Multilayer Modeling of a Market Covariance Matrix

Modeling a layer through its exposure to the previous layer.

Renato Staub

Asset allocation deals with the appropriate allocation of market risks. It relies on a sound understanding of the relationships among all equity, bond, currency, and alternative markets under consideration. In practice, this requires an underlying forward-looking market covariance matrix that is larger than 100×100, describing thousands of interrelationships. A large covariance matrix may easily become contaminated by internal contradictions. Hence, this research is about consistency.

How do we provide a potentially large but nonetheless consistent covariance matrix? In order to construct portfolios appropriately and efficiently, matrix consistency is indispensable.

Investors face a variety of key issues in modeling a market covariance matrix. I first discuss various empirical estimates of risks and correlations. Then I describe the problems we face when setting correlations directly, that is, item by item. Our first solution, a basic two-layer modeling approach, suffers from shortcomings that multilayer modeling can overcome.

There are plenty of empirical techniques for covariance matrix estimations. I feel it is important to add a forward-looking approach that affords flexibility in the face of structural change, inadequate or non-existent historical data, and a myriad of other considerations.

A multilayer modeling approach allows us to incorporate information from various sources, even elements from other risk frameworks, and enables a non-conflicting combination of them. In the end, the multilayer approach offers a solution benefiting from consistency and flexibility.

Renato Staub
is executive director of Asset Allocation & Currency, Global Investment Solutions, at UBS Global Asset Management in Chicago, IL.
renato.staub@ubs.com

HISTORICAL, FORWARD-LOOKING, AND EQUILIBRIUM RISK ESTIMATES

Inconsistency of a covariance matrix is caused exclusively by contradicting correlations. The key advantage of the framework I suggest is its ability to deliver a matrix with consistent correlations.

EXHIBIT 1
Monthly Rolling Volatility Estimates for S&P 500

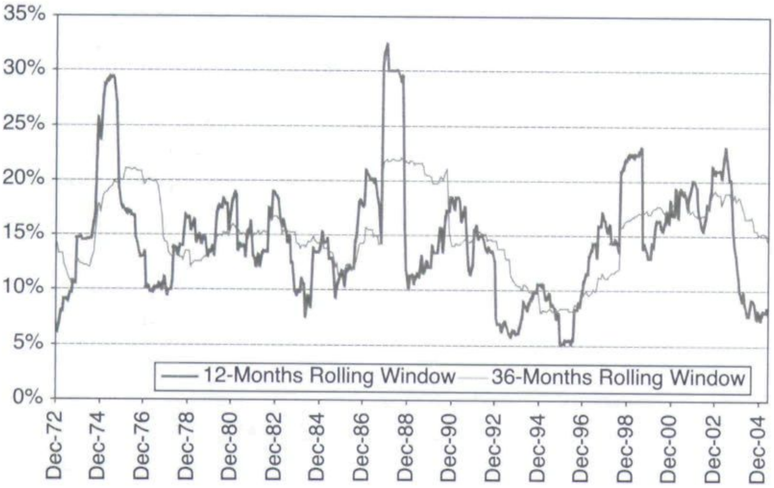

EXHIBIT 2
Monthly Weighted Rolling Volatility Estimates for S&P 500

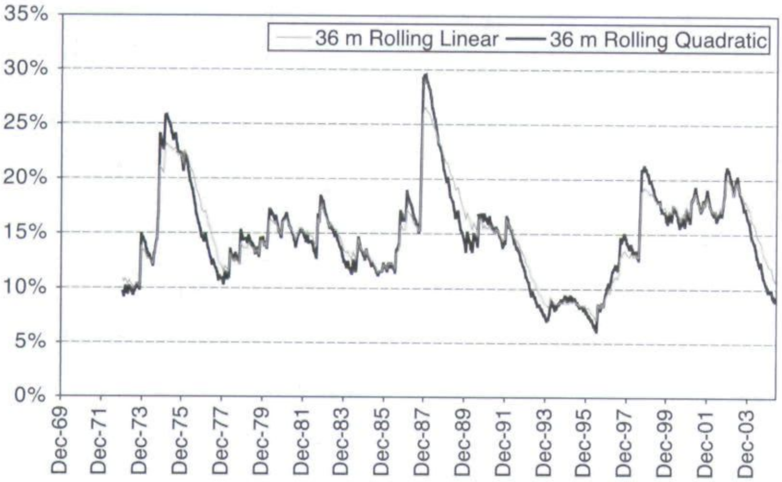

While risks can be of any size from a technical point of view, our approach cannot dispense us from having an idea about the magnitude of the involved risks. In practice, historical volatility estimates help to narrow the range of risk for selected markets.

Consider, for instance, the S&P 500 index for which we provide the estimates of historical volatility in Exhibit 1. *Rolling estimates* measure the return dispersion within a defined historical window. Each time the window is rolled forward, a new observation is included and the oldest one drops out.

But because a window covers a time *span*, it does not estimate *instantaneous* volatility. This conflict is reflected by the typical volatility plateaus, which persist until a return of unusual size drops out at the far end of the window. Rolling estimates do not have predictive power as they calculate the dispersion *within* a sample, but do not investigate serial relationships to be used out-of-sample.

Weighted rolling estimates such as in Exhibit 2 correct the volatility plateaus as an undesired by-product of standard rolling estimates. To that end, we calculate weighted volatilities with higher weights for more recent observations, which dampens the transitory effect of unusual returns.

