Lecture notes on risk management, public policy, and the financial system

Extreme market risk events

Allan M. Malz

Columbia University

Summary of the standard model

- The standard model predicts return behavior in several dimensions
  Distribution of returns: log asset returns normally distributed
  - ⇒Very large returns quite rare

  Behavior of return distribution over time: distributional family and parameters constant
  Return predictability: future returns cannot be forecast based on past returns
  - ⇒**Autocorrelations** of returns close to zero

  Cross-asset correlations are constant
  - CAPM further predicts that cross-asset correlations explained by beta to market portfolio
- Standard model actually a "baseline" rather than "accepted" model
  - Much effort devoted to identifying, modeling recognized departures from model
- Reality: most of the above not actually the case
  - Departures often summarized as "stylized facts of asset returns"

Standard model and reality

Non-normal distributions: asset returns far from normally distributed

- Very large returns rare, but not extremely so
- Negative extreme returns predominate for many assets
- **Tail risk:** loss arising from extreme events in the "tails"

Time-varying distributions: returns very noisy at times, docile at others

- Value-at-Risk typically assumes conditional volatility, parameters or quantiles reestimated periodically

Lack of return predictability: autocorrelations of daily and lower-frequency returns in fact close to zero

- High-frequency (intraday) returns display some predictability
- Some evidence of long-term equity market predictability via **priced factors**

Cross-asset correlations vary over time, and with other distributional characteristics

S&P 500 returns 1927-2020

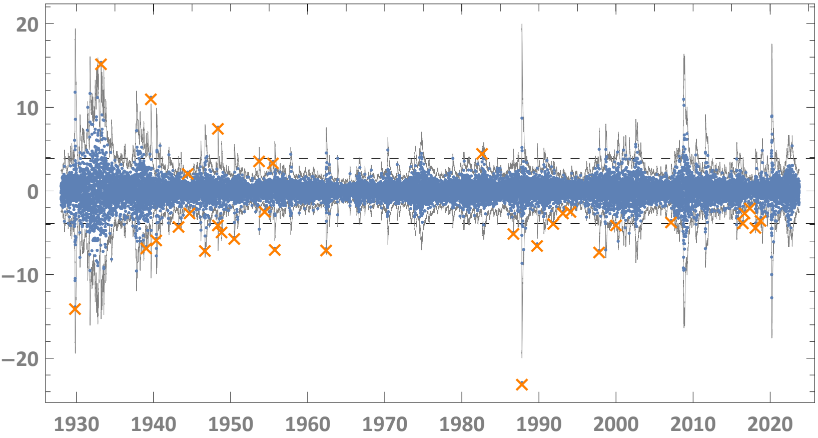

Points: daily realized returns (bps), solid lines EWMA 99.9% confidence interval (±3.3 vols), horizontal grid lines 99.9% CI using standard deviation of entire series. ×: returns outside the EWMA 99.9999% confidence interval (one-in-a-million chance). *Data source*: Bloomberg Financial L.P.

10-year Treasury Note yield fluctuations 1962-2020

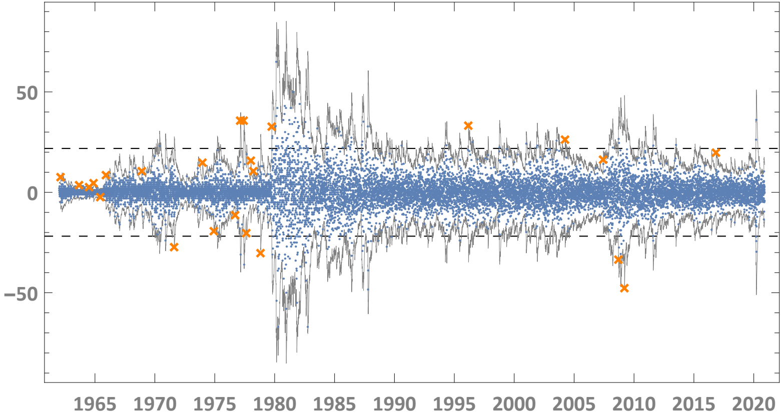

Points: daily realized change in yield (bps), solid lines EWMA 99.9% confidence interval ($\pm$3.3 vols), horizontal grid lines 99.9% CI using standard deviation of entire series. $\times$: returns outside the EWMA 99.9999% confidence interval (one-in-a-million chance). *Data source*: Bloomberg Financial L.P.

