Dynamic Relative Valuation

Liuren Wu, Baruch College

Joint work with *Peter Carr* from Morgan Stanley

Derivatives 2013: 40 Years after the Black-Scholes-Merton Model
October 11, 2013

The standard approach of no-arbitrage pricing

- Identify an "*instantaneous*" rate as the basis, *fully specify* the dynamics and pricing of this basis, and *represent* the value of actual contracts as some expectation of this basis over some periods.

    - **Options**: Choose the *instantaneous variance rate* as the basis, and fully specify its dynamics/pricing.

    - **Interest rates**: Choose the *instantaneous interest rate* as the basis, and fully specify its dynamics/pricing.

- The challenge: One needs to think far far into the future.

    - Under this approach, to price a 60-year option or bond, one needs to specify how the instantaneous rate moves and how the market prices its risk over the next 60 years.

- The market has a much better idea about how the prices of *many traded securities* move over the *next day*.

A new approach: Work with what the market knows better

Model what the market knows better — the *next move* of *many traded securities*.

1. Find commonalities of traded securities via some transformation:

$$B(y_t(C), t, X_t; C) = \text{price}$$

   - $y_t(C)$ — *transformation* of security prices that are similar and move together across securities with different contract details (C)
   - X_t — some dependence on other observables (such as the underlying security price for derivatives)
   - The market is reasonably good at finding this transformation

2. Specify how the transformations move together

$$dy_t(C) = \mu_t(C) + \omega_t(C)dZ_t$$

   - One Borwnian motion drives all $y_t(C)$
   - The levels of $(\mu_t(C), \omega_t(C))$ are known, but not their dynamics.

3. The new no-arbitrage pricing relation: What we assume about their *next move* dictates how their values ($y_t(C)$) compare right now.

I. Options market: Model BMS implied volatility

- Transform the option price via the Black-Merton-Scholes (BMS) formula

$$B(I_t(K, T), t, S_t; K, T) \equiv S_t N(d_1) - KN(d_2)$$

  - Zero rates for notional clarity.
  - BMS implied vol $I_t(K, T)$ is the commonality of the options market.
  - The implied volatilities of different option contracts (underlying the same security) share similar magnitudes and move together.
  - A positive quote excludes arbitrage against cash and the underlying.
  - It is the industry standard for quoting/managing options.
- Diffusion stock price dynamics: $dS_t/S_t = \sqrt{v_t}dW_t$.
- Leave the dynamics of v_t *unspecified*; instead, model the *near-future* dynamics of the BMS implied volatility under the risk-neutral ($\mathbb{Q}$) measure,

$$dI_t(K, T) = \mu_t dt + \omega_t dZ_t, \text{ for all } K > 0 \text{ and } T > t.$$

  - The drift (μ_t) and volvol (ω_t) *processes* can depend on K, T, and I_t.
  - Correlation between implied volatility and return is $\rho_t dt = \mathbb{E}[dW_t dZ_t]$.

From NDA to the fundamental PDE

We require that no dynamic arbitrage (NDA) be allowed between any option at (K, T) and a basis option at (K_0, T_0) and the stock.

- The three assets can be combined to neutralize exposure to dW or dZ.

- By Ito's lemma, each option in this portfolio has risk-neutral drift given by:

$$B_t + \mu_t B_\sigma + \frac{v_t}{2} S_t^2 B_{SS} + \rho_t \omega_t \sqrt{v_t} S_t B_{S\sigma} + \frac{\omega_t^2}{2} B_{\sigma\sigma}.$$

- No arbitrage and no rates imply that both option drifts must vanish, leading to the fundamental "PDE:"

$$-B_t = \mu_t B_\sigma + \frac{v_t}{2} S_t^2 B_{SS} + \rho_t \omega_t \sqrt{v_t} S_t B_{S\sigma} + \frac{\omega_t^2}{2} B_{\sigma\sigma}.$$

- When μ_t and ω_t are independent of (K, T), the "PDE" defines a linear relation between the theta (B_t) of the option and its vega (B_σ), dollar gamma ($S_t^2 B_{SS}$), dollar vanna ($S_t B_{S\sigma}$), and volga ($B_{\sigma\sigma}$).

