Building a Consistent Pricing Model from Observed Option Prices

Jean–Paul Laurent

Center for Research in Economics and Statistics, Finance Department, 15 Boulevard Gabriel Péri, 92245 Malakoff Cedex, France, email: jpl@ensae.fr

Dietmar P.J. Leisen

Stanford University, Hoover Institution, Stanford, CA 94305, U.S.A., email: leisen@hoover.stanford.edu

Abstract

This paper constructs a model for the evolution of a risky security that is consistent with a set of observed call option prices. It explicitly treats the fact that only a *discrete* data set can be observed in practice. The framework is general and allows for state dependent volatility and jumps. The theoretical properties are studied. An easy procedure to check for arbitrage opportunities in market data is proved and then used to ensure the feasibility of our approach. The implementation is discussed: testing on market data reveals a U–shaped form for the "local volatility" depending on the state and, surprisingly, a large probability for strong price movements.

Key words: Markov Chain, no–arbitrage, cross–entropy, model risk
JEL classification: C51, G12, G13

[1] *Date: January 1998. This version: 3 June 1999.* We are grateful to Marco Avellaneda, Phelim Boyle, Peter Carr, Emmanuel Derman, Christian Henning, Kenneth Judd and Dieter Sondermann for helpful discussions; also, participants in seminars at the University of Bonn, the Humboldt University in Berlin, the London School of Economics, in the Bonn–Aarhus workshop at the University of Bonn, Goldman Sachs, New York University, and at the University of Waterloo offered useful comments. The second author acknowledges financial support from the Deutsche Forschungsgemeinschaft, Sonderforschungsbereich 303 at Rheinische Friedrich–Wilhelms Universität Bonn and a Deutscher Akademischer Austauschdienst scholarship through HSP III (a joint program by the Federal and State Governments in Germany).

1 Introduction

The connection between the absence of arbitrage and equivalent martingale measures, first noted by Harrison and Kreps (1979), has led to a deep understanding of the pricing of derivatives *depending* on the proposed model. Models extending the now classical Black–Scholes setup become more and more sophisticated, and allow for time/state dependent and stochastic volatility or even for stock price jumps. These address "model risk" in the presence of "crashes," "smiles," "volatility surfaces" and other empirical shortcomings in the data (see Bakshi, Cao, and Chen (1997)).

In this paper we are interested in choosing the correct model from a second–level perspective. Such an approach is known by practitioners as "reverse engineering" or as "inverse problems" (see Chriss (1997)). Today's actively traded plain vanilla options can be seen as "correctly" priced. We take them as inputs to infer how the market prices risk. Here, we are interested in a model which explains the *observed* plain vanilla option prices in the market *and* is immediately *implementable* in practice to price other derivatives. Instead of estimating a continuous time model that will be further discretized to price complex derivatives, we think of a direct estimation of discrete models and restrict ourselves to those where time and asset space are discrete.

Here, we suppose that discounted asset prices can be modeled by a discrete Markov chain. This approach first appeared in a paper by Zipkin (1993) for the valuation of mortgage–backed securities. It has recently been used in a series of papers related to credit rating instruments (Jarrow, Lando, and Turnbull (1997) and Duffie and Singleton (1998)). A Markov chain approach is sufficiently flexible to handle a wide class of different models; it is a natural approximation of models with time and state dependent volatility and jumps (see Kushner and Dupuis (1992)). It is also a generalization of binomial and trinomial trees.

Breeden and Litzenberger (1978) were the first to infer information from option prices: the risk–neutral marginal density in terms of the second derivative of option prices with respect to the strike. Jackwerth and Rubinstein (1996), Dumas, Fleming, and Whaley (1998), Buchen and Kelly (1996), and Melick and Thomas (1997), among others, were interested in the practical implementation of this approach. More recently the problem to infer the dynamics of the underlying has been addressed: Dupire (1996) and Derman and Kani (1995) relate the local volatility of the risky asset price to partial derivatives of the option prices; Rubinstein (1994) constructed a binomial process consistent with observed prices at *one* date; Carr and Madan (1998) addressed this problem in a model with time and state dependent volatility.

Derman and Kani (1995) consider consistent binomial trees and determine both nodes and probabilities implicitly. Unfortunately their algorithm needs to negative probabilities. They claim this deficiency is due to arbitrage opportunities present in the data and override them by some artificial value (see Barle and Cakici (1995) for a discussion). In Derman, Kani, and Chriss (1996) a grid for a trinomial grid is ad hoc specified. However, even in this case, it may be impossible to fit all market prices, no matter what choice of state space is made, see Derman, Kani, and Chriss (1996), p. 18. This is due to the fact, that binomial and trinomial models are approximations to diffusion processes and will therefore have difficulties to handle option prices in the presence of jumps or large volatility.

The assumption of all these approaches is that call option prices are available, whatever strike and exercise date. In practice a continuous data set is usually obtained by numerical interpolation from a discrete set of observed option prices. As argued by Shimko (1993) special care must be taken in order to get proper probability densities. Also these approaches rely heavily on the choice of the smoothing technique.

In line with our view, Avellaneda, Friedman, Holmes, and Samperi (1997) assume only a discrete data set. However, their assumption of the underlying model in the form of a continuous one–dimensional diffusion seems to be too restrictive (see Dumas, Fleming, and Whaley (1998)). Also, in the diffusion setup of Avellaneda, Friedman, Holmes, and Samperi (1997) jumps are not included. However, their empirical works show that the local volatility often hits minimal and maximal levels which are set ad hod. The local discrepencies of the volatility show the difficulties of a continuous sample path setting. Our major criticism is that all previous approaches assume that the risky security price can be modeled through a scalar diffusion process, *excluding* jumps.

We assume, as did the previous authors, that markets are information efficient in the sense that there are no riskless profits for free, and use this "no–arbitrage" principle to price derivatives. However we have to guarantee that this assumption actually holds in the data. It is a well–known and necessary condition for absence of arbitrage that the option prices are decreasing and convex in the exercise price *and* are non–decreasing with the exercise date. A contribution of this paper is to prove the converse implication. This is a simple check for the presence of arbitrage opportunities in data sets.

We extend the notion of Arrow–Debreu securities to our dynamic case and show that the dynamic market can be transformed to a static one by introducing extra assets that correspond to the martingale restriction. This allows us to relate the "no–arbitrage" principle to the feasibility of our approach. Together with our characterization for the absence of arbitrage opportunities, this can be easily used to check that our approach is valid. Our analysis extensively uses superreplication ideas. Taking into account the consistency constraints with observed option prices narrows the bid–ask spread in our case compared to Cvitanic and Karatzas (1993), El Karoui and Quenez (1995), and Cvitanic, Pham, and Touzi (1998). Moreover our direct examination of the underlying probability structure allows for a thorough analysis of arbitrage violations and explicitly prevents the arbitrage problems, resulting from a simple implementation of the approach of Dupire (1996).

Finally, we address in detail the implementation of our approach. Once the feasibility of our approach has been ensured, we apply Bayesian estimation to infer a unique Markov Chain. We prove that standard results in the static case (see, e.g., Buchen and Kelly (1996)) still apply in our discrete framework. This leads to a simple expression of the risk–neutral probabilities and makes the implementation straightforward.

In FX markets, using a parametric approach, Bates (1996a), and Bates (1996b) find a jump–component whose intensity is small and does not depend on the exchange rate level. McCauley and Melick (1996) and Campa, Chang, and Reider (1997) document that in these markets skewness is strongly correlated with the spot rate. Looking at the data set of Avellaneda, Friedman, Holmes, and Samperi (1997), we find a U–shaped form for the "local volatility" and evidence that jumps are necessary to explain the data appropriately.

The paper is organized as follows. Section 2 introduces our setup and basic notation: it discusses joint, conditional and marginal probabilities, risk–neutral Markov chains, and their consistency with the set of observed option prices. Also, it considers the choice of a martingale measure consistent with option prices in a dynamic market. Section 3 studies tests for arbitrage opportunities, risk–neutral measure, and consistent superreplication prices in the static market case. Section 4 shows how the dynamic market can be transformed to a static one and discusses the set of attainable assets, dynamic arbitrage opportunities and their relation to superreplication. Section 5 presents a tractable way to test for the absence of arbitrage opportunities in the observed set of option prices. Section 6 implements our approach using the data set of Avellaneda, Friedman, Holmes, and Samperi (1997). The paper concludes with section 7. Proofs are postponed to the appendix.

