



Summary of the Joint Spring Seminar
INQUIRE UK and INQUIRE Europe

April 3-5, 2011

De Vere University Arms Hotel,
Cambridge, United Kingdom

Summarized by:
Maarten van Ravenswaaij
(AEGON Asset Management)



Table of Contents

Program of the Seminar	3
Shrinkage Regression for Multivariate Inference with Missing Data, and an Application to Portfolio Balancing (teach-in session) <i>Robert B. Gramacy and Ester Pantaleo</i>	4
Capital-Market Effects of Securities Regulation: the Role of Implementation and Enforcement <i>Hans B. Christensen, Luzi Hail and Christian Leuz</i>	5
To Trade or Not To Trade: the Strategic Trading of Insiders Around News Announcements <i>Adriana Korczak, Piotr Korczak and Meziane Lasfer</i>	6
Decoding Inside Information* <i>Lauren Cohen, Christopher Malloy and Lukasz Pomorski</i>	8
Does One Size Fit All? The Consequences of Switching Markets with Different Regulatory Standards <i>Tarun Ramadorai & Tim Jenkinson</i>	9
Improving Portfolio Selection Using Option-Implied Volatility and Skewness# <i>Victor DeMiguel, Yuliya Plyakha, Raman Uppal and Grigory Vilkov</i>	10
Econometric Measures of Systemic Risk in the Finance and Insurance Sectors# <i>Mila Getmansky, Liorana Pelizzon, Monica Billio and Andrew W. Lo</i>	11
Ex-post Portfolio Performance with Predictable Skewness and Kurtosis# <i>Giovanna Nicodano and Massimo Guidolin</i>	12
Do Foreigners Know Better? A Comparison of the Performance of Local and Foreign Mutual Fund Managers# <i>Miguel A. Ferreira, Pedro Matos and Joao Pereira</i>	14

* denotes paper sponsored by INQUIRE UK

denotes paper sponsored by INQUIRE Europe

Presenting authors are shown in italics



Program of the Seminar

Monday, April 4th, 2011

Shrinkage Regression for Multivariate Inference with Missing Data, and an Application to Portfolio Balancing (teach-in session)

Robert B. Gramacy and Ester Pantaleo

Capital-Market Effects of Securities Regulation: the Role of Implementation and Enforcement

Hans B. Christensen, Luzi Hail and Christian Leuz

To Trade or Not To Trade: the Strategic Trading of Insiders Around News Announcements

Adriana Korczak, *Piotr Korczak* and Meziane Lasfer

Decoding Inside Information*

Lauren Cohen, Christopher Malloy and *Lukasz Pomorski*

Does One Size Fit All? The Consequences of Switching Markets with Different Regulatory Standards

Tarun Ramadorai & *Tim Jenkinson*

Improving Portfolio Selection Using Option-Implied Volatility and Skewness#

Victor DeMiguel, Yuliya Plyakha, *Raman Uppal* and Grigory Vilkov

Tuesday, April 5th, 2011

Econometric Measures of Systemic Risk in the Finance and Insurance Sectors#

Mila Getmansky, Liorana Pelizzon, Monica Billio and Andrew W. Lo

Ex-post Portfolio Performance with Predictable Skewness and Kurtosis#

Giovanna Nicodano and Massimo Guidolin

Do Foreigners Know Better? A Comparison of the Performance of Local and Foreign Mutual Fund Managers#

Miguel A. Ferreira, Pedro Matos and *Joao Pereira*

* denotes paper sponsored by INQUIRE UK

denotes paper sponsored by INQUIRE Europe

Presenting authors are shown in *italics*



Shrinkage Regression for Multivariate Inference with Missing Data, and an Application to Portfolio Balancing (teach-in session)

Robert B. Gramacy

Robert Gramacy opened the Monday morning session with a teach-in on statistical inference of missing data. This concerns joint work with Ester Pantaleo. The goal of the analysis is to estimate the mean and variance parameters used in portfolio optimization. The sample data used are monthly equity returns from NYSE and AMEX from 1968 to 1997. In practice not all stocks have an observed return for every month as e.g. some firms got listed after the start of the sample period. Instead of limiting the estimation sample to the subset of stocks with relatively more observations or limiting the estimation sample to the subset of months with relatively more observations the goal of the analysis is to use all the available data to estimate the mean and covariances.

