Efficient Frontier

An Online Journal of Practical Asset Allocation

Edited by William J. Bernstein

January 1998

Table of Contents

Efficient Frontier

Efficient Frontiers Logo

The Bearded Ladies Investment Club

Paul W. Harvey

The Bearded Ladies have become famous in recent years for outpacing both the S&P 500 and most male-dominated investment clubs. I dropped in on them last Tuesday to see if they would give me some investment results.

I noticed two things when I showed up at Zelda Delderfield's modest bungalow in Stashabuck Hieghts:

1. These ladies all have *beards*, real ones. I asked if it was genetic, and Zelda replied airily that she sure hoped so, because there was good money to be made in sideshows from her condition. "Besides," she added, "I spend 120 hours a peek on investing, and wouldn't have time to shave even if I wanted to."

Nobody was wasting time making cookies. "We're obsessed with investing, and proud of it," said club member Velda Vivvelfarb.

Okay, so I noticed more than two things:

3. Unlike some groups that have 30 or more members, this group has thrived with a compact five members, of which only four are human (more about that later).

4. But despite having only five members, the group does not limit itself to only one investment method, as we shall soon see.

Each member of the group manages 20% of the portfolio, wi ointerference from the others Zelda is a dedicated Ben Graham/Warren Biffet investor: she scrutinizes annual reportf looking for companies managed by people named Benjamin, Warren, Graham, or Buffet (as surnames or given names, as the case may be).

Velda is a socially responsible investor -- sh gets herself invited to corporate social events (teas, Christmas parites, shareholder meetings), and then rates them according to quality and range of refreshments, décor, etc.

Imelda Weldenblatt likes momentum plays -- she writes the names of companies she is considering on bowling balls, and rolls them down the hill to test their momentum.

Rwelda Elverveldt is the club's most daring investor -- she uses astrology to guide her stock selection. Surprisingly, she has the besr record of the bunch. This amazed me at first, but the other ladies explaned that Rwelda had a Wharton MBA and 15 years experience in a bank trust department. "Plus, I get very lucky sometimes," she added sweetly.

Finally, there is Lady Luck, the club's pet chimpanzee, who chooses stocks by throwing darts at a copy of the NYSE listings that has been pasted to the wall. Lady Luck has done very well.

All of which goes to show that you can always buy cookies, but for serious dough you want to have the Bearded Ladies in your corner.

Paul Harvey exercizes his investing abilities and imagination in West Roxbury, Massachusetts.

Home E-Mail

The Intelligent Asset Allocator

William J. Bernstein

The January 1998 Coward's Portfolio

The January 1998 Coward's Review contains no surprises. As in previous issues, we compare the risk/return characteristics of our automatons with the performance of the "multiasset global" and "asset allocation" funds in the Morningstar universe. To review, the 4 portfolios we have studied in the past are fixed mixes of the following indexes/funds, rebalanced quarterly:

The Coward's Portfolio (CEI)

- **20% S&P 500**
- **20% US small stocks (DFA US 9-10 Portfolio)**
- **15% EAFE-Europe**
- **5%EAFE Pac. Ex Japan**
- **5% Japan Large (MSCI Japan)**
- **10% Continental Small (DFA Cont. Sm. Co. Portfolio)**
- **5%UK small (DFA UK Sm. Co. Portfolio**
- **5% Japan Small (DFA Jap. Sm. Co. Portfolio)**
- **5% Pac. EX Japan small (DFA Pac. Rim Sm. Co. Port., before 1/93 EAFE Pac. X J)**
- **10% Latin American (MSCI Lat. Am.)**

(This portfolio is mixed with the DFA 1 year corporate bond fund to produce a risk/return curve for the past 3, 5, and 10 years.)

The Small Investor's Coward's Portfolio (SICEI)

- **20% Vanguard Index Trust 500**
- **20% Vanguard Small Cap Index Fund**
- **15% Vanguard European Index Portfolio**
- **7% Vanguard Pacific Index Portfolio**
- **8% Vanguard Emerging Markets Index Portfolio**
- **5% Scudder Latin America Fund**
- **12.5% Tweedy Browne Global Value Fund**
- **12.5% Acorn International Fund**

(This portfolio is mixed with the Vanguard Short Term Corporate Bond Fund to produce a

risk/return curve for the past 3 years.)

The Academic Coward's Portfolio (ACEI)

- **25% DFA US Large Cap Value**
- **25% DFA US Small Cap Value**
- **25% DFA Int'l Value**
- **25% DFA Int'l Small Cap Value**

(This portfolio is mixed with the DFA 1 year corporate bond fund to produce a risk/return curve for the past 3 years.)

Lastly, in July I added a fourth entry, the Tweedy Coward, in attempt to see if a disclplined global Graham-and-Dodd approach can beat the robots. The fund which does this in the most rigorous fashion is the Tweedy Browne Global Value Fund, which is mixed with the Vanguard Short Term Corporate Bond Fund, to produce a return/risk spectrum accessible to the small investor.

Remember, the CEI and ACEI are meant as a benchmarks for institutional investors. The SICEI and Tweedy portfolios are provided as portfolios available to the small investor.

With all that out of the way, here are the results for 5 and 10 years (CEI only) and 3 years (all 4 indexes) I apologize for for the messiness of the 3 year graph. There are a lot of data points, and the Tweedy, CEI, and SICEI stock portfolios are located in the same place:

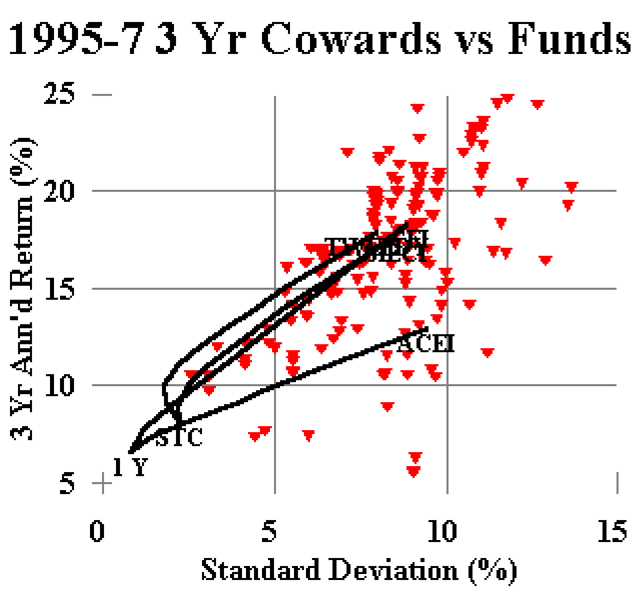

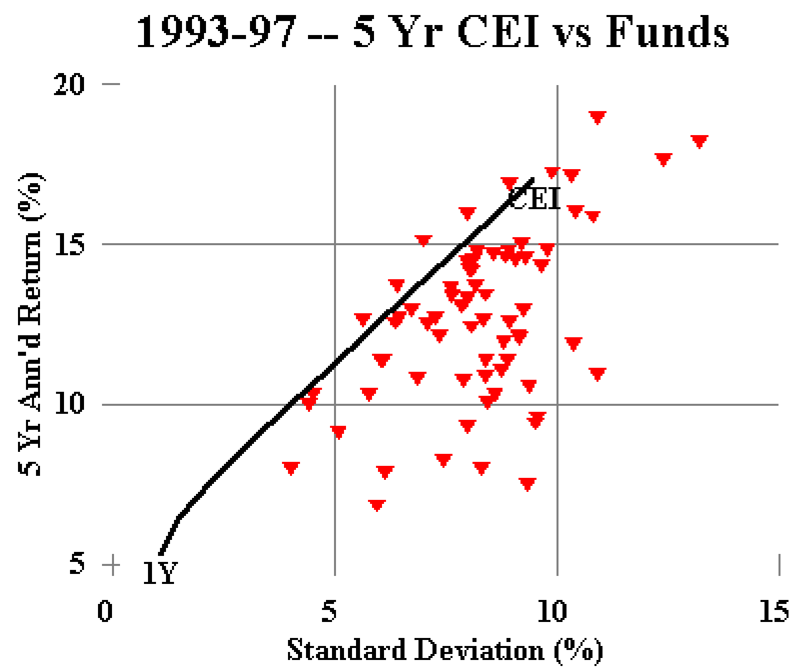

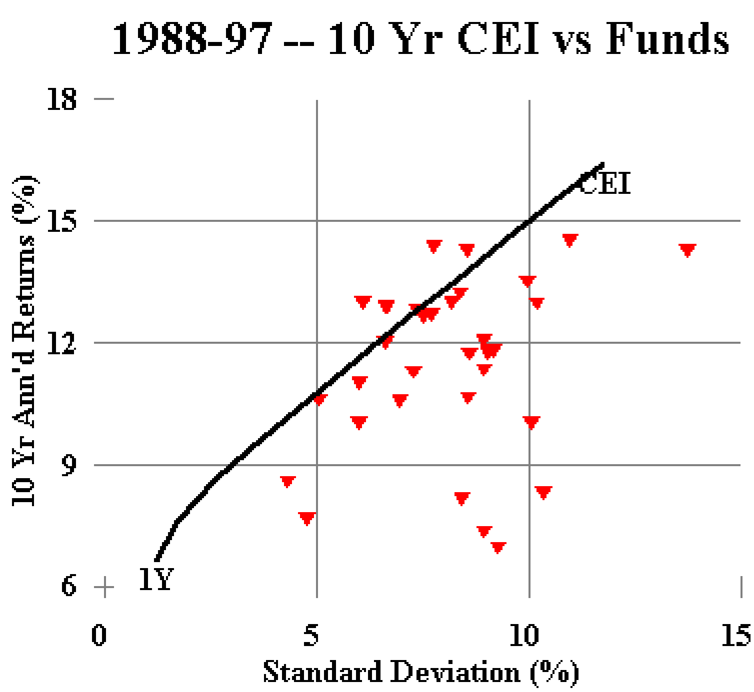

Once again, the cowards have done very well over the 5 and 10 year time frames, with only a few funds poking more than a percent above the curve (Income Fund of America for 5 and 10 years, and Pegasus Managed for 10 years.) Over the 5-10 year periods, the cowards provide risk adjusted return superior to greater than 80% of these funds. Over the 3 year time period the picture is not so pretty, with only the Tweedy Coward doing slightly better than average, the CEI and SICEI below average, and the ACEI near the cellar.