2 The Market Model

The basis of the paper is today's observation of the future contract and a set $\mathcal{D} = \{(\tau_l, K_l, C_l) | l = 1, \dots, L\}$ of observed option contracts written on asset S. Here C_l is the price of the call option with maturity τ_l; i.e., the contract paying $(S_{\tau_l} - K_l)^+$ at time τ_l if S_{τ_l} is the the price of the asset at that time.[2] We assume that the law of one price holds: for each K there is at most one $(t, K, C) \in \mathcal{D}$. We will also denote by $\mathcal{K}_t = \{K|(t, K, C) \in \mathcal{D}\}$ the set of strikes at date t; by $\mathcal{D}_t = \{(\tilde{K}, \tilde{C})|(\tilde{\tau}, \tilde{K}, \tilde{C}) \in \mathcal{D}, \tilde{\tau} = t\}$ the set of contracts with maturity t; by $\mathcal{D}_K = \{(\tilde{t}, \tilde{C}) \in \mathcal{D}|(\tilde{\tau}, \tilde{K}, \tilde{C}) \in \mathcal{D}, \tilde{K} = K\}$ the set of contracts with strike K; and by $\mathcal{T} = \{t|(t, K, C) \in \mathcal{D}\}$ the set of trading dates. To set apart interest rates, we assume that all observed prices are discounted, taking a bond as numeraire.

We adopt a Markov chain as our market model describing the evolution of the asset. Here, for the states of the Markov chain, we require that for any observed strike at t there is a node which is identical to it. We denote by Σ_t the set of nodes at date t. There are certainly many different ways to choose the nodes in relation to strikes. This does, however, not affect the results in this paper. Our choice is just a very simple one to keep track of them. Unless otherwise explicitly noted, $\mathcal{K}_t = \Sigma_t$ holds. The *asset grid* will be denoted $\mathcal{A} = \bigotimes_{t \in \mathcal{T}} \Sigma_t$. A *derivative asset* $\mathcal{X}$ is a random variable on $\mathcal{A}$.

Inferring the probability is our objective and will be discussed in the sequel of the paper. It is given through a sequence $\Pi = (\Pi_t)_{t \in \mathcal{T} \setminus T}$, where Π_t is a *stochastic matrix* on $\Sigma_t \times \Sigma_{t+1}$, i.e. a $|\Sigma_t| \times |\Sigma_{t+1}|$ matrix where elements are nonnegative and rows sum to one. It is also called a *transition matrix*, since it describes the dynamics of the risky asset process $S = (S_t)_{t \in \mathcal{T}}$ on $\mathcal{A}$ through the probability measure P^Π defined by

$$P^\Pi[S_{t+1} = s | S_t = s'] = \Pi_t(s, s') \ .$$

In practice $\Sigma_0 = \{S_0\}$, since today's asset price is observed, and Π_0 is a degenerated stochastic matrix. Nevertheless we introduce it for simplicity. Any time $t \in \mathcal{T}$, nodes are ordered, and we index them either by their corresponding number or directly by their corresponding node.

Pricing measures which are risk–neutral give rise to viable pricing models (in the Harrison–Kreps sense). They are characterized by the fact that the discounted (by numeraire price) asset price $(S_t)_t$ is a martingale under P^Π, i.e.

$$E^\Pi[S_{t'} | S_t = s] = s \quad \text{for all} \quad s \in \Sigma_t \quad \text{and} \quad t, t' \in \mathcal{T}, t > t' \ . \tag{1}$$

We are interested in those Markov chains which fulfill (1) and which fulfill for all $(t, K, C) \in \mathcal{D}$:

$$E^\Pi[(S_t - K)^+] = C \ .$$

[2] We could allow for put options. However by put–call parity all put options in the data set could be replaced by their corresponding call prices. So we adopt this assumption further without loss of generality.

We call them a $\mathcal{D}$ *Markov chain pricing model* ($\mathcal{D} - MCPM$ in short).

There are some interesting interpretations which will be useful in the sequel. The set of all derivative assets is a linear state space that can be identified with $\mathbb{R}^{|\mathcal{A}|}$, where its dimension is the number of linearly independent paths. The Euclidean basis of the state space can be seen as the set of Arrow–Debreu securities, i.e. the securities paying one unit of numeraire conditionally on the realization of some path. Moreover we can introduce the probability P^{Π} over sample paths by

$$P^{\Pi}(S_1 = \bar{s}_1, \dots, S_T = \bar{s}_T) = \prod_{t=0}^{T-1} \Pi_t(\bar{s}_t, \bar{s}_{t+1})$$

for any $(\bar{s}_1, \dots, \bar{s}_T) \in \mathcal{A}$, with $\bar{s}_0 \equiv S_0$. For $t, t' \in \mathcal{T}$, a stochastic matrix $\Pi_{t,t'}$ that describes the transition between dates t and t' can be defined by

$$\Pi_{t,t'} = \prod_{u=t}^{t'-1} \Pi_u \ .$$

These probabilities have a simple financial interpretation, which will be useful in the sequel: $P^{\Pi}(S_1 = \bar{s}_1, \dots, S_t = \bar{s}_T)$ can be viewed as the price of an Arrow–Debreu security paying one unit of numeraire conditionally on the realization of the path $(\bar{s}_1, \dots, \bar{s}_T)$. Similarly, entry $\Pi_{t,t'}(i,j)$ can be viewed as the price at time t in state i of an asset paying one unit of numeraire at time t' in state j or as the transition matrix between t and t'. In particular we will further denote by $\lambda_t = \Pi_{0,t}(S_0, \cdot)'$ the marginal distribution at date t. The i–th component of λ_t can be interpreted as the price of an asset paying one unit of numeraire if, and only if, the discounted price S is equal to s_i at time t (i.e. a path–independent Arrow–Debreu security).

For $T = |\mathcal{T}|$ dates, N states and a homogeneous grid, there are usually $N(2T-1)$ linear constraints, apart from the positivity constraints, and $N + (T-1)(N^2 - N)$ unknowns. The positivity constraint makes existence a difficult issue, especially since it has to be related somehow to a "no–arbitrage" condition on the data–set. We will study this issue through sections 4 and 5.

However, once the feasibility of our approach has been ensured, there is typically a multiplicity of $\mathcal{D}$–MCPM's. The Bayesian approach takes a prior, e.g. the Black–Scholes setup, and looks for the minimal departure from this model consistent with the observed option prices. In the financial literature, L^2–criteria have been used by Rubinstein (1994) and Jackwerth and Rubinstein (1996) among others for pricing options. On the statistical side, these criteria appear in Hansen and Jagannathan (1997) and Luttmer (1997). The cross–entropy criterion has been used by Buchen and Kelly (1996), Jackwerth and Rubinstein (1996), and Avellaneda, Friedman, Holmes, and Samperi (1997). Also known as the information optimization criterion, it is well established in probability theory and the statistical literature for its superiority in filling in missing information. We will therefore adopt it here.

We study the case of all dates by looking at each path over time as a different state. We assume that they are in some linear order and index them by i, $i \in \{1, \dots, \prod_{t \in \mathcal{T}} |\Sigma_t|\} \hat{=} \mathcal{A}$. Let us denote by $\mathcal{M}$ the set of all probability measures on the state space $\mathcal{A}$ associated with an MCMM and by $\mathcal{M}_{\mathcal{D}}$, the subset of those probability measures associated with a $\mathcal{D}$–MCPM. We are given an a priori probability measure $P \in \mathcal{M}$, allowing explicitly for the fact that P might not be consistent with the observed option price.

The cross–entropy of a probability measure $Q \in \mathcal{M}$ is defined by

$$H(Q) = E^P \left[\frac{dQ}{dP} \log \frac{dQ}{dP} \right] .$$

Here we are interested in the cross–entropy problem

$$\min_{Q \in \mathcal{M}_{\mathcal{D}}} H(Q) \ , \tag{2}$$

which corresponds to the optimization problem under the constraint to reproduce the calls in the observed data set.

Proposition 1 *Let us assume that $\mathcal{M}_{\mathcal{D}} \neq \emptyset$ and that there is $Q_0 \in \mathcal{M}_{\mathcal{D}}$, $Q_0 \sim P$; then the minimal cross-entropy $\mathcal{D}$-MCPM exists, and the associated probability measure Q^* is equivalent to P and is uniquely characterized by*

$$\frac{dQ^*}{dP} = \exp \left\{ \mu_0 + \sum_{t \in \mathcal{T}} \left[\lambda_t(S_t)(S_{t+1} - S_t) + \sum_{(t,K,C) \in \mathcal{D}} \mu_{t,K} (S_t - K)^+ \right] \right\} \ ,$$

where the parameters μ_0, $\mu_{t,K}$ (for $(t, K, C) \in \mathcal{D}$) and the functions $\lambda_t : \Sigma_t \to \mathbb{R}$ (for $t \in \mathcal{T} \setminus T$) are determined by

$$\begin{aligned} E^{Q^*}[1] &= 1 \ , \\ E^{Q^*}[S_{t+1} - S_t | S_t] &= 0 \quad \forall t \in \mathcal{T} \setminus T \ , \\ \textit{and} \quad E^{Q^*}\left[(S_t - K_{i,t})^+\right] &= C \quad \textit{for any } (t, K, C) \in \mathcal{D} \ . \end{aligned}$$

Existence of such a positive Q_0 follows: e.g., if all marginal distributions are strictly positive and the conditions of the previous section are fulfilled, we can simply redistribute part of the transition probability mass appropriately for any Q to ensure its strict positivity.

