Z)\mathrm{d}t - \mu(X_t, t)\mathrm{d}t. \tag{4.4}$$

That is, the instantaneous risk of X is a linear function of the volatility of X. Using the fact that $\mathbb{E}_s[M_s X_s] = X_s/B_s$ and the translation invariance of both coherent risk measures and the expectation operator, equation (4.2) can also be written as

$$-\psi\left(X_{t+s}/B_{t+s} - X_t/B_t\right) \leq \quad \mathbb{E}\left[M_{t+s}X_{t+s} - M_t X_t\right] \leq \psi\left(X_{t+s}/B_{t+s} - X_t/B_t\right).$$

Taking the limit $s \downarrow 0$ to obtain instantaneous restrictions, we find

$$-\psi\left(\mathrm{d}(X/B)_t\right) \leq \qquad \mathbb{E}\left[\mathrm{d}(MX)_t\right] \leq \psi\left(-\mathrm{d}(X/B)_t\right).$$

Suppressing the t-dependence (which will shortly be justified), writing $\sigma_X := \sigma(X_t, t)$ and $\mu_X := \mu(X_t, t)$, the restrictions can now be simplified as

$$\left(\mu_X - r_t/B_t - \|\sigma_X\|\psi(Z)\right)\mathrm{d}t \leq \left(\mu_X - r_t/B_t + \sigma_M'\sigma_X\right)\mathrm{d}t \leq \left(\mu_X - r_t/B_t + \|\sigma_X\|\psi(-Z)\right)\mathrm{d}t$$

or more simply

$$\left|\frac{\sigma_M'\sigma_X}{\|\sigma_X\|}\right| \leq |\psi(Z)|,$$

where we used (3.3), (3.4) and (4.4) in the first step, and the symmetry of Z for $\psi(-Z) = \psi(Z)$ in the second step. Because the choice of X was arbitrary, the above restriction must hold for any σ_X. By Schwartz's inequality, $\sigma_M'\sigma_X \leq \|\sigma_M\|\|\sigma_X\|$. However, if $\sigma_M = \sigma_X$, there is equality. Thus the left-hand side takes on its maximum value when $\sigma_X = \sigma_M$ independent of t, such that the restriction for all coherent risk measures can be expressed as

$$\|\sigma_M\| \leq |\psi(Z)|.$$

□

4.3 Restrictions from general risk measures

For inhomogeneous but cash-invariant risk measures the first steps of the proof of Theorem 2 can be followed to find the inequality

$$|\sigma_M'\sigma_X| \leq |\psi(\|\sigma_X\|Z)|.$$

If all market processes are considered attainable, and therefore all possible values for σ_X are known, then the above inequality can – at the least numerically – be solved for σ_M, from which restrictions on γ can be deduced. If some part of the market is incomplete, we have no choice but to solve the left-hand side for its maximum and the right-hand side for its minimum independently. For a given length $\alpha\|\sigma_M\|$ of σ_X, the left-hand side is still maximized if $\sigma_X = \alpha\sigma_M$, such that the restriction becomes

$$\|\sigma_M\| \leq \frac{|\psi(\|\sigma_X\|Z)|}{\|\sigma_X\|}.$$

The minimum of the right-hand side over $\|\sigma_X\|$ depends on the choice of ψ and the values attainable by the process σ_X. Not all incoherent risk measures will yield sensible bounds due to the loss of dependence between the left- and right-hand side.

Consider for example the convex risk measure $f(X) = \mathrm{Var}[X] - \mathbb{E}[X]$. The right-hand side of the last inequality becomes $|f(\|\sigma_X\|Z)|/\|\sigma_X\| = \|\sigma_X\|$, which attains its minimum 0 if $\|\sigma_X\| = 0$, which holds for $X = B$. This bound is not admissible since it is safe to assume that $\|\sigma_M\| > 0$ in practical situations. An investor who values risk according to f therefore will not enter any risky positions that contain products outside of the complete submarket, contrary to investors with coherent preferences. Generally put, not all investors with incoherent preferences can use their beliefs to restrict their expectations of the products outside of the complete submarket based on the complete market.

Chapter 5

Model analysis

Good deal bound pricing is by definition a subjective (albeit rational) pricing methodology due to its relation with risk measures. Understanding the interplay between parameters is therefore of paramount concern to choosing bound values in line with one's expectations. This chapter relates the model parameters to the price bound spread in section 5.1. In section 5.2 price bounds are discussed for investors that trade products from the complete markets, but cannot apply (near) continuous hedging strategies for practical reasons.

