

NBER WORKING PAPER SERIES

SIMPLE VARIANCE SWAPS

Ian Martin

Working Paper 16884

<http://www.nber.org/papers/w16884>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2011

I am very grateful to Peter Carr, Darrell Duffie, Stefan Hunt, and Myron Scholes for their comments. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Ian Martin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Simple Variance Swaps
Ian Martin
NBER Working Paper No. 16884
March 2011
JEL No. G01,G12,G13

ABSTRACT

The large asset price jumps that took place during 2008 and 2009 disrupted volatility derivatives markets and caused the single-name variance swap market to dry up completely. This paper defines and analyzes a simple variance swap, a relative of the variance swap that in several respects has more desirable properties. First, simple variance swaps are robust: they can be easily priced and hedged even if prices can jump. Second, simple variance swaps supply a more accurate measure of market-implied variance than do variance swaps or the VIX index. Third, simple variance swaps provide a better way to measure and to trade correlation. The paper also explains how to interpret VIX in the presence of jumps.

Ian Martin
Graduate School of Business
Stanford University
Stanford, CA 94305
and NBER
ian.martin@gsb.stanford.edu

In recent years, a large market in volatility derivatives has developed. An emblem of this market, the VIX index, is often described in the financial press as “the fear index”; its construction is based on theoretical results on the pricing of variance swaps. These derivatives permit investors and dealers to hedge and to speculate in volatility itself. They also play an informational role by providing evidence about perceptions of future volatility.

The events of 2008 and 2009 severely disrupted these markets and revealed certain undesirable features of variance swaps. Carr and Lee (2009) write, “The cataclysm that hit almost all financial markets in 2008 had particularly pronounced effects on volatility derivatives Dealers learned the hard way that the standard theory for pricing and hedging variance swaps is not nearly as model-free as previously supposed In particular, sharp moves in the underlying highlighted exposures to cubed and higher-order daily returns. The inability to take positions in deep OTM options when hedging a variance swap later affected the efficacy of the hedging strategy. As the underlying index or stock moved away from its initial level, dealers found themselves exposed to much more vega than a complete hedging strategy would permit. This issue was particularly acute for single names, as the options are not as liquid and the most extreme moves are bigger. As a result, the market for single-name variance swaps has evaporated in 2009.” In response to this sensitivity to extreme events, market participants have imposed caps on variance swap payoffs. These caps limit the maximum possible payoff on a variance swap, at the cost of complicating the pricing and interpretation of the original contract.

In this paper, I define and analyze a financial contract that I call a *simple variance swap*. Although it has, arguably, a simpler definition than that of a standard variance swap,¹ I show that simple variance swaps are robust to the issues mentioned in the quotation above, and explain why this is the case. Simple variance swaps can be priced and hedged under far weaker assumptions than are required for pricing variance swaps (or the recently introduced gamma swaps): in particular, they can be hedged in the presence of jumps. Moreover, from an informational perspective, simple variance swaps provide a more useful measure of market-implied volatility than do variance swaps.

Lastly, I consider the problem of computing market-implied correlations. In this respect, too, simple variance swaps are shown to be an improvement on variance swaps, both for observers interested in the market’s view of implied correlation and for investors seeking to implement dispersion trades.

Section 1 introduces the definition of a simple variance swap, and presents the main pricing and hedging results. Result 1 is a straightforward application of no-arbitrage

¹I will refer to this standard product as a variance swap, and to the proposed alternative as a simple variance swap.

logic. Results 2 and 3 derive an important simplification in the natural limiting case. Section 2 compares these results to those already known for standard variance swaps. Existing results are summarized in Result 4. These results depend on the assumption that the underlying asset’s price cannot jump. Next, I propose an index, SVIX, that is analogous to VIX but based on the strike on a simple variance swap rather than on a (standard) variance swap. Results 5 and 6 show how VIX can be interpreted in the presence of jumps, and compare to the much simpler interpretation of the proposed SVIX index. Section 3 discusses the potential application of simple variance swaps to the measurement of market-implied correlation. Section 4 concludes.

Related literature. The literature on variance swaps (Carr and Madan (1998), Demeterfi et al. (1999)) is based on papers by Breeden and Litzenberger (1978) and Neuberger (1990, 1994). Carr and Corso (2001) propose a contract related to the simple variance swap proposed below, but their pricing and hedging methodology is only valid if the underlying is a futures contract. Moreover, as the authors note, options on futures are almost invariably American-style, whereas the replicating portfolio calls for European-style options. Lee (2010a, 2010b) provides a brief summary of volatility derivative pricing in the absence of jumps. Finally, the paper of Carr and Lee (2009) quoted above provides an excellent survey of the state of the art in the area.

1 Pricing and hedging simple variance swaps

Consider an underlying asset whose price at time t is S_t . This price is assumed to include reinvested dividends, if the asset pays dividends. Today is time 0. I assume that the interest rate r is constant; this is the standard assumption in the related literature. At time 0, the asset’s forward price to time t is then $F_t \equiv S_0 e^{rt}$.

We can now define a simple variance swap on this underlying asset. This is an agreement to exchange, at time T , some prearranged payment V (the “strike” of the simple variance swap) for the amount

$$\left(\frac{S_\Delta - S_0}{S_0}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_\Delta}\right)^2 + \dots + \left(\frac{S_T - S_{T-\Delta}}{F_{T-\Delta}}\right)^2. \quad (1)$$

V is chosen so that no money changes hands at initiation of the contract. The choice to put forward prices in the denominators is important: below we will see that this choice leads to a dramatic simplification of the strike V , and of the associated hedging strategy, in the limit as the period length Δ goes to zero. In an idealized frictionless market, this simplification of the hedging strategy would merely be a matter of analytical convenience; in practice, with trade costs, it acquires far more significance.

The following result shows how to price a simple variance swap (i.e. how to choose V

so that no money need change hands initially) in terms of the prices of European call and put options on the underlying asset. I write $\text{call}_t(K)$ for the time-0 price of a European call option on the underlying asset, expiring at time t , with strike K , and $\text{put}_t(K)$ for the time-0 price of the corresponding put option.