3 The Static Market

This section studies both the transition between today and a *fixed* date $t \in \mathcal{T}$ and then the market with $\mathcal{D} = \mathcal{D}_t$ and $\mathcal{T} = \{0, t\}$. We call this the *static market*, as there is no conditional dynamics in the future. Many of the results here are well known. However we recall them here to introduce the main concepts useful in the dynamic case in an easy framework.

A $\mathcal{D}_t$–MCPM in this market is completely defined by a probability measure λ_t on Σ_t (i.e. $\lambda_t(s) \geq 0$ and $\sum_{s\in\Sigma_t} \lambda_t(s) = 1$) that is consistent with the prices of traded assets.

Property 2 *The observed option prices are decreasing and convex, i.e. for any (C, K), (C', K'), $(C'', K'') \in \mathcal{D}_t$ with $K'' > K' > K$:*

$$\frac{C - C'}{K - K'} < \frac{C - C'}{K' - K''} \ . \tag{3}$$

Please note that the property of decreasing call prices ensures positive option prices, since the call price at the highest nodes s_t^{max} was set equal to zero. For an inner node $K \in \Sigma_t \setminus \{s_t^{min}, s_t^{max}\}$, let us consider a so–called *butterfly spread* with payoff $B(\cdot)$ based on the three adjacent traded strikes $K_- < K < K_+$ ($K_- = \max_{(k,C)\in\mathcal{D}_t, k\leq K} k$, $K_+ = \min_{(k,C)\in\mathcal{D}_t, k\geq K} k$) and a payoff equal to one at K, i.e. $B(K) = 1$ and 0, otherwise. For $K = s_t^{min}$ and $K = s_t^{max}$, we consider call spreads. The positive payoff of the butterflies at inner nodes has the price

$$\frac{1}{K - K_-} C_- - \left(\frac{1}{K - K_-} - \frac{1}{K_+ - K} \right) C + \frac{1}{K_+ - K} C_+ \ .$$

This translates directly into property 2. If there is a complete set of observed option prices, i.e. $\Sigma_t = \mathcal{K}_t$, these butterflies form a basis of the payoff space: it is then straightforward to prove the following lemma 3. Otherwise, there will be nodes for which no option contract can be observed. Then we can introduce dummy call options for the missing strikes, with a price equal to the linear interpolation of the adjacent ones.

Lemma 3 [3] *There exists a risk–neutral marginal density λ_t at date t if, and only if, property 2 is fulfilled.*
Moreover, in a complete market with property 2, the prices $\lambda_t(s), s \in \Sigma_t$ of path–independent Arrow–Debreu securities with payment date t are uniquely

[3] This is a discrete analog of the results of Breeden and Litzenberger (1978), since the difference of the fractions is a natural approximation to the second derivative of call options with respect to the strike.

determined. For any $s \in \Sigma_t \setminus \{s_t^{min}, s_t^{max}\}$, we have

$$\lambda_t(s) = \frac{C_+ - C}{K_+ - K} - \frac{C - C_-}{K - K_-} \ .$$

λ_t^{min} and λ_t^{max} have a more complicated form due to boundary effects which we do not exhibit here. So property 2 allows check for arbitrage opportunities and for the existence of a risk–neutral measure λ_t in this static market.

When there are fewer traded strikes than nodes, there is not a unique price at which a call option with strike $K \in \Sigma_t \setminus \mathcal{K}_t$ can be consistently priced without introducing arbitrage opportunities. Similarly λ_t is not uniquely determined. However, it is straightforward to prove the following linear interpolation bounds

Lemma 4 *In the static case the superreplication price of a call option with strike $K \in \Sigma$ and maturity $t \in \mathcal{T}$ is given by the linear interpolation*

$$\frac{K_+ - K}{K_+ - K_-} C_- + \frac{K - K_-}{K_+ - K_-} C_+ \ .$$

The associated probability puts non–zero probability weights only on traded strikes and zero elsewhere.

This superreplication price is very different from that obtained in a standard stochastic volatility model, where it is equal to the trivial price S_0. The departure is due to the use of traded options; consistency with these option prices restricts the set of risk–neutral measures and, thus, the no–arbitrage bounds on calls are narrowed.

Remark 5 *In the case of a call option payoff, we have an explicit characterization of the superreplicating portfolio and price. This can be extended to the case of any concave payoff $\mathcal{X}$. The superreplicating portfolio is obtained as the linear interpolation of points $\{(K, C) | (C, t, K) \in \mathcal{D}\}$.*

4 The Dynamic Market and Arbitrage Opportunities

This section transforms the dynamic market into a static one. The martingale condition corresponds to constraints on assets we introduce. We also describe the linear subspace of all path dependent payoffs made of attainable claims by static or dynamic strategies. We restrict ourselves here to the case of a market with $\mathcal{T} = \{0, 1, 2\}$.

4.1 Transforming to a fictitious static market

Let us now introduce our *path–dependent static market.* It is a fictitious market in which the states are $\Sigma_1 \times \Sigma_2$. The payoff in state $(s, s') \in \Sigma_1 \times \Sigma_2$ depends on the joint occurrence of state s at date 1 and state s' at date 2. The payoff structure of all assets is then a matrix. The basic securities are the following: For a pair $(s, s') \in \Sigma_1 \times \Sigma_2$ let us consider the *path–dependent Arrow–Debreu security* $\delta^{dep}_{s,s'}$, paying 1 unit of numeraire at date 2 if the asset is in state i at date 1 *and* in state j at date 2, and 0 otherwise. Their price will be denoted by $P_{s,s'}$. This can be interpreted as the probability of occurrence for that specific path.

In the fictitious market, the (standard) *path–dependent* Arrow–Debreu security δ^1_s at date 1, paying 1 unit in state s is described by the payoff–matrix

$$\delta^1_s(\bar{s}_1, \bar{s}_2) = \begin{cases} 1 & \text{if } \bar{s}_1 = s \\ 0 & \text{otherwise} \end{cases} ,$$

i.e. has 1's in row s and 0's in all other rows. Similarly the (standard) *path–dependent* Arrow–Debreu security $\delta^2_{s'}$ at date 2, paying 1 unit in state s' is described by the payoff–matrix 1's in column s' and 0's in all other columns.

Let us further introduce for any $s \in \Sigma_1$ the contract σ_s between two parties: if the asset price is in state s at date 1, both exchange a bond with face value s and an asset, where it is specified that the "holder" receives the bond. At date 2 the option seller buys the asset back from the holder at the then–prevailing market price. This is just a trading strategy where the bond is used to transfer the payments such that there are neither initial nor intermediate payments. So in the absence of arbitrage, the price of the contract must be 0. This introduces $|\Sigma_1|$ new assets.

These conditions on the price of assets σ_s are actually sufficient for a probability measure to be a martingale measure. This can be seen as follows: The payoff at date 2 of contract σ_s is $(S_2 - s)1_{S_1=s}$. Its price can be rewritten as

$$\sum_{s\in\Sigma_2} P_{s,s'} \cdot (s' - s) = 0 \ , \tag{4}$$

which is precisely the martingale condition (1).

Definition 6 *The* set of investment strategies *is* $\mathcal{I} = I\!R^{|\Sigma_2|} \times I\!R^{|\Sigma_1|} \times I\!R^{|\Sigma_1|}$. *For any strategy* $(\alpha, \beta, \gamma) \in \mathcal{I}$, *the claim paying* $\alpha_{s'} + \beta_s + \gamma_s s'$ *in state* (s, s') *will be denoted by* $\mathcal{X}(\alpha, \beta, \gamma)$ *and its price at date* 0 *by* $\mathcal{P}(\alpha, \beta, \gamma)$.

Proposition 7 *Let us assume that we observe a set of options with* $\Sigma_1 = \mathcal{K}_1$

and $\Sigma_2 = \mathcal{K}_2$. The set of attainable claims through investment in calls, puts, and the asset and numeraire is the linear subspace

$$\{\mathcal{X}(\alpha,\beta,\gamma) | (\alpha,\beta,\gamma) \in \mathcal{I}\} \subset I\!R^{|\Sigma_1 \times \Sigma_2|} ,$$

i.e., equivalent to investment in $\delta^{dep}_{s,s'}$, δ^1_s and $\delta^2_{s'}$ ($(s,s') \in \Sigma_1 \times \Sigma_2$). In the absence of arbitrage opportunities, the price of the claim $\mathcal{X}(\alpha,\beta,\gamma)$ is uniquely determined by

$$\mathcal{P}(\alpha,\beta,\gamma) = \sum_{s\in\Sigma_1} (\beta_s + \gamma_s s)\lambda_1(s) + \sum_{s\in\Sigma_2} \alpha_s \lambda_2(s) \ .$$

This characterizes the claims attainable through static investments in traded options in more detail: any attainable claim can be replicated by buying $\alpha_{s'}$ units of the path–independent Arrow–Debreu security s' at date 2, β_s units of the path–independent Arrow–Debreu security s at date 1, and γ_s units of stocks conditional on being in state i at date 1. The set of attainable claims in the dynamic market is thus equal to the set of attainable claims in the static market where the extra assets have been introduced. In the sequel we will always adopt the static viewpoint to the dynamic problems.

