Orthogonalized Regressors and Spurious Precision, with an Application to Currency Exposures^{*}

Fang Liu[†], Piet Sercu[‡] and Martina Vandebroek[§]

November 29, 2014

*The authors thanks Christophe Croux, Geert Dhaene, Giovanni Urga, and participants at presentations in Leuven, Brussels, Chicago, and the University of Wisconsin for useful discussions, but accept full responsibility for any remaining errors.

[†]Central University of Finance and Economics, Beijing; liufang8000@aliyun.com

[‡]KU Leuven, Faculty of Economics and Business; piet.sercu@kuleuven.be.

[§]KU Leuven, Faculty of Economics and Business; martina.vandebroek@kuleuven.be.

Abstract

Regressions often use pre-orthogonalized regressors: prior to the main regression, an independent variable x_i is regressed upon the other regressor(s), and its residuals are used in the right-hand side of the main regression instead of the raw variable itself. For example, the exposure of a stock's return to exchange rate changes is conventionally estimated by a regression, and often the market return is included as an additional regressor. By first orthogonalizing the market return on the exchange rate, in a regression separate from the main one, one seems to have the best of both worlds: the market factor cannot subsume part of the exposure present in a stock's return, and the standard error (SE) of the estimate beats both the simple- and the multiple-regression SE's. This last effect is illusory: since the simple regression and its two-step variant, with the orthogonalization, produce the same exposure estimate, given the sample, their precision must be identical too. Technically, the source of the problem is that the uncertainty about the market's exposure estimate is left out of the calculated SE. In published work, the calculated error variances should be corrected upward by 20 to 100 percent.

Keywords: Financial econometrics, international finance, market model, currency exposure, orthogonal. JEL-codes: C3, C58.

Introduction

To assess the exposure of a stock's return to exchange-rate changes, one often removes from the market return the variation that is common with the exchange-rate changes. In fact, such a preorthogonalization step is far from uncommon in other contexts too: one is interested in some y's exposure to some x, and a regressor z' is added which is just the raw z orthogonalized on x. One objective may be to avoid the inflated standard errors (SES) in case of near-multicollinearity, but that is not really a problem in the exposure literature. Another consideration behind orthogonalization is that, this way, the exposure estimate is the same as if z were totally absent, while its reported standard error (SE) is still tightened because the addition of z' reduces the residual variance by as much as z would have done.

In both applications, shrinking the standard errors is the objective. However, the idea that adding a pre-orthogonalized regressor tightens the SE is an illusion, as the uncertainty of the orthogonalization vector—for instance, the market's currency exposure—is left out of the calculated SE. The main message, which holds in general, is: pre-orthogonalization should be avoided because it invalidates the usual apparatus of inference. Reported SEs are too small, yielding too many parameters that look significant.

In the remainder of the Introduction, we discuss the context and arguments in more detail.

In the currency exposure literature,¹ the practice of orthogonalization arose as follows. A stock's currency exposure is measured by the coefficients of the slope of the regression of the stock's return on the percentage changes in the exchange rates (Stein, 1960; Johnson, 1960; Dumas, 1978). After Jorion (1990), it is common practice to add the market return as an additional regressor. This reflects the 'market model' regression's standing as a returngenerating process in financial studies. One additional benefit from adding this regressor is that it lowers the residual variance, which, everything else remaining the same, reduces the standard error of the exposure estimate. On the downside, there may be some correlation between the market and exchange factors, and this multicollinearity could even undo the power gain from the reduced residual variance. But any correlation between the two regressors also affects the coefficient itself, not just its standard error: if the market return r_m itself is also exposed to the exchange return, the term $\beta_j r_m$ will already pick up part of the stock's

¹For a review of the theoretical and empirical literature, see, e.g., Boudt, Liu and Sercu (2013).

total currency exposure, leaving only a residual exposure to be captured by the coefficient for the exchange-rate regressor.² This way, currency exposure at the market level could kill the chances of finding convincing stock-specific currency effects, thus incorrectly suggesting low or zero exposure. Obtaining a lower residual noise without giving the market return the chance to subsume the individual currency effects is an attractive prospect. Our message is that the drop in the estimator's standard error is illusory: in reality the estimate has the same SE as the one from a simple regression.

Section 1 presents the analytical arguments. Section 2 presents the results of some Monte Carlo simulations. Section 3 presents some real-world results. The first of our two samples is a set of U.S. multinationals, similar to the firms studied in the seminal Jorion paper; we study it at daily and monthly frequencies. The second data set refers to Chinese exporters; these firms are much smaller and less diversified, cannot hedge forward, and are therefore clearly exposed. The two samples are nevertheless not very different in terms of the extent to which the market return subsumes the currency factor, the size of the raw exposure, and the amount of bias that is introduced by pre-orthogonalizing. We draw some conclusions in Section 4.

1 The effects of orthogonalizing one regressor

One needs to distinguish two cases. Some authors pre-orthogonalize the market return (the control variable) on the exchange rate (the variable of interest), while others orthogonalize the exchange rate on the market return. The first route has been adopted by, e.g., Allayannis (1996), Bartram and Bodnar (2007), Bodnar and Wong (2003), Bris and Koskinen (2002), Entorf, Moebert and Sonderhof (2006), Griffin and Stulz (2001), Kiymaz (2003), Priestley and Odegaard (2002), and Pritamani, Shome and Singa (2004). In contrast, Jorion (1990, 1991), followed by, e.g., Elton and Gruber (1991), Choi and Prasad (1995), Glaum, Brunner and Himmel (2000), and Hagelin and Pamborg (2002), orthogonalize the exchange rate on the market return. We show that (i) for the untransformed regressor, the standard error (SE) is underestimated and (ii) for the orthogonalized regressor, neither the coefficient nor the SE are affected. That is, if one's interest is in the currency exposure, then orthogonalizing the market return on the exchange rate leads to flawed inference tools, while the reverse procedure neither harms nor helps.

²See Section 1.1 if this does not sound familiar.

In the first strand of this literature, the relevant regressions are the following. We consider just one foreign currency, whose percentage changes are denoted by s; the generalization to multiple currencies is simple. Denote the simple returns on stock j and on the market index m by R_j and R_m , respectively. The (co)variances of the regressors are denoted by σ_s^2 , σ_m^2 and $\sigma_{m,s}$. We denote the correlation between R_m and s by $\rho_{m,s}$ and the conditional variance of R_m given s by $\sigma_{m|s}^2 = \sigma_m^2(1 - \rho_{m,s}^2)$. Lastly, we add hats whenever we refer to estimates rather than population values. All estimations are from time series. The market model regression is

Market Model:
$$R_j = \alpha_{0,j} + \beta_{0,j}R_m + u_{0,j}$$
. (1)

The Dumas regression defines the relative currency exposure, γ_1 , similarly:

Dumas regression:
$$R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j},$$
 (2)

$$R_m = \alpha_{1,m} + \gamma_{1,m} s + u_{1,m}.$$
(3)

The standard form for the equation nesting the Dumas and market models (1)-(2) is

Multivariate Model:
$$R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}.$$
 (4)

To obtain more precise estimates for γ , one often uses a hybrid version where the market regressor is replaced by $u_{1,m}$ from Equation (3). That is, one removes from R_m the variation that is common with s:

Hybrid Model:
$$R_j = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}.$$
 (5)

In the next section we show that if R_m is positively correlated with s, then $\gamma_{2,j}$ is below $\gamma_{1,j}$ and has a zero average across stocks.

