JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 59, No. 3, May 2024, pp. 933–955

doi:10.1017/S0022109023000145

Anomaly Discovery and Arbitrage Trading

Xi Dong
City University of New York Baruch College
xi.dong@baruch.cuny.edu

Qi Liu
Peking University Guanghua School of Management
qiliu@gsm.pku.edu.cn (corresponding author)

Lei Lu
University of Manitoba Asper School of Business
lei.lu@umanitoba.ca

Bo Sun
University of Virginia Darden School of Business
sunb@darden.virginia.edu

Hongjun Yan
DePaul University Driehaus College of Business
hongjun.yan.2011@gmail.com

Abstract

We analyze a model in which an anomaly is unknown to arbitrageurs until its discovery, and test the model implications on both asset prices and arbitrageurs' trading activities. Using data on 99 anomalies documented in the existing literature, we find that the discovery of an anomaly reduces the correlation between the returns of its decile-1 and decile-10 portfolios. This discovery effect is stronger if the aggregate wealth of hedge funds is more volatile. Finally, hedge funds increase (reverse) their positions in exploiting anomalies when their aggregate wealth increases (decreases), further suggesting that these discovery effects operate through arbitrage trading.

I. Introduction

A significant portion of the asset-pricing literature has been devoted to "anomalies," empirical patterns that appear inconsistent with existing benchmark

We thank Thummim Cho (the referee) and Jennifer Conrad (the editor) for their constructive comments. We also thank Nick Barberis, Bruno Biais, Alon Brav, David Brown, Bjorn Eraker, Will Goetzmann, Michael Gofman, Paul Goldsmith-Pinkham, David Hirshleifer, Jon Ingersoll, Wenxi Jiang, Marcin Kacperczyk, Andrew Karolyi, Leonid Kogan, Andrew Lo, Benjamin Loos, Steve Malliaris, David McLean, Alan Moreira, Justin Murfin, Lubos Pastor, Anna Pavlova, Jeff Pontiff, Mark Ready, Jialin Yu, Jianfeng Yu, and seminar participants at Boston University, DePaul University, Georgetown University, HKUST, Johns Hopkins University, PBCSF Tsinghua University, Peking University, Rutgers, SAIF, Temple University, SEM Tsinghua University, University of Florida, University of Toronto, University of Virginia, University of Wisconsin Madison, Yale, the 2019 American Economic Association Meetings (AEA), the 2016 European Finance Association Meetings (EFA), the 2016 European Summer Symposium in Financial Markets, the 2015 China International Conference in Finance (CICF), and the 2015 Northern Finance Association Meetings (NFA) for helpful discussions. Lu is grateful for financial support provided by Bryce Douglas Chair in Corporate Finance. An earlier version of this article was circulated under the title "A Model of Anomaly Discovery."

models. There are two contrasting views. The first assumes that investors have always known and understood the anomalies. Hence, the research agenda is to identify the risks that generate the anomalies and the reasons why investors care about those risks. The academic literature has predominantly focused on this view. Take the value premium as an example. Its discovery is often attributed to Basu (1983). Since then, numerous models have been proposed to explain why value stocks are indeed riskier (than what CAPM implies) and lead to higher risk premium. We argue that this view ignores the "discovery aspect." In those risk-based models, investors know that value stocks are riskier and demand higher returns. As expected, higher average returns are realized for value stocks. Hence, there is no real discovery: Professor Basu was the last one in the world to find out about the value premium. Investors knew about this return pattern all along.

The second view is that an anomaly is unknown to some market participants until its discovery. An immediate consequence of this view is that discoveries change the behavior of some investors and future anomaly returns. It seems natural to expect discoveries to have significant effects on investors' decisions and asset prices, as discoveries in academia have had increasingly important influences on the asset management industry. For example, according to the estimate by Morningstar, the total asset under management (AUM) of smart beta strategies, which are directly motivated by anomalies, was almost $1 trillion in 2017. Many prominent asset management companies regularly organize academic seminars and conferences. Some explicitly claim that they identify investment ideas from academic research.[1]

However, the second view has been largely ignored by the literature until recently. McLean and Pontiff (2016) find that anomaly returns tend to decrease significantly after the publication of the first academic study that examines the anomaly. Although the finding is consistent with the second view that anomalies are weakened by arbitrageurs after discovery, the evidence, especially the role of arbitrageurs, is indirect. In this article, we explicitly model the role of arbitrageurs after anomaly discoveries and empirically examine the implications on both asset prices and arbitrageurs' trading.

