Robust Replication of Volatility Derivatives

Peter Carr* and Roger Lee†

This version: May 31, 2009

Abstract

In a nonparametric setting, we develop trading strategies to replicate volatility derivatives – contracts which pay functions of the realized variance of an underlying asset's returns. The replicating portfolios trade the underlying asset and vanilla options, in quantities that we express in terms of vanilla option prices, not in terms of parameters of any particular model. Likewise, we find nonparametric formulas to *price* volatility derivatives, including volatility swaps and variance options. Our results are exactly valid, if volatility satisfies an independence condition. In case that condition does not hold, our formulas are moreover immunized, to first order, against nonzero correlation.

1 Introduction

The tradeoff between risk and return is a central theme of finance, and *volatility* and variance of returns are standard measures of risk. The volatility of a stock is revealed by the market price of an option on the stock, if one accepts the model of Black-Scholes [6], which does not require any assumptions on equity risk premium nor expected return.

However, that model does assume constant volatility, which contradicts the empirical observation that, for a given expiry, the market prices of option contracts at different strikes typically imply different Black-Scholes volatility parameters. In the presence of these non-constant implied volatilities across strikes – a phenomenon known as the smile or skew – the question arises: what information content, regarding the risk-neutral distribution of the path-dependent realized volatility and variance, can we extract from the profile of European option prices at a given expiry?

*Bloomberg LP and Courant Institute, NYU. `pcarr@nyc.rr.com`.

†University of Chicago. `RL@math.uchicago.edu`.

We answer this question with the help of an observation due to Hull-White [21]. Under some assumptions including an independence condition, the distribution of realized variance determines the value of a stock option. We invert this relationship in a more general setting. Analogously to how Latane-Rendleman [23] take as given the market price of a single option and invert the Black-Scholes equation to infer a constant volatility, we take as given the market prices of all options at a given expiry and invert a Hull-White-type relationship to infer the entire risk-neutral distribution of the random realized volatility.

The information in the profile of T-expiry option prices will, therefore, nonparametrically reveal the no-arbitrage prices of volatility derivatives – claims on payoffs contingent on realized volatility. This information will, moreover, allow us to replicate volatility derivatives, by dynamic trading in standard options and the underlying shares. Our valuations and our replication strategies will have explicit formulas in terms of observables, not the parameters of any model.

Our inference does not rely on any specification of the market price of volatility risk. Just as knowledge of the stock price sufficiently reflects the equity risk premium in the Black-Scholes framework, knowledge of option prices sufficiently reflects the volatility risk premium in our framework.

1.1 Variance

We define the realized variance of the returns on a positive underlying price S from time 0 to time T to be the quadratic variation of $\log S$ at time T. If S has an instantaneous volatility process σ_t, then realized variance equals integrated variance, meaning the time integral of σ_t^2. In practice, contracts written on realized variance typically define it discretely as the sample variance of daily or weekly log returns. Following the custom in the derivatives literature, we study the (continuously-sampled) quadratic variation / integrated variance, leaving tests of discrete sampling for future research.

Realized variance can be traded by means of a variance swap, a contract which pays at time T the difference between realized variance and an agreed fixed leg. The variance swap has become a leading tool – perhaps *the* leading tool – for portfolio managers to trade variance. As reported in the Financial Times [20] in 2006,

> Volatility is becoming an asset class in its own right. A range of structured derivative products, particularly those known as variance swaps, are now the preferred route for many hedge fund managers and proprietary traders to make bets on market volatility.

According to some estimates [1], the daily trading volume in equity index variance swaps reached

USD 4–5 million vega notional in 2006. On an annual basis, this corresponds to payments of more than USD 1 billion, per percentage point of volatility.

From a dealer's perspective, the variance swap admits replication by a T-expiry log contract (which decomposes into static positions in calls and puts on S), together with dynamic trading in S, as shown in Neuberger [25], Dupire [17], Carr-Madan [15], Derman et al [16], and Britten-Jones/Neuberger [11]. Perfect replication requires frictionless markets and continuity of the price process, but does not require the dynamics of instantaneous volatility to be specified. The variance swap's replicating portfolio became in 2003 the basis for how the Chicago Board Option Exchange (CBOE) calculates the VIX index, an indicator of short-term options-implied volatility. VIX implementation issues arising from data limitations are addressed in Jiang-Tian [22].

1.2 Volatility derivatives

More generally, *volatility derivatives*, which pay *functions* of realized variance, are of interest to portfolio managers who desire non-linear exposure to variance. Important examples include calls and puts on realized variance; and volatility swaps (popular especially in foreign exchange markets) which pay *realized volatility*, defined as the square root of realized variance.

In contrast to the variance swap's replicability by a log contract, general functions of variance present greater hedging difficulties to the dealer. In theory, *if* one specifies the dynamics of instantaneous volatility as a one-dimensional diffusion, then one can replicate a volatility derivative by trading the underlying shares and one option. Such simple stochastic volatility models are, however, misspecified according to empirical evidence, such as difficulties in fitting the observed cross-section of option prices, and pricing errors out-of-sample, as documented in Bakshi-Cao-Chen [2] and Bates [4]. Moreover, even if one could find a well-specified model, further error can arise in trying to calibrate or estimate the model's parameters, not directly observable from options prices.

Derivatives dealers have struggled with these issues. According to a 2003 article [26] in RiskNews,

> While variance swaps - where the underlying is volatility squared - can be perfectly replicated under classical derivatives pricing theory, this has not generally been thought to be possible with volatility swaps. So while a few equity derivatives desks are comfortable with taking on the risk associated with dealing volatility swaps, many are not.

A 2006 Financial Times article [20] quotes a derivatives trader:

> Variance is easier to hedge. Volatility can be a nightmare.

We challenge this conventional wisdom, by developing strategies to price and to replicate volatility derivatives – *without* specifying the dynamics of instantaneous volatility, hence without bearing the types of misspecification and misestimation risk discussed above.

The volatility derivatives studied in this paper (and referenced in the block quotations) are *realized* volatility contracts, which pay functions of *underlying price paths* – as opposed to the various types of *options-implied* volatility contracts, which pay functions of *option prices* prevailing at a specified time. For example, we do not explicitly study options on VIX (itself a function of vanilla option prices) nor options on straddles (Brenner-Ou-Zhang [9]); rather, we do study, for example, options and swaps on the variance and volatility actually *realized* by the underlying.

1.3 Our approach

We prove that general functions of variance, including volatility swaps, do admit valuation and replication using portfolios of the underlying shares and European options, dynamically traded according to strategies valid across all underlying dynamics specified in Section 2.

Our nonparametric exact hedging paradigm stands in marked contrast to previous treatments of volatility derivatives. In particular, consider the following features.

First, in contrast to analyses of particular models (such as Matytsin's [24] analysis of Heston and related dynamics), we take a nonparametric approach, both *robust* and *parameter-free*, in the sense that we do not specify the dynamics nor estimate the parameters of instantaneous volatility. Our *robust* pricing and hedging strategies remain valid across a whole class of models – including non-Markovian and discontinuous volatility processes as well as diffusive volatility – so we avoid the risk of misspecification and miscalibration present in any one model. Specifically, we define *robust* to mean that our strategies are valid across all underlying continuous price processes whose instantaneous volatility satisfies an independence assumption (and some technical conditions, designated below as $(\mathrm{B}, \mathrm{W}, \mathrm{I})$). Moreover, in case the independence condition does not hold, we immunize our schemes, to first order, against the presence of correlation; thus we can price approximately under dynamics which generate implied volatility skews – without relying on any particular model of volatility. Our *parameter-free* pricing formulas typically take the form of an equality of risk-neutral expectations of functions of realized variance $\langle X\rangle_T$ and price S_T respectively:

$$\mathbb{E}h(\langle X\rangle_T) = \mathbb{E}G(S_T), \tag{1.1}$$

where we find formulas for G, given various classes of payoff functions h, including the square root function which defines the volatility swap. The left-hand side is the value of the desired *volatility* or variance contract. The right-hand side is the value of a contract on a function of *price*, and is therefore model-independently given by the values of European options. Thus our formula for the volatility contract value is expressed not in terms of the parameters of any model, but rather in terms of prices directly observable, in principle, in the vanilla options market.

Second, in contrast to approximate methods (such as Carr-Lee's [13] use of a displaced lognormal to approximate the distribution of realized volatility) we find *exact* formulas for prices and hedges of volatility contracts. For example, the typical result (1.1) is exact under the independence condition.

Third, in contrast to studies of valuation without hedging (such as Carr-Geman-Madan-Yor's [12] model-dependent variance option valuations under pure jump dynamics), we cover not just valuation but also *replication*, by proving explicit option trading strategies which enforce the valuation results. The holdings in our replicating portfolios are rebalanced dynamically, but the quantity to hold, at each time, depends only on contemporaneously observable *prices*, not on the parameters of any model; this result arises because the observable prices already incorporate all quantities of possible relevance, such as instantaneous volatility, volatility-of-volatility, and market price of volatility risk. Indeed, to our knowledge, this paper is the first one to study nonparametrically the pricing restrictions induced by, and the volatility payoffs attainable by, the ability to trade options dynamically. Moreover, because perfectly hedging against a short (long) holding of some realized volatility payoff is equivalent to perfectly replicating a long (short) position in that volatility payoff, our replication strategies therefore provide explicit robust *hedges* of volatility risk.

Fourth, in contrast to treatments narrowly focused on particular payoff specifications, we develop valuation and replication methods for *general* functions of volatility. As Breeden-Litzenberger [8] showed, the information in the set of T-expiry option prices at all strikes, fully and model-independently reveals the risk-neutral distribution of S_T. We show that the same option price information, under our assumptions, fully and robustly reveals the risk-neutral *distribution of volatility*, and hence the valuations of arbitrary functions of volatility. This paper, moreover, breaks ground for ongoing research into general functions of volatility and price *jointly* – such as options on CPPI, constant proportion portfolio insurance [31].

Fifth, in contrast to valuation methods that rely crucially on continuity of the share price and the instantaneous volatility, this paper allows unspecified *jumps* in volatility. Moreover, our valuation methods have natural and far-reaching extensions to time-changed Lévy processes, including those with asymmetric skew-inducing *jumps* in price, which we develop in a companion paper.

2 Assumptions

Fix an arbitrary time horizon $T > 0$. Assume either that interest rates are zero, or alternatively that all prices are denominated relative to an asset (the "bond" or "cash") that pays 1 at time T. Thus the bond has price $B = 1$ at all times. Assume that markets are frictionless.

On a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ satisfying the usual conditions, assume there exists an equivalent probability measure $\mathbb{Q}$ such that the underlying share price S solves

$$dS_t = \sigma_t S_t dW_t, \qquad S_0 > 0 \tag{Assumption W}$$

for some $(\mathcal{F}_t, \mathbb{Q})$-Brownian motion W_t and some measurable $\mathcal{F}_t$-adapted process σ_t which satisfy

$$\int_0^T \sigma_t^2 dt \text{ is bounded by some } m \in \mathbb{R} \tag{Assumption B}$$

and

$$\sigma \text{ and } W \text{ are independent} \tag{Assumption I}$$

and such that $\mathbb{Q}$ is a risk-neutral pricing measure satisfying, in particular, that for all $\alpha, p \in \mathbb{C}$ and $t \leq T$, a power contract paying at time T the real part of αS_T^p has time-t price equal to the real part of $\alpha \mathbb{E}_t S_T^p$, where $\mathbb{E}_t$ denotes $\mathcal{F}_t$-conditional $\mathbb{Q}$-expectation. This assumption rules out arbitrage among the bond, stock, and power contracts.

Denote the logarithmic returns process by

$$X_t := \log(S_t/S_0) \tag{2.1}$$

and write $\langle X \rangle$ for the quadratic variation of X, also known as the *realized variance* of the returns on S. Under assumption (W),

$$\langle X \rangle_t = \int_0^t \sigma_u^2 du. \tag{2.2}$$

Unless otherwise stated, the assumptions (B, W, I) on S are in effect throughout this paper. These assumptions are sufficient for the validity of our methodology, but not necessary. Indeed each of the three assumptions can be relaxed:

Remark 2.1. In this paper we will relax our reliance on assumption (I), by finding results immunized – in a sense to be defined in Section 4 – against correlation between σ and W.

Assumptions (I) and (W) taken together imply that implied volatility skews are symmetric [3] – in contrast to typical implied volatility skews in equity markets, which slope downward. Therefore Section 4 on correlation-immunization has practical importance.

Remark 2.2. We drop assumption (B) in Section 8.

Remark 2.3. We drop assumption (W) in a companion paper, by introducing jumps in the price process. In particular, we allow asymmetries in the jump distribution, which can generate asymmetric volatility skews.

Remark 2.4. We need not and will not work under the actual physical probability measure $\mathbb{P}$. All expectations are with respect to risk-neutral measure $\mathbb{Q}$. Our typical result, of the form

$$\mathbb{E}h(\langle X\rangle_T) = \mathbb{E}G(S_T), \tag{2.3}$$

states nothing directly about the physical expectation of $h(\langle X\rangle_T)$.

Rather, it concludes that the *value* of the contract that pays $h(\langle X\rangle_T)$ equals the price of the contract that pays $G(S_T)$, by reasoning such as the following: the $G(S_T)$ claim, plus dynamic self-financing trading, replicates the $h(\langle X\rangle_T)$ payoff with risk-neutral probability 1, hence with physical probability 1, because $\mathbb{P}$ and $\mathbb{Q}$ agree on all events of probability 1. Thus, given the availability of the appropriate European-style contracts as hedging instruments, the variance payoff $h(\langle X\rangle_T)$ is dynamically spanned, and valuation result (2.3) follows, by absence of arbitrage.

The irrelevance of physical expectations (for this paper's valuation and replication purposes) renders also irrelevant the mapping between risk-neutral expectations and physical expectations. Thus we have no need of any assumptions about the volatility risk premia (nor indeed any other type of risk premia) which mediate between the risk-neutral and the physical probability measures. In particular, our results are valid regardless of the market's risk preferences, and regardless of whether volatility risk is priced or unpriced. Any effects of risk premia are already impounded in the prices of our hedging instruments.

Remark 2.5. Our replication strategy assumes frictionless trading in options. Of course, options trading incurs transaction costs in practice, but our results maintain relevance. First, transactions costs have decreased, and continue to decrease, as options markets become more liquid. Second, in practice a dealer typically manages a *portfolio* of volatility contracts, which mitigates trading costs, because offsetting trades (buying an option to hedge one volatility contract, selling that option to hedge another contract) need not actually be conducted. Third, our frictionless valuation can be regarded, in the presence of frictions, as a "central" valuation, relative to which a dealer planning to bid (offer) should make a downward (upward) adjustment dependent on transaction costs. Fourth, regardless of trading costs, our results are still implementable in non-trading contexts, such as the development of VIX-like indicators of expected volatility, as discussed in Remark 6.16.

3 Variance swap

A *variance swap* pays $\langle X \rangle_T$ minus an agreed fixed amount, which we take to be zero unless otherwise specified. Variance swap replication does not require assumption (I). As shown in Neuberger [25], Dupire [17], Carr-Madan [15], Derman et al [16], and Britten-Jones/Neuberger [11],

$$X_T = \log(S_T/S_0) = \int_0^T \frac{1}{S_u} dS_u + \frac{1}{2}\int_0^T \left(\frac{-1}{S_u^2}\right)\sigma_u^2 S_u^2 du.$$

by Itô's rule, so

$$\langle X \rangle_T = -2X_T + \int_0^T \frac{2}{S_u} dS_u. \tag{3.1}$$

Hence the following self-financing strategy replicates the $\langle X \rangle_T$ payoff. At each time $t \leq T$ hold a static position in the log contract, plus a dynamically traded share position, plus a bond position that finances the shares and accumulates the trading gains or losses:

$$\begin{array}{rl} 1 & \textit{log contract}\text{, which pays } -2\log(S_T/S_0) \\ \dfrac{2}{S_t} & \text{shares} \\ \displaystyle\int_0^t \frac{2}{S_u} dS_u - 2 & \text{bonds} \end{array}$$

By replication, therefore, the variance swap's time-0 value equals the price of the log contract. Alternatively, this may be derived by taking expectations of (3.1) to obtain

$$\mathbb{E}_0 \langle X \rangle_T = \mathbb{E}_0[-2\log(S_T/S_0)] = \mathbb{E}_0[-2\log(S_T/S_0) + 2(S_T/S_0) - 2]. \tag{3.2}$$

At general times $t \in [0, T]$, by similar reasoning,

$$\mathbb{E}_t \langle X \rangle_T - \langle X \rangle_t = \mathbb{E}_t[-2\log(S_T/S_t) + 2(S_T/S_t) - 2]. \tag{3.3}$$

The "delta-hedged" log contract in (3.3) may be regarded as a *synthetic* variance swap.

