FAMe

FINANCE & ACCOUNTING MeMOS

jagazine

Free E-Book www.fame-jagazine.com Free Web Online

Issue 3, 2016

Private Equity Performance

Investment Wisdom From Social Media

Investor Attention Through Advertising

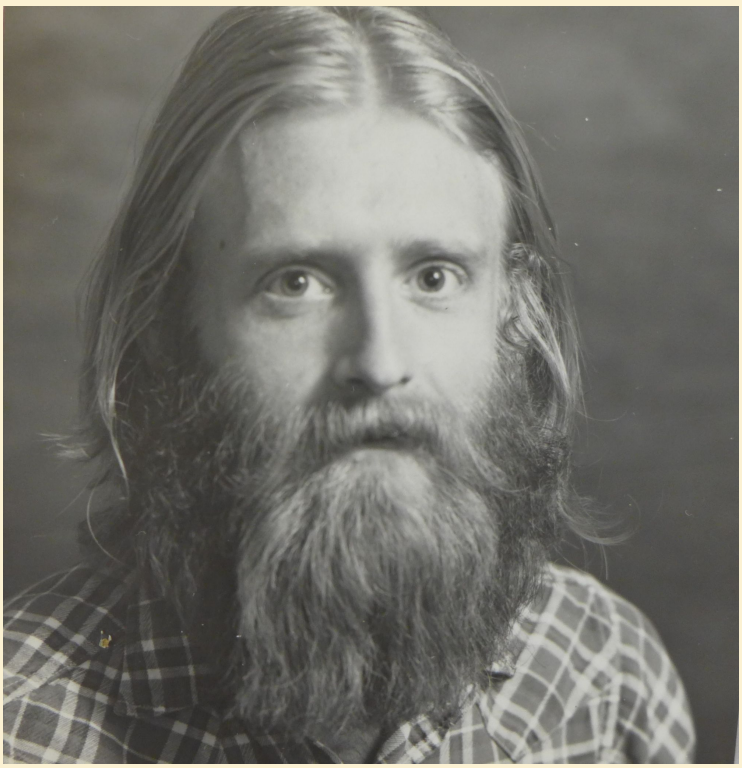

Social Good And Insider Trading

Genetics of Investment Biases

Housing Risk Diversification

and more

www.fame-jagazine.com

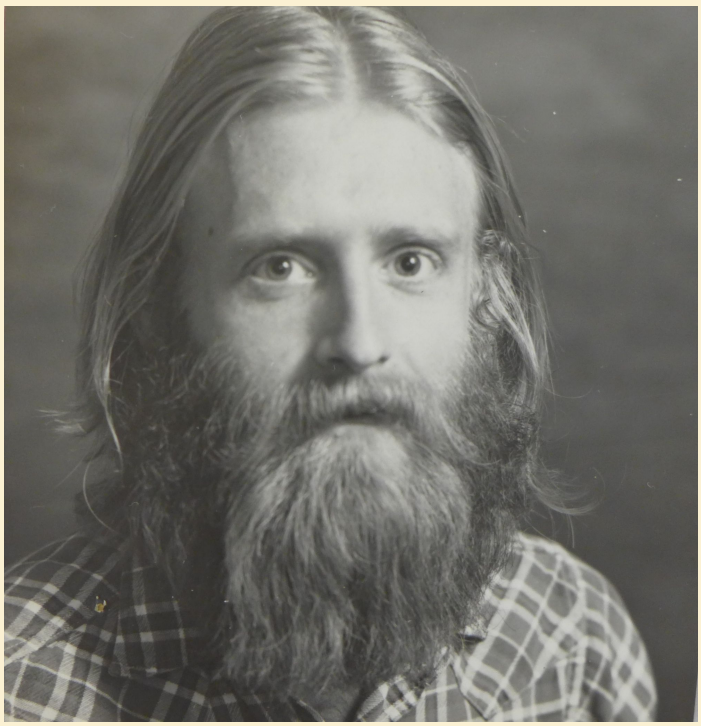

Rick Green

Advice to the (e-)reader

FAME is available in pdf form, html form, and epub form. Each format has advantages and disadvantages:

- Pdf is the best format for printing. The pdf layout is fixed. FAME can control both fonts and the placement of material. Pdf is not bad online, either, with very high-quality font rasterization on all platforms. However, it cannot re-flow the text according to the reader's font-size preference and viewing device. If you want to view a pdf page that is twice as wide, you need a canvas that is twice as wide.
- Html is the best format for on-line viewing on a high-speed-internet-connected *computer*. It is usually easy to change the browser font-size (using ctrl-minus or ctrl-plus) with a concomitant re-flow of the document. Unfortunately, the most common *tablet* web browsers (like Safari on iOS, Chrome on Android, and Internet Explorer on Windows-8) make font-size changes difficult. Pinching html text in tablets usually zooms into the document (just as it does in a pdf file), instead of re-flowing the text.
- Epub is the best format for tablets. Like pdf and unlike html, it works well for viewing on non-Internet connected devices, too. (Epub is basically pre-packaged html without javascript.) Unfortunately, many e-book readers are still "immature" (buggy). FAME has been tested to work well on Apple iBooks, Firefox's epubreader, and Adobe's Digital Editions. calibre and the Android Uiversal Book Reader mostly work, but have several bugs. (Amazon makes it intentionally difficult to transfer content that was not obtained through the Amazon store onto the Kindle reader.) Note that external links cannot be made to work reliably inside an ebook reader—epub is a self-contained package format. Some e-book readers will open a web browser; others will not.
- Source contains the LaTeX-dialect source from which all of the above have been created. This is of interest only to connoisseurs of typesetting systems.

In both browser html and ereader epub format, it is the client program and not the FAME team that handles font-size, line-breaking, and page-breaking. This is often but not always good. For example, tables and figures may be broken at *very* inopportune spots. Fortunately, when the reader resizes the content on a strange-looking page, the page often suddenly looks great. (linux users: please install msttcorefonts or ttf-mscorefonts-installer.)

Academic articles often include tables that require a wide page for comprehension. For this reason, FAME articles are not well-suited to reading on small-screen devices. A 10-inch diagonal high-resolution tablet screen is recommended, although a 7-8-inch high-resolution diagonal screen size may be acceptable. (Also, note that wide-screen is great if your e-reader does not decide to abuse it for two-column reformatting, in which case it becomes narrow-screen.) Please, do not think about reading FAME on a 3.3-inch diagonal iphone. It would be an exercise in frustration.

Small text font-sizes often work better with wide tables, but large font-sizes often look better on the main text in the absence of tables and figures. If you can easily switch font-size on the fly, you can get the best of both worlds.

FAMe editorial board

Submissions, advertising, and support

Advertisements and support

For advertisement inquiries, please email fame.jagazine@gmail.com.

FAMe's goal is to broaden the impact of academic finance and accounting research. Creating an issue of FAMe costs about $50,000, largely financed out of pocket by ourselves. We depend on the goodwill of authors, readers, and supporters. We cannot spread costs over a large subscriber base—we are not *The Economist*.

To help underwrite our efforts, please consider purchasing a $50 or $100 annual subscription, especially if you can cover this subscription expense from your research budget. There is a paypal button on our website (click here). The paypal receipt should make it easy to be reimbursed as a research expense. The **library** subscription fee is the same price as the individual subscription fee (and paid the same way). The $100 subscription fee entitles libraries to make the content available to their patrons online.

Instructions to authors

To submit a FAMe version of your article recently published or forthcoming in a "pre-approved" finance journal, please send your proposed short memo to: famesubmission@gmail.com. For papers from accounting journals, please use famesubmission+a@gmail.com, instead. For more information and detailed instructions on how to prepare a MeMO, please refer to FAMe guidelines for authors on the www.fame-jagazine.com website for more information. Read our most recent issue to understand our "flavor"—we are still evolving, too. Detailed instructions for accepted submissions are posted at http://www.fame-jagazine.info/.

Contact

If you have contacted us at one of the above email addresses (fame.jagazine@gmail.com), e.g., famesubmission@gmail.com, and have not heard back from us, please feel free to prod us again at ivo.welch@gmail.com and bhagwan@anderson.ucla.edu.

Featured MeMos

Bhagwan Chowdhry, Executive Editor

Editorial

The Rick Green Issue
(please cite only the original publication, not FAMe)

Issues 3 and 4 of FAMe are AFA Presidents issues. You will see pictures of past AFA presidents, both when they were young (with permission) and when they were adults. Click on the pictures to find out more about them.

The production of FAMe is a love of labor for us. We are looking for support, funding and sponsorship to be able to delegate some tasks to paid hired staff editors. Right now, we have to do everything ourselves. With more support, we could increase the frequency of publication for FAMe. Of course, this also depends on the desire of scholars to continue to write and submit memos.

We are dedicating Issue 3 of FAMe to the memory of Rick Green.

Bhagwan Chowdhry
Executive Editor

www.fame-jagazine.com
in pdf, html, and ebook format!

Robert S. Harris, Tim Jenkinson and Steven N. Kaplan

Private Equity Performance: What Do We Know?

Journal of Finance | Volume 69, Issue 5 (Oct 2014), 1851–1882
(please cite only the original publication, not FAME)

Prior research on the performance of private equity investing (including both buyout and venture capital funds) has led to mixed conclusions. This is perhaps not surprising given concerns about the lack of high quality data for research. For instance, some private equity (PE) data sets rely on voluntary reporting by funds or on Freedom of Information Act requests which may result in biased samples and incomplete data. We use a new research-quality data set of PE fund-level cash flows from Burgiss. The data are derived entirely from institutional investors (the limited partners) for whom Burgiss systems provide record-keeping and performance monitoring services. This feature results in detailed, verified, and cross-checked investment histories for nearly 1,400 PE funds derived from the holdings of over 200 institutional investors.

Our research highlights the importance of high quality data for understanding PE and the returns it provides to investors. Here we focus on results using the public market equivalent (PME) method of Kaplan-Schoar (JF 2005). PME is a market-adjusted multiple which compares how much a PE fund investor actually earned net of fees to what the investor would have earned in an equivalent investment in the public market. A PME above 1.0 signals higher returns to private compared to public equity. After summarizing conclusions, we present key empirical results.

Buyout fund returns have exceeded returns in public markets

Our results are markedly more positive for buyout funds than research has previously documented. It seems likely that buyout funds have outperformed public markets, particularly the S&P 500, net of fees and carried interest, in the 1980s, 1990s, and 2000s. Our estimates imply that each dollar invested in the average buyout fund returned at least 20% more than a dollar invested in the S&P 500. This works out to an out performance of at least 3% per year. The conclusion that there has been out-performance is relatively insensitive to assumptions about benchmark indices and systematic risk. For the more recent and less fully realized vintage funds, however, eventual performance will depend on the ultimate realization of their remaining investments. Our results (and those we estimate from other commercial data sets) imply that buyout funds outperformed public markets much more substantially gross of fees. Nailing down the sources of this out performance will be a fruitful subject.

The performance of venture capital funds has changed over time.

Venture capital (VC) funds outperformed public markets substantially until the vintages of the late 1990s, but have underperformed since. Since the 2000 vintage year, the average VC fund has underperformed public markets by about 5% over the life of the fund. Although disappointing, this

underperformance is less dramatic than the more commonly quoted absolute return measures. Extant research focuses on the earlier vintage years and inevitably obtains more positive results. Again, the qualitative conclusions do not appear sensitive to assumptions about systematic risk.

Private equity performance is negatively related to the inflow of capital into the asset class.

Vintage year performance for buyout and VC funds decreases with the amount of aggregate capital committed to the relevant asset class, particularly for absolute performance (internal rates of return and investment multiples), but also for performance relative to public markets. This suggests that a contrarian investment strategy in these asset classes would have been successful in the past. The magnitudes of these relations have been greater for VC funds. Why these patterns have persisted is something of a puzzle and an interesting topic for future research.

Estimated PMEs from other data sets confirm our results on private equity performance.

Using Burgiss data within a given vintage year, we find that fund PMEs are reliably related to the more generally available absolute performance measures, internal rates of return (IRRs) and investment multiples. For both buyout and VC funds, IRRs and investment multiples explain at least 90% of the variation of PMEs in most vintage years, with investment multiples explaining substantially more of the variation than IRRs. As a result, researchers and practitioners can use our models to estimate PMEs without having the underlying fund cash flows. We do this to estimate average PMEs from three other commercial data sets on private equity, Cambridge Associates, Preqin and Venture Economics (part of Thomson Reuters). The Burgiss, Cambridge Associates, and Preqin data sets yield qualitatively and quantitatively similar performance results. There is little reason to believe that the Burgiss and Preqin data sets, in particular, suffer from performance selection biases in the same direction. At the same time, consistent with Stucke (WP 2011), we find that performance, particularly for buyout funds, is markedly lower in the Venture Economics data. This confirms that academic research and practitioners should be cautious in relying on Venture Economics data.

Concluding Comments

Our findings strongly suggest that buyout funds have outperformed the public equity markets net of fees over most of our sample period. To invalidate that conclusion, all three reliable commercial data sets would have to be subject to a similar and large positive selection bias despite very different data collection and reporting methods. We view this as highly unlikely. Instead, we view the similar results to indicate that all three databases provide unbiased estimates of the overall performance of PE.

Because PE investments are illiquid, it is perhaps not surprising that they yield investors some premium relative to investing in public markets. As well as the relatively illiquid nature of PE investments, there is also uncertainty regarding how much to commit to PE funds to achieve a target portfolio allocation. This is due to the uncertain time profile of capital calls and realizations. Consequently, "commitment risk" exists when investing in PE. This contrasts with investing in public markets where there is no distinction between capital committed and invested, and trading is continuous. The cost of illiquidity or commitment is likely to vary across investors, and remains an important area for research.

Stephen L. Lenkey

Advance Disclosure of Insider Trading

Review of Financial Studies | Volume 27, Issue 8 (Dec 2015), 2504–2537
(please cite only the original publication, not FAME)

Current law requires certain corporate insiders (officers, directors, and large beneficial owners) to publicly disclose their trades within two business days *after* the trades are made. This mandate facilitates the flow of information to the rest of the market by creating an opportunity for ordinary investors to glean information from insiders' trades. However, the law appears to fall short in preventing insiders from earning insider-trading profits because the trades are revealed only after insiders have had an opportunity to capitalize on their information advantages. Might it not be better, then, to require insiders to disclose their trades *before* they occur? I show that welfare increases for both insiders and ordinary investors when insiders are required to disclose their trades in advance.

Theoretical framework

I evaluate the desirability of potential regulations that would require insiders to disclose their trades in advance using a strategic rational expectations equilibrium framework. In the model, there is a lone insider who possesses an information advantage over a continuum of less-informed investors. The insider can trade on the basis of her private information to earn insider-trading profits. Before she trades, however, the insider must publicly disclose her trade. This disclosure provides a noisy signal of the insider's private information from which ordinary investors can learn about the insider's information. A nontradeable random endowment received by the insider but unobservable to investors prevents the insider's private information from being fully revealed in equilibrium. As a benchmark against which to evaluate the impact of advance disclosure, I consider an alternative setting in which the insider is free to trade without providing a pretrade disclosure. In the benchmark setting, the stock price acts as a noisy signal of the insider's private information.

Advance disclosure improves risk sharing

Under an advance disclosure requirement, the insider must commit to a trade (and disclose it) without knowing the exact price at which the transaction will occur. This creates greater uncertainty for the insider in the form of price risk. To mitigate this risk, the insider trades less aggressively on her private information, which reduces adverse selection costs incurred by ordinary investors that arise from trading with a better-informed counterparty. Consequently, ordinary investors are more willing to trade when the insider predisclosures. Not only does this enhance the insider's ability to hedge her nontradeable endowment, but it also allows the risk in the economy to be shared more efficiently.

Figure 1: Risk Sharing

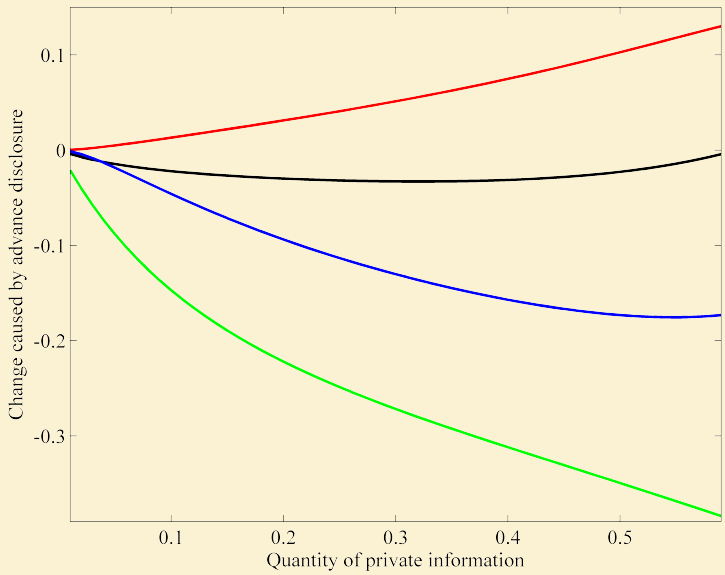

Changes in trading aggressiveness (black), hedging ability (red), and the price responsiveness to the insider's information (green) and nontradeable endowment (blue) caused by advance disclosure. We interpret this to mean that investors face lower adverse selection costs and risk sharing improves with advance disclosure.

Figure 1 decomposes the impact of advance disclosure on the price and allocations into informational and risk-sharing effects. The figure shows that the insider trades less aggressively on her private information but more aggressively hedges her nontradable endowment when she predisclosures her trade. Additionally, the price is less sensitive to both the insider's private information and her nontradeable endowment. This stems from the fact that the insider's hedging motive is stronger, whereas her information-based motive for trading is weaker with advance disclosure. The improvement in risk sharing is ultimately manifested in the equilibrium price and allocations. When the insider predisclosures her trade, the price and allocations more accurately reflect the optimal level of risk sharing that would be achieved under a symmetric-information equilibrium.

Advance disclosure also results in a less-efficient market. Even though investors have knowledge of the insider's trade before it occurs, the market is less efficient because less of the insider's information is revealed through her trade when she predisclosures. The insider's trade does not convey as much information because she trades less aggressively on her private information.

Better risk sharing enhances welfare

Advance disclosure tends to increase welfare for the insider, as shown in Figure 2. Despite the fact that making a pretrade disclosure reduces the insider's ability to capitalize on her information advantage, the insider's welfare rises because the improvement in risk sharing enables the insider to better hedge her nontradeable endowment and hold a portfolio that is better aligned with her risk preferences. The insider's welfare increases because the benefits of improved risk sharing outweigh the costs associated with her diminished ability to earn insider-trading profits.

Figure 2: Welfare

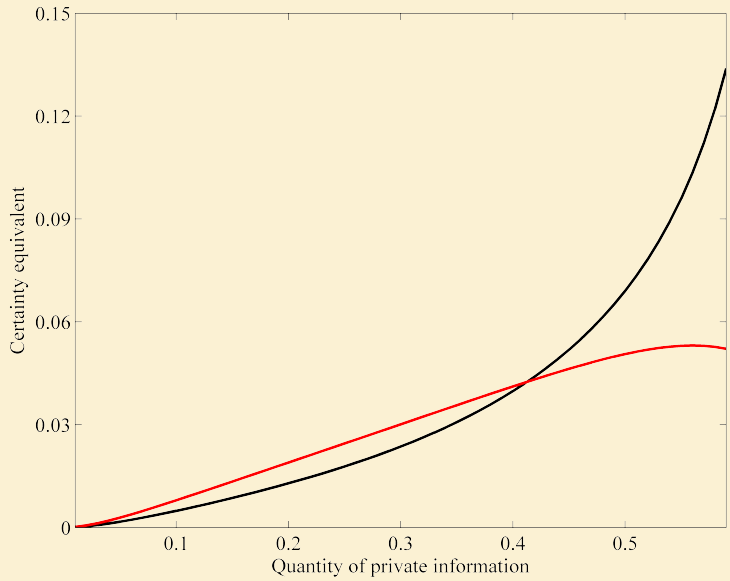

This shows certainty equivalents of advance disclosure for the insider in black and for ordinary investors in red. Advance disclosure increases welfare.

Ordinary investors also experience a rise in welfare when the insider prediscloses her trade, as indicated by Figure 2. This occurs because the improvement in risk sharing enables investors to hold portfolios that are better aligned with their risk preferences, even though the market is less efficient with advance disclosure. Investors' welfare increases because the benefits of improved risk sharing offset the greater uncertainty faced by investors.

Advance disclosure alters managerial incentives

Because advance disclosure alters both the information environment and the equilibrium allocations, requiring insiders to disclose their trades in advance affects managerial incentives in at least two ways. First, pretrade disclosure encourages risk taking. Advance disclosure strengthens the incentive for a manager to undertake risky projects because the market is less efficient and, therefore, private information is more valuable when insiders disclose their trades in advance. Second, pretrade disclosure may either encourage or discourage managerial effort. Advance disclosure either strengthens or weakens the incentive for a manager to exert costly effort because the personal benefits from exerting effort depend on an insider's stake in the equity of the firm, and an insider's equilibrium allocation may either rise or fall under advance disclosure.

Ulf von Lilienfeld-Toal and Stefan Ruenzi

CEO Ownership, Stock Market Performance, and Managerial Discretion

Journal of Finance | Volume 69, Issue 3 (Jun 2014), 1013–1050
(please cite only the original publication, not FAME)

What is the impact of CEO ownership on stock market returns? We examine this question by considering investment strategies which are based on publicly available CEO ownership information. Strategies which select firms into a portfolio based on CEO ownership outperform the market significantly by up to 10% per year. The cumulative abnormal equal- and value-weighted return of a strategy that invests in a high managerial ownership firms is shown in Figure 1. This outperformance is also observed after controlling for standard risk factors (e.g. Fama-French-Carhart factors). We can also show that this result is robust and holds in a variety of specifications, models, samples, and time horizons.

A puzzle?

These results are puzzling as we only use publicly available information to set up our investment strategy: The information on ownership information of CEOs has to be disclosed and we built our portfolios based on past disclosures, that is, based on public information. However, in an efficient market, all publicly available information should be reflected in the share price and—according to standard theories—we should thus see no impact of such information on future returns.

Possible explanations

In the second part of our paper we then investigate three possible explanations why we still see abnormal returns for owner-CEO firms. The first potential explanation uses the idea that owner-CEOs are better informed about the prospects of the firm and are only willing to hold large ownership stakes when they have positive private information. At the same time, this explanation assumes that the market is not fully efficient, that is, the market is not able to correctly interpret the positive information which is signaled to the market due the disclosure of the CEO ownership information. The second and third explanation build on the argument that managerial ownership is an incentive device and imply that owner-CEOs can be value increasing: it is assumed that owner-CEOs have strong incentives (due to their high ownership stake) to work hard, exert costly effort, and enough discretion to eventually increase the value of the firm. The two explanations differ with respect to the view on market efficiency. The second explanation builds on the assumption that the market is irrational and does not fully understand the incentive effects and the eventually value increasing impact of managerial ownership. Similarly as above, because the market does not interpret the signal of high disclosed ownership information correctly, stock prices at best partially reflect the positive effect of high CEO ownership. Thus, abnormal returns emerge only subsequently, when the consequences of the positive incentives

effects become visible in the form of an increased firm value. The third explanation in contrast assumes that the market is fully rational and understands the value increasing impact of managerial ownership.

This final explanation is motivated by recent theories about trading of value increasing shareholders developed in Lilienfeld-Toal (WP 2010) and Blonski and WP Lilienfeld-Toal (2009). These papers argue that if a value increasing shareholder can trade her shares on the open market prior to exerting effort, an equilibrium for firms with high CEO ownership will typically be characterized by positive abnormal returns. The result why the information is not fully reflected in prices even in an efficient market is based on the argument, that a CEO could sell her shares at the high price which already reflects her high future effort. Selling shares is possible for the CEO, because the high ownership fractions we focus on in this paper (typically requiring a cutoff of 5% or 10% of the firm's outstanding shares being held by the CEO) are overwhelmingly voluntary, i.e., the CEO is not restricted from selling. However, if she did sell, her incentives to work hard (which justified the high price) would be destroyed and the initially high price would not have been a rational equilibrium to start with. In contrast, a situation with CEO ownership and positive abnormal returns can be an equilibrium in this setting. The reason why the apparent underpricing does not get arbitraged away by investors in such a situation is that all investors are fully rational in this model and thus know that they would bid up the price by buying the underpriced shares, thereby eventually triggering the CEO to sell and thus making everybody worse off.

Asymmetric information or incentives?

In our further analysis, we try to differentiate between the three explanations, i) asymmetric information and irrational markets, ii) positive incentive effects of CEO ownership and irrational markets, and finally iii) positive incentive effects of CEO ownership and rational markets. We find no evidence for the first explanation: various proxies for asymmetric information have no impact on our results. However, we do find strong evidence that managerial incentives play an important role in explaining the positive impact of CEO ownership on stock returns, that is, we find strong support consistent with explanations ii) and iii). Specifically, we argue that managerial incentives only matter if the CEO has some discretion to increase the firm's value. Otherwise, exerting effort would not result in a higher firm value. Consequently, we split our overall sample of firms into firms in which we have reason to believe that CEOs have a lot of discretion and firms in which the CEO has no discretion. We find that the positive impact of managerial ownership is particularly high for firms in which the CEO has a lot of discretion. At the same time, high managerial discretion without high managerial ownership is not associated with abnormal returns. Similarly, high managerial ownership but little discretionary power of the CEO also does not lead to the large abnormal returns we observe when CEOs have high ownership and at the same time a lot of discretion. We think this is an interesting result in itself, as the proxies we use for managerial discretion (like her power, free cash flows, little pressure from outside investors) are often interpreted as signs of bad governance. We show that the impact of these proxies on firm performance can be reversed, if we at the same time observe high managerial ownership, that is, if CEOs are incentivized to use their freedom in the right way from the shareholders point of view.

Rational vs. irrational markets explanation

To differentiate between the second and third explanation we need to understand whether markets are rational and understand the incentive mechanisms underlying managerial ownership or whether the market is irrational and does not fully understand this. In an effort to do so, we conduct a series of additional tests: first, we show that there is not much evidence of learning over time as the abnormal returns due to managerial ownership are fairly evenly distributed over time. This is highlighted in Figure 1 and speaks against an irrational markets explanation, because even if the market might not have understood the positive message send by high ownership initially, investors should eventually learn about this over time. We find no evidence for this to be the case.

Figure 1: Cumulative abnormal returns of high CEO ownership firms

This figure plots the cumulative abnormal return of a portfolio which goes long in high managerial ownership and short in low managerial ownership firms, where high ownership firms are defined based on a cutoff of 10% of all outstanding shares and low ownership firms are defined as firms where the CEO owns 1% or less of the outstanding shares. This figure is taken from Lilienfeld-Toal and Ruenzi (2014). **Interpretation:** High CEO ownership firms consistently and strongly outperform the market

Next, we make sure that our results are not only driven by small and opaque firms and firms that are generally characterized by strong limits to arbitrage and where inefficient valuations can persist. Given that it is more costly to set up a trading strategy in small stocks, such limits would be more severe among small firms. However, the evidence in Figure 1 shows that our results are even stronger based on a value-weighted portfolio—which puts a larger weight on large firm observations—as compared to an equal-weighted portfolio. Finally, we look at earnings surprises and find that investors are at best very mildly surprised by earning announcements. If investors would simply be unaware that firms with high ownership will perform well, they should continuously be positively surprised by the earnings of these firms.

What do firms with high managerial ownership do differently?

In an attempt to better understand what drives the good performance of high managerial ownership firms, we also examine firm policies. We find that firms characterized by high CEO ownership engage less in empire building and are able to run their firm more efficiently. Given that it is cumbersome and requires a lot of effort to run a firm efficiently, these results are also consistent with the abnormal returns we document being driven by incentives of owner-CEOs.

To summarize, in our article, we show that high-managerial ownership is associated with high abnormal stock market returns. We find convincing evidence that managerial discretion is crucial for managerial ownership to lead to high abnormal returns and that the combination of strong incentives and managerial latitude seems to be optimal.

Giovanni Cespa and Thierry Foucault

Illiquidity Contagion and Liquidity Crashes

Review of Financial Studies | Volume 27, Issue 6 (Jun 2014), 1615–1660
(please cite only the original publication, not FAME)

Liquidity (the impact of trades on prices) fluctuates over time and these fluctuations are correlated across assets (see Holden-Jacobsen-Subrahmanyam (FTiF 2014) for a survey). This correlation ("co-movement in liquidity") is a source of non-diversifiable risk and this risk is priced (see, for instance, Acharya-Pedersen (JFE 2005) or Pastor-Stambaugh (JPE 2003)). There are yet very few theories explaining why liquidity co-moves across assets. We provide such a theory.

Learning from prices leads to interconnected liquidity

Liquidity suppliers (e.g., market makers) specialized in one asset often learn information from observing the price of related assets because asset prices contain information (Grossman-Stiglitz (AER 1980). Consider for instance a market maker in a stock part of the S&P500 index. This market maker will follow prices of futures on the S&P500 index or ETF on the same index because these prices contain information on risk factors common to the stock and these indexes. Similarly, the market maker in a bond or a stock of a given firm will extract information from prices of CDS written on this firm.

Liquidity suppliers in one asset, say X, face therefore less uncertainty when price movements of other assets are more informative. In this case, they bear less inventory risk and can provide liquidity at better terms. Price informativeness in turn depends on liquidity. Indeed, at a relatively high frequency, price movements are due both to news about fundamentals and demand shocks. For instance, a large sell market order pressures prices downward because liquidity suppliers require a compensation to take additional risk (see Hendershott-Menkveld (JFE 2014). As liquidity suppliers gradually unwind their position, this price pressure disappears. Price movements like this, due to transient demand or supply shocks, do not contain information on fundamentals and make prices noisy. The signal-to-noise ratio in prices is therefore smaller when the noise due to transient price pressures is higher, that is, when illiquidity is higher (see Ait-Sahalia-Yu (Arxiv 2009) for empirical evidence).

