



Stop Worrying about Your Portfolio

May 8, 2019

by Ben Inker
of GMO

EXECUTIVE SUMMARY

Investors have a tendency to focus on the characteristics of their portfolios almost to the exclusion of other factors that will lead to success or failure for the larger objective that the portfolio is intended to serve. By taking into account the characteristics of the assets and liabilities that exist outside of their investment portfolios, they could build portfolios that are a better match for the true problem they should be solving. Because the liabilities and assets outside of the portfolios do not generally have quantitatively well-estimated characteristics the way that traditional investment assets do, this type of analysis necessarily involves a certain amount of judgment rather than simpler historical return analysis. But this effort seems well worth the attempt, because most apparently rigorous attempts to build “optimal” investment portfolios are solving the wrong problem for most .

Investors have a tendency to obsess about their investment portfolios. On the surface, this is a perfectly reasonable focus given results in the portfolio are a crucial determinant of success for whatever purpose the portfolio is there to serve. But performance of one’s investment portfolio is not the only determinant of success, and it is almost certain that investors would achieve better overall outcomes if they recognized the risks outside of their portfolios that really matter and invested accordingly. Such a shift in mindset, while theoretically almost unarguable, does run into a couple of practical hurdles. The first is that the investment portfolio is the piece of the problem that is easiest to measure, and investment professionals are generally judged on measured outcomes. While a particular portfolio might be the right one for the specific assets and liabilities of a given investor, if that portfolio underperforms those of the investor’s peers, it is not clear whether an investment committee will care much about the theoretical superiority. The second issue is that once we move beyond the investment portfolio, estimates of risk and correlation are necessarily judgment calls and investors or risk managers cannot simply rely on a historical returns-based covariance matrix to estimate overall risk. That said, just because solving a particular problem is hard doesn’t mean the right thing to do is come up with an easier problem to solve that is less relevant. I’m going to make the argument as to why traditional ways of thinking about portfolios are flawed and can lead investors to make bad decisions.

The risks that matter

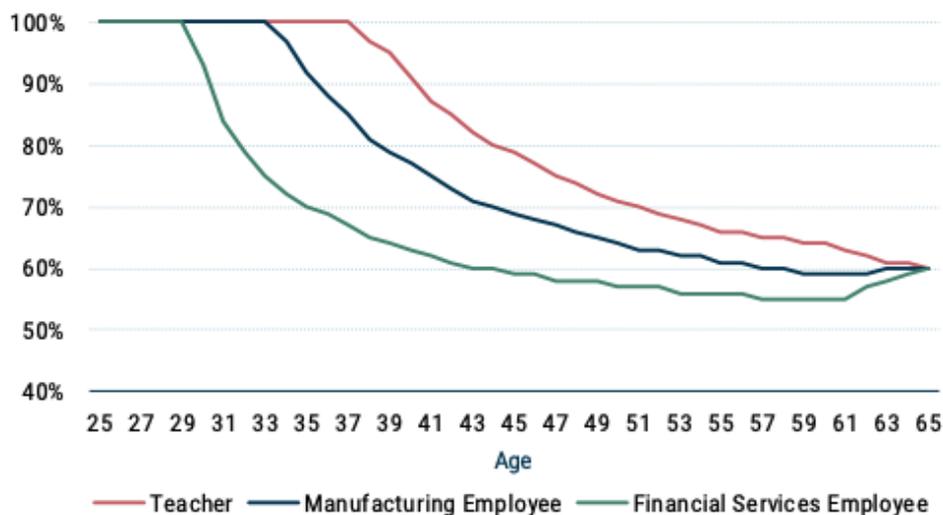
No investment portfolio exists in a vacuum. Every portfolio exists within a larger framework that includes both the purpose the portfolio serves – the liability – as well as any additional assets that will come into the portfolio over time. Risk is not merely a function of the volatility of the investment portfolio but also of the relationships between the investment portfolio, the liability, and the non-portfolio assets. Defined benefit pension fund managers are used to thinking about at least part of this conception of risk because accounting standards force them to recognize the way the pension liability varies with interest rates. In worrying about the funding ratio of the pension rather than simply the volatility of the portfolio, they are moving in the right direction. But most pension fund managers tend to stop there, failing to fully take into account the assets outside of the portfolio that are relevant to the overall problem – the potential of the fund sponsor to make additional contributions to the pension portfolio when needed. Because the sponsor’s ability to make these contributions varies depending on economic circumstances, all deteriorations in the funding ratio of the pension fund are not equal. A deterioration that occurs when the cash flow position of the sponsor is worse than average is significantly worse than one that occurs when sponsor cash flow is strong. Defining how much worse is a bit tricky, as it involves estimating the covariance between the sponsor’s cash flow and the investment portfolio as well as coming up with some estimate of the amount of disutility that stems from having to top up the pension fund at a time of stress. But the fact that it is tricky to estimate the parameters of the true problem doesn’t mean that ignoring it is the right answer. And the difference between looking at the true problem and an oversimplified version can be a big deal, which I hope to demonstrate with a few examples. I’m going to be using mean variance optimization (MVO) for my examples, because it’s a widely used tool and most professional investors are familiar with it. But the basic concepts in no way depend on MVO, or indeed quantitative portfolio construction methods in general, to be relevant.