Remark 3.1. By Breeden-Litzenberger [8] and Carr-Madan [15], the log contract, and indeed a claim on a general function $G(S_T)$, can be synthesized if we have bonds and T-expiry puts and calls at all strikes. Specifically, if $G : \mathbb{R}_+ \to \mathbb{R}$ is a difference of convex functions, then for any $\kappa \in \mathbb{R}_+$ we have for all $x \in \mathbb{R}_+$ the representation

$$G(x) = G(\kappa) + G'(\kappa)(x - \kappa) + \int_{K \geq \kappa} G''(K)(x - K)^+ dK + \int_{0 < K < \kappa} G''(K)(K - x)^+ dK \tag{3.4}$$

where G' denotes the left-derivative, and G'' the second derivative, which exists as a signed measure.

In practice, calls and puts do not trade at all strikes, but in liquid markets, such as S&P 500 options, they may trade at enough strikes to make satisfactory approximations to (3.4) for the contracts G that we will need. Nonparametric techniques of Bondarenko [7] estimate call/put prices at all strikes (hence the prices of $G(S_T)$ contracts), given a limited number of strikes.

4 Immunization against correlation

The typical pricing result in this paper has the following form. Given a desired function h of variance, we find a formula for a function G of price, such that

$$\mathbb{E}h(\langle X\rangle_T) = \mathbb{E}G(S_T). \tag{4.1}$$

Indeed, we will find an *infinite family* of G such that (4.1) holds for all processes S satisfying assumptions $(\mathrm{B}, \mathrm{W}, \mathrm{I})$. Now consider the following relaxation of (I). Fix some instantaneous volatility process σ_t. Let $\rho \in [-1, 1]$. Let price S have correlation ρ with volatility, in the sense that

$$dS_t = \sqrt{1-\rho^2}\sigma_t S_t dW_{1t} + \rho\sigma_t S_t dW_{2t},$$

where σ and W_1 are independent, as are the Brownian motions W_1, W_2. If $\rho = 0$, then we have (I), hence (4.1). Changing the correlation to some $\rho \neq 0$ has *no* effect on the left-hand side $\mathbb{E}h(\langle X\rangle_T)$, which depends only on the law of the σ process. From among the infinite family of G, we will find one such that the right-hand side $\mathbb{E}G(S_T)$ is *also* insensitive to ρ (at least locally). Thus we gain the benefit that (4.1) still holds approximately, even if condition (I) does not hold.

To quantify the impact of correlation, Proposition 4.1 will give a *mixing formula* that (without assuming independence) expresses the value of any European-style payoff (such as the $G(S_T)$ in (4.1)) as the expectation of the Black-Scholes formula for that payoff, evaluated at a randomized share price and random volatility. The parameter ρ appears explicitly in the mixing formula, enabling us to examine the formula's correlation-sensitivity and to choose a G such that $\mathbb{E}G(S_T)$ has zero sensitivity to correlation perturbations.

First we define what is meant by the Black-Scholes formula for a payoff.

Let $t \leq T$. Let $\mathcal{B}$ denote the Borel sets of $\mathbb{R}_+$ and let $m\mathcal{F}_t$ denote the set of $\mathcal{F}_t$-measurable random variables. Consider a time-t-contracted European *payoff function*, by which we mean a

$$F : \mathbb{R}_+ \times \Omega \to \mathbb{R}, \qquad F \text{ is } (\mathcal{B} \otimes \mathcal{F}_t)\text{-measurable}. \tag{4.2}$$

Think of F as a function which maps S_T to a European-style payout; for example, an at-the-money (ATM) call would have $F(\mathcal{S}) = (\mathcal{S} - S_t)^+$. The ω-dependence allows payoffs constructed at time t to depend on information in $\mathcal{F}_t$. Our notation may suppress this ω-dependence; for example, $F(\mathcal{S}) = (\mathcal{S} - S_t)^+$ is shorthand for $F(\mathcal{S}, \omega) = (\mathcal{S} - S_t(\omega))^+$.

Given payoff F, define the *Black-Scholes formula* by $F^{BS}(\mathsf{s}, 0, \omega) := F(\mathsf{s}, \omega)$ and, for $\sigma > 0$,

$$F^{BS}(\mathsf{s}, \sigma, \omega) := \int_0^\infty F(\mathsf{s}y, \omega) \frac{1}{\sqrt{2\pi}\sigma y} e^{-(y+\sigma^2/2)^2/(2\sigma^2)} dy.$$

The kernel in the integrand is a lognormal density with parameter σ. Note that the valuation F^{BS} is defined as a function of *today's* price (where "today" means the valuation date), unlike the payoff F which is defined as a function of *expiration* price. Notationally, we make a distinction: the placeholder for today's price is s, whereas the placeholder for expiration price is $\mathcal{S}$. Again, our notation may suppress the ω-dependence.

To prove the mixing formula, we recall the argument due to Romano-Touzi [28] and Willard [30], but in a slightly more general setting where we do not assume that instantaneous volatility follows a 1-factor diffusion.

Proposition 4.1 (Mixing formula)**.** *Without assuming* (I)*, let*

$$dS_t = \sqrt{1-\rho^2}\sigma_t S_t dW_{1t} + \rho\sigma_t S_t dW_{2t}$$

where $|\rho| \leq 1$, and W_1 and W_2 are $\mathcal{F}_t$-Brownian motions, and σ and W_2 are adapted to some filtration $\mathcal{H}_t \subseteq \mathcal{F}_t$, where $\mathcal{H}_T$ and $\mathcal{F}_T^{W_1}$ are independent. Then

$$\mathbb{E}_t F(S_T) = \mathbb{E}_t F^{BS}(S_t M_{t,T}(\rho),\ \bar{\sigma}_{t,T}\sqrt{1-\rho^2}), \tag{4.3}$$

where

$$M_{t,T}(\rho) := \exp\left(-\frac{\rho^2}{2} \int_t^T \sigma_u^2 du + \rho \int_t^T \sigma_u dW_{2u} \right) \tag{4.4}$$

and $\bar{\sigma}_{t,T} := (\int_t^T \sigma_u^2 du)^{1/2}$.

Remark 4.2. This setting includes the standard correlated stochastic volatility models, of the form

$$dS_t = \sigma_t S_t dW_{0t}$$
$$d\sigma_t = \alpha(\sigma_t)dt + \beta(\sigma_t)dW_{2t},$$

where W_2 and $W_0 := \sqrt{1-\rho^2}W_1 + \rho W_2$ have correlation ρ. Our setting also allows more general dynamics; for example, σ can have *jumps* independent of W_1.

Remark 4.3. Expanding (4.3) in a formal Taylor series about $\rho = 0$,

$$\begin{aligned}\mathbb{E}_t F(S_T) &= \mathbb{E}_t F^{BS}(S_t M_{t,T}(\rho),\ \bar{\sigma}_{t,T}\sqrt{1-\rho^2}) \\ &\approx \mathbb{E}_t F^{BS}(S_t, \bar{\sigma}_{t,T}) + \rho S_t \mathbb{E}_t\left[\frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t, \bar{\sigma}_{t,T})\int_t^T \sigma_u dW_{2u}\right] + O(\rho^2)\end{aligned}$$

The $\partial F^{BS}/\partial \mathsf{s}$ has randomness due to its argument $\bar{\sigma}_{t,T}$, so in general it cannot be pulled out of the $\mathbb{E}_t$. Suppose, however, that F has the property that $\partial F^{BS}/\partial \mathsf{s}$ does not depend on its second argument. Then the $\partial F^{BS}/\partial \mathsf{s}$ comes out of the expectation. What remains inside the expectation is a mean-zero integral, so the ρ term vanishes, leaving an error of only $O(\rho^2)$:

$$\mathbb{E}_t F(S_T) \approx \mathbb{E}_t F^{BS}(S_t, \bar{\sigma}_{t,T}) + O(\rho^2),$$

and we describe the F payoff as first-order *correlation-neutral* or *correlation-immune.*

In selecting hedging instruments and pricing benchmarks, we favor payoffs having this property, because of their valuations' immunity (in the sense of first-order invariance) to the presence of correlation. This motivates the following definition.

Definition 4.4. Let $t < T$. We say that a payoff function F is first-order *ρ-neutral* or *ρ-immune* or *correlation-neutral* or *correlation-immune* at time t if there exists $c \in m\mathcal{F}_t$ such that

$$\frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t, \sigma) = c \qquad \text{for } all \text{ constants } \sigma \geq 0,$$

almost surely. In other words, "the contract's Black-Scholes delta is constant across all volatility parameters".

Remark 4.5. Adding an affine function $\alpha S_T + \beta$ (where $\alpha, \beta \in m\mathcal{F}_t$) has no effect on whether or not a payoff is ρ-neutral, because the $\alpha S_T + \beta$ payoff is itself ρ-neutral.

Definition 4.6. Consider a *trading strategy* which holds at each time $t < T$ a portfolio of claims whose combined time-T payout is $F_t(S_T)$, where F_t is a payoff function (4.2). We say that the trading strategy is [first-order] *ρ-neutral* if for each $t < T$, the payoff function F_t is ρ-neutral.

5 Exponentials

Consider an exponential variance claim which pays $e^{\lambda\langle X\rangle_T}$ for some constant λ. Such payoffs will serve as building blocks, from which we will create more general functions of $\langle X\rangle_T$.

5.1 Basic replication

We introduce first a basic "correlation-sensitive" replication strategy for the exponential variance payoffs, relying on the independence assumption (I). In Section 5.2, we will improve this to a "correlation-immune" strategy, which neutralizes the first-order impact of departures from independence.

The fundamental pricing formula relates the value of an exponential claim on variance and the value of a power claim on price. The proof applies, to powers of S_T, the conditioning argument in Hull-White [21]. Intuitively, if S_T were lognormal, then the expectation of a power of S_T would be exponential in variance. In our case, S_T is a *mixture* of lognormals of various variances, so the expectation of a power of S_T is equal to the *expectation* of an exponential of a *random* variance.

Proposition 5.1 (Basic pricing of exponentials)**.** *For each $\lambda \in \mathbb{C}$ and $t \leq T$,*

$$\mathbb{E}_t e^{\lambda\langle X\rangle_T} = e^{\lambda\langle X\rangle_t}\mathbb{E}_t(S_T/S_t)^{1/2\pm\sqrt{1/4+2\lambda}}. \tag{5.1}$$

In particular, for $t=0$,

$$\mathbb{E}_0 e^{\lambda\langle X\rangle_T} = \mathbb{E}_0(S_T/S_0)^{1/2\pm\sqrt{1/4+2\lambda}}. \tag{5.2}$$

Remark 5.2. The distribution of $\langle X\rangle_T$ is (just as any distribution is) fully determined by its characteristic function, via the well-known inversion formula. In turn, the characteristic function of $\langle X\rangle_T$ is, via Proposition 5.1, determined by the values of $\mathbb{E}_t(S_T/S_t)^{1/2\pm\sqrt{1/4+2\lambda}}$ for λ imaginary, which are determined by the time-t prices of calls and puts (via (3.4) applied separately to the real and imaginary parts). Therefore, the information in T-expiry option prices fully reveals the risk-neutral distribution not only of price S_T, but also of variance $\langle X\rangle_T$.

Not only does the power claim on S_T correctly *price* the exponential variance claim, but indeed it dynamically *replicates* the exponential variance payoff.

Proposition 5.3 (Basic replication of exponentials)**.** *Let $\lambda \in \mathbb{R}$. If $p := 1/2 \pm \sqrt{1/4+2\lambda} \in \mathbb{R}$ then the payoff $e^{\lambda\langle X\rangle_T}$ admits replication by the self-financing strategy*

$$\begin{array}{rl} N_t & \textit{claims on } S_T^p \\ -pN_tP_{t-}/S_t & \textit{shares} \\ pN_tP_{t-} & \textit{bonds} \end{array} \tag{5.3}$$

where $N_t := e^{\lambda\langle X\rangle_t}/S_t^p$ and $P_t := \mathbb{E}_t S_T^p$.

Remark 5.4. The S and N are continuous. We can and do work with a right-continuous left-limits version of P. Although P may jump, we are free to replace the predictable process P_{t-} with the adapted process P_t everywhere in the statement and proof of Proposition 5.3, because the continuity of the relevant integrators (S, B, N) makes immaterial the distinction between P_t and P_{t-} in each integrand. Thus we have proved that the strategy

$$\begin{aligned} N_t &\quad \text{claims on } S_T^p \\ -pN_tP_t/S_t &\quad \text{shares} \\ pN_tP_t &\quad \text{bonds} \end{aligned} \tag{5.4}$$

replicates $e^{\lambda\langle X\rangle_T}$. Henceforth we follow the standard practice of allowing one-side-continuous adapted processes, as in (5.4), to serve as integrands (e.g. trading strategies) with respect to continuous integrators (e.g. continuous price processes).

Remark 5.5. If futures are available as hedging instruments, then they can replace the shares and bonds; the strategy to replicate the payoff $e^{\lambda\langle X\rangle_T}$ becomes

$$\begin{aligned} N_t &\quad \text{claims on } S_T^p \\ -pN_tP_t/S_t &\quad \text{futures} \end{aligned}$$

by similar reasoning.

Remark 5.6. For complex λ and p, and complex $\alpha = \alpha(\lambda)$,

$$\mathrm{Re}(\alpha N_T P_T) = \mathrm{Re}(\alpha P_0) + \int_0^T \mathrm{Re}(\alpha N_t) d\mathrm{Re}(P_t) - \int_0^T \mathrm{Im}(\alpha N_t) d\mathrm{Im}(P_t) - \int_0^T \frac{\mathrm{Re}(p\alpha N_t P_t)}{S_t} dS_t$$

so we can replicate $\mathrm{Re}(\alpha e^{\lambda\langle X\rangle_T})$ by trading cosine and sine claims. Specifically, at time t, hold

$$\begin{aligned} \mathrm{Re}(\alpha N_t) &\quad \text{claims on } \mathrm{Re}(e^{pX_T}) \\ -\mathrm{Im}(\alpha N_t) &\quad \text{claims on } \mathrm{Im}(e^{pX_T}) \\ -\mathrm{Re}(p\alpha N_t P_t)/S_t &\quad \text{shares} \\ \mathrm{Re}(p\alpha N_t P_t) &\quad \text{bonds.} \end{aligned}$$

Remark 5.7. In Proposition 5.3, the replicating portfolio's time-t holdings have a combined payoff function

$$F(\mathsf{S}) := N_t \mathsf{S}^p - \frac{pN_tP_t}{S_t}(\mathsf{S} - S_t),$$

which is "delta-neutral" in the sense that

$$\frac{\partial}{\partial \mathsf{s}}\bigg|_{\mathsf{s}=S_t} \mathbb{E}_t F(\mathsf{s}S_T/S_t) = N_t \mathbb{E}_t \left(\frac{\partial}{\partial \mathsf{s}} (\mathsf{s}S_T/S_t)^p \right)\bigg|_{\mathsf{s}=S_t} - \frac{pN_tP_t}{S_t} = 0.$$

Thus the share position $-pN_tP_t/S_t$ can be interpreted as a delta-hedge of the option position consisting of N_t claims on S_T^p. This agrees with intuition; in order to create a purely volatility-dependent payoff, we want zero net exposure to directional risk, hence we delta-neutralize.

Of course, this observation is neither necessary nor sufficient to prove the validity of our hedging strategy (for that purpose the Proposition 5.3 proof speaks for itself); but it can help us to understand and implement the strategy.

Remark 5.8. For pricing and replicating an exponential variance payoff, each "basic" strategy (there are two, due to the $\pm$) is but one member of an infinite family of strategies, all perfectly valid, under assumption (I). Specifically, Carr-Lee [14] show that (I) implies a general form of put-call symmetry: for *any* time-t-contracted payoff function f such that $f(S_T/S_t)$ is integrable,

$$\mathbb{E}_t f\left(\frac{S_T}{S_t}\right) = \mathbb{E}_t\left[\frac{S_T}{S_t} f\left(\frac{S_t}{S_T}\right)\right]. \tag{5.5}$$

Combining Proposition 5.1 and (5.5), we have an infinite *family* of European-style payoffs which correctly price the variance payoff: For all such f,

$$\mathbb{E}_t e^{\lambda\langle X\rangle_T} = e^{\lambda\langle X\rangle_t}\mathbb{E}_t\left[(S_T/S_t)^{1/2+\sqrt{1/4+2\lambda}} + f(S_T/S_t) - \frac{S_T}{S_t} f(S_t/S_T)\right]. \tag{5.6}$$

In particular, choosing $f(S) := \theta S^{1/2-\sqrt{1/4+2\lambda}}$ for $\theta \in m\mathcal{F}_t$ yields the sub-family of identities

$$\mathbb{E}_t e^{\lambda\langle X\rangle_T} = e^{\lambda\langle X\rangle_t}\mathbb{E}_t\left[(1-\theta)(S_T/S_t)^{1/2+\sqrt{1/4+2\lambda}} + \theta(S_T/S_t)^{1/2-\sqrt{1/4+2\lambda}}\right] \tag{5.7}$$

under (I). In the next section we choose θ in such a way as to achieve ρ-neutrality.

