Amherst College Endowment

The role of "asset allocation" and "manager selection" in endowment management
Why we don't have asset allocation targets

I was recently asked how we emphasize manager selection versus asset allocation in our investment process, and it gave me an excuse to finally put pen to paper on what is a complex and fascinating topic. Ultimately, I think separating asset allocation from manager selection results in suboptimal decision making. Our investment process should embrace the overlap between managers and asset allocation, and the central organizing principle for every investment in the portfolio is that it – *both* the manager and the opportunity set - must compete with the rest of the portfolio on a long-term risk/ return basis.

Let's talk about asset allocation, and start with the assertion that asset allocation is the greatest driver of returns. Lots of studies have shown this, and it feels intuitively true. The last couple of years have some good examples of this - if you had a large venture portfolio in 2021, you did great; if you didn't, you didn't do as well. And the long-time manager of the Yale endowment David Swensen literally wrote the book on this, which espouses the criticality of asset allocation decisions. The thing that isn't discussed enough in this conversation, is that these studies typically analyze absolute returns - the attribution is versus cash. If you instead analyze relative peer performance, or performance vs diversified asset benchmarks, manager selection jumps in importance as a return driver. And that's a key point, because our alternative for the Amherst endowment is not cash - it's a fully-invested portfolio of some kind. And if you go back again to the idea of 2021 being a good year to see the importance of asset allocation through venture where the delta between US equity and venture was 21%, you can also take the example of US equity in that year, where the delta between investing in US financials stocks and US energy stocks was 20%, showing that manager selection - between a manager that's skilled in financials versus one that focused on energy - is also very important *in that same year*. So yes, asset allocation is key, but there's a line of thought out there that it's the most important decision you make - and I very much disagree with that.

The second point I'd like to make is that asset allocation is very hard to get "right". Let's start with the easier piece of this first, which is how difficult it is to time asset allocation. It is very hard to make a judgment time and time again on whether real estate, credit, venture capital, china public equity, US tech stocks, you name it, will out- our under-perform. Too many factors, too complicated. Moreover, it is literally impossible to time private markets as an endowment because we don't actually have control over timing. There are firms that have tried to build businesses around timing asset classes - they have lots of data, fancy models, and impressive people. They haven't been able to do it. And there's no endowment that I know of that's been able to either. So, to put this in the nomenclature of our world, "tactical" asset allocation is a no- no. There are flashes of brilliance, for sure, but consistently good decision making around asset allocation based on what you think will do well — very very hard. So that's pretty simple.

The more complicated aspect of asset allocation decision-making is this notion of a policy portfolio and risk tolerance. When asset allocation policies are established, there are a couple of frameworks that people use to make decisions – there's a framework around long-term risk/ return profiles for asset classes, and there's a framework for selecting a risk tolerance or risk appetite, based on institutional considerations.

Let's start with the inputs to a long-term asset allocation policy- asset class return, volatility and correlation assumptions. You can make more or less reasonable assumptions for the basic asset classes - global equities, the domestic bond market. Although even those assumptions can be quite wrong for extended periods of time, and there are some broad asset classes, like emerging markets equities, where it's actually very hard to do. Moving to more complex asset classes like private equity, hedge funds - setting expected risk/ return profiles for these investments is very hard. These asset classes are comprised of managers that are competing with each other and the market to make money, where people have to innovate to survive. These highly intermediated investments constitute at least half of endowment portfolios. There are also examples of sectors that can be significantly impacted by externalities or a key innovation in the sector. Think real estate, oil and gas, commodities, credit. The point is, setting 20-year volatility, return and correlation assumptions is hard - this is the whole garbage in/ garbage out concept in modeling. And I'm afraid that when it comes to asset allocation policy setting, for the most part it's garbage in. Let's move on now to the other framework for setting asset allocation policy - risk tolerance.

When we talk about risk tolerance for an institution, most of the time we are talking about tolerance for volatility. This is a legitimate issue. Yes, we need to build a portfolio that the College can tolerate in terms of the returns it generates in *good and bad* years, and over the *long-term*. This question is typically translated into portfolio construction language by asking how much equity exposure I can have in the endowment. Should I have a 60/40 equity/bond- equivalent portfolio, or an 80/20 portfolio? This is fine, but there's a (small) leap there between what we care about - our tolerance for good and bad years, and our long-term returns - and our equity exposure. It's fine, but it's a short cut. Then, when we are asked to determine which risk profile is right, it's often discussed as if there's a risk profile that's inherent to an institution. There are elements about any given institution that need to be considered in this process, such as - time frame - is the institution perpetual?, general spending needs - are they sustainable or not?, debt burden, cash and liquidity outside the endowment, etc. These are all very important questions. But the most important question is not, is Amherst a perpetual institution? The most important question is, who is sitting around the table making decisions about the endowment? In this world, it is the people that determine the risk tolerance of an institution. This is the CIO and investment team, the Investment Committee, the Board, and the College leadership and other constituents. Yes, the college is perpetual, but it will only be reflected in the investment approach IF the people around the table are making decisions with that in mind. Yes, the college has a smoothed spending rule that protects the annual budget from swings in the endowment value, so we can take risk in the endowment, but only IF the people involved have bought into that approach and will support the approach when we have a year like we just had and are having. So the risk tolerance conversation is riddled with leaps and approximations - it's very messy, and my general take here is that endowments need to take risk to make the returns they need, but whether you're at a 60/40 or 80/20 "risk profile" or if you think about it that way at all is ultimately about the people around the table (more on this later).

