Examining Pairs Trading Profitability*

Xuanchi Zhu

April 3, 2024

Abstract

This paper establishes three important results regarding the performance of pairs trading, a popular statistical arbitrage strategy that identifies close-moving stocks and capitalizes on their temporary divergence. First, we replicate the results of Gatev et al. (2006) using data from the past twenty years and show that a distance-based pairs trading strategy can result in an average annual excess return of 6.2% and a Sharpe ratio of 1.35. Second, we provide empirical evidence for the risk nature of the strategy. Using a 6-factor model, we find that in our sample, a one-standard deviation increase in the momentum factor can nearly wipe out the return of our best-performing strategy. In the meantime, we show that trading profit is increasing in market risk premium, suggesting that arbitrageurs are compensated for enforcing the Law of One Price. Third, we construct and analyze a psychology-based model that generates these empirical findings. Taken together, these results help bridge the theory and practice of pairs trading.

1 Introduction

Univariate risk factors have drawn extensive attention in financial economics, and, in particular, empirical asset pricing. What is relatively unknown is the high dimensional interaction between stock prices both in the cross-section and time series. A growing body of literature on pairs trading, a bivariate, market-neutral generalization of well-studied contrarian strategies seeks to answer this question. Originally developed by traders at Morgan Stanley in 1980s, the strategy was first documented by Gatev et al. (2006). The rationale underlying the strategy is simple: for two stocks with similar payoffs in the *observed* states of the

*A senior essay submitted to the Department of Economics at Yale University. I thank my advisor, William Goetzmann, for his indispensable guidance and insight, and Nicholas Barberis for his invaluable comments. I am grateful for my family and friends for their care and support. All errors are mine.

world, it is natural to think that they should exhibit the same pattern in *future* states of the world. Thus, if their prices diverge, one can bet on a future convergence by shorting the overvalued stock and longing the undervalued one. Gatev et al. (2006) reported that such strategy generated an annualized excess return of about 11 percent between 1962 and 2002. Furthermore, the authors established that trading profitability is persistent under reasonable assumptions of short-selling constraints and transaction costs.

For an excess return of the size, it is only natural to ask why pairs trading is profitable. Gatev et al. (2006) regressed strategy returns on a five factor model, and found that the strategy's exposure to the known risk factors is small. They concluded that arbitrageurs are compensated for restoring the bivariate "Law of One Price" after perturbations to the equilibrium relationship. In a subsequent work, Engelberg et al. (2009) posited that pairs trading profits is related to information events. Intuitively, when a pair diverges because of firm-specific news, the fundamental shock might persist and thus render the divergence permanent and profits low. On the other hand, if pairs diverged because a common shock propagated through the constituents at different rates, it is very likely that the pair will converge, leading to better returns. The authors argued that arbitrageurs are compensated to provide liquidity when it is needed and take positions when stock prices react slowly to common information.

Seeking to obtain a better understanding of determinants of pairs trading returns, this paper contains three important results. First, we show that the benchmark pairs trading strategy specified in Gatev et al. (2006) has robust returns in the past twenty years, contrary to the findings of Rubesam (2021). Under the most conservative specifications, we find that increasing the pool of eligible stocks can result in an average annual excess return of 6.2% and a Sharpe ratio of 1.35. Second, we discover that the risk nature of the strategy has changed in the past 20 years. Using a 6-factor model to decompose the risk of the strategy, we find that in the new sample, a one-standard deviation increase in the momentum factor, as constructed in Jegadeesh and Titman (1993), can nearly wipe out the return of our best-performing strategy, which was not the case in the previous sample period studied. In the meantime, regressing pairs trading profitability on a set of macroeconomic variables allows us to identify aggregate risk premium as one of the determinants of pairs trading returns. This is consistent with the above view that arbitrageurs are compensated for providing liquidity and enforcing the Law of One Price. Third, after observing these results in the data, we construct a psychology-based model that generates equilibrium stock prices that enable an arbitrageur to generate consistent profits. Moreover, under a simplified pairs trading strategy, we show in simulation that this profit is increasing in the momentum factor and investor risk aversion. Taken together, these results help bridge the theory and practice of pairs trading from an angle unique in the literature.

The rest of the paper is organized as follows. Section 2 reviews two types of pairs trading literature. The first closely mimics Gatev et al. (2006)'s algorithm and focuses on explaining its persistent returns, while the second seeks to improve strategy performance using advanced statistical techniques and test the strategy in other markets. Section 3 describes the details Gatev et al. (2006)'s methodology, replicates their strategy using data from 2003 to 2023, and identifies risks that affect strategy returns. Inspired by our empirical results, Section 4 considers a simple model that recovers pairs trading profits and risks in simulations of equilibrium prices. Section 5 concludes the paper.

2 Literature Review

In this section, we summarize two themes of research on pairs trading. The first probes into the determinants of trading profit and seeks to understand the strategy through the lens of financial economics. The second type of research seeks to generate higher pairs trading profits using models such as cointegration. Although the latter is not the focus of the paper, it helps us understand the advantages and disadvantages of the Gatev et al. (2006)'s methodology, which we, following most papers in the first literature, assume to be the benchmark strategy.

2.1 Determinants of Trading Profit

To explain why pairs trading work, Gatev et al. (2006) posited that its profitability is linked to an omitted risk factor that compensates arbitrageurs for enforcing the Law of One Price. They presented three central pieces of evidence. First, regressing the monthly excess returns on Fama and French (1993)'s three factors augmented by a momentum and short-term reversal factor, they find that the risk-adjusted return is significantly positive, and only 0.1-0.2% below the raw excess return. Furthermore, although some of the regression coefficients with respect to momentum and reversals are statistically significant, a simple bootstrap replacing constituents of a pair with random securities with similar prior one-month returns yields a negative excess return. Lastly, there exists significant correlation (0.48) between excess returns of portfolios constituting top 20 and the 100-120 closest pairs, even though there is no overlap between the position of the two portfolios. It should be noted that this view was recently refuted by Chen et al. (2019), which asserts that pairs trading can be explained by the short-term reversal and one month version of industry momentum. We will come back to this key question in Section 3.

Although Gatev et al. (2006) provided a thorough description of a time-varying risk factor, it did not provide much insight on the profitability of the strategy in the cross-section. Further developing this idea, Do and Faff (2010) showed that returns are highest when strategy is restricted to within-industry pairs. They argue that homogeneity of the fundamentals underlying the securities ensure that they are close substitutes, which increases the strategy return by increasing the conditional probability of convergence. Consistent with Gatev et al. (2006)'s view that the common factor has become dormant since 1990s, Do and Faff (2010) concluded that the lower returns in recent years is caused by this lower probability of convergence.

