The Journal of Portfolio Management

INVITED EDITORIAL COMMENT

Factors—*Theory, Statistics, and Practice*

STEPHEN A. ROSS

We're in the middle of a frenzy of factor-focused investing. It seems as though every other article in the investment press talks about a new factor or how old ones are doing, or it reports a new mandate to manage factor tilts or factor-focused portfolios. I am certainly gratified to see the enormous amount of work inspired by the Arbitrage Pricing Theory (APT);[1] but with the estimated asset allocation to factors approaching $1 trillion, perhaps it is time to step back and think a bit about the economic foundations of this effort before we get to $2 trillion.

Identifying a small number of factors that influence or describe market returns formalizes our intuitive understanding of the market. Unfortunately, the tendency to keep adding factors is irresistible. If you are an analytics provider and a client believes that oil prices are a factor, what's the harm in adding an oil factor? And if you are a portfolio manager positioning your portfolio to take advantage of rising interest rates, shouldn't you have an interest rate factor? There is no end to such thinking—but it is undisciplined. There are now more mutual funds than individual stocks, and we could clearly have more factors if trends continue. To restore parsimony, we need to have some criteria for deciding what is a legitimate set of factors, let alone what is optimal.

For the stock market, there are two such desiderata. First, we want a set of factors that explains the differential returns of individual stocks. Why did some stocks rise and others fall—that is, what explains the cross section of realized stock returns, or, narrowly put, what explains the covariance matrix of returns? Second, we want to understand the expected returns on stocks, and to do so we need to measure factor risk premiums. We estimate expected stock returns from betas on the factors. A portfolio inherits its factor exposures or betas from its constituent stocks, and its expected excess return is the expected excess factor returns weighted by the betas of the stocks in the portfolio. This is the security market line of the APT.

It's important to understand that the APT differs from the capital asset pricing model (CAPM) on a fundamental level—and not just because it models many sources of risk that can be priced rather than a single one. The intuition that motivates the APT is based on the strongest force in economics, the absence of arbitrage, which differs from the traditional demand and supply equilibrium argument of the CAPM or its modern successor, the intertemporal CAPM (ICAPM). The CAPM is a one-period model in which the market portfolio is mean–variance efficient, and assets are priced by their covariance, or beta, on the market. In this one-period setting, all wealth at the end of the one-period investment horizon is consumed, so that wealth and consumption are really the same. The ICAPM is a dynamic version of the CAPM in which wealth is a stock variable, consumption is a flow, and all assets are priced by their betas with the growth of (aggregate) consumption. This means that consumption is (locally) mean–variance efficient, and it is the dynamic extension of wealth, the market portfolio in the CAPM.[2]

Although the market portfolio is almost always used as a factor in empirical work, the CAPM makes no use of a factor structure for returns. For example, suppose that both the CAPM and the APT hold and further suppose

that there are three factors driving returns. The APT says that assets with a beta on the second factor have the risk of that factor and enjoy the factor risk premium as the reward for that exposure. In the CAPM, there is only one beta that matters, the market beta.

How can these two views of the world coexist? The answer is that the market is a combination of the three factors, and in the CAPM, the risk of any factor is measured only by its contribution to the market beta; hence, factor risk premiums have no differential impacts on asset returns. The size factor might carry a risk premium that differs from the value factor premium, but the CAPM requires that each of those factors impact an asset's expected returns only through their covariance with the market portfolio.

Although stock returns are proportional to market betas in the CAPM, the market return isn't necessarily a factor in the sense that it explains the cross section of realized stock returns. Typically, though, some market index is used as a factor in multifactor models. This is appropriate, because with multiple factors the market portfolio is a diversified combination of the factors and can be substituted in the set of factors as a surrogate for one of them. Notice, too, that while the APT doesn't identify the factors that move stock returns, neither does the CAPM. The CAPM simply says that in equilibrium, to price an asset you only need to know the market portfolio and the covariance, or beta, of the asset on the market.

