swiss:finance:institute

sfi:practitioner roundups

Academic Insights on Finance Matters for the General Public

Summer 2016

SFI Practitioner Roundups aim to present the latest industry-oriented research findings and ideas from SFI faculty in a concise, focused manner.

Interested in reading SFI Practitioner Roundups online?
Read the latest roundups at www.sfi.ch/knowledge.

To subscribe to SFI Practitioner Roundups, e-mail your name, job title, and the name of your company to kc@sfi.ch.

Damir Filipovic and Loriano Mancini

Taming Volatility: A New Perspective on Variance Swaps

Page 4

Thorsten Hens

Is There "Swissness" in Investment Behavior?

Page 8

Semyon Malamud

How Do ETFs Influence Financial Markets?

Page 12

Michel Habib

Do Mutual Fund Fees Reduce Investor Returns?

Page 16

Karl Schmedders

Mean-Variance and the Carry Trade—An Ideal Match?

Page 20

Alberto Plazzi

Birds of a Feather—Do Hedge Fund Managers Flock Together?

Page 24

Any views expressed in these individual practitioner roundups are those of the authors of the papers cited, and those authors alone are responsible for each roundup's content.

Taming Volatility: A New Perspective on Variance Swaps

While the variance swaps market has grown significantly and is currently worth trillions of dollars, pricing these derivatives remains problematic. SFI's Damir Filipovic and Loriano Mancini, and former SFI PhD student Elise Gourier, provide a powerful new pricing model, revealing how to use variance swaps to construct and maintain an optimal investment portfolio.

By Damir Filipovic, Loriano Mancini, and Elise Gourier

Volatility is a fact of life in financial markets. The trading prices of assets, including—for example—stocks, vary over time. They move up and down on a daily basis; some more than others. At times, during economic crises or market panics for example, prices can fluctuate wildly. For some of those actively involved in financial markets volatility is welcomed as an opportunity to make money, trading the daily ebb and flow of price movements. For others though, volatility provides unwanted turbulence, making risk harder to assess, return curves harder to predict, and portfolio and wealth management more difficult. Fortunately for those who prefer a less volatile world or at least markets where volatility is more predictable, finance's seemingly never ending capacity for innovation has produced a number of solutions. To begin with, you can trade volatility using the Chicago Board Options Exchange Volatility Index, more commonly known as the VIX. By using a variety of strategies it is possible to speculate on, or minimize exposure to, volatility in the future. Or you can get even closer to volatility by trading variance swaps.

Volatility exposure

A variance swap is an over-the-counter financial derivative that allows two parties to take a position on volatility. Such derivatives underpin the VIX index. In effect a variance swap reflects the variation or degree of movement of an underlying asset, whether that is a stock price, currency, commodity, or an interest rate. One party purchases the variance swap contract from a counterparty at a fixed price reflecting the predicted volatility over time. The swap runs for a specific period of time—the term. At the end of this period the actual realized variance over the period is averaged out and converted into a monetary amount; this will either be more than or less than the strike price and the difference is settled between the parties accordingly. In this way it is possible for parties to protect themselves—to hedge—against a level of volatility over a particular time horizon. This is similar to buying insurance against a specific risk, such as a bad weather event for example.

"While the use of variance swaps may sound fairly straightforward, one considerable challenge is pricing the swaps correctly."

While the use of variance swaps may sound fairly straightforward, it is not. One considerable challenge for sellers, for example, is pricing the swaps correctly. And, while there are pricing models available, they do not necessarily properly factor in variance over time. However, recent research by finance academics including SFI's Damir Filipovic and Loriano Mancini has provided finance professionals with a new tool for assessing the value of these derivatives. The authors have created a new class of model for modeling variance swap rates, and therefore risk. A non-linear model that is practical to use, it also goes beyond existing industry standard models. In demonstrating the model's applicability, the authors show how it accurately reflects the variance swap rates, and thus the volatility, for the S&P 500 for the period from January 1996 to January 2010. At any one point in time the model shows the variance swap rate for a maturity of 2, 3, 6, 12, and 24 months. The sample time period includes the most recent financial crisis and, as would be expected, the model shows variance swap rates rising steeply post crisis reflecting increased anticipated volatility and investors' willingness to pay for volatility protection.

"Given that the model can be used to produce future scenarios, it can be used as a tool for helping to develop risk management strategies."

The authors' model makes possible more realistic pricing of variance swaps, and does this for any maturity, including those that might not be popularly quoted. One might, for example, have two parties where one wants to buy a ten-month variance swap from the other: the model can be used to work out a

reasonable variance swap rate for a ten-month time horizon. Also, given that the model can be used to produce future scenarios, it can be used as a tool for helping to develop risk management strategies—by financial institutions planning their capital requirements, for example.