As might be expected, the quadratic function allows an even faster risk reversion. At the same time, there are limits to the reversion speed, as rolling means become unstable due to an overemphasized front end. *GARCH estimates* try to *forecast* volatility on the basis of two equations, a *mean* equation (1-A) and a *variance* equation (1-B):[1]

$$y_t = \alpha x_t + \varepsilon_t \tag{1-A}$$

$$\sigma_t^2 = \beta\varepsilon_{t-1}^2 + \gamma\sigma_{t-1}^2 + \delta \tag{1-B}$$

That is, today's variance depends on three inputs:

- A constant.
- Previous news about volatility, proxied by the residual term of the mean equation.
- The previous variance forecast.

GARCH examines serial information, as there is empirical evidence that high volatility, expected as well as unexpected, is inherited by the next period. Granted statistical significance

EXHIBIT 3
Monthly GARCH Estimate for S&P 500

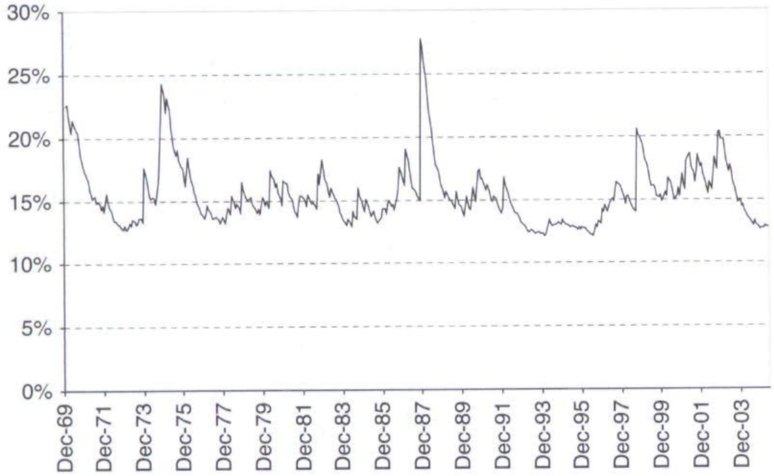

of β and γ, the variance equation predicts (partially) next-period's volatility. That is, GARCH has a competitive edge on forecasting (see Exhibit 3).

At this point, we need to be clear on what kind of risk our matrix needs to reflect. First and foremost, asset allocation is supposed to provide the covariance matrix for *equilibrium* conditions, i.e., balanced markets. These conditions determine the neutral risk allocation, i.e., the long-term policy portfolio, and the fair risk premium as part of the applicable discount rates. Hence, equilibrium covariance does not necessarily equal expected covariance. Considering our previous volatility estimates, most readers would probably agree that the equilibrium risk of the S&P 500 is in the neighborhood of 15%. As it is not driven by a mechanical process, there is simply no exact number, and the most appropriate level is anybody's guess. Setting 14.5%, 15.0%, or 17.0% is a subjective call, and reasonable subjective judgment sometimes requires a departure from historical values.

I can present some examples where this is the case. First, the credit quality of emerging market debt has improved steadily over the past few years. Various measures support this conclusion. A key reason is the introduction of flexible foreign exchange rates that led to less risk rather than higher risk. Historical volatility does not fully reflect this change, so it does not seem appropriate to equate future risks with their historical observation.

Second, several emerging equity markets have become less regulated. Further, the increased addition of foreign-educated professionals to the local work force and improved credit quality of local bond markets have created a more fertile ground for the local equity markets. Hence, we think historical data do not accurately reflect current emerging market conditions.

A third example is the Hong Kong dollar (HKD). As it is pegged to the U.S. dollar (USD), historical data show no HKD volatility versus the USD. We would hence infer no HKD risk from an empirical perspective. In fact, there is no guarantee that a peg will hold forever. Therefore the HKD realistically has some risk.

Fourth, since 1996 the U.K. Treasury bond market has steadily increased its duration from less than six years to over eight years. Hence, representing the risk of the U.K. bond market by historical data only would entail a significant downward bias.

Finally, it is a challenge just to get the past correct in alternative investments such as private equity, real estate, natural resources, and hedge funds. As these investments are not traded on a regular basis, according indexes are appraisal-based and usually significantly smoothed. As a result, any risk estimate based on index volatility is downward biased.

And as if this were not enough, the reporting lags of appraisals mean that alternative investments are not synchronized with liquid assets. Hence, historical estimates of correlations between illiquid and liquid markets are tilted toward zero.[2]

All these examples demonstrate quite clearly that we cannot accept historical volatilities by themselves. Their true power lies in their combination with qualitative assessment.

PROBLEMS WITH DIRECT CORRELATION SETTING

Historical data would support an estimated correlation of more than 0.40 between U.S. equity and U.S. bonds from the mid-1970s thorough 1998. As inflation is one common factor influencing stock and bond returns, and I do not feel that inflation shocks similar to the 1970s will happen again in the medium term, a realistic forward-looking equilibrium correlation would seem to be lower. In 1998, we reduced it to 0.30. And as inflation has been unusually low over the past few years, we did not lower this number as the most recent empirical data may suggest.[3]

As this example demonstrates, "setting" correlations does not mean arbitrariness. It simply means that a perception is translated into a number by qualitative judgment, rather than quantitatively.