Result 1 (Pricing). *The strike on a simple variance swap, i.e. an agreement to exchange V for the amount in (1) at time T , is given by*

$$V(\Delta, T) = \sum_{i=1}^{T/\Delta} \frac{e^{ri\Delta}}{F_{(i-1)\Delta}^2} [\Pi(i\Delta) - (2 - e^{-r\Delta})\Pi((i-1)\Delta)], \quad (2)$$

where $\Pi(t)$ is given by

$$\Pi(t) = 2 \int_0^{F_t} \text{put}_t(K) dK + 2 \int_{F_t}^{\infty} \text{call}_t(K) dK + S_0^2 e^{rt}, \quad \Pi(0) = S_0^2. \quad (3)$$

Proof. The absence of arbitrage implies that there exists a sequence of strictly positive stochastic discount factors $M_\Delta, M_{2\Delta}, \dots$ such that a payoff $X_{j\Delta}$ at time $j\Delta$ has price

$$\mathbb{E}_{i\Delta} [M_{(i+1)\Delta} M_{(i+2)\Delta} \cdots M_{j\Delta} X_{j\Delta}]$$

at time $i\Delta$. The subscript on the expectation operator indicates that the expectation is conditional on time- $i\Delta$ information. I abbreviate $M_{(j\Delta)} \equiv M_\Delta M_{2\Delta} \cdots M_{j\Delta}$.

V is chosen so that the swap has zero initial value, i.e.,

$$\mathbb{E} \left[M_{(T)} \left\{ \left(\frac{S_\Delta - S_0}{S_0} \right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_\Delta} \right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{T-\Delta}} \right)^2 - V \right\} \right] = 0. \quad (4)$$

We have

$$\begin{aligned} \mathbb{E} \left[M_{(T)} (S_{i\Delta} - S_{(i-1)\Delta})^2 \right] &= e^{-r(T-i\Delta)} \mathbb{E} \left[M_{(i\Delta)} (S_{i\Delta} - S_{(i-1)\Delta})^2 \right] \\ &= e^{-r(T-i\Delta)} \left\{ \mathbb{E} [M_{(i\Delta)} S_{i\Delta}^2] - (2 - e^{-r\Delta}) \mathbb{E} [M_{((i-1)\Delta)} S_{(i-1)\Delta}^2] \right\}, \end{aligned}$$

using the law of iterated expectations, the fact that $\mathbb{E}_{(i-1)\Delta} M_{i\Delta} S_{i\Delta} = S_{(i-1)\Delta}$, and the fact that the interest rate r is constant, so that $\mathbb{E}_{(i-1)\Delta} M_{i\Delta} = e^{-r\Delta}$. If we define $\Pi(i)$ to be the time-0 price of a claim to S_i^2 , paid at time i , then we have²

$$\mathbb{E} \left[M_{(T)} (S_{i\Delta} - S_{(i-1)\Delta})^2 \right] = e^{-r(T-i\Delta)} [\Pi(i\Delta) - (2 - e^{-r\Delta}) \Pi((i-1)\Delta)].$$

²Neuberger (1990), the working paper version of Neuberger (1994), briefly mentions the relevance of this ‘‘Squared contract’’ in a similar context.

Equation (2) follows on substituting this into (4), so it only remains to confirm that (3) holds. To see that it does, note that³

$$S_t^2 = 2 \int_0^\infty \max\{0, S_t - K\} dK.$$

The right-hand side is the time- t payoff on a portfolio of European call options of all strikes, so the absence of arbitrage implies that

$$\Pi(t) = 2 \int_0^\infty \text{call}_t(K) dK. \quad (5)$$

Equation (3) follows by put-call parity, which is the relationship $\text{call}_t(K) = \text{put}_t(K) + S_0 - Ke^{-rt}$. Although (3) is less concise than (5), it has the appealing feature that it expresses $\Pi(t)$ in terms of out-of-the-money options only. \square

The most important aspect of Result 1 is that it does not require the price process of the underlying asset to be continuous. The strike on a simple variance swap is dictated by the prices of options across all strikes and the whole range of expiry times $\Delta, 2\Delta, \dots, T$. Correspondingly, the hedging portfolio requires holding portfolios of options of each of these maturities, as in the examples considered by Lee (2010b). Although this is not a serious issue if Δ is large relative to T , it raises the concern that hedging a simple variance swap may be extremely costly in practice if Δ is very small relative to T . Fortunately, this concern is misplaced: both the pricing formula (2) and the hedging portfolio simplify nicely in the limit as $\Delta \rightarrow 0$, holding T constant.

Result 2 (Pricing, continued). *In the limit as $\Delta \rightarrow 0$, we have*

$$V(0, T) \equiv \lim_{\Delta \rightarrow 0} V(\Delta, T) = \frac{2e^{-rT}}{S_0^2} \left\{ \int_0^{F_T} \text{put}_T(K) dK + \int_{F_T}^\infty \text{call}_T(K) dK \right\}. \quad (6)$$

Proof. Observe that (2) can be rewritten

$$V(\Delta, T) = \sum_{i=1}^{T/\Delta} \left\{ \frac{e^{ri\Delta}}{F_{(i-1)\Delta}^2} [P(i\Delta) - (2 - e^{-r\Delta})P((i-1)\Delta)] \right\} + \frac{T}{\Delta} (e^{r\Delta} - 1)^2, \quad (7)$$

where

$$P(t) = 2 \int_0^{F_t} \text{put}_t(K) dK + 2 \int_{F_t}^\infty \text{call}_t(K) dK, \quad P(0) = 0.$$

³Darrell Duffie suggested this approach. A previous draft derived (5) via Breeden-Litzenberger (1978) logic and integration by parts: $\Pi(t) = \int_0^\infty K^2 \text{call}_t''(K) dK = 2 \int_0^\infty \text{call}_t(K) dK$. The latter approach has the advantage of being mechanical—it does not rely on a “trick”—but the disadvantage of requiring twice-differentiability of $\text{call}_t(K)$.

For $0 < j < T/\Delta$, the coefficient on $P(j\Delta)$ in equation (7) is

$$\frac{1}{F_{(j-1)\Delta}^2} e^{r\Delta j} - \frac{1}{F_{j\Delta}^2} e^{r\Delta(j+1)} (2 - e^{-r\Delta}) = \frac{(e^{r\Delta} - 1)^2}{S_0^2 e^{r\Delta j}}.$$

We can therefore rewrite (7) as

$$V(\Delta, T) = \frac{e^{rT}}{F_{T-\Delta}^2} P(T) + \underbrace{\sum_{j=1}^{T/\Delta-1} \frac{(e^{r\Delta} - 1)^2}{S_0^2 e^{rj\Delta}} P(j\Delta)}_{O(1/\Delta) \text{ terms of size } O(\Delta^2)} + \frac{T}{\Delta} (e^{r\Delta} - 1)^2.$$

The second term on the right-hand side is a sum of $T/\Delta - 1$ terms, each of which has size on the order of Δ^2 ; all in all, the sum is $O(\Delta)$. The third term on the right-hand side is also $O(\Delta)$, so both tend to zero as $\Delta \rightarrow 0$. The first term tends to $e^{rT} P(T)/F_T^2$, as required. \square

Result 2 is the motivation behind the choice of forward prices as the normalizing weights in the definition (1). In principle, we could have put any other constants known at time 0 in the denominators of the fractions in (1). Had we done so, we would have to face the unappealing prospect of a hedging portfolio requiring positions in options of all maturities between 0 and T . Using forward prices lets us sidestep this problem.

