Equity in the Time Series, Part 2

15.433 Financial Markets

October 3 & 5, 2017

This Version: September 26, 2017

Outline

- Volatility models and market risk measurement.
- Estimating volatility using financial time series:
  - SMA: simple moving average model (traditional approach).
  - EWMA: exponentially weighted moving average model (RiskMetrics).
  - ARCH and GARCH models (Nobel Prize).
- EWMA for covariances and correlations.
- Portfolio volatility and Value-at-Risk.

What have we learned about the aggregate stock market?

- It is pervasive, the single most important risk factor in the equity world.
- It yields a positive risk premium, but the risk premium is difficult to measure with precision because of
  - the "high" level of stock market volatility
  - and the limited length of the historical data.
- There is some evidence that the expected returns are time varying. The autocorrelation of the aggregate stock returns is slightly positive, and the dividend-to-price ratio has some predictability for future stock returns.
- Overall, only a small portion of future stock returns can be predicted (low R-squared's), and much of the uncertainty is unpredictable.

The volatility of the aggregate stock market

- Historical data can be used to measure volatility with much better precision. Between risk and return, risk is something we can collect more information about.
- In fact, we can learn about market volatility not only from the historical stock market data (backward looking), but also from derivatives prices (forward looking).
- Academics have made much progress in both directions, and practitioners have adopted many of the ideas developed by academics.
- We will study three volatility estimators:
  - SMA: simple moving average model (traditional approach).
  - EWMA: exponentially weighted moving average model (RiskMetrics).
  - ARCH and GARCH models (Nobel Prize).

The importance of measuring market volatility

- Portfolio managers performing optimal asset allocation.
- Risk managers assessing portfolio risk (e.g., Value-at-Risk).
- Derivatives investors trading non-linear contracts with values linked directly to market volatility.
- Increasingly, the level of market volatility (e.g., VIX) has become a market indicator ("the fear gauge") watched closely by almost all institutional investors, including those who are not trading directly in the U.S. equity or U.S. equity derivatives markets.