Characteristics of non-normal distributions

Typical departures from the normal distribution

Skewness: large-magnitude returns of one sign more frequent than the other

- **Skewness coefficient**: third standardized moment
- Normal distribution is symmetrical, skewness coefficient zero
- Mean higher than median of a distribution skewed to right
- Common in asset and portfolio returns
  - Returns bounded above (but not below)→negative skewness
  - Returns bounded below (but not above)→positive skewness

Kurtosis, "fat tails," or **leptokurtosis:** large returns more frequent than in normal

- **Kurtosis coefficient:** fourth standardized moment;
- Kurtosis of the normal distribution precisely 3
- →**kurtosis excess** defined as kurtosis coefficient minus 3.

Multimodal distributions: skewed and kurtotic distribution may also exhibit multiple modes or peaks

Moments of non-normal distributions

Statistical **moments** measure the shape of a distribution

Coefficient of skewness: ratio of the 3rd central moment to the $\frac{3}{2}$ power of the 2nd central moment

$$\frac{T^{-1}\sum_{t=1}^{T}(r_t-\bar{r})^3}{\left[T^{-1}\sum_{t=1}^{T}(r_t-\bar{r})^2\right]^{\frac{3}{2}}}=\frac{\sqrt{T}\sum_{t=1}^{T}(r_t-\bar{r})^3}{\left[\sum_{t=1}^{T}(r_t-\bar{r})^2\right]^{\frac{3}{2}}}$$

Coefficient of kurtosis: ratio of the 4th central moment to the square of the 2nd central moment

$$\frac{T^{-1}\sum_{t=1}^{T}(r_t-\bar{r})^4}{\left[T^{-1}\sum_{t=1}^{T}(r_t-\bar{r})^2\right]^{2}}=\frac{T\sum_{t=1}^{T}(r_t-\bar{r})^4}{\left[\sum_{t=1}^{T}(r_t-\bar{r})^2\right]^{2}}$$

- The excess kurtosis of a data set or distribution is its coefficient of kurtosis minus 3.0, the kurtosis of the normal distribution

Both use the number of observations rather than degrees of freedom

Normal and non-normal distributions

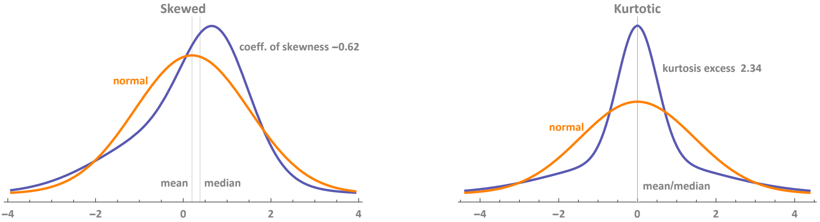

Left: skewed and normal distributions with identical means of 0.2 and variance 1.70875. The skewed distribution is a mixture of two normals, with distributions $\mathcal{N}(-0.35, 1.5)$ and $\mathcal{N}(0.75, 0.75)$, each with a probability of 50 percent of being realized. Right: kurtotic and normal distributions, both zero mean and with identical variance of 2.125. The kurtotic distribution is a mixture of two normals, with distributions $\mathcal{N}(0, 2.0)$ and $\mathcal{N}(0, 0.5)$, each with a probability of 50 percent of being realized.

Volatility asymmetry

- **Volatility asymmetry** or **feedback:** returns negatively correlated with increases in volatility
- Prominent feature of equity returns, leading to hypotheses regarding causes and corresponding terminology
  - Often called **leverage effect:** hypothesis that negative returns increase leverage, thus raising volatility
  - Alternatively, **volatility feedback:** hypothesis that higher volatility requires higher expected return, hence decline in price
- Can impact valuation and convexity
  - **Example**: reduces convertible bond convexity near conversion price
- **Conditional skewness**: degree of skewness in return distribution varies with current volatility