The reason for the poor 3 year showing is obvious -- over the past 3 years US

stocks, large and small, have trounced foreign assets, particularly those in the Pacific Rim, Japan, and Emerging Markets. This is demonstrated in the below table:

Stock Asset Class	3 Yr Returns (%)	5 Yr Returns (%)	10 Yr Returns (%)
US Large	31.13	20.25	18.04
US Small	24.77	19.38	16.44
European Large	22.16	19.24	14.04
Continental Small	8.5	12.19	N/A
UK Small	14.17	15.17	6.81
Japan Large	-13.4	-0.021	-2.9
Japan Small	-30.43	-13.03	-7.05
Pacific Rim Large	-2.06	7.69	10.75
Pacific Rim Small	-13.68	1.73	N/A
Latin America	8.69	12.64	28.12
MSCI Non US	6.78	11.56	6.35

As you can see, over all 3 time periods the place to be has been the good old USA. This has been particularly true in the past 3 years. What is truly remarkable is that the cowards have done as well as they have. Pay particular attention to the difference in returns between the S&P 500 (first row) and the MSCI-Non US index (last row). Over the past 10 years foreign stocks have returned 11.69% less than the S&P on an *annualized* basis, with the gap being 8.69% for 5 years and an astonishing 24.35% for 3 years. Now consider that the average asset allocation/global multiasset fund contains more than 80% US equity, versus 15%-50% US equity for the cowards. Given the data, it is perfectly understandable that the cowards have underperformed for the 3 year time period, and truly astonishing that they have managed to beat nearly all of these funds over 5 and 10 years.

Over the 29 year period beginning January 1, 1969, when Morgan Stanley initiated the foreign EAFE index, the returns of the S&P 500 (12.16%) and the EAFE (12.12%) have been nearly identical. There is no reason to believe that US and non US returns should be that different. If foreign stocks are riskier, then their returns should in fact be higher. Certainly, we have no reason to believe that the US-foreign discrepancies of the past decade are a premanent phenomenon. In fact, history shows that the dominant asset over the previous 5-10 year period tends to underperform in the next period.

So, the cowards are slighly bloodied, but unbowed. In July 1998 5 year data will become available for the SICEI and Tweedy Cowards. Stay tuned.

Efficient Frontier

William J. Bernstein

The Appropriate Use of the Mean Variance Optimizer

An Illustration

Few investment tools are more seductive than the Markowitz mean variance optimizer ("MVO"). For those of you unfamiliar with this beast, it will take the basic characteristics of the portfolio component assets (expected return, standard deviation of returns, and the correlations among all of the assets) and ground out the "efficient frontier" of portfolios -- a spectrum of asset allocations which are maximally efficient. In other words, each portfolio has the maximum return for a given degree of risk, or the minimum risk for a given return. Pretty heady stuff.

In fact, these things are so flashy that most of the commercially available ones are basically gee-whizzes for financial advisors to impress clients with, with relatively rudimentary analytical capabilities.

Only problem is, in order to obtain the future efficient frontier portfolios, you have to be able to predict future asset returns, SDs, and correlations. And, hey, if you can do that, you don't need an optimizer -- you need the ability to handle the tsuris which comes with being the world's wealthiest human.

What about simply using historical data? That's probably the dumbest thing you can do, for the simple reason that most national stock markets have a fairly strong tendency to mean revert. All other things being equal, the optimizer tends to pick those assets with the highest assigned returns. If you use historical returns you will wind up with a portfolio of the previously best performing assets, which are liable to be the future worst performing assets. Not cool. As a simple example, if you had optimized your portfolio at the end of 1989 using historical data, the outputted portfolios would have been very heavy weighted towards Japanese equity. Needless to say, had you actually done this you would be by now so far at the back of the class that your best chance of survival would be a search party.

So, what use is the thing? Well, first and foremost an MVO is a superb teaching tool. Play around with one for a few hours and you will begin to acquire a grasp of the rather counterintuitive way in which real portfolios behave. The process can be likened to one's first few hours of flight instruction -- "Golly, if I bank the airplane enough I gain airspeed and lose altitude. If I increase power the houses get smaller, and we go *slower,"* becomes "Geez, willya lookit that -- add in a few percent of Iranian equity and my portfolio becomes *less* risky."

Can you use an MVO to help you shape your portfolio? Yes, but you've got to be very careful. An MVO is like a chainsaw. Used appropriately, it is a powerful tool for clearing your backyard. Used inappropriately it will send your local surgeon's kids to college. Same thing with MVOs. Want to wind up in the financial version of intensive care? Just throw in some historical (or even plausible) returns and believe what comes out the other end.

What's the right way? In order to answer that question, you have to realize that the chances of your allocation, no matter how skillfully chosen, winding up exactly on the *future* efficient frontier are zero. In fact, your chances of winding up even within 1% of return of the EF are about the same as your chances of winning the Miss America contest. In order to do so, after all, you have to be able to accurately predict almost all of the MVO inputs. For a 10 asset portfolio, that's 65 parameters. Rotsa ruck.

Rather, the most rational way to use an MVO is to find a "reasonable" allocation (hereafter known as the "coward's portfolio," or "CP") which does fairly well under a wide range of scenarios. In other words, pick an allocation, and then figure out as many ways as you can of blowing that allocation up with adverse inputs.

In order to do this, I've taken the following 12 assets, and used some "baseline" returns/SD assumptions:

Asset	Return	SD
S&P 500	.07	.15
US Small Stocks	.09	.25
Latin American Stocks	.07	.40
Pacific Large Stocks	.07	.25
Pacific Small Stocks	.09	.30
European Large Stocks	.07	.20
European Small Stocks	.09	.25
REIT Stocks	.06	.20
Precious Metals Stocks	.03	.40
Natural Resources Stocks	.05	.20
Non US Bonds	.04	.09
Long Term US Bonds	.04	.09
One Year Corporate Bonds	.03	.015

(All returns/SDs are inflation adjusted, and decimailzed. E.g., the return of .07 for the S&P 500 denotes an inflation adjusted return of 7%. The European small and large asset classes are 2/3 continental European and 1/3 British, the Pacific small and large asset classes are equal parts Japan and EAFE-PACXJ.)

The correlation grid used was for quarterly returns from July 1988 to September

1997. Three parameters were varied:

1) Inflation could be "normal," "high," or "low." For high/low inflation the returns of precious metals equity are raised/lowered by 0.1, and for natural resources by 0.05. One year bond and REIT are not changed under any of these scenarios. All other stock and bond assets are lowered/raised by 0.04 for high/low inflation.

2) Foreign dominance. Under the high/low scenarios the returns of all foreign stock and bond assets are raised/lowered by 0.04.

3) Small/large dominance. Under high/low scenarios, the returns of all foreign and domestic small cap stock assets are raised/lowered by 0.04.

Each scenario is denoted by 3 letters. For example, scenario "lhn" corresponds to low inflation, high foreign returns, and a normal small cap premium. In this scenario the S&P would return 0.11 (0.07 + 0.04 for low inflation), and European and Pacific large caps 0.15 (0.07 + 0.04 for low inflation + 0.04 for foreign dominance).

There are thus 27 different scenarios, which would seem to cover most, but not all, economic and financial environments. For example, the current scenario is lll, the late 70s to early 80s hhh, the mid 80s nhn, the mid 60s lnh. The period of the great depression is not describable in this schema, as there was negative inflation and very low stock returns.

These 27 scenarios were tested against the following "reasonable" allocation:

Asset	Allocation
S&P 500	15%
US Small Stocks	15%
Latin American Stocks	5%
Pacific Large Stocks	5%
Pacific Small Stocks	6%
European Large Stocks	6%
European Small Stocks	6%
REITs	5%
Precious Metals Stocks	5%
Natural Resources Stocks	2%
Non US Bonds	10%
Long Term US Bonds	0%
One Year Corporate Bonds	20%

The return of this allocation was plotted against the unconstrained efficient frontier for all 27 scenarios, along with the S&P 500. The MVO used is "MvoPlus," a proprietary optimizer which closely approximates the returns of a rebalanced portfolio.

The plot below represents the nnn case -- that is, no adjustments to the "normal" scenario. This is a screenshot from MvoPlus, modified slightly to show the return/SD of the SP and S&P 500. .

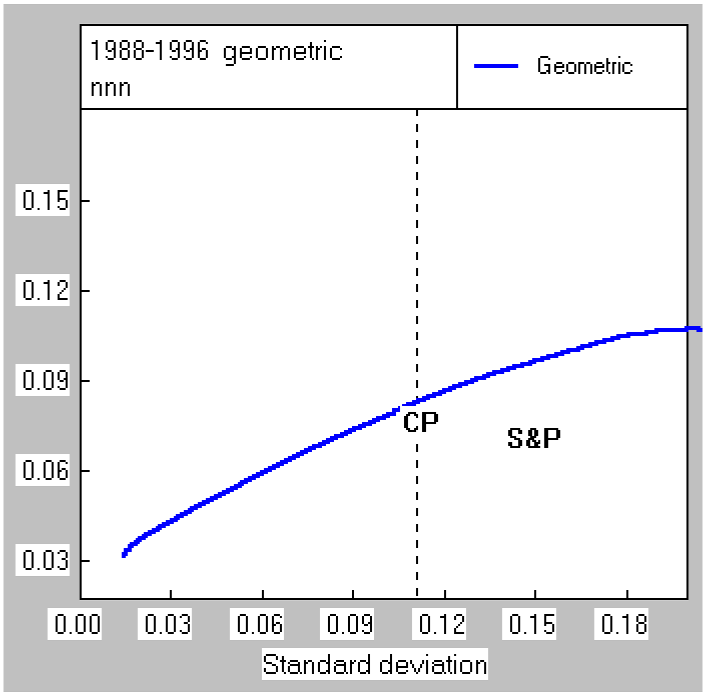

I've plotted the returns for the above "coward's portfolio" ("CP") as well as the S&P. Note that the CP is pretty close to the EF curve, and that it is considerably more efficient than the S&P 500. The "efficiency" of a portfolio is defined as the distance below the EF curve. The closer to the curve, the more efficient it is.

I've also plotted the same screenshots for the other 26 scenarios below. You can simultaneously view all 27 scenarios by clicking here, or view them individually by clicking on the individual scenarios below.

Scenario hhh
Scenario hhl
Scenario hhn
Scenario hlh
Scenario hll
Scenario hln
Scenario hnh
Scenario hnl
Scenario hnn

Scenario lhh
Scenario lhl
Scenario lhn
Scenario llh
Scenario lll
Scenario lln
Scenario lnh

Scenario lnl
Scenario lnn

Scenario nhh
Scenario nhl
Scenario nhn
Scenario nlh
Scenario nll
Scenario nln
Scenario nnh
Scenario nnl
Scenario nnn

It's a Low, Low, Low World

It turns out that of the 27 possible scenarios, the CP was more efficient than the S&P 500 20 times, less efficient 5 times (lhl, lll, lln, lhl, and nll), and about the same 2 times (llh and lnn).

The "worst case" for the CP is the lll portfolio -- with low inflation/high overall stock returns, but low foreign and small cap dominance. Buy the S&P, and you wind up almost smack dab on the sacred curve. That, unfortunately, is precisely the universe we've been living in for the past decade, and has encouraged the skeptics to proclaim the "Death of Diversification."

