Lecture notes on money, finance, risk management, and public policy

Risk, expectations and asset price behavior

Allan M. Malz

Columbia University

The difficulty of quantifying risk

- Quantitative analysis of risk involves understanding asset prices in two dimensions
  - Configuration of asset prices at a point in time: **equilibrium asset prices**
  - Behavior of asset prices over time: **asset price dynamics**
- Financial modeling typically assumes outcomes or prices have well-defined probability distributions
  - Typically in finance, assumption of normal distribution
- Artificial and possibly misleading procedure, assumes
  - Outcomes can be summarized quantitatively
  - Ability to assign probabilities to outcomes
  - Low dimensionality, lack of richness: limited number of things about which information would be useful
- Two types of problem arise:
  - Ignores role of fine-grained, qualitative information
  - Ability to unambiguously rank outcomes by desirability

Importance of qualitative information

- **Knightian uncertainty:** after Frank H. Knight, American economist, *Risk, uncertainty and profit* (1921)
- Distinction between

  Risk: pertains to situations or problems that are quantifiable, and which one can reasonably describe with a probability distribution.

  Uncertainty refers to situations in which no quantitative estimates can be made

  - The outcomes are unknown or can only be guessed at, "because the situation dealt with is in a high degree unique."
  - Entrepreneurial profit: compensation for bearing uncertainty, not risk
  - Knightian uncertainty, not risk, predominant in real-world situations

Arbitrage

- **Arbitrage:** profiting by buying and selling equivalent goods in two different markets
- What does "equivalent" mean
  - Identical goods
  - Goods with cash flows that are identical in size, risk, and timing
  - Weakest form of equivalence: same present value
- **Law of One Price** (**no-arbitrage principle**, Jevons' **Law of Indifference**) states that arbitrage opportunities are exceptional

Risk-neutral pricing

- **Risk-neutral pricing:** approach to finding **fair value** when market prices of securities not available
- Based on market inputs together with pricing models
  - For derivatives, use underlying asset price
  - For illiquid securities, use spreads of comparable securities
- Application of no-arbitrage principle: market price, if equal to risk-neutral price, would be consistent with such market inputs as we observe
- **Risk premiums** already embedded in market inputs to pricing models
  - Market price is a **certainty equivalent:** lower than value if uncertainty were absent
  - Thus increasing return and generating risk premium

Asset prices as forecasts

- To extent markets are ($\rightarrow$)efficient, asset prices clear at levels consistent with risk appetites and expectations for future
- Asset prices embed **risk-neutral expectations**: expectations that could be imputed to agents indifferent to risk
  - Provides a perspective for interpretation asset prices without a model of agent preferences
- Derivatives prices particularly suited to expressing expectations
- Derivatives delivery dates in the future$\rightarrow$prices or payoffs at those dates

Forwards and futures: mean forecasts, expectations of futures prices

Options: expectations of futures payoffs contingent on future prices$\rightarrow$estimate of entire distribution

Agreement and disagreement

- Models may assume disagreement or agreement among market participants about expected means and volatilities
- One common form of agreement is the **rational expectations hypothesis:**
  - Market participants each have a correct model of the economy and financial system
  - Implies agreement among market participants on model
  - Subjectively-expected moments then equal to statistical moments
    - Expected value of asset i's return $\mathbf{E}[r_i]$ equal to statistical mean μ_i
  - Distributions may then be time-varying, with expectations updated using most recent data

Profit and loss

- S_t the time-t, S_{t+1} next-period (1 day, week, month...) asset price
- A **position** in an asset is a holding of x units of the asset is
  - xS_t is the current (market or book) value of the position
- The **profit and loss** (P&L) on a position of x units of an asset is the change in value $x(S_{t+1} - S_t)$
- P&L may be

  Realized: based on actual purchase or sale that unwinds or adds to position

  Mark-to-market (MTM) or **fair value:** based on observed market price or assessment of what market price would be if there were an active market
- $x < 0$ for a **short position**: loss if price rises
- Asset price floored at $0 \Rightarrow$
  - Short P&L has no lower limit: risk of a short position unlimited
  - Long P&L has no upper limit

Mechanics of short positions

- Establishing a short position:
  - Borrow units of the asset itself, sell in spot market, buy and return to lender at future date
- Cash flows of short positions
  - Lender of asset retains right to interest and dividends
  - Interest and dividends must consequently be paid by borrower to lender of asset
  - Short seller—borrower of asset—receives interest on cash proceeds from sale
- Very different legal structure and mechanics for stocks, bonds and across countries
- Can be expressed through derivatives trades: sell a forward, futures or option on the asset, or engage in a swap
- Short position may be used to hedge a long position, reducing rather than increasing exposure and risk

The convention of risk-free assets

- Analysis may stipulate existence and availability of a **risk-free asset**
  - As part of the set of available investment choices
  - As a tool for valuation, e.g. discount factor used in measuring fair value of an asset
- Implied risk-free rate may be calculated within models, based on macroeconomic fundamentals, household time and risk preferences (→**stochastic discount factor**)
- Real-world investment choices include low-risk, but not truly risk-free assets
- Assets conventionally classified as riskless include

