NORTHFIELD | modeling financial markets worldwide

Portfolio Construction Under Economic Scenarios

Dan diBartolomeo

APAC Seminar Series

November 2019

Motivation

- Many investment organizations spend a lot of effort to forecast future economic conditions.

- These forecasts are then utilized in various ways (often *qualitative* in nature) to influence decisions such as tactical asset allocation, and “macro-driven” changes in active portfolio strategies.

- The puzzling question for these organization is how to accurately transform particular elements of economic scenarios such as forecasts of interest rates, exchange rates, trade levels, commodity prices, or consumer spending into explicit expectations of return and risk either whether for asset classes, factor bets, active management styles (e.g. “value”, “momentum), or individual securities.

Today's Topic

- In this webinar, we will present how our method of "optimized scenario analysis" in conjunction with our risk models can be used to translate economic forecasts directly into forecasts of return, volatility, skew and kurtosis for any individual security or any set of securities making up an asset class (i.e. a index portfolio).
  - Complex economic scenarios with many different elements are supported across countries, regions, or globally. Any number of entire scenarios (e.g. "recession", "expansion") can be combined over a user specified time horizon.
  - A key benefit of the process is that the variables being forecast can be any economic or financial market measure with available historical time series data. The scenarios are not limited to factors in any of our models.
  - The four moment descriptions of the asset return distributions can then be used as inputs to our Open Optimizer which has recently been enhanced to incorporate the influence of higher moments in determining portfolio allocations while retaining all existing Optimizer functionality.
  - *Our PRISM application will have this workflow automated*

What We Don’t Like About Scenario Methods

- Frequent use of scenarios is around “stress tests” that are hypothetical instantaneous events. In the real world, changes in economic conditions take finite time to play out.
- A better but still weak method of dealing with scenarios is to assume that the changes in conditions occur over a single defined time horizon.
- Most of these processes lose all possible *path dependence* during the interval while things are changing.
- Economic scenarios are defined as point forecasts like “inflation will increase 2%” with a 30% probability.
  - The likelihood that any specific economic variable will change by an *exact amount specified by a point estimate is zero.*
  - Realistic scenarios have to expressed in ranges like “inflation will increase between 2% and 4% with a 30% probability”.

More of What We Don't Like About Scenarios

- To use scenario methods for risk as well as expected return requires that the set of scenarios we put forward is exhaustive, and includes the full possible distribution of outcomes. *It's not "what if", it has to be "every if".*
- If the scenarios have multiple elements, assessing the probability of occurrence is a non-trivial exercise in conditional probability, so many proposed scenarios may have very little chance of actually taking place.
  - Kritzman, Czasonis, Pamir, and Turkington (2019) propose a novel use of the Mahalanobis distance (1936) as a way to estimate the likelihood that a given set of scenario elements is jointly plausible.
  - This measure is widely known as a means to identify outliers in multivariate distributions. Algebra is similar to the calculations of portfolio risk. If the "risk" of our scenario combination is high, probability of occurrence is low.
- There are a number of other issues detailed in past Northfield newsletter articles, https://www.northinfo.com/documents/202.pdf and https://www.northinfo.com/documents/213.pdf.

Our Approach to Forming Optimized Scenarios

- Force all scenario elements to be in ranges (e.g. “GDP rises 2 to 3% in 24 months”)

- Represent in a multi-period framework, so a three year horizon can have path dependence in the interval from now until the horizon.
  - This allows the calculation of portfolio risk conditional on one or many scenarios that need not be exhaustive.

- Allow many different scenario data types (levels, changes) and many separate elements in each scenario

- Allow all scenario elements to be represented as changes so as to capture the potential for hypothetical events outside the historical record.
  - If interest rates are now 2%, and our scenario element is -3% to -4% two years from now, that can be constructed from many random sequences of historical events where the cumulative change in the interest rate is -5% or more.

A Numerical Method We Like

- Since the future is unlikely to be exactly like the past, we should be interested in whether the sequence of past events we have lived through is typical or unusual, given available history.
- As described in our June 2013 newsletter, our preferred numerical simulation method for exploring the distribution of a set of outcomes is "bootstrap" resampling.
  - We can use bootstrap methods to answer the broader question of "what if things had been different" but drawn from a similar distribution. set of factor return experiences.
  - However, rather than using the actual sequence of events (e.g. factor returns) we will be using many sequences of randomized events drawn from an historic set of experiences.
  - In essence, we will assume that the future may follow *any one of an infinite number of paths that we might have experienced in the past.*

Basic Bootstrapping in Brief

- Mechanically, the process is easy and very, very fast.
  - We use any of our risk models to get the factor profile of the portfolio
  - Let's assume we want to make a period by period forecast of the return distribution for the next 12 months and that we have a 240 month history of factor returns.
  - To create our first sequence of synthetic history as our forecast, we draw random number N between 1 and 240. The factor returns for month N are now the first month of our first sequence of our forecast factor events. If we repeat the process 12 times, we will have one full sequence of potential future events.
  - Note that since the choice is random each time, not only is the order of events randomized but some observations may be omitted and some observations may be repeated more than once.
  - The probability of choosing each observation is 1/N at each moment
  - For each path we estimate the return on the portfolio for each month, assuming a random draw from the distribution of idiosyncratic risk.