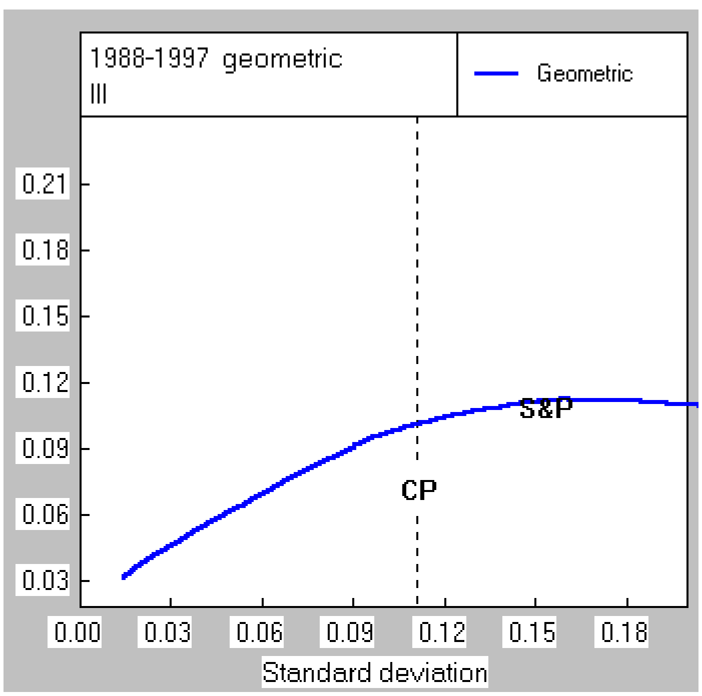

Maybe I'm just whistling past the graveyard, but it seems to me that it's pretty unlikely that the next decade will look anything like the last one. In fact, if the returns of global equity assets mean revert then over long enough time horizons the return scenario should look most like the nnn world, in which the CP beats the pants off of a domestic only strategy.

Anyway, the purpose of this piece is not to predict which of the above scenarios will best predict future returns, but rather how to properly use an MVO. VisualMvo, a freestanding Windows based MVO, will be available soon. If you are interested contact David Wilkinson via email. He may also be reached via snailmail or telephone at:

David Wilkinson
311 Ned's Mountain Road
Ridgefield CT 06877
1 203 778 1632

The MVO used in this piece, MvoPlus, with its mutliperiod rebalanced return formulation, will also be available from David sometime in early Spring. Happy Portfolio Hunting!

Home E-Mail

Efficient Frontier

William J. Bernstein

A Better International Benchmark

Unless you've just spent the last 10 years repairing vacuum tubes on Mir, you know that victory has been declared by the efficient marketeers at Vanguard. As of year's end, the fifteen year return of the Vanguard Index Trust 500 ranked 26th of 243 diversified mutual funds listed by Morningstar. Clearly, you're better off letting Lady Luck (the dart throwing chimpanzee in Paul Harvey's spoof of the "Beardstown Ladies" in this month' EF) pick your stocks than the average fund manager.

This is not entirely unexpected. After all, large cap US stocks exist in an informational fishbowl -- if the Efficient Market Hypothesis (EMH) holds anywhere, it should be here.

In fact, things are not so pretty elsewhere for the dart throwers. For example, 84 of 166 small cap funds managed to beat the Russell 2000 over the past 5 years. Similarly, over the same time period 66 of 128 funds beat the EAFE.

Seventy years ago Ben Graham and his colleagues electrified the financial world with the hostile takeover of Northern Pipeline. Graham had determined that NP held an amount of high quality bonds which greatly exceeded its market value. Today, you or I can determine a company's liquid assets with a few phone calls or a database search. If you subscribe to a stock database service such as Morningstar's *Stocktools* you can find this data with a few keystrokes. However, in Graham's time such data was not readily divulged to the investing public, and Graham's accomplishment lay more in prying the data loose than in the subsequent takeover. Paradoxically, today's transparent financial system makes such a coup impossible today -- our efficient markets instantaneously impound this sort of information into the price of a security. Graham allowed as much in the early 70s, just before his death. At that time he opined that the entire endeavor of security analysis was no longer worth the candle.

It would seem reasonable to expect that foreign markets, which are not nearly as informationally efficient as the US market, would present better opportunities for the astute investor. The Northern Pipelines may have long since disappeared from the US market, but perhaps they can still be found abroad. There is a problem, however, in measuring the performance of foreign money managers. Just what benchmark does one use? For years I've been unhappy with the ranking system used for foreign funds. The EAFE is heavily weighted towards Japanese equity (30%), and just about any foreign fund will look good next to it. Similarly, the past several years have also not been good to foreign small stocks; a well managed foreign small stock fund will likely underperform a less well managed large cap foreign fund. A mediocre fund with a primarily European large cap focus will find itself by default at the top of any general ranking system; an excellent fund with a focus on small companies in non-European markets will likely find itself at the bottom.

Morningstar Products does categorize some foreign funds as European, Pacific Rim, Latin American, or Emerging Markets, but among diversified foreign funds there is a large variation in regional allocation and market capitalization. How to adjust for this? Fortunately, our friends from Chicago provide us with the tools. It goes something like this:

Morningstar provides median market capitalization and regional allocation (Europe, Japan, Pacific Rim, and Latin America) for each fund. For all but Latin America both large and small cap benchmarks are available. (VIEIF-European/DFA Continental-UK, EAFE-PXJ/DFA Pacific Rim, EAFE-Japan/DFA Japan. For the European small cap component, a 2/1 mix of DFA's Continental/UK funds is used. Median market caps of $12B are assumed for the EAFE-Japan and EAFE-PXJ.)

Returns for each component are interpolated for size using the logarithm of the fund's median market cap. For example, assume the 5 year return for the European large cap component is 10%, and for the European small cap component is 0%. If the median market cap of the fund in question is the same as the EAFE-E, then the European bogie for the fund will be 10%. If the logarithm of the median market cap is half way between the small and large cap benchmarks, then the European bogie for the fund is 5%.

From the 4 bogies for each fund a benchmark return can now be computed from the arithmetically weighted regional bogies, which much more accurately reflects the fund makeup. For example, assume that the fund composition/bogie data looks something like this:

Region	Fund Composition	Fund Cap Weighted Regional Bogie
Europe	50%	15%
Japan	15%	-5%
Pacific Rim	10%	-10%
Latin America	25%	10%

Then, the "benchmark" return for the fund will be:

(.5 x .15) + (.15 x -.05) + (.1 x -.1) + (.25 x .1) =

$$.075 - .0075 - .01 + .025 = 0.825 = 8.25\%$$

It ain't perfect. For starters, Morningstar doesn't provide hedging info, so this critical parameter has to left out of the equation. Second, the regional allocation is measured only once, at the end of the period. Last, and most important, the benchmark by definition measures only security selection skill, and not asset allocation skill.

Still, it is a better mousetrap. So, Goldie, the envelope please:

Fund Name	Relative Performance (%)
59 Wall St European Equity	-4.05
59 Wall St Pac Basin Equity	1.19
Acorn International	3.85
Aetna International Grth Sel	0.20
AIM International Equity A	1.24
Alliance International A	-1.52
Alliance International B	-3.03
Alliance New Europe A	0.72
Alliance New Europe B	-0.11
American AAdvant Intl EqInst	4.02
American Cent-20thCIntlGrInv	0.58
Babson-Stewart Ivory Intl	-0.18
Bailard, Biehl Intl Equity	-5.39
Bartlett Value Intl A	-1.63
Bernstein Intl Value	4.33
BT Investment Intl Equity	4.29
Calvert World Value Intl EqA	-1.95
Capstone New Zealand	-1.14
Capstone Nikko Japan	3.97
Cat: Diversified Emerging Mkts	-3.56
Cat: Diversified Pacific/Asia	0.67
Cat: Europe Stock	-0.29
Cat: Foreign Stock	-1.23
Cat: Japan Stock	1.64
Cat: Latin America Stock	-1.44
Cat: Pacific/Asia ex-Japan Stk	-4.45
Colonial Newport Tiger T	-1.18
Columbia International Stock	-0.07
Compass Intl Equity Instl	-3.14
Compass Intl Equity Inv A	-3.75
Consulting Group Intl Equity	-1.98
CoreFund Intl Growth A	-1.35
CoreFund Intl Growth Y	-1.20
Dean Witter European Grwth B	3.45

Dean Witter Pacific Growth B	3.01
Delaware Intl Equity A	-0.54
Delaware Intl Equity Instl	-0.28
Delaware Pooled Intl Equity	0.47
Dreyfus Premier Intl GrowthA	-3.83
Enterprise Intl Growth A	-1.80
EuroPacific Growth	2.27
Excelsior International	-2.12
Excelsior Latin America	-0.85
Excelsior Pacific/Asia	-1.49
Excelsior Pan-European	-2.16
Fidelity Adv Overseas T	1.30
Fidelity Diversified Intl	5.57
Fidelity Emerging Markets	-8.32
Fidelity Europe	1.03
Fidelity Intl Growth & Inc	1.39
Fidelity Japan	3.92
Fidelity Overseas	0.90
Fidelity Pacific Basin	3.56
Flag Inv International A	0.92
GAM Europe A	-2.92
GAM International A	10.19
GAM Pacific Basin A	-0.34
Glenmede Instl International	2.25
Glenmede International	3.58
Goldman Sachs Intl Eqty A	1.78
Govett International Eqty A	-2.20
GT Global Emerging Mkts A	-6.45
GT Global Europe Growth A	-4.93
GT Global Intl Growth A	-6.29
GT Global New Pacific A	-7.45
GT Latin America Growth A	-4.12
Hancock Pacific Basin Eq A	-2.01
Harbor International	3.84
Hotchkis & Wiley Intl	3.07
IAI International	-1.74
IDS International A	-3.53
Invesco European	-2.29
Invesco International Growth	-4.64
Invesco Pacific Basin	-3.01
Ivy International A	2.79
Japan	2.03
JP Morgan Intl Equity	-2.15

Kemper International A	-1.85
Kent Intl Growth Instl	-1.23
Kent Intl Growth Invmt	-1.63
Keystone International	-2.97
Lazard Intl Equity Instl	-0.42
Lexington Worldwide EmergMkt	-5.71
Loomis Sayles Intl Eqty Inst	-0.81
MainStay Inst EAFE Instl	-5.12
Managers Intl Equity	0.69
Merrill Lynch Consults Intl	-2.24
Merrill Lynch Dev Cap Mkts A	-1.99
Merrill Lynch Dragon B	-5.06
Merrill Lynch Dragon D	-4.25
Merrill Lynch EuroFund A	1.70
Merrill Lynch EuroFund B	0.47
Merrill Lynch Latin Amer B	-2.83
Merrill Lynch Latin Amer D	-1.96
Merrill Lynch Pacific A	6.32
Merrill Lynch Pacific B	5.21
Montgomery Emerging Mkts R	-3.98
Morgan Stanley Inst Intl EqA	7.04
Munder International Eqty A	-6.06
Munder International Eqty K	-6.06
Munder International Eqty Y	-5.82
Nations Intl Equity Prim A	-5.45
Nations Intl Growth Inv A	-4.02
New England Intl Equity A	-5.38
Nomura Pacific Basin	-0.14
Oakmark International	3.49
One Group Intl Eqty Idx Fid	-1.97
Pacific European Growth A	-6.12
Parkstone Intl Disc Instl	-0.57
Phoenix International A	-2.91
PIMCo International A	-4.38
PIMCo International C	-5.17
Pioneer Europe A	3.91
Preferred International	1.06
Principal International A	1.80
Prudential Pacific Growth A	2.10
Prudential Pacific Growth B	1.58
Prudential World Intl Stk Z	0.85
Putnam Asia Pacific Growth A	6.28
Putnam Europe Growth A	1.44

