

Katholische Universität Eichstätt-Ingolstadt
Wirtschaftswissenschaftliche Fakultät Ingolstadt



Dissertation

Demystifying pairs trading

Stephanie Sarah Riedinger

Erstgutachter: Prof. Dr. mult. Anton Burger

Zweitgutachter: Prof. Dr. Max Götsche

München, im Juli 2022

Table of content

List of figures	IV
Index of tables.....	V
List of abbreviations.....	VI
1. Introduction	1
1.1 Why you should read this thesis.....	1
1.2 Introduction to core chapters.....	5
2. Methodology	8
2.1 The pairs trading algorithm	8
2.1.1 The identification period	9
2.1.2 The trading period.....	11
2.2 Return calculation	13
2.2.1 Return per trade calculation	13
2.2.2 Calendar Time Return calculation.....	14
3. Baseline results	19
3.1 Data selection	19
3.2 Coding the pairs trading algorithm.....	20
3.3 Empirical results	22
4. Core chapter 1: Demystifying pairs trading: The role of volatility and correlation	24
4.1 Motivation	24
4.2 Selected pairs trading literature	26
4.3 Research Hypothesis - Disentangling the pairs trading algorithm.....	27
4.3.1 The SSD selection criterion.....	27
4.3.2 Return per trade.....	29
4.3.3 Trading frequency	33
4.4 Modified methodology	37
4.5 Empirical results	39
4.5.1 Hypothesis 1 – Identification Period: A critical review of the SSD selection method in the identification period.....	39
4.5.2 Hypothesis 2 – Trading Period: The effect of volatility and correlation on the return per trade...	42
4.5.3 Hypothesis 3 – Trading Period: The effect of volatility and correlation on the trading frequency	45
4.6 Pairs Trading Return in calendar time	50
4.7 Robustness.....	54
4.7.1 Limits of arbitrage.....	54
4.7.2 Short Selling Constraints.....	58
5.6.1. Fama-MacBeth regression.....	60
4.8 Conclusion.....	61
5. Core chapter 2: Slow information diffusion as a driver for pairs trading.....	62

5.1	Motivation	62
5.2	Conditions to exploit information diffusion induced mispricing using pairs trading	64
5.3	Information diffusion exploiting trading strategies vs. pairs trading	65
5.4	The role of the leader-follower relationship of stock pairs in pairs trading	67
5.4.1	Aggregated cross-serial correlation of stock pairs.....	67
5.4.2	Empirical evidence from portfolios with differently cross-serial correlated pairs	68
5.4.3	Common risk factors and cross-serial correlated pairs.....	72
5.5	Slow information diffusion and trade types in pairs trading	74
5.5.1	Leader-follower trades: theoretical foundation	75
5.5.2	Leader-follower trades: empirical evidence	77
5.6	Persistent Leader-Follower Role Allocation	80
5.7	Conclusion.....	85
6.	Core chapter 3: idiosyncratic risk, costly arbitrage, and asymmetry: evidence from pairs trading.....	87
6.1	Motivation	87
6.2	Motivation to apply an alternative mispricing measure.....	88
6.3	Modified test design	90
6.4	Measuring mispricing with pairs trading	91
6.4.1	Pairs trading as an alternative mispricing measure	92
6.4.2	Empirical correlation between the different methods to measure mispricing	93
6.5	Idiosyncratic risk, mispricing, and pairs trading returns	97
6.5.1	Hypothesis 4 - Idiosyncratic risk puzzle and pairs trading returns.....	98
6.5.2	Empirical evidence.....	100
6.6	Arbitrage asymmetry and the idiosyncratic volatility effect in pairs trading	102
6.6.1	Hypothesis 5 - Idiosyncratic volatility effect and arbitrage asymmetry	102
6.6.2	Empirical evidence.....	104
6.7	Robustness - Fama-MacBeth analysis	106
6.8	Conclusion.....	108
7.	Conclusion and outlook	109
	References	VII

List of figures

Figure 1: Example for pairs trading with Coca-Cola and Pepsi.....	3
Figure 2: Pair selection procedure with the SSD selection measure.....	10
Figure 3: Calculation of the return per trade.	14
Figure 4: Calculation of pairs trading returns in calendar time.....	17
Figure 5: Interdependences of the technical decision rules and the return of pairs trading.....	25
Figure 6: Distribution of price spread X_t during the identification and the trading period.....	34
Figure 7: Types of return contribution	75

Index of tables

Table 1: Comparison with related studies	23
Table 2: Correlation and volatility levels over quintiles	38
Table 3: Decomposition of the classical selection criterion (SSD).....	40
Table 4: Return per trade.....	43
Table 5: Level shifts and frequency	48
Table 6: Influence between volatility and correlation level on beneficial volatility increases and correlation decreases	49
Table 7: Total number of trades	50
Table 8: Average monthly trading return.....	52
Table 9: Number of Closing Type No. 3 trades (overshooting trades).....	56
Table 10: Pairs trading in highly liquid markets	57
Table 11: Median percentage contribution of the long leg to the total return.....	58
Table 12: Returns of strong and weak cross-serial correlation portfolios.....	70
Table 13: CES score of cross-serial correlated portfolios.....	74
Table 14: Distribution of trade types in strong and weak cross-serial correlation portfolios ..	77
Table 15: Percentage of leader-follower trades in volatility correlation double-sorted portfolios	79
Table 16: Theoretical binomial distribution in case of a non-permanent role allocation	83
Table 17: χ^2 Goodness-of-fit test for binomial distribution.....	84
Table 18: Correlation and idiosyncratic volatility levels over quintiles	90
Table 19: Correlation between pairs trading and Stambaugh, Yu, and Yuan's (2015) mispricing measure	96
Table 20: Average monthly trading return.....	100
Table 21: Return contribution of the long leg to the total return	104
Table 22: Fama-MacBeth (1973) cross-sectional regression (first natural trade).....	107

List of abbreviations

A	example stock No. 1
B	example stock No. 2
bp	basis points
C	example stock No. 3
CES Score	close economic substitutes score
Corr_Qi	portfolio of pairs whose correlations are within the i^{th} quintile
$CP_{i,t}$	closing price of stock i at day t
D	example stock No. 4
(G)ARCH	(generalized) autoregressive conditional heteroscedasticity model
H_0	null-hypothesis in a significance test
IVOL	idiosyncratic volatility
IVOL_Qi	portfolio of pairs whose idiosyncratic volatilities are within the i^{th} quintile
μ	expected value
n	number of elements in a sample
P	example pair
$P_{i,t}$	normalized price of stock i at day t
\bar{P}_i	arithmetic mean of the price of stock i
$P(X=i p, n)$	probability that observed characteristic X equals i under the condition that the probability for the event is p with n repetitions of the experiment in a binominal test setting
Q_i	score of risk factor i
$\rho_{A,B}$	correlation of stock A and B
p_val	p-value
R_A^2	coefficient of determination of stock A
$r_{i,t}$	return of stock i at time t in %
SSD	sum of squared differences (= distance measure)
σ_i	standard deviation of stock price i
σ_i^2	volatility of stock price i
σ_{AB}	pair volatility of stock A and B
σ_{hist}	historic standard deviation of the spread between the normalized prices of stock A and B during the identification period (= trading signal)

$\sigma_{IVOL,i}$	idiosyncratic volatility of stock i
t	point in time
T	last point in time of a fixed time period
VIX	CBOE Volatility Index
Vola_Qi	portfolio of pairs whose volatilities' are within the i^{th} quintile
$V(X)$	expression for volatility of stock X in a mathematical derivation
$w_{i,t}$	weight of stock i in \$ at time t
X_t	normally distributed random variable that describes the price spread between two stocks on day t
z_i	element of the standard normal table

1. Introduction

1.1 Why you should read this thesis

In times of low interest rates, trading strategies that promise a high risk-adjusted return are of particular interest. In particular, intelligent arbitrage trading strategies that exploit mispricing to earn a risk-free premium are attractive to practitioners. The key success factor in such trading strategies is the identification of mispriced stocks.

In an efficient market, two assets that receive an identical payoff in each state of the world must have the same current price as arbitrage opportunities that arise otherwise. If one asset is more expensive than another, investors sell the more expensive asset and buy the cheaper asset. The investor receives the identical payoff in each state of the world with the new asset, but earns the price difference on top by the transaction. If enough investors use this arbitrage opportunity, the supply of the more expensive asset increases, which consequentially decreases the price. Likewise, the rising demand for the cheaper asset increases the price of the second asset. This process continues until both assets are equally priced. In a perfect world, there is a sufficient number of arbitrageurs in the market, and nothing prevents them from exploiting the profitable mispricing opportunity. However, for various reasons, markets are not always perfectly efficient, which creates opportunities for mispricing. The exploitation of this temporary mispricing generated by market inefficiencies is the main business of hedge funds with their arbitrage trading strategy.

One of the most profitable arbitrage trading strategies is pairs trading. The success of pairs trading on the market is a clear violation of the efficient market hypothesis. Studying the strategy is therefore not only interesting for practitioners, but also an interesting research field for the academic world. Nonetheless, relatively little academic research has been published about pairs trading until

now since the trading strategy was first introduced in academic literature by Gatev, Goetzmann, and Rouwenhorst (2006).

What is pairs trading? Pairs trading is a relative value trading strategy that exploits arbitrage opportunities with the following procedure: in the first step, pairs trading identifies stock pairs with an almost identical price pattern during a 12-month identification period. Natural candidates for pairs trading are, for instance, competitors in the same market, like Coca-Cola and Pepsi, or BMW and VW. For efficiency reasons, pairs trading as an algorithm-based strategy does not hand-pick pairs. To automate the selection process, Gatev, Goetzmann, and Rouwenhorst (2006) compute a KPI, the SSD, for all possible stock pair combinations at the end of the 12-month identification period. The SSD is smaller, the more similar the price development of both stocks is. The twenty pairs with the lowest SSD are included in a trading portfolio, and they are eligible for trading in a subsequent 6-month trading period.

During the subsequent trading period, the price development of each stock is compared daily with the development of the pair's second stock. The underlying argument of pairs trading is that price deviations of two close economic substitutes are temporary and dissolve in the near future, which can be profitably exploited. In order to do so, a new trade is initialized as soon as the price spread¹ between both stocks exceeds two standard deviations, as measured during the identification period. Gatev, Goetzmann, and Rouwenhorst (2006) assumes that the worse performing stock is temporary underpriced, and suggest initiating a \$1 long position in the worse performing "loser" stock. Accordingly, they assume that the better performing "winner" stock is overpriced, so they consequently initiate a \$1 short position in the winner stock. If the price difference is indeed the result of inefficient markets, then both stock prices revert and the mispricing disappears. Both positions are closed

¹ All prices are normalized to one at the beginning of the identification period and at the beginning of the trading period.

as soon as both prices cross. Pairs trading profits if the “loser” stock rises and the “winner” stock falls. Figure 1 illustrates one trade for Coca-Cola and Pepsi during March and May 2016.

Figure 1: Example for pairs trading with Coca-Cola and Pepsi

This figure shows the

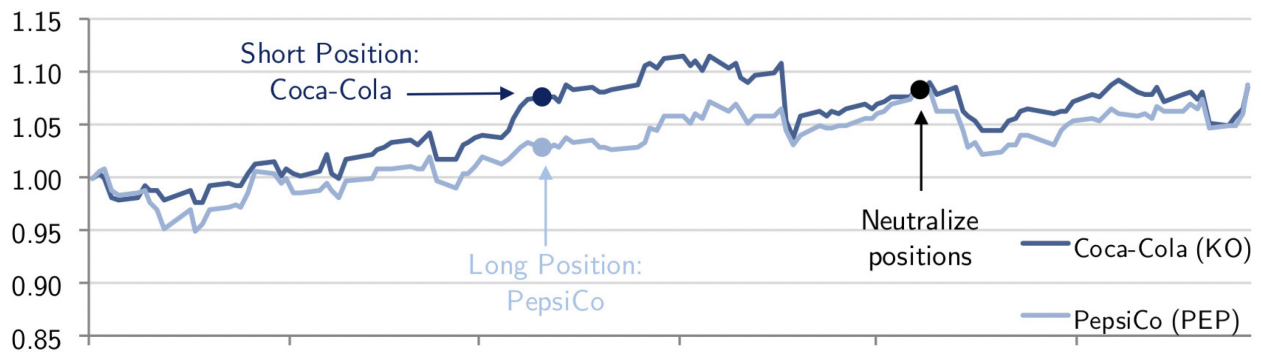


Figure 1 displays the price development of Coca-Cola and Pepsi and illustrates what trading activities pairs trading would perform during the trading period. Both stock prices are normalized to 1 on the first day of the trading period and accordingly only fractions of the stocks are bought or sold.

Pairs trading has become an interesting field for academic research since Gatev, Goetzmann, and Rouwenhorst (2006) for several reasons: in a perfect, frictionless markets, all assets are correctly priced. The evidence of positive pairs’ trading returns prove the existence of market inefficiencies. Therefore, studying pairs trading contributes new valuable insights to better understand the origin, the persistence, and the liquidation of mispricing. Learning more about the impact of responsible drivers, market conditions, and the inherent risk profile plays an important role in this research field. In addition, the pairs trading methodology can also be used as an alternative, out-of-the-box test method to measure a different type of mispricing. Likewise, pairs trading is of high interest for practitioners because of its self-financing character combined with a high return potential – in times of low interest rates now more than ever. Moreover, its market-neutral character makes it an interesting addition to every hedge fund portfolio. Pairs trading is, for instance, used by the American

Century Alternatives® Market Neutral Value Fund and many trading platforms offer pairs trading tools for users.

Previous papers examined the profitability of pairs trading in different markets and asset classes, alternative trading algorithms, and the source of mispricing. My interest in pairs trading developed while I coded a basic pairs trading program during for my master's thesis and found some anomalies that could not be explained by current literature. Far too little attention has been paid to the concrete impact of the strictly rule-based trading algorithm in the academic world. Up to now, hardly anything is known about whether the selection of pairs and the trading procedure are biased by the strict selection regime and the trading rules. The most difficult question, but at the same time one of the most important ones, is the question of which source generates pairs trading returns. However, only few studies have investigated return sources in a systematic way so far. The central purpose of this thesis is to provide new insights to close the identified knowledge gap and to contribute to a deeper understanding of market inefficiencies, arbitrage, and algorithm-based trading strategies.

The first part of this thesis explains the general methodology and the return calculation of pairs trading. Afterward, I introduce my data selection and summarize the base line results for pairs trading before I come to the main part of this thesis. A review of relevant, corresponding literature is provided in each chapter. The main part of this thesis consists of three major chapters that empirically investigate three novel aspects of pairs trading that are important for the academic world and practitioners alike. In the next section, I introduce and motivate the central research questions for each of the three chapters and explain the contribution of the outcome for the academic discussion and the resulting implications for practitioners. All three core chapters are stand-alone sections and can be read without the knowledge of the previous core chapter. Together, these three core chapters provide novel insights that are reviewed and interpreted in the final conclusion.

1.2 Introduction to core chapters

The first core chapter “Demystifying pairs trading: The role of volatility and correlation” explains why the profitability of the popular trading strategy “pairs trading,” as introduced by Gatev, Goetzmann, and Rouwenhorst (2006), varies across markets, industries, macroeconomic circumstances, and firm characteristics. I disentangle how the strategy’s two input parameters, stock volatility and pairs correlation, influence the rule-based pair selection, the trading algorithm, and the total return. I find that the traditionally applied pair selection fails to identify the most profitable pairs. In the US market, small algorithm modifications increase the average risk-adjusted monthly return from 37 bp to 76 bp between 1990 and 2014. The findings are robust regarding liquidity issues, bid-ask spread, and limits of arbitrage. The findings of this chapter complement the literature on pairs trading by providing empirical evidence on how the predefined trading algorithm has a technical impact on the total return of pairs trading. These insights are important, as the technical effect and the limitations of the algorithm might be partly the driver behind previous findings that were mistakenly explained by economic factors. The finding that the traditional pairs trading algorithm does not select the pairs with the best return-risk profile is especially important for practitioners. This insight demonstrates the importance of being cautious when interpreting the results of algorithm-based investment strategies.

The central purpose of the second core chapter, “Slow information diffusion as driver for pairs trading,” is to provide evidence that information diffusion is a major source of pairs trading returns. It is often claimed that the returns of technical strategies are a random product in general, as trading signals and positions are determined by an algorithm instead of fundamental data or economic events. The investigation of the economic source of the return of technical strategies is therefore an important field in the academic world. I provide several pieces of evidence that suggest that the return of pairs trading is not earned by chance. Gatev, Goetzmann, and Rouwenhorst (2006) argue

that pairs trading returns can be understood as premium for the exploitation of temporary stock mispricing. Previous pairs trading literature identifies a number of possible factors that result in mispricing, such as gradual information diffusion, among others. Relevant information is gradually processed by different stocks for numerous reasons. The gradual information diffusion creates temporary mispricing that is characterized by leader-follower relationships of stocks that gradually incorporate information. I argue that gradual information diffusion is a perfect candidate to explain pairs trading returns, because the pairs trading algorithm is designed to perfectly exploit the induced leader-follower relationship. In this core chapter, I therefore investigate the effect of leader-follower relationships of stocks on the return of pairs trading. Additionally, I contribute to the existing pairs trading literature by providing evidence that information diffusion is a major source of return.

Understanding the role of gradual information diffusion in pairs trading is likewise important for practitioners. A fundamental story behind the trading strategy facilitates explaining the idea of pairs trading to clients. Furthermore, the knowledge enables an active pairs trader to adequately manage the fundamental risk in times of turbulent markets when information processing is impeded.

The third core chapter, “Idiosyncratic risk, costly arbitrage, and asymmetry: Evidence from pairs trading,” applies pairs trading as a novel test setting to contribute to the growing literature about the idiosyncratic risk puzzle. The idiosyncratic risk puzzle describes the observation that stocks with high idiosyncratic volatility earn higher returns. This observation is puzzling, as the efficient market hypothesis states that idiosyncratic risk is not compensated because it can be diversified. Stambaugh, Yu, and Yuan (2015) explain the idiosyncratic risk puzzle with a combination of arbitrage risk and arbitrage asymmetry. They claim that the type of mispricing that could only be exploited with a short position is persistent, as short selling constraints impede arbitrage. In contrast, mispricing that can be exploited with a long position is not persistent, as arbitrageurs can easily exploit this mispricing type. In this chapter, I challenge Stambaugh, Yu, and Yuan’s (2015) line of arguments

and use pairs trading as an alternative proxy for the extent of mispricing. Analyzing the US stock market between 1990 and 2014, my findings confirm, consistent with Stambaugh, Yu, and Yuan (2015), a negative relation between idiosyncratic volatility and return among overpriced stocks and a positive relation between idiosyncratic volatility and return among underpriced stocks. As pairs trading measures a different type of mispricing compared to Stambaugh, Yu, and Yuan's (2015) applied mispricing index, this out-of-sample analysis tests whether the proposed idiosyncratic risk effect is determined by the type of mispricing. This analysis is a great example how pairs trading can be utilized as an alternative out-of-sample test setting to contribute additional evidence to any mispricing related research question.

Altogether, the findings in this thesis extend the knowledge about pairs trading, especially regarding the use of a rule-based trading process, the role of gradual information diffusion in pairs trading, and the usability of pairs trading in an out-of-sample test setting. Overall, the thesis contributes to the academic discussion about market efficiency, the use of algorithmic trading, and the limits of arbitrage.

2. Methodology

The most established approach to pairs is the multivariate distance approach as first introduced in academic literature by Gatev, Goetzmann, and Rouwenhorst (2006), and also used by Engelberg, Gao, and Jagannathan (2009), Jacobs and Weber (2015), among others. Besides this approach, several streams coexist in the literature that apply different algorithms for selecting pairs or applying modified trading rules. The most prominent approaches besides the distance approach are the cointegration approach (Vidyamurthy (2004), Lin, McCrae, and Gulati (2006)), the copula-based algorithm (Stander, Botha, and Marais (2013)), the stochastic approach (Tourin and Yan (2013)), and mixed forms of the listed approaches. For further reference, Krauss (2015) provides a comprehensive overview of different methods. Hauck and Afawubo (2015) and Huck (2013, 2015) empirically evaluate various approaches. This study closely follows Gatev, Goetzmann, and Rouwenhorst's (2006) original distance approach to implement pairs trading because of its high recognition. Despite its non-trivial coding, the multivariate distance approach of Gatev, Goetzmann, and Rouwenhorst (2006) is convenient for modifying the algorithm so that all requirements to test this thesis's research hypothesis are met.

The following Section 2.1 describes the trading rules in detail. It is important to understand the technical components of the algorithm to later comprehend the influence of the arbitrarily set threshold and decision rule. Afterward, Section 2.2 explains the return calculation.

2.1 The pairs trading algorithm

Following Gatev, Goetzmann, and Rouwenhorst (2006), a pairs trading period in this thesis consists of two periods – first, a twelve-month identification period, which is explained in further detail in Section 2.1.1, followed by a subsequent six-month trading period, which is explained in Section 2.1.2.

2.1.1 The identification period

The main objective of the identification period is to find two stocks that are close economic substitutes or more precisely, two stocks with a similar sensitivity to the arrival of new information. The rigorous pair selection algorithm ensures that pairs are not selected based on the investor's subjective opinion.

At the beginning of the identification period, all stock prices are normalized to one. The normalized price $P_{A,t}$ of stock A on day t with $t = 1, \dots, T$, is defined as follows:

$$P_{A,t} = \frac{CP_{A,t}}{CP_{A,1}}, \quad (1)$$

where $CP_{A,t}$ is the closing price of stock A on day t and $CP_{A,1}$ is the closing price of stock A on the first day of the identification period.

At the end of the twelve-month identification period, I compute four key figures for each possible pair combination with stock A and stock B in the stock universe:

1. Selection measure SSD: The traditional selection measure SSD (sum of squared differences) to identify close economic substitutes, as introduced by Gatev, Goetzmann, and Rouwenhorst (2006), is the sum of the squared deviations between the normalized price movements of stock $P_{A,t}$ and $P_{B,t}$ during the identification period:

$$\text{Distance Measure (SSD)} = \sum_{t=1}^T (P_{A,t} - P_{B,t})^2, \quad (2)$$

where $P_{A,t}$ and $P_{B,t}$ are the normalized prices of stock A and B (see formula (1)) on any day t during the identification period. The SSD is small if the price spreads $P_{A,t} - P_{B,t}$ is small, if both stocks move very similar, and so is the SSD.

2. Trading signal $2\sigma_{\text{hist}}$: Following Gatev, Goetzmann, and Rouwenhorst (2006), I determine daily

the spread between the two normalized prices $P_{A,t}$ and $P_{B,t}$ and calculate the standard deviation of this spread during the identification period. The trading signal in the subsequent trading period is defined as two times this particular standard deviation.

At the end of the identification period, all pairs are ranked in ascending order according to their selection measure SSD. Gatev, Goetzmann, and Rouwenhorst (2006) form a trading portfolio that consists of the twenty pairs with the lowest SSD. The price spread between stock A and B of a pair with a low SSD is small. Hence, pairs trading assumes that both stocks shared a similar price development in the past. Figure 2 illustrates the simplified pair selection process with the SSD selection measure.