Modern Finance

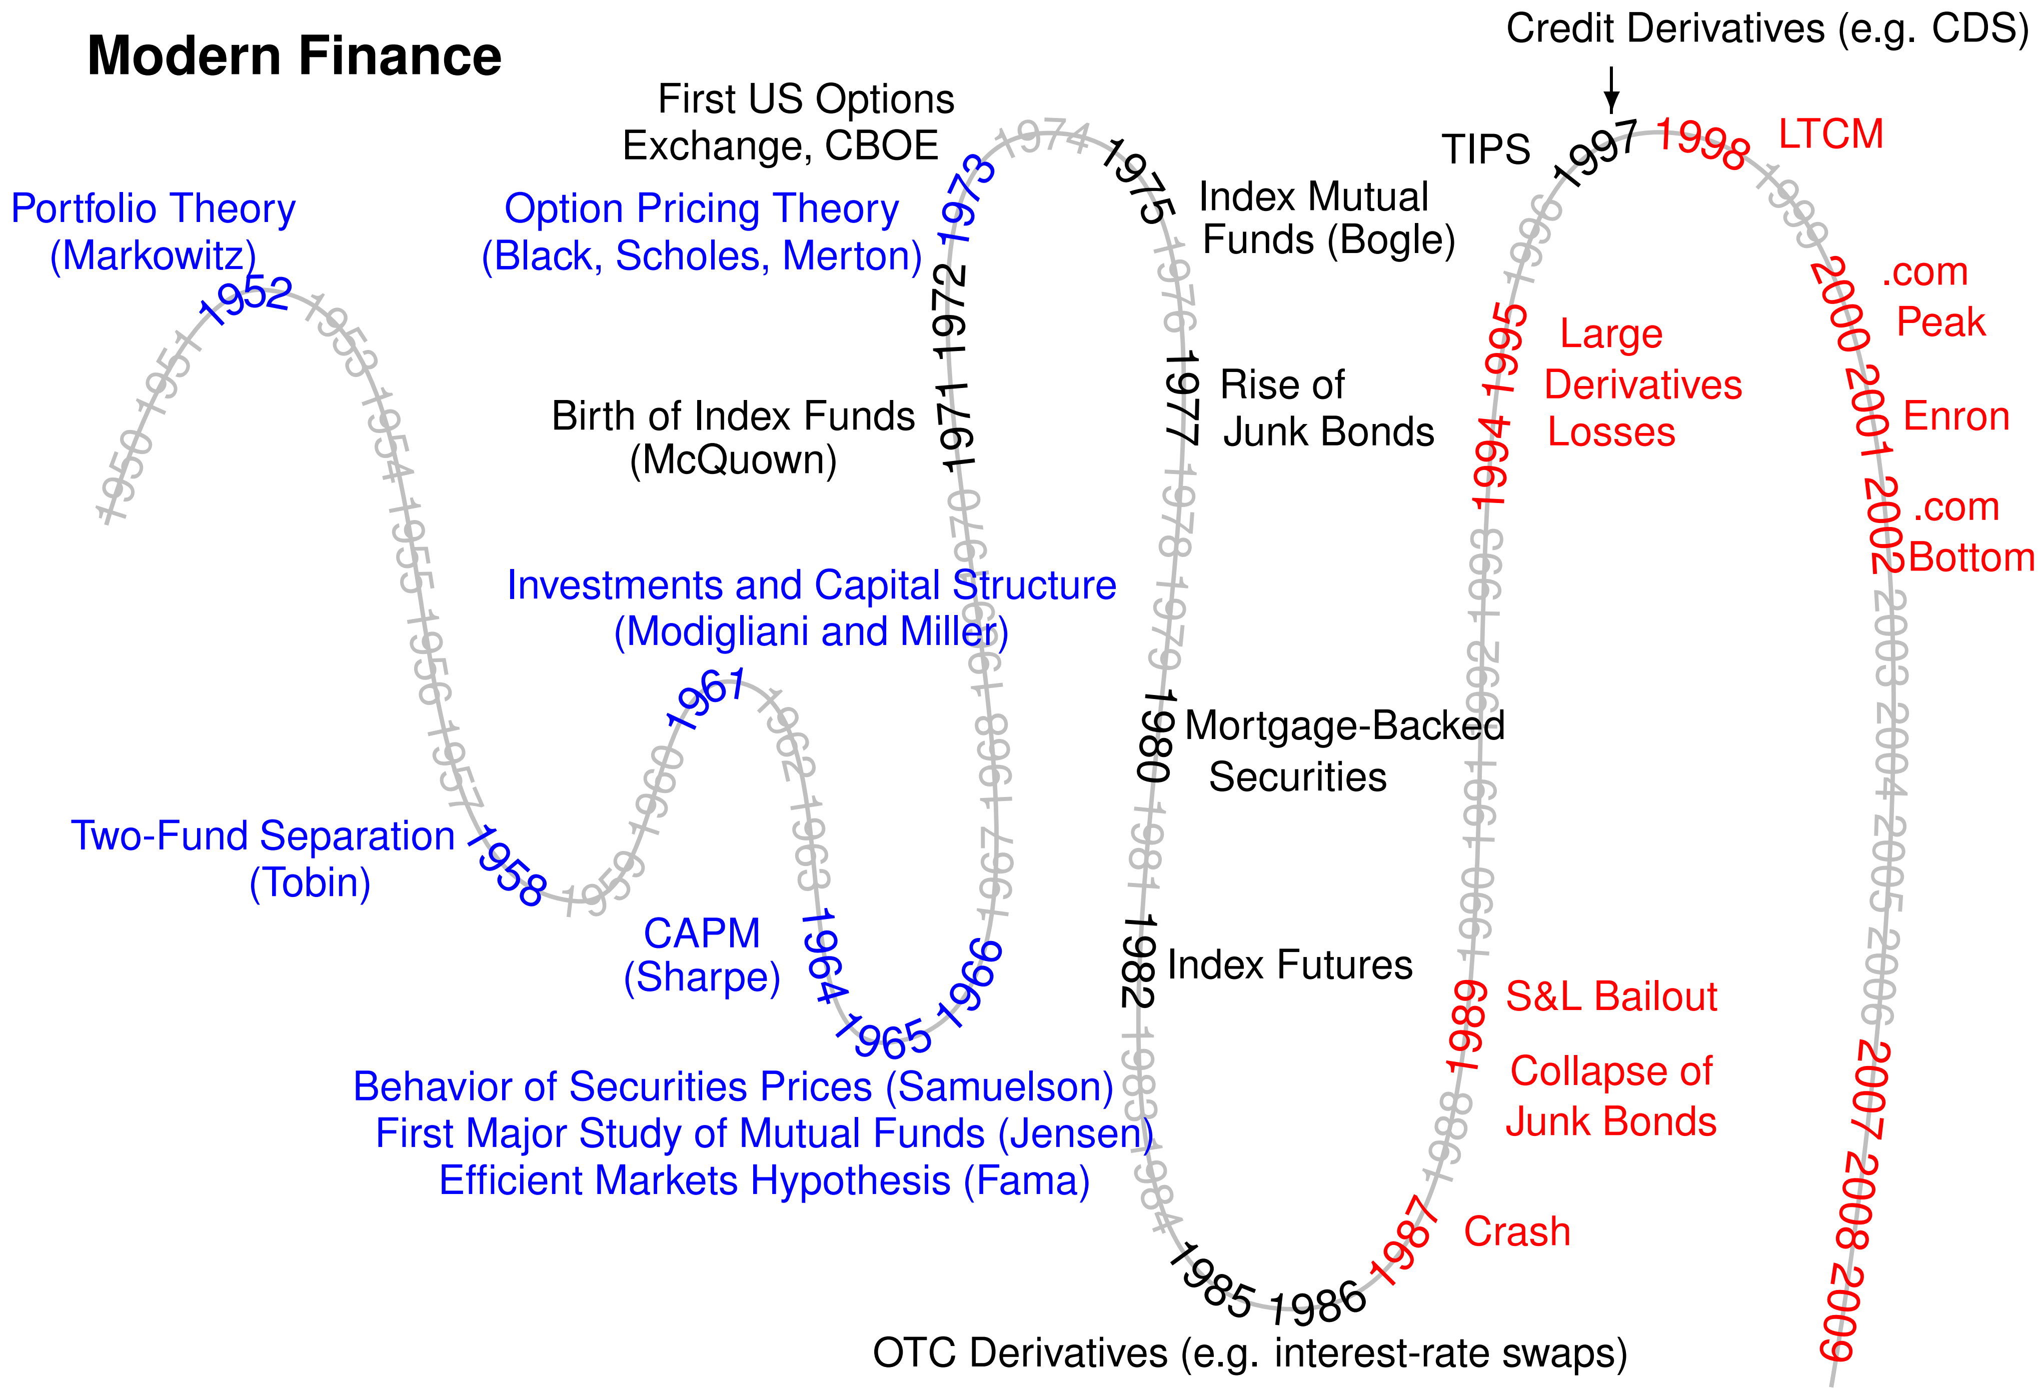

The Evolution of an Investment Bank

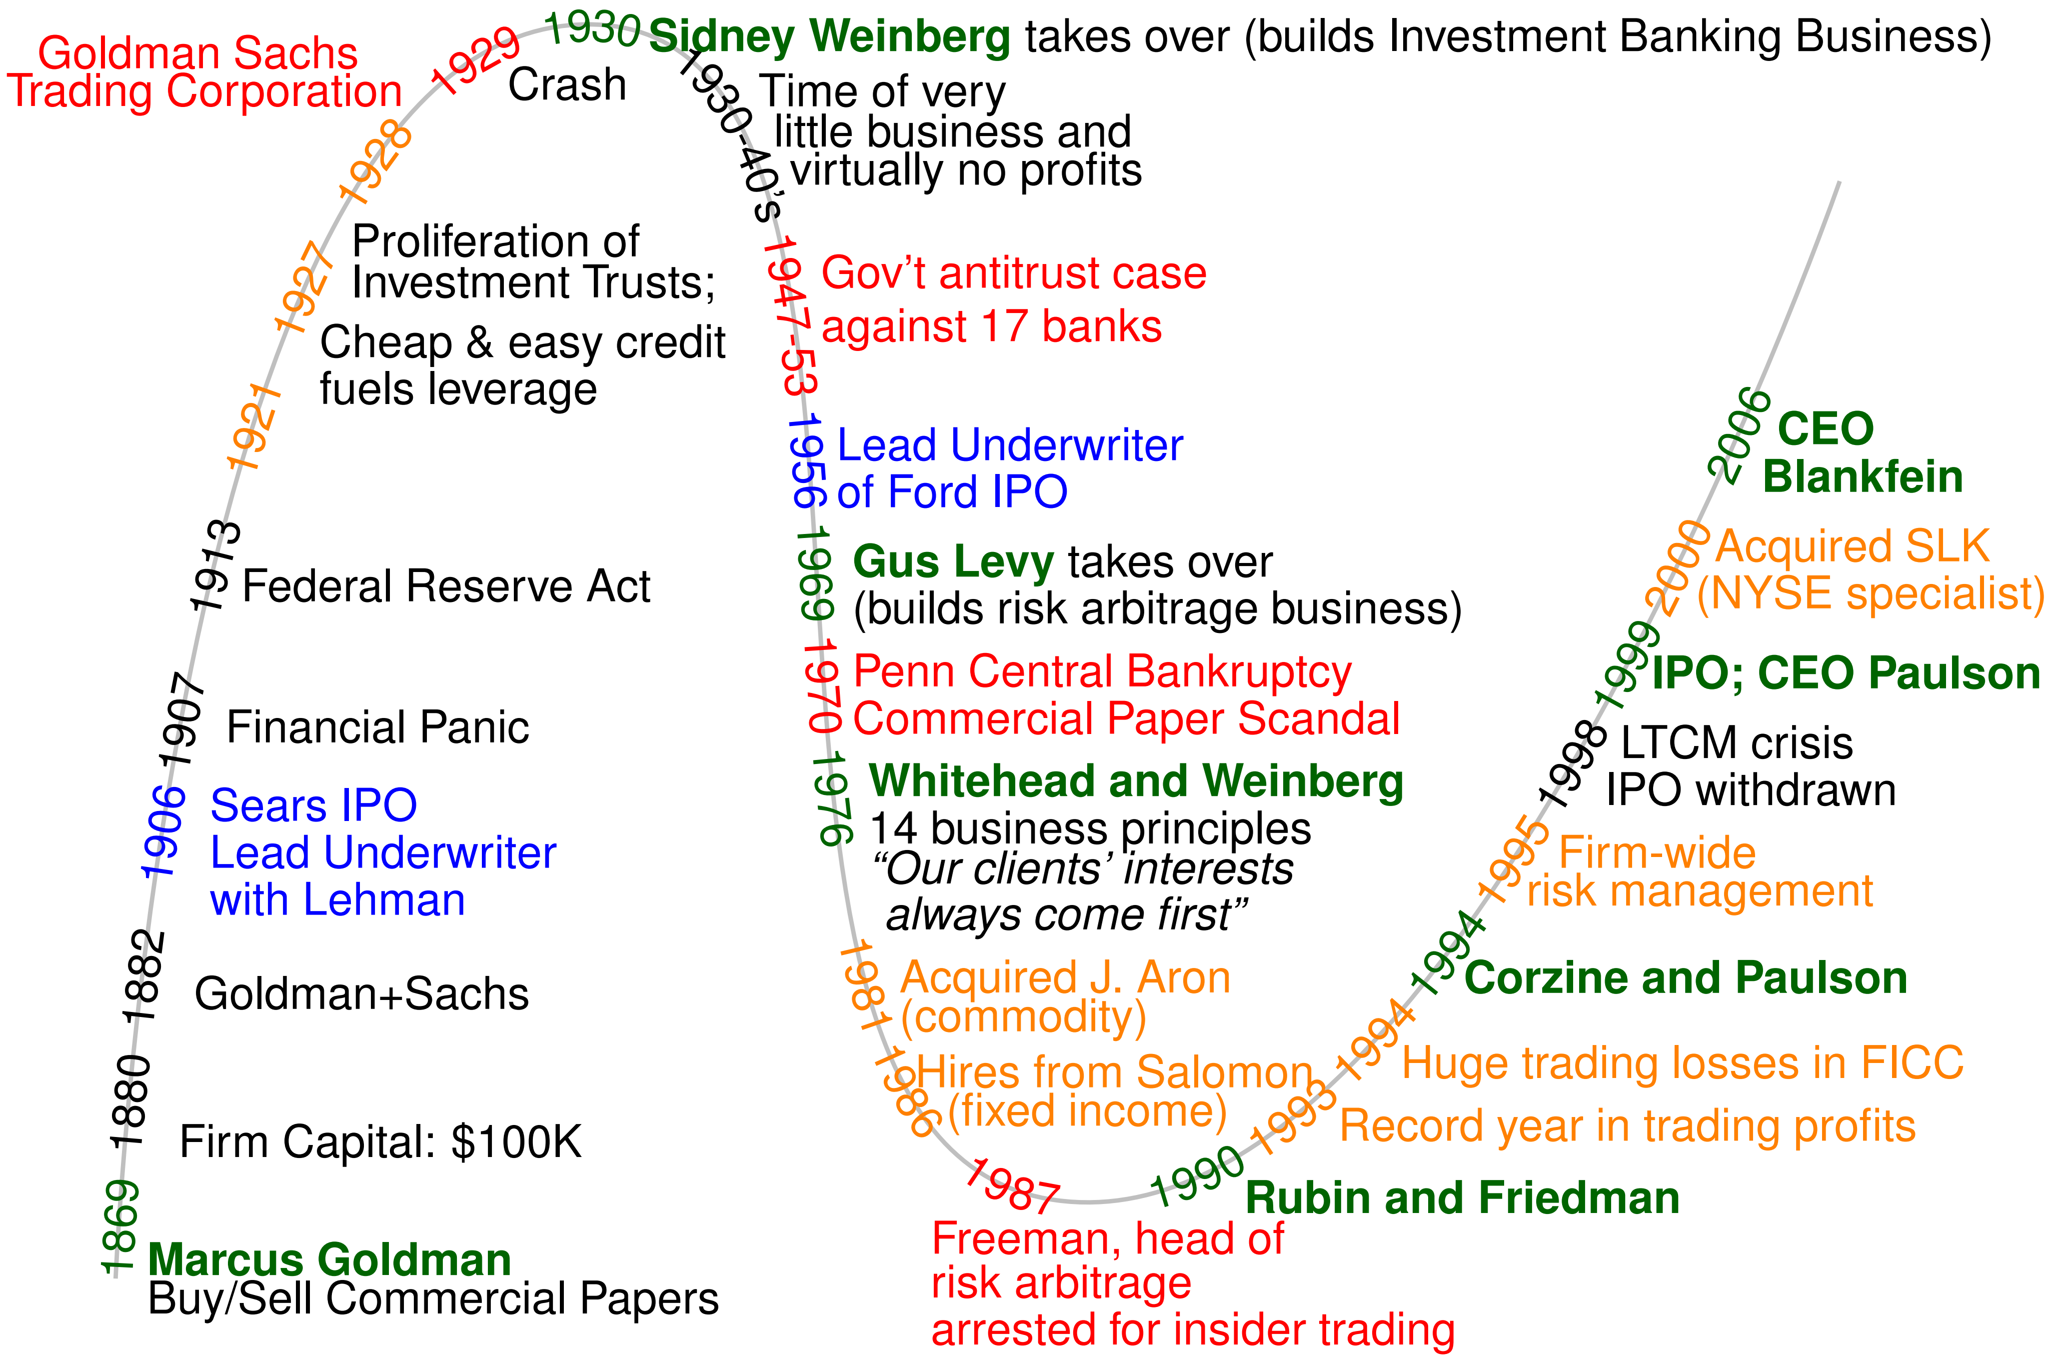

Some losses on derivatives positions by non-financial corporations in mid-1990s

Orange County: $1.7 billion, leverage (reverse repos) and structured notes

Showa Shell Sekiyu: $1.6 billion, currency derivatives

Metallgesellschaft: $1.3 billion, oil futures

Barings: $1 billion, equity and interest rate futures

Codelco: $200 million, metal derivatives

Proctor & Gamble: $157 million, leveraged currency swaps

Air Products & Chemicals: $113 million, leveraged interest rate and currency swaps

Dell Computer: $35 million, leveraged interest rate swaps

Louisiana State Retirees: $25 million, IOs/POs

Arco Employees Savings: $22 million, money market derivatives

Gibson Greetings: $20 million, leveraged interest rate swaps

Mead: $12 million, leveraged interest rate swaps

Measuring Market Risk

- By the early 1990s, the increasing activity in securitization and the increasing complexity in the financial instruments made the trading books of many investment banks too complex and diverse for the chief executives to understand the overall risk of their firms.
- Market risk management tools such as Value-at-Risk are ways to aggregate the firm-wide risk to a set of numbers that can be easily communicated to the chief executives. By the mid-1990s, most Wall Street firms have developed risk measurement into a firm-wide system.
- Daily estimates of market volatility, along with correlations across financial assets, constitute the key inputs to Value-at-Risk. JP Morgan's RiskMetrics uses exponentially weighted moving average (EWMA) model to estimate the volatilities and correlations of over 480 financial time series in order to construct a variance-covariance matrix of 480x480.

Table 9.5

Equity indices: sources*

Market	Exchange	Index Name	Weighting	% Mkt. cap.	Time, U.S. EST
Australia	Australian Stock Exchange	All Ordinaries	MC	96	1:10 a.m.
Hong Kong	Hong Kong Stock Exchange	Hang Seng	MC	77	12:30 a.m.
Indonesia	Jakarta Stock Exchange	JSE	MC		4:00 a.m.
Korea	Seoul Stock Exchange	KOPSI	MC		3:30 a.m.
Japan	Tokyo Stock Exchange	Nikei 225	MC	46	1:00 a.m.
Malaysia	Kuala Lumpur Stock Exchange	KLSE	MC		6:00 a.m.
New Zealand	New Zealand Stock Exchange	Capital 40	MC	—	10:30 p.m.
Philippines	Manila Stock Exchange	MSE Com'l &Inustil Price	MC		1:00 a.m.
Singapore	Stock Exchange of Singapore	Sing. All Share	MC	—	4:30 a.m.
Taiwan	Taipei Stock Exchange	TSE	MC		1:00 a.m.
Thailand	Bangkok Stock Exchange	SET	MC		5:00 a.m.
Austria	Vienna Stock Exchange	Creditanstalt	MC	—	7:30 a.m.
Belgium	Brussels Stock Exchange	BEL 20	MC	78	10:00 a.m.
Denmark	Copenhagen Stock Exchange	KFX	MC	44	9:30 a.m.
Finland	Helsinki Stock Exchange	Hex General	MC	—	10:00 a.m.
France	Paris Bourse	CAC 40	MC	55	11:00 a.m.
Germany	Frankfurt Stock Exchange	DAX	MC	57	10:00 a.m.
Ireland	Irish Stock Exchange	Irish SE ISEQ	—	—	12:30 p.m.
Italy	Milan Stock Exchange	MIB 30	MC	65	10:30 a.m.
Japan	Tokyo Stock Exchange	Nikei 225	MC	46	1:00 a.m.
Netherlands	Amsterdam Stock Exchange	AEX	MC	80	10:30 a.m.
Norway	Oslo Stock Exchange	Oslo SE General	—	—	9:00 a.m.
Portugal	Lisbon Stock Exchange	Banco Totta SI	—	—	11:00 a.m.
South Africa	Johannesburg Stock Exchange	JSE	MC		10:00 a.m.
Spain	Madrid Stock Exchange	IBEX 35	MC	80	11:00 a.m.
Sweden	Stockholm Stock Exchange	OMX	MC	61	10:00 a.m.
Switzerland	Zurich Stock Exchange	SMI	MC	56	10:00 a.m.
U.K.	London Stock Exchange	FTSE 100	MC	69	10:00 a.m.
Argentina	Buenos Aires Stock Exchange	Merval	Vol.		5:00 p.m.
Canada	Toronto Stock Exchange	TSE 100	MC	63	4:15 p.m.
Mexico	Mexico Stock Exchange	IPC	MC		3:00 p.m.
U.S.	New York Stock Exchange	Standard and Poor's 100	MC	60	4:15 a.m.

* Data sourced from DRI.

Table 9.1
Foreign exchange

Currency Codes					
Americas		**Asia Pacific**		**Europe and Africa**	
ARS	Argentine peso	AUD	Australian dollar	ATS	Austrian shilling
CAD	Canadian dollar	HKD	Hong Kong dollar	BEF	Belgian franc
MXN	Mexican peso	IDR	Indonesian rupiah	CHF	Swiss franc
USD	U.S. dollar	JPY	Japanese yen	DEM	Deutsche mark
EMB	EMBI+*	KRW	Korean won	DKK	Danish kroner
		MYR	Malaysian ringgit	ESP	Spanish peseta
		NZD	New Zealand dollar	FIM	Finnish mark
		PHP	Philippine peso	FRF	French franc
		SGD	Singapore dollar	GBP	Sterling
		THB	Thailand baht	IEP	Irish pound
		TWD	Taiwan dollar	ITL	Italian lira
				NLG	Dutch guilder
				NOK	Norwegian kroner
				PTE	Portuguese escudo
				SEK	Swedish krona
				XEU	ECU
				ZAR	South African rand

* EMBI+ stands for the J.P. Morgan Emerging Markets Bond Index Plus.

Table 9.2

Money market rates: sources and term structures

	Source		Time	Term Structure			
Market	J.P. Morgan	Third Party*	U.S. EST	1m	3m	6m	12m
Australia	•		11:00 a.m.	•	•	•	•
Hong Kong		•	10:00 p.m.	•	•	•	•
Indonesia†	•		5:00 a.m.	•	•	•	•
Japan	•		11:00 a.m.	•	•	•	•
Malaysia†	•		5:00 a.m.	•	•	•	•
New Zealand		•	12:00 a.m.	•	•	•	
Singapore		•	4:30 a.m.	•	•	•	•
Thailand†	•		5:00 a.m.	•	•	•	•
Austria		•	11:00 a.m.	•	•	•	•
Belgium	•		11:00 a.m.	•	•	•	•
Denmark	•		11:00 a.m.	•	•	•	•
Finland		•	11:00 a.m.	•	•	•	•
France	•		11:00 a.m.	•	•	•	•
Ireland		•	11:00 a.m.	•	•	•	•
Italy	•		11:00 a.m.	•	•	•	•
Netherlands	•		11:00 a.m.	•	•	•	•
Norway		•	11:00 a.m.	•	•	•	•
Portugal		•	11:00 a.m.	•	•	•	•
South Africa			11:00 a.m.	•	•	•	•
Spain	•		11:00 a.m.	•	•	•	•
Sweden	•		11:00 a.m.	•	•	•	•
Switzerland	•		11:00 a.m.	•	•	•	•
U.K.	•		11:00 a.m.	•	•	•	•
ECU	•		11:00 a.m.	•	•	•	•
Canada	•		11:00 a.m.	•	•	•	•
Mexico‡	•		12:00 p.m.	•	•	•	•
U.S.	•		11:00 a.m.	•	•	•	•

* Third party source data from Reuters Generic except for Hong Kong (Reuters HIBO), Singapore (Reuters MASX), and New Zealand (National Bank of New Zealand).

† Money market rates for Indonesia, Malaysia, and Thailand are calculated using foreign exchange forward-points.

‡ Mexican rates represent secondary trading in Cetes.

Table 9.3

Government bond zero rates: sources and term structures

	Source		**Time**	**Term structure**									
Market	J.P. Morgan	Third Party	U.S. EST	2y	3y	4y	5y	7y	9y	10y	15y	20y	30y
Australia	•		1:30 a.m.	•	•	•	•	•	•	•	•		
Japan	•		1:00 a.m.	•	•	•	•	•	•	•			
New Zealand		•	12:00 a.m.	•	•	•	•	•	•	•	•		
Belgium	•		11:00 a.m.	•	•	•	•	•	•	•	•	•	
Denmark		•	10:30 a.m.	•	•	•	•	•	•	•	•	•	•
France	•		10:30 a.m.	•	•	•	•	•	•	•	•	•	•
Germany	•		11:30 a.m.	•	•	•	•	•	•	•	•	•	•
Ireland		•	10:30 a.m.	•	•	•	•	•	•	•	•	•	
Italy	•		10:45 a.m.	•	•	•	•	•	•	•	•	•	•
Netherlands	•		11:00 a.m.	•	•	•	•	•	•	•	•	•	•
South Africa	•		11:00 a.m.	•	•	•	•	•	•	•	•	•	
Spain	•		11:00 a.m.	•	•	•	•	•	•	•	•		
Sweden		•	10:00 a.m.	•	•	•	•	•	•	•	•		
U.K.	•		11:45 a.m.	•	•	•	•	•	•	•	•	•	•
ECU	•		11:45 a.m.	•	•	•	•	•	•	•			
Canada	•		3:30 p.m.	•	•	•	•	•	•	•	•	•	•
U.S.	•		3:30 a.m.	•	•	•	•	•	•	•	•	•	•
Emerging Mkt.†	•		3:00 p.m.										

* Third party data sourced from Den Danske Bank (Denmark), NCB Stockbrokers (Ireland), National Bank of New Zealand (New Zealand), and SE Banken (Sweden).

† J. P. Morgan Emerging Markets Bond Index Plus (EMBI+).

Table 9.4

Swap zero rates: sources and term structures

Market	Source: J.P. Morgan	Source: Third Party*	Time: US EST	2y	3y	4y	5y	7y	10y
Australia	•		1:30 a.m.	•	•	•	•	•	•
Hong Kong		•	4:30 a.m.	•	•	•	•	•	•
Indonesia		•	4:00 a.m.	•	•	•	•		
Japan	•		1:00 a.m.	•	•	•	•	•	•
Malaysia		•	4:00 a.m.	•	•	•	•		
New Zealand		•	3:00 p.m.	•	•	•	•	•	
Thailand		•	4:00 a.m.	•	•	•	•		
Belgium	•		10:00 a.m.	•	•	•	•	•	•
Denmark	•		10:00 a.m.	•	•	•	•	•	•
Finland	•		10:00 a.m	•	•	•	•		
France	•		10:00 a.m.	•	•	•	•	•	•
Germany	•		10:00 p.m.	•	•	•	•	•	•
Ireland		•	11:00 a.m.	•	•	•	•		
Italy	•		10:00 a.m.	•	•	•	•	•	•
Netherlands	•		10:00 a.m.	•	•	•	•	•	•
Spain	•		10:00 a.m.	•	•	•	•	•	•
Sweden	•		10:00 a.m.	•	•	•	•	•	•
Switzerland	•		10:00 a.m.	•	•	•	•	•	•
U.K.	•		10:00 a.m.	•	•	•	•	•	•
ECU	•		10:00 a.m.	•	•	•	•	•	•
Canada	•		3:30 p.m.	•	•	•	•	•	•
U.S.	•		3:30 a.m.	•	•	•	•	•	•

* Third party source data from Reuters Generic except for Ireland (NCBI), Hong Kong (TFHK), and Indonesia, Malaysia, Thailand (EXOT).

Table 9.6

Commodities: sources and term structures

Commodity	Source	Time, U.S. EST	Term structure						
			Spot	1m	3m	6m	12m	15m	27m
WTI Light Sweet Crude	NYMEX[*]	3:10 p.m.		•	•	•	•		
Heating Oil	NYMEX	3:10 p.m.		•	•	•	•		
NY Harbor #2 unleaded gas	NYMEX	3:10 p.m.		•	•	•			
Natural gas	NYMEX	3:10 p.m.		•	•	•	•		
Aluminum	LME[†]	11:20 a.m.	•		•			•	•
Copper	LME	11:15 a.m.	•		•			•	•
Nickel	LME	11:10 a.m.	•		•			•	
Zinc	LME	11:30 a.m.	•		•			•	•
Gold	LME	11:00 a.m.	•						
Silver	LFOE[‡]	11:00 a.m.	•						
Platinum	LPPA[§]	11:00 a.m.	•						

* NYMEX (New York Mercantile Exchange)

† LME (London Metals Exchange)

‡ LFOE (London futures and Options Metal Exchange)

§ LPPA (London Platinum & Palladium Association)

Estimating volatility using financial time series

- SMA: simple moving average model (traditional approach).
- EWMA: exponentially weighted moving average model (RiskMetrics).
- ARCH and GARCH models (Nobel Prize).