1.1 The effect on exposure of adding a market regressor

When stocks tend to have similar exposures—for instance, in markets with many exporting firms listed, such as China, Japan, Germany, and the Netherlands—the market portfolio is exposed to the currency factor too. If one then runs a regression that includes both the exchange-rate change and the market return, the latter regressor is subsuming part of the stock's raw exposure. To see this, consider the familiar solution for a two-regressor slope coefficient:

$$\gamma_{2,j} = \frac{\gamma_{1,j} - \beta_{0,j}\gamma_{1,m}}{1 - \rho_{s,m}^2},\tag{6}$$



Key The multivariate exposure, $\gamma_{2,j} = [\gamma_{1,j} - \beta_{0,j}\gamma_{1,m}]/[1 - \rho_{s,m}^2]$, is plotted for various levels of beta (0.75, 1.00, 1.25 or 1.50) and $\rho_{s,m}$, at $\sigma_s = 0.1$, $\sigma_{r_m} = 0.2$ and $\gamma_{1,j} = 1$. A rising $\rho_{s,m}$ affects the market's $\gamma_{1,m}$ in the numerator and directly affects the denominator too, hence the slight non-linearity. The market's simple gamma, not shown, would be a positively sloped ray passing through the point (x = 0.5, y = 1).

Figure 1: The effect of $corr(R_m, s)$ on the multivariate currency exposure coefficient

where $\rho_{s,m}$ denotes the correlation between r_m and s. So if all stocks have similar Dumas exposures $\gamma_{1,j}$, the market's exposure $\gamma_{1,m}$ must be in the same range, and since betas are positive, the multivariate gamma will be lowered. Figure 1 illustrates the subsumption effect at the level of the individual stock as a function of the correlation between s and R_m , which affects $\gamma_{2,j}$ directly via the denominator and indirectly via $\gamma_{1,m}$. The curve is plotted for various levels of beta (0.75, 1.00, 1.25 or 1.50); the other parameters were set at $\sigma_s = 0.1$, $\sigma_{r_m} = 0.2$ and $\gamma_{1,j} = 1$. Predictably, the subsumption effect increase the higher the stock's beta and the market's exposure.

Being an estimate of exposure not present in the market return, the multivariate gamma is useless for detecting average exposures: whether that average is substantially positive or negative, the addition of R_m will generate about 50/50 positive/negative γ_2 coefficients for the individual stocks. This follows from Equation (6): denoting value weights by w_j , the weighted average multivariate exposure works out to be

$$\sum_{j} w_{j} \gamma_{2,j} = \frac{\left[\sum_{j} w_{j} \gamma_{1,j}\right] - \left[\sum_{j} w_{j} \beta_{0,j}\right] \gamma_{1,m}}{1 - \rho_{s,m}^{2}},$$

Orthogonalized regressors and spurious precision

$$= \frac{\gamma_{1,m} - 1 \times \gamma_{1,m}}{1 - \rho_{s,m}^2} = 0.$$
(7)

So this 50/50 sign distribution, in value-weighted terms, is a construct and has nothing to do with exports and imports roughly evening out at the aggregate level, as is sometimes claimed.

If the researcher's objective is to establish the existence of exposure, the above subsumption effect is not welcome. One solution is to not include the market return as a regressor, but often a compromise is sought by pre-orthogonalizing the market return on the currency returns instead.

1.2 The effects of first orthogonalizing the market regressor

An orthogonalization is just one particular linear transformation, and its effects are easily traced. Let us first assume we know the population moments, and consider the matrix

$$\mathbf{G} = \begin{bmatrix} 1 & -\gamma_{1,m} \\ 0 & 1 \end{bmatrix} \Rightarrow \mathbf{G}^{-1} = \begin{bmatrix} 1 & \gamma_{1,m} \\ 0 & 1 \end{bmatrix}.$$
 (8)

Without loss of generality, we may assume the variables to be demeaned; that is, there are no intercepts. If we write the multiple regression (4) in matrix form as $R_j = \mathbf{X}'_2 \mathbf{B}_2 + u_{2,j}$ where $\mathbf{X}_2 := (s, R_m)'$ and $\mathbf{B}_2 := (\gamma_{2,j}, \beta_{2,j})'$, we can insert $\mathbf{G}\mathbf{G}^{-1}$ and reinterpret:

$$R_j = \mathbf{X}_2' \mathbf{B}_2 + u_{2,j} = [\mathbf{X}_2' \mathbf{G}] [\mathbf{G}^{-1} \mathbf{B}_2] + u_{2,j}.$$
 (9)

The item $\mathbf{X}'_2 \mathbf{G}$ corresponds to the orthogonalized regressors of Equation (5), and matching it is an offsetting transformation in the regression coefficients: in the hybrid regression, the slopes are given by

$$\mathbf{B}_3 = \mathbf{G}^{-1} \mathbf{B}_2. \tag{10}$$

As proven in Technical Note A, it follows that regression (5) produces the same market beta as the regular multivariate regression (4) and, as intended, the same currency exposure as the simple Dumas equation (2). This, indeed, is why we call it a hybrid. Formally, we have

$$\begin{bmatrix} \gamma_{3,j} \\ \beta_{3,j} \end{bmatrix} = \begin{bmatrix} \gamma_{1,j} \\ \beta_{2,j} \end{bmatrix}.$$
(11)

All the above refers to population parameters and moments, but the OLS estimators are of the method-of-moments class: to compute them, one just replaces the population moments by the sample moments. Below, we use hats to denote sample estimates instead of population values. So also in every possible sample the exposure estimate $\hat{\gamma}_{j,3}$ of the hybrid regression (5) is identical to the one in the simple regression (2), $\hat{\gamma}_{1,j}$. It follows, logically, that these

equation to be estimated	$\hat{\gamma}$	conventional variance of estimate
(2) $R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j}$	$rac{\hat{\sigma}_{j,s}}{\hat{\sigma}_s^2}$	$\frac{\hat{\sigma}^2(u_{1,j})}{\sum_t (s_{j,t} - \overline{s})^2} = \frac{\hat{\sigma}_{j s,m}^2 + \beta_{3,j}^2 \hat{\sigma}_{m s}^2}{\sum_t (s_{j,t} - \overline{s})^2}$
(4) $R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$	$\frac{\hat{\gamma}_{1,j}-\hat{\beta}_{0,j}\hat{\gamma}_{1,m}}{1-\rho_{s,m}^2}$	$\frac{\hat{\sigma}_{j s,m}^2}{\sum_t (s_{j,t}-\overline{s})^2(1-\hat{\rho}_{s,m}^2)}$
(5) $R_j = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}$	$rac{\hat{\sigma}_{j,s}}{\hat{\sigma}_s^2}$	$\frac{\hat{\sigma}_{j s,m}^2}{\sum_t (s_{j,t}-\overline{s})^2}$

Key The table shows the three possible regressions, the OLS estimator of the corresponding γ coefficient, and the variance of the estimation error for each of these as implemented in standard software. In the first two cases, that implementation is consistent, in the last one it is not.

Table 1: The estimators for γ_1 , γ_2 and γ_3 and their standard errors

estimators should have the same standard error. Similarly, logic predicts that the hybrid and the regular multiple regressions, (5) and (4), should produce the same standard error for the market sensitivity.

The problem is that while the second requirement is met, the first seems not to hold if one adopts the standard methods. The estimators and conventional sampling errors for each gamma are given in Table 1, with the proofs relegated to Technical Note B. While the SE in the second line is hard to rank relative to the one in the first line, the third is unambiguously smaller. It actually seems that one can eat one's cake and have it: when going from the simple regression (2) to the hybrid (5), the exposure γ_j is not affected, but since the market return typically explains more of the stock return than does the exchange rate change and is far from perfectly correlated with it: the R^2 after adding the orthogonalized market return is substantially higher and the SE correspondingly lower. Our proposition is that the gain in precision is correct only with fixed regressors; conversely, whenever the regressors are random variables (like ours), the drop in the SEs is an illusion.

To see this, recall that much of basic regression theory starts with non-stochastic regressors: the treatments are assumed to be deliberately chosen by the researcher. As is well known, in the case of random regressors, most of that basic theory can be salvaged via an interim step: conditional on the observed values of the regressors, the estimator is unbiased, and the SE of a multiple regression coefficient is correct provided one knows the noise variance. Unconditionally, we need to realise that the slope estimates are only consistent: even though the estimated moments of the regressors are individually unbiased, the slope estimate uses a non-linear combination of sample moments, so the unbiasedness of the result is only obtained in an infinite sample, where all estimation errors disappear.³ For the SE, the same problem arises: the expression is a non-linear function of imperfect estimates, and the resulting bias disappears only in infinite samples. This is well known: while with random regressors every small-sample estimated SE would be noisy and biased, there is no systematic and strong bias in large samples. Our proposition is that the latter result no longer holds in the case of sample-based transformations of regressors, including orthogonalization: in our case, the conventional estimate of the SE is always biased downward, even in large samples, and often substantially so.