The proof of Result 1 implicitly supplies the dynamic trading strategy that replicates the payoff on a simple variance swap. Tables 1 and 2 in the Appendix describe the strategy in detail. Each row of Table 1 indicates a sequence of dollar cashflows that is attainable by investing in the asset indicated in the leftmost column. Negative quantities indicated that money must be invested; positive quantities indicate cash inflows. Thus, for example, the first row indicates a time-0 investment of $\$e^{-rT}$ in the riskless bond maturing at time T , which generates a time- T payoff of \$1. The second and third rows indicate a short position in the underlying asset, held from 0 to Δ and subsequently rolled into a short bond position. The fourth row represents a position in a portfolio of options expiring at time Δ that has simple return $S_\Delta^2/\Pi(\Delta)$ from time 0 to time Δ . (This is the portfolio whose price is provided in equations (3) or (5). It is perhaps easiest to think in terms of (5), so the portfolio consists of call options of all strikes.) The fifth, sixth, and seventh rows indicate how the proceeds of this option portfolio are used after time Δ . Part of the proceeds are immediately invested in the bond until time T ; another part is invested from Δ to 2Δ in the underlying asset, and subsequently from 2Δ to T in the bond. The replicating portfolio requires similar positions in options expiring at times $2\Delta, 3\Delta, \dots, T - 2\Delta$. These are omitted from Table 1, but the general such position is indicated in Table 2, together with the subsequent investment in bonds and underlying that each position requires.

The self-financing nature of the replicating strategy is reflected in the fact that the total of each of the intermediate columns from time Δ to time $T - \Delta$ is zero. The last column of Table 1 adds up to the desired payoff,

$$\left(\frac{S_\Delta - S_0}{S_0}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_\Delta}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{T-\Delta}}\right)^2 - V.$$

Therefore, the first column must add up to the cost of entering the simple variance swap. Equating this cost to zero, we find the value of V provided in Result 1.

The replicating strategy simplifies dramatically in the $\Delta \rightarrow 0$ limit. The dollar investment in each of the option portfolios expiring at times $\Delta, 2\Delta, \dots, T - \Delta$ goes to zero at rate $O(\Delta^2)$. We must account, however, for the dynamically adjusted position in the underlying, indicated in rows beginning with a U. As shown in Table 2, this calls for a short position of $2e^{-r(T-\Delta)}S_{j\Delta}/[F_{j\Delta}S_0]$ units of underlying at time $j\Delta$. In *dollar* terms this is a short position of $2e^{-r(T-\Delta)}S_{j\Delta}^2/[F_{j\Delta}S_0]$ in the underlying. In the limit as $\Delta \rightarrow 0$, holding $j\Delta = t$ constant, the dynamic strategy calls for a short position of $2e^{-rT}S_t/[F_tS_0]$ units of the underlying asset at time t , i.e. $2e^{-rT}S_t^2/[F_tS_0]$ in dollar terms.

The static position in options expiring at time T , shown in the penultimate line of Table 1, does not disappear in the $\Delta \rightarrow 0$ limit. We can think of the option portfolio as a collection of calls of all strikes, as in (5). It is perhaps more natural, though, to think of the position as a collection of calls with strikes above F_T and puts with strikes below F_T , together with a long position in $2/F_T$ units of the underlying asset and a bond position. Combining this static long position in the underlying with the previously discussed dynamic position, the overall position at time t is long $2/F_T - 2e^{-rT}S_t/[F_tS_0] = 2(1 - S_t/F_t)/F_T$ units of the asset and long out-of-the-money-forward calls and puts, all financed by borrowing. Initially, therefore, the direct position in the underlying asset is $2(1 - S_0/F_0)/F_T = 0$ at time 0; subsequently, if the underlying's price at time t exceeds $F_t = S_0e^{rt}$, the replicating portfolio is short the underlying in order to offset the effects of increasing delta as calls go in-the-money and puts go increasingly out-of-the-money. In the other direction, if the underlying asset's price declines, then the delta-hedge requires buying the underlying to offset the negative delta resulting from puts going in-the-money and calls going out-of-the-money.

Result 3 (Hedging). *In the $\Delta \rightarrow 0$ limit, the payoff on a simple variance swap can be replicated by holding a portfolio, financed by borrowing, of*

- (i) a static position in $2/F_T^2$ puts expiring at time T with strike K , for each $K \leq F_T$;
 - (ii) a static position in $2/F_T^2$ calls expiring at time T with strike K , for each $K \geq F_T$;
- and

(iii) a dynamic position that is long $2(1 - S_t/F_t)/F_T$ units of the underlying asset at time t .

2 Variance swaps, gamma swaps, and the VIX index

In contrast, a (standard) variance swap pays

$$\left(\log \frac{S_\Delta}{S_0}\right)^2 + \left(\log \frac{S_{2\Delta}}{S_\Delta}\right)^2 + \cdots + \left(\log \frac{S_T}{S_{T-\Delta}}\right)^2 - \tilde{V}$$

at time T . (This definition is very natural in a world in which asset prices follow diffusions.) To be fairly priced, we must have

$$\tilde{V} = \mathbb{E}^* \left[\left(\log \frac{S_\Delta}{S_0}\right)^2 + \left(\log \frac{S_{2\Delta}}{S_\Delta}\right)^2 + \cdots + \left(\log \frac{S_T}{S_{T-\Delta}}\right)^2 \right], \quad (8)$$

where the asterisk on the expectation indicates that it is taken with respect to the risk-neutral measure. To price the variance swap, i.e. to compute the expectation on the right-hand side of (8), it is generally assumed that the asset's price is an Itô process. When this is the case, we have the following result in the $\Delta \rightarrow 0$ limit; it is due to Carr and Madan (1998) and Demeterfi, Derman, Kamal, and Zou (1999), building on an idea of Neuberger (1994). From now on, \tilde{V} will always refer to the variance swap strike in the limiting case $\Delta \rightarrow 0$.

