13 Decomposing factor exposure for equity portfolios

David Tien, Paul Pfleiderer, Robert Maxim and Terry Marsh[*]

Abstract

This study addresses the problem of accurately forecasting and attributing risk in equity portfolios. We develop a hybrid methodology which takes advantage of the superior forecasting power of implicit factor models while also attributing portfolio risk to economic factors and firm-specific characteristics. We then compare the relative accuracy of risk attribution using our hybrid approach versus an explicit cross-sectional factor model. We present simulation results which suggest, given realistic parameter values, that the estimation efficiency gained by using the hybrid approach yields substantial improvements over explicit models.

13.1 Introduction

It is well known that the tendency of stock prices to move together is the primary source of return risk for equity portfolios containing more than just a few stocks. Factor models are used to describe and predict these price co-movements across stocks. The factor models can, roughly, be categorized as either implicit or structural; implicit models infer the common factors driving stock returns by looking at the factors' 'footprints' in observed returns, while the structural models specify the factors *a priori* in terms of observable characteristics of stocks or macroeconomic variables.

It seems generally agreed that the implicit models afford superior prediction of stock return risk, e.g. King et al. (1994) who show that changes in stock price volatilities can generally be better explained in terms of cross-sections of stock price changes than as a function of economic variables. Nevertheless, users of implicit risk models often want to decompose an equity portfolio's projected risk exposures with respect to economic variables or corporate characteristics. We begin, in section 13.2, by presenting such a decomposition procedure whereby conditional implicit factor exposures are projected onto a given set of cross-sectional characteristics of stocks that include industry classification, book-to-price, earnings-to-price, dividend yield, log market capitalization, and prior 12-month price momentum. Our procedure is similar in approach to that of

[*] Respectively, Santa Clara University and Quantal International Inc.; Stanford University and Quantal International Inc.; Quantal International Inc.; and UC Berkeley and Quantal International Inc.

hybrid models for default risk, for example Duffie and Singleton (1999) hazard-rate models for debt default that are mapped into the Merton (1974) structural model for default risk.

In section 13.3, we turn to an example in which industry classification is postulated to be an explicit cross-sectional characteristic that is associated with stocks' risk exposures. For illustration, we assume that there are only two industries, and that the *unlevered* returns on stocks in each of these two industries have the same true industry exposure. Although evidence[1], indeed casual observation, suggests that it is quite unrealistic to assume that the asset exposures of all firms in an industry are identical, this assumption (or something like it) is implicitly made by extant cross-sectional models which use zero-one dummy variables for industry. If, in our example, the unlevered companies in the same industry have identical exposures, their levered return exposures cannot be identical save in the unlikely case where their leverage ratios are identical. Our objective is to compare the estimation performance of the implicit factor model *cum* decomposition outlined in section 13.2 with that of a 'straight' cross-sectional model approach in which the levered returns are regressed cross-sectionally on a zero-one industry dummy. We find that the implicit factor model quite readily detects the differences in leverage-induced exposures, while the estimates of 'true' exposure in the cross-sectional approach are thrown off by the misspecification that the levered returns on stocks in the same industry have the same exposure. We also show that the specification error inherent in the explicit model is not corrected by adding leverage as an additional characteristic in the cross-sectional model.

Finally, we compare our procedure for decomposing conditional factor exposures in terms of cross-sectional characteristics of stocks with a recently proposed two-step approach, sometimes also called a 'hybrid' model. In this two-step approach, 'factors' that are specified to be explicit characteristics of stocks are 'taken out' as a risk exposure in a first step, and then an implicit model is fitted to the 'residual' co-movement in equity returns in a second step – all else equal, any significant residual co-movement indicates that there is a misspecification in the first-step explicit model. We show that specification errors such as the intra-industry differences among levered equity exposures are not corrected in this two-step procedure, in contrast to our hybrid model.

We include expected return as well as risk exposure parameters in our analysis since, in an arbitrage pricing theory framework, apparent 'alphas' may indicate misspecification in the risk model (e.g. MacKinlay and Pastor (2000)). Moreover, it is extremely important in practice to know how risk model misspecification affects expected return estimates, since the latter become the benchmark for measuring manager performance. Or, from the opposite point of view, if a manager wants to 'bet on' certain characteristics of stocks that, for example, look cheap, it is very important to know how much added risk is associated with that characteristic.