Throughout this chapter we assume that the underlying product of the modeled contingent claims can also be used to form hedging strategies, but with limited liquidity such that perfect hedging strategies are infeasible. This implies that the processes S and V have the same drift, but correlate imperfectly. Unless otherwise noted we fix $\rho_{S,V} = 0.75$.

5.1 Parameter stability

Figure 5.1 shows that the width of the option price bounds is a parabolic function of the correlation $\rho_{S,V}$ between S and V for fixed sizes $\|\sigma_V\|$ of the illiquid asset's volatility. This parameter models how well we can formulate a hedging strategy; a $|\rho_{S,V}|$ close to 1 means that investors can create hedging strategies using S that closely mimic the infeasible Black-Scholes hedging strategy for options on V. On the other hand, a $|\rho_{S,V}|$ close to 0 means that investors must almost purely rely on their expectation that good deals do not prevail within the market to price options on V. As $|\rho_{S,V}|$ decreases ceteris paribus, V will more strongly correlate with the risk component of the SDF priced by γ, and the price bounds will be wider. On the other hand, as $|\rho_{S,V}|$ approaches 1, the correlation between V and the γ_t component of the SDF volatility goes to 0 and the price interval closes. In fact, for a constant multiple c depending on the product being priced, the bound width relates to $\rho_{S,V}$ via $c\sqrt{1-\rho_{S,V}^2}$.

The upper bound attains its maximum when the total volatility of MV, $(\sigma_M+\sigma_V)'(\sigma_M+\sigma_V)$, is maximal. Consider a decomposition of σ_V as $\rho_{S,V}\sigma_S + \bar{\rho}_{S,V}\mathrm{orth}_{\sigma_S}(\sigma_V)$ for the (fixed) vectors σ_S, and $\mathrm{orth}_{\sigma_S}(\sigma_V)$, orthogonal to σ_S. Fixing everything except the correlation between V and S,

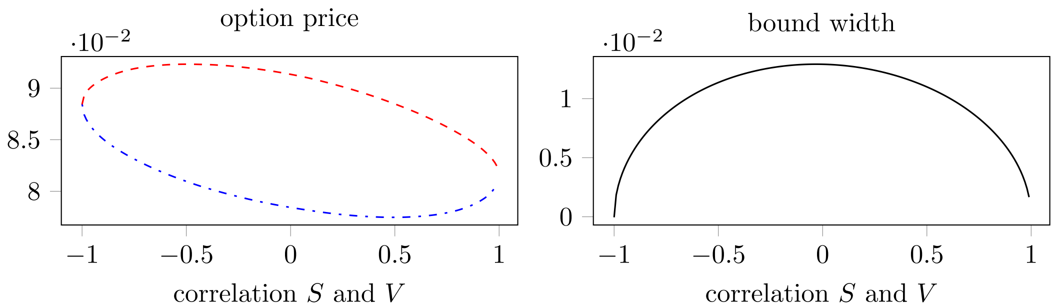

Figure 5.1: Call option price bounds as a function of the correlation between S and V. Option price bounds are average values over 100,000 simulated sample price paths.

$\|\sigma_M + \sigma_V\|$ is maximized when $\sigma_M' \sigma_V$ is maximal. Moreover, from the decomposition and Equations (3.3), (3.4) we can derive that the upper bound attains its maximum at time t when

$$-q_S \rho_{S,V} + \gamma_t \bar{\rho}_{S,V}$$

attains its maximum, which happens at $\rho_{S,V} = -q_S(q_S^2 + \gamma_t^2)^{-1/2}$. Similarly the lower bound attains its lowest value at $\rho_{S,V} = q_S(q_S^2 + \gamma_t^2)^{-1/2}$.

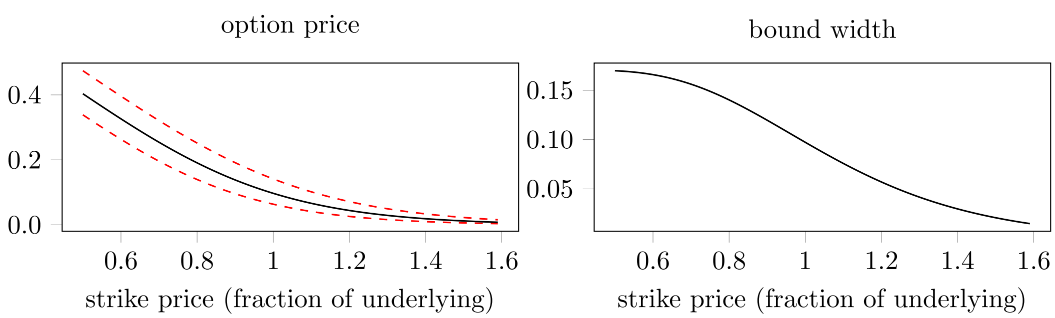

Figure 5.2: Call option price bounds as a function of the strike price. Option price bounds are average values over 100,000 simulated sample price paths.