  Cash or money balances: legal tender, or redeemable on very short notice at par value

  Government bonds: generally understood to include only issues of advanced market economies

  - Sovereign credit risk near-zero for bonds issued in local currency

Problems in the use of risk-free asset conventions

- Conventionally-defined riskless assets subject to several risks:
  - Currency appreciation or depreciation
  - Inflation
  - Liquidity risk: risk-on bond selloff, suspension of convertibility
  - Credit risk: sovereign default, bank failure, decline in collateral value
  - Interest-rate risk
- Yields on government bonds may contain (possibly negative) compensation for risk
- Extremely low or negative yields post-crisis
  - $\leftrightarrow$High bond values that fluctuate widely over time

Risks of cash balances

- Cash and cash substitutes include
  - Currency and other forms of **central bank** or **narrow money**
  - Short-term government or commercial paper
  - Bank deposits
  - MMMF accounts
- Even cash balances not riskless, are exposed to:
  - Theft, physical destruction
  - Disruption of payment or settlement systems
  - Currency reform: change in currency system
  - **Demurrage:** imposition of cost on holding currency
- **Examples** of recent risk events affecting central bank money:

  Demonetisation of largest-denomination notes in India
  08Nov2016, intended to thwart tax evasion

  Negative interest rates on some central bank deposits from 2009

Safe assets

- **Safe assets:** generally said to include debt with low credit, inflation and currency risk, high liquidity
  - Also characterized as **information-insensitive:** requires little or no credit analysis or research to be accepted
- Pre-crisis: included not-quite "riskless," senior tranches of mortgage-backed securities, other securitizations
- Post-crisis: primarily advanced-economy government issues, especially dollar-denominated
  - Current "shortage" exacerbated by disappearance or disqualification of some previously "safe" assets post-crisis
- Shortage of safe assets said to have played key role in depressing interest rates before crisis
  - (→)**Global savings glut hypothesis:** supply of safe assets lagging behind demand as world wealth grows

Government curves and swap curves

- Risk-free rates often used in financial contracts as base for setting prices of risky fixed-income securities
- Define **risk-free curve** or set of risk-free rates with different terms to maturity
- Two types of risk-free curves widely used as **benchmark** or **reference rates:**

  Government curves: rates on central government issues
  - Assumed to be issuer of currency in which bonds are denominated
  - **Example:** U.S. Treasury curve

  Swap curves: rates on interbank money market instruments and par swaps
  - **Example:** **London Interbank Offered Rate** (LIBOR) curve, with rates fixed privately daily
  - Contains credit, counterparty and other risk premiums likely absent from government curves
  - Generally higher than corresponding government curve, but anomalous behaviors after 2008 crisis

Choice of reference rates

- Wide agreement on choice of reference rates driven by network effects
- Alphabet soup of money market rate acronymss
- Government curves less frequently used as major-currency reference rates
  - Propensity for central government yields to decline sharply in financial stress (→**flight to quality**)
  - Possibility of dearth of issuance
- **LIBOR manipulation scandal:** submitted rates manipulated to affect valuations or avoid stigma
  - Some reference rates published by central banks, e.g. fed funds, and used as policy tools or indicators
  - Central banks and regulatory authorities now heavily involved in identifying substitute reference rates for LIBOR
  - Based on money market transactions
  - Some based on secured transactions and labeled **"nearly risk-free"**

Widely-used central bank reference rates

- Reference rates published by central banks based on interbank lending rates

  Effective federal funds rate (EFFR): uses U.S.-domiciled bank data

  Overnight bank funding rate (OBFR): includes overnight borrowing in Eurodollar market

  Euro short-term rate (€STR)

  Sterling overnight index average (SONIA)

- Newly-introduced "nearly risk-free" U.S. Treasury repo reference rates

  Secured Overnight Financing Rate (SOFR): index of overnight Treasury repo rates, blending rates on tri-party general collateral, GCF Repo, and DVP repo transactions

  Broad General Collateral Rate (BGCR): similar to SOFR index, but includes only tri-party and GCF

  Tri-Party General Collateral Rate (TGCR): similar to BGCR index, but includes only tri-party

Arithmetic and logarithmic returns

- Risk analysis focus usually on asset *returns*, based on price change in currency units or **capital gain**, rather than asset price *levels*
- Per-period **price return** can be defined as

Arithmetic or **simple** or **linear return:**

$$r_{t,t+1}^{\text{arith}} = \frac{S_{t+1} - S_t}{S_t} = \frac{S_{t+1}}{S_t} - 1$$

  - $1 + r_{t,t+1}^{\text{arith}}$ is called the **gross return**
  - P&L on a position of x units of an asset is $xS_t r_{t,t+1}^{\text{arith}} = x(S_{t+1} - S_t)$

Logarithmic or **geometric** or **continuously compounded return:** change in logarithm of price

$$r_{t,t+1} = \ln\left(\frac{S_{t+1}}{S_t}\right) = \ln(S_{t+1}) - \ln(S_t)$$

- t : date of observation r_t or S_t
  - Measured in any time unit, e.g. seconds, days, months, years,...