Volatility asymmetry in the U.S. stock market

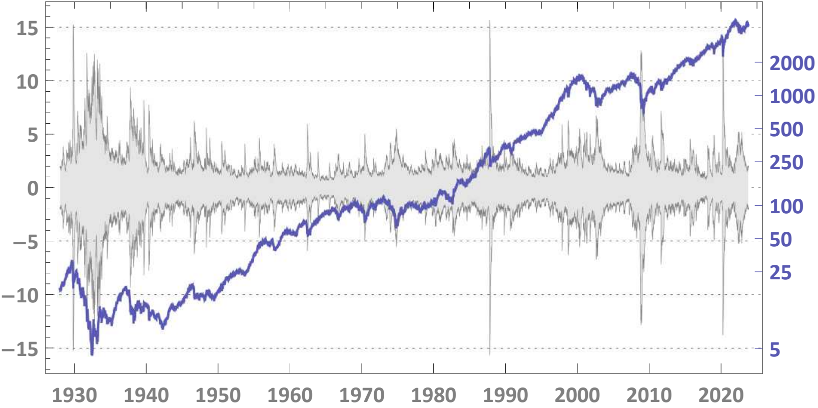

99.9% confidence interval using EWMA estimate of S&P 500 index return volatility, decay factor 0.94 (gray plot, left y-axis), and log of S&P 500 index closing price (purple plot, right y-axis), both daily 30Dec1927 to 11May2016. *Data source*: Bloomberg L.P.

Behavior of volatility across assets

Negative coskewness: sharp negative return surprises in some individual stocks coincide with very bad S&P 500 days

- Makes those stocks less effective diversifiers, but earn higher excess returns
- Can be rationalized by household aversion to skewness in portfolio returns→consistent with consumption-based approach
- Affects cross-section of returns→inconsistent with CAPM market beta-only model

Tail dependence: high probability of extreme return in one asset associated with extreme in others

- May be hard to distinguish from effect of high time-varying volatility on sample correlation
- Related to coskewness through **asymmetric correlations**: cross-sectional correlations increase in bear markets

Examining data

- History tells much but not everything
  - History especially problematic if short return time series
  - E.g. house prices, no observations of sustained “negative HPA”
- Visual tools
  - Return plots
  - Quantile-quantile plot
  - Kernel densities
- Outlier analysis

Extreme moves in the S&P 500 1927-2020

Confidence level (%)	99.9	99.99	99.9999	99.9999999
Odds: one in...	1,000	10,000	1,000,000	1,000,000,000
In the normal distribution				
Number of std dev	$\pm$ 3.291	$\pm$ 3.891	$\pm$ 4.892	$\pm$ 6.109
Expected no. exceedances	23.3	2.3	0.0	0.0
Unconditional volatility				
No. exceedances	326	199	106	49
No. negative	177	111	58	24
No. positive	149	88	48	25
Ratio negative/positive	1.2	1.3	1.2	1.0
Conditional (EWMA) volatility				
No. exceedances	178	92	33	12
No. negative	136	74	26	10
No. positive	42	18	7	2
Ratio negative/positive	3.2	4.1	3.7	5.0

Based on daily S&P 500 returns 1927-2012

Quantile-quantile plot

- A.k.a. **quantile** or **q-q plot**
- Used to visually compare
  - Distributions of two empirical data sets or theoretical distributions
  - Distribution of a data set to a theoretical distribution, generally the normal
- Each point in the plot is the pair consisting of corresponding quantiles of the two distributions
- Differences in the distributions are highlighting by drawing a straight line between points corresponding to 1st and 3rd quartiles
- Contrasts to comparison distribution reflected in shape and location
- If comparing to normal distribution,
  - Kurtosis leads to S shape, i.e. q-q plot below (above) straight line for low (high) quantiles
  - Skewness leads to asymmetric shape
  - Different means leads to position off straight line

Quantile plot of the S&P 500

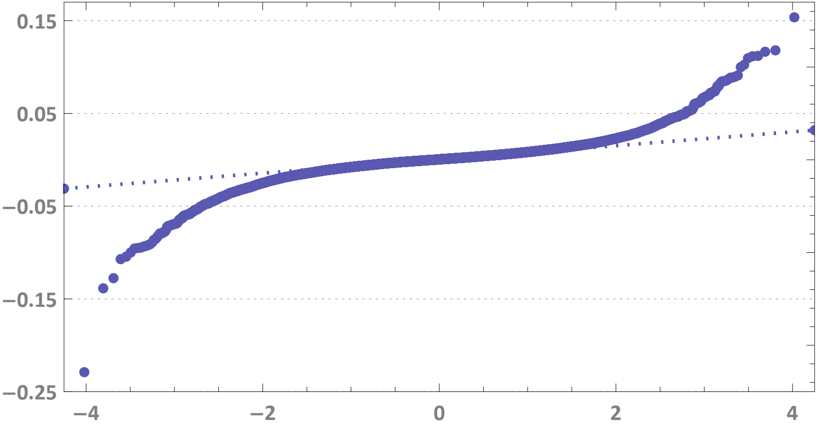

Quantiles of daily returns (as decimals) of the S&P 500 index, 1927-2011, plotted against quantiles of a standard normal distribution. *Data source*: Bloomberg Financial L.P.