Non-portfolio assets matter for retirement savings

One place where incorporating the existence of assets outside of the portfolio has important implications is in what I will call the retirement problem. For our purposes here, I am going to assume that workers need to accumulate 10 times their final salary at retirement, that they save 10% of their income every year that they are working, and that as a base case,

their salary will grow at 1% real per year throughout their working life. Their investment portfolios consist of stocks and bonds. I am going to assume that in a bad economic event there is a negative shock to both the returns to stocks and the income of the workers. It is this last piece that requires looking beyond the retirement portfolio itself in estimating true risk, and my goal is to show the implication of ignoring this covariance in building a retirement portfolio and why it is suboptimal. Exhibit 1 shows three different glide paths for three people saving for retirement – a teacher, a manufacturing employee, and a financial services employee.¹ The differences between them are driven solely by the level of covariance² assumed between employment income and the returns to stocks and bonds.

EXHIBIT 1: STOCKS AS PERCENT OF RETIREMENT PORTFOLIO



Source: GMO

All of these investors have the same level of risk aversion, which explains why at retirement (here assumed to be 65 years old) they have the same portfolio allocation to stocks.³ But depending on the assumption we make on the correlation between stock market returns and each saver’s labor income and the volatility of that labor income, the ideal portfolio moves around a good deal. If we assume a strong covariance between labor income and stock returns, such as might be the case for a worker in the financial services industry, we get the interesting result that workers in their 40s or older should have less in stocks than they should at retirement. While this runs counter to the standard advice and most traditional retirement glide paths, it does make plenty of sense when you think about how a depression impacts someone who is relatively close to retirement versus someone who has already retired.

For someone who has retired, there is nothing particularly special about a depression. It would certainly cause losses in his investment portfolio, but anything that caused similar losses would be similarly painful. For someone who was five years from retirement, however, a depression would be far worse than other negative portfolio events. A depression would have somewhere between a meaningful negative impact and a profound negative impact on that person’s expected future labor income. Under the best circumstances, where the person stays employed continuously over the full five-year period, it is almost certainly the case that his income will be at least mildly lower than it would have been in the absence of the depression. But the likelihood of that person winding up unemployed for some or perhaps all of the next five years has to be considered material. For the purpose of building the glide paths above, I made assumptions of savings relative to baseline over a five-year period in the event of a depression; those assumptions can be seen in Table 1.

TABLE 1: ESTIMATED IMPACT OF DEPRESSION EVENT ON EXPECTED 5-YEAR SAVINGS

	<i>Savings as Percent of Baseline</i>	PROBABILITY		
		<i>Teacher</i>	<i>Manufacturing Employee</i>	<i>Financial Services Employee</i>
Unemployed for Full Period	-200%	0%	7.5%	12.5%
Unemployed Half the Time	0%	0%	15.0%	17.5%
Meaningful Cut to Wages	75%	0%	27.5%	30.0%
Slight Cut to Wages	95%	0%	50.0%	40.0%
No Impact on Wages	100%	100%	0.0%	0.0%
Expected Savings as Percent of Baseline		100%	53.0%	36.0%

Source: GMO

The teacher case is the simple one. I assumed that a depression would have no impact on expected savings over the following five years. This is effectively what a standard optimization does.⁴ It is almost certainly an implausible assumption. Even in industries that are considered uncorrelated to the economy, such as education and health care, many more jobs are lost in bad economic times than in good ones. But it does show the extreme case of unvarying labor income, so it seemed a worthwhile example anyway. In the manufacturing employee case, the scenarios are more complex. I assumed a 50% probability that income takes a small hit of 5%, which reduces savings by a similar amount. I assumed a 27.5% probability of a more meaningful cut to income, which would cause the worker to only be able to save 75% of baseline expectations. I assumed a 15% probability that the worker would wind up out of work for about half of the five-year period, and net retirement savings would be zero over the period. Finally, I put a 7.5% probability on the worker losing his job for the whole period and effectively being forced to retire early, which would mean that instead of saving 10% of that baseline salary each year, he would be forced to withdraw 20% of baseline salary from his retirement account to make ends meet. For the financial services sector worker, I increased the probabilities of the most negative scenarios. The particular probabilities and scenarios I used are admittedly arbitrary, but they aren't obviously silly. And they hopefully bring home the point that the relationship between your retirement portfolio and your labor income can have a meaningful impact on the portfolio you should hold.