Let's Add a Little More Realism

- Given the simple computational process, we can repeat this entire procedure many thousands of times in a few minutes to *produce a very robust estimate of the future distribution.*
  - At each point in each path, we can estimate calculate the estimated mean, volatility, cumulative returns, maximum drawdown, etc.
  - We can also analyze the cross-section of paths at each moment in time to describe the period by period distribution of the statistics.

- We can also account for serial properties
  - If we believe that asset returns are serially correlated randomizing the sequences will fail to represent this aspect of the data.
  - To address this we can follow the procedure above, but build our sequences of future events from blocks of multi-month periods so as to capture most of the dependence from one month to the next.
  - The length of the blocks would relate to the number of lags in an autoregressive process.

"Even God Cannot Change the Past": Agathon

- So far, we are just sampling from an empirical distribution.
  - Any of the paths we generate are plausible
  - All of the statistical relationships between factors would hold together
  - We can see how typical or atypical the actual sequence of history within the range of the paths we generate.

- The results are not a lot different than if we did Monte Carlo simulations that incorporated the higher moments and serial properties of the expected distribution.
  - But the use of an empirical distribution at least ensures that effective distribution is realistic (it did actually happen).

- But a lot of things have changed since 450 BC. Even if we can't change the past we can pretend that we can.

Let's Try Playing God

- In terms of our risk simulations what we really want is to combine the rich distributional information of a numerical simulation with the "intuitive" nature of a set of explicit scenarios. Such a combined process is described in Meucci (2008)
  - Attilio Meucci and Dan diBartolomeo were part of a session on this at the Society of Actuaries conference in March of 2005.
- It is possible to "stress test" the projections by filtering the set of past observations from which our projected sequences are built.
  - We could include only months from periods of economic recession, or had rising interest rates or include only months that were perceived as particularly volatile. *Meucci refers to this as "crisp conditioning".*
  - If we have a "seed sample" of N observations and we filter out P observations, the probability of any observation being drawn to fill a position in a particular path is *1/(N-P) or zero.*
  - Now we have dense simulated data in both time series and cross-section, *conditional on the stressful or benign filtering.*

Now Let's Get Flexible

"Who is to say that truth is in the crystal and not in the mist?"
Kahlil Gibran

- We can set up a more flexible process where the probability of any particular observation being drawn for inclusion in a bootstrap path is explicitly defined by the user.
  - Instead of the probability of inclusion being 1/(N-P) we can choose a vector of explicit values for each observation.
  - Each probability p_t must be between zero and one
  - The sum of all values of p_t must equal one
  - Meucci refers to this as *flexible conditioning*
- While obviously feasible, it is not immediately obvious how an investor would decide what values should populate the probability vector.

Scenario Based Flexible Conditioning

- We would like combine bootstrap simulations with explicit scenarios.
  - We can build the probability vector for inclusion of observations so as to fulfill the some explicit scenario within a confidence interval.
  - For example, we could say "Do a bootstrap simulation where ***on every path,*** the 10 year interest rate rises between 297 and 303 basis points, and oil prices decline 11 to 13% over a 12 month interval".
  - Observations with increasing interest rates and declining oil prices get more weight and vice-versa.
  - We can specify any variable for which data exists for the seed sample. We are not limited to the factors of an underlying model.
  - We can generate several different scenarios and select the number of paths to be run for each to represent weights. We just do our cross-sectional statistics on the aggregated paths.
  - *The cross-sectional variation in the paths is an implicit measure of the likelihood of the scenario.* If all paths are similar we know that the only a small fraction of all feasible paths fulfill the scenario.

Shifting Time Scale

- It is also possible to shorten the apparent time length of a path to accommodate different forecast horizons irrespective of the time scale of the original factor return observations. Regulators often want scenarios to play out over short periods like one day or even instantaneously.
- Given an assumption of a particular fat tailed distribution, it is possible to functionally compress factor returns that have been observed over a given interval (e.g. month) into shorter intervals such as days, hours or minutes.
- For example, it is widely documented that high frequency financial return data has strong "fat tailed" characteristics.
- To convert an empirical distribution of monthly data to daily data, we assume the distribution has changed from normal to a T-5 distribution.
- For intraday horizons we assume no knowledge of the distribution and use the Chebyshev boundary probabilities. At ultra short horizons (seconds) we can invoke a stable Paretian distribution.
- See diBartolomeo "Fat Tails, Tall Tales, Puppy Dog Tails", Professional Investor (2007) for details

Calculating the Probability Vector

- Figuring out what probability vector best expresses a given scenario is an optimization problem. We want to find the vector of probabilities such that:
  - All values of p_t are equal to or greater than zero
  - All values of p_t are less than one
  - All values of p_t sum to one
  - The attributes described in the scenarios are fulfilled within the prescribed ranges.
  - We preserve maximum randomness by minimizing the sum of the differences (absolute or squared) between each value of p_t and 1/N

- If you use the Northfield Optimizer to solve the problem, it will also come up with the closest possible probability vector if the scenario is infeasible within the range of outcomes of the seed sample.