Kernel density estimate

- Estimate of probability density function
- *Nonparametric*: no distributional hypothesis
- Rather data-focused, smooth the histogram
  - Tradeoff between smoothness and adherence to the data
- Visualize *shape* of distribution, facilitate comparison
- Close-up of tails

Kernel density of the S&P 500

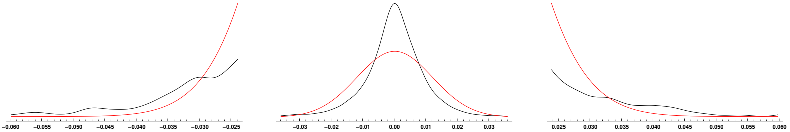

Kernel estimate of probability density function of daily returns of the S&P 500 index, 1927-2011. Gaussian kernel using optimal bandwidth. *Data source*: Bloomberg Financial L.P.

Approaches to the problem

- Identify alternative **stochastic process** to geometric Brownian motion
- Alternatively, seek valid generalizations about behavior of extreme returns→**extreme value theory**

Alternative stochastic processes

- Very wide variety of specifications: older option pricing literature
  - Stochastic volatility: focus on volatility regimes
  - Ornstein-Uhlenbeck: mean reversion
- Jump diffusion
  - Add Poisson-distributed jumps to geometric Brownian motion
  - Currencies, individual stocks

Jump diffusion and normal returns

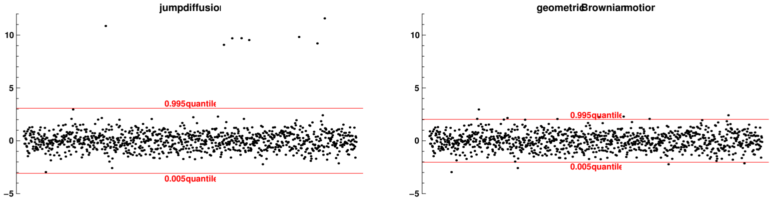

Simulation of 1000 daily sequential steps of a jump diffusion (left panel) and geometric Brownian motion process. The y-axis displays returns in percent. The horizontal grid lines mark the 99 percent confidence intervals. The jump diffusion follows

$$dS_t = \left(\mu + \frac{1}{2}\sigma^2 - \lambda \mathbf{E}\left[k_t\right]\right) S_t dt + \sigma S_t dW_t + k_t S_t dq_t$$

with

$$dq_t = \begin{Bmatrix} 1 \\ 0 \end{Bmatrix} \quad \text{with probability} \quad \begin{Bmatrix} \lambda dt \\ 1 - \lambda dt \end{Bmatrix}$$

Jumps are Poisson distributed with frequency parameter $\lambda = \frac{1}{252}$ and a nonstochastic jump size $k = 0.10$ (10 percent). For both the jump diffusion and normal processes, $\mu = 0$ and $\sigma = 0.12$ (12 percent) at an annual rate.

Example of a jumpy stock

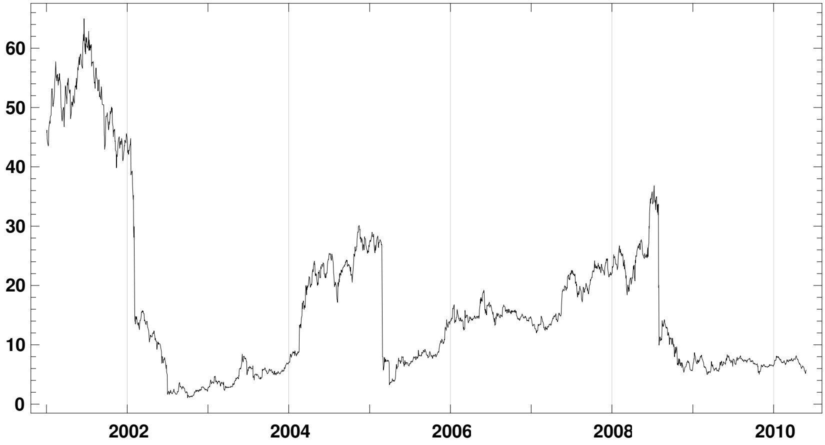

Price of Elan Corporation, PLC (ADR), NYSE ticker ELN. *Source*: Bloomberg Financial L.P.