Above we explained that the martingale condition translates into a consistency condition on the prices of these extra assets. It is well known that the existence of a state price density in the fictitious market consistent with $\mathcal{D}$ and additional assets is equivalent to the absence of arbitrage opportunities in the static market. This is a standard result in financial theory which can be found in any textbook (see, e.g., Duffie (1992), ch. 1). We have used this approach in the case of static markets. In the two–period case studied in this section, all probabilities in the fictitious market can be associated with a Markov chain: dependence on a path is exactly the dependence on the state at date 1. This implies the following:

Proposition 8 *Existence of a $\mathcal{D}$–MCPM is equivalent to no–arbitrage in the fictitious static market.*

4.2 α–Arbitrage Opportunities

The absence of arbitrage opportunities is equivalent to the positivity of the linear operator $\mathcal{P}$, i.e., the optimization problem

$$\begin{array}{ll} \min\limits_{\alpha,\beta,\gamma\in\mathcal{I}} & \mathcal{P}(\alpha,\beta,\gamma) \\ \text{s.t.} & \mathcal{X}(\alpha,\beta,\gamma) \geq 0 \end{array}$$

has a nonnegative minimum. Since we have a priori excluded straightforward static arbitrage opportunities, the only way an arbitrage opportunity can occur comes from a dynamic trade between 1 and 2. We split up this problem. First, we fix an arbitrary payoff $\alpha \in \mathbb{R}^{|\Sigma_2|}$ at date 2, and then we study the superreplication of α (i.e., we consider only strategies of the form $(0, \beta, \gamma)$ which dominate α). $(0, \beta, \gamma)$ is a self–financed dynamic strategy which corresponds to investing at date 1 through path–independent Arrow–Debreu securities and reinvest the proceeds between date 1 and date 2, conditional on the realized state at time 1 in the risky asset and the bond. This leads to the following:

Definition 9 *The* set of α–investment strategies *is*

$$\mathcal{I}^\alpha = \left\{ (\beta, \gamma) \in \mathbb{R}^{|\Sigma_1|} \times \mathbb{R}^{|\Sigma_1|} | \mathcal{X}(0, \beta, \gamma) \geq \mathcal{X}(\alpha, 0, 0) \right\} .$$

An α–arbitrage opportunity *is an investment* $(\beta, \gamma) \in \mathcal{I}^\alpha$ *such that* $\mathcal{P}(0, \beta, \gamma) < \mathcal{P}(\alpha, 0, 0)$.

The "no α–arbitrage" condition can then be interpreted in the sense that any dynamic strategy $(0, \beta, \gamma)$ synthesizes the $(\alpha, 0, 0)$ payoff at a higher price than the static strategy in Arrow–Debreu securities of date 2. The relation to the no–arbitrage condition is then expressed by

Proposition 10 *There are no arbitrage opportunities in the fictitious static market if, and only if, there are no α–arbitrage opportunities for any payoff $\alpha \in \mathbb{R}^{|\Sigma_2|}$.*

5 Testing for the Existence of a $\mathcal{D}$–MCPM

This section studies a condition on data set $\mathcal{D}$ to find an easy way to check for dynamic arbitrage opportunities. Clearly this can only hold if there are no static ones. Therefore we always assume that property 2 holds. We take as a starting point a homogeneous grid and study the general case only later in this section.

Assumption 11 *The option grid is homogeneous at all dates $t \in \mathcal{T}$.*

Similar to property 2, the following characterizes the absence of arbitrage opportunities.

Property 12 *Option prices are non–decreasing with the exercise date; i.e., for any time $t \in \mathcal{T}$, $K \in \mathcal{K}_t$ and $(t, C_t), (t+1, C_{t+1}) \in \mathcal{D}_K : C_t \leq C_{t+1}$.*

Merton (1973) states that under assumption 11 and in the absence of arbitrage opportunities, property 12 holds. It is straightforward to construct the arbitrage strategy if $C_t > C_{t+1}$: we sell the call with maturity t and buy the call with maturity $t+1$, a strategy known as selling a calendar spread. This generates a strictly positive inflow at date 0. At date t if the observed asset price is less than K, we do nothing. Otherwise, the short call will be exercised; in order to self–finance that transaction, we hold the asset short and put the received amount K into the bank account. At date $t+1$, this generates the payoff $(S_{t+1}-K)^+ - (S_{t+1}-K)1_{S_t>K} \geq 0$.

The striking fact is that property 12 is also sufficient for the absence of arbitrage opportunities and for ensuring existence of a $\mathcal{D}$–MCPM. We will prove this in several steps. First we prove it between two dates $t, t+1$ for $t \in \mathcal{T}$. Our aim will be to prove that increasing call option prices imply that there are no α–arbitrage opportunities; i.e.,

$$\forall \alpha \in \mathbb{R}^{|\Sigma_{t+1}|} : \inf_{(\beta,\gamma)\in\mathcal{I}^\alpha} \mathcal{P}(0,\beta,\gamma) \geq \mathcal{P}(\alpha,0,0) .$$

The proof is given in the appendix and proceeds through the following lemma. It proves that there are no α–arbitrage opportunities, when α represents the payoff of a short call with strike K and maturity $t+1$.

Lemma 13 *Under assumption 11 if $\mathcal{K}_t = \Sigma_t$, $\mathcal{K}_{t+1} = \Sigma_{t+1}$ and properties 2 and 12 hold, there are no "α–arbitrage" opportunities for $\alpha = (-(s-K)^+)_{s\in\Sigma_{t+1}}$.*

Any concave payoff is a linear combination with nonnegative coefficients of short calls for all positive strikes. Using theorem 10, it is then a small step to prove:

Theorem 14 *For any date $t \in \mathcal{T}$, if $\mathcal{K}_t = \Sigma_t$, $\mathcal{K}_{t+1} = \Sigma_{t+1}$ and under assumption 11, the existence of a $\mathcal{D}$–MCPM in the dynamic market consisting of trading dates $\{t, t+1\}$ is equivalent to the following two conditions:*

(1) Observed call option prices for a given arbitrary exercise date are positive, decreasing and convex in strike (property 2).
(2) For a given arbitrary exercise price, call option prices are increasing in exercise date (property 12).

By theorem 8 this also gives an equivalent characterization of the presence of arbitrage opportunities in the data set. We will now generalize the previous theorem to the case $\Sigma_t \subset \mathcal{K}_t$, $\Sigma_{t+1} \subset \mathcal{K}_{t+1}$. This will allow us to cover the case of the non–homogeneous node grid later, too. Furthermore, we now study the original market consisting of all dates in $\mathcal{T}$ at the same time. We first show the following

Proposition 15 *Under assumption 11 there exists a $\mathcal{D}$–MCPM if, and only if, $\mathcal{D}$ can be extended to a set of option prices $\tilde{\mathcal{D}}$ which is complete (i.e.,*

$\mathcal{K}_t = \Sigma_t$ for any $t \in \mathcal{T}$) and fulfills assumptions 2 and 11 such that $\mathcal{D} \subset \tilde{\mathcal{D}}$.

Proposition 15 states that there exists a risk–neutral measure if, and only if, we are able to "complete" the set of observed option prices while keeping the conditions that prices are decreasing and convex in the exercise price and monotone in the exercise date. Our remaining problem is, then, to find an easy procedure to check whether such an extension exists. We will address this in the remainder of this section and consider all observed options for all maturity dates in $\mathcal{T}$.

Definition 16 *For any strike $K \geq 0$ and any date $t \in \mathcal{T}$, the convex envelope $\mathcal{E}_t(K)$ of the discrete set of points in $\mathcal{D}$ is defined by*

$$\begin{aligned} \mathcal{S}_t &= \{(\lambda, \mu) \in I\!R^2 | \forall \tau \geq t, \tau \in \mathcal{T},\ (K', C') \in \mathcal{D}_t \ :\ \lambda + \mu K' \leq C'\} \ , \\ \mathcal{E}_t(K) &= \sup\{\lambda + \mu K | (\lambda, \mu) \in \mathcal{S}_t\} \ . \end{aligned}$$

Please note that $\mathcal{E}_t(K)$ fulfills properties 2 and 12, the latter since $\mathcal{S}_{t+1} \subset \mathcal{S}_t$ for any t. It is therefore a "first choice" for an extension of $\mathcal{D}$.

Proposition 17 *In the absence of arbitrage opportunities, the superreplication price $\bar{C}_t(K)$ of an option is less or equal than the convex envelope; i.e.,*

$$\begin{aligned} &\forall t \in \mathcal{T}\ \forall K \in [0, \infty[\ :\ \bar{C}_t(K) \leq \mathcal{E}_t(K) \ , \\ \textit{and} \qquad &\forall (t, K, C) \in \mathcal{D} \ :\ \bar{C}_t(K) = C = \mathcal{E}_t(K) \ . \end{aligned}$$

In the absence of arbitrage opportunities, according to proposition 17, the convex envelope $\mathcal{E}_t(K)$ interpolates the observed option prices with maturity t. Moreover, the superreplication price of an option is equal to the convex envelope. The following theorem proposes a simple and tractable way to check for the existence of a $\mathcal{D}$–MCPM in the general case.