While the previous two papers are built on the framework rational finance, Engelberg et al. (2009) sought to understand the difference between converging and non-converging pairs from the perspective of frictional finance. They theorize that information events create price divergence, for which arbitrageurs are compensated for providing liquidity. By this logic, if the news event is idiosyncratic to one of the stocks, then it is much more likely that the dislocation in prices is permanent. On the other hand, if the news event affects both stocks (yet is priced in at different rates due to market friction), then pairs trading profits are higher. This result has been further confirmed by Jacobs and Weber (2015), which attributes pairs trading profits to "distractions" that help shift investor's attention away from individual stocks to a group of stocks, broader market events, or simply happenings outside the market. From a rich behavioral literature on investor under-reaction and inattention such as Hirshleifer et al. (2007) and Dellavigna and Pollet (2009), they were able to construct seven variables that categorize these three main sources of distraction, and show that there is a strong link between investor inattention and trading profits.

2.2 Other Methods to Select Pairs

The above section presents conclusions reached from a class of models that largely replicate the original method put forth by Gatev et al. (2006), which select pairs based on the Euclidean distance of normalized recent prices. Econometricians and mathematicians have since proposed to fine-tune the strategy using models based on cointegration and Ornstein–Uhlenbeck processes, the first of which this section will explore in depth.

If the time series of two stocks are cointegrated, there then exists a stationary linear combination of the two series that is mean-reverting, meaning that any short-term deviation from this equilibrium is temporary *by construction*. In our problem, Vidyamurthy (2004) suggests that we test whether a pair exhibits such relationship using the two-step Engle-Granger approach. If $(R_{i,t})_{t \in T}$ and $(R_{j,t})_{t \in T}$ are two return series with order of integration $d = 1$, by definition, there exists β_0, β_1 such that

$$R_{i,t} = \beta_0 + \beta_1 R_{j,t} + u_t$$

where u_t is stationary. Immediately, we know that this model accommodates for more flexible strategies since the ratio between buying and selling is not one-to-one. Rad et al. (2015) propose to first estimate β using OLS, and calculate the residuals. Then, one can estimate the Error Correction Model of the following form

$$\Delta R_{j,t} = \gamma_0 + \gamma_1 \Delta R_{i,t} + \gamma_2 \hat{u}_{t-1} + v_t$$

If γ_2 is statistically significant, then the two series are cointegrated. After performing this test for all nominated pairs of stocks and discarding pairs that fail, they proceed to generate trading signals like Gatev et al. (2006). They demonstrate that the distance method shows a slightly higher monthly return than the cointegration method, though the latter exhibits a slightly better Sharpe Ratio. Krauss (2017) postulates that the suboptimal performance of the cointegration method could be compromised by the two-step selection approach as Rad et al. (2015) since the latter only tested the pairs whose Euclidean distance is small. Yet, without a screening step, performing $N(N-1)/2$ cointegration tests in a universe of N stocks is computationally expensive.

To address this, researchers have proposed two solutions. The first limits their stock universe to a subset of the larger pool previously considered. For example, Rudy et al. (2010) use data of EuroStoxx 50 index constituents and restrict pair formation to 10 industry groups, whereas Caldeira and Moura (2013) consider the 50 most liquid stocks of the Brazilian stock index IBovespa. After performing the aforementioned Engle-Granger test, the cointegrated pairs are sorted based on Sharpe Ratio in the pairs-formation period. The second approach seeks to identify other selection metrics in a large data set that maximizes profit opportunities and the ease of computation. Before performing Johansen cointegration tests on S&P500 constituents, Huck and Afawubo (2015) proposed to drop 80% of the pairs whose returns differ more than 10%. They then show that this cointegration method yields higher profit compared to the distance method, because it selects stocks whose price ratio is more volatile and less correlated as those of the distance method in addition to making sure the price spread is mean-reverting.

Beyond cointegration, researchers have explored more sophisticated time-series methods to ascertain the mean-reverting relationship of a pair (Elliott et al., 2005) and to establish dynamic trading signals more nuanced than the two standard deviation threshold (Chen et al., 2017). In recent years, with the burgeoning of machine learning, some authors generated profitable trading strategies using high-dimensional statistical approaches such as PCA (Avellaneda and Lee, 2008), and clustering (Han et al., 2021). The latter method is particularly interesting and deserves more attention in the field as not only does it use price information (including its derivatives such as momentum), but also it incorporates 78 firm characteristics in the identification of close-moving pairs. The authors find that these characteristics significantly reduce the volatility of the strategy and thus increase its Sharpe ratio.

3 Strategy Replication

As noted above, many authors have confirmed the economic and statistical significance of pairs trading returns. Yet, most work has not incorporated the stock data from the past ten years, during which macroeconomic conditions, market environment, and investor behaviors can be very different compared to before. For example, after Do and Faff (2010) extended Gatev et al. (2006)'s sample period to 2009, only Rubesam (2021) considered more recent data.[1] Furthermore, for the sake of computational ease, many restrict that close pairs can only form within the same industry groups (Engelberg et al., 2009), which precludes some profitable, yet fundamentally uncorrelated pairs. To address this, we first describe the algorithm in Gatev et al. (2006) in detail and then replicate their results. We find that pairs trading strategy is robust out of the original sample of Gatev et al. (2006). Then, we analyze

[1] The latter is a report on RPubs that shows that the monthly strategy returns after the Global Financial Crisis has sharply decreased from 0.22%-0.43% in the prior six-year period to 0.04%-0.07%, which we refute in this section. We do not discuss this work in depth due to its calculation errors.

the risk exposure of the strategy returns and find that pairs trading profit can be understood as risk premium to enforce the Law of One Price, similar to the interpretation of Gatev et al. (2006).

3.1 Methodology

From WRDS, we obtain daily stock price and volume in the CRSP universe from January 2003 to December 2023. To clean the data, we first filter out ETFs, trusts, and funds, many of which are perfectly correlated. Then, we exclude stocks that have zero trading volumes, leaving us with 13386 securities in the entire twenty-year sample period. Lastly, as we detected negative prices in minor cases, we corroborated the data series with online information. After cleaning, we construct a total price index by adding the daily closing price of each security and its dividend.