The logistic relation between bound width and strike price, shown in Figure 5.2, is up to a constant multiple given by the lognormal survival function for call options and by the lognormal cumulative distribution function for puts, both with volatility $\|\sigma_V\|$. As the strike increases, the probability that a call option matures in-the-money follows the lognormal survival function. The price bound decreases accordingly since the region where the option matures in-the-money also drives the option price.

The bound (4.3) denotes a hyperbolic relationship between the parameter γ_t and the maximum volatility of M. Because the value of V_T, $\mathbb{E}_T[MV] = \exp(\int_0^T q_V + \sqrt{\nu_t}\sigma_V'[\beta_t \sigma_S + \gamma_t \mathrm{orth}_{\sigma_S}(\sigma_V)]\mathrm{d}t)$, depends linearly on γ_t, the price bounds on V will also be hyperbolic with respect to the volatility bound. In particular, the price bounds will always fall within the asymptotes $\exp(\int_0^T q_V +$

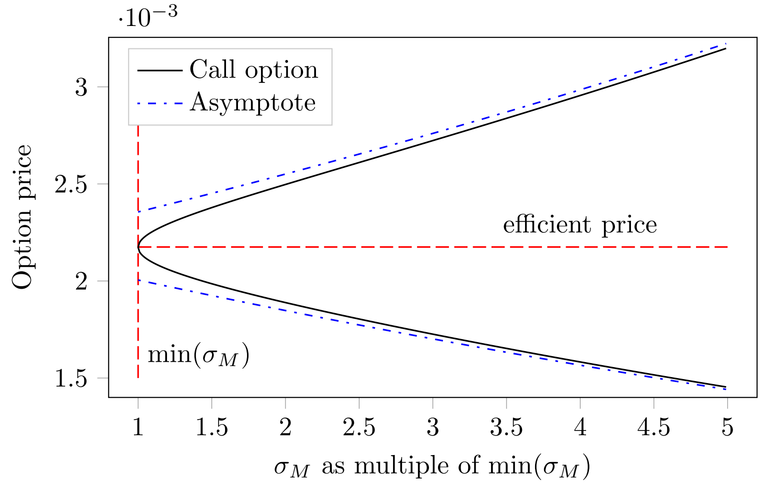

Figure 5.3: Call option price as a function of the volatility bound. The option shown has a strike of 150%. Note that both asymptotes are slightly upward-curving.

$\sqrt{\nu_t}\beta_t\sigma_S'\sigma_V \pm r_t\bar{\sigma}_M\mathrm{d}t)$. The same does not hold for contingent claims on V due to the nonlinear relation between the volatility of M and the option price, although nonlinear asymptotes can be derived by valuing options with $\gamma_t = \bar{\sigma}_M$.

Figure 5.3 shows how the volatility bound influences the option price bound. A call option with a strike of 150% is used because the nonlinearity is (somewhat) more pronounced for high strikes. For reasonable values of σ_M, the asymptote provides a decent approximation of the true option price bounds. This observation is particularly useful during model calibration and exploratory analysis. One can use the asymptote as approximate bounds during the calibration phase, allowing one to approximate good deal bounds before the definitive calibration for the price of volatility risk, or the exact specification of the stochastic interest rate- or stochastic volatility model is known. Moreover, simulated asymptotic discount factors remain valid when the parameters of the interest rate and volatility models are changed, which can be used to reduce the computational load when calibrating option bounds on real-world data.

5.2 Incompleteness through limited market access

In this section we demonstrate the implications of good deal bounds for an investor who has limited trading opportunities. A trader with limited market access cannot apply a continuous hedging strategy and could therefore resort to good deal bound pricing, even when issuers of the option can efficiently hedge their position. In such situations, good deal bounds are a means to quantify the risk an investor is willing to take on a single investment.

This section considers Expected Shortfall at the $100 \times \alpha$-percentile as the price bound generating risk measure, i.e. $\psi_\alpha(X) = \mathrm{ES}_\alpha(X) := \mathbb{E}[-X \mid X < \sup\{k : \mathbb{P}(X < k) < \alpha\}]$. This is a coherent risk measure that is decreasing as a function of α, such that lower values of α generate wider price bounds. An investor who measures risk with ES_α considers a risk-free sum of money as less attractive than a risky payoff if the expected payoff is higher than the sum of money, given that the payoff ends in the worst α-quantile.