To see the importance of the deterministic versus random regressors, compare the simple regression (2) and the hybrid, (5). If, in (3), R_m and s are non-random, then so are $\gamma_{1,m}$ and $u_{1,m}$. The hybrid regression would then correctly set the SE lower than the simple regression: if, in the first residual $u_{1,j} = \beta_{3,j}u_{1,m} + u_{3,j}$, the first term is not random, then it cannot really contribute to small-sample covariances between the residual and s: so it cannot contribute to the SE. So the SE of the simple regression would overstate the uncertainty. The hybrid regression would avoid that mistake: so dropping the term $\beta_{2,j}^2 \hat{\sigma}_{m|s}^2$ if we go from the first line of Table 1 to the last would be fully justified.

If all regressors become random, however, both terms on the right hand side of $u_{1,j} = \beta_{3,j}u_{1,m} + u_{3,j}$ vary across samples. Also $\beta_{3,j}u_{1,m}$ then contributes to small-sample covariances, and dropping them as a source of estimation errors would be wrong. Estimation errors in the numerator of $\hat{\sigma}_{j,s}/\hat{\sigma}_s^2$ now reflect small-sample covariances between s and the entire $u_{1,j}$, irrespective of whether they stem from R_m or from $u_{3,j}$. In short, with fixed regressors, the hybrid regression delivers asymptotically correct SEs, but with random regressors the correct estimate of the sampling error is the one from the simple regression.

The above argument does not quite say why the case of orthogonalized regressors differs from a *bona fide* regression and where exactly the standard treatment goes wrong as a result. To pinpoint exactly where the hybrid regression goes wrong, we add hats to our earlier results to denote estimators, and we consider the variance–covariance matrix of the estimates $\hat{\mathbf{B}}_3$ from (5), denoted $\mathbf{V}(\hat{\mathbf{B}}_3)$. In our discussion here, we focus on the implication of $u_{1,m}$'s being random, namely that $\hat{\gamma}_{1,m}$ is random too. First trace the link between $\mathbf{V}(\hat{\mathbf{B}}_3)$ and its counterpart in the standard regression (4) conditional on the regressors \mathbf{X}_2 . For given \mathbf{X}_2 , \mathbf{G} is non-random and

³A sum of unbiased estimates is an unbiased estimate of the sum, but the product or ratio of two unbiased estimators is not an unbiased estimator of the ratio, and similarly for products and other non-linear combinations.

can therefore be taken out of $\mathbf{V}(\hat{\mathbf{G}}^{-1}\hat{\mathbf{B}}|\mathbf{X}_2)$. It then suffices to fill out the variance–covariance matrix of $\hat{\mathbf{B}}_2$, and simplify:⁴

$$\mathbf{V}(\hat{\mathbf{B}}_{3}|\mathbf{X}) = \mathbf{V}(\hat{\mathbf{G}}^{-1}\hat{\mathbf{B}}_{2}|\mathbf{X}),
= \hat{\mathbf{G}}^{-1}\mathbf{V}(\hat{\mathbf{B}}_{2}|\mathbf{X})[\hat{\mathbf{G}}^{-1}]',$$

$$= \begin{bmatrix} 1 & \hat{\gamma}_{1,m} \\ 0 & 1 \end{bmatrix} \frac{1}{N-1} \begin{bmatrix} \hat{\sigma}_{s}^{2} & \hat{\sigma}_{m,s} \\ \hat{\sigma}_{m,s} & \hat{\sigma}_{m}^{2} \end{bmatrix}^{-1} \begin{bmatrix} 1 & 0 \\ \hat{\gamma}_{1,m} & 1 \end{bmatrix} \hat{\sigma}_{j|s,m}^{2},
= \frac{1}{N-1} \begin{bmatrix} \hat{\sigma}_{s}^{-2} & 0 \\ -\hat{\gamma}_{1,m}\hat{\sigma}_{m|s}^{-2} & \hat{\sigma}_{m|s}^{-2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ \hat{\gamma}_{1,m} & 1 \end{bmatrix} \hat{\sigma}_{j|s,m}^{2},
= \frac{1}{N-1} \begin{bmatrix} \hat{\sigma}_{s}^{-2} & 0 \\ 0 & \hat{\sigma}_{m|s}^{-2} \end{bmatrix} \hat{\sigma}_{j|s,m}^{2},
= \frac{1}{N-1} \begin{bmatrix} \hat{\sigma}_{s}^{2} & 0 \\ 0 & \hat{\sigma}_{m|s}^{2} \end{bmatrix}^{-1} \hat{\sigma}_{j|s,u_{m}}^{2}.$$
(13)

This indeed is the variance–covariance matrix of the estimates of the regression of R_j on sand the orthogonalized market return. Reversing the logic, we conclude that the conventional standard errors of the hybrid regression (5) are just those of the standard regression (4) with a linear transformation **G** applied to it as if **G** were deterministic. This is fine as long as we have in mind SEs conditional on s and R_m . But the usual next step fails: the above is not an unbiased estimate of the unconditional SE. When **X**, and therefore $\hat{\mathbf{G}}$, are random, $\hat{\mathbf{G}}$ can no longer be factored out of $\mathbf{V}(\hat{\mathbf{B}}_3)$ as we did in Equation (12). By nevertheless doing so, the standard regression procedure omits many items in the SE. In a regular regression, the only problem with random regressors is that all moments become imperfect estimates, not just those involving the epsilons. That is not a minor issue, but at least all relevant moments are present. In our case, the extra problem is that many relevant moments are completely missing from the computations of the SE. And the mistake occurs because the estimation of $\gamma_{1,m}$ is made in a separate step, so that the uncertainty about the estimate is in no way accounted for in the main regression.

To see the extent of the problem, note that from Equation (11) and the method-of-moments property of OLS, the sample slope coefficients satisfy $\hat{\beta}_{3,j} = \hat{\beta}_{2,j}$ and $\hat{\gamma}_{3,j} = \hat{\gamma}_{2,j} + \hat{\beta}_{2,j} \hat{\gamma}_{1,m}$, with the hat over $\gamma_{1,m}$ showing that this is just an estimate from the sample at hand. A proper standard error for $\hat{\gamma}_{3,j}$ takes into account that all of the three coefficients are now estimated with (probably correlated) errors, but OLS as implemented after orthogonalization ignores that: the researcher implicitly assumes that, in other samples, $\gamma_{1,m}$ will somehow never change. To

⁴To derive the fourth line from the third, follow Technical Note A until the last step but one.

see what should have been done, we write each estimate as the true value plus an error e, and then extract an expression for the estimation noise in $\gamma_{3,j}$. Again invoking, in line two, Equation (11), we get

$$\hat{\gamma}_{3,j} = (\gamma_{2,j} + e_{\gamma_{2,j}}) + (\beta_{2,j} + e_{\beta_{2,j}})(\gamma_{1,m} + e_{\gamma_{1,m}});$$

$$\Rightarrow \hat{\gamma}_{3,j} - [\gamma_{2,j} + \beta_{2,j}\gamma_{1,m}] = e_{\gamma_{2,j}} + \gamma_{1,m}e_{\beta_{2,j}} + \beta_{2,j}e_{\gamma_{1,m}} + e_{\beta_{2,j}}e_{\gamma_{1,m}}.$$
(14)

Standard OLS software sets $e_{\gamma_{1,m}}$ equal to zero, though, and accordingly deems the error to be just $e_{\gamma_{2,j}} + \gamma_{1,m}e_{\beta_{2,j}}$, thus ignoring the remaining part in the error, $\beta_{2,j}e_{\gamma_{1,m}} + e_{\beta_{2,j}}e_{\gamma_{1,m}}$.⁵

For the SE to be correct in large samples, all sixteen elements in the variance of this sum should be present in the calculation. But standard software calculates the SE on the basis of just four of these sixteen elements. Specifically, the commonly calculated total error variance ignores not only the variance of $\beta_{2,j} e_{\gamma_{1,m}}$, but also two times its covariances with the first two error terms on the right-hand side of the above equation, two times three covariances of the last term with the three preceding (six third moments, that is), and the variance of the last term (a fourth moment). Not surprisingly, then, the calculated SE is systematically inconsistent. A large sample will not solve the problem. To see this, note from Table 1 that the ratio of two error variances, one flawed and the other consistent, is

variance ratio :=
$$\frac{\hat{\sigma}_{j|s,m}^2 + \beta_{3,j}^2 \hat{\sigma}_{m|s}^2}{\hat{\sigma}_{j|s,m}^2}.$$
 (15)

In an infinite sample, the sample estimates converge to the population values, but as the latter are all non-zero, the sample counterparts do not disappear.