Theorem 18 *There exists a $\mathcal{D}$–MCPM if, and only if, the convex envelope $\mathcal{E}_t(K)$ interpolates the observed option prices, i.e., $\forall (t, K, C) \in \mathcal{D} :\ \mathcal{E}_t(K) = C$.*

This is a tractable way to test for arbitrage opportunities. The convex envelopes of observed option prices are easy to compute. In most standard cases in real applications, they are simply the linear interpolations of observed option prices of a given maturity. Then the simplest construction, where no further nodes have to be introduced is possible.

For example, we can consider the data set of exchange rate option prices used by Avellaneda, Friedman, Holmes, and Samperi (1997) (see table 1 and figure

maturity	type	strike	bid	ask	impl. volat.
30 days	Call	1.5421	0.0064	0.0076	14.9
	Call	1.5310	0.0086	0.0100	14.8
	Call	1.4872	0.0230	0.0238	14.0
	Put	1.4479	0.0085	0.0098	14.2
	Put	1.4371	0.0063	0.0074	14.4
60 days	Call	1.5621	0.0086	0.0102	14.4
	Call	1.5469	0.0116	0.0135	14.5
	Call	1.4866	0.0313	0.0325	13.8
	Put	1.4312	0.0118	0.0137	14.0
	Put	1.4178	0.0087	0.0113	14.2
90 days	Call	1.5764	0.0101	0.0122	14.1
	Call	1.5580	0.0137	0.0160	14.1
	Call	1.4856	0.0370	0.0385	13.5
	Put	1.4197	0.0141	0.0164	13.6
	Put	1.4038	0.0104	0.0124	13.6
180 days	Call	1.6025	0.0129	0.0152	13.1
	Call	1.5779	0.0175	0.0207	13.1
	Call	1.4823	0.0494	0.0515	13.1
	Put	1.3902	0.0200	0.0232	13.7
	Put	1.3682	0.0147	0.0176	13.7
270 days	Call	1.6297	0.0156	0.0190	13.3
	Call	1.5988	0.0211	0.0250	13.2
	Call	1.4793	0.0586	0.0609	13.0
	Put	1.3710	0.0234	0.0273	13.2
	Put	1.3455	0.0173	0.0206	13.2

Table 1
USD/DEM OTC market data set of Avellaneda, et al.; August 23, 1995

1). Here, straight lines connecting the different option prices do not intersect, so they are the convex envelope; and we deduce that there are no arbitrage opportunities in this specific dataset.

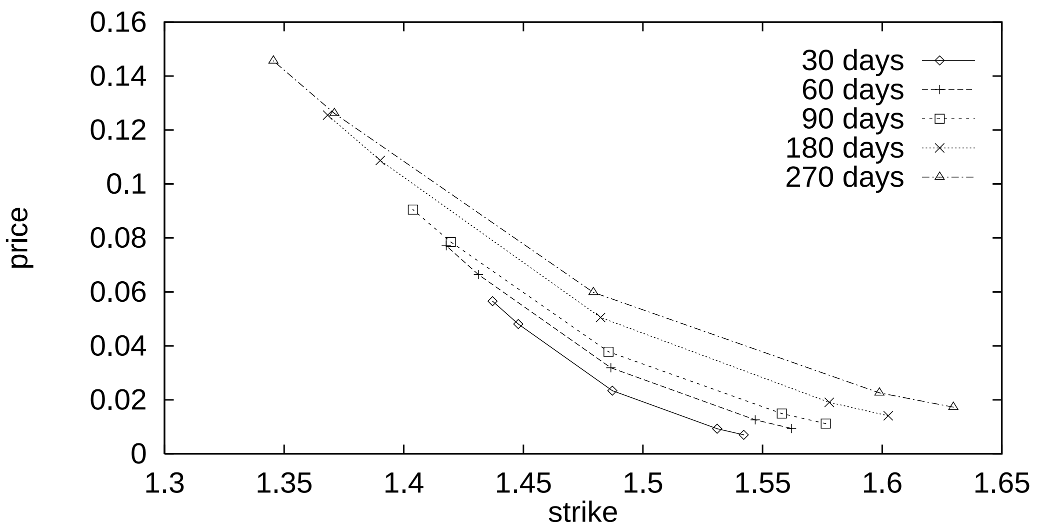

Fig. 1. Prices depending on the strike for different maturities

Remark 19 *Theorem 18 guarantees that once property 2 is satisfied, it is possible to calibrate option prices with a scalar Markov chain. In particular* *stochastic volatility or extra static variables are not required. Of course, one may think that stochastic volatility is a desirable feature and may consider processes consistent with an observed data set based on a larger state space provided that the properties on the call option prices are fulfilled.*

6 Implementation

We now explain the implementation of our approach using the dataset of Avellaneda, Friedman, Holmes, and Samperi (1997) — introduced in section 5 — in detail. We will draw some interesting conclusions from our implementation and compare them with the results obtained by Avellaneda, Friedman, Holmes, and Samperi (1997). We deduced from figure 1, together with theorem 18, that there are no arbitrage opportunities and that a $\mathcal{D}$–MCPM exists.

The following extension of the set of nodes is necessary. A call (put) option with strike equal to the highest (lowest) node s_t^{max} (s_t^{min}) has zero payoff. Thus in the absence of arbitrage, their price must be equal to zero. This can not be the consistent outcome of a $\mathcal{D}$–MCPM. Therefore we introduce two dummy nodes, s_t^{min} and s_t^{max}, with corresponding prices, $C(t, s_t^{max}) = 0$ and $C(t, s_t^{max}) = s_t^{max} - S_0$ [4]. Later in section 6 we explain in detail how to choose the additional strikes s_t^{max}, s_t^{min} in relation those already existing one's.

Besides the fact that we chose nodes identical to strikes, as explained in section 2, a full specification of our approach requires introducing "dummy"–nodes at each date above the highest node s_t^{max} and below the lowest node s_t^{min} of the grid. To do so, we first choose a constant *factor* μ. Then we study the difference between the highest ones and introduce the new node above the highest one so that the difference is exactly μ times the difference before. If s (s') denotes highest (second highest), then the new one is at $s + \mu(s - s')$. We proceed similarly for the new one below. It then remains to specify μ. The highest (lowest) node has to carry the probability weight for all possible higher (lower) movements. It should therefore not be chosen too small. We tested several factors, but only those with $\mu \geq 5$ worked well, producing adequate results in the marginal distribution. We therefore adopted $\mu = 5$ in this analysis.

The second choice is the prior. We adopted a trinomial prior where for inner

[4] These prices are such that put–call parity holds, since the interest rate was supposed to equal zero. For a time–homogeneous grid, the two extreme states s_t^{min} and s_t^{max} need to be absorbing for the process to be risk–neutral, i.e. $P^{\Pi}[S_{t'} = s_t^{min} | S_t = s_t^{min}] = P^{\Pi}[S_{t'} = s_t^{max} | S_t = s_t^{max}] = 1$ (see also Sondermann (1987) and Sondermann (1988)). In practical applications, however, the grid will spread out and the extreme states are not absorbing.

	1.2673	1.4178	1.4312	1.4866	1.5469	1.5621	1.6381	σ_{loc}
1.3831	0.38 (0.01)	0.43 (0.98)	0.04 (0.01)	0.06 (0.00)	0.03 (0.00)	0.02 (0.00)	0.04 (0.00)	0.36
1.4371	0.10 (0.00)	0.35 (0.30)	0.37 (0.62)	0.05 (0.08)	0.04 (0.00)	0.03 (0.00)	0.07 (0.00)	0.29
1.4479	0.02 (0.00)	0.28 (0.13)	0.34 (0.65)	0.34 (0.21)	0.01 (0.00)	0.00 (0.00)	0.01 (0.00)	0.15
1.4872	0.00 (0.00)	0.01 (0.00)	0.23 (0.09)	0.60 (0.87)	0.12 (0.04)	0.00 (0.00)	0.03 (0.00)	0.16
1.531	0.00 (0.00)	0.01 (0.00)	0.02 (0.00)	0.38 (0.29)	0.33 (0.63)	0.20 (0.08)	0.07 (0.00)	0.16
1.5421	0.06 (0.00)	0.05 (0.00)	0.06 (0.00)	0.05 (0.13)	0.31 (0.67)	0.21 (0.20)	0.26 (0.00)	0.33
1.5976	0.01 (0.00)	0.02 (0.00)	0.02 (0.00)	0.04 (0.00)	0.02 (0.02)	0.33 (0.60)	0.57 (0.38)	0.21