Note that ψ expresses the total risk on a position and not the 'hedging risk', the risk that remains after the best feasible hedging strategy is applied. To distill the hedging risk, we would need to be explicit on which hedging strategies are allowed, which is a different question altogether.

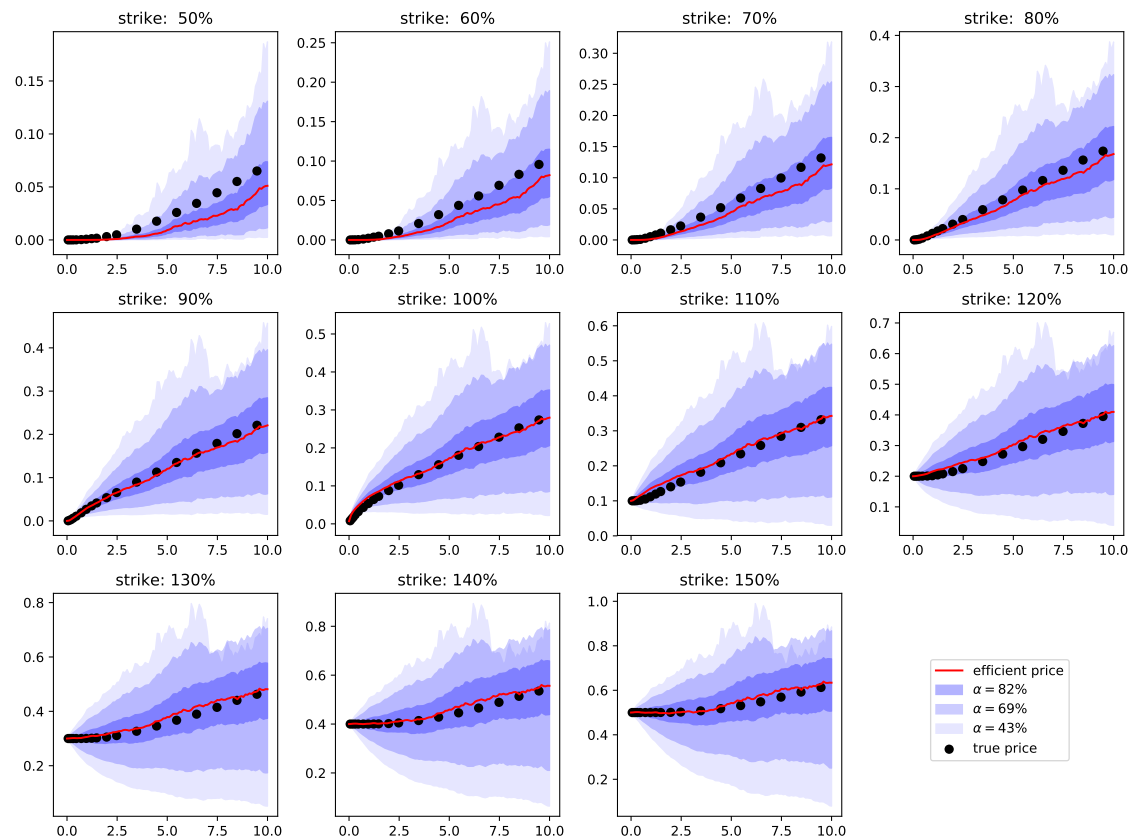

Figure 5.4: European put option prices (y axes) as a function of maturity (x axes) for different strikes. Price bounds based on the Expected Shortfall at α% is shown for different values of α. Note that all y-axes have different scales; the model-market pricing mismatch is around 0.015 for strikes of both 50% and 150% and lower for strikes closer to 100%. The efficient price and the price bounds are estimated from 10,000 Monte Carlo simulated paths. For low α (high σ_M) the simulation variance explodes giving volatile price bound estimates.