There is nothing inherently wrong about setting correlations directly, but direct correlation setting can easily lead to inconsistencies. While it is not necessarily straightforward, for instance, that three markets cannot have mutual correlations of −0.52, it is no more obvious that mutual correlations of −0.50 are feasible.

Risks and correlations can be represented geometrically—risks by straight lines and correlations by angles. The length of a line corresponds to the size of the risk, and the cosine of an angle between two lines represents the correlation between two risks. Then the risks of three markets, A, B, and C, can be depicted by three lines of proportional length in three-dimensional space as in Exhibit 4.

If all markets were perfectly correlated, all lines would point in the same direction, and all angles would equal zero. On the other hand, the lower the correlations, the more the three sides point away from each other. The border case is achieved when all three lines are in a two-dimensional plane, and hence the three angles, α, β, and γ sum to 360°; see Exhibit 5.

If all mutual correlations were −0.52, which corresponds to an angle of 121.33°, they would sum to 364°, which is impossible.

But while it is straightforward to identify and therefore avoid violations in a small number of markets, it becomes increasingly difficult when the covariance matrix increases in size. For an n × n correlation matrix, we would have to set the following number of individual items (nii)

$$\text{nii} = \frac{n^2 - n}{2} \tag{2}$$

For a 100 × 100 matrix, there would be 4,950 items, a number impossible to manage consistently by direct setting.

A second problem with direct correlation setting is that much effort is wasted in generating redundant observations or conclusions. For example, assume we set the correlation between U.S. equity and various western European equity markets. Most likely, all these correlations are quite similar. Only a relatively small portion of their risks are due to individual issues in these markets.

This emphasizes how important it is to capture common features across markets before looking at the idiosyncratic properties of individual markets, and leads us to look for a more practical approach that maintains consistency throughout the matrix.

EXHIBIT 4

Risks and Correlations in 3D Space

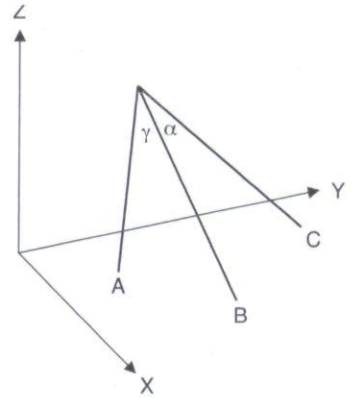

EXHIBIT 5

Border Case—All Three Sides in 2D Plane (Top down view)

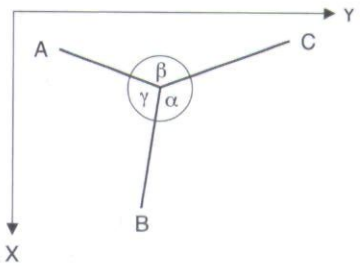

TWO-LAYER APPROACH

A two-layer approach can address the problems I note above. It is a methodology that divides the generation of the matrix into several manageable steps and ensures consistency. In this approach, we first set intuitive factors that drive markets and express their interrelationships. Next, we assign sensitivities or loadings of the various markets on these factors.

Here is an example of the two-layer approach. Assume that the risks of our investable universe are due to two factors only: global equities and global bonds. In this simple case, we start off the modeling process with a 2 × 2 covariance matrix, V_1, representing the global equity and global bond factor. Further, we assume the covariance structure as in Exhibit 6. This is called a *factor covariance matrix*, as it includes the factors only.[4]

Next, in order to derive our *market* covariance matrix, we need to know how the markets respond to factor movements. Their responses are reflected by the so-called sensitivity or loadings matrix, L. Essentially, L represents the betas of the markets to be modeled with regard to the factors.

If market A moves by 110 basis points in response to a 100 basis point move in the global equity factor, the corresponding position in L is 110%. In addition, every market has some risk that is not explained by the factors, so-called idiosyncratic or residual risk, Res.

Exhibit 7 describes a hypothetical example consisting of five markets only. In this example, market A is an equity market, as it loads only on global equity. Note that A's zero sensitivity to global bonds does not necessarily mean that it is uncorrelated with global bonds. Rather, this suggests only that the global bond factor is not one of market A's drivers. Through the positive correlation between global equity and global bonds, market A is still positively correlated with global bonds, although very moderately.

The resulting market covariance, V_2, is:

$$V_2 = LV_1L' + [\text{diag}(\text{Res})]^2 \tag{3}$$

where

$$V_{2,s} = LV_1L' \tag{3-A}$$

is its contribution to systematic risk through the markets' factor exposure, and

$$V_{2,u} = [\text{diag}(\text{Res})]^2 \tag{3-B}$$

is its contribution to unsystematic risk through the markets' unsystematic risk.

Given the structure as set in Exhibits 6, and 7, we get the covariance matrix across the five markets A, B, C, D, and E as in Exhibit 8.

A key advantage of the layer approach is that we formulate the *relations between markets and factors* rather than the *relations among the markets*. That is, we deal with many fewer relations than we would have to in the case of direct correlation setting.

An invaluable feature of the layer approach is that V_2 is consistent as long as V_1 is consistent, no matter how the sensitivities and the residual risks are set. At the same time, ensuring the consistency of V_1 is not a challenge, as V_1 proves to be very small.