Thus, when liquidity suppliers learn from the prices of other assets, the liquidity of different assets is interconnected: The liquidity of one asset is higher when the liquidity of other assets is higher because their prices are then more informative. We formalize this interconnection and analyze its implications.

Interconnected liquidity leads to contagion, amplification, and fragile liquidity

Interconnected liquidity generates liquidity spillovers. Consider Figure 1. An exogenous increase in the illiquidity of asset Y (e.g., a decrease in the risk bearing capacity of liquidity suppliers in this asset) triggers a drop in the informativeness of the price of this asset and thereby a drop in the liquidity of asset X, as if illiquidity was contagious. In turn, this reduces the informativeness of the price of asset X, which can reinforce the drop in liquidity for asset Y. This feedback loop can significantly amplify the initial drop in liquidity of assets Y and X, through a multiplier effect.

Figure 1: Liquidity spillovers and price informativeness

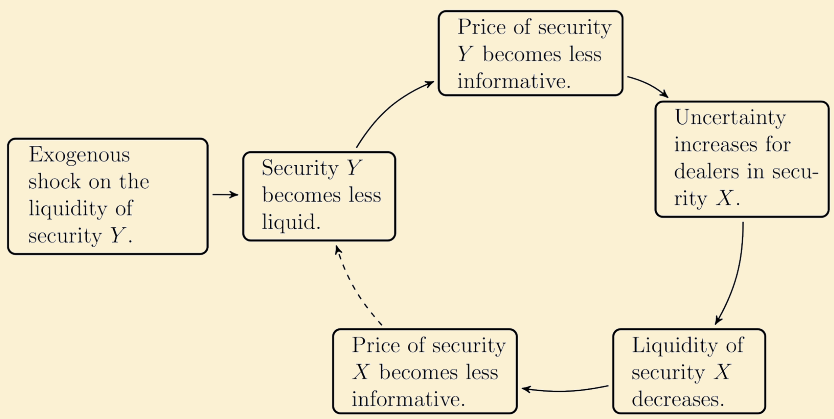

The liquidity of one asset affects the liquidity of another asset, which in turn again affects its own liquidity.

There exist situations in which the multiplier can be very large: a small initial shock to the liquidity of one asset can ultimately trigger a very large drop in liquidity of all assets. For instance, in the paper, we provide examples in which a 1% decrease in the risk tolerance of liquidity suppliers in asset Y triggers a very large percentage increase in the illiquidity of assets X and Y while such a shock would have negligible effects if dealers in each asset were not learning from each others' prices (see Table 1 in our RFS paper). Market liquidity can therefore be very fragile when market makers learn from other asset prices: a small increase in the illiquidity of one asset can ultimately generate a very large increase in the illiquidity of all assets. Thus, the model suggests the possibility of liquidity crashes, that is, market wide evaporations of liquidity, in the absence of noticeable changes in the economic environment, very much as was observed during the 2010 "flash crash."

Implications of broader price dissemination

In our theory, liquidity spillovers and co-movements are due to cross-asset learning: liquidity suppliers in one asset learn from prices of other assets. This mechanism has several implications. First, changes in market structure (e.g., market transparency) or regulations (e.g., a short-sale ban) that directly affect the liquidity of one asset class should also indirectly affect the liquidity of other classes, if dealers in the latter use prices of the former as a source of information. For instance, in 2002, the National Association of Securities Dealers began reporting transaction prices for a subset of corporate bonds ("eligible bonds"). Hence, prices of these bonds became easier to observe for market makers

acting in related assets, e.g., other corporate bonds or stocks of eligible bonds' issuers. Bessembinder-Maxwell-Venkataraman (JFE 2006) show that the liquidity of eligible bonds increased with reporting of their prices. Consistent with our model, this improvement propagated to non-eligible bonds (see Bessembinder-Maxwell-Venkataraman (JFE 2006)) and issuers' stocks (see Yin (IJEF 2011)).

Our theory also implies that the size of liquidity co-movements between assets should be stronger when access to information about asset prices is easier. Figure 2 shows the co-movement between two assets X and Y as a function of the fraction of dealers in asset X with information on the price of Y for specific parameters of our model. When more liquidity suppliers in asset X can observe the price of asset Y (e.g., because the market of this asset is more transparent or because of progress in information technologies), liquidity co-movements between the two assets become stronger. The reason is that the multiplier effect associated with the feedback loop highlighted in Figure 1 is larger.

Figure 2: Liquidity comovements and price dissemination

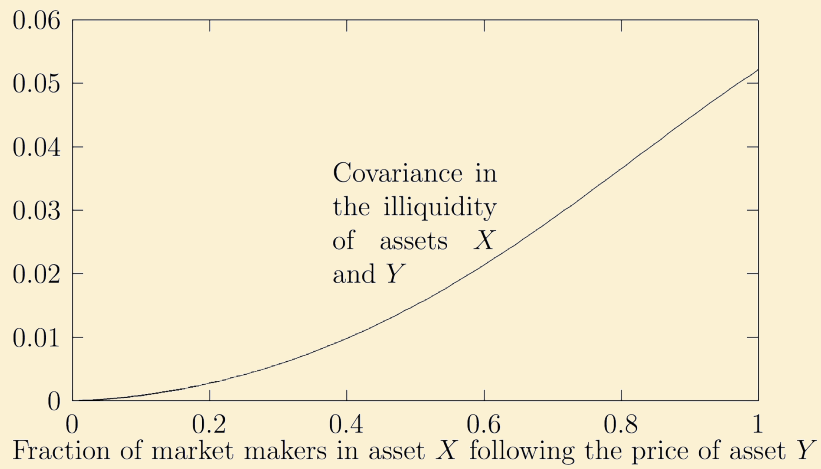

When more liquidity suppliers in asset X can observe the price the asset Y, liquidity comovements between the two assets become stronger.

Funding liquidity or cross-asset learning?

Funding liquidity (i.e., capital constraints) shocks and wealth effects for liquidity suppliers have been proposed as possible sources of covariation in liquidity (see Gromb-Vayanos (JFE 2002), Brunnermeier-Pedersen (RFS 2009), and Kyle-Xiong (JF 2001)). These mechanisms and the cross-asset learning channel described in our paper are not mutually exclusive. In reality, they could therefore work in tandem to generate comovements in liquidity. Their relative importance should vary according to the frequency at which comovements in liquidity are measured: shocks to liquidity suppliers' capital are more likely to play a role at relatively low frequency (daily, weekly, or monthly) whereas the cross-asset learning channel should play out even (and maybe predominantly) at high frequency (intra-daily). In any case, our cross-asset learning mechanism has implications distinct from other theories of comovements in liquidity. These could be used to test whether this channel plays a role in liquidity comovements beyond and above other possible channels.

Conclusion

Asset prices are signals and the informativeness of these signals increases with asset liquidity. In turn asset liquidity increases with the amount of information that liquidity suppliers can learn from prices. This joint determination of price informativeness and liquidity generates a feedback loop whereby small shocks affecting the liquidity of one asset can have large effects on the liquidity of all assets. This is a source of co-movement in liquidity and market fragility.

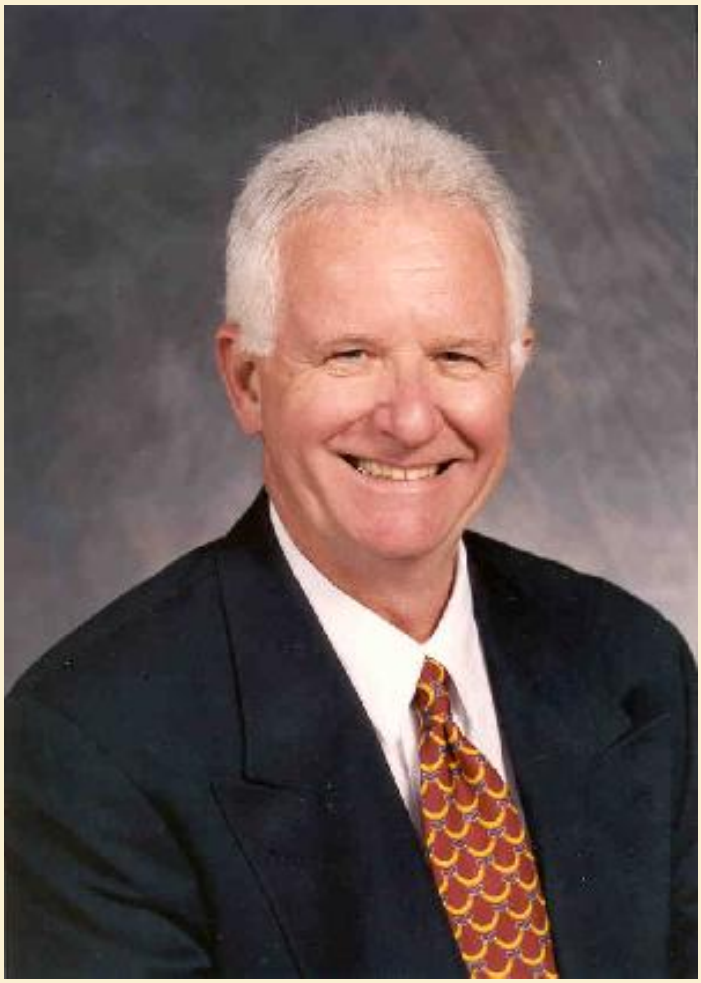

Hailiang Chen, Prabuddha De, Yu Jeffrey Hu, and Byoung-Hyoun Hwang

Wisdom of crowds: The value of stock opinions transmitted through social media

Review of Financial Studies | Volume 27, Issue 5 (May 2014), 1367–1403
(please cite only the original publication, not FAME)

In recent years, with the rapid development of social media, investors have begun conducting their own equity research and sharing their financial analysis of stocks with others on social media. In some cases, these reports are remarkably insightful.

Crowd-sourced stock analyses and opinions

For example, on March 17, 2011 a blogger using the pseudonym Alfred Little published an article about Deer Consumer Products Inc. (DEER) on the crowd-sourced investment research site Seeking Alpha. Based on his/her own research, the blogger claimed that the U.S.-listed Chinese producer of small kitchen appliances had inflated its sales and profit margins and overpaid for land purchases. In three months, the share price of DEER fell by approximately 40%; the stock was eventually delisted from Nasdaq in March 2013 for committing fraud.

In the context of the rapid growth of social media outlets specializing in financial markets, such anecdotal accounts pose an interesting question: Are stock opinions revealed on social media sites valuable, on average, to investors? Arguably, the openness and lack of regulation inherent in social media outlets imply that uninformed actors can easily spread erroneous "information" among market participants. Regulators also warn that e-mail and social media are increasingly used to pump and dump micro-cap stocks in the Internet age. At the same time, it is plausible that large crowds possess valuable information, particularly for stocks neglected by professional analysts and traditional media. Even for stocks covered by sell-side analysts, it is possible that the crowd sometimes conveys insights that are valuable yet not fully factored into sell-side analysts' earnings forecasts and recommendations. Moreover, financial analysts' jobs are fraught with built-in conflicts of interest and competing pressures. In Chen-De-Hu-Hwang (RFS 2014), we assess the performance of investors-turned-advisors and test whether investors can turn to their peers for truly useful investment advice.

Seeking Alpha

To examine the role of peer-based advice, we select Seeking Alpha (SA), one of the largest investment-related social-media websites in the United States, as our data source. Investors can voice their opinions and exchange investment ideas on this site through two channels: (a) Users can submit opinion articles to SA, which are generally reviewed by a panel and subject to editorial changes. If deemed of adequate quality, these articles are then published on the SA website. (b) In response to published articles, any interested user can write a commentary, sharing his or her own view, which may agree or disagree

with the author's view on the company in question. Over our 2005-2012 sample period, SA articles and SA commentaries were written by approximately 6,500 and 180,000 users, respectively, covering more than 7,000 firms.

Wisdom of crowds

To quantify the views disseminated through SA, we employ textual analysis and use the frequency of negative words employed in an article/commentary to capture its tone. Our analyses show that the fraction of negative words contained in SA articles and the fraction of negative words in SA commentaries both negatively predict stock returns over the ensuing three months. We obtain similar results over one-month, six-month, one-year, and three-year horizons. Figure 1 plots the coefficient estimates, along with the corresponding 95% confidence intervals, from regressions of abnormal returns on measures of the views reflected in SA articles and commentaries.

Figure 1: Seeking Alpha and abnormal returns over various holding periods

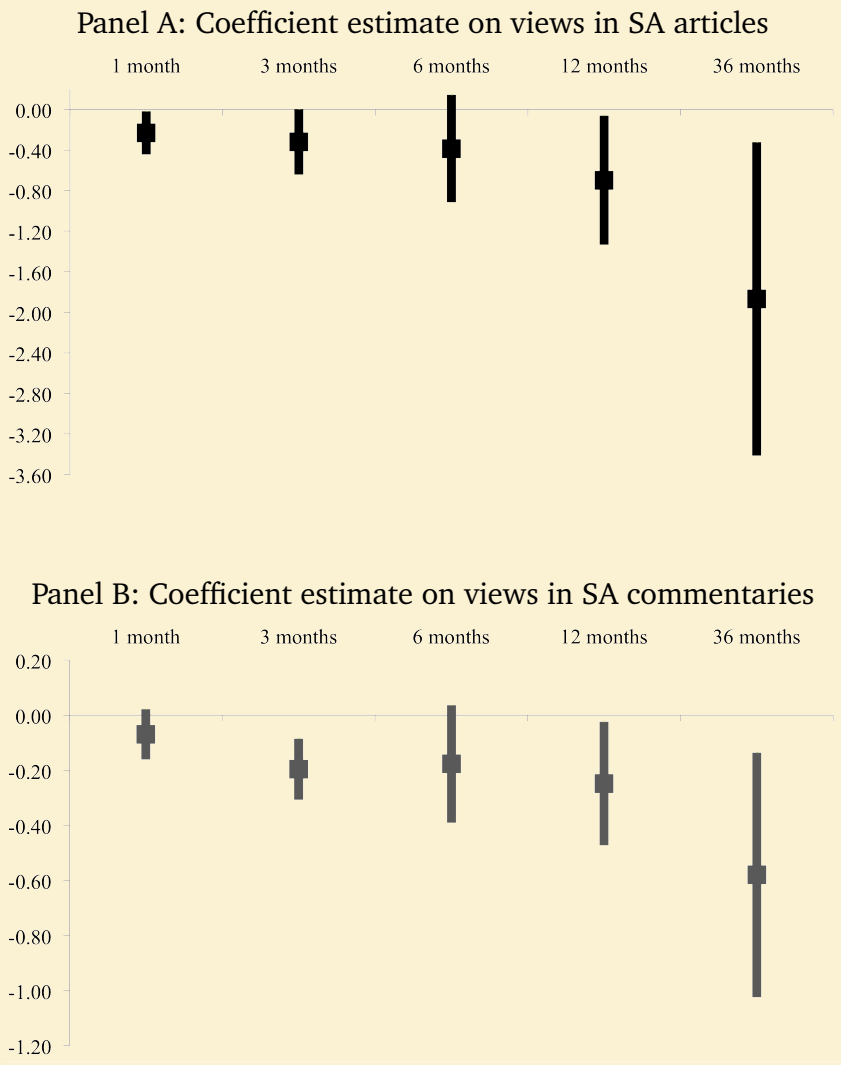

The are coefficient estimates and 95% confidence intervals from regressions of abnormal returns on measures of the views reflected SA articles and commentaries from 2005-2012.

One interpretation of our findings is that the views expressed in SA articles and SA commentaries contain some value-relevant information, which, as of the article publication date, is not fully factored into the price. As investors subsequently adopt the SA view, either through the SA platform itself or through news that arrives following the article publication, the price gradually adjusts. Such an interpretation would point to the usefulness of social media outlets as a source of genuine, value-relevant advice.

An alternative perspective is that SA views produce naive investor reactions. That is, SA views often reflect false or spurious information that nevertheless causes investors to trade in the direction of the underlying articles and commentaries and move prices accordingly. Our methodology (skipping the first two days after article publication and focusing on a three-month horizon) and our observed lack of a return reversal are somewhat at odds with this interpretation. Moreover, it is doubtful that followers of SA have enough capital of their own to cause market prices to move in the manner that we document in this study.

We also examine whether SA views predict subsequent earnings surprises, which is the difference between the reported earnings-per-share (EPS) and the average of financial analysts' EPS forecasts. If opinions expressed through SA that we observed were unrelated to firms' fundamentals or if the information was spurious and already fully incorporated by financial analysts into their reported EPS forecasts, then no association should be observed between our earnings-surprise variable and our measure of peer-based advice. In contrast, we find that the fraction of negative words in SA articles and commentaries strongly predict subsequent earnings surprises.

Mechanisms

How does the "wisdom of crowds" work in the domain of stock analyses and opinions? We provide empirical evidence along two channels. We first construct, for each author, a measure of historical consistency. Our measure captures the degree to which predominantly positive (predominantly negative) articles by the author in question are subsequently followed by positive (negative) abnormal returns.

Our first analysis shows that both the number of page views, which directly impacts the level of monetary compensation the author receives from SA, and the number of times an article is read to the end, increase with the author's historical level of consistency. This pattern suggests that followers can differentiate between authors who offer historically good advice versus those who offer historically bad advice and the "popularity" of these authors changes accordingly. This should motivate authors to provide good and honest advice.

Our second analysis computes the degree to which SA commentaries on an underlying SA article convey a tone that contrasts with that of the SA article, and we relate this measure of author/follower disagreement to the author's historical track record. Our evidence implies that followers disagree with authors to a greater extent when the authors' articles have been inconsistent in the past. For these historically inconsistent authors, our evidence also suggests that in instances where the tone of commentaries disagrees with the tone of the underlying article, it is the tone of the commentaries that more reliably predicts subsequent stock market performance. In other words, it appears that another reason social media works in the domain of stock analyses is the "social aspect," which creates diversity of opinions and corrects/discourages the involvement of malignant and uninformed contributors.

Conclusion

Our paper assesses the value relevance of stock analyses and opinions shared by investors on social media. Traditionally, the job of financial forecasting has been performed by highly paid professional analysts; in recent years, do-it-yourself (DIY) financial analysis has become a trend, and investors are adopting social media to publicize their research work. In a National Examination Risk Alert, the U.S. Securities and Exchange Commission stated that "social media is landscape-shifting." Our study is an initial attempt to investigate the role of the evolving social media phenomenon in financial markets. There are many opportunities for future research in this area.

John Cotter, Stuart Gabriel, and Richard Roll

Can Housing Risk be Diversified? A Cautionary Tale from the Housing Boom and Bust

Review of Financial Studies | Volume 28, Issue 3 (Nov 2015), 913–936
(please cite only the original publication, not FAME)

Does geographic diversification reduce housing investment risk? To characterize diversification potential, we estimate spatial correlation and integration among 401 US metropolitan housing markets. The 2000s boom brought a marked uptrend in housing market integration. Numerous factors contributed to that trend, including eased residential lending standards and rapid growth in private mortgage securitization. As boom turned to bust, macro factors, including employment and income fundamentals, became important contributors to enhanced integration. Portfolio simulations reveal substantially lower diversification potential and higher risk in the wake of increased market integration. High levels of systemic risk and a reduced importance of local influences made geographic diversification less effective.

Geographic Diversification as a Method of Risk Mitigation

Geographic diversification long has been fundamental in risk mitigation among investors and insurers of housing, mortgages, and mortgage-related derivatives. The housing government- sponsored enterprises (GSEs), the Federal National Mortgage Association, (Fannie Mae) and the Federal Home Loan Mortgage Corporation, (Freddie Mac), now both in government conservatorship, diversified geographically in an effort to reduce the risk of investment in a single asset class (residential mortgages.) Similar logic was employed during the 2000s by prominent Wall Street firms, including Bear Stearns, Merrill Lynch, and Citigroup. More recently, geographic diversification has become central to the investment strategies of multi-family real estate investment trusts (REITs) and single-family housing investment funds.

However, during the late-2000s meltdown, anecdotal evidence suggests that geographic diversification of housing and mortgage investments was less effective. Indeed, diversification has limited power when returns are highly correlated. It appears that the substantial losses incurred by housing and mortgage investors during the downturn period depended in part to unforeseen and unprecedented contemporaneous price declines across geographically-distinct markets.

The effectiveness of geographic diversification as a method of portfolio risk mitigation is also of importance to private and government-backed insurers of residential mortgages. Substantial geographic correlation of credit losses, when coupled with sizable insurer guarantee liabilities and constrained access to credit markets, may render private mortgage insurance less viable. In such circumstances,

policymakers may need alternative mechanisms, such as deeply-subordinated government-backed insurance on qualified mortgages, to assure the liquidity and stability of the housing finance system.

Metro Housing Returns were Highly Correlated during the Meltdown Period

Our study commences with an assessment of spatial correlation in housing returns. This includes an examination of contemporaneous and lagged return correlations among 401 metropolitan statistical areas (MSAs) over the 1985-2012 period. High levels of MSA return correlation raise concerns for mortgage or housing investors seeking to diversify risk associated with investment in this asset class. We find large numbers of MSA pairs with contemporaneous and lagged housing return correlations at high levels of statistical significance. Of the roughly 80,000 distinct MSA return pairs, over 53,000 pairs are statistically significant with a mean correlation level of 39 percent. Large MSA housing return correlations appear to be especially pronounced in California. In that state, 98 percent of MSA paired returns are significantly correlated, with a mean correlation level of about 77 percent.

Integration of MSA House Price Returns Trended Up during the Late-2000s

Given evidence of high levels of spatial correlation in returns, we turn to an assessment of the integration of housing markets. Our measure of integration is based on the proportion of a MSAs housing market returns that can be explained by an identical set of national factors (see Pukthuanthong-Le and Roll (JFE 2009)). The level of integration is indicated by the magnitude of R-square, with higher values representing higher levels of integration. Two MSAs are viewed as perfectly integrated if the same national factors fully explain housing market returns in both areas. In that case, the R-square would be 1.0, implying no diversification potential between the MSAs.

Figure 1: Housing Return Integration Trends for U.S. MSAs and California MSAs

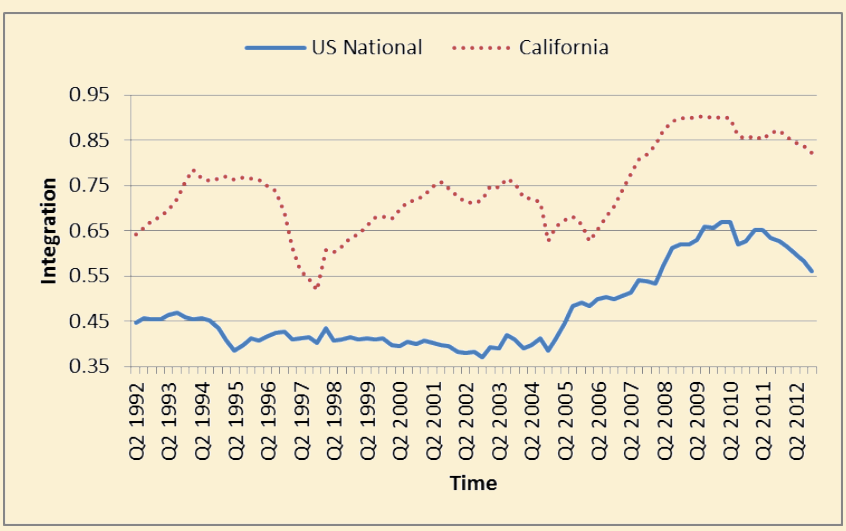

The level of integration is measured by the average R-squares from the multi-factor housing returns model fitted over a sequence of 30-quarter moving windows, 1992:Q2 - 2012:Q4 for 401 US MSAs and for 29 California MSAs.

As shown in Figure 1, there was a pronounced uptrend in US housing market integration over the period of boom and bust. Prior to the 2000s boom, average integration for the 401 MSAs held roughly steady at around 45 percent over the 1992-2004 period. However, starting in late 2004, average integration levels turned up and peaked at 67 percent in 2010. During 2011 and 2012, as the crisis abated, housing return integration trended back down to about 55 percent. Among California MSAs, integration is generally higher but the recent movements are similar, rising from about 63 percent in 2004 to around 90 percent late in the decade. By 2012, integration within California housing returns had declined to just over 80 percent. These recent movements in housing integration are robust to variations in MSA cohorts and estimation methods.

Drivers of Integration Trends

We are able to identify factors associated with the increased integration during the latter half of the 2000s. To do so, we compute the contribution to integration R-square associated with each factor. As shown in Figure 2, innovations in mortgage finance, notably including securitization of non-conforming mortgages and ease of mortgage underwriting, were strongly associated with higher integration during the 2004-2007 boom period. The economic significance associated with those factors also moved up substantially during the boom period. These results coincide with arguments in the literature (see, for example, Favilukis, Ludvigson, and Van Nieuwerburgh (WP 2013), Garriga, Manuelli, and Peralta-Alva (WP 2012), Duca, Muellbauer, and Murphy (WP 2012) and Mian and Sufi (QJE 2009)) that the boom in house prices was fueled in no small measure by widespread easing in mortgage qualification and in the provision of non-conforming secondary market liquidity.

Figure 2: Factor Contributions to MSA Housing Return Integration for U.S. MSAs 1992:Q2 through 2012:Q4

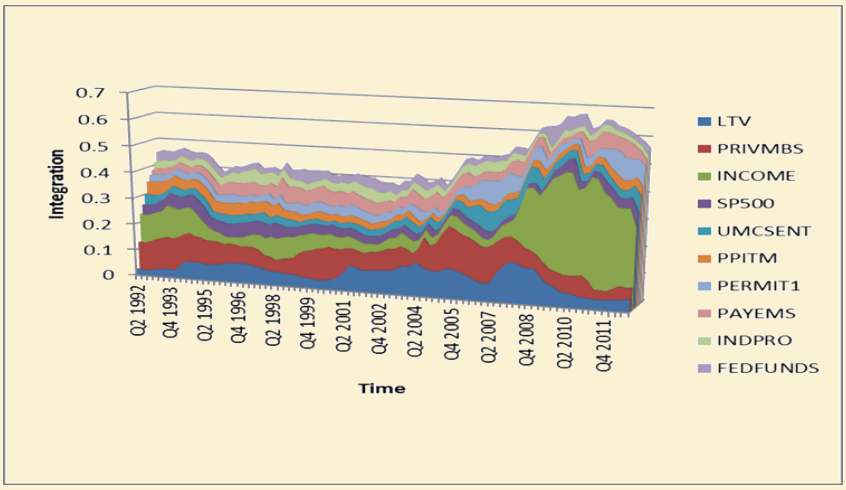

The level of integration is measured by the R-square from a multi-factor housing returns model estimated for sequential 30-quarter moving windows. The R-square contribution of each factor to the level of integration is plotted.

As boom turned to bust and the influence of mortgage liberalization waned, our results show that other macro factors, including employment and income fundamentals, contributed importantly to the ongoing trending up in housing return integration. Indeed, those factors were responsible for the

majority of the increment in U.S. housing market integration during the post-boom period. Similarly, the economic significance of those macro fundamentals was heightened during the crisis period. More recently, and in the wake of the attenuation of the crisis, those same macro factors accounted for much of the downward adjustment in metro return integration.

Does Geographic Diversification Reduce Housing Investment Risk?

Finally, to capture diversification potential for housing investors, we examine an equal-weighted portfolio of our U.S. metropolitan housing cohorts (1992:Q2-2012:Q4). This portfolio exhibits sharply rising levels of housing risk over the 2000s boom and bust. We then measure diversification potential by the difference between average individual MSA volatility and portfolio volatility.

Figure 3: Housing Return Integration, Portfolio Risk and Diversification for U.S. MSAs

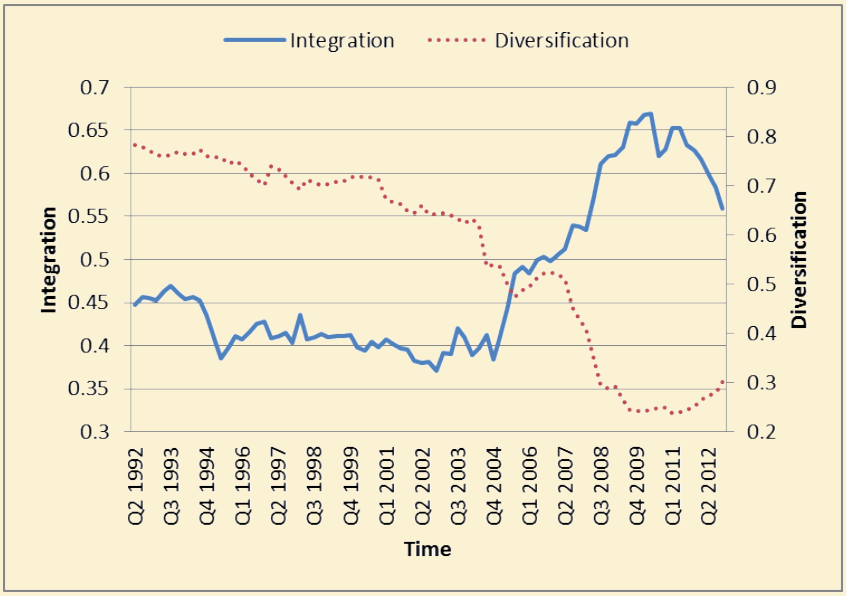

Integration is measured by the R-squares from the multi-factor housing returns model fit for a sequence of 30-quarter moving windows. Risk is measured by the standard deviation (volatility) of housing returns. Diversification is the average volatility of individual MSAs within an equal-weighted portfolio less the portfolio's volatility. The 30-quarter windows end on 1992:Q2 through 2012:Q4 for 401 US MSAs.

Changes in U.S. portfolio risk correlate strongly with the level of housing market integration. During the 2000s housing boom and bust, the simple correlation between the integration R-square and the standard deviation of portfolio returns is 0.96! As shown in Figure 3, during the crisis period, housing portfolio diversification provided only limited benefits in risk diversification. Indeed, the negative correlation between portfolio integration and diversification benefits averaged over –0.82 during the period of housing boom and bust. While integration slowed in 2011 and 2012, our findings still suggest substantial limitations to geographic diversification as a strategy for portfolio risk mitigation. Results suggest that investors and insurers of housing credit risk were rather exposed to the market downturn. Taken together, our findings offer a cautionary tale about geographic diversification as a mechanism to mitigate housing risk.