What information are we looking for?

- In efficient markets, asset prices impound information and points of view widely dispersed among market participants
- Asset prices inherently forward-looking
- Asset prices summarize not only dispersed information about current and future conditions, but also about future asset prices themselves
- Germane to risk measurement because may give us an external, market-adjusted assessment of future to compare with model-generated
- Chief challenge to extracting information about future asset prices from current asset prices: asset prices also heavily affected by risk

Information and asset types

Cash securities and exchange rates embed information about future price of the asset and on size and certainty of future cash flows

Interest rates contain information about future interest rates as well as current and future business conditions

Futures and forwards contain information about future prices

Options contain information about probability distributions of future prices

Information in prices of futures and forwards

- Futures and forwards have future maturity dates, appear to provide distinct information point regarding future underlying price
- For long a focus of efficient markets research
  - Basic conclusion: add little information about future prices to what we know from current prices
  - Equal to current price adjusted for carry, so little to add
- Interest rates contain forwards on future short-term rates

Risk aversion and asset prices

- Asset prices express “consensus” risk preferences
  - If market participants seek insurance, assets with high payoffs in bad times are dear
  - And even more so if market participants risk averse
- Expectations/views on probability distribution of future outcomes entangled with risk preferences
- But overall strong evidence of risk aversion
- Option prices most revealing because state-dependent

Risk premiums

- Asset prices embed **risk premiums**
- Risk premiums express difference between expected values of risky and risk-free securities
  - Expected returns on risky asset higher than risk-free rate
  - $\Leftrightarrow$Risky assets cheaper than safe ones with same average cash flows, since income discounted by risk-free rate *plus* positive risk premium
- Risk premiums not directly observable, must be estimated
- CAPM:
  - Simple explanation of risk premium as compensation for holding market portfolio, an asset with low payoffs in bad times
  - And simple way to measure, via beta to market portfolio

Risk-neutral and subjective distributions

Subjective distribution: also called **physical** or actuarial or simply "real-life" distribution

- It is either the distribution "the market" believes in or the actual distribution from the point of view of superior knowledge

Risk-neutral distribution is the distribution implied by asset prices

- Odd name, since expresses risk aversion
- Subjective distribution we would attribute to representative agent if she is assumed indifferent to risk

Option prices and implied volatility

- Options have asymmetric payoff function⇒option value a function of volatility
  - ⇒Option prices embed a volatility forecast
- **Black-Scholes formula**: benchmark for option valuation
- **Implied volatility** := volatility estimate that matches an observed option price to the Black-Scholes formula
- Implied volatility used by option traders as a pricing metric as well as a volatility estimate

Implied volatility and risk premiums

- General *level* of implied volatility contains a negative **variance risk premium**
- Options have high payoffs in bad times, hence dearer than "fair bet"
- Implied volatility consistently higher than expected (as measured by historical)
- Analogy to insurance: probability of fire 1 in 1,000,000, but priced as if 1 in 500,000

Variance risk premium

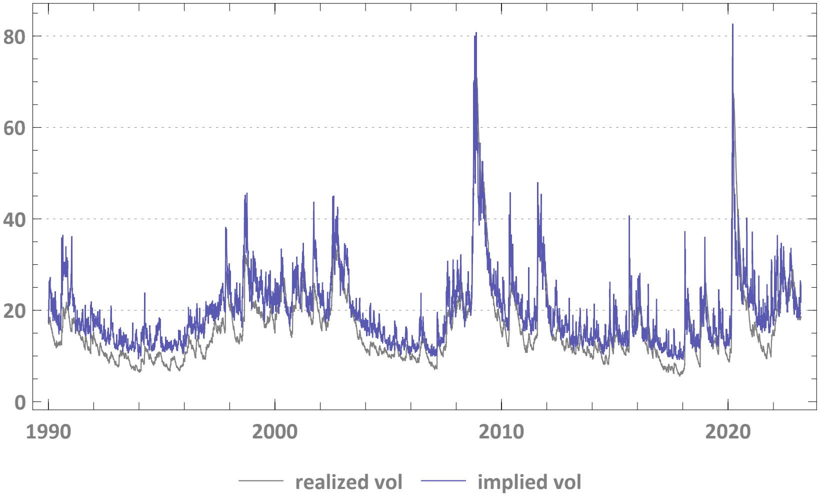

Black plot, estimated volatility (EWMA with decay factor $\lambda = 0.94$). Red plot: VIX. Both in percent at an annual rate.