Figure 2: Pair selection procedure with the SSD selection measure

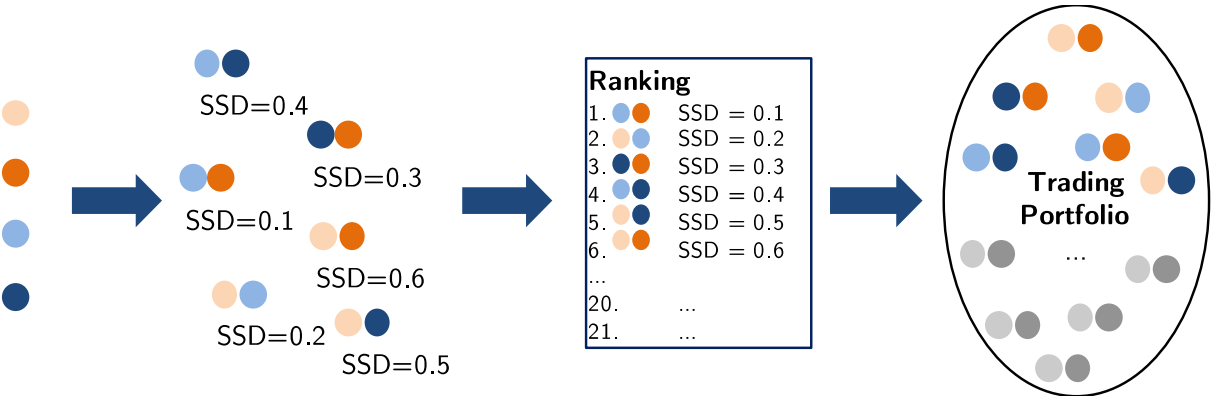


Figure 2 illustrates the process of calculating the SSD for each possible stock combination from the stock universe. Subsequently, the pair combinations are ranked according to their SSD. Finally, the best pairs are included in the trading portfolio on the basis of their ranking position.

We later see in chapter 4 that the SSD selection measure introduces a selection bias regarding the volatility and correlation of stocks. Therefore, I introduce a modified selection process to incorporate a different set of pairs in Section 4.4. For this purpose, I also determine two additional figures at the end of the twelve-month identification period for each possible pair combination

with stock A and B, the total pair volatility σ_{AB} and the pair correlation $\rho_{A,B}$.

3. Total pair volatility σ_{AB} : The total pair volatility is calculated as follows:

$$\sigma_{AB} = \sigma_A + \sigma_B, \quad (3)$$

where σ_A and σ_B are the standard deviations of the normalized stock prices $P_{A,t}$ and $P_{B,t}$ of stock A and B.

One obvious question is why not use the standard deviation of the price spread between $P_{A,t}$ and $P_{B,t}$? One of the following analyses examines the influence of pair volatility and correlation on the return separately. The measure for pair volatility and correlation must therefore be independent of each other. The spread between both prices over time is unfortunately not a suitable measure, as it is determined by the volatility as well as the correlation by both stocks. Therefore, I define total pair volatility and classify pairs accordingly.

4. Pair correlation $\rho_{A,B}$: The pair correlation is based on the simple correlation between the normalized stock prices $P_{A,t}$ and $P_{B,t}$.

2.1.2 The trading period

To avoid an in-sample bias, the trading period starts on the first day after the identification period. Again, all prices are normalized to one according to formula (1). The strategy later sells \$1 of stock² A or B and buys \$1 of the other stock with the short sale profit. The strategy is therefore self-financing.

² I apply the common assumption in financial research that shares can be bought and purchased in arbitrarily amounts. The assumption underlies the idea that especially institutional investors trade with high amounts that can be scaled accordingly.

At the end of each day during the trading period, the trading algorithm checks the status of each pair in the trading portfolio. A new trade is opened if:

$$|P_{A,t} - P_{B,t}| \geq 2\sigma_{\text{hist}}, \quad (4)$$

where $|P_{A,t} - P_{B,t}|$ is the spread between the two normalized prices and $2\sigma_{\text{hist}}$ represents the trading signal that was determined during the identification period as introduced in the previous section. I open a new self-financing trade as follows: I sell \$1 of the better performing “winner” stock and simultaneously buy \$1 of the less successful “loser” stock.

Furthermore, I check each open pair as to whether the two normalized stock prices completely converged during the day. If the answer is yes, the long and the short position are neutralized.

In general, a pair might open and close several times during the trading period. If prices do not fully converge until the end of the trading period, all open positions are closed. Furthermore, trades are also closed if the price spread exceeds $4\sigma_{\text{hist}}$ (overshooting trade), or one stock is delisted on any day while the trade is active (delisted trade). A new pairs trading period starts every month. At that time, the second portfolio has already been active for two months, the third one for three months and so on. Because of the overlapping trading periods, six portfolios are traded simultaneously at each point in time.

Gatev, Goetzmann, and Rouwenhorst (2006) and follow-up papers implement a one-day-waiting strategy to consider a potential upward bias in returns caused by the bid-ask bounce. In an upward trend, the closing price most likely represents the ask price. As pairs trading sells the increasing “winner stock,” an investor receives only the lower bid price. Likewise, the closing price of a decreasing stock is more likely to be the bid price, which must be purchased at the higher ask price. Hence, using closing prices might overestimate the real return. To overcome this problem, the one-day-waiting strategy waits for one day after the initial trading signal before a trade is executed.

I choose to avoid the one-day-waiting strategy as a proxy and use the exact last bid or last asking price to overcome this problem. At trade initialization, I buy the long position at the ask price and sell the short position at the bid price. I neutralize the long position at the bid price and the short position at the ask price. Similarly, I use these latter prices for the daily evaluation. Additional computations, like the average volatility or correlation, are executed with closing prices.

2.2 Return calculation

I closely follow the return calculation of Gatev, Goetzmann, and Rouwenhorst (2006) and Engelberg, Gao, and Jagannathan (2009) and differentiate two types of return: the return per trade and the monthly return, which I will briefly explain in the next two sections.

2.2.1 Return per trade calculation

The return per trade is defined as the absolute payoff of a pair's round trip, expressed as a monetary amount. This return is positive if the stock price of the short position decreases and/or the stock price of the long position increases.

For instance, the historical standard deviation of a pair's price spread is $\sigma_{\text{hist}} = 0.05$. Thus, I open the pair as soon as the absolute difference between the normalized price of stock A and the normalized price of stock B at time t equals or exceeds 0.1 (see formula 4: $|P_{A,t} - P_{B,t}| \geq 2\sigma_{\text{hist}}$).

For instance, $P_{A,1} = \$1$ and $P_{B,1} = \$1$ in $t = 1$, and $P_{A,2} = \$1.05$ and $P_{B,2} = \$0.95$ in $t = 2$. As the price spread of \$0.1 equals the trading signal $2\sigma_{\text{hist}}$, I open the trade in $t = 2$. I initiate a \$1 short position in the "winner" stock A with a capital commitment of $C_{A,2} = \$(-1)$ and a long position in the "loser" stock B with a capital commitment of $C_{B,2} = +\$1$ in $t = 2$. In $t = 3$, the normalized prices of stocks A and B fully converge into one: $P_{A,3} = \$1$ (decrease of -4.76%) and $P_{B,3} = \$1$ (increase of 5.26%). The capital commitments are: $C_{A,3} = -\$1 * (1 + (-0.0476)) = -\0.9524 , and $C_{B,3} = +\$1 * (1 +$

0.0526) = \$1.0526. As both normalized prices of stock A and B fully converge, I close the trade and receive the following payoff: $(C_{A,3} - C_{A,2}) + (C_{B,3} - C_{B,2}) = \$0.05 + \$0.05 = \0.10 . The return of \$0.10 represents the return per trade. Figure 3 illustrates the price movements of stock A and stock B (in black), as well as the corresponding pairs' trading positions ("PT" in green).

Figure 3: Calculation of the return per trade

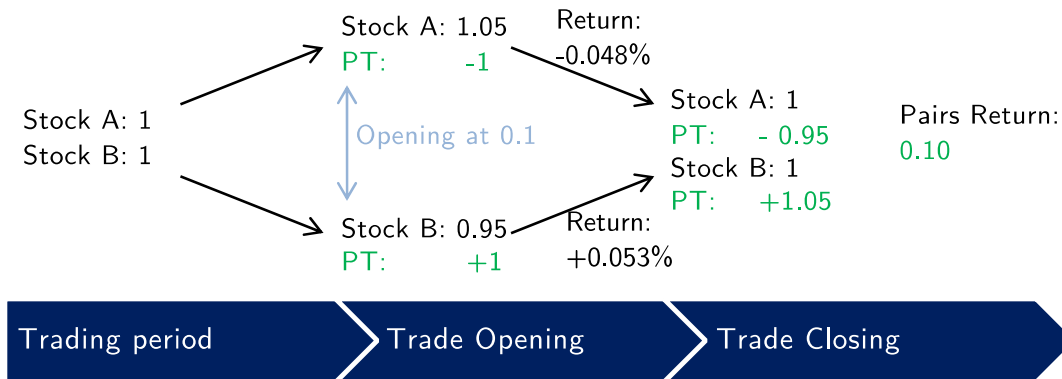


Figure 3 displays the stock price development of stock A and stock B during the trading period. 'PT' denotes to the absolute amount of \$ that a pairs trading invests in stock A and stock B at the time when the trade is opened and the amount he receives at the time when the trade is closed.

2.2.2 Calendar Time Return calculation

The return in calendar time is defined as a monthly return. Similar to Gatev, Goetzmann, and Rouwenhorst (2006), I calculate the monthly return as follows: all positions are marked-to-market daily. The return $r_{i,t}$ is the return that stock i realizes between day t-1 and t. The stock weight $w_{i,t}$ of stock i at day t can be interpreted as capital investment on day t-1 or as a buy-and-hold strategy that reinvests daily returns until day t-1. The stock weight $w_{i,1}$ on the day after the trade initialization is set as follows:

$$w_{i,1} = I = \begin{cases} +\$1 & \text{for a long position} \\ -\$1 & \text{for a short position} \\ 0 & \text{if the pair is closed} \end{cases} \quad (5)$$

On the following days $t > 1$, the stock weight $w_{i,t}$ is calculated as the product of the previous days' capital investment:

$$w_{i,t} = I * w_{i,t-1} * (1 + r_{i,t-1}) = I * \prod_{t=1}^{t-1} (1 + r_{i,t-1}) \quad (6)$$

The pairs trader's payoff of stock i on day t equals the weight (or capital investment in $t-1$) $w_{i,t}$ multiplied with the stock return $r_{i,t}$ between $t-1$ and t :

$$\text{Payoff}_{i,t} = w_{i,t} * r_{i,t} \quad (7)$$

For instance, the stock commitment in stock A in $t-1$ is +\$1.25 (long position), and stock A increases from \$20 to \$21 (5% increase). Hence, the stock commitment in A in t is +\$1.25 * (1 + 0.05) = \$1.3125. The payoff is: \$1.3125 - \$1.25 = \$1.25 * 0.05 = \$0.0625. Likewise, another example for a short position: The stock commitment in stock B in $t-1$ is -\$2 (short position) and stock B decreases from \$40 to \$37.60 (6% decrease). Hence, the stock commitment in B in t is: -\$2*(1+(-0.06))= -\$1.88. The payoff is: -\$1.88-(-\$2)= -\$2*-\$0.06=\$0.12.

The daily return of a pair P , which includes the two stocks i , is:

$$r_{P,t} = \frac{\sum_{i \in P} w_{i,t} * r_{i,t}}{\sum_{i \in P} |w_{i,t}|} = \frac{\sum_{i \in P} \text{payoff}_{i,t}}{\sum_{i \in P} |w_{i,t}|} \quad (8)$$

The daily pair return can be understood as the sum of the daily payoffs divided by the total pair capital commitment.

The pair return of our example pair P , including stock A and B, equals: $r_{P,t} =$

$$\frac{\$1.25 * 0.05 + (-\$2 * -0.06)}{|\$1.25| + |-\$2|} = \frac{\$0.0625 + \$0.12}{|\$1.25| + |-\$2|} = \frac{\$0.1812}{\$3.25} = 0.0562 \text{ at time } t.$$

The daily returns are cumulated to monthly returns afterward. Following Gatev, Goetzmann, and Rouwenhorst (2006) and Do and Faff (2010), I report the return on "committed capital." The term "committed capital" might be confusing, as pairs trading is a self-financing strategy with a \$1 long

and a \$1 short position. In order to calculate a return, Gatev, Goetzmann, and Rouwenhorst (2006) scale the portfolio's payoff by the number of actively traded pairs. The authors argue that computing the excess return relative to the actual capital employed considers the opportunity costs for committing capital to pairs trading if the strategy does not trade. As hedge funds are typically free in their decisions about the source and the use their funds, Gatev, Goetzmann, and Rouwenhorst (2006) argue that the measure "committed capital" provides a more realistic measure to compare different investment strategies. Do/Faff (2010) follow their approach.

The previously outlined procedure implicates that at any time six portfolios, each starting lagged by one month, are traded simultaneously. Trades are more likely in the last months in the trading period. All prices were normalized at the first day of the trading period. A significant price divergence of over $2\sigma_{\text{hist}}$ and the trade triggered thereby is more likely at the end of the trading period when incremental price divergences cumulate. To avoid this bias in the return calculation, I average the monthly return of the six simultaneously traded portfolios. Figure 4 illustrates this process.

Figure 4: Calculation of pairs trading returns in calendar time




Figure 4 illustrates the method to calculate the average return of the six simultaneously traded portfolios during the trading period.

Gatev, Goetzmann, and Rouwenhorst (2006) interpret the return of pairs trading as a payoff to a proprietary trading desk where different traders manage six pairs trading portfolios whose identification and trading periods are each staggered by one month. Alternatively, the return can be interpreted as the average return of an active pair across all open pairs within the same portfolio and across six portfolios that trade simultaneously, but were started in consecutive months. I refer to Gatev, Goetzmann, and Rouwenhorst (2006) for further details on the return calculation.

The calendar time return might differ considerably from the return per trade for three reasons. First, the monthly return reflects the combined effect of trading frequency and return per trade. Second, the monthly return is scaled to the number of included active pairs in the portfolio. Third, a trade

might be open for several months. Therefore, the pair earns only a fraction of the return per trade within one month. The interested practitioner can find more information about the return calculation in Vidyamurthy's (2004) implementation book.

3. Baseline results

3.1 Data selection

I obtained the historical index constituents of the S&P 1500 between January 1990 and December 2014 from the Compustat Monthly Updates - Index Constituents file. Furthermore, I drew daily stock price data from the WRDS CCM merged database. I exclude stocks for a particular pairs trading cycle if any of the three – bid, ask or closing price – is unavailable on more than one day during the identification period. After this preselection, approximately 1,900 stocks on average are eligible for trading each month, forming around 1.8 million possible pair combinations.

In contrast to related papers, I keep stocks in my sample that are not traded on one day during the identification period. I argue that trading is theoretically possible even if a stock is not traded on one particular day during the identification period if bid and ask quotes are available. At trade initialization, I buy the long position at the ask price and sell the short position at the bid price. I neutralize the long position at the bid price and the short position at the ask price. Similarly, I use these latter prices for the daily evaluation. Additional computations, like the average volatility or correlation, are based on closing prices. This modified approach allows me to refrain from approximating the bid-ask bounce and trade with the exact bid and ask prices, and instead to work with higher accuracy. Even more important, the modified approach allows me to include less liquid, highly volatile stocks that were not traded on one day during the identification period, although it was theoretically possible. The inclusion of these less liquid, but highly volatile stocks is crucial to evaluating my research hypotheses, as the exclusion of such stocks results in a biased sample.

The risk factors for the commonly used five factor model are obtained from Kenneth French's and Lubos Pastor's homepage. Further data is hand-collected from the homepages of the respective authors.

3.2 Coding the pairs trading algorithm

There are numerous tools on the market that allow institutional investors to implement pairs trading. Analyzing the outlined research questions demands, however, a high flexibility of the algorithm and the possibility of modifying some trading rules. Therefore, I coded the pairs trading algorithm myself.

The complex coding is based on a simplified pairs trading program that I coded during my master's thesis. The major part of the program is coded in Stata, as Stata is one of the most commonly used statistics programs in the field of financial research. Unfortunately, Stata does not meet the requirements to process the extensive amount of data. In particular, the limitation of elements in a matrix is an unfeasible problem in Stata. To overcome this problem, I switched to Mata, a matrix language that is integrated into Stata and similar to Java and C within the code. The additional effort of data transfer between both programming languages is compensated by the fast data processing in Mata.

The first part of the algorithm imports three data sets that include bid prices, ask prices, and closing prices for each stock in the trading universe, and each trading day between January 1990 and December 2014. The stock universe also includes stocks with an incomplete price history that went public or were delisted during this time. In the first step, the bid, the ask, and the closing price data set must be normalized to 1 at the beginning of the identification phase. In the next step, the algorithm computes the SSD, the volatilities of both stocks σ_A and σ_B and the pair correlation $\rho_{A,B}$ in a loop over all days of the identification period. This loop is repeated for each of the approximately 1.8 million possible pair combinations. Afterward, the algorithm ranks all pair combinations according to the SSD in ascending order and identifies the twenty pairs with the lowest SSD. In addition, the algorithm forms twenty-five portfolios in a modified test setting, which will be explained

in chapter 4.4.

To test the research hypotheses introduced later adequately, the algorithm repeats these steps in a rolling window for each identification period that starts with a one-month time lag. In the end, the data in the matrix includes the following key figures: individual stock volatility σ_A and σ_B , pair volatility σ_{AB} , pair correlation $\rho_{A,B}$, and the SSD for each possible pair combination and for each identification period. This part of the computation is the most time-consuming computation part, which took up to six weeks on a standard computer in 2015.

In the second part of the pairs trading program, the algorithm imports the three price data sets (bid, ask, and closing price) for trading and the previously introduced data matrix with the key figures of the first algorithm part. Again, the algorithm normalizes all prices of the previously selected stocks to 1 on the first day of the trading period. Afterward, the algorithm checks the price differences in a loop on each day of the trading period for each selected pair and trades as described in section 3.1.2. In particular, the identification of the winner/loser stock and the consequentially correct use of the bid price for the winner stock or the ask price for the loser stock make this step complex. The trading program is highly flexible: it can trade with bid/ask prices or the closing price, it can trade at the day of the trading signal or with a one-day-delay to address the bid-ask bounce, it can deal with stocks that were delisted during the trading period, and it includes several more options for modified trading.

For each trading period, my algorithm computes the cumulative return of all pairs included in the trading portfolio for each of the six months within the trading period. Furthermore, it stores the return of each trade together with several key figures, for instance volatility and pair correlation during the trading period. As in the identification period, this procedure is repeated with a rolling

window over all trading periods. Thereafter, my algorithm uses several loops to perform the complex return computation to average the returns over all open pairs and over all six simultaneously traded portfolios within one month, as illustrated in Figure 4.

At the end of both parts, the algorithm generated a data set that includes monthly pairs trading returns and a trade data set that includes the returns of each trade. To shed light on my research hypotheses, I merge both datasets with information about the stocks and pair specific characteristics. These two datasets form the foundation for all of the following analyses. All analyses, regressions, and further statistical investigations in this thesis are executed in Stata using the monthly return dataset or the trade dataset.

3.3 Empirical results

To evaluate the quality of my computation, I benchmark my results with pairs trading studies from recognized journals. I exclude papers that apply major modifications, conceal raw returns, or one-day-waiting returns. Do and Faff (2010) detect a declining trend in pairs trading returns over time. Therefore, I also consider the covered period when comparing returns. Furthermore, minor adaptations in the original algorithm might cause small differences, like the preselection of pairs as in Papadakis and Wysocki (2007). Table 1 summarizes the monthly US returns for similar stock universes.

Table 1: Comparison with related studies

Table 1 compares the monthly excess returns of my portfolio selected based on the SSD with the excess returns of related papers. The column Period shows the period of time investigated in the respective study.

Paper	Period	Excess Return
This study	1990-2014	0.0037
Gatev, Goetzmann, and Rouwenhorst (2006)	07/1963-12/2002	0.009
Do and Faff (2010)	1962–1988	0.0086
	1989–2002	0.0037
	2003–2009	0.0024
Do and Faff (2012)	1963-2009	0.0085
Papadakis and Wysocki (2007)	1981-2006	0.0062
	1981-1993	0.0087
	1994-2006	0.0036

My return of 37 bp is consistent with the 37 bp return in Do and Faff (2010) for 1989-2002 and 24 bp for 2003-2009, as well as the 36 bp in Papadakis and Wysocki (2007) for 1994-2006. These results indicate that my computation and the bid-ask price approach is reliable.

4. Core chapter 1: Demystifying pairs trading: The role of volatility and correlation

4.1 Motivation

Numerous papers find that pairs trading returns are higher in markets with high volatility and strong correlation. Varying market conditions and factors across individual stocks, industries, and countries determine the levels of volatility and correlation among stocks that affect the profitability of pairs trading. For instance, Gatev, Goetzmann, and Rouwenhorst (2006) find that certain industry portfolios outperform others. Do and Faff (2010) discover a superior return in times of market distress between January 2000-December 2002 and July 2007-June 2009. Jacobs and Weber (2015) observe varying returns across different countries. These variations over time, across markets, and industries are often explained with different levels of market efficiency by Do and Faff (2010), and the number of eligible pairs and limited investor's attention by Jacobs and Weber (2015). Less efficient markets or industries and investor's inattention allow more mispricing opportunities, which in turn increase the strategy's return.

However, until now, little is known about the channel that translates these factors into higher pairs trading returns. I extend the current literature about pairs trading by explaining the concrete mechanism behind the observation. My insights help to better understand the influence of external factors on the profitability of pairs trading. Furthermore, my results encourage a critical evaluation of the chosen procedure to select pairs for trading based on the SSD.

Pairs trading is a technical, rule-based investment strategy. As such, its return is a product of the interaction between different input parameters and more or less arbitrarily set rules and thresholds. Figure 5 disentangles the technical interdependencies of the rule-based pair selection (SSD measure) and the trading algorithm (2σ -trading rule) with the strategy's return generating components, return per trade and trading frequency.

Figure 5: Interdependences of the technical decision rules and the return of pairs trading

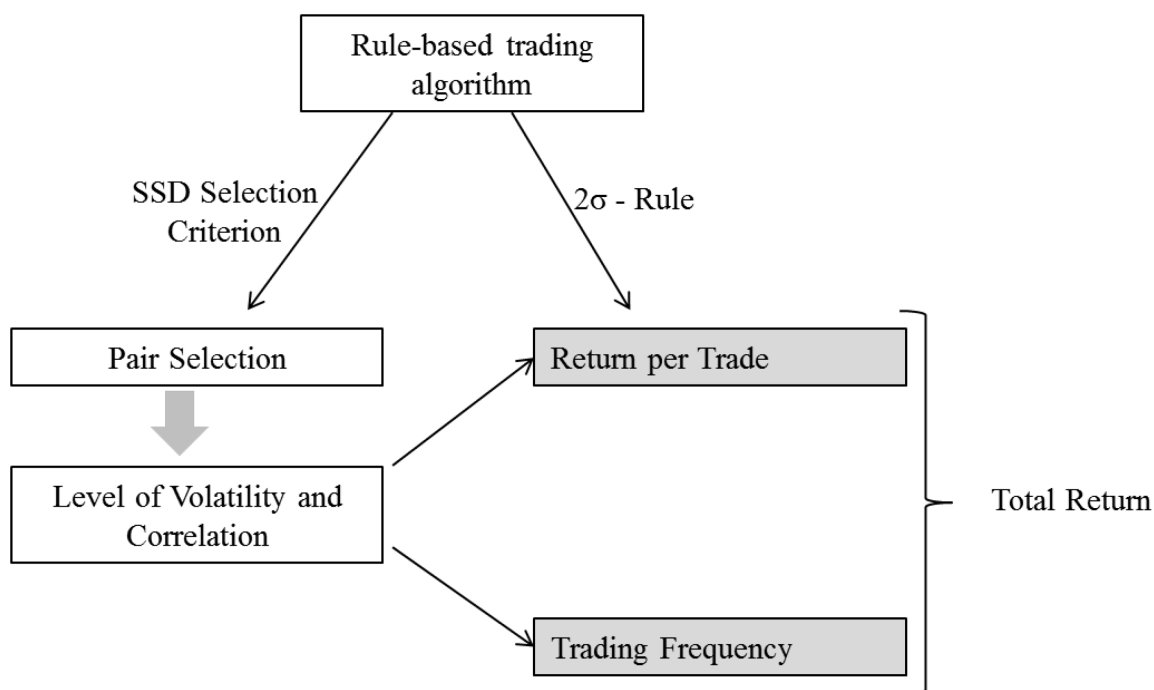


Figure 5 illustrates the interdependencies of pairs trading. The rule-based trading algorithm influences the pair selection and the maximum return per trade. The pair selection determines the correlation and the pair volatility level. These levels, in turn, influence the return per trade and the trading frequency. The strategy's total return is a function of the return per trade and the trading frequency.