The benchmark model considered by Gatev et al. (2006) specifies two key steps. First, in the 12-month estimation period, it uses Euclidean distance as the selection criterion for pair formation. Then, in the 6-month trading period, the strategy opens when the difference in normalized prices exceeds two standard deviations, as measured in the estimation period. Formally, consider a universe of N stocks, and denote the trading days of the estimation period by $t \in \{1, 2, \ldots, T\} = [T]$. For a pair of stocks (i, j), we can normalize their (total) price $P'_i = (P'_{i,1}, \ldots, P'_{i,T})$ and $P'_j = (P'_{j,1}, \ldots, P'_{j,T})$ via

$$P_{s,t} = \frac{P'_{s,t}}{P'_{s,0}}, \quad s \in \{i, j\}, t \in [T]$$

so that $P_{i,0} = P_{j,0} = 1$. The resulting price vector, then, represents the value of a \$1 investment in stock s at time t. Then, define $D_{i,j}$ to be the Euclidean distance (or the sum of squared differences) between P_i and P_j

$$D_{i,j} = \frac{1}{T} \sum_{t=1}^{T} (P_{i,t} - P_{j,t})^2$$

with sample variance

$$s_{i,j}^2 = \frac{1}{T}\sum_{t=1}^{T}\left[(P_{i,t} - P_{j,t}) - \frac{1}{T}\sum_{t=1}^{T}(P_{i,t} - P_{j,t})\right] \tag{1}$$

After calculating this pairwise distance for all $N(N-1)/2$ pairs, they considered the top M pairs eligible in the subsequent trading period, which has T' periods. Gatev et al. (2006)'s benchmark model set $M = 20$, and we consider the effect of increasing this to 100 and 500. Then, at $t \in [T'-1]$, the trading algorithm initiates a \$1 long-short position on the pair (i, j) if $|P_{i,t-1} - P_{j,t-1}| > 2 \cdot s_{i,j}$, that is, if their spread exceeds twice the historical standard deviation. Let t_1 denote this starting date, and $\epsilon = \text{sgn}(P_{i,t_1-1} - P_{j,t_1-1}) = 1$ if stock i is the leader and -1 otherwise. Note that the positions are opened one day after the initial divergence of prices. This "wait-one-day" trading rule minimizes the effect of bid-ask bounce – the phenomenon that the arbitrageur is likely selling at an ask price and buying at a bid price conditional on price divergence – thus yielding a more conservative estimate of pairs trading returns. There are three ways in which the position on a pair can close at time $t_2 \in \{t_1+1, t_1+2, \ldots, T'\}$: if the pair converges, that is, when $t_2 = \min\{t : t > t_1, \text{sgn}(P_{i,t-1} - P_{j,t-1}) \neq \epsilon\}$; if the trading period ends, that is, when $t_2 = T'$; if one of the stocks, say i, is delisted at $t_2 - 1$, yielding $P_{i,t}$ to be `nan` for all $t \in \{t_2-1, \ldots, T'\}$. If the pair converges, we wait for it to diverge again to potentially engage in many "round-trip" trades over the trading period.

Although the trading rule is intuitive, the calculation of the excess return of the strategy is more nuanced. Gatev et al. (2006) points out that the stream of cash flows is randomly distributed throughout the trading period, and for each pair we can have multiple cash flows. Because the nature of the strategy entails potentially frequent trades, the authors propose to calculate the following buy-and-hold returns of a long-short portfolio to avoid transaction cost associated with daily rebalancing. We start by calculating the return for an open trade in the k-th pair, as a function of the aforementioned leader indicator ϵ_t^k at time

$t \in \{t_1^k, \ldots, t_2^k\}$

$$R_t^k(\epsilon_t^k) = -\epsilon_t^k(R_{i,t} - R_{j,t})$$

where i, j are the constituents of the k-th pair and $R_{i,t}, R_{j,t}$ denote their return at time t. Note that we use the superscript to denote variables associated with pairs and subscript to denote those associated with stocks. Also, note that we need the minus sign since, in the case of $\epsilon = 1$, we are shorting stock i and longing stock j. The return to a portfolio comprising the top M closest pairs, R_t^P, can be written as a weighted sum of the returns of individual pairs

$$R_t^P = \frac{\sum_{k=1}^M w_t^k R_t^k(\epsilon_t^k)}{\sum_{k=1}^M w_t^k}, \quad w_t^k = \prod_{\tau=t_1^k}^{t-1}(1 + R_\tau^k(\epsilon_t^k)) \tag{2}$$

Intuitively, the returns of the individual pairs are scaled by their compounded value since the opening of the trade, w_t^k, normalized by the sum of those weights across all M stocks eligible for trading.

At the start of every month between January 2004 and June 2023, we initiate a six-month trading period, which results in overlapping series of R_t^P. In fact, on each day between June 2004 and June 2023, we record six such returns, corresponding to six strategies opened in the current month and each of the past five months. After taking the average between these returns (which Gatev et al. (2006) interpret as the returns from six traders of the same desk) and compounding, we obtain a monthly series that span the past 20 years. This fully describes our benchmark strategy. Below, we provide brief descriptions and motivations of all the strategies considered in our study: (1) [T20], our benchmark, considers the Top 20 closest pairs; (2) [T100] considers the top 100 closest pairs. This is a mix of the Gatev et al. (2006) strategy that employs top 100-120 pairs (to show that pairs trading constitutes a risk factor), and the Engelberg et al. (2009) strategy to use top 200 pairs; (3) [T500] includes the top 500 closest pairs and tries to further reduce the variance of returns in the presence of large outliers in data; (4) [R20] also considers 20 closest pairs, yet they are Restricted to the same industry groups based on SICCD[2]; (5) [L50] considers 20 closest pairs among the stocks whose *L*iquidity, as measured by the average daily trading volume in the week prior to pair formation, is above the 50^{th} percentile. In addition to the aforementioned "wait one-day rule", this minimizes transaction costs; and (6) [L75] considers very liquid stocks ranked in the top 25%. We report the results of our replication in the next section.

3.2 Results

À la Table 1 in Gatev et al. (2006), we show the monthly return distribution for the above strategies below. Note that this return denotes excess return, where we retrieve the risk-free rate, along with the factors employed in computing the information ratio, from Ken French's data library.