Result 4 (Neuberger (1994); Carr and Madan (1998); Demeterfi et al. (1999)). *If the underlying asset's price follows an Itô process $dS_t = rS_t dt + \sigma_t S_t dZ_t$ under the risk-neutral measure, then the strike on a variance swap is*

$$\tilde{V} = 2e^{rT} \left\{ \int_0^{F_T} \frac{1}{K^2} \text{put}_T(K) dK + \int_{F_T}^\infty \frac{1}{K^2} \text{call}_T(K) dK \right\}, \quad (9)$$

which has the interpretation

$$\tilde{V} = \mathbb{E}^* \left[\int_0^T \sigma_t^2 dt \right]. \quad (10)$$

The variance swap can be hedged by holding a portfolio, financed by borrowing, of

- (i) a static position in $2/K^2$ puts expiring at time T with strike K , for every $K \leq F_T$;
- (ii) a static position in $2/K^2$ calls expiring at time T with strike K , for every $K \geq F_T$;
- and
- (iii) a dynamic position that is long $2(F_t/S_t - 1)/F_T$ units of the underlying asset at time t .

Sketch of proof. In the $\Delta \rightarrow 0$ limit, the expectation (8) converges to⁴

$$\tilde{V} = \mathbb{E}^* \left[\int_0^T (d \log S_t)^2 \right].$$

Given the Itô process assumption, $d \log S_t = (r - \frac{1}{2} \sigma_t^2) dt + \sigma_t dZ_t$ under the risk-neutral measure, by Itô's lemma, so $(d \log S_t)^2 = \sigma_t^2 dt$, and

$$\begin{aligned} \tilde{V} &= \mathbb{E}^* \left[\int_0^T \sigma_t^2 dt \right] \\ &= 2 \mathbb{E}^* \left[\int_0^T \frac{1}{S_t} dS_t - \int_0^T d \log S_t \right] \\ &= 2rT - 2 \mathbb{E}^* \log \frac{S_T}{S_0}. \end{aligned} \quad (11)$$

This shows that the strike on a variance swap is determined by pricing a notional contract that pays, at time T , the logarithm of the underlying asset's simple return $R_T = S_T/S_0$. Using Breeden-Litzenberger (1978) logic, the price of this contract, P_{\log} , can be computed explicitly in terms of the prices of European call and put options on the underlying asset:

$$P_{\log} \equiv e^{-rT} \mathbb{E}^* \log R_T = rT e^{-rT} - \int_0^{F_T} \frac{1}{K^2} \text{put}_T(K) dK - \int_{F_T}^{\infty} \frac{1}{K^2} \text{call}_T(K) dK. \quad (12)$$

For completeness, a derivation is provided in the Appendix. Substituting (12) back into (11), we have the result. \square

Another recent innovation, the gamma swap, is closely related to a variance swap. At time T , a gamma swap pays

$$\frac{S_{\Delta}}{S_0} \left(\log \frac{S_{\Delta}}{S_0} \right)^2 + \frac{S_{2\Delta}}{S_0} \left(\log \frac{S_{2\Delta}}{S_{\Delta}} \right)^2 + \dots + \frac{S_T}{S_0} \left(\log \frac{S_T}{S_{T-\Delta}} \right)^2 - V_{\gamma}.$$

Under the Itô process assumption, gamma swaps can be replicated in a similar way to variance swaps; see Lee (2010a). The replicating portfolio consists of a static position in out-of-the-money-forward calls and puts, in amounts inversely proportional to the strike (together with a dynamic delta-hedge, and money market positions for financing). In contrast, the standard variance swap holds options in amounts inversely proportional to the *square* of the strike.

Thus we have a sequence: variance swaps are hedged with a static portfolio of options with strikes K , held in amounts proportional to $1/K^2$; gamma swaps are hedged with a static portfolio of options with strikes K , held in amounts proportional to $1/K$; and

⁴To be more formal about it, \tilde{V} converges, under technical conditions established by Jarrow et al. (2010), to $\mathbb{E}^* [\langle \log S \rangle_T]$, where $\langle \log S \rangle_t$ is the quadratic variation process of $\log S_t$.

simple variance swaps are hedged with a static portfolio of options with strikes K , held in amounts proportional to 1. Only for simple variance swaps, though, is replication possible in the presence of jumps.

Figure 1 makes this point graphically. It shows the required payoff and associated hedge portfolio payoff over a particular sample path for the underlying asset, for a 3-month variance swap and for a 3-month simple variance swap.

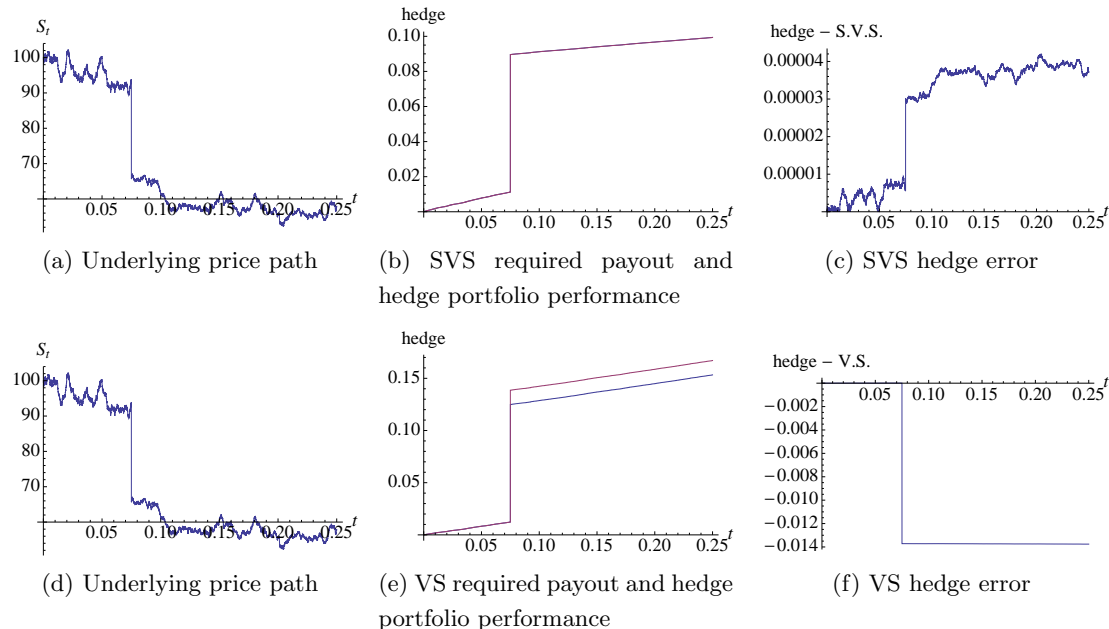


Figure 1: A sample path of the required payout on a simple variance swap (SVS), and on a variance swap (VS), given the underlying price path shown in panels (a) and (d).

