Tobias Tip: Dynamics and VaR

This week's topic is going to cover two of the more complex or confusing topics from the material: coherent risk measures, and the actual mechanics of the Barnstable proposals.

Barnstable Proposals

You'll note from the case that Barnstable strongly believes that the market is highly unlikely to underperform the risk-free rate over the long term. Specifically, they assumed that stocks would continue to return a 13% annualized return, with a 16% annualized volatility, compared to bonds returning 6%.

Selling Puts

The first proposal is relatively straightforward: selling put options on the S&P 500 to generate income. The payoff of a put option is:

$$\text{Payoff}_{\text{Put}} = \max(K - S_T, 0)$$

Where K is the strike price, and S_T is the stock (index) price at maturity. The idea here is that if a put option is out-of-the-money, i.e., when the index price is above the strike price, then if the index continues to rise (or even stays flat!), then the option will expire worthless, and Barnstable will keep the premium collected from selling the option.

More specifically, Barnstable is investing their money in the S&P 500, and then selling put options on the index. These put options have a sensitivity to the index price, known as Δ, which tells us how much the option price will change for a small change in the index price. For example, a Δ of -0.3 means that for every 1 point increase in the index price, the option price will decrease by 0.3. Additionally, note that the underlying index has a Δ of 1, since for every 1 point increase in the index price, the index price itself increases by 1.

So really, this is in effect doubling down on the index going up. Why? Because we're *selling* an option that goes down in value as the index goes up. Therefore, if the index goes up, the option that we sold goes down in value (so we make money), and additionally our index holdings go up in value. However, if the index goes down, then our index holdings go down in value, and the option that we sold goes up in value (so we lose money). This is in effect a *levered* bet on the index going up.

The Trust

The second proposal is a little more complicated. Here, Barnstable is proposing putting their assets into a trust. The trust will then invest in the S&P 500: so far so good – this is the same as before (only the assets are in a trust instead of directly held by Barnstable). In exchange, Barnstable will receive two classes of shares from the trust: Preferred and Common.

The Preferred shares pay out as follows: if, after 30 years, the S&P 500 had returned less than or equal to 6% annually, then the Preferred shares would get whatever the trust assets are worth after 30 years. However, if the S&P 500 had returned more than 6% annually, then the Preferred shares would get back their initial investment plus 6% annualized return.

The Common shares pay out as follows: if after 30 years the S&P 500 had returned less than or equal to 6% annually, then the Common shares would get nothing. However, if the S&P 500 had returned more than 6% annually, then the Common shares would get whatever is left over after the Preferred shares have been paid out (so in effect, the Common shares get all the upside above 6% annualized return).

Again, so far so good. If Barnstable just held these two share types, then the return profile would be *identical* to just holding the S&P 500 directly.

However, this is where things get weird. The second proposal then involves selling the Preferred shares to raise cash – which the case doesn't explicitly say what they plan to do with. The idea is that these Preferred shares should be "relatively safe", since they are the first to be paid out from the trust assets, and should therefore command a relatively high price. Let's suppose that Barnstable re-invests the cash raised from selling the Preferred shares back into the S&P 500.

Then, what are they really doing? Well, they are in effect taking out a loan (by selling the Preferred shares) to invest even more money into the S&P 500! So again, this is a levered bet on the S&P 500 going up.

If they don't re-invest the cash raised from selling the Preferred shares, then really what they're doing is just collecting the return that the S&P 500 gives them (supposing it's more than 6%), minus 6%. That is, they get some cash, but what they're left with is really just a bet on the S&P 500 going up above 6%.

Comparing the Two Proposals

I would personally describe this case as having three proposals: the first is selling puts, the second is re-investing the cash from selling Preferred shares, and the third is just holding the Common shares and not re-investing the cash from selling Preferred shares.

The first proposal is quite simple: it's a levered bet on the S&P 500 going up, with the risk of large losses if the index goes down significantly. However, their upside from the puts is also limited to the premiums collected.