So back to the central question we're dealing with - how do you set long-term asset allocation? Let's review: you have OK data about stocks and bonds, you have pretty poor data for over 50% of your portfolio, and your overall risk tolerance is not some inherent profile you can set and forget, it's determined by people around a table that change over time. Now you can see why "long term" asset allocation targets are changed almost every year. Or why decision-making around those targets is muddied by long-term vs short-term thinking. Or why there's constant friction around this idea that there's "policy" which is separate from "implementation" which basically gets back to the separation between asset allocation and manager selection. And my answer to this question is - you can do it this way, and it's a decent solution for many endowments, but I think it's suboptimal. Which, finally, gets us to the most important question:

Is there a better way? Yes, but it's not easy.

The way we Amherst solve this problem is to reject the set up that pits asset allocation against manager selection, and fully embrace the overlap between these two. Every decision we make in the endowment is a decision about the long-term merits of the investment under review, vs the rest of the portfolio. The merits and consideration of every investment is a function of both the manager itself - how skilled is the manager, what's the organizational risk and alignment, etc - and the strategy they employ. We are trying to build a portfolio that is fully integrated - we have to understand how each investment interacts and stacks up against each other. And there's an effort to have diverse sources of return across the entire portfolio.

We "look through" to every underlying holding and track it over time, to better understand overlap and adjust if necessary. We don't have to be everywhere - there's no requirement to have real estate, or private credit, or agriculture - we have very little exposure to those strategies today. And the burden of decision-making is on us, the investment team - we need to consider the full universe of investment options and make decisions about what we think will do best for Amherst over the next decade or two, accounting for all our advantages and considerations, including manager quality and access, intelligence, the incumbent portfolio, and our assessment of the opportunity set.

And we do this within a general framework of generating high long-term real returns, but not taking full equity risk - we need sufficient liquidity for capital calls and spending, and we need relative protection in the portfolio to survive a GFC or some equivalent market crisis.

A couple notes on this approach. The first harkens back to the original question. I know I said no endowment has been able to consistently time asset classes. And this is true, but I left out a key part of the endowment story when I framed it that way. Over the last 30 years, endowments have been at the cutting edge of new "asset classes" emerging, and this has been a spectacular source of returns for the industry. This was true for venture capital, leveraged buyouts, long/short equity, private credit, oil & gas, the opening up of China. There are many examples of endowments making bank on investing in new categories, and my argument here is that these are very much manager- *and* asset class-driven decisions. No endowment does this on its own - these are by definition new, and for new you *need* partners. These are collaborations between the asset owners and asset managers, based on trusted relationships that allow for experimentation in unproven strategies. It is the perfect example of why trying to disentangle asset allocation from manager selection is the wrong approach. Our managers are our guides, and we want to find extraordinary people looking to innovate and make extraordinary returns for the College in these constantly evolving markets.

Second, and this is where the discomfort I mentioned at the very beginning comes in. We may say we're not making asset allocation "calls" - that we partner with managers, we mash these two decisions together, we're building the portfolio from the bottom-up, etc. But the fact is, every day, our portfolio is expressing a "view" on what's going to work. By sector, geography, asset class, capital structure, company maturity. If you look at the Amherst portfolio today, you would say: you love tech, early and late stage private companies, and US public equities; and you do not like buyouts, real estate, credit, European equities, or emerging market equities outside China. That's a valid observation - you could certainly say our portfolio does express some strong views. But that's an output of our decision-making process, and it reflects decisions that incorporate asset class considerations, but not exclusively. I'd like to make a few points on this:

The shape of the portfolio informs our work process, in that we're constantly checking "do we still love our venture partners and do they continue to earn a place in our portfolio", to which the answer is a resounding yes.

The role of "asset allocation" and "manager selection" in endowment management
Why we don't have asset allocation targets

The shape of the portfolio also helps determine project priorities. For example, we are currently doing a buyout project and a credit project to "test" whether or not we should be investing more heavily in those areas. If real estate becomes a truly big mess, we'll do a project on real estate.

Finally, and perhaps most importantly, we are not in the business of "optimizing" this portfolio day by day or even year by year. We are comfortable with the fact that there are times when our portfolio will look quite lumpy. Usually, this will be after good or bad runs - after which we'll try to use cash flow needs to incrementally rebalance the portfolio. This won't always be possible, like in the case of private investments. And, remember, we have a hefty amount of skepticism around our ability to call what's going to do well or poorly over shorter periods of time. And that's OK. We are in the business of building decades-long partnerships with extraordinary managers trying to generate extraordinary returns for Amherst through the cycle, in a variety of ways.

There are two key complications to this approach.

First, when you push these two things together - asset allocation and manager selection - you eliminate a set of "rules" - the asset allocation targets - which means your understanding of the portfolio and managers becomes absolutely critical. You can't just say: go fill the US equity bucket and build me a portfolio of the best performing US equity managers, and we'll know we've succeeded if it beats the S&P 500 over rolling 5-year periods. This is a very straight forward goal and exercise. It's not easy to do, but it's easy to articulate and evaluate. A CIO doesn't necessarily have to deeply understand every one of their US equity managers, they just need to trust their colleague that's building that portfolio. In our world, we have no US equity target, and our US equity manager has to earn its place versus our venture capital manager and our open mandate manager on a risk/return/fit basis. The way we solve for this is to have deep manager relationships with a very concentrated set of managers, so that we as a team can make these cross-portfolio determinations. Today, we have less than 40 partners, with an average relationship of over 15 years. Most endowments have closer to 100, and I'd guess the average relationship duration is under 10.

The second complication has to do with oversight and governance, which is — this is *not* a typical approach. The investment world understands conceptually that generating excellent returns requires a differentiated approach, but it introduces risk when things aren't going well. It's much harder to defend an underperforming strategy when you're standing alone. Which comes full circle back to the people around the table. A successful investment approach requires conviction and endurance, which at the end of the day, is all about people.

Letitia Johnson, May 2023