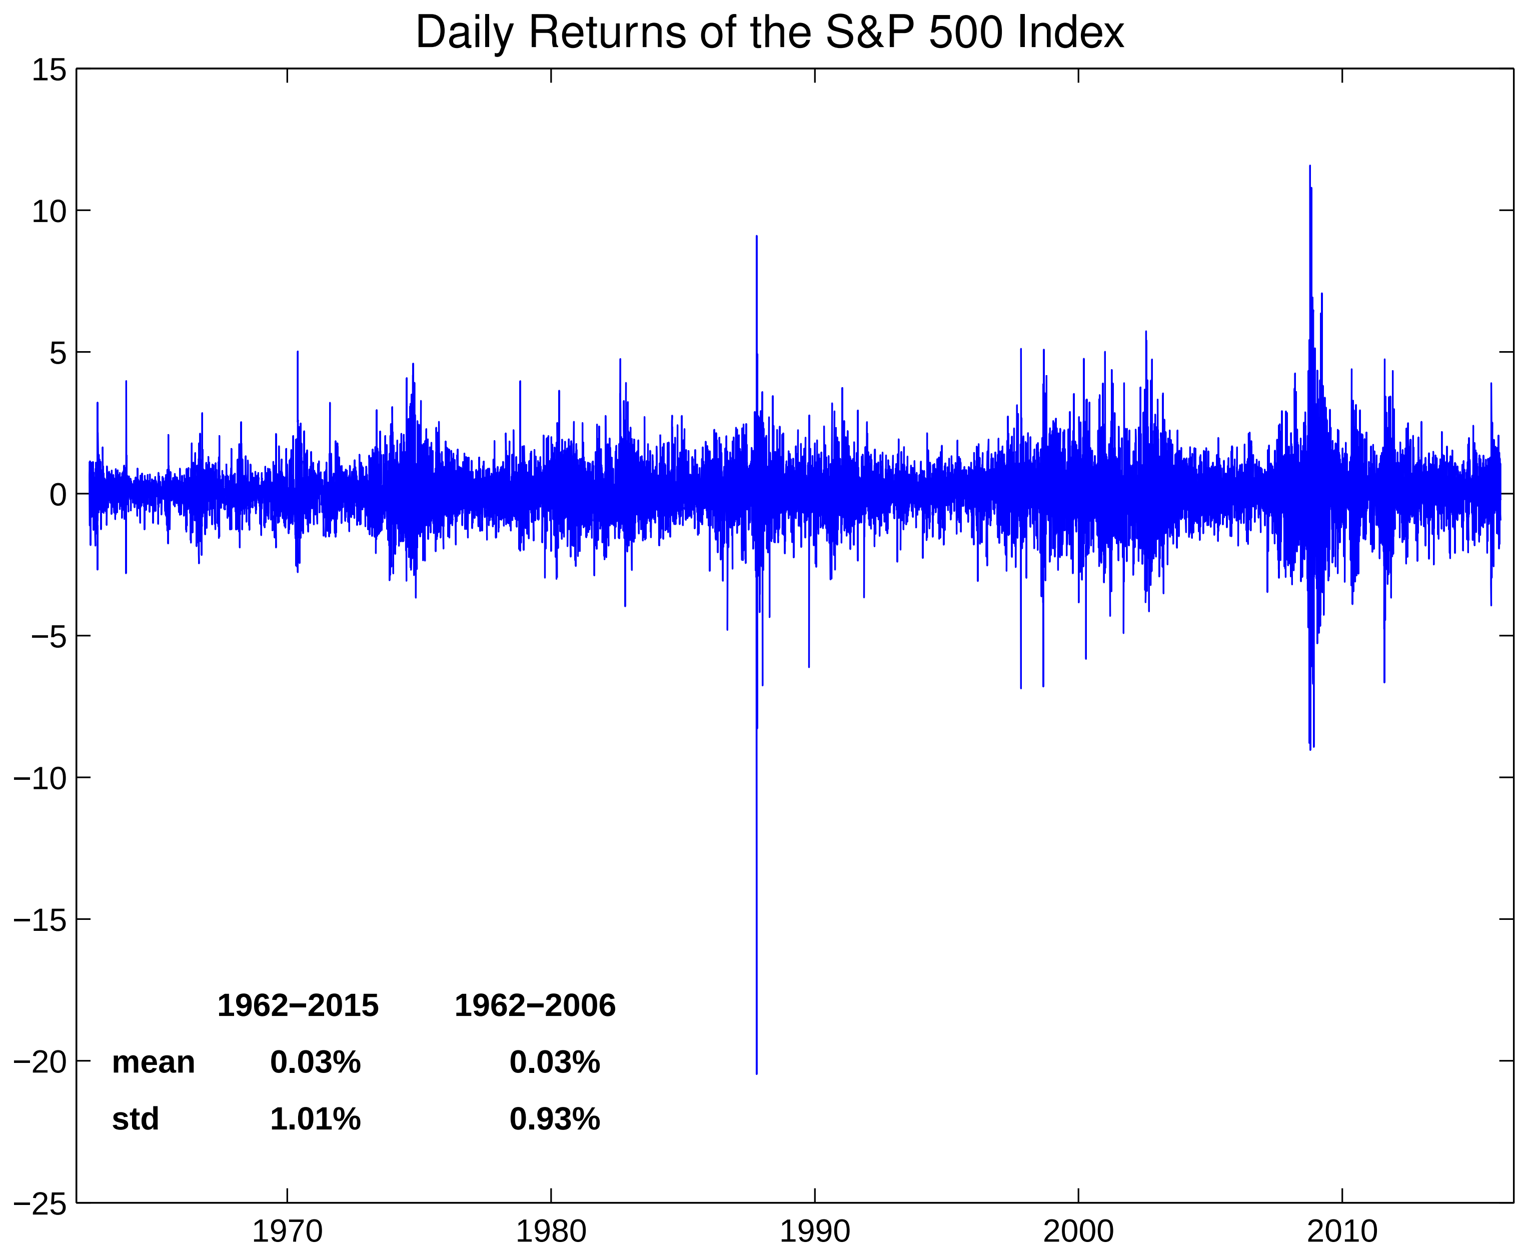
Daily Returns of the S&P 500 Index
15
10
5
0
−5
−10
−15
−20
−25
1970
1980
1990
2000
2010
1962−2015
1962−2006
mean
0.03%
0.03%
std
1.01%
0.93%

The Simple Moving Average Model

- Unlike expected returns, volatility can be measured with better precision using higher frequency data. So let's use daily data.
- Some have gone into higher frequency by using intra-day data. But micro-structure noises such as bid/ask bounce start to dominate in the intra-day domain. So let's not go there in this class.
- Suppose in month t, there are N trading days, with R_n denoting n-th day return. The simple moving average (SMA) model:

$$\sigma = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(R_n\right)^2}$$

- To get an annualized number: $\sigma \times \sqrt{252}$. (252 trading days per year).

Whether or not to take out μ?

- The industry convention is such that $(R_t - \mu)^2$ is replaced by R_t^2 in the volatility calculation.
- The reason is that, at daily frequency, μ^2 is too small compared with $E(R^2)$. Recall, μ is several basis points while σ is close to 1%.
- So instead of going through the trouble of doing $E(R^2) - \mu^2$, people just do $E(R^2)$.

Volatility estimates using the simple moving average (SMA) model

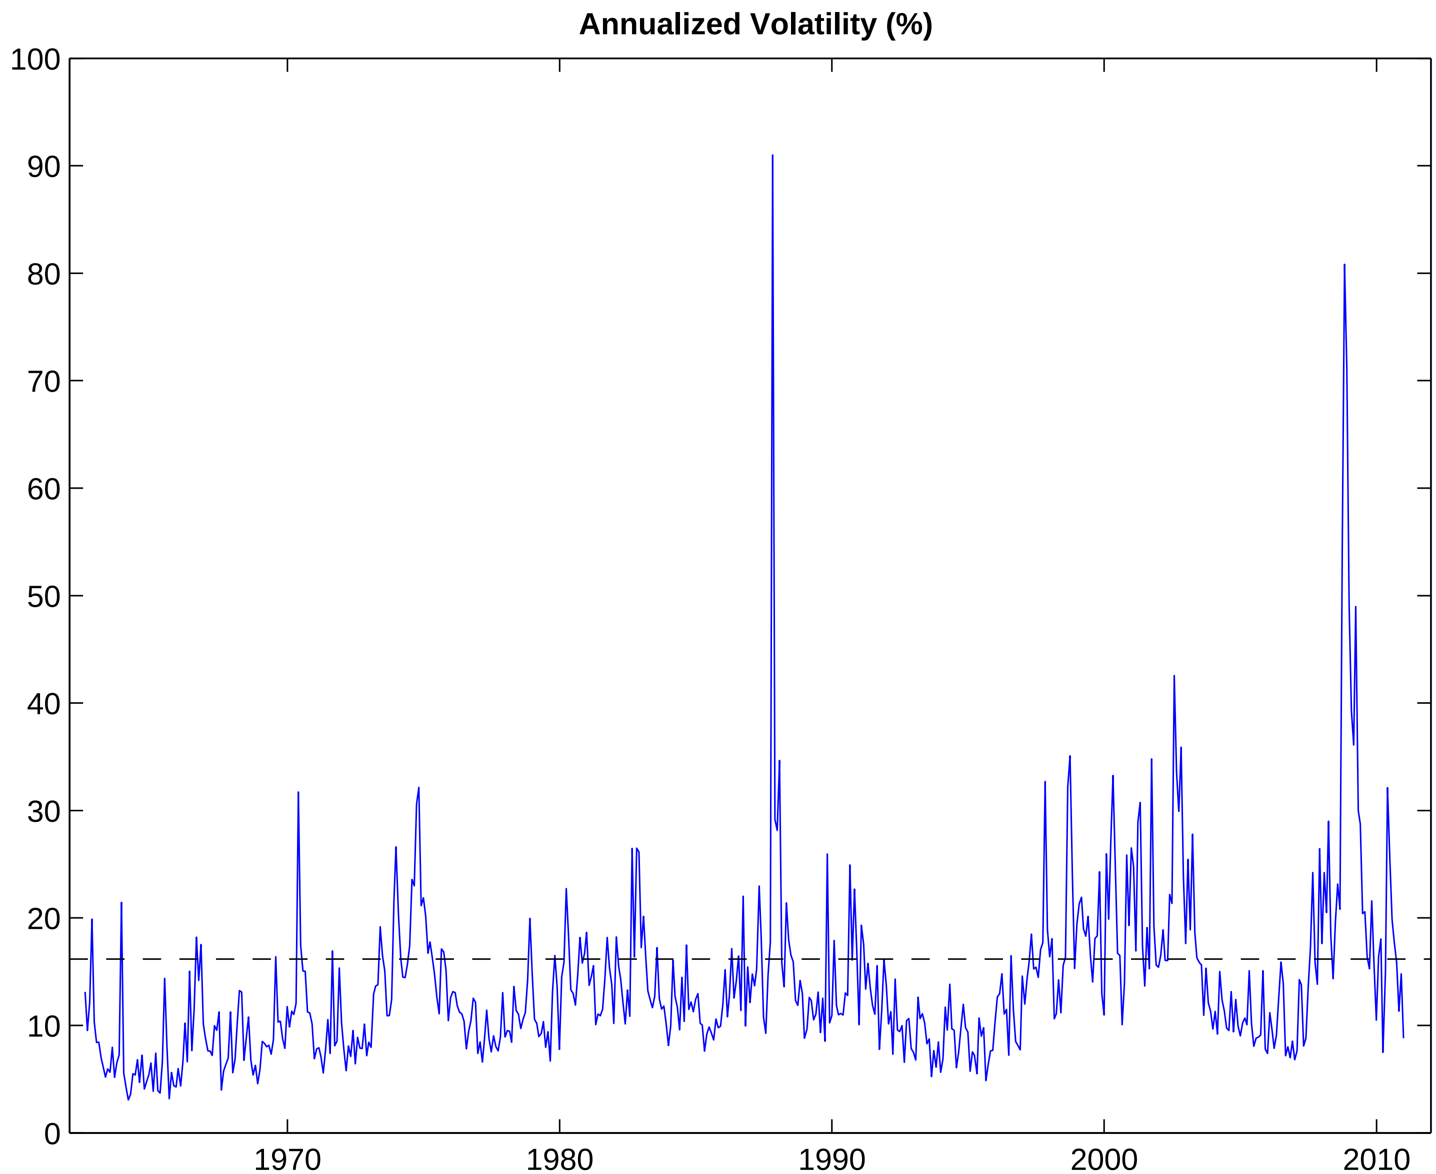

How precise are SMA volatility estimates?

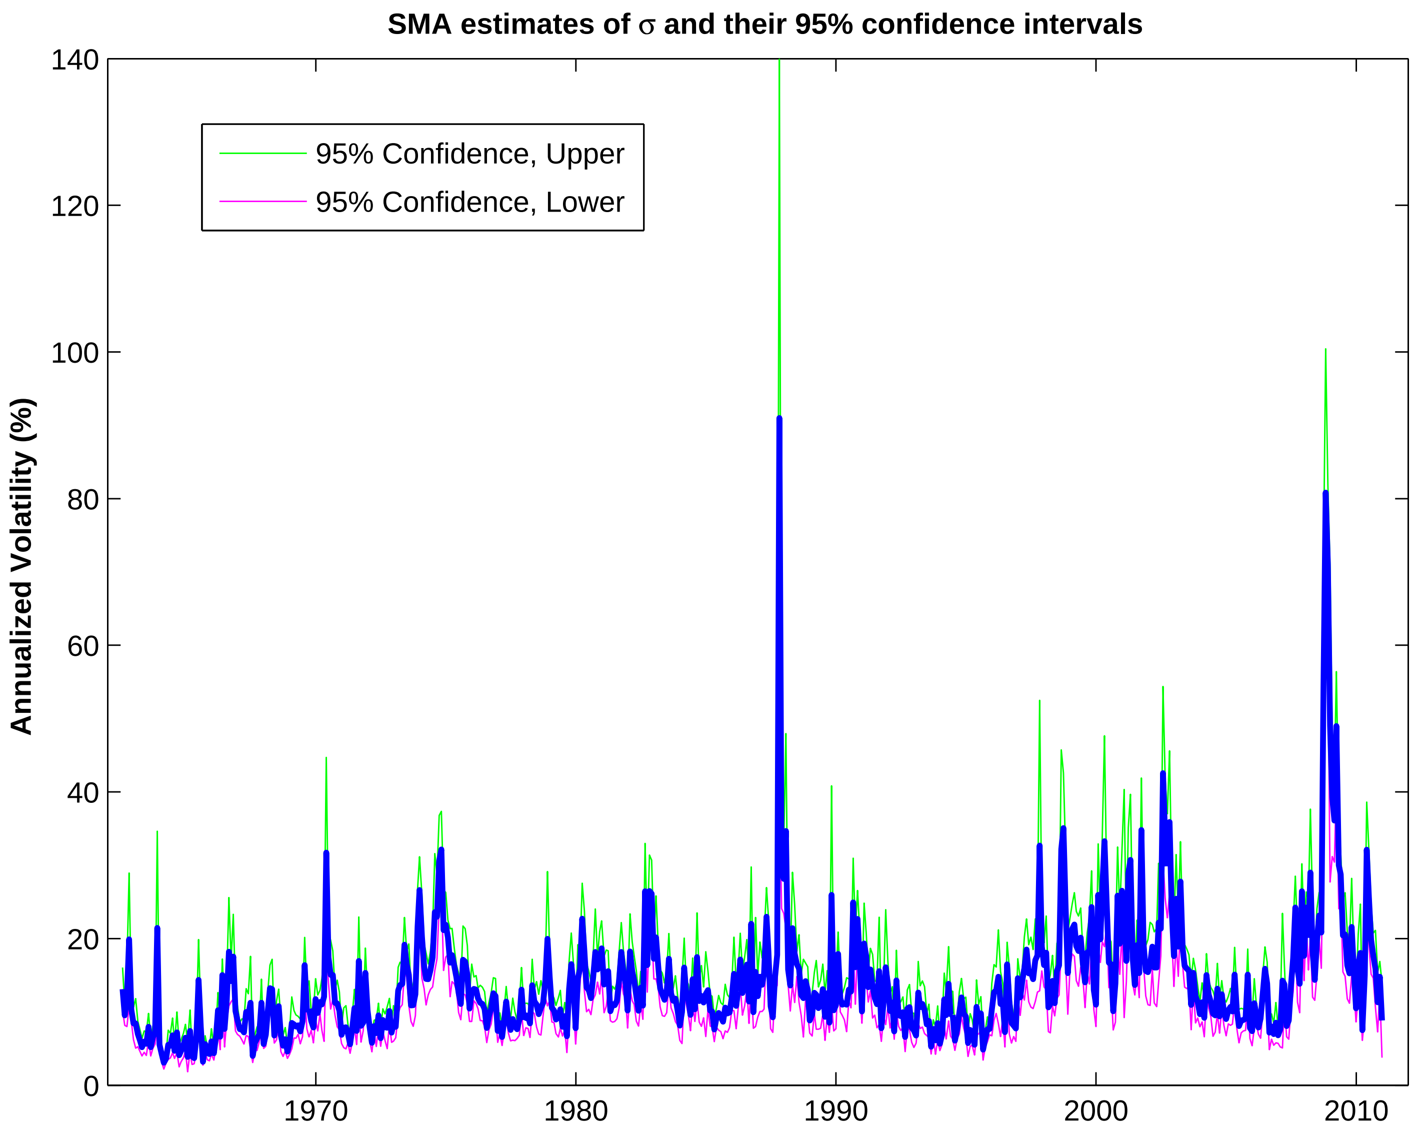

What about SMA mean estimates?