The second proposal is also a levered bet on the S&P 500 going up, but with potentially more upside since there's no cap on the returns from the index. That is, you're effectively borrowing money at a 6% interest rate to re-invest in the index. You also have the risk of large losses if the index goes down, too, since you need to pay back the Preferred shareholders 6% before you get any returns from the Common shares.

The third proposal makes zero sense to me: you're just holding the Common shares, which is identical to holding the S&P 500 directly, but you're losing 6% annualized return to the Preferred shareholders. Why would you do that? Unless Barnstable really needs cash, it makes zero sense to me why they would choose this option. Remember, their *core belief* is that the market will outperform the risk-free rate over the long term – so why would they willingly give up 6% annualized return to the Preferred shareholders? Just borrow money and invest it directly in the S&P!

In my opinion, the second option seems about equivalent to taking out a margin loan to invest more money in the S&P 500.

> *Really, the only difference between the trust structure and a margin loan is that your investment bankers get to charge you juicy fees for setting up a complex trust structure.*

Moreover, there are also many more ways to have leverage with options or other derivatives that don't involve this complex trust structure. Specifically, selling puts (as discussed), entering into a total return swap, buying call options, or even buying index futures would all be much simpler ways to achieve the same levered exposure to the S&P 500.

One thing that is tangentially related to this is the advent of buffer ETFs, where we can package options strategies and treasuries to create a product that has some downside protection and capped upside. There's a Money Stuff article **here** that discusses this in more detail.

VaR and Time Diversification

I recommend reviewing the TA Review on VaR, as the notebook covers many of the concepts in more detail, specifically regarding rolling volatility estimates and actually calculating VaR.

This is going to focus on *coherent* risk measures, and specifically why VaR is not one. Note that the appendix of the textbook covers this in detail, too.

Let's start off with why we care about coherence. The big benefit of coherent risk measures is that they satisfy a property known as *subadditivity*. This means that the risk of a combined portfolio should be less than or equal to the sum of the individual risks. Mathematically, this is expressed as:

$$\rho(X + Y) \leq \rho(X) + \rho(Y)$$

Where ρ is the risk measure, and X and Y are two portfolios. The intuition here is that diversification should not increase risk: combining two portfolios should not lead to a risk measure that is greater than the sum of their individual risks.

> *The single most important formula for risk mangement is, given two portfolios X and Y:*
>
> $$\sigma^2_{(X+Y)} = \sigma^2_X + \sigma^2_Y + 2 \cdot Covariance(X, Y)$$
>
> *Where:*
>
> $$Covariance(X, Y) = \rho_{X,Y} \cdot \sigma_X \cdot \sigma_Y$$

Let's start with an example of a risk measure where subadditivity doesn't hold: variance. Note that we just gave the setup that a measure is coherent if it satisfies $\rho(X + Y) \leq \rho(X) + \rho(Y)$. However, variance doesn't satisfy this property – from the box above! If two assets are positively correlated, then the covariance term is positive, and therefore:

$$\text{Var}(X+Y) > \text{Var}(X) + \text{Var}(Y)$$

Therefore, variance is not a coherent risk measure.

Let's give an example of a coherent risk measure: standard deviation. To see why standard deviation (the square root of variance) works, we can calculate it assuming $\rho = 1$, where variance is maximized:

$$\begin{aligned}
\sigma^2_{(X+Y)} &= \sigma^2_X + \sigma^2_Y + 2 \cdot \text{Covariance}(X, Y) \\
&= \sigma^2_X + \sigma^2_Y + 2\rho_{X,Y}\sigma_X\sigma_Y \\
&\leq \sigma^2_X + \sigma^2_Y + 2\sigma_X\sigma_Y \\
&= (\sigma_X + \sigma_Y)^2 \\
\therefore \sigma_{(X+Y)} &\leq \sigma_X + \sigma_Y
\end{aligned}$$

Thus, standard deviation is a coherent risk measure, specifically it states that for any two portfolios X and Y, the standard deviation of the combined portfolio is less than or equal to the sum of the individual standard deviations.