Optimal portfolio strategy

Furthermore, and importantly, the variance swap model can be used by those involved in selecting and managing investment portfolios. The authors investigate the challenge of optimal portfolio choice, and the role of variance swaps as a component in an optimal mix of investments. In particular, they look at a portfolio investing in a stock index, a mixture of variance swaps (three month and two-year horizons), index put options, and risk free bonds.

> "Investors can optimize the relative weightings of their portfolios dynamically, depending on their risk profile. This is a method that would have outperformed the S&P 500 index over time."

In doing so, their analysis reveals that an optimal portfolio involves adopting a short–long strategy with respect to variance swaps, with weighting depending on the risk profile of the investor. On average the long-term strike price is greater than the realized variance, so shorting the long-term variance swap allows the investor to pocket any difference. At the same time by buying short-term variance swaps the investor partially hedges portfolio losses against imminent volatility in the underlying asset. Indeed, using the authors' methods investors can optimize the relative weightings of their portfolios dynamically, at any given time, depending on their risk profile and the maturity of the various components of their portfolio. As the authors demonstrate over a fourteen-year period from 1996 to 2010, this is a method that—with the portfolio rebalanced daily—would have outperformed the S&P 500 index over time for a risk-tolerant investor. While for a risk-averse investor it would have produced steady wealth growth, avoiding volatility and reducing the risk of significant losses during a market crisis. All this during a period in which there were two substantial market crashes.

The full paper

—

http://bit.ly/1H8fA5L

About the Authors

Damir Filipovic

Damir Filipovic holds the Swissquote Chair in Quantitative Finance at the École Polytechnique Fédérale de Lausanne (EPFL) and an SFI Senior Chair. He also acts as head of SFI at EPFL. His research interests lie in quantitative finance and risk management and he received his PhD in Mathematics from ETH Zurich.

Loriano Mancini

Loriano Mancini is Assistant Professor of Finance at EPFL and holds an SFI Junior Chair. He received his PhD in Economics and Finance from the University of Lugano. His primary research interests are volatility modeling and asset pricing.

Elise Gourier

Lecturer at Queen Mary, University of London and former SFI PhD student.

Key words

—

Stochastic volatility
Volatility is a measure of the variation in asset returns, which changes randomly over time. Stochastic volatility is modeled as a stochastic process.

Variance swap
A derivative contract to trade volatility.

Quadratic term structure
A nonlinear stochastic factor model for pricing variance swaps over different terms.

Dynamic optimal portfolio
The optimal portfolio resulting from maximizing an investor's utility over all possible dynamic trading strategies.

"Given that the model can be used to produce future scenarios, it can be used as a tool for helping to develop risk management strategies."

Is There "Swissness" in Investment Behavior?

Knowledge and emotionality determine investment behavior. Are Swiss investors better prepared to avoid investment mistakes than their closest neighbors abroad?

By Thorsten Hens and Kremena Bachmann

In Switzerland the term "Swissness" is used to indicate how consumer products are distinguishable from similar products from other countries. While Swissness is established as a successful brand for products, little is known about the question of whether Swissness exists in other areas.

The authors analyze whether there is Swissness in the behavior surrounding investment decisions. Since such decisions are determined by financial literacy and investment emotionality, the authors analyze whether the Swiss have different financial knowledge and a different ability to deal with the emotions that bring about investment mistakes. The study measures financial knowledge with questions that address the long-term rewards of different asset classes, the size of well-diversified portfolios, and the relative importance of different performance drivers. The occurrence of emotionally motivated mistakes is addressed with questions on risk taking and strategy following gains and losses.

"The Swiss decide less emotionally and most have better financial knowledge than their neighbors abroad."

The study is based on a large survey completed by more than 2,000 participants spanning the three linguistic regions of Switzerland and the closest neighboring countries (Germany, France, and Italy). The authors find that Swiss make decisions differently from non-Swiss. For most questions, even Swiss who speak different languages to one another display greater similarities to one another than they do to their "neighbors" abroad who speak the same language. Moreover, the authors find that Swiss in all language regions are more likely to avoid emotionally motivated investment mistakes, even if only the German- and the French-speaking regions of Switzerland can boast higher levels of financial knowledge than are evident elsewhere. The authors conclude that there is Swissness, in behavior regarding investment decisions and in emotional investment competence, that cannot be explained by proxies for industry differences or by differences in the demographic and socioeconomic characteristics of the regional samples.