Underlying the CAPM are investors with standard utility functions and preferences—which is to say, not preferences in the new "behavioral" style in which framing and myopia and all the other human psychological idiosyncrasies hold sway. I have always been surprised that on occasion I'm accused of being an efficient-market fanatic because of my work on the APT and option-pricing theories. The very essence of the APT, unlike the CAPM, is that it isn't based on strong assumptions of human rationality—that is, of *homo economicus*. All that's required for the APT to hold is that a few greedy investors have access to adequate capital to take advantage of arbitrage opportunities.

Nor is this the loose kind of arbitrage that goes by such names as "risk arbitrage." APT arbitrage is the kind of arbitrage in which you get to borrow at 3% and lend at 4% without any risk. The old story about the $100 bill lying on the ground but not really being there, because if it was someone else would have picked it up before you saw it, means that there are no such arbitrages in traded markets.[3] Nor are these arbitrages that are easily thwarted by the all-too-popular so-called limits to arbitrage. The enthusiasm in behavioral finance for such limitations is silly, and the supportive examples cited typically have nothing to do with arbitrage and much more to do with market frictions that affect bid–ask spreads but not the overall level of market pricing. But that is a topic for another essay.

Although the APT provides a framework for measuring the impact of factors on returns, it does not identify them, and the CAPM tells us to not even worry about APT factors because it only requires that we measure market betas and the market risk premium. Unfortunately, empirical support for market betas explaining expected stock returns is widely regarded as quite weak (and that for consumption is even weaker). This is hardly surprising because even without going so far as behavioral finance in disparaging the traditional expected utility theories, the CAPM places daunting demands on investors to be both lions of rationality and computational demons able to quickly process enormous swaths of information.

If we do decide to use the CAPM and market betas to compute expected stock returns, we then still have to confront a difficult measurement problem. What exactly is the market portfolio in the CAPM? As Roll [1977] pointed out in his famous critique, the market portfolio isn't just the S&P 500 Index or some value-weighted index of stocks; rather, it is every asset, including those not easily observed—for example, it certainly includes real estate and human capital because most of the world's wealth is likely between our ears and not under the ground or on the factory floor. Testing whether the S&P 500 is mean–variance efficient—that is, whether betas on the S&P 500 explain expected returns—has nothing to do with the CAPM. Although it would be very useful if betas on this (or any other market index) explain asset returns, other than if there is only one factor, I know of no reason why this should be so. But even if the index is not a good measure of wealth as is required by the CAPM, it can still be a good factor in the APT.

Although we focus here on equity markets, by any measure the sovereign bond market is the poster child

for breaking returns into factors and using the betas to explain returns. If we leave aside the shortest end of the term structure, three factors at most are sufficient to explain over 96% of the cross correlation. In other words, a regression of the returns of any bond on the three factors has an *R*-squared of more than 96%.

What are the three factors? With such a high fit for any choice of the factors, there will always be a linear combination of the factors with the same explanatory power. Typically, the first factor summarizes the overall level of the term structure; this is the fixed-income equivalent of the market factor, and it accounts for more than 80% of bond return volatility. The second factor is the slope of the term structure or the spread between the long and short ends, and it brings the total explanatory power up to about 94%. The remaining 2% or so comes from the third factor, which is typically taken as a measure of the convexity or curvature of the term structure. For example, it could be measured by the returns on a butterfly portfolio made up of long positions in, say, the 2- and 20-year bonds and two offsetting short positions in the 7-year bond; or alternatively, we could use the implied swaption volatility as the third factor.

This still leaves us mute in identifying the underlying external forces that move interest rates; but for pricing, it doesn't matter what theory we form of the bond market. Whatever the true drivers of interest rates are, they must work through the level, slope, and curvature of the term structure. In practice, all major bond portfolios are controlled by adjusting duration—which is just the beta of returns on the level of rates—by controlling the exposure to changes in the slope, and lastly by controlling the impact of curvature. And, like the market portfolio in the equity APT, the yield level and slope have great economic intuition to justify them as representing underlying economic factors. Unfortunately, though, the degree of precision in the sovereign bond market is an unattainable dream for stock market returns.