An Example: Base Case from Historical Data

Null Scenario					
		Mean	StDev	Skew	Kurtosis
	S&P TSX Composite	3.87	18.07	0.31	3.29
	S&P 500	6.01	14.23	0.06	2.97
	MSCI EAFE	7.46	17.59	-0.05	3.02
	S&P 600	6.68	19.08	0.23	3.1
	MSCI EM	10.07	20.46	-0.05	3.02
	JPM Global Bond Index	-0.21	17.27	0.17	3.7
	BAFI US Govt Inflation Linked	1.8	10.32	0.2	4.43
	Portfolio	4.34	11.19	-0.02	3.02

An Example: Inflation Moves 2-3% in Two Years

Inflation Up						
		Mean	StDev	Skew	Kurtosis	
	S&P TSX Composite	3.2	17.8	0.27	3.26	
	S&P 500	5.68	14.56	0.07	2.98	
	MSCI EAFE	7.34	17.41	-0.04	3.02	
	S&P 600	6.76	19.18	0.22	3.01	
	MSCI EM	9.86	20.53	0.06	3.06	
	JPM Global Bond Index	0	17.59	0.34	4.23	
	BAFI US Govt Inflation Linked	1.96	10.36	0.27	4.81	
	Portfolio	4.44	11.39	-0.04	3.12	
Inflation Down						
		Mean	StDev	Skew	Kurtosis	
	S&P TSX Composite	-18.68	23.02	0.21	2.72	
	S&P 500	-9.41	18.52	0.04	2.67	
	MSCI EAFE	-15.68	23.29	0.15	2.56	
	S&P 600	-11.81	22.86	0.27	2.94	
	MSCI EM	-16.94	26.51	0.24	2.56	
	JPM Global Bond Index	-7.66	16.28	0.48	4.13	
	BAFI US Govt Inflation Linked	-10.28	15.9	-0.05	2.56	
	Portfolio	-9.82	15.69	-0.04	2.72	

Forming and "Stacking" Scenarios

- To make things easier, each scenario can have up to 10 elements and our software provides automated downloads from FRED, the time series database of the St. Louis Federal Reserve Bank.
  - As previously noted, very low probability scenarios are noted by the cross-sectional standard deviation of the path returns being much smaller than the portfolio volatility in a standard risk report.

- In this example, we have three scenarios, Base, Inflation Up, Inflation Down
  - We can assign each probability and sum over a proportional number of paths.
  - For example if we want 50% probability on the Base Case, 35% probability on "Inflation Up" and 15% probability on "Inflation Down"
  - We run 5000 paths of "Base Case", 3500 paths of "Inflation Up" and 1500 paths of "Inflation Down".
  - Note that this combination will produce negative skew and positive kurtosis

- Calculation of the four moments of events over the 10,000 paths is automated.

A Closed Form Way to "Stack"

- If you don't want to run lots of paths you can combine the scenario outcomes algebraically as a "mixture of normals" problem.

$$\mu = \sum_{i=1}^{n} p_i \mu_i$$

$$\sigma^2 = \sum_{i=1}^{n} p_i (\sigma_i^2 + \mu_i^2) - \mu^2$$

$$skew = \frac{1}{\sigma^3} \sum_{i=1}^{n} p_i (\mu_i - \mu)[3\sigma_i^2 + (\mu_i - \mu)^2]$$

$$kurtosis = \frac{1}{\sigma^4} \sum_{i=1}^{n} p_i [3\sigma_i^4 + 6(\mu_i - \mu)^2 \sigma_i^2 + (\mu_i - \mu)^4]$$

- If you want to include all four moments in a stacking the relevant algebra appears in Satchell and Hall (2013).

Optimizing

- Now we should have four moment data for every asset in our portfolio, and the initial portfolio as a whole.
- We have been working for several years on the best way to incorporate four moment distributions into the Northfield Optimizer.
  - The next optimizer release will include our process for optimizing portfolios in the presence of asset level higher moments.
  - All existing optimizer functionality is retained including transaction costs, taxes, and a wide variety portfolio and process constraints.
  - All of economic intuition that arises from mean-variance analysis (efficient frontiers, Sharpe Ratio, etc.) continues to be applicable.
- The early background was given in this 2014 webinar, https://www.northinfo.com/documents/608.pdf and updated information will appear in the lead article in our next newsletter.

Conclusions

- Building portfolios on the basis of forward looking economic scenarios is a popular but often poorly implemented way of doing asset allocation, particularly at the "tactical" time horizon.
- Effective use of this technique requires that methods require scenario elements defined in ranges, which will unfold over a multi-period and path dependent process.
- Forming rational probabilities requires that we can estimate the likelihood that a given scenario will occur.
- Once we have "stacked" multiple scenarios, we can easily obtain ex-ante four moment distributions for each asset and the existing portfolio (if any).
- The recent enhancement of the Northfield optimizer to incorporate higher moments in both risk estimates and portfolio construction allows us to properly represent scenario dependent expectations.