Henrik Cronqvist and Stephan Siegel

The Genetics of Investment Biases

Journal of Financial Economics | Volume 113, Issue 2 (Aug 2014), 215–234
(please cite only the original publication, not FAME)

Why do some investors have an internationally well-diversified portfolio while others invest solely in their home country? Are investment biases primarily learned from others or are they genetic? Does work experience in finance moderate the impact of genetic determinants? What about education?

In an effort to better understand why investors exhibit investment biases, and the degree to which they are genetic or can be influenced by external factors such as education and upbringing, Cronqvist-Siegel (JFE 2014) conduct a research study using data from the world's largest twin registry, the Swedish Twin Registry. Until 2007, taxpayers in Sweden were subject to a wealth tax. Prior to the abolishment of this tax, all Swedish banks, brokerage firms, and other financial institutions were required by law to report to the Swedish Tax Authority information about individuals' portfolios (i.e., stocks, bonds, mutual funds, derivatives, and other securities) held as of December 31 and also all sales transactions during the year. We matched 15,208 adult twin pairs from the Swedish Twin Registry with their portfolio and sales transaction data from the Tax Authority between 1999 and 2007, which gave us a detailed picture of their investment behavior.

Using empirical methodology adopted from quantitative behavioral genetics research (Neale-Maes (Kluwer 2004)), which has recently been used also in finance research, we match data from the twin registry with detailed data on the twins' investment behaviors and decomposed their differences into genetic and environmental components. We base our methodology on an intuitive insight: Identical twins share 100% of their genes, while the average proportion of shared genes is only 50% for fraternal twins. If identical twins exhibit more similarity with respect to investment biases than do fraternal twins, then this is evidence that these behaviors are influenced, at least in part, by genetic factors.

A long list of investment biases

We look at the following investment biases: Diversification, Home Bias (favoring home-country stocks), Turnover, Disposition Effect (reluctance to sell losers), Performance Chasing, and Skewness Preference (preferring 'lottery' or more volatile stocks). We measure Diversification for direct stock holdings as the number of distinct stocks held in an individual's portfolio at the end of a given year. For holdings of stocks and mutual funds, we follow Calvet-Campbell-Sodini (AER 2009) and define diversification as the proportion of equity investments invested in mutual funds as opposed to individual stocks. To reduce measurement error, we calculate the equally weighted average diversification across all years the individual is in the data set.

We measure Home Bias by the average proportion invested in Swedish securities. We measure Turnover, i.e. an individual's propensity to trade and turnover the portfolio, following Barber-Odean (JF 2000, QJE 2001). Specifically, for direct stock holdings, we divide, for each individual investor

and year, the sales volume (in Swedish krona) during the year by the value of directly held stocks at the beginning of the year. Since we do not have sales prices for mutual funds, we also construct a Turnover measure using the number of sales transactions during the year divided by the number of equity securities in the investor's portfolio at the beginning of the year. For each measure, we compute the average turnover using all years with available data. We measure the Disposition Effect in the spirit of Odean (JF 1998) and Dhar-Zhu (MS 2006). Specifically, at the end of each year during which we observe a sales transaction, we classify securities in an investor's portfolio as winners or losers based on the security's price relative to the approximate price at which the investor acquired the security. Using data across all years with sales transactions, we calculate for each investor the proportion of gains realized to the total number of realized and unrealized gains (PGR) as well as the proportion of losses realized to total losses (PLR). The larger the difference between PGR and PLR, the more reluctant a given investor is to realize losses.

We measure Performance Chasing by an individual's propensity to purchase securities that have performed well in the recent past. More specifically, each year we sort stocks and equity mutual funds separately into return deciles using the returns during the year. For each investor and year with net increases in holdings of stocks or mutual funds, we calculate the fraction of purchased securities with returns in the top two deciles. The higher that fraction, the more the individual chases performance by overweighting securities with higher recent performance.

We measure an individual's Skewness Preference as in Kumar (JF 2009). For each investor and year we calculate the proportion of the portfolio that is invested in "lottery" securities, ie, securities with a below median price as well as above median idiosyncratic volatility and above median skewness. Skewness preference is the fraction of lottery security securities averaged over all years with portfolio data.

Figure 1: Correlations of Investment Biases by Genetic Similarity

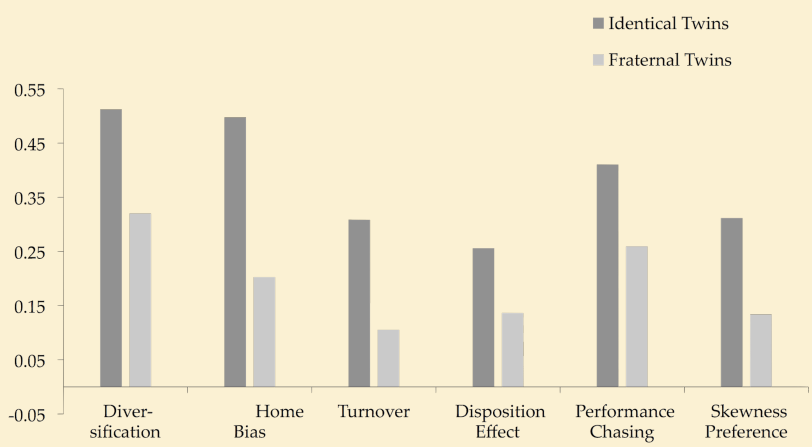

This figure contrasts correlations between identical twins with those between fraternal twins for several investment behaviors. **Interpretation: Identical twins behave in more similar ways compared to fraternal twins.**

Genetic differences explain up to 45% of the variation in investment biases across investors

Figure 1 illustrates one of the key findings of the study and shows that twins in an identical pair display much more similar investment biases compared to twins in a fraternal pair. This finding suggests that investment biases are partly genetic.

The importance of environmental influences

While our results are consistent with several behavioral genetic studies that have shown significant heritability of human behavior, they provide the first direct evidence from real-world, non-experimental data that persistent investment biases are to a significant extent, determined by genetic factors. Our results also show that even genetically identical investors who grew up in the same family environment can differ substantially in terms of their investment behaviors. Therefore, individual-specific environments, experiences, or events also play an important role in shaping individuals' investment behaviors. Our findings suggest that at least 55% of the unexplained variation in investment behaviors is due to environmental factors.

Among the environmental factors we explore are work experience in finance and education. We would like to know whether work experience in, for example, a bank or a corporate treasury department might reduce the influence of genetic tendencies towards investment biases. We were able to identify twins with finance-related work experience, and when we analyze their portfolio data we find that the relative importance of genetic factors on their investment biases is substantially smaller.

Education is another potentially important moderator of genetic effects. Here, we find that an increase in general education seems to have little effect on reducing genetic tendencies towards investment biases.

What explains the genetic effects we find and what are the implications of our findings?

Our evidence is consistent with evolutionary arguments of behavior as in Brennan-Lo (QJF 2011). Nature selects fitness maximizing behaviors, i.e., behaviors associated with a reproductive advantage relative to alternative behaviors. What in finance research is referred to as "biases" may indeed be manifestations of fitness maximizing psychological mechanisms. Consistent with this view of investment biases as partly innate features of human behavior, we find that the genetic factors that influence investment biases also affect behaviors in other, non-investment, domains.

Indeed, recent research in behavioral genetics has related specific genes to several of the psychological mechanisms that may manifest themselves as investment biases. That is, some individuals are endowed with genes related to familiarity (e.g., Chew-Ebstein-Zhong, JRU 2011), overconfidence (e.g., Cesarini-Johannesson-Lichtenstein-Wallace [JEEA 2009], or sensation-seeking (e.g., Derringer-etal [PS 2010]), and these genes may manifest themselves in the individual's investment behavior as well as the individual's behavior in other, non-investment domains.

An additional explanation for some of our results, which is consistent with recent work in finance (e.g., Grinblatt-Keloharju-Linnainmaa, JF 2011, JFE 2012), is that variation in IQ is genetic which results in genetic variation in investment biases.

Finally, our results have implications for the design of public policy in the domain of financial literacy. Specifically, the evidence suggests that policy should be designed accounting for the existence of genetic predispositions to investment biases and considering the challenges in reducing such biases.

Diane Del Guercio and Jonathan Reuter

Mutual Fund Performance and the Incentive to Generate Alpha

Journal of Finance | Volume 69, Issue 4 (Aug 2014), 1673–1704
(please cite only the original publication, not FAME)

The underperformance of the average actively managed mutual fund has been documented in numerous academic studies, and reported on regularly in investor-oriented publications, such as Forbes and the Wall Street Journal. Because this finding implies that investors are unable to earn back all of the fees they pay to mutual funds for active management, personal finance articles and American Finance Association presidential addresses alike have concluded that investors would be better off investing in passive index funds (French, JF 2008). The fact that the actively managed fund industry continues to thrive despite this persistent underperformance is a long-standing puzzle that several researchers have sought to explain.

Martin Gruber's American Finance Association presidential address first highlighted the puzzle of active management (Gruber, JF 1996). He conjectured that it might be driven by "disadvantaged" investors who are either ignorant of the average underperformance or behaving irrationally. Several recent studies posit that investors might rationally favor actively managed funds because they provide a form of insurance during recessionary times. Under this explanation, investors accept average underperformance in order to benefit from outperformance during periods when they would value it the most, when market returns are poor (Glode, JFE 2011).

In our paper, we propose an explanation related to Gruber's disadvantaged investor hypothesis. Namely, we show that the majority of actively managed funds are sold through brokers, and that investor flow within this segment of the fund industry responds to raw but not risk-adjusted returns (alpha). Because broker-sold fund investors do not reward high-alpha funds with additional investment dollars, broker-sold funds have little incentive to generate alpha. As a result, the underperformance of the average mutual fund in pooled samples is driven by this segment of broker-sold funds. At the same time, funds in the other retail fund segment conform quite closely to what theory would predict. Self-directed investors reward high-alpha funds, which incentivizes asset managers to generate alpha. Importantly, we find no evidence of average underperformance in this segment and thus no puzzle of active management. This implies that the underperformance of the average actively managed fund tells us more about manager and broker incentives than it does about market efficiency.

Many of our insights into this puzzle emerge from an ability to empirically distinguish between retail mutual funds directed at two distinct groups of investors: self-directed investors who are comfortable making their own investment allocation decisions, and those who require the advice of a broker in order to select a fund. This is important because we show that the mutual fund industry is largely segmented along these lines. Fund families that offer their retail funds direct to investors tend not to distribute any of their funds through brokers. Similarly, fund families that sell their funds through brokers tend

not to allow investors to purchase funds direct from the family. For example, investors who wish to buy one of the largest funds, the Investment Company of America fund offered by the American Funds family, can only do so through a registered financial advisor. Under this type of segmentation, fund families can tailor their products and their marketing and operational strategies to cater to the needs of their distinct clientele.

Do direct-sold funds face a stronger incentive to generate alpha?

It has long been recognized that fund families have a strong incentive to increase assets under management since funds primarily charge fees that are a percentage of assets. Thus, families are incentivized to deliver funds with whatever product features will ultimately generate more dollar flow into their funds. We show that differences in investor preferences and expertise in the two segments has important implications for the nature of competition for flow within each segment.

According to two surveys, investors who rely on the advice of their broker report that they are not comfortable making their own allocation decisions. They tend to be inexperienced investors who highly value the face-to-face contact and repeated interaction with a trusted advisor. In contrast, surveyed investors who do not purchase mutual funds through a financial adviser state that they "want to be in control of own investments" and already "have access to all of the resources needed to invest on my own". Furthermore, several papers have documented that commissions paid to brokers provide powerful incentives to direct flow to those products, suggesting that broker-sold funds can better compete for assets through larger payments to brokers rather than through increased efforts toward improving fund performance. For these reasons, we predict that investor flows in the broker-sold segment are less sensitive to alpha than in the direct-sold segment. We find strong empirical support for this prediction.

Table 1: Monthly Flow — Performance Sensitivity Across Market Segments, Actively Managed Funds (1993 to 2004)

Dependent Variable:	Net Flow (t)	Net Flow (t)	
Sample:	Both Segments	Direct-Sold	Broker-Sold
Net flow (t-1)	0.206**	0.189**	0.229**
	(0.033)	(0.048)	(0.026)
Net return (t-1)	0.077**	0.050	0.135**
	(0.034)	(0.045)	(0.023)
Four-factor alpha (t-1)	0.107**	0.176**	0.021
	(0.032)	(0.049)	(0.020)
R^2	0.0784		0.0887

The dependent variable is monthly net percentage fund flow, using the standard definition of flow, growth in the fund's Total Net Assets less capital appreciation. The sample size is 122,111.

In Table 1, we report the results of a standard method testing for which fund characteristics are the main drivers of investor flow. We report the results of a regression with fund flow as the dependent variable and two measures of past performance, last period's four-factor alpha and last period's raw return unadjusted for any risk, as the main explanatory variables of interest. We also include in the regression other standard fund-level control variables, such as fund size and expense ratio, but do not report these in the table. Column 1 contains the regression coefficients for the pooled sample that includes funds from both segments, and columns 2 and 3 contain the coefficients for regressions run separately on direct-sold-only and broker-sold-only funds. The key finding is that alpha is a significant driver of flow only in the direct-sold segment, while raw return is a driver of flow only in the broker-sold segment. This suggests that only direct-sold funds are rewarded with more flow when they are able to generate alpha. Broker-sold investor flows do not respond to alpha. The important implication is that only direct-sold funds have an incentive to invest resources in the skilled managers, analysts, and trading desk personnel who will ultimately generate alpha. Note that we were only able to uncover these incentive differences by allowing flow-performance relations to vary across the two market segments. Standard pooled regressions like those in column 1 erroneously suggest that flow into the typical retail mutual fund responds to both raw and risk-adjusted returns.

Do direct-sold funds respond to their stronger incentive?

Although we cannot directly observe whether direct-sold funds invest more resources into the inputs that will generate alpha, we can test whether there are significant differences in the outcomes of such efforts. To show this, we test whether actively-managed direct-sold funds have better measures of value-added performance than actively-managed broker-sold funds. Because the investment universe of small-cap stocks is expected to be less efficient and therefore has greater room for a skilled fund manager to profit, we also separately examine small-cap funds within each segment. Figure 1 summarizes our findings for two of the measures we report in Del Guercio-Reuter (JF 2014): the four-factor alpha of Carhart (JF 1997) and the return gap measure of Kacperczyk-Sialm-Zheng (RFS 2008). The latter measure is designed to capture unobservable aspects of a fund's performance, such as whether trade executions are poor or stellar or whether the fund manager timed their stock purchases or sales well. In regressions controlling for many fund characteristics, we find that the average four-factor alpha of direct-sold funds is 9.6 basis points per month higher relative to broker-sold funds when we consider all equity funds, and 17.4 basis points per month higher when we restrict the sample to small-cap funds. Similarly, we find that the return gap measure is 6.1 basis points per month higher performance for direct-sold funds when we consider all equity funds, and 12 basis points per month higher for small-cap funds. Additional tests reveal that direct-sold families are more likely to make operational decisions associated with higher fund returns. For example, they offer a narrower range of fund styles and rely less on subadvisors than broker-sold families. In sum, we find that direct-sold families internalize the preferences of their target clientele for alpha.

Figure 1: Difference in Performance: Direct-sold vs. Broker-sold funds (in basis points per month)

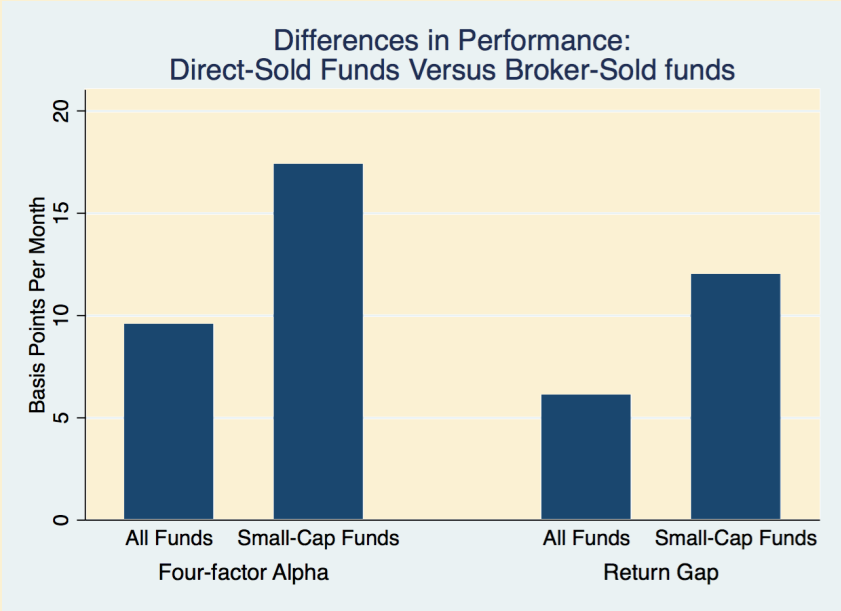

Direct-sold funds performed better than broker-sold funds.

No puzzle of active management in the direct-sold segment

Finally, we show that the puzzle of active management is confined to the broker-sold (disadvantaged investor) segment of the fund industry.

Table 2 shows the results of tests for whether index funds outperform actively-managed funds, or equivalently, have higher alpha. In column 1, we pool all funds together. Index funds perform better than actively managed funds. However, examining the two segments separately reveals that this underperformance is driven by the broker-sold segment, where index funds outperform active funds by 1.12% per year! In contrast, there is no statistical or economically meaningful difference between the performance of the average direct-sold index fund and the average direct-sold actively managed fund. Because there is no underperformance, there is no puzzle of active management within the direct-sold segment; managers are able to earn high enough abnormal returns to justify their fees. This finding conforms to theoretical arguments in Grossman-Stiglitz (AER 1980) and Berk-Green (JPE 2004), and undercuts claims that the U.S. equity market is too efficient for active managers to earn back their fees.

Our findings suggest that in order to better understand the demand for underperforming active management we need to better understand the continued demand for broker services. One possibility is that brokers provide other valuable services to their clients that compensate for the underperformance, for example, by helping them avoid other costly mistakes. Another possibility, however, is that brokers direct clients to the funds that offer them the greatest commission payments, and that their inexperienced clients are unable to distinguish good advice from bad. The fact that we find the overwhelming majority of assets in the broker-sold segment are invested in actively managed funds is suggestive of an agency conflict, but further research is needed to provide a definitive answer.

Table 2: Monthly Fund Four-Factor Alphas of Actively Managed and Index Funds Across Market Segments (1993 to 2004)

	Dependent Variable: Four-Factor Alpha (t)		
	(1)	(2)	(3)
Sample:	Both Segments	Direct-Sold	Broker-Sold
Index fund dummy (t)	0.073**		
	(0.034)		
Active fund dummy (t)	Omitted Category		
Direct-sold dummy (t) × Index fund (t)		0.018	
		(0.035)	
Direct-sold dummy (t) × Active fund (t)		Omitted Category	
Broker-sold dummy (t) × Index fund (t)			0.093
			(0.039)
Broker-sold dummy (t) × Active fund (t)			Omitted Category
Sample size	122,833	51,469	71,364
R^2	0.1135	0.0974	0.1458

Investors are better off in index funds, as the average index fund outperforms an otherwise similar actively-managed fund by 7.3 basis points per month (0.88% per year). This is driven by the broker-sold segment. Standard errors are in parentheses.

Jason T. Greene and David Rakowski

A Note on the Sources of Portfolio Returns: Underlying Stock Returns and the Excess Growth Rate

Critical Finance Review | Volume 4 (2015), 117–138
(please cite only the original publication, not FAME)

It is tempting to assume that the performance of a portfolio is equivalent to the performance of its underlying stocks. However, when returns are compounded over time and across firms, the resulting compound return for a portfolio is not necessarily equal to the average compound return of that portfolio's underlying securities. While this mathematical fact is well-known, researchers and investors often ignore its implications when interpreting the sources of returns to portfolios of stocks. We illustrate a method of parsing portfolio returns into their underlying sources drawn from stock-level returns and stock-level variance-covariance properties. We demonstrate the usefulness of this method for analyzing the cross-section of stock returns by showing how some portfolios formed from popular stock characteristics display compound return patterns that are quite different from the average compound returns of the underlying stocks in those portfolios, while other portfolios display return patterns that closely match those of their underlying stocks.

Table 1: Examples of the sources of returns for stocks and portfolios

Raw Returns				
Period 1	Period 2	Average	Variance	
A_1	+100%	–50%	+25%	0.5625
A_2	–50%	+100%	+25%	0.5625
EW_A	**+25%**	**+25%**		
B_1	+25%	+25%	+25%	0.0000
B_2	+25%	+25%	+25%	0.0000
EW_B	**+25%**	**+25%**	**+25%**	**0.0000**

Log Returns				
Period 1	Period 2	Average	Variance	
A_1	+0.6931	–0.6931	0.0000	0.4805
A_1	–0.6931	+0.6931	0.0000	0.4805
EW_A	**+0.2231**	**+0.2231**	**+0.2231**	**0.0000**
B_1	+0.2231	+0.2231	+0.2231	0.0000
B_2	+0.2231	+0.2231	+0.2231	0.0000
EW_B	**+0.2231**	**+0.2231**	**+0.2231**	**0.0000**

This illustrates how compound portfolio returns are not necessarily related to average compound individual constituent returns. This example shows that portfolio returns can be high even when constituent returns are zero.

Portfolio compound returns and average stock compound returns

The following simple example illustrates the underlying economic issue that is the focus of our paper. Suppose there are four stocks with returns over two periods as presented in Table 1. For each stock we present both its raw return and its compound (log) return. First, consider two stocks A_1 and A_2. Each stock doubles in value one period and loses half its value in the other period. Both stocks will have a 25% average raw return over the two periods, but an average log return of zero. Even though the average log return is zero for each stock, an equal-weighted portfolio of the two stocks earns a 22.31% log return per period.

Now consider stocks B_1 B_2 and , which both have raw returns of 25% (log returns of 22.31%) each period. An equally-weighted portfolio (EW_B) of these two stocks will have a raw return of 25% (log return of 22.31%) per period. Thus, portfolio EW_B performs identically to portfolio EW_A, but the sources of the returns are strikingly different. In portfolio EW_A the compound return is driven by the variance of the underlying stocks. However, the compound return of portfolio EW_B is a direct result of the compound returns of the two stocks. We provide a method for parsing portfolio returns into these two distinct sources. Fernholz-Shay (1982), hereafter FS, offer what we believe to be the first rigorous mathematical analysis to identify the sources of the longterm performance of a portfolio. Specifically,

FS show that a portfolio that is rebalanced to the same constant weights has a compound return that is a function of the underlying stocks' 3 compound returns and an "excess growth rate" that is due to the difference between the stocks' variances and the (diversified) portfolio's variance.

Portfolios formed from market capitalization and book-to-market ratios

We illustrate the importance of the FS model by empirically estimating the sources of portfolios' returns for several characteristic-based portfolios that appear frequently in the literature. Figure 1 displays portfolio compound returns, average stock compound returns, and the excess growth rate for portfolios formed from deciles of stocks' market capitalizations. Our results reveal that a portfolio of stocks from the smallest market cap decile has higher average monthly compound returns than a large stock portfolio, but the average compound returns of stocks in the small-cap portfolio are lower than the average compound returns of stocks in the large-cap portfolio. The difference between these two estimates derives from the excess growth rate that is induced by the higher variance of stocks in the small-cap portfolio.

Figure 1 illustrates portfolio compound returns, and the components of portfolio compound returns, for decile portfolios formed from stocks' book-to-market ratios. Portfolio compound returns increase consistently with the book-to-market ratio. In contrast to market-cap portfolios, the difference in portfolio compound returns is driven primarily by the difference in the average underlying stock compound returns for book-to-market portfolios, as the average underlying stock returns also increase with the book-to-market ratio.

Figure 1: Returns for Portfolios formed from firm size

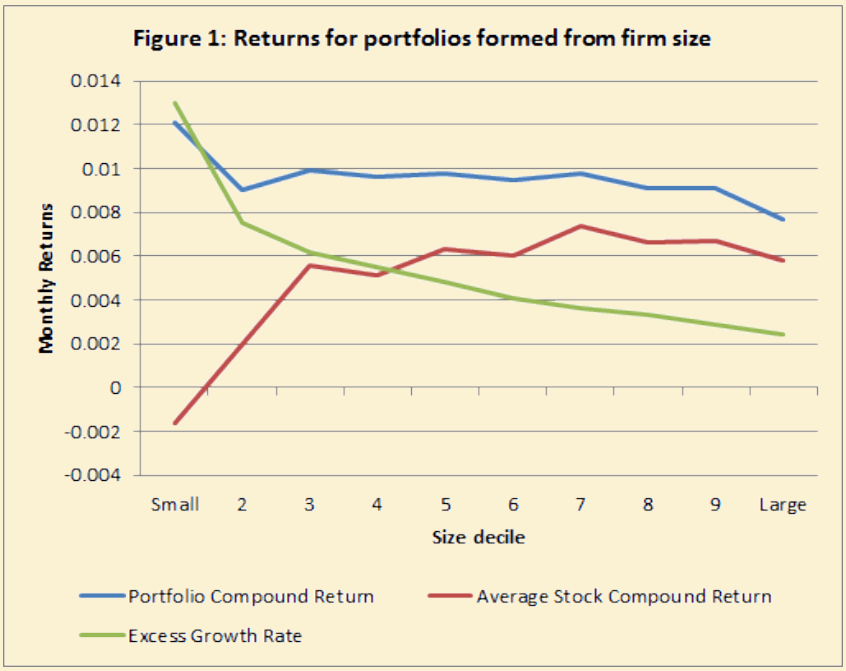

This displays average monthly portfolio compound returns, and the components of those returns, for a sample of all valid US-listed common stocks over the 1960-2012 period. Portfolios are formed based on firms' market capitalization. **Interpretation:** Figure 1 shows that average stock returns increase as firm size increases, but portfolio returns decrease.

Extensions and implications

Our procedure of decomposing portfolio returns can be used to shed light on the sources of portfolio returns for a variety of common portfolio-formation variables. For example, portfolios formed based on long-term past performance (i.e., typical contrarian portfolios) display portfolio-level effects in returns, like size-based portfolios, that are the opposite of their stock-level average returns. On the other hand, portfolios formed based on short-term past performance (such as momentum portfolios) or liquidity characteristics, display portfolio patterns, like B-M portfolios consistent with their underlying stocks.

We extend our analysis by re-examining the impact of variance on the analysis of the cross-section of stock returns. Cross-sectional regressions of monthly raw returns and log returns confirm the results from the decomposition of portfolio returns. Specifically, the apparent cross-sectional relationship between size and stock raw returns is not robust to the inclusion of stock variance in the regression model, while the book-to-market effect is unchanged when variance is included. Furthermore, no cross-sectional relationship is evident between stock log returns and either size or own-return variance, but a firm's book-to-market ratio is positively associated with its average log return.

The implications and applications of decomposing portfolio returns into effects due to stock returns and portfolio excess growth rates are wide-ranging. Whenever the returns of portfolios composed of high-volatility stocks are compared to benchmarks or control portfolios of lower-volatility stocks, we can expect to observe differences due to the excess growth rates of the portfolios. Even in the wider context of other social sciences that study phenomena subject to compound growth rates, any time the growth rate of a high-volatility treatment group is compared to a lower-volatility control group, there could be differences that may be interpreted as an excess growth effect.

Feng Gao, Ling Lei Lisic, and Ivy Zhang

Commitment to Social Good and Insider Trading

Journal of Accounting and Economics | Volume 57, Issue 2-3 (Apr-May 2014), 149–175
(please cite only the original publication, not FAME)

Corporate social responsibility (CSR) has gained tremendous attention from regulators, corporate executives, investors, and various other stakeholders in recent years. Our study documents an unintended consequence of these increasingly important corporate activities. We find that a firm's commitment to CSR constrains executive self-serving behavior as manifested in insider trading.

CSR Reputation and Insider Trading

When a firm commits to be socially responsible, the positive image imposes costs on managers engaging in self-serving activities, such as insider trading, that conflict with the appearance of doing social good. Informed trading is prohibited by law and often considered greedy and unfair expropriation of uninformed investors. When informed trading is detected, the negative publicity likely impairs the firm's positive social image and reduces the associated benefits, eventually causing damage to managers' personal interests that are tied to the image of the firm. We thus expect that a firm's socially responsible image constrains profitable insider trading.

A firm's CSR engagement can also be negatively associated with insider trading if both corporate CSR policies and executives' insider trading reflect executives' personal preferences to be a good citizen. Socially responsible executives are less likely to pursue self-serving actions such as informed trading.

Using the social ratings data issued by MSCI ESG STATS from 1991 to 2010, we measure a firm's commitment to social good with a summary CSR score that reflects various aspects of social responsibility, including environment, community affairs, employee relations, diversity, customer relations and involvement in controversial industries. A firm is considered CSR-conscious if it has more CSR strengths than concerns. We focus on insider purchases only because on average executives profit from purchase transactions but not from sale transactions. We measure executives' insider trading profits as potential profits gained from purchases.

Table 1 reports the trading profits of executive purchases in CSR-conscious firms versus non-CSR-conscious firms. The mean daily trading profits from purchases made by executives of non-CSR-conscious firms are 0.0564% for the 180-day window, 0.0813% for the 120-day window, and 0.0916% for the 90-day window. In contrast, the mean daily trading profits from purchases made by executives of CSR-conscious firms are 0.0402%, 0.0604%, and 0.0621%, respectively. The trading profits from purchases made by executives of non-CSR-conscious firms are at least 35% higher than those from purchases made by executives of CSR-conscious firms, and all differences are statistically significantly at the 1% level. These results suggest that executives of CSR-conscious firms profit significantly less

than executives of non-CSR-conscious firms from insider trading. Further analyses of changes in a firm's CSR orientation suggest that executives make lower trading profits as a firm becomes CSR conscious and they make higher trading profits as a firm becomes non-CSR-conscious.

