Modern Portfolio Theory, and Why I Have a Job

chris bemis
Whitebox Advisors

September 1, 2009

Modern portfolio theory establishes

- A necessity for a forecasted reward
- ...and a specified risk over this same period.
- These two quantities are balanced against each other in some 'optimal' way.

Modern portfolio theory establishes

- A necessity for a forecasted reward
- ...and a specified risk over this same period.
- These two quantities are balanced against each other in some 'optimal' way.

Modern portfolio theory establishes

- A necessity for a forecasted reward
- ...and a specified risk over this same period.
- These two quantities are balanced against each other in some 'optimal' way.

Modern portfolio theory establishes

- A necessity for a forecasted reward
- ...and a specified risk over this same period.
- These two quantities are balanced against each other in some 'optimal' way.

Modern portfolio theory establishes

- A necessity for a forecasted reward
- ...and a specified risk over this same period.
- These two quantities are balanced against each other in some ’optimal’ way.

In the final analysis, the vagueness of the above is why portfolio optimization is a tool rather than a solution.

Markowitz in 1952 established what is likely the first portfolio optimization problem; namely

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

A QP problem with linear constraints.
It was a significant achievement, and it has been widely adopted. Enough so that we know it doesn't work.

Markowitz in 1952 established what is likely the first portfolio optimization problem; namely

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

A QP problem with linear constraints.
It was a significant achievement, and it has been widely adopted. Enough so that we know it doesn't work.

Markowitz in 1952 established what is likely the first portfolio optimization problem; namely

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

A QP problem with linear constraints.
It was a significant achievement, and it has been widely adopted. Enough so that we know it doesn't work.

Markowitz in 1952 established what is likely the first portfolio optimization problem; namely

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

A QP problem with linear constraints.
It was a significant achievement, and it has been widely adopted. Enough so that we know it doesn't work.

Markowitz in 1952 established what is likely the first portfolio optimization problem; namely

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

A QP problem with linear constraints.
It was a significant achievement, and it has been widely adopted. Enough so that we know it doesn't work.

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

This formulation is normative and not empirical.

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

This formulation is normative and not empirical.
Financial data are nonstationary.

The above is generally referred to as mean-variance optimization.
In terms of the original discussion, we are equating

- Reward with expected return
- Risk with portfolio variance
- We are minimizing risk subject to a minimum expected return

This formulation is normative and not empirical.
Financial data are nonstationary.
Additionally, measurements have error.

Why use the variance and mean?

The normal distribution is both widely witnessed in the physical sciences, and is also tractable mathematically. Not only is it determined by only two parameters, it is *the* distribution obtained in summed i.i.d. experiments. It looks like:

Why use the variance and mean?
The normal distribution is both widely witnessed in the physical sciences, and is also tractable mathematically.
Not only is it determined by only two parameters, it is *the* distribution obtained in summed i.i.d. experiments. It looks like:

Why use the variance and mean?

The normal distribution is both widely witnessed in the physical sciences, and is also tractable mathematically.

Not only is it determined by only two parameters, it is *the* distribution obtained in summed i.i.d. experiments. It looks like:

Why use the variance and mean?
The normal distribution is both widely witnessed in the physical sciences, and is also tractable mathematically.
Not only is it determined by only two parameters, it is *the* distribution obtained in summed i.i.d. experiments. It looks like:

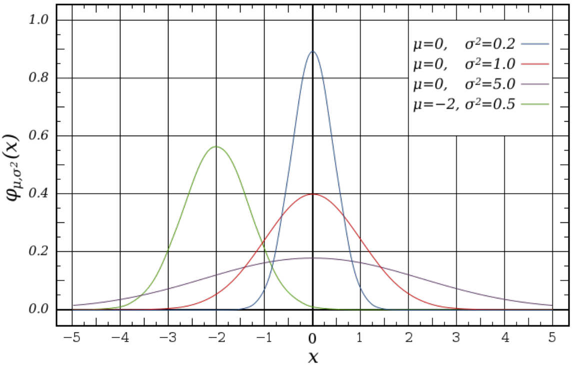

So what types of things are independent, identically distributed experiments?