The other question is how important a depression event is to the covariance matrix. This seems like a hard problem, but I was able to take a short cut to make it much easier. It is not directly the covariance matrix that matters, but the combination of the covariance matrix and risk aversion. I built a covariance matrix based on the three events we generally use for long-only portfolios – depression, unanticipated inflation, and a liquidity shock. Those are the events that will cause a big enough loss in your portfolio to really matter, so to our minds, those are the events that your risk model should focus on. I then came up with a risk aversion level that resulted in a reasonable answer to the simplest problem – how much equity should you hold at retirement. Everything else flowed naturally from setting up the right problem – recognizing the existence of human capital properly and coming up with a set of covariances in the depression event.⁵ I didn't have to tell the optimization about changing risk aversion given a shrinking time horizon, or make assumptions about the impact of mean reversion on long-term equity risk.⁶ Simply by setting up to solve the actual problem rather than an oversimplified one that missed important aspects of risk, I was able to get reasonable solutions without having to build in the added complexities generally used to generate retirement glide paths.

Non-portfolio assets matter for sovereign wealth funds

Other types of investors may be trying to solve different problems, but they still can benefit from solving the whole problem instead of just a piece of it. Let's take the example of two different sovereign wealth funds. In the first case, let's assume that the portfolio is funded by commodity revenues and is intended to cushion the country's economy for a time in the future when the commodities run out. In the second case, let's assume that the portfolio is funded by a current account surplus stemming from manufactured goods trade surplus and is intended to cushion the economy for a time in the future when an aging population will reduce savings rates and output. In both cases the liability is similar to the retirement problem – replacement of future income funded by savings for a period of time up until then. But the nature of the non-portfolio asset is meaningfully different in the two cases and should lead to different portfolios.

For the commodity-driven wealth fund, the outside asset has three notable features that mean incorporating it into the portfolio construction problem will have a real impact. First, like human capital in the case of the retirement problem, the commodity-linked inflows should be expected to have significant downside sensitivity to a depression event. The revenues of commodity producers have a strong cyclical component given that both prices and volumes are impacted by the economy. Second, unlike human capital, there is a reasonable possibility that in an unexpected inflation event, cash flow

will improve in real terms. This is not a certainty, but rising real commodity prices are one of the plausible drivers of an unexpected inflation in the first place, and in that circumstance, cash flow into the portfolio would improve. Third, the existence of the commodity-linked asset creates another downside scenario that would matter for this investor even though it would not be an important event for most other investors. A scenario in which real commodity prices fall without a corresponding economic downturn matters for the country, even though it would not be expected to have a material impact on the pure investment portfolio.

For the portfolio funded out of manufactured goods current account surpluses, the risks to the cash flows are quite different. A depression would probably improve the current account situation as imports tend to fall by more than exports in economic downturns.⁷ As long as the true purpose of the portfolio is about replacing a future income loss driven by demographic shifts and not generally cushioning the economy in the event of cyclical downturns, this means this cash flow might actually be better in a depression event. On the other hand, a commodity price boom would hurt the current account position of a resource-light, manufacturing-heavy economy, making this event a negative for this investor above and beyond the impact of generally rising inflation, whereas an idiosyncratic fall in commodity prices would be a benefit to the investor's current account balance.

If these two sovereign wealth funds had otherwise similar risk aversions, they should wind up with materially different portfolios due to the different nature of their nonportfolio assets, as we can see in Table 2. I have added a third fund as well, for another sovereign wealth fund whose non-portfolio asset is a fixed real cash flow regardless of economic circumstance.⁸

TABLE 2: OPTIMAL PORTFOLIOS FOR SOVEREIGN WEALTH FUNDS⁹

	<i>Commodity Country</i>	<i>Aging Manufacturing Country</i>	<i>Neutral Country</i>
Non-commodity Equities	56%	51%	59%
Commodity Equities	0%	24%	10%
Bonds	44%	26%	30%
Commodity Weight in Equity Portfolio	0%	32%	15%
Equity Weight in Portfolio	56%	74%	70%

Source: GMO

The most striking difference among the three portfolios is the weight in commodity equities. The commodity-backed sovereign wealth fund should have no weight in commodity equities, the neutral country winds up with a market capitalization weight in them, and the aging manufacturing country should have more than a double weight in such equities given how well commodity equities fit with that country's overall risks. But beyond commodity equities themselves, we can also see that the overall equity level is meaningfully different across the three funds. Given the economic sensitivity of cash flows into the commodity-backed fund, the overall equity weight should be a relatively low 56%. For the neutral country, 70% turns out to be the optimal weight. And for the aging manufacturer, an even more equity-heavy 74% is called for, as a depression has a bit of a silver lining for the export-driven country.