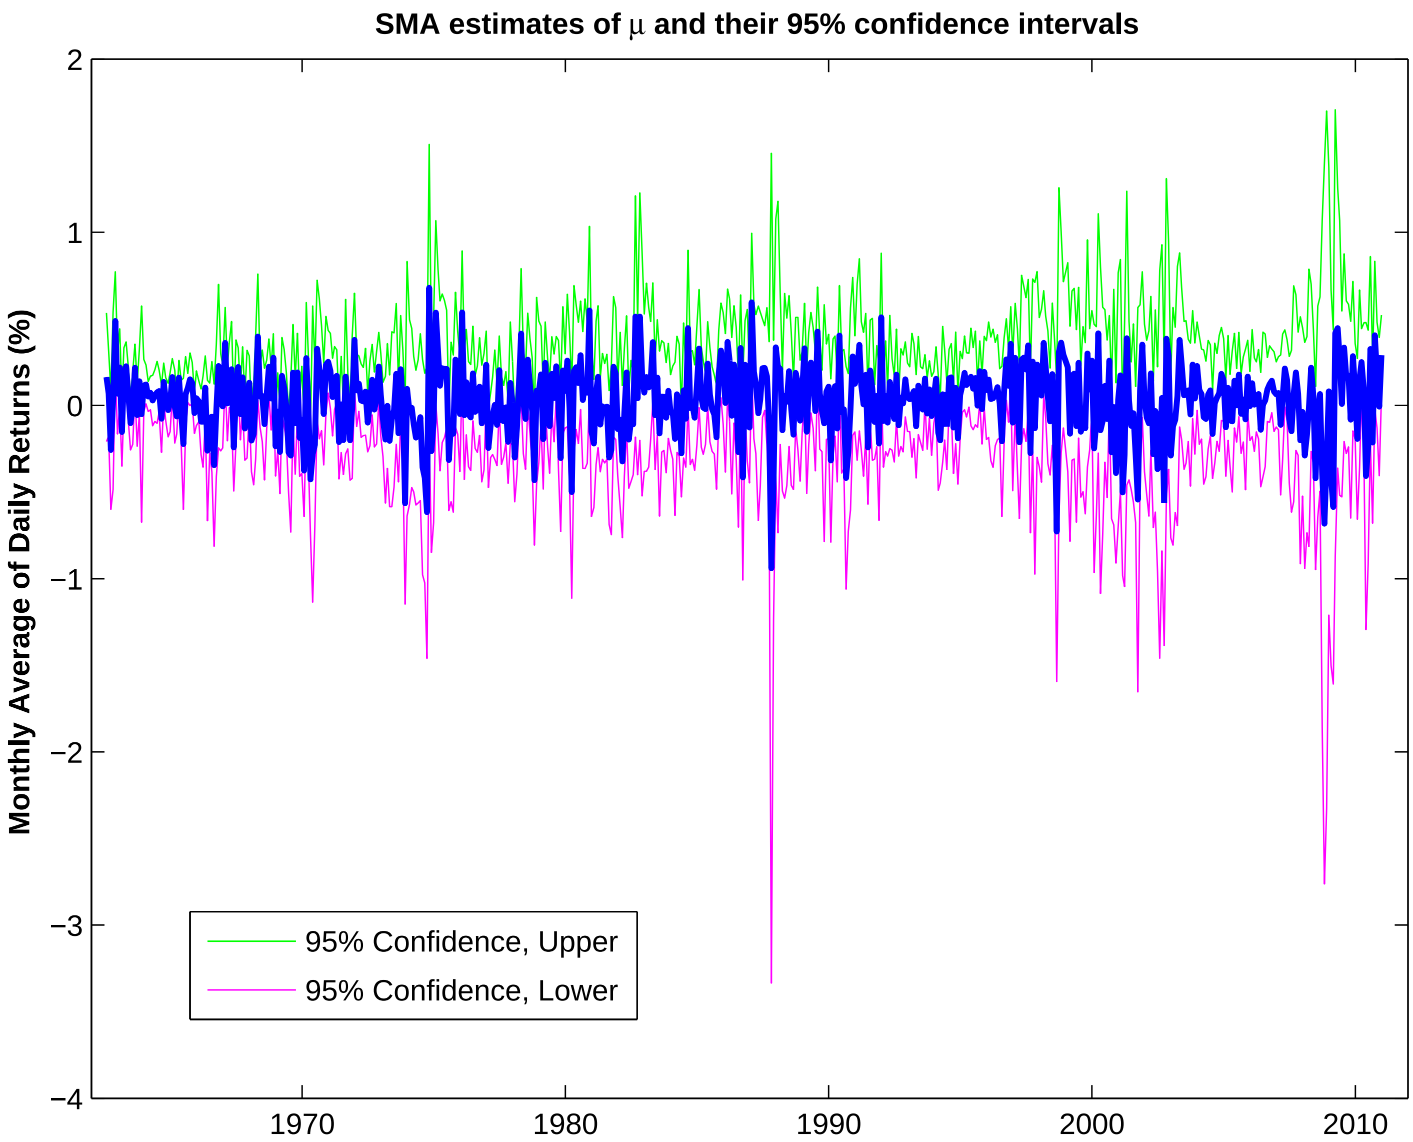

Why does volatility change over time?

- If the rate of information arrival is time-varying, so is the rate of price adjustment, causing volatility to change over time.
- The time-varying volatility of the market return is related to the time-varying volatility of a variety of economic variables, including inflation, unemployment rate, money growth and industrial production.
- Stock market volatility increases with financial leverage: a decrease in stock price causes an increase in financial leverage, causing volatility to increase.
- Investors' sudden changes of risk attitudes, changes in market liquidity, and temporary imbalance of supply and demand could all cause market volatility to change over time.

Time-varying volatility and business cycles

(The shaded areas are the NBER dated peak to trough)