Figure 5.4 shows the calibration results of the market model, calibrated on SPX implied volatility quotes for June 30, 2017 from ICE Data Services. To limit the computational complexity the pairwise correlation between interest rates, the volatility process and equity is set to 0. The calibration was performed by optimizing the model parameters over Monte Carlo simulated price paths, interest rate paths and volatility paths using a basin hopping algorithm to limit the risk of finding local minima. The objective function used in the optimization is

$$\sum_{\substack{k\in K\\ t\in T}} f(t,k) := \sum_{\substack{k\in K\\ t\in T}} \frac{(p(k,t)-p_m(k,t))^2}{p_m(k,t)},$$

where K is the set of strike prices, T is the set of maturities and $p(k,t)$ (resp. $p_m(k,t)$) is the efficient model (resp. market-observed) price for strike k and maturity t. This objective function seeks the middle ground between squared absolute pricing errors[1] and squared relative pricing errors[2]. It gives a better fit for out-of-the-money options than the squared absolute pricing error, since large relative errors for out-of-the-money options are often small in absolute terms due to the low price of such options. For the same reason the squared relative pricing error tends to underfit options with a high strike price.

The efficient market values risk as the exprected shortfall at $\alpha = 0.834$, which is denoted in Figure 5.4 by the efficient price. We assume that the inefficient investor can feasibly obtain a correlation of 0.75 between her own hedging strategies and efficient hedging strategies. Under these conditions the price bounds are shown for $\alpha \in \{0.43, 0.69, 0.82\}$ in Figure 5.4. As the strike price increases the price bounds become more pronounced and more easily capture the variation in the observed prices. Put options with low strike prices on average deviate more, measuring deviation as the decrease in α required to capture observed prices within the price bounds. This is largely due to the low value of such options, which leads to a high relative difference between model price and market price.

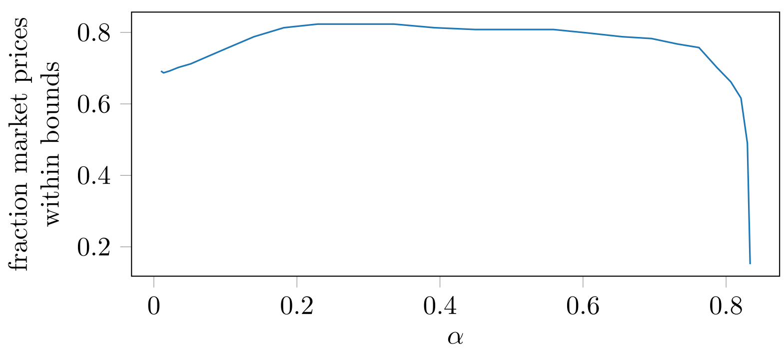

Figure 5.5: The number of market prices within bounds as a function of α, where the bounds are given by $\sigma_M \le \mathrm{ES}_\alpha$. The efficient market values products at most at $\alpha = 0.834$. As α approaches 0, the volatility of the discounted processes explode such that simulated bounds haphazardly cross the true data, causing the fraction of market prices within bounds to decrease.

As α is decreased and the price bound widens, the variance of the stochastic discount factor inherently increases. This increases the noise in the Monte Carlo-estimated price bounds, which can clearly be recognised in Figure 5.4 as the increased ragedness of the $\alpha = .43$ price bound with respect to the $\alpha = .69$ and $\alpha = .82$ price bounds. For high strike prices and long maturities the

[1] The squared absolute pricing error is given by $f(t,k) = (p(k,t) - p_m(k,t))^2$.
[2] The squared relative pricing error is given by $f(t,k) = (p(k,t) - p_m(k,t))^2 / p_m(k,t)$.

upper bound at $\alpha = .43$ even drops below the more tight upper bound at $\alpha = .69$. Such irregularities could be removed by increasing the sample size; however, as α approaches 0, the variance σ_M grows without bounds and with it, the required sample size for a stable estimate for the price bounds.

Another question of interest is how many market observed prices fall within the price bounds for a choice of α. Figure 5.5 shows that values of α close to the efficient market α yield bounds that contain the lion's share of market observed prices. Only very short maturity options or low strike price options fall outside reasonably tight price bounds. As α decreases further, the fraction of market prices captured within the price bounds only marginally increases and for very low values of α the fraction decreases again. The latter is the result of unstable estimates due to the increase in the sample path variance of the stochastic discount factor.

Chapter 6

Concluding remarks

This thesis demonstrates how preferences expressed over risk can be used to obtain option price bounds in a general continuous-time pricing framework that allows for stochastic interest and stochastic volatility, while excluding arbitrage opportunities. The model is formulated as a generalization of continuous-time complete market pricing models with stochastic volatility and stochastic interest such that results are consistent with findings in a complete market setting.