Specifically, we consider a model with two risky assets, asset 1 and asset 2. The cash flows from these two assets have the same distribution. There are two types of investors and both are mean–variance maximizers. We refer to the first type as "consumers." They find asset 1 riskier because their endowment is positively correlated with asset 1's cash flow. Due to this hedging demand, asset 1 has a lower price and a higher expected return than asset 2 in equilibrium. We call this return pattern an "anomaly" because it is inconsistent with a model that does not account for consumers' hedging demand.

After this anomaly is discovered, the second type of investors, whom we refer to as "arbitrageurs," become aware of the return pattern. Importantly, arbitrageurs do not have the above-mentioned hedging demand, perhaps because they have a different labor income profile and do not face the endowment risk of consumers.

[1]Take Dimensional Fund Advisors as an example. According to its website, as of June 30, 2014, it managed $378 billion. Academic research appears to have a deep influence on its operation, as its website states: "Working closely with leading financial academics, we identify new ideas that may benefit investors."

As a result, arbitrageurs find the return pattern worth exploiting and their trading would alter the equilibrium, leading to the discovery effect.

Specifically, to analyze the discovery effect, we construct an equilibrium without these arbitrageurs, which we call the "pre-discovery equilibrium," and an equilibrium with these arbitrageurs, which we call the "post-discovery equilibrium." The discovery effect is captured by the difference between the pre- and post-discovery equilibria.

Our model has implications on both asset prices and arbitrageurs' trading. We show that, under reasonable conditions, the anomaly discovery reduces the correlation between the returns of assets 1 and 2, and this effect is stronger when arbitrageurs' wealth is more volatile. This is because arbitrageurs increase (reverse) their positions in exploiting the anomaly when their wealth increases (decreases). Specifically, after the discovery, arbitrageurs have a long–short position in assets 1 and 2. Suppose the arbitrageurs' wealth increases due to, say, capital flows from their investors. They will buy asset 1 and sell asset 2. This increases asset 1's return but decreases asset 2's. Similarly, when arbitrageurs' wealth decreases, they will unwind some of their long–short positions, that is, sell asset 1 and buy asset 2, which decreases asset 1's return but increases asset 2's. In both cases, arbitrageurs' portfolio re-balancing pushes the returns of the two assets in opposite directions, reducing their correlation. This intuition also suggests that the effect is stronger when arbitrageurs' wealth is more volatile.

We empirically test these implications based on 99 anomalies that can be constructed based on widely accessible public data. For each anomaly, we construct a dummy variable that takes the value of 0 before its "discovery" and 1 afterward. We use the publication time of the article that documented the anomaly (or latest working article dates for unpublished articles) as a proxy for the discovery time.

We first test the discovery effect on the correlation between the long and short leg returns. Specifically, for each anomaly, we use a 5-year rolling window to estimate the correlation coefficient between the monthly excess returns of deciles 1 and 10. To control for its potential time trend, our analysis focuses on excess correlation: the correlation between deciles 1 and 10 minus the correlation between deciles 5 and 6. The idea is that arbitrageurs are likely to take larger long–short positions in deciles 1 and 10 than in deciles 5 and 6. Hence, the correlation between deciles 5 and 6 should have little discovery effect, but should share the common time trend with the correlation between deciles 1 and 10.

We then regress the excess correlation on the discovery dummy with anomaly fixed effects. The coefficients for the discovery dummy are significantly negative and imply that the discovery of an anomaly reduces the excess correlation measure by 4% to 10%, which represents 33% to 83% of the standard deviation of the measure.

We then link the discovery effect to arbitrageurs. Specifically, our model implies that after the discovery of an anomaly, the excess correlation between its deciles 1 and 10 becomes more negatively correlated with the volatility of arbitrageurs' wealth. To test this prediction, we use the aggregate AUM of U.S. equity hedge funds as a proxy for arbitrageur's wealth. We run a panel regression of the excess correlation measure on the interaction term of the discovery dummy and the hedge fund AUM volatility. Consistent with the model prediction, the estimated interaction coefficients are significantly negative. A 1-standard-deviation increase in the AUM volatility increases the magnitude of the discovery effect on the excess correlation by 3% to 10%. Note that this prediction is the opposite of the conventional wisdom that the correlations among asset returns increase with the market volatility.[2] Hence, our empirical evidence suggests that the discovery effect in our model dominates the effect implied by the conventional intuition.

