Peter Carr and the variance contract

Jim Gatheral

Baruch College
Weissman School of Arts and Sciences
Master of Financial Engineering

The Peter Carr Memorial Conference
NYU Tandon School of Engineering, June 4, 2022

Outline of this talk

- The Carr-Madan spanning formula
- Variance and gamma contracts
- The VIX
- Optimal trading under log utility
- Volatility arbitrage
- The variance premium

Towards a theory of volatility trading

29

Towards a Theory of Volatility Trading

Peter Carr and Dilip Madan
Morgan Stanley; University of Maryland

We review three methods which have emerged for trading realised volatility. The first method involves taking static positions in options, such as straddles, while the second method involves delta-hedging an option position. If the investor is successful in hedging away the price risk, then a prime determinant of the profit or loss from this strategy is the difference between the realised volatility and the anticipated volatility used in pricing and hedging the option. The final method involves buying or selling a "vol swap". This contract pays the buyer the difference between the realised volatility and the fixed swap rate determined at the outset of the contract. We also uncover the link between volatility contracts and some recent work on local volatility.

The Carr-Madan spanning formula

- Assume that European options with all possible strikes and expirations are traded.
- The Carr-Madan spanning formula shows that any twice-differentiable payoff at time T may be statically hedged using a portfolio of European options expiring at time T.

Lemma (Carr-Madan spanning formula)

Consider a generalized European payoff $g(S_T)$. Then

$$\begin{aligned} g(S_T) &= \int_0^F g''(K)\,(K - S_T)^+\,dK + \int_F^\infty g''(K)\,(S_T - K)^+\,dK \\ &\quad + g(F) + g'(F)\,(S_T - F). \end{aligned} \tag{1}$$

Proof

- Following [CM1998], the value of a claim with a generalized payoff $g(S_T)$ at time T is given by

$$\begin{aligned} g(S_T) &= \int_0^\infty g(K)\,\delta(S_T - K)\,dK \\ &= \int_0^F g(K)\,\delta(S_T - K)\,dK + \int_F^\infty g(K)\,\delta(S_T - K)\,dK. \end{aligned}$$

- Integrating by parts gives

$$\begin{aligned} g(S_T) &= g(F) - \int_0^F g'(K)\,\theta(K - S_T)\,dK \\ &\quad + \int_F^\infty g'(K)\,\theta(S_T - K)\,dK. \end{aligned}$$

- ... and integrating by parts again gives

$$
\begin{aligned}
& g(S_T) \\
= \; & \int_0^F g''(K)\,(K - S_T)^+\,dK + \int_F^\infty g''(K)\,(S_T - K)^+\,dK \\
& + g(F) - g'(F)\,\left[(F - S_T)^+ - (S_T - F)^+\right] \\
= \; & \int_0^F g''(K)\,(K - S_T)^+\,dK + \int_F^\infty g''(K)\,(S_T - K)^+\,dK \\
& \qquad + g(F) + g'(F)\,(S_T - F).
\end{aligned}
$$

- Equation (1) shows how to decompose any payoff $g(S_T)$ into hockey-stick payoffs.
  - In particular, any such payoff can be hedged with a *static* position in European vanilla options and forward contracts.

The fair value of the payoff $g(S_T)$

With $F = \mathbb{E}[S_T]$,

$$\mathbb{E}\left[g(S_T)\right] = g(F) + \int_0^F dK\, \tilde{P}(K)\, g''(K) + \int_F^\infty dK\, \tilde{C}(K)\, g''(K) \qquad (2)$$

where $\tilde{P}$ and $\tilde{C}$ represent undiscounted put and call prices.

- The fair value of $g(S_T)$ is thus expressed in terms of an (in general) infinite strip of puts and calls.

Remarks on spanning of European-style payoffs

- From equations (1) and (2), we see that any European-style twice-differentiable payoff may be replicated using a portfolio of European options with strikes from 0 to ∞.
  - The weight of each option equal to the second derivative of the payoff at the strike price of the option.
- This portfolio of European options is a static hedge because the weight of an option with a particular strike depends only on the strike price and the form of the payoff function and not on time or the level of the stock price.
- Note further that (1) is *completely model-independent* (assuming continuity of paths of the underlying).

The log-contract and quadratic variation

Now consider a contract whose payoff at time T is $X_T = \log(S_T/F)$. Then $g''(K) = -\left.1/{S_T}^2\right|_{S_T=K}$ and it follows from equation (2) that

$$\mathbb{E}\left[\log\left(\frac{S_T}{F}\right)\right] = -\int_0^F \frac{dK}{K^2}\,\tilde{P}(K) - \int_F^\infty \frac{dK}{K^2}\,\tilde{C}(K) \quad (3)$$

Rewriting (3) in terms of the log-strike variable $k := \log(K/F)$, we get the promising-looking expression

$$\begin{aligned}\mathbb{E}\left[\log\left(\frac{S_T}{F}\right)\right] &= -\int_{-\infty}^{0} dk\, p(k) - \int_{0}^{\infty} dk\, c(k) \\ &:= -\int_{-\infty}^{\infty} dk\, q(k) \qquad (4)\end{aligned}$$

with

$$q(k) := \frac{\min[\tilde{C}(F e^k), \tilde{P}(F e^k)]}{F e^k}$$

representing out-of–the-money option prices expressed in terms of percentage of the strike price.