> *In the extreme, if the two portfolios are perfectly positively correlated, then the standard deviation of the combined portfolio is exactly equal to the sum of the individual standard deviations, and, if* $-1 < \rho_{X,Y} < 1$, *then the standard deviation of the combined portfolio is strictly less than the sum of the individual standard deviations (diversification!).*

I suggest re-reading the above carefully, as there is an exact translation from this finding to time diversification. Specifically, so long as returns are not perfectly positively correlated across time, then the standard deviation of returns over multiple periods will be less than the sum of the individual period standard deviations.

Now back to VaR! Why is VaR not coherent? It's pretty easy to come up with an example of where it doesn't satisfy subadditivity. VaR of our portfolio is actually dependent on the *ordering* of returns.

Let's consider 2 assets, A and B. Asset A has a 4% chance of losing \$100 on any given day, and a 96% chance of making \$1. This gives it a VaR of \$1 at 95% confidence level (that is, the cutoff to be a worst 5%

day is still making \$1). Let's say that B has identical payoffs to A, but is totally independent of A. Then the VaR for B is also \$1 at 95% confidence level.

What happens if we combine A and B into a portfolio? Well, now we have to consider the joint distribution of returns. There is a 0.16% chance that both A and B lose \$100 $(0.04 \cdot 0.04)$, a 0.0768% chance that one loses \$100 and the other makes \$1 $(2 \cdot 0.04 \cdot 0.96)$, and a 0.9216 chance that both make \$1 $(0.96 \cdot 0.96)$. Therefore, the 95% VaR for the combined portfolio is \$99 (since the worst 5% of outcomes include the two cases where at least one asset loses money).

So we just combined two uncorrelated assets, each with a VaR of \$1, and ended up with a combined portfolio with a VaR of \$99! Therefore, VaR is not coherent.

However, just because VaR isn't coherent doesn't mean it isn't useful. In fact, it is a staple across the industry and clearing houses for measuring risk, setting margin requirements, etc. You can also get quite sophisticated with VaR, such as using Monte Carlo simulations to capture complex interactions between assets (final portion of the TA Review).

Make Your Graphs Less Ugly

A quick note on making your graphs less ugly. The default matplotlib styling is pretty bad, and you should consider using a different style – I personally recommend picking a style you like and then doing:

```
import matplotlib.pyplot as plt
plt.style.use('bmh') # Or whatever style you like
```

This can already elevate your graphs:

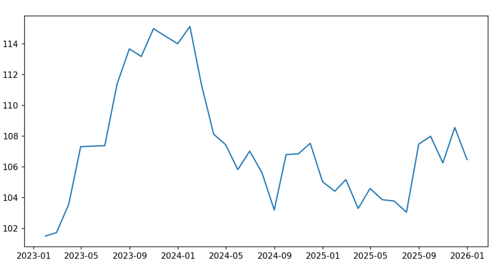

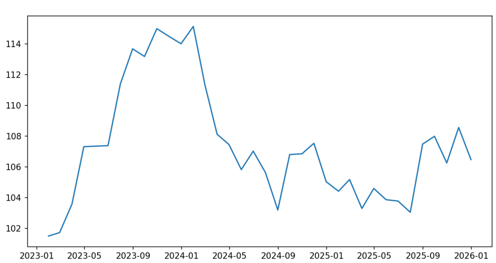

However, that's just the beginning; here are some other features you can use to make your graphs look better:

```
import matplotlib.pyplot as plt
# 1. Start with a theme you like
plt.style.use('bmh')

# 2. Use larger figure sizes and higher DPI for clarity
fig, ax = plt.subplots(figsize=(10, 5), dpi=150)

# 3. Removing some spines for a cleaner look
ax.spines['top'].set_visible(False)
ax.spines['right'].set_visible(False)

# 4. Customize tick parameters
from matplotlib.mdates import DateFormatter, MonthLocator
ax.xaxis.set_major_locator(MonthLocator(interval=3))
ax.xaxis.set_major_formatter(DateFormatter('%b %Y'))
ax.tick_params(axis='x', rotation=45)
ax.yaxis.set_major_formatter('${x:,.0f}')

# 5. Add titles and labels with custom styling
ax.set_title('My Improved Plot', fontsize=16, fontweight='bold')
ax.set_xlabel('Date', fontsize=12)
ax.set_ylabel('Value', fontsize=12)
ax.legend(loc='upper left', bbox_to_anchor=(1.01, 1.0), frameon=False)

# 6. You can use grid lines (if appropriate, but make them light)
# Note that zorder=0 puts the grid behind the data
ax.grid(axis="y", color="#ecf0f1", linestyle="-", linewidth=0.8, zorder=0)
```

These small changes can lead to significant improvements in the look of your plots – this is a sample demo:

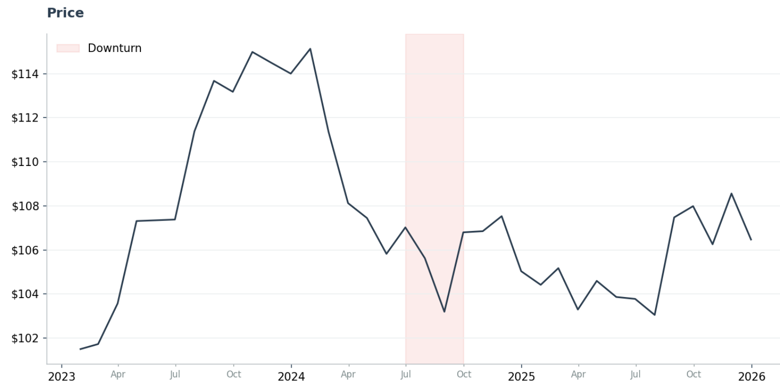

Additionally, you can create your own custom matplotlib styles by creating a .mplstyle file. And, a lot of publications, namely the Financial Times and The Economist, make all of their plots in R using custom themes! If you want to learn more about making nice graphs, there are plenty of resources on it online, but Edward Tufte's "The Visual Display of Quantitative Information" is a classic book on the topic. It's also worth noting that with the latest AI tools, you can directly inject charting guidance into your prompts (e.g., `CLAUDE.md`, `GEMINI.md`).

Bloomberg Opinion

Stocks that only go up (a little)

Buffer funds are probably fine? (Not tax advice!)

I wrote on Monday about a hypothetical, bad financial product:

1. You give me $100 today.
2. In a year, I give you your $100 back, plus any return on the S&P 500 stock index, capped at 5% (i.e. you get $105 back if the S&P is up more than 5%) and floored at 0% (i.e. you get $100 back if the S&P goes down).

I intended to make fun of this product; this product is a little derivatives structuring joke. The way I hedge this product — the way I produce returns that allow me to pay you back the contracted amount — has nothing to do with the S&P 500. I don't buy any stocks. I just buy a one-year Treasury bill, which pays me about 5.1% for the year. Then, if stocks are up, I pay you the stock return, up to 5%. And I keep the extra 0.1%. If stocks are flat or down, I keep the whole 5.1%.

Today Bloomberg's Emily Graffeo reports:

> A new ETF is promising a ... place to ride it out for the next six months, with none of the losses if the stock market swoons.
>
> The Innovator Equity Defined Protection ETF - 6 Mo Jan/July (ticker JAJL) will be the shortest-term 100% buffer fund, an expanding universe of those that use options to deliver some of the stock market's gains while promising protection against the downside.
>
> Over a six-month holding period that begins on the fund's first trading day on July 1, JAJL will match the price return of the SPDR S&P 500 ETF trust up to a cap of roughly 4.8%. Investors will need to hold the fund for the entire period to receive the full risk protection. At the end of the year, they can redeem their shares or roll them into the next cycle.