Table 2
Transition between date 1 and 2

	1.1833	1.4038	1.4197	1.4856	1.558	1.5764	1.6684	σ_{loc}
1.2673	0.65 (0.97)	0.30 (0.03)	0.02 (0.00)	0.02 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.40
1.4178	0.03 (0.00)	0.55 (0.44)	0.32 (0.55)	0.05 (0.02)	0.02 (0.00)	0.01 (0.00)	0.02 (0.00)	0.22
1.4312	0.01 (0.00)	0.46 (0.18)	0.25 (0.73)	0.27 (0.08)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.15
1.4866	0.00 (0.00)	0.03 (0.00)	0.15 (0.05)	0.68 (0.93)	0.13 (0.02)	0.00 (0.00)	0.01 (0.00)	0.15
1.5469	0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.32 (0.13)	0.37 (0.76)	0.22 (0.11)	0.08 (0.00)	0.19
1.5621	0.02 (0.00)	0.07 (0.00)	0.05 (0.00)	0.06 (0.03)	0.34 (0.64)	0.22 (0.33)	0.25 (0.00)	0.33
1.6381	0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.02 (0.00)	0.02 (0.00)	0.27 (0.28)	0.68 (0.72)	0.18

Table 3
Transition between date 2 and 3

	0.9697	1.3682	1.3902	1.4823	1.5779	1.6025	1.7255	σ_{loc}
1.1833	0.53 (0.44)	0.36 (0.56)	0.02 (0.00)	0.03 (0.00)	0.02 (0.00)	0.02 (0.00)	0.02 (0.00)	0.46
1.4038	0.03 (0.00)	0.56 (0.39)	0.24 (0.40)	0.06 (0.21)	0.03 (0.01)	0.01 (0.00)	0.06 (0.00)	0.24
1.4197	0.02 (0.00)	0.39 (0.28)	0.22 (0.41)	0.34 (0.30)	0.01 (0.01)	0.01 (0.00)	0.01 (0.00)	0.17
1.4856	0.00 (0.00)	0.04 (0.04)	0.15 (0.20)	0.64 (0.60)	0.15 (0.14)	0.00 (0.02)	0.01 (0.00)	0.14
1.558	0.00 (0.00)	0.02 (0.00)	0.01 (0.02)	0.35 (0.36)	0.36 (0.38)	0.20 (0.21)	0.06 (0.03)	0.15
1.5764	0.05 (0.00)	0.05 (0.00)	0.04 (0.01)	0.04 (0.26)	0.33 (0.39)	0.22 (0.29)	0.27 (0.05)	0.34
1.6684	0.00 (0.00)	0.02 (0.00)	0.01 (0.00)	0.03 (0.02)	0.02 (0.12)	0.35 (0.41)	0.57 (0.45)	0.17

Table 4
Transition between date 3 and 4

nodes i

$$P_{i,j} = \begin{cases} \frac{1}{3} & ;\ i = j \\ \frac{1}{3} - \frac{1}{10(\nu_t - 3)} & ;\ j = i \pm 1 \\ \frac{1}{10(\nu_t - 3)} & ;\ \text{otherwise} \end{cases} ,$$

and for the outer most nodes add the probability of the missing node to the one corresponding to the starting node. Here $\nu_t = |\Sigma_t|$ denotes the number of nodes at date t. This prior will, in general, not be consistent with the observed option prices. However, it supports good the fact that we would like a departure as small as possible from the Black–Scholes setup to make it compatible with the observed option prices. Furthermore it is similar to a trinomial model which represents the fact that it is only a discrete approximation to some continuous time model. We nevertheless scrambled the probabilities slightly to allow for non–zero probability weights also in non–adjoint nodes. A lognormal prior — the Black–Scholes model — yields similar results in the following implementation..

We then proceeded exactly as described in the previous subsection. Tables 2 to 5 contain the transition matrices in the form found throughout the paper.

	0.8335	1.3455	1.371	1.4793	1.5988	1.6297	1.7842	σ_{loc}
0.9697	0.75 (1.00)	0.21 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.00 (0.00)	0.47
1.3682	0.01 (0.00)	0.63 (0.50)	0.28 (0.40)	0.05 (0.10)	0.02 (0.00)	0.00 (0.00)	0.01 (0.00)	0.16
1.3902	0.02 (0.00)	0.43 (0.33)	0.27 (0.47)	0.28 (0.20)	0.01 (0.00)	0.00 (0.00)	0.00 (0.00)	0.19
1.4823	0.00 (0.00)	0.04 (0.02)	0.16 (0.18)	0.66 (0.69)	0.14 (0.11)	0.00 (0.01)	0.01 (0.00)	0.14
1.5779	0.00 (0.00)	0.01 (0.00)	0.01 (0.01)	0.35 (0.31)	0.43 (0.47)	0.12 (0.20)	0.08 (0.01)	0.18
1.6025	0.02 (0.00)	0.07 (0.00)	0.05 (0.00)	0.06 (0.18)	0.40 (0.46)	0.16 (0.32)	0.24 (0.03)	0.34
1.7255	0.00 (0.00)	0.01 (0.00)	0.01 (0.00)	0.03 (0.00)	0.03 (0.05)	0.25 (0.40)	0.66 (0.54)	0.21

Table 5
Transition between date 4 and 5

To make them more easily accessible, however, we index them directly by the nodes. In parentheses we have put the corresponding probabilities resulting from a Black–Scholes model with constant volatility of 0.15 p.a. The last column contains the local volatility of the risk–adjusted process.

Except for the column furthest to the left, we see that the lower left hand part is always different from the Black–Scholes case, where it is zero. The market anticipates a strictly positive probability for large downward movements. The fear of "crashes" seems to be present in the observed option prices or, differently, an appropriate model should allow for jumps in the underlying prices. We observe a similar pattern in the upper right–hand part, indicating that there is a strictly higher probability for large upward movements. Although this effect seems to be similar to "fat tails," however, here it is much more pronounced since it attributes "high" probability (between 0.01 and 0.07) on zero events, and we see it in the conditional evolution in the future. The observation that the first column is zero might be attributable to the fact that μ is too large. With the exception of the last row we find a U–shaped form for the "local volatility".

7 Conclusion

We calibrated a risk–neutral asset process to observed call option prices under the assumption that the discounted asset prices follow a Markov chain. We proved that existence of such a Markov chain is equivalent to the condition that call option prices are decreasing and convex in the strike price and increasing in the exercise date. An important technique we used throughout was to reduce our dynamic market to a static one with extra constraints due to the possibility of dynamically trading in the bond and the risky asset. We characterized the superreplication price of a call in both a static and dynamic framework. The bid–ask spread was shown to be reduced due to the trade–ability of other options. We applied the Bayesian approach to infer the "optimal" measure and revealed surprising results.

A Proofs

Proof of proposition 1 Let us for the moment consider an arbitrary finite set of couples $\tilde{\mathcal{D}} = \{(\mathcal{X}_j, C_j),\ 1 \leq j \leq J\}$, where $\mathcal{X}_j$ is a derivative asset and C_j, its price. We denote by $\hat{\mathcal{M}}$ the set of probability measures on the state space $\mathcal{A} = \bigotimes_{t\in\mathcal{T}} \Sigma_t$ consistent with $\tilde{\mathcal{D}}$; they do not necessarily have to be associated to Markov chains. We have $\mathcal{M}_{\mathcal{D}} \subset \mathcal{M} \subset \hat{\mathcal{M}}$. Then we study the following optimization over the set of probability measures absolutely continuous w.r.t. to P and consistent with $\tilde{\mathcal{D}}$:

$$\min_{Q\in\hat{\mathcal{M}}_{\tilde{\mathcal{D}}}, Q\ll P} E^P\left[\frac{dQ}{dP}\log\frac{dQ}{dP}\right] .$$

We identify the elements of $\hat{\mathcal{M}} \subset \tilde{\mathcal{D}}$ with the elements of the simplex of $\mathbb{R}^{|\mathcal{A}|}$ and, similarly, $\mathcal{X}_j$ with an element in $\mathbb{R}^{|\mathcal{A}|}$. The set of absolutely continuous measure w.r.t P and consistent with $\tilde{\mathcal{D}}$ is a closed bounded set of $\mathbb{R}^{|\mathcal{A}|}$ and the optimization criterion is continuous; since $\mathcal{M}_{\mathcal{D}} \neq \emptyset$, there exists an optimal probability measure Q^* absolutely continuous w.r.t. P.

Let us assume that Q^* is not equivalent to P. We denote by q_i^* the probability being in state i under Q^*. For any $\epsilon \geq 0$, we define a (signed) measure $Q(\epsilon)$ by $q_i(\epsilon) = q_i^* + \epsilon(q_{0,i} - q_i^*)$, where $q_{0,i}$ denotes the state probability i under Q_0. This measure satisfies all equality constraints. Furthermore with $q^*_{\min} := \inf_{q_i^*>0} q_i^*$, $\delta := \sup_i |q_{0,i} - q_i^*|$, we can check that we have $q_i(\epsilon) \geq 0$ for all $\epsilon \in [0, q^*_{\min}/\delta[$ and $i \in \mathcal{A}$. Now, let us consider for any such $\epsilon \in [0, q^*_{\min}/\delta[$:

$$\begin{aligned} &\frac{H(Q(\epsilon)) - H(Q^*)}{\epsilon} \\ &= \sum_{q_i^*\neq 0} \frac{q_i(\epsilon)\log q_i(\epsilon)/p_i - q_i^*\log q_i^*/p_i}{\epsilon} + \sum_{q_i^*=0} q_{0,i}\log\frac{\epsilon q_{0,i}}{p_i} . \end{aligned}$$

The first term admits a limit when $\epsilon \longrightarrow 0$, while the second term tends to $-\infty$. Thus there exists some ϵ such that $H(Q(\epsilon)) < H(Q^*)$ which contradicts the assumption and proves that Q^* is an equivalent measure to P.