- Henceforth we assume zero interest rates and dividends, so $\mathbb{E}[S_T] = S_0$, $\tilde{C} = C$, $\tilde{P} = P$, and so on.

The variance contract

With zero rates and dividends, $F = S_0$ and applying Itô's Lemma, path-by-path

$$\begin{aligned} \log\left(\frac{S_T}{F}\right) &= \log\left(\frac{S_T}{S_0}\right) \\ &= \int_0^T d\log\left(S_t\right) \\ &= \int_0^T \frac{dS_t}{S_t} - \int_0^T \frac{v_t}{2}dt \end{aligned} \tag{5}$$

Hedging the variance contract

- The second term on the RHS of (5) is immediately recognizable as half the quadratic variation $\langle X \rangle_T$ over the interval $[0, T]$.
- The first term on the RHS represents the payoff of a hedging strategy which involves maintaining a constant dollar amount in stock (if the stock price increases, sell stock; if the stock price decreases, buy stock so as to maintain a constant dollar value of stock).

Hedging the variance contract

- Since the log payoff on the LHS can be hedged using a portfolio of European options as noted earlier, it follows that quadratic variation may be replicated in a completely model-independent way so long as the stock price process is a diffusion.
- In particular, volatility may be stochastic or deterministic and (5) still applies.

The log-strip hedge for a variance contract

- Now taking the risk-neutral expectation of (5) and comparing with equation (4), we obtain

$$\begin{aligned}\mathbb{E}\left[\int_0^T v_t dt\right] &= -2\,\mathbb{E}\left[\log\left(\frac{S_T}{F}\right)\right] \\ &= 2\int_{-\infty}^{\infty} q(k)\,dk. \qquad (6)\end{aligned}$$

- We see that the fair value of quadratic variation is given by the value of an infinite strip of European options in a completely *model-independent* way so long as the underlying process is a diffusion.

One application: The VIX index

- In 2004, the CBOE listed futures on the VIX - an implied volatility index.
- Originally, the VIX computation was designed to mimic the implied volatility of an at-the-money 1 month option on the OEX index. It did this by averaging volatilities from 8 options (puts and calls from the closest to ATM strikes in the nearest and next to nearest months).
- The CBOE changed the VIX computation: "CBOE is changing VIX to provide a more precise and robust measure of expected market volatility and to create a viable underlying index for tradable volatility products."

The VIX formula

- Here is the revised VIX definition (converted to our notation) as specified in the CBOE white paper:

$$VIX^2 = \frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} Q_i(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2$$

  where Q_i is the price of the out-of-the-money option with strike K_i and K_0 is the highest strike below the forward price F .
- We recognize this formula as a straightforward discretization of the variance log-strip (6) and makes clear the reason why the CBOE implies that the new index permits replication of volatility.

History of the VIX

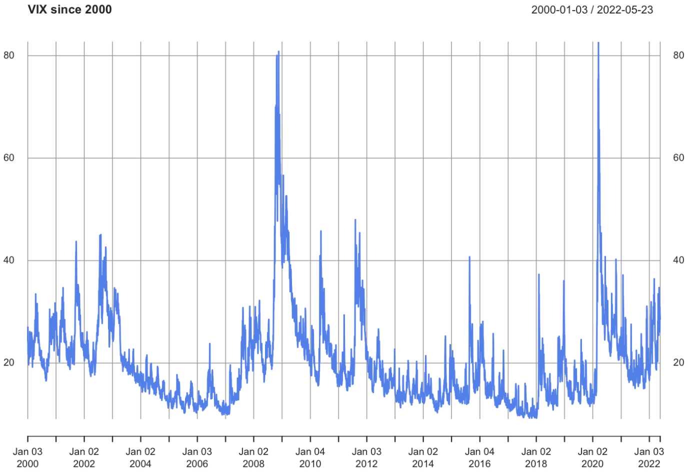

Figure 1: The VIX index since 2000. (SPX and VIX data from Yahoo!)

VIX futures and options

- The variance log-strip underlies the construction of the VIX index.
  - Since 2013, more vega is traded in VIX futures and options than is traded in SPX.

VIX open interest and volume

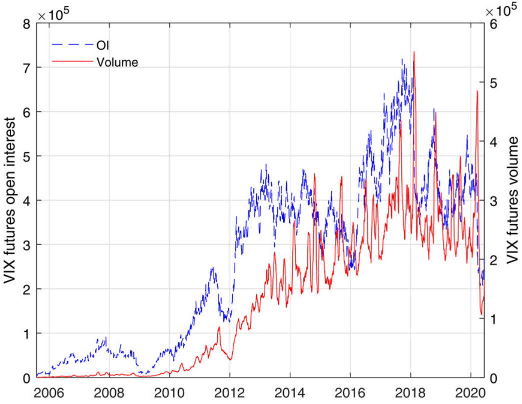

Figure 2: VIX futures open interest and volume (from [Pos2021])