Table 1: CSR and Insider Trading Profits

	Purchases	
	Mean (%)	Median (%)
CSR-conscious firms		
Profit (t + 180)	0.0402	0.0374
Profit (t + 120)	0.0604	0.0577
Profit (t + 90)	0.0621	0.0537
Non-CSR-conscious,firms		
Profit (t + 180)	0.0564	0.0524
Profit (t + 120)	0.0813	0.0739
Profit (t + 90)	0.0916	0.0861
Difference		
Profit (t + 180)	–0.0163	–0.0149
Profit (t + 120)	–0.0209	–0.0162
Profit (t + 90)	–0.0296	–0.0325

This reports the mean and median of trade-specific daily returns on four common factors over the 180, 120, and 90 days following executives' insider purchases. The difference in trading profits between CSR-conscious and non-CSR-conscious firms (i.e., the bolded numbers) are statistically significant at 1% level.

We obtain similar inferences with the likelihood of executive trades prior to future corporate news. We find that executives of non-CSR-conscious firms trade on future corporate news contained in earnings and stock returns. However, we do not find such evidence for executives of CSR-conscious firms. Further analyses of changes in CSR orientations also indicate that executives trade on future corporate news before a firm becomes CSR conscious, but not afterwards, and after a firm ceases to be CSR conscious, but not beforehand.

Channels of CSR Impact

We next explore channels through which damage to a firm's CSR image translates into personal costs for executives and thus discourages informed trading. First, executives gain personal reputational capital from corporate activities when executives are vocal (or visible) with regard to the firm's CSR activities. We thus expect the firm's CSR orientation to have a stronger constraining effect on these executives' trading. In addition, if being vocal about CSR reveals that these executives are more

altruistic or less greedy, we would also expect them to refrain more from informed trading. We determine whether executives are vocal about CSR via a comprehensive search of the internet, news media, press releases, and CSR reports. Consistent with our expectation, we find the constraining effect of CSR on insider trading profits to be stronger for CSR advocates.

Second, damage to the CSR image of the firm can hurt executives through their stock ownership. We find the constraining effects of CSR to be more pronounced when executives' wealth is more closely tied to firm value, i.e. when they hold more company shares.

Personal Preferences?

The negative relation between CSR and insider trading could be explained by both CSR image constraining insider trading and managerial preferences for good citizenship. After controlling for executives' personal styles, our results suggest that the constraining effects of CSR image are not subsumed by individual preferences.

Concluding Remarks

Our paper advances the understanding of the impact of committing to social good on self-serving managerial behavior. We document a robust negative association between a firm's CSR consciousness and executive insider trading. Various analyses suggest that the CSR image can act as an unintended governance mechanism restricting informed trading. We add to the growing literature on the implications of CSR engagement. We also contribute to the insider trading literature by proposing an additional constraint on informed trading other than regulation and corporate insider trading policies.

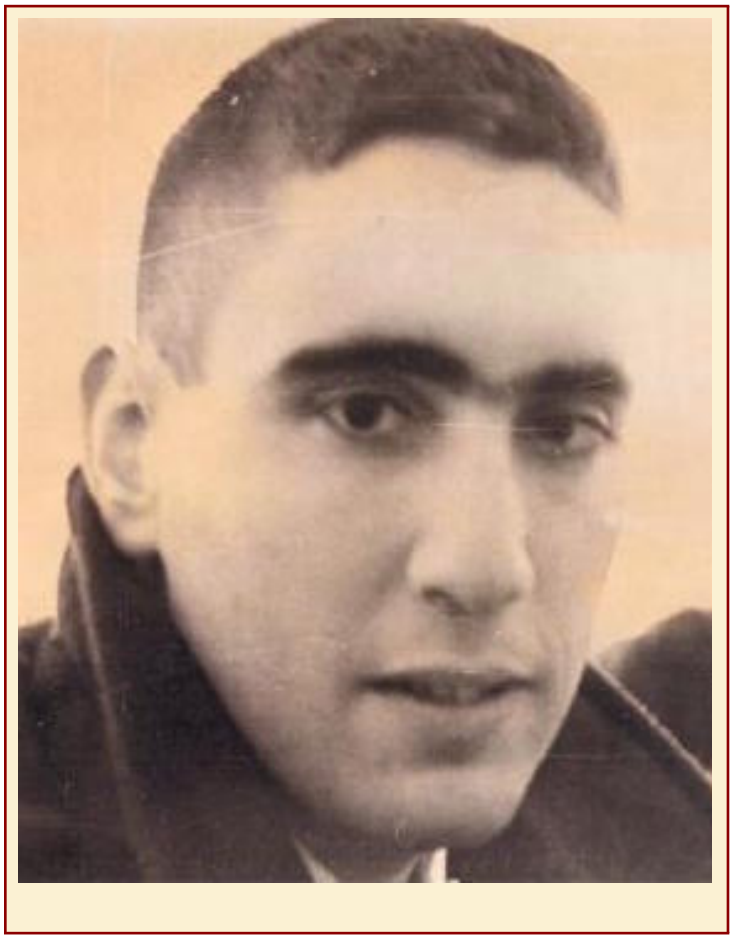

Ningzhong Li, Scott Richardson, and Irem Tuna

Macro to Micro: Country Exposures, Firm Fundamentals and Stock Returns

Journal of Accounting and Economics | Volume 58, Issue 1 (Aug 2014), 1-20
(please cite only the original publication, not FAME)

In this paper we examine whether information about a company's geographic (macroeconomic) exposure is useful for forecasting firm fundamentals and stock returns. While the link between firm operating and investing decisions and broader macroeconomic features seems relevant for forecasting, surprisingly little archival, empirical research has examined these relations. Indeed, with an increasingly inter-connected system of economic and financial markets across developed and developing countries, understanding the macroeconomic landscape is important for security valuation.

We outline a systematic approach to incorporate macroeconomic information into firm level forecasting from the perspective of an equity investor. Using a global sample of 198,315 firm-years over the 1998-2010 time period, we find that combining firm level exposures to countries (via geographic segment data) with forecasts of country level performance, generates superior forecasts for firm fundamentals. This result is particularly evident for purely domestic firms. We further find that this forecasting benefit is associated with future excess stock returns. These relations are stronger after periods of higher dispersion in expected country level performance.

Combining country exposures to form a firm level forecast

We consider how each company is exposed to its home country and other countries. We identify country exposures via our own manual coding of the geographic segment sales disclosures in Compustat and FactSet Fundamentals and the geographic sale data collated by FactSet Geographic Revenue Exposures. For each firm-year observation we disaggregate total sales into country level sales based on the geographic segment data in the most recent annual report. We retain companies with a purely domestic footprint (i.e., those companies with zero foreign sales) as this allows us to more cleanly assess the importance of macroeconomic information. For example, if firm A has 50 percent of its sales in Germany and 50 percent of its sales in Greece and Firm B has 100 percent of its sales in Greece, and Greece is expected to outperform Germany, then holding all else equal, the 'best' portfolio exposure to express that view would be via Firm B, the purely domestic firm.

After gathering the sales data for firm i, for each country c, at each point in time t, $\text{Sales}_{i,t,c}$, we standardize these sales measures so that they sum to one. We then use our forecasts of expected country level performance for each county c at each point in time t, $\text{E[Performance]}_{c,t}$ to construct a firm level measure of exposure to expected macroeconomic performance. We use forecasts of real GDP growth from Consensus Economics as our primary measure of expected country level performance. To generate a company specific fundamental forecast, we take the sum-product of $\text{Sales}_{i,t,c}$ and $\text{E[Performance]}_{c,t}$ which we label $\text{MACRO}_{i,t}$(i.e, $\text{MACRO}_{i,t} = \sum_{c}^{1} \text{Sales}_{i,t,c} \cdot \text{E[Performance]}_{c,t}$ This measure captures both

cross sectional and time series variation in firm level sensitivities to macroeconomic (country level) performance drivers.

In our full sample, we group together the 'domestic' and 'non-domestic' firms. However, we also separately examine the importance of macroeconomic information for these two groups. A potential benefit of examining 'domestic' and 'non-domestic' firms separately is to highlight two related effects. First, investors and analysts may be ignoring macroeconomic information in general. If this is true, then we should see strong predictive ability for the domestic only sample. Second, investors and analysts may be ignoring information in the differential geographic reach of companies. If this is true, then we should also see predictive ability for the 'non-domestic' sample.

Predictive power of MACRO for future firm profitability

For a sample of 198,315 US and non-US firm-year observations over the 1998-2010 period, we find that combining country exposures with country level forecasts ($MACRO_{i,t}$) improves forecasts of firm fundamentals. The predictive power of $MACRO_{i,t}$ is evident in annual regressions, which suggest that a one percentage point increase in expectations of real GDP growth translates to an additional 27 basis points of return on net operating assets (RNOA) over the next year. The predictive power of $MACRO_{i,t}$ is robust to including a wide set of explanatory variables, including sell-side analyst forecasts.

Our sample includes 135,974 'domestic' firm-year observations with exposure only to their home country, and 62,341 'non-domestic' firm-year observations that have exposures to multiple countries. We separately estimate the usefulness of country exposures and country level forecasts to improve forecasts of firm fundamentals for 'domestic' and 'non-domestic' firms. For domestic firms we find that $MACRO_{i,t}$ is strongly associated with future RNOA: a one percentage point increase in expectations of real GDP growth translates to an additional 31 basis points of RNOA. For non-domestic firms we find that $MACRO_{i,t}$ is weakly associated with future RNOA: a one percentage point increase in expectations of real GDP growth translates to an additional 20 basis points of RNOA.

There are two distinct effects driving the positive relation between $MACRO_{i,t}$ and future firm performance. First, we find that forecasts of real GDP growth are useful in forecasting future firm performance for domestic firms. This result does not require the use of potentially noisy geographic segment data. To the best of our knowledge this is a new result in the literature. Second, we find that forecasts of real GDP growth are also useful (albeit less strongly) in forecasting future firm performance for non-domestic firms. This result does require the use of potentially noisy geographic segment data. However, additional tests suggest that despite the potential measurement error in the geographic segment disclosure data, they are still useful to equity investors interested in forecasting future firm performance. Specifically, we compare our measure of geographic exposure to 'naive' alternatives which ignore the information about the countries from which sales are sourced. We show that these naive measures have no ability to forecast future firm fundamentals. Thus, despite the measurement error in our geographic exposures based on subjective, and potentially inconsistent, geographic categories by management, they are superior to ignoring the level of disaggregation provided in geographic segment sales disclosures.

Predictive power of MACRO for future stock return

We also show evidence that stock returns appear to incorporate the information in geographic exposures with a lag. We sort firms into quintiles each month based on $MACRO_{i,t}$ and compute value

weighted raw returns, as well as value weighted characteristics adjusted returns, to a dollar neutral hedge portfolio. As shown in Table 1, in the full sample, the average monthly raw (characteristics adjusted) return to this hedge portfolio is 1.28% (1.15%), which translates to an annualized Sharpe ratio of 0.85 (0.93).

Table 1: Monthly Returns to Hedge Portfolios Formed on MACRO

	Raw Hedge Return			Characteristics-Adjusted		
	All	High Disp	Low Disp	All	High Disp	Low Disp
Avg Monthly Return	1.28%	2.08%	0.47%	1.15%	1.61%	0.69%
T-stat	3.00	4.18	1.45	3.26	4.17	2.23
Ann. SR	0.85	1.18	0.41	0.93	1.19	0.63
N	148	73	73	148	73	73

Monthly returns and Sharpe ratios for various hedge portfolios.

When is the predictive power stronger?

We also show that the fundamental and return predictability of information contained in current country exposures and country level forecasts is greater after periods of increased dispersion in real GDP growth forecasts across countries. This suggests that when the information content of MACRO is high ex ante, there is a stronger predictive content. As indicated in Table 1, the return to the hedge portfolio formed on MACRO is much larger in the subsample of months with higher dispersion in real GDP growth forecasts across countries than that in lower dispersion months. We also show that the return predictability of $MACRO_{i,t}$ is greater when there is greater ex post information content. For this test we split our sample based on the ex post forecast accuracy improvement of including our $MACRO_{i,t}$ measure into forecasts of RNOA. Mechanically this partition identifies sub-samples where $MACRO_{i,t}$ is (is not) useful for forecasting RNOA. Interestingly, the return predictability is only evident in the sub-sample where there is an ex post improvement in forecasting RNOA. This suggests that the return predictability we document is attributable to improved forecasts of firm fundamentals that the market does not price correctly.

A key contribution of this paper is to introduce a simple framework to identify and exploit linkages between firm performance and its potential macroeconomic drivers. Our paper is different from prior studies in that we (i) examine a different basis for identification of linkages (geographic location as opposed to industry membership or supply chain links), (ii) use explicit forecasts based on the links we identify (i.e., we focus on forecasts of real GDP growth), and (iii) demonstrate predictive ability for firm fundamentals and stock returns and link these results together.

Jerchern Lin and Wayne Ferson

Alpha and Performance Measurement: The Effects of Investor Disagreement and Heterogeneity

Journal of Finance | Volume 69, Issue 4 (August 2014), 1565–1596
(please cite only the original publication, not FAME)

Finance researchers have an easy familiarity with alpha, the most well-known measure of the expected abnormal return of an investment. Studies refer to CAPM alpha, three-factor alpha or four-factor alpha, assuming the reader hardly requires a definition. Investment practitioners discuss their strategies in terms of the quest for alpha. Alpha can be active, conditional or portable. The number of investment firms with alpha in their names is staggering.

Despite the apparent familiarity with alpha, the current literature too often fails to think rigorously about how alphas can be interpreted. The literature has provided examples and counterexamples, leading to the conclusion that, given a fund with a positive traditional alpha, an investor would not necessarily wish to buy the fund.

In this paper, we propose a client-specific alpha that unambiguously signals the attractiveness of a fund to a client. We also derive and estimate an upper bound on the extent to which a client is likely to disagree with a traditional alpha about a fund's performance. We find that disagreement based on client preferences is likely to be of a similar economic magnitude as the errors in the estimation of alpha, and roughly as important as the choice of the performance benchmark. We provide empirical proxies for the expected disagreement with a traditional alpha and for heterogeneity, which is the variation in disagreement across a fund's investors. We study empirically the relation of disagreement and heterogeneity to investor flows of new money in and out of mutual funds. We find that investor flows do respond to disagreement and heterogeneity, and conclude that disagreement and heterogeneity are economically important issues for mutual funds.

A General Model for a Client's Alpha

We incorporate client preferences in a portfolio optimization problem to derive a client-specific alpha without the assumption in utility functions and normally distributed returns. We then study how the client adjusts consumption and portfolio choices when facing a new investment. We analytically show that the client optimally buys (sells) a fund with a positive (negative) client-specific alpha.

Investor Disagreement and Heterogeneity

Investors will not in general agree about the "right" alpha, so the same fund may look attractive to one investor but not to another. We thus define a measure of the expected disagreement a client will

have with a traditional alpha and derive bounds on its magnitude. We also define heterogeneity as the variance of disagreement across investors. We investigate the implications for mutual fund flows of the expected disagreement and heterogeneity.

Main Results

Our results are summarized as follows. First, we derive a well-specified alpha using a multi-period model under general returns distributions, and show that the client-specific alpha does provide a reliable buy or sell indication. Second, we find that disagreement with traditional alphas can be similar in importance to the choice of the performance benchmark or to the statistical imprecision in estimates of alphas. Third, investor disagreement and heterogeneity are economically significant in the behavior of fund investors. Funds whose clients' alphas are expected to be larger than a traditional alpha have positive average disagreement. We find that such funds experience larger inflows, controlling for the traditional alpha, than funds with low average disagreement. Funds with a large cross-sectional variance of investor disagreement have greater investor heterogeneity. We find that funds with more heterogeneity across investors experience lower average flows. These effects are separate from uncertainty about the true value of the traditional alpha.

Implications

The importance of investor disagreement with traditional alphas and investor heterogeneity are likely to be greater than our empirical results indicate, because we abstract from important sources of additional disagreement and heterogeneity. For example, we do not consider differences in investor beliefs, taxes or institutional setting. If investors disagree enough, then even a fund with a negative traditional alpha can thrive. A fund with a positive traditional alpha can find itself spurned in the marketplace. These results are potentially important in practice. If our empirical results are conservative, the economic magnitudes of disagreement and heterogeneity for funds and investors might be larger than the benchmark choice decision and the uncertainty in estimates of traditional alphas. These two issues have received much research, but the effects of disagreement and heterogeneity have been little explored in the literature.

Sanjiv Das, Madhu Kalimipalli, and Subhankar Nayak

Did CDS Trading Improve the Market for Corporate Bonds?

Journal of Financial Economics | Volume 111, Issue 2 (February 2014), 495–525
(please cite only the original publication, not FAME)

Financial innovation is a double-edged sword. The creation of new markets and new securities may complete markets, provide new investment opportunities and risk hedging alternatives, and favorably impact information generation and dissemination; yet such innovations may also have negative externalities if the gains accrue to only a few market participants and cause adverse impact on rest of the market.

A salient innovation in the fixed-income and credit markets since the turn of the century is the introduction of the credit default swap (CDS), a credit insurance contract that provides payoffs contingent on the default or change (particularly, deterioration) in credit characteristics of an underlying reference bond or issuer. In our JFE paper, we examine whether the advent of CDS trading was beneficial to the underlying secondary market for corporate bonds. We explore this objective by tracking the efficiency, market quality, and liquidity of an issuer's bonds after CDS trading was instituted on the bonds of the issuer, and also by comparing the bonds of firms with traded CDS contracts to the bonds of firms without any CDS contracts. We find that the advent of CDS was largely detrimental to corporate bond markets, particularly to its efficiency and market quality.

Data

We focus on US-domestic, dollar-denominated, non-convertible corporate bonds of publicly traded firms that witnessed CDS introduction between 2002 and 2008. Our sample consists of 1,545 bonds issued by 350 firms and comprises 1,365,381 bond transactions. In addition, we also collect various issue-, issuer-, and transaction-specific attributes, issuer's equity returns, CDS spreads, systematic VIX values and benchmark interest rate swap rates. We classify the bond transactions into pre- and post-CDS sub-samples based on whether the bond trades occurred before or after the introduction of CDS. We study the consequences of CDS introduction by comparing the efficiency, market quality, and liquidity of bonds in the pre- and post-CDS periods.

CDS introduction adversely affected bond efficiency, ...

We test for bond efficiency by ascertaining whether delays exist in relevant information being incorporated into bond prices. To this end, we determine the extent to which bond prices depend on a lagged information set (relative to contemporaneous information set). Greater dependence on lagged information denotes higher pricing inefficiency, because information already incorporated into other firm-related securities only enters bond prices with a time lag.

We regress contemporaneous bond returns on contemporaneous and lagged values of stock returns, benchmark swap returns, changes in VIX, CDS returns, and lagged bond returns. The regressions include interaction terms that enable comparison between pre- and post-CDS periods (in joint panel regressions) and also between CDS firms and a (matched or pooled) control sample of non-CDS firms (in difference-in-difference regressions). In each regression, we compute the joint significance of incremental lagged variables in order to determine the extent to which current bond returns depend on the lagged variables in the post-CDS period relative to (i.e., over and above) that in the pre-CDS period.

In all regressions, we find that bond returns rely on lagged information to a greater extent after CDS are introduced than before, and this increased dependence persists even when benchmarked against control samples. Similar results obtain in various sub-samples and alternate variations of regression specifications. Incorporation of relevant information into bond prices gets delayed in the post-CDS period. **Conclusion:** The advent of CDS market had a deleterious effect on the efficiency of corporate bond market.

... and bond market quality, ...

How did the inception of CDS trading impact the accuracy of bond prices (which we refer to as the bond market quality)? We develop a measure of market quality called the q measure based on an extension of Hasbrouck's (RFS 1993) model. Market quality q is defined as one minus normalized pricing error. The value of q ranges between zero and one, and higher q denotes better market quality, i.e., lower risk of deviation of prices from their efficient levels. Table 1 reports the values of bond market quality measure q in the pre- and post-CDS periods.

Table 1: Market quality measure q before and after the introduction of CDS

	Pre-CDS mean q	Post-CDS mean q
For pooled sample of all observations		
CDS sample	0.91	0.87
Control sample	0.90	0.91
For 82 pairs of matched CDS and non-CDS bonds		
CDS bonds	0.90	0.88
Non-CDS bonds	0.85	0.92

Pre-CDS q measure is computed using bond transactions over the two years prior to CDS introduction; post-CDS q measure uses two years of observations after the CDS inception date. Control sample (non-CDS bonds) refer to bond issues of firms with no CDS introduction. Sample period is 2002–2008.

For the pooled panel data of all bond transactions, the quality of bonds of CDS issuers decreases by 0.04 and that of control sample bonds slightly increases by 0.01. When the sample of 82 pairs of matched CDS and non-CDS bonds are considered, the quality of bonds of CDS issuers declines by 0.02 but the quality of bonds of CDS non-issuers increases substantially by 0.07. The difference-in-difference

value of post-CDS decline in quality relative to control sample equals $[(0.92-0.85)-(0.88-0.90)]$ or 0.09. In addition, when we track the values of q for individual bonds, we find that a greater fraction of bonds of CDS issuers experience a post-CDS decline in the value of q, whereas a larger fraction of matched control sample bonds demonstrate an increase in the value of q.

In conclusion, on a comparative basis, CDS introduction appears to have a detrimental impact on the market quality of the underlying bonds.

... with no improvement in bond liquidity

A likely consequence of CDS trading is that fixed-income traders no longer need to use bond markets to speculate on or hedge credit risk. Did liquidity in the bond market also suffer following CDS introduction? Figure 1 plots the mean size of trades and Figure 2 the mean turnover for bonds of CDS issuers and bonds of CDS non-issuers over a four-year (500 trading days) window around the CDS introduction date. We observe that trade size as well as turnover of bonds of issuers with CDS contracts fall in the two years following CDS introduction, whereas there are no appreciable changes for control sample bonds.

Figure 1: Mean trade size before and after introduction of CDS, 2002–2008

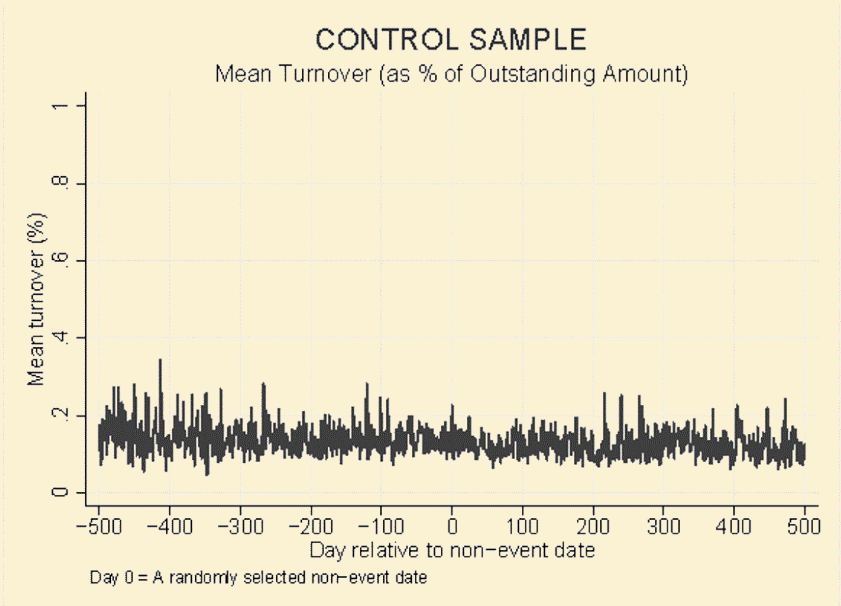

There is post-CDS decline in secondary bond market trading activity. The comparable control sample shows no pattern.

Figure 2: Mean turnover before and after introduction of CDS, 2002–2008

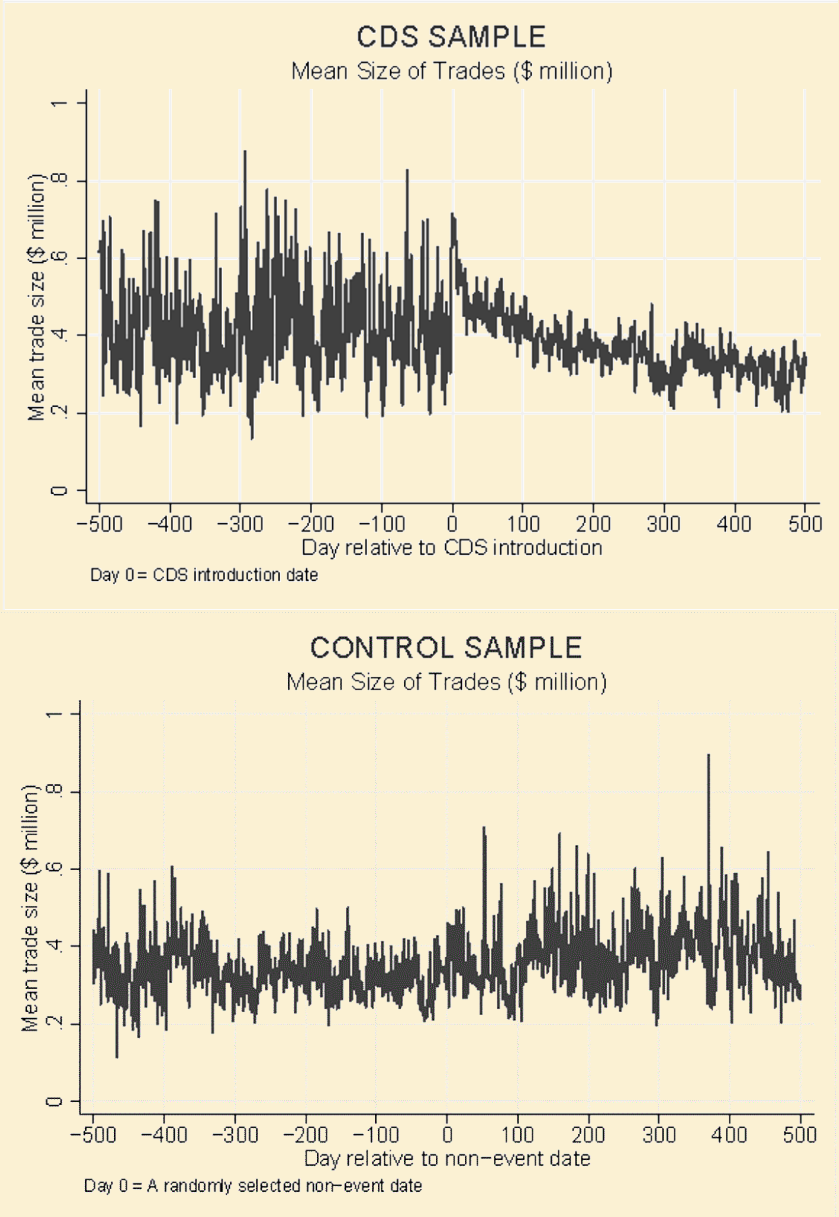

There is post-CDS decline in secondary bond market trading activity.

For formal assessment of post-CDS impact on bond liquidity, we compute ten different measures that are either proxies for liquidity or may be highly correlated to liquidity, and compare their values in the pre- and post-CDS sub-periods. Although results are largely mixed, when all ten liquidity measures are considered, more liquidity attributes deteriorated than improved after the inception of CDS markets. Hence, there is no evidence that CDS introduction improved the liquidity of the bonds underlying the CDS entity; if at all, liquidity likely deteriorated.

A likely explanation for adverse CDS impact: migration of institutional traders

One possible explanation for the decline in efficiency and quality of bond markets subsequent to CDS introduction is the likely migration of institutional traders from trading bonds to trading CDS in order to implement their credit views. Underlying corporate debt often do not traded actively, and institutions likely use CDS markets to incur synthetic exposures to the debt market.

To explore this issue, we track the large institutional bond trades in the TRACE database and the bond transactions by insurance companies in the NAIC database. We find that, from the pre-CDS period to the post-CDS period, the number, volume, and turnover of institutional trades decreased and the LOT illiquidity measure increased relative to the control sample of non-CDS bonds. We also implement the liquidity tests adopted by Bessembinder-Maxwell-Venkataraman (JFE 2006). For trades by insurance companies, we decompose the price changes (i.e., the effective bid-ask spreads obtained from signed order flows) into two components: an informational component that indicates the effect of private information, and a non-informational component that reflects one-way trade execution costs. We find that there is no change in the role of private information on bond price evolution after CDS introduction. However, the post-CDS trade execution costs increase; this reconciles with the decrease in trading activity by insurance companies. Hence, the introduction of CDS increased bond illiquidity for institutional transactions.

In short, a demographic shift in bond trading appears the likely driver of the empirical results we obtain, namely, that the introduction of CDS trading was detrimental to bond market efficiency, quality, and liquidity.