Panels 1a, 1b, and 1f correspond to the simple variance swap, and panels 1d, 1e, and 1f to the variance swap. The underlying price follows the same path in each case. The middle panels compare the required payout on the simple variance swap and variance swap to the payout of the hedge portfolio in each case.⁵ In Panel 1b, only one line is visible: the simple variance swap is essentially perfectly hedged. In contrast, Panel 1e shows that the hedge portfolio (lower line) substantially underperforms the required payout on the variance swap (upper line) due to the downward jump in the price of the underlying. Panels 1c and 1f plot the difference between required payout and hedge portfolio performance in each case. In the case of the variance swap, the hedge error is enormous—on the order of 10% of the required payoff—while in the case of the simple

⁵For times t prior to expiry ($T = 0.25$), I compute the required payout and hedge portfolio performance on the assumption that volatility goes to zero after time t , i.e. that the underlying asset price grows deterministically at the riskless rate between time t and time T .

variance swap, the corresponding error is three orders of magnitude smaller.⁶

Result 4 motivates the definition of the VIX index, which is calculated based on option prices using an annualized and discretized version of (9), and is generally interpreted as a measure of risk-neutral variance in the sense of (10). Working with the idealized version of VIX (i.e. not discretizing), we have

$$\text{VIX}^2 \equiv \frac{2e^{rT}}{T} \left\{ \int_0^{F_T} \frac{1}{K^2} \text{put}_T(K) dK + \int_{F_T}^{\infty} \frac{1}{K^2} \text{call}_T(K) dK \right\}.$$

This is a *definition*, not a statement about pricing. Analogously, we can define an index, SVIX, based on the annualized strike of a simple variance swap. Based on (6), let SVIX be defined by

$$\text{SVIX}^2 \equiv \frac{e^{2rT}}{T} V(0, T) = \frac{2e^{rT}}{T} \left\{ \int_0^{F_T} \frac{1}{S_0^2} \text{put}_T(K) dK + \int_{F_T}^{\infty} \frac{1}{S_0^2} \text{call}_T(K) dK \right\}.$$

This definition annualizes the strike on a simple variance swap, $V(0, T)$, and scales it by e^{2rT} . The scaling has two benefits: it makes SVIX more directly comparable to VIX, and it ensures that SVIX has a clean interpretation, as will be shown in Result 5.

Under the Itô process assumption, VIX^2 corresponds to the (annualized) strike on a variance swap, and has the interpretation (10). But this result leans heavily on the Itô process assumption: if there are jumps, the correctly priced strike \tilde{V} will not be given by (9); the replicating portfolio implied by the above analysis will not replicate the variance swap payoff, as Figure 1 shows; and neither \tilde{V} nor VIX^2 has the interpretation (10).⁷ The next result shows what VIX *does* measure, and contrasts this with the much simpler interpretation of SVIX.

Result 5 (The interpretation of VIX and SVIX). *Whether or not there are jumps, VIX measures the risk-neutral entropy of the simple return:*

$$\text{VIX}^2 = \frac{2}{T} L(R_T), \tag{13}$$

where the entropy $L(X) \equiv \log \mathbb{E}^* X - \mathbb{E}^* \log X$ for positive random variables X . If the simple return R_T is lognormal, then

$$\text{VIX}^2 = \frac{1}{T} \text{var}^* \log R_T \approx \frac{1}{T} \text{var}^* R_T, \tag{14}$$

⁶This example was generated using a discretization both in time— Δ strictly greater than zero—and in the gap between strikes of options in the hedging portfolio. In the absence of this discretization, the hedge error on a simple variance swap would be *exactly* zero, as shown in Result 3.

⁷Up to a third-order approximation, Carr and Lee (2009) show that the fair variance swap strike is higher than that given in Result 4 if there are jumps and risk-neutral returns are negatively skewed.

where the approximation is accurate over short time horizons. But, in general, with jumps and/or time-varying volatility, VIX depends on all of the (annualized, risk-neutral) cumulants of log returns,

$$VIX^2 = 2 \sum_{n=2}^{\infty} \frac{\kappa_n^*}{n!} = \kappa_2^* + \frac{\kappa_3^*}{3} + \frac{\kappa_4^*}{12} + \frac{\kappa_5^*}{60} + \dots, \quad (15)$$

where $\kappa_n^* \equiv \frac{1}{T} \tilde{\kappa}_n^*$, and $\tilde{\kappa}_n^*$ is the n th cumulant of $\log R_T$.⁸

In contrast, SVIX measures the risk-neutral variance of the simple return:

$$SVIX^2 = \frac{1}{T} \text{var}^* R_T. \quad (16)$$

Proof. Equation (13) follows from the definition of VIX^2 and (12), together with the fact that $\mathbb{E}^* R_T = e^{rT}$.

Equation (14) follows from (13) because $\log E^* R_T = \mathbb{E}^* \log R_T + \frac{1}{2} \text{var}^* \log R_T$ if R_T is lognormal. For the approximation, write $\mu = \mathbb{E}^* \log R_T$ and $\sigma^2 = \text{var}^* \log R_T$; over short time horizons, $\text{var}^* R_T = e^{2\mu} (e^{2\sigma^2} - e^{\sigma^2}) = e^{2\mu} \sigma^2 + O(\sigma^4) \approx \sigma^2 = \text{var}^* \log R_T$.

For the general result (15), we introduce the function $\kappa^*(\theta) = \log \mathbb{E}^* [e^{\theta \log R_T}]$. This function can be expanded as a power series in θ ,

$$\kappa^*(\theta) = \sum_{n=1}^{\infty} \frac{\tilde{\kappa}_n^* \theta^n}{n!},$$

where $\tilde{\kappa}_n^*$ is the n th risk-neutral cumulant of $\log R_T$. The definition of entropy implies that $L(R_T) = \kappa^*(1) - \kappa^{*\prime}(0)$, from which (15) follows after annualizing the cumulants: $\kappa_n^* \equiv \frac{1}{T} \tilde{\kappa}_n^*$. For Normally distributed random variables, all cumulants above the variance are zero, so skewness and excess kurtosis (and so on) drop out in the lognormal case.

Finally, using asterisks to indicate variances and expectations with respect to the risk-neutral measure, we have

$$\text{var}^* R_T = \mathbb{E}^* \left[\left(\frac{S_T}{S_0} \right)^2 \right] - \left[\mathbb{E}^* \left(\frac{S_T}{S_0} \right) \right]^2 = \frac{e^{rT} \Pi(T)}{S_0^2} - e^{2rT}.$$

From (3), this implies

$$\text{var}^* R_T = \frac{2e^{rT}}{S_0^2} \left\{ \int_0^{F_T} \text{put}_T(K) dK + \int_{F_T}^{\infty} \text{call}_T(K) dK \right\},$$

and (16) follows, as required. \square

⁸So, for example, $\tilde{\kappa}_1^* = \mathbb{E}^* \log R_T$; $\tilde{\kappa}_2^* = \text{var}^* \log R_T$; $\tilde{\kappa}_3^*$ is the skewness of $\log R_T$ multiplied by $(\tilde{\kappa}_2^*)^{3/2}$; $\tilde{\kappa}_4^*$ is the excess kurtosis, multiplied by $(\tilde{\kappa}_2^*)^2$; and so on.