These findings are not particularly surprising, particularly to the managers of sovereign wealth funds. The Norwegian sovereign wealth fund famously announced that it would divest from oil and gas exploration companies earlier this year. While the above analysis suggests an oil-backed sovereign wealth fund might want to go further than Norway's recent move and remove all companies tied to oil production, my point is not to state with high conviction exactly what the right answer is. It is rather that by setting up the problem the right way in the first place, we are able to see that these three different sovereign wealth funds should have meaningfully different portfolios even though their risk aversion is identical in all three cases.

Liabilities matter too!

The idea that the nature of your liabilities should impact the portfolio you run will certainly not come as any surprise to the manager of a defined benefit pension fund. But other investors have a tendency to think less about the liability side of things in building their portfolios. I'm going to use an example of two different charitable foundations to look at a way in which the liability side of the equation arguably should impact portfolio decisions. Let's assume we have two identical private foundations that will not receive any future donations. We, therefore, do not have to concern ourselves with assets outside of the portfolio the way we did in our previous examples. But these two foundations exist for different purposes.

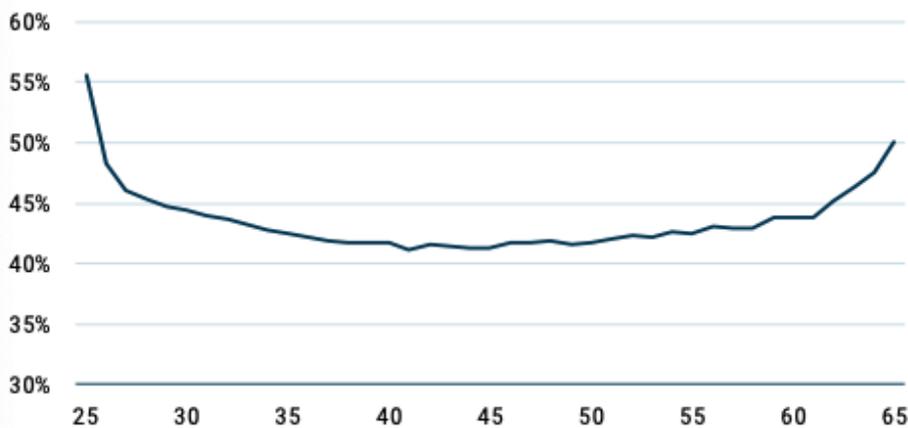
One is dedicated to funding research to cure cancer. The other is dedicated to feeding the poor in its local community. If it is safe to assume that the need for curing cancer and the cost to do so is more or less unaffected by economic circumstance but that the need to feed the hungry grows materially in a depression, then the nature of those liabilities strongly suggests the two portfolios should have different allocations to stocks and bonds. The same risk aversion that would give the cancer-curing foundation a 75% stock/25% bond portfolio would suggest the hunter oriented foundation should have 49% in stocks and 51% in bonds.¹⁰

Focusing on the right risks

Over the years, a number of economists have looked at the nature of outside assets and liabilities and their implications for investment decision making. Some of these have been quite intriguing. Brainard and Tobin wrote about how the correlation between labor incomes and the returns on domestic stocks argue for an anti-home bias in stocks for retirement portfolios.¹¹ But in what almost amounts to a bad joke about how many economists it takes to make a disagreement, Bottazzi, Pesenti, and van Wincoop a few years later wrote a paper making the exact opposite argument¹² – that workers should have a domestic bias in the equities in their retirement portfolios! While Brainard and Tobin were looking at the sensitivity of labor incomes and stock market returns to overall economic circumstances, Bottazzi et al. were looking at another potential source of disappointing labor income, a rise in the profit share of GDP. I have to give the authors full points for prescience, given that they wrote about the potential damage to labor incomes stemming from a rising profit share in 1996, a few years before the largest sustained increase in profit share and fall in labor share in U.S. history.

Both papers' effects seem economically intuitive, but the arguments can't both be correct, at least for a given investor. So how can we determine which effect should dominate for a worker saving for retirement? One simple way is by putting both scenarios into our covariance matrix and seeing which is more meaningful from a risk perspective. Exhibit 2 shows the ideal percentage of domestic stocks for a worker with a moderate covariance between stocks and labor income by age.