The importance of finding the right factors explains why it so absorbs the attention of the equity investment management profession. The first attempt to measure stock market factors was that of Roll and Ross [1980], who estimated the principal components of the covariance matrix of stock returns and then estimated the risk premiums on these factors and tested their ability to explain stock returns. Interestingly, the first factor extracted by this approach is typically not the value-weighted market but rather the equally weighted portfolio. (The second factor looks a lot like a long position on the high betas on the first factor offset by a short position on the low betas.) As of this writing, this work is ongoing by a number of researchers.

Recently, Ait-Sahalia and Xiu [2015a, 2015b] used principal components analysis to estimate a factor model with high-frequency data and obtain median *R*-squareds for individual stocks of around 40%. This contrasts with about 20% on daily data and is quite a strong result. Although this approach provides the most direct way of estimating the factors, it is somewhat unsatisfying in that it offers so little in the way of economic intuition. At the least, we would like to be able to simply give names to the factors it finds.

Another less purely statistical approach was initiated by Chen, Roll, and Ross [1986], who perturbed the variables in a simple Gordon growth-type valuation model for pricing stocks as the discounted dividends and earnings and showed that these perturbed variables—for example, interest rates and earnings—could be used as factors. Success at this would allow us to identify the factors as macroeconomic variables driving stock returns and would conceptually be an improvement over the purely statistical approach that leaves an intuitive identification of the factors unresolved. More than 25 years later, this work is continuing, recently by Fama and French [2015] and, ultimately, as part of the holy grail of identifying the (relatively) exogenous forces that move the stock market.

By far, though, most activity in the area of factor research is devoted to constructing portfolios of assets, such as market indices, whose returns are used as factors. This has become a busy cottage industry, as exemplified by Fama and French with a retinue of followers. The most popular factors currently being used are the market return itself minus the riskless rate, a size portfolio that is long small-cap stocks and short large-cap stocks, a value portfolio that is long stocks with high earnings/price or some other measure of value and short low-value stocks, and a momentum portfolio that is long the stocks that have increased the most in recent history and short the losers. Ait-Sahalia and Xiu [2015a, 2015b] report that when these factor choices are combined with some

others, the total set does a reasonable job of spanning the covariance matrix of returns.

However, whether these or other portfolio factors have risk premiums with alphas relative to the market portfolio is still an area of some dispute. Personally, I share the skepticism neatly voiced in Harvey, Liu, and Zhu [2016], who call into question the statistical strength of the supportive results. After all, some skepticism is warranted if only because we have spent so much time at this and have looked at so many candidates. If we look hard enough to find rewarded "factors" and, particularly, if we only publish the positive results, we will indeed find factors with alphas. But do these really have premiums or is this just a financial version of a Rorschach test in which the ink blots start to look like animals or your parents?

We test factor premiums by requiring that they be "statistically significant," which typically means that they have *t*-statistics over 2. Harvey, Liu, and Zhu [2016] argue that we should have higher standards because we have looked at so many factor candidates. When physicists observed the Higg's boson, they insisted that it had to be a five sigma event. Harvey, Liu, and Zhu [2016] are less demanding; but if five sigma is a good hurdle rate for the physicists, I think that it should be good enough for financial researchers too.

