Workshop on Robust Methods in Probability & Finance

ICERM, Brown University, June 19 -23, 2017

What is accomplished by successful non-stationary stochastic prediction?

Glenn Shafer, Rutgers University,
www.glennshafer.com

Answer: It tells us nothing about the future.

But it permits market efficiency.

Probability and Finance: It's Only a Game!
Glenn Shafer and Vladimir Vovk
Wiley, 2001

Working papers at
www.probabilityandfinance.com

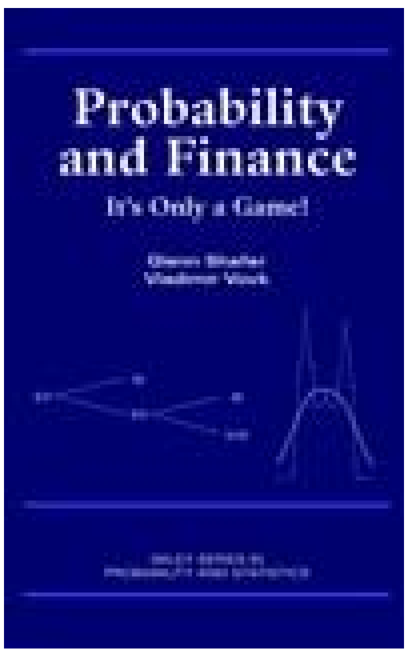

Game-theoretic understanding of probability, testing, and prediction

- Reality is a player in the game.
- **When forecaster has feedback**, good probabilistic prediction is possible, regardless of what Reality does.
- So successful non-stationary prediction with feedback says **nothing** about the future.
- The game is **not** a generative model. We are **not** modelling Reality.
  - Don't say *true probability law*.
  - Don't say *robust*.
- As Rama said this morning, "get rid of probability altogether".

Probability Game

- Forecaster sets prices.
- Skeptic selects bet.
- Reality decides outcome.

…repeat

Perfect information game (prediction with feedback = online prediction)
Players move in order; each sees the others' moves; many rounds.

Probability = betting rate
$P(A)=p$ means Skeptic must risk p to get 1 if A happens.

Statistical test = strategy for Skeptic
Skeptic tests Forecaster by trying to multiply money risked by large factor.

Probability Game

- Forecaster sets prices.
- Skeptic selects bet.
- Reality decides outcome.

…repeat

In financial applications, the market is both Forecaster and Reality.

Game-theoretic definition of market efficiency: Skeptic will not multiply capital risked by large factor.

Surprising result:

Forecaster can pass Skeptic's tests regardless of how Reality moves.

Consequences:

1. Adaptive prediction tells us about the past, not the future.
2. Speculation can make markets efficient.

Some details…

1. Game theory as mathematical foundation for probability

2. Game-theoretic upper probabilities

3. Game-theoretic significance testing

4. Predictions that pass statistical tests (defensive forecasting)

5. Implications for nonstationary prediction (e.g., macroeconomics)

6. Implications for market efficiency

References

1. Game theory as mathematical foundation for probability
   Probability and Finance: It's Only a Game!, Shafer and Vovk, Wiley, 2001.

2. Game-theoretic upper probabilities
   Chapter 6 by Vovk and Shafer in *Introduction to Imprecise Probabilities*, edited by Thomas Augustin et al., Wiley 2014.

3. Game-theoretic significance testing
   Working Paper #49, www.probabilityandfinance.com.

4. Predictions that pass statistical tests (defensive forecasting)
   Working paper #8 at www.probabilityandfinance.com.

5. Implications for nonstationary prediction

6. Implications for market efficiency
   Working Paper #47, www.probabilityandfinance.com.

1. Game theory as mathematical foundation for probability

Example: Betting at even odds

Skeptic announces $\mathcal{K}_0 \in \mathbb{R}$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces $L_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + L_n(y_n - p_n)$.