2 Monte Carlo simulations of the orthogonalization issue

The Monte Carlo simulations in Table 2 illustrate all of the above. In each set of simulations, we generated one million series of 50 records $\{R_j, R_m, s\}$ each, as follows. The independent random variables are s, $u_{1,m}$ and $u_{2,j}$. From these, we generated sample values for the other variables, namely $R_m = \gamma_{1,m}s + u_{1,m}$ and $R_j = \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$, for subsequent use in sample regressions. To get the sample–orthogonalized noise term, lastly, we calculated the sample coefficients in $R_m = \hat{\alpha}_{1,m} + \hat{\gamma}_{1,m} s + \hat{u}_{1,m}$ and retrieved $\hat{u}_{1,m}$, the estimates of $u_{1,m}$.

⁵For $\beta_{3,j}$, in contrast, the no-randomness assumption about $\gamma_{1,m}$ has no impact, and the standard error is not invalidated by that issue.

With these ingredients, we then ran the four key regressions. Three are clean: the single regressor Dumas equation (2), the multiple regression (4) that also includes the market return, and the hybrid (5) that instead adds the orthogonalized market return—the correct $u_{1,m}$ numbers, in this stage, not the calculated residuals. (Equivalently, in the third regression, we orthogonalized using the true $\gamma_{1,m}$, not its sample estimate.) In all these regressions the usual OLS SE should work fine. The fourth regression, lastly, is the realistic version of the third regression, where the estimate $\hat{\gamma}_{1,m}$ has been used to orthogonalize, giving us the sample residuals $\hat{u}_{1,m}$, not the true errors $u_{1,m}$, as a regressor. Here the SE obtained should be too small.

For each of the four regressions we carried out three sets of simulations. In the first, we chose a realistic set of parameters producing a moderate bias in the estimated SE, another producing a much stronger bias, and another, a much weaker bias. The assumed *per annum* parameter values, along with some implied numbers, are shown in the Key to the Table. In simulation S1, the numbers are calibrated to what one gets with monthly data: *per annum* volatilities 0.20 and 0.40 for market and stock, respectively; a market model that explains about one-quarter of the return variability; and a weak exposure effect. In simulation S2, the market factor has a high variance, while the stock has little idiosyncratic variance. Thus, the market model does quite well, R^2 -wise, and adding R_m to the exposure regression gives us very precise estimates for γ . Adding the sample-orthogonalized market, instead, seems to give us about the same precision, but this would be quite misleading. So here we should see a true variability that swamps the reported one. In the third simulation, S3, the explanatory powers for the regressors are swapped, and the market explains very little. So assuming away that part of the noise should not matter much, here, and the SES for the hybrid regression's exposure should still be reasonably accurate.

For each equation we show the mean of the $\hat{\gamma}_{.,j}$ s, the average of the error variances predicted by the regression program, and then the variance of the estimated gammas across the one million experiments, a trustworthy estimate for the true unconditional variance of the estimate. The last column gives us the ratio of this true variance of the exposure estimate to the average variance produced by OLS; any value far away from unity of course indicates that the regression output cannot be trusted, not even on average.

In the clean regressions (the first three), the OLS-computed variances match the true variability across samples nearly perfectly. In the hybrid regression with the constructed data, the

Equation	true γ_j	$\overline{\hat{\gamma_j}}$	$\overline{\mathrm{var}}_{\mathrm{\scriptscriptstyle OLS}}\left(\hat{\gamma}_{j} ight)$	$\widehat{\operatorname{var}}(\hat{\gamma}_j)$	var ratio				
Simulation 1									
$R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j}$	2.25	2.250	.0864	.0865	1.00				
$R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$	1.00	1.000	.1705	.1704	1.00				
$R_{j} = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}$	2.25	2.250	.0665	.0666	1.00				
$R_{j} = \alpha_{4,j} + \gamma_{4,j}s + \beta_{4,j}\hat{u}_{1,m} + u_{4,j}$	2.25	2.250	.0651	.0865	1.33				
5	Simulati	on 2							
$R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j}$	2.25	2.250	.1383	.1384	1.00				
$R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$	1.00	1.000	.0068	.0068	1.00				
$R_{j} = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}$	2.25	2.250	.0054	.0054	1.00				
$R_{j} = \alpha_{4,j} + \gamma_{4,j}s + \beta_{4,j}\hat{u}_{1,m} + u_{4,j}$	2.25	2.250	.0053	.1384	26.03				
Simulation 3									
$R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j}$	2.25	2.250	.1382	.1383	1.00				
$R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$	1.00	1.000	.9846	.9643	1.00				
$R_{j} = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}$	2.25	2.250	.1358	.1359	1.00				
$R_{j} = \alpha_{4,j} + \gamma_{4,j}s + \beta_{4,j}\hat{u}_{1,m} + u_{4,j}$	2.25	2.250	.1329	.1383	1.04				

Key In each simulation, we generated 10^6 samples, each with 50 records $\{R_j, R_m, s\}$, as follows. We generated numbers for s, $u_{1,m}$ and $u_{2,j}$, and derived values for $R_m = \gamma_{1,m}s + u_{1,m}$ and $R_j = \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$. We next calculated the sample coefficients in $R_m = \hat{\alpha}_{1,m} + \hat{\gamma}_{1,m}s + \hat{u}_{1,m}$ and retrieved $\hat{u}_{1,m}$. We lastly ran the four key regressions: (i) the single regressor Dumas equation (2), (ii) the multiple regression (4) including R_m , (iii) the hybrid (5) that instead includes the correct $u_{1,m}$ values, and (iv) the hybrid that uses the estimated $\hat{u}_{1,m}$. Below are the *p.a.* parameter values that we chose in each of the experiments S1–S3, and then the parameters whose values are postulated.) We have $\operatorname{cov}(s, r_m) = \gamma_{1m}\operatorname{var}(s)$ from the familiar definition of a simple regression slope. The Dumas regression (3) then implies $\operatorname{var}(r_m) = \gamma_{1m}^2 \operatorname{var}(s) + \operatorname{var}(r_m | s)$. From the multivariate equation (4), one obtains $\operatorname{var}(r_j) = \gamma_{2j}^2 \operatorname{var}(s) + \beta_{jw}^2 \operatorname{var}(r_m) + 2\gamma_{2j}\beta_{jw}\operatorname{cov}(s, r_m) + \operatorname{var}(u_j | m, s)$; $\operatorname{cov}(r_j, r_m) = \gamma_{2,j}\operatorname{cov}(r_m, s) + \beta_{2j}\operatorname{var}(r_m)$; and $\operatorname{cov}(r_j, s) = \gamma_{2,j}\operatorname{var}(r_m) + \beta_{2j}\operatorname{cov}(r_m, s)$. From these covariances one then obtains the simple slopes $\beta_{0j} = \operatorname{cov}(r_j, r_m)/\operatorname{var}(r_m)$ and the squared correlations $\rho_{ms}^2 = \operatorname{cov}(r_m, s)^2/[\operatorname{var}(m)\operatorname{var}(s)]$; $\rho_{js}^2 = \operatorname{cov}(r_j, \operatorname{var}(s)]$; and $\rho_{jm}^2 = \operatorname{cov}(r_j, s)^2/[\operatorname{var}(r_j)\operatorname{var}(r_m)]$.