When the stocks are sorted by historical data availability, the stocks decrease in observability when scanning the return matrix from left to right. The pattern that emerges is referred to as a monotone missingness pattern. Andersen showed in 1957 how to use iterative OLS regressions to get the maximum likelihood estimator for the mean vector and the covariance matrix for this type of data. The algorithm starts with the best observed stock and ends with the least observed stock. In each step the mean and covariances of a particular stock are estimated using data of all the other stocks that have observed returns in the same period as the particular stock. The algorithm is implemented in an R package which was demonstrated to the audience.

Bayesian inference can be used as an alternative to the maximum likelihood approach. In this case a non-informative prior is used which enables sampling from the posterior distribution. Conditional on the data the distribution for both parameters is centered around the maximum likelihood estimate. Samples from these distributions can easily be converted into estimates of the mean and covariances using similar transformations as in the maximum likelihood approach. Nice features of this approach are that it takes estimation risk into account as documented by Stambaugh (1997) and that it can handle missingness patterns that contain glitches (i.e. are not perfectly monotone).

Unfortunately neither approach can be applied when there are fewer observations than assets or when there is an asset in the sample with fewer returns than the number of assets with more returns. This problem can be solved by replacing the OLS regressions by ridge regressions or lasso regressions. These solutions don't account for estimation risk analytically but a Bayesian approach using Gibbs sampling will allow to do so.

In the setup thus far, missing asset returns are estimated using returns of other assets only. But market factors such as the Fama French factors can be added to form a joint model. This leads to even better estimates of the mean vector and the covariance matrix. This is supported by results from an empirical analysis that allows for a ranking of the different approaches by Sharpe ratio. Approaches that do not use all the available data perform worst. Of the approaches that can deal with the missingness pattern, those that take parameter uncertainty into account outperform the other approaches. The lasso models seem more suitable to model stock returns than the ridge models.



Capital-Market Effects of Securities Regulation: the Role of Implementation and Enforcement

Hans B. Christensen

The presentation by Hans Christensen on capital markets regulation is joint work with Luzi Hail and Christian Leuz. All developed markets have extensive securities regulation but the debate on the costs and benefits of regulation is ongoing. Prior evidence is decidedly mixed and often negative on the benefits of regulation. Previous studies are characterized by focusing on the implementation of a single law in a single country or by focusing on cross-sectional differences between countries. In the research presented, two laws are studied which were implemented at different dates in the 30 member states of the European Union (EU) as part of the Financial Services Action Plan: the Market Abuse Directive and the Transparency Directive. Apart from variation in the implementation date, the country implementations differ mainly by the severity of the enforcement (penalty) and by the amount of resources used to implement the directives.

Existing theory suggests securities regulation can be beneficial to capital markets but information effects, political resistance and regulatory capture can stand in the way. In the EU the cross sectional effect could be “catching up” i.e. countries with weaker regulation improve more relative to countries that already have good regulation in place. An alternative prediction is the absence of a “catching up” effect in case the cause of weaker regulation is still present and prevents the weaker countries to properly implement the new directives.

The market abuse directive (MAD) passed in 2003 and was implemented in 2004. It is a directive against market manipulation and insider trading to ensure investor confidence. It uses a clear definition of inside information which made it easier to identify market abuse and requires member states to create or designate an enforcement agency with sufficient powers to investigate possible breaches.

The transparency directive (TPD) passed in 2004 and was implemented in 2007. It is a directive to ensure appropriate transparency for investors by requiring disclosure of accurate and timely information such as financial statements. Each member state needed to designate or create an enforcement agency with appropriate enforcement tools.

To examine the effect of the implementation of both directives, quarterly data is used. The benefits of regulation is proxied by market liquidity which is measured by the bid-ask spread. Costs of regulation are not included. For each country the MAD and TPD implementation dates are known and these differ among countries which helps to isolate the implementation effect. Explanatory variables are the change in the amount of resources allocated to supervisory powers, the strength of the supervisory powers and the maximum degree of enforcement. The estimated increase in liquidity is 14% for the MAD and about 20% for the TPD after controlling for time trends, country effects, macro economic effects and other EU directives. Alternative model specifications also display a significant positive effect albeit with different magnitudes.

Although the EU intended to cause “catching up” of the weaker countries the results indicate that the implementation benefits are strongest in the countries which had strong regulation prior to implementation. Only for these countries a higher degree of implementation and a higher amount of resources allocated to supervisory increases the benefits even more. Why weaker countries fail to catch up is subject of further research.