Think of gambling:

- Coin flips
- Spins of a roulette wheel

In general, what we are looking for by saying some random variables are iid is that they all share the same distribution.

So what types of things are independent, identically distributed experiments?
Think of gambling:

- Coin flips
- Spins of a roulette wheel

In general, what we are looking for by saying some random variables are iid is that they all share the same distribution.

So what types of things are independent, identically distributed experiments?
Think of gambling:

- Coin flips
- Spins of a roulette wheel

In general, what we are looking for by saying some random variables are iid is that they all share the same distribution.

So what types of things are independent, identically distributed experiments?
Think of gambling:

- Coin flips
- Spins of a roulette wheel

In general, what we are looking for by saying some random variables are iid is that they all share the same distribution.

So what types of things are independent, identically distributed experiments?
Think of gambling:

- Coin flips
- Spins of a roulette wheel

In general, what we are looking for by saying some random variables are iid is that they all share the same distribution.

In time series, we often think of observations through time as random variables.

The time series extension of the iid idea is to want to have the joint distribution of time lagged variables to only depend on the lag - not the time.

For example, we might expect that the correlation between two stocks' returns might be the same over distinct 200 day periods.

Or, a covariance matrix might be assumed to remain the same in the future.

Stopping at the second moment is an example of weak stationarity.

In time series, we often think of observations through time as random variables.

The time series extension of the iid idea is to want to have the joint distribution of time lagged variables to only depend on the lag - not the time.

For example, we might expect that the correlation between two stocks' returns might be the same over distinct 200 day periods.

Or, a covariance matrix might be assumed to remain the same in the future.

Stopping at the second moment is an example of weak stationarity.

In time series, we often think of observations through time as random variables.

The time series extension of the iid idea is to want to have the joint distribution of time lagged variables to only depend on the lag - not the time.

For example, we might expect that the correlation between two stocks' returns might be the same over distinct 200 day periods.

Or, a covariance matrix might be assumed to remain the same in the future.

Stopping at the second moment is an example of weak stationarity.

In time series, we often think of observations through time as random variables.
The time series extension of the iid idea is to want to have the joint distribution of time lagged variables to only depend on the lag - not the time.
For example, we might expect that the correlation between two stocks' returns might be the same over distinct 200 day periods.
Or, a covariance matrix might be assumed to remain the same in the future.
Stopping at the second moment is an example of weak stationarity.

The question arises...

Can we test if financial time series are weakly stationary?

We don't have to look very far:

The question arises...

Can we test if financial time series are weakly stationary?

We don't have to look very far:

The question arises...
Can we test if financial time series are weakly stationary?
We don't have to look very far:

The question arises...
Can we test if financial time series are weakly stationary?
We don't have to look very far:

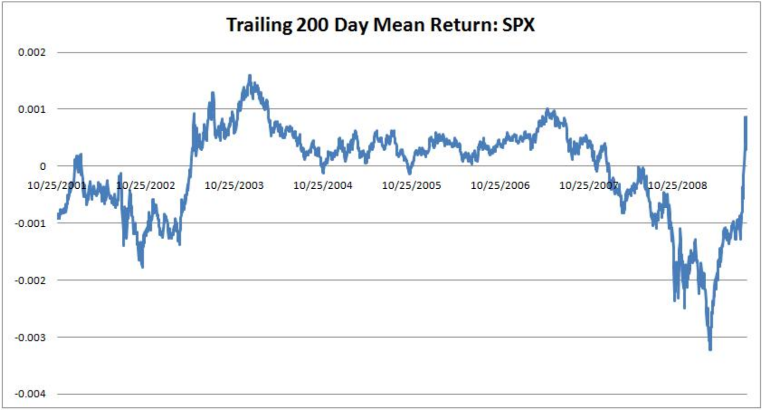

Putting aside practical questions for a moment, there are interesting features of the problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

like, what happens when we solve the problem for various α's?

Putting aside practical questions for a moment, there are interesting features of the problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

like, what happens when we solve the problem for various α's?