5.2 Correlation-immune replication of exponentials

The functions of S_T given in Proposition 5.1 are not correlation-immune, but we will exploit their "non-uniqueness" to achieve correlation-immunity. There exist infinitely many functions of S_T, all of which perfectly replicate (hence price) the exponential variance payoff under assumption (I). From this infinite family, we choose a strategy which *is* correlation-immune, and hence still prices the variance claim approximately, in case (I) does not hold. The idea is to take a weighted combination, with weights $\theta_\pm$, of the power claims, with exponents $p_\pm$, where

$$\begin{aligned} \theta_\pm(\lambda) &:= \frac{1}{2} \mp \frac{1}{2\sqrt{1+8\lambda}} \\ p_\pm(\lambda) &:= \frac{1}{2} \pm \frac{1}{2}\sqrt{1+8\lambda}. \end{aligned} \tag{5.8}$$

Figure 5.1: Exponential variance claims $e^{\lambda \langle X \rangle_T}$ on the left, and their European-style synthetic counterparts $G_{\exp}(S_T; S_0; \langle X \rangle_0; \lambda)$ on the right, for $\langle X \rangle_0 = 0$ and $\lambda \in \{-4, -3, \ldots, 3, 4\}$.

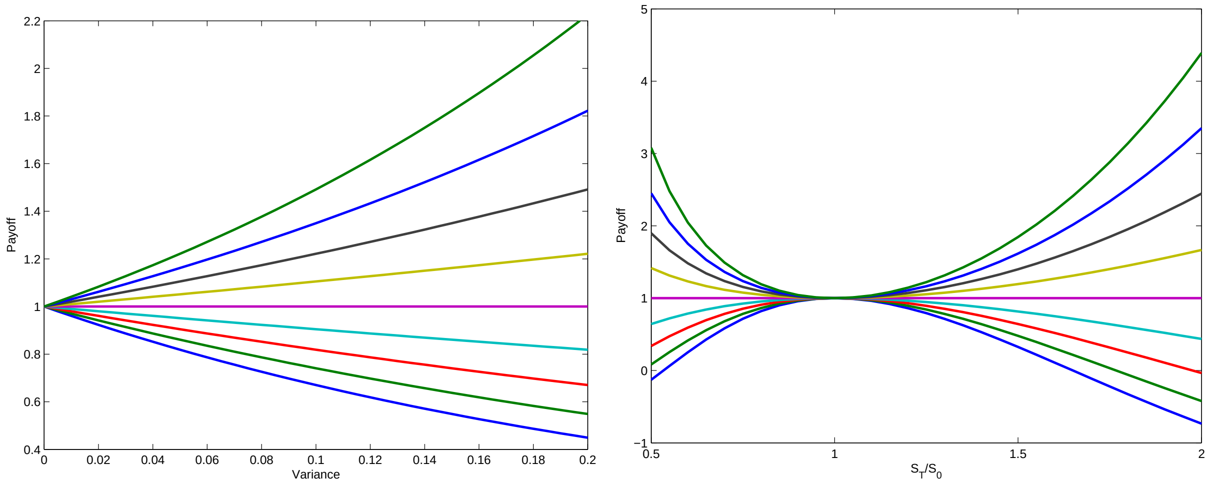

Proposition 5.9 (Correlation-immune pricing of exponentials)**.** *Let $t \leq T$. For any $\lambda \in \mathbb{C}$,*

$$\mathbb{E}_t e^{\lambda \langle X \rangle_T} = \mathbb{E}_t G_{\exp}(S_T, S_t, \langle X \rangle_t; \lambda), \tag{5.9}$$

where

$$G_{\exp}(S, u, q; \lambda) := e^{\lambda q}\big[\theta_+(S/u)^{p_+} + \theta_-(S/u)^{p_-}\big] \tag{5.10}$$

For each t, the payoff function $F(S) := G_{\exp}(S, S_t, \langle X \rangle_t; \lambda)$ is ρ-neutral.

Remark 5.10. Therefore the relationship

$$\mathbb{E}_t e^{\lambda \langle X \rangle_T} = e^{\lambda \langle X \rangle_t} \mathbb{E}_t \Big[\theta_+(S_T/S_t)^{p_+} + \theta_-(S_T/S_t)^{p_-}\Big] \tag{5.11}$$

holds exactly under independence (I), and is first-order *immune* to the presence of correlation. Figure 5.1 plots the payoff functions appearing in the left and right-hand sides. Note that at the valuation date t, every variable in the right-hand side is determined and observable, except S_T.

Like the basic methodology, the correlation-immune methodology provides not only valuation, but also replication of exponential variance payoffs.

Proposition 5.11 (Correlation-immune replication of exponentials)**.** *Define $p_\pm, \theta_\pm$ by* (5.8). *Let*

$$N_t^\pm := e^{\lambda \langle X \rangle_t} / S_t^{p_\pm}$$

$$P_t^\pm := \mathbb{E}_t S_T^{p_\pm}$$

If $\lambda \in \mathbb{R}$ and $p_\pm \in \mathbb{R}$, then the self-financing strategy

$$\begin{aligned} &\theta_+ N_t^+ \quad \textit{claims on} \quad S_T^{p+} \\ &\theta_- N_t^- \quad \textit{claims on} \quad S_T^{p-} \end{aligned}$$

replicates the payoff $e^{\lambda\langle X\rangle_T}$. Moreover, the strategy is ρ-neutral.

6 Volatility swap

A *volatility swap* pays $\sqrt{\langle X\rangle_T}$ minus some agreed fixed amount, which we take to be 0 unless otherwise specified.

6.1 Bounds and approximations

For $F_{\text{atmc}}(S) := (S - S_0)^+$, a direct computation shows that

$$F^{BS}_{\text{atmc}}(S_0, \sigma) = S_0(N(\sigma/2) - N(-\sigma/2))$$

which is strictly increasing and concave in σ.

Define the unannualized at-the-money *implied volatility* IV_0 as the unique solution to

$$F^{BS}_{\text{atmc}}(S_0, \text{IV}_0) = \mathbb{E}_0 F_{\text{atmc}}(S_T). \tag{6.1}$$

Let VOL_0 denote the time-0 volatility swap value, and VAR_0 denote the square root of the time-0 variance swap value.

$$\text{VAR}_0 := \sqrt{\mathbb{E}_0\langle X\rangle_T} \tag{6.2}$$

$$\text{VOL}_0 := \mathbb{E}_0\sqrt{\langle X\rangle_T}. \tag{6.3}$$

These values are model-independently determined by prices of European options, according to Sections 3 and 6.2 respectively. In particular, VAR_0 equals the square root of the value of the log contract; VAR_0 is what the VIX attempts to approximate, and is sometimes described as a *model-free implied volatility.*

Proposition 6.1. *We have the following observable lower and upper bounds on* VOL_0

$$\frac{\sqrt{2\pi}}{S_0}\mathbb{E}_0(S_T - S_0)^+ \underset{(a)}{\leq} \text{IV}_0 \underset{(b)}{\leq} \text{VOL}_0 \underset{(c)}{\leq} \text{VAR}_0 = \sqrt{-2\mathbb{E}_0 \log(S_T/S_0)}.$$

Inequalities (a) and (c) do not assume (I).

Figure 6.1: Proofs of inequalities (a) and (c). Left side (a): The ATM BS formula is concave and nearly linear in σ. Right side (c): The volatility swap payoff admits model-independent superreplication by variance swaps.

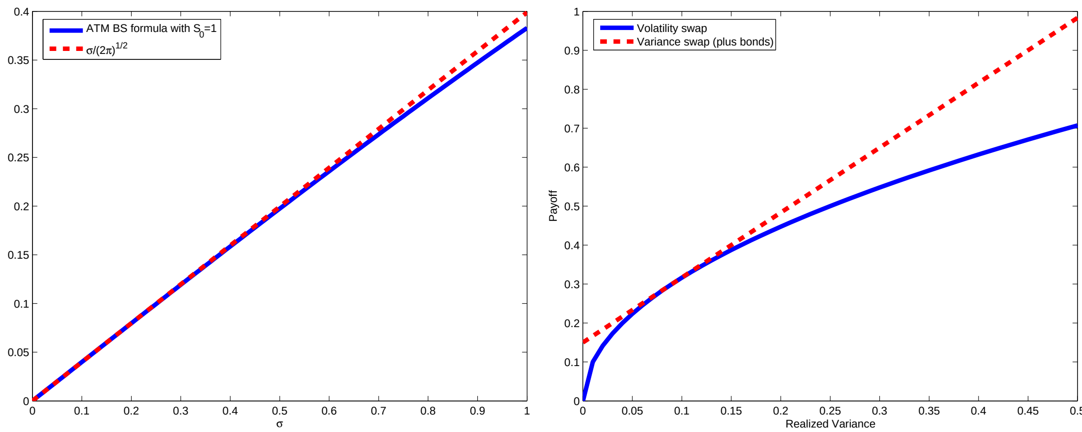

Remark 6.2. The (c) proof given in the appendix can be enforced by model-independent arbitrage.

A portfolio of $1/(2\,\mathrm{VAR}_0)$ variance swaps, plus $\mathrm{VAR}_0/2$ in bonds, has total time-0 value VAR_0, and superreplicates the $\langle X\rangle_T^{1/2}$ payoff. Essentially this portfolio realizes Jensen's inequality, by constructing the appropriate tangent, as shown in Figure 6.1. If variance and volatility swap values fail to respect (c), then going long the superreplicating portfolio, and short a volatility swap, model-independently locks in an arbitrage profit.

In Remarks 6.3 and 6.4, we include some approximations, mainly to provide reference points and context for our theory. We emphasize that we do not actually advocate the use of these two approximations, because our theory is more powerful and robust, in ways described in Remark 6.5.

Remark 6.3. Although $F^{BS}_{\mathrm{atmc}}(S_0,\cdot)$ is concave, it is nearly linear – indeed, linear to a *second* order approximation near 0, because its second derivative vanishes at 0. Thus the inequality in (A.1) is an approximate equality (as shown by Feinstein [18] and Poteshman [27]); and the inequality in (A.2) is an approximate equality (as shown by Brenner and Subrahmanyam [10]). Therefore, the lower bounds of Proposition 6.1 are indeed *approximately equal* to the volatility swap value:

$$\mathrm{VOL}_0 \approx \frac{\sqrt{2\pi}}{S_0}\mathbb{E}_0(S_T - S_0)^+ \approx \mathrm{IV}_0 \tag{6.4}$$

where the first $\approx$ assumes (I), but the second does not. Under the independence assumption, therefore, ATM implied volatility approximates the initial value of a volatility swap – but see Remark 6.5.

Remark 6.4. Under assumption (I), the approximation (6.4) can be refined, to the following simple approximation using ATM implied volatility and the variance swap value:

$$\mathrm{VOL}_0 \approx \mathrm{IV}_0 \left(1 + \frac{\mathrm{VAR}_0^2 - \mathrm{IV}_0^2}{8 + 2\,\mathrm{IV}_0^2}\right). \tag{6.5}$$

Remark 6.5. We do not endorse the approximations (6.4) and (6.5). They do not establish how to *replicate* realized volatility, they do not apply at times *after* inception, they do not value general *functions of* volatility, and they do not suggest what to do in the presence of *correlation.* Our theory does all of the above. Regarding the last point in particular, Section 6.5 will illustrate the benefits of our correlation-immunized approach, compared to the naive approximation (6.4).

6.2 Basic (correlation-sensitive) methodology

We introduce first a basic "correlation-sensitive" valuation strategy for the volatility swap, relying on the independence assumption (I). In Section 6.3, we will improve this to a "correlation-immune" strategy, which neutralizes the first-order impact of correlation.

For our correlation-immune strategy, we will give a full treatment, including seasoned volatility swaps at times $t > 0$, and including the replication argument. For our basic strategy, however, we restrict our coverage to the valuation of volatility swaps at inception $t = 0$, because we do not advocate the basic strategy; for the basic case we include only enough material to draw some connections with other representations/approximations, in Remarks 6.7 and 6.14 and Section 6.5.

Proposition 6.6 (Pricing a volatility swap using the basic synthetic volatility swap)**.** *We have*

$$\mathbb{E}_0\sqrt{\langle X\rangle_T} = \mathbb{E}_0 g_\pm(S_T/S_0)$$

where

$$g_\pm(x) := \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{e^{(1/2\pm 1/2)\log x} - \mathrm{Re}[e^{(1/2\pm\sqrt{1/4-2z})\log x}]}{z^{3/2}} dz \tag{6.6}$$

In particular, we prove the convergence of the integral.

Remark 6.7. Figure 6.2 plots the functions $g_\pm$. They closely resemble $\sqrt{2\pi}/S_0$ at-the-money puts and calls, respectively. Our result is consistent with the naive approximation (6.4), but as discussed in Remark 6.5, our theory has implications far beyond the naive approximations.

We call a claim on $g_+(S_T/S_0)$ the *basic* (or *correlation-sensitive*) *synthetic volatility swap.*

Figure 6.2: European-style payoffs $g_-(S_T)$ and $g_+(S_T)$ = the basic synthetic volatility swap.

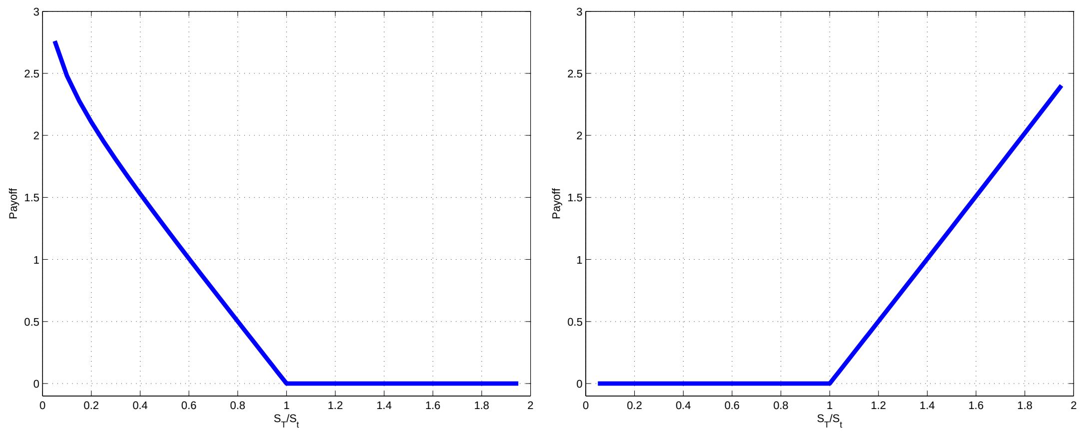

Figure 6.3: European-style payoff $G_{\mathrm{svs}}(S_T, S_t, 0)$, the correlation-immune synthetic volatility swap.