To sum up, we have presented here a two-layer approach with an original layer including the factors and a second layer that depends on the first layer and includes the markets to be modeled.[5]

MULTILAYER APPROACH

In practice, a two-layer approach is not sufficient. It does often not let us achieve the correlation pattern that we have in mind. To demonstrate this limitation, assume we want to model various U.S. real estate sectors such as apartment, industry, office, and retail on the basis of a two-layer approach, with U.S. equity and U.S. bonds as the two available factors.

Our prior knowledge of the real estate markets would suggest a correlation structure as follows:

EXHIBIT 6
Sample Factor Covariance Matrix, V_1

	Global Equity	Global Bonds	Risk
Global Equity	1.00	0.30	15.0%
Global Bonds	0.30	1.00	5.1%

EXHIBIT 7
Exposures, L, and Residual Risks of Sample Market Universe

	Global Equity	Global Bonds	Residual Risk
Market A	110%	0%	10.0%
Market B	105%	0%	8.0%
Market C	90%	0%	7.0%
Market D	0%	103%	1.2%
Market E	0%	99%	0.9%

EXHIBIT 8
Resulting Covariance Structure

	Market A	Market B	Market C	Market D	Market E	Risk
Market A	1.00	0.76	0.76	0.25	0.25	19.3%
Market B	0.76	1.00	0.79	0.26	0.26	17.7%
Market C	0.76	0.79	1.00	0.26	0.26	15.2%
Market D	0.25	0.26	0.26	1.00	0.96	5.4%
Market E	0.25	0.26	0.26	0.96	1.00	5.1%

- Moderate correlations between the RE sectors and U.S. equity.
- Weak correlations between the RE sectors and U.S. bonds.
- Strong correlations among the RE sectors.

We conclude that in order to achieve a moderate correlation (only) between the RE sectors and U.S. equity and U.S. bonds, we need to add substantial residual risks to the RE sectors. This leads to a data structure like that in Exhibit 9.

The problem here is that the large amounts of residual risk on top of each RE sector differentiate the RE sectors *among themselves*. As a result, their mutual correlations are in stark contrast to what we want to achieve, as they range between 0.12 and 0.15 only.

Technically, we could resolve this problem by setting off-diagonal elements (ODE) in the residual matrix, diag(Res). That is, we would introduce *correlated residual risks* between the RE sectors. This way, we would increase the correlations only among the RE sectors, but the correlations between the RE sectors and other asset classes would remain unchanged. Unfortunately, although introducing ODE improves the correlation structure, there are problems with this approach:

- Correlated residuals imply a missing factor. In other words, this is a poorly specified model structure.
- Introducing ODE increases the chances of creating an inconsistent covariance matrix. In essence, it throws us back to the limitations we face with direct correlation setting.

In short, the challenge of modeling the desired structure cannot be mastered by a two-layer approach. Rather, we need to enhance it. The idea is to broaden the two-layer approach to a more general solution.

First Step

We model U.S. real estate as a *whole* as in Exhibit 10. On the basis of this input structure, we get the covariance matrix in Exhibit 11. This structure matches the desired one, because the correlation between:

- U.S. real estate and U.S. equity is moderate.
- U.S. real estate and U.S. bonds is weak

Second Step

This intermediate matrix serves as an input structure for modeling the extended covariance matrix including the real estate sectors (Exhibit 12). This gives us

EXHIBIT 9

First Approach Modeling RE Sectors. Note the cell shading of the Exhibits, which gets darker in the following order: factor correlation matrix, factor standard deviations, exposures, and residual risks (darkest)

	U.S. Equity	U.S. Bonds	Risk
U.S. Equity	1.00	0.30	15.5%
U.S. Bonds	0.30	1.00	4.5%

	U.S. Equity	U.S. Bonds	Residual Risk
U.S. Equity	100%	0%	0.0%
U.S. Bonds	0%	100%	0.0%
R.E. Apt.	20%	11%	9.1%
R.E. Ind.	21%	19%	9.0%
R.E. Off.	24%	23%	9.7%
R.E. Ret.	22%	16%	9.8%

EXHIBIT 10

Viable Approach to Modeling RE Sectors in order to achieve our target correlation structure: first step, modeling US Real Estate as a whole. Note, the cell shading is explained in Exhibit 9

	U.S. Equity	U.S. Bonds	Risk
U.S. Equity	1.00	0.30	15.5%
U.S. Bonds	0.30	1.00	4.5%

	U.S. Equity	U.S. Bonds	Residual Risk
U.S. Equity	100%	0%	0.0%
U.S. Bonds	0%	100%	0.0%
U.S. R.E.	22%	16%	8.0%

EXHIBIT 11

First Step of Second Approach: Results

	U.S. Equity	U.S. Bonds	U.S. R.E.	Risk
U.S. Equity	1.00	0.30	0.41	15.5%
U.S. Bonds	0.30	1.00	0.20	4.5%
U.S. R.E.	0.41	0.20	1.00	8.8%

EXHIBIT 12

Viable Approach to Modeling RE Sectors in order to achieve our target correlation structure: second step: Note, the cell shading is explained in Exhibit 9