Like our hybrid implicit factor *cum* decomposition model, the two-step explicit-implicit 'hybrid' model can provide more accurate estimates of total risk exposure and expected return than the explicit cross-sectional risk model. However, the decomposition of that risk with respect to the cross-sectional characteristics can be misleading.

[1] For example, Tufano (1997) looks at stocks in the gold mining industry, commonly considered one of the more homogeneous industry groups, and finds that their exposures in fact vary considerably and also vary across time.

Indeed, there is something of an internal contradiction in the two-step procedure, viz. if the specification for the explicit cross-sectional characteristics in the first step is well specified, then there is no need for a second implicit-factor step. On the other hand, if the first explicit-factor step is misspecified so that indeed the second step is internally consistent, the cross-sectional attribution of risk to characteristics in the first step is potentially flawed and thus the first step attribution is not useful. A recent paper by Asgharian and Hansson (2003) suggests that our results may be quite general, i.e. they find that orthogonal implicit factors better capture risk than a two-step procedure where pre-specified market or industry indexes are fit in a first step, and then implicit factors are fit to the residuals.

13.2 Risk decomposition: cross-sectional characteristics

In this section we show how common factor exposures in a conditional factor model estimated using implicit techniques can be projected onto a set of cross-sectional characteristics of stocks. There are two goals to the exercise:

(a) decomposition of a portfolio's forecasted risk in terms of cross-sectional characteristics of stocks such as industry classification, P/E and/or P/B as measures of value-growth, (log of) market capitalization as a size measure, and momentum; and
(b) use of the estimated parameters for the conditional factor model to compute risk parameters in the cross-sectional 'space', such as the variance-covariance matrix of returns associated with the cross-sectional characteristics.

We show how the variance-covariance matrix of returns associated with the cross-sectional characteristics can be computed from the conditional factor exposures at a point in time, rather than from a time series of coefficients from successive cross-sectional regressions.

Let B_t be the $N \times k$ matrix of time-t conditional risk exposures for N stocks to k factors in the conditional factor model (or, more generally, any model with a set of k factors). Under this formulation the return vector of the N stocks, $\widetilde{R}_t$, is given by

$$\underset{N\times 1}{\widetilde{R}_t} = \underset{N\times 1}{E_t(\widetilde{R}_t)} + \underset{N\times k}{B_t}\ \underset{k\times 1}{\widetilde{F}_t} + \underset{N\times 1}{\tilde{\varepsilon}_t} \tag{13.1}$$

In this specification we assume without loss of generality that the factors themselves are uncorrelated and have unit variance. This means that the variance-covariance matrix of $\tilde{R}_t$ is $B_t B_t' + var(\tilde{\varepsilon}_t)$. Now let H be an $N \times j$ matrix with N rows of stocks and j columns of explicit cross-sectional characteristics for the stocks, e.g. industry classification with a dummy variable equal to unity if a stock is in a given industry, and zero otherwise. Other columns of H might contain characteristics such as P/B and P/E as value/growth measures, dividend yield, log market capitalization, and prior 12-month price momentum. We assume that H has full column rank and in particular $H'H$ is invertible.

To decompose the factor exposures B_t with respect to the cross-sectional characteristics H, we formally project these factor exposures onto the space of the explicit characteristics. Consider the regression of B_t on H:

$$B_t = HQ_t + \Lambda \tag{13.2}$$

where the $(j \times k)$ matrix Q is a mapping from the explicit factor exposures H to the factor exposures B_t. The 'residuals' (Λ) in (13.2) are risk components picked up by the conditional factor exposures B_t that cannot be accounted for by the pre-specified list of characteristics in H. Then we have

$$Q_t = (H'H)^{-1}H'B_t \tag{13.3}$$

and

$$\widehat{B}_t = H(H'H)^{-1}H'B_t \tag{13.4}$$

where $\widehat{B}_t$ is a (linear) projection of B_t onto the space spanned by the characteristics H. Note that if H spans exactly the same space as B_t (possibly with a rotation but without noise), then the H characteristics 'explain' 100% of the B_t and $\widehat{B}_t = B_t$. Also note that if there are characteristics in H that are not related to the conditional systematic risk, B_t, then these spurious characteristics 'drop out', i.e. they do not distort the estimated mapping.