Putting aside practical questions for a moment, there are interesting features of the problem

$$\begin{aligned} \text{minimize} \quad & \mathbf{w}'\Sigma\mathbf{w} \\ \text{subject to} \quad & \mu'\mathbf{w} \geqslant \alpha \\ & \mathbf{1}'\mathbf{w} = 1 \\ & \mathbf{w} \geqslant 0, \end{aligned}$$

like, what happens when we solve the problem for various α's?

We get something called the efficient frontier.

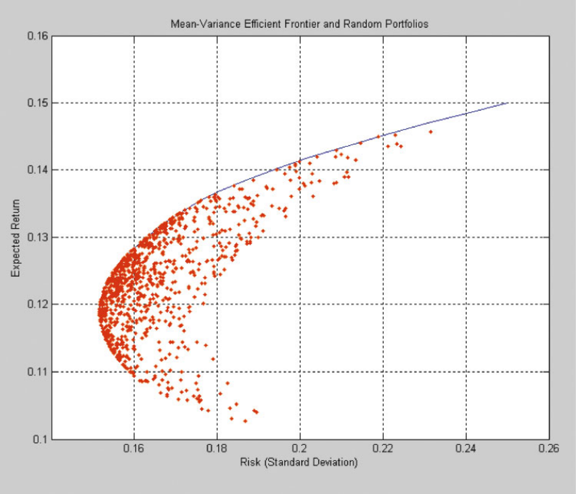

Now, what if there is a risk free asset. One with no volatility. We get the 'market portfolio' and β

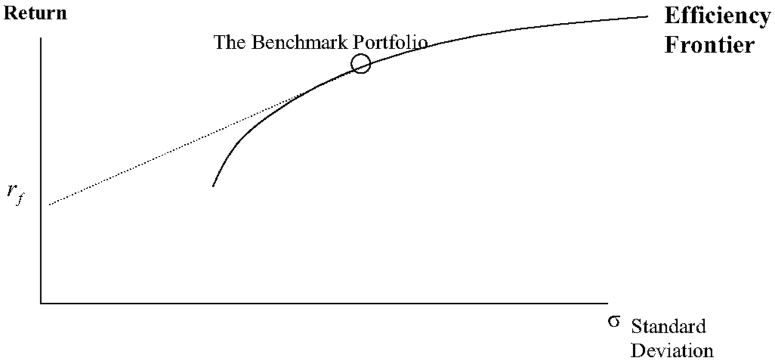

If we put on our normative hats, we start making claims:

- You'd always want a portfolio *on* the efficient frontier
- If there is a risk free asset, you'd always want to be on the market line
- Ergo, all risk is systemic and you just need to figure out how much of the market portfolio you want
- huh? Well,

$$r_i - r_f = \beta(r_m - r_f)$$

If we put on our normative hats, we start making claims:

- You'd always want a portfolio *on* the efficient frontier
- If there is a risk free asset, you'd always want to be on the market line
- Ergo, all risk is systemic and you just need to figure out how much of the market portfolio you want
- huh? Well,

$$r_i - r_f = \beta(r_m - r_f)$$

If we put on our normative hats, we start making claims:

- You'd always want a portfolio *on* the efficient frontier
- If there is a risk free asset, you'd always want to be on the market line
- Ergo, all risk is systemic and you just need to figure out how much of the market portfolio you want
- huh? Well,

$$r_i - r_f = \beta(r_m - r_f)$$

If we put on our normative hats, we start making claims:

- You'd always want a portfolio *on* the efficient frontier
- If there is a risk free asset, you'd always want to be on the market line
- Ergo, all risk is systemic and you just need to figure out how much of the market portfolio you want
- huh? Well,

$$r_i - r_f = \beta(r_m - r_f)$$

If we put on our normative hats, we start making claims:

- You'd always want a portfolio *on* the efficient frontier
- If there is a risk free asset, you'd always want to be on the market line
- Ergo, all risk is systemic and you just need to figure out how much of the market portfolio you want
- huh? Well,

$$r_i - r_f = \beta(r_m - r_f)$$

The result is neat.

You have to ask yourself, though, what made such an oversimplified statement possible?

Everything in the setting above, the Capital Asset Pricing Model, assumes that means are stable through time. And so are covariances.

The good news for the future practitioner: they're not. At all.