Table 1: Return Distribution and Performance of Pairs Trading Strategies

	T20	T100	T500	R20	L50	L75
Return Distribution						
Mean	0.461	0.426	0.498	0.474	0.281	0.188
t-stat	3.647	4.618	4.471	3.459	3.038	2.434
Median	0.290	0.331	0.341	0.271	0.161	0.146
Standard Deviation	1.961	1.254	1.282	2.057	1.273	1.160
Skewness	7.331	1.895	2.263	6.391	1.848	-0.364
Minimum	-6.696	-3.631	-3.146	-8.652	-3.914	-5.853
Maximum	24.285	7.428	9.316	24.285	9.330	3.512
Performance Metrics						
Sharpe Ratio	0.814	1.177	1.345	0.798	0.766	0.560
Information Ratio	0.240	0.381	0.438	0.253	0.258	0.187

More than 20 years after Gatev et al. (2006) brought pairs trading into the attention of academics and investors alike, the profit of a "vanilla" algorithm remains both economically and statistically significant. [T20], the benchmark considered in the original paper, generated

[2]Short for Standard Industrial Classification Code, SICCD divides companies into the following industries: 1) Agriculture, Forestry, Fishing; 2) Mining and Construction; 3) Manufacturing; 4) Transportation & Public Utilities; 5) Wholesale/Retail Trade; 6) Finance, Insurance, Real Estate; 7) Services; and 8) Public Administration

an average monthly excess return of 0.46%, not too much lower than the 0.52% originally reported in Panel B of Table 1 in Gatev et al. (2006). To retrieve a robust standard error, recall that the intercept term of an OLS model equals the unconditional expectation. Thus, we can regress strategy return on a vector of 1's, then estimate the form of heteroskedasticity and autocorrelation consistent (HAC) standard errors for the intercept (Newey and West, 1987). For the benchmark model, the t-statistic based on standard errors recovered by this method is 3.65, resulting in a p-value smaller than 0.001. In fact, the mean of the monthly return distribution is significantly positive at the 5% level for all six strategies. We see that strategy [T500] achieves the highest average monthly return of near 0.5%, even though at every trading period it involves a bigger pool of stocks, some of which did not comove as closely as others during pair formation.

Examining the excess return distribution, we observe that its overall shape is positively skewed. [T20] and [R20] are particularly affected by positive outliers, the largest of which results in a monthly return of 24.3%. Increasing M, the number of close stocks eligible for trading, we achieve excellent variance reduction as characterized by both the standard deviation and extreme points. Interestingly, to a certain degree, imposing the floor for the daily trading volume also has such effects: the standard deviation of [T100], [T500], [L50], [L75] are very similar. This is because empirically, the abnormal movement of prices driving extreme returns almost always results from either a stock split or a delisting, and the latter particularly common for less liquid stock-like securities such as SPACs.[3] Hence, by largely excluding the latter category from the pool of eligible stocks, [L50] and [L75] effectively achieve variance reduction while keeping $M = 20$. In the data, we find that this effect is not monotonic: if we set too high of a liquidity requirement, the effect of contracting the universe of stocks eligible for formation outweighs the effect of trading "well-behaved" stocks with high liquidity, thus resulting in a smaller mean and a higher standard deviation.

[3]Short for Special-purpose acquisition company, SPAC is a publicly listed company designed solely to acquire one or more privately held companies. The boom of this class of blank-check companies in recent years has inspired work such as Gahng et al. (2023), which analyzes the agency problems that certain SPAC features address, and Chen et al. (2022), which investigates hedge fund ownership in SPACs.

Abstracting from the distribution of returns, we present two standard metrics that allow us to evaluate and compare the performance of the two strategies. First, we calculate the Sharpe Ratio of the strategy using plug-in estimates of the mean and standard deviation calculated above. In accordance to the discussion above, [T500] achieves a Sharpe Ratio of 1.35, the highest out of all the strategies considered. Also, as [L50] reduces the spread of the empirical distribution, its Sharpe Ratio is very close to that of [T20] and [R20], all around 0.8. This provides further evidence to Gatev et al. (2006)'s finding that pairs trading profit is robust after accounting for transaction costs. Second, we can calculate the information ratio IR of the strategy

$$IR = \frac{\alpha}{\sigma_\epsilon}$$

where α, σ_ϵ are the intercept and standard deviation of the residuals in the six-factor model

$$R_t = \alpha + \beta^T F_t + \epsilon_t \tag{3}$$
$$F_t = [MKT_t,\ SMB_t,\ HML_t,\ MOM_t,\ SRV_t,\ LRV_t]^T$$

Similar to previous studies, we decompose our excess return into the three factors of Fama and French (1993), augmented by a medium-term momentum factor (MOM), a short-term reversal factor (SRV), and a long-term reversal factor (LRV). Intuitively, the information ratio scales the α of a strategy by the noisy-ness of the residuals. Table 1 shows that this yields a similar ranking of the strategies compared to the Sharpe Ratio.

3.3 Risk Analysis

To understand the determinants of profitability, we regress the monthly series of pairs trading returns on two sets of variables. We first estimate the 6-factor model in (3) using Newey-West standard errors with 6 lags, which we report in parentheses below the point estimates in Table 2. The coefficients significant at the 5%, 1%, and 0.1% level are highlighted with 1, 2, 3 stars, respectively.

Table 2: Risk Decomposition of Pairs Trading Strategies

	T20	T100	T500	R20	L50	L75
alpha	0.455***	0.419***	0.497***	0.505***	0.307**	0.204**
	(0.110)	(0.082)	(0.100)	(0.123)	(0.099)	(0.069)
MKT	-0.001	0.004	0.004	-0.036	-0.031	-0.012
	(0.025)	(0.025)	(0.022)	(0.033)	(0.033)	(0.020)
SMB	0.172	0.097*	0.060	0.171	0.070*	0.047
	(0.130)	(0.046)	(0.034)	(0.129)	(0.035)	(0.031)
HML	0.015	-0.030	-0.060	-0.022	-0.118**	-0.063
	(0.102)	(0.052)	(0.059)	(0.103)	(0.041)	(0.040)
MOM	-0.051	-0.091***	-0.095***	-0.070*	-0.041*	-0.067***
	(0.031)	(0.016)	(0.017)	(0.028)	(0.019)	(0.019)
SRV	-0.023	0.068*	0.073**	-0.066	0.076*	0.033
	(0.088)	(0.029)	(0.028)	(0.088)	(0.031)	(0.023)
LRV	-0.034	0.008	0.053	0.013	0.098*	0.070
	(0.096)	(0.046)	(0.048)	(0.098)	(0.043)	(0.036)
R^2	0.060	0.230	0.217	0.059	0.124	0.113

On a high level, our regression shows that pairs trading returns cannot be explained by existing risk factors. First, not only are the strategy alphas almost always significant at the 0.1% level, but its magnitude (between 0.4 and 0.5% a month) also shows this deviation from a conventional asset pricing model is economically significant. Second, conditional on medium-term momentum, strategy returns are orthogonal to the excess market returns, the size factor, the value factor, and both short and long-term reversal as the coefficients are either mostly statistically insignificant or have mixed signs. Although this result largely aligns with Gatev et al. (2006) and Jacobs and Weber (2015), the loading on SRV in our data is only statistically significant for [T100], [T500], and [L50]. This is somewhat counter-intuitive since pairs trading can be thought of as a contrarian strategy, and one would expect that periods when short-term reversal profits are high should correlate with those with high pairs trading profits. One explanation to this is that, while a short-term reversal strategy buys past "losers" and sells past "winners" blindly, our pairs trading strategy does the same only after verifying the two stocks are close.