Quantitative Intl Eqty Ord	-3.38
RBB-BEA Instl Intl Equity	-3.33
Rembrandt Intl Equity Comm	-3.13
RSI Retrmnt Intl Equity	-1.85
Schroder International Inv	-1.13
Scudder International	-1.59
Scudder Pacific Opport	-7.34
SEI International Equity A	-5.84
Seligman Henderson Intl A	-2.59
Sit International Growth	-0.84
Smith Barney Intl Equity A	-3.27
Smith Barney Intl Equity C	-4.23
T. Rowe Price European Stock	0.15
T. Rowe Price Intl Discovery	-1.48
T. Rowe Price Intl Stock	-0.24
T. Rowe Price Japan	-0.66
T. Rowe Price New Asia	-4.64
TCW/DW Latin American Grth B	-5.33
Templeton Developing Mkts I	1.61
Templeton Foreign I	0.20
Templeton Foreign Smaller Co	1.59
Templeton Instl Forgn Equity	1.41
Templeton Pacific Growth I	-3.13
UAM TS&W Intl Equity	-3.75
United International Grth A	1.17
USAA International	2.61
Vanguard International Value	-3.97
Vanguard Intl Eqty European	0.26
Vanguard Intl Eqty Pacific	-0.11
Vanguard Intl Growth	2.93
Victory International Grth A	-3.59
Vontobel Intl Equity	-0.71
Warburg Pincus Adv Intl Eqty	0.62
Warburg Pincus Inst Intl Eq	3.28
Warburg Pincus Intl Eq Comm	1.12
William Blair Intl Growth	3.22
WPG International	-3.99
Wright EquiFund-Hong Kong	-5.85
Wright EquiFund-Netherlands	0.97
Wright Intl Blue Chip Equity	-2.29

(Note: "Relative performance" is the 5 year annualized return in excess of the fund's custon benchmark, as calculated above.)

Some of the results are nonsensical, such as for many of the single country funds (e.g., Capstone New Zealand, and all of the Wright Equifunds). Still, I feel that this system provides a fairer appraisal of the highly diversified funds. For example, the Scudder and T. Rowe Price diversified International stock funds are usually at the top of many recommended lists, and for reasons which have never been clear to me, so are the INVESCO Pacific and European Funds. The above table shows all of these to be lackluster. On the other hand, although both Acorn International Fund has been at best in the middle of the pack in the past 5 years, but when their small market cap is taken into account it can be seen that they outperform. By any standards, the Morgan Stanley Institutional International Fund is a star. Unfortunately, it's closed.

Are any of the fund results statistically significant? Possibly, but to find out you would have to calculate the fund's return and benchmark for multiple periods, then determine whether the series of realtive returns was statistically different from zero. That's a heck of a lot of work, and requires a much bigger database than I have. Perhaps Morningstar could be interested in such a project. However, for the small investor the message is clear: Before you swallow a foreign fund's superior performance whole, season it with an analysis of its regional allocation and market cap.

The Intelligent Asset Allocator

William J. Bernstein

Diversification, Rebalancing, and the Geometric Mean Frontier
An Improved MVO Algorithm

A few comments about this article. David Wilkinson and I decided to post this piece on the mathematics of rebalancing and a "fix" for the return formulation of classical MVO. It's a pretty techincal piece, and gets quite dense at places. Not for everybody.

It's also in .pdf format, so it will require Adobe Acrobat 3.0, which may be downloaded here. Get Acrobat Reader

.

That said, here it is:

Diversification, Rebalancing, and the Geometric Mean Frontier
An Improved MVO Algorithm

Home E-Mail

[an error occurred while processing this directive]

DIVERSIFICATION, REBALANCING, AND THE GEOMETRIC MEAN FRONTIER

William J. Bernstein

wbern@mail.coos.or.us

and

David Wilkinson

DaveWilk@worldnet.att.net

November 24, 1997

Abstract

The effective (geometric mean) return of a periodically rebalanced portfolio always exceeds the weighted sum of the component geometric means. Some approximate formulae for estimating this effective return are derived and tested. One special case of these formulae is shown to be particularly simple, and is used to provide easily computed estimates of the benefits of diversification and rebalancing. The results are also used to show how classical Mean-Variance Optimization may be modified to generate the Geometric Mean Frontier, the analog of the efficient frontier when the geometric mean is used as the measure of portfolio return.

1 Introduction

The calculation of the true long term, or effective, return is an often ignored part of the portfolio optimization process. The case of U.S. common stocks and long term corporate bonds in Table 1 provides a well known example. Here the arithmetic mean return is simply the average of the 69 yearly returns, while the geometric mean return is the effective, or annualized, return over the entire period.

Let us consider an annually rebalanced portfolio consisting of equal parts of stocks and corporate bonds. Using a simple 50/50 average of the individual returns one obtains an anticipated portfolio return of 9.00 percent using the arithmetic mean returns and 7.85 percent using the geometric mean returns. In fact, neither is correct: A 50/50 portfolio, rebalanced annually, has an annualized return of 8.34 percent. To complicate matters further, if one had purchased the 50/50 mix on January 1, 1926 and not rebalanced, then by December 31, 1994 an almost 100 percent stock portfolio would have resulted, with an annualized return of 9.17 percent.

MacBeth (1995) recognized that the uncritical use of a weighted geometric mean results in an anticipated portfolio return which is too low, and suggested instead a formulation using the arithmetic mean corrected for variance. For the above example, this gives a return of 8.30 percent, very close to the actual rebalanced return of 8.34 percent.

The question of whether or not rebalancing benefits portfolio return is more complex. Perold and Sharpe (1995) examined the problem from the perspective of historical stock/bill returns and concluded:

> In general, a constant-mix (rebalanced) approach will underperform a comparable buy-and-hold (unrebalanced) strategy when there are no reversals. This will be the case in strong bull or bear markets, when reversals are small and relatively infrequent, because more of the marginal purchase and sell decisions will turn out to have been poorly timed.

We feel that this can only be part of the answer, and that another important criterion is whether or not the individual assets have similar long term return. Common experience demonstrates that rebalancing often yields significant excess returns when the return differences are small. Contrariwise, rebalancing penalizes the investor when asset return differences are large.

In this paper, we investigate the questions of rebalancing and long term portfolio return in a quantitative manner. We begin by discussing Mean-Variance Optimization (Markowitz (1952, 1991)) in the context of multi-period portfolio optimization. We then derive two families of approximate Portfolio Return Formulae, the first of which uses the individual arithmetic means as input, and the second the geometric means. As a special case of the latter, we obtain a simple approximate formula for the Diversification Bonus, the amount by which the geometric mean return of a rebalanced portfolio exceeds the weighted sum of the individual geometric means. This result is used to obtain a corresponding formula for the Rebalancing Bonus, the amount by which the return of the rebalanced portfolio exceeds that of the corresponding unrebalanced one. Lastly, we show how Mean-Variance Optimization may be modified to obtain the Geometric Mean Frontier (GMF), the analog of the efficient frontier when the geometric mean is used as the measure of portfolio return.

Maximization of the geometric mean return has been discussed extensively in the literature: Latane (1959), Hakansson (1971), Elton and Gruber (1974a, 1974b), Fernholz and Shay (1982). However most of this work was concerned with exact results, the question of whether maximizing the geometric mean can be justified on the basis of utility theory, or with the case of continuous time rebalancing. We take a more pragmatic and general approach: (a) we will be satisfied with approximate formulae for the portfolio geometric mean, (b) we consider the geometric mean in the context of actual historical data with a finite, but arbitrary, rebalancing interval, and (c) we focus on the entire Geometric Mean Frontier, as opposed to just the single portfolio which maximizes the geometric mean.

2 MVO – Past and Future

Mean-Variance Optimization (MVO) is designed to produce return/variance efficient portfolios. The standard Markowitz (1952, 1991) analysis is applicable to a single period, and the inputs are the expected returns R_i and covariance matrix V_{ij} for the individual assets over this period. The latter are related to the standard deviations s_i and correlation matrix ρ_{ij} by

$$V_{ij} = s_i s_j \rho_{ij} \quad , \tag{1}$$

where $\rho_{ij} = 1$ for $i = j$.

For a portfolio with fraction X_i assigned to asset i, with $\sum_i X_i = 1$, the expected return R and its variance V are given by

$$R = \sum_i X_i R_i \quad , \tag{2}$$

and

$$V = \sum_{ij} X_i X_j V_{ij} \quad . \tag{3}$$

The risk of the portfolio is taken to be the standard deviation $s = \sqrt{V}$.

One way to supply the inputs R_i and V_{ij} is to use historical returns, and to assume that the upcoming period will resemble one of the previous N periods, each with equal probability $1/N$. In this case, the expected return R_i becomes the arithmetic mean return of asset i over the N periods

$$R_i = \frac{1}{N} \sum_{k=1}^{N} r_i^{(k)} \quad , \tag{4}$$

where $r_i^{(k)}$ is the return of asset i in period k. For the portfolio, the return R in Eq. (2) becomes the arithmetic mean of the returns of a portfolio which is *rebalanced* to the mix specified by the X_i at the beginning of each period

$$R = \frac{1}{N} \sum_{k=1}^{N} r^{(k)} \quad , \tag{5}$$

where $r^{(k)} = \sum_i X_i r_i^{(k)}$ is the return of the rebalanced portfolio in period k. Note that while the period represented by the returns, e.g. annual or quarterly, is arbitrary, the rebalancing interval must always be the same as the measurement interval.

When viewed in this multi-period context, the Markowitz analysis is unsatisfactory because the long term return of an asset with returns $r^{(k)}$ in the different periods is given not by the arithmetic mean, Eq. (5), but rather by the geometric mean

$$G = \left[\prod_{k=1}^{N} (1 + r^{(k)}) \right]^{1/N} - 1 \quad . \tag{6}$$

Since the arithmetic mean of any return series is always greater than the geometric mean, the return predicted by the Markowitz analysis is always greater than the true long term return that would have been obtained by using the actual rebalanced allocation. For this reason the geometric mean returns of the individual assets G_i given as

$$G_i = \left[\prod_{k=1}^{N} (1 + r_i^{(k)}) \right]^{1/N} - 1 \tag{7}$$

are often used as inputs to the Markowitz analysis in preference to the arithmetic means R_i. However, as pointed out by MacBeth (1995), this is not correct either. As we shall see explicitly in the next section, this prescription always underestimates the true return of the rebalanced portfolio.