Figure 2: If the prices of call and put options expiring at time T are as shown, then the annualized risk-neutral variance of the underlying asset's simple return equals the shaded area under the curves multiplied by $2e^{rT}/(TS_0^2)$.

No part of Result 5 requires the Itô process assumption, because the derivation of equation (12) only depends on the static Breeden-Litzenberger logic. The entropy operator $L(\cdot)$ provides a measure of the variability of positive random variables. Like variance, it is nonnegative by Jensen's inequality, and like variance it measures variability by the extent to which a concave function of an expectation of a random variable exceeds an expectation of a concave function of a random variable. It has been applied in the finance literature by Alvarez and Jermann (2005), who refer to it as Theil's (1967) second entropy measure, and by Backus, Chernov and Martin (2010).⁹

Equation (16) has a nice graphical implication that is illustrated in Figure 2, which shows how to calculate the risk-neutral variance of the underlying asset's simple return to time T , given the prices of call and put options of all strikes expiring at time T . Calls and puts have equal value when the strike equals the forward price, so the two lines intersect at $K = F_T$. The annualized risk-neutral variance equals the shaded area under the two curves multiplied by $2e^{rT}/(TS_0^2)$. SVIX is the square root of this quantity, so measures risk-neutral volatility.

Equation (15) implies that if returns are more negatively skewed then, all else equal, VIX will be *lower*; while, intuitively, one might have expected that negative skewness would drive VIX *higher*. This logic, of course, is based on intuition about real-world, not risk-neutral, cumulants, and it can be clarified by considering an equilibrium model that supplies a link between the two. The model also helps to bring out the distinction between the information in a simple variance swap and the information in VIX.

⁹Entropy is an overloaded term. I use it here because of the link to Theil (1967). Backus, Chernov and Martin (2010) refer to $L(M)$ as the entropy of a stochastic discount factor M because, in a complete market, $L(M)$ can be shown to equal the *relative entropy*, in the information-theoretic sense, of risk-neutral probabilities with respect to real-world probabilities.

Result 6 (Interpretation of VIX and SVIX in an equilibrium model). *If there is a representative agent with log utility, then VIX can be expressed in terms of the cumulants of $\log R_T$ under the real-world probabilities, i.e. $\kappa_1 = \frac{1}{T}\mathbb{E} \log R_T$, $\kappa_2 = \frac{1}{T} \text{var} \log R_T$, and so on,*

$$\text{VIX}^2 = 2 \sum_{n=2}^{\infty} (-1)^n (n-1) \frac{\kappa_n}{n!} = \kappa_2 - \frac{2}{3}\kappa_3 + \frac{\kappa_4}{4} - \frac{\kappa_5}{15} + \dots, \quad (17)$$

while SVIX measures the risk premium under the real-world probabilities,

$$\text{SVIX}^2 = \frac{e^{rT}}{T} \mathbb{E} [R_T - e^{rT}]. \quad (18)$$

Proof. To derive (17), let $\kappa(\theta) = \log \mathbb{E} [e^{\theta \cdot \log R_T}] = \sum_{n=1}^{\infty} \frac{\tilde{\kappa}_n \theta^n}{n!}$ (the generating function of the real-world cumulants $\tilde{\kappa}_n$), and let $\kappa^*(\theta)$ be the corresponding CGF calculated with respect to risk-neutral probabilities. In this notation, equation (13) becomes

$$\text{VIX}^2 = \frac{2}{T} [\kappa^*(1) - \kappa^{*'}(0)]. \quad (19)$$

If the representative investor—a holder of the market portfolio, which is assumed to be the asset underlying VIX and SVIX—has log utility, then the reciprocal of the market return, $1/R_T$, is a stochastic discount factor, so for any time- T payoff X , $\mathbb{E}^* X = e^{rT} \mathbb{E} [X/R_T]$. With $X = 1$, this implies that $rT = -\kappa(-1)$; and from this it follows, on setting $X = R_T^\theta$, that $\kappa^*(\theta) = \kappa(\theta - 1) - \kappa(-1)$. Using these observations, (19) becomes

$$\begin{aligned} \text{VIX}^2 &= \frac{2}{T} [-\kappa(-1) - \kappa'(-1)] \\ &= -\frac{2}{T} \left[\sum_1^{\infty} \frac{(-1)^n \tilde{\kappa}_n}{n!} + \sum_1^{\infty} \frac{(-1)^{n-1} \tilde{\kappa}_n}{(n-1)!} \right], \end{aligned}$$

which simplifies to (17) after annualizing the cumulants: $\kappa_n \equiv \frac{1}{T} \tilde{\kappa}_n$.

Similarly, equation (18) follows immediately from equation (16). (Note, incidentally, that from equation (13) we have

$$\text{VIX}^2 = 2r + \frac{2e^{rT}}{T} \mathbb{E} \left[\frac{1}{R_T} \log \frac{1}{R_T} \right]; \quad (20)$$

this is an alternative to (17) that is more directly comparable to (18). \square

Equation (20) shows that VIX does not have an obvious interpretation within the model, though the extreme sensitivity to the possibility of bad outcomes is evident once again. We can also reevaluate the intuition that making skewness more negative should drive VIX up. Equation (17) shows that this is true when cumulants are calculated with respect to real-world probabilities; equation (15) shows it is false when cumulants are calculated with respect to risk-neutral probabilities.

In contrast, the interpretation of SVIX within the model is clear: it gives us a direct measure of the expected risk premium on the market under the true, not the risk-neutral, probability distribution. This is a natural—perhaps *the* natural—measure of risk.