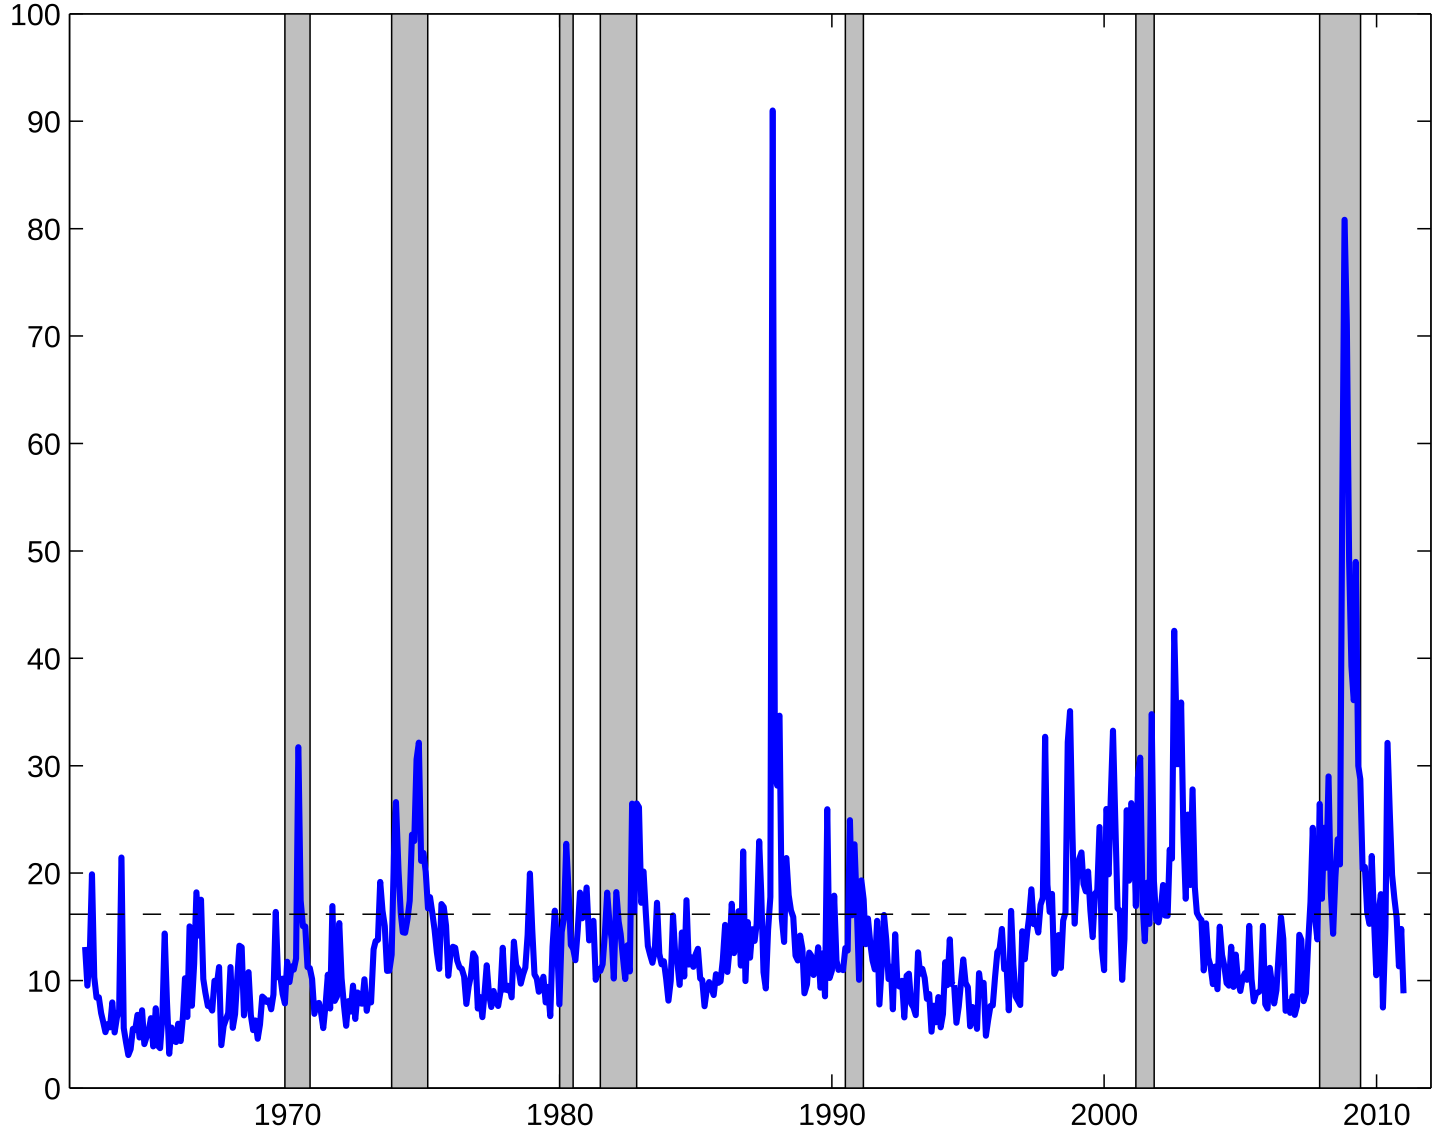

SMA vs. Option-Implied

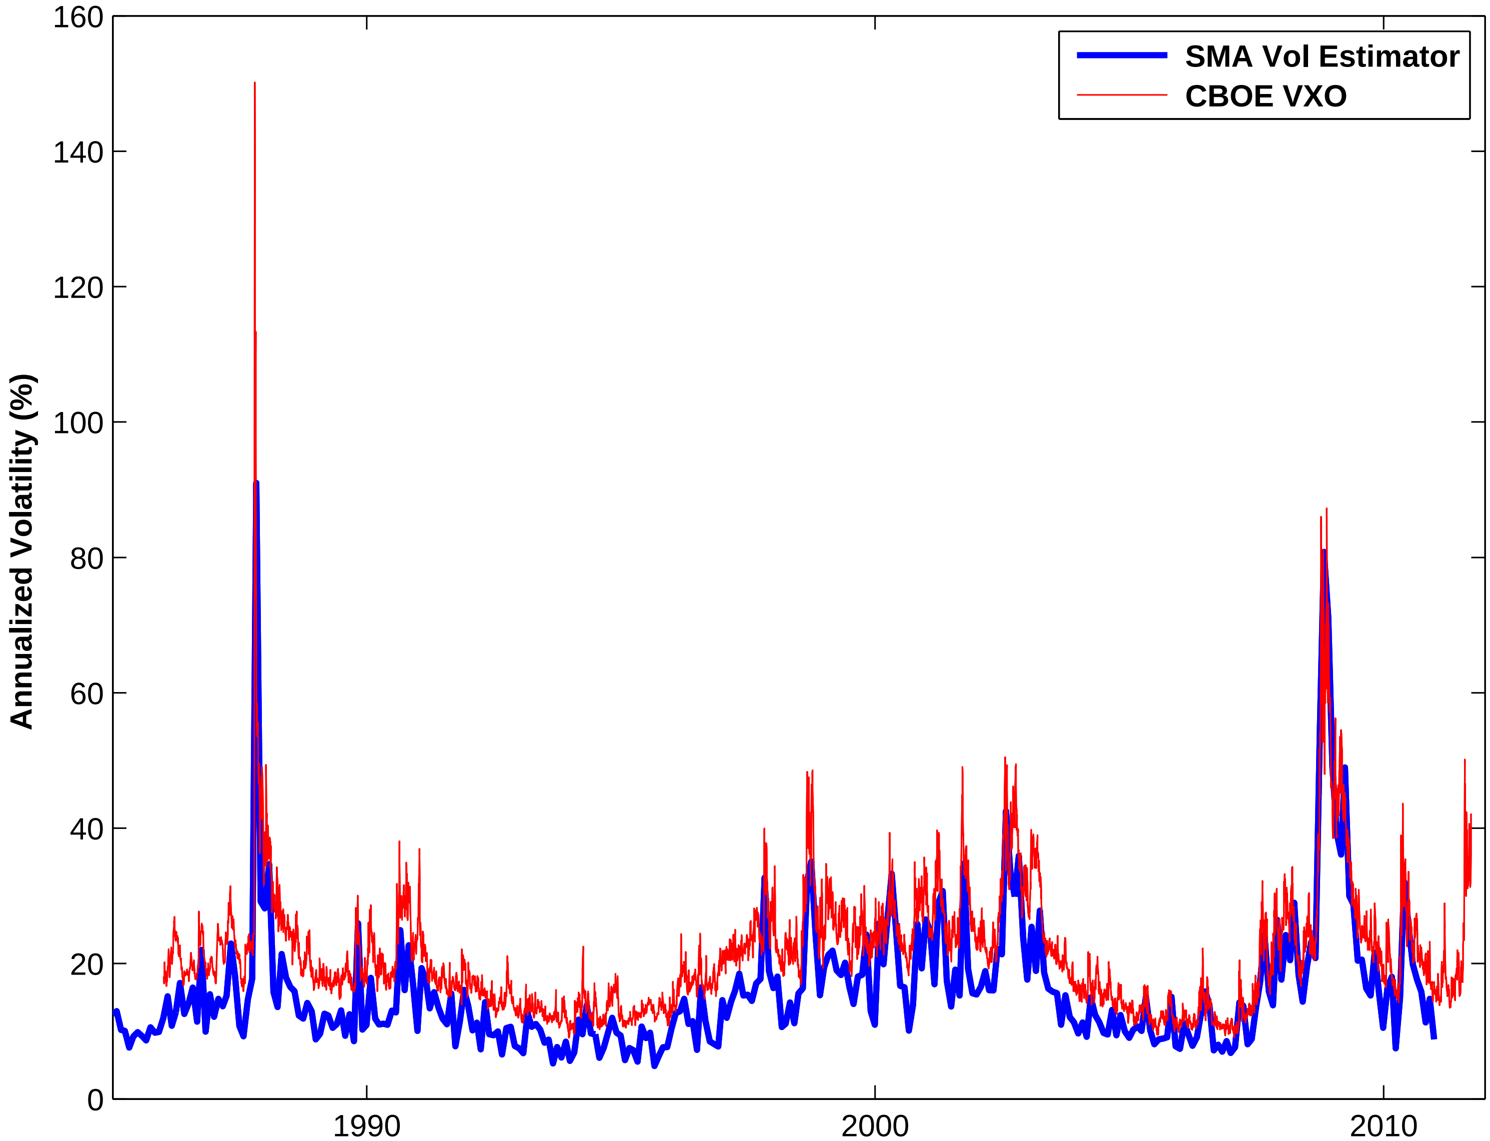

VXO vs. VIX

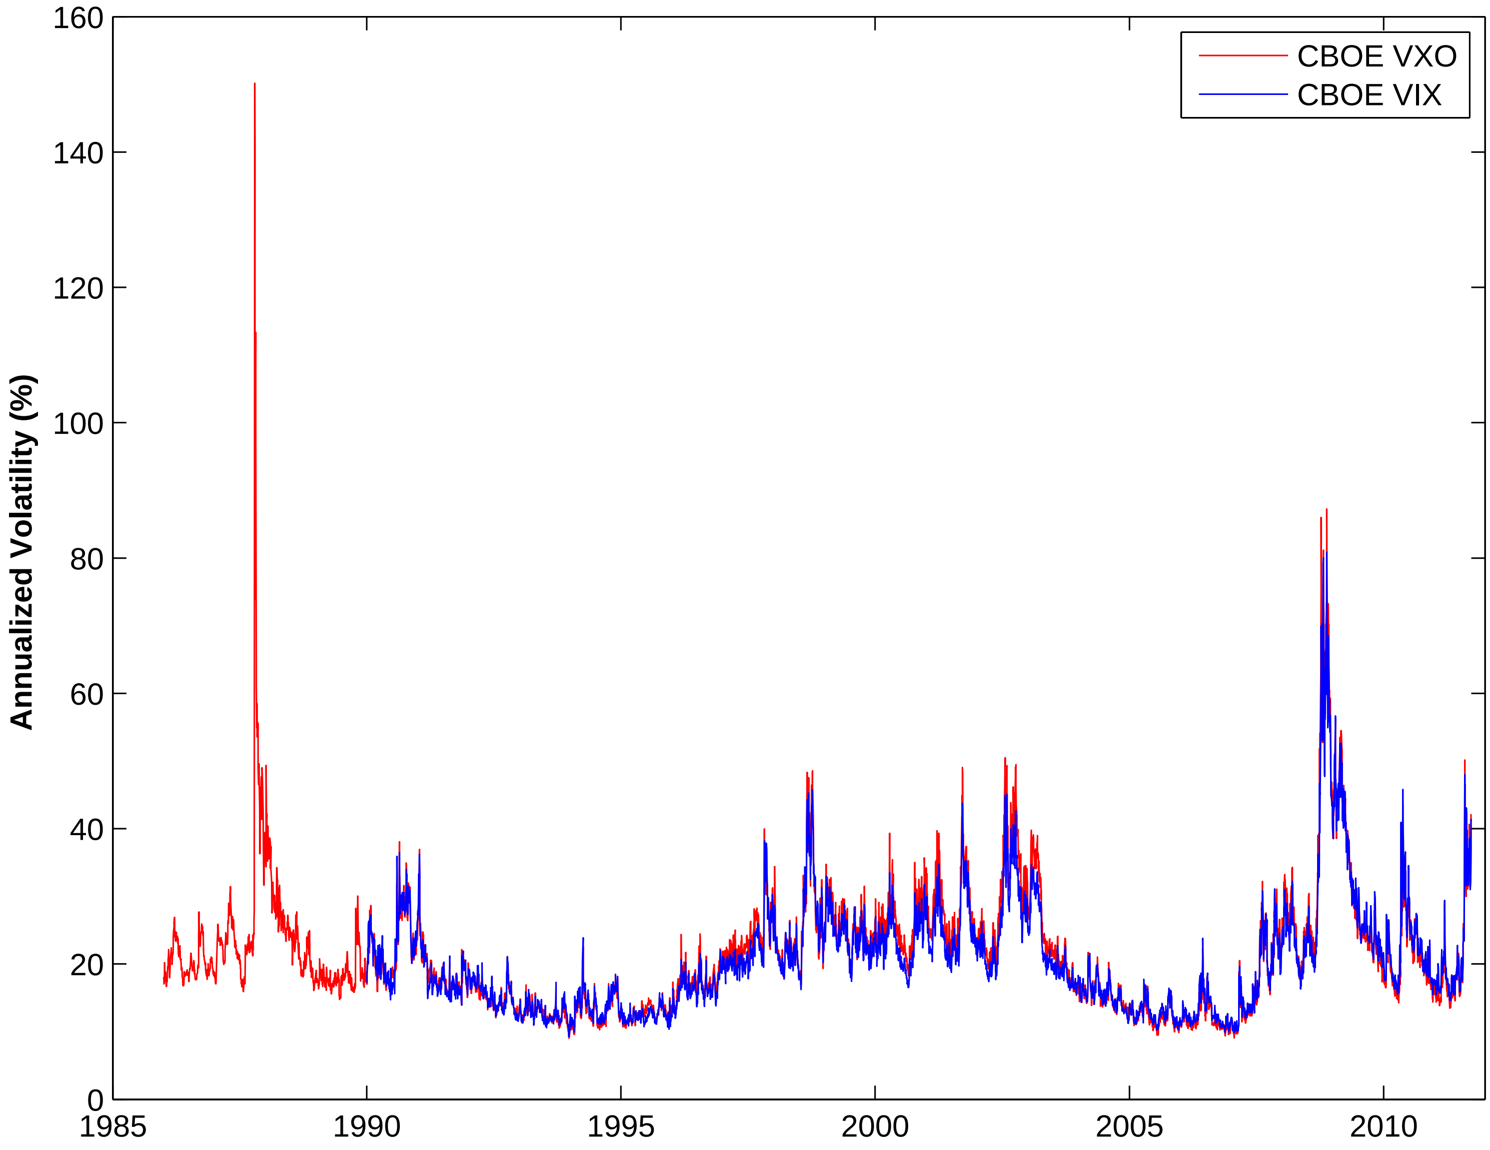
160
140
120
100
80
60
40
20
0
Annualized Volatility (%)
1985
1990
1995
2000
2005
2010
CBOE VXO
CBOE VIX

Exponentially weighted moving average model

- The simple moving average (SMA) model fixes a time window and applies equal weight to all observations within the window.
- In the exponentially weighted moving average (EWMA) model, the more recent observation carries a higher weight in the volatility estimate.
- The relative weight is controlled by a decay factor λ.
- Suppose R_t is today's realized return, R_{t-1} is yesterday's, and R_{t-n} is the daily return realized n days ago. Volatility estimate σ:

Equally Weighted	**Exponentially Weighted**
$\sqrt{\frac{1}{N}\sum_{n=0}^{N-1}\left(R_{t-n}\right)^2}$	$\sqrt{(1-\lambda)\sum_{n=0}^{N-1}\lambda^n\left(R_{t-n}\right)^2}$

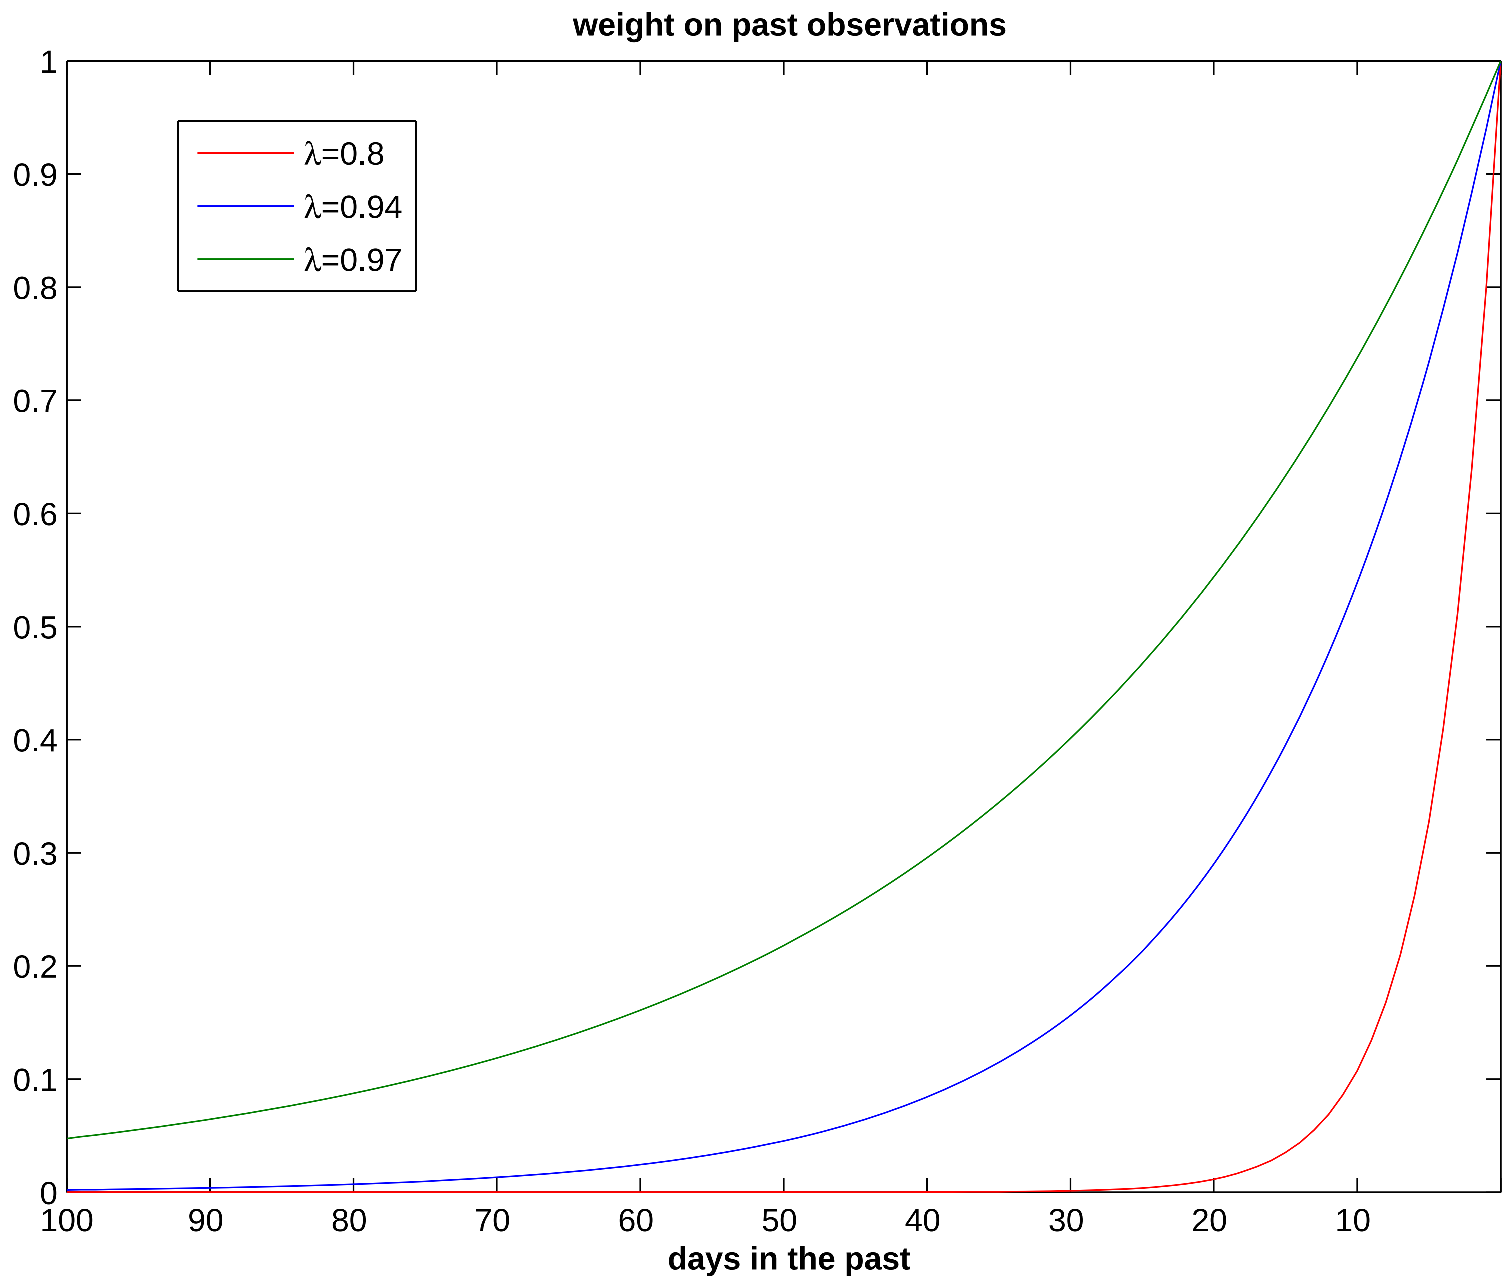
weight on past observations
λ=0.8
λ=0.94
λ=0.97
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
100
90
80
70
60
50
40
30
20
10
days in the past

Calculating equally and exponentially weighted volatility

Date	A: Return USD/DEM (%)	B: Return squared (%)	C: Equal weight ($T = 20$)	D: Exponential weight ($\lambda = 0.94$)	Volatility: Equally weighted, $B \times C$	Volatility: Exponentially weighted, $B \times D$
28-Mar-96	0.634	0.402	0.05	0.019	0.020	0.007
29-Mar-96	0.115	0.013	0.05	0.020	0.001	0.000
1-Apr-96	-0.460	0.211	0.05	0.021	0.011	0.004
2-Apr-96	0.094	0.009	0.05	0.022	0.000	0.000
3-Apr-96	0.176	0.031	0.05	0.024	0.002	0.001
4-Apr-96	-0.088	0.008	0.05	0.025	0.000	0.000
5-Apr-96	-0.142	0.020	0.05	0.027	0.001	0.001
8-Apr-96	0.324	0.105	0.05	0.029	0.005	0.003
9-Apr-96	-0.943	0.889	0.05	0.030	0.044	0.027
10-Apr-96	-0.528	0.279	0.05	0.032	0.014	0.009
11-Apr-96	-0.107	0.011	0.05	0.034	0.001	0.000
12-Apr-96	-0.160	0.026	0.05	0.037	0.001	0.001
15-Apr-96	-0.445	0.198	0.05	0.039	0.010	0.008
16-Apr-96	0.053	0.003	0.05	0.041	0.000	0.000
17-Apr-96	0.152	0.023	0.05	0.044	0.001	0.001
18-Apr-96	-0.318	0.101	0.05	0.047	0.005	0.005
19-Apr-96	0.424	0.180	0.05	0.050	0.009	0.009
22-Apr-96	-0.708	0.501	0.05	0.053	0.025	0.027
23-Apr-96	-0.105	0.011	0.05	0.056	0.001	0.001
24-Apr-96	-0.257	0.066	0.05	0.060	0.003	0.004
			Standard deviation:	Equally weighted	0.393	
				Exponentially weighted		0.333

Source: RiskMetrics—Technical Document

SMA and EWMA Estimates after a Crash

Chart 5.2
Log price changes in GBP/DEM and VaR estimates (1.65σ)

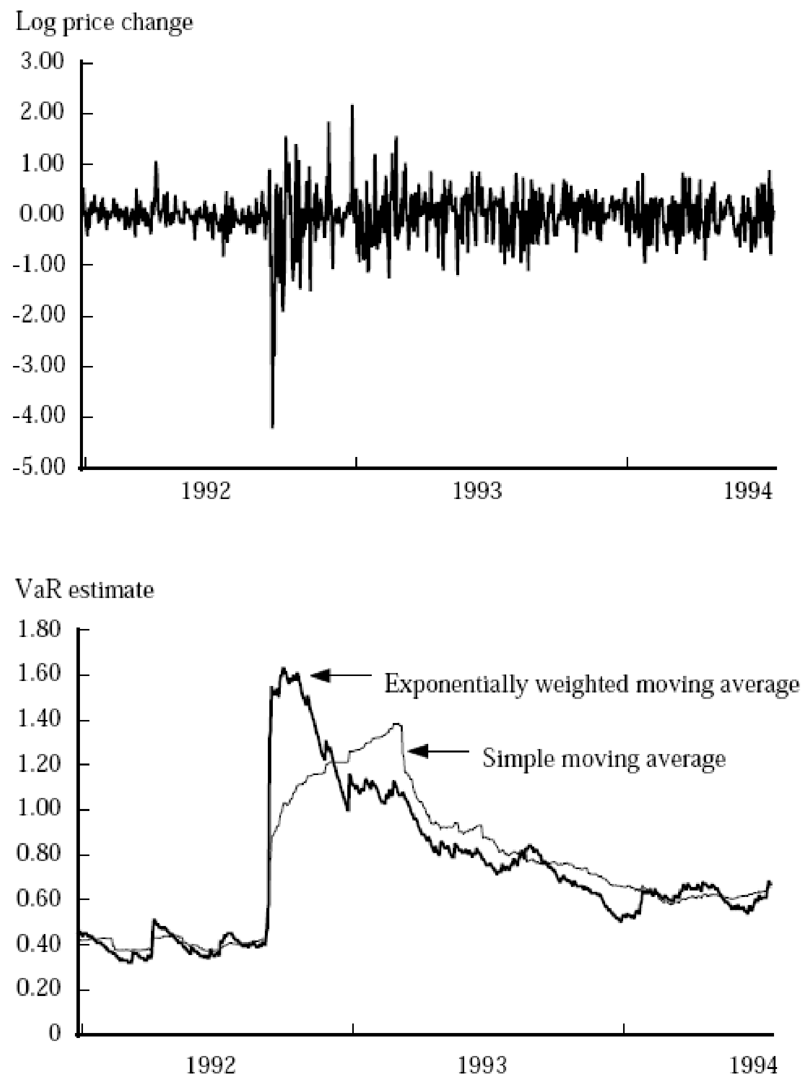

Source: J.P.Morgan/Reuters RiskMetrics — Technical Document, 1996

Computing EWMA recursively

- One attractive feature of the exponentially weighted estimator is that it can be computed recursively.
- You will appreciate this convenience if you have to compute the EWMA volatility estimator in Excel.
- Let σ_t be the EWMA volatility estimator using all the information available on day $t-1$ for the purpose of forecasting the volatility on day t.
- Moving one day forward, it's now day t. After the day is over, we observe the realized return R_t.
- We now need to update our EWMA volatility estimator σ_{t+1} using the newly arrived information (i.e. R_t). It turns out that we can do so by

$$\sigma_{t+1}^2 = \lambda\,\sigma_t^2 + (1-\lambda)\,R_t^2$$

What about the first observation?