It is shown that, given a risk measure restriction, three discount factors can be readily identified which correspond to the efficient price, the lower price bound and the upper price bound respectively, for all securities. There is therefore no need to numerically search for a minimum or maximum price amongst the range of admissible stochastic discount factors. This result is particularly useful in a Monte Carlo pricing context where risk-neutral scenarios are used to evaluate the price of complex derivatives. Because price bounds are determined by the on beforehand identified stochastic discount factors, the real-world price paths need only be sampled once and can simply be discounted with the appropriate discount factors to obtain price bounds. The additional computational load with respect to complete market Monte Carlo pricing is thus limited to computing the additional discount factors.

The entire class of coherent risk measures is shown to generate bounds that are independent of model parameters. Restrictions from this class limit the total volatility of of the stochastic discount factor by the risk measured on a standard normally distributed variable. An interesting track for future research would be to generalize this result to other risk measures or to other market models – e.g. in the presence of jumps.

Taking a queue from Cochrane and Saa-Requejo (2000) who reversed the Hansen-Jagannathan bound to obtain good deal bounds, the good deal bounds derived in Theorem 2 can be reversed to obtain a lower bound on a measure for market risk. It might be very interesting to compare this measure to other risk- or stress measures of the state of the world. There might be interesting analogies that can be used to choose a specific risk measure or to obtain selection criteria for risk measures eligible for the calculation of good deal bounds.

Determining risk-neutral price processes in an incomplete market Heston model presented an issue that is lacking in the complete market. If excess returns are bounded away from zero (i.e.

always greater than some $\epsilon > 0$) the 'market price of risk' can become arbitrarily large because risk can become arbitrarily small in the Heston model. In the complete market the unbounded market price of risk does not pose a problem because the exposure to risk and the market price of risk cancel each other out by the completeness assumption. This thesis resolved the unboundedness issue by enforcing a minimum market risk in the model. An alternative solution would be to force the excess return of risky assets to scale with the vol-of-vol ν. Since in this case both excess return and risk exposure scale with ν, the market price of risk will be independent of ν. On the other hand, the market price of risk does not wholly explain the excess return in this model, which is why it has not been used in this thesis. A qualitative comparison between different stochastic volatility models suitable for incomplete market valuation would be an interesting pursuit for further research.

The framework presented in this thesis can also be used to explain the bid-ask spread if one assumes that all trades are performed in the incomplete market. This assumption poses the problem that the parameters related to the complete submarket are no longer identified. If this issue can be tackled, the methods employed in this thesis could provide new insights in the behavior of the bid-ask spread.

Bibliography

Acciaio, B. and Penner, I. (2011). *Dynamic Risk Measures*, pages 1–34. Springer Berlin Heidelberg, Berlin, Heidelberg.

Andersen, L. B. (2007). Efficient simulation of the heston stochastic volatility model.

Arai, T. (2017). Good deal bounds with convex constraints. *International Journal of Theoretical and Applied Finance*, 20(02):1750011.

Arai, T. and Fukasawa, M. (2014). Convex risk measures for good deal bounds. *Mathematical Finance*, 24(3):464–484.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.

Back, K. (2010). *Asset pricing and portfolio choice theory.* Oxford University Press.

Bajeux-Besnainou, I. and Rochet, J.-C. (1996). Dynamic spanning: Are options an appropriate instrument? *Mathematical Finance*, 6(1):1–16.

Baxter, M. and Rennie, A. (1996). *Financial calculus: an introduction to derivative pricing.* Cambridge university press.

Bernardo, A. E. and Ledoit, O. (2000). Gain, loss, and asset pricing. *The Journal of Political Economy*, 108(1):144–172.

Björk, T. and Slinko, I. (2005). Towards a general theory of good deal bounds. *Review of Finance*, 10(2):221–260.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Bondarenko, O. and Longarela, I. R. (2009). A general framework for the derivation of asset price bounds: an application to stochastic volatility option models. *Review of Derivatives Research*, 12(2):81–107.

Cheyette, O. (1992). Term structure dynamics and mortgage valuation. *The Journal of Fixed Income*, 1(4):28–41.

Cochrane, J. H. (2009). *Asset Pricing.* Princeton university press.

Cochrane, J. H. and Saa-Requejo, J. (2000). Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, 108(1):79–119.

Davis, M. H. (2006). Optimal hedging with basis risk. In Kabanov, Y., Liptser, R., and Stoyanov, J., editors, *From stochastic calculus to mathematical finance*, pages 169–187. Springer Berlin Heidelberg.

Delbaen, F. (2002). Coherent risk measures on general probability spaces. In Sandmann, K. and Schönbucher, P. J., editors, *Advances in Finance and Stochastics: Essays in Honour of Dieter Sondermann*, pages 1–37. Springer Berlin Heidelberg.