Therefore there exists an interior solution $Q^* > 0$; i.e., the inequality constraints are not binding. Applying theorems 28.2 and 28.3 of Rockafellar (1970), we find that there exists $\lambda_j \in \mathbb{R},\ j = 1, \ldots, J$ such that

$$\frac{dQ^*}{dP}\exp\{\sum_{j=1}^{J}\lambda_j\mathcal{X}_j\} ,$$

which proves the desired result.

Proof of proposition 7 The set of attainable claims is clearly $\{\gamma(S_1)(S_2 - S_1) + \sum_{(t,K,C\in\mathcal{D})} a_{t,K}(S_t - K)^+ | a \in \mathbb{R}^{|\Sigma_1|} \times \mathbb{R}^{|\Sigma_2|}, \gamma \in \mathbb{R}^{|\Sigma_1|}\}$. For a payoff $V_2 = \gamma(S_1)(S_2 - S_1) + \sum_{s\in\Sigma_1} a_s(S_1 - s)^+ + \sum_{s\in\Sigma_2} b_s(S_2 - s)^+$, the term $-\gamma(S_1)S_1 + \sum_{s\in\Sigma_1} a_s(S_1 - s)^+$ can be written on the basis of path–independent Arrow–Debreu securities as $\sum_{s\in\Sigma_1} \beta_s 1_{S_1=s}, \beta \in \mathbb{R}^{|\Sigma_1|}$; similarly, there exist some $\alpha_s \in \mathbb{R}^{|\Sigma_2|}$ such that $\sum_{s\in\Sigma_1} b_s(S_2 - s)^+ = \sum_{s\in\Sigma_2} \beta_s 1_{S_2=s}$. This proves that an investment through calls, puts, and the asset and numeraire can be achieved as an investment in $\delta^{dep}_{s,s'}$, δ^1_s and $\delta^2_{s'}$ $((s,s') \in \Sigma_1 \times \Sigma_2)$. The converse follows in a similar way. Using equation (4), it is straightforward to prove the form for the price functional $\mathcal{P}$.

Proof of proposition 10 If there exists an arbitrage opportunity $\bar{\alpha}, \bar{\beta}, \bar{\gamma}$ such that $\mathcal{X}(\bar{\alpha}, \bar{\beta}, \bar{\gamma}) \geq 0$ and $\mathcal{P}(\bar{\alpha}, \bar{\beta}, \bar{\gamma}) < 0$, then by linearity of $\mathcal{X}$ and $\mathcal{P}$ we get $\mathcal{X}(0, \bar{\beta}, \bar{\gamma}) \geq \mathcal{X}(-\bar{\alpha}, 0, 0)$ and $\mathcal{P}(0, \bar{\beta}, \bar{\gamma}) < \mathcal{P}(-\bar{\alpha}, 0, 0)$, which means that the superreplication price is strictly below the replication price.

Conversely, any α–arbitrage opportunity is obviously an arbitrage opportunity in the strict sense.

Proof of lemma 13 Let us first fix a call with strike K and maturity date $t+1$ and denote by C_t, C_{t+1} the two call prices and by i_K the node corresponding to the strike. We address the problem by backward induction and consider in state ι at date t the minimum amount of numeraire to be held in order to superreplicate the short call payoff at date $t+1$:

$$
\begin{aligned}
V_\iota = \inf_{\beta_\iota, \gamma_\iota} \; & \beta_\iota + \gamma_\iota s_\iota \\
\text{s.t.} \quad & \forall s \in \Sigma_{t+1} : \beta_\iota + \gamma_\iota s \geq -(s-K)^+ \; .
\end{aligned}
$$

The optimal superreplication strategy splits up into two cases:

(1) For $\iota \leq i_K$: From $\beta_\iota + \gamma_\iota S_\iota \geq -(S_\iota - K)^+ = 0$, we deduce that $V_\iota \geq 0$. Since $\beta_\iota = \gamma_\iota = 0$ satisfies the constraint $\beta_\iota + \gamma_\iota s = 0 \geq -(s-K)^+$ for any $s \in \Sigma_{t+1}$, we deduce that the optimum is attained at $V_\iota = 0$.

(2) For $\iota > i_K$: Similar to the previous case, we deduce from $\beta_\iota + \gamma_\iota s_\iota \geq -(s_\iota - K)^+ = K - s_\iota$, that $V_\iota \geq K - S_\iota$. For $\beta_\iota = +K$, $\gamma_\iota = -1$ the constraint $\beta_\iota + \gamma_\iota s = K - s \geq -(s-K)^+$ $(s \in \Sigma_2)$ is satisfied and so the optimum $V_\iota = K - S_\iota$ is attained.

In both cases, the inflow $V_\iota = -(s_\iota - K)^+$ for the superreplication strategy at time t is simply a short call payoff with strike K. By the dynamic programming principle, $\inf_{(\beta,\gamma)\in\mathcal{I}^\alpha} \mathcal{P}(0, \beta, \gamma)$ corresponds to the superreplication price of $V_\iota = -(S_\iota - K)^+$, and so, by property 2, it is equal to the replication price of the short call with maturity date one, i.e. $-C_t$. So for α with $\alpha_s = -(s-K)^+$, we have

$$\inf_{(\beta,\gamma)\in\mathcal{I}^\alpha} \mathcal{P}(0,\beta,\gamma) = -C_t \ .$$

By assumption, the price of the short call with strike K and maturity t is larger than $-C_{t+1} = \mathcal{P}(\alpha,0,0)$; this proves the assertion.

Proof of theorem 14 By the result of Merton (1973), cited at the beginning of section 5, we only need to prove sufficiency. First, let us assume that α is concave. Short calls form a basis of the space of (path–independent) payoffs, so there exists $\tilde{\alpha}_l \in \mathbb{R}$ s.t. $\alpha_s = -\sum_l a_l(s-K_l)^+$. We quickly check that $\tilde{\alpha}_l = (\alpha_l - \alpha_{l-1}) - (\alpha_{l+1} - \alpha_l) \geq 0$. Thus the superreplication price of $\mathcal{X}(\alpha,0,0)$ can be written as

$$\tilde{\alpha}_0 S_0 - \sum_l \tilde{\alpha}_l C_{t,l} \geq \tilde{\alpha}_0 S_0 - \sum_l \tilde{\alpha}_l C_{t+1,l} = \mathcal{P}(\alpha,0,0) \ .$$

Together with property 12, this proves that there are no α–arbitrage opportunities for any concave payoff.

In the case of a general payoff α at date $t+1$, we will denote by $\bar{\alpha}$ the concave envelope of α. Similar to lemma 13, we test for α–arbitrage and study the optimization problem in every state ι at time t:

$$\begin{aligned} V_\iota = \inf_{\beta_\iota,\gamma_\iota} & \ \beta_\iota + \gamma_\iota s_\iota \\ \text{s.t.} & \ \ \forall j : \beta_\iota + \gamma_\iota s_j \geq \alpha_j \ . \end{aligned}$$

The concave envelope $\bar{\alpha}$ is by definition the solution of this optimization problem. Thus, we have $V_\iota = \bar{\alpha}_\iota$ and so

$$\inf_{(\beta,\gamma)\in\mathcal{I}^{\overline{\alpha}}} \mathcal{P}(0,\beta,\gamma) = \inf_{(\beta,\gamma)\in\mathcal{I}^\alpha} \mathcal{P}(0,\beta,\gamma) \ .$$

In other words, the superreplication price of α is equal to the superreplication price of its concave envelope. The absence of arbitrage opportunities in the static market at time $t+1$ and $\bar{\alpha} \geq \alpha$ imply $\mathcal{P}(\bar{\alpha},0,0) \geq \mathcal{P}(\alpha,0,0)$. By the preceding study of concave payoffs, we get

$$\inf_{(\beta,\gamma)\in\mathcal{I}^{\bar{\alpha}}} \mathcal{P}(0,\beta,\gamma) \geq \mathcal{P}(\bar{\alpha},0,0) \ ,$$

which guarantees the result. The proof now follows with proposition 10 and proposition 8.

Proof of proposition 15 According to proposition 8, we only need to check for arbitrage opportunities. Necessity is the result of Merton (1973) stated at the beginning of section 5.