The discussion of the remainder of this paper will mostly be in terms of historical data, and the questions we will address are whether, given only partial information, it is possible to (a) estimate the true long term geometric mean return of a given rebalanced portfolio, (b) compare this return with that of the corresponding unrebalanced portfolio, and (c) estimate the composition of the rebalanced portfolio which maximizes the geometric mean return for a given level of risk. The partial information we consider is the MVO data alone, which for the case of historical data we will generalize to mean *either* the arithmetic mean return R_i *or* geometric mean return G_i for each asset, together with the covariance matrix V_{ij}. As in the classical situation, we identify the risk with the standard deviation, which for the case of historical data is a consequence of the fluctuation of the individual single period returns about their (arithmetic) mean value. The general conclusion is that, to a good approximation, all three goals above may be accomplished, and the reason we use the historical perspective is that it enables us to validate these conclusions. However, the purpose of portfolio theory is, of course, to provide a basis for decision-making for the future. In the case of the standard single period analysis, the input data are necessarily statistical in nature, and the covariance matrix represents the uncertainty in the returns for the single upcoming period. However, for multi-period forcasting, two entirely different viewpoints are possible.

In the first viewpoint, the investor seeks the best course of action over a chosen number of upcoming periods, based on the assumption that the hypothesized MVO data are actually realized. In this case the problem is conceptually no different from that with partial historical data. This viewpoint has no meaning in the case of the usual single period analysis, because if the return of each asset is assumed known, then the covariance matrix is zero, and the optimum strategy is just to select the asset with the highest return. However, in the multi-period case, the problem is not so simple. Suppose, as will usually be the case, that the input returns are chosen to be the geometric mean returns G_i. Firstly, the input covariance matrix represents not the uncertainty in these values, but rather the fluctuations of, and correlations between, the unspecified individual period returns which go to make up these values. Thus there is no contradiction between having both specified values for the returns and a non-zero correlation matrix. Secondly, the computational problem itself is non-trivial. Generally, as we shall see explicitly in Section 6, the rebalanced strategy with the highest geometric mean return is not to invest 100 percent in the asset with the highest geometric mean return; often the highest return strategy is to invest in a diversified portfolio containing several of the assets.

The second viewpoint is statistical in nature, and superficially more similar to the conventional single period one. It is, however, necessary to make some additional hypothesis about the time correlation of the returns. The most natural approach, which is in accord with the simplest form of the random walk theory of asset price movement, is to asssume that the distribution of returns is stationary, i.e. the same in each period, with each period being independent of the others. In this case the input return and covariance matrix represent the expected value and uncertainty of this unique single period distribution. The only new feature is that in the multi-period application we allow the possibilty of specifying the geometric means G_i of this distribution instead of the arithmetic means R_i. In this viewpoint, the significance of the geometric mean G (both for the individual assets, and for any rebalanced portfolio of them) is that, as the number of periods becomes large, it becomes increasingly probable that the actual long term return will lie close to G. This property is quite general, and is analogous to the observation that if a fair coin is thrown a large number of times, then it becomes increasingly likely that the fraction of heads will lie close to one half.

Which of these two viewpoints to adopt is perhaps a matter of personal taste. The actual mathematical development is the same in either case. We tend to prefer the first viewpoint, because it does not have to make any independent assumption about the time correlation of events. For example, if the proposed rebalancing period is annual, then the covariance matrix should correspond to annual returns; if it is quarterly, then the covariance matrix should correspond to quarterly returns. There is no need for any relationship to exist between these two covariance matrices. In the second viewpoint, however, consistency requires that the two covariance matrices be related by a factor of four, a relationship which is not necessarily satisfied in reality. For the purposes of this paper, it will certainly be simpler to think in terms of the first viewpoint, because there the future is treated in exactly the same way as the past.

3 Portfolio Return Formulae

In this section we derive and discuss a variety of approximate formulae for the geometric mean return of a balanced portfolio. These formulae express the portfolio geometric mean in terms of the arithmetic or geometric mean of the individual assets, together with their covariance matrix. In order to validate these formulae we will use an actual 9-asset global data set of annual index returns over the years 1970-1996. The arithmetic mean return, geometric mean return, and standard deviation of each asset are listed in Table 2; the correlation matrix was also computed but is not shown. Note that the asset with the highest arithmetic mean return (gold) is not the one which has the highest geometric mean return. In fact, both Japan and U.S. small stocks have higher long term return than gold over this period. Note also that, since these are annual returns, the rebalancing period we consider in the examples is also annual.

We first consider the weighted arithmetic mean and weighted geometric mean as possible portfolio return formulae. The results for the arithmetic mean are shown graphically in the top left plot of Fig. 1. As expected, we see that the weighted arithmetic mean overestimates the actual portfolio return in all cases. The corresponding results for the weighted geometric mean are shown in in the top right plot of Fig. 1, where we see that, in contrast, the geometric mean of the portfolio is always underestimated. Neither the weighted arithmetic mean or weighted geometric mean are adequate models for the portfolio geometric mean.

Improved geometric mean formulae may be based on the well known approximate relationship between the geometric mean G and arithmetic mean R of any set of returns

$$G \approx R - \frac{V}{2(1+R)} \quad , \tag{8}$$

where V is the variance of the returns (see, for example, Markowitz (1991)). This formula strictly holds in the limit where the variance is small, but it is surprisingly accurate even outside this range. If, in addition, the arithmetic mean return is small compared to unity then this reduces to

$$G \approx R - \frac{V}{2} \quad . \tag{9}$$

It will be useful in the following to regard both these formulae as special cases of

$$G \approx R - \frac{\alpha V}{2(1+\beta R)} \quad , \tag{10}$$

which we will refer to as the (α, β) formula. In this paper we will only consider the cases $(1, 0)$ and $(1, 1)$, but the general (α, β) case is no more difficult to analyze than $(1, 1)$, so most of the results will apply to this general case.

A useful feature of the (α, β) formula is that it may be inverted to give

$$R = \frac{2G + \alpha V}{1 - \beta G + \sqrt{(1+\beta G)^2 + 2\alpha\beta V}} \quad . \tag{11}$$

When $\beta = 0$ this simplifies to

$$R = G + \frac{\alpha V}{2} \quad . \tag{12}$$

We first explore the strategy of simply applying the (α, β) formula to the portfolio. We will call this method A(α, β), where the 'A' signifies that the Arithmetic means R_i of the individual assets are used to construct the arithmetic mean $R = \sum_i X_i R_i$ of the portfolio. The portfolio variance is obtained from

$$V = \sum_{ij} X_i X_j V_{ij} \quad . \tag{13}$$

The results for methods A(1,0) and A(1,1) are shown in the left column of Fig. 1. It is seen that in both cases there is an improvement over the simple weighted arithmetic or geometric means. The A(1,1) formula in particular does extremely well for the two-asset and randomly selected portfolios, and only fails for the high return single asset portfolios. Even here, however, the results for both A(1,0) and A(1,1) are much better than those obtained using simple weighted arithmetic or geometric means. Note that the weighted arithmetic mean may be considered as the A(0,0) case of A(α, β), so that the three plots on the left side of Fig. 1 may be viewed as three cases of the same formula, with the accuracy impoving from top to bottom.

Two drawbacks of the A(α, β) method are (a) the formula uses the arithmetic means of the individual assets rather than their geometric means, and (b) the formula does not give the correct result in the simplest case where the portfolio consists of only a single asset. Both these defects may be remedied by applying the (α, β) formula not only to the portfolio but also to the individual assets. To do this we first use the inverse relationship on the individual assets to obtain "pseudo-arithmetic means"

$$R_i^* = \frac{2G_i + \alpha V_{ii}}{1 - \beta G_i + \sqrt{(1 + \beta G_i)^2 + 2\alpha\beta V_{ii}}} \quad , \tag{14}$$

and then compute the portfolio arithmetic mean using $R = \sum_i X_i R_i^*$ before using the (α, β) formula for the portfolio. We call this method G(α, β), where the 'G' signifies that the geometric means of the individual assets are used as inputs. The results for G(1,0) and G(1,1) are shown in the right column of

Fig. 1. It is seen that while the results for the two-asset and random portfolios are not as good as for the corresponding A(1,0) and A(1,1) cases, the results for the single asset portfolios are now exact. The weighted geometric mean may be considered as the G(0,0) case of G(α, β), so that the three plots on the right side of Fig. 1 may be viewed as three cases of the same formula, with the accuracy improving from top to bottom.

In this section we have derived two families A(α, β) and G(α, β) of portfolio return formulae, the first of which requires the individual arithmetic means as inputs and the second their geometric means. While we have focused on the two cases (1,1) and (1,0), which have some theoretical foundation, it is possible that, empirically, values other than these might prove to be more accurate.

4 The Diversification Bonus

The general G(α, β) method succeeds in expressing the geometric mean return of the portfolio in terms of the geometric mean returns of the individual assets. Though simple to evaluate numerically, the expession is rather cumbersome and not very intuitive. However, when $\beta = 0$ the expression simplifies because the inverse relation in Eq. (11) reduces to the simpler Eq. (12). In the following, we will also set $\alpha = 1$, though the case of general α is no more difficult. In this case, G$(1, 0)$, we find

$$G \approx \sum_i X_i G_i + \frac{1}{2} \left(\sum_i X_i V_{ii} - \sum_{i,j} X_i X_j V_{ij} \right) \quad . \tag{15}$$

The analogous result in the limit of continuous time rebalancing has been obtained previously by Fernholz and Shay (1982), who showed that it is exact when the prices follow geometric Brownian motion. A different form of the same result may be obtained by using $\sum_i X_i = 1$ in the first term inside the parentheses and symmetrizing to give

$$G \approx \sum_i X_i G_i + \sum_{i<j} X_i X_j \left(\frac{V_{ii}}{2} + \frac{V_{jj}}{2} - V_{ij} \right) \quad , \tag{16}$$

where the sum is over all pairs i and j with $i < j$.

This form of the G$(1,0)$ formula has been obtained previously using an empirical argument[1]. The interesting feature of this formula, not shared by the general G(α,β) with $\beta \neq 0$, is that it separates into two terms, the first of which depends only on the geometric mean returns, and the second only on the covariance matrix. Eq. (16) demonstrates that, for given values of the individual returns G_i, it is possible for the portfolio return G to be an increasing function of the volatility (standard deviation) of the individual assets. In particular, if all the off-diagonal elements of the correlation matrix ρ_{ij} are zero or negative, then the portfoloio return is an increasing function of each of the standard deviations s_i. The benefits of a rebalancing strategy become much greater for assets which are volatile and poorly correlated.

Let us write Eq. (16) in the form

$$G - \sum_i X_i G_i \approx \sum_{i<j} X_i X_j \left(\frac{V_{ii}}{2} + \frac{V_{jj}}{2} - V_{ij} \right) \quad . \tag{17}$$

We will refer to the left side of Eq. (17), which is always positive, as the Diversification Bonus[2], because it represents the effective return in excess of that calculated from the simple combination of the individual geometric mean returns. The right hand side of Eq. (17) provides an approximation to this quantity which is expressed solely in terms of the covariance matrix of the assets. On the left side of Fig. 2 we compare the approximate formula, Eq. (17), with the true value of $G - \sum_i X_i G_i$. It is seen that some of the portfolios have a Diversification Bonus in excess of 200 basis points. Although Eq. (17) tends to overestimate the precise value of $G - \sum_i X_i G_i$, as we would expect from the plot for formula G(1,0) in Fig. 1, it does a good job of predicting which portfolios will have significant Diversification Bonus. On the right side of Fig. 2 we show the corresponding plot for the case G(1,1); while this formula does not lead to a simple form like Eq. (17), it gives a better approximation to the true Diversification Bonus, as again we would expect from Fig. 1.