The corridor variance contract

- Now consider a contract whose payoff at time T is

$$\begin{aligned} g_C(S_T) &= -2\log\frac{S_T}{S_0}\,\mathbb{1}_{\{K_-\le S_T\le K_+\}} \\ &\quad +2\left[-\frac{S_T}{{K_-}^2}-\log\frac{K_-}{S_0}+\frac{1}{K_-}\right]\mathbb{1}_{\{S_T<K_-\}} \\ &\quad +2\left[-\frac{S_T}{{K_+}^2}-\log\frac{K_+}{S_0}+\frac{1}{K_+}\right]\mathbb{1}_{\{S_T>K_+\}}. \end{aligned}$$

  - Just the log-contract in the inner interval, linearly extrapolated.
- Then $g_C(\cdot)$ and $g_C'(\cdot)$ are continuous. Moreover

$$g_C''(S_T)=\frac{2}{{S_T}^2}\,\mathbb{1}_{\{K_-\le S_T\le K_+\}}.$$

The payoff g_C

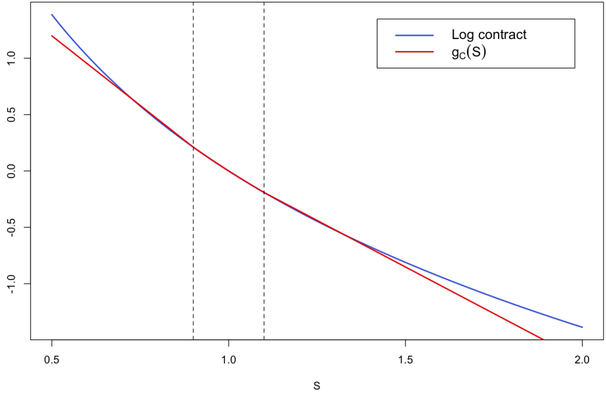

Figure 3: The log contract payoff $-2 \log S/S_0$ in blue; the payoff $g_C(S)$ in red.

Hedging the corridor variance contract

Similarly, the model-free hedging strategy is to:

- Hold a strip of options with strikes between K_- and K_+.
- Maintain a constant dollar amount in stock, if the stock price is between K_- and K_+.
  - No rehedging if the stock price is outside this interval.

- Applying Itô's Lemma path-by-path and taking expectations as before gives

$$\begin{aligned} & \mathbb{E}\left[\int_0^T v_t \, \mathbb{1}_{\{K_- \leq S_t \leq K_+\}} \, dt\right] \\ = \;& \mathbb{E}\left[g_C(S_T)\right] = 2 \int_{k_-}^{k_+} q(k) \, dk. \end{aligned} \tag{7}$$

  where $k_\pm := \log K_\pm / S_0$.

- We see that the fair value of a corridor variance contract that pays only when $K_- \leq S_t \leq K_+$ is given by a strip of options with strikes above K_- and below K_+.

Forward starting corridor variance contracts

- We get the forward-starting version by trading the spread:

$$
\begin{aligned}
& \mathbb{E}\left[\int_{T_-}^{T_+} v_t \, \mathbb{1}_{\{K_- \le S_t \le K_+\}} \, dt\right] \\
= \; & \mathbb{E}\left[\int_{0}^{T_+} v_t \, \mathbb{1}_{\{K_- \le S_t \le K_+\}} \, dt\right] - \mathbb{E}\left[\int_{0}^{T_-} v_t \, \mathbb{1}_{\{K_- \le S_t \le K_+\}} \, dt\right] \\
= \; & 2 \int_{k_-}^{k_+} \left[q(k, T_+) - q(k, T_-)\right] dk. \qquad (8)
\end{aligned}
$$

- In this way, we can trade instantaneous variance v_t conditional on $S_t \in (K_-, K_+)$ and $t \in (T_-, T_+)$.
  - A *localized* variance contract.
- An immediate practical corollary is that local variance surfaces should be as smooth as possible.
  - Localized variance spreads can in principle be traded.

Dupire's local variance

- Recall that Dupire's local variance is given by

$$v_\ell(K,T) = \mathbb{E}\left[v_T \middle| S_T = K\right].$$

- In the limit $T_\pm \to T$, $K_\pm \to K$, we obtain

$$\begin{aligned}
\mathbb{E}\left[\int_{T_-}^{T_+} v_t\, \mathbb{1}_{\{K_- \le S_t \le K_+\}}\, dt\right] &\approx \mathbb{E}\left[v_T\, \mathbb{1}_{\{K_- \le S_T \le K_+\}}\right] dT \\
&\approx v_\ell(K,T)\, \mathbb{E}\left[\mathbb{1}_{\{K_- \le S_T \le K_+\}}\right] dT \\
&\approx v_\ell(K,T)\, p(K,T)\, dK\, dT, \qquad (9)
\end{aligned}$$

where $p(K,T) = \partial_{K,K} C(K,T)$ is the risk-neutral density.