- We call the class of BMS implied volatility surfaces defined by the above fundamental PDE as the **Vega-Gamma-Vanna-Volga (VGVV)** model.

$$-B_t = \mu_t B_\sigma + \frac{v_t}{2} S_t^2 B_{SS} + \rho_t \omega_t \sqrt{v_t} S_t B_{S\sigma} + \frac{\omega_t^2}{2} B_{\sigma\sigma}.$$

- Traditionally, PDE is specified to solve the value function. In our case, the value function $B(I_t, t, S_t)$ is definitional and it is simply the BMS formula that we use the transform the option price into implied volatility.

- The coefficients on traditional PDEs are deterministic; they are *stochastic processes* in our "PDE."

- Our "PDE" is not derived to solve the value function, but rather it is used to show that the various stochastic quantities have to satisfy this particular relation to exclude NDA.

  ⇒ Our "PDE" defines an *NDA constraint* on how the different stochastic quantities should relate to each other.

From the "PDE" to an algebraic restriction

$$-B_t = \mu_t B_\sigma + \frac{v_t}{2} S_t^2 B_{SS} + \rho_t \omega_t \sqrt{v_t} S_t B_{S\sigma} + \frac{\omega_t^2}{2} B_{\sigma\sigma}.$$

- The value function B is well known, so are its various partial derivatives:

$$
\begin{array}{rclcrcl}
B_t & = & -\frac{\sigma^2}{2} S^2 B_{SS}, & & B_\sigma & = & \sigma\tau S^2 B_{SS}, \\
S B_{\sigma S} & = & -d_2 \sqrt{\tau} S^2 B_{SS}, & & B_{\sigma\sigma} & = & d_1 d_2 \tau S^2 B_{SS},
\end{array}
$$

where dollar gamma is the common denominator of all the partial derivatives, a nice feature of the BMS formula as a transformation.

- The "PDE" constraint on B is reduced to an algebraic restriction on the shape of the implied volatility surface $I_t(K, T)$,

$$\frac{I_t^2}{2} - \mu_t I_t \tau - \left[\frac{v_t}{2} - \rho_t \omega_t \sqrt{v_t \tau} d_2 + \frac{\omega_t^2}{2} d_1 d_2 \tau \right] = 0.$$

  - If (μ_t, ω_t) do not depend on $I_t(K, T)$, we can solve the whole implied volatility surface as the solution to a quadratic equation.
  - We just need to know the current values of the *processes* $(\mu_t, \omega_t, \rho_t, v_t)$, which dictate the *next move* of the implied volatility surface, but we do not need to specify their future dynamics.

Proportional volatility dynamics, as an example

- (μ_t, ω_t) are just generic representations of the drift and diffusion processes of the implied volatility, we can be more specific if we think we know more.
- As an example,

$$dI_t(K, T)/I_t(K, T) = e^{-\eta_t(T-t)}\left(m_t dt + w_t dZ_t\right),$$

with $\eta_t, w_t > 0$ and (η_t, m_t, w_t) independent of K, T.

  - A proportional specification has more empirical support than a square-root variance specification.
  - The exponential dampening makes long-term implied volatility less volatile and more persistent.
  - (m_t, w_t) just specify our views on the trend and uncertainty over the next instant.