- The recursive formula has to start from the beginning:

$$\sigma_2^2 = \lambda\, \sigma_1^2 + (1-\lambda)\, R_1^2$$

  So what to use for σ_1?

- In practice, the choice of σ_1 does not matter in any significant way after running the iterative process long enough:

$$\begin{aligned}
\sigma_3^2 &= \lambda\, \sigma_2^2 + (1-\lambda)\, R_2^2 \\
&= \lambda^2\, \sigma_1^2 + (1-\lambda)\left(\lambda R_1^2 + R_2^2\right) \\
\sigma_4^2 &= \lambda\, \sigma_3^2 + (1-\lambda)\, R_3^2 \\
&= \lambda^3\, \sigma_1^2 + (1-\lambda)\left(\lambda^2 R_1^2 + \lambda R_2^2 + R_3^2\right) \\
\ldots & \\
\sigma_t^2 &= \lambda^{t-1}\, \sigma_1^2 + (1-\lambda)\left(\lambda^{t-2} R_1^2 + \ldots + R_{t-1}^2\right)
\end{aligned}$$

- A good idea is to have the iterative process run for a while (say a few months) before recording volatility estimates.
- (Prof. Pan's Choice:) I like to set $\sigma_1 =$std(R), which is the "unconditional" or sample standard deviation of R. The logic is that if I don't have any information about σ_1 at the beginning of the volatility estimation, I might as well use the unconditional estimate of σ.
- (The industry practice:) It is typical to set $\sigma_2^2 = R_1^2$ and start the recursive process from σ_3. The rationale is that σ_1 is unknowable and the only data we have at time 1 is R_1. So R_1^2 is our best estimate for σ_2^2. This approach is adopted by most of the practitioners, including RiskMetrics.

Dating Convention for σ_t

- The dating convention adopted by most people:

$$\sigma_{t+1}^2 = \lambda\,\sigma_t^2 + (1-\lambda)\,R_t^2$$

  The rationale is that this σ is estimated for the purpose of forecasting the next period's volatility. So it should be dated as σ_{t+1}.

- (Prof. Pan's Choice:) I actually like to use

$$\sigma_t^2 = \lambda\,\sigma_{t-1}^2 + (1-\lambda)\,R_t^2$$

  The rationale is that at time t, I am forming an estimate σ_t using all of the information available to me at time t.

- I will always use the main-stream approach and date it by σ_{t+1}.

Decay factor, Strong or Weak?

- A strong decay factor (that is, small λ) underweights the far away events more strongly, making the effective sample size smaller.
- A strong decay factor improves on the timeliness of the volatility estimate, but that estimate could be noisy and suffers in precision.
- On the other hand, a weak decay factor improves on the smoothness and precision, but that estimate could be sluggish and slow in response to changing market conditions.
- So there is a tradeoff.

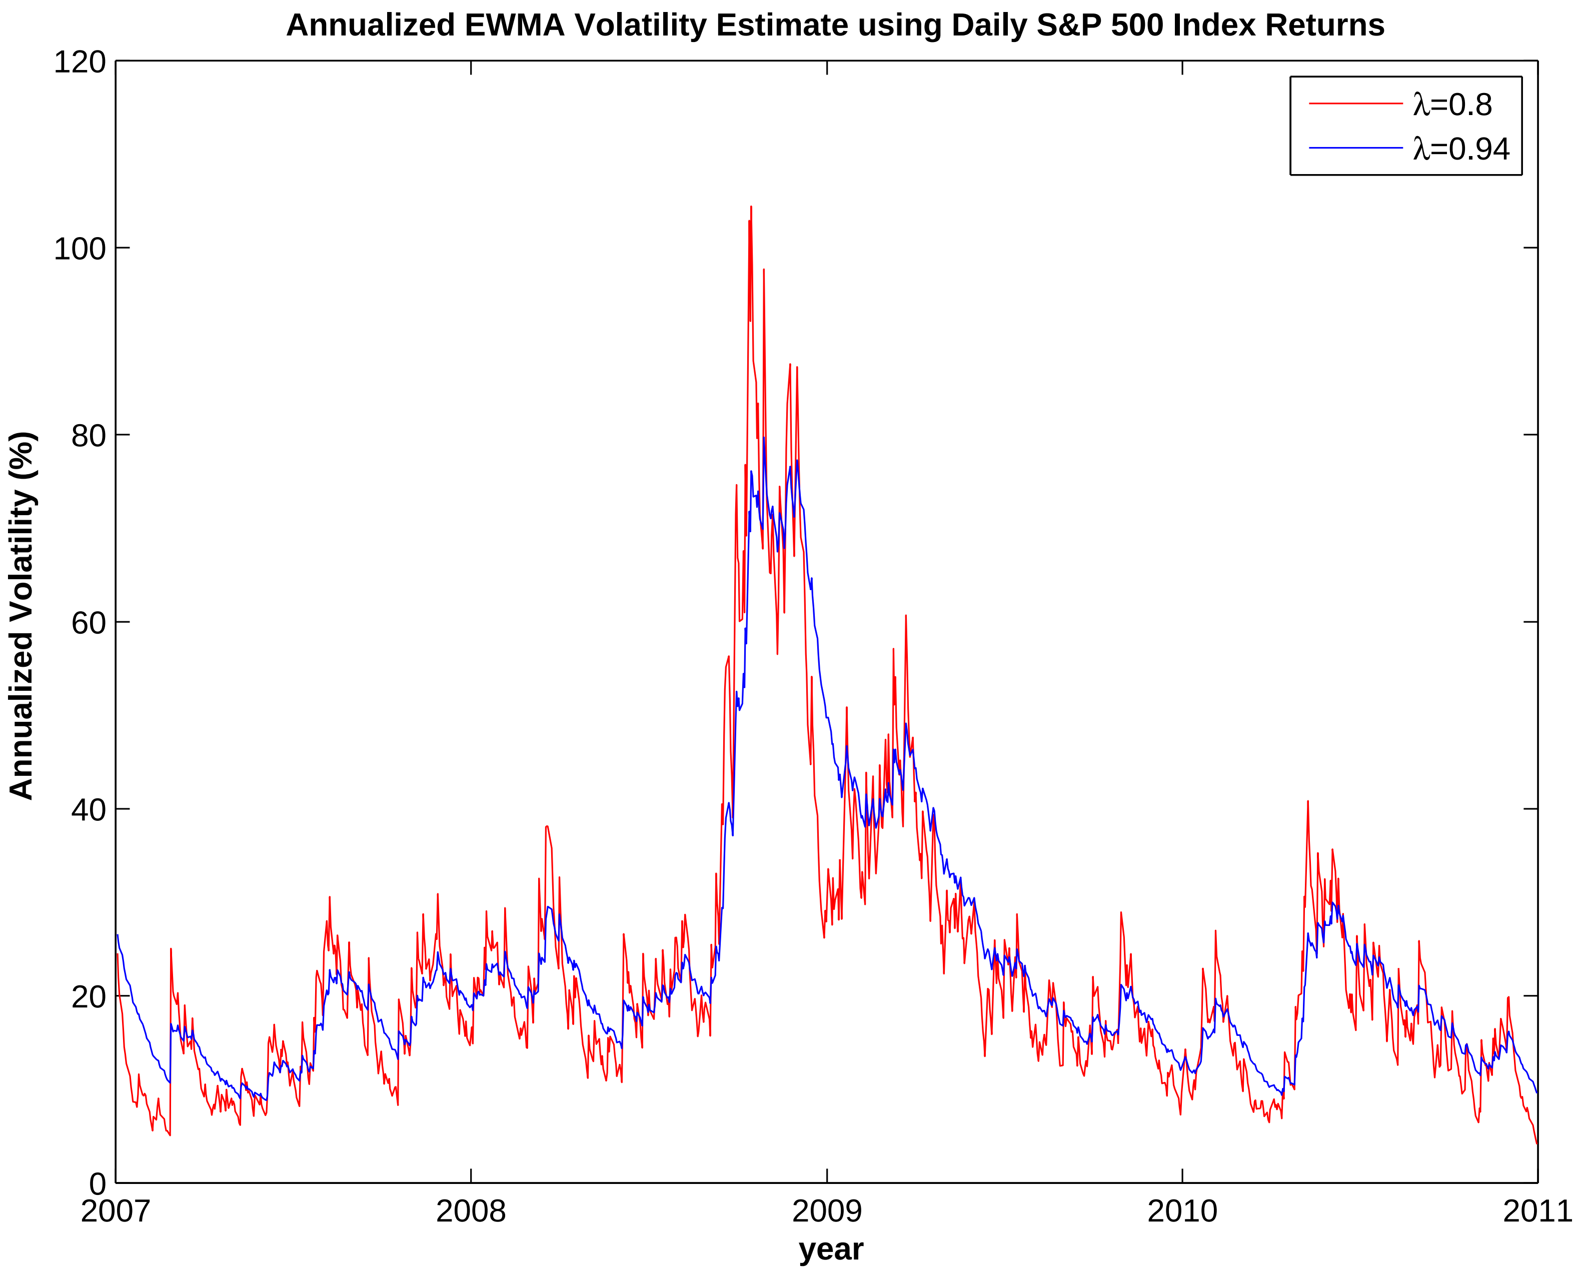
Annualized EWMA Volatility Estimate using Daily S&P 500 Index Returns
Annualized Volatility (%)
year
λ=0.8
λ=0.94
0
20
40
60
80
100
120
2007
2008
2009
2010
2011

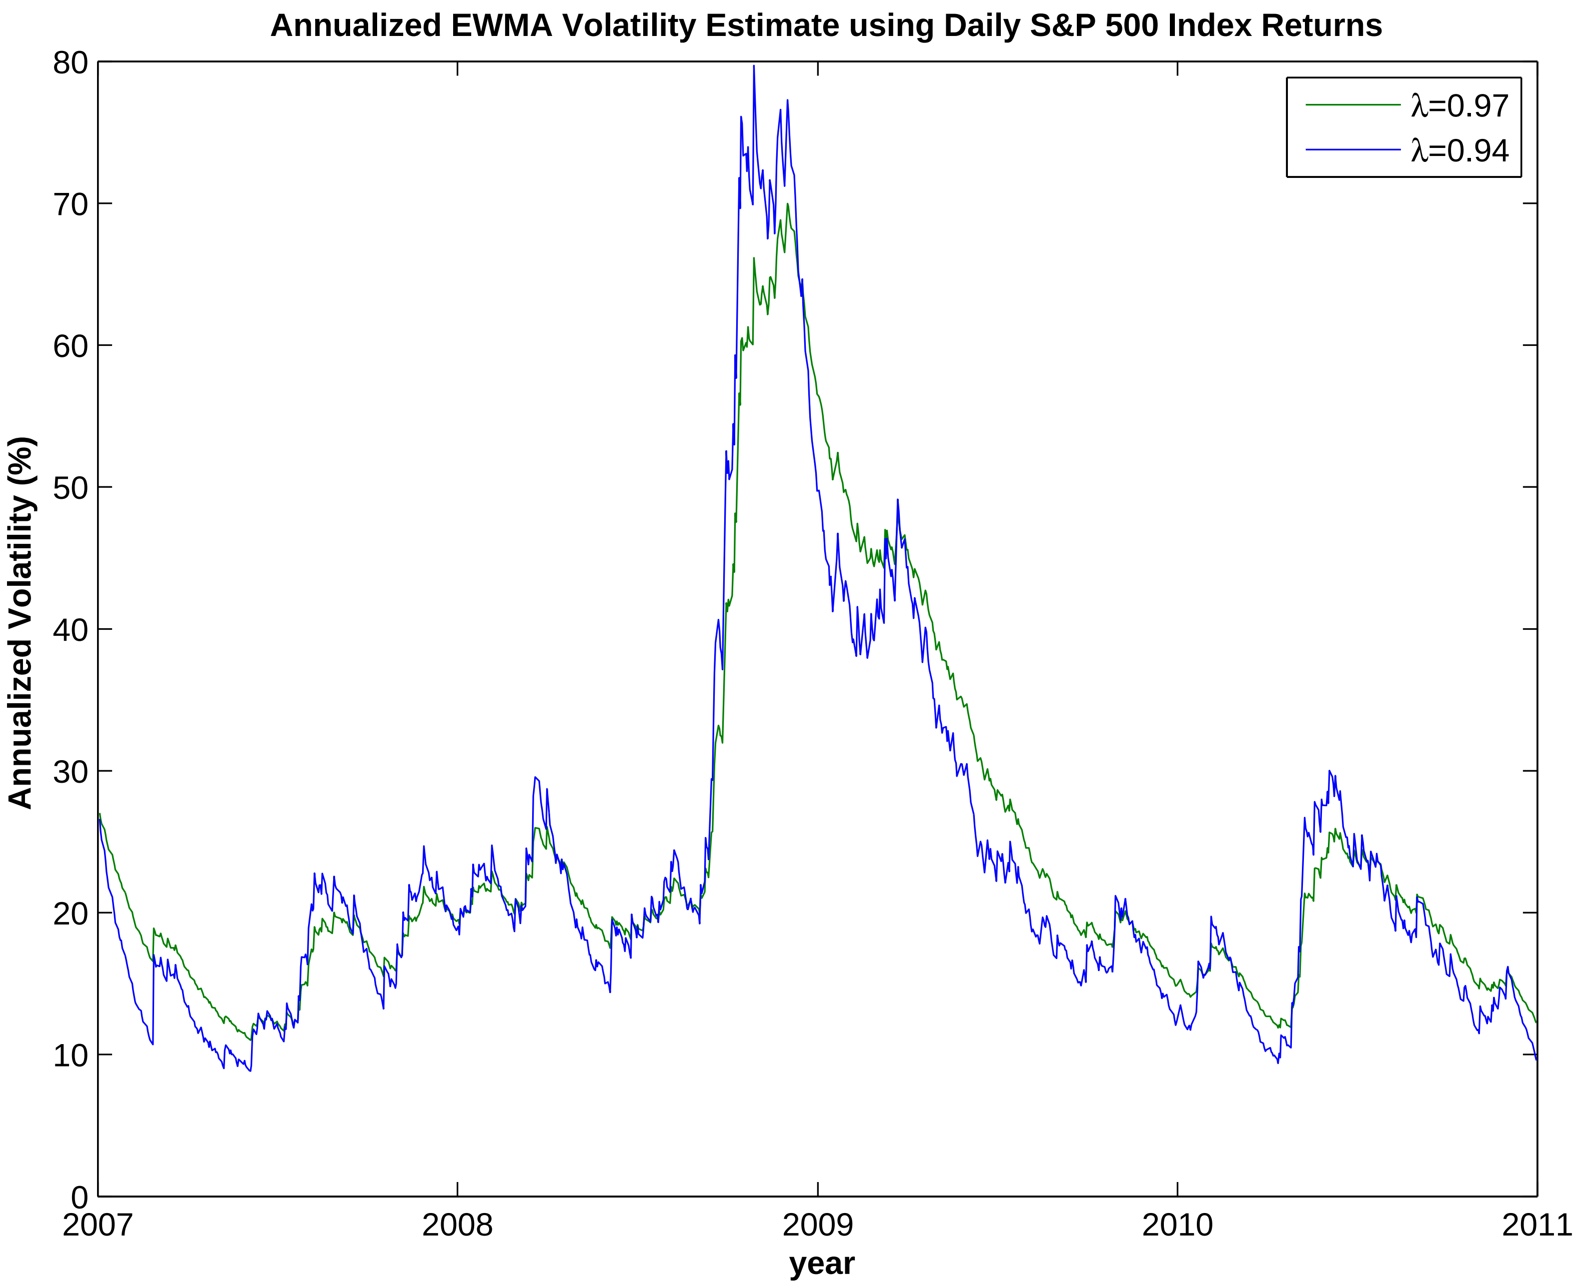
Annualized EWMA Volatility Estimate using Daily S&P 500 Index Returns
Annualized Volatility (%)
year
λ=0.97
λ=0.94
0
10
20
30
40
50
60
70
80
2007
2008
2009
2010
2011

Picking the optimal decay factor based on volatility forecast

- RiskMetrics sets $\lambda = 0.94$ in estimating volatility and correlation. One of their key criteria is to minimize the forecast error.
- We form σ_{t+1} on day t in order to forecast the next-day volatility. So after observing R_{t+1}, we can check how good σ_{t+1} is in doing its job.
- This leads to the daily root mean squared prediction error

$$RMSE = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(R_{t+1}^2 - \sigma_{t+1}^2\right)^2}$$

- The deciding factor of RMSE is our choice of λ. For my running example (daily S&P 500 index returns 2007-2010):

λ	0.80	0.9075*	0.94	0.97
RMSE	8.1844	8.0124	8.0544	8.2444

Maximum Likelihood Estimation

- The gold standard in any estimation is maximum likelihood estimation, because it is the most efficient method. So let's see what MLE has to say about the optimal λ.