The result is neat.
You have to ask yourself, though, what made such an oversimplified statement possible?
Everything in the setting above, the Capital Asset Pricing Model, assumes that means are stable through time. And so are covariances.
The good news for the future practitioner: they're not. At all.

The result is neat.

You have to ask yourself, though, what made such an oversimplified statement possible?

Everything in the setting above, the Capital Asset Pricing Model, assumes that means are stable through time. And so are covariances.

The good news for the future practitioner: they're not. At all.

The result is neat.
You have to ask yourself, though, what made such an oversimplified statement possible?
Everything in the setting above, the Capital Asset Pricing Model, assumes that means are stable through time. And so are covariances.
The good news for the future practitioner: they're not. At all.

A disclaimer: this view on the origins of modern portfolio theory may not be the norm. In fact, Markowitz was quoted as saying "Diversifying sufficiently among uncorrelated risks can reduce portfolio risk toward zero" as recently as 2008. However,

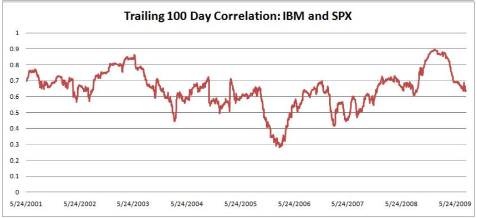

And in case that doesn't seem fair:

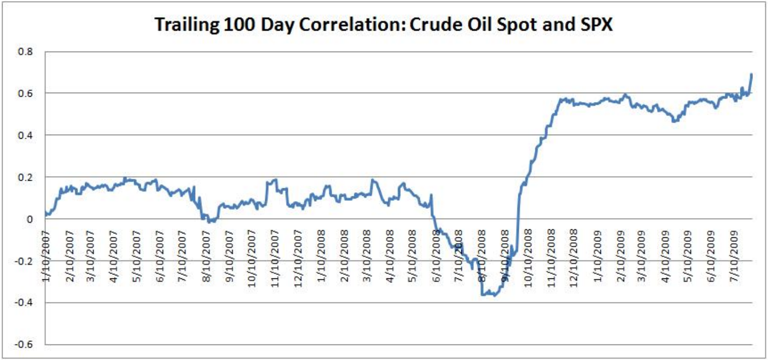

There are times when ”correlation goes to one.” Hence true diversification may fail exactly when you need it most.

We see here the inherit importance of understanding our underlying; e.g., beginning with the premise that our data is nonstationary.

We also see the need to mitigate definable risks according to our understanding (model).

This is an empirical exercise and not a normative one.

The opportunity set for those willing to understand finance through the lens of an empirical mathematician is riper than ever.

There are times when "correlation goes to one." Hence true diversification may fail exactly when you need it most.

We see here the inherit importance of understanding our underlying; e.g., beginning with the premise that our data is nonstationary.

We also see the need to mitigate definable risks according to our understanding (model).

This is an empirical exercise and not a normative one.

The opportunity set for those willing to understand finance through the lens of an empirical mathematician is riper than ever.

There are times when "correlation goes to one." Hence true diversification may fail exactly when you need it most.

We see here the inherit importance of understanding our underlying; e.g., beginning with the premise that our data is nonstationary.

We also see the need to mitigate definable risks according to our understanding (model).

This is an empirical exercise and not a normative one.

The opportunity set for those willing to understand finance through the lens of an empirical mathematician is riper than ever.

There are times when "correlation goes to one." Hence true diversification may fail exactly when you need it most.
We see here the inherit importance of understanding our underlying; e.g., beginning with the premise that our data is nonstationary.
We also see the need to mitigate definable risks according to our understanding (model).
This is an empirical exercise and not a normative one.
The opportunity set for those willing to understand finance through the lens of an empirical mathematician is riper than ever.

There are times when "correlation goes to one." Hence true diversification may fail exactly when you need it most.
We see here the inherit importance of understanding our underlying; e.g., beginning with the premise that our data is nonstationary.
We also see the need to mitigate definable risks according to our understanding (model).
This is an empirical exercise and not a normative one.
The opportunity set for those willing to understand finance through the lens of an empirical mathematician is riper than ever.

fin.