	U.S. Equity	U.S. Bonds	U.S. R.E.	Risk
U.S. Equity	1.00	0.30	0.41	15.5%
U.S. Bonds	0.30	1.00	0.20	4.5%
U.S. R.E.	0.41	0.20	1.00	8.8%
U.S. Equity	100.0%	0.0%	0.0%	0.0%
U.S. Bonds	0.0%	100.0%	0.0%	0.0%
R.E. Apt.	0.0%	9.3%	99.1%	5.3%
R.E. Ind.	0.0%	2.8%	106.5%	4.5%
R.E. Off.	0.0%	0.0%	115.8%	4.5%
R.E. Ret.	0.0%	-5.6%	110.2%	5.4%
	U.S. Equity	U.S. Bonds	U.S. R.E.	Residual Risk

EXHIBIT 13

Second Step of Second Approach: Results

	U.S. Equity	U.S. Bonds	R.E. Apt.	R.E. Ind.	R.E. Off.	R.E. Ret.	Risk
U.S. Equity	1.00	0.30	0.36	0.37	0.38	0.35	15.5%
U.S. Bonds	0.30	1.00	0.21	0.19	0.18	0.15	4.5%
R.E. Apt.	0.36	0.21	1.00	0.77	0.78	0.75	10.3%
R.E. Ind.	0.37	0.19	0.77	1.00	0.83	0.79	10.4%
R.E. Off.	0.38	0.18	0.78	0.83	1.00	0.80	11.2%
R.E. Ret.	0.35	0.15	0.75	0.79	0.80	1.00	11.1%

the correlation matrix in Exhibit 13. This time, the resulting structure is aligned with our fundamental expectations. The correlation between U.S. RE and U.S. equity is moderate, and the correlation between U.S. RE and U.S. bonds is weak. The correlations among the RE sectors are high.

The trick is that the 8% residual risk of U.S. RE (Exhibit 10) differentiates our U.S. RE factor from U.S. equity and U.S. bonds. Yet the differences among the RE sectors are not driven by the 8% residual risk as this is "inherited" as a joint factor by all RE sectors.

Instead, the extent to which the RE sectors differ from one another is driven by their individual residual risks—as added in the second layer (Exhibit 12)—of:

- 5.3% for RE apartment.
- 4.5% for RE industrial.
- 4.5% for RE office.
- 5.4% for RE retail.

Practically speaking, the 8% of residual risk is bundled. Only then do the residual risks of the *individual* real estate sectors spread out, resembling a tree.

Upon first inspection, the final structure shown in Exhibit 13 appears to be a two-layer structure comprising:

- U.S. equity and U.S. bonds on the first layer.
- Apartment, industrial, office, and retail on the second layer.

Yet the portrayal of the new matrix incorporating RE as a two-layer approach omits the fact that, in an intermediate step, U.S. RE was derived on the basis of U.S. equity and U.S. bonds. What matters is the number of laps we need to go through in achieving the final structure.

In this example, we have created a three-layer structure with:

- 1st layer: U.S. equity, U.S. bonds.
- 2nd layer: U.S. RE.
- 3rd layer: Apartment, industrial, office, retail.

It would be impossible to derive the appropriate correlation structure without applying a three-layer approach.

To extend the approach, we need to remember that at one end of the investment spectrum there are very broad factors, and at the other, there is granularity. Ultimately, the world is driven by more than just two factors, but the perception of a factor depends on our initial perspective—that is, which layer we are looking from.

Two factors are sufficient, but only for modeling the next level of granularity after global equity and global bonds, which is still a very broad aggregate. Global equity is sufficient to model global sectors, but not individual stocks. Clearly, we cannot go all the way from the broadest aggregates to individual components with one step only. Rather, we need to go through several layers.

Therefore, we need to replace the two-layer approach with a *multilayer* approach, which is a simple extension of the two-layer approach. First, we capture the broadest aggregates, such as

- Global equity.
- Global bonds.

- Broad regional factors such as U.S., Europe, or Japan in the base covariance matrix, V_1. The elements of the base covariance matrix are *set* rather than modeled:

$$V_1 = \begin{bmatrix} V_{11} & \cdots & V_{1n} \\ \vdots & & \vdots \\ V_{n1} & \cdots & V_{nn} \end{bmatrix} \tag{4-A}$$

EXHIBIT 14
Assignment of Aggregates in Modeling Approach

Layer	# of Aggregates	Aggregates
6	66	National early stage venture capital markets National late stage venture capital markets National LBO markets National mezzanine markets National distressed debt markets U.S. hedge fund strategies U.S. private real estate sectors U.S. REIT sectors
5	22	National corporate bond markets National inflation protected bond markets National real estate markets Regional timber factors Regional farmland factors
4	82	National equity markets National bond markets National cash markets National currency markets
3	25	10 equity sectors Regional currency factors Country factors
2	13	Broad regional bond factors Broad equity sectors Global currency factors Global real estate factor Global timber factor Global farmland factor
1	5	Global equity factor Global bond factor Broad regional factors

Given the consistency problem with larger matrices, this is the prime reason why V_1 tends to be small. Then we model the covariance matrix, V_2, containing the aggregates on the next layer, layer two. Layer two has more granularity, or detail, than the first, but it is still quite broad.