3 Market-implied correlation

The strikes of simple variance swaps on the market and of simple variance swaps on the constituents of the market can be combined to supply a measure of market-implied (i.e. risk-neutral) correlation. Write σ_M^2 for the risk-neutral variance of the *simple* return on the market from time 0 to time T ; write σ_i^2 for the risk-neutral variance of the *simple* return on stock i from time 0 to time T ; write ρ_{ij} for the correlation between stocks i and j ; and write w_i for the market weights:

$$\sigma_M^2 = \text{var}^* \left(\sum_{i=1}^N w_i R_i \right) = \sum_{i=1}^N w_i^2 \sigma_i^2 + \sum_{i \neq j} w_i w_j \rho_{ij} \sigma_i \sigma_j. \quad (21)$$

Note that for the purposes of backing out a correlation measure, the variances we are interested in, σ_M^2 and σ_i^2 , are the variances of *simple returns*, not log returns: because the quantity $\text{var}^* \log \sum w_i R_i$ does not decompose nicely into a sum of individual variances $\text{var}^* \log R_i$, there is no analogue of (21) for log returns.

Based on (21), we can define the *implied correlation* measure

$$\hat{\rho} = \frac{\sigma_M^2 - \sum w_i^2 \sigma_i^2}{\sum_{i \neq j} w_i w_j \sigma_i \sigma_j}.$$

The implied correlation $\hat{\rho}$ can then be calculated directly from the strike on a simple variance swap on the market (which reveals σ_M^2 , via equations (6) and (16)) and from the strikes on simple variance swaps on the index constituents (which reveal $\{\sigma_i^2\}_{i=1, \dots, N}$), together with the observable market weights w_i .

Computing a measure of market-implied correlation from (standard) variance swaps is considerably more challenging. Even under the assumption that prices are continuous, the strike on a variance swap would not reveal risk-neutral variance, but the risk-neutral expectation of integrated instantaneous variance, as in equation (10). And even if the index constituents had lognormal simple returns $R_T^{(i)}$, so that variance swaps on the underlying assets revealed the risk-neutral variance of log returns, $\text{var}^* \log R_T^{(i)}$, as in (14), the market return itself—a sum of lognormals—would not be lognormal. In the presence of jumps, these problems are even more acute.

4 Conclusion

The market turmoil of 2008 and 2009 should have been the volatility derivatives market's moment in the sun. Unfortunately, precisely when the ability to hedge and speculate in volatility and variance would have been particularly valuable, these markets dried up.

This decline in liquidity can be attributed to the critical assumption that underpins the theory of pricing and hedging of variance swaps: that prices follow diffusions, and hence cannot jump. The dependence on this assumption means that variance swaps are hardest to price at times when they are needed most. This problem is particularly severe for variance swaps on single names: Carr and Lee (2009) observe that “the contractual payoffs that appear in thousands of term sheets become literally infinite if the underlying ever closes at zero.” As a workaround, it is now the market convention to cap the payoffs on variance swaps, though doing so complicates both the pricing and the interpretation of the contract.

This paper has analyzed a financial contract, a *simple variance swap*, that is closely related to variance swaps. Unlike variance swaps, however, simple variance swaps can be priced and hedged even in the presence of jumps. The weighting by forward price in the definition (1) leads to an important simplification of the hedging strategy that is critical to make the contract potentially tradable in practice.

Simple variance swaps have a natural interpretation: whether or not there are jumps, they reveal the (risk-neutral) variance of the simple return on the underlying asset, so permit hedging and speculation by market participants with views on this quantity.

The contrast between simple variance swaps and variance swaps can be seen most directly by examining their respective hedging portfolios. The hedge portfolio for a simple variance swap holds equal amounts of options of all different strikes, while variance swaps require increasingly large positions in puts with increasingly low strikes. This makes explicit the dependence of variance swaps on extreme events. Even if one is prepared to assert that there are no jumps (as is required to legitimize the theory of variance swap and gamma swap pricing), variance swaps are harder to hedge than simple variance swaps, since they load more strongly on deep-out-of-the-money puts.

As a result of these unfortunate properties of variance swaps, gamma swaps have recently started trading. While the hedge portfolio for variance swaps holds portfolios of options with strikes K in amounts proportional to $1/K^2$, the hedge for a gamma swap holds portfolios of options with strikes K in amounts proportional to $1/K$. However, as with variance swaps, this only holds if prices cannot jump. In the sense that the hedge portfolio on a simple variance swap holds options with strikes K in amounts proportional to 1, simple variance swaps are the answer to the question: ‘What is the next member of the sequence “variance swap, gamma swap,”?’ But simple variance swaps

are distinguished from the other two by the fact that they can be priced and hedged in the presence of jumps.

By comparing the strikes on simple variance swaps on an index to the strikes on simple variance swaps on the index constituents, a measure of implied correlation can be calculated. What one needs here are the variances of *simple* returns, not log returns. Thus simple variance swaps would provide the natural way to measure implied correlation even in a world in which jumps did not occur. With jumps, their advantages become even clearer. Most obviously—even setting aside the difficulties in pricing variance swaps in the presence of jumps, and in using the resulting prices to compute a correlation measure—it is critical that the single-name market should not evaporate at times of stress, since the variances of single names are needed to compute correlation.

5 Bibliography

- Alvarez, F., and U. J. Jermann (2005), “Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth,” *Econometrica*, 73:6:1977–2016.
- Backus, D., M. Chernov, and I. Martin (2010), “Disasters Implied by Equity Index Options”, NBER working paper.
- Breedon, D. T., and R. H. Litzenberger (1978), “Prices of State-Contingent Claims Implicit in Option Prices,” *Journal of Business*, 51:4:621–651.
- Carr, P., and A. Corso (2001), “Covariance Contracting for Commodities,” *EPRM* April 2001.
- Carr, P., and R. Lee (2009), “Volatility Derivatives,” *Annual Review of Financial Economics*, 1:1–21.
- Carr, P., and D. Madan (1998), “Towards a Theory of Volatility Trading,” in R. Jarrow, ed., *Volatility: New Estimation Techniques for Pricing Derivatives*, London: Risk Books, pp. 417–427.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999), “More Than You Ever Wanted to Know about Volatility Swaps,” Goldman Sachs Quantitative Strategies Research Notes.
- Jarrow, R. A., P. Protter, M. Larsson, and Y. Kchia (2010), “Variance and Volatility Swaps: Bubbles and Fundamental Prices,” Cornell University, Johnson School Research Paper.
- Lee, R. (2010a), “Gamma Swap,” in Rama Cont, ed., *Encyclopedia of Quantitative Finance*, Wiley.
- Lee, R. (2010b), “Weighted Variance Swap,” in Rama Cont, ed., *Encyclopedia of Quantitative Finance*, Wiley.
- Neuberger, A. (1990), “Volatility trading,” working paper, London Business School.
- Neuberger, A. (1994), “The Log Contract,” *Journal of Portfolio Management*, 20:2:74–80.

Theil, H. (1967), *Economics and Information Theory*, Amsterdam: North-Holland.