In the case of only two exercise dates, sufficiency follows directly from theorem 14 since a $\tilde{\mathcal{D}}$–MCPM is also a $\mathcal{D}$–MCPM. With several exercise dates, let us first note that the marginal densities of S_t, λ_t are uniquely determined from the option prices $(t, K, \tilde{C}(t, K)) \in \tilde{\mathcal{D}}$. From the above argument, looking at only two dates, follows the existence of a set of joint probabilities $P[S_t = x, S_{t+1} = y], (x, y) \in \Sigma_t \times \Sigma_{t+1}$ such that: $\forall (x, y) \in \Sigma_t \times \Sigma_{t+1}$: $P[S_t = x] = \lambda_t(x), P[S_{t+1} = y] = \lambda_{t+1}(y)$ and $\forall x \in \Sigma_t$: $\sum_{y \in \Sigma_t} P[S_t = x, S_{t+1} = y](y - x) = 0$. We can define a stochastic matrix Π_t by

$$\Pi_t(x, y) = \frac{P[S_t = x, S_{t+1} = y]}{\lambda_t(x)} = P[S_{t+1} = y | S_t = x] \ .$$

Now, for $t \in \mathcal{T}$, the MCMM defined by Π_t has marginal densities equal to λ_t and is thus both consistent with the option prices $\tilde{C}(t, K), \ K \in \Sigma$ and with the smaller set of observed option prices $C(t, K), \ t \in \mathcal{T}, K \in \Sigma_\tau$. Moreover, it fulfills the martingale restriction and so it is a $\mathcal{D}$–MCPM.

Proof of proposition 17 Absence of arbitrage opportunities implies there is a $\mathcal{D}$–MCPM which we represent here by P^Π. Proposition 15 implies that the option prices $C_\Pi(t, K), K \in \Sigma_t, t \in \mathcal{T}$ generated by P^Π are convex and fulfill

$$\forall \tau \geq t \ \forall K \in \Sigma_\tau \ : \ C_\Pi(t, K) \ . \leq C_\Pi(\tau, K) = C(\tau, K)$$

The superreplication price $\bar{C}(t, K)$ is the supremum of C_Π taken over all possible P^Π given by $\mathcal{D}$–MCPM's. $\bar{C}(t, K)$ inherits the convexity from $C_\Pi(t, K)$. It also clearly satisfies by definition $\forall \tau \geq t \ \forall K \in \Sigma_\tau \ : \ \bar{C}(t, K) \leq \mathcal{E}_t(K)$. This proves directly the first part of the assertion.

Whatever $\mathcal{D}$–MCPM P^Π, we have $\forall K \in \tilde{\Sigma}_t \ : \ C_\Pi(t, K) = C(t, K)$. Thus $\bar{C}(t, K) = C(t, K)$. From the first part follows that $C(t, K) = \bar{C}(t, K) \leq \mathcal{E}_t(K)$. On the other hand, $\mathcal{E}_t(K) \leq C(t, K)$; this implies $C = \mathcal{E}_t(K)$. By proposition 15, there exists a $\mathcal{D}$–MCPM which then has, by definition, $\bar{C} = C$. This proves the second assertion.

Proof of theorem 18 To prove the theorem, we apply proposition 8. First we prove the theorem in the homogeneous case of assumption 11. By proposition 17, we need only to study sufficiency: The convex envelope $\{\mathcal{E}_t(K) | t \in \mathcal{T}, K \in \Sigma\}$ provides a set of option prices consistent with observed option prices:

$$\forall t \in \mathcal{T}, K \in \Sigma_t \ : \ \mathcal{E}_t(K) = C(t, K) \ .$$

By definition $\mathcal{E}_t(K)$ is convex in K and is also clearly increasing in t. Thus, we can apply proposition 15 and prove the absence of arbitrage opportunities.

Now we address the general inhomogeneous case. First we prove necessity by contradiction: Assume the condition would not hold. Then, by definition of the convex envelope, there exists a $(\tilde{C}, \tilde{t}, \tilde{K}) \in \mathcal{D}$ such that $\mathcal{E}_t(K) < \mathcal{L}_t(K)$, where $\mathcal{L}_t(K)$ is the linearly interpolated value corresponding to the next traded strike to the left and right. However this constitutes an arbitrage opportunity as in Merton (1973), cited at the beginning of section 5. This contradiction proves the first part.

Now we prove sufficiency and study a fictitious model with the thinnest grid corresponding to all nodes in which we set the call option prices corresponding to the additional strikes equal to their price resulting from interpolating the option prices of the next two traded options strikes. This will result in the same $\mathcal{E}$. The above argument for the homogeneous case ensures existence. By construction of this fictitious model, the additional nodes have zero probability weight. Therefore we can translate the fictitious model directly back into our original model.

References

Avellaneda, M., C. Friedman, R. Holmes, and D. Samperi (1997): "Calibrating Volatility Surfaces via Relative–Entropy Minimization," *Applied Mathematical Finance*, 4, 37–64.

Bakshi, G., C. Cao, and Z. Chen (1997): "Empirical performance of Alternative Option Pricing Models," *The Journal of Finance*, 52, 2003–2049.

Barle, S., and N. Cakici (1995): "Growing a Smiling Tree," *RISK*, 8, 76–81.

Bates, D. S. (1996a): "Dollar Jump Fears, 1984–1992: Distributional Abnormalities Implicit in Currency Futures Options," *Journal of International Money and Finance*, 15(1), 65–93.

——— (1996b): "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in PHLX Deutschemark Options," *The Review of Financial Studies*, 9(1), 69–107.

Breeden, D., and R. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621–651.

Buchen, P., and M. Kelly (1996): "The Maximum Entropy Distribution of an Asset Inferred from Options Prices," *Journal of Financial and Quantitative Analysis*, 31, 143–159.

Campa, J., K. Chang, and L. Reider (1997): "Implied Exchange Rate Distribution : Evidence from OTC Option Markets," Discussion Paper, New York University.

Carr, P., and D. Madan (1998): "Determining Volatility Surfaces and Option Values From an Implied Volatility Smile," *Disussion Paper*.

Chriss, N. A. (1997): *Black–Scholes and Beyond — Option Pricing Models*. Irwin.

Cvitanic, J., and I. Karatzas (1993): "Hedging Contingent Claims with constrained portfolios," *The Annals of Applied Probability*, 3, 652–681.

Cvitanic, J., H. Pham, and N. Touzi (1998): "Superreplication in Stochastic Volatility Models under Portfolio Constraints," *to appear: Applied Probability Journals*.

Derman, E., and I. Kani (1995): "Riding on the Smile," *RISK*, 7, 32–39.

Derman, E., I. Kani, and N. Chriss (1996): "Implied Trinomial Tree of The Volatility Smile," *The Journal of Derivatives*, 3(4), 7–22.

Duffie, D. (1992): *Dynamic Asset Pricing Theory*. Princeton University Press.

Duffie, D., and K. Singleton (1998): "Rating–Based Term Structure of Credit Spreads," Discussion Paper, Stanford University.

Dumas, B., J. Fleming, and R. Whaley (1998): "Implied Volatility Functions: Empirical Tests," *The Journal of Finance*, 53(6), 2059–2106.

Dupire, B. (1996): "Pricing with a Smile," *RISK*, 7, 18–20.

El Karoui, N., and M. C. Quenez (1995): "Dynamic Programming and Pricing of Contingent Claims in Incomplete Markets," *SIAM Journal of Control and Optimization*, 33.

Hansen, L. P., and R. Jagannathan (1997): "Assessing Specification Errors in Stochastic Factor Models," *The Journal of Finance*, 52(2), 557–90.

Harrison, J. M., and D. M. Kreps (1979): "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

Jackwerth, J.-C., and M. Rubinstein (1996): "Recovering Probability Distributions from Option Prices," *The Journal of Finance*, 51, 1611–1631.

Jarrow, R., D. Lando, and Turnbull (1997): "A Markov Model of the Term Structure of Credit Spreads," *The Review of Financial Studies*, 10, 481–523.

Kushner, H., and G. Dupuis (1992): *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer–Verlag.

Luttmer, E. G. (1997): "Asset Pricing in Economics with Frictions," *Econometrica*, 64(6), 1439–76.

McCauley, R., and W. R. Melick (1996): "Risk Reversal Risk," *RISK*, 9(11), 54–57.

MELICK, W., AND C. THOMAS (1997): "Recovering an Asset's PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis," *Journal of Financial and Quantitative Analysis*, 32(1), 91–115.

MERTON, R. C. (1973): "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

ROCKAFELLAR, R. T. (1970): *Convex Analysis*. Princeton University Press.

RUBINSTEIN, M. (1994): "Implied Binomial Trees," *The Journal of Finance*, 49, 771–818.

SHIMKO, D. (1993): "Bounds of Probability," *RISK*, 6, 33–37.

SONDERMANN, D. (1987): "Currency Options: Hedging and Social Value," *European Economic Review*, 31, 246–256.

——— (1988): "Option Pricing with bounds on the underlying securities," in *Bankpolitik, finanzielle Unternehmensführung und die Theorie der Finanzmärkte, Festschrift für Hans Krümmel*, pp. 421–442. Duncker & Humblot.

ZIPKIN, P. (1993): "Mortgages and Markov Chains: A Simplified Evaluation Model," *Management Science*, 39, 683–691.