What Practitioners Need To Know . . .

. . . About Factor Methods

Mark Kritzman
Windham Capital Management

Financial analysts are concerned with factors, or common sources of risk that contribute to changes in security prices. By identifying such factors, analysts may be able to control a portfolio's risk more efficiently and perhaps even improve its return.

I will describe in general terms two approaches often used to identify factors. The first approach, called factor analysis, allows analysts to isolate factors by observing common variations in the returns of different securities. These factors are merely statistical constructs that represent some underlying source of risk; that source may or may not be observable. The second approach, called cross-sectional regression analysis, requires that we define a set of security attributes that measure exposure to an underlying factor and determine whether or not differences across security returns correspond to differences in these security attributes.

Factor Analysis

I begin with a nonfinancial, hopefully intuitive, example. I will apply the insights gained by this example to show how we might go about identifying the factors that underlie the stock market.

Suppose we wish to determine whether or not there are common sources of scholastic aptitude, based upon the grades of 100 students in the following nine courses—algebra, biology, calculus, chemistry, composition, French, geometry, literature and physics. We proceed as follows. First, we compute the correlation between the algebra grades of students 1 through 100 and their grades in each of the other eight courses. Then we compute the correlations between their biology grades and their grades in each of the seven other courses. We continue until we have computed the correlations between the grades of every pair of courses—36 correlations in all. Table I displays these hypothetical correlations.

That all these correlations are positive suggests the presence of a pervasive factor, which is probably related to intelligence or study habits. In addition to this pervasive factor, there appear to be three other factors, or commonalities, in performance.

First, the variation in algebra grades is highly correlated with the variation in calculus and geometry grades. Moreover, performance in calculus is highly correlated with performance in geometry. The grades in these three courses, however, are not nearly as highly correlated with the grades in any of the other six courses. We might thus conclude that there is a common aptitude that underlies performance in these three courses.

Second, performance in biology is highly correlated with performance in chemistry and physics, and performance in chemistry is highly correlated with performance in physics. Again, performance in these courses does not correspond as closely with performance in any of the other courses. We might conclude that there is a common source of aptitude associated with biology, chemistry and physics.

Finally, the grades in composition, French and literature are all highly correlated with each other but not with the grades in any of the other courses. This leads us to deduce the presence of a third factor.

Our next task is to identify these factors. Here we must rely on our intuition. We might reasonably conclude that one of the common sources of scholastic aptitude is skill in mathematics or quantitative methods, because we observe high correlations between performances in the three math courses. Aptitude in science appears to be another common factor, given the high correlations in the three science courses. The remaining source of common variation in course grades pertains to composition, French and literature; we might label this factor verbal aptitude.

We do not actually observe the underlying factors. We merely observe that a student who performs well in algebra is more likely to perform well in geometry and calculus than in French. From this observation, we infer that there is a particular aptitude that helps to explain performance in algebra, calculus and geometry but not in French. This aptitude is the factor.

We should note that these results do not imply that performance in a given course is explained by a single factor. If such were the case, we would observe only correlations of 1 and 0. This point is underscored by the fact that the variation in physics grades is more highly correlated with performances in algebra, calculus and geometry than it is with performances in composition, French and literature. This result is intuitively pleasing, in that physics depends more on mathematics than do composition, French and literature. We might therefore conclude that perfor-

Table I Correlations of Student Grades

	Bio.	Calc.	Chem.	Comp.	Fre.	Geo.	Lit.	Phy.
Algebra	.41	.93	.52	.31	.35	.88	.29	.59
Biology		.39	.94	.49	.44	.50	.31	.90
Calculus			.42	.29	.33	.95	.38	.60
Chemistry				.37	.41	.47	.40	.91
Composition					.87	.28	.94	.35
French						.32	.89	.46
Geometry							.38	.55
Literature								.43

mance in physics is primarily explained by aptitude in science, but that it is also somewhat dependent on math skills.

Now we will substitute stock performance for scholastic performance.

Factors in Stock Returns

Suppose we wish to determine the factors that underlie performance in the stock market. We begin by calculating the daily returns of a representative sample of stocks during some period. In this study, the stocks are analogous to courses, the days in the period are analogous to students, and the returns are analogous to grades.

To isolate the factors that underlie stock market performance, we begin by computing the correlations between the daily returns of each stock and the returns on every other stock. Then we seek out groups consisting of stocks that are highly correlated with each other but not with the stocks outside the group.

For example, we might observe that stock 1's returns are highly correlated with the returns of stocks 12, 21, 39, 47, 55, 70 and 92, and that the returns of the other stocks in this group are all highly correlated with each other. From this observation, we might conclude that the returns of these stocks are explained at least in part by a common factor. We proceed to isolate other groups of stocks whose returns are highly correlated with each other, until we isolate all the groups that seem to respond to a common source of risk.