My analysis of the algorithm, how volatility and correlation translate into higher or lower pairs trading returns, complements the literature on pairs trading by bridging the gap between previous empirical findings and the mechanism of the technical trading algorithm. Furthermore, the insights provide a solid foundation to critically evaluate the influence of pairs trading's arbitrarily set trading rules.

The remaining parts of the chapter are organized as follows: Section 4.2 summarizes related research studies. In Section 4.3, I derive three research hypotheses about the technical drivers of pairs trading. To adequately evaluate these research hypotheses, Section 4.4 introduces a modified test setting.

Afterward, Section 4.5 tests the proposed hypotheses, before Section 4.6 addresses robustness issues. Finally, Section 8 concludes the chapter.

4.2 Selected pairs trading literature

My research is motivated by previous papers that indirectly examine the influence of pair volatility and correlation. In this section, I summarize the most important related findings.

First, out of fifty pairs with the lowest distance measure (SSD), Do and Faff (2010) select the twenty pairs with the highest number of price crossings during the identification period and observe a higher profitability for these pairs. The authors argue that multiple divergences and convergences of a pair within the identification period predict more trading opportunities. I therefore analyze which level of pair volatility and correlation maximize the trading frequency. Moreover, Do and Faff (2010) find an increased profitability during times of market distress from January 2000-December 2002 and July 2007-June 2009. These two bear market phases were accompanied by high market volatility. Hence, I suspect that high volatility increases the return.

Second, the success of pairs trading was explored in many countries, for instance Andrade, di Pietro, and Seasholes (2005), Perlin (2009) and Bolgün, Kurun, and Güven (2010), among others. Most notably, Jacobs and Weber (2015) provide a broad overview for 34 countries worldwide. They observe that less developed markets usually realize a higher return. These markets are normally less liquid, and hence possess higher market volatility.

Third, Engelberg, Gao, and Jagannathan (2009) provide a number of in-depth analyses about the drivers of opening, horizon, and divergence risks. Among other findings, they demonstrate the positive influence of idiosyncratic volatility on the opening probability, and a negative impact of idiosyncratic volatility on the time until convergence and also on the divergence probability. Similarly, Jacobs and Weber (2015) examine the influence of idiosyncratic volatility as a proxy for limits to

arbitrage. They find a positive influence of volatility on the total return. Finally, Huck (2015) investigates the effect of total market volatility over time. He initializes positions if the VIX is categorized into a certain regime, also to the traditional opening signal. However, the return is only significant at times of an increasing or high 3-month moving average VIX. These exemplary findings indicate that volatility and correlation are significant drivers for pairs of trading returns.

4.3 Research Hypothesis - Disentangling the pairs trading algorithm

Based on the previously introduced findings in the literature, I derive three research hypotheses to identify the channel that translates the levels of volatility and correlation into higher or lower pairs trading returns. In Section 4.3.1, I derive why I expect that pairs, which are picked by the traditional selection criterion SSD, include stocks with low pair volatility and high correlation levels. These levels, in turn, influence the return per trade and the trading frequency. I formulate two research hypotheses on how the two drivers, pair volatility and correlation, influence the return per trade in Section 4.3.2 and the trading frequency in Section 4.3.3.

4.3.1 The SSD selection criterion

I now take a closer look on the distance measure (SSD), and its implicit selection of pairs with low pair volatility and high correlation. A broad stream of literature applies the SSD to identify close economic substitutes during the identification period. These pairs are eligible for trading in the subsequent trading period. The SSD computes the sum of squared differences of the normalized prices $P_{i,t}$ of two stocks $i = A$, and B at time t , and selects the twenty pairs with the lowest SSD:

$$\text{Distance Measure (SSD)} = \sum_{t=1}^T (P_{A,t} - P_{B,t})^2 = \tag{9}$$

$$= \sum_{t=1}^T (P_{A,t}^2 + P_{B,t}^2 - 2\rho_{A,B}), \tag{10}$$

where $\rho_{A,B}$ is the Pearson correlation coefficient of the standardized price time series. The Pearson correlation coefficient equals:

$$\rho_{A,B} = \frac{1}{t-1} \sum_{t=1}^T \left(\frac{P_{A,t} - \bar{P}_A}{s_{P_A}} \right) \left(\frac{P_{B,t} - \bar{P}_B}{s_{P_B}} \right), \quad (11)$$

where $P_{i,t} = \frac{P_{i,t} - \bar{P}_i}{s_{P_i}}$ are the daily z-transformed normalized prices.

From the formula in (10), I derive that the correlation coefficient is negatively correlated with the SSD. The pairs trading algorithm selects the pairs with the lowest SSD. Based on this, I conclude that pairs trading that selects pairs based on the SSD trades with highly correlated pairs.

Furthermore, I argue that low stock volatility is also crucial to minimize the SSD. Formula (10) reveals that high stock prices for $P_{A,t}$ or $P_{B,t}$ increase the SSD. Both prices are normalized to one in $t = 0$. Thus, high values for $P_{A,t}$ and $P_{B,t}$ only arise if stock volatility is high.

Whether the minimizing influence of strong correlation with the SSD or the increasing effect of high stock volatility dominates depends on the magnitude of both factors. Formula (9) provides a good clue: let's consider three stocks A, B, and C, whose prices P_A , P_B , and P_C are all normalized to one in $t = 0$. The returns of stock A and C are uncorrelated ($\rho_{A,C} = 0$), and the returns of both stocks are independently and normally distributed with the same standard deviation σ . Hence, the confidence intervals for the normally distributed price realizations P_A and P_C are identical with:

$\left[1 - z_{1-\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}}, 1 + z_{1-\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}} \right]$. The return of stock B is perfectly correlated with the return of stock A ($\rho_{A,B} = 1$). However, a Δ change in the return of stock A translates into a 4Δ change of stock B's return. For instance, if P_A increases from \$1 to \$1.01, P_B increases from \$1 to \$1.04. As a consequence, stock B's standard deviation is higher, and so is the confidence interval of P_B . The daily price difference between P_A and P_B might therefore be higher than the price difference between P_A

and P_C , although $\rho_{A,B} = 1$ and $\rho_{A,C} = 0$. For example, in $t = 1$, stock A increases from $P_{A,0} = 1$ to $P_{A,1} = 1.01$. Hence, stock B increases from $P_{B,0} = 1$ to $P_{B,1} = 1.04$. The price difference between stock A and B is: $|P_{A,1} - P_{B,1}| = 0.04$. Stock C is uncorrelated with stock A and does not move at all, $P_{C,0} = P_{C,1} = 1$. The price difference between stock A and C is therefore: $|P_{A,1} - P_{C,1}| = 0.01$. If we cumulate all daily squared price differences as in formula (9), I come to the following conclusion: the SSD for two uncorrelated stocks (A and C) can be smaller than the SSD of two highly correlated stocks (A and B) if the highly correlated pair includes at least one highly volatile stock.

Altogether, I expect a negative relationship between the SSD and correlation, and a positive one between SSD and pair volatility.

Research proposition 1: Pairs with high correlation and low pair volatility are associated with a low SSD.

A validation of research proposition one implies that applying the SSD results in trading with highly correlated pairs that exhibit little volatility. However, it is not clear whether this pair volatility and correlation combination is beneficial for the return per trade and the trading frequency.

4.3.2 Return per trade

This section takes a closer look at the influence of pair volatility and correlation on the return per trade. During the trading period, I initialize positions as soon as the stock pair's price difference exceeds two historic standard deviations $2\sigma_{\text{hist}}$ as calculated during the identification period. The return is continuously earned while the pair is "open" until stock prices fully converge. At that point, I neutralize positions. I distinguish four different closing types.

- First, trades that fully converge during the trading period are denoted as "natural trades" (closing type 1).
- Second, pairs that are not converged until the last day of the trading period are forcefully

closed and categorized as “incomplete trades” (closing type 2).

- Third, as I trade with highly volatile stocks, investors might be concerned about the asymmetric return profile of pairs trading. The positive return per trade is limited to 2σ , while pairs that further diverge might generate an infinite loss. Therefore, I close a pair if the price difference exceeds $4\sigma_{\text{hist}}$ on any day while a trade is open. This procedure is similar to Engelberg, Gao, and Jagannathan’s (2009) 10-day maximum strategy to close open stock positions if they do not converge within ten days after the initial trade opening. The return potential is symmetrical within the range of $-2\sigma_{\text{hist}}$ and $+2\sigma_{\text{hist}}$. These trades are categorized as “overshooting trades” (closing type 3). Affected pairs are blocked for the rest of the trading period. Otherwise, a new trade would open on the next day, as the price spread is still above $2\sigma_{\text{hist}}$ because prices are not set back to one during the trading period. I also assume that a price spread of $4\sigma_{\text{hist}}$ indicates a permanent price spread and thus refrain from further trading with these pairs.
- Fourth, “delisted trades” are trades that automatically close if one stock is delisted (closing type 4).

I now take a closer look on the return potential of these four trade closing types:

- Closing type 1 (natural trades): I conjecture that a pair always earns a positive return if prices fully converge (natural trade). A possible unprofitable price development of one stock is always overcompensated by the return of the other stock that overcomes the previous price difference. For instance, if the stock price of the short position is rising, the return of the long position will overcompensate for the negative return of the short position, as stock prices would otherwise not fully converge.

The return per trade equals the stock price difference at the time of the opening. Positions

are initialized as soon as prices diverge by more than $2\sigma_{\text{hist}}$. Thus, the return of a natural trade equals $2\sigma_{\text{hist}}$ plus an overshooting component. The overshooting component occurs if the price difference exceeds the $2\sigma_{\text{hist}}$ threshold during the day and prices further diverge until closing prices are set. I derive:

$$\text{Return per natural trade} = 2\sigma_{\text{hist}} + \text{overshooting.} \quad (12)$$

Accordingly, the return per natural trade is higher if the historical standard deviation $2\sigma_{\text{hist}}$ is higher.

Pairs trading can be understood as a portfolio that consists of stock A (long) and stock B (short). The pair's historic standard deviation σ_{hist} can easily be computed in this case. A simple mathematic conversion of formula (12) reveals the relationship between trade return, volatility of stock A and B, and the pair's correlation coefficient ρ_{AB} :

$$\text{Return per natural trade} = 2\sqrt{\sigma_A^2 + \sigma_B^2 - 2\rho_{AB}\sigma_A\sigma_B} + \text{overshooting,} \quad (13)$$

The formula validates my conjecture that the return of a successful trade is always positive. The overshooting component is always positive as otherwise, the total price difference would not exceed $2\sigma_{\text{hist}}$. If the price difference is smaller than $2\sigma_{\text{hist}}$, the trade would not be opened. It applies:

$$(\sigma_A - \sigma_B)^2 \geq 0. \quad (14)$$

Hence, the influence of volatility on the return per trade is always positive or zero, as formula (14) can be transformed as follows:

$$\sigma_A^2 + \sigma_B^2 \geq 2\rho_{AB}\sigma_A\sigma_B, \quad (15)$$

with $\rho_{AB} \in [-1, 1]$. Based upon (15), I argue that a higher volatility level of stock A or stock B

(or both) increases the historic standard deviation, and thus the return per trade. On the contrary, higher pair correlation decreases $2\sigma_{\text{hist}}$. Neglecting the overshooting component, the first differentiation of the function regarding the historical correlation coefficient equals:

$$\frac{\partial f(\sigma_A, \sigma_B, \rho_{AB})}{\partial \rho_{AB}} = - \frac{2\sigma_A\sigma_B}{\sqrt{\sigma_A^2 + \sigma_B^2 - 2\rho_{AB}\sigma_A\sigma_B}}, \quad (16)$$

The first differentiation is always negative, as the numerator and the denominator (as shown in (15)) are positive. The return should therefore increase with declining historical correlation. I conclude that more volatile stocks and a negative or low stock correlation increase the return per trade.

- Closing type 3 (overshooting trades): In contrast, overshooting trades, which are closed if the price spread exceeds $4\sigma_{\text{hist}}$, always generates a negative return. Recall that positions are initialized at a price spread of $2\sigma_{\text{hist}}$ and are closed at $4\sigma_{\text{hist}}$. The potential loss is limited to $2\sigma_{\text{hist}}$, and therefore symmetrical to the return potential of natural trades. The effect of pair volatility and correlation with overshooting trades is the reverse of natural trades. High volatility and low correlation increase the risk of exceeding the price difference of $4\sigma_{\text{hist}}$. Hence, low pair volatility and high correlation reduce the downside risk.
- Closing type 2 (incomplete trades) and closing type 4 (delisted trades): These trades neither fully converge nor diverge by more than $4\sigma_{\text{hist}}$. These trades earn a positive return if prices converge, and generate a loss if prices diverge. In conclusion, the link depends upon the direction of the stock's movement.

Altogether, I summarize these propositions in the next research hypothesis:

Research Hypothesis 2: Highly volatile pairs and negatively correlated pairs increase the return per trade of converging trades. In contrast, low volatility and high correlation reduce the downside

risk of diverging trades.

4.3.3 Trading frequency

I now turn to the second return dimension, the trading frequency, and argue that a high trading frequency is beneficial for the strategy's total return. The trading frequency is defined as the number of trades of all pairs within a portfolio over one trading period. Obviously, more trades with positive returns are beneficial, whereas many trades with a negative return are negative for the overall return. A pair can generate several natural trades within one trading period. Earning the historic price spread several times during one trading period successively increases the total return. In contrast, incomplete trades, overshooting trades, and delisted trades are always the last trade within a trading period. A pair cannot open after an incomplete trade, as the incomplete trade is closed on the last day of the trading period. Pairs are blocked for further trades after an overshooting trade, as pairs are otherwise opened and closed daily, while the price spread exceeds $4\sigma_{\text{hist}}$. Likewise, a pair can obviously not open after the delisting of one included stock. As possibly unprofitable trades are limited to one, increasing the trading frequency only increases the number of natural trades – the closing trade type that always yields a positive return.

The trading frequency is determined by the probability of opening a pair and the time until convergence. The probability to generate loss-making overshooting trades is determined by the probability of a further price divergence while the pair is open. Engelberg, Gao, and Jagannathan (2009) explore the influence of various variables on the former probabilities including average mean idiosyncratic volatility. The authors observe a positive influence of idiosyncratic volatility on the opening probability, a negative influence of idiosyncratic volatility on the time until convergence, and on the probability of further price divergence.

In this thesis, I concentrate on the probability of a pair opening and argue that a pair volatility increase or a correlation decrease between the identification and the trading period increases the opening probability. Let X_t be a normally distributed random variable, which describes the price spread that is the difference between the normalized prices of stock A and B on day t . This is actually a simplified assumption to demonstrate the idea. X_t is rather a function of X_{t-1} , the return of stock A, and the return of stock B.

Figure 6 displays the density function of X_t during the identification period with $X_t \sim n, \mu, \sigma_{\text{hist}}$: Let the expected value of price spread X_t be μ . The distribution of X_t is symmetrical, as it is equally likely that stock A outperforms stock B as it is that stock B outperforms stock A. Following the traditional pairs trading algorithm, the threshold to open a trade is $2\sigma_{\text{hist}}$. More precisely, a pair is opened if X_t falls either below $\mu - 2\sigma_{\text{hist}}$ or exceeds $\mu + 2\sigma_{\text{hist}}$.

Figure 6: Distribution of price spread X_t during the identification and the trading period

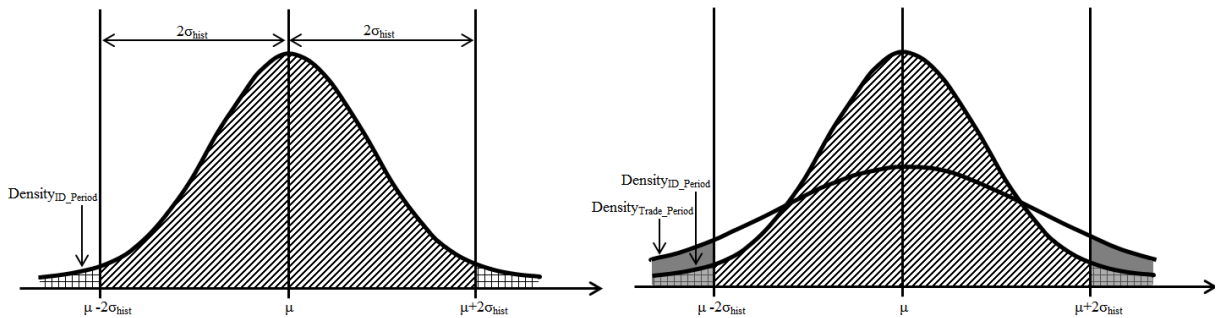


Figure 6 shows the probability density function of X_t during the identification period ($Density_{\text{ID_Period}}$) and during the trading period ($Density_{\text{Trade_Period}}$). X_t is the spread between two normalized prices of a pair on day t . $2\sigma_{\text{hist}}$ is calculated as two times the standard deviation of X_t during the identification period. Pairs trading initiates a new trade if X_t is smaller than $\mu - \sigma$ or bigger than $\mu + \sigma$ during the trading period. The figure on the right-hand side displays the density function of X_t during the identification period ($Density_{\text{ID_Period}}$) and the density function of X_t during the trading period ($Density_{\text{Trade_Period}}$).

The figure on the right-hand side displays the case that the volatility of X_t increases between the identification and the trading period. The probability density function $\text{Density}_{\text{Trade_Period}}$ is now flatter and with fatter tails compared to $\text{Density}_{\text{ID_Period}}$. The trading signal of $2\sigma_{\text{hist}}$, however, remains unchanged. The fatter tails of the $\text{Density}_{\text{Trade_Period}}$ increase the chance of initiating a new trade.

The dashed area in the left figure between $\mu - 2\sigma_{\text{hist}}$ and $\mu + 2\sigma_{\text{hist}}$ represents the density to not exceed the trading thresholds. I derive from the density function that this probability is 95.45% if X_t is normally distributed.

Hence, the probability to open a new trade is 4.55% if the volatility does not change between the identification and the trading period. However, the opening probability changes significantly if X_t 's volatility increases to $\sigma_{\text{Trade_Period}}$ at the beginning of the trading period. However, the former thresholds to open the trade of $\mu \pm 2\sigma_{\text{hist}}$, defined based on the volatility as measured during the identification period σ_{hist} , are still active. The probability of X_t to exceed the trigger points, illustrated by the grey area in the right figure, is significantly higher than before with $X_t \sim n, 0, \sigma_{\text{trade}}$. If the volatility of X_t doubles ($\sigma_{\text{Trade_Period}} \rightarrow 2\sigma_{\text{hist}}$), then the opening probability increases from 4.55% to 31.73%.³ On the contrary, a drop in the volatility of X_t ($\sigma_{\text{Trade_Period}} < \sigma_{\text{hist}}$) decreases the probability to open a new trade.

I now turn to the drivers of X_t 's volatility. The volatility of the price spread X_t (σ_{hist} and $\sigma_{\text{Trade_Period}}$) is directly influenced by the stock volatilities σ_A and σ_B . Furthermore, X_t 's volatility is negatively influenced by $\rho_{A,B}$, the correlation of stock A and B.⁴ I derive these results from the following mathematical conversion:

³ The probability of a realization within the $\pm 2\sigma$ area around μ decreases from 95.45% to 68.27% for a realization within the $\pm \sigma$ threshold.

⁴ I derive these results from the following mathematical conversion: $V(X) = V(P_A - P_B) = V(P_A) + V(P_B) - 2\text{Cov}(P_A, P_B) = V(P_A) + V(P_B) - 2\rho_{A,B}\sqrt{V(P_A)}\sqrt{V(P_B)}$.

$$V(X) = V(P_A - P_B) = V(P_A) + V(P_B) - 2\text{Cov}(P_A, P_B) = V(P_A) + V(P_B) - 2\rho_{A,B}\sqrt{V(P_A)}\sqrt{V(P_B)}.$$

(17)

I conjecture that an increase in stock volatility σ_A or σ_B , or a decrease in pair correlation $\rho_{A,B}$ between the identification and trading period increases the volatility of X_t . In turn, the increase of X_t 's volatility increases the chance to open a new trade, and hence increases the trading frequency. In contrast, a decrease in stock volatility or a correlation increase reduces the trading frequency. I conclude that selecting promising pairs is also a matter of correctly forecasting volatility and correlation changes.

Correctly forecasting volatilities is one of the major fields in finance research. Volatility clustering over time implies that a period of low or high volatility is usually followed by a more medium volatility level. Most volatility forecasting models, therefore, assume that volatility follows a mean-reverting process (e.g., Ornstein-Uhlenbeck (1930) process in Stein and Stein (1991) in Heston's model (1993), and in ARCH/GARCH models (Engle (1982), Bollerslev (1986))). Please refer to Andersen, Bollerslev, Christoffersen, and Diebold (2006) for a comprehensive overview of models and literature on volatility forecasting.

If volatility reverts to its mean, the volatility of low volatility pairs should only be temporarily low. I expect more volatility increases for low volatility pairs, and hence, more trades by these pairs. In contrast, the volatility of high volatility pairs is more likely to decrease, which would generate fewer trades. I argue similarly for correlation and expect to observe more correlation decreases for pairs with a strong correlation compared to pairs with a weak correlation. I claim:

Research Hypothesis 3: The number of natural trades is increased by low stock volatility and high pair correlation during the identification period, coupled with higher volatility and lower correlation during the trading period. Pair volatility increases are more likely for pairs with currently low

pair volatility, while correlation decreases are more likely for pairs with a currently high correlation.

4.4 Modified methodology

I apply the methodology and return calculation as explained in chapter 2 with minor modifications that allow me to evaluate my research hypotheses.

To systematically analyze the influence of volatility and correlation, I analyze pairs trading returns separately for different pair volatility and correlation levels. I modify the classical pair selection algorithm as follows:

First, I calculate the pair correlation $\rho_{A,B}$ and pair volatility $(\sigma_A + \sigma_B)$, a proxy for the combined volatility of the pair, during the twelve-month identification period.

Second, I define five pair correlation and five pair volatility quintiles and classify pairs accordingly. Quintile includes the pairs with the weakest pair correlation, Quintile Corr_Q5 includes the pairs with the strongest pair correlation. Quintile Vola_Q1 includes the pairs with the lowest level of pair volatility, Quintile Vola_Q5 includes the pairs with the highest level of pair volatility.

Third, I construct twenty-five pair groups from the intersection of the five pair volatility and five pair correlation quintile groups. For example, Corr_Q1/Vola_Q1 includes the pairs, whose pair correlation is among the 20% of the weakest correlated pairs and whose pair volatility is among the 20% of the lowest pair volatility compared to all other pair combinations. The quintile and group affiliation of a pair is updated every identification period.

Fourth, I randomly pick twenty pairs out of each group and include them into one of the twenty-five portfolios, double-sorted on pair volatility and pair correlation. A pair remains in the trading portfolio during the subsequent six-month trading period and is eligible for trading.