- A strategy for Skeptic *risks bankruptcy* if it allows Forecaster and Reality to make $\mathcal{K}_n < 0$ for some n.
- A strategy for Skeptic is *safe* if it does not risk bankruptcy.

Strong law of large numbers Skeptic has a safe strategy that guarantees that

$$\frac{\sum_{i=1}^{n}(y_i - p_i)}{n} \to 0 \text{ or else } \mathcal{K}_n \to \infty.$$

Weak law of large numbers Given $K > 1$, $\epsilon > 0$, there exists N such that for $n \geq N$, Skeptic has a safe strategy that guarantees

$$\left|\frac{\sum_{i=1}^{n}(y_i - p_i)}{n}\right| > \epsilon \implies \mathcal{K}_n > K.$$

This and other standard probability theorems proven in 2001 book.

2. Upper and lower probabilities

Skeptic announces $\mathcal{K}_0 \in \mathbb{R}$.
FOR $n = 1, 2, \ldots$:
    Forecaster announces $p_n \in [0, 1]$.
    Skeptic announces $L_n \in \mathbb{R}$.
    Reality announces $y_n \in \{0, 1\}$.
    $\mathcal{K}_n := \mathcal{K}_{n-1} + L_n(y_n - p_n)$.

Definition

$$\overline{\mathbb{P}}(E) := \inf\{\mathcal{K}_0 | \text{ Skeptic can guarantee } \lim_{n\to\infty} \mathcal{K}_n \geq \mathbf{1}_E\}.$$

Strong law of large numbers

$$\overline{\mathbb{P}}\left(\frac{\sum_{i=1}^{n}(y_i - p_i)}{n} \not\to 0\right) = 0.$$

Weak law of large numbers Given $\epsilon > 0$ and $n \in \mathbb{N}$,

$$\overline{\mathbb{P}}\left(\frac{\sum_{i=1}^{n}(y_i - p_i)}{n} \geq \epsilon\right) \leq \frac{1}{4\epsilon^2 n}.$$

2. Game-theoretic upper probabilities (and expected values)

A functional $\overline{\mathbf{E}} : \overline{\mathbb{R}}^{\mathcal{X}} \to \overline{\mathbb{R}}$ is an *upper expectation* if:

Axiom E1. If $f_1, f_2 \in \overline{\mathbb{R}}^{\mathcal{X}}$, then $\overline{\mathbf{E}}(f_1 + f_2) \leq \overline{\mathbf{E}}(f_1) + \overline{\mathbf{E}}(f_2)$.

Axiom E2. If $f \in \overline{\mathbb{R}}^{\mathcal{X}}$ and $c \in (0, \infty)$, then $\overline{\mathbf{E}}(cf) = c\overline{\mathbf{E}}(f)$.

Axiom E3. If $f_1, f_2 \in \overline{\mathbb{R}}^{\mathcal{X}}$ and $f_1 \leq f_2$, then $\overline{\mathbf{E}}(f_1) \leq \overline{\mathbf{E}}(f_2)$.

Axiom E4. For each $c \in \mathbb{R}$, $\overline{\mathbf{E}}(c) = c$.

Skeptic announces $\mathcal{K}_0 \in \overline{\mathbb{R}}$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces an upper expectation $\overline{\mathbf{E}}_n$ on $\mathcal{X}$.
  Skeptic announces $f_n \in \overline{\mathbb{R}}^{\mathcal{X}}$ such that $\overline{\mathbf{E}}_n(f_n) \leq \mathcal{K}_{n-1}$.
  Reality announces $x_n \in \mathcal{X}$.
  $\mathcal{K}_n := f_n(x_n)$.

Global upper probability is a special case of global upper expected value:

$$\overline{\mathbb{E}}(X) := \inf\{\mathcal{K}_0 \mid \text{Skeptic can guarantee } \lim_{n\to\infty} \mathcal{K}_n \geq X\}$$

$$\overline{\mathbb{P}}(E) := \inf\{\mathcal{K}_0 \mid \text{Skeptic can guarantee } \lim_{n\to\infty} \mathcal{K}_n \geq \mathbf{1}_E\}$$

Thus defined global upper expectation also satisfies Axioms E1-E4.
Law of large numbers and other theorems hold in this general context.