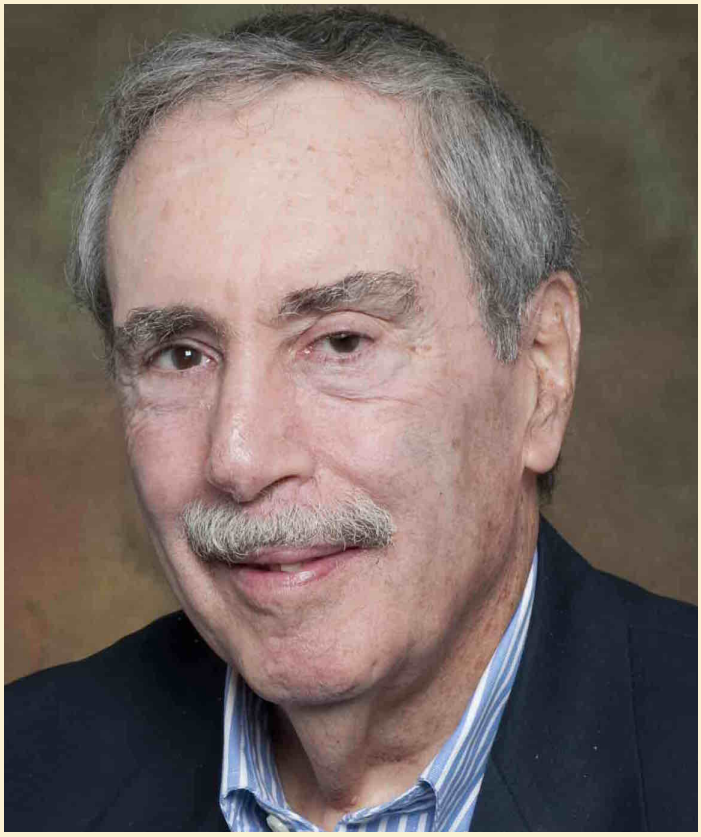

Kee H. Chunga and Chairat Chuwonganant

Uncertainty, market structure, and liquidity

Journal of Financial Economics | Volume 113, Issue 3 (September 2014), 476-499
(please cite only the original publication, not FAME)

It is well known that the liquidity of individual assets tends to move together. Liquidity co-movements could arise because liquidity providers share common capital and information and thus respond similarly to outside shocks. Liquidity commonality could also arise from the trading behavior of liquidity demanders. Our study sheds light on the cause of liquidity commonality by analyzing the relation between market uncertainty and liquidity.

We show that market uncertainty [measured by the Chicago Board Options Exchange Market Volatility Index (VIX)] exerts a large market-wide impact on liquidity, which gives rise to co-movements in individual asset liquidity. Our study also shows that the uncertainty elasticity of liquidity (UEL: percent change in liquidity given a 1% change in VIX) has increased significantly around regulatory changes in the US markets that increased the role of public traders in liquidity provision, reduced the minimum allowable price variation, weakened the affirmative obligation of NASDAQ dealers, and abolished the specialist system on the NYSE.

Market volatility and liquidity commonality

We conjecture that an important source of liquidity commonality is overall market uncertainty using the Chicago Board Options Exchange Market Volatility Index. This index, often referred to as the fear index or the fear gauge, is a measure of the implied volatility of Standard & Poor's (S&P) 500 index options. Indeed, we find that uncertainty plays an important role in liquidity commonality after controlling for aggregate market liquidity, industry-wide liquidity, and well known determinants of individual stock liquidity, such as return volatility, trading volume, and price.

For NASDAQ stocks, we find that a 1% increase in VIX leads to a 0.46%, 0.49%, and 0.43% increase in the quoted spread, effective spread, and price impact, respectively, and a 0.41% and 0.49% decrease in the quoted depth and market quality index, respectively. We find similar results for NYSE stocks. We also show that the effect of VIX on stock liquidity is greater than the combined effects of all other common determinants of stock liquidity. Apparently, liquidity providers react strongly to general market uncertainty (in addition to individual asset risk), and this behavior generates a commonality in liquidity.

Our results support the view that market liquidity decreases with VIX because higher volatility tightens funding constraints on market makers and thereby reduces their liquidity-provision capacity. Liquidity could disappear during periods of financial market turmoil because liquidity providers demand a higher expected return from liquidity provision during such times. Our results are also consistent with anecdotal evidence that liquidity providers adjust their positions uniformly across stocks according to market volatility as reflected in VIX before they make adjustments based on individual asset risks.

Stock attributes and the uncertainty elasticity of liquidity

Our paper shows that UEL is positively and significantly related to share price, trading volume, and return volatility and negatively and significantly related to firm size (market value of equity) when liquidity is measured by the quoted spread, effective spread, or price impact. When liquidity is measured by the dollar depth or market quality index, the signs of the estimated coefficients on share price, trading volume, firm size, and return volatility are the opposite of those when liquidity is measured by the spreads and price impact: UEL is positively and significantly related to firm size and negatively and significantly related to share price, trading volume, and return volatility. These results suggest that uncertainty exerts a greater impact on the liquidity of stocks with higher price, larger trading volume, greater return volatility, or smaller market capitalizations. The larger UEL of higher priced stocks could arise because the tick size is less likely to be a binding constraint on price quote changes. The larger UEL of high-volume or high-volatility stocks could be due to the fact that these stocks are more likely to exhibit information-based trading, prompting greater reactions from liquidity providers. Uncertainty may exert a smaller impact on the liquidity of larger companies because these companies have lower information-based trading as more information is available about them and, thus, liquidity providers in these stocks may be less sensitive to changes in market volatility. Uncertainty exerts a smaller impact on the liquidity of stocks with higher analyst following or institutional ownership, or both. This could be because liquidity providers in these stocks are less sensitive to changes in market volatility because more information is available through analysts' information collection and dissemination as well as institutional monitoring. We find that uncertainty exerts a larger impact on the liquidity of firms with higher market-to-book ratios. Such firms may have larger intangible assets (e.g., higher proportions of these firms' market values are accounted for by future growth options) and changes in market volatility lead to larger swings in liquidity when firm value is highly subject to future managerial actions.

Risk premium structure and the uncertainty elasticity of liquidity

We conjecture that uncertainty exerts a larger impact on liquidity when public traders play a greater role in liquidity provision, when the minimum price variation (i.e., tick size) is smaller, and when market makers play a smaller role in liquidity provision. We test these conjectures using the following four regulatory changes, which serve as natural experiments, in market structure: (1) the implementation of the new order handling rules (OHR) on NASDAQ in 1997, (2) the reduction of tick size from $1/8 to $1/16 in 1997 and from $1/16 to $0.01 (decimalization) in 2001, (3) the amendment of NASDAQ Rule 4613(c) in 2007 that dealer quotes must be reasonably related to the prevailing market, and (4) the implementation of the designated market maker (DMM) system on the NYSE in 2008.

We show that the effects of VIX on liquidity have changed over time across different market structures. Specifically, we show that the uncertainty elasticity of liquidity (UEL) has increased dramatically around regulatory changes in the US markets that increased the role of public traders in liquidity provision, reduced the minimum allowable price variation, weakened the affirmative obligation of NASDAQ dealers, and abolished the specialist system on the NYSE. These results support the idea that a direct reflection of expected volatility in prices and quotes, without the filtering by market intermediaries, could increase the volatility of market liquidity. Although some of these regulatory changes have been shown to increase pricing efficiency and reduce trading costs, they might also have the unintended consequence of increasing liquidity volatility, leading to an increase in liquidity risk premiums in asset returns.

Conclusion

Prior research shows that the return premium is related to commonality in liquidity with market liquidity, return sensitivity to market liquidity, and liquidity sensitivity to market returns. To the extent that the market volatility-induced liquidity risk cannot be diversified away, such systematic liquidity risk would also be built into the required return of investors. It would be interesting to find out whether stocks with greater UELs command higher return premiums after controlling for other liquidity-related risk factors shown in prior research. Our results show that the uncertainty elasticity of liquidity is greater than the liquidity commonality with the market liquidity, suggesting that the risk premium associated with the former could be even greater than the risk premium associated with the latter.

Joseph Gerakos and Juhani T. Linnainmaa

Market reactions to tangible and intangible information revisited

Critical Finance Review | Volume 5, Issue 1 (2016), tbd
(please cite only the original publication, not FAME)

Daniel and Titman(2006) posit that the value premium is due to investors overreacting to intangible information. They therefore decompose five-year changes in book-to-market ratios into stock returns and a residual that is a proxy for tangible information based on accounting performance ("book returns"). That is, their decomposition is the following identity:

$$\underbrace{dbm(t,t-5)}_{\text{change in log-BE/ME}} \equiv \underbrace{r^B(t,t-5)}_{\text{book return}} - \underbrace{r(t,t-5)}_{\text{stock return}} . \tag{1}$$

Consistent with their thesis, they find that it is only when stock prices move in response to intangible information—the return component orthogonal to book returns, obtained by regressing stock returns on book return and lagged book-to-market—that returns reverse. They conclude that the book-to-market ratio forecasts returns because it is a good proxy for the intangible return.

Problems with the decomposition

We show that changes in book-to-market ratios cannot be decomposed into book and stock returns. A simple example illustrates the problem. Suppose a firm has \$1 in book value of equity and \$2 in market value of equity, and so its book-to-market ratio is 1/2. If the firm now issues \$1 worth of additional equity, both the book and market values of equity increase by \$1, and the book-to-market ratio increases to 2/3. Because the Daniel-Titman decomposition factors out stock returns—which are neutral with respect to net issuances and dividends—from changes in the book-to-market ratio, book return captures this increase in full by being distorted upward by $\log(2/3) - \log(1/2) \approx 29\%$.

Book return combines return-on-equity with a nonlinear bundle of past book-to-market ratios and contemporaneous stock returns, net issuances, and dividends. This problem occurs because changes in market and book values of equity are additive and not multiplicative when firms pay dividends or issue or retire equity. Hence, it is not possible to factor out stock returns from changes in the book-to-market ratio without creating a book return correlated with past book-to-market ratios, stock returns, net issuances, and dividends.

Approximation of the error in the decomposition

A cross-sectional regression of the book return on the return on equity measures the approximation error in book returns. For the Daniel-Titman decomposition to hold, such a regression should have an R^2 of one. Instead, we find an average R^2 is 0.57, implying that two-fifths of the cross-sectional

variation in book returns is due to the approximation error induced by factoring out stock returns from changes in the book-to-market ratio. These results are not surprising given that firms rarely conform to the model under which the Daniel-Titman decomposition would hold. A majority of firms issue or repurchase equity or pay dividends during the typical five-year period, thereby introducing approximation error into book returns.

How big is the approximation error if we view book return as being equal to the return on equity plus noise?

$$r^B(t, t-5) = roe(t, t-5) + \text{approximation error}. \tag{2}$$

Among growth firms—the lowest quintile of firms in the BE/ME distribution—the average five-year return on equity is 0.77, the average approximation error is –0.15, and the standard deviations of return on equity and the approximation error are 0.71 and 0.60. Among value firms the average five-year return on equity is 0.33, the average approximation error is 0.08, and the standard deviations are 0.40 and 0.32. These numbers show, first, that the average approximation error flips sign as we move from growth to value and, second, that the standard deviation of the approximation error nearly equals that of the return on equity.

Net issuances and dividends

To illustrate the importance of net issuances and dividends in the evolution of book-to-market ratios, we decompose changes in the book-to-market ratio into changes in the book and market values of equity. Among all-but-microcaps, net issuances and dividends explain one-third of the cross-sectional variation in changes in the book value of equity, and one-quarter of the variation in changes in the market value of equity. These sources of variation are important in return regressions. Without a control for net issuances, changes in the book value of equity barely correlate with future returns. The reason is that the book value of equity can increase either because the firm is profitable or because it issues new equity. But profitable firms have high returns and firms issuing equity have low returns. Therefore, changes in the book value of equity *alone* are not informative about future returns.

The approximation error in book returns makes Daniel and Titman (JF 2006) projection of stock returns on book returns ill-suited for delineating between tangible and intangible returns. Even in an economy in which accounting performance is fully disconnected from firm valuations, stock returns *are* correlated with book returns in the Daniel and Titman decomposition. Furthermore, net issuances and dividends are counted as book returns, and their signs depend on whether a firm is classified as value or growth. If a growth firm, for example, issues equity, this action shows up as a high book return. That is, the firm appears to have good tangible performance just because it issued equity. Value firms that issue equity, in contrast, look worse than their peers that do not issue equity. Given that book returns are substantially polluted, cross-sectional differences in book returns have little economic content and provide no insight into the source of the value premium.

Tangible versus intangible information

Questions regarding the effects of tangible and intangible information are interesting but difficult to resolve empirically. We show that the data are at odds with the argument that book-to-market predicts returns because "it is a good proxy for the intangible return" (Daniel and Titman (2006) p. 1605). Current book-to-market fully subsumes intangible-return proxy's ability to predict the cross-section of average returns. This analysis suggests that book-to-market does not predict future returns because it

is a good proxy for the intangible return. Instead, it appears that the intangible return is a good proxy for the current book-to-market ratio. It also suggests that there is more to book-to-market's predictive power than its correlation with what Daniel and Titman call the intangible return. We also show that Daniel and Titman conclusions reverse when we change the tangible-information proxy from book return to the change in the book value of equity. Under this specification the *tangible* returns reverse. We do not suggest that this is the correct conclusion. Rather, we view it as an example of the difficulty in separating the effects of tangible and intangible information.

The question of why book-to-market ratios correlate with future returns remains important. Fama and French (JF 2008) show that recent changes in book-to-market ratios carry more information about expected returns than historical book-to-market ratios, but their results shed no light on what this information might entail. Gerakos and Linnainmaa (2013) use Fama and French (JF 2008) decomposition to divide the HML factor into two parts, and suggest that these parts carry different prices of risk. They show that book-to-market ratios' ability to explain variation in future returns traces back to changes in the market value of equity, and suggest that these changes are important because they pick up *changes* in expected returns. If a stock's expected return increases, its instantaneous return is negative, thus turning it into a value stock, and vice versa.

Clifton Green, Russell Jame, Stanimir Markov, and Musa Subasi

Broker-hosted investor conferences

Journal of Accounting and Economics | Volume 58, Issue 1 (Aug 2014), 142–166
(please cite only the original publication, not FAME)

Institutional investors allocate billions of dollars in commissions each year as a payment for brokerage research. Conventional wisdom equates brokerage research with the distribution of stock recommendations and earnings forecasts, yet institutional investors routinely rate "access to management" over all published forms of research. One way brokerage firms meet investors' demand for access to management is by organizing conferences which connect their institutional clients with company executives. Broker-hosted investor conferences are invitation-only events that include formal company presentations as well as private breakout sessions with select investors and management.

We address three questions. What types of firms attend broker-hosted investor conferences? Do institutional investors reward brokerages for facilitating access to management? Do firms benefit from participation in investor conferences?

What types of firms participate in broker-hosted investor conferences?

Our conference data are obtained from Bloomberg Corporate Events Database, and they include information on the broker host and the firms in attendance for the period 2004-2008. We merge the conference data with brokerage stock recommendation data from IBES, institutional transaction data from Ancerno, and company information from CRSP and Compustat. Of the 108 research brokers in our sample, 66 brokers host conferences, with a broker average of 7 conferences per year attended by roughly 24 companies. Hosting brokers tend to be larger, and analyst hosts tend to be more experienced and reputable.

Conferences often have an industry focus, and we first explore whether brokers emphasize certain industries. Table 1 shows healthcare firms attend more conferences than firms in other industries. We also create a hosting-to-covering ratio to examine brokers' relative propensity to host rather than cover firms through published research. We find brokers are most likely to emphasize conferences relative to traditional research for healthcare and drug firms as well firms in the business equipment industry (which includes tech firms). On the other end of the spectrum, they are relatively least likely to host utility or energy firms.

Table 1: Brokerage Research by Industry

	Brokers Hosting at Conferences	Brokers Covering with Research	Hosting/Covering Ratio
Healthcare & Drugs	2.75	3.70	0.74
Business Equipment	1.92	4.55	0.42
Telephone & Television	1.69	4.55	0.37
Manufacturing	1.11	3.06	0.36
Other	1.00	2.94	0.34
Wholesale & Retail	1.43	4.24	0.33
Consumer Durables	0.88	3.01	0.29
Consumer Non-Durables	0.85	3.07	0.28
Utilities	0.97	3.62	0.26
Oil, Gas, & Coal	1.13	5.36	0.21

Healthcare firms attend more conferences than other firms.

Brokers Hosting at Conferences is the firm-level average of the number of brokers hosting the firm. Brokers Covering with Research is defined similarly. For example, the average Healthcare firm attends 2.75 conferences per year and is covered with published research by 3.7 brokers. Hosting/Covering is the ratio of these two variables and shows the relative propensity to host a firm at a conference.

We conjecture that the industry variation in hosting activity is driven by differences in valuation complexity. For example, biotech firms are harder to value than utility firms, and therefore they garner more invitations to conferences where investors can meet with management and ask questions to flesh out their understanding of the company. We explore this idea using two proxies for valuation difficulty at the firm level: intangible assets (e.g., patents) and R&D plus advertising expenditures, which are harder to value than other capital expenditures (e.g., to purchase property).

We examine the relation between these and other firm characteristics and four possible categories of broker research output: (1) no research at all, (2) published research only, (3) hosting at conferences only, and (4) both published research and hosting. Not surprisingly, large firms with high institutional ownership and greater turnover are more likely to receive both types of research. Our interest is in contrasting hosting with publishing. Consistent with there being greater client demand for management access among hard-to-value firms, we find firms with higher levels of intangible assets and greater R&D and advertising expenses are more likely to be hosted at a conference than covered with published research. This relation also holds when we examine the number of conferences attended relative to the number of reports published. Finally, we find young companies and those that are planning to issue equity are more likely to attend broker-hosted conferences, suggesting firm incentives to seek investor recognition also play an important role in conference attendance.

Do institutional investors reward brokers for hosting conferences?

Institutional investors pay for brokerage research through trading commissions, and we examine whether brokers attract greater commissions in conference stocks. Our commission data are from Ancerno, a consulting firm which helps institutional investors monitor transaction costs. Their clients account for roughly 8information on transaction volume and price, as well as which broker executed the trade and the level of commission paid. Because the data allow us to measure a broker's commission share in a given stock over a particular period, we are able to empirically investigate the incremental benefits of hosting a firm at a conference.

Figure 1 shows average monthly commission share for a stock after hosting the firm at a conference. Host commission share peaks at roughly 5.5% in the month of the event and gradually 3 falls to 4.6% a year after the conference. The 4-5% level of commissions is more than twice the average stock-level commission share among conference hosting brokers (1.8%), which suggests conference participation may proxy for additional research services such as a company visits or non-deal road shows.

Figure 1: Monthly commission share following investor conferences

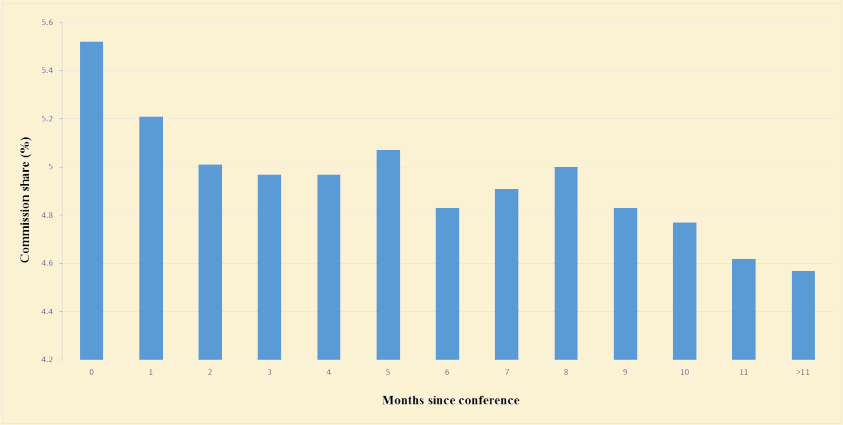

Monthly commission share for a stock after hosting a firm at a conference leaks in the month of the event and gradually falls afterwards.

Focusing on the month of the conference, we find commission share is largest during the first week after the conference, suggesting conferences have an immediate effect on commissions. Moreover, informative conferences, as measured by price changes or trading volume in the days around the event, lead to greater increases in weekly commission share. This result suggests institutional investors compensate brokers for facilitating the transfer of market moving information.

Table 2 examines the importance of investor conferences for broker commissions relative to traditional published research. The regression coefficients in the first column show that hosting a company at a conference increases commission share by 3.4%, which is similar to the effect of covering a company with stock recommendations (2.2%, with each recommendation raising commission share by 1.4%). One concern with our approach is that conference hosts tend to be large brokers, so comparing hosts with non-hosts may be picking up difference in broker size or reputation. We address this by including broker fixed effects. The last column of Table 2 compares commission share for conference and non-conference stocks within a given broker. We observe that the effect of a hosting a firm is weakened but still comparable to the effect of coverage with traditional research.

The effects of investor conferences are economically meaningful, with conferences having a sizable impact on hosting brokers' annual revenues. Based on conservative estimates, we calculate that hosting an additional conference is associated with a $260,000 increase in annual commissions directly in conference stocks, and we note that investors may reward brokers with additional trading in non-conference stocks as well.

Table 2: Explaining monthly commision share with brokerage research

Conference in the last year	3.36	1.89
	[11.48]	[7.20]
Coverage in the last year	2.15	1.19
	[7.85]	[9.38]
Number of Recommendations	1.37	0.88
	[15.93]	[11.03]
Fixed Effects	Intcpt	Broker
Adjusted R-squared	2.48%	8.25%

Hosting a company at a conference increases commission shares. The T-statistics are in parentheses.

We regress broker commission share (at the individual stock level) on measures of brokerage services and control variables. Conference in the last year is equal to one if the firm attended a conference hosted by the broker in the past year. Coverage in the last year is equal to one if the broker issued a stock recommendation on the firm in the last year. Number of Recommendations is the total number of recommendations issued by the broker on the firm over the past year. Bother regressions include controls not tabulated for brevity.

How do firms benefit from participating at conferences?

Anecdotal evidence suggests firms view broker-hosted investor conferences as a valuable tool for communicating with the investment community. Conferences provide unique opportunities to discuss the company's value proposition with current and potential investors and build connections with equity analysts. According to a 2009 Investor Relations survey by Hanley & Associates, broker-hosted investor conferences are the second most significant channel of investor outreach, after non-deal road shows and ahead of investor visits of company headquarters (Gedvila, 2010).

We first consider whether conference attendance improves investor recognition. Table 3 reports the coefficients of regressing analyst following, institutional ownership, and a measure of stock liquidity (bid-ask spread) on variables related to conference attendance. We find that firms experience an increase in analyst coverage of 0.34 analysts and a 1.24% increase in the level of institutional ownership. These effects are economically meaningful relative to the mean coverage of 4.9 analysts and institutional ownership of 52%. The effects are also significantly larger if the conference was the first one attended

by the firm in over a year. The improvements in investor recognition following conferences also translate into enhanced liquidity. Bid-ask spreads fall by 6%, with an additional 3% if the conference was the first in a year.

To gauge the effect of conferences on the market valuation of a firm, we examine changes in the cost of equity capital and Tobin's q. Our cost of capital measure, standard in the accounting literature, is implied from estimates of future earnings based on past earnings and other financial statement variables. The last two columns of Table 3 show the results. We find that conference participation is associated with a decline in cost of equity capital of 0.87%. Relative value, as measured by Tobin's q (market value of assets over book value of assets), increases by 0.03, although neither measure changes significantly more so when the conference is the first in a year. Taken together, our findings suggest conferences convey important benefits to firm participants.

Table 3: Firm Benefits of Conference Participation

	Change in Analyst Following	Change in Institutional Ownership	Change in Bid-Ask Spread	Change in Cost of Capital	Change in Tobin's q
Conference	0.34	1.24	–0.06	–0.87	0.03
	[13.06]	[9.64]	[–12.49]	[–10.45]	[3.17]
First Conf. in a Year	0.08	1.51	–0.03	–0.07	0.00
	[2.17]	[7.87]	[–4.03]	[–0.72]	[0.08]
(Industry and year fixed effects are always included)					
Adjusted R^2	9.41	15.48	43.27	28.91	34.44

Conference participation is associated with a decreased cost of capital and higher Tobin's Q. An increase in analyst following an increase in institutional ownership. The T-statistics are in parentheses.

The table shows the coefficients from regressions of annual changes in firm outcomes on variables related to conference participation. Conference is an indicator variable that is one if the company participates at a conference in the current quarter (zero otherwise). First Conference in a Year is one if the company did not present in the previous three quarters (zero otherwise). Each regression includes controls that we do not report for brevity.

Conclusion

Broker-hosted investor conferences are invitation-only events that allow access to management for select institutional clients. We find that firms with high institutional ownership and high intangibles assets attend more conferences, consistent with brokers catering to institutional clients and greater client demand for management access when valuation difficulty is high. We also find that younger firms and those that issue equity in the future attend more conferences, consistent with firms viewing conference

participation as a vehicle for increasing investor recognition. Finally, we document substantial benefits for brokers, in the form of increased trading revenues, and for firms, in the form of improved investor recognition, liquidity, and valuation. Our findings point towards broker-hosted conferences as an equilibrium arrangement benefitting both the broker-host and firm-participants, and underscore the interconnected nature of corporate investor relations activities and brokerage research activities.

Han Ozsoylev, Johan Walden, M. Deniz Yavuz, and Recep Bildik

Investor Networks in the Stock Market

Review of Financial Studies | Volume 27, Issue 5 (May 2014), 1323-1366
(please cite only the original publication, not FAME)

Heterogeneous and decentralized spread of information is known to influence the trading behavior of investors, see, e.g., Shiller and Pound (JEBO 1989), Ivkovic and Weisbenner (RFS 2007), and Hong, Kubik, and Stein (JF 2004). Such so-called *information diffusion* may help explain several fundamental stylized facts of stock markets. First, investors are known to hold vastly different portfolios. Such heterogeneous behavior is what one would expect when there is decentralized information diffusion in the market, but is contrary to the prediction of classical models that everyone should hold the market portfolio. Second, stock markets are known to experience large movements that are unrelated to public news, as documented in Cutler, Poterba, and Summers (JPM 1989) and Fair (JB 2002), suggesting other channels than public information channels through which information is incorporated into asset prices. Third, the dynamics of asset returns and trading volume are known to be very rich, see Karpoff (JFQA 1987), Gallant, Rossi and Tauchen (RFS 1992), Bollerslev and Jubinski (JBES 1999), Lobato and Velasco (JBES 2000), and Gabaix, et al. (NATURE 2003). Returns and volume in many markets are "heavy-tailed" (meaning that the risk for extreme events is much higher than in standard models), time varying, have "long memory" (meaning that shocks are very persistent), and are related to each other in a complex manner. Lumpy, heterogeneous information diffusion provides a potential explanation for such rich behavior.

Model and Results

We use a network approach to model information diffusion in a stock market, and study how the network position of an investor affects performance and trading behavior with respect to information events in the market. Loosely speaking, an information network describes how diverse information spreads over time among a population of investors, as shown in Figure 1. A fundamental property of these models is that more *central* agents in the network tend to access and trade on information earlier than less central agents, as well as make higher profits, see Ozsoylev and Walden (JET 2011) and Walden (WP 2014). Centrality here captures the important property that it is not only how many people you know that is important, but also *who* you know. An investor's centrality therefore does not only depend on how many neighbors that investor has but also on the centrality of those neighbors, in an endogenous fashion. In Figure 1, for example, the implication is that investor 1 is more central than investor 2, although investor 2 actually has more neighbors than investor 1.

Figure 1: [Information network]

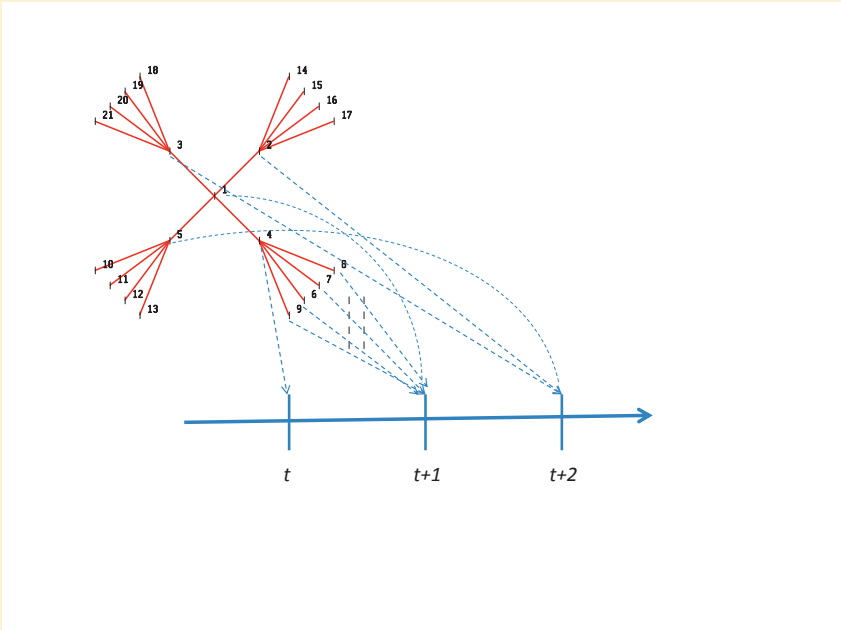

The figure shows an information network of 21 investors in a market. Investors are represented by circles and links between investors by lines. In the Figure, investor 4 receives a valuable piece of information at time t, trades, and then shares it with his neighbors at $t+1$, who then trade and share the information with their neighbors at $t+2$. Information thus diffuses through the network in a well-defined manner.

Investors who are centrally placed in the network tend to receive information signals early, whereas investors who are in the periphery tend to receive them late. As a result, the trading behavior and profitability of individual investors are influenced by their position in the network, and the dynamics of aggregate asset prices depend on the network's structure.

We develop a method to approximate the market's information network, using observable data. The general idea is that information links may be identified from realized trades, since investors who are directly linked in the network tend to trade in the same direction in the same stock at a similar point in time. Using this approach, we identify an Empirical Investor Network (EIN), and in simulations show that the true information network is indeed well estimated by the EIN. We calculate the EIN using account level data that cover all trades on the Istanbul Stock Exchange in 2005. We find that our method is quite robust to omitting a significant fraction of investors from the sample, suggesting that our method can also be used with nonexhaustive datasets.

We then study the relationship between investor centrality and returns, and find substantial support for a positive relationship. A one-standard deviation increase in centrality, all else equal, leads to a 0.7%-1.8% increase in returns for an investor over a 30-day period, depending on the specification. These results are obtained after controlling for other variables, such as trading volume. We also document that centrality is directly related to acting early on information. We identify several idiosyncratic information events that were associated with large stock price movements, and find that central investors in the network tended to trade—in the right direction—before peripheral investors. Our results suggest that information diffusion is an important determinant of investors' trading behavior and profitability.