Implied volatility surface

- **Black-Scholes model**—as opposed to the *formula*—incorporates standard asset return model
  - Envisions a single unvarying implied volatility
  - Across time, option tenors, strike prices
- But standard model doesn't hold precisely
- Variation over time
- $\rightarrow$**Implied volatility surface**: variation of implied volatility by *option maturity* and by *exercise price*

  Term structure of implied volatility: long-dated options generally closer to "forever volatility"

  Implied volatility skew Out-of-the-money call option volatilities$\neq$those of equally out-of-the-money puts

  Implied volatility smile Out-of-the-money options volatilities$>$those of at-the-money options

Options and the risk-neutral distribution

- Options state-dependent, so express risk aversion to specific outcomes
- **Empirical pricing kernel** "distributes" risk premium across price axis
  - Defined as ratio of subjective to risk neutral probability of a specific range of outcomes
- Leads to a forecast of entire return distribution (**Breeden-Litzenberger formula**)
- Estimate of risk-neutral distribution
- $\rightarrow$Alternatives to the standard return model

Breeden-Litzenberger formula

- Option prices embed a risk-neutral distribution (RND)
- $c(t, X, \tau) \equiv$ implied volatility surface, the market's time-t price schedule of European calls with tenor τ, struck at X
- $\tilde{\Pi}(X) \equiv$ time-t risk-neutral CDF, looking τ years ahead

$$\tilde{\Pi}(X) = 1 + e^{r\tau} \frac{\partial c}{\partial X}$$

- Intuitively, CDF $\approx$ 1 minus the slope of the call value as a function of strike
- Equivalently via Black-Scholes implied volatility smile $\sigma(t, X, \tau)$ and relationship to Black-Scholes formula
  $c(t, X, \tau) = v[X, \tau, \sigma(t, X, \tau), \ldots]$

Challenges arise from limited data

- In principle not hard to compute as discrete approximation
- Ideal data requirements to construct RND for specific asset/horizon
  - Dense set of plain-vanilla European options with same tenor but varying strikes, covering support of distribution
  - Ross theorem may permit using less data
- Real-world limitations: paucity and poor quality of data
  - Gaps in both strike and maturity dimensions
  - Sometimes no data, but even when available, need to fit or interpolate
  - Tails generally poorly represented or not at all
  - Liquidity: at any time, trading focuses on just a few strikes
  - Data problems→violations of no-arbitrage conditions

Risk-neutral distribution of S&P 500 index

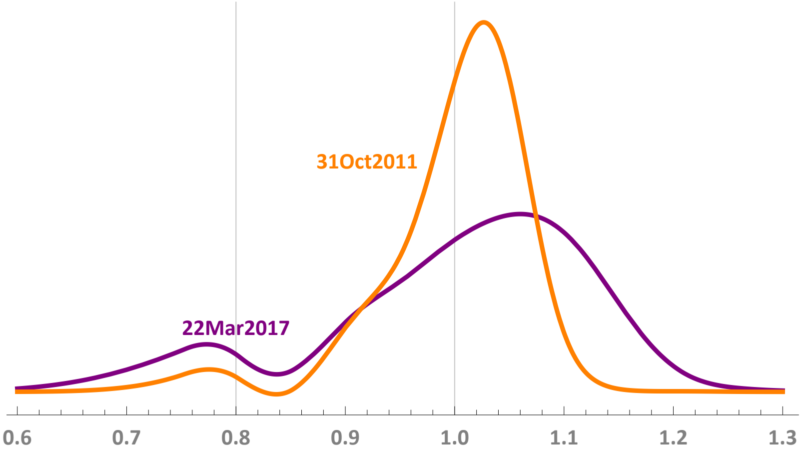

Density function of arithmetic return of the S&P 500 index 3 months hence. Black plot, using the estimated volatility (EWMA with decay factor $\lambda = 0.97$) on Oct. 31, 2011. Red (blue) plot: risk-neutral, estimated using the 3-month implied volatility smiles on Oct. 31, 2011 (Aug. 29, 2012).