	Implied parameters													
Assumed values							vola	tilities	marke	t model	Eq (2)) for m	Eq (2)) for j
	σ_s	$\sigma_{m s}$	$\sigma_{u_{j m,s}}$	γ_{1m}	γ_{2j}	β_{2j}	σ_j	σ_m	$\beta_{0,j}$	$ ho_{jm}^2$	γ_{1m}	$ ho_{ms}^2$	γ_{1j}	ρ_{js}^2
S1	.20	.20	.35	1.25	1.00	1.00	.41	.21	1.06	.27	1.25	.06	2.25	.06
S2	.20	.50	.10	1.25	1.00	1.00	.56	.68	1.16	.91	1.25	.20	2.25	.44
S3	.20	.10	.50	1.25	1.00	1.00	.68	.27	1.70	.45	1.25	.86	2.25	.44

For each equation we show, in the main table, the mean of the $\hat{\gamma}_{.,j}$ s, the average of the error variances predicted by the regression program, and the actually observed cross-sectional variance (var) of the estimated gammas. The ratio of the last two is then shown under the heading "var ratio."

Table 2: Monte Carlo simulation results

last one in each panel, the regression program claims to come up with a SE that is even better than the multivariate while preserving the simple-regression estimate, but this SE underestimates the true one. The theoretical ratio of the actual estimation variance to the calculated variance for the third regression, which from Table 1 equals $1 + \beta_{2,j}^2 \operatorname{var}(u_{1,m})/\operatorname{var}(u_{2,j})$, equals 1.33 in S1, 26 in S2, and 1.04 in S3 when calculated from the (known) population parameter. This is exactly what we see in the variance ratios (the ratios of the average observed variances 1073

8044

519

5320

3.22

29.25

1.18

1.49

	Table 3: Data: Some descriptive statistics											
Chinese exporters sample						U.S. multinationals sample						
	BV	MV	Q	$\beta_0 m$	FSR		BV	MV	Q	$\beta_0 m$	$\beta_0 d$	
	566	1068	2.76	1.12	37.4		30415	63429	3.23	1.09	1.18	
	854	1447	2.85	0.11	22.0		48707	96268	2.51	0.52	1.66	
	5	154	0.62	0.75	9.6		82	4	0.57	-1.50	-9.41	
	135	308	1.52	1.08	18.2		3636	8664	1.65	0.77	0.68	
	273	511	2.09	1.13	32.4		11833	24822	2.53	1.03	1.13	

26552

221890

63153

626550

3.65

19.95

1.42

2.98

1.64

8.78

Key Our China sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. Data are monthly, the main regressor is the CNY Real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return refers to the Shanghai A market index. For the same period we also study daily and monthly returns for the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as multinational. For these stocks we downloaded both daily and monthly returns.

50.2

93.9

 $MV = market value, in USD 10^6; BV = book value, in USD 10^6; Q = MV/BV; FSR = foreign sales ratio; <math>\beta_0 m$ $(\beta_0 m) =$ simple market sensitivity, estimated via OLS from monthly (daily) data; stdev = standard deviation; $\min/\max = \text{lowest/highest}$ value; Q1, Q3 - first and third quartile value.

to the true variance, the rightmost column) for the hybrid regressions.

3 Tests with real data

3.1Data

average stdev \min Q1 Median

Q3

max

We evaluated the size of the bias in four samples extracted from two sets of stocks. Our original sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. (June 2005 is the date when China stopped pegging the Yuan to the USD.) Data are monthly, the main regressor is the Yuan (CNY) real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return refers to the Shanghai A market index. For comparison and to study the intervalling effect⁶ we added a U.S. sample similar to the one in the seminal Jorion study. For the same period as the above, this includes the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as partly multinational (there always is some domestic activity too). For these stocks we downloaded both the daily and monthly returns. The U.S. market factor for these stocks is the Compustat value-weighted market return, and the exchange rate regressor is the effective nominal exchange rate change computed for the U.S. by the Bank of England.

⁶We thank the anonymous referee for this suggestion.

The samples are different in potentially relevant ways. The U.S. firms are roughly 30–100 times bigger (Table 3). From the table, we see that the U.S. firms also have higher Market-to-Book ratios, with a somewhat lower dispersion. Consistently with their relative size, Chinese firms have above-unit betas (average 1.12, median 1.13). These betas exhibit a relatively low dispersion: the total range is 0.75-1.49, with 50% of them being in the range 1.08-1.18. The U.S. sample consisting of larger firms, the mean beta is still surprisingly high (average 1.18; median 1.13). More strikingly, their dispersion is many times larger, whether we consider cross-sectional standard deviation, total range, or interquartile range. Regardless of whether this reflects measurement error or true heterogeneity in beta, it should lead to much larger differentials across stocks in the effect on γ of introducing the market return as a new regressor.

3.2 Total exposure

Let's consider our priors about the firms' total exposures. Being unhedged committed exporters, our Chinese firms should be positively exposed. The U.S. multinationals may be different in many ways. While they very often export, they typically also import, for instance, and they have access to financial hedges (unlike the Chinese firms). They also often use operational hedges by producing abroad in or near their foreign markets. All this makes it hard to make a clear a priori prediction about the headquarter's exposure. The exposures of any given foreign subsidiary could likewise be positive or negative.⁷ In light of all this, we are again not *a priori* sure whether the group as a whole would gain or lose from appreciating foreign currencies.

To get an empirical answer, in each of the three samples the stocks' returns are regressed on the percentage change in the currency index with and without the orthogonalized market return as an additional regressor. The results are summarized in Table 4 and the nine plots shown at the back of the paper, Figures 3–5.

First consider the distribution of the simple gammas, summarized by the first and third columns of figures of Table 4. Visual information is found in Figure 3, where the γ_1 estimates are plotted on the horizontal axis. As expected for exporting firms, most of the total ('Dumas') exposures for the Chinese sample are positive, and half of these are significant. But it turns out

⁷For instance, if a foreign plant produces for sales in that same country, the parent still makes a loss (on the net profits) if that currency depreciates, while if the plant produces for third countries or the parent's home market, a depreciation of the host's currency would be good news.

sample	measure	γ_1	γ_2	$t(\gamma_1)$	$t(\gamma_2)$	$t(\gamma_3)$	$rac{\mathrm{t}(\gamma_3)}{\mathrm{t}(\gamma_1)}$
S&P daily	mean	0.54	0.15	3.68	0.73	4.78	1.16
	\min	-15.40	-16.90	-1.71	-4.06	-1.82	0.87
	quartile 1	0.16	-0.12	1.43	-1.01	1.55	1.03
	median	0.39	-0.03	3.54	-0.15	4.52	1.10
	quartile 3	0.70	0.19	5.25	1.77	7.00	1.20
	max	32.92	31.55	11.80	10.03	16.44	5.17
	std	3.53	3.47	3.15	2.90	4.19	0.39
	signif >0			80	27	81	
	signif < 0			0	10	0	
S&P monthly	mean	2.37	-0.68	2.39	0.07	2.72	1.22
	min	-28.05	-77.19	-2.28	-9.83	-11.80	1.00
	quartile 1	0.68	-0.58	1.15	-0.91	1.23	1.09
	median	1.16	-0.02	2.32	0.01	2.64	1.20
	quartile 3	1.97	0.46	3.70	1.05	4.34	1.36
	max	142.21	29.31	7.09	4.39	8.68	1.67
	std	13.49	8.11	1.82	1.67	2.55	0.16
	signif >0			68	11	74	
	signif < 0			1	4	1	
Chinese Exporters	mean	2.35	0.38	1.72	-0.27	2.45	1.36
	min	-1.42	-4.07	-1.09	-2.27	-1.14	1.03
	quartile 1	1.58	-0.61	1.21	-0.95	1.39	1.20
	median	2.47	0.36	1.79	-0.28	2.34	1.31
	quartile 3	3.33	1.19	2.49	0.48	3.54	1.52
	max	6.28	3.75	4.18	2.87	7.44	2.11
	std	1.35	1.31	0.99	1.05	1.61	0.24
	signif >0			56	3	79	
	signif < 0			0	8	0	

Summary measures for all regression results

The stocks' returns are regressed on the percentage change in the currency index with and without the market return as an additional regressor. The table shows summary measures for the distributions of the simple Dumas gamma (γ_1) and its *t*-statistic, the Jorion multivariate gamma (γ_2) and its *t*-statistic, and the hybrid gamma's (γ_3) *t*-statistic.