Now we can solve for the variance-covariance matrix G of the returns associated with the cross-sectional coefficients that is implied by the projection of those characteristics on the conditional factor exposures B_t. That is, we can solve for G where:

$$\begin{aligned} HGH' &= \widehat{B}_t\widehat{B}'_t \\ &= H(H'H)^{-1}H'B_tB_t'H(H'H)^{-1}H' \end{aligned} \tag{13.5}$$

We obtain:

$$G = (H'H)^{-1}H'B_tB_t'H(H'H)^{-1} \tag{13.6}$$

We now illustrate the projection (13.4) for some actual 'real-world' portfolios where the cross-sectional characteristics are those often used in industry risk models, specifically: 59 dummy variables for (GIC) industrial classification, plus book-to-price, earnings-to-price, dividend yield, log market capitalization, and prior 12-month price momentum.

We will use (2.4) to decompose a portfolio's forecasted factor risk with respect to the cross-sectional characteristics. As is well known, there is no unambiguous way to decompose a portfolio's volatility in terms of various risk attributes if those attributes are not uncorrelated. We use the following procedure to decompose the portfolio volatility across the factors implied by the cross-sectional characteristics. To begin we identify the cross-sectional characteristic that taken on its own explains the largest portion of the portfolio's systematic risk. Call this characteristic 'a' and let the associated column in H be h_a. We then take all the other characteristics and make them orthogonal relative to characteristic 'a'. This means that column i of H becomes $h_i - h_a(h_a'h_a)^{-1}h_a'h_i$. We then identify the characteristic among the remaining characteristics (orthogonalized relative to 'a') that explains the greatest portion of the risk that is unexplained by characteristic 'a'. We call this characteristic 'b'. We then take all of the remaining characteristics and orthogonalize them relative to both characteristic 'a' and characteristic 'b' and determine which among these remaining characteristics explains the greatest portion of the portfolio risk unaccounted for by both 'a' and 'b'. We proceed in this fashion until all the characteristics have been used. Note that after this procedure is completed there may still be some systematic risk left unexplained. This will occur if the conditional factor exposures are not completely spanned by the cross-sectional characteristics.

We did a cross-sectional decomposition of the sort described above for a set of domestic portfolios using the conditional factor exposure matrix B_t estimated as of 30 July 2003. Summary results are presented in Table 13.1 for the following:

(a) a market cap-weighted portfolio of the largest 5000 exchange-listed US stocks;
(b) an equally weighted portfolio of these 5000 stocks;
(c) a 'small cap-tilted' portfolio where the portfolio weights are inversely proportional to the 5000 stocks' market caps.

As can be seen, the cross-sectional characteristics explain roughly half of the conditional factor exposure for these three portfolios on 30 July – slightly less than

	Largest 5000 listed stocks			**Largest 300 stocks**	**Smallest 300 stocks**
	Mkt cap weights	**Equal weights**	**Small cap weights**	**Equally weighted**	**Equally weighted**
% of factor exposure explained	48	54	53	47	52
Breakdown: (fundamental first)					
Expl. by fundamentals	39	44	43	38	42
Expl. by sector	9	10	10	9	10
(Sector first)					
Expl. by sector	45	51	50	44	49
Expl. by fundamentals	3	3	3	3	3

Attribution of the decomposition given by the projection:

$\widehat{B}_t \quad H(H'H)^1 H'B_t$

where: H is the $n \times m$ matrix with m fundamental characteristics and industrial classifications for the n stocks, where $n = 5000$ or 300, and B_t is the $n \times k$ matrix of Quantal conditional factor exposures on 31 July 2003.

Table 13.1 Cross-sectional attribution of conditional factor exposures estimated by Quantal for a selection of portfolios of US stocks on 30 July 2003

50% when stocks are market cap weighted and slightly more than 50% when stocks are equally or 'small cap' weighted. As just discussed, the decomposition of the portfolio risk exposure is not unique since the fundamental characteristics (e.g. size and E/P) are cross-sectionally correlated with industry classification and with each other. If we first attribute the risk cross-sectionally to fundamentals, and then to orthogonalized industry classifications, roughly 40 percentage points are associated with fundamentals, and 10 percentage points with industry classification.

Although not shown in the table, the most important of the fundamental characteristics in the conditional risk exposure projection is size, i.e. log market cap. This result may not hold in all time periods but we do find that it holds for the projection on 30 July 2003. The result does support the researchers who argue that so-called size effect anomalies in historically observed average returns are due to a higher market risk for small firms. The measures E/P, B/P, and D/P are cross-sectionally correlated with each other and with size (they are all 'inverse of price' characteristics). Once the log of market cap is 'taken out' as a characteristic, only the residual factor exposure remains to be attributed to characteristics such as E/P, B/P, and D/P, and thus the risk attributed to them is less than it would be if size had not been 'taken out' first.