To further analyze the underlying mechanism of the discovery effect, we directly examine the trading activities of hedge funds. Specifically, we identify hedge funds in the 13F institutional holdings filings. For each anomaly, a quarterly measure of trading intensity by hedge funds is constructed as their aggregate trading of the decile-1 stocks of the anomaly minus that of the decile-10 stocks. A positive value suggests that, in aggregate, hedge funds appear to trade to exploit the anomaly. A negative value suggests that hedge funds appear to trade in the "wrong" direction for the anomaly.

Consistent with our model predictions, the trading intensity measure for an anomaly is significantly positive only after the anomaly's discovery. The average of the post-discovery trading intensity measure is over 12 times that of its pre-discovery level. The implied trading activity is 0.85% of the total shares outstanding of the traded stocks. Our evidence also suggests that, after the discovery of an anomaly, hedge funds expand (unwind) their positions in the anomaly when their aggregate AUM increases (decreases). A 1-standard-deviation increase in the aggregate AUM of all hedge funds leads to an increase in their quarterly trading intensity in the anomaly by up to 0.8% of the total shares outstanding. These effects are economically significant, especially given recent evidence that hedge fund transactions are most effective, among all investors, in affecting asset prices (Dong, Kang, and Peress (2020), Koijen, Richmond, and Yogo (2022)). We also separately analyze the effects on the long and short legs of the anomalies. Consistent with the model prediction, we find that the discovery of an anomaly affects hedge fund trading in both deciles 1 and 10 of the anomaly, but in opposite directions.

Our interpretation suggests that the discovery effect should be stronger for anomalies that attract more attention from arbitrageurs. To test this, for each anomaly, we use the Google citation count of the original study that discovered the anomaly as a proxy for this attention. The idea is that anomalies that are highly cited by both academic and practitioner journals are more likely to be robust and attract more attention from arbitrageurs.

We consider two citation measures: the raw citation counts and the citation counts per year. Then, we run weighted least square (WLS) regressions that assign each anomaly a weight according to its citation measure. Consistent with the interpretations that highly cited anomalies are more likely to be robust and attract arbitrage trading, for all the previously described results, the estimates based on citation-weighted regressions are indeed stronger, providing further support to the underlying mechanism in the model.

[2]Conventional wisdom is that arbitrageurs' wealth tends to be more volatile when the market is more volatile (e.g., in a financial crisis). Since stocks tend to be more correlated when the market is more volatile, this intuition suggests that the correlation between deciles 1 and 10 should be *increasing* in the volatility of arbitrageurs' wealth, the opposite of our model prediction.

Our article contributes to the recent literature on anomaly discovery. McLean and Pontiff (2016) find that anomaly returns tend to decrease significantly and become more correlated with the returns of other existing anomalies after the publication of the first academic study that examines the anomaly. Since then, there has been growing interest in conducting meta-analysis on the common properties of a large number of anomalies.[3] Our stylized model not only generates implications that are consistent with their empirical findings, but also leads to new predictions on both asset prices and arbitrageurs' trading, which are supported by our empirical evidence. Moreover, note that the correlation analyzed in McLean and Pontiff (2016) is between the return of a newly discovered anomaly and the returns of other anomalies. In our study, the correlation is between the long and short legs of an anomaly. Although these two are completely different variables, as illustrated in the model, they both are influenced by arbitrageurs' trading.

Our analysis highlights the unique role of anomaly discovery and arbitrageurs for asset pricing in general. The discovery effect implies that we should not expect a single asset pricing model to explain asset returns in the entire sample of the modern stock market. As anomalies are discovered over time, asset pricing factors also evolve, and therefore cannot be explained by the same known factors throughout the entire sample. This also cautions against the practice of splitting a sample by time to use the first-half sample to estimate parameter values and use those parameters for the analysis of the second-half sample, which essentially treats sub-samples as drawn from the same population of distributions.

Our article is also related to the literature on the comovement caused by arbitrageurs. Lou and Polk (2022) use the high-frequency correlation among stocks to infer the size of arbitrage capital. Liao (2020) shows that arbitrage trading on one anomaly affects an anomaly in another market. In contrast, our analysis focuses on the correlations between the long and short legs of an anomaly. Our article is closely related to Cho (2020), which focuses on the abnormal short selling and shows that arbitrage trading turns an anomaly alpha into a higher beta. Our study is based on hedge fund transactions and focuses on the return correlations between the extreme decile portfolios.