3. Game-theoretic significance testing

Skeptic announces $\mathcal{K}_0 \in \mathbb{R}$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces $L_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + L_n(y_n - p_n)$.

Skeptic announces $\mathcal{K}_0 \in \overline{\mathbb{R}}$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces an upper expectation $\overline{\mathbf{E}}_n$ on $\mathcal{X}$.
  Skeptic announces $f_n \in \overline{\mathbb{R}}^{\mathcal{X}}$ such that $\overline{\mathbf{E}}_n(f_n) \leq \mathcal{K}_{n-1}$.
  Reality announces $x_n \in \mathcal{X}$.
  $\mathcal{K}_n := f_n(x_n)$.

Skeptic tests Forecaster by trying to multiply his capital without risking bankruptcy.

- The factor by which he multiplies his capital is the *game-theoretic test score.*
- A test score of 1,000, for example, is interpreted like a conventional significance level of 0.001.
- Assume $\mathcal{K}_0 = 1$. So the test score at time n is $\mathcal{K}_n$.
- If Forecaster chooses a probability distribution $\mathbb{P}$ for Reality's moves $y_1, y_2, \ldots$ and uses it as his strategy, then Skeptic is testing $\mathbb{P}$.

4. Predictions that pass statistical tests (defensive forecasting)

Skeptic announces $\mathcal{K}_0 \in \mathbb{R}$.
FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces $L_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + L_n(y_n - p_n)$.

- Forecaster predicts with feedback.
- Skeptic tests Forecaster with safe strategy (trying to multiply capital risked by large factor).

Fix Skeptic's strategy, taking him out of the game.

Given a safe strategy $\mathcal{S}$ for Skeptic and moves $p_1, y_1, \ldots, p_{n-1}, y_{n-1}$ by Skeptic's opponents, define $\mathcal{S}_n : [0, 1] \to \mathbb{R}$ by

$$\mathcal{S}_n(p) := \mathcal{S}(p_1, y_1, \ldots, p_{n-1}, y_{n-1}, p).$$

FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Reality announces $y_n \in \{0, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + \mathcal{S}_n(p_n)(y_n - p_n)$.

FOR $n = 1, 2, \ldots$:

Forecaster announces $p_n \in [0, 1]$.

Reality announces $y_n \in \{0, 1\}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + \mathcal{S}_n(p_n)(y_n - p_n)$.

Takemura's lemma says Forecaster can block any particular continuous strategy for Skeptic.

Takemura's lemma: If $\mathcal{S}_n$ is always continuous, then Forecaster has a strategy that ensures $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \geq \cdots$.

Proof By the intermediate-value theorem, the continuous function $\mathcal{S}_n$ is always positive, always negative, or else satisfies $\mathcal{S}_n(p) = 0$ for some $p \in [0, 1]$. So Forecaster can use the following strategy:

- if $\mathcal{S}_n$ is always positive, take $p_n := 1$;
- if $\mathcal{S}_n$ is always negative, take $p_n := 0$;
- otherwise, choose p_n so that $\mathcal{S}_n(p_n) = 0$.

This guarantees that $\mathcal{S}_n(p_n)(y_n - p_n) \leq 0$, so that $\mathcal{K}_n \leq \mathcal{K}_{n-1}$. ∎

FOR $n = 1, 2, \ldots$:

Forecaster announces $p_n \in [0, 1]$.

Reality announces $y_n \in \{0, 1\}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + \mathcal{S}_n(p_n)(y_n - p_n)$.

Takemura's lemma says Forecaster can block any particular continuous strategy for Skeptic.

