OCCAM Financial Technology

Last revised: 5th Oct 1998

OCCAM Application Note 1
Constant Group Models in POW!

A. Introduction to Constant Group Models

A1. The Need for Constant Group Models

This note describes a class of risk models called constant group (CG) models and how the POW! optimiser can used to apply them. Conceptually, they are extremely simple, even crude; in practice, as Elton and Gruber (1973 etc) have demonstrated, they can be surprisingly effective. As we will show, they are really a form of factor model, and like factor models they address the problem that although historic risk data may be one of the best guides we have to the future, it is certainly not perfect.

Imagine that on a given day there happens to be a program trade, involving the purchase of large number of company A shares; and on the same day, for personal reasons, a director of company B decides to sell a large part of his holding. These two independent acts may well give rise to a spurious negative correlation between the shares of A and B. However, since the price-impact of these two trades, and similar microstructural activity, is likely to be reversed within a few days, fund managers can eliminate their effect by looking at the correlation not of daily, but of monthly, returns - in effect averaging the returns on a monthly basis[1].

But now imagine the CEO of transport company A has a heart attack on the same day as oil company B announces that a promising hole off the coast of Greenland is in fact dry: this is likely to cause a simultaneous downward movement in the price of both shares, and in this case one which may not be reversed. So, although the correlation effect is just as spurious as the first example, it will not be eliminated simply by choosing a longer time frame. We do not want just to ignore the correlation, as we would certainly expect there to be a relationship between oil companies and transport companies - in fact a negative correlation. So we need to find some cross-sectional technique, which eliminates individual chances but preserves more general relationships.

[1] Fund managers can use this technique because their time horizon is typically measured in weeks or months rather than days; it will not be open to option hedgers and others for whom daily movements are important.

A2. Constant group Correlation Models

The solution proposed by Elton and Gruber was to replace the correlation between two companies in different sectors, say oil and transport, by the average correlation between every oil company and every transport company, so that instead of

Group		Mining		Transport		Oil	
	Share	A	B	C	D	E	F
Mining	A	1.0	0.8	0.1	0.2	0.3	0.4
	B	0.8	1.0	0.2	0.3	0.7	0.2
Transport	C	0.1	0.2	1.0	0.6	0.1	-0.4
	D	0.2	0.3	0.6	1.0	0.0	-0.5
Oil	E	0.3	0.7	0.1	0.0	1.0	0.7
	F	0.4	0.2	-0.4	-0.5	0.7	1.0

we would use

Group		Mining		Transport		Oil	
	Share	A	B	C	D	E	F
Mining	A	1.0	0.8	0.2	0.2	0.4	0.4
	B	0.8	1.0	0.2	0.2	0.4	0.4
Transport	C	0.2	0.2	1.0	0.6	-0.2	-0.2
	D	0.2	0.2	0.6	1.0	-0.2	-0.2
Oil	E	0.4	0.4	-0.2	-0.2	1.0	0.7
	F	0.4	0.4	-0.2	-0.2	0.7	1.0

Notice that the same principle also applies to companies in the same sector, except that the correlation of each company with itself has to be 1. Taken as it is, the need for special treatment of the diagonal is rather tiresome, because it means that each time we wish to present or perform calculations on the correlation matrix, we have to use the entire matrix. However we can avoid this by separating the diagonal 1s into their industry component and their specific component as follows:

Group		Mining		Transport		Oil	
	Share	A	B	C	D	E	F
Min-ing	A	0.8	0.8	0.2	0.2	0.4	0.4
	B	0.8	0.8	0.2	0.2	0.4	0.4
Trans-port	C	0.2	0.2	0.6	0.6	-0.2	-0.2
	D	0.2	0.2	0.6	0.6	-0.2	-0.2
Oil	E	0.4	0.4	-0.2	-0.2	0.7	0.7
	F	0.4	0.4	-0.2	-0.2	0.7	0.7

+

Mining		Transport		Oil	
A	B	C	D	E	F
0.2					
	0.2				
		0.4			
			0.4		
				0.3	
					0.3

which we can show more compactly as:

Group	Mining	Transport	Oil		Additional correlation on diagonal
Mining	0.8	0.2	0.5		0.2
Transport	0.2	0.6	-0.2	+	0.4
Oil	0.4	-0.2	0.7		0.3

This is how a CG correlation model is represented on the U.CGC sheet of POW!.