Finally, I compute the average monthly pairs trading returns and compare the performance of portfolios of stocks with different pair volatility levels. I form the twenty-five portfolios based on five total pair volatility quintiles. I considered that the Vola_Q5 portfolio might contain outliers with extreme volatility as the volatility level of these pairs is not capped. The risk of picking an outlier within a given group exists in all Corr_Vola combination groups. A possible solution to overcome this issue is to simulate pairs trading with each possible pair combination in the Corr_Vola portfolio and determine the median of each group. I have decided against this approach because it is highly computational and time intensive. Instead, I eliminate the possible sample bias by picking twenty stocks out of each group and repeating the random pair selection ten times. Afterward, I average the returns of my ten repetitions.

Table 2 reports the mean pair correlation and pair volatility per quintile during the identification period.

Table 2: Correlation and volatility levels over quintiles

Table 2 shows the median correlation, the median pair volatility, and the median pair volatility per quintile for all selected pairs during all identification periods between January 1990 and June 2014. Pair volatility is defined as sum of stock A's and B's volatility.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Correlation	-0.4885	-0.0471	0.2973	0.5738	0.8026
	Vola_Q1	Vola_Q2	Vola_Q3	Vola_Q4	Vola_Q5

Pair Volatility	0.0134	0.0269	0.0434	0.0727	0.1892
-----------------	--------	--------	--------	--------	--------

For instance, the median pair correlation of all Corr_Q5 pairs (highest pair correlation quintile) is 0.8026. The median pair correlation ranges from -0.4885 (Corr_Q1) to 0.8024 (Corr_Q5). The average pair volatility ranges from 0.0134 (Vola_Q1) to 0.1892 (Vola_Q5).

In the next section, I evaluate my research hypotheses based on this modified test design with twenty-five portfolios.

4.5 Empirical results

4.5.1 Hypothesis 1 – Identification Period: A critical review of the SSD selection method in the identification period

What are the implications of applying the traditional pair selection criterion (SSD)? This section investigates the extent to which pair volatility and correlation influence the SSD as proclaimed in research hypothesis no. 1.

Research Hypothesis 1: Pairs with high correlation and low pair volatility are associated with a low SSD.

A significant influence of volatility and correlation implies that the use of the SSD equals selecting pairs with a specific pair volatility correlation combination. In turn, this pair selection with a specific volatility correlation combination may influence the trading behavior as proposed in research hypotheses No. 2 and No. 3, and hence, impacts the total return potential.

The classical distance measure SSD was introduced by Gatev, Goetzmann, and Rouwenhorst (2006)

and applied by several other authors. The appeal to use the SSD is its simplicity and the possibility of automatizing the process. The SSD calculates the sum of the squared deviations between the normalized price movements of two stocks. The twenty pairs with the lowest SSD are selected for trading in the subsequent trading period. To evaluate research hypothesis no. 1, I regress the SSD on pair volatility and correlation to investigate the explanation power of both factors. My sample includes all possible pair combinations across all identification periods.

Table 3 reports the regression results.

Table 3: Decomposition of the classical selection criterion (SSD)

*Table 3 displays the results of a panel regression of the SSD distance measure on the standardized pair volatility and the standardized correlation coefficient. I control for pair fixed and time fixed effects. Standard errors are clustered at the pair level and p-values (p_val) are reported in parentheses and adjusted for heteroscedasticity and clustered by pair combination. ***, **, * denote significance at the 1%, 5%, and 10% level, respectively.*

Variables	Pair Volatility (std)	Correlation (std)	Constant
Coefficient	604.2384***	-27.8919***	81.6122***

p_val	(0.000)	(0.000)	(0.000)
Observations	1,440,000		
R-squared	0.88		

Pair volatility and correlation explain 88% of the SSD's total variation, which confirms an overall strong influence of both factors. As expected, the SSD is positively correlated with pair volatility and negatively correlated with pair correlation. The SSD formula (2) does not reveal whether the effect of pair correlation or the volatility effect dominates. The standardized regression coefficients suggest that low pair volatility dominates the effect of high correlation. To check the robustness of my results, I repeat the regression individually for each of my ten data sets. The results confirm the previous results separately for each set. These findings add further evidence that the SSD selects pairs with low volatility that are strongly correlated.

As an additional test, I examine into which of my volatility and correlation double-sorted quintiles the twenty pairs with the lowest SSD are classified. Almost 67% of all SSD selected pairs are classified into Corr_Q5/Vola_Q1, the portfolio with the lowest pair volatility and highest correlation. The remaining pairs are all allocated to one of the other Corr_Q5 or Vola_Q1 portfolios. I conclude that either their pair volatility is very low or their correlation is high.

The next section investigates the isolated effect of pair volatility and correlation on the trading algorithm. The results revealed that the traditional selection procedure picks pairs with low pair volatility and high correlation. Therefore, I do not restrict my analysis to SSD-selected pairs and use the modified test design as introduced in Section 4.4.

4.5.2 Hypothesis 2 – Trading Period: The effect of volatility and correlation on the return per trade

Research hypothesis no. 2 derived the direction of the influence of pair volatility and correlation on the return per trade:

Research Hypothesis 2: Highly volatile pairs and negatively correlated pairs increase the return per trade of converging trades. In contrast, low volatility and strong correlation reduce the downside risk of diverging trades.

After the mathematical derivation of the hypothesis in section 4.3.2, this section empirically reviews the established hypothesis.

Table 4 reports the average return per trade for each closing type and portfolio.

Table 4: Return per trade

Table 4 shows the average return per trade for double sorted portfolios on pair volatility and correlation between January 1991 and December 2014. Panel A displays the average return per trade for natural trades that close after the full convergence of both stocks within the trading period. Panel B displays the average return per trade for incomplete trades that are forcefully closed on the last day of the trading period. Panel C displays the average return per trade for overshooting trades that are closed if prices diverge by more than $4\sigma_{hist}$. Panel D displays the average return per trade for delisted trades, which are closed if one stock is delisted while the pair is open. The trade universe includes all trades by pairs over time. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

Panel A: Closing type 1 (natural trades) - Mean return per trade

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	0.3022***	0.2482***	0.2180***	0.1841***	0.1441***
Vola_Q2	0.4269***	0.3658***	0.3236***	0.2666***	0.1979***
Vola_Q3	0.5207***	0.4584***	0.4004***	0.3274***	0.2442***
Vola_Q4	0.5912***	0.5592***	0.4847***	0.4042***	0.3027***
Vola_Q5	0.7256***	0.6580***	0.6103***	0.5376***	0.3861***

Panel B: Closing type 2 (incomplete trades) - Mean return per trade

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	0.0173***	0.0167***	0.0141***	0.0156***	0.0118***
Vola_Q2	0.0195***	0.0182***	0.0186***	0.0175***	0.0154***
Vola_Q3	0.0241***	0.0187***	0.0191***	0.0212***	0.0180***
Vola_Q4	0.0265***	0.0242***	0.0218***	0.0208***	0.0208***
Vola_Q5	0.0338***	0.0323***	0.0295***	0.0294***	0.0274***

Panel C: Closing type 3 (overshooting trades) - Mean return per trade

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	-0.2343***	-0.1998***	-0.1705***	-0.1448***	-0.1125***
Vola_Q2	-0.3098***	-0.2733***	-0.2398***	-0.2015***	-0.1529***
Vola_Q3	-0.3554***	-0.3232***	-0.2834***	-0.2452***	-0.1846***
Vola_Q4	-0.3904***	-0.3563***	-0.3310***	-0.2866***	-0.2211***
Vola_Q5	-0.4668***	-0.4137***	-0.3808***	-0.3463***	-0.2664***

Panel D: Closing type 4 (delisted trades) - Mean return per trade

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	0.0095*	0.0134***	0.0138***	0.0153***	0.0145***
Vola_Q2	0.0008	0.0126**	0.0174***	0.0190***	0.0151***
Vola_Q3	-0.0175	-0.0048	-0.0016	0.0162***	0.0167***
Vola_Q4	-0.1107***	-0.0395***	-0.0095	-0.0107	0.0088
Vola_Q5	-0.1322***	-0.1110***	-0.0169	-0.0766***	-0.0550***

The empirical findings demonstrate:

- Closing type 1 (natural trades): Panel A reports that the average return per natural trade of a Vola_Q1/Corr_Q5 pair (lowest pair volatility and highest correlation quintiles) is \$0.14 for a \$1 commitment in the long and the short position. The average return per trade ranges between \$0.1441 for Vola_Q1/Corr_Q5 pairs and \$0.7256 for Vola_Q5/Corr_Q1 pairs. As conjectured in research hypothesis no. 2, the return per trade for natural trades increases for higher pair volatility levels within each correlation quintile. Furthermore, it decreases for higher correlation (within each volatility quintile).
- Closing type 2 (incomplete trades): Although incomplete trades are forcefully closed on the last day, panel B reports an average positive return per trade. Similar to natural trades, higher pair

volatility and lower correlation significantly enhance the profitability of incomplete successful trades. The return ranges between \$0.0118 for Vola_Q1/Corr_Q5 and \$0.0338 for Vola_Q5/Corr_Q1.

- Closing type 3 (overshooting trades): Panel C shows that higher pair volatility and lower correlation increase the loss per trade for overshooting trades. The return per trade ranges between -\$0.4668 for Vola_Q5/Corr_Q1, and -\$0.1125 for Vola_Q1/Corr_Q5. As expected, the link between volatility, correlation and return of overshooting trades is the reverse of natural trades.
- Closing type 4 (delisted trades): Panel D reports that the return of delisted trades is on average negative for highly volatile and weakly correlated pairs. Yet, the return is increasing and positive for less volatile and highly correlated pairs. The average return per trade ranges between -\$0.1322 for Vola_Q5/Corr_Q1 and \$0.0190 for Vola_Q2/Corr_4. However, the return of eight out of twenty-five portfolios is insignificant, indicating that there is no clear volatility and correlation effect for delisted trades.

Altogether, the empirical findings support research hypothesis no. 2. High pair volatility and low correlation levels increase the return of converging pairs, but also increase the loss of diverging pairs. The volatility and correlation levels during the identification period determine the average return level per trade. Hence, pairs in more volatile markets or in times of general higher market volatility should earn a higher return per trade. My findings also imply that the SSD's preselection of strong correlated pairs with low volatility does not exploit the full potential compared to weakly or negative correlated pairs with higher volatility.

4.5.3 Hypothesis 3 – Trading Period: The effect of volatility and correlation on the trading frequency

After examining the return per trade in the previous section, I now turn to the second return building block, the trading frequency. This section explores the link between pair volatility, correlation, and

trading frequency to evaluate research hypothesis no. 3:

Research proposition 3: The number of natural trades is increased by low stock volatility and high pair correlation during the identification period, coupled with higher volatility and lower correlation during the trading period. Pair volatility increases are more likely for pairs with currently low pair volatility, while correlation decreases are more likely for pairs with currently high correlation.

I argue that raising the trading frequency leads to earning the return per trade multiple times during one trading period. Recall that a high trading frequency increases only the number of profitable, closing type 1 trades (natural trades). Closing types 2, 3 and 4 are always the last trades within a specific trading period. Hence, the number of these potentially loss-making trades is unaffected by a higher trading frequency. The argument is subject to the common assumption in financial research that the return of the additional trades is not eaten up by the transaction costs. Research on pairs trading (e.g. Do/Faff (2012)) reveals evidence that the return is persistent to trading costs.

In summary, pairs traders should identify the type of volatility and correlation level that increases the probability to open new trades. In other words: pairs traders should strive to increase the trading frequency.

The first analysis investigates the influence of pair volatility and correlation on the probability of a level shift. A level shift means that the pair volatility or correlation of a pair changes between the identification period and the trading period as explained in section 4.3.3. Price spread X_t is defined as the difference between the two normalized stock prices of the pair on day t and determines whether a new trade is initialized. The opening probability is higher if the volatility of X_t is low during the identification period and high during the trading period. As derived in research hypothesis no. 3, I conjecture that volatility increases and correlation decreases increase the volatility of X_t ,

which, in turn, increases the desired opening probability. I expect to observe more volatility increases for pairs with low pair volatility (Vola_Q1 quintile) and less volatility increases for pairs with high volatility (Vola_Q5 quintile). Likewise, I argue that correlation decreases are more likely for strongly correlated pairs (Corr_Q5) and less likely for weakly correlated pairs (Corr_Q1). I first compute the average pair volatility and the correlation coefficient of each pair during the identification period, and separately for the subsequent trading period. Based on this, I generate two dummies for each pair: the volatility dummy indicates whether the pair volatility increases between the identification and the trading period ($D = 1$) or not ($D = 0$). Likewise, I introduce a correlation dummy that indicates whether a correlation decrease happened between the identification and the trading period. I determine the volatility and the correlation dummy for each trade. Afterward, I calculate the average percentage of trades that are generated by pairs with a beneficial volatility increase or a correlation decrease for each volatility correlation double-sorted portfolio.

Table 5 reports the results.

Table 5: Level shifts and frequency

*Table 5, Panel A (B) shows the percentage of pairs with a volatility increase (correlation decrease) between the identification and the trading period within the twenty-five double-sorted portfolios. The percentage reveals that a volatility increase or a correlation decrease, which are both beneficial for the total return, are more likely for some portfolios than for others. The analysis includes all trades of the 57,600 pairs (twenty pairs per portfolio*10 sets*288 identification periods) during January 1991 and December 2014.*

Panel A: Average % of trades whose pairs realized a volatility increase

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	61.08%	60.69%	55.25%	46.78%	33.06%
Vola_Q2	53.12%	49.41%	43.60%	36.43%	25.05%
Vola_Q3	46.81%	42.58%	36.24%	30.08%	20.47%
Vola_Q4	42.17%	35.73%	29.66%	24.21%	16.43%
Vola_Q5	36.09%	28.48%	22.02%	18.74%	12.11%

Panel B: Average % of trades whose pairs realized a correlation decrease

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	23.68%	41.61%	57.26%	71.59%	85.53%
Vola_Q2	24.83%	45.33%	61.61%	75.30%	88.65%
Vola_Q3	25.47%	48.20%	64.09%	77.33%	90.47%
Vola_Q4	26.58%	51.43%	66.87%	79.51%	92.52%
Vola_Q5	26.68%	53.04%	69.39%	81.44%	93.78%

Table 5, panel A reports the results for pair volatility increases. 61.08% of all trades in portfolio

Corr_Q1/Vola_Q1 are generated by pairs with a volatility increase. Consistent with mean-reverting volatility (Engle and Patton (2001)), volatility increases are more likely for low pair volatility levels (Vola_Q1) across all correlation quintiles. Furthermore, pair volatility increases are more likely for reversely correlated pairs (Corr_Q1 quintile). Overall, panel A suggests that desired volatility increases are most frequently generated by pairs with low pair volatility and weak or negative correlation. Similar as in panel A, panel B shows the percentage of trades by pairs with a beneficial correlation decrease. For instance, 23.68% of all trades in portfolio Corr_Q1/ Vola_Q1 are generated by pairs whose correlation level decreased. Panel B indicates that desired correlation decreases are most often generated by pairs with high volatility and high correlation. Table 6 summarizes the insights from panels A and B. These results confirm research hypothesis No. 2.

Table 6: Influence between volatility and correlation level on beneficial volatility increases and correlation decreases

Table 6 provides an overview of the directions of the effects of high/low volatility and strong/weak correlation on the probability of a volatility increase or correlation decrease between the identification and the trading period.

Influence of	High volatility	Low volatility	Strong correlation	Weak correlation
Volatility increase	↓	↑	↓	↑
Correlation decrease	↑	↓	↑	↓

Pairs with a high volatility and strong correlation are disadvantageous for the volatility increase effect, but beneficial for the correlation decrease effect. Likewise, pairs with low volatility and weak or negative correlation are beneficial for the volatility increase effect, but disadvantageous for the correlation decrease effect. Yet, it is not clear whether the volatility increase effect dominates the correlation decrease effect or vice versa.

To shed light on this question whether the volatility increase effect or the correlation decrease effect

dominates, I analyze the absolute number of trades in each of the twenty-five double-sorted portfolios. The result is displayed in Table 7.

Table 7: Total number of trades

Table 7 reports the total number of trades of all pairs within the double sorted trading portfolio on pair volatility and correlation between January 1991 and December 2014.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	29238	39390	46821	53662	61083
Vola_Q2	22933	31443	38007	45015	53477
Vola_Q3	18814	25949	32051	38480	47633
Vola_Q4	14378	19762	24873	30808	40435
Vola_Q5	7576	10955	13740	17306	25206

Overall, Table 7 confirms the results of Table 6. The trading frequency is especially high for pairs with little volatility and a strong correlation. Pairs with high volatility and a weak or negative correlation trade relatively little. The number of trades generated by pairs with weak or negative correlation and low volatility are relatively higher than the number of trades generated by pairs with strong correlation and high volatility. These results suggest that the effect of high correlation in the correlation decrease effect is stronger than the effect of low correlation in the volatility increase effect.

4.6 Pairs Trading Return in calendar time

Two major results emerge from the previous two sections: first, the SSD selects pairs with low pair volatility and high correlation. This pair selection is a convincing selection to increase the trading

frequency. Second, the level of pair volatility and correlation determine the return per trade and the trading frequency in a complex relationship. For the predominant closing type 1 (natural trades), high pair volatility and negative correlation increase the return per trade on the one hand, but reduce the trading frequency on the other. To identify which effect dominates, I study the aggregated effect. Therefore, I take a closer look at the monthly pairs of trading returns. The monthly return represents the payoff to an investor, and it considers the aggregated effect of pair volatility and correlation on the return per trade and the trading frequency. It provides not only the possibility of computing a risk-adjusted return, but also allows us to compare the monthly return of a traditionally formed SSD portfolio and portfolios with alternating volatility correlation combinations. Additionally, the results help us to ultimately judge whether the SSD is superior in selecting profitable pairs.

The analysis of the monthly returns between January 1991 and December 2014 uncovers whether the positive effect of high pair volatility and low correlation on the return per trade, or the negative effect of high volatility and low correlation on the trading frequency dominates.

Table 8 shows the results of the analysis.

Table 8: Average monthly trading return

Table 8, panel A reports the average monthly raw returns of each double sorted trading portfolio on pair volatility and correlation between January 1991 and December 2014. To account for well-established return patterns (Fama and French (1993); Jegadeesh and Titman (1993); Carhart (1997); Da, Liu and Schaumburg (2014)), I regress monthly excess returns on a six-factor model. Panel B displays the alphas from a time series regression of the one-month pairs trading returns on the six-factor model including Fama and French's (1993) three factor model, a momentum factor, a short-term reversal factor, and Pastor and Stambaugh's (2003) liquidity factor. I use Newey-West (1987) standard errors with lag six. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

Panel A: Pairs trading return in calendar time

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5	Q5-Q1	Q5-Q2
Vola_Q1	.0046***	.0037***	.0040***	.0051***	.0047***	0.0000	0.0010***
Vola_Q2	.0049***	.0035***	.0046***	.0052***	.0050***	0.0000	0.0015**
Vola_Q3	.0033***	.0038***	.0058***	.0057***	.0065***	0.0033***	0.0027***
Vola_Q4	.0044***	.0058***	.0062***	.0065***	.0076***	0.0032**	0.0018
Vola_Q5	.0019*	.0043***	.0046***	.0056***	.0234*	0.0215	0.0191
Q4-Q1	-0.0002	0.0021*	.0023**	.0014*	.0030***		

*** p<0.01, ** p<0.05, * p<0.1

Panel B: Alphas of each correlation and pair volatility quintile portfolio

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5	Q5-Q1	Q5-Q2
Vola_Q1	.0046***	.0036***	.0039***	.0052***	.0047***	0.0011	0.0001**
Vola_Q2	.0047***	.0031***	.0045***	.0052***	.0048***	0.0017	0.0001***
Vola_Q3	.0028**	.0037***	.0059***	.0056***	.0066***	0.0029***	0.0038***
Vola_Q4	.0040*	.0059***	.0066***	.0068***	.0076***	0.0017**	0.0036
Vola_Q5	0.0018	.0047***	.0053*	.0057***	0.0209	0.0162	0.0191
Q4-Q1	-0.0006	0.0023	0.0028**	0.0016*	0.0029***		

*** p<0.01, ** p<0.05, * p<0.1

In panel A, all returns are positive, significant and range between 19 bp to 234 bp. For instance, portfolio Vola_Q1/Corr_Q5 (low pair volatility/high correlation) yields an average monthly return

of 47 bp. The most striking result to emerge from panel A is the successive increase in monthly returns for higher volatility levels from Vola_Q1 to Vola_Q4 in correlation quintiles Corr_Q2 to Corr_Q5. The return increase is particularly stronger at higher correlation levels. The return differences between Corr_Q5 and Corr_Q1 portfolios are insignificant in three out of five cases, indicating equally high portfolio returns. Furthermore, twenty-one out of twenty-five volatility correlation double-sorted portfolios outperform the SSD portfolio with a monthly return of 37 bp. Even the direct peer portfolio Vola_Q1/Corr_Q5 with low pair volatility and high correlation earns 10 bp more. I speculate that the superior performance of the Vola_Q1/Corr_Q5 portfolio results from randomly selecting twenty out of approximately 72,000 pairs within the particular quintile intersection group. These selected pairs have an average level of correlation and volatility within the volatility correlation double-sorted group. In contrast, the pairs of the SSD portfolio observe the absolutely lowest pair volatility and the highest correlation. I conclude that even slightly higher volatility and lower correlation levels increase the pairs' trading return.

To control for well-established return patterns, I regress monthly returns on a six-factor model, including Fama and French's three-factor model, a momentum factor and a short-term-reversal factor. All data is obtained from Kenneth French's website. As my sample includes stocks with different liquidity levels, I further extend my factor model with Pastor and Stambaugh's (2003) liquidity factor.

Panel B reports the alphas of the regression for each volatility correlation double-sorted portfolio. Like related papers, I use Newey-West Standard errors with lag 6. The results confirm my previous findings. All alphas are positive and almost always significant. Again, high pair volatility and strong correlation increase alpha.

Altogether, the results suggest that pairs with high pair volatility and high correlation maximize the

return. I conclude that the positive pair volatility effect on the return per trade and the positive effect of strong correlation on the trading frequency dominate. This implicates that the SSD, which selects pairs with low pair volatility, is not a good instrument to identify the most profitable pairs. The monthly return of the traditionally selected portfolio with 37 bp is considerably lower than the risk adjusted return of most alternatively formed volatility correlation double-sorted portfolios.

4.7 Robustness

Trading strategies that successfully exploit anomalies always raise several robustness questions. Therefore, I first address the limits of arbitrage in the next section before I test my results in a cross-sectional Fama-MacBeth (1973) test design afterward.

4.7.1 Limits of arbitrage

One obvious question is whether pairs trading returns are persistent to trading costs. Do and Faff (2012) dedicate a full paper to the question and find that pairs trading returns are decreasing but still significant and positive after controlling for commissions, market impact, and short selling constraints. Their result is consistent with Gatev, Goetzmann, and Rouwenhorst (2006), who also find significant positive returns after controlling for transaction costs.

My previously chosen highly volatile pairs might, however, demand higher trading costs. I address the most critical issue, the bid-ask bounce, by directly trading with bid and ask prices, which is more accurate than the traditional one-day waiting procedure. As my alphas are positive and significant, I conclude that transaction costs are not an impediment, especially not for large institutional traders. Furthermore, an investor must be able to hold the position until the pair converges. While the trade is open, investors might be confronted with margin calls if the short stock diverges further. Subsequently, some arbitrageurs must liquidate their positions to meet the required additional capital commitment as addressed by Shleifer and Vishny (1997). Recalls are more likely when prices are falling.