Hmm: Stock returns with a floor of 0% and a cap of 4.8%. Sorry Inovator! But, no, I'm kidding. My product is bad; Innovator's product is fine. The trick is that my product offers the S&P return up to a cap of 5% for the year, which is what you'd get on T-bills. Innovator's offers the S&P return up to a cap of 4.8% for six months. And then, after the six months, it does it again (with a new cap). If you hold it for a year and the S&P goes up 30%, you probably get a return of like 9% or so, which is much better than my dumb hypothetical.

This is fine. Conceptually the way you do this is something like:

1. You give them $100 today.
2. They spend about $97.40 to buy a Treasury bill that will mature in six months at $100.
3. They use the remaining $2.60 to buy a six-month call spread on the S&P 500. That is, they buy a six-month call option on the S&P, struck at today's price, and sell a six-month call option on the S&P, struck at 104.8% of today's price. This gives them the return on the S&P index above today's price, but capped at 4.8%.

That's economically what's happening, but the actual mechanics are a bit different: They do all of it through Flex options, essentially buying the entire return profile through equity options rather than Treasury bills plus equity options.

That's an important point too. When I made fun of my bad hypothetical product on Monday, I said: "You should not buy it; you should just buy the Treasury bill yourself instead." But several readers emailed to disagree. Their point was: If you buy a Treasury bill, you get 5.1% interest, but you pay taxes on that interest at ordinary income rates. If you buy my dumb product, you get a return of between 0% and 5%, but that return is (probably) in the form of capital gains. If my product runs for a year and a day, it's long-term capital gains, which are taxed at a lower rate than ordinary income, so my 5% return cap is worth more, to you, as a taxpayer, than a 5.1% return on T-bills.

Even better, if my product runs for longer than that — if, say, instead of giving you back the $105 in a year, I roll it over into a new bet for you — then you (probably) don't pay the capital gains taxes until you sell out of my product. If you keep this bet on for 20 years — each year, I give you the return on the S&P 500, floored at 0% and capped at 5% (or whatever the T-bill rate is that year) — then you don't pay the taxes until the end of the 20 years. And then, when you do, you pay capital gains rates. Graffeo writes:

> Elevated interest rates – particularly at the short end of the Treasury curve –- are helping ETF issuers generate the income needed to offer these kind of funds. While that also means investors can get high risk-free payouts on Treasury bills, Bond said the ETF can still provide bigger returns over the six-month period, particularly when the tax advantages are factored in.

I think that this is part of the point of these buffer funds. The proposition they offer is:

1. You don't take equity risk: If you put in $100, you always get back at least $100.
2. You get some return that is sort of centered around the risk-free rate: You might get a bit more than 5% per year (if stocks are up), or you might get a bit less or even zero, but you're not going to lose money and you're not going to make 30%.
3. That return is taxed more favorably than just investing in Treasury bills. Instead of earning interest on Treasuries, you are taxed like you hold stocks for the long term.

Notice how magical this is. I wrote down, above, how you could structure this product (not how it is actually structured): You buy a Treasury bill and some stock options. The Treasury bill accounts for roughly 97% of the value of the product; the stock options account for a bit less than 3%. But the returns on the product are taxed like stock returns, not Treasury bill interest. That's good financial engineering!

We talked earlier this year about another exchange-traded fund, called BOXX (or more formally Alpha Architect 1-3 Month Box ETF), which also used a collection of stock options to achieve relatively risk-free returns. The explicit proposition of BOXX is that you get paid the Treasury bill rate, but it is taxed as long-term capital gains, and you pay taxes only when you sell out of the ETF and realize your gains.

Since then, various tax experts and Bloomberg Opinion columnists have written about it, but buffer funds are probably fine? (Not tax advice!) Buffer funds buy and sell different stock options, different enough that your return is not the T-bill rate but something a bit riskier, a bit more linked to stock returns. Not that much more linked to stock returns — your return here is between 0% and 4.8%, when T-bills would return about 2.7% — but enough.