Ekström, E. (2004). Properties of american option prices. *Stochastic Processes and their Applications*, 114(2):265–278.

Černý, A. (2003). Generalised sharpe ratios and asset pricing in incomplete markets. *Review of Finance*, 7(2):191–233.

Floroiu, O. and Pelsser, A. (2013). Closed-form solutions for options in incomplete markets. Technical report, Netspar.

Gikhman, I. I. (2011). A short remark on feller's square root condition. Available at SSRN: https://ssrn.com/abstract=1756450 or http://dx.doi.org/10.2139/ssrn.1756450.

Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262.

Heath, D., Jarrow, R., and Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, pages 77–105.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327.

Hoorens, B. (2011). On the cheyette short rate model with stochastic volatility. (Master's thesis).

Huber, P. J. (1981). *Robust statistics*. Wiley.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300.

Hull, J. and White, A. (1990). Pricing interest-rate-derivative securities. *The Review of Financial Studies*, 3(4):573–592.

Janson, S. and Tysk, J. (2003). Volatility time and properties of option prices. *The Annals of Applied Probability*, 13(3):890–913.

Klöppel, S. and Schweizer, M. (2007). Dynamic utility-based good deal bounds. *Statistics & Risk Modeling*, 25(4):285–309.

Kreps, D. M. (1981). Arbitrage and equilibrium in economies with infinitely many commodities. *Journal of Mathematical Economics*, 8(1):15–35.

Ritchken, P. and Sankarasubramanian, L. (1995). Volatility structures of forward rates and the dynamics of the term structure. *Mathematical Finance*, 5(1):55–72.

Rockafellar, R. T. (1997). *Convex analysis*. Princeton university press.

Romano, M. and Touzi, N. (1997). Contingent claims and market completeness in a stochastic volatility model. *Mathematical Finance*, 7(4):399–412.

Ruszczyński, A. and Shapiro, A. (2006). Optimization of convex risk functions. *Mathematics of Operations Research*, 31(3):433–452.

Rynne, B. P. and Youngson, M. A. (2008). *Linear functional analysis*. Springer-Verlag London.

Schilling, R. L. (2017). *Measures, integrals and martingales*. Cambridge University Press.

Sircar, R. K. and Papanicolaou, G. C. (1999). Stochastic volatility, smile & asymptotics. *Applied Mathematical Finance*, 6(2):107–145.

Xu, M. (2006). Risk measure pricing and hedging in incomplete markets. *Annals of Finance*, 2(1):51–71.

Appendix A

Proofs

Lemma 1. *Let $\psi : \mathcal{L} \to \mathbb{R}$ be a function on a linear space $\mathcal{L}$ that extends $\mathbb{R}$ in some way such that the following statement makes sense: Let ψ be translation invariant: $\psi(X + a) = \psi(X) - a$ for all a in $\mathbb{R}$. ψ is convex if and only if $\mathcal{A}_\psi := \{X \in \mathcal{L} \mid \psi(X) \leq 0\}$ is convex. ψ is positive homogeneous if and only if $\mathcal{A}_\psi$ is a cone.*

Proof. Let $\{X, Y\} \subset \mathcal{A}_\psi$ and $\lambda \in [0, 1]$. If ψ is convex, then

$$\psi(\lambda X + (1 - \lambda)Y) \leq \lambda\psi(X) + (1 - \lambda)\psi(Y) \leq 0$$

by convexity, such that $\lambda X + (1 - \lambda)Y \in \mathcal{A}_\psi$. Conversely, assume that $\mathcal{A}_\psi$ is convex, and note that $\psi(X+\psi(X)) = \psi(Y+\psi(Y)) = 0$. Then by the convexity of $\mathcal{A}_\psi$, $\lambda(X+\psi(X))+(1-\lambda)(Y+\psi(Y)) \in \mathcal{A}_\psi$, hence

$$0 \geq \psi(\lambda(X + \psi(X)) + (1 - \lambda)(Y + \psi(Y))) = \psi(\lambda X + (1 - \lambda)Y) - \lambda\psi(X) - (1 - \lambda)\psi(Y).$$

Reordering gives $\psi(\lambda X + (1 - \lambda)Y) \leq \lambda\psi(X) + (1 - \lambda)\psi(Y)$, proving that ψ is convex.

Now let $\lambda \geq 0$. First assume that ψ is positive homogeneous. Then $\psi(\lambda X) = \lambda\psi(X)$, hence if $X \in \mathcal{A}_\psi$, it must be that

$$0 \geq \lambda\psi(X) = \psi(\lambda X),$$

and therefore $\lambda X \in \mathcal{A}_\psi$, which in turn implies that $\mathcal{A}_\psi$ is a cone.