- We can re-cast the implied volatility surface in terms of log relative strike and time to maturity, $I_t(k, \tau) \equiv I_t(K, T)$, with $k = \ln K/S_t$ and $\tau = T - t$.
- The implied variance surface $(I_t^2(k, \tau))$ solves a *quadratic equation*:

$$
\begin{aligned}
0 = \; & \tfrac{1}{4}e^{-2\eta_t\tau}w_t^2\tau^2 I_t^4(k, \tau) + \left(1 - 2e^{-\eta_t\tau}m_t\tau - e^{-\eta_t\tau}w_t\rho_t\sqrt{v_t}\tau\right)I_t^2(k, \tau) \\
& - \left(v_t + 2e^{-\eta_t\tau}w_t\rho_t\sqrt{v_t}k + e^{-2\eta_t\tau}w_t^2 k^2\right).
\end{aligned}
$$

Unspanned dynamics

$$
\begin{aligned}
0 \;=\; & \tfrac{1}{4} e^{-2\eta_t \tau} w_t^2 \tau^2 I_t^4(k,\tau) + \left(1 - 2 e^{-\eta_t \tau} m_t \tau - e^{-\eta_t \tau} w_t \rho_t \sqrt{v_t} \tau \right) I_t^2(k,\tau) \\
& - \left(v_t + 2 e^{-\eta_t \tau} w_t \rho_t \sqrt{v_t} k + e^{-2\eta_t \tau} w_t^2 k^2 \right).
\end{aligned}
$$

- Given the current levels of the five stochastic processes $(\rho_t, v_t, m_t, \eta_t, w_t)$, the current shape of the implied volatility surface must satisfy the above quadratic equation to exclude dynamic arbitrage.
- The current shape of the surface does NOT depend on the exact dynamics of five stochastic processes.
  - The dynamics of the five processes are not spanned by the current shape of the implied volatility surface.
  - The current volatility surface shape is only linked to the *near-future* dynamics of the security return and implied volatility.
- The dynamics of the five processes will affect the future dynamics of the surface, but not its current shape.
- The current shape of the implied volatility surface is determined by *5 state variables*, but with *no parameters*!

The implied volatility smile and the term structure

- At a fixed maturity, the implied variance smile can be solved as

$$I^2(k, \tau) = a_t + \frac{2}{\tau} \sqrt{\left(k + \frac{\rho \sqrt{v_t}}{e^{-\eta_t \tau} w_t} \right)^2 + c_t}.$$

  - In the limit of $\tau = 0$, $I_t^2(k, 0) = v_t + 2\rho_t \sqrt{v_t} wk + w^2 k^2$.
  - The smile is driven by vol of vol (convexity), and the skew is driven by the return-vol correlation.

- At $d_2 = 0$, the at-the-money implied variance term structure is given by,

$$A_t^2(\tau) = \frac{v_t}{1 - 2e^{-\eta_t \tau} m_t \tau}.$$

  - The slope of the term structure is dictated by the drift of the dynamics.

- The Heston (1993) model generates the implied vol surface as a function of 1 state variable (v_t) and 4 parameters ($\kappa, \theta, \omega, \rho$). Performing daily calibration on Heston would result in the same degrees of freedom, but the calculation is much more complicated and the process is inconsistent.

Sequential multi-layer valuation

- Suppose you have generated valuations on these option contracts based on either the "instantaneous rate" approach or some other method (including the earlier example of our VGVV models).

- One way to gauge the virtue of a valuation model is to assess how fast market prices converge to the model value when the two deviate.

- One can specify a *co-integrating relation* between the market price and the model value,

$$dI_t(K, T) = \kappa_t \left(V_t(K, T) - I(K, T) \right) dt + w_t I(K, T) dZ_t,$$

  where $V(K, T)$ denotes the model valuation for the option contract, represented in the implied vol space.

    - κ_t measures how fast the market converges to the model value.

- A similar quadratic relation holds: The valuation $V_t(K, T)$ is regarded as a number that the market price converges to, regardless of its source.

    - Our approach can be integrated with the traditional approach by providing a second layer of valuation on top of the standard valuation.

    - It can also be used to perform sequential *self-improving* valuations!

Applications

1. Implied volatility surface interpolation and extrapolation:
   - Much faster, much easier, much more numerically stable than any existing stochastic volatility models.
   - Performs better than lower-dimensional stochastic volatility models, without the traditional parameter identification issues associated with high-dimensional models.
   - Reaps the benefits of both worlds: First estimate a low-dimensional stochastic vol model, with which one can perform long-run simulations, and add the layers of VGVV structure on top of it.