- We assume that conditioning on the volatility estimate σ_{t+1}, the stock return R_{t+1} is normally distributed:

$$f\left(R_{t+1}|\sigma_{t+1}\right) = \frac{1}{\sqrt{2\pi}\sigma_{t+1}} e^{-\frac{R_{t+1}^2}{2\sigma_{t+1}^2}}$$

- Take natural log of f:

$$\ln f\left(R_{t+1}|\sigma_{t+1}\right) = -\ln \sigma_{t+1} - \frac{R_{t+1}^2}{2\sigma_{t+1}^2}$$

I dropped 2π since it will not affect anything we will do later.

- We now add them up to get what econometricians call log-likelihood (llk):

$$\text{llk} = -\sum_{t=1}^{T} \left(\ln \sigma_{t+1} + \frac{R_{t+1}^2}{2\sigma_{t+1}^2} \right)$$

- The only deciding factor in llk is our choice of λ. It turns out that the best λ is the one that maximizes llk.

- In practice, we take -llk and minimize -llk instead of maximizing llk.

- For my running example (daily S&P 500 index return 2007-2010), I find the optimal λ that minimizes -llk is 0.9320. Not exactly the same as the optimal λ of 0.9075 that minimizes RMSE, but these two are reasonably close.

The Surface of Planet MLE

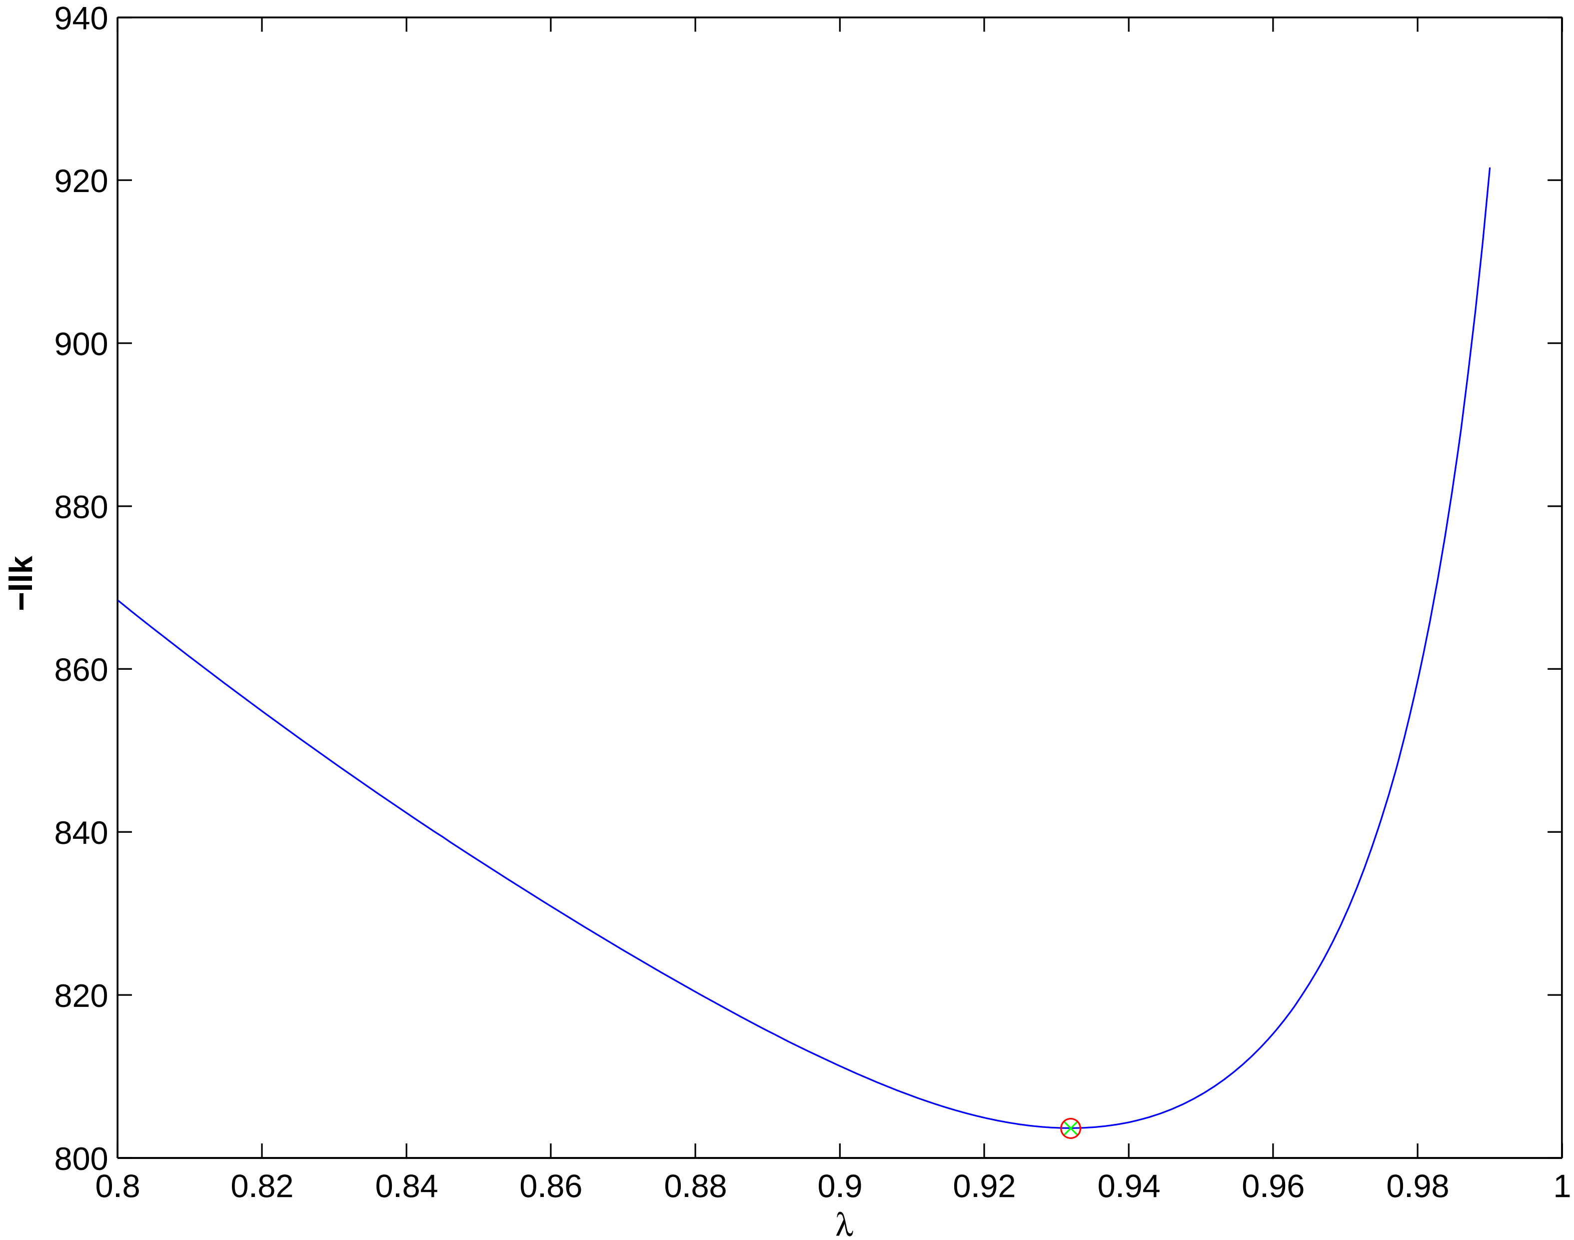

The ARCH and GARCH models

- The ARCH model, autoregressive conditional heteroskedasticity, was proposed by Professor Robert Engle in 1982. The GARCH model is a generalized version of ARCH.
- ARCH and GARCH are statistical models that capture the time-varying volatility:
$$\sigma_{t+1}^2 = a_0 + a_1\, R_t^2 + a_2\, \sigma_t^2$$
- As you can see, it is very similar to the EWMA model. In fact, if we set $a_0 = 0$, $a_2 = \lambda$, and $a_1 = 1 - \lambda$, we are doing the EWMA model.
- So what's the value added? This model has three parameters while the EWMA has only one. So it offers more flexibility (e.g., allows for mean reversion and better captures volatility clustering).
- But I think EWMA is good enough for us, for now.

KUNGL. VETENSKAPSAKADEMIEN
THE ROYAL SWEDISH ACADEMY OF SCIENCES

English
French
German
Swedish

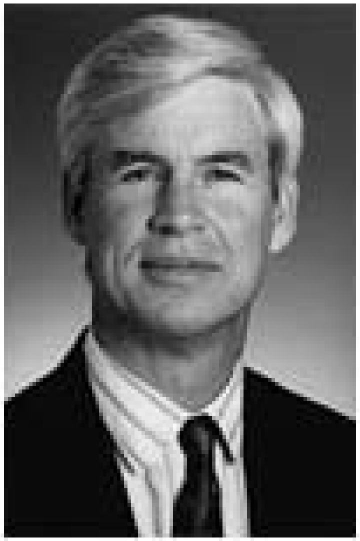

Robert F. Engle III

1/2 of the prize

USA

New York University
New York, NY, USA

b. 1942

Press Release: The Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel 2003

8 October 2003

The Royal Swedish Academy of Sciences has decided that the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel, 2003, is to be shared between

Robert F. Engle
New York University, USA

"for methods of analyzing economic time series with time-varying volatility (ARCH)"

and

Clive W. J. Granger
University of California at San Diego, USA

"for methods of analyzing economic time series with common trends (cointegration)".

EWMA covariances and correlations

- Our goal is to create the variance-covariance matrix for the key risk factors influencing our portfolio.
- For the moment, let's suppose that there are only two risk factors affecting our portfolio.
- Let R_t^A and R_t^B be the day-t realized returns of these two risk factors. The covariance between A and B:

$$\text{cov}_{t+1} = \lambda \, \text{cov}_t + (1-\lambda)\, R_t^A \times R_t^B$$

- And their correlation:

$$\text{corr}_{t+1} = \frac{\text{cov}_{t+1}}{\sigma_{t+1}^A \sigma_{t+1}^B}\,,$$

where σ_{t+1}^A and σ_{t+1}^B are the EWMA volatility estimates.

The negative correlation between R^M and ΔVIX

- Monthly returns R_t^M on the stock market portfolio is highly negatively correlated with monthly changes in VIX: -69.41%.
- Now let's apply our EWMA approach, which will give us a time-series of correlations between these two risk factors.
- We see an interesting time-series pattern of the negative correlation between daily stock returns and daily changes in VIX.
- In particular, this correlation has become more negative in recent years.
- (CBOE started to offer futures trading on VIX on March 26, 2004.)

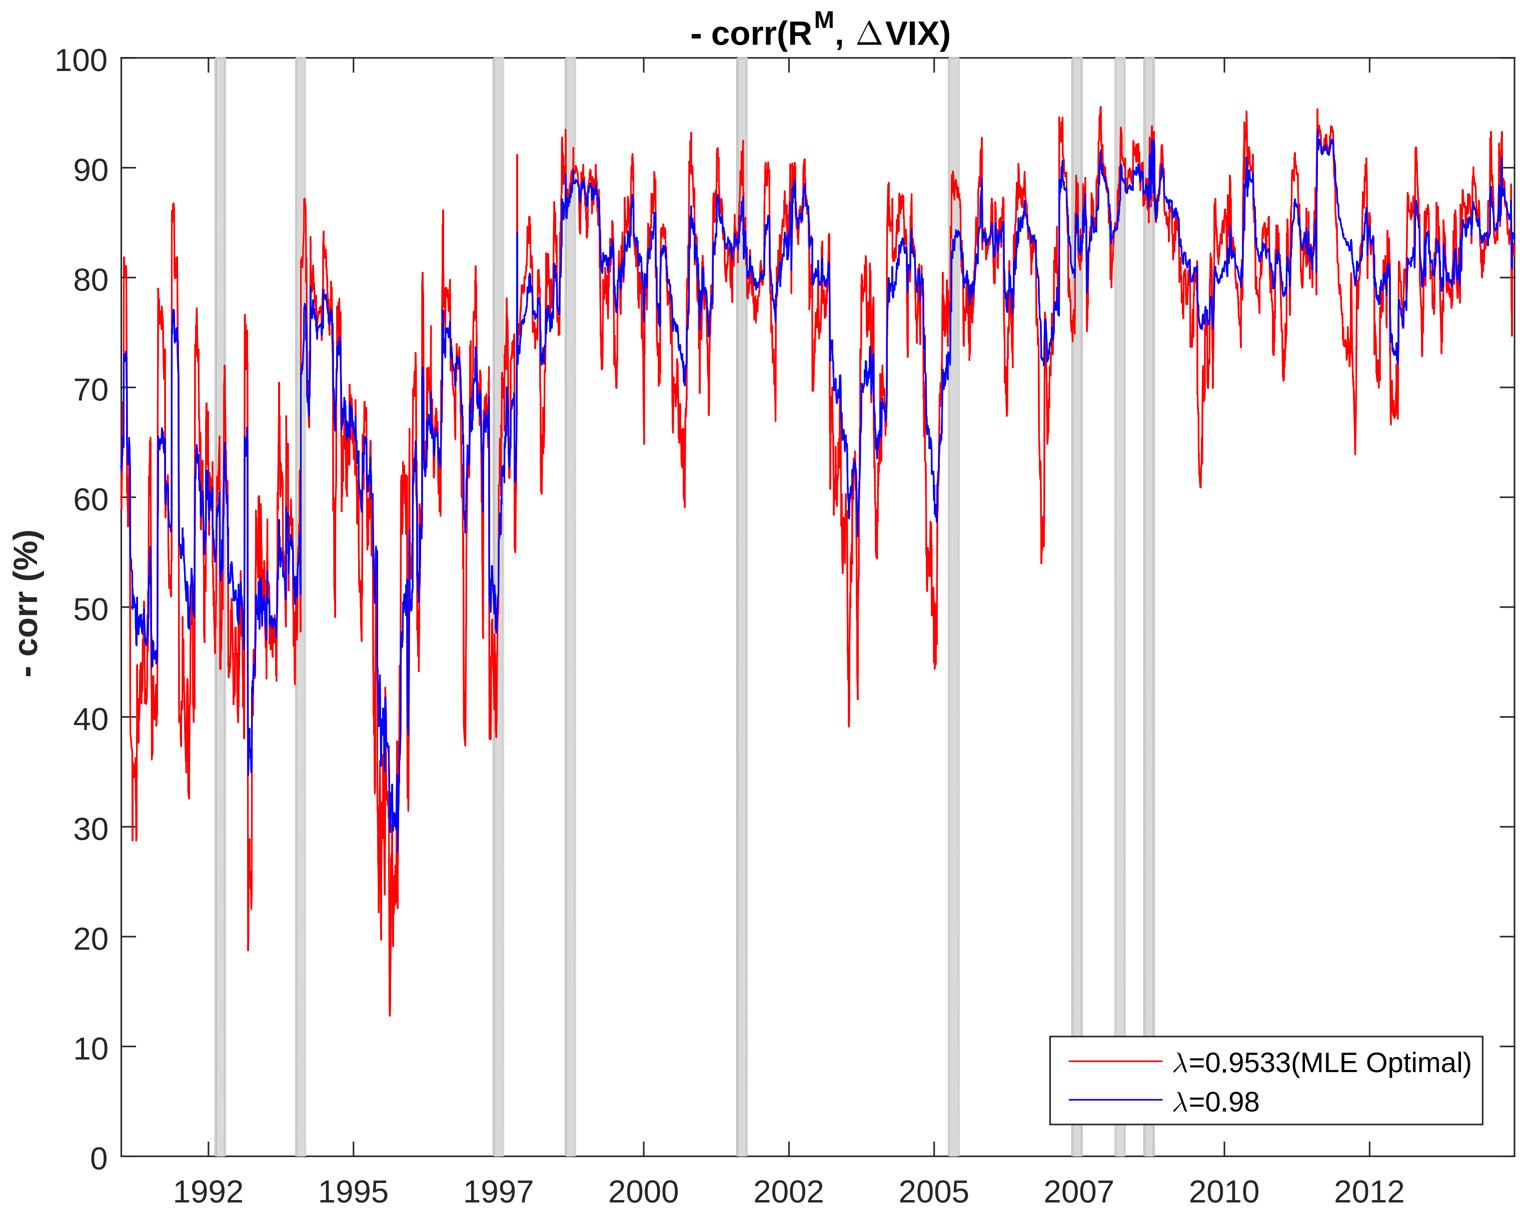

- corr(R^M, ΔVIX)
- corr (%)
100
90
80
70
60
50
40
30
20
10
0
1992
1995
1997
2000
2002
2005
2007
2010
2012
λ=0.9533(MLE Optimal)
λ=0.98