Concluding Remarks

Our study reinforces a view of the stock market as a place where information is gradually incorporated into asset prices. Information networks provide an intermediate information channel, in between the public arena, where news events and prices themselves make some information available to all investors, and the completely local arena of private signals and inside information.

Our results suggest a decentralized diffusion mechanism through localized channels, e.g., social networks, in line with several recent studies that focus on specific investor groups. Our knowledge is still limited, however. Which factors determine a market's information network? Geographical location? Social networks? Other channels? Given datasets with more detailed information about investors in the market, further research may shed light on this important question.

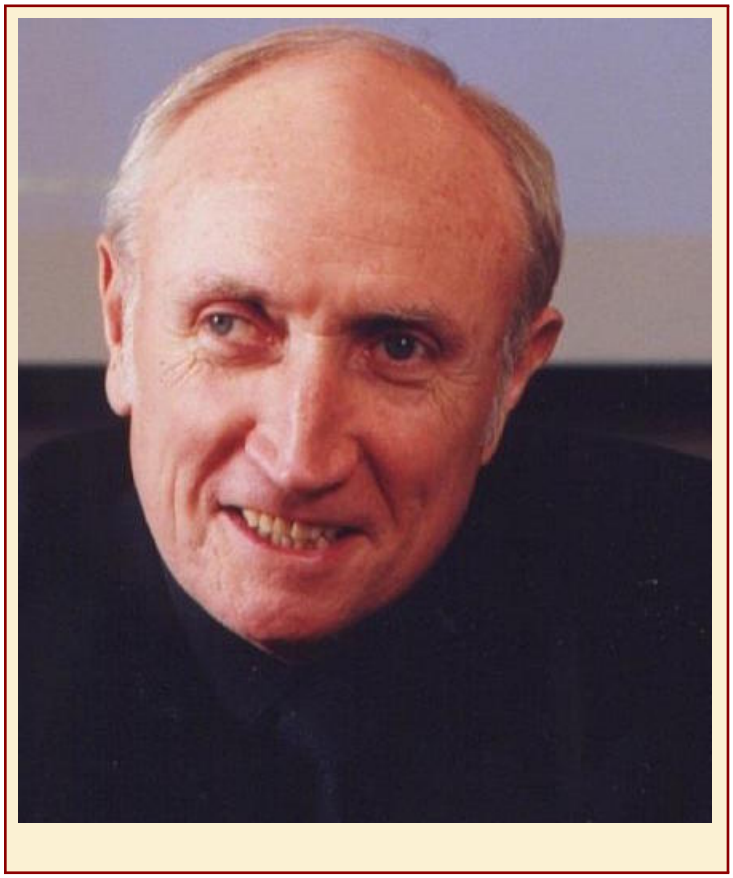

Marcin Kacperczyk, Stijn van Nieuwerburgh, and Laura Veldkamp

Time-Varying Fund Manager Skill

Journal of Finance | Volume 69, Issue 4 (Aug 2014), 1455-1484
(please cite only the original publication, not FAME)

Do fund managers have investment skill? If so, what is the nature of this skill? Previous literature thought about market timing and stock picking as separate and immutable skills that once present, should manifest themselves in every trading period. The evidence of such immutable skill is controversial. Contrary to this premise, we show that the ability of funds to time the market and pick stocks varies over the business cycle. The findings suggest that skill is a general cognitive ability that can be used in different ways - to pick profitable stocks or to time the market - at different times. When we measure skill, allowing its manifestation to change in booms and recessions, we find stronger evidence that a subset of managers indeed possesses skill.

Measuring Skill and Its Business Cycle Dependence

Investors use skill to form portfolios that outperform the average investor. We decompose this outperformance into two categories: If an investor times the market, he earns excess returns by being more exposed to the market portfolio in periods when the realized market return ends up being high and being less exposed when the realized market return is low. Stock picking means the investor earns excess returns by holding more of a stock in periods when its realized return ends up high. Thus, Timing (Picking) skill is measured as a covariance of the portfolio weights today with the systematic (idiosyncratic) component of realized returns in the next period.

The main premise of our paper is that skill manifests itself in a time-dependent fashion. To test this claim empirically, we compare fund manager skill over the business cycle: during booms and recessions. Our definition of business cycles follows that of National Bureau of Economic Research (NBER) but the results we report are qualitatively similar for other measures of economic conditions, including measures that are available in real time.

Using the data on portfolio holdings and returns of the universe of U.S. equity mutual funds (1980-2005), we compute skill measures for each fund in each month. The data set, from Thomson Reuters and CRSP, includes 3477 distinct funds and 250,219 fund-month observations. The number of funds in each month varies between 158 in May 1980 and 1670 in July 2001.

To test the significance of the relationship between skill and business cycles, we estimate the linear regression model, in which we relate each measure of skill (Timing and Picking) individually to an indicator variable (Recession), equal to one any time the economy enters the NBER recession period, and zero otherwise. We control for several fund features that might correlate with business cycles and skill, such as: fund age, assets under management, expenses, turnover, net flow, load fee, and style measures (size, value, and momentum).

Main Results

The main results document the highly cyclical use of skill. On average, Timing is 1.67 percentage points per year higher in recessions than in expansions, whereas Picking is 1.75 percentage points per year lower in recessions. In both cases, this amount of skill represents excess fund returns of about 14 basis points per month. The results remain similar if we measure skill at the fund manager rather than fund level, which suggests that the skill is largely a fund-manager attribute. Even though our results point to strong time-variation in measured skills for an average fund in our sample, a closer look reveals that the effect is largest for the most successful funds: The effect of Recession on Timing for extremely successful managers is about four times larger than that for the median manager, with a (return) difference of 2.3% per year. The effect on *Picking* doubles.

In another test, we find that it is actually the same manager who has a high level of Timing skill in recession that has also a high level of Picking skill in booms. This result indicates that applying skills indeed changes in a time-dependent fashion.

Implications for Fund Investors

The last result also suggests that some managers may outperform other funds in a consistent fashion. We test this claim formally by selecting a subset of 25% of funds with the highest level of Picking skill in expansions and comparing the performance of this group to a passive investment strategy. Depending on the passive benchmark specification, we find that these funds outperform their passive benchmarks by 48 to 82 basis points per year, an economically significant margin.

The group of outperforming funds is signified by being younger, of smaller AUM, and more active in its investment strategies. Its managers are also more likely having an MBA degree, and ultimately are more likely to be hired by hedge fund companies, thus underscoring their success in the mutual fund business.

In our final set of results, we attempt to predict who a successful fund manager might be. Following the idea that the most successful managers display picking skill in expansion and timing skill in recession, we define an individual Skill Index that weighs each manager's Timing and Picking according to the real-time probability of being in a recession. Thus, Timing receives greater weight in the Index when a recession is more likely. We show that our Skill Index can identify funds/managers which will outperform the market in the subsequent 1-12 months: A fund with a Skill Index value that is one standard deviation above the average value outperforms the average fund by 1-2% per year.

Economic Interpretation

Our findings raise the following questions: Why does picking skill seem to be present in booms, while in recessions any ability to pick stocks vanishes? Why do managers who are effective market-timers in recessions lose that ability in booms? In a theoretical companion paper entitled "A Rational Theory of Mutual Funds' Attention Allocation," we examine how a fund manager with limited ability to observe or process information should allocate his attention between stock-specific and aggregate shocks. The model teaches us that paying attention to and learning about the aggregate state is most valuable in recessions because there are times when aggregate shocks are relatively more volatile and the market price of aggregate risk is high. The model predicts that, in booms, a skilled manager should allocate more attention to stock-specific shock and will therefore be more successful at stock picking.

Similar attention allocation theories have been used to explain a variety of other economic phenomena, including portfolio under-diversification (Van Nieuwerburgh and Veldkamp, RES 2010), portfolio home bias (Van Nieuwerburgh and Veldkamp, JF 2009), and capital income inequality (Kacperczyk, Nosal, and Stevens, NBER 2014).

Conclusions

The findings in our paper touch on fundamental issues such as market efficiency, the dynamics and cross-section of equity returns, as well as what practical portfolio advice to give households. The results also point to a new understanding of what funds do and how they add value for investors. The theoretical framework paints a new picture of what a fund manager who has limited ability to process financial information should do to maximize investor value. Of course, there are many frictions that could distort managers' incentives. But it is still useful to have a clear idea of what the undistorted, efficient manager's investment strategy would look like. The theory offers a host of new testable predictions to contribute to the debate on fund skill and the role of delegated portfolio management in financial markets. In future work, our findings and our framework can be used to analyze the effects of fund entry, rising information processing capacity, or long-run shifts in aggregate or idiosyncratic volatility.

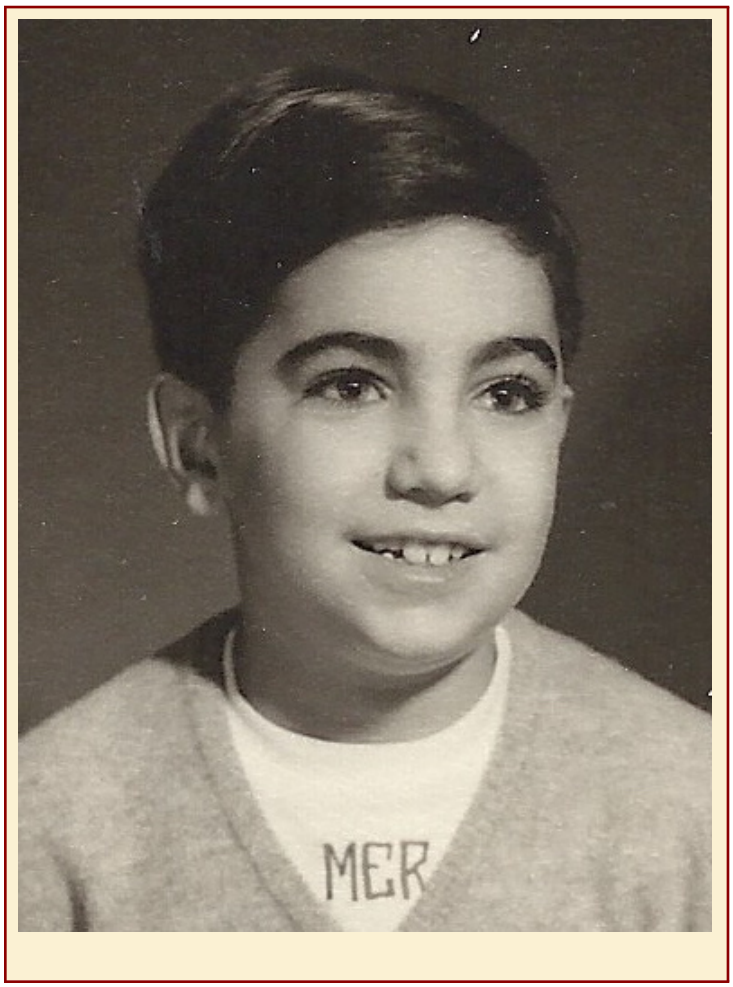

Dong Lou

Attracting Investor Attention through Advertising

Review of Financial Studies | Volume 27, Issue 6 (Jun 2014), 1797–1829
(please cite only the original publication, not FAME)

Recent research has found that advertising has an important impact on the liquidity and breadth of ownership of stocks. This is intriguing as advertising is intended to increase the awareness of a firm's products rather than its securities. Nevertheless, there appears to be a spillover effect. I start by providing evidence of the spillover effect of advertising on firm valuation. In particular, I show that an increase in advertising spending is accompanied by a contemporaneous rise in retail buying and higher abnormal stock returns, and is followed by lower future returns. I then examine whether firm managers are aware of this spillover effect of product-market advertising. Evidence from insider sales, as well as seasoned equity offerings and stock-financed acquisitions, appears consistent with the view that managers opportunistically adjust advertising spending, in part, to influence short-term stock prices.

The temporary stock-return effect of advertising

There are good reasons to believe that advertising can have a temporary impact on firm valuation. First, a typical investor has to search through thousands of stocks when making a buy decision but only through the limited number of stocks he already holds when making a sell decision. To the extent that attention is a scarce resource, investors are more likely to buy attention-grabbing stocks than to sell them. As such, advertising, which is designed to attract attention, can temporarily boost stock value by generating more buy orders than sell orders. Second, although advertising almost never portrays the product or firm in a comprehensive, objective manner, investors with limited attention may take advertisements at face value and react overly optimistically, thus resulting in a temporary stock price overshooting.

This prediction of a spillover effect from product market advertising to financial market valuation is borne out in the data. Firms in the top decile ranked by year-to-year changes in advertising spending outperform those in the bottom decile by almost 13% in the ranking year, but underperform by a combined 15% in the following two years. Adjusting the portfolio returns for common risk factors has virtually no impact on this return pattern. In addition, the documented return effect is significantly stronger for the subsets of firms with lower analyst coverage, lower institutional ownership, higher retail trading intensity, as well as for firms whose brand names are more reminiscent of the firm name.

To further pin down the underlying mechanism of the documented return pattern, I examine trading behavior of retail vs. institutional investors. Since retail investors are usually more attention and resource constrained than their institutional counterparts, we expect that retail investors are

more susceptible to the influence of advertising and thus the net buyers of firms with increased advertising expenditures. Regression results confirm this prediction: a one-standard-deviation increase in advertising spending is associated with a 3.4% increase in small-trade imbalance (which measures retail investors' propensity to buy a stock) in the same year. Interestingly, there is also a 4.3% increase in short interest (i.e., the amount of shares shorted divided by the total number of shares outstanding) in the same period, suggesting that institutions are taking the other side of the trades.

Advertising around equity sales

If product-market advertising can attract investor attention and temporarily boost firm value, a natural question to ask is whether firm managers are aware of this temporary return effect and the extent to which managers adjust firm advertising to exploit investors' bounded rationality. Anecdotal evidence suggests that managers indeed use advertising to influence market perceptions and firm valuation. An October 2003 issue of the Wall Street Journal reports: "United Technologies Corp has launched an advertising campaign focused on the Wall Street area and a Times Square building looking into a Morgan Stanley trading room [...] seeks to overcome the view that it is steady, but not a star and to correct what it believes is a 20% discount in its share price against those of peers."

To empirically test managers' opportunistic behavior, I examine variation in advertising spending around times when short-term stock prices matter the most—i.e., around equity sales by both top executives and the firm itself. The main prediction is that we should observe a sharp increase in advertising spending before equity sales (in order to pump up the stock price), and a significant decrease in advertising spending in the subsequent year.

There indeed is an inverted V-shaped pattern in advertising spending around share sales by top executives in the firm (e.g., the chief executive officer, president, chairman of the board). The average advertising spending in the years prior, contemporaneous, and subsequent to insider sales is 5.3% higher, 6.9% higher, and 3.9% lower than that in other years, respectively. Taking the mean annual advertising spending of $42 million, these coefficients imply that firms in the years prior, contemporaneous, and subsequent to insider sales spend $2.2 million more, $2.9 million more, and $1.6 million less on advertising relative to other years, respectively. Interestingly, if we instead focus on share sales by lower-tier executives (e.g., the chief financial officer, chief investment officer, and chief technology officer) who have little control over firm investment decisions, we do not observe any significant variation in advertising spending around their stock transactions.

I conduct similar analyses of advertising spending around firm equity issues and stock-financed acquisitions. There also appears an inverted V-shaped pattern in advertising around these corporate events. For example, advertising spending in years of equity issuance and stock-finance acquisitions is 7.7% and 14% higher than that in other years, respectively; in other words, firms spend an additional $3.2 million and $5.9 million on advertising in these years. Not surprisingly, there is no similar pattern in advertising spending around firm debt issuance or cash-financed acquisitions, as equity valuation has minimal impact on these corporate events.

Asaf Manela

The value of diffusing information

Journal of Financial Economics | Volume 111, Issue 1 (Jan 2014), 181–199
(please cite only the original publication, not FAME)

How does the speed by which information diffuses affect its value to a stock market investor? Consider two drugs: Viagra treats erectile dysfunction while Allegra fights nasal allergies. Both drugs, if approved by the US Food and Drug Administration (FDA), would generate similar revenues to their publicly traded developers. When Viagra is approved, news of the approval diffuses fast because it makes for good conversation, while news of Allegra's approval travels slower. As an investor, would you pay more today to know that Viagra or Allegra will be approved tomorrow?

Faster-diffusing information can be more or less valuable

The intuitive answer to this question is that faster-diffusing information is more valuable because it translates into a quicker and less random realization of profits. But this line of reasoning ignores the fact that traders' decision to collect private information would take into account its future diffusion rate. For example, even before the two drugs are approved, one might expect news of Viagra's approval to diffuse faster than Allegra's. The model I develop accounts for this realistic feature of the world, and shows that competing informed investors trade more aggressively on faster-diffusing information. The increased private information collection impounds more information into prices, thus decreasing the returns to informational trading.

In my model, the transmission rate of information determines equilibrium asset prices and volume. The model is a four-period noisy rational expectations model, ala Grossman-Stiglitz (AER 1980), with asymmetric information about a risky asset that pays off in the post-announcement period. Information diffuses through a large population of risk-averse investors. The transmission rate controls the probability that uninformed investors in the pre-announcement period become informed for free in the announcement period. Pre-announcement, investors make their endogenous information choice, taking into account its future diffusion through both direct communication and indirect learning from prices.

The main theoretical result provides a closed-form expression for the value of information given the transmission rate that has eluded previous studies of similar settings (e.g. Hirshleifer, Subrahmanyam, and Titman, JF 1994). It is the sum of three terms. The first, and empirically dominant, term is positive and increasing in the intertemporal decline in uninformed investors' uncertainty. The transmission rate has two contrasting effects on this term. The first and more intuitive effect is that informational gains realized earlier are better than those realized only in the future that are subject to additional randomness. However, a second effect of this potential gain is that informed investors trade more aggressively, which makes pre-announcement prices more informative. The resulting reduction in pre-announcement uncertainty reduces the intertemporal decline in uncertainty and lowers the equilibrium value of information. This second effect happens holding the early informed fraction fixed. Unlike in

Grossman-Stiglitz where strategic substitutability in information choice results from endogenous price informativeness, here price informativeness affects the value of information through the aggressiveness of informed investors.

The second term of the value of information is positive as well and has to do with the intertemporal decline in uncertainty of the informed relative to this decline by the uninformed. The third term is negative and represents the extent of information spillover to uninformed investors who do not pay for the signal. The relative contribution of each of these terms to the total value depends on the parameters of the model, mainly the prior variances of the noisy supply and of the terminal payoff, and remains an empirical question.

Measuring the diffusion rate of information

In order to estimate the model, we need a setting where we can measure the diffusion rate of market moving information. I estimate the parameters of the model and the magnitudes of the three terms of the value of information in a panel of FDA drug approvals, using media coverage as a proxy for the approval-specific transmission rate of information.

Specifically, I calculate *media exposure* for each approval as the sum of all news articles that report the approval on the official approval day and the next, weighted by the price of their adjacent advertising space. The idea is fairly intuitive. The more interesting the news, the more prominently news outlets feature its coverage, on pages where the price of advertising is high due to its high exposure to readers. One benefit of this measure over an indicator media coverage variable is that it includes moderately interesting news as well, which turns out to be most valuable to acquire.

Drug approvals provide a particularly clean laboratory for examining stock market reaction to news that varies in its prominence for several reasons, but most importantly, the unique drug names and active ingredients allow for a free-text article search that is likely to produce only articles that discuss the approval story. This is not the case for other well-studied events, such as earnings announcements, which are harder to classify and often include additional profitability-relevant soft information.

Figure 1: Stock returns around drug approvals: media exposure subsamples

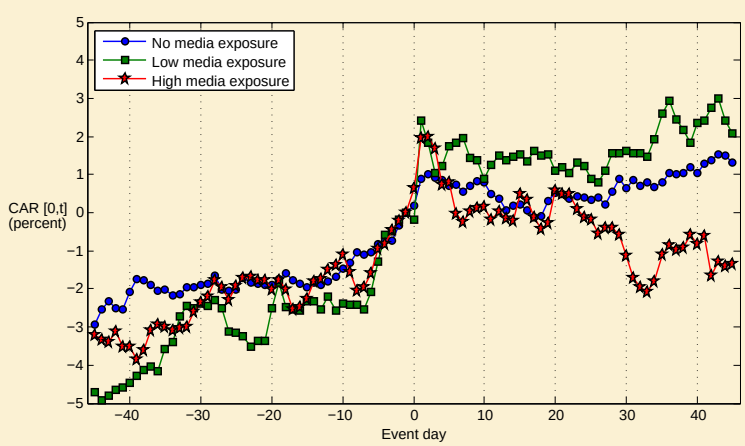

The sample includes 320 Original New Drug Approvals from 1990 to 2007. *Media Exposure* is the sum of all articles on the approval day and the following day. The high media exposure sub-sample are the top half of drug approvals with positive media exposure (56 observations). The low exposure contains the bottom 57 observations. No media exposure are the remaining 207 observations. Variation in media exposure is related to the path of price adjustment to news.

In Figure 1, I split the sample into high, low, and zero media exposure sub-samples. All three subsamples feature a price increase in the days before the approval. The pre-approval return seems higher for drugs that would later appear in the news. This suggests insider-trading activity is increasing in future media exposure. At approval time, drugs covered by the media exhibit a higher price increase than the rest. Post-approval, the stock price of drug sponsors that received no initial media exposure continue to appreciate while low media exposure firms maintain their valuation. Interestingly, high media exposure approvals exhibit a negative drift following the approval, which continues even at a longer horizon than the one I test below. These results suggest that consistent with the model, variation in media exposure is related to the path of price adjustment to news.

The most valuable information diffuses at a moderate speed

Empirically, the value of drug approval information turns out to be hump-shaped in its diffusion speed. The most valuable information diffuses at a moderate speed. The answer to the question of how much to pay to know today that Viagra as opposed to Allegra will be approved tomorrow turns out to depend on how fast exactly news of each approval diffuses, through its effect on the intertemporal decline in uncertainty.

As mentioned above, the value of information has three terms. It turns out that the value of information stems entirely from the first term, which captures the intertemporal decline in uninformed agents' uncertainty. Moreover, the transmission rate of information has a quantitatively important effect on the value of information.

The structural estimation exercise in the paper, one of its key contributions to the literature, embarks on the ambitious goal of fitting the model to the quantitative features of the data. The results show that uninteresting news that propagates slowly is not pursued by anyone before the official announcement, because the fixed cost of information is prohibitively high. Faster-diffusing information is purchased at a higher rate, while the fastest diffusing news is somewhat less valuable.

The structural estimation approach taken here of focusing on a particular anticipated news release, namely drug approvals, using media coverage as a proxy for the transmission rate of information is novel. It can hopefully be used by future research to further improve our understanding of the diffusion of information in financial markets.

Blake Phillips, Kuntara Pukthuanthong, and P. Raghavendra Rau

Past performance may be an illusion: Performance, flows, and fees in mutual funds

Critical Finance Review | Volume 6, Issue 2 (2016), tbd
(please cite only the original publication, not FAME)

One of the most persistent and robust patterns documented in the mutual fund literature is return-chasing by investors. The finding that investors allocate wealth disproportionally to funds with prior superior performance transcends mutual fund asset classes, country boundaries, and investment objectives. In this paper, we examine a different type of return-chasing behavior, whether investors chase stale fund returns, returns arising from horizon effects in holding period returns (HPR). Specifically, when advertising past performance, Rule 482 of the Securities Act of 1933 requires investment companies to report past performance in the form of an HPR over horizons of 1, 5, and 10 years for funds in existence over those horizons. In addition, the 3 year horizon is also commonly reported but is not required.

The focus of our analysis is how investors interpret the change in reported HPRs. For example, consider the following 5 quarter return time series:

Period	-1	-2	-3	-4	-5
Return	–2%	3%	4%	5%	–4%

The annual HPR for quarters -2 to -5 is 8% and the corresponding annual HPR for quarters -1 to 4 is 10%:

$$\text{HPR}_{t-1} = [(1 + -0.02) \cdot (1 + 0.03) \cdot (1 + 0.04) \cdot (1 + 0.05)] - 1 = 0.102$$

$$\text{HPR}_{t-2} = [(1 + 0.03) \cdot (1 + 0.04) \cdot (1 + 0.05) \cdot (1 + -0.04)] - 1 = 0.079$$

Even though the fund experienced a negative return in the most recent period ($t = -1$), the HPR increased as the end-return which dropped from the sample was more negative. The change in the HPR is therefore a function of the most recent return (–2%) which enters the horizon and the end-return (–4%) which drops from the horizon. As all other intervening returns are common in the return sequences, they have no influence on the change in the HPR. Reacting to the new information conveyed in the most recent return is arguably rational to the extent it is related manager skill and future fund performance. In contrast, end-returns convey no new information and correspondingly should not influence rational investor preferences.

We study whether mutual fund investors differentiate between current and end-return influences on changes in reported mutual fund performance. We call investor reaction to end-returns the "stale return chasing" effect. We then examine the determinants of this behavior and the mechanisms by which mutual fund managers propagate and benefit from investor reaction to stale returns.

Investor Reaction to Stale Performance: Identification Strategy

Our data comes from the Center for Research in Security Prices (CRSP) Mutual Fund Database for the period of 1991-2010. To measure the stale performance effect, we relate proxies for investor allocation preferences to the first return lag plus successive return lags ($n = 2$ to 73), in 72 separate regressions:

$$\Delta m_t = \alpha + \beta_1 R_{t-1} + \beta_{t-1} R_{t-1} + \epsilon_t \tag{1}$$

where n signifies the lag of the second return included as an independent variable in month t. The primary measure of investor preferences is the change in market share held by the fund, defined as:

$$\Delta m_{i,t} = \frac{n_{i,t}}{\hat{N}_t} - \frac{n_{i,t-1}}{\hat{N}_{t-1}} \tag{2}$$

where $n_{i,t}$ is total net assets under management for fund i in month t, and $\hat{N}_t$ is aggregate total net assets for all funds in the sample at time t. We also cluster standard errors by fund and date (month-year). To motivate our use of the first equation, we note that the change in HPR has only two influences, the magnitude of the return in the current period and the magnitude of the end-return which drops from the calculation. These two returns have an equivalent impact on the HPR, though only the former is new information. Modeling investor response to the change in HPR, it then follows that Δmt becomes a function of the new and end-return, linearly approximated by equation (1). All intervening returns are common between adjacent HPRs and have no influence on ΔHPR. If investors interpret the stale information component of the change in HPR as information regarding future fund performance, we expect the β_n coefficient related to return observations at the end of standard HPR reporting periods to be negative. This relation follows from the inverse relation between the magnitude of the end return which drops from the sample and the change in the HPR.

Do investors chase stale returns?

The βn coefficients from our pooled, OLS regression series in the first equation show that they do. In particular, the magnitudes and signs of the variables of interest (the coefficients on the 13th, 37th, and 61st monthly return lag which are the end-returns related to the reported 1, 3, and 5 year HPR) are strikingly different from all the other variables. They are larger than all but one of the other coefficients. The average absolute value of these three coefficients is 0.35, twice the size of the absolute mean coefficient (significant at the 5% level). Almost none of the other coefficients are significantly different from both zero and the absolute mean of the coefficient sample ($\alpha = 10\%$). Interestingly, the magnitudes of return-chasing on the stale return chasing lags are of equal or greater magnitude to return-chasing on the first return lag. Investor asset allocations appear equally or more sensitive to the stale information reflected in HPR end-returns relative to the new information reflected in the most recent return. These relations are robust to a range of alternative specifications including alternative investor preference proxies, alternative adjustment of standard errors and inclusion of controls for investor preferences from the literature.

Does fund advertising take advantage of stale return chasing behavior?

We use two proxies for fund marketing expenditures, the 12b – 1 expenditure of the fund and investment company advertising data compiled by Kantar Media. We then relate the change in annual fund-specific advertising expenditures to determinant variables. We focus on advertising that reports HPR related information, either the actual return or analyst rating by holding period. We show that funds spend more money advertising HPRs the year after the 1, 3, and 5 year HPRs are individually higher, consistent with funds seeking to draw attention to prior strong performance. We also partition the sample into quartiles by the 1-year HPR which dropped from the sample (i.e. the HPR^{1Year}_{T-4} or HPR^{1Year}_{T-6} for the 3 and 5 year HPR, respectively). If funds are preferentially advertising stale performance, we expect greater spending on HPR advertising when the return which dropped from the sample is small (more negative), resulting in an increase in the reported HPR. That is precisely what we find. Funds indeed spend a greater amount on HPR related advertising when reported HPRs are higher due a low return dropping from the horizon of assessment.

What does stale performance chasing depend on?

To explore the determinants of stale performance chasing at the fund-level, we next estimate the first equation by fund and year, producing an annual time series of β_n coefficients for each fund. We then relate the fund-level β_{13}, β_{37}, and β_{61} coefficient time series to determinant variables using a pooled, cross-sectional panel. We find that stale performance chasing tends to be lower for larger (greater TNA) and older funds and funds from larger families (greater family TNA). Funds from families which offer a larger number of funds also tend to have less return-chasing, though the relation is not significant at conventional levels. Stale performance chasing tends to be higher for funds with greater value uncertainty as proxied by the standard deviation of mutual fund returns over the prior year. Advertising broadly increases stale performance chasing, both at the fund and family-level. Ads that report HPR info induce the largest investor response to stale performance signals, followed by advertising that features analyst ratings. The magnitude of the response of investors to ads that feature HPR info is typically twice that of ads that feature analyst rating with no reference to HPR info. Finally we examine how investor sophistication affects the stale return chasing behavior. Sophisticated investors should understand the mechanical influence of time on HPRs and recognize that HPR signals provide no new information in excess of the most recent return observation. As a proxy for investor sophistication, we contrast retail with institutional funds. Institutional funds are marketed to high net worth individuals and large pension funds managed by professionals who are often viewed as being more sophisticated than retail investors. Across all three stale return horizons we find that, if anything, institutional investors are more prone to stale return chasing than retail investors.

Do mutual funds change their fees to take advantage of stale performance chasing?