EXHIBIT 2: DOMESTIC STOCKS AS PERCENT OF TOTAL STOCKS



Source: GMO

There are a couple of interesting things that this chart tells us. First, for the retired worker who no longer has any human capital to worry about, the ideal allocation between domestic and international stocks is 50/50. This makes sense as the two assets have the same expected return and risk apart from their covariance with domestic human capital, and a retiree has no human capital. Second, for a very young worker with almost nothing but human capital, the profit shock scenario is a big enough deal that it actually causes a slight home bias. This is admittedly driven by my assumptions about what these scenarios actually will entail – that a profit shock means a permanent small fall in savings whereas a depression means a large temporary fall in savings. But given those assumptions, it is still the case that for any worker older than about 25, the fact that domestic stocks do worse in a local depression and a purely international depression has no impact on domestic human capital means that domestic stocks look riskier than international ones and an anti-home bias is the rational portfolio assuming expected returns are otherwise identical between the groups of stocks.

Don't look for your keys under the lamp post if you lost them in the alley

While I used MVO in all the investment examples above, I am not claiming that it, or quantitative tools generally, are the best way to build a portfolio. At heart, my point has little to do with quantitative methods at all. I was simply using a few straightforward examples to show ignoring the risks that lie outside of your investment portfolio can lead to running the wrong portfolio, and that incorporating the risk to your liabilities and/or non-portfolio assets can make answering some questions that are otherwise difficult to judge significantly more tractable.

Solving for the right portfolio irrespective of the factors you face outside your investment portfolio is solving the wrong problem. A simpler problem, to be sure, but the wrong one. Whether or not you use quantitative tools to help guide your portfolio construction, the act of trying to determine the nature of the true problem you should be solving for is likely to lead to insights that otherwise wouldn't occur to you. And in a number of cases, the implications of looking beyond your investment portfolio into the larger problem your investment portfolio exists to solve suggest that the conventional wisdom in portfolio construction is flawed.

Disclaimer

The views expressed are the views of Ben Inker through the period ending May 2019, and are subject to change at any time based on market and other conditions. This is not an offer or solicitation for the purchase or sale of any security and should not be construed as such. References to specific securities and issuers are for illustrative purposes only and are not intended to be, and should not be interpreted as, recommendations to purchase or sell such securities.

Copyright © 2019 by GMO LLC. All rights reserved.

¹ These are intended more as caricatures of worker types whose labor incomes have meaningfully differing sensitivities to economic conditions.

² The covariance is a measure of the joint variability of the assets, and you can think of it as being a function of both the volatility of the assets and their correlation.

³ Because at retirement there is no more human capital, the covariance between human capital and stock and bond returns becomes irrelevant.

⁴ By this, I mean a standard optimization considering only the investment portfolio.

⁵ I assumed that neither unexpected inflation nor a liquidity shock had an impact on expected savings from wage income.

⁶ Time horizon and mean reversion are meaningful issues, and the portfolio for a 65-year-old and an 85-year-old, for example, should generally be different. But part of the reason why the managers of the "glide paths" for retirement portfolios rely on time horizon to justify their shifting portfolios is because they don't have a way of directly considering human capital in building those portfolios.

⁷ This is almost certainly true for any economic event that is primarily domestic in nature – a domestic recession reduces the demand for imports without having much impact on the demand for a country's exports. For a small open economy, however, it might well be that in any global downturn, the fall in its exports will be larger than the drop in imports, particularly if the imports are of necessities such as food or energy. The historical data for the tradedominated economies in Asia is mixed on this.

⁸ There are a number of other assumptions I'm making in setting up the problem. Notably, I'm assuming the present value of future cash flows into the fund is equal to the fund's current value. Commodity equities and non-commodity equities are assumed to have meaningfully different expected returns in the scenarios discussed, and the "importance" of the commodity boom and bust events that are not associated with either overall inflation or depression is half as important as the depression and inflation events for the purposes of building the covariance matrix.

⁹ Due to rounding, portfolio weights may not sum exactly to 100%.

¹⁰ I assumed that the liability for the hunger foundation grew 20% in real terms in the event of a depression, whereas a depression had no impact on the real value of the cancercuring liability.

¹¹ William Brainard and James Tobin, "On the Internationalization of Portfolios," Cowles Foundation Discussion Paper 991.

¹² Laura Bottazzi, Paolo Pesenti, and Eric van Wincoop, "Wages, Profits and The International Portfolio Puzzle," European Economic Review, Volume 40 Issue 2.