Comparing arithmetic and logarithmic returns

- Arithmetic and logarithmic returns are related by

$$r^{\text{arith}}_{t,t+1} = e^{r_{t,t+1}} - 1$$
$$r_{t,t+1} = \ln\left(1 + r^{\text{arith}}_{t,t+1}\right)$$

- Arithmetic returns are greater than logarithmic

$$r^{\text{arith}}_{t,t+1} > r_{t,t+1},$$

apart from trivial case where both are zero

- If $r_{t,t+1}$ and $r^{\text{arith}}_{t,t+1}$ relatively small,

$$r^{\text{arith}}_{t,t+1} \approx r_{t,t+1}$$

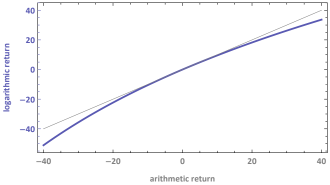

In percent. The vertical distance is the difference between the two return definitions for a given initial level and change in asset prices.

When to use arithmetic and logarithmic returns?

- Returns **compound** over time
  - As returns are added to the previous period's principal
- Log returns most useful when measuring *multiperiod* returns over
  - Computed as if returns accrete uniformly in tiny increments at tiny intervals
  - Based on cumulative principal at that instant
  - Log returns also more convenient when computing returns encompassing several sources of risk, e.g. ($\rightarrow$)total or real returns
- Arithmetic returns most useful when measuring *portfolio returns*
- Risk models generally framed in terms of logarithmic rather than arithmetic returns
  - ($\rightarrow$)Volatility typically defined for log returns

Multiperiod logarithmic returns

- Log returns aggregate nicely across time: τ-period **compound** or **cumulative log return** is sum of shorter-term log returns

$$\ln\left(\frac{S_{t+\tau}}{S_t}\right) = \ln\left(\frac{S_{t+1}}{S_t}\cdots\frac{S_{t+\tau}}{S_{t+\tau-1}}\right) = \sum_{\theta=1}^{\tau} r_{t+\theta-1,t+\theta}$$

  - τ is the number of periods and θ the index of summation
- With t measured in months, annual related to monthly returns by

$$\ln\left(\frac{S_{t+12}}{S_t}\right) = \sum_{\theta=1}^{12} r_{t+\theta-1,t+\theta}$$

- With t measured in years, logarithmic **return at an annual rate** or annualized τ-year return is simple average of yearly rates:

$$\frac{1}{\tau} r_{t,t+\tau} = \frac{1}{\tau}\sum_{\theta=1}^{\tau} r_{t+\theta-1,t+\theta}$$

Multiple-period arithmetic returns

- Arithmetic returns don't aggregate so nicely across time
  - Inconvenient in modeling or working with return time series
- τ-period **compound arithmetic return**:

$$\frac{S_{t+\tau}}{S_t} - 1 = \frac{S_{t+1}}{S_t} \cdots \frac{S_{t+\tau-1}}{S_{t+\tau-2}} \cdot \frac{S_{t+\tau}}{S_{t+\tau-1}} - 1$$
$$= \prod_{\theta=1}^{\tau} \frac{S_{t+\theta}}{S_{t+\theta-1}} - 1 = \prod_{\theta=1}^{\tau} (1 + r^{\text{arith}}_{t+\theta-1,t+\theta}) - 1$$

- $\rightarrow$ τ-year arithmetic return at an annual rate is a *geometric* average:

$$(1 + r^{\text{arith}}_{t,t+\tau})^{\frac{1}{\tau}} - 1 = \left(\frac{S_{t+\tau}}{S_t}\right)^{\frac{1}{\tau}} - 1$$

- Realized cumulative return from time 0 to time τ can be expressed as an index, using historical short-term returns and setting $S_0 = 100$:

$$S_t = S_0(1 + r^{\text{arith}}_{0,\tau}) = 100 \prod_{\theta=1}^{\tau} (1 + r^{\text{arith}}_{\theta-1,\theta})$$

Portfolios and constituent asset weights

- **Portfolios** are combinations of asset positions
- With x_i the number of units of asset i and $S_{i,t}$ its time-t price,
    - The value of each position is $x_i S_{i,t}$
    - The value of the portfolio is $\sum_i x_i S_{i,t}$
    - **Example:** with two positions, the portfolio value is $x_1 S_{1,t} + x_2 S_{2,t}$
- For a portfolio consisting of long positions only, $x_i > 0, \forall i$
    - Asset i's **weight**:

$$w_i = \frac{x_i S_{i,t}}{\sum_i x_i S_{i,t}}$$

    - $w_i > 0, \forall i$ and $\sum_i w_i = 1$
- Use of constituent weights more complicated for long-short portfolio
    - Major issue in measuring portfolio $(\rightarrow)$leverage
- For a long-short portfolio with zero net market value, $x_i \gtrless 0$ and $\sum_i x_i S_{i,t} = 0$
    - Weights defined in alternative ways to avoid division by 0, e.g.