Table 4: Summary of regression results for the three samples

that in actual fact the U.S. multinationals, despite the *a priori* differences we just reviewed, behave in the same way. Even more of the estimates are significant, and the average size of the exposure is quite similar to the one for the exporters sample, at least as long as we stick to monthly data. Still for monthly data, the median for China is close to the mean while U.S. one is much lower, indicating right skewness in the exposures.

If one sample looks radically different from the other two in terms of estimated exposures, it would not be the Chinese set but the daily data, whose results in Figure 3 are shown most

Key Our China sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. Data are monthly, the main regressor is the CNY Real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return refers to the Shanghai A market index. For the same period we also study daily and monthly returns for the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as multinational. For these stocks we downloaded both daily and monthly returns.

to the left and whose summary numbers are in the top part of Table 4. It does have more significant estimates $\hat{\gamma}_1$ (or, equivalently, $\hat{\gamma}_3$), in line with the much larger sample size, but its exposures are less often positive and are on average much lower. All the evidence suggests exposures below unity, unlike for the monthly data. This is unexpected, as there is a consensus that value should be convex in the exchange rate, and in an iso-elastic model $V_j = k_j S^{\gamma_j} Z_j$ this requires an elasticity exceeding unity. But this seemingly strong 'intervalling effect' has a disarmingly simple potential explanation: the non-synchronicity of the variables in the right and the left. Our U.S. stock data, from Compustat, are based on New York closing prices while the exchange-rate level is the Bank of England's currency basket for the U.S., based on London late-afternoon prices, which is before New York's lunchtime. So there is a non-synchronicity issue that is especially crucial in daily returns.⁸ Corrections for this bias are discussed in the last section.

3.3 Adding the market return as a regressor

First consider the orthodox procedure, where the raw r_m is added as a regressor. Compared to what we saw for the total exposures, the columns for γ_2 and its *t*-ratio in Table 4 show a very different picture. The shifts are pictured directly in the three cross-plots of γ_2 against γ_1 in Figure 3 or those of $t(\gamma_2)$ against $t(\gamma_1)$ in Figure 4. As expected, the by and large positive exposures of the original Dumas regression become about 50/50 positive/negative after adding the market regressor, reflecting the fact that in the multiple regression the market's exposure, weighted by the stock's beta, is subtracted from the simple-regression exposure. As could be expected on the basis of the distribution of the simple market betas, in the Chinese sample with its narrow range of betas the downward shift is much more orderly than in the U.S. samples.

Simultaneously, the new t-statistics drop: instead of being mostly positive and often significant, they now are about 50/50 positive/negative, with few being significantly different from zero and some of these even significantly negative. For the researcher hunting for evidence of exposure all this is bad news. But the main question is what happens if the researcher's reaction is to first orthogonalize the market return on the regressor of interest, s.

In Table 4, the results are found in the two rightmost columns, and plots of the two

⁸The exchange-rate movements that occurred before 11:00 in New York are reported as part of the day-t change, while the post-11:00 events are part of s_{t+1} , which is not in the regression. In daily data the five-hour non-overlap is quite important, but it is less so in monthly data.



Key Histograms for the ratio of the two *t*-statistics plotted in Figure 5. The numbers are relative frequencies in percent, not absolute frequencies (numbers of stocks). The numbers on the horizontal axis show the lower bound of each interval. Most intervals are 0.125 wide, except for the extreme ones on the right.

Figure 2: Plots of biased against unbiased *t*-statistics (131 Shanghai-listed companies, 2005–2013 and subperiods)

competing t-statistics for the same estimate ($\gamma_1 = \gamma_3$) are shown in Figure 5 at the end of the paper. Each of these three graphs now shows a plot of the t-ratios calculated without (horizontal axis) and with (vertical axis) the constructed market regressor. As expected, the t-ratios from the hybrid regression virtually all lie above the 45-degree line, meaning that these statistics not only maintain their sign but even rise when the orthogonalized market return is added as a regressor. The contrast with the previous regression is very stark. Figure 2 illustrates the same numbers in a slightly different way, viz., as the ratios of the two ts while the rightmost column of Table 4 provides numerical information on that ratio. We start with the table.

The means for the ratios of t-statistics in Table 4 range from 1.16 (daily data) to 1.22 and to 1.36 (Chinese sample). The medians are similar (1.10, 1.20, 1.30) with the difference vis a vis the means again suggesting some right skewness. Graphically, that right skewness was already visible in the cross-plots of Figure 5: while virtually none of the ratios is below unity, the area close to the 45-degree line is packed more tightly. In the histograms for the ratios of ts shown in Figure 2, that skewness is quite obvious, and most pronounced in the Chinese data. Any such asymmetry means that the mean could be quite unrepresentative. More than a quarter of the regression ts are inflated by at least 20% in the daily sample, rising to 36% in the U.S. monthly data and 52% in the Chinese sample. Given that monthly date dominate in this literature, this is not good news. In one case, the ratio is overstated by a factor of five.

It is true that in all three samples the relation to exchange rates is so strong that lowering the t-ratios would not affect the conclusion that exposures are positive, by and large. Even in the Chinese sample, with the correct test, 56 of the coefficients retain their significance, compared to 79 with the flawed ratio, so the broad conclusions are unaffected. But that does not mean the bias is generally a non-issue. In a lower power application, the choice of the standard error is much more decisive. If the average simple t-statistic is already 5 and the inflation factor 1.5, then conclusions are not much affected; but if the average simple t-statistic is just 1.5, then an inflation factor of 1.5 generates an average inflated t-ratio of 2.25, lifting the average stock to a misleading semblance of significance.

The example is somewhat unsatisfactory in the sense that it just discusses means, and ignores the distribution around the mean. To get a somewhat more complete impression of the potential effect in samples with less power, we carry out a modest thought experiment along the lines of the above example. We lower all simple *t*-statistics for the 116 Chinese stocks to levels one would see if the R^2 s had all been lower,⁹ and we compute what the inflated *t* would have been for the same ratio $SE_{j,3}/SE_{j,1}$ as before. That is, from a lowered value $t'_{j,1}$ we compute a corresponding $t'_{3,1} = SE_{j,3}/SE_{j,1} \times t'_{j,1}$. So for every degree of shrinkage in the R^2 s, we end up with a corresponding distribution of 116 pairs of *t*-statistics, one reliable and one inflated. Then we count how many *t*s are above 1.96.

Table 5 summarizes the results. The lower the R^2 in the simple regressions, the fewer *t*-statistics we see, an utterly unsurprising result, but also the more decisive is any given inflation factor. That is, the lower the R^2 , the higher the chance of incorrectly concluding that the estimate is significant. For instance, if we shrink all R^2 s so that the mean becomes 0.05, 61 of the inflated *t*s seem significant whereas only 27 are if we rely on the better test. For an even lower R^2 , 22 regressions still seem to produce a statistically clear answer, but in reality

⁹There is a one-to-one link between R^2 and the *t*-ratio, in a simple regression. So we start from the distribution of 116 observed R^2 s from the simple ('Dumas') regressions, and shrink them all by a factor between 20 and 80 percent. That is, we compute four distributions of lowered coefficients, $R_{j,1}^{2'} = f R_{j,1}^2$ with $f = \{0.8, 0.6, 0.4, 0.2\}$. For each of the new R^2 s we then compute the *t*-statistic that would have been observed in such a regression, using the known relation $t'_{j,1} = \sqrt{(n-2)R_{j,1}^{2'}/(1-R_{j,1}^{2'})}$.

f	1	0.8	0.6	0.4	0.2
new average R^2	0.12	0.10	0.07	0.05	0.02
Actually significant	96	87	64	27	1
Deemed significant	111	105	95	61	22

Key In this thought experiment, we imagine that, in each regression, R^2 is down to a fraction $f = \{0.8, 0.6, 0.4, 0.2\}$ of the actually observed level; we recalculate the Dumas t corresponding to each new R^2 ; we calculate a new inflated t using the same $SE_{j,3}/SE_{j,1}$ ratio as in the real data; and we count how many ts are above 1.96.