Our next task is to identify the underlying source of risk for each group. Suppose that a particular group consists of utility companies, financial companies and a few other companies that come from miscellaneous industries but that all have especially high debt-to-equity ratios. We might reasonably conclude that interest rate risk is a common source of variation in the returns of this group of stocks. Another group might be dominated by stocks whose earnings depend on the level of energy prices; we might thus hypothesize that the price of energy is another source of risk. Yet another group might include companies across many different industries that have in common the fact that they derive a large fraction of their earnings from foreign operations; we might conclude that exchange rate risk is yet another factor.

We must first rely on our intuition to identify the factor that underlies the common variation in returns among the member stocks. Then we can test our intuition as follows. We define a variable that serves as a proxy for the *unanticipated* change in the factor value. We regress the returns of stocks that seem to depend on our hypothesized factor with the unanticipated component of the factor value. It is important that we isolate the unanticipated component of the factor value, because stock prices should not respond to an anticipated change in a factor. It is new information that cause investors to reappraise the prospects of a company.

Suppose, for example, we identify inflation as a factor. If the Consumer Price Index is expected to rise 0.5% in a given month and it rises precisely by that amount, then the prices of inflation-sensitive stocks should not change in response. If the CPI rises 1.5%, however, then the prices of these stocks should change in response. In order to test whether or not a particular time series represents a factor, we must therefore model the unanticipated component of its changes.

A reasonable approach for modeling the unanticipated component of inflation is to regress inflation on its prior values under the assumption that the market's outlook is conditioned by past experience. The errors, or residuals, from this regression represent the unanticipated component of inflation. We thus regress these residuals on the returns of the stocks we believe to be dependent on an inflation factor to determine if inflation is indeed a factor.

The approach I have just described is heuristic. It is designed to expose factors by identifying groups of stocks with common price variations. Its intuitive appeal is offset by the fact that it produces factors that explain only part of the variation in returns. Moreover, these factors are not necessarily independent of each other.

With a more advanced mathematical technique—called maximum likelihood factor analysis—we can identify several linear combinations of securities, comprised of both long and short positions, that explain virtually all the covariation in the returns of a sample of securities.[1] These linear functions are called ***eigenvectors***, and the sensitivity of a particular security to an eigenvector is called an *eigenvalue*.[2]

Instead of groups of highly correlated stocks, this approach yields precise linear combinations of stocks that represent independent sources of common variation in returns. In effect, the eigenvectors are the factors. Not only are the factors derived in this fashion independent of each other, but we can derive as many factors as necessary to explain as much of the covariation in a portfolio as we would like.

In order to label these factors, we proceed as described earlier. We determine whether or not the returns of these linear combinations of stocks correlate with the unanticipated changes in the variables that proxy for the factors. Within this context, we represent a security's return as follows:

$$R_i = \alpha_i + b_{i1} \cdot F_1 + b_{i2} \cdot F_2 + \ldots + b_{in} \cdot F_n + e_i,$$

where

- R_i = the return of security i,
- α_i = a constant,
- b_{i1} = the sensitivity of security i to factor 1,
- b_{i2} = the sensitivity of security i to factor 2,
- b_{in} = the sensitivity of security i to factor n,
- F_1 = the first factor representing common variation in security returns,
- F_2 = the second factor representing common variation in security returns,
- F_n = the nth factor representing common variation in security returns, and
- e_i = variation in return that is specific to the ith security.

Issues of Interpretation

Factors derived through factor analysis, whether we employ the heuristic approach described earlier or the more formal approach, are not always amenable to interpretation. It may be that a particular factor cannot be proxied by a measurable economic or financial variable. Instead, the factor may reflect a combination of several influences, some perhaps offsetting, that came together in a particular way unique to the selected measurement period and the chosen sample of securities. In short, factors may not be definable. Moreover, factors derived through factor analysis may not persist through time, or factor 1 in one test may be factor 5 in another test.

We thus face the following tradeoff with factor analysis. Although we can account for nearly all a sample's common variation in return with independent factors, we may not be able to assign meaning to these factors, or even know if they represent the same sources of risk from period to period or sample to sample. Below we consider an alternative procedure called cross-sectional regression analysis.

Cross-Sectional Regression Analysis

Whereas factor analysis reveals covariation in returns and challenges us to identify the sources of this covariation, cross-sectional regression analysis requires us to specify the sources of return covariation and challenges us to affirm that these sources do indeed correspond to differences in return.