To Trade or Not To Trade: the Strategic Trading of Insiders Around News Announcements

Piotr Korczak

The presentation by Piotr Korczak on insider trading is joint work with Adriana Korczak and Meziane Lasfer. The focus of the paper is to examine the relation between reported trades by corporate insiders and corporate news releases. There is widespread evidence that corporate insiders earn abnormal returns on trades in their companies' shares. This suggests they have superior knowledge about the firm's prospects. This link between reported insider trading and subsequent corporate events is supported by empirical research especially for the long-term but less so for the short term, perhaps because trading on inside information can be easier detected in the short term by regulators.

In the UK, insider trading regulation consists of legislation and codes of good practice governed by the FSA. Trading on inside information is prohibited. Inside information is defined as specific and precise information that is likely to have a significant effect on the stock price when made public. The definition of a significant effect is not quantified and may differ from case to case. This also leaves some judgment to the corporate insider. UK insider trading regulations also define trading bans of 30 to 60 days for corporate insiders around earnings announcements or during any period when there exists a matter which constitutes inside information in relation to the company.

The authors study the reported trading by insiders in the 30 day period before news releases where private information is short lived. It is assumed insiders can correctly predict the significance of the news announcement yet they face the trade-off between capitalising on private information and the penalties of getting caught (regulatory scrutiny and reputational damage). To estimate how insiders deal with this trade-off data of all 120 thousand news releases by FTSE All Share firms published from 1999 up to 2002 were collected and categorized as "banned" for earnings announcements and "not banned" for various other type of news releases. Based on the abnormal returns, the sample contains about 50% good news events and 50% bad news events.

More than 8000 trades by corporate insiders are linked to these news events in the 30 day period prior to the release. The net trading (buys minus sells) is measured as well as an indicator variable that equals one when the direction of the insider trading is consistent with the content of the news and zero otherwise. About 10% of the good news events is preceded by insider trading versus 3% of the bad news events. OLS and Logit regression models are estimated separately for good and bad news events to explain which factors impact the net trading and the probability of trading.

For the good news events the relation between the cumulative abnormal return (CAR) of the news event and the net amount of trading by insiders is inverted U-shaped. This supports the hypothesis that insiders have an incentive to trade more when the expected return is larger, but after a certain level the potential penalties of getting caught reduce the trading activity. This result is especially observed in the post December 2001 sub period after the implementation of tougher insider trading regulations. Of the controlling variables the past return comes out significant which is in line with other papers that report contrarian trading by insiders.



For the bad news events the inverted U-shape relation between CAR and net trading by insiders is not observed. However, the coefficient estimates of CAR as explanatory variables of the probability of trading (as estimated by the Logit model) come out significant. Perhaps the disincentives (such as increased risk of litigation) causes insiders to attach more importance to the trading decision rather than the actual trading amount.



Decoding Inside Information

Lukasz Pomorski

The paper presented by Lukasz Pomorski is joint work with Lauren Cohen and Christopher Malloy. It aims to discover trading patterns by corporate insiders. An example of a possible pattern is buying by insiders after a cash bonus is received and when a discount plan in the company stock is in place. Patterns in sells can happen when insiders routinely get rid of an overweight in the company stock in order to diversify their portfolio.

Insider data from Thomson Reuters insider filings database is used. It contains the identity of the insider, the name of the company and the date of the trade. Stock data from CRSP/Compustat/IBES and news data including Dow Jones Newswires is added to complete the dataset. Based on the past 3 years of trading, an insider is classified as a routine trader if he/she placed a trade in the same calendar month each year. All other insiders are classified as opportunistic. As the average insider trades about three times during the sample period, this classification rule cuts down the sample of insiders. Of the remaining trades approximately 50% is classified as routine.

A regression model is specified where the stock return of the subsequent month depends on opportunistic and routine buys and sells as well as common determinants of returns in the current month. From the estimates there does not appear to be a significant relation between routine trades and stock returns, but for opportunistic trades the subsequent stock return increases significantly by 0.57% for buys and decreases significantly by 0.67%. These results suggest opportunistic trades are informative and routine trades are not. The difference between opportunistic and routine trades is 0.76% for buys and 0.82% for sells. Portfolios that are setup to exploit this information earn about 1.80% excess return on a monthly basis and these returns persist for about a year.