Although the language an individual speaks might be more relevant to that individual's self than their country of residence, the results suggest that in countries with multiple identities there might be some traits on the national level that exert an influence on emotional investment competence.

"Investment experience improves financial knowledge but not the ability to deal with emotions after gains and losses."

Beyond the regional differences in the investment competence, the authors find that individuals who believe they have a greater investment experience also have better financial knowledge. This observation supports the current practice of using clients' investment experience levels in judgments regarding the suitability of investment products. However, individuals with greater investment experience are also more likely to decide emotionally once they have made an investment. Investors with greater investment experience may have a better understanding of the financial risks involved in their investments, but the study results reveal that they are even less prepared to deal with the related emotional risks than are other investors. Similarly, unhelpful is information on financial wealth or on higher education as a proxy for competence. While less wealthy or less educated clients tend to have lower levels of financial knowledge than wealthy clients who have undergone higher education, there are no differences between the emotional competence of these two groups. It seems that the current education system helps to improve financial knowledge but it does not necessary improve individuals' ability to deal with emotions when making investment decisions. Such ability improves with age, income, and professional responsibility. Income is also a much better proxy for financial knowledge than is financial wealth.

“Investment experience improves financial knowledge but not the ability to deal with emotions after gains and losses.”

"The Swissness we find in investment behavior suggests that Swiss investors may enjoy greater freedom with regards to their investment decisions and thus need less investor protection."

These results have important policy implications, for example for the design of laws that aim to protect investors, as does FIDLEG. The results show that Swiss investors may enjoy greater freedom with regards to their decisions and thus need less protection than investors in neighboring European countries. In all countries, however, educational measures put in place to prepare investors to act responsibly should consider the negative impact of emotions on the investment decision making process, and should not confine themselves to the dissemination of financial knowledge alone.

About the Authors

Thorsten Hens
Thorsten Hens is SFI Professor of Financial Economics at the University of Zurich and Adjunct Professor of Finance at the Norwegian School of Economics in Bergen. He studied in Bonn and Paris and previously held professorships at Stanford and Bielefeld. His main research area is behavioral finance.

Kremena Bachmann
Senior Research Associate, University of Zurich.

Key words

—

Swissness
High quality of goods and services that one finds in Switzerland and not so much abroad.

Financial literacy
Ability to understand and communicate about financial terms.

Behavioral finance
An area of research that combines psychology and finance in order to get a better understanding of individuals' deviations from rational decision making and of market deviations from market efficiency.

Investment competence
Ability to take suitable investment decision.

The full paper

—

http://bit.ly/1PDcQLb

How Do ETFs Influence Financial Markets?

ETFs have attracted considerable attention in recent years. Yet despite growing concerns about the impact of ETF trading on the behavior of financial markets not one theoretical model exists to analyze that impact. SFI's Semyon Malamud has developed such a powerful tool.

By Semyon Malamud

An exchange-traded fund (ETF) is an investment fund traded on a stock exchange. Most ETFs are passive, index-tracking funds that mimic the performance of a specific index. A wide variety of ETFs—using different ETF holdings including stocks, commodities, or bonds—are actively traded in financial markets.

"In 2015, total ETF assets surpassed USD 3 trillion… average annual growth rate since the early 1990s has been 26 percent."

Two major features distinguish ETFs from open-end mutual funds. First, retail investors can buy and sell ETF shares on a stock exchange through a broker-dealer. Second, ETF shares can only be created by financial institutions called "authorized participants" (APs). These serve as broker-dealers and ETF market makers on the stock exchange and have a special agreement with the ETF issuer. A key ETF variable is a fund's net asset value (NAV)—the sum of all its assets less any liabilities, all divided by the number of shares outstanding. ETFs and other exchange-traded products have experienced tremendous growth in recent years, attracting considerable attention from investors, regulators, and academics alike. In 2015, total ETF assets surpassed USD 3 trillion, according to Markit. The average annual growth rate of ETF assets since the early 1990s has been 26 percent—twice that of actively managed assets.

"Do ETFs increase volatility and systemic risk? Should ETF trading be regulated or even prohibited?"

All this naturally raises concerns about the impact of ETF trading on the behavior of financial markets. Do ETFs increase volatility and systemic risk? Does ETF trading "steal" liquidity from the underlying "simple" securities? Should ETF trading and the creation of new ETF products be regulated, or even prohibited?

Numerous research papers have empirically investigated both the pricing of ETFs and their impact on market risk and liquidity. Most argue that ETF trading increases both volatility and systemic risk by channeling new demand and supply shocks into the ETF basket securities. A particular topic of ongoing debate has been regulatory concerns with regard to the hidden risks to which ETF investors are exposed and the threat that ETFs pose to market stability. Indeed, the US Securities and Exchange Commission has begun investigating the role of ETFs in raising market volatility.