If the factor exposures are first associated with industry classification rather than fundamentals, the fundamentals explain only about 3% of the 30 July 2003 factor risk for the 5000 stocks that is not explained by industry classification.

The last two columns in Table 13.1 provide a straightforward robustness check on the results for the 5000 stocks. They show the decomposition of the factor exposure for (a) an equally weighted portfolio of the largest 300 US stocks; and (b) a portfolio of the smallest 300 US stocks. As can be seen, the decomposition remains roughly the same when applied to these 300-stock portfolios.

To further analyze how conditional risk breaks down in terms of the cross-sectional characteristics of stocks, we formed 59 portfolios of stocks where each portfolio contains only stocks within a single GIC industry classification, equally weighted. We examine the breakdown of conditional systematic risk exposures for these portfolios. Again, the decomposition will not be unique because industry classification and fundamentals are cross-sectionally correlated. To analyze the relation between an industry portfolio's risk exposures on 30 July and its respective industrial classification as a cross-sectional explanatory characteristic, we first 'took out' cross-sectional size (log market capitalization) before projecting conditional factor exposure onto industry classification.

Figure 13.1(a) presents results for the portfolio of Energy Equipment stocks, and Figure 13.1(b) the results for the portfolio of Oil and Gas Industry stocks. As can be seen, the energy equipment classification explains about 8% of the total conditional risk that can be attributed to all industry classifications. Interestingly, about 13% of the energy industry's conditional risk exposure appears to be explained by the semiconductor industrial classification, i.e. since industrial classification is a dummy, the systematic risk of a portfolio of energy stocks might appear to have more to do with semiconductors than with energy! Further (not shown), the Semiconductor, Communications, Software, and Biotech classification appear to explain much of the risk of a number of other industry portfolios as well.

As noted earlier, the cross-sectional analysis was performed as of 30 July 2003. During the preceding months, stocks in the Semiconductor, Communications, Internet Software, and to a lesser extent Biotech and Banks classifications were widely cited as

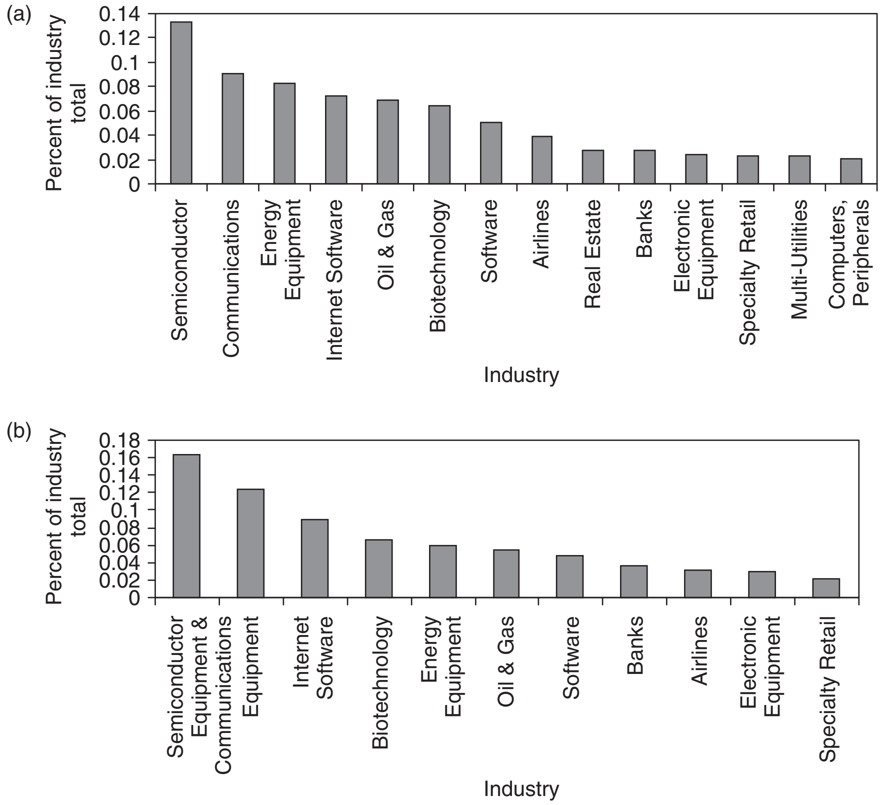

Figure 13.1 (a) Decomposition of energy equipment portfolio risk exposure by industry. (b) Decomposition of Oil and Gas portfolio risk exposure by industry

'driving the market'[2]. Thus, one might hypothesize that the Figure 13.1 results suggest that these 'technology stocks' are in fact proxying for 'the market' in the 30 July environment, and thus that (loosely) a market-wide risk impact is being confounded with an industry effect.