When prices start to converge in pairs trading, the shorted stock of a pair is expected to fall, and is hence more likely to be recalled. However, D'Avolio (2002) observes a low recall rate of around 2% per month. Furthermore, pairs traders can still earn a profit if the price already converges until the stock is recalled. To shed further light on this issue, I examine the volatility and correlation effect on the trading frequency of overshooting trades (closing type No. 3). These were the trades that are forcefully closed during the trading period as the difference between the standardized prices exceeds 4σ . Table 9 reports the absolute number of overshooting trades for each volatility correlation double-sorted portfolio.

Table 9: Number of Closing Type No. 3 trades (overshooting trades)

Table 9 reports the absolute number of overshooting trades for each volatility correlation double-sorted portfolio. For instance, 6152 overshooting trades were observed for the Corr_Q1/Vola_Q1 portfolio. This means that all stock pairs belonging to the group of stocks with the 20% lowest correlation and the 20% lowest volatility generated 6152 overshooting trades.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	6152	10751	15135	19344	23677
Vola_Q2	3999	7320	10781	14588	19621
Vola_Q3	3093	5519	7979	11246	16490
Vola_Q4	2282	3920	5688	8125	13256
Vola_Q5	1259	2177	2972	4302	7812

The results indicate that high pair volatility and low correlation reduce the number of overshooting trades, and hence, divergence risk. This finding is consistent with Engelberg, Gao, and Jagannathan (2009), who find that high idiosyncratic volatility decreases the divergence probability.

In addition, I address the previously mentioned arbitrage problems by replicating my analyses with stocks that are included at the time in the S&P 100 or the NASDAQ 100 indices. To consider possible liquidity declines, I exclude stocks if they lose their index membership. This test setting challenges the robustness of my findings in an almost limits of arbitrage-free setting, as all stocks are liquid, easy-to-short, cheap-to-trade, and highly efficient due to high analyst coverage. Table 10 summarizes the results.

Table 10: Pairs trading in highly liquid markets

Table 10, panel A reports the average monthly raw returns for each double-sorted trading portfolio on volatility and correlation between January 1991 and December 2014. The sample includes only the, at that time, current members of the S&P 100. Likewise, panel B displays the average monthly returns of NASDAQ 100 members between January 1995⁵ and December 2014. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***.

Panel A: Average monthly returns of the S&P 100 (Jan 1990-Dec 2014)

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	0.0038***	0.0030***	0.0025***	0.0025***	0.0019***
Vola_Q2	0.0055***	0.0040***	0.0040***	0.0028***	0.0031***
Vola_Q3	0.0037**	0.0049***	0.0036***	0.0039***	0.0030***
Vola_Q4	0.0035**	0.0039***	0.0046***	0.0048***	0.0042***
Vola_Q5	0.0023	0.0025*	0.0027**	0.0033***	0.0034***
Q5-Q1	-0.0015	-0.0005	0.0002	0.0008	0.0015**
Q4-Q1	-0.0003	0.0009	0.0021*	0.0023**	0.0023***

Panel B: Average monthly returns of the NASDAQ 100 (Jan 1995-Dec 2014)

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	0.0032**	0.0046***	0.0051***	0.0045***	0.0020***
Vola_Q2	0.0022	0.0048***	0.0055***	0.0046***	0.0057***
Vola_Q3	0.0024	0.0065**	0.0065***	0.0055***	0.0072***
Vola_Q4	0.0026*	0.0076***	0.0112***	0.0066***	0.0067***
Vola_Q5	0.0070**	0.0142	0.0108*	0.0063***	0.0074***
Q5-Q1	0.0038	0.0096	0.0057	0.0018	0.0055***
Q4-Q1	-0.0006	0.0030	0.0061**	0.0021	0.0048***

⁵ The reporting of bid and ask prices for the NASDAQ 100 starts in Jan 1995.

All returns are smaller compared to previous returns. Thus, part of the previous returns can be explained with short selling constraints, a liquidity premium, and efficiency issues. However, all returns are still positive and significant. The NASDAQ 100 portfolios outperform the corresponding S&P 100 portfolios. This result is consistent with my previous results, as NASDAQ 100 stocks are more volatile than S&P 100 stocks.

4.7.2 Short Selling Constraints

Another issue is whether it is difficult to short stocks for pairs trading. Stambaugh, Yu, and Yuan (2015) argue that the return asymmetry of long-short strategies might originate from short selling constraints of overpriced stocks compared to easily exploitable underpriced stocks. In this case, the isolated short leg return exceeds the long leg return. Alternatively, Gatev, Goetzmann, and Rouwenhorst (2006) claim that long leg profits might represent compensation for an unrealized bankruptcy. In this alternative case, the long leg contributes more to the total return than the short leg. To shed further light on this topic, I separately compute the isolated return generated by the short and the long leg, similar to Gatev, Goetzmann, and Rouwenhorst (2006), and determine the following contribution ratio:

$$\text{Contribution Ratio} = \frac{\text{Long leg return per trade}}{\text{Total return per trade}} \quad (18)$$

Table 11 reports the results.

Table 11: Median percentage contribution of the long leg to the total return

Table 11 investigates the contribution of the short and the long leg separately, and displays the median percentage contribution of the long leg to the total return between January 1991 and December 2014. A contribution rate above 0% but below 50% indicates that the long leg contributes to the total return, but less than the short leg.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	51.29%	50.07%	49.96%	49.45%	49.80%
Vola_Q2	49.01%	48.91%	48.15%	48.66%	49.21%
Vola_Q3	47.69%	45.75%	46.59%	46.85%	48.03%
Vola_Q4	41.87%	41.36%	43.05%	43.97%	46.21%
Vola_Q5	31.64%	32.60%	35.17%	36.45%	42.12%

Table 11 reveals that the percentage is below 50% in twenty-four out of twenty-five portfolios. This finding indicates a higher contribution of the short leg. The long legs of Vola_Q5/Corr_Q1&2 portfolios contribute relatively little with a share of 31.64%-32.60%. As the return asymmetry is especially strong for highly volatile stocks, the short leg contribution might originate from mispricing that is practically not exploitable. However, D’Avolio (2002) finds that only 16% of all stocks included in the monthly CRSP file can eventually not be shorted. 91% of all stocks, including almost all S&P 500 constituents, cost less than 1% to borrow, and their value-weighted average fee is approximately 17%. The remaining 9% of stocks, which are called “special stocks,” have an average fee of 4.5% per annum. Less than 1% of the special stocks demand negative rebate rates and charge a fee of up to 50%. Not surprisingly, smaller stocks demand higher fees. However, these special stocks account for less than 1% of the market by market value. I assume that the probability of “special stocks” in my dataset is relatively low for two reasons. First, my stock universe is restricted to current or former members of the S&P 1500. These stocks cover around 90% of the total market capitalization and should be under regular investor’s attention according to Wurgler and

Zhuravskaya (2002). Second, my stock universe is restricted to stocks with a listed bid and ask price. These restrictions secure the liquidity and tradability of my stocks. Even if shorting is impossible, a modified pairs trading strategy could only invest in the “loser” stock (long position) and use the “winner” stock as a trading signal without physically selling the stock. Alternatively, investors could include only stocks that they already hold in their portfolios.

5.6.1. Fama-MacBeth regression

In addition to my analyses based on portfolio sorts, I conduct a Fama-MacBeth (1973) regression to investigate the role of pair volatility and correlation in the cross-section. I regress pair volatility and correlation on the return per trade. This approach allows me to evaluate how the technical drivers influence the return per trade (research hypothesis No. 2) in the cross-section. I use standardized coefficients for pair volatility and correlation and apply Newey-West (1987) standard errors with lag six. I control for the duration of a trade to consider a potentially wider price spread at the opening.

A pair can open multiple times during one trading period. The returns of these trades are similar because the opening signal, which determines the return per trade, is always identical in these cases. Therefore, I only consider the first trade of each pair within a trading period. The volatility coefficient and the correlation coefficient are both significant at a 1% significance level. A one standard deviation increase in pair volatility increases the trade return by 13 cents, whereas a one standard deviation increase in correlation decreases the trade return by 7.5 cents. The result adds further evidence to my previous findings. For further robustness, I repeat the analysis and choose one random trade of each pair within a trading period. I also repeat the analysis with non-overlapping trading periods. The results remain robust in both additional tests.

4.8 Conclusion

I contribute to the academic discussion about pairs trading by demonstrating that pair volatility and correlation are the important levers that influence the two return building blocks of pairs trading: the return per trade and the trading frequency. High pair volatility and low correlation are beneficial for the return per trade, but not beneficial for the trading frequency.

The key learning for active pairs traders is that choosing stocks of one specific industry, one country, or applying a certain selection procedure like the SSD results in a preselection of eligible pairs for trading in terms of volatility and correlation levels. In consequence, this preselection of pairs determines a given pair volatility and correlation combination. The selection of pair volatility and correlation levels, in turn, affects the total return via the return per trade and the trading frequency. Hence, smart investors should optimize the pair volatility correlation level of their stock selection to maximize their returns.

5. Core chapter 2: Slow information diffusion as a driver for pairs trading

5.1 Motivation

Pairs traders are confronted with a major problem: One can only observe the deviation of two stocks from each other, while the reason why these stocks deviate is veiled. The unobservable deviation driver implicates the risk of a permanent stock price deviation. The lack of an economic explanation and the unknown return source, which creates the market inefficiency, makes it difficult to sell the trading idea to clients.

Gatev, Goetzmann, and Rouwenhorst (2006) motivate pairs trading with the exploitation of temporary mispricing according to the law of one price. Prices are initially driven away from their fair value by any sort of market frictions. Natural candidates for pairs are pairs with a lead-lag relationship as their prices regularly diverge and later converge. In recent years, a considerable amount of literature introduces gradual information diffusion as the explanation for lead-lag relationships between stocks. For instance, Engelberg, Gao, and Jagannathan (2009) reveal that pairs, whose industry information diffusion speed differs, outperform pairs with a similar information diffusion speed. Furthermore, they find that pairs trading is more profitable for pairs with few common analysts or few common institutional holdings. Chen, Chen, and Li (2013) document that an environment with slower information diffusion is beneficial for pairs trading, e.g., small companies without media coverage, low analyst coverage, or low investor recognition. Moreover, they find that pairs trading returns have decreased over the last few years, consistent with improved information efficiency over time. They also show that pairs trading is more profitable for highly correlated pairs based on known determinants that drive return co-movement. Jacobs and Weber (2015) provide evidence that limited investor attention, which provokes slower information diffusion, increases the strategy's profitability.

The idea of gradual information diffusion is basically that one stock incorporates common news faster than the other one. The first stock, the “leader stock,” processes information faster and reacts more quickly to the arrival of new information than the second stock. The second stock, the “follower stock,” incorporates information more slowly. It moves into the same direction like the “leader stock” – however, with a time lag. Throughout this section, the term “leader-follower relationship” will be used to refer to this relationship of stocks, generated by gradual information diffusion. This leader-follower relationship creates a temporary mispricing of the “follower stock,” which can be exploited by appropriate trading strategies – for instance, pairs trading.

The literature provides numerous explanations as to why stocks process price relevant information faster or slower, which leads to gradual information diffusion across the stock universe. For instance, Hong and Stein (1999) theoretically model dynamic information diffusion across sophisticated investors and momentum traders over time. Thereafter, many papers identify sources that induce gradual information diffusion including size (Lo and MacKinlay (1990)), number of analyst following the company (Brennan, Jegadeesh, and Swaminathan (1993)), the level of institutional ownership (Badrinath, Kale, and Noe (1995)), and trading volume (Chordia and Swaminathan (2000)). There are also company links that indirectly delay the incorporation of information so that follower-leader stock relationships develop. For instance, Cohen and Lou (2012) find that less complex companies, which only concentrate on one core business, incorporate industry information faster into prices than more complex companies with many subsidiaries and business fields in the same industry. Other factors include intra-industry leadership (Hou (2007)), industries that lead the market (Hong, Torous, and Valkanov (2007)), economically-linked companies (Cohen and Frazzini (2008)), supplier and customer industries (Menzly and Ozbas (2010)), and mutual ownership (Antón and Polk (2014)).

The central purpose of this section is to provide evidence that gradual information diffusion is a major source of pairs trading returns. I start with a closer look on the connection, which both stocks of a pair must have, and explain how pairs trading exploits information diffusion-generated mispricing. Afterward, I reason why pairs trading should be included in the circle of promising investment strategies to exploit that type of mispricing. The next section investigates the leader-follower relationship of pairs more closely and links it to common risk factors (Fama/French three factor model, short-term reversal factor, and momentum factor). In the final section, I analyze whether each trade shows the typical price development that would be expected for a leader-follower relationship and whether the leader-follower role allocation is persistent. The following empirical analyses are based on the data set of the previous core chapter.

5.2 Conditions to exploit information diffusion induced mispricing using pairs trading

Gradual information diffusion can generate temporary mispricing. Pairs trading is a perfect candidate to exploit this type of mispricing. However, both stocks of the pair must share some similarities but must also have different attributes:

First, both stocks must be economically connected in a certain way. This link generates a similar price reaction upon the arrival of link-specific news. Broad literature on co-movement proposes multiple links that cause similar price reactions to exogenous impulses. Obvious factors include industry affiliation, regional location, but also veiled links like mutual ownership. This link generates a similar price reaction. I refer to this type of news as **common news**.

Second, there must be a source or specific stock characteristic that initiates the gradual information processing. I refer to this component as the **deviation factor**.

To understand how pairs trading exploits market inefficiencies generated by gradual information diffusion, one can think of two stock exchange-listed companies A and B. Both companies operate

in the same industry and their stock prices are consequently influenced by similar factors. However, stock A processes information faster than stock B, as stock A is covered by more analysts that closely observe the business development and react accordingly. Upon the arrival of new positive industry news, stock A incorporates the positive information almost immediately and rises. A's stock price increase causes the initial price deviation of 2σ , the threshold for pairs trading. \$1 of the normalized price of stock A is sold and \$1 of the normalized price of stock B is purchased. As stock B is also affected by the particular positive industry information, stock B also reacts to the novel positive industry information. However, stock B is covered by less analysts, so the stock price of B rises slower and it takes a while until the price fully reflects the piece of information. At a certain point, stock B catches up with stock A and both pairs' trading positions are neutralized. In this setting, stock A is the leader stock and stock B the follower stock. The key to understanding the source of pairs trading returns is to understand gradual information diffusion and in particular, the resulting leader-follower relationship of a stock pair. On the chart, stock B follows stock A in a steep movement, until the information is completely incorporated into stock B, so that both prices intersect. At that time, both positions are neutralized.

5.3 Information diffusion exploiting trading strategies vs. pairs trading

Many papers apply a simple procedure to exploiting temporary mispricing that is generated by gradual information diffusion. First, traders identify and form economically linked pairs, for instance Cohen and Frazzini (2008) define pairs based on a customer-supplier relationship. Afterward, stocks are classified as fast reacting "leader stocks" or as information slowly processing "follower stocks," based on an economical derived factor. Second, based on the leader's performance in the previous month, traders buy or sell the corresponding follower stocks thereafter. This type of trading strategy requires a clear identification of the leader and the follower stock. The classification is usually based

on one particular factor, which induces gradual information processing, like size or analyst coverage.

However, this generalized procedure fails to grasp the total complexity of pair links. Two stocks might have a mutual link, as they are affected by the same influence factors. Yet, some factors have a positive influence on both stocks, while other factors have a positive influence on one stock but a negative influence on the other stock. Mutual influence factors might intensify, neutralize or contradict each other. Let me give you an example: Company A and company B share two links. Company A is a supplier of company B. This link suggests that stock A predicts the return of stock B. However, stock B is more liquid than stock A, which suggests that stock B incorporates relevant common information more quickly than stock A. The leader-follower role allocation is unclear in this setting. Stock A predicts stock B upon the arrival of company-specific information (Cohen and Frazzini (2008)), whereas stock B predicts stock A regarding market news. It is unclear whether the supplier link or the liquidity link dominates overall. In the worst case, stock B is erroneously included in a follower portfolio, although it predicts stock A if the liquidity link is stronger than the supplier link.

In contrast, the holistic approach of pairs trading allows the coexistence of contradicting links. A permanent leader follower role allocation is not necessary for pairs trading. If company specific news arrives, positions are set so that stock A is the leader and stock B the follower. The positions are set in the opposite direction if market news arrives and stock B is the leader and stock A the follower. This attribute of pairs trading is also useful if the leader-follower relationship changes over time.

Another potential concern about the classic procedure to exploit gradual information diffusion is the timing. Most strategies update their portfolio with follower stocks every month. The follower stocks

are bought at the beginning of the month and sold at the end of the month. This procedure implicitly presumes that the follower stock's information processing lasts from the first day until the last day of the stock's portfolio affiliation months. This timing is presumably not synchronized with the actual temporary mispricing generated by gradual information diffusion. For instance, stock B is included in the follower portfolio based on its leader stock A's development in the previous month. However, it is possible that stock B already completely incorporated the information in the previous month or does not or only partly processes the information in this particular month. This imprecise timing of the strategy introduces additional risk to the return. Pairs trading overcomes this timing problem as positions are not set before a sufficient price difference occurs. The case that stock B already fully incorporated the relevant information is therefore excluded. Similarly, the positions are held until the price divergence completely disappears if it does not exceed 4σ . Pairs trading accordingly also captures the incorporation of news beyond one month. For all these reasons, pairs trading should be included in the circle of promising investment strategies to exploit information diffusion induced mispricing.

5.4 The role of the leader-follower relationship of stock pairs in pairs trading

The previous section demonstrated the importance of links that induce a leader-follower relationship for pairs trading. The objective of the next section is to take a closer look at the prediction power of the leader-follower relationship in the next two sections and the role of common risk factors in the last section.

5.4.1 Aggregated cross-serial correlation of stock pairs

The success of pairs trading heavily depends on the existence of leader-follower relationships. To learn more about the prediction power of leader-follower relationship, I determine the total prediction power of one stock's lagged returns on the second stock and vice versa in the first step. The

term “cross-serial correlation” refers to the degree of prediction power in the following. In the next step, I form a portfolio with strongly and weakly cross-serial correlated pairs and back-test pairs trading for the defined portfolio.

If gradual information diffusion drives pairs trading returns, pairs trading with strongly cross-serial correlated pairs should be less risky to exploit because of the high prediction power of the leader-follower relationship. If the leader stock accurately predicts the price development of the follower stock, then there must have been a sufficient number of fundamental traders who incorporated news into prices in the past. Hence, temporary price deviations should disappear after a short while and open trades will close more frequently. Thus, the pairs trading portfolio should yield a higher trade return.

In contrast, price deviations of cross-serial uncorrelated pairs are riskier. Here, the successful elimination of a temporary price deviation requires the presence of sufficient arbitrageurs in the market and the absence of limits to arbitrage. Therefore, a non-economical initiated price deviation is more likely to be persistent. I claim that a portfolio that includes strongly cross-serial correlated pairs outperforms a portfolio that includes weakly cross-serial correlated pairs.

5.4.2 Empirical evidence from portfolios with differently cross-serial correlated pairs

To verify the impact of cross-serial correlation, I compare the return of strong and weak cross-serial pairs in three steps:

First, I regress the return of stock A on stock B’s lagged returns to measure the degree of cross-serial correlation in the first step. I consider twenty lags that approximate a cross-serial correlation of up to twenty trading days or one month. This time period was also used by Hong, Torous, and Valkonov (2007). R_A^2 describes how much of the volatility of stock A is described by the lagged returns of stock B. It can be understood as explanation power of stock B (leader) to explain the

return of stock A (follower), and is a proxy for the strength of the leader-follower relationship between stock A and B. After repeating the regression for A as the leader and stock B as the follower, I determine the stock with the higher explanation power R_{\max}^2 (with $R_{\max}^2 = \{R_A^2; R_B^2\}$). Overall, each pair combination is described by two R^2 (R_A^2 and R_B^2) that describe the magnitude of the leader-follower relationship between both stocks.

In the second step, I form two portfolios – one portfolio that includes pairs with a strong leader-follower relationship, and a second portfolio that includes pairs with a low leader-follower relationship. For this portfolio, I rank each pair based on its R_{\max}^2 and include the pairs with the twenty highest R_{\max}^2 . For the strong cross-serial correlated portfolio, it is not important whether R_{\max}^2 is R_A^2 or R_B^2 or whether both R^2 are high, as long as one strong cross predictability link exists. In contrast, the weak cross-serial correlation portfolio must contain pairs with two low R_{low}^2 . Therefore, I rank R_A^2 and R_B^2 separately in ascending order and compute the sum of both ranks. I include the twenty pairs with the lowest rank sum in the weak cross-serial correlated portfolio.

In the final step, I simulate the returns for pairs trading for both portfolios and compare the monthly returns. To do that, I keep updating the portfolio each month based on the R^2 .

One has to keep in mind that stocks are often influenced by common risk factors like size, book-to-market value and others. As I am also interested in influence factors that link both companies, for instance industry or ownership affiliated links, and I also repeat the previously described procedure with fitted values. For this purpose, I use the residuals of the previously introduced five-factor model (including the three Fama-French factors, a momentum, and a short-term reversal factor) to perform the cross-serial correlation regressions to get the $R_{\text{fitted}}^2 \cdot R_{\text{fitted}}^2$, which can be understood as correlation measure, orthogonal of similar exposure to common risk factors. This measure is important, as many influence factors like mutual ownership are independent of these common risk factors. The

portfolios are labeled as fitted cross-serial correlation portfolios with strong or weak cross-serial correlation.

Table 12 shows the average pairs trading returns for all four portfolios (strong vs. weak cross-serial correlation; not-fitted vs. fitted).

Table 12: Returns of strong and weak cross-serial correlation portfolios

Table 12 displays the average pairs trading return for stocks with a strong and a weak cross-serial correlation. I determine the extent of cross-serial correlation by regressing the time-lagged returns of one stock on the returns of the other stock and vice versa. The respective two R^2 show the explanation power. The pairs with the highest individual R^2 are included in the strong cross-serial correlation sample and pairs with the two combined lowest R^2 are included in the weak cross-serial

correlation portfolio. Fitted vs. not-fitted differentiates whether the five-factor model was added in the regression to determine R^2 (fitted) or not (not-fitted). The column 'Strong – Weak' tests whether the returns between both samples significantly differ.

Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively. P_val reports the p-values from whether the return is significantly greater than 0.

	Strong cross-serial correlation	Weak cross-serial correlation	Strong - Weak
Not-Fitted	0.0100***	0.0036**	0.0064**
p_val	(0.0000)	(0.0234)	(0.0036)
Fitted	0.0118***	0.0043	0.0075**
p_val	(0.0000)	(0.1248)	(0.0391)

Both strong cross-serial correlation portfolios earn a positive pairs trading return, and the positive return is significant. Furthermore, they both significantly outperform the corresponding weak cross-serial correlated portfolio by 64 bp (not-fitted) and 75 bp (fitted), respectively. This result is interesting, as it provides evidence that pairs trading is more profitable with strongly cross-serial correlated pairs than with weakly cross-correlated pairs. Notably, the pairs trading return of the fitted, strong cross-serial correlated portfolio is also positive and significant. I conclude that a strong leader-follower relationship does not depend on common risk factors.

On the contrary, the pairs trading return of the weak cross-serial correlation portfolio is only significant and positive for non-fitted pairs, but not for fitted pairs. This result is interesting, as it allows two conclusions. Firstly, it demonstrates that company-specific links between the stock pairs are necessary to earn a positive pairs trading return in the absence of exposure to common risk factors. Secondly, in the absence of company-specific links between a stock pair, exposure to common risk factors creates a profitable leader-follower relationship.