$$w_i = \frac{|x_i| S_{i,t}}{\sum_i |x_i| S_{i,t}}$$

Arithmetic and logarithmic portfolio returns

- Period-end value is more simply related to arithmetic than to logarithmic return:

$$S_{t+1} = \left(1 + r^{\text{arith}}_{t,t+1}\right) S_t = e^{r_{t,t+1}} S_t$$

- $\rightarrow$Arithmetic returns aggregate nicely across assets:

$$\begin{aligned} r^{\text{arith}}_{p,t,t+1} &\equiv \frac{\sum_i x_i S_{i,t+1} - \sum_i x_i S_{i,t}}{\sum_i x_i S_{i,t}} \\ &= \frac{\sum_i x_i \left(1 + r^{\text{arith}}_{i,t,t+1}\right) S_{i,t} - \sum_i x_i S_{i,t}}{\sum_i x_i S_{i,t}} = \sum_i w_i r^{\text{arith}}_{i,t,t+1} \end{aligned}$$

- Portfolio P&L is equal to the arithmetic portfolio return times the initial portfolio value: $r^{\text{arith}}_{p,t,t+1} \sum_i x_i S_{i,t}$
- Logarithmic returns don't aggregate so nicely across assets:

$$r_{p,t,t+1} = \ln\left(\frac{\sum_i x_i S_{i,t+1}}{\sum_i x_i S_{i,t}}\right) = \ln\left(1 + \sum_i w_i r_{i,t,t+1}\right)$$

Total return

- **Total return** includes cash flows thrown off by assets, such as interest, dividends or rent, in addition to capital gain
- With time-$t+1$ cash flows d_{t+1}
  - Arithmetic total return is $\frac{S_{t+1}+d_{t+1}}{S_t} - 1$
  - Logarithmic total return is $\ln\left(\frac{S_{t+1}+d_{t+1}}{S_t}\right)$
- If d_{t+1} is a coupon payment and the period is one year, then $\frac{d_{t+1}}{S_t}$ is the bond's **current yield**
- For assets such as real estate and commodities
  - Estimated physical storage and maintenance costs can be represented as $d_{t+1} < 0$
  - Possibly offset by the **convenience yield** of readily available inventory or of housing services
- Over longer periods of time, with compounding, the differences between total and price return can be very large

Price and total return of S&P 500 1970–2024

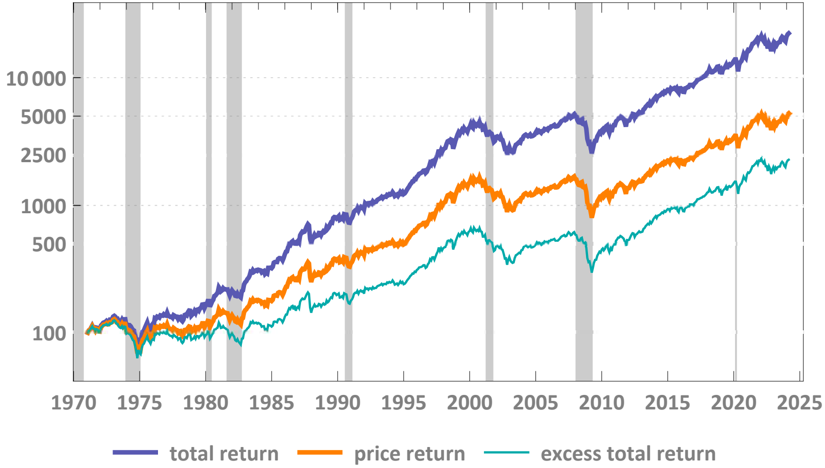

Index of cumulative price and total, monthly, logarithmic scale, Nov. 1970 (=100). Trailing 12-month dividend yield; excess returns over 3-month U.S. T-bill yield. Vertical shading represents NBER recession dates. *Data source*: Bloomberg LP.

Excess return

- **Excess return:** difference between the return on an asset and the so-called (→)**risk-free** or **riskless** return r^f
  - Risk-free rate typically proxied by yield or return on short-term central government debt denominated in same currency
  - For long-term investment, longer-term proxy, e.g. government bond, for risk-free rate may be appropriate for some purposes
- Logarithmic excess total return is $\ln\left(\frac{S_{t+1}+d_{t+1}-r_t^f}{S_t}\right)$
- Risk-free return should be known at time t
  - Desirable to match term to return measurement period
  - Longer-term proxies subject to price fluctuations within return measurement period→"riskless return" uncertain
- Over longer periods of time, with compounding, the differences between excess return and other return measures can be very large