Takemura's lemma: If $\mathcal{S}_n$ is always continuous, then Forecaster has a strategy that ensures $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \geq \cdots$.

Question 1. Why assume continuity in Forecaster's last move?

- Skeptic can test all the classical probability properties with continuous strategies.
- If you don't like continuity, just let Forecaster hide p_n's zillionith decimal place by randomizing a tad.

Question 2. Why is it enough for Forecaster to defeat a single particular strategy for Skeptic?

- For the probabilities to look good, it is enough to pass a few dozen tests (e.g., y=1 about 40% of the times when p≈0.4). Forecaster can average these few dozen strategies and make sure that the average does not make money.

More general formulation

1. Auxiliary information

3. Skeptic chooses any payoff with expected value 0 or less.

DEFENSIVE FORECASTING PROTOCOL

FOR $n = 1, 2, \dots$:

Reality announces $x_n \in \mathbf{X}$.

Skeptic announces a lower semicontinuous $F_n : \mathbf{Y} \times \mathcal{P}(\mathbf{Y}) \to \mathbb{R}$ such that $\int_{\mathbf{Y}} F_n(y, P) P(\mathrm{d}y) \leq 0$ for all $P \in \mathcal{P}(\mathbf{Y})$.

Forecaster announces $P_n \in \mathcal{P}(\mathbf{Y})$.

Reality announces $y_n \in \mathbf{Y}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + F_n(y_n, P_n)$.

END FOR.

2. Forecaster announces probability distribution on outcome space $\mathbf{Y}$.

Lemma (Takemura) *Let* $\mathbf{Y}$ *be a metric compact. In the defensive forecasting protocol, Forecaster can play in such a way that Skeptic's capital never increases, no matter how he and Reality play.*

From Working Paper 17, www.probabilityandfinance.com.

Why we thought successful probability forecasting is not always possible.

FOR $n = 1, 2, \ldots$
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  Skeptic's profit $:= s_n(y_n - p_n)$.

Reality can make Forecaster uncalibrated by setting

$$y_n := \begin{cases} 1 & \text{if } p_n < 0.5 \\ 0 & \text{if } p_n \geq 0.5 \end{cases}$$

Skeptic can then make steady money with

$$s_n := \begin{cases} 1 & \text{if } p < 0.5 \\ -1 & \text{if } p \geq 0.5 \end{cases}$$

But here Skeptic's strategy is not continuous.

5. Implications for nonstationary prediction

Defensive forecasting shows that successful on-line prediction tells us about the past, not the future.

So what should we think about the recurrent efforts to make it work?

Randomly selected work on nonstationary prediction

- Vitaly Kuznetsov and Mehryar Mohri. Learning theory and algorithms for forecasting non-stationary time series. *Advances in Neural Information Processing Systems (NIPS 2015)*. Montreal, Canada, December 2015. Machine learning.

- Piotr Fryzlewicz, Sébastien Van Bellegem, and Rainer von Sachs. Forecasting non-stationary time series by wavelet process modelling, *Annals of the Institute of Statistical Mathematics* 55(4):737-764, 2003. Wavelets.

- Simon Haykin and Liang Li. Nonlinear adaptive prediction of nonstationary signals, *IEEE Transactions on Signal Processing*, 43(2):526-535, 1995. Neural networks.

Example: non-stationary macroeconomic forecasts

Recurrent failure to predict the business cycle:

1. 1929: Business cycle institutes folded across the globe.
2. 1950s: Cowles commission quietly gave up.
3. 1970s: Large simultaneous equation models failed.
   (Simple Box-Jenkins time-series models predict as well or better.)
4. 2008: Modern Bayesian DSGE (dynamic stochastic general equilibrium) models failed spectacularly.

History of econometrics

Mary Morgan, *The History of Econometric Ideas*, Cambridge. 1990
Early history, culminating in formation of the Econometric Society and *Econometrica* in the 1930s and Haavelmo's 1944 article on the probability approach.