In the next step, we model a further level, V_3, with even more detail:

$$V_2 = L_2 V_1 L_2' + \left[\mathrm{diag}(\mathrm{Res}_2)\right]^2 \tag{4-B}$$

$$V_3 = L_3 V_2 L_3' + \left[\mathrm{diag}(\mathrm{Res}_3)\right]^2 \tag{4-C}$$

$$\vdots$$

$$V_i = L_i V_{i-1} L_i' + \left[\mathrm{diag}(\mathrm{Res}_i)\right]^2 \tag{4-D}$$

That is, the *output* covariance matrix, V_k, as derived on the basis of V_{k-1}, is considered the (*input*) factor covariance matrix for generating the covariance matrix of the next layer, V_{k+1}. In order to do this, we need to set the according loadings, L_{k+1}, and the residual risks, Res_{k+1}. Theoretically, such a chain scheme can be endless. Practically, we need four to six layers for achieving the desired level of detail of the final layer.

Assigning aggregates to selected layers may be a bigger challenge than it would seem at first sight, as this decision requires some judgment. On the other hand, there are constraints. As a national real estate market, for instance, is driven (to some extent) by its national stock market, it must be modeled on a layer following the layer that models the stock market. Our covariance matrix is based on a six-layer approach describing the aggregates in Exhibit 14.

For purposes of illustration, let us follow the modeling line all the way up from the root, i.e., the first layer, to the U.S. leveraged buyout (LBO) market (Exhibit 15). Arrows indicate the output/input relationship. Further, "U.S. → U.S.," for instance, indicates that the U.S. factor is carried forward—that is, unchanged—to a consecutive layer.

Ultimately, the resulting covariance matrix, V_i, is considered a *market* covariance matrix. It is not practical to present the entire market covariance matrix, however, as it includes too much detail and has, in our case, a dimension of

EXHIBIT 15
Modeling Line to U.S. LBO

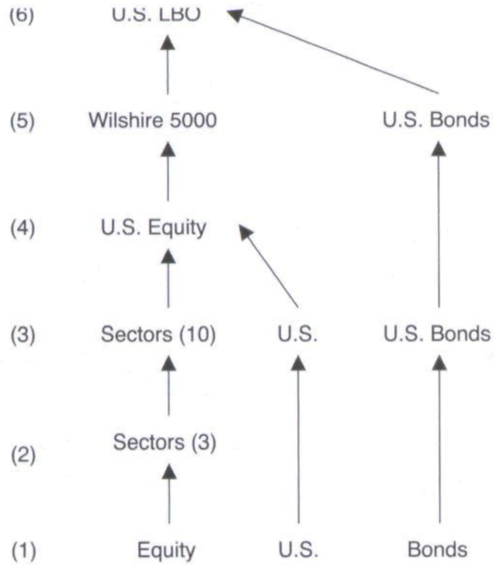

approximately 250 × 250. Typically the user is more interested in the covariance matrix on the asset class level. Creating a matrix on the asset class level, V_{AC}, is obtained through *aggregation*:

$$V_{AC} = w'V_i w \tag{5}$$

where w is the matrix including the weights that sum to 100% for every asset class.

STATISTICAL FACTORS VERSUS PORTFOLIO FACTORS

Assume the three aggregates: U.S. equity (USE), all other equities outside (except) the U.S. (EXE), and global equity (GLE) as the sum of USE and EXE. If we run two ordinary least square regressions, one with USE versus GLE and the other one with EXE versus GLE, the unexplained portions are reported as residuals. Yet Exhibit 16 proves that this is a misperception.

Namely, as GLE is the aggregation of USE and EXE, *all three aggregates are on a two-dimensional plane*, and the residuals of the two regressions are *perfectly negatively correlated*. If we assume them to be uncorrelated when we model them, this would portray a three-dimensional structure as in Exhibit 17. This would suggest that USE and EXE are more strongly correlated than they actually are.[6]

Assuming for simplicity that:

$$\sigma_{USE} = \sigma_{EXE} = \sigma \tag{6-A}$$

$$Res_{USE} = Res_{EXE} = Res \tag{6-B}$$

we conclude that the true angle φ:

$$\varphi = 2\arcsin\left(\frac{Res}{\sigma}\right) \tag{7-A}$$

is understated as follows:

$$\varphi^* = 2\arcsin\left(\frac{Res}{\sqrt{2}\sigma}\right) \tag{7-B}$$

The point is that GLE in Exhibit 16 is a portfolio *combined* of USE and EXE. In Exhibit 17, by contrast, GLE is a purely statistical factor. It is not an investable portfolio, and is not perfectly replicable by USE and EXE. Rather, it is an *orientation in risk space*, pulling USE and EXE into the desired direction, like a guyline that that shapes the envelop of a tent but is not part of the envelop itself.

EXHIBIT 16
True Risk Constellation

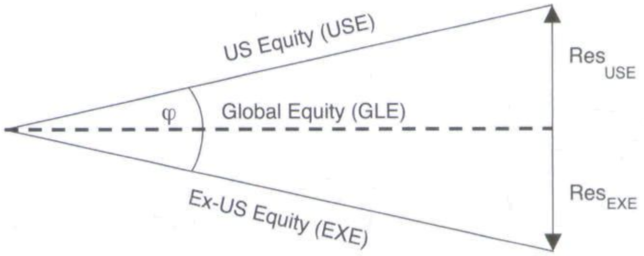

EXHIBIT 17
Represented Risk Constellation

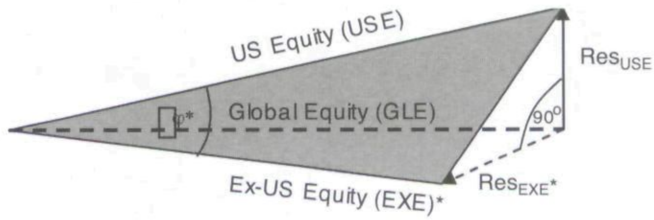

Given assumptions (6) and (7) and interpreting GLE as a statistical factor, the residual risks would have to be increased by a factor $\sqrt{2}$ in order to achieve the same correlation between USE and EXE as if we interpreted GLE as a portfolio.