For those used to factor models, it may seem a little odd to have anything other than 1 on the group diagonal; to get round this, we can think of each group member as having a beta to its group not of 1 but the square root of the reciprocal of the group correlation. We will then need to adjust the off-diagonals, so that the group grid above would be represented as:

Group		Mining	Transport	Oil
	Beta	1.12	1.29	1.20
Mining	1.12	1.00	0.14	0.30
Transport	1.29	0.14	1.00	-0.13
Oil	1.20	0.30	-0.13	1.00

The user can, of course, insert an extra sheet into his workbook to provide such a presentation, but to avoid a proliferation of matrices this is not supplied automatically.

Similarly, the user can think of the additional correlation as being 1 multiplied by a specific beta squared, where the beta is square root of the correlation. Such a presentation *is* provided in U.Main; this should not be overwritten.

The risks applied to both the group correlations and the additional (specific) correlations to produce the covariance matrix are those the user has entered as specific risks in U. Specific. (The user has the theoretical ability also to enter a full residual correlation matrix in U.Specific; this is not a feature of the standard CG Correlation model, and is not recommended).

A2. <u>Constant Group Covariance Models</u>

It may seem odd that we are applying the same risks to both group and additional (specific) correlations; one might expect there to be (a) a quasi-factor risk and (b) a separate specific risk. In a similar vein, it might be argued that just as historic correlations need to be subjected to some averaging process, so too should risks.

The CG Covariance model addresses these issues by applying the same group-averaging technique as the CG Correlation model, but to the covariance matrix, rather than the correlation matrix. Although the analogy seems at first sight straightforward, there are a few complications which will make it worth going through the process in full.

In the CG Covariance model we replace the covariance between two companies in different sectors, say oil and transport, by the average covariance between every oil company and every transport company, so that instead of

Group	Share	Mining A	Mining B	Transport C	Transport D	Oil E	Oil F
Mining	A	0.1000	0.0092	0.0015	0.0026	0.0042	0.0076
	B	0.0092	0.0132	0.0035	0.0045	0.0113	0.0044
Transport	C	0.0015	0.0035	0.0225	0.0117	0.0021	-0.0114
	D	0.0026	0.0045	0.0117	0.0169	0.0000	-0.0124
Oil	E	0.0042	0.0113	0.0021	0.0000	0.0196	0.0186
	F	0.0076	0.0044	-0.0114	-0.0124	0.0186	0.0361

we would use

Group	Share	Mining A	Mining B	Transport C	Transport D	Oil E	Oil F
Mining	A	0.1000	0.0092	0.0030	0.0030	0.0069	0.0069
	B	0.0092	0.0132	0.0030	0.0030	0.0069	0.0069
Transport	C	0.0030	0.0030	0.0225	0.0117	-0.0054	-0.0054
	D	0.0030	0.0030	0.0117	0.0169	-0.0054	-0.0054
Oil	E	0.0069	0.0069	-0.0054	0.0054	0.0196	0.0186
	F	0.0069	0.0069	-0.0054	-0.0054	0.0186	0.0361

The same principle also applies to companies in the same sector, except that the covariance of each company with itself, i.e. its variance, remains unaltered. As in the CG Correlation model, the need for special treatment of the diagonal is rather tiresome, because it means that each time we wish to present or perform calculations on the matrix, we have to use it in its entirety. However we can avoid this by separating the diagonal variances into their industry component and their specific component as follows:

Group	Share	Mining A	Mining B	Transport C	Transport D	Oil E	Oil F
Mining	A	.0092	.0092	.0030	.0030	.0069	.0069
	B	.0092	.0092	.0030	.0030	.0069	.0069
Trans-	C	.0030	.0030	.0117	.0117	-.0054	-.0054
port	D	.0030	.0030	.0117	.0117	-.0054	-.0054
Oil	E	.0069	.0069	-.0054	-.0054	.0186	.0186
	F	.0069	.0069	-.0054	-.0054	.0186	.0186

+

Mining A	Mining B	Transport C	Transport D	Oil E	Oil F
.0008					
	.0040				
		.0108			
			.0052		
				.0010	
					.0175

which we can show more compactly as:

Group	Mining	Transport	Oil
Mining	0.0092	0.0030	0.0069
Transport	0.0030	0.0117	-0.0054
Oil	0.0069	-0.0054	0.0186

+

Additional covariance on diagonal	
A: 0.0008	B: 0.0040
C: 0.0108	D: 0.0052
E: 0.0010	F: 0.0175

The left hand side of the table is how the group covariance matrix is represented in U.CGC. In addition the implied "group risk", that is the square root of the compacted matrix's diagonal, is shown on the left of the matrix as "implied SD on group diagonal"; it should not be overwritten.