On the other hand, when $\lambda X \in \mathcal{A}_\psi$ for each $X \in \mathcal{A}_\psi$, then by similar logic, again $X+\psi(X) \in \mathcal{A}_\psi$ and therefore

$$0 \geq \psi(\lambda(X + \psi(X)) = \psi(\lambda X) - \lambda\psi(X), \tag{A.1}$$

which shows

$$\psi(X) = \psi(\lambda^{-1}\lambda X) \leq \lambda^{-1}\psi(\lambda X) \leq \lambda^{-1}\lambda\psi(X) = \psi(X),$$

from which equality in (A.1), and with it positive homogeneity follows. □

Lemma 2. *If ψ is a coherent risk measure on $\mathcal{L}^2$, then its convex conjugate ψ^* acts as an indicator function on the set*

$$L := \{Y \mid \mathbb{E}[XY] \leq \psi(X) \textit{ for all } X \in \mathcal{L}^2\}.$$

This implies in particular that

$$\psi(X) = \sup_{Y \in \mathcal{L}^2} \{\mathbb{E}[XY] - \psi^*(Y)\} = \sup_{Y \in L} \{\mathbb{E}[XY]\}. \tag{A.2}$$

Proof. Let $\lambda \geq 0$. We consider three cases: $\psi^*(Y)$ is positive, negative or zero. First, let Y such that $\psi^*(Y) > 0$. Then by the properties of the supremum there exists an X and an $\epsilon > 0$ such that $\mathbb{E}(\lambda XY) - \psi(\lambda X) = \lambda(\mathbb{E}[XY] - \psi(X)) = \lambda\epsilon$. But then $\psi^*(Y) \geq \lambda\epsilon$ for all $\lambda \geq 0$, and thus $\psi^*(Y) = \infty$.

If on the other hand $\psi^*(Y) < 0$, then the preceding holds for an $\epsilon < 0$, in which case $0 > \psi^*(Y) \geq \lambda\epsilon$ for all $\lambda > 0$, which yields the contradiction $\psi^*(Y) = 0$. This shows that $\psi^*(Y) \in \{0, \infty\}$ for all Y.

Then, note that if $\psi^*(Y) \leq 0$,

$$0 \geq \psi^*(Y) \geq \mathbb{E}[XY] - \psi(X)$$

such that $\psi(X) \geq \mathbb{E}[XY]$, implying that $Y \in L$. Conversely, $\psi^*(Y) > 0$ implies $\mathbb{E}[XY] > \psi(X)$ for some X such that $Y \notin L$. This demonstrates that $\psi(Y) = 0$ for all $Y \in L$, and ∞ otherwise. Clearly, the left supremum in Equation (A.2) is not obtained where $\psi^*(Y) = \infty$, validating the restriction from $\mathcal{L}^2$ to L. □

Lemma 3. *Let $\{\psi_t\}_{t\in[0,\infty)}$ be a dynamic risk measure $\{\mathcal{L}_t^2\}_{t\in[0,\infty)} \to \{\mathbb{R}_t\}_{t\in[0,\infty)}$. The following characterizations of time consistency are equivalent:*

(monotonicity) $\psi_{t+s}(X) \geq \psi_{t+s}(Y) \Rightarrow \psi_t(X) \geq \psi_t(Y)$ *for all* $t, s \geq 0$ *and* $X, Y \in \mathcal{L}^2$*,*

(recursive) $\psi_t(-\psi_{t+s}(X)) = \psi_t(X)$ *for all* $t, s \geq 0$ *and* $X \in \mathcal{L}^2$*.*

Proof. First, assume that $\{\psi_t\}_{t\in[0,\infty)}$ has the recursive property. Because ψ_t is a risk measure for all t, it is monotone, and therefore $\psi_{t+s}(X) \geq \psi_{t+s}(Y)$ implies

$$\psi_t(X) = \psi_t(-\psi_{t+s}(X)) \geq \psi_t(-\psi_{t+s}(Y)) = \psi_t(Y).$$

Now assume that $\{\psi_t\}_{t\in[0,\infty)}$ has the monotonicity property. Set $Y = -\psi_{t+s}(X)$. Then by cash invariance, $\psi_{t+s}(Y) = \psi_{t+s}(X)$, and with the monotonicity property follows $\psi_t(X) = \psi_t(Y) = \psi_t(-\psi_{t+s}(X))$. □