Roy Epstein, *A History of Econometrics*, North-Holland. 1987
Failed efforts to predict the business cycle from Cowles Commission in the 1940s through the 1970s.

Duo Qin, *A History Econometrics: The Reformation from the 1970s*, Oxford. 2013
Three threads of thought coming out of the failures of 1970s:

- VAR (vector autoregression); rational expectations; Christopher Sims.
- Bayes. First championed for model selection, then applied to DSGE.
- LSE school. David Hendry. Closer to Cowles tradition.

Macro-econometrics in the 2000s

The chief economist for the world bank declares modern macroeconomic theory (DSGE) to be Bayesian nonsense: so many parameters that the prior dominates.
The trouble with macroeconomics, Paul Romer, 2016.

DSGE models could not predict the 2008 crisis or its aftermath.
Challenges for Central Banks' Macro Models, Jesper Lindé, Frank Smets, and Rafael Wouters, 2016.

Hendry claims that nonstationary modelling is the solution.
All Change! The Implications of Non-stationarity for Empirical Modelling, Forecasting and Policy, David F. Hendry and Felix Pretis, 2016.

Does the failure of stationary prediction imply a nonstationary stochastic “generative” mechanism?

My answers:

- There is no justification for “generative” talk.
- Better to say that there is no “generative” mechanism at all.

- We are observing the results of a complex game.
- Outcomes may or may not have certain emergent regularities.

6. Implications for market efficiency

Recent work in game-theoretic probability (see especially the summary in Working Paper 47), shows that we can reconstruct the Black Scholes model (modulo a change in time) starting merely from the assumption that the market index (e.g., the S&P 500) is efficient in the game-theoretic sense (see slides in Appendix).

This can provide a foundation for Platen and Heath‘s real world pricing or Föllmer‘s pathwise pricing.

The success of defensive forecasting suggests how the game-theoretic efficiency of a market index might arise. Can this be substantiated, theoretically or experimentally?

This is a call for research.

References

A Benchmark Approach to Quantitative Finance. Eckhard Platen and David Heath, Springer, 2006.

We ask whether there exists a strictly positive process, for instance, a market index, which when used as numeraire or benchmark, generates realistic benchmarked derivative price processes that are martingales with respect to the real world probability measure.

Pathwise integration with respect to paths of finite quadratic variation. Anna Ananova and Rama Cont. *Journal de Mathématiques Pures et Appliquées*, 2016.

In his seminal paper "Calcul d'Ito sans probabilités" [12], Hans Föllmer proved a change of variable formula for smooth functions of paths with infinite variation, using the concept of quadratic variation along a sequence of partitions.

Appendix: Game-Theoretic Explanation of Equity Premium

The equity premium puzzle

- Returns from stocks are about 6 percentage points better than returns from bonds.
- Risk aversion can account for only about 1 percentage point.

Game-theoretic explanation

- Speculation causes volatility.
- Speculation makes market efficient.
- Speculation forces an efficient market to appreciate in proportion to the square of its volatility.

Three roles of speculation

- Speculation causes volatility.
  Traders know this, though some academic literature wants to attribute volatility to information.

- Speculation makes market efficient.
  Conventional wisdom, even in academia.

- Speculation forces an efficient market to appreciate in proportion to $(\text{volatility})^2$.
  This is our theoretical contribution.

Three roles of speculation

- Speculation causes volatility. Traders and experts in option pricing agree.
- Speculation makes the market efficient by exhausting opportunities for low-risk profit. An investor can rarely do better than hold all tradables in proportion to their capitalization.
- Assuming that you can trade an index that holds all tradables in proportion to their capitalization, speculation forces this index to appreciate in proportion to the square of its volatility.

John Hull, author of leading textbook on option pricing:

What Causes Volatility?

It is natural to assume that the volatility of a stock is caused by new information reaching the market. This new information causes people to revise their opinions about the value of the stock. The price of the stock changes and volatility results. This view of what causes volatility is not supported by research.

