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

Extreme market risk events

Allan M. Malz

Columbia University

Limitations of the standard model

Alternatives to the standard model

Limitations of the standard model

- Problems with the standard model
- Tools for understanding data

Alternatives to the standard model

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

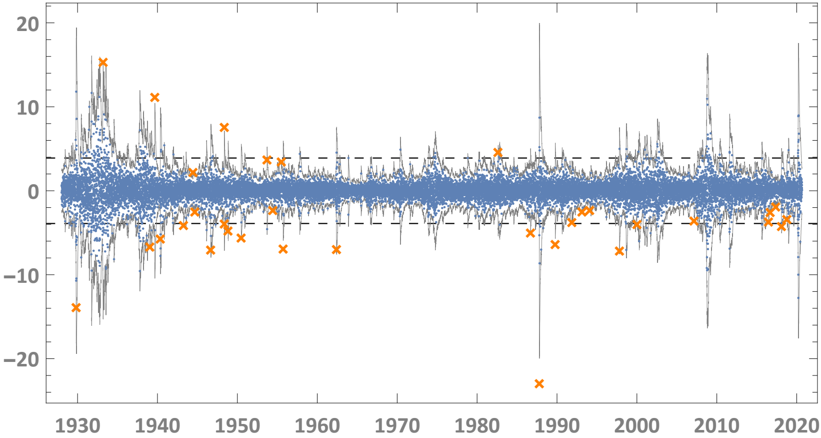

Points: daily realized returns (bps), solid lines EWMA 99.9% confidence interval ($\pm$3.3 vols), horizontal grid lines 99.9% CI using standard deviation of entire series. $\times$: returns outside the EWMA 99.9999% confidence interval (one-in-a-million chance). *Data source*: Bloomberg Financial L.P.

10-year Treasury Note yield fluctuations 1962-2020

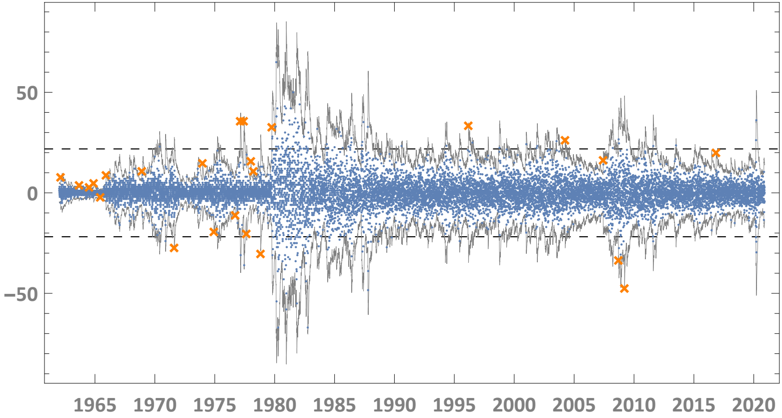

Points: daily realized change in yield (bps), solid lines EWMA 99.9% confidence interval (±3.3 vols), horizontal grid lines 99.9% CI using standard deviation of entire series. ×: returns outside the EWMA 99.9999% confidence interval (one-in-a-million chance). *Data source*: Bloomberg Financial L.P.

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
- **Conditional skewness**: degree of skewness depends on current volatility

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

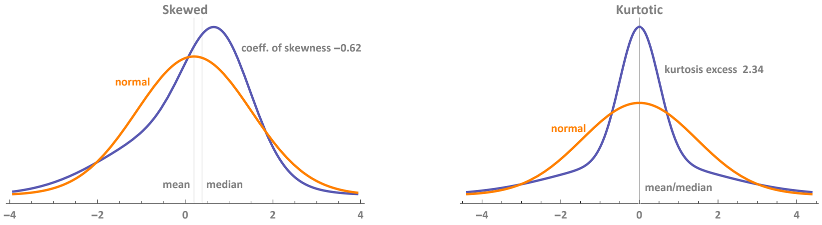

Left: skewed and normal distributions with identical means of 0.2 and variance 1.70875. The skewed distribution is a mixture of two normals, with distributions $\mathcal{N}(-0.35, 1.5)$ and $\mathcal{N}(0.75, 0.75)$, each with a probability of 50 percent of being realized. Right: kurtotic and normal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