Risk-neutral and subjective tail risk

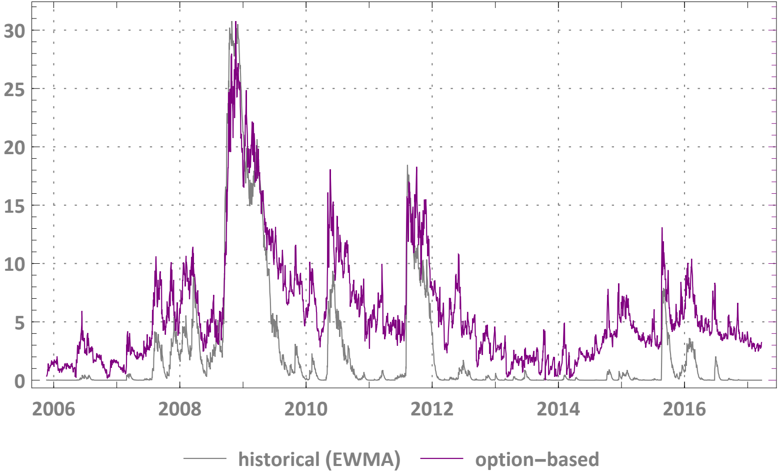

Probability of a decline of at least 20 percent in the S&P 500 index over the subsequent 3 months. Black plot, based on estimated volatility (EWMA, decay factor $\lambda = 0.94$). Red plot: risk-neutral, estimated using the 3-month implied volatility smile.

Correlation in market prices

Equity prices: correlation between returns on individual stocks

Interest rates: correlation between rates for different tenors⇔likelihood of changes in term spread

Credit spreads: default correlation

Definition and asset class coverage

- Correlations implied by derivatives prices
- Option-implied correlations
  - Equity markets: return correlation implied by index and single-stock vols
  - Analogous construction for FX: infer correlations of 2 USD currency pairs from cross-currency options
  - Fixed-income markets: correlation of changes in yields of two swap maturities, implied by swaption and curve option prices
- Implied default correlations
  - Correlation of default for two different firms
  - Implied by prices of two standard tranches

Interpretation and limitations

- Related to market assessment of relative importance of idiosyncratic and systematic risk
- Index vols high relative to typical single-stock vol⇒market participants (risk-neutrally) eager to cover generic equity exposure→high implied correlation
  - For any set of firms, expectation of large common shocks⇒higher correlation
  - Generally, high implied equity correlation in crises
- Data limitations
  - Some "constant pairwise correlation" coverage, from options on indexes
  - Very little on pairwise return and default correlation exposures

Construction of equity option-implied correlation

- Definition of stock portfolio volatility (k constituents):

$$\sigma^2_{\text{index},t} = \sum_k \omega^2_{kt}\sigma^2_{kt} + 2\rho_t \sum_k \sum_{j<k} \omega_{kt}\omega_{jt}\sigma_{kt}\sigma_{jt}$$

- Construction of implied correlation
  - Substitute observed index and single-firm implied vols $\sigma_{\text{index},t}$ and σ^2_{kt}, market cap weights ω_{kt}
  - Solve for ρ_t, constant (in cross-section) pairwise correlation

S&P 500 option-implied correlation

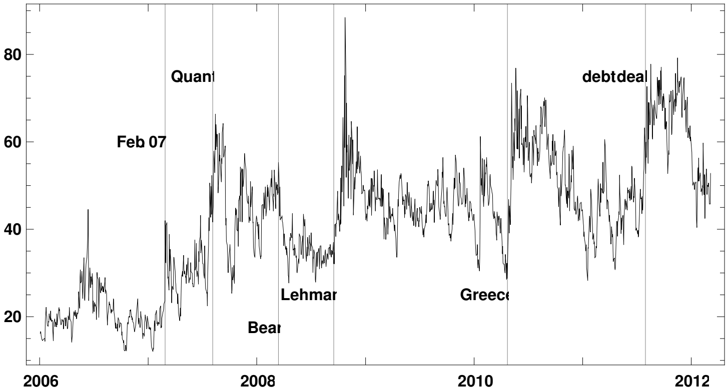

In percent, daily, 03Jan2006 to 31Oct2012. Uses the 100 largest-cap constituents (as of end.-Oct 2012) that have been in the S&P 500 index since at least Jan. 3, 2006, accounting for about 72 percent of the the total market cap of the index. *Data source*: Bloomberg Financial L.P.