The only reasonable conclusion is that volatility is to a large extent caused by trading itself. (Traders usually have no difficulty accepting this conclusion.)

What is an efficient market?

- Fama 1965: Prices incorporate all information.
- Shafer/Vovk 2001: No strategy selected in advance multiplies capital risked by large factor.

Why should a market be efficient?

- Fama: Speculators use each bit of new information.
- Shafer/Vovk: Speculators are using every trick to multiply their capital, not merely exogenous information.

How do we test whether a market is efficient?

- Fama: Postulate a model and test it statistically.
- Shafer/Vovk: Try to multiply your capital in the market.

How do we test whether a market is efficient?

Try to multiply your capital in the market.

- Define a trading strategy and implement it.
- If you multiply your money by 1000, reject the hypothesis of efficiency.
- Confidence of rejection same as when you reject a hypothesis at significance 0.001.

THE EFFICIENT INDEX HYPOTHESIS (EIH)

You will not multiply the capital you risk by a large factor relative to an index defined by the total value of all the readily tradable assets.

To fix ideas, suppose the index is the S&P500.

ETF Symbol	ETF Name	Fees, per year
IVV	iShares Core S&P 500	4 bps
SPY	SPDR S&P 500	11 bps
VOO	Vanguard S&P 500	5 bps

Our mathematical story

We have argued that speculation causes volatility, and that speculation makes the market efficient, in the sense that the market index will not be beat.

This is the efficient index hypothesis.

Using the efficient market hypothesis, we now prove mathematically that the market index must grow in proportion to the variance of the index.

Assume zero interest rate.

For traders, “cash” is a money-market account that pays the short-term risk-free interest rate.

Use the accumulated value of $1 in such an account as the ***numéraire*** for measuring the value of other financial instruments.

Mathematically, this is equivalent to assuming that the interest rate is zero.

Efficient Index Hypothesis (EIH)

You will not multiply the capital you risk by a large factor relative to an index I.

Volatility and Variance

Suppose the value of the index over N times periods is $I_0, I_1, \ldots, I_N$. The *returns* $m_1, m_2, \ldots, m_N$ are defined by

$$m_n := \frac{I_n}{I_{n-1}} - 1 = \frac{I_n - I_{n-1}}{I_{n-1}}.$$

for $n = 1, \ldots, N$. The *relative quadratic variation* is

$$\Sigma_N := \sum_{n=1}^{N} m_n^2.$$

The *cumulative volatility* is $\sqrt{\Sigma_N}$.

Measure time by accumulated variance.

To fix ideas, consider daily returns, so N is the number of days.

The relative quadratic variation Σ_N is usually approximately proportional to the amount of time elapsed.

- Explanation by probability theory: The m_n are random and independent. Each is random with mean 0 and standard deviation σ. So $\Sigma_N := \sum_{n=1}^{N} m_n^2$ is an estimate of $N\sigma^2$.
- Same conclusion from EIH without probability assumptions.

But volatility does vary (σ changes). It is greater when traders get excited, for whatever reason.

To keep the mathematics simple, we use Σ_N as our clock!! In other words, we measure time by the amount of trading.

Pass to continuous time

Makes picture mathematically elegant.

- Mathematical finance now uses measure-theoretic continuous-time probability.

- Instead, we use game-theoretic continuous-time probability.

Assuming continuous time...

- Measure time by cumulative variance $\Sigma = \sum m_n^2$.
- Write I_s for the value of index at time s, where $s = \Sigma$.
- Assume $I_0 = 1$.

Then the EIH implies that I_s will look like

$$I_s = \exp\left(\frac{s}{2} + W_s\right),$$

where W_s is Brownian motion.

The EIH implies that I_s will look like

$$I_s = \exp\left(\frac{s}{2} + W_s\right),$$

where W_s is Brownian motion.