- Also on the RHS,

$$2\int_{k_-}^{k_+} \left[q(k,T_+) - q(k,T_-)\right] dk \approx 2\, \partial_T q(k,T)\, dk\, dT \qquad (10)$$

The Dupire Formula

Equating the RHS (10) with the LHS (9) and using that

$$q(k,T) = \frac{C(K,T)}{K} \text{ and } dk = \frac{dK}{K},$$

we get

The Dupire Formula

$$v_\ell(K,T) = \frac{2\,q(k,T)}{K\,p(K,T)} = \frac{\partial_T C(K,T)}{\frac{1}{2}\,K^2 \partial_{K,K} C(K,T)}.$$

Weighted variance contracts

- Consider the weighted variance contract $\int_0^T \alpha(S_t)\, v_t\, dt$.
- Following [Fuk2014], an application of Itô's Lemma gives the quasi-static hedge:

$$\int_0^T \alpha(S_t)\, v_t\, dt = A(S_T) - A(S_0) - \int_0^T A'(S_u)\, dS_u \quad (11)$$

with

$$A(x) = 2 \int_1^x dy \int_1^y \frac{\alpha(z)}{z^2}\, dz.$$

- The LHS of (11) is the payoff to be hedged.
- The last term on the RHS corresponds to rebalancing.
- The first term on the RHS corresponds to a static position in options given by the spanning formula (1).

The gamma contract

- The payoff of a gamma contract is $\frac{1}{S_0}\int_0^T S_t\, v_t\, dt$.
- Thus $\alpha(x) = x$ and

$$A(x) = \frac{2}{S_0}\int_1^x dy \int_1^y \frac{z}{z^2}\, dz = \frac{2}{S_0}\left\{1 - x + x \log x\right\}.$$

- The static options hedge is the spanning strip for $2\,\frac{S_T}{S_0}\,\log\frac{S_T}{S_0}$
  - The contract with payoff $\frac{S_T}{S_0}\,\log\frac{S_T}{S_0}$ is known as the *entropy contract*.

A cool formula for the variance contract

- Define

$$d_{\pm} = -\frac{k}{\sigma_{BS}(k)\sqrt{T}} \pm \frac{\sigma_{BS}(k)\sqrt{T}}{2}$$

and further define the inverse functions $g_{\pm}(z) = d_{\pm}^{-1}(z)$. Intuitively, z measures the log-moneyness of an option in implied standard deviations.

- Then, as a corollary of result of Matytsin's,

$$\mathbb{E}\left[\int_0^T v_t dt\right] = -2\,\mathbb{E}\left[\log \frac{S_T}{F}\right] = \int_{-\infty}^{\infty} dz\, N'(z)\, \sigma_{BS}^2\left(g_-(z)\right) T.$$

A cool formula for the gamma contract

- Fukasawa [Fuk2012] derives an expression for the value of a generalized European payoff in terms of implied volatilities.
- As one application, he derives the following expression for the value of a gamma contract.

$$\mathbb{E}\left[\int_0^T \frac{S_t}{F}\, v_t dt\right] = 2\,\mathbb{E}\left[\frac{S_T}{F} \log \frac{S_T}{F}\right] = \int_{-\infty}^{\infty} dz\, N'(z)\, \sigma^2_{BS}\left(g_+(z)\right) T.$$

(note g_+ instead of g_- in the variance contract case).

- In particular, if we have a parameterization of the volatility smile (such as SVI), computing the fair value of the covariance contract is straightforward.

Optimal positioning in derivative securities

QUANTITATIVE FINANCE VOLUME 1 (2001) 19–37
INSTITUTE OF PHYSICS PUBLISHING

RESEARCH PAPER
quant.iop.org

Optimal positioning in derivative securities

Peter Carr[1] **and Dilip Madan**[2]

[1] Banc of America Securities, 40th Floor, 9 West 57th Street, New York, NY 10019, USA
[2] Robert H Smith School of Business, University of Maryland, College Park, MD 20742, USA

E-mail: pcarr@bofasecurities.com and dbm@rhsmith.umd.edu

Received 21 June 2000, in final form 13 November 2000

Abstract
We consider a simple single period economy in which agents invest so as to maximize expected utility of terminal wealth. We assume the existence of three asset classes, namely a riskless asset (the bond), a single risky asset (the stock), and European options of all strikes (derivatives). In this setting, the inability to trade continuously potentially induces investment in all three asset classes. We consider both a partial equilibrium where all asset prices are initially given, and a more general equilibrium where all asset prices are endogenously determined. By restricting investor beliefs and preferences in each case, we solve for the optimal position for each investor in the three asset classes. We find that in partial or general equilibrium, heterogeneity in preferences or beliefs induces investors to hold derivatives individually, even though derivatives are not held in aggregate.

Optimal investing with options

- Following [CM2001], consider investing initial wealth W_t in a derivative claim with payoff $f(S_T)$. By definition of the pricing measure, $W_t = \mathbb{E}^{\mathbb{Q}}\left[f(S_T) | \mathcal{F}_t \right]$.
- Expected utility of terminal wealth is given by

$$\mathcal{U}_t[f] = \mathbb{E}^{\mathbb{P}}\left[U(W_T) | \mathcal{F}_t \right] = \int \rho_{\mathbb{P}}(S_T)\, U\left(f(S_T) \right) dS_T,$$

  where U is the investor's utility function.
- Now find the f that maximizes $\mathcal{U}_t[f]$ subject to $W_t = \mathbb{E}^{\mathbb{Q}}\left[f(S_T) | \mathcal{F}_t \right]$.