"The key difference between ETFs and other financial securities is the presence of the creation/redemption mechanism. Malamud's model allows the effects of that mechanism on the underlying prices, volatility, and liquidity to be explicitly characterized."

Despite all these concerns, not one single theoretical model that could be used to analyze the impact of ETFs on the functioning of financial markets has been developed. Without such an equilibrium model, it is difficult to assess the role and interactions of all the complex mechanisms that shape the underlying market structure. SFI's Semyon Malamud remedies this state of affairs with his recently developed and powerful modeling approach.

The key difference between ETFs and other financial securities is the presence of the creation/redemption mechanism. A realistic model of ETFs must include this mechanism. The

analytical solution developed by Malamud in his model allows the effects of the mechanism on the underlying prices, volatility, and liquidity to be explicitly characterized. The ETF creation/redemption mechanism serves two purposes. First, ETF issuers effectively offer an additional primary market to APs—a market that serves as a source of complementary liquidity and therefore improves liquidity in the secondary ETF market. Second, the mechanism ensures that ETF shares trade in-line with the underlying NAV: if the ETF price deviates from the NAV—that is, if an ETF trades at a premium or at a discount, APs can exploit this arbitrage opportunity by taking opposite positions in the ETF and the underlying basket and then offloading this inventory to the ETF issuer at the end of the trading day.

"The ETF creation/redemption mechanism may serve as a shock propagation channel."

Despite this built-in arbitrage mechanism, ETF prices constantly exhibit puzzling deviations from their NAVs. For example, the iPath S&P GSCI Crude Oil Total Return Index exchange-traded note's (OIL) price rose to a 48 percent premium over its NAV in January 2016. In his paper, Malamud shows that—surprisingly—the creation/redemption mechanism itself may be the cause of such inefficiencies. In fact, one of Malamud's key conclusions is that the mechanism may serve as a shock propagation channel that transfers temporary demand shocks into the future. Sometimes it may be optimal—in terms of lower volatility, higher liquidity, or higher social welfare—to increase creation/redemption fees in order to control ETF growth and limit this shock propagation channel.

The constant growth of the ETF universe naturally raises the question of whether the number of ETFs is excessive and whether the introduction of new ETFs, which may only destabilize the existing market structure, should be regulated. Malamud shows that this may, indeed, be the case.

"ETFs may be both a blessing and a curse."

Introducing new, properly designed ETFs may be welfare improving, reduce overall volatility and systemic risk, and improve liquidity. At the same time, if new ETFs do not span a sufficient number of useful new risk dimensions for the different ETF investor clienteles, their effect may be detrimental for welfare.

About the Author

Semyon Malamud
Semyon Malamud is SFI Associate Professor of Finance at EPFL and holds an SFI Senior Chair. He obtained his PhD in Mathematics from ETH and his main research interest lies in asset pricing.

Key words

—

Exchange-traded fund (ETF)
A mutual fund traded on a stock exchange. Most ETFs track an index and are passively managed but meanwhile there are also ETFs that try to outperform their benchmark.

Liquidity
The degree to which an asset or security can be quickly bought or sold in the market without affecting the asset's price.

Volatility
A measure to evaluate the risk of a security. It describes the average fluctuation of a time-varying variable such as the price of an asset.

Systemic risk
Refers to the risk of a breakdown or severe instability of a market.

The full paper

—

http://bit.ly/1TnQ9kJ

"The key difference between ETFs and other financial securities is the presence of the creation/redemption mechanism."

Do Mutual Fund Fees Reduce Investor Returns?

As pundits and regulators repeat calls to drive down the fees charged by actively managed open-end mutual funds, a new research study argues that such a move could be devastating for investors, removing the incentive to manage actively and compromising investor trust.

By Michel Habib and D. Bruce Johnsen

It is a commonly held belief that actively managed open-end mutual fund fees are excessive and reduce investor returns. According to the authors, this belief is wrong except under the narrowest of circumstances. It is correct only if all things other than fees remain equal. Yet fund assets under management (AUM) vary with the fee per franc of assets; AUM in turn affect investor returns.

"For a wide range of fees, fees in equilibrium affect neither investor returns nor manager payoff."