Geometric Brownian motion with drift 1 & volatility 1.

$$\mathbf{E}(\ln I_s) = \frac{s}{2}$$

$$\mathbf{s.\,d.}(\ln I_s) = \sqrt{s}$$

For large s,

$$\ln I_s \approx \frac{s}{2}.$$

When time is measured by cumulative variance Σ,

$$\ln I_s \approx \frac{s}{2}, \text{ where } s = \Sigma.$$

or

$$\ln I_\Sigma \approx \frac{\Sigma}{2}.$$

In terms of calendar time t.

$$\ln I_t \approx \frac{\Sigma_t}{2},$$

where Σ_t is the cumulated variance at time t.

Average return overestimates growth in value by half the variance.

Suppose m_n is the return for period n,

$$\mathrm{M} = \sum_n m_n, \qquad \Sigma = \sum_n m_n^2,$$

and I is the accumulated value when you start with one unit. Then

$$\ln I = \ln \prod_n (1 + m_n) = \sum_n \ln(1 + m_n).$$

By the Taylor expansion $\ln(1 + x) \approx x - \frac{1}{2}x^2$,

$$\ln I \approx \sum_n \left(m_n - \frac{1}{2} m_n^2 \right) = M - \frac{\Sigma}{2}.$$

By our theory (still to be explained),

$$\ln I_t \approx \frac{\Sigma_t}{2},$$

By the properties of the logarithm,

$$M_t \approx \ln I_t + \frac{\Sigma_t}{2}.$$

So

$$M_t \approx \Sigma_t$$

This is the equity premium.

$$M_t \approx \Sigma_t.$$

The annualized volatility of the S&P 500 is approximately 20% ([10], page 8). Squaring this, we obtain an equity premium of 4%. This is closer to the empirical estimates than the 1% obtained from standard theory, and GTP38 shows that it is within (3)'s anticipated error of approximation (Section 4).

$$\mathrm{M}_t \approx \Sigma_t.$$

How does the EIH implies this equity premium?

Answer: There are strategies that can beat the index (multiply your capital by a large factor relative to the index) if the approximation does not hold.

The trading strategy

Suppose I grows faster than our theory predicts:

$$\mathrm{M}_t >> \Sigma_t.$$

To make money, you invest all you have in I and borrow money to invest even more.

Say you always invest $(1+\epsilon) \times$ (current capital) in I. Then on round n, when I is multiplied by $1+m_n$, your capital is multiplied by $1+(1+\epsilon)m_n$. Relative to I, your capital is multiplied by

$$\frac{1+(1+\epsilon)m_n}{1+m_n}.$$

Use Taylor's series for the logarithm again:

$$\begin{aligned}\ln \frac{1+(1+\epsilon)m_n}{1+m_n} &= \ln(1+(1+\epsilon)m_n) - \ln(1+m_n) \\ &\approx \epsilon m_n - \epsilon m_n^2 - \frac{\epsilon^2}{2} m_n^2.\end{aligned}$$

multiplied by $1 + (1+\epsilon)m_k$. So your capital will increase (or decrease) relative to I by the factor

$$\frac{1 + (1+\epsilon)m_k}{1 + m_k}.$$

Using Taylor's series for the logarithm, we obtain the approximation

$$\ln \frac{1 + (1+\epsilon)m_k}{1 + m_k} \approx \epsilon m_k - \epsilon m_k^2 - \frac{\epsilon^2}{2} m_k^2.$$

So over K rounds, your capital will grow relative to I by a factor whose logarithm is approximately

$$\begin{aligned} \epsilon \sum_{k=1}^{K} m_k - \epsilon \sum_{k=1}^{K} m_k^2 - \frac{\epsilon^2}{2} \sum_{k=1}^{K} m_k^2 &= \epsilon \mathrm{M}_t - \epsilon \Sigma_t - \frac{\epsilon^2}{2} \Sigma_t \\ &= \epsilon(\mathrm{M}_t - \Sigma_t) - \frac{\epsilon^2}{2} \Sigma_t, \qquad (5) \end{aligned}$$