A comparison of the returns of these portfolios, built on the strength of the follower-leader relationship with the portfolio in Section 4.6, double-sorted on volatility and correlation, is worth mentioning. Both strong follower-leader portfolios outperform all but one volatility correlation-sorted portfolio of Table 8, panel A. The returns of 100 bp (not-fitted) and 118 bp (fitted) are considerable higher than the 76 bp of the second most profitable portfolio (Vola_Q4/Corr_Q5). Only the Vola_Q5/Corr_Q5 portfolio earns more. This finding is highly remarkable, as it suggests that pairs should not be selected based on the classic SSD distance measure, but on the strength of their leader-follower relationship, approximated by their cross-serial correlation. Furthermore, the finding demonstrates that selecting promising pairs should not be based on the at-that-day price difference as the classical correlation formula does. Pairs traders should rather focus on cross-serial correlation, as this type of correlation identifies gradual information diffusion across stocks more accurately than classic correlation.

5.4.3 Common risk factors and cross-serial correlated pairs

Until now, the evaluation of the leader-follower relationship in this thesis is purely based on the statistical prediction power without considering the economic source of the common link. This section overcomes the weakness and investigates the economical source of the common news factor that initiates the leader-follower relationship.

Both stocks must share an exposure to similar economic factors as otherwise, the arrival of new information would not affect both stocks. Previous studies on pairs trading concentrate on similarities of firm characteristics as in Chen, Chen, and Li (2013). However, the results of the previous section suggest a leader-follower relationship can partly be explained by the exposure of stocks to common risk factors.

I develop a close economic substitute score (CES score) to measure the economic closeness of two stocks. The CES score measures the coexisting exposure of stocks to common risk factors. I choose the highly acknowledged five-factor model as previously introduced. All data is obtained from Kenneth French's website. The CES score is computed as follows:

First, I regress the daily closing prices of each stock on the factors of the five-factor model to determine each stock's exposure to each of the five common risk factors. I compute standardized coefficients for each stock in the sample. I do not use these coefficients for an economical interpretation. Rather, scaling the coefficients by calculating standardized coefficients allows me to compare the difference in exposure to a common risk factor between two stocks across all risk factors.

Second, I calculate the absolute difference of each standardized coefficient between stocks A and B for all possible pair combinations in each identification period. This coefficient difference indicates whether both stocks have a similar exposure to the common risk factor. If the price difference is small, both stocks react similarly to news shocks of the particular common factor.

Third, I rank all pairs in ascending order based on their coefficient difference for each risk factor and divide them into five quintiles accordingly. Each pair combination is sorted into one quintile for each risk factor i . A score Q_i is assigned based upon that. For instance: if a pair is sorted into the lowest quintile for risk factor i , then it reacts more similarly to a shock of risk factor i compared to at least 80% of the other stocks. The Q_i score of the stock pair is 1 for risk factor i .

In the next step, I compute the CES score for each pair. The CES score sums up the values of all five pricing factors i (Q_i). To fit the CES score into a range between zero and twenty, the sum is deducted by five.

$$\text{CES Score} = \sum_{i=1}^5 Q_i - 5 \quad (19)$$

The CES score is defined in the interval between zero (extremely close substitutes) and twenty (no systematic connection). A CES Score of zero indicates that both stocks react more alike upon the arrival of a risk factor shock than at least 80% of the other stock pairs would. The score is incremented by one point if the quintile of any pricing factor is increased by one. For example, the score is equal to ten if the quintiles of all risk factors equal three.

Table 13 reports the average CES score across pairs and time for my cross-serial correlated portfolios.

Table 13: CES score of cross-serial correlated portfolios

Table 13 displays the average CES score for portfolios of strong and of weak cross-serial correlated pairs. The CES score shows whether both stocks in the pair react with a similar sensitivity to five pricing factors (Fama and French's three-factor model, momentum factor, liquidity factor). The lower the CES Score, the more similar the exposure to the risk factors.

	Strong cross-serial correlation	Weak cross-serial correlation
Average CES Score	7.8481	10.4906

The average CES Score is 7.85 in the strong cross-serial correlation portfolio and 10.49 in the weak correlation portfolio. The CES score suggests that stocks with a strong cross-serial correlation have more similar exposure to common risk factors than stocks with a weak cross-serial correlation. The previous section revealed that strong cross-serial correlated stocks earn higher pairs trading returns. Although it is not a causal evidence, both findings combined suggest that pairs with stronger common links earn higher pairs trading returns. This relationship is important, as it supports the idea that gradual information diffusion is a major mispricing source that pairs trading exploits.

5.5 Slow information diffusion and trade types in pairs trading

The previous sections provided the theoretical foundation that portfolios that include pairs with a strong leader-follower relationship yield higher returns. I now stress this finding and examine whether each trade shows the typical price development that would be usual for a leader-follower relationship.

5.5.1 Leader-follower trades: theoretical foundation

The price development during the trade provides insights into the underlying economic reason why the two stocks deviate and converge. Analyzing the underlying price patterns helps to understand the economic source of the market inefficiency that allows pairs trading returns. For this purpose, I transfer an approach from a paper about mispricing of dual class shares, written by Schultz and Shieve (2010). They argue that asynchronous price movements arise if information is incorporated more quickly into the price of one share class than the other.

I argue that my close economic substitute pairs can be regarded as dual class shares in a broader sense. Hence, I classify trades that earn a positive return into three different types based on their price development. All three types are illustrated in Figure 7.

Figure 7: Types of return contribution

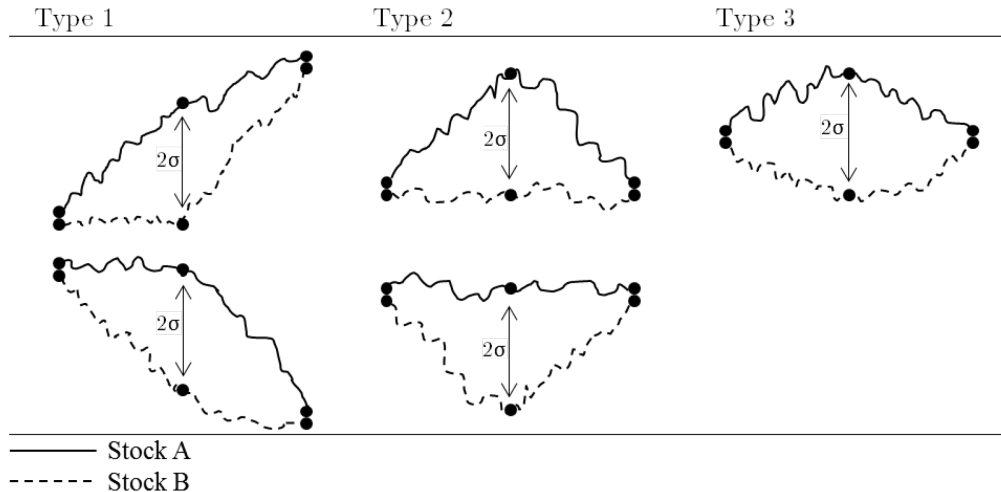


Figure 7 displays the three different characteristically types of trades that may occur during a trading period. Both stock A as well as stock B can be the initiator of a stock price deviation. Accordingly, the upper figure shows a variant in which stock A is the initiator and the lower figure shows a variant with stock B as initiator.

Type 1 (leader-follower): stock A significantly deviates from stock B. After the price difference exceeds 2σ and pairs trading positions are opened, the price of stock A continues to deviate in the same direction. However, stock B starts to move in the same direction after a while until both stock prices intersect. In this stock constellation, pairs trading earns a positive return if stock A declines or stock B rises after the initiation of the trade. Follower stock B earns at least 100% of the trade return. If stock A moves further away after the trade initiation, stock B overcompensates the loss of A and hence contributes over 100%. Likewise, stock B could also lead the trade.

Type 2 (single mispricing): Similar to type 1 trades, stock A significantly deviates from stock B. However, stock A moves in the opposite direction after pairs trading positions were opened until both stocks prices intersect again. The return of type 2 trades is mainly earned by the diverging stock that converges back afterward.

Type 3 (joint mispricing): Both stocks deviate from each other and positions are set as soon as the price difference exceeds 2σ . At a certain point, both positions reverse until they intersect. The return of type 3 trades is more or less equally earned by both stocks.

Type 1 trades are generated by leader-follower relationships of stocks, in particular if one stock incorporates common news faster than the other one. In contrast, type 2 and type 3 trades are generated if one (type 2) or both (type 3) stocks are temporarily mispriced or over-/underreacting.

5.5.2 Leader-follower trades: empirical evidence

To empirically investigate which source contributes most to the return of pairs trading, I analyze the distribution of trades across the three trade types. For each positive trade, I separately identify the return of stock A and stock B. Afterward, I determine whether it is a type 1, type 2, or type 3 trade. For this purpose, I compute how much each stock contributes individually to the total trade return. The stock that contributes more to the total return is defined as the dominant leg. I define a trade as a type 1 trade if the contribution exceeds 105%. If the dominant leg contributes between 77.5% and 105%, the trade is defined as type 2 trade. A trade is classified as a type 3 trade if the dominant leg contributes between 50% and 77.5%.

In the next step, I analyze the distribution of trades across the three trade types for the strong and the weak cross-serial correlation portfolio. As pairs in the strong cross-serial correlation portfolio have a stronger leader-follower relationship, I expect to observe more type 1 (leader-follower trades) for this portfolio. Table 14 shows the results.

Table 14: Distribution of trade types in strong and weak cross-serial correlation portfolios

Table 14 shows the percentage of pairs that are classified as trade Type 1 (Leader-Follower), Type 2 (Single mispricing), and Type 3 (Joint mispricing) differentiated by strong and weak cross-serial correlation. A trade is classified as Type 1, if the return of the follower stock contributes more than 105% to the total trade return (the follower stock overcompensates a loss from the leader stock in this case). The trade is classified as Type 2 (Type 3), if the return of the follower stock contributes more between 77.5% and 105% (50% - 77.5%) to the total trade return.

	% of contribution of follower stock to the total return	Strong cross-serial correlation	Weak cross-serial correlation
Type 1 (Leader-Follower)	> 105%	65.85%	45.79%
Type 2 (Single mispricing)	77.5% - 105%	16.66%	27.64%
Type 3 (Joint mispricing)	50% - 77.5%	17.49%	26.57%
Total number of trades		4826	2352

As expected, the percentage of leader-follower trades is very high with 65.85% in the portfolio with strong cross-serial correlated pairs. In contrast, only 45.79% of the trades in the low cross-serial correlation portfolio are leader-follower trades.

Moreover, the strong cross-serial correlation portfolio has twice as many trades as the weak cross-serial correlation portfolio. This finding is consistent with the finding of Section 4.3.3 that shows that portfolios with a high trading frequency have higher pairs trading returns.

As both portfolios were constructed based on their cross-serial correlation, the percentage of leader-follower trades might be disproportionately high (strong) and low (weak). Therefore, I provide an out-of-sample analysis with the percentage of type 1 trades for each volatility correlation double-sorted portfolio in Table 15.

Table 15: Percentage of leader-follower trades in volatility correlation double-sorted portfolios

Table 15 shows the percentage of trades that are classified as trade Type 1 (leader-follower) for each volatility correlation double-sorted portfolio. A trade is classified as leader-follower trade if the follower stock contributes more than 105% to the total return of the trade.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Vola_Q1	49.45%	48.33%	47.11%	46.11%	46.54%
Vola_Q2	50.24%	49.98%	48.81%	47.41%	46.76%
Vola_Q3	50.31%	49.89%	49.75%	48.40%	47.21%
Vola_Q4	50.35%	50.71%	50.26%	49.36%	48.37%
Vola_Q5	51.06%	51.26%	51.06%	50.64%	49.64%

Table 15 shows that trades of portfolios with highly volatile and less correlated pairs are classified as type 1 (leader-follower) trades in approximately 50% of all cases. The remaining trades are caused by type 2 or type 3 trades.

The most striking aspect that emerges from the data is that the percentage of leader-follower trades is relatively similar across correlation quintiles. This finding is fairly surprising. The SSD claims to be an ideal candidate to identify pairs that are close substitutes. These pairs reacted similarly to the arrival of news in the past. In this line of argument, traders expect more leader-follower trades for these pairs. However, Table 15 reveals that the percentage of leader-follower trades for highly correlated pairs hardly differs from less correlated pairs.

Finally and most notable, my portfolio with strong cross-serial correlated pairs has a significantly higher percentage of leader-follower trades with $\approx 15\%$ points more compared to the classic portfolios with the highest correlation (Corr_Q1). This finding suggests that cross-serial correlation is

more accurate in identifying promising pairs than mathematical correlation or the SSD. It substantiates the hypothesis that cross-serial correlation among stocks is more important than normal correlation, as it allows the trader to exploit gradual information diffusion-generated trading opportunities more precisely.

5.6 Persistent Leader-Follower Role Allocation

As previously introduced, trading on gradual information processing requires two links:

- Linking factor - so that both stocks are similarly affected by the arrival of common news.
- Deviation factor - there must be a certain factor or stock characteristic that drives the different speed of information processing. The deviation factor determines which stock moves first (leader stock) and which stock moves second (follower stock).

The existence of both factors are the result of a certain stock characteristic. These stock characteristics determine, for instance, stock liquidity or the number of analysts covering the stock, and usually remain relatively persistent over time. Hence, the role allocation of the leader and the follower stock is relatively constant over time. Frequent changes of the leader-follower roles indicate that a pair exploits an unsystematic lead-lag relationship that is presumably not created by gradual information processing. Therefore, evidence of a persistent role allocation further strengthens the hypothesis that gradual information diffusion is a major source of pairs trading returns.

To test the persistence of the leader-follower role allocation, I analyze pairs that trade multiple times within one trading period. A leader-follower trade is characterized by the observation that the majority of the return is earned by the follower stock. Hence, the follower stock is identified by identifying the stock with a higher proportional return contribution than the total trade return. I pick one stock (stock A) of each pair and count the number of trades in which this particular stock takes the leader role within one trading period.

Each trade of a pair is basically a Bernoulli experiment with two possible outcomes: “ $X=1$ ” if stock A is the leader stock in the particular trade and “ $X=0$ ” is otherwise. In the absence of a common link, it is equally likely that stock A is the leader or the follower stock. In this case, the probability p of the event is $P(X=1)=0.5$. If stock A is more likely to be the leader stock upon the arrival of new common information (because A is e.g. more liquid), then the probability should be $P(X=1)>0.5$. Accordingly, $P(X=1)<0.5$ if stock A is more likely to be the follower stock.

As each trade is a Bernoulli experiment, multiple trades of one pair within one trading period can be regarded as n repetitions of the Bernoulli experiment. Therefore, X should follow a Binomial distribution with the probability of $p=0.5$. However, if the leader-follower role allocation is persistent as a common factor and a constant deviation factor exists, then a statistical test should reject the hypothesis that the outcome $P(X=1 \mid p=0.5, n)$ follows a binomial distribution with a probability of $p=0.5$ for a pair that trades n times during one trading period.

I test whether the leader-follower relationship is persistent over time as follows:

- In the first step, I form five subsamples consisting of pairs that trade $n = 2, 3, 4, 5, 6$ times⁶ during one trading period. That allows me to test the hypothesis with five independent samples. The observation unit in this analysis is one pair in one trading period.
- Afterward, I apply an χ^2 goodness-of-fit test for a binomial distribution separately for each subsample. Thereby, I test whether the empirical joint distribution of the quantity, how often stock A leads, significantly differs from a theoretical binomial distribution with $p = 0.5$ and n observations.
- If H_0 cannot be rejected, no clear follower-leader relationship exists. In this case, I conclude that there is no permanent common factor and deviation factor that determines the leader-follower

⁶ A test with more than six trades within a trading period cannot be conducted, as the number of observations is insufficiently small.

role allocation. If I find more observations in the tails compared to the theoretical distribution and reject H_0 , I conclude that stock A either more frequently leads (high number of realized events) or follows (low number of realized events).

Let's have a closer look at the results for all pairs that trade exactly four times during one trading period to exemplify the analysis. The sample includes $n_4 = 981$ pairs. The probability of stock A to be leading in zero out of five trades is 6.25% with $P(X = 0 | p = 0.5, n = 4) = \binom{4}{0} 0.5^0 * (1 - 0.5)^{4-0} = 0.0625$. If there is no clear leader-follower role allocation, statistically stock A leads in zero trades in 61.3 out of the 981 cases.

Table 16 lists the expected outcomes in case there is no permanent role allocation.

Table 16: Theoretical binomial distribution in case of a non-permanent role allocation

Table 16 reports the theoretical number of cases where stock A is the leader if the role allocation of stock A and stock B is non-permanent. The observation unit is one pair that trades exactly n times during one trading period. Table 16 shows the results only for pairs that trade exactly four times. In total, 981 pairs traded exactly four times, so my sample comprises 981 observations. I have analogously computed the results for four further samples where pairs trade exactly two times, three times, five times or six times during one trading period.

The first column of the table shows X , the number of trades, where stock A is the leader stock. The second column shows the theoretical binomial distribution: $P(X|p, n)$. p is the theoretical probability that stock A is the leader stock in a Bernoulli experiment. If we assume a non-permanent role allocation, $p = 0.5$ as it is equally likely that stock A is the leader stock as that stock B is the leader. n stands for the total number of trades within a trading period of the particular pairs and is 4 as Table 16 only includes pairs that trade four times during one trading period. The last column shows the theoretical number of cases in which stock A should be the leader stock in X out of four trades, if the role allocation between stock A and stock B is not permanent.

X (# trades, where A is the leader stock)		Probability	Cases for $n_4 = 981$
0	$P(X=0 p=0.5, n=4)$	0.0625	61.3
1	$P(X=1 p=0.5, n=4)$	0.25	245.3
2	$P(X=2 p=0.5, n=4)$	0.375	367.9
3	$P(X=3 p=0.5, n=4)$	0.25	245.3
4	$P(X=4 p=0.5, n=4)$	0.0625	61.3

In the next step, I conduct a chi-squared test to test whether the observed distribution follows the theoretical binominal distribution. Table 17 reports the results of the chi-squared test.

Table 17: χ^2 Goodness-of-fit test for binomial distribution

Table 17 reports the chi-squared statistic for pairs that trade $n = 2, 3, 4, 5$ or 6 times during one trading period. The chi-squared statistic indicates whether to reject the hypothesis that the observed distribution follows the expected theoretical binomial distribution. In this table, the chi-squared goodness-of-fit test compares the observed number of cases where stock A was the leader stock with the theoretical number of cases if the role allocation of stock A and B is not permanent. Table 16 above shows an example of the theoretical distribution for $n = 4$. Statistical significance that the hypothesis of a random leader-follower role allocation in the sample can be rejected at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

n	2	3	4	5	6
χ^2	2313.32***	744.84***	193.35***	51.21***	56.35***

The chi-squared goodness-of-fit test statistic is 193.35 for $n = 4$ with a p-value of 0.000. Hence, H_0 is rejected. H_0 can also be rejected for the other four samples where pairs trade exactly $i = 2, 3, 5$ and 6 times within one particular trading period. This result suggests that the probability that the role allocation of the leader and the follower stock is not random.

A closer look to the tails reveals further information. In this test design, more observations in the tails compared to the theoretical binomial distribution means that for a given pair, it is more likely that stock A is either always or never the follower stock. This is exactly what I expect to find if a pair has a persistent leader-follower relationship. For $n=4$, I find that the observed distribution for $X=0$ (stock A is always the follower) and $X=4$ (stock A is always the leader) is indeed higher than the theoretical binomial distribution. This result indicates that the distribution has fat tails or, in other words, that stock A is more often always the leader stock or always the follower stock. This result proves that the leader-follower role allocation is persistent in this sample.

5.7 Conclusion

This section investigated the role of gradual information diffusion in pairs trading to establish gradual information diffusion as an important return source for the trading strategy. Pairs must share a common factor so that a new piece of information affects both stocks. Furthermore, a deviation factor ensures that both stocks process the relevant information at a different speed. The time lag in information processing creates temporary mispricing that is exploited by pairs trading. The leader-follower relationship of the pair exploits the arbitrage opportunity. Pairs trading is a natural candidate with some appealing features that overcome downsides of classic strategies to exploit gradual information diffusion.

My first empirical analysis of this section demonstrated that pairs trading with a strong leader-follower relationship, approximated by the cross-serial correlation, is more successful than trading stocks with a weak leader-follower relationship. Most notably, pairs trading with strong cross-serial correlated pairs is more profitable than any of the pairs selected based on the classical Distance Measure or their volatility correlation profile as shown in Chapter 4. Furthermore, my results provide evidence that these pairs are close economic substitutes, as they react more similarly to common risk factors than other stock pairs do. The next results demonstrated that the majority of trades of strong cross-serial correlated pairs follow a typical price development that is typical for a leader-follower relationship. The percentage of these leader-follower trades is significantly smaller for weakly cross-serial correlated pairs. The final analysis provided evidence that the leader-follower role allocation of a pair is persistent. This finding is consistent with a stable deviation factor as a condition to successful gradual information diffusion. Altogether, these findings provide significant evidence that strengthens the idea of gradual information diffusion as an important source of pairs trading returns. The study contributes to prove that pairs trading returns are not the random product

of a technical stock analysis. Rather, the source of pairs trading returns can be economically explained by the exploitation of market inefficiencies created by gradual information diffusion.

6. Core chapter 3: idiosyncratic risk, costly arbitrage, and asymmetry: evidence from pairs trading

The last two chapters analyzed the mechanics and the return source of pairs trading. In the following chapter I use pairs trading as an alternative measure for mispricing to investigate the role of idiosyncratic risk on the return.

6.1 Motivation

Can we earn a risk premium for bearing idiosyncratic risk? Consistent with the capital asset pricing model, early studies find no relationship (Fama and MacBeth (1973) and Bali and Cakici (2008)) between returns and idiosyncratic risk, measured by idiosyncratic volatility (IVOL). However, some papers reveal a positive idiosyncratic volatility effect. High idiosyncratic volatility stocks outperform low idiosyncratic volatility stocks. The positive idiosyncratic volatility effect is explained by investors' demand for a risk premium or the impossibility of holding perfectly diversified portfolios (Merton (1987)). Further empirical evidence of a positive idiosyncratic volatility effect is provided by, among others, Malkiel and Xu (2002), Wurgler and Zhuravskaya (2002), Goyal and Santa-Clara (2003), and Fu (2009).

In contrast, Ang, Hodrick, Xing, and Zhang (2006, 2009) document a negative idiosyncratic volatility effect in the US and in an international sample. As the literature disagrees on whether idiosyncratic volatility has a negative or a positive effect on average returns, and empirical findings are contradictory, a growing body of literature, including this thesis, now strains to learn more about the economics underlying the idiosyncratic volatility puzzle.

Hou and Loh (2016) provide a comprehensive overview of different explanation approaches and compare their explanation power. They find that approaches based on expected idiosyncratic skewness (Boyer, Mitton, and Vorkink (2010)), maximum daily return (Bali, Cakici, and Whitelaw

(2011)), retail trading proportion (Han and Kumar (2013)), one-month return reversals (Fu (2009); Huang, Liu, Rhee, and Zhang (2009)), bid-ask spread (Han and Lesmond (2011)), and past earnings surprises (Jiang, Xu, and Yao (2009)) are promising candidates to explain a large part of the puzzle.