Cumulative total and excess return of S&P 500

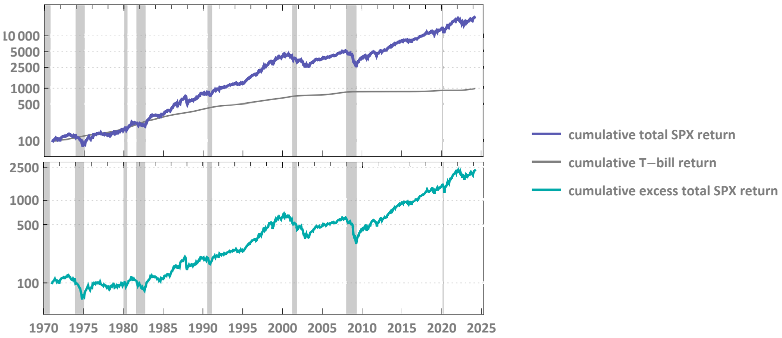

Index of cumulative total, price and excess returns, monthly, logarithmic scale, Nov. 1970 (=100). Trailing 12-month dividend yield; excess returns over 3-month U.S. T-bill yield. Vertical shading represents NBER recession dates. *Data source*: Bloomberg LP.

Nominal and real returns

- Asset prices generally expressed in **nominal** terms, i.e. money units
- Asset prices and returns can also be expressed in **real** terms, i.e. units of goods or purchasing power
- Real price levels computed by applying a **price index** or **deflator**
- If P_t represents price index level at time t, real asset price is $\frac{S_t}{P_t}$
- If P_t rising over time, inflation rate > 0 and real return$<$nominal
- In U.S., two widely-used deflators of consumer goods—not asset—prices, differ in what goods included and how computed:
  Consumer Price Index for All Urban Consumers (CPI-U), used for indexing Social Security
  Price Index for Personal Consumption Expenditures (PCE), defines Federal Reserve's 2 percent inflation goal
- Price indexes have "core" variants excluding volatile food and energy

Measures of real return

- Logarithmic **real return**:

$$\ln\left[\frac{S_{t+1}}{P_{t+1}}\left(\frac{S_t}{P_t}\right)^{-1}\right] = \ln\left[\frac{S_{t+1}}{S_t}\left(\frac{P_{t+1}}{P_t}\right)^{-1}\right] = r_{t,t+1} - \pi_{t,t+1}$$

  and $\pi_{t,t+1}$ is the logarithmic **rate of inflation** from t to $t+1$
    - Definition can include asset cash flows d_{t+1} →**total real return**
- Arithmetic real return is a slightly more complicated expression
    - But the approximation is generally just fine

$$\frac{\frac{S_{t+1}}{P_{t+1}}}{\frac{S_t}{P_t}} - 1 = \frac{S_{t+1}}{P_{t+1}}\frac{P_t}{S_t} - 1 = \frac{S_{t+1}}{S_t}\frac{P_t}{P_{t+1}} - 1 = \frac{1 + r^{\text{arith}}_{t,t+1}}{1 + \pi^{\text{arith}}_{t,t+1}} - 1$$

$$\approx r^{\text{arith}}_{t,t+1} - \pi^{\text{arith}}_{t,t+1}$$

Nominal and real return of T-bills 1970–2024

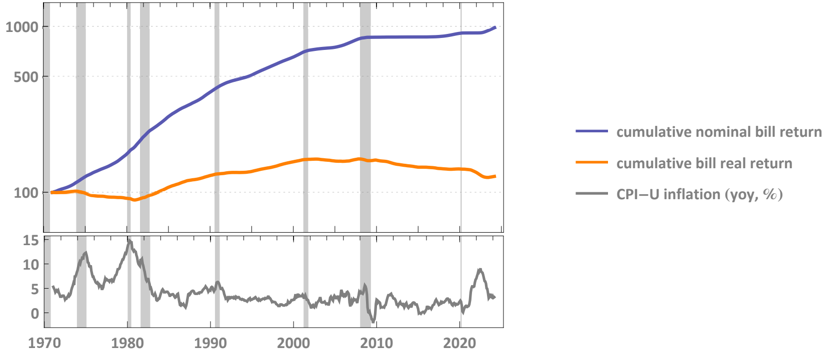

Index of cumulative total nominal and real returns, monthly, logarithmic scale, Nov. 1970 (=100). Real returns computed by subtracting year-over-year change in CPI-U, calculated as geometric average of monthly rates. This slightly smooths the month-to-month effect of inflation. Vertical shading represents NBER recession dates. *Data source*: Bloomberg LP.

Nominal and real return of S&P 500 1970–2024

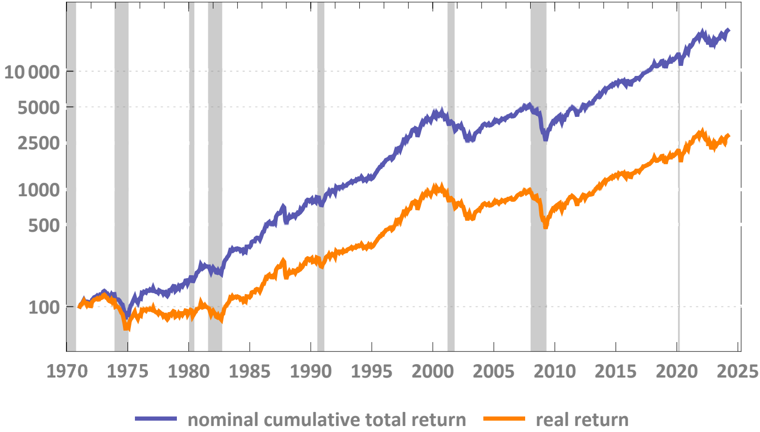

Index of cumulative total nominal and real returns, monthly, logarithmic scale, Nov. 1970 (=100). Real returns computed by subtracting year-over-year change in CPI-U, calculated as geometric average of monthly rates. This slightly smooths the month-to-month effect of inflation. Vertical shading represents NBER recession dates. *Data source*: Bloomberg LP.