If some of the dimensions of market exposure are being confounded with industry in the Figure 13.1 results, a traditional solution in cross-sectional risk models is to include a historical, say five-year, beta as an additional cross-sectional characteristic in explaining the conditional risk exposures. However, including historical beta leaves the results substantially unchanged. It is perhaps not surprising that unconditional historical beta appears to be a poor proxy for factor exposure given the evidence that it is a reasonably poor predictor of cross-sectional expected returns. To control for 'the market' in a different way, we tried to mimic a market-neutral (and dollar-neutral) energy equipment industry portfolio that is long energy stocks and short NASDAQ 100 stocks where the latter is a one-dimensional proxy for the market. The apparent association between the factor exposures of this long-short energy equipment stock

[2] In Semiconductors: INTC, KLAC, AMAT, NVLS, QCOM, QLGC; in Communication: CISCO, JNPR, RBAK, and NT; in Internet Software: EBAY, YHOO, AMZN; in Banks: C, JPM, BAC, COF, GS, MDW.

portfolio and the technology stock dummy variable is considerably reduced, though not eliminated. Again, this result itself should not be too surprising when the multiple factor exposures are not perfectly captured by the single NASDAQ 100 index. Also, 'interest rate factors' *qua* time-series variables, which are potentially important on 30 July, are not taken into account in the cross-sectional decomposition framework.

The results in Figure 13.1(a) do, however, suggest that industry classification is somewhat informative about the conditional risk exposures of stocks in the respective industry, even though market exposures are confounded with the industry dummies. The Energy Equipment Industry classification is the industry characteristic with the highest explanatory power outside of the groups of stocks that were 'driving the market' in the second quarter of 2003. The results in Figure 13.1(a) are roughly the same as those for the portfolio of Oil and Gas Industry stocks, given in Figure 13.1(b). Both Energy Equipment and Oil and Gas are arguably more homogeneous than most industry groupings, and thus the relative importance of the own-industry classification in explaining factor exposures does not extend across all industry classes. For example, the most important industrial classification in the decomposition of factor exposure for a portfolio of Construction Materials stocks (equally weighted) is not the respective Construction Materials industry dummy variable, but the Real Estate classification dummy variable – it seems completely plausible that real estate would explain the market exposure of construction materials stocks.

Taken together, we think these results point to the zero-one industry classification scheme as a relatively poor instrument for conditional factor exposure. Marsh et al. (1997) show that a classification of stocks into nine market sectors explains roughly the same amount of *ex post* stock return volatility as do much finer gradations in industry classification, e.g. into 29 industries within the nine sectors, a result that is also consistent with industrial classification as being a 'very noisy' indicator of equity risk.

In the next section, we examine a best case in which there is no noise in industrial classification as an instrument for unlevered equity risk. That is the noise consists of differences in leverage across firms in an industry where it is assumed *a priori* that, leverage aside, all firms in an industry have exactly the same all-equity (or asset) factor exposure.

13.3 Decomposition and misspecification in the cross-sectional model: a simple example

13.3.1 Industry classification projected onto factor exposures

In this section, we consider a very simple example of risk decomposition in a cross-sectional model that is misspecified. We consider a case in which all stocks in one of two industries would have the same factor exposure if none of the firms were levered, i.e. their 'asset' or unlevered factor exposures are identical. We make this quite extreme assumption not because it is likely to be realistic, but rather to show that even in this best of possible cases for the cross-sectional model, differences in leverage among the otherwise identical exposure firms is a problem. We show that the implicit model of estimating the factor exposures quite readily detects these differences in exposures – it is the cross-sectional decomposition that is thrown off by the misspecification. Finally, we show that adding leverage as a cross-sectional characteristic in addition to the dummy variables for industrial classification *will not* compensate for the incorrect assumption in the cross-sectional model that within-industry exposures are homogeneous.