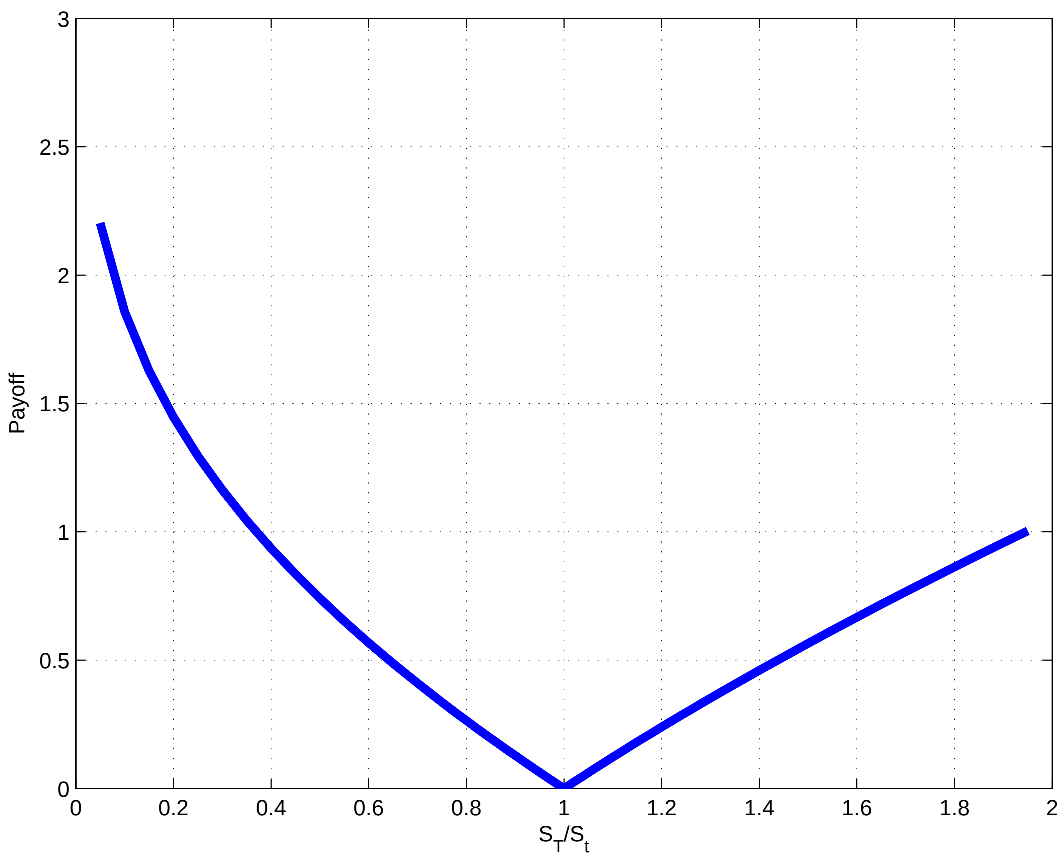

The correlation-immune SVS has some resemblance to a straddle, but its arms are not straight: the left arm is convex, and the right arm is concave.

6.3 Correlation-immune methodology

We improve the previous section's basic synthetic volatility swap to a correlation-immune synthetic volatility swap (SVS), which neutralizes the first-order impact of correlation.

Moreover, for hedging purposes, we will need valuations at all times $t \in [0, T]$, so "today" is now a generic time t instead of time 0.

Proposition 6.8 (Pricing a volatility swap using the correlation-immune SVS)**.** *For all* $t \in [0, T]$*,*

$$\mathbb{E}_t \sqrt{\langle X \rangle_T} = \mathbb{E}_t G_{\text{svs}}(S_T, S_t, \langle X \rangle_t) \tag{6.7}$$

where

$$G_{\text{svs}}(\mathsf{S}, \mathsf{u}, q) := \frac{1}{2\sqrt{\pi}} \int_0^\infty \theta_+ \frac{1 - e^{-zq}(\mathsf{S}/\mathsf{u})^{p_+}}{z^{3/2}} + \theta_- \frac{1 - e^{-zq}(\mathsf{S}/\mathsf{u})^{p_-}}{z^{3/2}} dz. \tag{6.8}$$

$$\theta_\pm := \theta_\pm(-z) := \frac{1}{2} \mp \frac{1}{2\sqrt{1 - 8z}} \qquad p_\pm := p_\pm(-z) := \frac{1}{2} \pm \frac{1}{2}\sqrt{1 - 8z} \tag{6.9}$$

In particular, we prove the convergence and integrability of G_{svs}*.*

For each t*, the payoff function* $F(\mathsf{S}) := G_{\text{svs}}(\mathsf{S}, S_t, \langle X \rangle_t)$ *is* ρ*-neutral.*

Remark 6.9*.* We call a claim on $G_{\text{svs}}(S_T, S_t, \langle X \rangle_t)$ the time-t *correlation-immune synthetic volatility swap* (SVS). If we simply say "synthetic volatility swap" or "SVS," we mean the correlation-immune variety, not the basic variety. Let SVS_t denote $\mathbb{E}_t G_{\text{svs}}(S_T, S_t, \langle X \rangle_t)$, the time-$t$ value of the SVS contract. Proposition 6.8 shows that SVS_t reveals the volatility swap value. Corollaries 6.12 and 6.13 will make explicit the observability of SVS_t, given call and put prices.

Remark 6.10*.* The correlation-immune SVS is *not* simply a linear combination of the put-like and call-like basic synthetic volatility swaps (6.6), because the linear combinations are taken *inside* the z-integral, and the weights $\theta_\pm$ depend on z. As shown in Figures 6.3–6.6, the SVS does resemble a straddle, but its arms are curved, not straight. Indeed, the three arguments of the payoff function $G_{\text{svs}}(\mathsf{S}, \mathsf{u}, q)$ have the following interpretation: S stands for the terminal share price; u represents the "strike" of the curved straddle; and q controls the "curvature" of the curved straddle. Proposition 6.8 shows that the "strike" should be chosen at-the-money and that the "curvature" should be chosen to depend on how much variance has been already accumulated.

At inception, the correlation-immune synthetic volatility swap may be written concisely in terms of Bessel functions. Let I_ν denote the modified Bessel function of order ν.

Corollary 6.11 (Payoff of newly-issued synthetic volatility swap: Bessel formula)**.** *We have*

$$\mathbb{E}_0 \sqrt{\langle X \rangle_T} = \mathbb{E}_0 \psi(S_T) \tag{6.10}$$

where $\psi(S) := \phi(\log(S/S_0))$, *where*

$$\phi(x) := \sqrt{\frac{\pi}{2}} e^{x/2} \Big| x I_0(x/2) - x I_1(x/2) \Big|. \tag{6.11}$$

The payoff is ρ*-neutral.*

Instead of expressing the synthetic volatility swap as a payoff function, we may express it as a mixture of put and call payoffs. We treat separately the case of a newly-issued volatility swap and the case of a seasoned volatility swap.

Corollary 6.12 (Put/call decomposition of newly-issued synthetic volatility swap: Bessel formula)**.** *The initial* ($\langle X \rangle_t = 0$) *correlation-immune synthetic volatility swap decomposes into the payoffs of*

$$\begin{aligned} &\sqrt{\pi/2}/S_0 \quad \textit{straddles at strike } K = S_0 \\ \sqrt{\frac{\pi}{8K^3 S_0}} \Big[I_1(\log\sqrt{K/S_0}) - I_0(\log\sqrt{K/S_0}) \Big] &dK \quad \textit{calls at strikes } K > S_0 \\ \sqrt{\frac{\pi}{8K^3 S_0}} \Big[I_0(\log\sqrt{K/S_0}) - I_1(\log\sqrt{K/S_0}) \Big] &dK \quad \textit{puts at strikes } K < S_0 \end{aligned} \tag{6.12}$$

Corollary 6.13 (Put/call decomposition of seasoned synthetic volatility swap)**.** *The seasoned* ($\langle X \rangle_t > 0$) *correlation-immune synthetic volatility swap decomposes into the payoffs of*

$$\begin{aligned} \frac{dK}{\sqrt{\pi}} \int_0^\infty \frac{e^{-z\langle X \rangle_t}}{K^2 z^{1/2}} \big[\theta_+ (K/S_t)^{p_+} + \theta_- (K/S_t)^{p_-} \big] dz \quad &\textit{calls at strikes } K > S_t, \textit{ puts at } K < S_t \\ \langle X \rangle_t^{1/2} \quad &\textit{bonds.} \end{aligned} \tag{6.13}$$

Remark 6.14. Our *basic* volatility valuation formula (6.6) is transformed by Friz-Gatheral [19] into one Bessel representation of the *basic* synthetic volatility swap. In contrast, in this section, we transform our *correlation-immune* volatility valuation formula (6.8) into two Bessel representations of our *correlation-immune* synthetic volatility swap (SVS), in Corollaries 6.11 (Bessel formula for payoff) and 6.12 (Bessel formula for put/call decomposition).

Our SVS provides not only valuation, but also replication of the volatility swap. Indeed, holding at each time t a claim on $G_{\text{svs}}(S_T, S_t, \langle X \rangle_t)$ replicates the volatility swap.

Proposition 6.15 (Synthetic volatility swap replicates the volatility swap)**.** *Holding at each time* t *a claim on* $G_{\text{svs}}(S_T, S_t, \langle X \rangle_t)$ *replicates the volatility swap. In other words:*

Choose an arbitrary constant $\kappa > 0$ *as a put/call separator. For* $K \in (0, \kappa)$ *let* $P_t(K)$ *denote the time-*t *value of a* K*-strike* T*-expiry binary put. For* $K \geq \kappa$ *let* $P_t(K)$ *denote the time-*t *value of a* K*-strike* T*-expiry binary call.*

Let the time-t binary option holdings (puts at strikes below κ, calls at strikes above κ) be given by the signed measure φ_t defined by the density function $K \mapsto \pm\partial G_{\mathrm{svs}}/\partial S(K; S_t, \langle X\rangle_t)$ on the domain $K \in (0,\infty)$, where the $+$ and $-$ correspond to $K > \kappa$ and $K < \kappa$ respectively.

Then the self-financing strategy of holding at each time t

$$\begin{aligned} \varphi_t &\quad \textit{options} \\ G_{\mathrm{svs}}(\kappa, S_t, \langle X\rangle_t) &\quad \textit{bonds} \end{aligned} \tag{6.14}$$

replicates the payoff $\sqrt{\langle X\rangle_T}$. Moreover, the strategy is ρ-neutral.

6.4 Evolution of the synthetic volatility swap

As variance accumulates during the life of the synthetic volatility swap, its payoff profile evolves. Proposition 6.8 makes this precise, but here let us give some intuition.

The initial payoff resembles a straddle struck at-the-money. The dynamics of the payoff depend on two factors. First, as the spot moves, the "strike" of the "straddle" floats to stay at-the-money. Second, as quadratic variation (an increasing process) accumulates, the "straddle" smooths out, losing its kink; indeed, only when $\langle X\rangle_t = 0$ does the kink literally exist.

We can, moreover, understand the limiting shape approached by the payoff. At time t, decompose $\langle X\rangle_T$ into the already-revealed portion $\langle X\rangle_t > 0$, and the random remaining variance $R_{t,T} := \langle X\rangle_T - \langle X\rangle_t$. By the square root function's concavity and (3.3),

$$\mathbb{E}\sqrt{\langle X\rangle_T} = \mathbb{E}\sqrt{\langle X\rangle_t + R_{t,T}} \le \sqrt{\langle X\rangle_t} + \frac{1}{2\sqrt{\langle X\rangle_t}}\mathbb{E}R_{t,T} \tag{6.15}$$

$$= \sqrt{\langle X\rangle_t} + \frac{1}{2\sqrt{\langle X\rangle_t}}\mathbb{E}_t\big[-2\log(S_T/S_t) + 2(S_T/S_t - 1)\big]. \tag{6.16}$$

As $\langle X\rangle_t$ increases, the intuition is that the square root function on $[\langle X\rangle_t, \infty)$ becomes less concave and more linear, hence the inequality (6.15) becomes an approximate equality. In view of (6.16), then, we expect that as time t rolls forward and $\langle X\rangle_t$ accumulates, the synthetic volatility swap will *evolve toward* a combination of synthetic *variance* swaps (3.3) and cash, with total time-T payoff

$$\sqrt{\langle X\rangle_t} + \frac{1}{\sqrt{\langle X\rangle_t}}(S_T/S_t - 1 - \log(S_T/S_t)). \tag{6.17}$$

This is visually confirmed in the right side of Figure 6.6, which compares the two time-T payoff functions (contracted at time t): the SVS $G_{\mathrm{svs}}(S_T, S_t, \langle X\rangle_t)$ and the log-contract-plus-cash (6.17).

Figure 6.4: At initiation ($\langle X \rangle_t = 0.0$), the volatility swap and synthetic volatility swap (SVS)

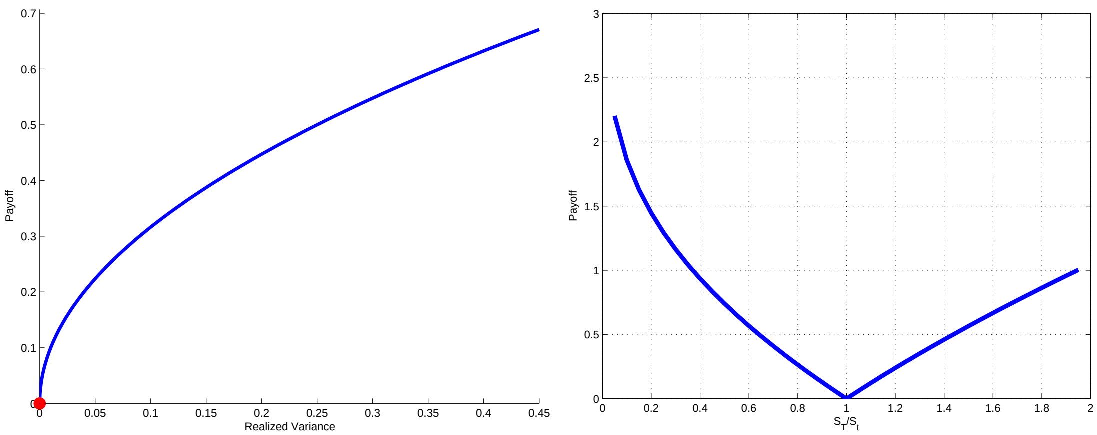

Figure 6.5: Seasoned ($\langle X \rangle_t = 0.1$) volatility swap and synthetic volatility swap (SVS)

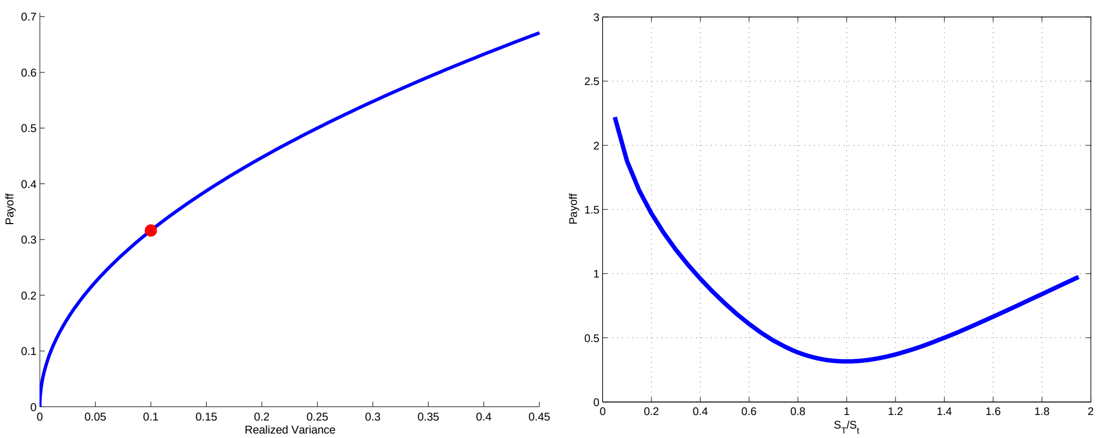

Figure 6.6: Seasoned ($\langle X \rangle_t = 0.25$) volatility swap and SVS, compared to variance swaps

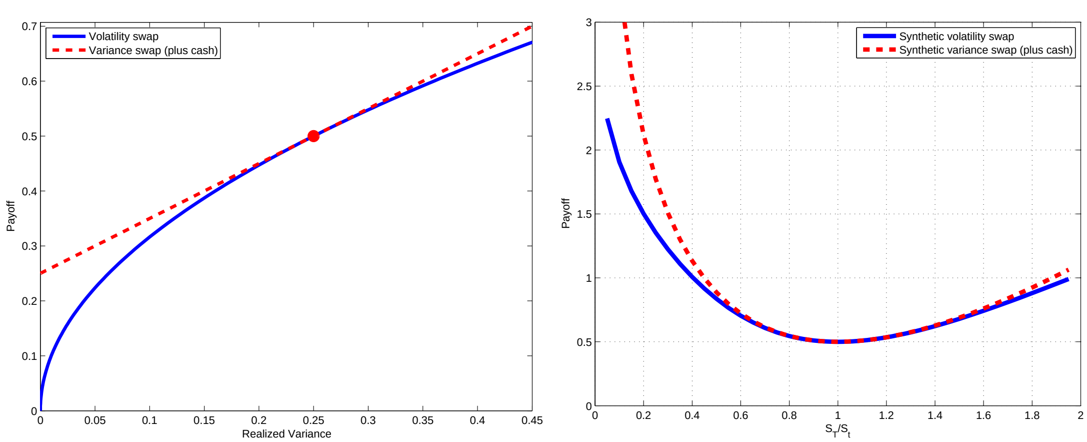

6.5 Accuracy of the ρ-neutral synthetic volatility swap

Figure 6.7 shows how closely the time-0 ρ-neutral synthetic volatility swap (SVS) price approximates the true volatility swap fair value, under Heston dynamics with parameters from Bakshi-Cao-Chen [2]. For comparison, we plot also the ATM implied volatility, and the basic (correlation-sensitive) synthetic volatility swap price.