In fact, the resulting difference between the two approaches is not critical as long as we deal with aggregates of *several* markets rather than two markets only. This is probably the right moment to stress again that—if we decide on a statistical factors approach—the process requires qualitative judgment and is highly circular. The resulting matrix must be cross-checked, which usually leads to an input recalibration and an additional run. The process usually goes through several feedback loops until the result is acceptable.

The portfolio approach suffers from the drawback that aggregated markets must be identical with factors, which means that *residual* risks are not truly residual. The two-market case in Exhibit 16 demonstrates this impressively, although a *multi*market context, of course, is much less constraining. As a result, residual risks cannot be set independently, since they are, nonetheless, correlated and hence not truly residual versus each other. This is overly complicated and leads quickly to an excessive number of constraints.

Complication is the opposite of what we want to achieve, namely, to provide a reasonable and consistent correlation structure on the basis of a process that emphasizes

- Intuition.
- Contents.
- Simplification.

rather than delivering a mathematical tool that generates overly complicated solutions.

IMPLEMENTING A CURRENCY PERSPECTIVE

From a U.S. dollar perspective, the U.S. dollar (USD) is risk-free. Hence, we set it at the intersection of the two risk axes. Then we can introduce two other currencies, C1 and C2, which are risky from a USD perspective (see Exhibit 18).

In turn, C2 has another risk from a C1 perspective than from a USD perspective. The vector subtraction in Exhibit 18 shows that the risk between C1 and C2 is the risk of a portfolio with the two positions: 100% in C2, and −100% in C1.

EXHIBIT 18
Risk Triangle Between USD and Two Hypothetical Currencies

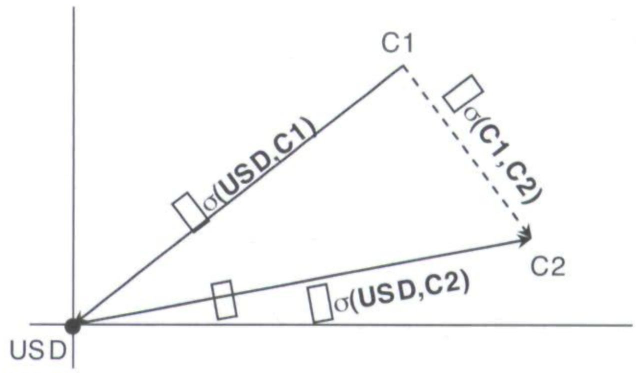

EXHIBIT 19
Data Set of Translation from USD into JPY Perspective

	EUR	JPY	GBP	USD	Risk
EUR	1.00	0.45	0.75	0.00	10.5%
JPY	0.45	1.00	0.34	0.00	10.5%
GBP	0.75	0.34	1.00	0.00	9.1%
USD	0.00	0.00	0.00	1.00	0.0%
EUR	100%	0%	0%	0%	
JPY	−100%	0%	−100%	−100%	
GBP	0%	0%	100%	0%	
USD	0%	0%	0%	100%	
	EUR	JPY	GBP	USD	

Assume we have the currency covariance matrix in Exhibit 19 from a USD perspective, V_{USD}. Now we can easily translate this matrix into another perspective. To do so, we need to go through two steps:

- Set the four currency portfolios EUR, JPY, GBP, and USD.
- Assuming, we want to establish a JPY point of view, we subtract a JPY portfolio from these currency portfolios. The resulting portfolio is called v.

This results in the data set in Exhibit 19. Then the resulting covariance matrix from the selected currency

E X H I B I T 20
Currency Covariance Matrix from JPY Perspective

	EUR	JPY	GBP	USD	Risk
EUR	1.00	0.00	0.80	0.52	11.0%
JPY	0.00	1.00	0.00	0.00	0.0%
GBP	0.80	0.00	1.00	0.65	11.3%
USD	0.52	0.00	0.65	1.00	10.5%

perspective is:

$$V_{i,JPY} = v' V_{i,USD} v \tag{8}$$

This multiplication applies only to the currency part of the covariance matrix, however, as asset returns are in local terms and considered uncorrelated with currencies (Exhibit 20).

SUMMARY

The discipline of asset allocation requires an underlying market covariance matrix as an input. Yet estimating a large matrix is far from straightforward, and there are various alternatives available for construction. A matrix can, of course, be set directly, that is, item by item. But setting a large matrix consistently is almost impossible, as it is highly likely that such a matrix would be subject to internal contradictions. A further challenge is that direct matrix setting is time-consuming and most likely an exercise involving considerable redundancy.

A layer approach addresses all these issues appropriately, as it separates the generation of the matrix into manageable steps, thus ensuring consistency. The practical implications of this are twofold. First, the key advantage is that we deal with relations between markets and factors, rather than relations among markets directly, which means dealing with fewer items. Essentially, only the layer approach keeps the modeling process manageable. And second, the layer approach delivers a consistent matrix, that is, a matrix without internal contradictions. An inconsistent matrix would allow the creation of portfolios with significantly wrong or even negative risks. Although long-only portfolios are less vulnerable to this kind of shortcoming, the potential impact for long-short portfolios should not be underestimated.