- The Lagrangian of this optimization problem is

$$\mathcal{L}[f] = \int \rho_{\mathbb{P}}(S_T)\, U\left(f(S_T)\right)\, dS_T - \lambda \left[\mathbb{E}^{\mathbb{Q}}\left[f(S_T) \middle| \mathcal{F}_t \right] - W_t\right]$$

where

$$\mathbb{E}^{\mathbb{Q}}\left[f(S_T) \middle| \mathcal{F}_t \right] = \int \rho_{\mathbb{Q}}(S_T)\, f(S_T)\, dS_T.$$

- The first order condition is then

$$\rho_{\mathbb{P}}(S_T)\, U'(f(S_T)) - \lambda\, \rho_{\mathbb{Q}}(S_T) = 0. \tag{12}$$

- Integrating wrt S_T gives

$$\lambda = \int_{\mathbb{R}_{\geqslant 0}} \rho_{\mathbb{P}}(S_T)\, U'(f(S_T))\, dS_T.$$

- We then obtain a risk-adjusted measure

$$\rho_{\mathbb{Q}}(S_T) = \frac{\rho_{\mathbb{P}}(S_T)\, U'(f(S_T))}{\int_{\mathbb{R}_{\geqslant 0}} \rho_{\mathbb{P}}(S_T)\, U'(f(S_T))\, dS_T}.$$

- Thus, the optimal payoff f is such that the risk-adjusted physical measure $\rho_{\mathbb{P}}$ equals the pricing measure $\rho_{\mathbb{Q}}$.
- Solving (12) for $f(\cdot)$ gives

$$f(S) = U'^{-1}\left(\lambda \frac{\rho_{\mathbb{Q}}(S)}{\rho_{\mathbb{P}}(S)}\right) \tag{13}$$

which is Equation (10) of [CM2001].

Marco Avellaneda quote

- Let's now suppose we know the utility function of the representative investor...

Giordano Bruno (1585)

Se non è vero, è ben trovato[a].

[a] From Wiktionary: Even if it is not true, it is a good story

Log utility

- Suppose that the representative investor $\mathcal{I}$ maximizes $\log W$.
  - *i.e.* that $\mathcal{I}$ is a Kelly investor.
- If $U(W) = \log W$, $U'(W) = 1/W$ and (13) becomes

$$\frac{1}{f(S_T)} = \lambda \, \frac{\rho_{\mathbb{Q}}(S_T)}{\rho_{\mathbb{P}}(S_T)}.$$

- Rearranging and integrating gives

$$\lambda \, \mathbb{E}^{\mathbb{Q}} \left[f(S_T) \middle| \mathcal{F}_t \right] = \int_{\mathbb{R}_{\geqslant 0}} \rho_{\mathbb{P}}(S_T) \, dS_T = 1.$$

- It follows that the Kelly-optimal derivative payoff is given by

$$f(S_T) = \mathbb{E}^{\mathbb{Q}} \left[f(S_T) \middle| \mathcal{F}_t \right] \frac{\rho_{\mathbb{P}}(S_T)}{\rho_{\mathbb{Q}}(S_T)}. \tag{14}$$

An equilibrium argument

- From now on, let S_t be the price of the market portfolio (SPX say).
- In equilibrium, the market portfolio is optimal for the representative investor $\mathcal{I}$.
- Thus

$$\frac{d\mathbb{P}}{d\mathbb{Q}} = \frac{\rho_{\mathbb{P}}(S_T)}{\rho_{\mathbb{Q}}(S_T)} = \frac{S_T}{S_t} \tag{15}$$

  – the *Long numeraire portfolio* or *growth optimal portfolio*.
- The change of measure is just the stock price.
  - The change of measure is a $\mathbb{Q}$-martingale since S is a martingale under $\mathbb{Q}$.

Another argument

- The same result follows from the argument, originally due to Mark Davis, that in equilibrium, market option prices should be such that the utility of an optimal stock portfolio cannot be increased by trading options.

The Kelly optimal investment policy and drift under $\mathbb{P}$

- Dynamics of the stock price under $\mathbb{P}$ are

$$\frac{dS_t}{S_t} = \mu_t\,dt + \sqrt{v_t}\,dZ_t^{\mathbb{P}}.$$

- What proportional π_t of wealth W should be held in stock at time t?
- We have

$$dW_t = \pi_t\,W_t\,\frac{dS_t}{S_t} = \pi_t\,W_t\left\{\mu_t\,dt + \sqrt{v_t}\,dZ_t^{\mathbb{P}}\right\}.$$

- Applying Itô's Formula,

$$\log W_T = \log W_t + \int_t^T \pi_s \, \mu_s \, du + \int_t^T \pi_s \sqrt{v_s} \, dZ_s - \frac{1}{2} \int_t^T \pi_s^2 \, v_s \, du.$$

    - It follows that utility is maximized *pathwise* if $\pi_u = \mu_u / v_u$.
- We found that $\pi_t = 1$ is optimal, so we must have $\mu_t = v_t$.
    - The *equity risk premium*, the extra return that investors require for taking on risk, is equal to instantaneous variance.

- Taking $\sigma = 0.15$, we get an equity risk premium of around 2.25% per annum, which seems not unreasonable.

Ross recovery

- In this simple framework, we know the change of measure $d\mathbb{P}/d\mathbb{Q}$.
- So we can get the $\mathbb{P}$ distribution knowing the $\mathbb{Q}$ distribution.
- This is reminiscent of Ross recovery, on which Carr and Yu [CY2012] wrote a fascinating paper.