As approximations of the true volatility swap value, our correlation-immune SVS outperforms ATM implied volatility and outperforms our basic (correlation-sensitive) replication – across essentially all correlation assumptions. In the case $\rho = 0$, both of our methods are (as promised) exact and the implied volatility approximation is nearly exact; but more importantly, in the empirically relevant case of $\rho \neq 0$, our correlation-immune SVS's relative "flatness" with respect to ρ results in its greater accuracy. This illustrates why, in equity markets, we do not recommend any method or approximation which relies on assumption (I), unless it has the additional correlation-immunity present in our SVS.

Figure 6.7: Heston dynamics: Volatility swap valuations as functions of correlation

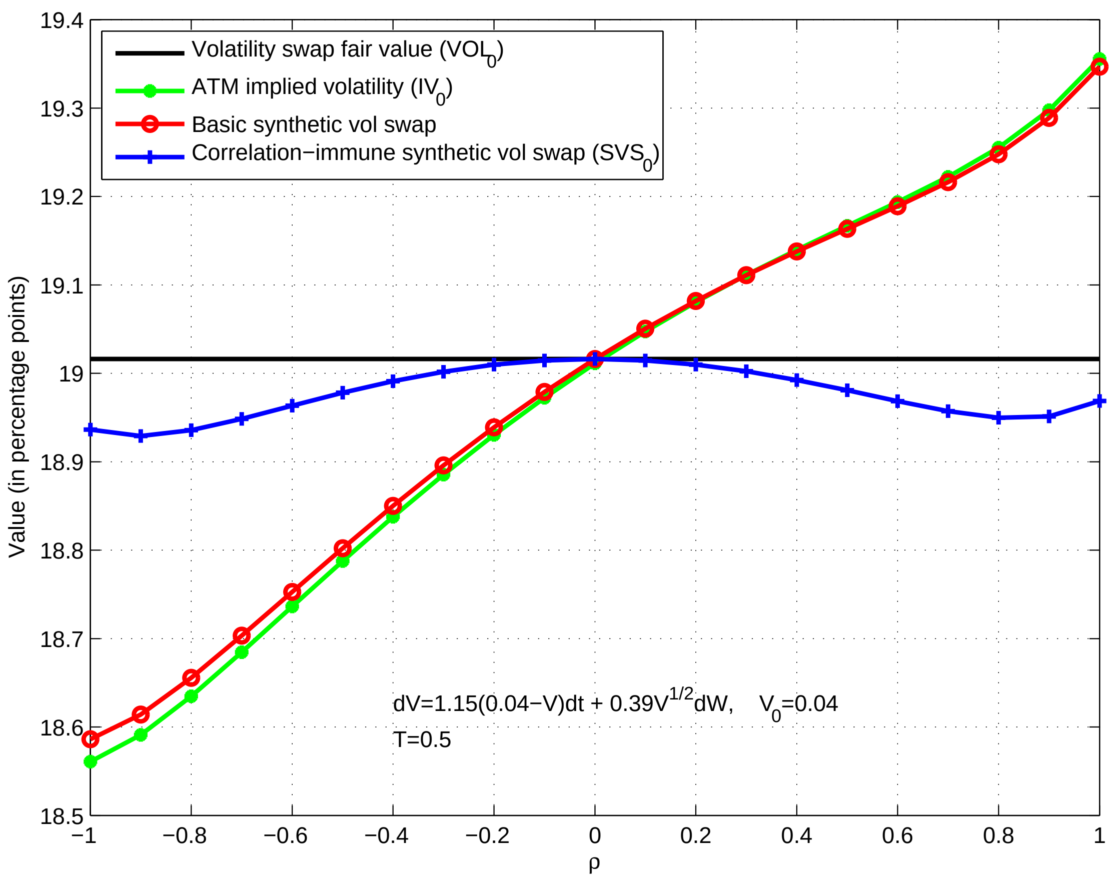

We comment on each curve in greater detail.

The volatility swap fair value (denoted by $\mathrm{VOL}_0 := \mathbb{E}_0\sqrt{\langle X\rangle_T}$ as in Section 6.1) equals the expectation of realized volatility. It is determined by the distribution of realized variance $\int V_t dt$, which is determined entirely by the given dynamics

$$dV_t = 1.15(0.04 - V_t)dt + 0.39V_t^{1/2}dW_t, \quad V_0 = 0.04 \tag{6.18}$$

of instantaneous variance $V_t = \sigma_t^2$. So the correlation ρ is irrelevant to VOL_0, which therefore plots as a horizontal line. Its level 0.1902 is computable via the known distribution of $\int V_t dt$ given (6.18).

The basic (correlation-sensitive) synthetic volatility swap payoff is approximately the payoff of $\sqrt{2\pi}/S_0$ calls, as noted in Remark 6.7. Therefore its value and the ATM Black-Scholes implied volatility IV_0 are nearly equal, due to (6.4). The plots confirm this across the full range of ρ. More importantly, the plots confirm that VOL_0 is well-approximated by these two values *if* $\rho = 0$, but due to the correlation-sensitivity of IV_0 and of the basic synthetic volatility swap, both values underestimate VOL_0 by more than 40 basis points, for certain values of ρ.

Our correlation-immune SVS has value SVS_0 which, as promised, exactly matches VOL_0 if $\rho = 0$. Furthermore, as intended by its design, SVS_0 is ρ-invariant to first-order, at $\rho = 0$. There is no guarantee that this flatness will extend to ρ far from 0, but for these parameters the ρ-neutrality does indeed result in accuracy gains across the entire range of ρ, as confirmed in the plot.

Finally we comment on a benchmark not plotted in the figure. The variance swap value (which equals the log-contract value) is 0.04; and its square root (which we denote by $\mathrm{VAR}_0 = \sqrt{\mathbb{E}_0\langle X\rangle_T}$ as in Section 6.1, and which the VIX seeks to approximate) is 0.20, regardless of ρ. Therefore, a plot of VAR_0 would be a horizontal line far above the upper boundary of Figure 6.7, and would not be a competitive approximation to $\mathrm{VOL}_0 = 0.1902$.

To summarize, in this example the *best* approximation of VOL_0, for essentially all $\rho \in [-1, 1]$, is given by our correlation-immune SVS value (SVS_0), and the worst is given by the VIX-style quantity VAR_0. The other approximations – ATM implied volatility IV_0 and the basic (correlation-sensitive) volatility swap value – are accurate for $\rho = 0$ but lose accuracy for ρ nonzero.

Remark 6.16. Figure 6.7 can be regarded as a numerical comparison of two notions of "model-free implied volatility" (MFIV). When defined in the "VIX-style," MFIV is understood to mean VAR_0, the square root of the variance swap (or log contract) value. Here we have introduced the correlation-immune *synthetic volatility swap*, whose observable value we regard as an alternative notion of MFIV. Indeed, let us define "SVS-style" MFIV to be SVS_0, the time-0 value of our SVS.

Our SVS-style MFIV is truly an expected *volatility*, because it does indeed equal VOL_0, by Proposition 6.8 – in contrast to the VIX-style MFIV which equals VAR_0, the square root of expected *variance*. Moreover, although Proposition 6.8 assumes (I), we observe that even in the (I)-violating $\rho \neq 0$ dynamics of Figure 6.7, the expected volatility VOL_0 is still approximated much more accurately by our SVS-style MFIV (with errors of only 9 basis points even in the worst cases near $\rho = -1$) than by the VIX-style MFIV (with errors of 98 basis points).

7 Pricing other volatility derivatives

Using exponential variance payoffs, we can price general variance payoffs.

7.1 Fractional or negative power payoffs

Our volatility swap formula is the $r = 1/2$ case of the following generalization to powers in $(0, 1)$.

Proposition 7.1. *For $0 < r < 1$,*

$$\mathbb{E}_t \langle X \rangle_T^r = \mathbb{E}_t G_{\mathrm{pow}(r)}(S_T, S_t, \langle X \rangle_t)$$

where

$$G_{\mathrm{pow}(r)}(\mathsf{S}, \mathsf{u}, q) := \frac{r}{\Gamma(1-r)} \int_0^\infty \theta_+ \frac{1 - e^{-zq}(\mathsf{S}/\mathsf{u})^{p_+}}{z^{r+1}} + \theta_- \frac{1 - e^{-zq}(\mathsf{S}/\mathsf{u})^{p_-}}{z^{r+1}} dz \tag{7.1}$$

$$\theta_\pm := \theta_\pm(-z) := \frac{1}{2} \mp \frac{1}{2\sqrt{1-8z}} \qquad p_\pm := p_\pm(-z) := \frac{1}{2} \pm \frac{1}{2}\sqrt{1-8z} \tag{7.2}$$

For each t, the payoff function $\mathsf{S} \mapsto G_{\mathrm{pow}(r)}(\mathsf{S}, S_t, \langle X \rangle_t)$ is ρ-neutral.

For arbitrary negative powers, we have the following formula for "inverse variance" claims.

Proposition 7.2. *For any $r > 0$ and any ε such that $\langle X \rangle_t + \varepsilon > 0$,*

$$\mathbb{E}_t(\langle X \rangle_T + \varepsilon)^{-r} = \mathbb{E}_t G_{\mathrm{pow}(-r)}(S_T, S_t, \langle X \rangle_t + \varepsilon)$$

where

$$G_{\mathrm{pow}(-r)}(\mathsf{S}, \mathsf{u}, q) := \frac{1}{r\Gamma(r)} \int_0^\infty (\theta_+(\mathsf{S}/\mathsf{u})^{p_+} + \theta_-(\mathsf{S}/\mathsf{u})^{p_-}) e^{-z^{1/r} q} dz$$

$$\theta_\pm := \theta_\pm(-z^{1/r}) := \frac{1}{2} \mp \frac{1}{2\sqrt{1-8z^{1/r}}} \qquad p_\pm := p_\pm(-z^{1/r}) := \frac{1}{2} \pm \sqrt{1/4 - 2z^{1/r}}.$$

For each t, the payoff function $F(\mathsf{S}) := G_{\mathrm{pow}(-r)}(\mathsf{S}, S_t, \langle X \rangle_t)$ is ρ-neutral.

Figure 7.1: Polynomial variance claims $\langle X\rangle_T^n$ on the left, and their European-style synthetic counterparts $G_{\mathrm{pow}(n)}(S_T, S_0, \langle X\rangle_0)$ on the right, for $n = 1, 2, 3$ and $\langle X\rangle_0 = 0$

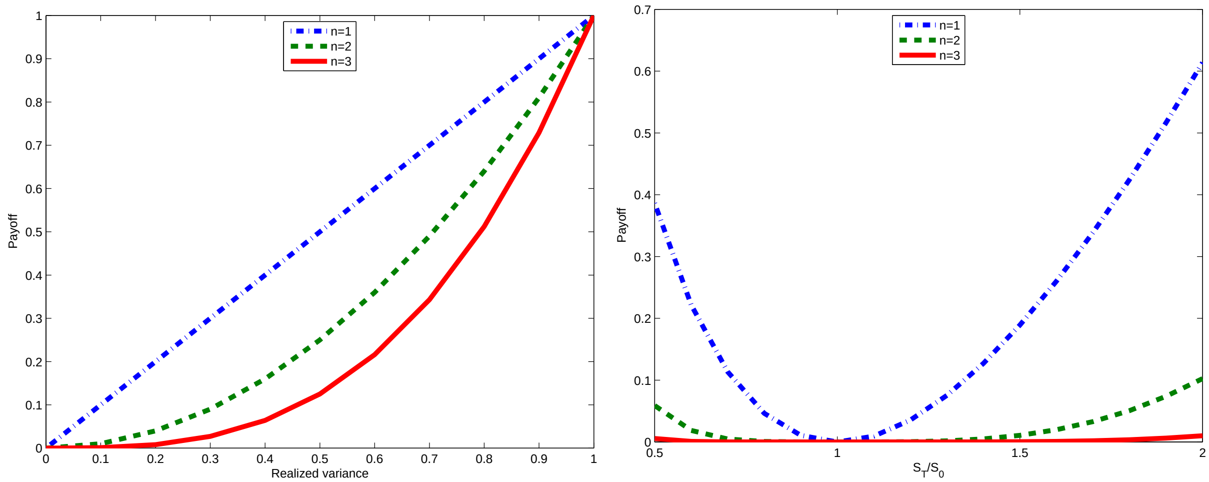

7.2 Polynomial payoffs

We obtain polynomials in variance by differentiating, in λ, the exponential of $\lambda\langle X\rangle_T$.

Proposition 7.3. *For each positive integer n,*

$$\mathbb{E}_t\langle X\rangle_T^n = \mathbb{E}_t G_{\mathrm{pow}(n)}(S_T, S_t, \langle X\rangle_t)$$

where

$$G_{\mathrm{pow}(n)}(S, u, q) := \partial_\lambda^n G_{\exp}(S, u, q, \lambda)\big|_{\lambda=0} \tag{7.3}$$

with $G_{\exp}$ defined in (5.10). In particular, for $n = 1, 2, 3$:

$$\mathbb{E}_0\langle X\rangle_T = \mathbb{E}_0(-2X_T + 2e^{X_T} - 2)$$

$$\mathbb{E}_0\langle X\rangle_T^2 = \mathbb{E}_0(4X_T^2 + 16X_T + 8X_Te^{X_T} - 24e^{X_T} + 24)$$

$$\mathbb{E}_0\langle X\rangle_T^3 = \mathbb{E}_0(-8X_T^3 + 24X_T^2e^{X_T} - 72X_T^2 - 192X_Te^{X_T} - 288X_T + 480e^{X_T} - 480).$$

For each t, the payoff function $F(S) := G_{\mathrm{pow}(n)}(S, S_t, \langle X\rangle_t)$ is ρ-neutral.

Note that $n = 1$ recovers the usual valuation of the variance swap using a hedged log contract. Figure 7.1 plots $G_{\mathrm{pow}(n)}$ for $n = 1, 2, 3$.

7.3 Payoffs whose transforms decay exponentially

In Sections 7.3 to 7.5 we make use of exponential variance payoffs as *basis* functions, to span a space of general variance payoff functions h.

Definition 7.4 (Bilateral Laplace transform)**.** For any continuous $h : \mathbb{R} \to \mathbb{R}$, and any $\alpha \in \mathbb{R}$ such that $\int_{-\infty}^{\infty} e^{-\alpha q} h(q) dq < \infty$, define for $\mathrm{Re}(z) = \alpha$

$$H(z) := \int_{-\infty}^{\infty} e^{-zq} h(q) dq. \tag{7.4}$$

Proposition 7.5 (Variance contracts in terms of Europeans, under decay conditions)**.** *Under Definition 7.4, assume that $|H(\alpha + \beta i)| = O(e^{-|\beta|\mu})$ as $|\beta| \to \infty$ for some $\mu > m/2$. Then*

$$\mathbb{E}_t h(\langle X \rangle_T) = \mathbb{E}_t G_h(S_T, S_t, \langle X \rangle_t)$$

where

$$G_h(\mathsf{S}, \mathsf{u}, q) := \frac{1}{2\pi i} \int_{\alpha - \infty i}^{\alpha + \infty i} H(z) e^{zq} [\theta_+ (\mathsf{S}/\mathsf{u})^{p_+} + \theta_- (\mathsf{S}/\mathsf{u})^{p_-}] dz \tag{7.5}$$

$$\theta_\pm := \theta_\pm(z) := \frac{1}{2} \mp \frac{1}{2\sqrt{1+8z}} \qquad p_\pm := p_\pm(z) := \frac{1}{2} \pm \sqrt{1/4 + 2z}.$$

In particular, we prove the convergence and finite expectation of G_h.

For each t, the payoff function $\mathsf{S} \mapsto G_h(\mathsf{S}, S_t, \langle X \rangle_t)$ is ρ-neutral.

Remark 7.6. Recall the heuristic that the smoother a function, the more rapid the decay of its transform. For insufficiently smooth h (such as payoffs of puts/calls on volatility), the transform H does not decay rapidly enough to satisfy the assumption of Proposition 7.5. Such payoffs can be treated by Propositions 7.7 through 7.13, which weaken the assumptions on h.