Once we account for fund flows it follows that, for a wide range of fees, fees in equilibrium affect neither investor returns nor manager payoff. Investors receive a normal rate of return in expectation—that is, the return obtained by investing in an index fund of similar risk. The reason for this is simple, if not necessarily obvious. Consider a mutual fund manager whose investment skills are such that he or she can be expected to generate a positive after-fees alpha. Naturally, investors eager to partake in the fund's over-performance will invest in the fund. As AUM grow, the manager's ability to generate a positive after-fees alpha diminishes. Still, investors can be expected to continue investing in the fund until the after-fees alpha equals zero and the return is normal. This occurs at a low level of AUM if fees are high and at a high level if they are low; but in either case investors receive a zero after-fees alpha.

And the manager? For the manager too the fee per franc of assets is irrelevant. Conversely to investors, the manager receives the entirety of the added value created in the form of total fees. Since the manager creates this value by possessing the ability to beat the market, none of it accrues to investors, because they invest in the fund until they have driven the rate of return down to the normal rate. Regardless of whether fees are high and AUM low, or fees low and AUM high, their product remains the same and exactly equals the value created by the manager.

"Since the manager creates value by possessing the ability to beat the market, none of it accrues to investors, because they invest in the fund until they have driven the rate of return down to the normal rate."

This says nothing more than that, under competition, fund managers can be expected to earn what economists have long understood as their Ricardian rents. Three assumptions underlie this reasoning. The first is that of competitive equilibrium: AUM will adjust to ensure that fund returns equal the normal rate. While this generally will not be exactly true, it is nonetheless a relatively good and unbiased approximation of reality.

The second assumption is that fees are neither too low nor too high. Fees that are too low are those that fail to cover those costs the manager cannot charge to the fund, such as the opportunity cost of his or her own effort; fees that are too high are those that leave AUM below the level at which the manager can exploit all the value-creating investment opportunities he or she has identified.

Most important is the third assumption, which is that there is no managerial moral hazard in the sense that the manager can be trusted to actively seek all value-creating investment opportunities. This ensures investors do not end up with consistent below-normal returns. In reality, the third assumption fails, but at the same time allows us to better understand the incentive structure of management fees. Identifying value-creating investment opportunities is very difficult. The manager might promise to engage in active management and instead index the entire portfolio while charging a high per-franc fee well in excess of the cost of passive management. This is well known, and is referred to as "closet indexing".

"For a wide range of fees, fees in equilibrium affect neither investor returns nor manager payoff."

"A premium fee ensures that the manager captures the value he or she stands to create through active management."

Even in noisy markets closet indexing will eventually be discovered, but that may take a while, during which investors will earn the below normal returns that are the combined consequence of the zero alpha of passive management and the high fees of active management. One way to avoid such shirking is for the manager to charge a per-franc fee sufficiently high to bond the provision of active management. A premium fee ensures that the manager captures the value he or she stands to create through active management. This is the sense in which fees matter. They need to be at least as high as the minimum necessary to deter shirking. Because of the payout structure of fund fees, a per-franc premium fee—even if not very high—will deter the manager from shirking. This is because mutual fund fees are "back-end loaded". Managers are paid a recurring share of AUM. A one hundred franc increase in AUM as a result of active management yields, say, an additional fifty centimes per year as long as investors stay with the fund. Investors who believe they have been cheated can withdraw and leave the manager without his or her trailing fees. The discipline imposed by investors' threat to withdraw on discovering managerial shirking is a potent one because it imposes a capital loss on the manager. Thus, only a relatively modest fee premium is needed to deter shirking.

Why does all this matter? If heeded, repeated calls to drive down fees would at best be neutral; at worst such a move could be devastating for investors. Mandated fee reductions would be neutral when lower fees simply increase AUM without compromising managers' incentive to engage in active management. In such a case, investors would still receive the normal rate of return. Mandatory fee reductions would be detrimental when fees are driven down so low that the manager loses the incentive to engage in active management. Trust would be lost between investors and active managers, passive management would prevail, and price discovery would suffer.

About the Authors

Michel Habib
Michel Habib is Professor of Finance at the University of Zurich and has been an SFI faculty member since 2006. After graduating from the Wharton School of Business, he taught at the London Business School. His primary research interest lies in corporate finance.

D. Bruce Johnsen
Professor of Law at George Mason University School of Law.

Key words

—

Advisory fees
Fees paid to the investment advisor of a mutual fund.

Alpha
Return received in excess of that required as compensation for bearing the market risk of an investment.

Mutual fund
Pools numerous investors' monies for the purpose of investment in stocks, bonds, cash, or other securities.

General equilibrium
Describes an idealized situation in which prices of goods and services are such that supply equals demand in all markets.

The full paper

—

http://bit.ly/1TwxnYW

Mean-Variance and the Carry Trade—An Ideal Match?

Despite its longstanding prominence, mean-variance analysis still leads to portfolios that are outperformed. Is it, then, a blind alley? A new paper focusing on the carry trade suggests not.