Definition and interpretation

- Equivalent to spread of bond yield over swap/Treasury curve
    - But CDS floating-payment trigger may not be identical to a bond's default event
    - Similar to forwards in being issued daily with uniform tenor rather than fixed maturity date
- Indicator of risk-neutral default expectations
    - CDS curve can be used to extract time-varying hazard rates
    - But requires assumptions/estimates of recovery rate
- Spread to congruent bond yield an indicator of funding liquidity conditions
    - Exposure can be maintained with leverage
- Comparison to equity tail risk can be revealing

Financial CDS spreads

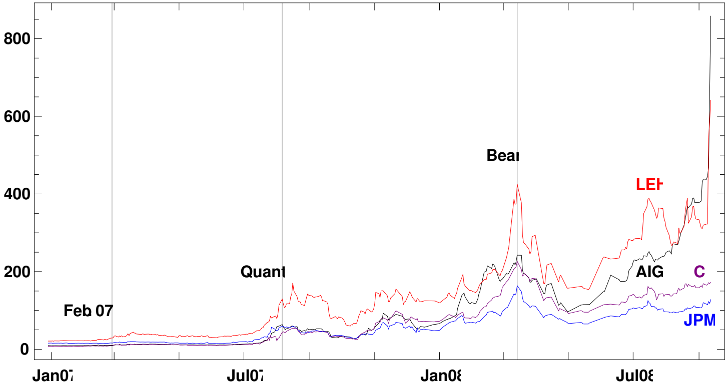

American International Group, Lehman Brothers Holdings Inc., JPMorgan Chase & Co. and Citigroup Inc. Senior unsecured, 5-year tenor, in bps. Daily end-2006 to 12Sep08. *Source*: Bloomberg Financial L.P.

Definition and interpretation

- Underlying an equally-weighted basket of CDS, e.g. CDX.NA.IG or iTraxx
  - Spread an indicator of average/overall risk-neutral default expectations
- Standard tranches
  - Exposure to segments of loss: 0–3 percent, 3–7 percent, etc.
  - **Base correlation:** copula correlation that matches model to market price
    - May be interpreted as an “index” of risk-neutral default correlation
    - **Correlation skew:** base correlation varies with tranche
  - Hedge ratios: notional amount of index that hedges tranche vs. small change in index spread
    - Level and volatility related to convexity in leveraged credit trades
    - Indicator of credit-market stress (2005 downgrade of Ford, GM)
    - Indicator of large positions (early 2012)

Base correlation 2004–2012

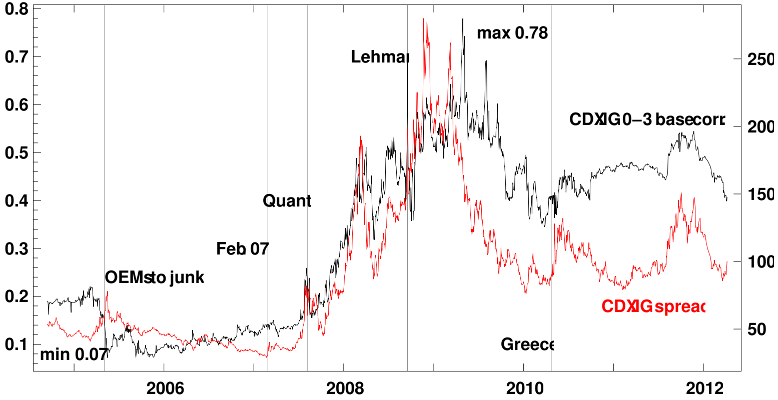

Black line (left axis) plots the equity (0-3 percent) base correlation of the on-the-run 5-year CDX.NA.IG series. Purple line (right axis) plots the on-the-run 5-year CDX.NA.IG series spread. Daily, 15Sep2004 to 31Oct2012. *Source*: JPMorgan.

Overview

- Spread between a derivatives-based and a spot rate. Examples:
  - LOIS
  - Swap to Treasury
  - Foreign-exchange forward-implied to Libor