For payoff functions h smooth enough to satisfy the Proposition 7.5 assumption, we have proved that the volatility contract has the same value as the European contract with payoff $G_h(S_T, S_t, \langle X \rangle_t)$, defined by the convergent integral in (7.5). Although this payoff G_h may be oscillatory in S_T, Proposition 7.5 guarantees that the payoff has a well-defined price, in the sense that the payoff's positive and negative components each have finite expectation.

Observation of the G_h payoff's price, from a practical standpoint, may be a non-trivial issue, if the G_h payoff profile has significant curvature at price levels which happen to lack liquid vanilla option strikes. In such cases, regularization of the payoff profile can be achieved by projecting h onto a finite set of basis functions, as we do in Proposition 7.13. Alternatively, in contrast to this payoff replication approach, a different approach, by Friz-Gatheral [19], conducts distributional inference, using a finite set of pricing benchmarks, in conjunction with Tikhonov-style regularization.

7.4 Payoffs whose transforms are integrable

If instead of having exponential decay, the payoff's transform is merely integrable, then our usual pricing formulas of the form $\mathbb{E}h(\langle X\rangle_T) = \mathbb{E}G(S_T)$ may not be available by the Laplace transform method. Nonetheless, the prices of claims on S_T do still determine the price of the $h(\langle X\rangle_T)$ contract.

Proposition 7.7 (Inverting an integrable transform)**.** *Under Definition 7.4, assume that H is integrable along* $\mathrm{Re}(z) = \alpha$. *Let V_T be a random variable. If $\mathbb{E}_t e^{\alpha V_T} < \infty$, then*

$$\mathbb{E}_t h(V_T) = \frac{1}{2\pi i}\int_{\alpha-\infty i}^{\alpha+\infty i} H(z)\mathbb{E}_t e^{zV_T} dz. \tag{7.6}$$

Corollary 7.8 (Variance and volatility puts, without assuming $(\mathrm{B}, \mathrm{W}, \mathrm{I})$)**.** *Let V_T be the quadratic variation of an* arbitrary *semimartingale (not necessarily X). For a Q-strike* realized variance put *where $h(q) := (Q-q)^+$ hence*

$$H(z) = \frac{e^{-Qz}}{z^2}, \tag{7.7}$$

or for a $\sqrt{Q}$-strike realized volatility put *where $h(q) := (\sqrt{Q} - \sqrt{q^+})^+$ hence*

$$H(z) = -\frac{\sqrt{\pi}\,\mathrm{erf}(\sqrt{zQ})}{2z^{3/2}}, \tag{7.8}$$

we have for all $\alpha < 0$ the formula (7.6) *for the put price $\mathbb{E}_t h(V_T)$.*

Variance and volatility call prices follow from put-call parity.

One application is in cases where $\mathbb{E}_t e^{zV_T}$ has an explicit formula, such as in affine diffusion or jump-diffusion models, including Heston. Then (7.6) with (7.7) or (7.8) gives explicit formulas for variance and volatility options respectively.

Another application of Proposition 7.7 is in cases where $\mathbb{E}_t e^{zV_T}$ has no explicit formula, but can be inferred model-independently from Europeans, such as the case $V_T := \langle X\rangle_T = \langle \log S\rangle_T$, for any process S that satisfies $(\mathrm{B}, \mathrm{W}, \mathrm{I})$. The next two corollaries pursue this.

Corollary 7.9 (Variance contracts in terms of Europeans)**.** *Under Definition 7.4, assume that H is integrable along* $\mathrm{Re}(z) = \alpha$ *where $\alpha \in \mathbb{R}$. Then*

$$\mathbb{E}_t h(\langle X\rangle_T) = \frac{1}{2\pi i}\int_{\alpha-\infty i}^{\alpha+\infty i} H(z) e^{z\langle X\rangle_t}\mathbb{E}_t[\theta_+(S_T/S_t)^{p_+} + \theta_-(S_T/S_t)^{p_-}]dz \tag{7.9}$$

where

$$\theta_\pm := \theta_\pm(z) := \frac{1}{2} \mp \frac{1}{2\sqrt{1+8z}} \qquad p_\pm := p_\pm(z) := \frac{1}{2} \pm \sqrt{1/4 + 2z}.$$

In particular, we prove the convergence of the integral.

Corollary 7.10 (Variance/volatility puts/calls in terms of Europeans). *For the variance put* $h(q) := (Q - q)^+$ *or the volatility put* $h(q) := (\sqrt{Q} - \sqrt{q^+})^+$*, define* H *by* (7.7) *or* (7.8) *respectively. Then we have for all* $\alpha < 0$ *the formula* (7.9) *for the put price* $\mathbb{E}_t h(\langle X\rangle_T)$*.*

For a Q*-strike* realized variance call *where* $h(q) = (q - Q)^+$ *hence*

$$H(z) = \frac{e^{-Qz}}{z^2}, \tag{7.10}$$

or for a $\sqrt{Q}$*-strike* realized volatility call *where* $h(q) = (\sqrt{q^+} - \sqrt{Q})^+$ *hence*

$$H(z) = \frac{\sqrt{\pi}\,\mathrm{erfc}(\sqrt{zQ})}{2z^{3/2}}, \tag{7.11}$$

we have for all $\alpha > 0$ *the formula* (7.9) *for the call price* $\mathbb{E}_t h(\langle X\rangle_T)$*.*

Remark 7.11. Relative to the results of previous sections, Corollary 7.9 has greater generality, but also has a possible drawback: To price a variance contract exactly using Corollary 7.9 requires the valuation of infinitely many different functions of S_T (one for each z). In contrast, using Propositions 6.8, 7.1, 7.2, 7.3, 7.5, to price one variance contract exactly requires the valuation of a *single* function of S_T.

If, instead of an exact formula, we accept (a sequence of) approximate prices which converge to the exact price, then an even more general class of variance contracts can be priced using (a sequence of) *single* functions of S_T. That is the subject of the next section.

7.5 General payoffs continuous on $[0, \infty]$

Let $C[0,\infty]$ denote the set of continuous $h : [0,\infty) \to \mathbb{R}$ such that $h(\infty) := \lim_{q\to\infty} h(q)$ exists in $\mathbb{R}$. For example, the *variance put* payoff $h(q) = (Q - q)^+$ belongs to $C[0,\infty]$.

This section gives two ways to determine prices of general payoffs in $C[0,\infty]$. The first will take limits of uniform approximations, and the second will take limits of mean-square approximations.

Although call payoffs do not belong to $C[0,\infty]$, they can still be priced by the methods of this section, using put-call parity: a variance call equals a variance put plus a variance swap.

In this section let $h \in C[0,\infty]$ and let $c > 0$ be an arbitrary constant.

Proposition 7.12 (Prices as limits of uniform approximations' prices). *Define* $h^* : [0,1] \to \mathbb{R}$ *by* $h^*(0) := h(\infty)$ *and* $h^*(x) := h(-(1/c)\log x)$ *for* $x > 0$*. For integers* $n \geq k \geq 0$*, let*

$$b_{n,k} := \sum_{j=0}^{k} h^*(j/n)\binom{n}{k}\binom{k}{j}(-1)^{k-j}. \tag{7.12}$$

Then

$$\mathbb{E}_t h(\langle X\rangle_T) = \lim_{n\to\infty} \mathbb{E}_t \sum_{k=0}^{n} b_{n,k} e^{-ck\langle X\rangle_t}[\theta_+(S_T/S_t)^{p_+} + \theta_-(S_T/S_t)^{p_-}], \tag{7.13}$$

where

$$\theta_\pm := \frac{1}{2} \mp \frac{1}{2\sqrt{1-8ck}} \qquad p_\pm := \frac{1}{2} \pm \sqrt{1/4 - 2ck}. \tag{7.14}$$

In particular, we prove the existence of the limit.

Proposition 7.13 (Prices as limits of L^2 projections' prices)**.** *Let μ be a finite measure on $[0,\infty)$. Let*

$$a_{n,n}e^{-cnq} + a_{n,n-1}e^{-c(n-1)q} + \cdots + a_{n,0} =: A_n(q)$$

be the $L^2(\mu)$ projection of h onto $\operatorname{span}\{1, e^{-cq}, \ldots, e^{-cnq}\}$*. Let P denote the $\mathbb{Q}$-distribution of $\langle X\rangle_T$, conditional on $\mathcal{F}_t$. Assume P is absolutely continuous with respect to μ and $d\mathsf{P}/d\mu \in L^2(\mu)$. Then*

$$\mathbb{E}_t h(\langle X\rangle_T) = \lim_{n\to\infty} \mathbb{E}_t \sum_{k=0}^{n} a_{n,k} e^{-ck\langle X\rangle_t}[\theta_+(S_T/S_t)^{p_+} + \theta_-(S_T/S_t)^{p_-}] \tag{7.15}$$

where

$$\theta_\pm := \frac{1}{2} \mp \frac{1}{2\sqrt{1-8ck}} \qquad p_\pm := \frac{1}{2} \pm \sqrt{1/4 - 2ck}. \tag{7.16}$$

In particular, we prove the existence of the limit.

Remark 7.14. For each n, the $a_{n,k}$ $(k = 0, \ldots, n)$ are given by the solution to the linear system

$$\sum_{k=0}^{n} a_{n,k}\langle e^{-cjq}, e^{-ckq}\rangle = \langle h(q), e^{-cjq}\rangle, \qquad j = 0, \ldots, n \tag{7.17}$$

of normal equations, where $\langle \alpha(q), \beta(q)\rangle := \int_0^\infty \alpha(q)\beta(q)\, d\mu(q)$. In practice, one can compute $a_{n,k}$ as the coefficients in a weighted least squares regression of the $h(q)$ function on the regressors $\{q \mapsto e^{-ckq} : k = 0, \ldots, n\}$, with weights given by the measure μ.

For example, consider the variance put payoff $h(\langle X\rangle_T) = (0.04 - \langle X\rangle_T)^+$ with expiry $T = 1$. Under the Heston variance dynamics specified in Figure 6.7 with $\rho = 0$, let us compare the put's true time-0 value $\mathbb{E}h(\langle X\rangle_T)$ against the sequence of European prices in the right-hand side of (7.15). For example, let $c = 0.5$, and let μ be the lognormal distribution whose parameters are consistent with the values of T-expiry variance and volatility swaps (which are observable from European options, by Propositions 6.8 and 7.3). We compute:

$\mathbb{E}A_3(\langle X\rangle_T)$	$\mathbb{E}A_4(\langle X\rangle_T)$	$\mathbb{E}A_5(\langle X\rangle_T)$	$\cdots$	$\mathbb{E}h(\langle X\rangle_T)$
0.01108	0.01133	0.01147	$\cdots$	0.01149

(7.18)

Here small values of n have sufficed to produce an accurate approximation of $\mathbb{E}h(\langle X\rangle_T)$.

Remark 7.15. In principle, each A_n and B_n function admits perfect pricing by European options, via (7.13) and (7.15) respectively; in practice, the convergence benefits of incrementing n must be considered in the context of whether the available European options data (which may have noisy or missing observations) can provide sufficient resolution.

Remark 7.16. Each A_n and B_n function is a linear combination of exponentials, hence admits perfect *replication* by European options, according to Proposition 5.11. Consequently, by the explicit uniform approximation (A.10), any variance payoff continuous on $[0, \infty]$ can be replicated to within an arbitrarily small uniform error.

8 Extension to unbounded quadratic variation

Here we show how to drop the assumption (B) that $\langle X \rangle_T \leqslant m$ for some constant m.

For practical purposes, it could be argued that a bound of, say, $m = 10^{10}T$ may be an acceptable assumption for an equity index. However, for dynamics such as the Heston model, (B) does not hold for any m. This section extends our framework to include such dynamics.

Proposition 8.1 (Unbounded quadratic variation)**.** *Assume the measurable functions h and G satisfy*

$$\mathbb{E}h(\langle X \rangle_T) = \mathbb{E}G(S_T) \tag{8.1}$$

for all S which satisfy (B, W, I)*.*

Assume that h is bounded or *that h is nonnegative and increasing.*

Assume that G has a decomposition $G = G_1 - G_2$, where $G_{1,2}$ are convex and $\mathbb{E}G_{1,2}(S_T) < \infty$.

Then (8.1) *holds, more generally, for all S which satisfy* (W) *and* (I) *and $\mathbb{E}\langle X \rangle_T < \infty$.*

Remark 8.2. The finiteness of $\mathbb{E}h(\langle X \rangle_T)$ is a conclusion, not an assumption.

Remark 8.3. The assumptions on G are very mild, in the following sense: They are satisfied by any payoff function which can be represented as a mixture of calls and puts at all strikes, such that the long and short positions have finite values.

Corollary 8.4. *Propositions 5.1, 5.9 on exponential variance valuation, Propositions 6.6, 6.8 on volatility swap valuation, Propositions 7.1, 7.2, 7.3 on valuation of fractional and integer powers of variance, and Proposition 7.5 on valuation by Laplace transform, all hold* without *assuming* (B) *– provided that the long and short positions in calls and puts in the replicating portfolios have finite values.*

9 Conclusion

Contracts on general functions of realized variance, which allow investors to manage their exposure to volatility risk, have presented to dealers significant challenges in pricing and hedging. For pricing purposes, we derive explicit valuation formulas for such contracts, in terms of vanilla option prices – not in terms of the parameters of any model. The formulas are exact under an independence condition, and they are first-order immunized against the presence of correlation. For hedging purposes, we enforce these valuation formulas by replicating the variance payoffs using explicit trading strategies in vanilla options and the underlying shares.

Future research can extend the dynamics we study and the risks we hedge. This paper, which already allows unspecified jumps in the instantaneous volatility, moreover lays the foundation for the addition of jumps to the price paths; and this paper's analysis of volatility risk contributes to a broad research program which nonparametrically utilizes European options, to extract information about path-dependent risks, and to hedge those risks robustly.

A Appendix: Proofs

Proof of Proposition 4.1. We have

$$\begin{aligned} dX_t &= -\frac{1}{2}\sigma_t^2 dt + \sqrt{1-\rho^2}\sigma_t dW_{1t} + \rho\sigma_t dW_{2t} \\ &= -\frac{1-\rho^2}{2}\sigma_t^2 dt + \sqrt{1-\rho^2}\sigma_t dW_{1t} - \frac{\rho^2}{2}\sigma_t^2 dt + \rho\sigma_t dW_{2t} \end{aligned}$$

So conditional on $\mathcal{H}_T \vee \mathcal{F}_t$,

$$X_T \sim \text{Normal}\Big(X_t + \log M_{t,T}(\rho) - \bar{\sigma}_{t,T}^2 \frac{1-\rho^2}{2},\ \bar{\sigma}_{t,T}\sqrt{1-\rho^2}\Big).$$

Hence

$$\mathbb{E}_t F(S_T) = \mathbb{E}_t(\mathbb{E}(F(S_T)|\mathcal{H}_T \vee \mathcal{F}_t)) = \mathbb{E}_t F^{BS}(S_t M_{t,T}(\rho),\ \bar{\sigma}_{t,T}\sqrt{1-\rho^2})$$

as desired. □

Proof of Proposition 5.1. We apply a general version of Hull-White's [21] conditioning argument. Conditional on $\mathcal{F}_T^\sigma$, the W is still a Brownian motion, by independence. So conditional on $\mathcal{F}_t \vee \mathcal{F}_T^\sigma$,

$$X_T - X_t = \int_t^T \sigma_u dW_u - \frac{1}{2}(\langle X\rangle_T - \langle X\rangle_t) \quad \sim \quad \text{Normal}\Big(-\frac{\langle X\rangle_T - \langle X\rangle_t}{2},\ \langle X\rangle_T - \langle X\rangle_t\Big).$$

For each $p \in \mathbb{C}$, therefore,

$$\begin{aligned}
\mathbb{E}_t e^{p(X_T - X_t)} &= \mathbb{E}_t\Big[\mathbb{E}(e^{p(X_T-X_t)}|\mathcal{F}_t \vee \mathcal{F}_T^\sigma)\Big] \\
&= \mathbb{E}_t\Big[e^{\mathbb{E}(pX_T - pX_t|\mathcal{F}_t\vee\mathcal{F}_T^\sigma)+\mathrm{Var}(pX_T-pX_t|\mathcal{F}_t\vee\mathcal{F}_T^\sigma)/2}\Big] \\
&= \mathbb{E}_t e^{(p^2/2-p/2)(\langle X\rangle_T - \langle X\rangle_t)} = \mathbb{E}_t e^{\lambda(\langle X\rangle_T - \langle X\rangle_t)},
\end{aligned}$$

where $\lambda = p^2/2 - p/2$. Equivalently, $p = \frac{1}{2} \pm \sqrt{\frac{1}{4} + 2\lambda}$. □

Proof of Proposition 5.3. Our portfolio at each time t has value $N_tP_t - (pN_tP_{t-}/S_t)S_t + pN_tP_{t-} = N_tP_t$. In particular it has the desired time-T value $N_TP_T = e^{\lambda\langle X\rangle_T}$. To prove that it self-finances,

$$\begin{aligned}
d(N_tP_t) &= N_t dP_t + P_{t-}dN_t + d[P,N]_t \\
&= N_t dP_t + P_{t-}\left(\frac{-pN_t}{S_t}dS_t\right) + dA_t,
\end{aligned}$$

where A has finite variation. The continuity of S implies the continuity of N, hence $[P,N]$, hence A. Moreover, A is a local martingale because N_tP_t ($= \mathbb{E}_t e^{\lambda\langle X\rangle_T}$ by Proposition 5.1) and the stochastic integrals with respect to P and S are all local martingales. So dA vanishes. Therefore

$$d(N_tP_t) = N_t dP_t - (pN_tP_{t-}/S_t)dS_t + pN_tP_{t-}dB_t$$

because $dB = 0$. This proves self-financing. □

Proof of Proposition 5.9. The weights $\theta_\pm$ have the properties that $\theta_+ + \theta_- = 1$ and $\theta_+p_+ + \theta_-p_- = 0$.