5 The Rebalancing Bonus

The weighted geometric mean $\sum_i X_i G_i$ provides a naive estimate of effective portfolio return, but it does not correspond to the return obtainable by any particular investment strategy. A more meaningful comparison is that between the return G of the rebalanced portfolio and the return G' of the corresponding unrebalanced portfolio. The latter is given by

$$G' = \left[\sum_i X_i (1+G_i)^N\right]^{1/N} - 1 \quad . \tag{18}$$

Subtracting Eq. (18) from Eq. (16) we obtain

$$\begin{aligned} G - G' \quad &\approx \quad \left[\sum_i X_i(1+G_i)\right] - \left[\sum_i X_i (1+G_i)^N\right]^{1/N} \\ &\qquad + \sum_{i<j} X_i X_j \left(\frac{V_{ii}}{2} + \frac{V_{jj}}{2} - V_{ij}\right) \quad . \end{aligned} \tag{19}$$

We will refer to the left side of Eq. (19) as the Rebalancing Bonus [3], because it represents the difference between the returns of the rebalanced and unrebalanvced portfolios. In the approximation on the right hand side, the term on the first line always gives a negative contribution, thus favoring the unrebalanced portfolio. This term is large when the differences between the long term returns G_i are large, and vanishes when all the returns G_i are the same. The term in the second line always gives a positive contribution, thus favoring the rebalanced portfolio. This term is large when the individual assets have large variances, and low or negative correlation.

For the data set considered in Section 3, it is found that the majority of portfolios perform better when rebalanced. This is illustrated in the left side of Fig. 3, in which we compare Eq. (19) with the true value of $G - G'$. It is seen that the great majority of the portfolios have positive values for this quantity. Examination of the data underlying Fig. 3 shows that the only portfolios giving a significant negative value are 50/50 2-asset portfolios in which one of the assets is Treasury Bills. This is in accordance with the above discussion – only in this case is the return difference sufficient to overcome the second term in Eq. (19). We also see that the approximate formula in Eq. (19)

always gives the correct sign for $G-G'$, although it tends to overestimate the precise benefit of rebalancing, as we would expect from the plot for formula G(1,0) in Fig. 1. On the right side of Fig. 3 we show the corresponding plot for the case G(1,1); while this formula does not lead to the simplifying separation of Eq. (19), it gives a better model value for $G - G'$, as again we would expect from Fig. 1.

The first line in Eq. (19) vanishes when all the returns G_i are equal, but there is no simple way to express Eq. (19) in terms of return differences alone. Over a long enough time period, however, the unrebalanced portfolio becomes dominated by the highest return asset, independent of the initial allocation, and we obtain

$$G' \approx G_{\max} \quad , \tag{20}$$

where $G_{\max} = \max_i G_i$ is the maximum return obtainable from any unrebalanced portfolio. In this case the Rebalancing Bonus becomes

$$G - G' \approx \sum_i X_i (G_i - G_{\max}) + \sum_{i<j} X_i X_j \left(\frac{V_{ii}}{2} + \frac{V_{jj}}{2} - V_{ij} \right) \quad . \tag{21}$$

In this case it is clear that if all the G_i are the same, then the rebalanced portfolio is always superior to the unrebalanced one.

6 The Geometric Mean Frontier

When used with historical data, the traditional Markowitz efficient frontier designates those portfolios with greater (arithmetic mean) return than any other with the same or lesser risk, and lesser risk than any other with the same or greater return. The Markowitz return R and risk s are defined by

$$R = \sum_i X_i R_i \quad , \tag{22}$$

and

$$s^2 = \sum_{ij} X_i X_j V_{ij} \quad . \tag{23}$$

In this section we show that it is possible to modify the MVO analysis to construct the analogous efficient frontier when the geometric mean is used instead of Eq. (22) as the measure of return, while keeping Eq. (23) as the measure of risk. More precisely, we show that this Geometric Mean Frontier (GMF) may be computed exactly when any of the approximate geometric mean formulae A(α, β) or G(α, β) of Section 3 is used, and that in each case the efficient frontier so obtained provides a good approximation to the efficient frontier of the true geometic mean, the latter being obtained using the optimization feature of a standard spreadsheet package[4].

In each of the methods A(α, β) or G(α, β), the portfolio geometric mean has the form

$$G = R - \frac{\alpha s^2}{2(1+\beta R)} \quad , \tag{24}$$

where R and s are given in terms of the individual assets by expressions of the form of Eqs. (22) and (23). The only difference between the cases is that for A(α, β) the individual returns R_i are the true arithmetic mean returns, while for G(α, β) the R_i are pseudo-arithmetic means R_i^* given in Eq. (14). To simplify the notation we will omit the superscript in this latter case. Also, in this section we will use the symbol 'G' to denote the appropriate model geometric mean, so that Eq. (24) should be thought of as a definition, rather than as an approximate relation for the true geometric mean.

The key to the analysis is to consider an auxiliary classical Markowitz optimization which uses the *arithmetic* means as inputs. We emphasize that even in the case of G(α, β), which expresses the portfolio geometric mean in terms of the individual geometric means, the inputs to this auxiliary Markowitz problem are the individual pseudo-arithmetic means, not the geometric means.

We begin by generalizing an argument first given by Elton and Gruber (1974b) to show that any portfolio which is (G, s) efficient must also be Markowitz (R, s) efficient. For any (α, β), the geometric mean expression in Eq. (24) has the properties

$$\frac{\partial G}{\partial R} > 0 \quad , \tag{25}$$

$$\frac{\partial G}{\partial s} < 0 \quad . \tag{26}$$

Suppose that (R_1, s_1), with corresponding geometric mean G_1, is not Markowitz efficient. Then there exists a portfolio (R_2, s_2), with corresponding geometric mean G_2, with either $R_2 > R_1$ and $s_2 \leq s_1$, or $R_2 \geq R_1$ and $s_2 < s_1$, or both. In either case it follows from Eqs. (25) and (26) that $G_2 > G_1$ and $s_2 \leq s_1$. Thus no portfolio which is not Markowitz efficient can be geometric mean efficient; equivalently, every geometric mean efficient portfolio must be Markowitz efficient. Note that the above argument does not assume the existence of a continuous path of portfolios between portfolios 1 and 2.

The Markowitz efficient frontier extends between the minimum variance portfolio at $s = s_{\rm min}$ and the maximum return portfolio at $s = s_{\rm max}$. For each s in the range $s_{\rm min} \leq s \leq s_{\rm max}$ the frontier gives a corresponding return $R_f(s)$, and this function $R_f(s)$ has the properties

$$\frac{\mathrm{d}R_f}{\mathrm{d}s} > 0 \quad , \tag{27}$$

$$\frac{\mathrm{d}^2 R_f}{\mathrm{d}s^2} < 0 \quad . \tag{28}$$

Let us define $G_f(s)$ by evaluating Eq. (24) on the Markowitz frontier, i.e.

$$G_f(s) = R_f(s) - \frac{\alpha s^2}{2(1+\beta R_f(s))} \quad . \tag{29}$$

From the above discussion, the (G, s) efficient frontier must be a sub-graph of the graph of $G_f(s)$. Thus the multi-dimensional space of candidate (G, s) efficient portfolios has been reduced to a one-dimensional one.

It can be shown using Eqs. (27) and (28) that, under very reasonable assumptions on the nature of the function $R_f(s)$, the function $G_f(s)$ can have only one maximum in the range $s_{\rm min} \leq s \leq s_{\rm max}$. In general there are three possibilities: (a) the maximum occurs at $s_{\rm min}$; in this case the geometric mean frontier consists of a single portfolio, the minimum variance portfolio, (b) the maximum occurs at $s_{\rm max}$; in this case the entire Markowitz frontier is geometric mean efficient, or (c) the maximum occurs at some intermediate point; in this case only that part of the Markowitz frontier to the left of this maximum is geometric mean efficient. This third possibility is the one which occurs for our chosen example.

The portfolio with the maximum geometric mean may be algorithmically determined by starting from the minimum variance portfolio at $s = s_{\min}$ and moving to the right along the Markowitz frontier until the corresponding $G = G_f(s)$ first decreases. The standard MVO output consists of a set of "corner portfolios" located on the Markowitz frontier, between each adjacent pair of which the frontier is obtained by linear combination. The task of testing for a maximum of $G_f(s)$ between any pair of corner portfolios is easily accomplished by using the composition, rather than the risk, as the independent variable. Since $G_f(s)$ has only one maximum, a maximum will be found in at most one corner interval, and that interval will always be one which is adjacent to the corner portfolio with the highest geometric mean return.

We examine the validity of these ideas using the same 9-asset data set considered in Section 3. In Fig. 4 we plot the Markowitz efficient frontier, and associated model and actual geometric means, for the four cases A(1,0), G(1,0), A(1,1) and G(1,1). The result of the exact spreadsheet optimization of the true geometric mean is also shown. The four cases are qualitatively very similar. In each, the maximum of the geometric mean (model or actual) occurs at an intermediate value of risk; the Markowitz frontier greatly overestimates the true reward of increasing risk. Note also that, while the model returns in the four cases show some differences, as we would expect from the discussion of Section 3, the true geometric mean evaluated on the frontier is for each case in excellent agreement with the result of the full optimization of the true geometric mean.

In Table 3 we present in tabular form three portfolios on the efficient frontier for each of the the four methods A(1,0), G(1,0), A(1,1) and G(1,1), together with the exact numerical result. Table 3a shows the portfolio which maximizes the model geometric mean. We see that the composition of the optimum portfolio is in all four model cases in reasonable agreement with the exact one, and that despite the fact that the corresponding model geometric means are somewhat different from each other, the true geometric means are close not only to each other, but also to the correct maximum of 16.85 percent. An even better result may be obtained by choosing the optimal portfolio to be that which maximizes (over the Markowitz frontier) the true geometric mean return, rather than the model geometric mean. The corresponding results are shown in Table 3b. In this case it is seen that the portfolio compositions are in closer agreement with the exact ones, and that, to the quoted precision, all four approximate methods obtain the correct maximum geometric mean of 16.85 percent. It is also possible to use the efficient frontier of the auxiliary Markowitz problem to determine the geometric mean efficient portfolio corresponding to a given value of risk s, or model geometric mean G, or true geometric mean. As an example, we show in Table 3c the results obtained by fixing s to be 10.0 percent. Again, we see that the portfolio compositions are in good agreement with the exact ones, and that the values for the true geometric mean reproduce the exact value of 13.54 percent to within at most one basis point.

The extremely close agreement when the true geometric mean return is used follows from a general result of perturbation theory, namely that when the objective function of an optimization is perturbed by a small amount, the change in the value of the unperturbed objective function in going to the new maximum is second order in the perturbation. The excellent agreement is thus due to the fact that all the $\mathrm{A}(\alpha, \beta)$ and $\mathrm{G}(\alpha, \beta)$ formulae are reasonable approximations to the true geometric mean.