To study the effect of varying levels of financial leverage across firms, we simulate a stylized market where asset returns are solely driven by orthogonal industry factors and all firms belong to one of two industries. In this setting a firm's asset return would have an exposure of one to its own industry factor and zero to the other industry factor. Mathematically, the true data generating process for a firm's asset returns is:

$$\tilde{r}^a_{it} = \begin{cases} \tilde{f}_{1t} + \tilde{\varepsilon}_{it}, & \text{if the firm is in industry 1} \\ \tilde{f}_{2t} + \tilde{\varepsilon}_{it}, & \text{if the firm is in industry 2} \end{cases} \tag{13.7}$$

where r^a_{it} is the asset return on firm i at time t, $\tilde{f}_{jt}$ is the return on the industry factor (for industry $j = 1$, or 2) in period t (industry factor returns are assumed to be uncorrelated and assumed be normally distributed), and $\tilde{\varepsilon}_{it}$ is the idiosyncratic return.

Despite the fact that asset returns on all firms have the same exposure to the industry factor, differing levels of financial leverage will cause *equity* returns to have varying industry factor exposures. We have

$$\begin{aligned} \beta_A &= \frac{E}{V}\beta_E + \frac{D}{V}\beta_D \\ \beta_E &= \frac{V}{E}\beta_A - \frac{D}{E}\beta_D \end{aligned} \tag{13.8}$$

where E is the equity market capitalization of the firm, D is the market value of the firm's debt, V is the total value of the firm, and $\beta_D(\beta_E)$ is the firm's debt (equity) beta. If the firm's debt has a beta close to zero, i.e. $\beta_D \approx 0$, then $\beta_E \approx (V/E)\beta_A$. Thus, if we assume, as above, that $\beta_A = 1$, the exposure of a firm's *equity* returns is approximately V/E, not 1.

13.3.2 Incorporating expected return information

Standard approaches to asset pricing would suggest that the conditional exposures B_t will be priced if they are sources of risk to investors, e.g. if the factors reflect the systematic impact across stock returns of shifts over time in perceived investment opportunities (e.g. Merton (1973)). There is also evidence that when stocks are ranked on the basis of cross-sectional characteristics such as relative market capitalization, dividend yield, and value-growth measures (price-earnings or price-book), their average returns have historically differed from those predicted by at least simple unconditional CAPM predictions. One might, then, infer that the cross-sectional characteristics can be linked directly with compensated risk in the factor exposures, and indeed some researchers have taken these historical 'left-hand-side' return anomalies to the 'right-hand side' of asset pricing models by simply labeling the cross-sectional characteristics as 'factors'.

We now include expected returns in addition to factor exposures as parameters to be estimated in the presence of the misspecification in the cross-sectional model due to intra-industry leverage differences. MacKinlay and Pastor (2000) take a similar approach in examining a situation where stock returns have an exact factor structure but a factor is omitted. In an exact factor model, and with the original no-arbitrage reasoning from Ross (1976), an 'alpha' would appear in average returns reflecting the unobserved exposure to the factor that is not correctly included in the model. In our example, the misspecification involves the effect of leverage on intra-industry risk exposures.

We show that, when the standard error of noise in expected returns is comparable to that in practice, incorporating the restrictions on expected returns implied by the factor model is of little help in overcoming the cross-sectional misspecification. Thus, the distortion in the link running from the cross-sectional characteristics, here industry classification and leverage, through equity factor exposures and subsequent expected returns causes a severe problem in measuring alphas, in assessing portfolio manager performance, and in enabling managers to understand where their active 'bets' are incurring risk.

Having generated simulated stock returns using the model outlined above, we compare the performance of our hybrid (implicit *cum* decomposition) model, an explicit model, and a two-step explicit/implicit 'hybrid' model in estimating expected returns on individual stocks. The implicit factor approach essentially allows the data to guide the creation and selection of factors. In this approach, the modeler analyzes the variance-covariance matrix of stock returns using the principal components decomposition, selecting some subset of eigenvectors from the decomposition to serve as implicit factors. Individual stock returns can then be regressed onto these constructed implicit factors (which are orthogonal by construction) to get the implicit factor exposures. The implicit factor exposures are then mapped onto a given set of cross-sectional characteristics as in (13.4).

The explicit cross-sectional approach takes a subset of observed firm characteristics and treats these as exposures to fundamental factors believed to be driving stock returns. The explicit approach takes a long time series of characteristics and runs cross-sectional regressions each time period of returns onto the characteristics (factor exposures in the model) to estimate fundamental factor returns.