The first property, together with Remark 5.8, implies (5.9). To see that the second property implies ρ-neutrality, let ϕ_v be the lognormal density with parameters $(-v/2, v)$. Then

$$\begin{aligned}
\frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t) &= e^{\lambda\langle X\rangle_t}\frac{\partial}{\partial \mathsf{s}}\bigg|_{\mathsf{s}=S_t}\int_0^\infty \Big[\theta_+(\mathsf{s}/S_t)^{p_+}y^{p_+} + \theta_-(\mathsf{s}/S_t)^{p_-}y^{p_-}\Big]\phi_v(y)dy \\
&= e^{\lambda\langle X\rangle_t}\int_0^\infty\left(\theta_+\frac{p_+}{S_t}y^{p_+} + \theta_-\frac{p_-}{S_t}y^{p_-}\right)\phi_v(y)dy = e^{\lambda\langle X\rangle_t}\frac{\theta_+p_+ + \theta_-p_-}{S_t}\int_0^\infty y^{p_+}\phi_v(y)dy = 0
\end{aligned}$$

using the equality of integrals of $y^{p_+}\phi_v(y)$ and $y^{p_-}\phi_v(y)$. □

Proof of Proposition 5.11. The strategy is a linear combination of the two strategies $(+,-)$ specified in Proposition 5.3, with constant weights θ_+ and θ_- which sum to 1. Each strategy self-finances and replicates $e^{\lambda\langle X\rangle_T}$, so the combination does also. Proposition 5.9 implies ρ-neutrality. □

Proof of Proposition 6.1. The upper bound (c) is known (Britten-Jones/Neuberger [11]) to hold, by Jensen's inequality.

For (b), we have by Proposition 4.1 and the concavity of F^{BS}_{atmc},

$$F^{BS}_{\text{atmc}}(S_0, \text{IV}_0) = \mathbb{E}_0 F_{\text{atmc}}(S_T) = \mathbb{E}_0 F^{BS}_{\text{atmc}}(S_0, \bar{\sigma}_{0,T}) \le F^{BS}_{\text{atmc}}(S_0, \mathbb{E}_0 \bar{\sigma}_{0,T}). \tag{A.1}$$

By the monotonicity of F^{BS}_{atmc}, therefore, $\text{IV}_0 \le \mathbb{E}_0 \bar{\sigma}_{0,T}$. For (a),

$$\frac{\sqrt{2\pi}}{S_0}\mathbb{E}_0(S_T - S_0)^+ = \frac{\sqrt{2\pi}}{S_0} F^{BS}_{\text{atmc}}(S_0, \text{IV}_0) \le \frac{\sqrt{2\pi}}{S_0}\frac{S_0\,\text{IV}_0}{\sqrt{2\pi}} = \text{IV}_0 \tag{A.2}$$

because concavity implies that $F^{BS}_{\text{atmc}}(S_0, \cdot)$ lies everywhere below its tangent at 0. □

Proof of Proposition 6.6. Of the $\pm$, we prove the $+$ equation; the $-$ equation is similar. We have

$$\sqrt{q} = \frac{1}{2\sqrt{\pi}} \int_0^\infty \frac{1 - e^{-zq}}{z^{3/2}} dz \qquad \text{for all } q \ge 0,$$

as shown in sources such as Schürger [29]. So

$$\begin{aligned} \mathbb{E}_0 \sqrt{\langle X \rangle_T} &= \frac{1}{2\sqrt{\pi}} \int_0^\infty \mathbb{E}_0 \frac{1 - e^{-z\langle X \rangle_T}}{z^{3/2}} dz = \frac{1}{2\sqrt{\pi}} \int_0^\infty \mathbb{E}_0 \frac{1 - e^{(1/2 - \sqrt{1/4 - 2z})X_T}}{z^{3/2}} dz \\ &= \frac{1}{2\sqrt{\pi}} \mathbb{E}_0 \int_0^\infty \frac{1 - e^{(1/2 - \sqrt{1/4 - 2z})X_T}}{z^{3/2}} dz. \end{aligned}$$

and take real parts. The first application of Fubini is justified by $|1 - e^{-z\langle X \rangle_T}| < 1 - e^{-zm}$. The second application of Fubini is justified by $\mathbb{E}_0|1 - e^{(1/2 - \sqrt{1/4 - 2z})X_T}| = O(1)$ as $z \to \infty$; and on the other hand for z sufficiently small,

$$\begin{aligned} (\mathbb{E}_0|1 - e^{(1/2 - \sqrt{1/4 - 2z})X_T}|)^2 &\le \mathbb{E}_0(|1 - e^{(1/2 - \sqrt{1/4 - 2z})X_T}|^2) \\ &= \mathbb{E}_0(1 - 2e^{(1/2 - \sqrt{1/4 - 2z})X_T} + e^{(1/2 - \sqrt{1/4 - 2z})2X_T}) \\ &= 1 - 2\mathbb{E}_0 e^{-z\langle X \rangle_T} + \mathbb{E}_0 e^{(\frac{1 - 8z - \sqrt{1 - 8z}}{2})\langle X \rangle_T} \\ &= 1 - 2(1 - zf'(0) + O(z^2)) + 1 - 2zf'(0) + O(z^2) \\ &= O(z^2) \qquad \text{as } z \to 0 \end{aligned}$$

using the analyticity of the moment generating function $f(\xi) := e^{\xi\langle X \rangle_T}$, which follows from (B). □

Proof of Proposition 6.8. For arbitrary $\mathcal{F}_t$-measurable $q \ge 0$ we have

$$\begin{aligned} \mathbb{E}_t \sqrt{\langle X \rangle_T - \langle X \rangle_t + q} &= \frac{1}{2\sqrt{\pi}} \mathbb{E}_t \int_0^\infty \frac{1 - e^{-z(\langle X \rangle_T - \langle X \rangle_t + q)}}{z^{3/2}} dz & \text{(A.3)} \\ &= \frac{1}{2\sqrt{\pi}} \int_0^\infty (\theta_+ + \theta_-) \frac{1 - \mathbb{E}_t e^{-z(\langle X \rangle_T - \langle X \rangle_t + q)}}{z^{3/2}} dz & \text{(A.4)} \\ &= \frac{1}{2\sqrt{\pi}} \int_0^\infty \sum_\pm \theta_\pm \frac{1 - e^{-zq}\mathbb{E}_t e^{p_\pm(X_T - X_t)}}{z^{3/2}} dz & \text{(A.5)} \\ &= \frac{1}{2\sqrt{\pi}} \mathbb{E}_t \int_0^\infty \sum_\pm \theta_\pm \frac{1 - e^{-zq} e^{p_\pm(X_T - X_t)}}{z^{3/2}} dz & \text{(A.6)} \end{aligned}$$

Taking $q := \langle X\rangle_t$ yields the conclusion (6.7). The application of Fubini in (A.4) is justified by $|1 - e^{-z(\langle X\rangle_T + q)}| < 1 - e^{-z(m+q)}$. The application of Fubini in (A.6) is justified by

$$A_\pm := (\mathbb{E}_t|1 - e^{-qz+(1/2\pm\sqrt{1/4-2z})(X_T - X_t)}|)^2 \le \mathbb{E}_t(|(1 - e^{-qz+(1/2\pm\sqrt{1/4-2z})(X_T - X_t)}|^2) \tag{A.7}$$

which is $O(1)$ as $z \to \infty$, hence

$$\mathbb{E}_t \frac{|\theta_\pm(1 - e^{-qz+(1/2\pm\sqrt{1/4-2z})(X_T - X_t)})|}{z^{3/2}} = O(z^{-3/2}) \qquad z \to \infty.$$

On the other hand, for z sufficiently small, the term in the absolute values in (A.7) is real, so

$$\begin{aligned} A_\pm &\le \mathbb{E}_t\big[1 - 2e^{-qz+(1/2\pm\sqrt{1/4-2z})(X_T - X_t)} + e^{-2qz+2(1/2\pm\sqrt{1/4-2z})(X_T - X_t)}\big] \\ &= 1 - 2e^{-qz}\mathbb{E}_t e^{-z(\langle X\rangle_T - \langle X\rangle_t)} + e^{-2qz}\mathbb{E}_t e^{(\frac{1-8z\pm\sqrt{1-8z}}{2})(\langle X\rangle_T - \langle X\rangle_t)}. \end{aligned}$$

Hence as $z \to 0$, we have $A_+ = O(1)$ and

$$A_- = 1 - 2(1 - zf'(0) - qz + O(z^2)) + 1 - 2zf'(0) - 2qz + O(z^2) = O(z^2) \tag{A.8}$$

using analyticity of the moment generating function $f(\xi) := e^{\xi\langle X\rangle_T}$, which follows from (B). Combining this with $\theta_- = O(1)$ and $\theta_+ = O(z)$ as $z \to 0$, we have

$$\mathbb{E}_t \frac{|\theta_\pm(1 - e^{-qz+(1/2\pm\sqrt{1/4-2z})(X_T - X_t)})|}{z^{3/2}} = \frac{|\theta_\pm| A_\pm^{1/2}}{z^{3/2}} = O(z^{-1/2}) \qquad z \to 0,$$

which allows the interchange in (A.6).

To establish ρ-neutrality, let ϕ_v be the lognormal density with parameters $(-v/2, v)$. Then

$$\begin{aligned} \frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t) &= \frac{1}{2\sqrt{\pi}} \frac{\partial}{\partial \mathsf{s}}\bigg|_{\mathsf{s}=S_t} \int_0^\infty \int_0^\infty \left[\theta_+ \frac{1 - e^{-z\langle X\rangle_t}(\mathsf{s}y/S_t)^{p_+}}{z^{3/2}} + \theta_- \frac{1 - e^{-z\langle X\rangle_t}(\mathsf{s}y/S_t)^{p_-}}{z^{3/2}} dz\right] \phi_v(y) dy \\ &= \frac{1}{2\sqrt{\pi}} \int_0^\infty \int_0^\infty \frac{-e^{-z\langle X\rangle_t}(\theta_+ p_+ y^{p_+} + \theta_- p_- y^{p_-})}{S_t z^{3/2}} \phi_v(y) dz dy = 0 \end{aligned}$$

using the equality of integrals of $y^{p_+}\phi_v(y)$ and $y^{p_-}\phi_v(y)$, and the identity $\theta_+ p_+ + \theta_- p_- = 0$. ☐

Proof of Corollary 6.11. By a Mathematica computation,

$$\frac{1}{2\sqrt{\pi}} \int_0^\infty \theta_+ \frac{1 - e^{p_+ X_T}}{z^{3/2}} + \theta_- \frac{1 - e^{p_- X_T}}{z^{3/2}} dz = \sqrt{\frac{\pi}{2}} e^{X_T/2} |X_T I_0(X_T/2) - X_T I_1(X_T/2)|.$$

The result now follows from Proposition 6.8. ☐

Proof of Corollary 6.12. From (6.11), compute $\psi''(K)$, and apply Remark 3.1. ☐

Proof of Corollary 6.13. From (6.8), compute $\partial^2 G_{\text{svs}}/\partial \mathsf{S}^2(K, S_t, \langle X\rangle_t)$, and apply Remark 3.1. ☐

Proof of Proposition 6.15. For background in measure-valued trading strategies, see [5]. The trading strategy at each time t has value

$$
\begin{aligned}
V_t &= \int P_t(K)\varphi_t(dK) + G_{\mathrm{svs}}(\kappa, S_t, \langle X\rangle_t) \\
&= \int P_t(K)(-1)^{\mathbb{I}_{K<\kappa}} \frac{\partial G_{\mathrm{svs}}}{\partial S}(K; S_t, \langle X\rangle_t)dK + G_{\mathrm{svs}}(\kappa, S_t, \langle X\rangle_t) \\
&= \mathbb{E}_t G_{\mathrm{svs}}(S_T, S_t, \langle X\rangle_t) = \mathbb{E}_t\sqrt{\langle X\rangle_T}
\end{aligned}
$$

by Proposition 6.8. In particular it has at time $t = T$ the desired terminal value.

To prove that it self-finances, we have

$$
\begin{aligned}
dV_t &= \varphi_t dP_t + \left[\int_0^\infty P_t(K)(-1)^{\mathbb{I}_{K<\kappa}} \frac{\partial^2 G_{\mathrm{svs}}}{\partial S \partial u}(K, S_t, \langle X\rangle_t)dK\right] dS_t + d\tilde{A}_t + dG_{\mathrm{svs}}(\kappa, S_t, \langle X\rangle_t) \\
&= \varphi_t dP_t + \left[\int_0^\infty P_t(K)(-1)^{\mathbb{I}_{K<\kappa}} \frac{\partial^2 G_{\mathrm{svs}}}{\partial S \partial u}(K, S_t, \langle X\rangle_t)dK + \frac{\partial G_{\mathrm{svs}}}{\partial u}(\kappa, S_t, \langle X\rangle_t)\right] dS_t + dA_t \\
&= \varphi_t dP_t + \left[\mathbb{E}_t \frac{\partial G_{\mathrm{svs}}}{\partial u}(S_T, S_t, \langle X\rangle_t)\right] dS_t + dA_t = \varphi_t dP_t + dA_t
\end{aligned}
$$

where $\tilde{A}$ and A denote time-continuous finite-variation processes. The last step follows from $\mathbb{E}_t(S_T/S_t)^{p+} = \mathbb{E}_t(S_T/S_t)^{p-}$ and $\theta_+ p_+ + \theta_- p_- = 0$.

Moreover, A is a local martingale because $\varphi_t P_t$ and the integrals with respect to P and S are local martingales. Therefore dA vanishes. Because $dB = 0$, we have

$$
dL_t = \varphi_t dP_t + G_{\mathrm{svs}}(\kappa, S_t, \langle X\rangle_t)dB_t,
$$

which is the self-financing condition. The ρ-neutrality is proved in Proposition 6.8. □

Proof of Proposition 7.1. Using the identity [29]

$$
q^r = \frac{r}{\Gamma(1-r)} \int_0^\infty \frac{1 - e^{-zq}}{z^{r+1}} dz \qquad\qquad 0 < r < 1,\ q \geq 0
$$

follow the proof of Proposition 6.8. □

Proof of Proposition 7.2. Using the identity [29]

$$
q^{-r} = \frac{1}{r\Gamma(r)} \int_0^\infty e^{-z^{1/r}q} dz, \qquad\qquad r > 0,\ q > 0
$$

we have

$$
\begin{aligned}
\mathbb{E}_t(\langle X\rangle_T + \varepsilon)^{-r} &= \frac{1}{r\Gamma(r)} \mathbb{E}_t \int_0^\infty e^{-z^{1/r}(\langle X\rangle_T + \varepsilon)} dz \\
&= \frac{1}{r\Gamma(r)} \int_0^\infty e^{-z^{1/r}(\langle X\rangle_t + \varepsilon)} \mathbb{E}_t e^{-z^{1/r}(\langle X\rangle_T - \langle X\rangle_t)} dz \\
&= \frac{1}{r\Gamma(r)} \mathbb{E}_t \int_0^\infty (\theta_+ e^{p_+(X_T - X_t)} + \theta_- e^{p_-(X_T - X_t)}) e^{-z^{1/r}(\langle X\rangle_t + \varepsilon)} dz
\end{aligned}
$$

where the two uses of Fubini are justified by (B) and $|e^{(1/2\pm\sqrt{1/4-2z^{1/r}})(X_T-X_t)}| \leq 1$ respectively.