The only meaningful risk consistent across all specifications of our strategy, then, is momentum, which suggests that past stock returns between 2 and 12 months predicts pairs trading return with a positive sign. In particular, in [T100] and [T500], in which the momentum effect is the strongest, pairs trading returns would increase by around 0.1% if the return difference between a portfolio of winner stocks and that of loser stocks gets 1% smaller. This effect is certainly non-trivial considering the standard deviation of MOM is 4.4 and that the mean of monthly returns for these two strategies are between 0.4 and 0.5%. All else equal, this suggests that a one standard deviation increase in MOM can wipe out our entire strategy return.

Next, we examine which, if any, economic indicators can explain pairs trading profits. To start, we select the same set of variables as Hutchinson and O'Brien (2020), which studies the momentum factor in different macroeconomic conditions: (1) DEF: default spread measured as the difference between yields on BAA-rated and AAA-rated corporate bonds; (2) DIV: aggregate dividend yield on the S&P 500 Index; (3) GDP: monthly growth of U.S GDP per capita; (4) INF: monthly inflation rate based on the consumer price index; (5) MKT: excess market return; (6) $RREL$: relative T-bill rate, defined as the difference between the three-month T-bill rate and its 12-month backward moving average; (7) $TERM$: term spread measured as the difference between yields on ten-year and three-month Treasury securities; and (8) $UNEMP$: the U.S. monthly unemployment rate. We download DIV from Robert Shiller's website and the rest from Federal Reserve Economic Data (FRED). Regressing monthly pairs trading returns on these 8 covariates, we present the point estimates and Newey-West standard errors in Table 3.

Our model selects the default spread, GDP growth, and excess market return as the statistically significant predictors, of which the first is uniformly significant at the 0.1%

Table 3: Pairs Trading and Macroeconomic Risk

	T20	T100	T500	R20	L50	L75
Constant	-0.331 (0.536)	-1.090 (0.556)	-1.721** (0.659)	-0.723 (0.638)	-1.371* (0.617)	-0.625 (0.410)
DEF	0.563*** (0.138)	0.798*** (0.154)	0.810*** (0.185)	0.805*** (0.167)	0.705*** (0.159)	0.485** (0.157)
RREL	0.013 (0.143)	0.090 (0.127)	0.061 (0.145)	-0.019 (0.171)	0.191 (0.133)	-0.044 (0.107)
INF	0.011 (0.058)	0.069 (0.060)	0.109 (0.070)	0.030 (0.077)	0.067 (0.054)	-0.033 (0.037)
TERM	-0.159 (0.178)	0.056 (0.101)	0.065 (0.097)	-0.153 (0.190)	0.147 (0.099)	0.238** (0.077)
UNEMP	0.035 (0.051)	0.006 (0.053)	0.054 (0.051)	0.040 (0.056)	0.024 (0.056)	-0.037 (0.055)
DIV	0.005 (0.009)	0.007 (0.006)	0.016* (0.007)	0.008 (0.010)	0.005 (0.007)	0.008 (0.005)
MKT	0.044 (0.023)	0.093*** (0.023)	0.090*** (0.027)	0.006 (0.028)	0.041* (0.021)	0.042* (0.020)
GDP	-0.061* (0.027)	-0.153** (0.053)	-0.188** (0.065)	-0.082* (0.036)	-0.152*** (0.042)	-0.114** (0.035)
R^2	0.038	0.206	0.254	0.047	0.148	0.133

level across all strategies except for [L75]. Holding other variables constant, for every 1% widening in the gap between investment-grade and high-yield bond yield, the monthly return of [T500] increases by 0.8%. If we see default spread as a measure of aggregate risk premium conditional on the current macroeconomic condition spanned by the other covariates, it implies that pairs trading strategy is more profitable if investor risk-aversion is high. Indeed, if we replace DEF by the VIX index, another market-based measure of investor uncertainty, the results in Table 3 remain largely the same.

Thus far, we have established that pairs trading profit is decreasing in the size of the momentum factor and increasing in aggregate risk premium. Although some papers such as Hong and Stein (1999) show that stock momentum increases in the risk tolerance of traders, including default spread DEF as one of our factors in (3) does not drastically change the t-stat or the magnitude of the loading on MOM. For example, consider [T500], the point estimate (standard error) of β_{MOM} before including DEF is -0.095 (0.017) while the one after including DEF is -0.0709 (0.019). This suggests that, to a certain extent, the two risks are "separate". To further substantiate this empirical finding, we turn to a simple model in the next section that explores the relationship between pairs trading expected payoff and the two key risks using price vectors derived in equilibrium.

4 A Psychology-based Model

In this section, we write down a simple psychology-based model that helps ground our empirical results. Although it is certainly possible to generate similar results using a model from rational finance, recent research on the source of comovement between stock prices has taken a behavioral view (Barberis, 2018). Most similar to Hong and Stein (1999), our model deviates from those of rational finance by introducing a positive weight of noise traders that are subject to categorical thinking, the notion that investors like to group stocks under broad labels, such as tech stocks and inflation-proof stocks, and allocate funds across stocks on a category level. This common demand of the two stocks, even in the absence of correlated fundamentals, generates a common factor in equilibrium prices. In numerical simulations, we find that not only is pairs trading profitable, but also the expected payoff increases in risk aversion and decreases in the size of price momentum.

4.1 Setup

Similar to Barberis et al. (2002), consider an economy with $T+1$ dates from $t=0$ to $t=T$. There is a risk-free asset and two risky assets, each with Q shares of fixed supply. Asset i is a claim to a single cash flow $D_{i,T}$ paid at time T satisfying the following bivariate relationship

$$D_T = D_0 + \epsilon_1 + \ldots + \epsilon_T, \quad \epsilon_t \sim \mathcal{N}(0, \Sigma) \text{ i.i.d over time}$$

where $D_0 = [D_{1,0}, D_{2,0}]^T$ are the initial dividends realized and publicly announced at time 0, and $\epsilon_t = [\epsilon_{1,t}, \epsilon_{2,t}]^T$ are the shocks realized at time t. For simplicity, we use P_t and $P_t - P_{t-1}$ to denote the price and return vector at time t.