The net result of this section is that, although the $\mathrm{A}(\alpha, \beta)$ and $\mathrm{G}(\alpha, \beta)$ formulae are not perfect models for the true geometric mean, in each case the efficient frontier of the auxiliary Markowitz problem does contain portfolios whose true geometric mean return is extremely close to the real maximum, either in an absolute sense, or when the risk is specified to take a given value. Thus, finding the optimal geometric mean portfolio becomes a tractable numerical task, even for a large number of assets, because the space of candidate portfolios becomes one-dimensional.

7 Discussion

This paper contains four significant results, each of which can be easily used by investors who are able to express their data in the form commonly used by a standard Mean-Variance Optimization package – namely the arithmetic mean return R_i or geometric mean return G_i of each asset, and the covariance matrix V_{ij} between the assets. These results are:

1. Two families $\mathrm{A}(\alpha,\beta)$ and $\mathrm{G}(\alpha,\beta)$ of approximate Portfolio Return Formulae for estimating the geometric mean return of a rebalanced portfolio. The family $\mathrm{A}(\alpha,\beta)$ requires the arithmetic means of the individual assets, and $\mathrm{G}(\alpha,\beta)$ the geometric means.

2. For the case G(1,0), a particularly simple approximate formula, Eq. (17), for the Diversification Bonus, the amount by which the portfolio geometric mean exceeds the weighted sum of the individual geometric means.

3. A corresponding approximate formula, Eq. (19), for the Rebalancing Bonus, the difference between the returns of the rebalanced and unrebalanced portfolios.

4. An extension of classical Mean-Variance Optimization which allows construction of a good approximation to the entire Geometric Mean Frontier (GMF), and in particular that portfolio which maximizes the geometric mean return.

The validity of these results was demonstrated using historical data. The chosen data set, containing as it does some extremely volatile assets, is a stringent test of the theory; much more accurate results would be obtained by applying the methodology to a less volatile group of assets, such as different classes of U.S. stocks and bonds. However we believe that the major application of these ideas is precisely in generating diversified portfolios of very volatile assets, because there lies the greatest potential benefit of diversification and rebalancing, and of optimizing specifically for long term return.

The formalism presented here is not limited to annual returns, but may be applied for any rebalancing period. For example, if the rebalancing period is quarterly, then the arithmetic and geometric mean returns should be quarterly returns, and the covariance matrix should be constructed from the historical quarterly returns. If desired, the resulting portfolio return G may be converted to annual return at the end of the calculation. The methodology may thus be used to investigate the effect of varying the rebalancing period.

While our focus has been on the analysis of historical data, the goal of portfolio theory is of course to generate optimal strategies for the future. In the historical case, the simplest and most accurate of the MVO methods presented here is to use the arithmetic means as inputs to the optimizer, and use the A(1,1) formula or, even better, the actual returns, to evaluate the portfolio geometric mean along the frontier. This will always generate an excellent approximation to the true Geometric Mean Frontier. There is no need to consider the $G(\alpha, \beta)$ formulae and pseudo-arithmetic means. However, when considering long term future strategies the situation is very different (for the interpretation of the formalism in this situation, recall the discussion at the end of Section 2). Here it is unlikely that investors will simply adopt the historical geometric mean, arithmetic mean and covariance matrix. At the least, they will provide their own estimate of return, and this return will be the effective long term, i.e. geometric mean, return. In this case the $G(\alpha, \beta)$ family of Portfolio Return Formulae, and the associated use of pseudo-arithmetic means as MVO inputs, become indispensible. The G(1,1) formula will give the best results when used in the MVO analysis, while the simpler G(1,0) formula can provide good intuitive understanding of the performance of different portfolio allocations. We emphasize once again that simply using the individual geometric means as MVO inputs always underestimates the true long term return, even when the input data for the individual assets prove to be correct.

Notes

1. Bernstein, William J., "The Rebalancing Bonus: Theory and Practice", Efficient Frontier (September 1996), http://www.coos.or.us/~wbern/ef.

2. Previously (see Note 1) the left side of Eq. (17) was termed the Rebalancing Bonus; here we use this expression for the left side of Eq. (19).

3. See Note 2.

4. Quattro Pro, Version 6.0.

References

[1] Elton, Edwin J., and Martin J. Gruber, 1974a. "On the Optimality of Some Multi-Period Selection Criteria", Journal of Business, Vol. 47, No. 2, pp. 231-243.

[2] Elton, Edwin J., and Martin J. Gruber, 1974b. "On the Maximization of the Geometric mean with Lognormal Return Distribution", Management Science, Vol. 21, No. 4, pp. 483-488.

[3] Fernholz, Robert, and Brian Shay, 1982. "Stochastic Portfolio Theory and Stock Market Equilibrium", Journal of Finance, Vol. 27, No. 2, pp. 615-624.

[4] Hakansson, Nils H., 1971. "Capital Growth and the Mean-Variance Approach to Portfolio Selection", Journal of Financial and Quantitative Analysis, Vol. VI, No. 1, pp. 517-558.

[5] Latane, Henry A., 1959. "Criteria for Choice Among Risky Ventures", Journal of Political Economy (April).

[6] MacBeth, James D., 1995. "What's the Long-Term Expected Return to Your Portfolio?", Financial Analysts Journal (September/October).

[7] Markowitz, Harry M., 1952. "Portfolio Selection", Journal of Finance, Vol. 3, No. 1.

[8] Markowitz, Harry M., 1991. *Portfolio Selection: Efficient Diversification of Investments*, 2nd edition, Basil Blackwell Ltd, Cambridge, Massachusetts.

[9] Perold, Andre F. and William F. Sharpe, 1995. "Dynamic Strategies for Asset Allocation", Financial Analysts Journal (January/February).

	Arithmetic mean	Geometric mean	Standard deviation
STOCKS	0.1216	0.1019	0.2020
BONDS	0.0583	0.0551	0.0850

Table 1: Arithmetic mean, geometric mean and standard deviation for U.S. Common Stocks and Long Term Corporate Bonds for 1926-94.

	Arithmetic mean	Geometric mean	Standard deviation
S&P 500	0.1346	0.1227	0.1585
SMALL US (9-10)	0.1659	0.1415	0.2293
EAFE-E	0.1486	0.1305	0.2095
EAFE-PXJ	0.1641	0.1226	0.3084
JAPAN	0.1895	0.1454	0.3368
GOLD	0.2014	0.1370	0.4299
20 Y TREAS	0.0989	0.0927	0.1189
5 Y TREAS	0.0950	0.0928	0.0686
T BILL	0.0692	0.0688	0.0267

Table 2: Annual arithmetic mean return, geometric mean return, and standard deviation for nine asset classes over the period 1970-1996. *S&P 500, U.S. 9-10, and Treasury security retrurns from SBBI, Ibbotson and Associates. MSCI-EAFE-E, MSCI-EAFE-PXJ and MSCI Japan from Morgan Stanley. Gold from Morningstar Inc, and Van Eck Group.*

	A(1,0)	G(1,0)	A(1,1)	G(1,1)	Exact
SMALL US (9-10)	0.3530	0.2190	0.3049	0.2204	0.2767
JAPAN	0.3385	0.3778	0.3633	0.3866	0.3533
GOLD	0.3085	0.4032	0.3318	0.3931	0.3699
Markowitz return	0.1848	0.2056	0.1863	0.1946	0.1874
Standard deviation	0.1969	0.2242	0.2047	0.2228	0.2120
Model geometric mean	0.1655	0.1805	0.1686	0.1738	0.1685
True geometric mean	0.1681	0.1683	0.1684	0.1684	0.1685

(a) Maximum model geometric mean

	A(1,0)	G(1,0)	A(1,1)	G(1,1)	Exact
SMALL US (9-10)	0.2649	0.2769	0.2649	0.2732	0.2767
JAPAN	0.3839	0.3526	0.3839	0.3623	0.3533
GOLD	0.3512	0.3705	0.3512	0.3646	0.3699
Markowitz return	0.1874	0.2027	0.1874	0.1925	0.1874
Standard deviation	0.2120	0.2120	0.2120	0.2119	0.2120
Model geometric mean	0.1650	0.1803	0.1685	0.1736	0.1685
True geometric mean	0.1685	0.1685	0.1685	0.1685	0.1685

(b) Maximum true geometric mean

	A(1,0)	G(1,0)	A(1,1)	G(1,1)	Exact
SMALL US (9-10)	0.2407	0.2218	0.2407	0.2270	0.2396
EAFE-E	0.0172	0.0096	0.0172	0.0186	0.0166
JAPAN	0.1402	0.1445	0.1402	0.1414	0.1400
GOLD	0.1290	0.1403	0.1290	0.1359	0.1303
5 Y TREAS	0.4729	0.4839	0.4729	0.4772	0.4736
Markowitz return	0.1399	0.1461	0.1399	0.1416	0.1399
Standard deviation	0.1000	0.1000	0.1000	0.1000	0.1000
Model geometric mean	0.1349	0.1411	0.1355	0.1372	0.1354
True geometric mean	0.1354	0.1353	0.1354	0.1354	0.1354

(c) Portfolio with standard deviation $s = 0.10$

Table 3: Three portfolios on the Markowitz frontier for each of the four cases A(1,0), G(1,0), A(1,1) and G(1,1). The final column shows the corresponding exact result.

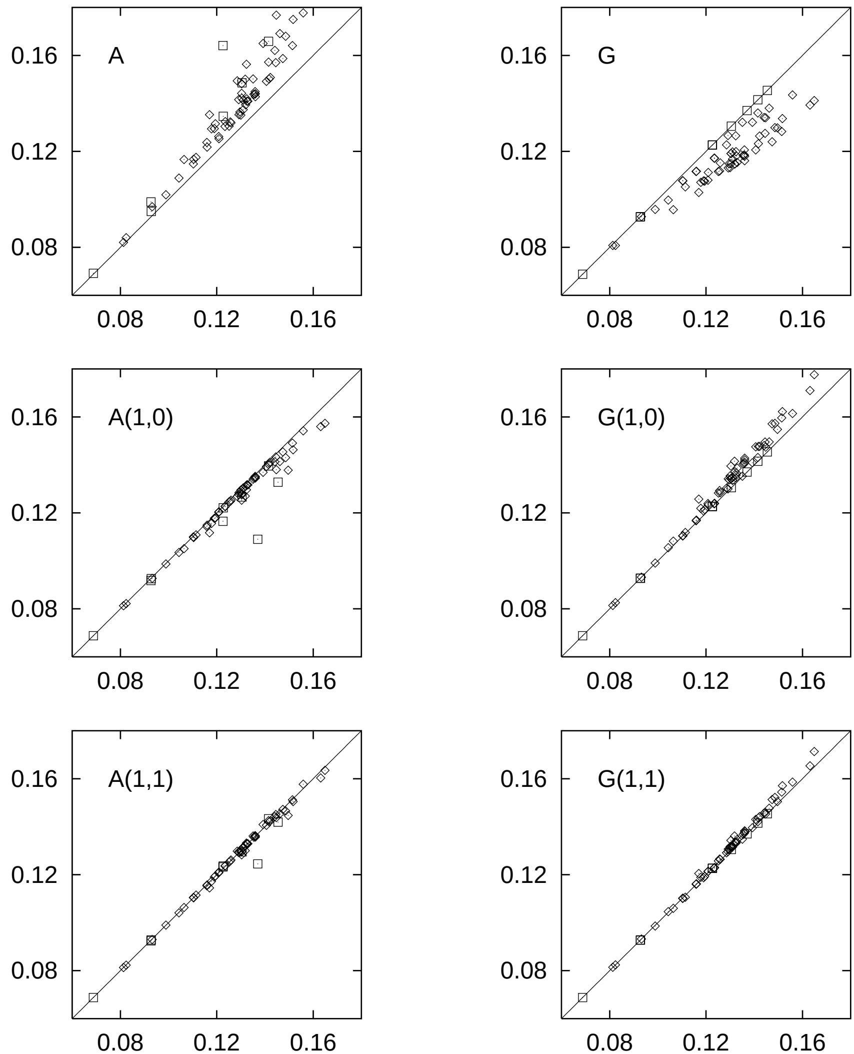

Figure 1: Cross-plot of model (vertical axis) against actual (horizontal axis) geometric mean return, for weighted arithmetic mean (A), weighted geometric mean (G), and formulae A(1,0), G(1,0), A(1,1) and G(1,1). Each plot shows the 9 single-asset portfolios (squares), the 36 two-asset portfolios with 50-50 mix, and 25 randomly generated portfolios.