Risk, Return, and Ross Recovery

PETER CARR AND JIMING YU

PETER CARR
is the executive director of the Masters in Math Finance Program at Courant Institute, New York University in New York, NY.
pcarr@nyc.rr.com

JIMING YU
is a vice president at a large financial institution in New York, NY.

The risk return relation is a staple of modern finance. When risk is measured by volatility, it is well known that option prices convey risk. One of the more influential ideas in the last twenty years is that the conditional volatility of an asset price can also be inferred from the market prices of options written on that asset. Under a Markovian restriction, it follows that risk-neutral transition probabilities can also be determined from option prices. Recently, Ross has shown that real-world transition probabilities of a Markovian state variable can be recovered from its risk-neutral transition probabilities along with a restriction on preferences. In this article, we show how to recover real-world transition probabilities in a bounded diffusion context in a preference-free manner. Our approach is instead based on restricting the form and dynamics of the numeraire portfolio.

market forecasts. Yet when it comes to predicting the average return, the conventional wisdom is that option prices are silent in this respect.

Recently, Stephen Ross has written a working paper [2011], that challenges this conventional wisdom. Under the assumptions of his model, option prices forecast not only the average return, but also the entire return distribution. Further tweaking the nose of conventional wisdom, option prices even convey the conditional return distribution,when the conditioning variable is a Markovian state variable that determines aggregate consumption.

Those of us raised on the Black–Merton–Scholes (BMS) paradigm find Ross's claims to be startling. If one can value options without knowledge of expected return, then how can one use option prices to infer expected return?

Carr and Yu

- In that paper too, we have that the equity risk premium associated with the numeraire portfolio is the instantaneous variance.
- ... but that is a story for another day ...
- Instead, we will focus on the variance risk premium.

Variance risk premiums

Variance Risk Premiums

Peter Carr
Bloomberg LP and Courant Institute, New York University

Liuren Wu
Zicklin School of Business, Baruch College

We propose a direct and robust method for quantifying the variance risk premium on financial assets. We show that the risk-neutral expected value of return variance, also known as the variance swap rate, is well approximated by the value of a particular portfolio of options. We propose to use the difference between the realized variance and this synthetic variance swap rate to quantify the variance risk premium. Using a large options data set, we synthesize variance swap rates and investigate the historical behavior of variance risk premiums on five stock indexes and 35 individual stocks. (*JEL* G10, G12, G13)

Why trade variance?

- Banks buy variance to hedge short vega risk.
- Hedge funds sell variance to capture the variance risk premium.
  - If realized variance is less than the strike, this trade makes money - so-called *volatility arbitrage*.
  - In a top-cited paper, Carr and Wu [CW2009] studied the *variance premium*, the amount that this short variance trade is expected to make.

History of variance contracts

- Variance contracts took off as a product in the aftermath of the LTCM meltdown in late 1998 when implied stock index volatility levels rose to unprecedented levels.
  - LTCM had dominated the market as sellers of variance.
  - Probably the first and biggest 'volatility arbitrageurs' ever.
- Hedge funds took advantage of this squeeze by selling variance contracts at historically incredibly high levels.
  - The key to their willingness to sell variance contracts rather than sell options was that a variance contract is a pure play on realized variance – no labor-intensive delta hedging or other path dependency is involved.

Definition of the variance risk premium

- According to our definition (which is the negative of Carr and Wu's), the variance risk premium is simply

$$\mathbb{E}^{\mathbb{Q}}\left[\int_t^T v_u\,du\middle|\mathcal{F}_t\right]-\mathbb{E}^{\mathbb{P}}\left[\int_t^T v_u\,du\middle|\mathcal{F}_t\right].$$

  - How much we expect to make selling a variance contract.

Variance risk premium under log utility

- With the change of measure

$$\frac{d\mathbb{P}}{d\mathbb{Q}} = \frac{\rho_{\mathbb{P}}(S_T)}{\rho_{\mathbb{Q}}(S_T)} = \frac{S_T}{S_t},$$

computation of the variance risk premium is straightforward.

- The fair value of the variance contract is given by

$$\mathcal{V}_t(T) := \mathbb{E}^{\mathbb{Q}}\left[\int_t^T v_u\,du\middle|\mathcal{F}_t\right] = -2\,\mathbb{E}^{\mathbb{Q}}\left[X_{t,T}\middle|\mathcal{F}_t\right].$$

- On the other hand, expected integrated variance under $\mathbb{P}$ is given by

$$\mathbb{E}^{\mathbb{P}}\left[\int_t^T v_u\,du\middle|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}}\left[\int_t^T \frac{S_u}{S_t}\,v_u\,du\middle|\mathcal{F}_t\right] =: \mathcal{G}_t(T).$$

- Thus the variance risk premium is given by

$$\begin{aligned} & \mathbb{E}^{\mathbb{Q}}\left[\int_t^T v_u\,du \middle| \mathcal{F}_t\right] - \mathbb{E}^{\mathbb{P}}\left[\int_t^T v_u\,du \middle| \mathcal{F}_t\right] \\ = \; & \mathcal{V}_t(T) - \mathcal{G}_t(T) =: -\mathcal{L}_t(T), \end{aligned}$$

  where $\mathcal{L}_t(T)$ denotes the value of the leverage contract.
- Both $\mathcal{V}_t(T)$ and $\mathcal{G}_t(T)$ have model-free prices in terms of the log- and entropy-contracts respectively.
  - Thus, under log-utility, we have model-free expressions for both equity risk and variance risk premia!
- What is the empirical variance risk premium?