The investors in the economy are divided into two categories. The fundamentalists, the first type, constitute a fraction $1 - \mu$ of the population, and choose their per-capita demand N_t^f at time t by maximizing their expected CARA utility

$$\max_{N_t^f} E_t[-\exp(-\gamma(W_t + N_t^{f'}(P_{t+1} - P_t)))]$$

where W_t is their wealth at time t. At time $T-1$, these fundamental traders forecast the price in the next period to be exactly the dividend amount at the current period.

$$E_{T-1}[P_T] = D_{T-1} \tag{4}$$

The other type of investors are noise traders with demand linear in the most recent return

$$N_t^c = a\mathbf{1} + b(P_t - P_{t-1}) \tag{5}$$

where we use the superscript to denote *c*ategorical thinkers. This can be seen as the bivariate version of the demand function of momentum traders in Hong and Stein (1999), although we assume the elasticity parameter b to be exogenous so that we can derive the close-form of our equilibrium price without having to solve a dynamic programming problem. By Proposition 1 in Hong and Stein (1999), we assume $b > 0$, that is, noise traders rationally behave as trend-chasers. The sign of a is not important as it will be eliminated in equilirbium. As a whole, we can interpret the noise trader demand on the *pair* level to be benchmarked to a certain level a, though they adjust their demand on the *stock* level based on the most recent price return, $P_t - P_{t-1}$. Though simplified, this demand function is realistic for institutional and retail traders who invest in particular categories in a mostly passive yet occasionally "irrationally exuberant" way.

4.2 Equilibrium

We can now derive the equilibrium price vector by calculating the demand function of fundamental traders and imposing market clearing. The first order condition for the fundamental trader is the following

$$\frac{d}{dN_t^f} E_t[-\exp(-\gamma(W_t + N_t^{f'}(P_{t+1} - P_t)))] = 0 \tag{6}$$

For simplicity, we assume the conditional distribution of returns at time t to be normal, $P_{t+1} - P_t \sim \mathcal{N}(E_t(P_{t+1} - P_t), \Sigma)$. Then, we have that

$$-\gamma(W_t + N_t^{f'}(P_{t+1} - P_t)) \sim \mathcal{N}(-\gamma N_t^{f'} E(P_{t+1} - P_t), \gamma^2 N_t^{f'} \Sigma N_t^f)$$

Recall for a univariate Gaussian X, we have that $E[e^X] = M_X(1) = \exp(\mu + \sigma^2/2)$. Using this fact, we can simplify our objective in (6) to get

$$N_t^f = \frac{1}{\gamma} \Sigma^{-1}(E_t[P_{t+1}] - P_t) \tag{7}$$

Then, imposing market clearing $\mu N_t^c + (1-\mu) N_t^f = Q\mathbf{1}$, we have that

$$P_t = E_t[P_{t+1}] + \frac{\gamma\mu}{1-\mu} \Sigma N_t^c - \frac{Q\gamma}{1-\mu} \Sigma \mathbf{1} \tag{8}$$

Now, we set $t = T - 1$ so we can plug (4) and (5) into (8) to get

$$P_{T-1} = (I - \phi\Sigma)^{-1} \left[D_{T-1} - \phi\Sigma P_{T-2} + c_{T-1}\Sigma\mathbf{1}\right] \tag{9}$$

where I is the identity, $\phi = \frac{b\gamma\mu}{1-\mu}$, and $c_{T-1} = \frac{\gamma}{1-\mu}(a\mu - Q)$. Note that in order for $I - \phi\Sigma$ to be invertible, we require $\rho\sigma_1^2\sigma_2^2 \neq (1 - \sigma_1^2)(1 - \sigma_2^2)$.

Based on this, we can roll backward to derive the previous price vectors. Keeping $t = T-1$ and taking the (fundamentalist's) expectation of both sides of (8) at time $T - 2$, we get

$$E_{T-2}[P_{T-1}] = D_{T-2} + \frac{Q\gamma}{1-\mu}(\mu - 1)\Sigma\mathbf{1} \tag{10}$$

where we simply set $E_{T-2}[N^c_{T-1}] = Q$ due to bounded rationality as many related models in the literature. This imposes that fundamentalists assume that noise traders will demand the risky asset in proportion to the weight in future periods, since they have short horizons and do not know the demand function (5). Equation (10) extends (4) and represents a fundamentalist's forecast of price movement at time $T - 2$. To derive the equilibrium price at $T - 2$, we follow the same steps as above, first writing down the fundamentalist's demand

$$N^f_{T-2} = \frac{1}{\gamma}\Sigma^{-1}[D_{T-2} + \frac{Q\gamma}{1-\mu}(\mu - 1)\Sigma\mathbf{1} - P_{T-2}]$$

and then impose market clearing $\mu N^c_{T-2} + (1-\mu)N^f_{T-2} = Q\mathbf{1}$. The price vector in time $T - 2$ is thus

$$P_{T-2} = (I - \phi\Sigma)^{-1} \left[D_{T-2} - \phi\Sigma P_{T-3} + c_{T-2}\Sigma\mathbf{1}\right] \tag{11}$$

where $c_{T-2} = \frac{\gamma}{1-\mu}(a\mu - Q - (\mu - 1)Q)$. Iteratively, for a fixed P_0, we can write down the price vector at time t to be

$$P_t = (I - \phi\Sigma)^{-1} \left[D_t - \phi\Sigma P_{t-1} + c_t\Sigma\mathbf{1} \right] \tag{12}$$

where $c(t) = \frac{\gamma}{1-\mu}[a\mu - Q - (T-t-1)(\mu-1)Q]$. If we impose that the covariance matrix of the fundamental shocks is diagonal, that is, $\Sigma = \sigma^2 I$, we get

$$P_t = \frac{1}{1-\phi\sigma^2} D_t + \frac{-\phi\sigma^2}{1-\phi\sigma^2} P_{t-1} + c_t\sigma^2\mathbf{1} \tag{13}$$

We interpret each term in this equilibrium object. Similar to the case of return extrapolation in Barberis (2018), the first term in (13) resembles the prediction of rational finance – the price of a risky asset should be equal to the expected cashflow – only with a slight distortion. Moreover, if we assume the cashflow at time t should positively contribute to the price of the security, that is, $1 - \phi\sigma^2 > 0$, this distorting factor is increasing in idiosyncratic risk, risk aversion, and momentum reaction. Observe that although the first two appear in the optimal demand of fundamental traders, it is the extra holdings of noise traders that cause the price series to diverge from the fundamentals. Intuitively, in addition to the risk aversion and idiosyncratic risk, the momentum reaction parameter exacerbates this upward bias. Also, in accordance with results from classical economics, the third term in (13) suggests that equilibrium price is decreasing in supply Q and increasing in time t. The latter reflects the fact that the rational investors put higher pressure on the demand when the liquidation date is near, thus lifting the prices. This term is interesting since it shows that stock prices can comove even if their fundamentals are uncorrelated.