Volatility asymmetry in the U.S. stock market

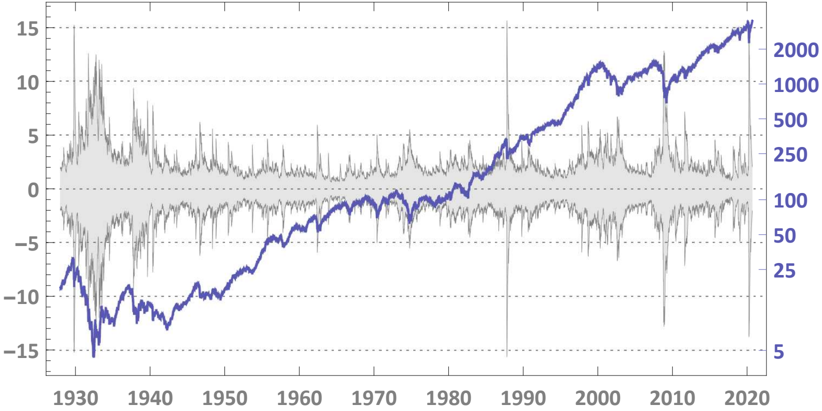

99.9% confidence interval using EWMA estimate of S&P 500 index return volatility, decay factor 0.94 (gray plot, left *y*-axis), and log of S&P 500 index closing price (purple plot, right *y*-axis), both daily 30Dec1927 to 11May2016. *Data source*: Bloomberg L.P.

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

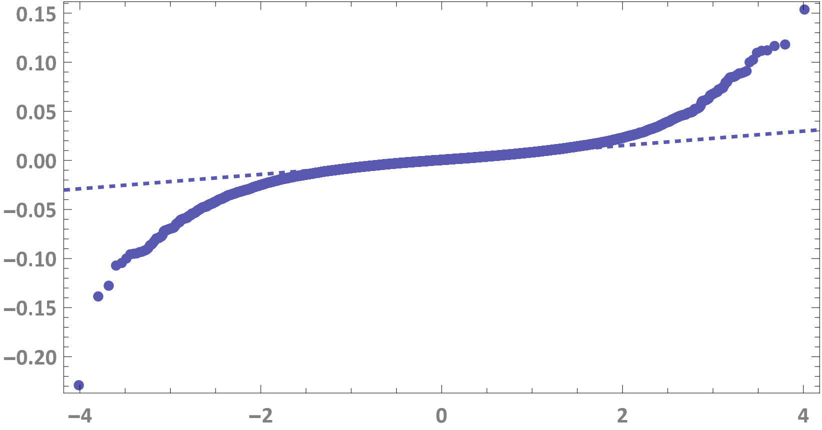

Quantiles of daily returns (as decimals) of the S&P 500 index, 1927-2011, plotted against quantiles of a standard normal distribution. *Data source*: Bloomberg Financial L.P.

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

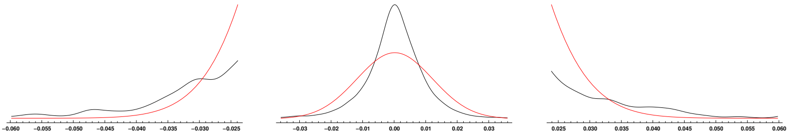

Kernel estimate of probability density function of daily returns of the S&P 500 index, 1927-2011. Gaussian kernel using optimal bandwidth. *Data source*: Bloomberg Financial L.P.

Limitations of the standard model

Alternatives to the standard model

Alternative distributional hypotheses

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

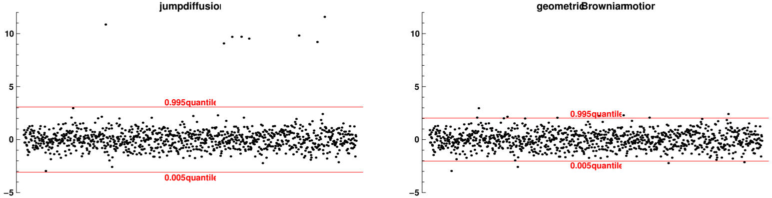

Simulation of 1000 daily sequential steps of a jump diffusion (left panel) and geometric Brownian motion process. The *y*-axis displays returns in percent. The horizontal grid lines mark the 99 percent confidence intervals. The jump diffusion follows

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

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.