The implied additional covariance on the diagonal is not presented by POW! as such, but converted in U.Main to a beta to specific risk. The latter is entered in U.Specific as the full specific risk of the asset in question, as in column 3 below, and the betas in our example would accordingly appear in U.Main as in column 6, where they should not be overwritten.

1	2	3	4	5	6
Share	Full covariance on diagonal	Full Specific risk	Group covariance on diagonal	Additional covariance on diagonal	Implied beta to full specific risk
A	0.1000	10.0%	0.0092	0.0008	0.28
B	0.0132	11.5%	0.0092	0.0040	0.55
C	0.0225	15.0%	0.0117	0.0108	0.69
D	0.0169	13.0%	0.0117	0.0052	0.55
E	0.0196	14.0%	0.0186	0.0010	0.23
F	0.0361	19.0%	0.0186	0.0175	0.70

The user has the theoretical ability also to enter a full residual correlation matrix in U.Specific; this is not a feature of the standard CG Covariance, and is not recommended.

The user should ensure that the full specific risk of an asset is never less than the square root of the group variance; otherwise there will negative additional covariance.

The procedure described, with some provisions to avoid negative additional covariances, is that used in the QUANTEC IRAS single country models.

A4. CG Returns

POW! allows group returns to be entered. These are effectively factor returns where the factors are group memberships with all betas equal to one.

A5. Possible Extensions

The averaging procedure commonly used in estimating CG models has the laudable quasi-Bayesian effect of damping the specific asset data towards the cross-sectional mean; but in doing so it reduces the off-diagonal estimate of portfolio risk, while preserving the on-diagonal risk.

The consequences of this and their severity will vary according to the type of fund and benchmark under consideration; but they can be most clearly understood in the context of a full replication index fund, which should have zero risk relative to its benchmark, with which by construction it is identical. The off-diagonal risk should exactly offset the on-diagonal risk, but because of their differential treatment by the averaging procedure model there will be a discrepancy, with the result that the fund may appear to have a significant level of tracking error. Rice (1992) shows that in the CG correlation model this can be avoided by replacing the simple average correlation by an average weighted by the index weights times the relevant standard deviations. This presupposes that the index and the benchmark are one and the same; the technique will fail to the extent that the two are different.

There is no reason why the constant group structure should not be used for correlations or covariances that are not derived directly or wholly from an averaging process of this kind. One problem with simple averaging, whether weighted or otherwise, is that one may throw away more of the asset-specific information than one would wish; so why not use say 50% of the asset-specific correlations and 50% of the averaged correlations ? If you like the idea, but feel that 50% is a bit unscientific, then Ledoit (1997) may give you some ideas!

Facilities of this kind will be available in the next version of POW! In the meantime, once you have mastered the basic set-up described in sections B and C, you may wish to experiment by setting the switches provided at the head of the relevant columns in rows 13 and 14 of U.Main to something other than the default value of 1; if the rows are hidden, first select the adjacent row numbers and do Format Row Unhide.

B. Applying POW! to Constant Group Correlation Models

This section explains how to use POW! to set up a CG Correlation model, to the extent that the procedure differs from that for other factor models; If you are not familiar with factor models in POW!, you may find it helpful to refer to the User Manual first. Setting up a CG Covariance Model is described in section C.

B1. Set up Model in Dialog Box 2

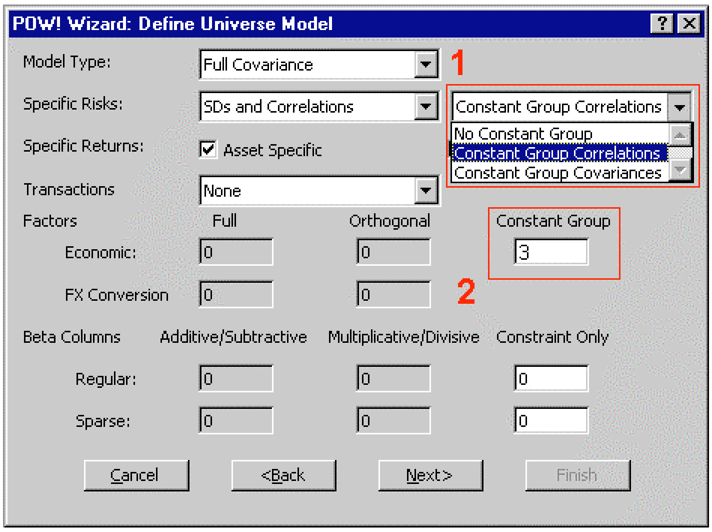

1. Select the Model Type. If the CG process is being applied to the total covariance matrix, select Full Covariance; if to residuals only, select Custom Factor.