Table 5: Numbers of instances of apparent versus genuine significance for various R^2 s.

only one coefficient is significant.

Quite apart from the above calculations, one cannot accept that there is no problem with using a 30–40 percent inflated *t*-ratios on the basis of the argument that in many samples the conclusion may not be affected. The divergence is comparable to the difference between, e.g., Dickey–Fuller and regular critical values in a unit root test and should not be ignored. And the misleading conclusion could be much worse if a diagnosis is based on a single regression rather than the set of 116 or 131 tests we use here, or if the relation is less clear than it is here.

Besides the technical issue discussed in this section, there are also some interesting research design questions and findings raised by this study. That discussion is relegated to the concluding section.

4 Conclusion

Before summarizing the statistical issue that gave rise to this paper, we open this section with a few less technical insights that, we hope, are useful for the exposure literature. First, when estimating exposure, the data frequency may matter: daily data may provide different answers than monthly. But our results show that it is by no means necessary to adopt long holding periods to find a meaningful link. In fact, despite a huge non-synchronicity bias that makes us lose about three-quarters of the currency effect, the daily sample still has more significant coefficients than the two monthly-return samples.

A related practical conclusion is that when daily data are studied, it pays to be careful about the time of the day at which the variables are recorded. Given our objective to expose orthogonalization as an undesirable ingredient in one's research, we took a conveniently available existing aggregate, but if one's ambition is to to study the exposure itself rather than the consistency of a SE, one needs to be more careful. The Federal Reserve at St. Louis now has intra-day data, for instance, that allow one to construct a currency index fully synchronous with the close of the stock market. The alternative is to somehow correct for the missing information by including, one way or another, a leading s observation.¹⁰ While this removes the bias from general non-synchronicity, it comes at the cost of substantially increased standard errors (Sercu, Vandebroek and Vinaimont, 2008). Constructing a better synchronized database may, therefore, be worth the effort.

A third finding is that multinationals may be as exposed as very committed exporters. In our sample, they certainly seem to be. This is mildly surprising as many of their activities are to some extent hedged operationally (by producing in the country of selling, for instance) and they are thought to be rather good at hedging in financial markets. Our observation suggests that their ability or willingness to hedge, financially or otherwise, may be lower than what many researchers think.

The above comments are all based on simple exposures, Dumas' γ_1 ; the difference from the findings in a substantial part of the literature to a large extent derives from the widespread practice of adding the market return as a regressor. One comment related to this issue is that this choice is far from innocent. Its impact is quite big, as we showed. That impact is related to the special nature of the market return. In many research areas it makes sense to add many 'control' variables so as to reduce the chance that an apparent link may actually be due to the regressor proxying for an omitted third variable. One has to realise, however, that the market return is different: it is not an exogenous third variable, but an aggregate of the very stock returns that act as dependent variables in the exposure regressions. If most stocks are positively exposed, then so will be the market. As a result, the market return ends up proxying for the currency effect rather than the other way around. The conclusion in a regular context would have been that what looks like currency exposure really is market exposure, but here

$$\hat{\gamma}_{1,j}^{IV(L)} := \frac{\operatorname{cov}(r_j, z)}{\operatorname{cov}(s, z)}, \text{ with } z_t := \sum_{\tau = -L}^{\tau = +L} s_{t-\tau},$$
(16)

with the familiar 2SLS generalization if there are other regressors.

¹⁰In the thin-trading literature on the market model, Dimson (1979) proposes using the sum of leading and lagged regression coefficients as the estimator of the total impact. Fowler and Rorke (1983) show that this generates a small bias and propose a generalized Scholes–Williams (1977) estimator. Apte, Kane and Sercu (1994) show that the same estimator also works outside the market model framework, for example if one wants to measure the total effect of commodity price changes on exchange rates when the former affect CPI inflation with one or more lags. The estimator is basically an instrumental variable one, where the instrument is a moving sum of xs from lag L to lead L. In the case of the exchange rate return as the sole regressor, we have

the opposite holds.

As a last non-technical comment, then, a researcher should perhaps think why the market return is to be added as a regressor, if it is added at all. The answer very much depends on the purpose. Again, if one wants to document the existence of exposure *tout court*, then adding the market return is a bad idea. The situation is different if the regression is part of an asset pricing test. If one were to find that once r_m is added, s no longer has any effect, then the useful conclusion is that the simple CAPM suffices: the market return is taking enough of the currency factor into account, in the same way as we expect it to take into account, say, interest risk or business cycle risks. More generally, in an asset pricing context, we know *a priori* that r_m has to be there, and the question whether it has subsumed risks like currency risk is not a threat to the validity of the findings. For other purposes, the answer may again depend on one's objectives. For a hedger who adopts just currency forwards as a risk management instrument, the total exposure γ_1 matters, and adding r_m as a control would be quite misguided. But the answer changes if the hedger decides to adopt also index futures as a tool: then for the currency hedge γ_2 is needed, not γ_1 .

We conclude with the statistical issue associated with orthogonalized regressors. This is a not infrequent practice in economic and business research. By orthogonalizing the market return on the exchange rate, one seems to have the best of both worlds: the market factor cannot subsume part of the forex exposure present in a stock's return, and the SE of the estimate beats both the simple- and the multiple-regression SEs. In other contexts, the objective of orthogonalization is to compress the SEs in case of near-collinear regressors.

But this SE-reducing effect is illusory: since in any particular sample the simple and the hybrid regression coefficients are always equal to one another, their precisions must be identical too. The reported standard errors of the hybrid regression are biased, and often substantially so, even in large samples. Technically, the source of the problem is that the uncertainty about the market's exposure estimate is left out of the calculated SE: when, as in traditional regression theory, the regressors are deemed to be non-random, the regression relations between them are therefore also deemed to be non-random. This assumption is the feature that makes the SE estimates inconsistent. And of course the problem is related to the fact that the orthogonalization happens outside the main regression, so that the uncertainties created in that first step cannot be taken into account in the main regression.

In practice, then, one can*not* have the coefficient of the simple regression with the precision

of a multiple one. If one does not want a control variable to subsume the role of the main variable of interest, then the only solution is to exclude it as a regressor. The only valid alternative to omitting the control variable is to run a regular multiple regression, and take the subsumption effect as inevitable. There is no in-between solution.