Before commenting on the autoregressive term in the middle of (13), we first calculate the price spread s_t so that the third term vanishes while the coefficients in front of first and second term remain unchanged

$$
\begin{aligned}
s_t &= \frac{1}{1-\phi\sigma^2} d_t + \frac{-\phi\sigma^2}{1-\phi\sigma^2} s_{t-1} \\
&= \frac{1}{1-\phi\sigma^2} d_{t-1} + \frac{-\phi\sigma^2}{1-\phi\sigma^2} s_{t-1} + e_t
\end{aligned} \tag{14}
$$

where $d_{t-1} = D_{1,t-1} - D_{2,t-1}$ is the difference in current fundamentals at time $t-1$, and $e_t = \epsilon_{1,t} - \epsilon_{2,t} \sim \mathcal{N}(0, 2\sigma^2)$ is the difference between the cashflow shocks. Figure 1 gives a graphical representation of this relationship.

Figure 1: Evolution of Price Spread

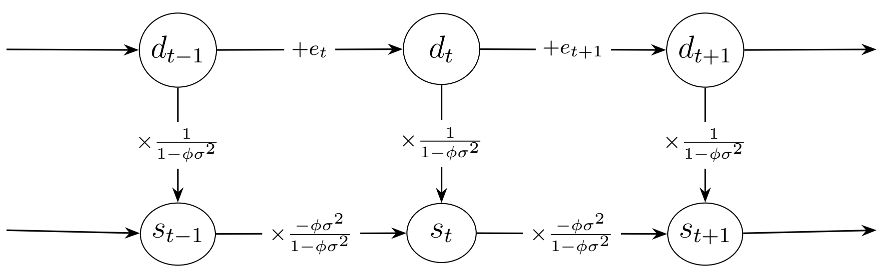

Note: At $t > 1$, the value of each node on the bottom row (s_t) is the sum of the value in the node above (d_t), scaled by $1/(1-\phi\sigma^2)$, and the value of the node on the left (s_{t-1}), scaled by $-\phi\sigma^2/(1-\phi\sigma^2)$. At $t = 1$, we assume $s_0 = 0$ and thus $s_1 = \frac{1}{1-\phi\sigma^2} d_1$. By definition of d_t, there is no d_0.

Since the coefficient in front of the first two terms add up to one, we can interpret this similar to a normalization. To see this, we consider two cases. On one hand, if $\phi\sigma^2$ is near 0, then price spread is almost one-to-one to the difference in current fundamentals. This needs no adjustment from s_{t-1} in the horizontal direction (in Figure 1). On the other hand, if $\phi\sigma^2$ is large, then we would have "double-counted" the contribution of d_t. Due to the evolution of the dividend series, we would have also exaggerated the contribution of $d_{t-1}, \ldots, d_0$ in the vertical direction, from which we need to add a negatively-weighted s_{t-1} in the horizontal direction. Mathematically, we can write (14) as a weighted sum of previous innovations

$$s_t = \frac{1}{1-\phi\sigma^2}\sum_{\tau=0}^{t}\left(\frac{-\phi\sigma^2}{1-\phi\sigma^2}\right)^{\tau} d_{t-\tau}$$
$$= \frac{1}{1-\phi\sigma^2}\sum_{\tau=1}^{t}\left[\sum_{k=0}^{t-\tau}(\frac{-\phi\sigma^2}{1-\phi\sigma^2})^{k}\right] e_{\tau}$$

From here, we can deduce the momentum effect is weaker when $\phi\sigma^2$ is bigger, since the negative terms in the bracket sum above will decrease s_t's dependency on recent innovations. In other words, since $\phi\sigma^2$ is dependent on idiosyncratic risk σ^2, the proportion of noise traders μ, and price elasticity b, increasing risk aversion decreases, but is not the only way to decrease, the momentum factor. We want to show this ultimately increases pairs trading profits, thus in accordance with our empirical results above. There are two ways to verify this. First, we could check if our model in (14) predicts mean reversion explicitly. This is hard since the series of price spread is unconditionally heteroskedastic, yet not impossible, as shown in both theoretical (Bollerslev, 1986) and empirical work (Poterba and Summers, 1987). Alternatively, we can abstract from the technicality of time-series econometrics and analyze the comparative statics through simulation.

4.3 Numerical Comparative Statics

In this section, we are interested in how pairs trading realized profit, defined to be π, changes in relation to $\phi\sigma^2$, which the previous section establishes that is increases risk aversion γ and decreases in the strength of momentum. For simplicity, since the latter is more nebulous, we center our discussion on γ. First, following Gatev et al. (2006)'s algorithm introduced earlier, we write down a simpler version of pairs trading payoff, ignoring the "wait-one-day" rule and the long-short return calculation

$$\pi(\phi) = \frac{1}{M}\sum_{i=1}^{M}\left(c + \sum_{t\in\mathcal{I}} |s_t|\right)$$

where $\mathcal{I}$ is the index set whose elements are dates in which pairs diverge before converging, and c is the potential payoff from forcibly closing an open position if spread does not converge. In other words, for every pair, we sum up the spread difference at each divergence provided convergence occurs, and adjust it by the (likely negative) payoff of the last trade if it does not converge. After this, we can compute the average payoff from trading M pairs, and run this T times to approximate the distribution of realized profit. Bridging this model with the execution of our strategy, we can interpret T as the number of overlapping trading periods.

Figure 2: Empirical distribution of trading profit for different values of risk aversion

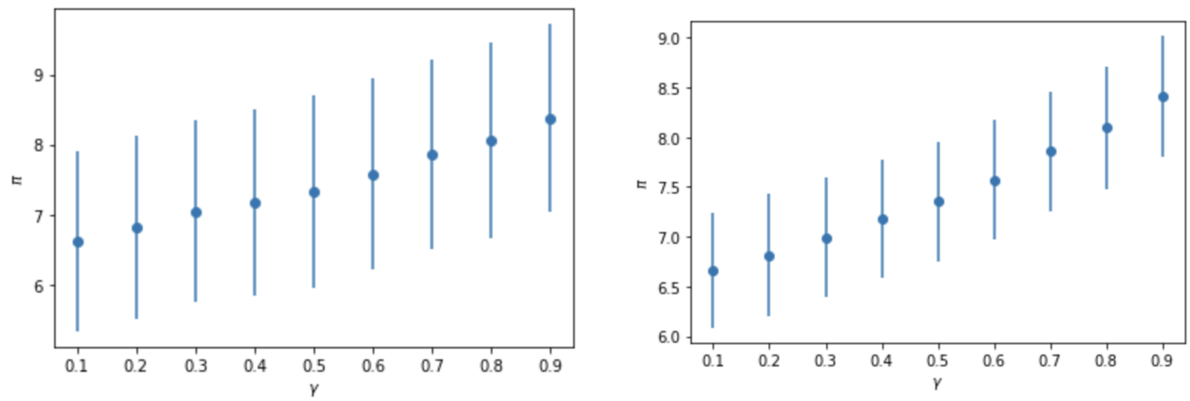

We set $n = 500$, $\sigma = 1$, $b = 2$, $\mu = 0.5$, $T = 100$, $\gamma \in \{0.05, 0.10, \ldots, 0.95\}$. The left figure sets $M = 100$ and the right sets $M = 500$. The dots show the point estimates of profit, averaged over M pairs per T-day trading period. The error bands denote the interquartile range, spanning 25^{th} to 75^{th} percentile of profits across n trading periods.