2. Select the Asset Risks option: usually SDs only for CG models.

3. Select Constant Group Correlations from the drop-down menu on the right (box 1 in the figure opposite).

4. Check Asset Specific and/or Group Specific Returns (box 3) as needed.

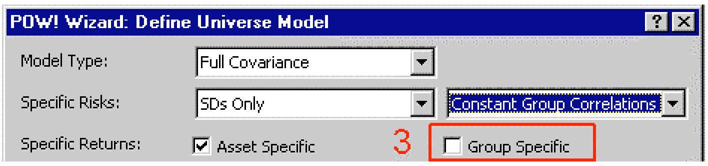

5. Specify the number of groups (box 2).

6. Specify other parts of the model in the usual way.

B2. Enter the CG Names and Correlations in U.CGC

The POW! Wizard creates a worksheet called U.CGC (see figure below) in which the CG names and correlations are entered. The layout of U.CGC is similar to that of U.Factor with the dimension of the risk matrix equal to the number of groups.

	A	B	C	D	E	F	G
3	Report:	Constant Group Risks and Correlations					
4	Date:	September 1998					
5							
6							
7		Return	Additional correl	Correlations between group members			
8			on asset diag.				
9				Mining	Transport	Chemicals	
10	Mining	N/A	0.20	0.80	0.20	0.50	
11	Transport	N/A	0.40	0.20	0.60	-0.10	
12	Chemicals	N/A	0.30	0.50	-0.10	0.70	

The figure shows a three group constant correlation matrix. Note that the correlation of each asset with itself must be 1, so the additional correlation on the asset diagonal is automatically set to 1 minus the relevant diagonal group correlation; it should not be overwritten.

Column B could contain group specific returns, but they are not used in this model so the cells are disabled.

B3. Set Group Membership for the Assets in U.Main

Group membership is set using the same mechanism as for sparse factors. When a CG model is created, the POW! Wizard will create a column in the U.Main sheet, to the left of any other factors or constraints, so that you can specify the group membership of the assets.

Assuming you have already entered names for the groups in the U.CGC sheet, double-clicking on one of the pink cells in column D will bring up a helper listing the names of the groups to which the asset can be assigned. The group beta in column E should normally be set to one.

The specific return betas in column B are fixed at 1. The specific risk betas in column C are automatically set to the square root of the addition correlations shown in the U.CGC sheet; they should not be overwritten.

B4. Factor Decomposition in CG Models

For purposes of factor risk and return decomposition, groups can be treated like other factors. When the POW! Wizard creates a factor decomposition worksheet, it will create a line or column for each of the groups.

C. Applying POW! to Constant Group Covariance Models

This section explains how to use POW! to set up a CG Covariance model, to the extent that the procedure differs from that for other factor models. If you are not familiar with factor models in POW!, you may find it helpful to refer to the User Manual first. Setting up a CG Correlation Model is described in section B.

C1. Set up the Model in Dialog Box 2

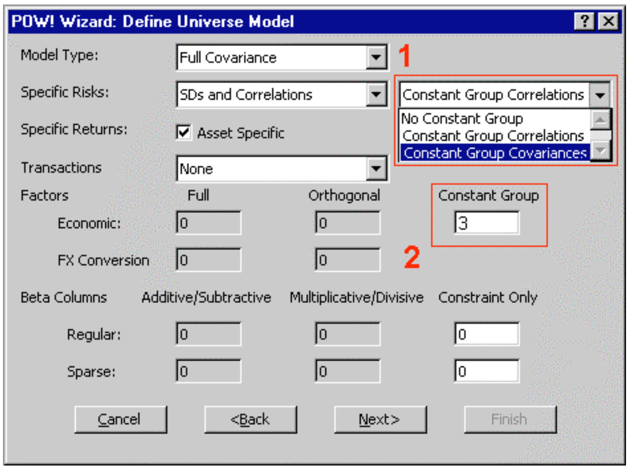

1. Select the Model Type. If the CG process is being applied to the total covariance matrix, select Full Covariance; if to residuals only, select Custom Factor.