By Karl Schmedders, Fabian Ackermann, and Walt Pohl

Mean-variance analysis is the highest-profile application of mathematical optimization in the practice of finance. Introduced by Markowitz in 1952, it provides a simple answer to the question of how to construct a diversified portfolio of risky assets. Using optimization, it builds portfolios by trading off the assets' returns and risks. Despite its prominence however, the empirical verdict of asset managers after decades of experience in using mean-variance analysis to choose stocks has been largely negative. Considerable evidence shows that mean-variance portfolios often perform worse than even a naive strategy of simply holding equal positions in every asset. So was the optimization approach to portfolio choice a blind alley?

"The carry trade offers considerable scope for diversification—a mean-variance optimal strategy outperforms naive diversification strategies."

A new paper by three authors, including SFI's Karl Schmedders, shows that such a conclusion may be premature. The paper considers a different setting, where mean-variance analysis works well. The carry trade–borrowing in currencies where the interest rate is low and investing where the interest rate is high–is a high-risk, high-reward strategy. The authors show that the carry trade offers considerable scope for diversification, and that a mean-variance optimal strategy outperforms naive diversification strategies.

The mechanism behind the carry trade is simple. For example, in June 2011 the interbank interest rate for a one-year, yen-denominated loan was 0.56 percent, while that for an Australian dollar loan was 5.70 percent. If the exchange rate remains unchanged, an investor will make a return of 5.14 percent in one year. But how successful is this approach once we take into account exchange rate risk? The economic theory, *uncovered interest parity*, suggests that it should not be successful at all. It predicts that exchange rates will move to close up any opportunities for profit. But this theory has not fared well when confronted with data. Instead, much empirical evidence shows that in the short run exchange rate movements are unpredictable, and resemble a so-called random walk.

These results imply that while the carry trade investor still faces exchange rate risk, on average the strategy has a positive expected return. The carry trade has long had a disreputable reputation in both the financial press and the academic literature. Ever since Japan's economic woes pushed its interest rate below 1 percent for almost two decades, borrowing in yen to invest in high-rate currencies has provided an easy route to investment wealth. But did this strategy carry with it a commensurately high risk?

"The carry trade 'goes up by the stairs, but down by the elevator'."

In the wake of the Lehman brothers bankruptcy a worldwide flight to quality caused the yen to appreciate by more than 28 percent in October 2008 against the Australian dollar. This experience suggests that the carry trade strategy carries with it considerable systematic risk. In the academic literature this risk has been identified as "crash risk"–the risk that many small upward moves are paired with the occasional large downward move, such as those that occurred in the immediate wake of the Lehman bankruptcy. This notion is captured by a piece of market folk wisdom: the carry trade "goes up by the stairs, but down by the elevator". This observation suggests a simple explanation for the high returns of the carry trade: they are a compensation for bearing this crash risk. The historical performance of the typical simple carry trade—borrowing in a single, low-interest-rate currency and investing in a single, high-interest-rate currency—is consistent with this explanation, producing Sharpe ratios comparable to those of the S&P 500 stock market index.

In stark contrast to the crash risk theory, the optimal portfolio of 11 common currencies presented in the new paper, which the authors construct using mean-variance analysis, signifi-

cantly outperforms both the simple carry trade strategy and the S&P 500, producing a Sharpe ratio of 0.91 over the period 1990–2015. A naive diversification strategy of holding the five highest interest rate currencies and shorting the five lowest interest rate currencies improves over the simple (1 long, 1 short) strategy but falls well short of the optimal strategy, with a Sharpe ratio of 0.62.

The results in optimizing the carry trade provide a sharp contrast to the performance of mean-variance analysis in the stock market. A critical difference between mean-variance analysis in the stock market and mean-variance analysis for the carry trade is that—for the latter—expected returns are not estimated from prior return data. Instead, under the random walk hypothesis, they are given by the interest rates themselves. Therefore, unlike the stock market case, in the currency setting expected returns are not affected by statistical estimation error. While the covariance matrices must be estimated, the paper finds that the ex ante prediction of the volatility of the optimal portfolio provides a reasonable guide to the ex post realized volatility.

"Applying mean-variance analysis to currency markets can work—three elements in particular contributing to the strategy's success."