Here is how we proceed. Based upon our intuition and prior research, we hypothesize attributes that we believe correspond to differences in stock returns. For example, we might believe that highly leveraged companies perform differently from companies with low debt, or that performance varies according to industry affiliation. In either case, we are defining an attribute—not a factor. The factor that causes low-debt companies to perform differently from high-debt companies most likely has something to do with interest rates. Industry affiliation, of course, measures sensitivity to factors that affect industry performance, such as military spending or competition.

Once we specify a set of attributes that we feel measure sensitivity to the common sources of risk, we perform the following regression. We regress the returns across a large sample of stocks during a given period—say a month—on the attribute values for each of the stocks as of the beginning of that month. Then we repeat this regression over many different periods. If the coefficients of the attribute values are not zero and are significant in a sufficiently high number of the regressions, we conclude that differences in return across the stocks relate to differences in their attribute values.

According to this approach, a security's return in a particular period equals:

$$R_i = \alpha + \lambda_1 \cdot b_{i1} + \lambda_2 \cdot b_{i2} + \ldots + \lambda_n \cdot b_{in} + e_i,$$

where

- R_i = the return of security i,
- α = a constant,
- λ_1 = the marginal return to attribute 1,
- λ_2 = the marginal return to attribute 2,
- λ_n = the marginal return to attribute n,
- b_{i1} = attribute 1 of security i,
- b_{i2} = attribute 2 of security i,
- b_{in} = attribute n of security i and
- e_i = the unexplained component of security i's return.

It is not necessary for the coefficient λ in the above formula to be significantly positive or negative on average over all the regressions. The attribute b_i is a measure of sensitivity to some underlying factor. Suppose the attribute is affiliation with industries that benefit from military spending. If there is an unexpected increase in military spending in a particular period, the coefficient of this attribute will be positive. If military spending declines, the coefficient will be negative. The average value for the coefficient over many regressions may be zero, but the attribute would still be important if the coefficient were not zero in a large number of the regressions.

We can measure the extent to which a coefficient is significant in a particular regression by its t-statistic. The t-statistic equals the value of the coefficient divided by its standard error. A t-statistic of 1.96 implies that the likelihood of observing a significant coefficient by chance is only 5%. In order to be confident that a particular attribute helps to explain differences across security returns, we should observe a t-statistic for its coefficient of 1.96 or greater in more than 5% of the regressions. Otherwise, it is possible that the attribute occasionally appears significant merely by chance.

Which Approach is Better?

I have described two approaches for identifying common sources of variation in stock performance—factor analysis and cross-sectional regression analysis. There are pros and cons with both approaches. Through factor analysis, we can isolate independent sources of common variation in returns that explain nearly all of a portfolio's risk. It is not always possible, however, to attach meaning to these sources of risk. They may represent accidental and temporary confluences of myriad factors. Because we cannot precisely define these factors, it is difficult to know whether they are stable or simply an artifact of the chosen measurement period or sample.

As an alternative to factor analysis, we can define a set of security attributes we know are observable and readily measurable and, through cross-sectional regression analysis, test them to determine if they help explain differences in returns across securities. With this approach we know the identity of the attributes, but we are limited in the amount of return variation we are able to explain. Moreover, because the attributes are typically codependent, it is difficult to understand the true relationship between each attribute and the return. Which approach is more appropriate depends on the importance we attach to the identity of the factors versus the amount of return variation we hope to explain with independent factors.

Why Bother with Factors?

At this point you may question why we should bother to search for factors or attributes in the first place. Why not address risk by considering the entire covariance matrix, as originally prescribed by Markowitz?[3]

There are two reasons why we might prefer to address risk through a limited number of factors. A security's sensitivity to a common source of risk may be more stable than its sensitivity to the returns of all the other securities in the portfolio. If this is true, then we can control a portfolio's risk more reliably by managing its exposure to these common sources.

The second reason has to do with parsimony. If we can limit the number of sources of risk, we might find that it is easier to control risk and to improve return simply because we are faced with fewer parameters to estimate.

I have attempted to provide a flavor for the statistical methodology that underlies the search for common sources of return variation without prejudice toward one method or the other. The choice of a particular approach should depend on one's specific needs, one's biases, and a thorough understanding of the merits and limitations of each approach.

Footnotes

1. *For a review of this methodology, see K. G. Joreskog,* Statistical Estimation in Factor Analysis *(Stockholm: Almqvist & Wiksell, 1963).*
2. *For a discussion of eigenvectors and eigenvalues, see A. C. Chiang,* Fundamental Methods of Mathematical Economics *(New York: McGraw-Hill, 1974), pp. 340–45.*
3. *For a review of Markowitz' approach for estimating portfolio risk, see M. Kritzman, "What Practitioners Need to Know About the Nobel Prize,"* Financial Analysts Journal, *January/February 1991.*