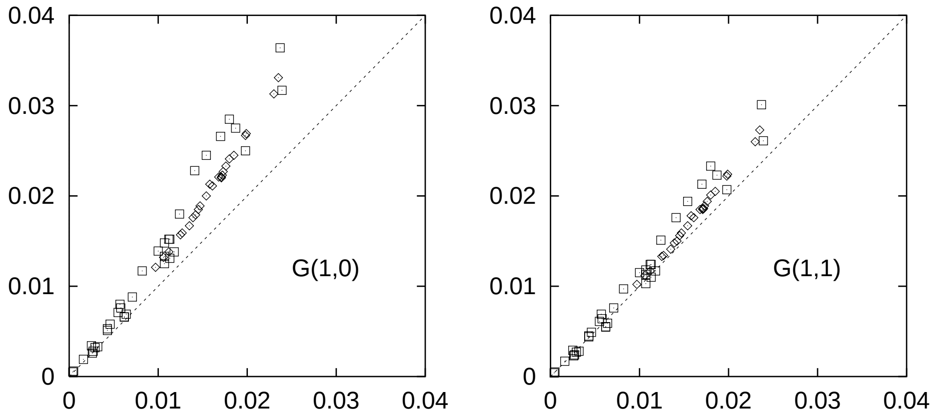

Figure 2: Cross-plot of model (vertical axis) against actual (horizontal axis) Diversification Bonus, $G - \sum_i X_i G_i$, using formulae G(1,0) and G(1,1) for the model return. Each plot shows the 36 two-asset portfolios with 50-50 mix (squares), and 25 randomly generated portfolios (diamonds).

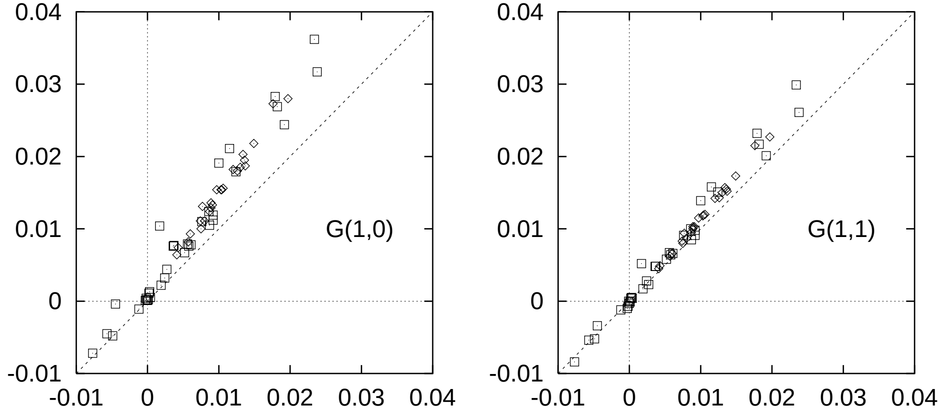

Figure 3: Cross-plot of model (vertical axis) against actual (horizontal axis) Rebalancing Bonus, $G - G'$, using formulae G(1,0) and G(1,1) for the model rebalanced return. Each plot shows the 36 two-asset portfolios with 50-50 mix (squares), and 25 randomly generated portfolios (diamonds).

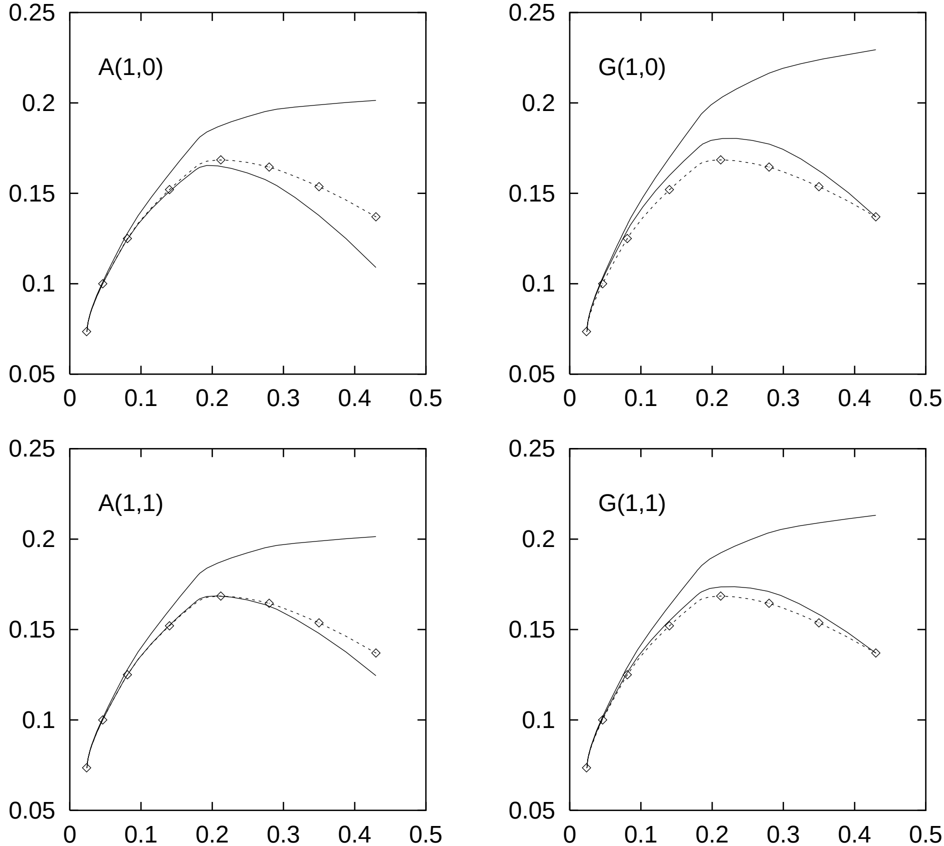

Figure 4: Risk-return plot for each of the four cases A(1,0), G(1,0), A(1,1) and G(1,1). In each plot the upper solid curve is the Markowitz efficient frontier, the lower solid curve the corresponding value of the model geometric mean, and the dotted curve the corresponding value of the true geometric mean. The open diamonds denote the exact optimized return for each value of risk.

Optimizer Update

In September's ***EF*** I promised a working, practical, easy to use freestanding mean variance optimizer for under $100, from my colleague David Wilkinson. Perfectionist that he is, David will sell no optimizer before its time, and is currently putting the software through rigorous beta testing at the hands of a few dozen volunteers.

Due to the large number of excellent suggestions made by the testers, David has decided to delay sending out a finished product for another few weeks. David will also have a higher level product ready at some point in the not too distant future, which will be able to optimize for long term (geometric mean) return of rebalanced portfolios, as described in our paper in this month's EF. If you are interested in either of these MVO products, please contact David via email at DaveWilk@worldnet.att.net. He may also be reached via snailmail or telephone at:

David Wilkinson
311 Ned's Mountain Road
Ridgefield CT 06877
1 203 778 1632

Efficient Frontier

William J. Bernstein

What is the Very Long Term Return of Equity?

Say you had a very, very thoughtful and considerate ancestor living in Octavian's Rome, who invested the purchasing equivalent of one modern dollar for you at a constant real (inflation adjusted) rate of return. After 2000 years, you'd have the following amounts at the given rates:

Real Rate of Return	Principal in 1997	Radius of Gold Sphere (at $280/oz.)
1.0%	$440,000,000	8.1 meters
3.0%	\$4.7 x 10^{25}	3,800 kilometers
5.0%	\$2.4 x 10^{42}	1,400,000,000 kilometers
7.0%	\$5.9 X 10^{58}	4.1 x 10^{14} kilometers

(Thanks to my daughters for lending me their scientific calculator; my financial calculator proved inadequate to the task.)

Even at 1%, you are now wealthy beyond most of your rational desires. At 3% your accumulated wealth buys a sphere of gold larger than the moon. At 5% the gold sphere has a radius reaching to Saturn's orbit, and at 7% it is 43 light years in radius, a tidy sum indeed. (To give due credit, the above paradigm isn't mine, it's Peter Bernstein's.)

What is the point of this exercise? Well, perhaps we need to question the current universal assumption of high expected equity returns. Over the past 200 years English speaking investors have been very lucky indeed, earning real equity returns of about 6%. If were you a stock investor elsewhere in the world (say, Germany, Japan, or South America) you weren't so lucky, with long term returns not greatly exceeding inflation.

Unless you dropped out of school long before you reached puberty (or perhaps spent too much time watching MTV and Fox) the speedbumps on your road to the above theoretical wealth are pretty obvious. Little things like the sacking of Rome, dark ages, plague bacillus, and assorted European and global Armageddons.

Have we gotten any smarter in the interim? For starters, publicly traded equity is a relatively modern invention. The first recorded publicly traded stock was probably that of a water mill in southern France, at Bazacle. Built about 800 AD, its ownership was divided up into shares around 1100, with more or less continuous pricing from 1400 on, until the French government nationalized the mill a few decades ago. (But then again, what do you expect from a culture which finds subtext in Jerry Lewis movies?) So maybe it is different this time. One can only hope.

The long term economic growth of the world economy can be looked at from another perspective. Let's be wildly optimistic and assume that the productivity of the average citizen had increased by a factor of 1,000 since Roman times. This translates into an annualized per capital growth rate of 0.35% over the past 2 millennia.

What's the bottom line here? Take another look at the above table. Even at a real rate of return of 3%, the result of a 2000 year investment of one dollar is billions of times greater than the world's current wealth. In other words, over the long term a real rate of return in excess of 1% is mathematically impossible. If you're planning for 5%-7% real returns over the next 30 years to fund your retirement, you may or may not wind up being very unpleasantly surprised.

You say you're planning a 10% return to fund your retirement? Sure, go ahead. Just don't set up any trust funds for your distant descendants.