In summary, extending the mean-variance analysis to the currency market can work well, the following three aspects contributing to the success of the strategy described: For each asset in isolation, the procedure takes into account the risk–return trade-off between assets. For example, over the sample period, the three lowest interest rate currencies are the yen, the Singapore dollar, and the Swiss franc. The yen generally has the lowest interest rate of the three, but both the franc and Singapore dollar have had lower exchange rate volatility. This makes these last two currencies potentially superior choices on a risk-adjusted basis. Second, the mean-variance strategy exploits the correlation between the assets. In the sample, the two highest interest rate currencies are the Australian and New Zealand dollars. They are both highly correlated, which makes them close substitutes for one another. An optimal strategy can use one as a hedge against the other. Finally, the optimal portfolio can use the aggregate risk–return to choose the total exposure. There is considerable time variation in interest rate spreads. A target mean criterion, for example, will automatically decrease exposure when the spreads are wide, and increase it when the spreads narrow.

About the Authors

Karl Schmedders
Karl Schmedders is Professor of Quantitative Business Administration at the University of Zurich and heads the SFI Knowledge Center. He holds a PhD in Operations Research from Stanford University. His research interest focuses on computational economics and finance.

Fabian Ackermann
Zürcher Kantonalbank.

Walt Pohl
University of Zurich.

Key words

—

Carry trade
An investment strategy whereby an investor borrows money in low-interest-rate currencies and invests them in high-interest-rate currencies.

Mean-variance analysis
An optimization framework for choosing portfolios of assets by trading off the assets' expected returns and risks.

Portfolio optimization
The application of mathematical optimization to portfolio choice problems in finance.

Sharpe ratio
A financial metric that measures how well an investment compensates investors for the riskiness of their investment.

The full paper

—

http://bit.ly/1NF2bQh

"The carry trade offers considerable scope for diversification—a mean-variance optimal strategy outperforms naive diversification strategies."

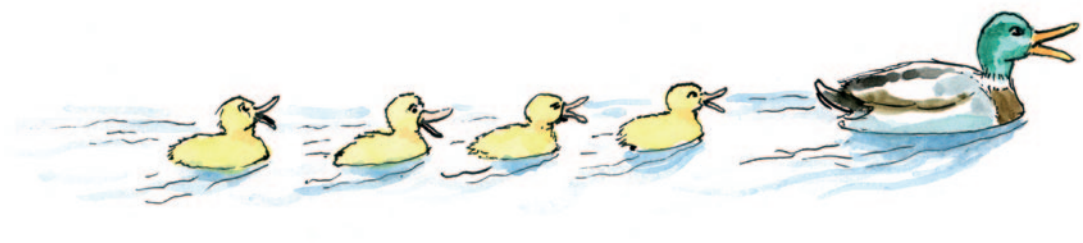

Birds of a Feather—Do Hedge Fund Managers Flock Together?

Only a small portion of hedge funds' alpha can be explained by risk models. Could a change of perspective provide answers, and potential investors with a valuable addition to their due-diligence arsenal?

By Alberto Plazzi, Marc Gerritzen, and Jens Carsten Jackwerth

The spectacular growth of the hedge fund industry in recent decades has stimulated a great interest in understanding the roots of this success. In the academic literature, this question has been tackled mainly by developing increasingly rich factor models. This approach tries to explain hedge funds' stellar performance through exposure to primitive risk factors—that is to say, portfolios deliver a compensation for bearing some systematic risk, such as the market or liquidity. Despite this wealth of research, however, much remains to be understood about the determinants of hedge funds' returns. The average fund still delivers a significant abnormal return, or alpha, and the amount of variability in funds' returns that is not explained by these models remains quite sizeable.

Three authors, including SFI's Alberto Plazzi, adopt a different modeling perspective. Their study looks at hedge funds from the standpoint of a network, and asks whether personal connections that link hedge fund managers together matter when it comes to explaining performance. The authors investigate this question for the UK hedge fund industry, where mandatory filings allow a nearly ideal setting for such research. More specifically, starting in 2002 insurance, investment, and banking companies that operate in the UK (i.e., onshore) are required to report detailed information on current and past employment of their key employees. The resulting data set is maintained and made publicly available by the Financial Conduct Authority, which regulates the UK financial system, with full disclosure of company and employee names.

"Social ties, in the form of prior employment experience, may lead to similarities among funds' returns."

How might prior employment history ultimately affect hedge fund managers' investment decisions, and lead to similarities in funds' returns? Several channels come to mind. First, managers who share a common experience in an industry are likely to have been exposed to similar training. For example, managers who worked in the life insurance sector may develop an attitude with regard to risk that is different from that of employees in the banking sector. This attitude could manifest itself later, as the managers set up their funds with similar levels of risk exposure. For any given industry, having worked for the same employer is likely to exert an additional effect. Managers may have learned portable skills at their former workplace that guide their current investment strategies. To the extent that these skills allow managers to take better decisions, they have the potential to explain some of the abovementioned abnormal performance. Finally, employees may establish personal connections. These connections are likely to be stronger for managers that overlap in their prior experience—that is, managers who worked for the sample employer at the same time. Through the sharing of views and information, these personal connections may lead to correlation (coordination) in managers' trading behavior. In sum, social ties, in the form of prior employment experience, may lead to similarities among funds' returns that show up in the various components of performance: exposures to risk factors (i.e., beta), abnormal performance (i.e., alpha), and the unexplained (mean-zero) idiosyncratic component.