Return experience of the S&P 500 1970–2020

	Annual	Cumulative
Price return	7.84	5 530
Total return	10.81	23 458
Bill return	4.40	988
Excess total return	6.15	2 384
Inflation	3.96	790
Real total return	6.59	2 973
Real bill return	0.42	125

Annual and cumulative arithmetic returns, percent, Nov. 1970 through Mar. 2024. *Data source*: Bloomberg LP.

Hypotheses about return distributions

- Estimates of the probability distribution of *future* returns (or prices) a fundamental component of risk measurement
- Normal distribution usually a good first approximation to return behavior
  - **Central limit theorem:** averages behave like normals, even if the underlying data non-normal
  - But normal-based may be poor estimate of **tail events**
  - And assumes symmetry between positive and negative returns (→tail risk, higher moments)
- Distributional hypotheses are often needed, but highly restrictive
  - E.g. underestimate of risk due to normality assumption
- Payoffs and risk of asset captured by characteristics (usually **parameters** or **moments**) of return distribution
- Models may assume distribution and its moments invariant over time, or **time-varying** moments

What we forecast when we forecast asset returns

Expected return or **mean**

- Average or likeliest or "central tendency" of future return/asset value
- Fundamentals-oriented, trade-ideas seeking analysts: focus on expected return

Volatility or **standard deviation**

- How far might future return stray from recent or expected?

Extreme events or **tail risks**

- Likelihood of extremely large returns, clusters of defaults

Correlation of returns

- Do values of different assets tend to move in the same or in opposite direction?

Expected values and expectations

- "Standard" models, e.g. normal, typically focus on
  Expected return (μ_i): expected value of the rise in asset i's value (plus cash flows it yields over time)
  Return variance or its square root (σ_i)—the return **standard deviation** or **volatility**—measures how much asset i's value "bounces around" over time
- How do we understand volatility?
  Realized volatility over recent past
  - Can vary widely depending on frequency and observation interval

  Expected based on realized volatility or model
  Implied by option prices
- Higher moments shed light on $\rightarrow$tail risk, extreme events

Where do distributional hypotheses come from?

Statistical estimates of the return distribution, sometimes called the **actuarial**, **physical**, or simply **real-world** distribution

Subjectively-expected distribution: that expected by market participants

- Unobservable, can only be inferred using models and surveys

Risk-neutral return distribution is that implied by market prices

- Can be estimated if prices of derivatives—futures, forwards, and options—in addition to cash assets are observed

Comparison of more and less volatile stocks

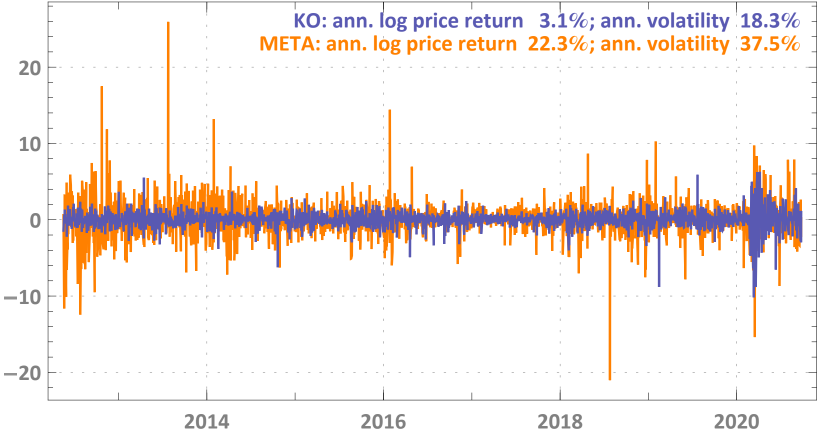

Daily logarithmic stock price returns of Coca-Cola Co. (ticker KO) and Meta Inc. (META), 18May2012 to 24Sep2020. Volatility: daily return standard deviation over the entire period, measured at an annual rate.

Multivariate return distributions

- Portfolio payoffs depend on combined or joint behavior of returns of constituents
- In simple normal model, joint behavior of returns on several assets described by **multivariate normal distribution**
- Parameters include **return correlations** in addition to means and volatilities

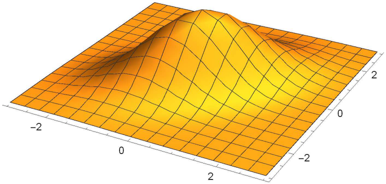

Density of a bivariate standard normal distribution with correlation coefficient $\rho = 0.5$.