In the so-called two-step hybrid risk model, a cross-sectional explicit exposure is estimated first, then the return residuals after adjusting for those first step exposures are analyzed using implicit factor methods. Mathematically, the model below summarizes the gist of the hybrid approach:

$$
\begin{aligned}
r_{it} &= \beta_{i1} f_{1t} + \beta_{i2} f_{2t} + \varepsilon_{it} \\
\varepsilon_{it} &= \gamma_{i1} g_{1t} + \gamma_{i2} g_{2t} + \eta_{it}
\end{aligned}
\qquad (13.9)
$$

where β_{ij} represent observable firm characteristics used to estimate the factor returns, f_{jt}, in the cross-section. The residual, ε_{it}, is then decomposed into (in this example) two implicit factors, g_1 and g_2, with exposures γ_i. Note that this process effectively orthogonalizes the residuals from the explicit model estimated in the first step. As MacKinlay and Pastor (2000) show, the second step in the hybrid procedure can also potentially improve the estimation of expected returns, and thus of portfolio managers' true alphas, by fully exploiting information contained in the covariance among returns on stocks in estimating expected returns.

Table 13.2 outlines the performance of the three approaches along various metrics. In estimating expected returns for industry as a whole, none of the modeling approaches suffers from biased estimates of expected returns – intuitively, the biases in factor exposures and expected returns for individual stocks aggregate away at the industry level. The models diverge somewhat when looking at the efficiency of expected return estimation. The two-step model is roughly on a par with that for the hybrid implicit-*cum*-decomposition approach – like the implicit-decomposition model, the two-step approach benefits from the flexibility of specification that allows it to approximate the nonlinear relation between leverage and industry betas. In contrast, the explicit model seriously lags the competition, since it is a 'prisoner' of its misspecification.

Expected return estimation

	Annualized bias (%)	Mean squared error (%)
Hybrid method	−0.04	1.99
Explicit method	0.02	4.53
Two-step method	0.03	2.01

Industry factor risk estimation

	Annualized bias			Mean squared error		
	Industry 1 (%)	Industry 2 (%)	Leverage (%)	Industry 1 (%)	Industry 2 (%)	Leverage (%)
Hybrid method	4.08	4.08	13.08	0.44	0.45	0.95
Explicit method	4.18	4.20	14.05	0.44	0.45	0.74
Two-step method	4.18	4.20	14.05	0.44	0.45	0.74

Systematic risk, stock level

	Annualized bias (%)	Mean squared error (%)
Hybrid method	2.82	3.71
Explicit method	12.86	16.57
Two-step method	12.86	16.57

Simulated industry returns were generated iid from a normal distribution with mean return of 10% annually and 20% annualized volatility. Firm leverage was drawn from a uniform distribution on [0,1). Results are based on 1000 trials.

Table 13.2

Though the hybrid and two-step models perform similarly in the sense of estimation accuracy, they do differ dramatically in their ability to correctly attribute risk exposures across stocks. The two-step model assumes that the modeled explicit characteristics capture all observable factors driving returns. Since the portion of risk captured by the implicit step is orthogonal to the explicit factors by construction, there is no way for the second implicit factor step to make up for the first-step specification error in equally attributing risk (assumed to be equal exposures) to the industry classification. In essence, the two-step procedure faces a Catch-22: if the first-step explicit factor exposure step is correct, there is no need for the second implicit factor step; if the explicit factor exposure is misspecified so that there is a need for the implicit factor step, then the implicit factor step will not repair the misspecification in the first step. The fully implicit-*cum*-decomposition approach, however, does not suffer from this weakness. This advantage is apparent when looking at how the models perform when estimating systematic risk at the stock level. The implicit-*cum*-decomposition approach clearly dominates the other estimation strategies. This is not surprising since the explicit and two-step methods are restricted in the sense that all firms are assumed to be equally exposed to their respective industry.

Lastly, we examine the ability of the three models to estimate the association between systematic factor risk exposure and the respective zero-one industry classification. This association is estimated 'directly' in both the explicit and two-step approaches, while we use equation (13.6) to estimate the association in the implicit-*cum*-decomposition approach. Here, we see that the models perform quite closely with the hybrid implicit-*cum*-decomposition approach just beating the explicit and hybrid models.

13.4 Summary and discussion

We have analyzed the decomposition of conditional factor exposures for equities with respect to cross-sectional characteristics of the stocks at a point in time. We found that the decomposition in terms of commonly used characteristics explains roughly one-half of the factor exposure captured by the conditional implicit model. We also showed that the decomposition is critically dependent upon the ordering of the cross-sectionally correlated characteristics. We compared the estimates in our decomposition with that of a 'plain vanilla' explicit cross-sectional risk model, and with a so-called two-step hybrid model which combines the plain vanilla cross-sectional model with a second-step implicit model, when there is plausible misspecification in the industry dummy of the cross-sectional model due to intra-industry leverage differences. We find that both of the cross-sectional and two-step hybrid models are 'thrown off' by the misspecification.

