Journal of International Money and Finance 51 (2015) 245–263

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Orthogonalized regressors and spurious precision, with an application to currency exposures☆

Fang Liu [a], Piet Sercu [b,*], Martina Vandebroek [b]

[a] Central University of Finance and Economics, Beijing, China
[b] KU Leuven, Faculty of Economics and Business, Leuven, Belgium

ARTICLE INFO

Article history:
Available online 3 December 2014

JEL-codes:
C3
C58

Keywords:
Financial econometrics
International finance
Market model
Currency exposure
Orthogonal

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 20–100 percent.

☆ 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.

* Corresponding author.
E-mail addresses: liufang8000@aliyun.com (F. Liu), piet.sercu@kuleuven.be (P. Sercu), martina.vandebroek@kuleuven.be (M. Vandebroek).

http://dx.doi.org/10.1016/j.jimonfin.2014.11.008
0261-5606/

1. 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 pre-orthogonalization 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, 1961; 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 return-generating 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 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.

2. 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 (1997), Bartram and Bodnar (2012), Bodnar and Wong (2003), Bris and Koskinen (2002), Entorf et al. (2006), Griffin and Stulz

[1] For a review of the theoretical and empirical literature, see, *e.g.*, Boudt, Liu and Sercu (2014).
[2] See Section 1.1 if this does not sound familiar.