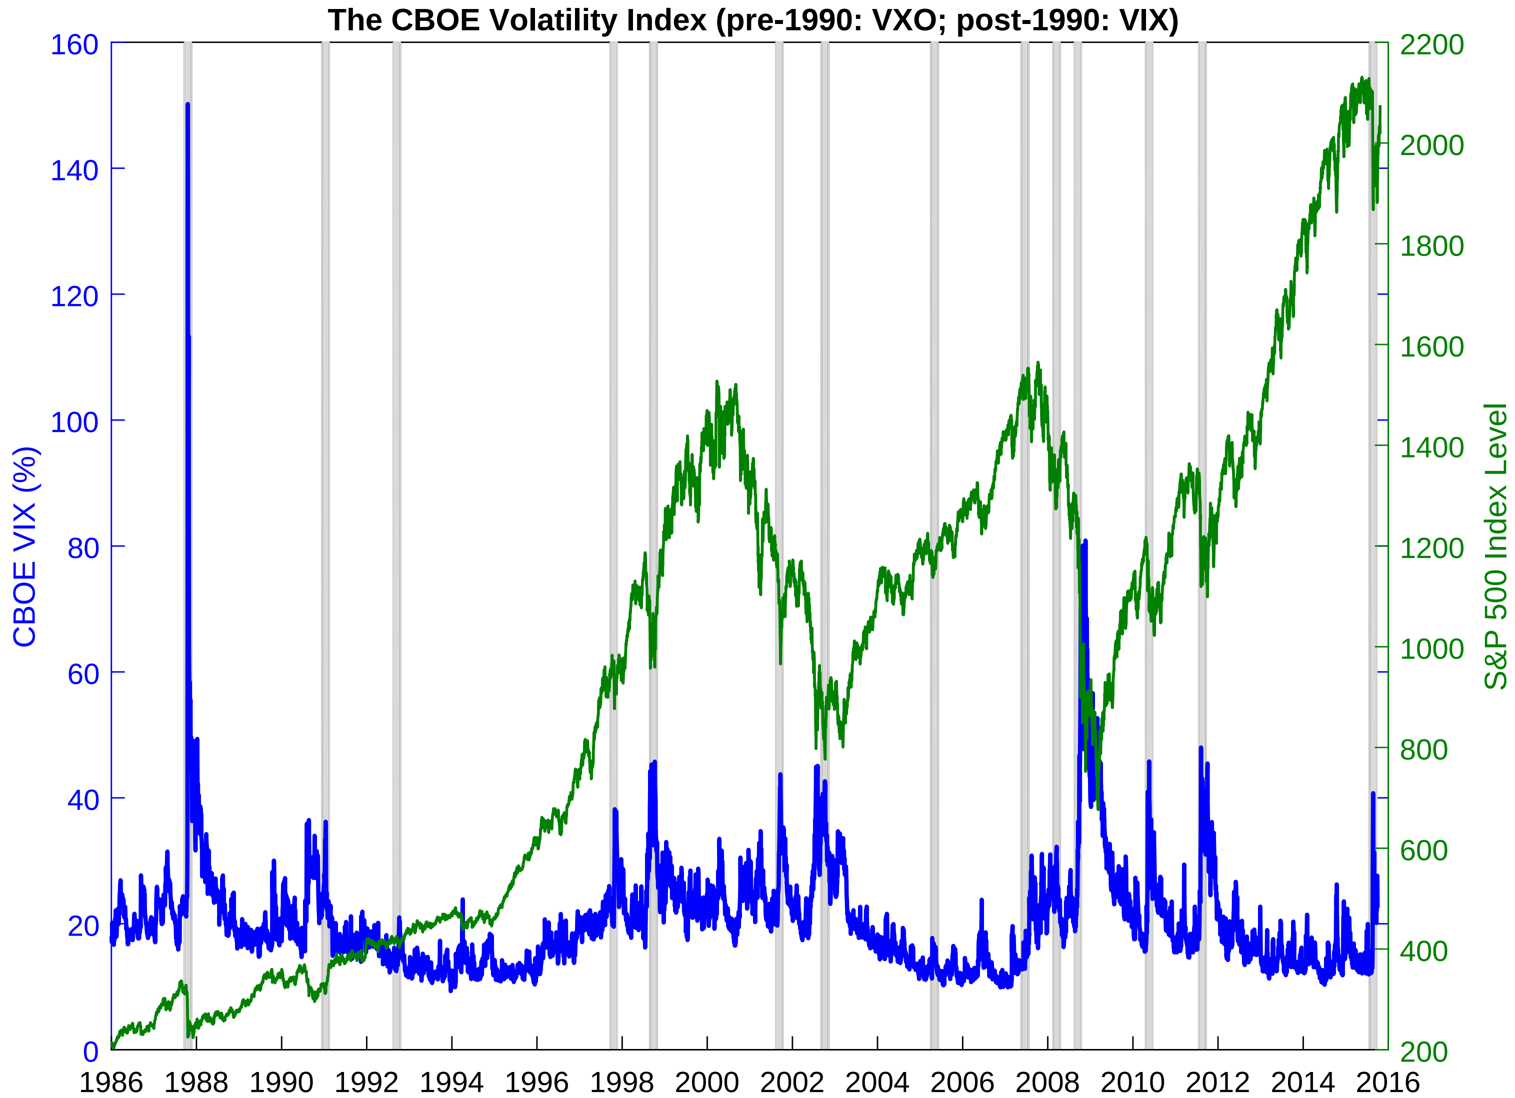
The CBOE Volatility Index (pre-1990: VXO; post-1990: VIX)
CBOE VIX (%)
S&P 500 Index Level
0
20
40
60
80
100
120
140
160
200
400
600
800
1000
1200
1400
1600
1800
2000
2200
1986
1988
1990
1992
1994
1996
1998
2000
2002
2004
2006
2008
2010
2012
2014
2016

Calculating Volatility for a Portfolio

- Suppose that our portfolio has two important risk factors, whose daily returns are R^A and R^B, respectively.
- Performing risk mapping using individual positions, the portfolio weights on these two risk factors are w_A and w_B.
- Let's focus only on the risky part of our portfolio and leave out the cash part. So let's normalize the weights so that $w^A + w^B = 1$. Let's assume our risk portfolio has a market value of $100 million today.
- We apply EWMA to get time-series of their volatility estimates σ_t^A and σ_t^B, and correlation estimates ρ_t^{AB}. And our portfolio volatility is

$$\sigma_t^2 = w_A^2 \times (\sigma_t^A)^2 + w_B^2 \times (\sigma_t^B)^2 + 2 \times w_A \times w_B \times \rho_t^{AB} \times \sigma_t^A \times \sigma_t^B$$

- It is in fact easier to do this calculation using matrix operations, especially when you have to deal with hundreds of risk factors.

Variance-Covariance Matrix

- We construct a variance-covariance matrix for risk factors A and B:

$$\Sigma_t = \begin{pmatrix} (\sigma_t^A)^2 & \rho_t^{AB}\sigma_t^A\sigma_t^B \\ \rho_t^{AB}\sigma_t^A\sigma_t^B & \left(\sigma_t^B\right)^2 \end{pmatrix}$$

- It is a 2×2 matrix, since we have only two risk factors. If you have 100 risk factors in your portfolio, then you will have a 100×100 matrix. For example, in JPMorgan's RiskMetrics, 480 risk factors were used. In Goldman's annual report, 70,000 risk factors were mentioned.

- A risk manager deals with this type of matrices everyday and the dimension of the matrix can easily be more than 100, given the institution's portfolio holdings and risk exposures.

- Notice the timing: for σ_t, we use all returns up to day $t-1$ for the purpose of forecasting volatility for day t.

Portfolio Volatility

- Let's write our weights in vector form, time stamped by today, t-1,

$$w_{t-1} = \begin{pmatrix} w_{t-1}^A \\ w_{t-1}^B \end{pmatrix}$$

- Our portfolio volatility is

$$\sigma_t^2 = \begin{pmatrix} w_{t-1}^A & w_{t-1}^B \end{pmatrix} \times \begin{pmatrix} (\sigma_t^A)^2 & \rho_t^{AB}\sigma_t^A\sigma_t^B \\ \rho_t^{AB}\sigma_t^A\sigma_t^B & \left(\sigma_t^B\right)^2 \end{pmatrix} \times \begin{pmatrix} w_{t-1}^A \\ w_{t-1}^B \end{pmatrix}$$

- Using the notation we've developed so far, we can also write

$$\sigma_t^2 = w_{t-1}' \times \Sigma_t \times w_{t-1} \,,$$

which involves using mmult and transpose in Excel.

Portfolio VaR

- Let σ be the daily volatility estimate of the portfolio. Then the 95% one-day VaR is,

$$\text{VaR} = \text{portfolio value} \times 1.645 \times \text{sigma}$$

- The 99% tail event corresponds to a -2.326 σ move away from the mean. The 95% tail event corresponds to -1.645 σ.

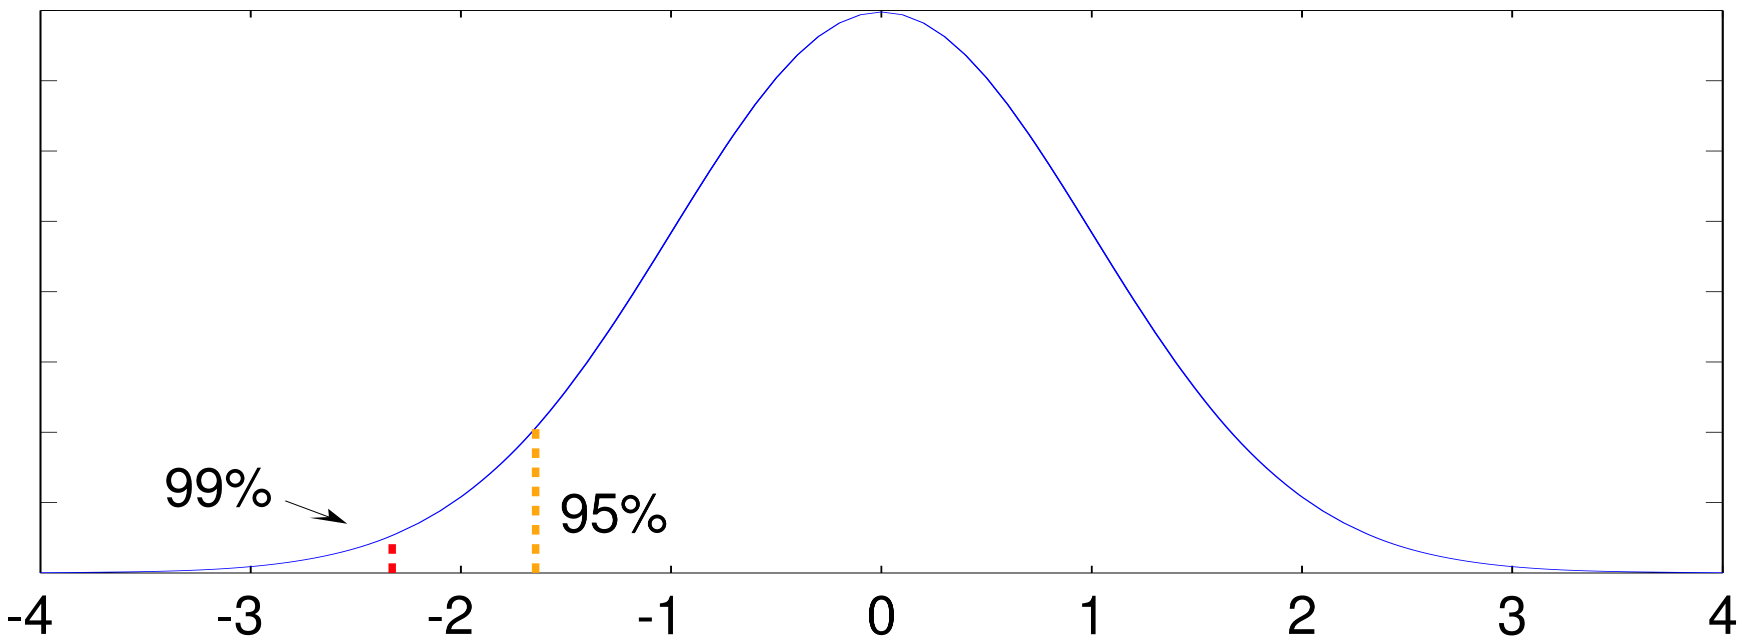

- Assuming the market value of our risk portfolio is \$100 million, the one-day loss in portfolio value associated with the 5% worst-case scenario is

$$\$100\text{M} \times 1.645 \times \sigma$$

- Suppose that we have only one risk factor, which is the S&P 500 index. If today is a normal day with an average volatility around 1%, then the one-day 95% VaR is \$1.645M. For the same portfolio value, if the reported VaR is much higher than \$1.645M, then today must be a volatile day.

- Overall, if we fix our VaR estimate to a certain horizon, say daily, then the main drivers to the VaR estimates are: the market value and volatility of our portfolio. A reduction in VaR could be caused by a reduction in the market value (either by active risk reduction or passive loss in market value) or a reduction in market volatility.

Key Asset Classes for Market Risk Management

- What JP Morgan RiskMetrics had to offer (free of charge) back in 1996 gives a good overall picture of what kind of asset classes are involved in calculating the market risk exposure of an investment bank.
- RiskMetrics data sets: Two sets of daily estimates of future volatilities and correlations of approximately 480 rates and prices, with each data set totaling 115,000+ data points. One set is for computing short-term trading risks, the other for medium term investment risks. The data sets cover foreign exchange, government bond, swap, and equity markets in up to 31 currencies. Eleven commodities are also included.
- This set of data (equity, currency, interest rates, and commodity) is very much the domain of Market Risk Management. In addition, **Credit** and **Liquidity Risk Management** have become increasingly important. For this, good data, models, and talents on credit and liquidity are in need.

Broad Asset Classes for Market Risk Management

from Goldman Sachs 2010 10-K form

Average Daily VaR

in millions Risk Categories	Year Ended December 2010	December 2009	November 2008
Interest rates	$ 93	$176	$ 142
Equity prices	68	66	72
Currency rates	32	36	30
Commodity prices	33	36	44
Diversification effect [1]	(92)	(96)	(108)
Total	**$134**	**$218**	**$ 180**

1. Equals the difference between total VaR and the sum of the VaRs for the four risk categories. This effect arises because the four market risk categories are not perfectly correlated.