2. Select the Asset Risks option: usually SDs only for CG models.

3. Select Constant Group Covariances from the drop-down menu on the right (box 1 in the figure opposite).

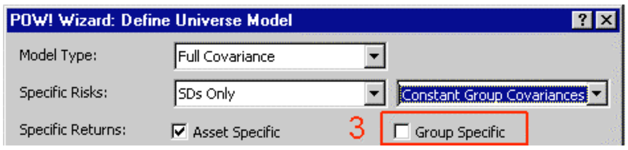

4. Check Asset Specific and/or Group Specific Returns (box 3) as needed

5. Specify the number of groups (box 2).

6. Specify other parts of the model in the usual way.

C2. Enter the CG Names and Covariances in U.CGC

The POW! Wizard creates a worksheet called U.CGC (see figure below) in which the CG names and covariances are entered. The layout of U.CGC is similar to that of U.Factor with the dimension of the risk matrix equal to the number of groups.

	A	B	C	D	E	F	G
3	Report:	Constant Group Covariances					
4	Date:	September 1998					
5							
6							
7		Return	Implied SD	Covariances between group members			
8			on group diag.				
9				Mining	Transport	Chemicals	
10	Mining	N/A	10.95%	0.0120	0.0020	0.0005	
11	Transport	N/A	20.25%	0.0020	0.0410	-0.0001	
12	Chemicals	N/A	14.56%	0.0005	-0.0001	0.0212	

The figure shows a three group constant covariance matrix. Column C shows, for information only, the implied SD of each group, ie the square root of the figure appearing on the diagonal of the covariance matrix. It should not be overwritten.

Column B could contain group returns, but they are not used in this model so the cells are disabled.

C3. Set Group Membership for the Assets in U.Main

Group membership is set using the same mechanism as for sparse factors. When a CG model is created, the POW! Wizard will create a column in the U.Main sheet, to the left of any other factors or constraints so that you can specify the group membership of the assets.

Assuming you have already entered names for the groups in the U.CGC sheet, double-clicking on one of the pink cells in column D will bring up a helper listing the names of the groups to which the asset can be assigned. The group beta in column E should normally be set to one.

	A	B	C	D	E	F	G
3	Report:	Assets					
4	Date:	September 1998				Currency:	Not Given
5							
6							
7		Specific	Specific	Constant Group		Constraint	
8		Return	Risk				
9				CONST.GRP		G.CONSTR	
10	OPERATOR	+	+	+		N/A	
15				Factor	Beta	Factor	Beta
16	A	1.00	0.45	Mining	1.0	Constraint	1.0
17	B	1.00	0.45	Mining	1.0	Constraint	1.0
18	C	1.00	0.71	Transport	1.0	Constraint	1.0
19	D	1.00	0.71	Transport	1.0	Constraint	1.0
20	E	1.00	0.55	Chemicals	1.0	Constraint	1.0
21	F	1.00	0.55	Chemicals	1.0	Constraint	1.0

The specific return betas in column B are fixed at 1. The specific risk betas in column C are automatically set to a value which ensures that when reassembled from its group and non-group components, the specific risk used by the optimiser still adds up to the specific risk originally entered in U.Specific. These betas should not be overwritten.

C4. Factor Decomposition in Constant Group Models

For purposes of factor risk and return decomposition, constant groups can be treated like other factors. When the POW! Wizard creates a factor decomposition worksheet, it will create a line or column for each of the groups.

D. Bibliography

Elton, Edwin J. and Gruber, Martin J.(1973): *Estimating the Dependence Structure of Share Prices: Implications for Portfolio Selection*, JOF VIII no.5 (Dec 1973), pp 1203-1232
- presents tests of results from CG correlation models

Elton, Edwin J. and Gruber, Martin J.(1995): *The Correlation Structure of Security Returns: Multi-Index Models and Grouping Technique*, Chapter 8 of 5th edition of Modern Portfolio Theory and Investment Analysis, Chapter 6 of 3rd edition, Wiley
- the chapter gives a useful balanced discussion, and the bibliography at the end contains many useful references, especially to Elton and Gruber's own articles

Ledoit, Olivier (1997): *Improved estimation of the covariance matrix of stock returns with an application to portfolio selection*, Anderson School UCLA Working Paper #6-97
- proposes the estimation of the covariance matrix by taking an optimally weighted average of the single-index covariance matrix and the sample covariance matrix

Rice, Robert (1992): *Towards a General Model of Stock Returns*, QUANTEC Quarterly 1992 Q1, esp. pp23f
- suggests a technique for adjusting a CG matrix, to improve the accuracy of the estimate of portfolio risk