In the remainder of the talk results of additional tests were presented to explain the observed trading patterns and their impact on returns. For a one standard deviation increase in the number of opportunistic trades there is a 20% increase in the number of important news events about the firm next month. Returns on announcement days are positively related to opportunistic buys in the previous month and negatively related to opportunistic sells in the previous month. For routine trades there does not appear to be such an information effect.

Since opportunistic trades predict future firm level returns and news it is plausible that opportunistic traders might be especially sensitive to the potential costs and penalties associated with illegal insider trading. To test this idea the fraction of insiders trading in a given month who are opportunistic is regressed on the number of recent SEC releases regarding litigation cases against illegal insider trading. Although the estimated effect is not very large, the results indicate that opportunistic trading is reduced in times of high SEC activity regarding insider trading.

Within the group of insiders the trades by local insiders are especially informative for future information events. Trades of other insiders (independent directors, senior insiders and inside directors) appear to have much less signalling power.



Does One Size Fit All? The Consequences of Switching Markets with Different Regulatory Standards

Tim Jenkinson

The presentation by Tim Jenkinson is on joint work with Tarun Ramadorai. Before the financial crisis there was a debate concerning the up rise of less regulated markets. In this paper two U.K. markets with different levels of regulatory standards are compared. On the one hand regulation protects investors and improves the corporate governance standards but on the other hand it costs money and may limit the flexibility firms have to operate.

After recent moves in developed markets to raise regulatory standards, companies have been fleeing regulatory burdens as illustrated by the collapse of international listings in the U.S. since the passing of the Sarbanes-Oxley act. Lightly regulated market segments have become more popular such as the Alternative Investment Market (AIM) segment in London.

In the 2004-2007 period, when most developed markets didn't attract more than 50 IPO's per year, AIM attracted a few hundred new admissions each year. As many firms switched trading from the Main Market (MM) in London to AIM in this period it caused quite a controversial debate within the investor community. The decision to switch "down" from MM to AIM could be made by company management without shareholder approval. Many companies motivated their decision to switch down to bring down the regulatory costs of being listed: both admission requirements (minimum shares in public hands, available trading record, documents pre-vetted by the UK Listing Authority and minimum market capitalisation) and continuing obligations (shareholder approval for corporate actions) are absent or less strict on AIM and the annual fee is lower as well.

The authors apply an event study methodology to the switching firms using data from Datastream and the Regulatory News Service of the LSE. Firms with major other corporate actions during the switch period were excluded from the sample, resulting in 220 firms in the sample that switched down and about 50 that switched up. On average 30 days elapsed between the announcement date of the switch and the switch event.

In the two week period surrounding the announcement date of the down-switch the average firm suffered a price decrease of -5.4%. The up switching companies experienced a 5.4% price increase. Only for the down switchers there is a surprising long term bounce back effect: the average one year return effect is about +15% to +25%. Of all the potential causes of this phenomenon only changes in operating performance helps to explain this bounce in stock returns. Possibly the future returns were not fully anticipated at the time of the announcement. As operating performance increased, investors gradually learned about the positive effects of the switch. Despite the controversial debates the switch to the less regulated AIM market does not appear to be so bad for investors after all.



Improving Portfolio Selection Using Option-Implied Volatility and Skewness

Raman Uppal

The presentation held by Raman Uppal is joint work with Victor DeMiguel, Yuliya Plyakha and Grigory Vilkov. When performing a mean-variance analysis to construct an optimal portfolio, one needs to estimate means, volatilities and correlations. Traditionally these moments were estimated from historical data but several papers have documented poor out of sample performance of such portfolios due to estimation error. Various solutions have been proposed in the literature to deal with this issue such as Bayesian methods and the use of additional stock return characteristics. In this paper the authors use option implied information that reflect the market's perception of the future rather than just historical data.

The dataset consists of daily stock returns from CRSP from January 1995 to June 2007, intraday data from NYSE's TAQ, options data from IvyDB (OptionMetrics) that has all U.S. exchange traded index and equity options. Only out of the money implied volatilities are used as in the money options may contain a premium for early exercise which could affect the analysis. The merge of the stock and options data results in two sets of data: one small sample of 100 stocks with no missing observations and a larger sample of 561 stocks with missing observations. The results for these sets were similar.