In the end, an inconsistent matrix means higher model risk, and an odd risk structure as a result of inconsistency may lead to odd risk premiums and discount rates that are modeled on the basis of the covariance matrix. That is, matrix inconsistencies may be translated all the way into the market valuation. Further, these inconsistencies may be *mean* in that their impact is not immediately obvious, but adds up steadily over time.

Our first approach, a two-layer approach, is a useful didactical device, but it is insufficient in practice. To support a smooth transition from the level of highest aggregation to the level of most detail, we need intermediate layers. This is why we use a multilayer approach. It models a layer through its exposure to its previous layer. Theoretically, such a modeling chain can be endless, but in practice we need four to six layers for achieving the desired level of detail.

Ultimately, we want to provide a consistent correlation structure on the basis of intuition rather than a lot of math. At the same time, of course, the correlation inputs must be credible. We believe the multilayer approach meets all these requirements.

APPENDIX

Geometric Visualization of Risk: Derivation[7]

From statistics, recall the relationship, where B and P are random variables:

$$\sigma^2_{aP+bB} = a^2\sigma^2_P + 2\,rab\,\sigma_P\sigma_B + b^2\sigma^2_B \tag{A-1}$$

In case of a = b = 1 and r = 0, we get the well-known formula for adding variance of uncorrelated variables:

$$\sigma^2_{P+B} = \sigma^2_P + \sigma^2_B \tag{A-1$'$}$$

Further, let's assume a = 1 and b = −1, in which case we get:

$$\sigma^2_{P-B} = \sigma^2_P - 2\,r\,\sigma_P\sigma_B + \sigma^2_B \tag{A-1$''$}$$

And solving for, we find:

$$r = \frac{\sigma^2_P + \sigma^2_B - \sigma^2_{P-B}}{2\sigma_P\sigma_B} \tag{A-1$'''$}$$

In a next step, recall the *cosine rule*. Given the triangle in the appendix exhibit, the cosine rule states:

$$\cos(\varphi) = \frac{P^2 + B^2 - T^2}{2PB} \tag{A-2}$$

Appendix Exhibit

Cosine Rule

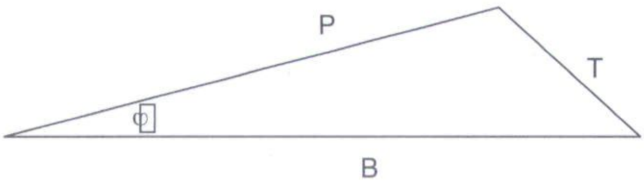

Thinking in vector space and substituting for

$$T = P - B \tag{A-3}$$

in (A-2), we get

$$\cos(\varphi) = \frac{P^2 + B^2 - (P - B)^2}{2PB} \tag{A-2'}$$

The structural identity of (A-1‴) and (A-2′) prompts an important observation: The risks of two portfolios, P and B, and the risk of the difference portfolio between P and B can be represented by a triangle with side lengths proportional to the risks of the portfolios. Then the cosine of each angle equals the correlation between the portfolios as represented by the adjacent sides.

Given three elements of a triangle's three sides and three angles, we can determine the three other pieces. (apart from three angles at the same time, which determines the shape of the triangle but not the length of its sides). If, for instance, the portfolio risk, the benchmark risk, and the correlation between portfolio and benchmark are given, then the size of the tracking error (or relative risk) is given.

ENDNOTES

[1]This is a GARCH(1, 1) estimate, i.e., one lagged ε term and one lagged σ term.

[2]Terhaar, Staub, and Singer [2003] elaborate on the effects of numerical analysis of historical illiquid data.

[3]The correlation between U.S. equity and U.S. bonds has been the subject of an investigation on its own. At its core, there is, again, a factor analysis. For more detail, see Staub [2002].

[4]We always represent a covariance matrix as the combination of a correlation matrix and a column of standard deviations. This is more intuitive.

[5]Essentially, the two-layer concept goes back to Grinold and Kahn [1995], who use it for stock modeling. Terhaar, Staub, and Singer [2003] use a two-layer approach to include illiquid assets.

[6]Because we suggest a relationship between USE and EXE as shown in Exhibit 17 (EXE versus USE*) rather than in Exhibit 16 (EXE to USE).

[7]For geometric visualization of risk, see Zerolis [1996] and Singer, Terhaar, and Zerolis [1998].

REFERENCES

Grinold, Richard C., and Ronald Kahn. *Active Portfolio Management.* Chicago: Probus, 1995.

Singer, Brian, Kevin Terhaar, and John Zerolis. "Maintaining Consistent Global Asset Views (with a Little Help from Euclid)." *Financial Analysts Journal*, January/February 1998, pp. 63–71.

Staub, Renato. "The Correlation between U.S. Equity and U.S. Bonds." Working paper, UBS Global Asset Management, March 2002.

Terhaar, Kevin, Renato Staub, and Brian Singer. "The Appropriate Policy Allocation for Alternative Investments." *The Journal of Portfolio Management*, (2003), pp. 101–110.

Zerolis, John. "Triangulating Risk." *Risk*, Vol. 9, No. 12 (April 1996), pp. 36–38.