It might appear at first glance that the two-step model would be more robust to misspecification problems. Alas, the problem lies with the two-step model's schizophrenia, at least with respect to a wide range of the misspecification problems against which it is intended as protection. The objective of the two-step model's first step – a cross-sectional regression of returns against predetermined characteristics – is the decomposition of risk with respect to these observable characteristics. If the cross-sectional step is correctly specified, there is no need for the second-step implicit procedure. If, however, the cross-sectional step does contain specification error, one might hope that the second step will correct for this, thus better capturing the 'total risk' of a stock or portfolio. But the misspecification in the first step throws off the decomposition with respect to the cross-sectional characteristics, so that it is not obvious what is gained at that step. Decomposition of our (implicit) conditional factor exposures with respect to the cross-sectional characteristics is not thrown off by misspecification in the characteristics, as is the two-step model, while at the same time we suffer no relative loss of prediction in the factor exposures.

Of course, a correct model connecting factor exposures to cross-sectional characteristics could be quite useful, just as are correctly specified structural models and rosetta stones in general. For example, suppose that one knew that 'truth' is that market-to-book is the one and only linear risk factor. Then clearly we know immediately that a manager whose projected returns do literally 'line up' with respect to market-to-book has no portfolio problem – by construction, there is no alpha. If, on the other hand, the manager's projected returns depend on other characteristics as well as market-to-book, it would be very useful to be guaranteed that market-to-book is the only source of systematic risk, since in this case the manager is in a position to form a portfolio that 'zeros out' exposure to that one 'factor' (this is feasible if the alpha characteristics are not perfectly co-linear in market-to-book), and asymptotically the portfolio will have a positive alpha and be risk free.

Unfortunately, we believe that the more realistic situation is one where the portfolio manager has alphas that he or she might know are related to market-to-book (e.g. the manager has discovered a downward bias in earnings expectations for high-tech, high market-to-book stocks), but he or she doesn't know the degree to which these apparent alphas are real, i.e. he/she doesn't know if the market-to-book completely captures the factor risk inherent in this alpha strategy. For concreteness, suppose that a second tech/Internet/'bubble' factor has emerged, and the earnings-expectation-biased/high-tech stocks have an exposure to this second factor and thus a market exposure beyond market-to-book – indeed, even if market-to-book were known with certainty to be related to risk exposure, high-tech stocks which tend to have high market-to-book would look *less* risky along the market-to-book *cum* risk dimension. The argument in this chapter is that in situations like this, using an implicit model with up-to-date conditional exposures can accurately detect the hidden risk and produce better decompositions of that conditional risk with respect to characteristics like market-to-book when high-tech industry is an 'omitted' factor. The manager can then intelligently decide how good a deal the stock is – he or she does not do this using a risk model that is defined along the same dimensions on which alphas are being constructed!

References

Asgharian, H. and Hansson, B. (2003). Investment Strategies using Orthogonal Portfolios. Working Paper, July, Department of Economics, Lund University.

Chen, C.-J. and Panjer, H. (2002). Unifying Discrete Structural Credit Risk Models and Reduced Form Models. Working Paper, 15 July, University of Waterloo.

Duffie, D. and Singleton, K. J. (1999). Modelling Term Structures of Defaultable Bonds. *Review of Financial Studies*, 12:687–720.

King, M., Sentana, E., and Wadhwani, S. (1994). Volatility and Links between National Stock Markets. *Econometrica*, 62(4), July:901–933.

MacKinlay, A. C. and Pastor, L. (2000). Asset Pricing Models: Implications for Expected Returns and Portfolio Selection. *Review of Financial Studies*, 13(4):883–916.

Marsh, T., Pfleiderer, P., and Tien, D. (1997). The Role of Country and Industry Effects in Explaining Global Stock Returns. Working Paper, September 16.

Merton, R. C. (1973). An Intertemporal Capital Asset Pricing Model. *Econometrica*, 41(5), September:867–887.

Merton, R. C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance*, 29:449–470.

Ross, S. A. (1976). The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory*, 13:341–360.

Tufano, P. (1997). The Determinants of Stock Price Exposure: Financial Engineering and the Gold Mining Industry. Working Paper 97-040, Harvard Business School.