From the data the following implied characteristics are estimated:

- The model-free option implied volatility (see Bakshi, Kapadia and Madan, 2003)
- The historical volatility risk premium (HVRP) as the ratio of implied volatility to expected volatility. The latter is proxied by realized volatility.
- The implied minus the realized volatility spread (IRVS) measured as the spread between the Black-Scholes implied volatility and the realized stock return volatility for the past month
- Implied correlations
- The model-free option implied skewness (MFIS)
- The spread between Black-Scholes implied volatility for pairs of calls and puts (CPVS). From the put-call parity relation it can be seen this spread is related to skewness.

Regression results indicate that the implied volatilities and correlations can predict future volatilities and correlations significantly better than their historical counterparts although for correlations the improvement is very small. The measures IRVS, MFIS and CPVS significantly predict stock returns on the daily, weekly and monthly horizon. In the next step all measures are used to construct portfolios and the out of sample performance is compared to various benchmarks. To reduce the effect of estimation error mean-variance parametric portfolios are constructed as in Brandt, SantaClara and Valkanov (2005).

Option implied correlations do not help to improve the portfolio performance and implied volatilities improve portfolio risk but not the Sharpe ratio. The volatility risk premium and the option implied skewness improve the Share ratio but worsen turnover: when these variables are added to the Fama French factors and momentum the out of sample Sharpe ratio increases from 1.03 to 1.39. This construction strategy is sufficiently profitable to offset transaction costs when executed on a weekly basis, and even more so when executed at the monthly frequency.



Econometric Measures of Systemic Risk in the Finance and Insurance Sectors

Mila Getmansky

The presentation on systemic risk by Mila Getmansky is joint work with Lorian Pelizzon, Monica Billio and Andrew Lo. It studies the degree of interconnectedness between financial institutions as that amplified the problems during the financial crisis of 2007-2009. Systemic risk is a function of connections and the network structure of financial institutions. It captures the risk of a spill-over of large losses, large negative returns, or distress from one institution or a group of financial institutions to almost all other institutions that belong to the system, threatening the system as a whole. The spill-over results in an increase in return correlations.

As the spill-over can be instantaneous and lagged (a cascade), different techniques are used for systemic risk in this study: principal components and Granger causality tests. The financial institutions under consideration are the 100 largest from the categories: hedge funds, banks, brokers and insurance companies although the distinctions between these have blurred due to financial innovation and deregulation. For hedge funds the monthly net returns come from the TASS database and for the other institutions the monthly returns come from CRSP. The sample period runs from 1994 to 2008.

The principal components systemic risk measure (PCAS) captures the contribution and the exposure of each institution to the overall risk of the system given a strong commonality across returns of institutions. In the 2006-2008 crisis period the 10 largest principal components capture a cumulative risk fraction of 89% while in the tranquil period of 1994 to 1996 the cumulative risk fraction of these components is only 71%. This suggests increased interconnectedness during the recent crisis. A similar increase is observed for the period in which the LTCM crisis occurred. In both crisis periods the PCAS score for hedge funds is relatively small compared to the other institutions, suggesting hedge funds didn't have much impact during the crisis.

The linear Granger causality tests are to explain causality relations in (lagged) returns. As informationally efficient markets should not exhibit Granger causality, the degree of Granger causality can be used as a proxy for spill-over effects as argued in other papers. The number of significant causality relations as a fraction of all possible connections increased from 6% in the tranquil 1994-1996 period to 13% in the 2006-2008 crisis period. In the latter period banks and insurance companies are more important sources of systemic risk as their returns affect hedge fund returns and broker returns significantly, but this relation is asymmetric i.e. not reciprocal.

Regression results show that the maximum losses incurred by financial institutions during the crisis of 2007-2008 were predictable by the PCAS measure and network measures of interconnectedness observed in the period prior to the crisis (2002-2005). So these new measures could be used as early warning signals of financial crises.



Ex-post Portfolio Performance with Predictable Skewness and Kurtosis

Giovanna Nicodano

The presentation by Giovanna Nicodano is about work with Massimo Guidolin. Sector and country stock returns display excess skewness and kurtosis relative to the normal distribution. In this paper the added value of using this information in portfolio construction is evaluated on an out of sample basis. In a framework with bull and bear dynamics the additional performance gains from timing higher order moments are compared to a strategy with timing of the bull and bear market regimes.

The number of states in the regime switching framework is optimized, not predetermined. Each state has its own set of returns and covariances. Preferences are moment based which facilitates the inference of the optimal portfolio weights and allows to disentangle preferences for mean, variance, skewness and kurtosis of terminal wealth.