2. Extracting variance risks and risk premiums:
   - The unspanned nature allows us to extract the levels of the 5 states from the implied volatility surface without fully specifying the dynamics.
   - One can also extract variance risk premiums by estimating an analogous contract-specific *option expected volatility* surface ...

Extensions

1. Non-zero *financing rates*: Treat S and B as the forward values of the underlying and the option, respectively.

2. Options on single names with *potential defaults*, upon which stock price drops to zero (Merton (1976))

   - The BMS implied volatility transformation is no longer as tractable:

   $$0 = B_t + \mu_t B_\sigma + \frac{v_t}{2} S_t^2 B_{SS} + \rho_t \omega_t \sqrt{v_t} S_t B_{S\sigma} + \frac{\omega_t^2}{2} B_{\sigma\sigma} + \lambda K N(d_2).$$

   The last term induced by default does not cancel.

   - We can choose the transformation through the Merton (1976) model: Let M_t (Merton implied volatility) as the σ input in the Merton model to match the market price, conditional on λ being known:

   $$B(S_t, t, M_t) = S_t N(d_1) - e^{-\lambda \tau} K N(d_2), \quad d_{1/2} = \frac{\ln S_t / K + \lambda \tau \pm \frac{1}{2} M_t^2 \tau}{M_t \sqrt{\tau}}.$$

   $B_t = -\frac{\sigma^2}{2} S^2 B_{SS} - \lambda e^{-\lambda \tau} K N(d_2)$ has an extra term that cancels out the default-induced term

3. It is a matter of choosing the right transformation...

Contract versus instantaneous specification

- Traditional modeling strives to identify an instantaneous rate as the basis for all traded securities.

- The consistency is naturally guaranteed relative to this basis if all securities are priced from it.

- The challenge is to fully specify the instantaneous rate dynamics that can reasonably price short-term as well as long-term contracts, with reasonable tractability and stability in parameter identification.

- Our approach directly models contract-specific quantities (such as the implied volatility of an option contract) and link the contract values together without going through the basis dynamics.

    - We can focus on what we know better: the near-term moves of market contracts instead of the dynamics of some unobservable instantaneous rate over the next few decades.

    - The derived no-arbitrage relation shows up in extremely simple terms.

    - No parameters need to be estimated, only states need to be extracted.

II. Bonds market: Model YTM

- For bond pricing, standard approach starts with the full dynamics and pricing of the instantaneous interest rate.

- We directly model the *yield to maturity* (y_t) of each bond as the commonality transformation of the bond price,

$$B(t, y_t) \equiv \sum_j C_j e^{y_t \tau_j}$$

where C_j denotes the cash flow (coupon or principal) at time $t + \tau_j$.

- Assume the following risk-neutral dynamics for the yields of bond m,

$$dy_t^m = \mu_t^m dt + \sigma_t^m dW_t$$

where we use the superscript m to denote the potentially bond-specific nature of the yield and its dynamics.

- Again, (μ_t, σ_t) is just a generic representation, one can be more specific in terms of which process is global and which is contract-specific...

Dynamic no-arbitrage constraints on the *yield* curve

- Assume dynamic no-arbitrage between bonds and the money market account, we obtain the following PDE,

$$r_t B = B_t + B_y \mu_t + \frac{1}{2} B_{yy} \sigma_t^2,$$

- With the solution to $B(t, y)$ known, the PDE can be reduced to a simple algebraic relation on yield:

$$y_t^m - r_t = \mu_t^m \tau_m - \frac{1}{2} (\sigma_t^m)^2 \tau_m^2,$$

where τ_m and τ_m^2 denote cash-flow weighted average maturity and maturity squared, respectively,

$$\tau_m = \sum_j \frac{C_j e^{-y_t \tau_j}}{B^m} \tau_j, \quad \tau_m^2 = \sum_j \frac{C_j e^{-y_t \tau_j}}{B^m} \tau_j^2.$$

- The relation is simple and intuitive: The yield curve goes up with its risk-neutral drift, but comes down due to convexity.