A Appendix

This appendix provides a derivation of (12). Using the result of Breeden and Litzenberger (1978), we have

$$P_{\log} = \int_0^{\infty} \log \frac{K}{S_0} \text{call}'_T(K) dK.$$

Evaluating this integral is a straightforward exercise in integration by parts, though we do need one small trick right at the beginning, splitting the range of integration into two parts and using the observation that $\text{put}''_T(K) \equiv \text{call}''_T(K)$, which follows from put-call parity; and, half-way through, to use the fact that $\text{put}'_T(K) - \text{call}'_T(K) = e^{-rT}$, which again follows from put-call parity.

$$\begin{aligned} P_{\log} &= \int_0^{F_T} \log \frac{K}{S_0} \text{put}''_T(K) dK + \int_{F_T}^{\infty} \log \frac{K}{S_0} \text{call}''_T(K) dK \\ &= \left[\log \frac{K}{S_0} \cdot \text{put}'_T(K) \right]_0^{F_T} - \int_0^{F_T} \frac{1}{K} \text{put}'_T(K) dK + \left[\log \frac{K}{S_0} \cdot \text{call}'_T(K) \right]_{F_T}^{\infty} - \int_{F_T}^{\infty} \frac{1}{K} \text{call}'_T(K) dK \\ &= rT e^{-rT} - \int_0^{F_T} \frac{1}{K} \text{put}'_T(K) dK - \int_{F_T}^{\infty} \frac{1}{K} \text{call}'_T(K) dK \\ &= rT e^{-rT} + \left[\frac{-\text{put}_T(K)}{K} \right]_0^{F_T} - \int_0^{F_T} \frac{1}{K^2} \text{put}_T(K) dK + \left[\frac{-\text{call}_T(K)}{K} \right]_{F_T}^{\infty} - \int_{F_T}^{\infty} \frac{1}{K^2} \text{call}_T(K) dK \\ &= rT e^{-rT} - \int_0^{F_T} \frac{1}{K^2} \text{put}_T(K) dK - \int_{F_T}^{\infty} \frac{1}{K^2} \text{call}_T(K) dK. \end{aligned}$$

asset	0	Δ	2Δ	\dots	$T - \Delta$	T
B	$-e^{-rT}$			\dots		$\frac{S_0^2}{S_0^2}$
U	$2e^{-r(T-\Delta)}$	$-2e^{-r(T-\Delta)} \frac{S_\Delta}{S_0}$		\dots		
B		$2e^{-r(T-\Delta)} \frac{S_\Delta}{S_0}$		\dots		$-2 \frac{S_0 S_\Delta}{S_0^2}$
Δ	$-\frac{(e^{r\Delta}-1)^2 \Pi_\Delta}{e^{r(T-\Delta)} F_\Delta^2}$	$\frac{(e^{r\Delta}-1)^2 S_\Delta^2}{e^{r(T-\Delta)} F_\Delta^2}$		\dots		
B		$-e^{-r(T-\Delta)} \left(\frac{S_\Delta}{S_0} + \frac{S_\Delta^2}{F_\Delta^2} \right)$		\dots		$\frac{S_\Delta^2}{S_0^2} + \frac{S_\Delta^2}{F_\Delta^2}$
U		$\frac{2e^{-r(T-2\Delta)} S_\Delta^2}{F_\Delta^2}$	$-\frac{2e^{-r(T-2\Delta)} S_\Delta S_{2\Delta}}{F_\Delta^2}$	\dots		
B			$\frac{2e^{-r(T-2\Delta)} S_\Delta S_{2\Delta}}{F_\Delta^2}$	\dots		$-\frac{2S_\Delta S_{2\Delta}}{F_\Delta^2}$
\vdots	\vdots			\dots		\vdots
$T - \Delta$	$-\frac{(e^{r\Delta}-1)^2 \Pi_{T-\Delta}}{e^{r\Delta} F_{T-\Delta}^2}$			\dots	$\frac{e^{-r\Delta} (e^{r\Delta}-1)^2 S_{T-\Delta}^2}{F_{T-\Delta}^2}$	
B				\dots	$-e^{-r\Delta} \left(\frac{S_{T-\Delta}^2}{F_{T-2\Delta}^2} + \frac{S_{T-\Delta}^2}{F_{T-\Delta}^2} \right)$	$\frac{S_{T-\Delta}^2}{F_{T-2\Delta}^2} + \frac{S_{T-\Delta}^2}{F_{T-\Delta}^2}$
U				\dots	$\frac{2S_{T-\Delta}^2}{F_{T-\Delta}^2}$	$-\frac{2S_{T-\Delta} S_T}{F_{T-\Delta}^2}$
T		$-\frac{\Pi_T}{F_{T-\Delta}^2}$		\dots		$\frac{S_T^2}{F_{T-\Delta}^2}$
B		$V e^{-rT}$		\dots		$-V$

Table 1: Replicating the simple variance swap. In the left column, B indicates positions in the bond, U indicates positions in the underlying, and $j\Delta$ indicates a position in the portfolio of options expiring at time $j\Delta$ that replicates the payoff $S_{j\Delta}^2$.

asset	0	$j\Delta$	$(j+1)\Delta$	T
$j\Delta$	$-\frac{(e^{r\Delta}-1)^2 \Pi_{j\Delta}}{e^{r(T-j\Delta)} F_{j\Delta}^2}$	$\frac{(e^{r\Delta}-1)^2 S_{j\Delta}^2}{e^{r(T-j\Delta)} F_{j\Delta}^2}$		
B		$-e^{-r(T-j\Delta)} \left(\frac{S_{j\Delta}^2}{F_{(j-1)\Delta}^2} + \frac{S_{j\Delta}^2}{F_{j\Delta}^2} \right)$		$\frac{S_{j\Delta}^2}{F_{(j-1)\Delta}^2} + \frac{S_{j\Delta}^2}{F_{j\Delta}^2}$
U		$\frac{2S_{j\Delta}^2}{e^{r(T-(j+1)\Delta)} F_{j\Delta}^2}$	$\frac{-2S_{j\Delta} S_{(j+1)\Delta}}{e^{r(T-(j+1)\Delta)} F_{j\Delta}^2}$	
B			$\frac{2S_{j\Delta} S_{(j+1)\Delta}}{e^{r(T-(j+1)\Delta)} F_{j\Delta}^2}$	$\frac{-2S_{j\Delta} S_{(j+1)\Delta}}{F_{j\Delta}^2}$

Table 2: Replicating the simple variance swap. The generic position in options of intermediate maturity, expiring at time $j\Delta$, together with the associated trades required after expiry. In the left column, B indicates positions in the bond, U indicates positions in the underlying, and $j\Delta$ indicates positions in options expiring at $j\Delta$.