The following strategies are considered:

- A single state model with no predictability and mean-variance preferences (MV1)
- A multi state model with predictability and mean-variance preferences (MV2)
- A multi state model with regime and moment predictability and mean-variance-skewness preferences (MVS2)
- A multi state model with regime and moment predictability and mean-variance-skewness-kurtosis preferences (MVSK2)

The out of sample performance of the strategies is evaluated using the Sharpe ratio, the Sortino ratio and certainty equivalent returns. The equally weighted portfolio serves as a benchmark. Three equity datasets at the monthly frequency are used: MSCI international in USD from 1998 to 2008, CRSP US industry returns from 1926 to 2008 and MSCI/Bloomberg international book to market portfolio returns from 1975 to 2007.

In the two state model, the weights of the optimal portfolios differ substantially between the bull and bear state. Optimally the investor goes long cyclical industries in the bull state and shorts these in the bear state. Moving from the single state model to the multi state model increases the weights assigned to North American stocks from 0% to over 25% so the choice of model has a considerable impact. The investor with MVSK2 preferences will optimally seek exposure to many international regions compared to an investor with MV2 or MVS2 preferences. The weights in the two state strategies fluctuate quite a lot while the weights of the MV1 strategy are quite persistent.

For the subset of international returns, where non-normalities are weak ex-ante, there is support for embedding regimes in the optimization as for some investment horizons these strategies beat the MV1 and the equally weighted benchmark on most metrics, especially the certainty equivalent return measure. For the 5 year horizon the top performing strategy outperforms the equally weighted benchmark by 5.86% per year.

For the subset of industry returns, where non-normalities are more important, the MV2 and the MVSK2 strategies perform well and the equally weighted benchmark performs worst. For the 5 year horizon the top performing strategy outperforms the equally weighted benchmark by 3.18% per year.



For the subset of book-to-market returns, MVSK dominates at the 1 month horizon, but closely competes with MVS at longer horizons. For the 5 year horizon the top performing strategy outperforms the equally weighted benchmark by 6.4% per year.

Planned extensions to this paper are incorporating transaction costs, extending the sample period to include the financial crisis and tests for statistical differences of the performance measures.



Do Foreigners Know Better? A Comparison of the Performance of Local and Foreign Mutual Fund Managers

Joao Pereira

The presentation by Joao Pereira is joint work with Miguel Ferreira and Pedro Matos. The authors study the following setting: suppose an investor from country X wants to invest in stocks from country Y. Can he/she better invest in a fund based in country Y (a domestic fund) or in a fund outside of country Y (a foreign fund)? Evidence from research on this topic is mixed and differs per country.

Monthly returns (in USD) and quarterly fund holdings are used in the analysis for open end mutual funds of 29 countries from 2000 to 2007. Holdings are retrieved from the FactSet/LionShares database and is merged with stock data from Lipper Hindsight and Datastream/Worldscope. Considering the S&P500, the sample period can be summarized by a bear market from 2000 until 2002 followed by a bull market.

Each individual portfolio holding is classified as domestic or foreign based on the location of the fund management company. Alternative classifications will be considered as well. On average the fund holdings amount to 10% of the market capitalisation. For the U.S. 80% of these 10% are classified as domestic and 20% as foreign.

The return on domestic holdings are compared to the return on foreign holdings using various excess return models such as the global four factor model of Carhart (1997) and the global characteristics adjusted return model of Daniel, Grinblatt, Titman and Wermers (1997). At the country level the results are inconclusive which is in line with the mixed evidence found in the existing literature. More or less regardless the excess return model, the excess return of foreign holdings is about 0.5% larger (at the monthly level) than the excess returns of domestic holdings. This result is mostly driven by the bull market period.

At the country level, foreign holdings perform better in countries with common law (instead of civil law), in countries with higher GDP per capita, in countries with the English language and in countries with insider trading enforcement. In countries without insider trading enforcement, domestic holdings perform better. At the stock level, foreign holdings perform better in “visible” stocks: in large stocks, in MSCI constituent stocks, in stocks with a large number of analysts and stocks with high foreign sales.

Perhaps financial globalization has contributed to level the playing field for international investors. Large foreign investors may have more resources, more skills, better payed managers and more sophisticated models to achieve better performance. To test the comparative advantage of global companies, holdings are classified as global (if the parent company has more than 5 funds and is present in more than two countries) and as local otherwise. Indeed, globally held assets outperform local holdings by 0.15% per month, again especially for the “safer” countries and the more visible stocks.