Applications

1. Decomposing the yield curve:

   - Assume a mean-reverting square-root dynamics on the YTM:

     $$dy_t^m = \kappa_t(\theta_t - y_t^m)dt + \sigma_t\sqrt{y_t^m}dW_t$$

   $\Rightarrow$ The no-arbitrage yield curve

     $$y_t^m = \frac{r_t + \kappa_t\theta_t\tau_m}{1 + \kappa_t\tau_m + \frac{1}{2}\sigma_t^2\tau_m^2}$$

   - The yield curve starts at r_t at $\tau_m = 0$ and moves toward its risk-neutral target θ_t as average maturity increases, subject to a convexity effect that drives the yield curve downward in the very long run.

   - The target θ_t is a combination of expectation and risk premium.

   - No fixed parameters to be estimated, only the current states of some stochastic processes to be extracted.

Applications

1. Decomposing the yield curve:

2. Extracting risk premium from long bonds:
   - Assuming long-dated bonds move like a random walk under $\mathbb{P}$, the risk-neutral dynamics become

$$dy_t^m = \gamma_t \sigma_t^m dt + \sigma_t^m dW_t,$$

   with γ_t denoting the market price of bond risk.

   - The no-arbitrage relation is:

$$r_t = y_t^m - \gamma_t \sigma_t^m \tau_m + \frac{1}{2}(\sigma_t^m)^2 \tau_m^2,$$

   - Based on the observed yield-to-maturity of each bond and a volatility estimator for each bond, we can infer the market price of bond risk as

$$\hat{\gamma}_t^m = \frac{y_t^m - r_t + \frac{1}{2}(\sigma_t^m)^2 \tau_m^2}{\sigma_t^m \tau_m}$$

   - To the extent that the market price estimator $\hat{\gamma}_t^m$ differs across different bonds, the difference represents relative value opportunities.

III. Defaultable bonds

- For corporate bonds, directly modeling yield to maturity should work the same as for default-free bonds.

$$r_t B = B_t + B_y \mu_t + \frac{1}{2} B_{yy} \sigma_t^2 + (R - 1) B \lambda.$$

- The PDE can be reduced to an algebraic equation that defines the defaultable yield term structure:

$$y_t^m - (r + (1 - R)\lambda) = \mu_t^m \tau_m - \frac{1}{2} (\sigma_t^m)^2 \tau_m^2.$$

- One can also directly model the credit spread if one is willing to assume deterministic interest rates...

Conclusion: Choose the right representation

At the end of the day, we are choosing some transformation of the prices of traded securities to obtain a better understanding of where the valuation comes from.

- Academics often try to go bottom up in search for the "building blocks" that can be used to value anything and everything.

    - Arrow-Debrew securities, state prices, instantaneous rates...

- Practitioners tend to be more defensive-minded, often trying to reduce the dimensionality of the problem via localization and facilitate management/monitoring via stabilization

    - Most transformations, such as implied volatility, YTM, CDS are contract-specific — very localized.

    - They also stabilize the movements, standardize the value for cross-contract comparison, and preclude some arbitrage possibilities.

    - We start with these contract-specific transformations, and derive cross-contract linkages by assuming how these transformed quantities move together in the next instant, but nothing further.

Summary: Nice features and future efforts

- Nice features
  - We obtain no-arbitrage cross-contract constraints while only specifying what we know better.
  - The relations we obtain are extremely simple.
  - Empirical work only involves state extraction, no parameter estimation.
  - If you have built well-performing models with full dynamics, we can embed them into our framework by letting market prices reverting to the model value.
  - We can also build several layers of models sequentially that are consistent with each other, and satisfy the conflicting needs of different groups: Market makers v. prop. traders targeting different horizons.
- Future efforts:
  - Theoretical efforts: General representation across markets, multiple source of shocks, non-diffusion shocks.
  - Empirical work: lots to be done.