For me, perhaps even more troubling than the empirical evidence is the lack of a strong economic foundation for many of the factor candidates. Even if we accept statistical evidence that momentum, say, is a useful factor with an associated risk premium, why should this be true? Is there some compelling economic argument supporting this? The overall market is obviously a concern for investors and the level of the yield curve is clearly a source of risk for bond portfolios, but why would momentum be a concern? Is there some underlying significant risk it expresses? More to the point, is there a strong argument for why a momentum factor should contribute to a stock's expected return? After decades of searching without compelling answers to such questions, we cannot be entirely comfortable that any empirical results are either valid or enduring. Of course, there are others who take issue with this view and argue that the current state of theory in support of these factors is adequate.[4]

We are still on this journey of identifying the factors driving stock returns. To do a satisfactory job, we will need theory as well as statistics; this cannot be an expedition in the world of data analysis alone. Statistics can only take us so far toward identifying the factors and extracting them from return data. The APT tells us that if we find the right factors, they will carry risk premiums we can use to estimate a stock's expected return. If the factor proxies are portfolios of stocks, we need solid theoretical analysis to give us confidence that they are economically meaningful. For some factor choices—for example, the level of yields or the market index—this is clear just as they stand and it is easy to justify their risk premiums. But many other proxies are not so easy to justify. Nor is it sufficient to allude to some market imperfection or behavioral aberration and use it as an after-the-fact ex post rationale for supporting the use of a particular factor.

The reality is that unlike the sovereign bond markets, stock markets are simply too noisy for the statistical data alone to unequivocally justify the choice of a particular set of factors or their risk premiums. We now need to pay more attention to economics and somewhat less to constructing yet another combination of stock returns to try as a factor. Surely finding yet a "seventh factor" will not add much to either our understanding of stock markets or our ability to construct portfolios that better serve our goals.

ENDNOTES

I wish to thank my colleagues at Sloan and RJA for their helpful comments.

[1]See Ross [1976a, b].

[2]What follows applies to the ICAPM as well as to the CAPM, and the reader can freely substitute ICAPM for CAPM and consumption for the market to take this more modern perspective.

[3]This does not mean that prices are always perfect; rather, it is the remorseless drive, e.g., high-frequency trading, to take advantage of any aberrations that pushes prices to perfection.

[4]See, for example, Asness et al. [2014].

REFERENCES

Ait-Sahalia, Y., and D. Xiu. "Principal Component Analysis of High-Frequency Data." Chicago Booth Research Paper No. 15-39, 2015a.

——. "Using Principal Component Analysis to Estimate a High Dimensional Factor Model with High-Frequency Data." Chicago Booth Research Paper No. 15-43, 2015b.

Asness, C., A. Frazzini, R. Israel, and T.J. Moskowitz. "Fact, Fiction and Momentum Investing." *The Journal of Portfolio Management*, Vol. 40, No. 5 (2014), pp. 75-92.

Chen, N., R. Roll, and S. Ross. "Economic Forces and the Stock Market: Testing the APT and Alternative Asset Pricing Theories." *Journal of Business*, Vol. 59, No. 3 (1986), pp. 383-403.

Fama, E., and K. French. "A Five-Factor Asset Pricing Model." *Journal of Financial Economics*, Vol. 116, No. 1 (2015), pp. 1-22.

Harvey, C., Y. Liu, and H. Zhu. "…and the Cross-Section of Expected Returns." *Review of Financial Studies*, Vol. 29, No. 1 (2016), pp. 5-68.

Roll, R. "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory." *Journal of Financial Economics*, Vol. 4, No. 2 (1977), pp. 129-176.

Roll, R., and S. Ross. "An Empirical Investigation of the Arbitrage Pricing Theory." *The Journal of Finance*, Vol. 35, No. 5 (1980), pp. 1073-1103.

Ross, S. "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory*, Vol. 13, No. 3 (1976a), pp. 341-360.

Ross, S. "Return, Risk and Arbitrage," Wharton Discussion Paper, 1973, published in I. Friend and J. Bicksler, eds., *Risk and Return in Finance*. (Cambridge: Ballinger), 1976b, pp. 189-217.

Stephen A. Ross *is the Franco Modigliani Professor of Financial Economics at MIT Sloan School of Management and managing partner at Ross, Jeffrey & Antle LLC (RJA).*
sross@mit.edu