So over the entire time period, your capital will grow relative to I by a factor whose logarithm is approximately

$$\begin{aligned}\epsilon\sum_n m_n - \epsilon\sum_n m_n^2 - \frac{\epsilon^2}{2}\sum_n m_n^2 &= \epsilon \mathrm{M}_t - \epsilon\Sigma_t - \frac{\epsilon^2}{2}\Sigma_t \\ &= \epsilon(\mathrm{M}_t - \Sigma_t) - \frac{\epsilon^2}{2}\Sigma_t.\end{aligned}$$

This factor will be large if you continue until Σ_t is so large that $\epsilon\Sigma_t$ is large even though ϵ is small, and if M_t then exceeds Σ_t substantially.

Example: $\epsilon = 0.01$, $\epsilon\Sigma_t = 3$, $\mathrm{M}_t \approx 1.5\Sigma_t$. Then

$$\epsilon(\mathrm{M}_t - \Sigma_t) - \frac{\epsilon^2}{2}\Sigma_t \approx 1.5,$$

so that you have multiplied your capital relative to I by $e^{1.5} \approx 4.5$.

To similarly make money if I grows too slowly—i.e., if M_t is substantially less than Σ_t, you can take ϵ in the preceding argument to be a small negative number. In other words, you keep a small fixed fraction of your capital in the risk-free bond on each round, investing the rest in I.

You can implement the two strategies simultaneously: put half your initial capital on one of them and half on the other. So you have a strategy that will mulitply its initial capital substantially relative to I unless $\mathrm{M}_t \approx \Sigma_t$. (We promised a strategy that multiplies *the capital it risks*, so you need to implement the strategy just sketched in a way that risks no more than its initial capital. You can do this by stopping the strategy if its capital gets close to zero. In GTP44 we rely on the assumption that the price path is continuous to make sure we can stop in time. Weaker assumptions can also be accommodated.)

What are the macroeconomic implications?

The market represents one portion of society's productive capital—the portion that is so liquid that we can speak of its volatility. How do changes in the valuation of this portion of society's capital relative to cash[7] affect its valuation relative to the portion of society's productive capital that is not so liquid?

Consider an extended period in which the publicly traded portion of the economy is exceptionally productive, so that the value of I is growing because of economic fundamentals at a rate exceeding its volatility. A naive expectation is that this exceptional growth will draw investors into the market, creating a demand for cash. The speculative strategy described in Section 4.3, if widely played, would reinforce this demand for cash. Half of its initial capital is invested according to the strategy (4), and since its current capital will grow, the amount $\epsilon \times$ (current capital) that it borrows will grow. The other half of its initial capital will be invested according to

$$(1 - \epsilon) \times (\text{current capital}),$$

but this current capital will grow more slowly, and so the amount $\epsilon \times$ (current capital) that it keeps in cash will be growing more slowly. The demand for cash will presumably drive up the interest rate. Since the higher interest rate will not be justifed by the productivity of capital outside the market, this might spur inflation. Restoration of stability might require slowing the productivity of the publicly-traded portion of the economy relative to the privately held portion, perhaps by taking some corporations private.

In an extended period of slow productivity for the publicly traded portion of the economy, which falls short of justifying an increase in capitalization commensurate with volatility, we would see pressures in the opposite direction. As the growth in I lags its volatility, the strategy described in Section 4.3, perhaps together with more aggressive strategies that sell all or parts of the market short, could produce an excess supply of cash, driving down the interest rate and creating deflationary pressure. As this process continues, wealth would be increasingly concentrated in the hands of those who own the assets in the market. Barring an increase in productivity, these pressures might eventually be released by a financial crisis that violates the EIH, suspending consequences such as $\mathrm{M} \approx \Sigma$ and perhaps durably destroying market capitalization that is not producing a commensurate flow of goods and services.