Visualizing joint distributions via density contours

- **Constant probability density contours:** set of points with constant density
- Contours are computed by equating standardized realizations to a constant:

$$c = (\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}),$$

    - $\mathbf{x}$ represents realizations of the jointly normal pair
        - $\mathbf{x}$ are the pairs on the "floor" of the bivariate distribution graph
    - $\boldsymbol{\mu}$ the mean and Σ the covariance matrix of the distribution
- To generate contours such that realizations of the pair of random variables fall within each contour with probability α
    - Set the constant $c = \chi_2^2(1 - \alpha)$
    - $\chi_2^2(1 - \alpha)$ is the $1 - \alpha$ quantile of the **chi-square distribution** with 2 degrees of freedom,

Correlated and uncorrelated returns

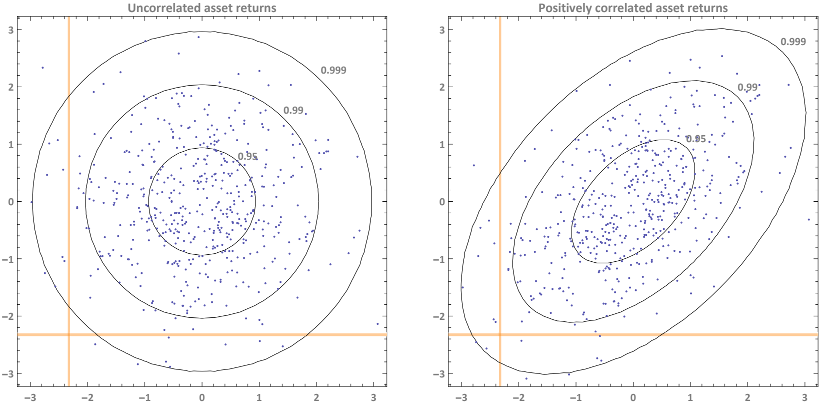

Constant probability density contours of a bivariate standard normal distribution and random sample of 500 realizations. Realizations of the joint normal pair **x** fall within each contour with the indicated probability. Orange grid lines indicate 0.01-quantile of each marginal distribution. Left panel: correlation coefficient $\rho = 0$. Right panel: correlation coefficient $\rho = 0.50$.

Impact of correlation on joint return distributions

- For a bivariate normal distribution, a positive (negative) return correlation increases the probability of pairs of large-magnitude returns with the same (opposite) sign
- For a bivariate normal constant density contours are ellipsoids
  - Multivariate normal belongs to family of **elliptical distributions**
  - Multivariate normal with pairwise zero correlation and identical variances belong to family of **spherical distributions**

Risk and uncertainty

Asset return distributions

Standard model of asset price dynamics

- Forecasting asset return: random walks
- The geometric Brownian motion model of asset price dynamics

Basic characteristics of the standard model

- Successive changes in asset prices are **independent** and approximately follow a **random walk**
  - But just an approximation: possible first-moment dependency, e.g. positive drift in prices
  - And well-established serial dependence in higher moments
- Quantitative models often assume log returns normally distributed
  - Many asset prices and factor returns empirically distributed close to normally
- Tractable, easy to estimate, and few parameters
  - And what do you do if no alternative model is unambiguously better?
- Define behavior of price or return via a few properties, then derive many more
- Consistent with (→)**efficient markets hypothesis**
  - Zero or low return **autocorrelation:** correlation of succesive returns

Defining properties of a random walk

- Function of time, starts at $(0, 0)$: time 0, position 0
- Discrete: adds increments at time steps of length Δt
- Two possible sizes of increments: $\pm\sqrt{\Delta t}$
- With probabilities:
    - Positive increments π
    - Negative increments $1 - \pi$
- It can't stay in the same position two consecutive time steps
    - But it can return to a position after at least one more step
- With passing of time, random walk may arrive at several different values
    - But in general several **sample paths** lead to each value
    - Each sample path has same probability

One sample path of a random walk

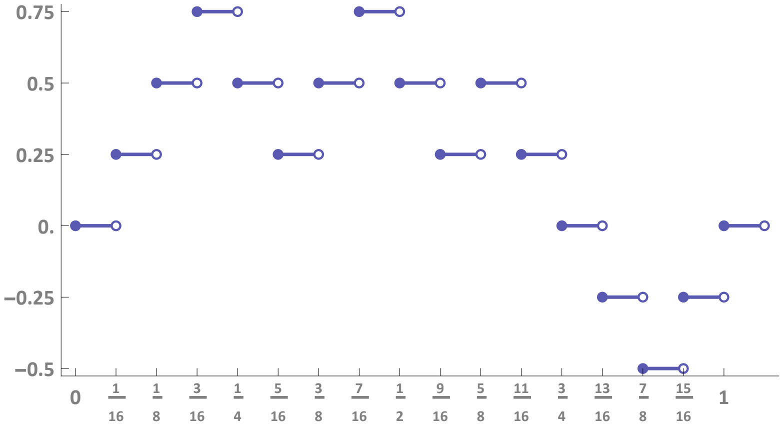

Plot shows the first 16 positions of one simulation of a random walk over 1 time unit. The time interval between steps is $\Delta t = \frac{1}{16}$ and the magnitude of the increments is $\sqrt{\Delta t} = \frac{1}{\sqrt{16}} = \frac{1}{4}$. The dots show the value the position takes on at the start of each time interval.