Variance risk premium

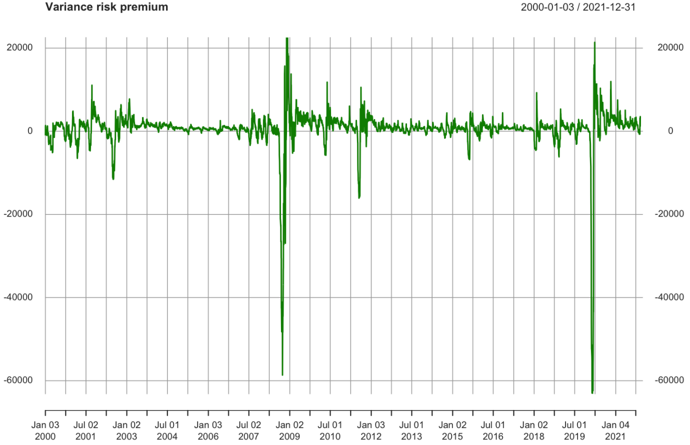

Figure 4: Realized variance risk premium from 2000 to 2021 per $1mm notional.

Variance and gamma contract estimates use data from OptionMetrics.

Cumulative P&L

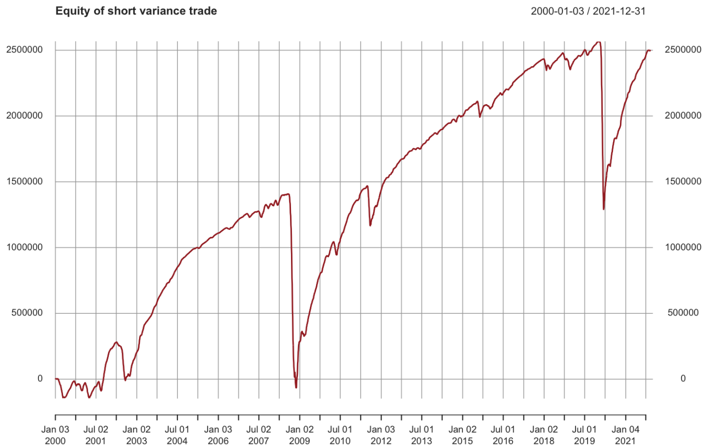

Figure 5: Cumulative P&L of the short variance trade.

Profitability of the short variance trade

- Figure 5 clearly shows that the short variance trade is profitable on average.
  - But with huge drawdowns.
- Assuming we traded a 1-month \$1mm variance contract contract every day, our average P&L would have been $454 \pm 5,283$ dollars per day.
- The equivalent average risk premium according to our formula

$$\mathbb{E}^{\mathbb{Q}}\left[\int_t^T v_u\,du \middle| \mathcal{F}_t\right] - \mathbb{E}^{\mathbb{P}}\left[\int_t^T v_u\,du \middle| \mathcal{F}_t\right] = -\mathcal{L}_t(T)$$

was 228 dollars per day.

Regressions

- Let V denote the price of the 30-day variance contract and RV denote the subsequently realized 30-day variance.
- Carr and Wu estimated the following relations:

$$\begin{aligned} RV &= a + b\,V + \epsilon \\ \log RV &= a + b \log RV + \epsilon. \end{aligned}$$

- Performing linear regressions on our dataset gives

$$\begin{aligned} RV &= -0.0001635 + 0.9234\,V + \epsilon \\ \log RV &= -0.04774 + 1.07221 \log V + \epsilon. \end{aligned}$$

  The first regression has an R^2 of 0.37 and the second has an R^2 of 0.56.

  - In other words, the variance contract is a good predictor of realized variance – see Figure 6.

VIX^2 vs actual variance

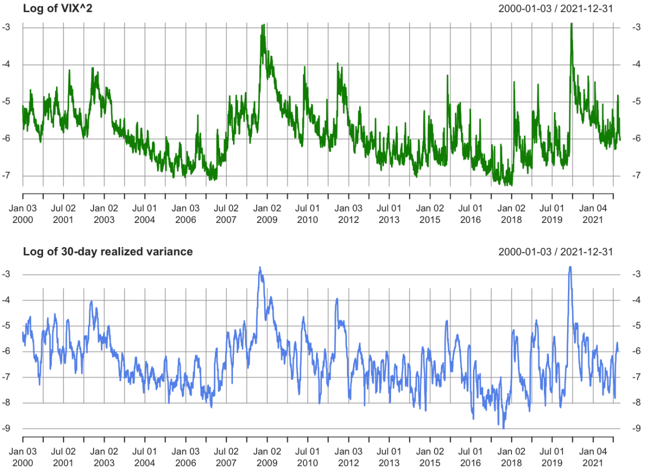

Figure 6: VIX^2 is a good predictor of realized variance.