"Having worked in the same industry captures a significant portion of the differences in funds' risk exposures and especially alpha."

Guided by these arguments the authors show that components of the UK hedge fund market are densely linked through such ties, which are found to be important determinants of proximities in any two hedge fund pairs. In particular, having worked in the same (finance) industry and, to a greater extent, having worked for the same employer in the past capture a significant portion of the differences in funds' risk exposures and especially alpha. In contrast, social connections measured by an overlap in prior employment experience explain only differences in the idiosyncratic component of returns. Interestingly, these connections play a much greater role for funds that

invest in styles that are particularly sensitive to the exchange of relevant information, such *event driven* and *merger arbitrage*. Can we conclude that social connections are ultimately responsible for similarities in hedge funds' performance? In order to draw conclusions about causality, several competing channels should be addressed. Managers may self-select and find themselves working for the same prior employer because of similar preferences or risk profiles. Alternatively, there may be other network-related conduits, such as access to local information, that are responsible for similarities in trades. Adding managers' personal characteristics or controlling for geography does not, however, dissipate the effect of prior employment connections.

> "The exchange of information through social ties ultimately has a positive effect on performance."

A potentially more challenging task is to control for managers' skills. The argument here is that (past) employers may hire individuals with similar levels of skills. To account for this, the authors exploit the fact that a subset of managers in the data have previous experience in the hedge fund industry. They then use the abnormal performance in the previously managed fund as a control variable. If skilled managers tend to outperform their peers consistently over time, this should go a long way toward absorbing the effect of innate intelligence. It turns out that the results are also robust with regard to this test. The evidence that fund pairs of connected managers are closest in performance begs the intriguing question whether differences exist in the average returns to hedge funds (not pairs) that are grouped based on the extent of their connectedness. Indeed, the authors show that loading on portfolios of connected hedge funds generates a positive spread in terms of risk-adjusted performance compared to unconnected funds. They conclude that the exchange of information through social ties has ultimately a positive effect on performance.

Overall, these results have clear implications for the industry. They imply that managers' social ties should be considered when evaluating the performance of a single hedge fund or a portfolio, for example via funds of funds. They also stress that social ties should be an important aspect of investors' due diligence processes when deciding which fund to invest in.

About the Authors

Alberto Plazzi
Alberto Plazzi is Assistant Professor of Finance at USI Lugano and an SFI faculty member. He holds a PhD in Finance from the UCLA Anderson School of Management. His research interests lie in empirical asset pricing, real estate, asset management, and financial econometrics. He teaches Master's and executive education courses on risk management.

Marc Gerritzen
Berenberg Bank.

Jens Carsten Jackwerth
University of Konstanz.

Key words
—

Hedge funds
A class of investment funds that invest in a myriad of asset classes, and are (legally) unconstrained in the level of leverage they can undertake.

Alpha
The portion of a fund performance that cannot be explained through exposures to risk factors, and is therefore often referred to as "abnormal".

Social ties
Connections that link individuals that are part of the same social network, for example because they belong to the same community such as a city, a neighborhood, or the workplace.

Financial Conduct Authority
The authority that regulates the UK financial system and ensures the well-functioning of the financial industry.

The full paper
—
http://bit.ly/2ovoOtW

“Social ties, in the form of prior employment experience, may lead to similarities among funds’ returns.”

swiss:finance:institute

Swiss Finance Institute

Swiss Finance Institute (SFI) strives for excellence in research and doctoral training, knowledge transfer, and continuing education in the fields of banking and finance, as befits Switzerland's international reputation as a leading financial center. Created in 2006 as a public–private partnership, SFI is a common initiative of the Swiss finance industry, leading Swiss universities, and the Swiss Confederation.

SFI Knowledge Center

The SFI Knowledge Center promotes and supports dialogue, information flow, and collaboration between academia and the financial services industry. The center performs an active intermediary role in gathering and disseminating financial knowledge, helping find solutions to current problems, and identifying new trends in banking and finance.

Swiss Finance Institute – Walchestrasse 9 – CH-8006 Zurich
KnowledgeCenter@sfi.ch

www.sfi.ch/knowledge