Convergence of random walk to Brownian motion

- Let length of time step $\Delta t \to 0$
- Random walk$\to$**Brownian motion** or **Wiener process** or **diffusion process**
- $\to$There's only one, unique Brownian motion W (though infinitely many possible paths)
  - Regardless of probability π of positive increment to random walk
  - Though infinitely many possible paths
- Discrete time approximation of asset price behavior:

$$\Delta S_t = z\sqrt{\Delta t},$$

with S_t an asset price and z a standard normal variate

Defining properties of Brownian motion

- Position or level of asset price or risk factor S_t a function of time t, starts at $S_0 = 0$
- Every sample path is continuous (but awfully jagged)
- Position after t time units (but uncountably many steps) a standard normal variate

$$S_t \sim \mathcal{N}(0, t)$$

- Implying **martingale property**, useful in modeling asset prices:
    - Successive and non-overlapping increments to S_t are *independent* of one another and of initial position

From Brownian to geometric Brownian motion

- Brownian motion has unfortunate property for modeling asset prices: S_t can be negative
- So apply it to tiny increments of *logarithm* rather than *level* of asset price S_t, *and*
  - Add **drift** term: deterministic increase at rate μ per unit of time along every sample path of S_t
  - Scale variance of $\ln(S_t)$ by a **volatility** parameter σ
- Logarithm of S_t then follows a **geometric Brownian motion**
- Discrete time approximation:

$$\Delta \ln(S_t) \approx \frac{\Delta S_t}{S_t} = \mu\Delta t + \sigma z\sqrt{\Delta t}$$

Itô's Lemma and the distribution of log returns

- **Itô's Lemma**⇒*level* of asset price S_t follows **lognormal distribution**
- ⇒Log returns normally distributed:

$$r_{t,t+\tau} \equiv \ln(S_{t+\tau}) - \ln(S_t) \sim \mathcal{N}\left[\left(\mu - \frac{1}{2}\sigma^2\right)\tau, \sigma^2\tau\right]$$

$$\Rightarrow \frac{r_{t,t+\tau} - \left(\mu - \frac{1}{2}\sigma^2\right)\tau}{\sigma\sqrt{\tau}} \sim \mathcal{N}(0,1)$$

Volatility drain

- Itô's Lemma—normality of log returns—implies (statistical) expected value of logarithmic—continuously compounded—return of

$$\mathbf{E}\left[r_{t,t+\tau}\right] = \mathbf{E}\left[\ln(S_{t+\tau})\right] - \ln(S_t) = \left(\mu - \frac{1}{2}\sigma^2\right)\tau$$

- But expected value of *level* of S_t higher by **Jensen's Inequality term** $\frac{1}{2}\sigma^2\tau$:

$$\mathbf{E}\left[S_{t+\tau}\right] = S_t e^{\mu\tau}$$

- Result: **volatility drain** or **drag**: mean arithmetic exceeds mean logarithmic rate of return
- Impact higher
  - Over longer periods
  - If volatility higher
- Practical implication: reported average annual return exceeds cumulative long-term return to investors

Sample paths of geometric Brownian motion

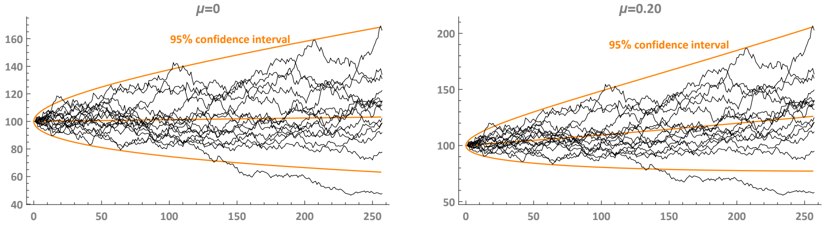

Fifteen simulations of the path of the price level over time of an asset following a geometric Brownian motion process with annualized volatility 25 percent, initial asset price 100. At each point in time, the levels of the paths are lognormally distributed. Orange rays and hyperbolas plot the mean paths and 95 percent confidence intervals over time.

Properties of geometric Brownian motion

- If zero drift ($\mu = 0$), asset price almost (Jensen's Inequality) equally likely to move up or down
  - Like random walk, but always near 50-50
- The further into the future we look, the likelier it is that the asset price will be far from its current level
- The higher the volatility, the likelier it is that the asset price will be far from its current level within a short time
- The time series of daily logarithmic asset returns are independent draws from a normal distribution
- Variation properties of geometric Brownian motion:

  Total variation The total distance traveled over even a tiny time interval is *infinite* in every sample path

  Quadratic variation But the variance of the distance traveled in different sample paths is *finite*
- Asset price S_t is never negative