References

- Allayannis, G., 1997. The time-variation of exchange rate exposure: An industry analysis. Working Paper, NYU.
- Apte, P., M. Kane, and P. Sercu, 1994: Evidence of PPP in the medium run, Journal of International Money and Finance, 601–622
- Bartram, S. M., and G. M. Bodnar, 2012. Crossing the lines: The conditional relation between exchange rate exposure and stock returns in emerging and developed markets, *Journal of International Money and Finance*, 31(4), 766–792,
- Bodnar, G. M., and M. H. F. Wong, 2003. Estimating Exchange Rate Exposures: Issues in Model Structure, *Financial Management* 32(1), 35–63.
- Bris A., Y. Koskinen and V. P. Pons-Sanz, 2002. Corporate Financial Policies and Performance Around Currency Crises, Yale ICF Working Paper No. 00-61
- Dimson, E., 1979, Risk measurement when shares are subject to infrequent trading, Journal of Financial Economics 7, 197–226
- Dumas, B., 1978. The theory of the trading firm revisited, *Journal of Finance* 33(3), 1019–1029.
- Entorf, H., J. Moebert and K. Sonderhof, 2006. The Foreign Exchange Exposure of Nations, Darmstadt UT, Working paper.
- Fowler, D. J., and H. C. Rorke, 1983, Risk Measurement when shares are subject to infrequent trading: Comment, Journal of Financial Economics 12, 297–283
- Frisch, R., and F. Waugh, 1933, Partial time regressions as compared with individual trends, *Econo*metrica 45, 939–953.
- Glaum, M., M. Brunner and H. Himmel, 2000. the DAX and the Dollar: The economic exchange rate exposure of German corporations, *Journal of International Business Studies* 31(4), 715–724
- Griffin, J. M., and R. M. Stulz, 2001. International competition and exchange rate shocks: A crosscountry industry analysis of stock returns. *Review of Financial Studies* 14(1), 215–241.
- Hagelin N., and B. Pramborg, 2002. Hedging Foreign Exchange Exposure: Risk Reduction from Transaction and Translation Hedging. *Journal of International Financial Management and Accounting* 15, 1–20.
- Jayasinghe P., and B. Pramborg, 2002. Hedging Foreign Exchange Exposure: Risk Reduction from Transaction and Translation Hedging, Journal of International Financial Management and Accounting 15, 1–20
- Johnson, L. L., 1960: The theory of hedging and speculation in commodity futures, Review of Economic Studies 27, 139–151
- Jorion, P., 1990: The Exchange Rate Exposure of US Multinationals, Journal of Business (July), 331–345.
- Kiymaz, H., 2003. Estimation of foreign exchange exposure: An emerging market application. Journal of Multinational Financial Management 13, 71–84.
- Priestley R., and B. A. Odegaard, 2002: New Evidence on Exchange Rate Exposure, Norwegian School of Management working paper.

٦

0

- Pritamani, M. D., D. K. Shome, and V. S. Singal, 2004. Foreign exchange exposure of exporting and importing firms, *Journal of Banking and Finance* 28(7), 1697–1710.
- Scholes, M., and J. Williams, 1977: Estimating betas from non-synchronous data, Journal of Financial Economics 5, 308–328
- Sercu, P., M. Vandebroek and T. Vinaimont, 2008. Thin-Trading Effects in Beta: Bias v. Estimation Error. Journal of Business Finance and Accounting 35(9–10), 1196–1219
- Stein, J. L., 1961, The simultaneous determination of spot and futures prices, American Economic Review 51, 1012–1025
- Stoll, H., and R. Whaley, 1990: The dynamics of stock index and stock index futures returns, Journal of Financial and Quantitative Analysis 25, 441–368

Technical Note A: Proof of Equation (11)

The proof, for completeness, is as follows:

$$\begin{aligned} \mathbf{G}^{-1}\mathbf{B}_{2} &= \begin{bmatrix} 1 & \gamma_{1,m} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sigma_{s}^{2} & \sigma_{m,s} \\ \sigma_{m,s} & \sigma_{m}^{2} \end{bmatrix}^{-1} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} 1 & \gamma_{1,m} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sigma_{m,s}^{2} & -\sigma_{m,s} \\ -\sigma_{m,s} & \sigma_{s}^{2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{m}^{2}\sigma_{s}^{2}(1-\rho_{m,s}^{2}) \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{m}^{2}\sigma_{s}^{2}(1-\rho_{m,s}^{2}) \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{s}^{-2} & 0 \\ -\gamma_{1,m}\sigma_{m,s}^{-2} & \sigma_{m,s}^{-2} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,s} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} = \begin{bmatrix} \sigma_{j,m} \\ \sigma_{j,m} \\ \sigma_{j,m} \\ \sigma_{j,m} \\ \sigma_{j,m} \end{bmatrix} \begin{bmatrix} \sigma_{j,m} \\ \sigma_$$

The last line uses the familiar definitions of simple and multivariate regression coefficients.

Technical Note B: Proof of the propositions in Table (13).

The slopes follow directly from the main arguments, especially Equations (6) and (13), except that sample moments replace the population counterparts. For the error (co)variances, the general OLS solution is

$$\sigma_u(\mathbf{X}'\mathbf{X})^{-1} = \frac{\hat{\sigma}_u}{N-1} \frac{\begin{bmatrix} \hat{\sigma}_m^2 & -\hat{\sigma}_{m,s} \\ -\hat{\sigma}_{m,s} & \hat{\sigma}_s^2 \end{bmatrix}}{\hat{\sigma}_m^2 \hat{\sigma}_s^2 (1-\hat{\rho}_{m,s}^2)},$$

and the error variance of the exposure relies on the first diagonal element of that matrix. That element simplifies to $1/[(N-1)\hat{\sigma}_s^2(1-\hat{\rho}_{m,s}^2)]$, the inverse of the sum of squares (SS) of the currency regressor conditional on the other regressor. The error variance in the first line in the table follows because in the Dumas regression there is no other regressor, so the relevant SS is just the raw sum of squares. In the second line, the table follows the above, except that $(N-1)\hat{\sigma}_s^2$ is again written as a SS, while in the third line we also use the fact that the sample correlation is exactly zero, the second regressor having been orthogonalized first. Lines two and three share the same estimated variance of the residuals because their residuals are identical. In Line 1, the variance of the residuals is larger because that regression error contains, next to u_2 , also $\beta_{3,j}u_{1,m}$. That relation is made explicit in the second version of the error variance shown in the table.



Key Our China sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. Data are monthly, the main regressor is the CNY Real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return refers to the Shanghai A market index. For the same period we also study daily and monthly returns for the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as multinational. For these stocks we downloaded both daily and monthly returns.

The stocks' returns are regressed on the percentage change in the currency index with and without the market return as an additional regressor. The graphs in the top part plot the multivariate gamma against the simple one, alongside a 45-degree line for reference. So the relevant regressions are (2) ('Dumas') and the multivariate (4):

Dumas regression: $R_j = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j};$ multivariate Model:

 $R_{j} = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_{m} + u_{2,j}.$

Figure 3: Comparing simple against multiple gammas



Key Our China sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. Data are monthly, the main regressor is the CNY Real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return refers to the Shanghai A market index. For the same period we also study daily and monthly returns for the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as multinational. For these stocks we downloaded both daily and monthly returns.

The stocks' returns are regressed on the percentage change in the currency index with and without the market return as an additional regressor. The graphs in the top part plot the multivariate gamma's t-statistic against that of the simple one, alongside a 45-degree line for reference. So the relevant regressions are (2) ('Dumas') and the multivariate (4):

Dumas regression: $R_j = \alpha_{1,j} + \gamma_{1,js} + u_{1,j};$ m

multivariate Model: $R_j = \alpha_{2,j} + \gamma_{2,j}s + \beta_{2,j}R_m + u_{2,j}$.

Figure 4: Comparing simple against multiple *t*-statistics



Key Our China sample consists of 131 Shanghai- or Shenzen-listed stocks that have full data since June 2005 and report at least 10 percent foreign sales. Data are monthly, the main regressor is the CNY Real effective exchange rate based on CPI (IMF) re-expressed as the value of the basket in CNY, and the market return The stocks' returns are regressed on the percentage change in the currency index with and without the orthogonalized market return as an additional regressor (the refers to the Shanghai A market index. For the same period we also study daily and monthly returns for the 116 companies that are in the S&P 500 index and that, in Compustat, are marked as multinational. For these stocks we downloaded both daily and monthly returns.

hybrid regression). The graphs in the top part plot the hybrid gamma's t-ratio against that of the simple one, alongside a 45-degree line for reference. So the relevant regressions are (2) ('Dumas') and the hybrid one (5):

 $R_{j} = \alpha_{3,j} + \gamma_{3,j}s + \beta_{3,j}u_{1,m} + u_{3,j}.$ Hybrid Model: $R_{j} = \alpha_{1,j} + \gamma_{1,j}s + u_{1,j};$ Dumas regression:

Figure 5: Comparing simple against multiple gammas, and similarly for the t-statistics