Time series of the leverage contract

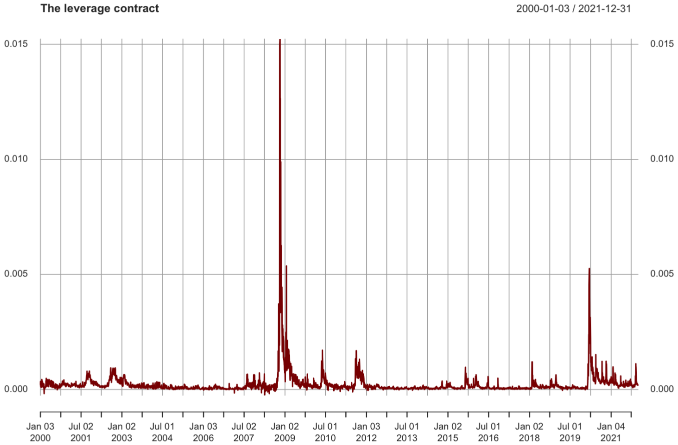

Figure 7: The leverage contract is very spiky!

Variance risk premium and the leverage contract

- Recall that in our simple framework, the variance risk premium is given by the leverage contract.
- We (probably) can't predict spikes in volatility, but we can measure leverage!
  - Why not sell variance contracts only after a spike in leverage?
  - In other words, sell only when the variance risk premium is high.
- For example, suppose we only sell variance on days when the leverage contract is greater than 0.0015. What does the equity plot look like?

Time series of the leverage contract with sell barrier

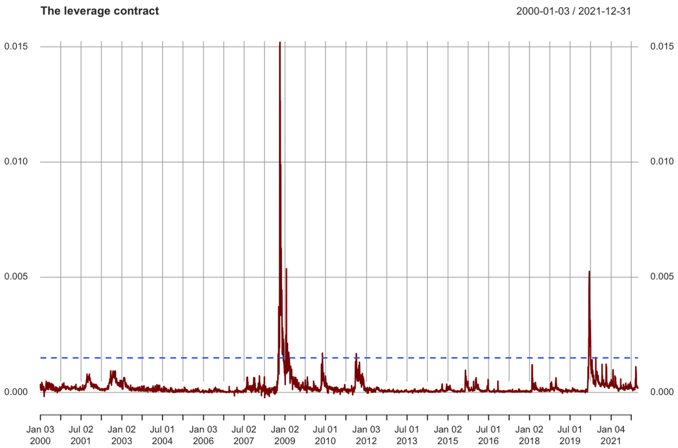

Figure 8: The leverage contract with a sell level at 0.0015.

Cumulative P&L of the conditional short variance trade

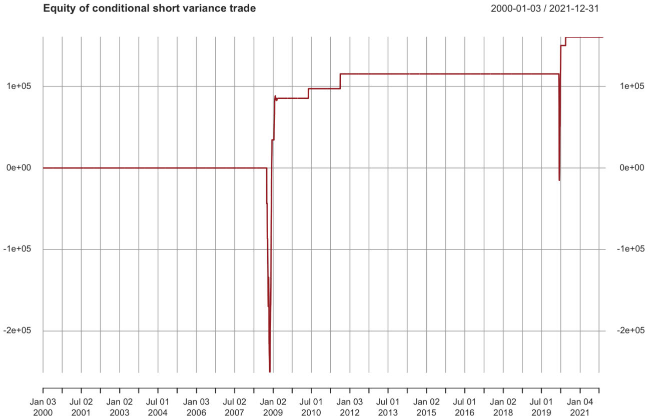

Figure 9: Cumulative P&L of the conditional short variance trade. Back to the drawing board!

Summary

- Using the Carr-Madan spanning formula, we computed various model-free quantities including
  - The variance contract.
  - The gamma contract.
- Under log utility, the Carr-Madan optimal portfolio is the market portfolio, which we take to be SPX.
  - The change of measure $d\mathbb{P}/d\mathbb{Q} = S_T/S_t$.
  - The equity risk premium is just instantaneous variance.
- It follows that the variance risk premium of Carr and Wu is given by the leverage swap.

Epilog

- Be curious.
- Look for beauty.
- Work very hard.

- ... be like Peter Carr.

References

Peter Carr and Dilip Madan
Towards a theory of volatility trading.
In R. Jarrow, editor, *Volatility: New estimation techniques for pricing derivatives*, Risk Publications, 417–427, 1998.

Peter Carr and Dilip Madan
Optimal positioning in derivative securities.
Quantitative Finance, 1(1):19–37, 2001.

Peter Carr and Liuren Wu.
Variance risk premiums.
Review of Financial Studies, 22(3):1311–1341, 2009.

Peter Carr and Jiming Yu.
Risk, return, and Ross recovery.
The Journal of Derivatives, 20(1):38–59, 2012.

Masaaki Fukasawa.
The normalizing transformation of the implied volatility smile.
Mathematical Finance, 22(4):753–762, 2012.

Masaaki Fukasawa.
Volatility derivatives and model-free implied leverage.
International Journal of Theoretical and Applied Finance, 17(1):1450002, 2014.

Anders Merrild Posselt.
Dynamics in the VIX complex.
Journal of Futures Markets, online, 2021.