Figure 2 plots the mean and interquartile range of the empirical distribution of π for different values of γ. Although the spread of the distribution for each γ is quite large, it is clear that π increases in γ. We confirm that the sign of $\Delta\bar{\pi}/\Delta\gamma$, where $\Delta\bar{\pi}$ is the sample average, is robust independent of our choices of parameter value given large n, M, although we cannot determine the shape of the curve connecting the dots, which represent the average realized profit. Furthermore, comparing the two panels, we can see the benefit of increasing M: it slightly increases the mean at each γ and significantly narrows the band. This simulation shows that not only is pairs trading profitable, but also the expected payoff increases in the size of the pool and risk aversion, and decreases in the strength of momentum, all of which mirrors our results in the previous section.

5 Conclusion

In this paper, we replicate the results of pairs trading strategy initially proposed by Gatev et al. (2006) using stock price data from the past twenty years. We find that pairs trading results are still robust under transaction cost despite the changing market environment, with our top performing strategy producing a compounded annual excess return of 6.2%. Compared to most papers in the literature, our best performing strategy uses a larger pool of stocks to average out the influence of outliers whose price jumped "overnight" in response to delisting or stock split news. We then analyze two determinants of pairs trading profit: the medium-term momentum factor and the default spread, which closely tracks the investor risk premium holding other macroeconomic variables constant. This supports Gatev et al. (2006)'s hypothesis that arbitrageurs are compensated for restoring the "Law of One Price". Our empirical results are then confirmed in simulation, using equilibrium prices predicted by a psychology-based model.

We highlight two directions of future studies. First, on the empirical front, we can explore and analyze many variants of the distance approach, such as the "cream-skimming" strategy of Engelberg et al. (2009), which prematurely closes positions on divergent pairs, and the cointegration step of Rad et al. (2015), which checks for mean-reverting behavior of the spread of close-moving stocks. We can also better proxy macroeconomic risk and economic uncertainty, similar to Bali et al. (2014). On the theoretical front, we can consider a richer model by adding the speed of information diffusion and explicitly specifying the information events that cause the pair to diverge. In particular, based on the limited attention hypothesis of Jacobs and Weber (2015), it is possible that we can generate category-learning behavior and investor overreaction endogenously (Peng and Xiong, 2006). This will help unify our results with not only Engelberg et al. (2009) and Jacobs and Weber (2015), but also the rich behavioral finance literature.

References

Avellaneda, M. and Lee, J.-H. (2008). Statistical arbitrage in the U.S. equities market.

Bali, T. G., Brown, S. J., and Caglayan, M. O. (2014). Macroeconomic risk and hedge fund returns. *Journal of Financial Economics*, 114:1–19.

Barberis, N. (2018). Psychology-based models of asset prices and trading volume.

Barberis, N., Shleifer, A., and Wurgler, J. (2002). Comovement.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327.

Caldeira, J. F. and Moura, G. V. (2013). Selection of a portfolio of pairs based on cointegration: A statistical arbitrage strategy.

Chen, C. W., Wang, Z., Sriboonchitta, S., and Lee, S. (2017). Pair trading based on quantile forecasting of smooth transition garch models. *North American Journal of Economics and Finance*, 39:38–55.

Chen, H. J., Chen, S. J., Chen, Z., and Li, F. (2019). Empirical investigation of an equity pairs trading strategy. *Management Science*, 65:370–389.

Chen, Y., Goetzmann, W. N., and Liang, B. (2022). Are they mafias? hedge funds in SPACs.

Dellavigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64:709–749.

Do, B. and Faff, R. (2010). Does simple pairs trading still work? *Financial Analysts Journal*, 66:83–95.

Elliott, R. J., Hoek, J. V. D., and Malcolm, W. P. (2005). Pairs trading. *Quantitative Finance*, 5:271–276.

Engelberg, J., Gao, P., and Jagannathan, R. (2009). An anatomy of pairs trading: the role of idiosyncratic news, common information and liquidity.

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

Gahng, M., Ritter, J. R., and Zhang, D. (2023). SPACs.

Gatev, E., Rouwenhorst, K. G., and Goetzmann, W. (2006). Pairs trading: Performance of a relative value arbitrage rule.

Han, C., He, Z., Jun, A., and Toh, W. (2021). Pairs trading via unsupervised learning.

Hirshleifer, D., Lim, S. S., and Teoh, S. H. (2007). Driven to distraction: Extraneous events and underreaction to earnings news.

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

Huck, N. and Afawubo, K. (2015). Pairs trading and selection methods: is cointegration superior? *Applied Economics*, 47:599–613.

Hutchinson, M. C. and O'Brien, J. (2020). Time series momentum and macroeconomic risk. *International Review of Financial Analysis*, 69:101469.

Jacobs, H. and Weber, M. (2015). On the determinants of pairs trading profitability. *Journal of Financial Markets*, 23:75–97.

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

Krauss, C. (2017). Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of Economic Surveys*, 31:513–545.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708.

Peng, L. and Xiong, W. (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80:563–602.

Poterba, J. M. and Summers, L. H. (1987). Mean reversion in stock prices: Evidence and implications.

Rad, H., Low, R. K. Y., and Faff, R. (2015). The profitability of pairs trading strategies: distance, cointegration, and copula methods.

Rubesam, A. (2021). Pairs Trading: Replicating Gatev, Goetzmann and Rouwenhorst (2006).

Rudy, J., Dunis, C. L., Giorgioni, G., and Laws, J. (2010). Statistical arbitrage and high-frequency data with an application to eurostoxx 50 equities.

Vidyamurthy, G. (2004). *Pairs Trading: Quantitative Methods and Analysis*. John Wiley & Sons, Inc.