Our final analysis examines the implications of stale performance chasing for fees. Mutual fund fees are set in a competitive environment. Funds that wish to increase fees are likely to time the increase to special situations. For example, Bris-Gulen-Kadiyala-Rau (RFS 2007) show that mutual funds choose to raise fees at points in time when they close their funds to new investors. Thus, stale performance

chasing may provide mutual funds with the opportunity to increase fees, in essence cashing in on the mechanical effects of HPRs.

To test this hypothesis, we relate fees to our return-chasing estimates. We find that funds with higher stale return-chasing in the prior period typically charge higher fees as a percentage of TNA. Although contractual fees are typically time invariant and may only be changed with shareholder consent, mutual funds routinely voluntarily waive fees during periods of poor performance in order to retain performance-sensitive investors (Christoffersen [JF 2001]). We obtain fee waiver data from Morningstar, which reports the percentage of TNA waived from fees in a given year. We then relate the change in waived fees to the stale performance chasing estimates and control variables. The stale performance chasing coefficients are all positive and significant, suggesting that funds reduce waived fees on the heels of stale performance chasing flow. The absolute sizes of the coefficient estimates across the model sets are highly similar, suggesting that adjustment to waivers explains the majority of the change in fees associated with stale performance chasing. Thus, our results suggest a channel by which managers are able to realize opportunities to return fees to prior levels without having to generate exceptional returns, i.e. by waiting for mechanical increases in HPRs realized as time passes.

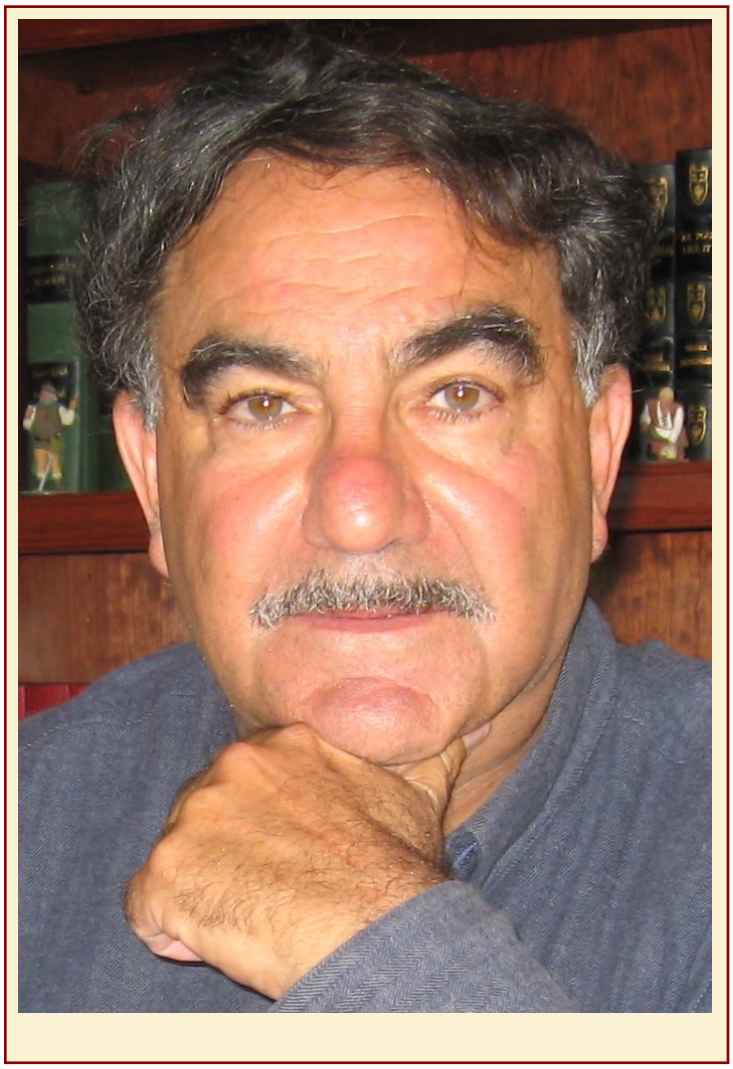

Bo Becker and Victoria Ivashina

Reaching for yield in the bond market

Journal of Finance | Volume 70, Issue 5 (Oct 2015), 1863–1902
(please cite only the original publication, not FAME)

A key principle of finance is that evaluations and comparisons of returns are only meaningful after adjusting for risk. However, risk is often hard to measure. This creates an important limitation in the delegation of investment decisions. Financial intermediaries that are evaluated based on imperfect risk metrics face an incentive to buy assets that comply with a set benchmark but have "hidden risk." In other words, imperfect benchmarks may create incentives to "reach for yield" in the context of fixed-income investing, or to "search for alpha" more generally. This could lead to excess risk-taking in financial institutions, persistent distortion of investments and, potentially, amplification of the overall risk in the economy. Indeed, reaching-for-yield is believed to be one of the core factors contributing to the buildup of credit that preceded the recent financial crisis (Yellen [2011] and Rajan [2010]).

We study "reaching for yield" in the corporate bond market. We show that insurance companies—the largest investor in bonds—reach for yield in choosing their investments.

Insurance companies' capital requirements allow reaching for yield within ratings buckets

Insurance companies in the U.S. face capital requirements that depend on the credit ratings of fixed income securities in their portfolios. Bonds are grouped in broad buckets based upon credit ratings; bonds in these groupings face similar capital requirements. For the best bonds, rated AAA down to A-, the requirement is 0.30% of face value (an insurance company needs 30 cents of equity for every $100 of book value of bonds in this range of ratings). For bonds rated BBB, the capital requirement is 0.96% of book value; for BB, 3.39%; for B 7.38%; for CCC 16.96%; for with a rating below CCC, the capital requirement is 19.50

Given these rules, insurers could reach for yield by selectively buying the riskiest bonds within a given category, thus increasing the risk of the portfolio without raising capital requirements. We demonstrate that this is indeed what happens. We first show that in thirteen quarters leading to the financial crisis (2004:Q2 to 2007:Q2), insurers—as compared to pension funds and mutual funds—exhibit a strong preference for safer bonds, representing over 70% of institutional purchases of investment grade bonds. (See Figure 1.)

That is, across risk categories, insurers exhibit the risk aversion that capital requirements are meant to induce. Yet, within capital requirement buckets, the risk preferences are reversed. The share of newly-issued corporate bonds acquired by insurers in the same time frame increases in risk (as measured by yield to maturity at issue, or the CDS spread) within the AAA to A category (the same is true for BBB, but we lack a precise estimate for high-yield bonds, of which there are far fewer). (See Figure 2.) Thus, insurers appear to reach for yield in a way that is invisible to the standard metric on

which they are evaluated. It is also characteristic of firms with poor corporate governance, and firms for which the regulatory capital requirement is more binding.

Figure 1: Insurance companies' holdings of newly issued corporate bonds

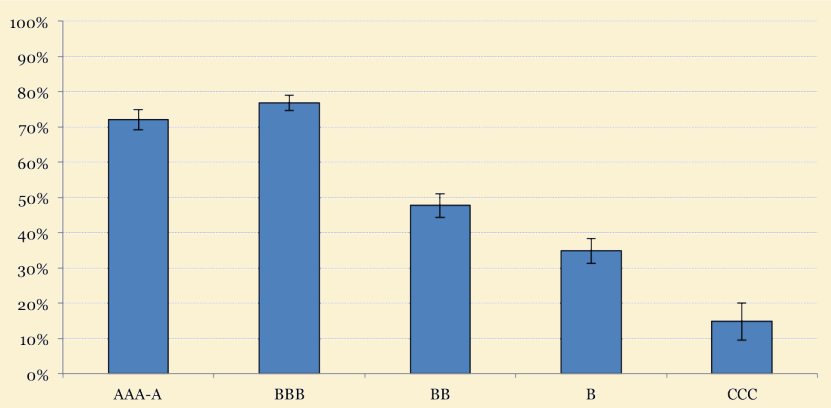

This shows corporate bond holdings 2004:Q2-2007:Q2, by rating category, for newly issued bonds. The horizontal axis shows holdings of insurance companies, relative to the sum of insurance companies, pension and mutual funds' holdings. The bars indicate 95% confidence intervals. Insurers buy very large shares of newly issued safe bonds (rated BBB and above) but avoid riskier (lower rated) bonds.

This result does not depend on comparing insurance firms to other investors. In the paper, we show that a similar pattern emerges when we look only at insurance firms' investment in a given bond over time, as its yield and CDS changes but its credit rating does not. The pattern is robust to inclusion of duration and liquidity controls. Importantly, we also show that this risk-taking is not "alpha," i.e., superior investment ability. With various choices of benchmarks, we document that the abnormal return on the aggregate corporate bond portfolio of insurers is negative or zero.

Business Cycle and Crisis

Consistent with concerns that reaching for yield contributed to buildup of risk in the financial system and its consequent contraction, we show that insurers' reaching behavior is pro-cyclical. In particular, reaching-for-yield behavior disappears during the most recent financial crisis and comes back in the second half of 2009, as documented in Figure 2.

The Impact of Reaching for Yield

The impact of reaching for yield is to increase risk taking. We find that the firms with the most reaching tended to do worse in the financial crisis of 2008. For example, AIG, which was noted for its aggressive managing of regulatory capital, and which—as is well known—was at the center of the financial crisis, had one of the higher portfolio yields for corporate bonds among large insurance companies. In our regression analysis, we document that during the financial crisis, the drop in stock prices for publicly traded insurance companies was deeper for those with less owner oversight (as measured by the number of block holders). This suggests that the risk-taking induced by reaching for yield may be undesirable for owners (which likely implies that it is not good for society at large either), and may contribute to cyclicality in credit markets.

On the other hand, reaching for yield has broad implications for the credit supply. Firms that happen to belong to the favored "buckets" (high-risk firms with "A" ratings on senior debt, for example) would be able to borrow at better terms relative to other firms. Indeed, we document that bond issuance by riskier firms coincides with times of pronounced reaching for yield by insurance firms.

Figure 2: Reaching for yield over time

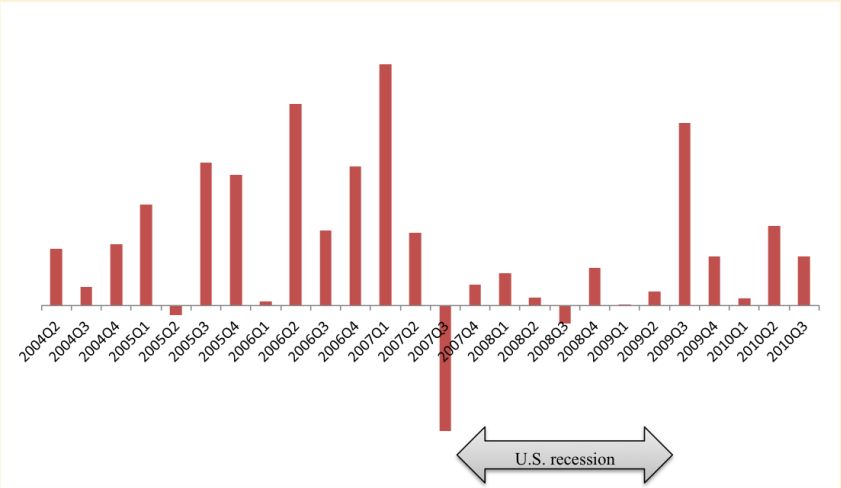

This figure shows the reaching for yield by US insurance companies (the strength of insurance companies' preference for higher-yielding investment grade bonds within ratings categories) by quarter. Reaching for yield peaked in 2006 and early 2007, disappeared in the crisis, and then returned as financial markets recovered in 2009.

Final Thoughts

Just like regulation based on ratings, delegated management—which is based on credit ratings—may incite reaching for yield within rating buckets. Portfolio managers can raise the promised yield on a bond portfolio by taking more credit risk. As long as securities do not default, this additional yield improves performance relative to benchmarks, and reported profits are raised. To the extent that reported profits are more visible than the credit risk in bond portfolios, owners may misunderstand actual performance of managers that reach for yield.

While we study insurance companies and risk assessment based on credit ratings, in all likelihood, the reaching for yield we document for insurance firms and corporate bonds happens in other parts of the financial system. All in all, reaching for yield is an ongoing phenomenon hardwired into financial intermediation by imperfect risk measurement. Illiquid and complex securities are particularly prone to this issue. Weak incentives and regulation that reduces visibility or discourages assessment of risk by outside claimholders exacerbate this problem.

Kristian Rydqvist, Joshua Spizman, and Ilya Strebulaev

Government policy and ownership of equity securities

Journal of Financial Economics | Volume 111, Issue 1 (Jan 2014), 70–85
(please cite only the original publication, not FAME)

Since the end of World War II, direct stock (also referred to as equity) ownership has declined. In the United States, households held 90% of equity directly just following World War II. By 2010, the amount of direct equity held by households declined to under 30%. Equity ownership migrated to financial institutions, who own almost 50% of U.S. stocks in 2010. Where historical data is available, this pattern is similar worldwide. Between 2005 and 2010, the average direct household ownership of equities is approximately 17%. We show that tax policy was a major factor in causing this worldwide shift in equity ownership.

The Tax Benefits of Pension Funds

Governments around the world allow retirement savings to be favored by the tax system. For example, the United States defined benefit pension funds and 401(k) accounts both provide tax incentives to contribute to them. For brevity, we refer to all of these plans as pension plans. Investment returns from assets held inside a pension plan accrue tax free. When holding stocks directly (outside of a pension plan), both dividends and capital gains are subject to personal income tax. This difference is the first tax benefit of contributing to a pension plan. A second tax benefit of pension plans is that employers and employees contribute pre-tax income. Therefore, households can reduce their marginal tax rates while working by shifting taxable income into lower-income retirement years.

The Evolution of Stock Ownership

Figure 1f show the evolution of stock ownership for the eight countries (Canada, Finland, France, Germany, Japan, Sweden, United Kingdom, and United State) in our sample from 1945-2010. The solid diamonds and lines represent direct household ownership of stocks and the open diamonds and dashed line represent tax-preferred institutional holdings (pension funds, investment funds, and insurance companies). There are several important observations. First, the decline in direct household ownership was large across all countries, with the average decline being 40%. The second observation is that household direct ownership of stock was largely replaced by financial institutions, which increased by 24%. Finally, there was significant variation both across time and across countries in the shifts in stock ownership. The decline was larger for most countries in the 1970s and 1980s and varied across countries in all time periods.

Figure 1: The evolution of stock ownership: US

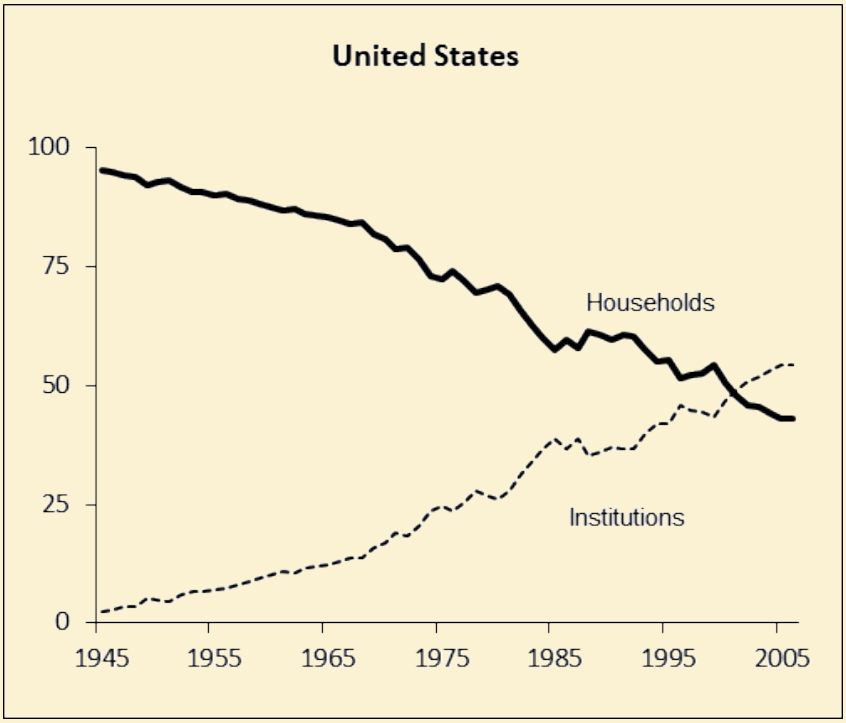

This and subsequent figures show how stock ownership has evolved over time. The sample period is from 1945 to 2010. Solid diamonds and lines represent direct household ownership of stocks while open diamonds and the dashed line represent tax-preferred institutional holdings (defined as pension funds, investment funds, and insurance companies). **Interpretation:** Households have shifted from holding stock directly to indirectly in tax-preferred accounts across the world.

Figure 2: The evolution of stock ownership: Canada

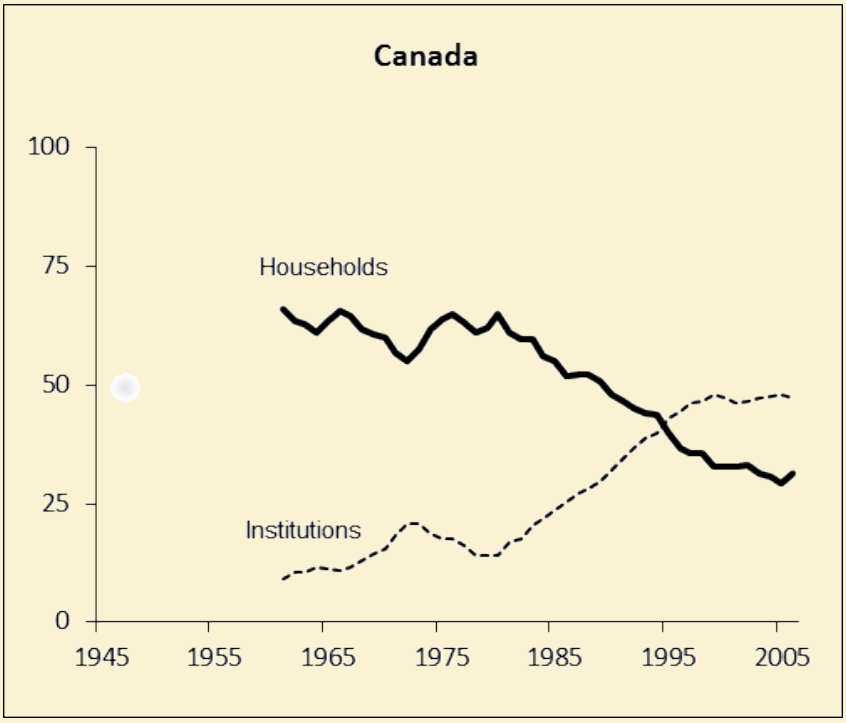

Figure 3: The evolution of stock ownership: Japan

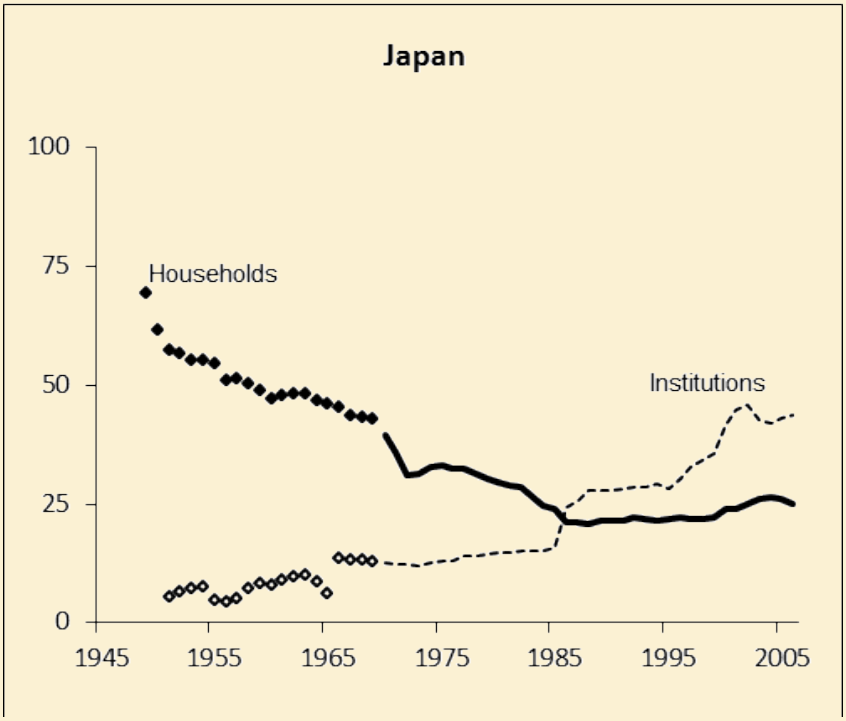

Figure 4: The evolution of stock ownership: Germany

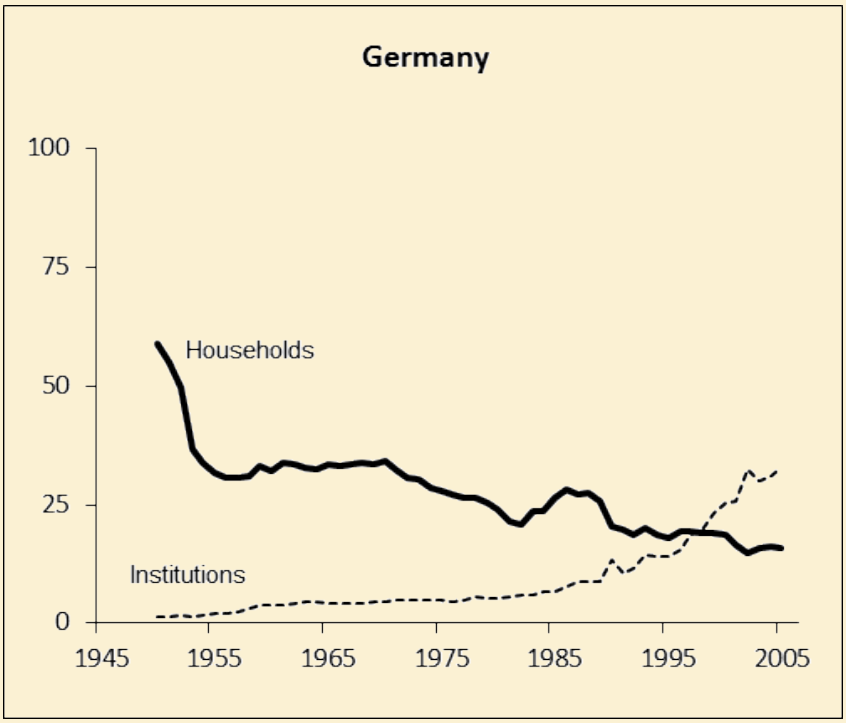

Figure 5: The evolution of stock ownership: UK

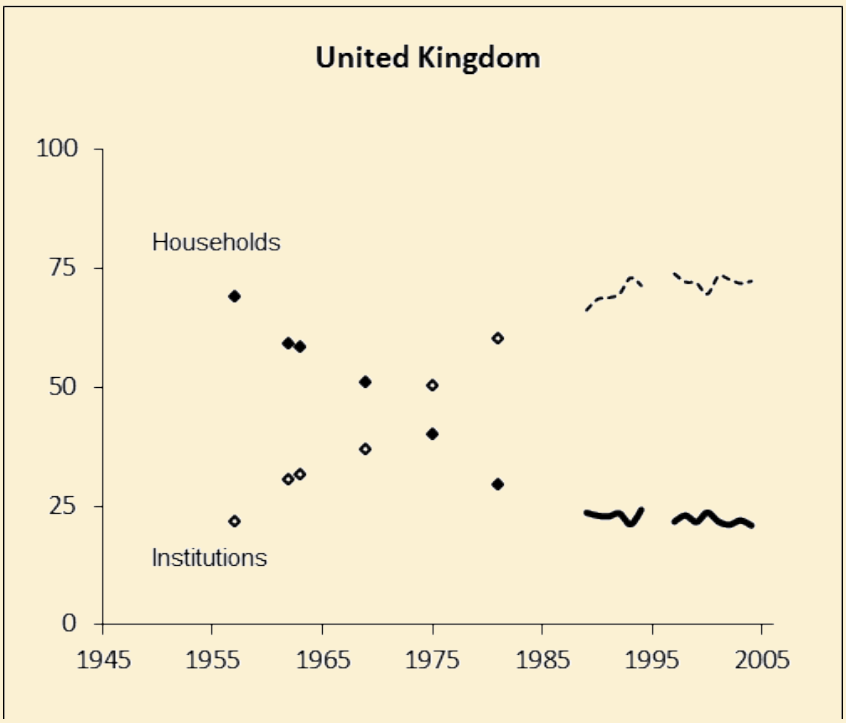

Figure 6: The evolution of stock ownership: Sweden

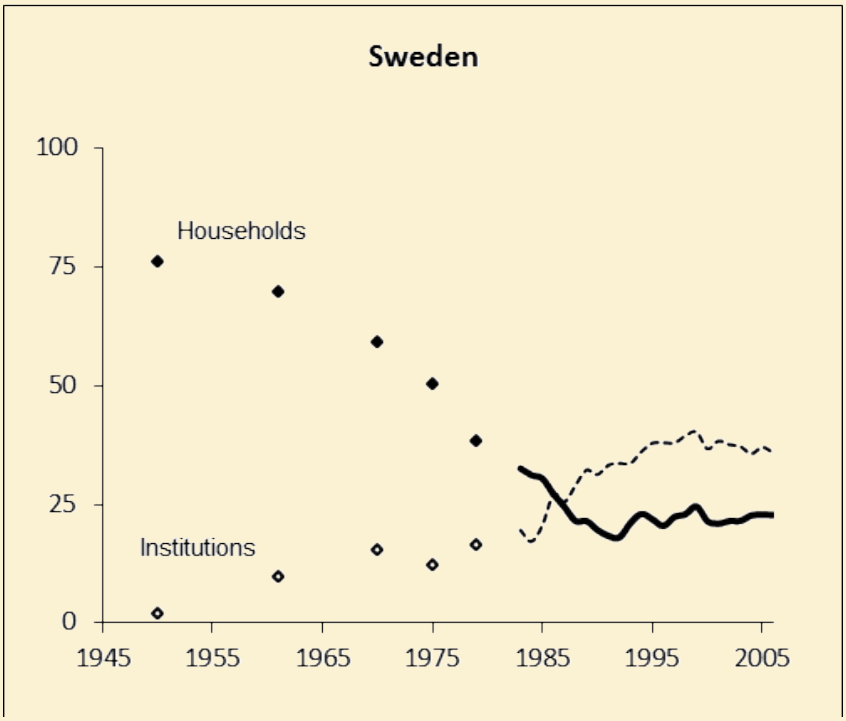

Figure 7: The evolution of stock ownership: France

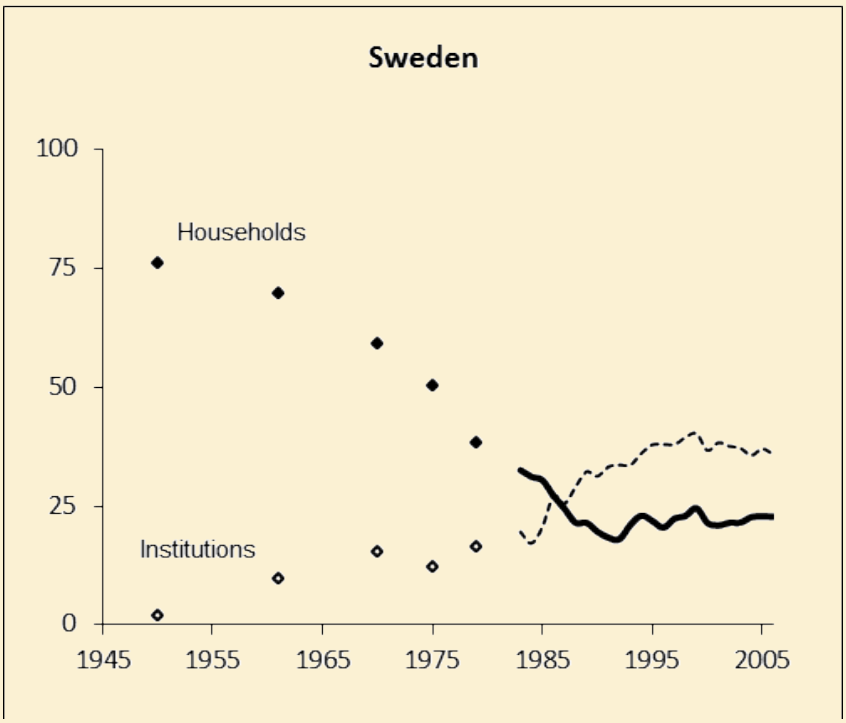

Figure 8: The evolution of stock ownership: Finland

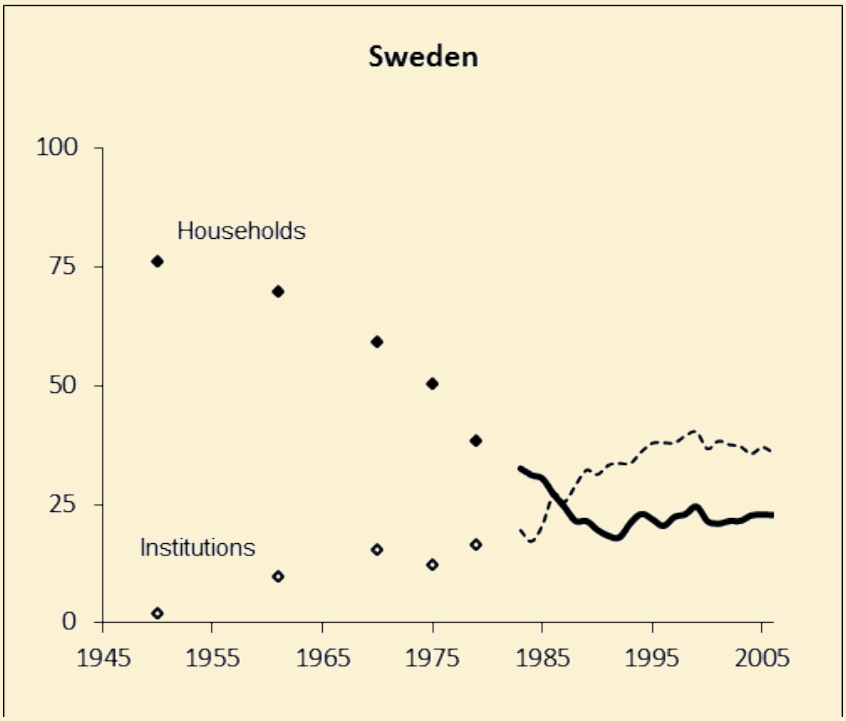

The Tax Explanation: The United States

We argue that tax policy is a major factor in explaining the shifting ownership structure. The development of the United States pension and mutual fund industries provides an example of how government tax policy helped shape the structure of institutional ownership. Figures 9f report the evolution of the U.S. pension fund and mutual fund industries. They show pension fund and mutual fund holdings of stocks and bonds. The figures show that since the 1940s, pension funds ownership of both stocks and bonds has grown; and by the 1980s, pension funds owned a large fraction of both stocks and bonds. Interestingly, mutual funds owned an insignificant portion of these assets during this time period. However, mutual funds began to grow in the 1980s, replacing pension funds as the dominant form of equity savings. This phenomenon can be explained by government tax policy.