I concentrate on the idiosyncratic volatility effect among mispriced stocks and differentiate between under- and overpriced stocks, similar to Stambaugh, Yu, and Yuan (2015) and Cao and Han (2016). Both papers find a positive idiosyncratic volatility effect among underpriced stocks and a negative idiosyncratic volatility effect among overpriced stocks, caused by a combination of arbitrage risk and arbitrage asymmetry. I extend their findings with a novel mispricing measure and broadly confirm their findings in the US stock market between January 1991 and December 2014. Therefore, I utilize the established relative value arbitrage strategy pairs trading, introduced by Gatev, Goetzmann, and Rouwenhorst (2006), to measure a different kind of stock mispricing. Most important, my analysis provides insights into whether the type of mispricing influences the direction of the idiosyncratic volatility effect. This will be explained in more detail in the next section.

6.2 Motivation to apply an alternative mispricing measure

Stambaugh, Yu, and Yuan (2015) and Cao and Han (2016) approximate the extent of mispricing for each stock with a composite mispricing index, which combines a stock's individual ranks among selected anomaly strategies that exploit mispricing. Stambaugh, Yu, and Yuan (2015) include financial distress, O-score bankruptcy probability, net stock issues, composite equity issues, total accruals, net operating assets, momentum, gross profitability, asset growth, return on assets, and investment-to-assets, while Cao and Han (2016) consider short-term return reversals, size, book-to-market, and momentum.

Both mispricing rankings are mostly based on firm characteristics and accounting data, which are updated on a monthly or quarterly basis. The index ranks a stock in each of the anomaly strategies

relative to the stock universe. However, the extent of mispricing between two adjacent ranks differs within and across the eleven categories because of the rankings' ordinal nature. Stocks that are correctly priced according to the composite index can still be temporary mispriced. Furthermore, the combination of mispricing categories indicates that moderate mispricing in some categories can outweigh extreme mispricing in another category, or that specific mispricing information is overrated as it affects more than one anomaly ranking.

In contrast, pairs trading measures the aggregated mispricing of one stock relative to a historical close economic substitute. I adopt the rule-based pairs trading strategy of Gatev, Goetzmann, and Rouwenhorst (2006). The strategy is based on the idea that price deviations of two close economic substitute stocks represent temporary mispricing that dissolves in the near future. Market- and trading-related mispricing, like gradual information diffusion and short-term reversals, induce temporary price divergences that are responsible for the return of pairs trading. Although my mispricing measure is positively correlated with Stambaugh, Yu, and Yuan's composite index in an empirical test, the correlation is rather weak and often insignificant. I conclude that my mispricing proxy measures a different kind of mispricing, which allows me to challenge the outlined research hypotheses in a novel test setting.

The analyses in this chapter are related to several ongoing academic discussions. It contributes additional empirical evidence to establish the idiosyncratic volatility effect in a novel test setting with an alternative mispricing measure. Moreover, I introduce pairs trading as a successful investment strategy to exploit the idiosyncratic volatility effect. My study is also related to literature on gradual information diffusion and lead-lag relationships, as both concepts potentially generate the initial mispricing required for pairs trading.

6.3 Modified test design

For the following analyses in this chapter, I use the data set and the pairs trading algorithm as previously described in chapter 2. The previous results revealed that the traditional pairs trading algorithm picks pairs with low volatility. To investigate the returns for different idiosyncratic volatility levels, I therefore slightly modify my procedure as follows:

First, in addition to “normal” volatility, I also determine each stock’s idiosyncratic volatility during the identification period. Idiosyncratic volatility is the volatility of the residuals, extracted from a time series regression of daily stock returns on Fama and French’s (1993) three-factor model, following Ang, Hodrick, Xing, and Zhang (2006). Afterward, I compute a proxy for the combined idiosyncratic volatility level of both stocks that equals the sum of both idiosyncratic pair volatilities ($\sigma_{IVOL,A} + \sigma_{IVOL,B}$). Thereafter, I assign each possible pair combination to one of five pair correlations⁷ and five pair idiosyncratic volatility quintiles. Quintile Corr_Q1 (Corr_Q5) includes the pairs with the lowest (highest) pair correlation, while quintile IVOL_Q1 (IVOL_Q5) includes the pairs with the lowest (highest) level of pair idiosyncratic volatility. Please refer to Section 6.3 for further details. I update the pair group affiliation in each identification period. In the final step, I arbitrarily select twenty pairs out of each of the twenty-five pair groups double-sorted on the five pair idiosyncratic volatility and five pair correlation quintile groups.

Table 18, panel A reports the median pair correlation and idiosyncratic pair volatility per quintile during the identification period between January 1990 and June 2014.

Table 18: Correlation and idiosyncratic volatility levels over quintiles

Table 18 reports the median correlation, the median pair volatility, and the median idiosyncratic pair volatility (IVOL) per quintile for all selected pairs during all identification periods between

⁷ The correlation is calculated based on the normal volatility, not the idiosyncratic volatility.

January 1990 and June 2014. The (idiosyncratic) volatility of a pair is defined as sum of stock A's and B's (idiosyncratic) volatility. IVOL is defined as volatility of the residuals from regressing the daily returns on Fama and French's (1993) three factor model as in Ang et al. (2006).

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
Correlation	-0.4885	-0.0471	0.2973	0.5738	0.8026
	IVOL_Q1	IVOL_Q2	IVOL_Q3	IVOL_Q4	IVOL_Q5
Volatility	0.0130	0.0271	0.0435	0.0728	0.1930
IVOL	0.0008	0.0029	0.0120	0.0523	0.2546

For instance, the median pair correlation of all Corr_Q1 pairs is -0.4885, and in IVOL_Q1, the median pair volatility is 0.0130 while the median pair idiosyncratic volatility is 0.0008. Please note that the pair volatility slightly differs from the reported median pair volatility in Table 2, as the quintiles are now formed based on idiosyncratic volatility, not on volatility. The median pair correlation ranges from -.4885 to 0.8026 and includes correlated (Corr_Q3-Corr_Q5), almost uncorrelated (Corr_Q2), and negatively correlated pairs (Corr_Q1). The median idiosyncratic pair volatility ranges from 0.0008 to 0.2546.

6.4 Measuring mispricing with pairs trading

A key aspect of Stambaugh, Yu, and Yuan (2015) and Cao and Han's (2016) approach to solving the idiosyncratic volatility puzzle is their method for measuring mispricing. Different methods

measure different forms of mispricing and capture mispricing static or dynamic over time. This study contributes to the existing literature by confirming the idiosyncratic volatility effect among mispriced stocks with an alternative mispricing measure. My results clarify whether the type of mispricing, which influences the level of idiosyncratic volatility, is the driver behind the idiosyncratic volatility effect.

6.4.1 Pairs trading as an alternative mispricing measure

I start with a closer look at mispricing and compare how my approach to evaluating mispricing differs from Stambaugh, Yu, and Yuan (2015) and Cao and Han (2016). Both papers rank each stock in eleven (Stambaugh, Yu, and Yuan) or four (Cao and Han) anomaly categories and compute a composite mispricing index based thereupon. Each anomaly ranking is ordinal, so differences in index values between two adjacent ranks cannot be compared in a meaningful way 1) between any two ranks within an anomaly ranking and 2) across anomaly categories.

Let me give you two examples to explain these problems in detail:

- Limitation 1: I rank four stocks (total sample size: 1,000 stocks) according to their mispricing of a random anomaly. Stock A (500th rank) is 1% less mispriced than stock B (501th rank) with a rank difference of one. Also with a rank difference of one, stock C (999th rank) is 5% less mispriced than stock D (1,000th rank). Although stock D is far more mispriced to stock C than stock B to A, the composite index considers a one rank difference in both cases.
- Limitation 2: The difference in mispricing between rank 500 and 501 is 1% in anomaly No. 1. However, it is possible that the difference in mispricing between rank 500 and 501 is 5% in anomaly No. 2. The composite index only considers the one-unit rank difference. The mispricing of stocks that are extremely mispriced in one category might be undervalued if the stock obtains medium ranks in all other categories.

These limitations particularly affect the ranking's tails, which are especially interesting for the comparison between the most underpriced and overpriced quintiles. The mispricing rankings are mainly determined by accounting data that are updated on a monthly or quarterly basis. The index technically aggregates information about company and stock characteristics that reflect a stock's probability to be mispriced. However, it fails to reflect the current actual mispricing level.

I now turn to the question of how pairs trading measures a different kind of mispricing. Pairs trading overcomes the previously listed limitations. The strategy measures mispricing more accurately, as it does not use an ordinal ranking. Moreover, pairs trading measures mispricing more dynamically than the composite index. For instance, two stocks deviate strongly from each other but converge soon afterward. The mispricing is only temporary, so the stocks' composite mispricing index ranking remains unchanged. However, pairs trading allows one to exploit the temporary mispricing.

In this next section I compare the empirical correlation between Stambaugh, Yu, and Yuan's (2015) mispricing measure and my proposal to measure mispricing – pairs trading.

6.4.2 Empirical correlation between the different methods to measure mispricing

A high rank in Stambaugh, Yu, and Yuan's (2015) mispricing ranking indicates stronger mispricing. Likewise, higher pairs' trading returns identify higher mispricing levels. To test whether both mispricing measures are statistically correlated, I specifically test whether pairs that are classified as highly mispriced by Stambaugh, Yu, and Yuan (2015) also earn higher pairs trading returns. If this is the case, I conclude that both approaches are correlated. As some stock characteristics influence the information processing speed as well as the composite mispricing ranking, I expect a positive correlation. However, I expect a weak connection because pairs trading measures a different, more dynamic type of mispricing.

The analysis is executed as follows:

- In the first step, I merge Stambaugh, Yu, and Yuan's mispricing index with my dataset to get a combined sample with stocks that are included in Stambaugh, Yu, and Yuan's, but also in my sample. I exclude pairs whose stocks are classified into different Stambaugh, Yu, and Yuan mispricing deciles because I cannot determine whether the overall mispricing of the pair is high or low. Furthermore, I am only interested in pairs with extremely high or extremely low mispricing in the following analysis. A pair consisting of differently mispriced stocks is always in the middle mispricing range.
- In the next step, I form mispricing deciles (1st decile = lowest mispricing, 10th decile = highest mispricing) based on Stambaugh, Yu, and Yuan's total ranking.
- Then, I calculate the pairs trading returns separately for the 1st and 10th mispricing decile trading portfolio. Both mispricing measures are positively correlated if Stambaugh, Yu, and Yuan's 10th decile pairs trading portfolio outperforms the 1st decile pairs trading portfolio.
- As stock pair volatility and correlation influence pairs trading returns, I control for both stock characteristics in this step. Therefore, I separately classify all stocks into one of the ten mispricing deciles 1) based on Stambaugh, Yu, and Yuan's mispricing measure and 2) into one of my twenty-five volatility correlation double-sorted portfolios.

A pair that is classified into the 1st decile comprises two stocks that are individually ranked among the top 10% of least mispriced stocks.

- Afterward, I compute the pairs trading returns for the 1st and the 10th mispricing decile for each of the twenty-five double-sorted portfolios. Within each portfolio, we select the twenty pairs with the lowest SSD for trading to further assimilate both portfolios in terms of volatility and correlation. Overall, I get the time series of 50 pairs of trading portfolios (the 1st and 10th Stambaugh, Yu, and Yuan mispricing decile for each of the twenty-five double-sorted pairs trading portfolios).

- In the next step, I regress the monthly returns of my 50 portfolios on Fama and French's (1993) three factors, the momentum factor, the short-term reversal factor and Pastor and Stambaugh's (2003) liquidity factor.
- Finally, I compare the risk-adjusted returns between Stambaugh, Yu, and Yuan's 1st and their 10th mispricing decile for each of the twenty-five portfolios. Both measures are positively correlated if the pairs trading return of the 10th portfolio (Stambaugh, Yu, and Yuan: highest mispricing) is significantly higher than the pairs trading return of the 1st portfolio (Stambaugh, Yu, and Yuan: lowest mispricing).

Table 19 reports the risk-adjusted pairs trading returns and the return from buying the 10th decile portfolio and selling the 1st decile portfolio.

Table 19: Correlation between pairs trading and Stambaugh, Yu, and Yuan's (2015) mispricing measure

Table 19 reports the risk-adjusted pairs trading returns and the return from buying the 10th decile portfolio and selling the 1st decile portfolio. H-L (high minus low mispricing) shows the return from buying the 10th decile portfolio and selling the 1st decile portfolio. P_val reports the p-values from whether the H-L return is significantly greater than 0, in other words whether the returns of both portfolios significantly differ. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

		Correlation Quintiles					
Volatility Quintiles	SY Y classification	Q1 (low)	Q2	Q3	Q4	Q5 (high)	
Q1 (low)	10 th High	0.0039***	0.0051***	0.0042***	0.0028***	0.0030***	
	1 st Low	0.0044***	0.0032***	0.0025***	0.0021***	0.0020***	
	H-L	-0.0005	0.0478	0.0017*	0.0007	0.0010	
	p_val	(0.7969)	(0.2756)	(0.0915)	(0.388)	(0.1548)	
Q2	10 th High	0.0089***	0.0079***	0.0072***	0.0036***	0.0047***	
	1 st Low	0.0043***	0.0037**	0.0040***	0.0038***	0.0034***	
	H-L	0.0046**	0.0042**	0.0032*	-0.0002	0.0013*	
	p_val	(0.0459)	(0.0910)	(0.0529)	(0.8751)	(0.0998)	
Q3	10 th High	0.0070***	0.0076***	0.0082***	0.0086***	0.0046***	
	1 st Low	0.0045*	0.0020	0.0032***	0.0046***	0.0033***	
	H-L	0.0024	0.0056**	0.0050**	0.0039*	0.0013	
	p_val	(0.4266)	(0.0389)	(0.0164)	(0.0547)	(0.1739)	
Q4	10 th High	0.0075**	0.0076***	0.0068***	0.0041**	0.0059***	
	1 st Low	0.0019	0.0029*	0.0029**	0.0025*	0.0033***	
	H-L	0.0056*	0.0047*	0.0039	0.0016	0.0026**	
	p_val	(0.0833)	(0.0919)	(0.1146)	(0.4403)	(0.0307)	
Q5 (high)	10 th High	0.0071**	0.0048***	0.0054*	0.0070***	0.0065***	
	1 st Low	0.0026***	0.0012	0.0005	0.0021	0.0058***	
	H-L	0.0045	0.0036	0.0049	0.0049	0.0007	
	p_val	(0.1818)	(0.3368)	(0.1332)	(0.1085)	(0.7227)	

The difference is positive in twenty-three out of twenty-five portfolios, which indicates that pairs that are classified as mispriced by Stambaugh, Yu, and Yuan (2015) also earn higher pairs trading returns. However, the difference is only significant in eleven out of twenty-five portfolios. I conclude that pairs trading successfully identifies a broad type of mispricing, which, however, differs from Stambaugh, Yu, and Yuan's mispricing definition to some extent.

6.5 Idiosyncratic risk, mispricing, and pairs trading returns

I would now like to address the question of whether high idiosyncratic risk is compensated.

Standard asset pricing theory reasons that idiosyncratic risk can theoretically be diversified and is therefore not compensated. However, recent papers find mixed evidence. Consistent with the capital asset pricing model, Fama and MacBeth (1973) and Bali and Cakici (2008) find no relationship between idiosyncratic risk and returns.

Pontiff (2006) disagrees and emphasizes the role of idiosyncratic risk as practically non-diversifiable arbitrage risk in the form of holding costs that force investors to avoid or limit these risky positions. Arbitrage risk includes, among others, market liquidity risk, funding liquidity risk and trading costs (e.g. DeLong, Shleifer, Summers, and Waldmann (1990), Shleifer and Vishny (1997), Mitchell, Pulvino, and Stafford (2002)). Lehmann (1990), Goyal and Santa-Clara (2003), and Fu (2009) justify a positive relationship with investors' demand for a risk premium and the impossibility of holding perfectly diversified portfolios. In a seminar paper, Ang, Hodrick, Xing, and Zhang (2006, 2009) document negative idiosyncratic volatility effects in the US market and in international samples. Hou and Loh (2016) provide a comprehensive overview of different explanation approaches to explain the idiosyncratic volatility puzzle and compare their power to explain it. They find that approaches based on expected idiosyncratic skewness (Boyer, Mitton, and Vorkink (2010)), maximum daily return (Bali, Cakici, and Whitelaw (2011)), retail trading proportion (Han and Kumar

(2013)), one-month return reversals (Fu (2009); Huang, Liu, Rhee, and Zhang (2009)), bid-ask spread (Han and Lesmond (2011)), and past earnings surprises (Jiang, Xu, and Yao (2009)) are promising candidates to explain a large part of the idiosyncratic volatility puzzle.

Overall, higher levels of idiosyncratic risk impose higher arbitrage risk and hence stronger mispricing (Doukas, Kim, and Pantzalis (2010)). Investors who are willing to bear arbitrage risk or are less exposed to it must earn higher compensation for purchasing or short selling mispriced stocks.

6.5.1 Hypothesis 4 - Idiosyncratic risk puzzle and pairs trading returns

Stambaugh, Yu, and Yuan's (2015) and Cao and Han's (2016) provide a differentiated explanation approach for the different findings. I follow their approach and apply their line of arguments on pairs trading. Stocks with high idiosyncratic risk face higher arbitrage costs that deter market participants from exploiting mispricing. As investors hold smaller proportions of stocks with high arbitrage risk (Pontiff (2006)), a sudden revision of one investor's expectation causes only slight price pressure. Hence, mispricing persists. The return potential among strongly mispriced stocks is naturally higher than among less mispriced stocks. As a consequence, anomaly effects are often stronger among stocks with high idiosyncratic volatility. For instance, Ali, Hwang, and Trombley (2003) report that high idiosyncratic volatility stocks have a stronger book-to-market effect. McLean (2010) cannot confirm the effect for the momentum strategy, but finds prevalent long-term reversals only among high idiosyncratic volatility stocks. The higher anomaly effect among stocks with high idiosyncratic volatility is a promising opportunity for investors who are willing to bear idiosyncratic risk or are less exposed to it. For instance, some investors might be less exposed to arbitrage risk if they already hold the stock that pairs trading suggests to sell short, as they are unaffected by the risk of margin calls.

While the stock's mispricing is gradually reduced, the price and the return of an underpriced stock

rises and the price and the return of an overpriced stock falls. Enough active arbitrageurs in the market push prices back to their fundamental value so that underpriced stocks rise and overpriced stocks fall. In this setting, informed investors fill the position of rational arbitrageurs. They bear the risk that noise traders drive prices further away from their fair value (DeLong, Shleifer, Summers, and Waldmann (1990)). This leads to a positive relation between idiosyncratic risk and expected returns among underpriced stocks and a negative relation between idiosyncratic risk and expected returns among overpriced stocks.

I now transfer the line of arguments to pairs trading. Does idiosyncratic volatility also increase the type of mispricing that pairs trading exploits? Highly volatile stocks incorporate information more slowly into prices (e.g., Mech (1993), Chorida and Swaminathan (2000)). The price adjustment to new information is therefore slower, which imposes additional risk and creates more mispricing. Slower information diffusion leads to wider price spreads, which increases the return per trade. Furthermore, sufficient price spreads for pairs trading appear more frequently, which also increases the trading frequency, and hence, the return.

In my first research hypothesis, I combine pairs trading with the idiosyncratic volatility theory. Pairs trading is technically a long-short strategy that holds a long position in the underpriced stock and a short position in the overpriced stock. I expect to observe a positive idiosyncratic volatility effect among underpriced and a negative idiosyncratic volatility effect among overpriced stocks. The long position participates directly from the performance of the underpriced stock that realizes a positive return. The short position participates inversely in the performance of the overpriced stock that realizes a negative return. Hence, a pairs trader profits from the positive idiosyncratic volatility effect among underpriced stocks and additionally from the inverse negative effect among overpriced stocks. Overall, I expect higher pairs trading returns for higher idiosyncratic volatility quintiles.

Research hypothesis No. 4: Portfolios that include high idiosyncratic volatility stock pairs earn higher pairs trading returns compared to portfolios that include lower idiosyncratic volatility pairs.

6.5.2 Empirical evidence

This section takes a closer look at the monthly pairs trading returns to assess research proposition No. 4. Table 20, panel A reports the average monthly raw returns of each double-sorted trading portfolio on pair idiosyncratic volatility and correlation between January 1991 and December 2014. Panel B shows the alphas of these regressions for each volatility correlation-sorted portfolio.

Table 20: Average monthly trading return

Table 20, panel A reports the average monthly raw returns of each double-sorted trading portfolio on pair idiosyncratic volatility and correlation between January 1991 and December 2014. To account for well-established return patterns, I regress monthly excess returns on a six-factor model. The model includes Fama and French's three factors, a momentum factor, and a short-term reversal factor. It is possible that some illiquid stocks in the highest idiosyncratic volatility quintile bias the portfolio returns, as they might include a liquidity premium. I therefore extend the factor model with Pastor and Stambaugh's (2003) liquidity factor to control for a liquidity premium. Panel B reports the alphas of these regressions for each volatility correlation-sorted portfolio. Like related papers, I use Newey-West (1987) standard errors with lag six. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

Panel A: Pairs trading profitability in calendar time

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
IVOL_Q1	0.0033***	0.0031**	0.0026***	0.0052***	0.0046***
IVOL_Q2	0.0072***	0.0033**	0.0045***	0.0067***	0.0047***
IVOL_Q3	0.0080***	0.0054**	0.0045**	0.0081***	0.0071***
IVOL_Q4	0.0057**	0.0064**	0.0051**	0.0080***	0.0088***
IVOL_Q5	0.0037*	0.0009	0.0030	0.0061***	0.0052***

Panel B: Alphas of each correlation and idiosyncratic volatility quintile portfolio

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
IVOL_Q1	0.0031*	0.0031***	0.0026***	0.0049***	0.0046***
IVOL_Q2	0.0072**	0.0029**	0.0039***	0.0064***	0.0044***
IVOL_Q3	0.0078**	0.0045***	0.0057***	0.0076***	0.0069***
IVOL_Q4	0.0053*	0.0061**	0.0057*	0.0081***	0.0091***
IVOL_Q5	0.0035	-0.0003	0.0024	0.0058**	0.0044**

The returns of twenty-three out of twenty-five portfolios are positive, significant, and they range between 26 bp and 88 bp. For instance, portfolio IVOL_Q1/Corr_Q5 (low idiosyncratic volatility/high correlation) yields an average monthly return of 46 bp. In line with research hypothesis no. 4, the return increases for higher levels of pair idiosyncratic volatility, however, not monotonically as in Cao and Han (2016). Surprisingly, the highest returns within each correlation quintile are earned by the second highest idiosyncratic volatility pair quintile IVOL_Q4. The inter-quintile difference between IVOL_Q4 and IVOL_Q1 is positive for each correlation quintile and ranges between 25 bp and 42 bp. Stocks with higher idiosyncratic volatility levels as in IVOL_Q5 bear higher arbitrage risk and should therefore realize higher returns. Panel B shows that all alphas are positive and significant for twenty-two out of twenty-five portfolios. Consistent with research hypothesis no. 4 and with the results in panel A, the risk-adjusted returns increase for higher levels of pair idiosyncratic volatility.

I previously justified my portfolio sort by correlation with the control for diversification effects. As pairs trading is a long-short strategy, a portfolio with positively correlated pairs is relatively diversified, whereas negatively correlated pairs amplify risk. This effect is especially strong in the risky, high idiosyncratic volatility quintiles. Within each idiosyncratic volatility quintile, the return mon-

otonically increases for stronger (positive) pair correlation quintiles. Corr_Q2 (median pair correlation -0.0471), which includes uncorrelated pairs, is outperformed by Corr_Q5 (median pair correlation 0.8026) with highly correlated pairs. Yet, it is surprising that within IVOL_Q2 and IVOL_Q3, the Corr_Q1 portfolio, which includes negatively correlated pairs, outperforms more diversified portfolios. However, the inter-quintile differences are not significant.