To establish ρ-neutrality, let ϕ_v be the lognormal density with parameters $(-v/2, v)$. Then

$$\frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t) = \frac{1}{r\Gamma(r)}\frac{\partial}{\partial \mathsf{s}}\Big|_{\mathsf{s}=S_t}\int_0^\infty \int_0^\infty \Big[\theta_+(\mathsf{s}y/S_t)^{p_+} + \theta_-(\mathsf{s}y/S_t)^{p_-}dz\Big]e^{-z^{1/r}\langle X\rangle_t}\phi_v(y)dy$$
$$= \frac{1}{r\Gamma(r)S_t}\int_0^\infty \int_0^\infty e^{-z^{1/r}\langle X\rangle_t}(\theta_+p_+y^{p_+} + \theta_-p_-y^{p_-})\phi_v(y)dzdy = 0$$

using the equality of integrals of $y^{p_+}\phi_v(y)$ and $y^{p_-}\phi_v(y)$, and the identity $\theta_+p_+ + \theta_-p_- = 0$. □

Proof of Proposition 7.3. Take the nth derivative of (5.9) with respect to λ, and evaluate at $\lambda = 0$:

$$\mathbb{E}_t\partial_\lambda^n e^{\lambda\langle X\rangle_T}\big|_{\lambda=0} = \mathbb{E}_t\partial_\lambda^n G_{\exp}(S_T, S_t, \langle X\rangle_t; \lambda)\big|_{\lambda=0} \tag{A.9}$$

Differentiation through the expectations is justified by the boundedness of $\langle X\rangle_T$ and the analyticity of the moment generating function of X_T.

To establish ρ-neutrality, let ϕ_v be the lognormal density with parameters $(-v/2, v)$. Then

$$\frac{\partial F^{BS}}{\partial \mathsf{s}}(S_t) = \frac{\partial}{\partial \mathsf{s}}\Big|_{\mathsf{s}=S_t} G_{\mathrm{pow}(n)}(\mathsf{s}y/S_t, S_t, \langle X\rangle_t)\phi_v(y)dy$$
$$= \frac{\partial^n}{\partial\lambda^n}\Big|_{\lambda=0}\frac{\partial}{\partial \mathsf{s}}\Big|_{\mathsf{s}=S_t} G_{\exp}(\mathsf{s}y/S_t, S_t, \langle X\rangle_t, \lambda)\phi_v(y)dy = 0$$

by the ρ-neutrality of $G_{\exp}$. □

Proof of Proposition 7.5. Inverting the Laplace transform,

$$h(q) = \frac{1}{2\pi i}\int_{\alpha-\infty i}^{\alpha+\infty i} H(z)e^{zq}dz$$

Therefore

$$\mathbb{E}_t h(\langle X\rangle_T) = \frac{1}{2\pi i}\mathbb{E}_t\int_{\alpha-\infty i}^{\alpha+\infty i} H(z)e^{z\langle X\rangle_T}dz = \frac{1}{2\pi i}\int_{\alpha-\infty i}^{\alpha+\infty i} H(z)\mathbb{E}_t e^{z\langle X\rangle_T}dz$$
$$= \frac{1}{2\pi i}\int_{\alpha-\infty i}^{\alpha+\infty i} H(z)e^{z\langle X\rangle_t}\mathbb{E}_t[\theta_+e^{p_+(X_T-X_t)} + \theta_-e^{p_-(X_T-X_t)}]dz$$
$$= \mathbb{E}_t G_h(S_T, S_t, \langle X\rangle_t)$$

where the two applications of Fubini (and, in particular, the convergence of the integral in (7.5)) are justified respectively by assumption (B) and by

$$\mathbb{E}_t|e^{p_\pm(X_T-X_t)}| = \mathbb{E}_t e^{\mathrm{Re}(1/2\pm\sqrt{1/4+2(\alpha+\beta i)})(X_T-X_t)} = \mathbb{E}_t e^{(1/2\pm\sqrt{|\beta|}+O(|\beta|^{-1/2}))(X_T-X_t)}$$
$$= \mathbb{E}_t e^{(|\beta|/2+O(1))(\langle X\rangle_T-\langle X\rangle_t)} = O(e^{|\beta|m/2})$$

and $|H(z)e^{z\langle X\rangle_t}\theta_\pm(z)| = O(e^{-|\beta|\mu})$ as $|\beta| \to \infty$.

Proof of ρ-neutrality is by calculation similar to the proof of Proposition 6.8. □

Proof of Proposition 7.7. By integrability of H, apply Bromwich inversion to obtain $h(V_T)$. By finiteness of $\mathbb{E}_t e^{\alpha V_T}$, apply Fubini to obtain $\mathbb{E}_t h(V_T)$. $\square$

Proof of Corollary 7.8. For $\alpha < 0$, both types of put payoffs h imply integrability of $e^{-\alpha q}h(q)$. Computation of (7.4) implies (7.7) and (7.8). Moreover, $\mathbb{E}_t e^{\alpha V_T} \leq 1$ and each H is integrable along $\mathrm{Re}(z) = \alpha$, so Proposition 7.7 applies. $\square$

Proof of Corollary 7.9. Assumption (B) implies that Proposition 7.7 applies for arbitrary $\alpha \in \mathbb{R}$. Substitute (5.11) into the convergent integral (7.6) to conclude. $\square$

Proof of Corollary 7.10. In the case of a put payoff h and $\alpha < 0$, we have $e^{-\alpha q}h(q)$ integrable, and H integrable along $\mathrm{Re}(z) = \alpha$. In the case of a call payoff h and $\alpha > 0$, we have $e^{-\alpha q}h(q)$ integrable, and H integrable along $\mathrm{Re}(z) = \alpha$. Hence Corollary 7.9 applies. $\square$

Proof of Proposition 7.12. The nth Bernstein approximation for h^* is defined by

$$B_n(x) := b_{n,n}x^n + b_{n,n-1}x^{n-1} + \cdots + b_{n,0}$$

and satisfies $h^*(x) = \lim_{n\to\infty} B_n(x)$ uniformly in $x \in [0,1]$. Therefore

$$h(q) = \lim_{n\to\infty} B_n(e^{-cq}) \tag{A.10}$$

uniformly in $q \in [0,\infty)$. Hence

$$\mathbb{E}_t h(\langle X\rangle_T) = \lim_{n\to\infty} \mathbb{E}_t B_n(e^{-c\langle X\rangle_T}) = \lim_{n\to\infty} \mathbb{E}_t \sum_{k=0}^{n} b_{n,k} e^{-ck\langle X\rangle_t}[\theta_+ e^{p_+(X_T - X_t)} + \theta_- e^{p_-(X_T - X_t)}]$$

as claimed. $\square$

Proof of Proposition 7.13. The span of the polynomials $\{1, x, x^2, \ldots\}$ is dense in $C[0,1]$ with respect to the uniform norm. By the transformation $q = -(1/c)\log x$, the span of exponential functions $\{1, e^{-cq}, e^{-2cq}, \ldots\}$ is dense in $C[0,\infty]$ with respect to the uniform norm, hence dense in $C[0,\infty]$ with respect to the $L^2(\mu)$ norm. Then $h = \lim_{n\to\infty} A_n$ in the $L^2(\mu)$ sense, hence

$$\begin{aligned}\Big(\mathbb{E}\big[h(\langle X\rangle_T) - A_n(\langle X\rangle_T)\big]\Big)^2 &= \Big(\int \frac{dP}{d\mu}\big[h(q) - A_n(q)\big]d\mu(q)\Big)^2 \\ &\leq \int \Big(\frac{dP}{d\mu}\Big)^2 d\mu \int \big[h(q) - A_n(q)\big]^2 d\mu(q) \longrightarrow 0\end{aligned}$$

as $n \to \infty$. Thus

$$\mathbb{E}_t h(\langle X\rangle_T) = \lim_{n\to\infty} \mathbb{E}_t A_n(\langle X\rangle_T) = \lim_{n\to\infty} \mathbb{E}_t \sum_{k=0}^{n} a_{n,k} e^{-ck\langle X\rangle_t}[\theta_+ e^{p_+(X_T - X_t)} + \theta_- e^{p_-(X_T - X_t)}] \tag{A.11}$$

as desired. $\square$

Proof of Proposition 8.1. For each positive integer m, define the process $\sigma_t^m := \sigma_t \mathbb{I}(\langle X \rangle_t \le m)$.

Define the process S_t^m by $dS_t^m = \sigma_t^m S_t^m dW_t$. Let $X_t^m := \log(S_t^m)$.

Then S^m satisfies (B), so

$$\mathbb{E}h(\langle X^m \rangle_T) = \mathbb{E}G(S_T^m). \tag{A.12}$$

Now let $m \to \infty$. The left-hand side approaches $\mathbb{E}h(\langle X \rangle_T)$ because $\langle X^m \rangle_T \to \langle X \rangle_T$ almost surely, and either monotone convergence or dominated convergence applies.

It remains to show that the right-hand side of (A.12) approaches $\mathbb{E}G(S_T)$. There exist constants α, β and convex nonnegative functions G_+, G_- such that $G_\pm(S_0) = 0$ and $\mathbb{E}G_\pm(S_T) < \infty$ and

$$G(S) = G_+(S) - G_-(S) + \alpha S + \beta \qquad \text{for all } S \ge 0.$$

We need only to show that $\mathbb{E}G_+(S_T^m) \to \mathbb{E}G_+(S_T)$; convergence proofs for the other terms are then trivial. It suffices to show that the family

$$\{G_+(S_T^m) : \ m \ge 1\}$$

is uniformly integrable. Since $\mathbb{E}G_+(S_T) < \infty$, it is enough to show that for all m and all $A > 0$,

$$\mathbb{E}G_+(S_T^m)\mathbb{I}(G_+(S_T^m) > A) \le \mathbb{E}G_+(S_T)\mathbb{I}(G_+(S_T) > A).$$

By the convexity of G_+, there exist $a, b \in [0, \infty]$ such that

$$\mathbb{I}(G_+(S) > A) = \mathbb{I}(S < S_0 - a) + \mathbb{I}(S > S_0 + b) \qquad \text{for all } S > 0$$

Moreover, the function

$$U(S) := G_+(S)\mathbb{I}(G_+(S) > A) - \frac{A}{b}(S - S_0)\mathbb{I}(S > S_0 + b) - \frac{A}{a}(S_0 - S)\mathbb{I}(S < S_0 - a)$$

is convex. We have

$$\begin{aligned}
&\mathbb{E}G_+(S_T^m)\mathbb{I}(G_+(S_T^m) > A) \\
&= \mathbb{E}\left[\frac{A}{a}(S_0 - S_T^m)\mathbb{I}(S_T^m < S_0 - a) + \frac{A}{b}(S_T^m - S_0)\mathbb{I}(S_T^m > S_0 + b) + U(S_T^m)\right] \\
&= \mathbb{E}\left[\frac{A}{a}\mathbb{E}[(S_0 - S_T^m)\mathbb{I}(S_T^m < S_0 - a)|\langle X \rangle_T] + \frac{A}{b}\mathbb{E}[(S_T^m - S_0)\mathbb{I}(S_T^m > S_0 + b)|\langle X \rangle_T] + \mathbb{E}[U(S_T^m)|\langle X \rangle_T]\right] \\
&\le \mathbb{E}\left[\frac{A}{a}\mathbb{E}[(S_0 - S_T)\mathbb{I}(S_T < S_0 - a)|\langle X \rangle_T] + \frac{A}{b}\mathbb{E}[(S_T - S_0)\mathbb{I}(S_T > S_0 + b)|\langle X \rangle_T] + \mathbb{E}[U(S_T)|\langle X \rangle_T]\right] \\
&= \mathbb{E}G_+(S_T)\mathbb{I}(G_+(S_T) > A)
\end{aligned}$$

where the inequality holds because if Z is a mean-S_0 lognormal with $\operatorname{Var}\log(Z) = \sigma^2$, then each of $\mathbb{E}[(S_0 - Z)\mathbb{I}(Z < S_0 - a)]$, $\mathbb{E}[(Z - S_0)\mathbb{I}(Z > S_0 + b)]$, and $\mathbb{E}U(Z)$ is increasing in σ. For the first two expectations, this comes from direct calculation; for $\mathbb{E}U(Z)$, it follows from Jensen's inequality. $\square$

References

[1] Vexed by variance. *Risk*, August 2006.

[2] Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical performance of alternative option pricing models. *Journal of Finance*, 52:2003–2049, 1997.

[3] David Bates. The skewness premium: Option pricing under asymmetric processes. *Advances in Futures and Options Research*, 9:51–82, 1997.

[4] David Bates. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94:181–238, 2000.

[5] Tomas Björk, Giovanni Di Masi, Yuri Kabanov, and Wolfgang Runggaldier. Towards a general theory of bond markets. *Finance And Stochastics*, 1:141–174, 1997.

[6] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659, 1973.

[7] Oleg Bondarenko. Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics*, 116:85–112, 2003.

[8] David Breeden and Robert Litzenberger. Prices of state contingent claims implicit in options prices. *Journal of Business*, 51:621–651, 1978.

[9] Menachem Brenner, Ernest Ou, and Jin Zhang. Hedging volatility risk. *Journal of Banking & Finance*, 30:811–821, 2006.

[10] Menachem Brenner and Marti Subrahmanyam. A simple formula to compute the implied standard deviation. *Financial Analysts Journal*, 44(5):80–83, 1988.

[11] Mark Britten-Jones and Anthony Neuberger. Option prices, implied price processes, and stochastic volatility. *Journal of Finance*, 55(2):839–866, 2000.

[12] Peter Carr, Hélyette Geman, Dilip Madan, and Marc Yor. Pricing options on realized variance. *Finance and Stochastics*, 9(4):453–475, 2005.

[13] Peter Carr and Roger Lee. Realized volatility and variance: Options via swaps. *Risk*, 20(5):76–83, 2007.

[14] Peter Carr and Roger Lee. Put-call symmetry: Extensions and applications. *Mathematical Finance*, 2009. Forthcoming.

[15] Peter Carr and Dilip Madan. Towards a theory of volatility trading. In R. Jarrow, editor, *Volatility*, pages 417–427. Risk Publications, 1998.

[16] Emanuel Derman, Kresimir Demeterfi, Michael Kamal, and Joseph Zou. A guide to volatility and variance swaps. *Journal of Derivatives*, 6(4):9–32, 1999.

[17] Bruno Dupire. Model art. *Risk*, 6(9):118–124, 1993.

[18] Steven P. Feinstein. The Black-Scholes formula is nearly linear in sigma for at-the-money options: Therefore implied volatilities from at-the-money options are virtually unbiased. Federal Reserve Bank of Atlanta, 1989.

[19] Peter Friz and Jim Gatheral. Valuation of volatility derivatives as an inverse problem. *Quantitative Finance*, 5(6):531–542, 2005.

[20] Anuj Gangahar. Volatility becomes an asset class. *Financial Times*, May 23, 2006.

[21] John Hull and Alan White. The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42(2):281–300, June 1987.

[22] George J. Jiang and Yisong S. Tian. The model-free implied volatility and its information content. *Review of Financial Studies*, 18:1305–1342, 2005.

[23] Henry A. Latane and Richard J. Rendleman. Standard deviations of stock price ratios implied in option prices. *Journal of Finance*, 31(2):369–381, 1976.

[24] Andrew Matytsin. Modelling volatility and volatility derivatives. JP Morgan, 1999.

[25] Anthony Neuberger. The log contract. *Journal of Portfolio Management*, 20(2):74–80, 1994.

[26] Navroz Patel. *RiskNews*, June 12, 2003.

[27] Allen Poteshman. Forecasting future volatility from option prices. University of Illinois, 2000.

[28] Marc Romano and Nizar Touzi. Contingent claims and market completeness in a stochastic volatility model. *Mathematical Finance*, 7(4):399–410, 1997.

[29] Klaus Schürger. Laplace transforms and suprema of stochastic processes. University of Bonn, 2002.

[30] Gregory A. Willard. Calculating prices and sensitivities for path-independent derivative securities in multifactor models. *Journal of Derivatives*, pages 45–61, Fall 1997.

[31] Guo Dong Zhu. Pricing options on trading strategies. PhD thesis, NYU, 2007.