Figure 9: Pension fund stock holdings in the US

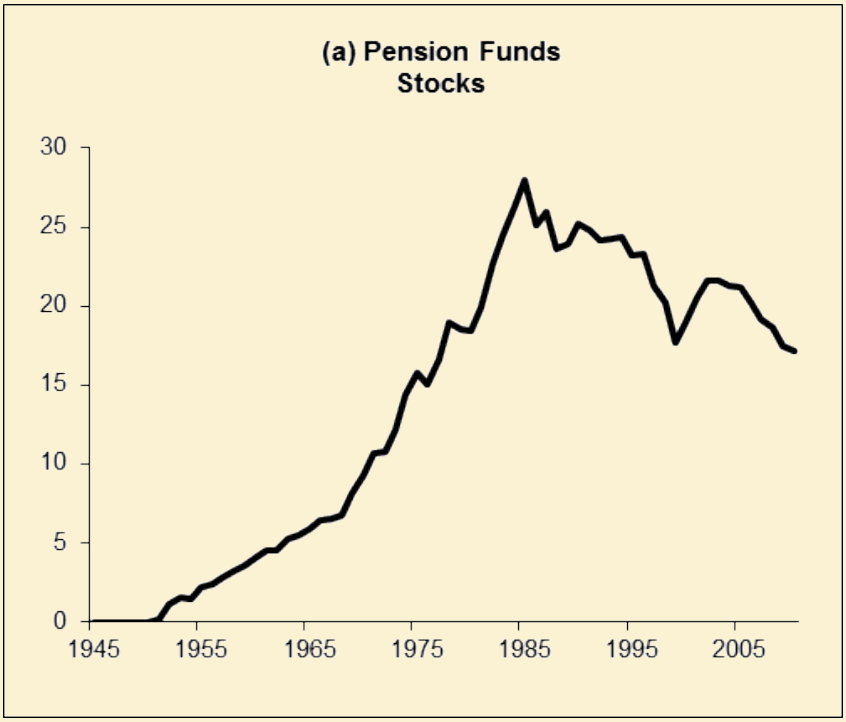

This and subsequent figures show how institutional holdings of stocks and bonds in pension and mutual funds has evolved over time in the United States. The sample period is from 1945 to 2010. **Interpretation:** Institutional holdings of stocks and bonds in the United States grew depending on tax incentives.

Figure 10: Pension fund bond holdings in the US

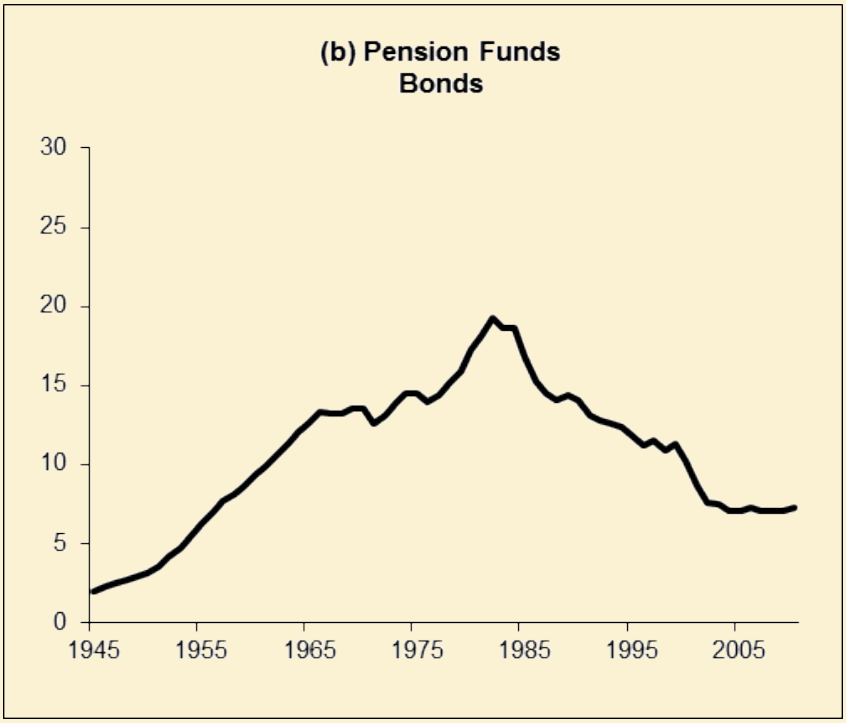

Figure 11: Mutual fund stock holdings in the US

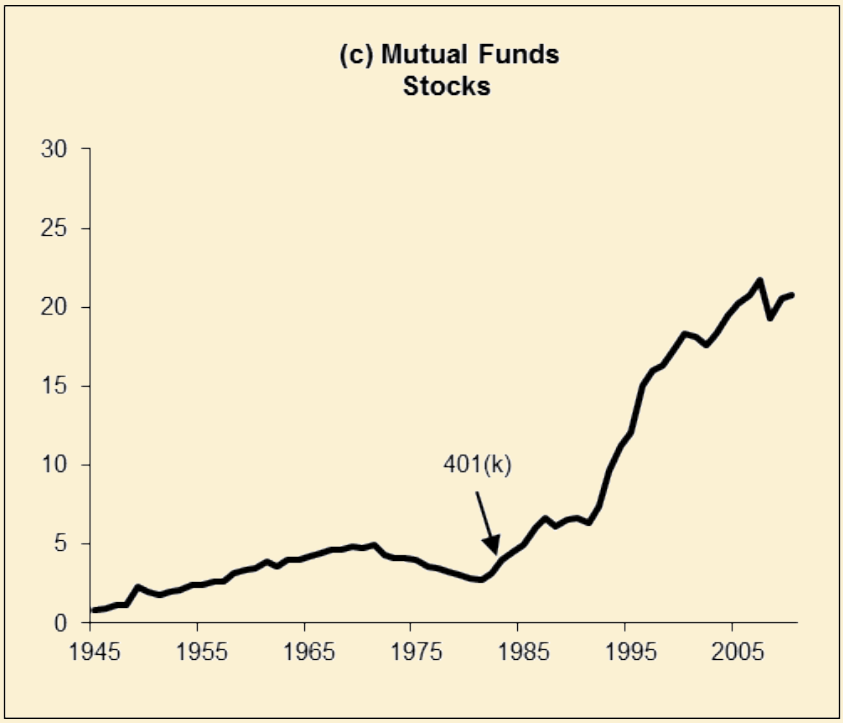

Figure 12: Mutual fund bond holdings in the US

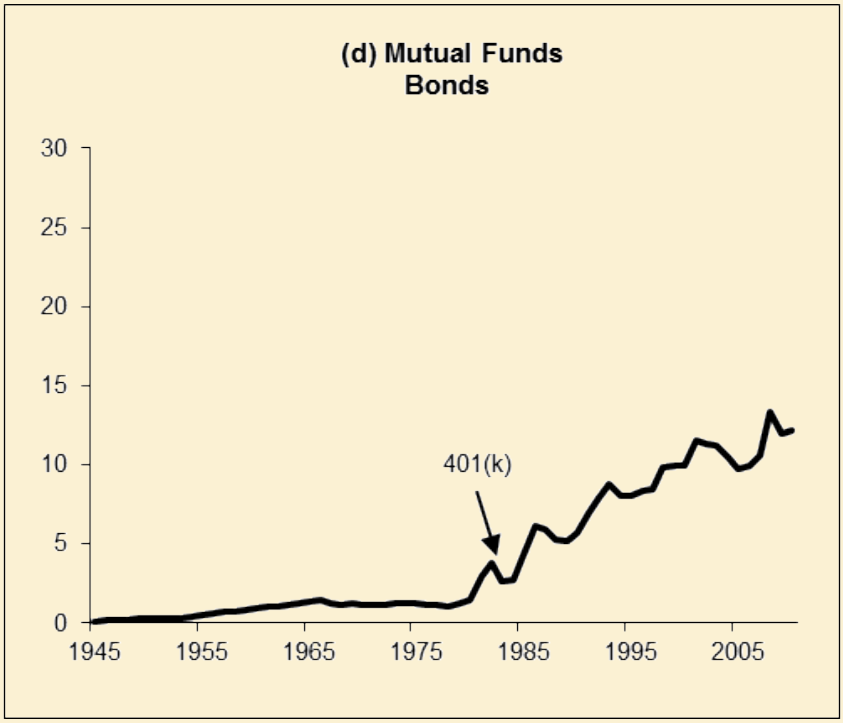

During the 1940-1980 time periods, pension funds have a tax advantage that mutual funds do not have, due to the Revenue Act of 1921. In 1978, the new 401(k) legislation allows for adoption of defined-contribution plans. However, it is not until the Economic Recovery Tax Act of 1981 where defined contribution plans became operational, by specifying contribution limits. The combination of the 1978 and 1981 reforms create a new demand for defined contribution plans that are managed by mutual funds. Dickson, Shoven, and Sialm (NTJ 2000) show that roughly half of equity mutual funds are held inside a tax-deferred account. Obviously, taxes are not the only reason for the growth of the mutual fund industry and, while this evidence is telling, we must provide systematic evidence relating the ownership of financial asset to government tax policy, using our panel of eight countries.

The Tax Explanation: Worldwide Evidence

We construct two proxies to capture the tax incentives from holding stock inside a pension plan for all eight countries in our study. GAP measures the difference in investment returns between holding stocks inside a pension plan versus outside a pension plan. SMOOTH measures the benefit of shifting taxable income from higher income work years into lower-income retirement years. Meaningful measures of both GAP and SMOOTH rely on understanding details of how dividend, capital gains, and ordinary income are taxed. We use a married-filing jointly income of five times GDP per capita to construct all tax related variables.

Table 1: Effect of tax incentives on direct household stock ownership

Direct Household Stock Ownership		
Constant	–0.73**	–0.19
GAP5		–32.3**
SMOOTH5		0.5
N	396	396
R^2	0.000	0.027

This table reports the results of regressing the annual change in households' annual, direct stock ownership on GAP, which measures the rate of return difference between saving inside and outside a pension plan and SMOOTH, which measures the tax benefit of income smoothing. The sample period is from 1945 to 2010. **Interpretation:** A three percentage point difference between saving inside and outside the pension plan associates with an annual decline in the fraction of household direct equity ownership of about one percentage point.

We test whether or not tax incentives (GAP and SMOOTH) affect the change in direct household stock ownership. Table 1 reports our regression results. Specification (1) reports only the average annual change in household ownership across the eight countries (0.73%). When we add our tax incentive variables, the coefficient on GAP is significantly different from zero, while the coefficient on SMOOTH is not. Most importantly, the intercept term is no longer statistically significant. From an economic standpoint, our results mean a three percentage point difference between saving inside and outside a pension plan results in an annual decline in the fraction of household direct equity ownership by about one percentage point.

Table 2: Effect of tax incentives on direct household stock ownership

	1950-1959	1960-1969	1970-1979	1980-1989	1990-1999	2000-2010
Constant	1.03	0.59	0.21	0.46	–0.84	2.65**
GAP5	–19.4	–39.4**	–22.1	–41.9**	–31.7	–147.6
SMOOTH5	–94.9	–31.8	–21.4	–13.2	28.2	–35.3
N	37	68	72	80	80	53
R^2	0.098	0.132	0.083	0.101	0.002	0.050

This table reports the results of regressing the annual change in households' annual, direct stock ownership on proxy variables for the tax benefits of saving inside a pension plan on a decade-by-decade basis. The coefficient on GAP is generally negative. Our statistical results are driven by the cross-sectional variation in tax policy across our eight countries.

Table 2 explores the time-series impact by estimating our model on a decade by decade basis. The coefficient on GAP is generally negative and of similar economic magnitude to the full panel estimates of Table 1. Statistical significance varies because the number of observations in each decade is small.

In conclusion, we show that government tax policy has an important impact on the evolution of stock ownership. Over time, households moved from holding stock directly to holding it indirectly inside a pension fund or mutual fund.

Berk A. Sensoy, Yingdi Wang, and Michael S. Weisbach

Limited Partner Performance and the Maturing of the Private Equity Industry

Journal of Financial Economics | Volume 112, Issue 3 (Jun 2014), 320–343
(please cite only the original publication, not FAME)

A fundamental question in delegated asset management is whether and why some investors, or classes of investors, have systematically different performance. One place where differences in performance between investors are potentially present is the private equity industry. In 1990, private equity was a little-known alternative asset class, with $6.7 billion in funds raised. By 2008, just prior to the financial crisis, the industry raised almost 40 times this amount, $261.9 billion. Over this relatively short period, the industry transformed from a niche investment to a standardized asset class that is important component of institutional investors' portfolios.

In this paper, we construct a large sample of LP investments in private equity funds, and investigate how different types of LPs performed over time. Our analysis focuses on the way relative performance has changed as private equity matured as an asset class. In addition, we evaluate the importance of access to the best funds, as opposed to investment selection skill, in explaining differences in performance. Did the changes in the private equity industry in recent years lead to corresponding changes in the fundamental relationships between limited partners (LPs), the investors in private equity funds, and general partners (GPs), the ones who manage private equity funds? How did these changes affect the relative performance of different types of investors? More generally, how does the relative performance of different types of investors depend on whether the investments in question are new innovations or established asset classes?

Industry changes and limited partner performance

Our sample consists of 14,380 investments by 1,852 LPs in 1,250 buyout and venture capital funds raised between 1991 and 2006. These LPs include endowments, pension funds, banks, insurance companies, etc. Consistent with the trends observed in prior work (see Robinson and Sensoy (RFS 2011), Harris et al. (WP 2014)), since the late 1990s there has been an industry-wide decline in both absolute and relative returns to venture capital funds. In addition, the cross-sectional variation in returns and the relations between GP experience and return have decreased.

During the 1990s, endowments investing in private equity funds did extremely well, and endowments outperformed other types of institutional private equity investors. On their investments in private equity funds raised between 1991 and 1998, endowments enjoyed an average 35.7% internal rate of return and 2.43 implied PME, the highest of any LP type. This finding replicates that of Lerner et al. (HBS 2007), who find that endowments outperformed other institutional investor types in their private equity investments.

However, in funds raised in the more recent 1999-2006 period, endowments did not outperform other types of investors. In this later period, there are no statistically significant differences in returns across different LP types. However, we do find using the implied PME that all investor types had private equity returns exceeding that of the S&P 500 in both periods. Broadly speaking, there are two possible reasons why endowments could have outperformed during the earlier period. First, endowments had could have had superior investment skill. Such skill could potentially lead them to be better at evaluating alternative investments, such as private equity, that were unfamiliar at the time to other investors. Second, endowments could have had superior access to the best-performing fund families. Historically, rather than expand or raise fees to market-clearing levels, the best private equity partnerships have rationed access to their funds, accepting investments from favored investors, such as prestigious educational and other nonprofit endowments.

We present evidence suggesting that the abnormal performance of endowments in the 1990s occurred largely because of their access to the best venture capital funds. The performance gap between endowments and other LP types in this period is driven entirely by endowments' investments in the venture capital industry, which benefited most from the 1990s technology boom. To the extent that endowments have special investment skill, one would expect it to affect both venture capital and buyout returns. However, endowment performance in buyout funds during the 1991-1998 period was no different from that of other types of investors.

A direct test of investor skill that is unaffected by differential access involves comparing the quality of reinvestment decisions, since all LPs are usually given the option of investing the subsequent funds of the partnerships in which they invest. Even within the venture capital space, LPs' reinvestment decisions do not suggest that endowments or any other any other type of investor has superior skill during any subperiod. Since endowments do not make consistently better reinvestment decisions compared to other LP types, their superior performance during the 1990s is likely explained by their access to the best venture capital funds.

We also show that endowments are no more likely to be able to pick out the best-performing first-time funds than any other type of investor, either before the technology crash or afterward. First-time funds do not tend to restrict access to LPs, as they have yet to establish a track record. Therefore, they represent a pure test of investors' selection skills. Compared with other types of institutional investors, endowments were more likely to invest in older partnerships, which not only were more likely to restrict access but also earned higher returns.

Finally, in the 1991-1998 period, endowments were more likely to invest in venture capital funds whose increase in size from the firm's prior fund was lower than would be expected based on the prior fund's performance. Such funds were likely limiting access and were particularly successful during that period. In the 1999-2006 period, such funds no longer outperform and endowments no longer invest unusually in them.

Implications

In 1978, the Prudent Man rule was modified to allow institutional investors to allocate part of their portfolios to alternative assets. Since then, the private equity industry has changed substantially. In 1980, the largest fund raised was the Golder-Thoma $60 million fund that invested in many different kinds of deals, including both venture capital and buyouts. At the time, institutional investors were somewhat skeptical of the industry; GPs, LPs, and portfolio firms were experimenting with different

contractual structures; and "private equity" itself was not an accepted term. By the time of the 2008 financial crisis, individual funds of over $20 billion were being raised, and funds became specialized in particular types of investments so that renewable energy or infrastructure funds were commonplace. Contracts have become standardized, and private equity has become an accepted part of the financial world in which most major business schools teach courses.

We argue that the industry's maturing has had implications for the relations between GPs and LPs. The evidence presented here suggests that the huge inflows of capital and commoditization of the industry has lowered the rents to GPs. In addition, the evidence suggests that because limited access reflected the sharing of these rents, the importance of limited access decreased as well.

The private equity industry has become an important part of institutional portfolios. Yet, it has always been an industry that has been evolving at a rapid rate. Going forward, it is important for industry participants to understand the current state of the industry. Because of the industry's changes, past performance is unlikely to predict future performance. The enormous inflows of capital together with the increasing commonality of experience and knowledge of GPs have likely permanently changed the relations between GPs and LPs, and potentially investors' expected returns as well.

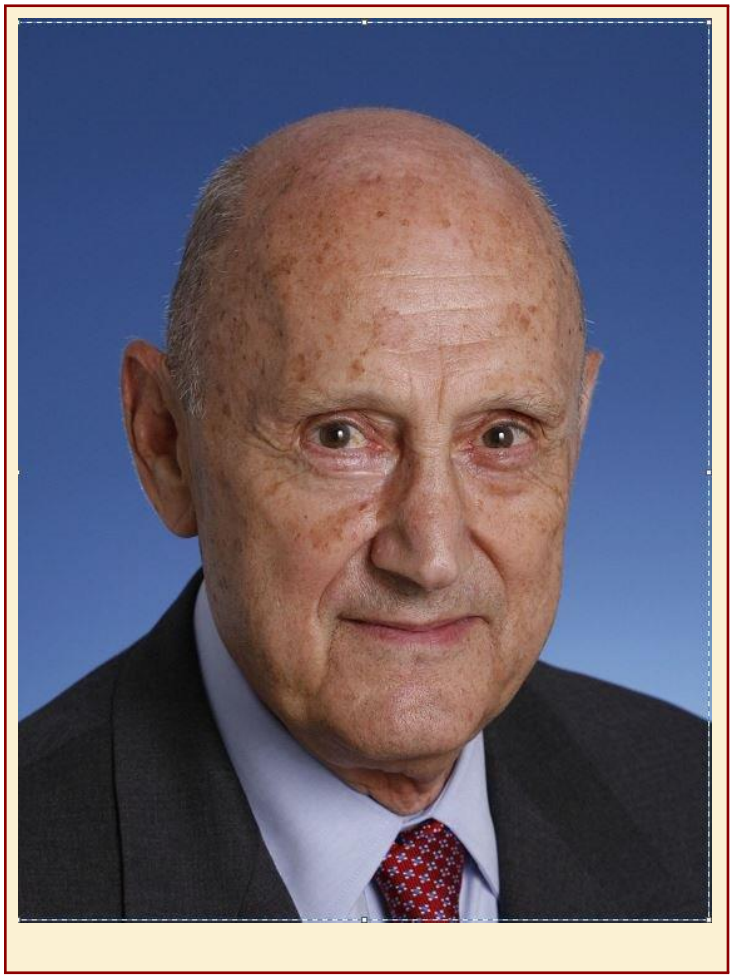

Tarun Chordia, Avanidhar Subrahmanyam, and Qing Tong

Have Capital Market Anomalies Attenuated in the Recent Era of High Liquidity and Trading Activity?

Journal of Accounting and Economics | Volume 58, Issue 1 (August 2014), Pages 41-58
(please cite only the original publication, not FAME)

Our paper explores how a host of capital market anomalies have evolved in recent years, as stocks have become more liquid and more actively traded, and hedge funds that exploit these anomalies have become more popular. We empirically explore how the Fama and MacBeth (FM) (1973) cross-sectional coefficient estimates and the decile-based hedge portfolio returns, have changed over time. We find that most of the hedge portfolio returns and regression coefficients for the anomalies attenuate towards zero over time.

We conduct additional analysis to identify the reason behind the attenuation of the anomaly profits. Specifically, we try different identification schemes, including (i) the decrease in the tick size due to decimalization, (ii) the impact of hedge fund assets under management (AUM), (iii) the impact of the aggregate short interest, and (iv) aggregate share turnover. All of above variables are proxies for arbitrage activity.

We find that the characteristic premiums (i.e., FM coefficients) of almost all anomalies have attenuated from before to after decimalization. The average returns as well as the Sharpe ratio from a comprehensive anomaly-based trading strategy have more than halved after the shift to decimal pricing. Further, the returns to several anomalies are negatively related to hedge funds' AUM, short interest, and aggregate trading activity, indicating a link between arbitrage proxies and attenuation in anomalies.

Characteristics that capture the anomalies

The included firm characteristics to capture equity market anomalies are

SIZE: The natural logarithm of the market value of the firm's equity.

BM: Book equity for the fiscal year-end in a calendar year divided by market equity at the end of December of that year.

TURN: The logarithm of the firm's share turnover, measured as the trading volume divided by the total number of shares outstanding.

R1: The lagged one month return.

R212: The cumulative return on the stock over the eleven months ending at the beginning of the previous month.

ILLIQ: The Amihud (2002) measure of illiquidity. It is the average daily price impact of order flow and is computed as the absolute price change per dollar of daily trading volume:

$$\mathrm{ILLIQ}_{it} = \left[\sum_{d=1}^{D_{it}} (|R_{itd}|)/\mathrm{DVOL}_{itd} \times 10^6 \right]$$

where R_{itd} is the return for stock i, on day d of month, DVOL_{itd} is the dollar trading volume of stock i, on day d of month t, and D_{it} represents the number of trading days for stock i in month t.

ACC: Accounting accruals—the change in non-cash current assets, less the change in current liabilities (exclusive of short-term debt and taxes payable), less depreciation expense, all divided by average total assets.

AG: Asset growth—the year-on-year percentage change in total assets.

ISSUE: New issues—the change in shares outstanding from the eleven months ago.

IVOL: Idiosyncratic volatility—the standard deviation of the regression residual of the Fama and French (1993) three-factor model using daily data within a month.

PROFIT: Profitability—earnings divided by book equity, where earnings is income before extraordinary items.

SUE: Standardized unexpected earnings—the most recently announced quarterly earnings less the earnings four quarters ago, standardized by its standard deviation estimated over the prior eight quarters. This is used to proxy for earnings surprises, in order to analyze post-earnings-announcement-drift (PEAD).

Anomaly related profits are declining over time

The base sample includes common stocks listed on the NYSE/AMEX (NYAM) from 1976 to 2011. A second Nasdaq sample begins in 1983, because trading volume on Nasdaq, required for computation of turnover and the illiquidity measure, is not available prior to this date.

We analyze the economic and statistical significance of anomalies in two ways.

The first way constructs extreme decile portfolios that are long the high characteristic values and short the low characteristic values. Table 1 provides the coefficients of the exponential time trend for the hedge portfolio returns (all coefficients are multiplied by 10). The coefficient estimate on time, $\hat{b}$, for the NYAM portfolio returns formed by sorting on the past one month return (reversal strategy) is 0.0124. An attenuation obtains when the trend is in the opposite direction of the baseline effect in Table 1. Because the return from buying (selling) stocks with low (high) values of the past month's return is –0.5% (from Table 1) per month, a positive trend coefficient is consistent with a decline in profits to a reversal strategy over time. The coefficient on cumulative returns over the past two to twelve months (momentum strategy) is –0.0142. Because the return to the momentum strategy is

positive, a negative coefficient signifies a decline in profits over time. Similarly, the negative trend coefficient suggests a decline in profits to supplying liquidity. The signs of the NYAM trend coefficient estimates suggest an attenuation in anomaly-based trading profits for ten of twelve anomalies.

There is a significant decline (10% level cutoff) in the profitability of eight of twelve of the anomalies for NYAM stocks. In the case of Nasdaq stocks, ten of twelve anomalies attenuate, with four (value, momentum, profitability, and PEAD) demonstrating significant attenuation. In Table 1, we also show the number of significant accentuations (at the 10% level, i.e., the number of trend coefficients with a -value exceeding 0.9). There is a strong asymmetry in significant attenuations and accentuations. While eight anomalies have significantly attenuated for NYAM stocks (four for Nasdaq stocks), only one anomaly has accentuated for NYAM stocks, and no anomaly has significantly accentuated for Nasdaq stocks. The overall picture is quite consistent with attenuation in anomaly profits over time.

In the last row of Panel A of Table 1, we also present the result of fitting the exponential trend to the portfolio return obtained by equally weighting the twelve individual anomaly-based hedge portfolios (henceforth termed the "EW hedge portfolio"). The returns on this portfolio show a significant attenuation for both NYAM and Nasdaq stocks.

The second way show that these results also obtain in Fama-Macbeth specifications. In eleven of twelve cases, the characteristic premiums attenuate for NYAM stocks. Six of the twelve NYAM characteristic premiums (those for size, monthly reversals, momentum, accruals, and profitability IVOL, and PEAD) exhibit a significant declining trend (five of these six cases attenuate with -values of 0.05 or less). The half-lives for reversals, accruals, and SUE, are 5.4, 9.3, 10.2, and 6.7 years, respectively. Size, R1, R212, IVOL, and SUE had predicted coefficients in 2010 that were just about zero.

Increased liquidity and arbitrage activity make it difficult to earn anomaly related profits

Our paper studies several equity market anomalies over more than three decades. It finds that the regime of increased liquidity, trading activity, and hedge funds' assets under management have resulted in a decrease in the economic and statistical significance of these anomalies.

In order to establish a link between increased arbitrage activity and the decline in the profitability of the anomaly based trading strategies we examine (i) the impact of the decline in the tick size due to decimalization and (ii) the impact of hedge fund assets under management, short interest, and trading activity on the anomaly based predictability. The decrease in the tick size is associated with improvements in liquidity and a decline in trading costs. We find that the characteristic premiums have declined towards zero for several anomalies in the post-decimal period, and the profits to a comprehensive anomaly-based portfolio have declined by about half in the post-decimal period. Moreover, the impact of many anomalies such as size, reversals, momentum, and PEAD, as well as the return to a composite portfolio have declined with an increase in hedge fund assets, short interest, and aggregate share turnover, suggesting that arbitrage activity has indeed led to a decline in the profitability of the anomaly based trading strategies.

These results are relevant because they indicate that it may be challenging to attain consistent profits from well-documented anomalies in the future. Note, however, that while anomaly profits, based on a comprehensive hedge portfolio, decline significantly in the recent high-liquidity era, they remain

statistically significant. Looking to the future, these profits may not disappear completely because of limits to arbitrage or imperfect competition amongst arbitrageurs that preserves some rents.

Our analysis suggests that it might be fruitful to explore the effect of mechanisms and policies that remove trading frictions and improve liquidity in markets. The results suggest that cross-sectional return predictability would diminish to a greater extent in countries that have experienced greater enhancements in trading technologies and larger increases in trading activity and liquidity. This hypothesis awaits rigorous testing in an international context.

Table 1: Exponential Trend Fits to Hedge-Portfolio Returns, 1976-2011.

	Attenuation	
	b	p-value
SIZE	0.082	0.045**
BM	–0.047	0.076**
R1	0.124	0.001**
R212	–0.142	0.012**
TURN	0.046	0.103
ILLIQ	–0.058	0.074**
ACC	0.028	0.082**
AG	0.027	0.123
ISSUE	0.036	0.046**
IVOL	–0.026	0.687
PROFIT	0.058	0.961**
SUE	–0.048	0.006**
Number of attenuations	10 of 12	
(significant attenuations)	(8 of 12)	
(significant accentuations)	(1 of 12)	
EW hedge portfolio return	–0.045	0.001**

Table 1 regresses the extreme-decile hedge portfolio returns on an exponential time trend. The fitted model is $Y = \alpha \cdot \exp(b \cdot t + u)$, where Y is one plus the quantity of interest and t is time, scaled to between –1 and +1. The EW hedge portfolio return weights the twelve anomalies-based hedge portfolio returns equally. The coefficients estimates are multiplied by 10.