Altogether, my results confirm research hypothesis no. 4. Higher idiosyncratic volatility induces more mispricing and allows higher returns, consistent with prior findings from various alternative investment strategies.

6.6 Arbitrage asymmetry and the idiosyncratic volatility effect in pairs trading

In the next step, I extend my test setting and introduce arbitrage asymmetry. Arbitrage asymmetry proclaims that mispricing is stronger among overpriced stocks than among underpriced stocks at the same idiosyncratic volatility level because of short-selling restrictions. While buying a stock is easy, the action of short selling is far more complicated. This leads to several implications that impose additional arbitrage costs, but only on overpriced stocks. In this section, I first explain why the idiosyncratic volatility effect among overpriced stocks is higher than among underpriced stocks, while the idiosyncratic volatility effect is equally strong for low idiosyncratic volatility stocks. Afterward, I examine this hypothesis empirically.

6.6.1 Hypothesis 5 - Idiosyncratic volatility effect and arbitrage asymmetry

In a seminal paper, Miller (1977) hypothesizes that divergent opinions about the intrinsic value exert upward price pressure on short-selling restricted stocks. Empirical evidence is provided, among others, by Figlewski (1981); Chen, Hong, and Stein (2002); and Boehme, Danielsen, and Sorescu (2006). Stambaugh, Yu, and Yuan (2015) note that short selling impediments exist at the stock level, but also on the investor level. D'Avolio (2002) provides a comprehensive overview of factors that

influence the ability and the costs to short a stock. For example, short selling is more expensive and overpricing more likely to persist for stocks with low institutional ownership due to the lack of sophisticated stock owners who sell or supply stocks (Nagel (2005)). Additionally, fund managers can be constrained by investment policy restrictions or have incentives to avoid short selling. Almazan, Brown, Carlson, and Chapman (2004) report that only 32.9% of 1,838 polled funds in 2,000 were permitted to short, whereas only 9.7% in fact shorted stocks. Overpricing can also be increased by firms that thwart short selling by legal threats, physical stock certificates, and lawsuits (Lamont (2012)). Also, Stambaugh, Yu, and Yuan (2012) find that mispricing in the short leg (overpriced stocks) is higher compared to the long leg when investors' sentiment is high. Together, these studies provide convincing evidence that investors are confronted with higher arbitrage costs for exploiting overpriced stocks. Mispricing, and hence, arbitrage returns should therefore be higher among overpriced stocks compared to underpriced stocks.

With arbitrage asymmetry in their mind, Stambaugh, Yu, and Yuan (2015) reason that the aggregated idiosyncratic volatility effect is negative in a sample that is equally composed of over- and underpriced stocks. The negative idiosyncratic volatility effect of overpriced stocks dominates the positive idiosyncratic volatility effect among underpriced stocks because of arbitrage asymmetry. As previously explained, pairs trading profits from a positive idiosyncratic volatility effect in both legs, so I cannot expect to find a negative effect in the overall sample. To overcome this difficulty, I test arbitrage asymmetry by computing the return per trade separately for the long leg (underpriced stock) and the short leg (overpriced stock) in the first step. Afterward, I determine the percentage contribution of the long leg to the total return per trade (contribution ratio). The contribution ratio reveals whether the effect of underpriced stocks is stronger than among overpriced stocks.

As arbitrage risk is low among stocks with low idiosyncratic volatility, sufficient arbitrageurs in the market should secure efficient prices. Among low idiosyncratic volatility stocks, I therefore expect

a contribution ratio around 50%, which indicates that both legs contribute equally to the total return. On the contrary, mispricing among high idiosyncratic volatility stocks should be stronger especially in the short leg. I therefore expect to find a contribution ratio in the long leg far below 50%, which indicates that the short leg contributes more.

Research hypothesis no. 5: The short leg contributes more to the return per trade for pairs with high idiosyncratic volatility, whereas both legs contribute on average the same return among low idiosyncratic volatility stocks.

6.6.2 Empirical evidence

The fifth research proposition argues that the idiosyncratic volatility effect of the overpriced stock (short leg) dominates the idiosyncratic volatility effect compared to the underpriced stock (long leg) because of arbitrage asymmetry. Overpriced stocks are affected by additional arbitrage costs that impede short-leg arbitrage trades.

In section 4.7.2, Table 11, I demonstrated that the long leg also contributes to the total return in portfolios clustered by “normal” volatility. I repeat the analysis with portfolios clustered by idiosyncratic volatility.

Table 21 shows the median contribution of the long leg to the total return per trade between January 1991 and December 2014.

Table 21: Return contribution of the long leg to the total return

Table 21 shows the median contribution of the long leg to the total return per trade between January 1991 and December 2014. Each trade consists of a long position (long leg) in the loser stock and a

short position (short leg) in the winner stock. The contribution ratio of the long leg in this table is calculated as in formula (18) in section 4.7.2: $\frac{\text{long leg return}}{\text{total return per trade}}$.

	Corr_Q1	Corr_Q2	Corr_Q3	Corr_Q4	Corr_Q5
IVOL_Q1	51.02%	50.05%	50.75%	49.96%	50.41%
IVOL_Q2	49.15%	49.58%	50.11%	49.92%	50.16%
IVOL_Q3	49.13%	49.10%	49.31%	49.52%	49.51%
IVOL_Q4	48.16%	47.51%	47.94%	47.76%	49.15%
IVOL_Q5	43.59%	46.47%	45.20%	45.37%	46.89%

The table shows three interesting results:

First, the correlation level seems to have no impact on the return contribution of the long and the short leg, as the table does not show a clear pattern within each idiosyncratic volatility level.

Second, the contribution ratio is close to 50% within the low to medium idiosyncratic volatility quintiles (IVOL_Q1-IVOL_Q3). This means that the return of exploiting overpriced stocks with a shorted stock is similar to the return of exploiting underpriced stocks with a long position.

Third, the contribution ratio is significantly below 50% in the highest idiosyncratic volatility quintile. The long leg contributes less to the total return than the short leg. In other words, the return from exploiting the overpriced stock with a short position is more profitable than the return from exploiting the underpriced stock with a long position. If restriction impedes short selling the overpriced “winner” stock, pairs trading could be restricted to the long leg. The positive contribution of the long leg indicates that this modified approach is also successful.

These results confirm research hypothesis no. 5. Furthermore, they are consistent with Stambaugh, Yu, and Yuan’s (2015) and Cao and Han’s (2016) line of arguments. Limits to arbitrage deters the exploitation of overpriced stocks. Hence, overpricing is stronger than easy-to-exploit underpricing.

Another important finding is that the long leg also contributes a significant share to the total return. Hence, if restrictions impede short selling, investors could also use the trading signals of pairs trading, but only trade with the long leg.

Altogether, my findings suggest that a negative idiosyncratic volatility effect in one sample contradicts neither the observation of a positive idiosyncratic volatility effect in another sample, nor the economic explanations for a positive idiosyncratic volatility effect. The overall negative idiosyncratic volatility effect is simply the result of a positive idiosyncratic volatility effect among underpriced stocks and an overlaying dominant negative idiosyncratic volatility effect among overpriced stocks.

6.7 Robustness - Fama-MacBeth analysis

In addition to my analyses based on portfolio sorts, I investigate the influence of arbitrage asymmetry on the return in the cross-section with a Fama-MacBeth (1973) regression.

Table 22 reports the coefficients of a Fama-MacBeth (1973) cross-sectional regression. Two observations are made, extracted from each trade: one for the return of the long leg, and one for the short leg. The dependent variable is the leg return, which is the return earned by either the short position or the long position of a pair. The independent variables include: pair volatility, correlation, a dummy for the short leg, and the two interaction terms to measure the influence of idiosyncratic volatility and correlation on the short leg. The data sample includes only the first trade of a pair within a trading period due to reasons of robustness. Considering all trades of a pair within one trading period would bias the analysis. The returns of these trades are significantly determined by the opening signal, which is identical for each trade of one particular pair within one specific trading period.

Table 22: Fama-MacBeth (1973) cross-sectional regression (first natural trade)

Table 22 reports the coefficients of a Fama-MacBeth (1973) cross sectional regression. The dependent variable is the leg return. The leg return is the return earned by either the short position or the long position of a pair. The independent variables include: Pair volatility (Volatility), correlation (Correlation), a dummy (D_Short) for the short leg in model (II)-(IV), and the two interaction terms (I_Short×Vola and I_Short×Corr) to measure the influence of volatility and correlation on the short leg. The control variable includes the duration of a trade. I standardize pair volatility and correlation in the data set separately for each cross-sectional regression to obtain standardized coefficients. I use Newey-West (1987) standard errors with lag six and report the p-values in brackets. Statistical significance at the 10%, 5%, and 1% level is indicated by *, **, and ***, respectively.

Model	(I)	(II)	(III)	(IV)
Volatility	0.0665*** (0.000)	0.0665*** (0.000)	0.0609*** (0.000)	0.0597*** (0.000)
Correlation	-0.0361*** (0.000)	-0.0361*** (0.000)	-0.0366*** (0.000)	-0.0351*** (0.000)
D_Short		0.0268*** (0.000)	0.0280*** (0.000)	0.0280*** (0.000)
I_Short×Vola			0.0136*** (0.007)	0.0136*** (0.007)
I_Short×Corr			-0.0021 (0.442)	-0.0021 (0.442)
Control	yes	yes	no	yes
R-squared	0.116	0.173	0.195	0.200

Table 22 broadly confirms research hypothesis no. 5. A one standard deviation increase in idiosyncratic pair volatility increases the average leg return by 6.65 cents (model I) and is significant. Thus, higher idiosyncratic risk is compensated with a positive premium. Consistent with short selling restrictions, the short leg earns 2.68 cents on average more (model II). This effect is even stronger among stock pairs with high idiosyncratic risk.

For additional robustness, I replicate the analysis and select a random trade if a pair trades more than once within a given trading period. Furthermore, I conduct the Fama-MacBeth regression with non-overlapping trading periods. The results in both additional robustness tests confirm the previous findings.

6.8 Conclusion

I use a novel, innovative test setting to provide evidence that high idiosyncratic volatility increases the return of underpriced stocks and decreases the return of overpriced stocks. This effect was introduced by Stambaugh, Yu, and Yuan (2015) and Cao and Han (2016) with a composite mispricing index that is based on static firm-specific data to a great extent. To overcome the weaknesses of their methodology, I apply pairs trading as an alternative mispricing measure. My mispricing measure identifies dynamic mispricing that is mainly originated by market frictions. Moreover, I introduce pairs trading as an investment strategy that allows investors to exploit the idiosyncratic volatility effect. My results confirm the arbitrage asymmetry argument underlying the idiosyncratic volatility effect, as proposed by Stambaugh, Yu, and Yuan (2015) and Cao and Han (2016). Higher levels of idiosyncratic risk generate a positive return premium among underpriced stocks, and a negative premium among overpriced stocks. Altogether, my study demonstrates that the effect is not driven by the type of mispricing, but mispricing in general.

7. Conclusion and outlook

This thesis contributes to the literature on algorithm-based hedge fund strategies with new insights into pairs trading to further demystify the rule-based trading strategy. Thereby, I closely follow the popular trading approach of Gatev, Goetzmann, and Rouwenhorst (2006). After introducing the methodology, the return calculation, and previous academic work, I investigate three research questions to extend the current knowledge of researchers and practitioners about the trading strategy.

In the first core chapter, I study the influence of volatility and correlation on the pairs trading algorithm. My research reveals that high pair volatility and high correlation increase the overall return, because the rule-based trading algorithm works in favor of these two stock characteristics. Unfortunately, the traditional SSD pair selection criterion fails to identify the most promising pairs. The key learning of this chapter is that pairs traders should pay close attention to the drivers of their chosen trading algorithm. Furthermore, smart investors should optimize their pair selection based on volatility and correlation levels to maximize the risk-adjusted return.

The second core chapter examines the role of gradual information diffusion as a potential return source of pairs trading. Based on previous findings in the literature, I provide a comprehensive explanation of how exactly pairs trading exploits inefficiencies that were generated by gradual information diffusion. My findings about cross-serial correlation, economic links, and persistent leader-follower relationships among pairs significantly contribute to existing research. Institutional traders can apply the knowledge about gradual information diffusion to improve their individual trading approach and use the fundamental story to explain the source of pairs trading returns to clients.

The final core chapter contributes to the field of research that studies whether idiosyncratic volatility is compensated with a risk premium. I confirm Stambaugh, Yu, and Yuan's (2015) and Cao and

Han's (2016) research with pairs trading as a novel, innovative mispricing measure. Like them, I find that idiosyncratic volatility generates a positive return premium among underpriced stocks and a negative premium among overpriced stocks. As pairs trading measures a different type of mispricing compared to previous work, my out-of-sample analysis contributes to the academic discussion with evidence that the idiosyncratic risk effect exists, and that the effect does not depend on the source of mispricing.

The major limitation of this study is the application of the specific pairs trading methodology of Gatev, Goetzmann, and Rouwenhorst (2006). Future research could evaluate my findings with an alternative pairs trading algorithm, as well as different markets with other asset classes. Notwithstanding this limitation, this thesis certainly adds to our understanding of pairs trading.

References

- Ali, A., Hwang, L., Trombley, M., 2003. Arbitrage risk and the book-to-market anomaly. *Journal of Financial Economics* 69, 355–373.
- Almazan, A., Brown, K., Carlson, M., Chapman, D., 2004. Why constrain your mutual fund manager? *Journal of Financial Economics* 73, 289–321.
- Andersen, T., Bollerslev, T., Christoffersen, P., Diebold, F., 2006. Volatility and Correlation Forecasting. In: *Handbook of Economic Forecasting*, Vol. 1, 777-878.
- Ang, A., Hodrick, R., Xing, Y., Zhang, X., 2006. The cross-section of volatility and expected returns. *The Journal of Finance* 51, 259–299.
- Ang, A., Hodrick, R., Xing, Y., Zhang, X., 2009. High idiosyncratic volatility and low returns: International and further U.S. evidence. *Journal of Financial Economics* 91, 1–23.
- Andrade, S. C., di Pietro, V., Seasholes, M. S., 2005. Understanding the Profitability of Pairs Trading. Working paper, UC Berkeley Haas School and Northwestern University.
- Antón, M., Polk, C., 2014. Connected Stocks. *The Journal of Finance* 69, 1099-1127.
- Badrinath, S., Kale, J., Noe, T., 1995. Of shepherds, sheep, and the cross-autocorrelations in equity returns. *Review of Financial Studies* 8, 401-430.
- Bali, T.G., Cakici, N., 2008. Idiosyncratic volatility and the cross section of expected returns. *Journal of Financial and Quantitative Analysis* 43, 29–58.
- Bali, T.G., Cakici, N., Whitelaw, R.F., 2011. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics* 99, 427–446.

Boehme, R., Danielsen, B., Sorescu, S., 2006. Short-sale constraints, differences of opinion, and overvaluation. *Journal of Financial and Quantitative Analysis* 41, 455-487.

Bolgün, K. E., Kurun, E., Güven, S., 2010. Dynamic Pairs Trading Strategy for the Companies Listed in the Istanbul Stock Exchange. *International Review of Applied Financial Issues and Economics*, 2, 37-57.

Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31, 307-327.

Boyer, B., Mitton, T., Vorkink, K., 2010. Expected idiosyncratic skewness. *Review of Financial Studies* 23, 169–202.

Brennan, M., Jegadeesh, N., Swaminathan, B., 1993. Investment analysis and the adjustment of stock prices to common information. *Review of Financial Studies* 6, 799-824.

Cao, J., Han, B., 2016. Idiosyncratic risk, costly arbitrage, and the cross-section of stock returns. *Journal of Banking and Finance* 73, 1-15.

Carhart, M., 1997. On persistence in mutual fund performance. *The Journal of Finance* 52, 57–82.

Chen, H., Chen, S., Li, F., 2013. Empirical investigation of an equity pairs trading strategy. Working Paper, University of British Columbia, University of Michigan.

Chen, J., Hong, H., Stein, J., 2002. Breadth of ownership and stock returns. *Journal of Financial Economics* 66, 171–205.

Chorida, T., Swaminathan, B., 2000. Trading volume and cross-autocorrelations in stock returns. *The Journal of Finance* 22, 913-935.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. *The Journal of Finance* 63,

1977-2011.

Cohen, L., Lou, D., 2012. Complicated firms. *Journal of Financial Economics* 104, 383–400.

Da, Z., Liu, Q., Schaumburg, E., 2014. A closer look at the short-term return reversal. *Management Science* 60, 658-674.

D'Avolio, G., 2002. The market for borrowing stock. *Journal of Financial Economics* 66, 271-306.

DeLong, B., Shleifer, A., Summers, L., Waldmann, R., 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98, 703-738.

Do, B., Faff, R., 2010. Does Simple Pairs Trading Still Work? *Financial Analysts Journal*, 66, 83–95.

Do, B., Faff, R., 2012. Are Pairs Trading Profits Robust to Trading Costs? *The Journal of Financial Research*. 2012, 35, 261–287.

Doukas, J., Kim, C., Pantzalis, C., 2010. Arbitrage risk and stock mispricing. *Journal of Financial and Quantitative Analysis* 45, 907-934.

Engelberg, J., Gao, P., Jagannathan, R., 2009. An anatomy of pairs trading: The role of idiosyncratic news, common information and liquidity. Working Paper, University of North Carolina, University of Notre Dame, Northwestern University.

Engle, R., 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50, 987-1007.

Engle, R., Patton, A., What good is a volatility model? *Quantitative Finance*, 2001, 1, 237-245.

Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. *Journal of*

Financial Economics 33, 3–56.

Fama, E., MacBeth, J., 1973. Risk return, and equilibrium: Empirical tests. *Journal of Political Economy* 71, 607–636.

Figlewski, S., 1981. Futures trading and volatility in the GNMA market. *The Journal of Finance* 36, 445–456.

Fu, F., 2009. Idiosyncratic risk and the cross-section of expected stock returns. *Journal of Financial Economics* 91, 24–37.

Gatev, E., Goetzmann, W. N., Rouwenhorst, K. G., 2006. Pairs trading: Performance of a relative-value arbitrage rule. *Review of Financial Studies* 19, 797–827.

Goyal, A., Santa-Clara, P., 2003. Idiosyncratic risk matters! *The Journal of Finance* 58, 975–1008.

Han, B., Kumar, A., 2013. Speculative trading and asset prices. *Journal of Financial and Quantitative Analysis* 48, 377–404.

Han, Y., Lesmond, D., 2011. Liquidity biases and the pricing of cross-sectional idiosyncratic volatility. *Review of Financial Studies* 24, 1590–1629.

Heston, S., 1993. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies* 6, 327-343.

Hong, H., Stein, J., 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance* 54, 2143-2184.

Hong, H., Torous, W., Valkanov, R., 2007. Do industries lead stock markets? *Journal of Financial Economics* 83, 367–396.

Hou, K., 2007. Industry Information Diffusion and the Lead-lag Effect in Stock Returns. *The Review of Financial Studies* 20, 1113-1138.

Hou, K., Loh, R., 2016. Have we solved the idiosyncratic volatility puzzle? *Journal of Financial Economics* 121, 167–194.

Huang, W., Liu, Q., Rhee, S.G., Zhang, L., 2009. Return reversals, idiosyncratic risk, and expected returns. *Review of Financial Studies* 23, 147– 168.

Huck, N., 2013. The high sensitivity of pairs trading returns. *Applied Economics Letters* 20, 1301 - 1304.

Huck, N., 2015. Pairs trading: Does volatility timing matter? *Applied Economics* 47, 1-18.

Huck, N., Afawubo, K., 2015. Pairs trading and selection methods: Is cointegration superior? *Applied Economics* 47, 599-613.

Jacobs, H., Weber, M., On the determinants of pairs trading profitability. *Journal of Financial Markets*, 2015, 23, 75-97.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance* 48, 65–91.

Jiang, G.J., Xu, D., Yao, T., 2009. The information content of idiosyncratic volatility. *Journal of Financial and Quantitative Analysis* 44, 1–28.

Krauss, C., 2015. Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of economic surveys*, forthcoming.

Lamont, O., 2012. Go down fighting: Short sellers vs. firms. *Review of Asset Pricing Studies* 2, 1-30.

Lehmann, B., 1990. Fads, martingales, and market efficiency. *The quarterly Journal of Economics* 105, 1-28.

Lo, A., MacKinlay, C., 1990. When are contrarian profits due to stock market overreaction. *Review of Financial Studies* 3, 175-205.

Malkiel, B., Xu, Y., 2002. Idiosyncratic risk and security returns. Princeton University, University of Texas at Dallas.

Lin, Y., McCrae, M., Gulati, C., 2006. Loss protection in pairs trading through minimum profit bounds: A cointegration approach. *Journal of Applied Mathematics and Decision Science*, 1-14.

McLean, D., 2010. Idiosyncratic risk, long-term reversal, and momentum. *Journal of Financial and Quantitative Analysis* 45, 883-906.

Mech, T., 1993. Portfolio return autocorrelation. *Journal of Financial Economics* 34, 307-334.

Merton, R., 1987. A simple model of capital market equilibrium with incomplete information. *The Journal of Finance* 42, 483-510.

Menzly, L., Ozbas, O., 2010. Market segmentation and cross-predictability of returns. *The Journal of Finance* 65, 1555-1580.

Miller, E., 1977. Risk, uncertainty, and divergence of opinion. *The Journal of Finance* 32, 1151-1168.

Mitchell, M., Pulvino, T., Staffo, E., 2002. Limited arbitrage in equity markets. *The Journal of Finance* 57, 551-584.

Nagel, S., 2005. Short sales, institutional investors and the cross-section of stock returns. *Journal of Financial Economics* 78, 277-309.

Newey, W.K., West, K.D., 1987. A simple positive-definite heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.

Ornstein, L., Uhlenbeck, G., 1930. On the theory of Brownian Motion. *Physical Review* 36, 823–841.

Papadakis, G., Wysocki, P., 2007. Pairs Trading and Accounting Information. Working Paper, Boston University School of Management/MIT Sloan School of Management.

Pastor, L., Stambaugh, R., 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* 111, 642-685.

Perlin, M. S., 2009. Evaluation of pairs-trading strategy at the Brazilian financial market. *Journal of Derivatives & Hedge Funds*, 15, 122-136.

Pontiff, J., 2006. Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics* 42, 35–52.

Schulz, P., Shive, S., 2010. Mispricing of dual-class shares: Profit opportunities, arbitrage, and trading. *Journal of Financial Economics* 98, 524–549.

Stein, E., Stein, J., 1991. Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *The Review of Financial Studies* 4, 727-752.

Shleifer, A., Vishny, R., 1997. The Limits of Arbitrage. *The Journal of Finance* 52, 35-55.

Stambaugh, R., Yu, J., Yuan, Y., 2012. The short of it: Investor sentiment and anomalies. *Journal of Financial Economics* 104, 288-302.

Stambaugh, R., Yu, J., Yuan, Y., 2015. Arbitrage asymmetry and the idiosyncratic volatility puzzle. *The Journal of Finance* 70, 1903-1948.

Stander, Y., Botha, I., Marais, D., 2013. Trading strategies with copulas. *Journal of Economic and Financial Sciences*. *Journal of Economic and Financial Sciences* 6(1), 102-126.

Tourin, A., Yan, R., 2013. Dynamic pairs trading using the stochastic control approach. *Journal of Economic Dynamics and Control* 37, 1972–1981.

Vidyamurthy, G., 2004. *Pairs trading: Quantitative methods and analysis*. John Wiley & Sons, Hoboken.

Wurgler, J., Zhurayskaya, E., 2002. Does arbitrage flatten demand curves for stocks? *Journal of Business* 75, 583-608.