There is a vast theoretical literature on the role of arbitrageurs.[4] A more closely related recent contribution is Hanson and Sunderam (2014), which studies two anomalies and infers the amount of arbitrage capital devoted to each strategy. In contrast, we focus on the effects of the aggregate (rather than strategy-specific) arbitrage capital and its effects on a large number of anomalies.

The rest of the article is as follows: Section II analyzes a model of anomaly discovery. Empirical analysis is reported in Section III, and Section IV concludes. All derivations are in the Appendix.

[3]Examples of such studies are Harvey, Liu, and Zhu (2016), Green, Hand, and Zhang (2017), Engelberg, McLean, and Pontiff (2018), Feng, Giglio, and Xiu (2020), Gu, Kelly, and Xiu (2020), Guo, Li, and Wei (2020), Hou, Xue, and Zhang (2020), Karolyi and Van Nieuwerburgh (2020), and Dong, Li, Rapach, and Zhou (2022).

[4]Early seminal contributions include Dow and Gorton (1994), Pontiff (1996), Shleifer and Vishny (1997), Kyle and Xiong (2001), Gromb and Vayanos (2002), and Pontiff (2006).

II. A Model of Anomaly Discovery

To motivate and organize our empirical analysis, we explore a stylized model in this section.[5] Consider a two-period model, with time $t = 0, 1, 2$. There is one risk-free asset, and its interest rate is normalized to 0. Two risky assets, assets 1 and 2, are in 0 net supply. Asset i, for $i = 1, 2$, pays a dividend of $d_{i,t}$ at time t for $t = 1, 2$. We specify the dividends as follows:

$$d_{1,t+1} = v_t + z_{t+1}, \tag{1}$$

$$d_{2,t+1} = v_t - z_{t+1}, \tag{2}$$

for $t = 0, 1$. Note that v_t is known at t, while z_{t+1} is realized at $t+1$ and has a zero mean and unit variance. This specification is meant to capture the observation that the long and short legs of an anomaly typically have opposite loadings on a systematic factor. For example, value and growth stocks have opposite loadings on the "HML" factor.

At time t, a representative "consumer" is born with wealth $w_t \geq 0$, for $t = 0, 1$. He invests for one period and consumes all his wealth at $t+1$ and his objective function is

$$\max\left\{E_t[c_{t+1}] - \frac{1}{2}\mathrm{var}_t(c_{t+1})\right\}. \tag{3}$$

The specification in equations (1) and (2) implies that the dividends from the two assets always have the same conditional mean and variance. Hence, given the objective function (3), the two assets should always have the same price, unless the consumer prefers one asset over the other due to hedging or behavioral biases. Hence, if the two assets have different prices, we can refer to this as an "anomaly."

We model the anomaly in a reduced form by assuming that the consumer born at time t faces an endowment shock of $w_t z_{t+1}$ at $t+1$. As a result, the consumer treats these two assets differently, leading to an anomaly.

In order to analyze the discovery effect, we assume that at time $t = 0, 1$, a representative "arbitrageur" is born. She lives for one period and consumes all her wealth at $t+1$. Her objective function is

$$\max\left\{E_t\left[c^a_{t+1}\right] - \frac{1}{2}A_t \mathrm{var}_t\left(c^a_{t+1}\right)\right\}, \tag{4}$$

where A_t is the arbitrageur's risk aversion. In the case of $A_t = \infty$, the arbitrageur completely avoids the two risky assets. Hence, it reflects the "pre-discovery case." If $A_t < \infty$, however, the arbitrageur exploits the anomaly and hence it reflects the "post-discovery case." To analyze the discovery effect, we can simply compare the equilibria across the two cases.

Note that the above mean–variance preference (4) implies that the arbitrageur's wealth level does not play a role. Instead, the fluctuations of the arbitrageur's

[5]We thank an anonymous referee for detailed suggestions of the entire setup.

investment intensity are captured by the changes in the risk aversion A_t. Hence, in this setup, the effect of the fluctuations of the arbitrageur's risk-bearing capacity manifests itself through the changes in risk aversion A_t.[6]

For $i=1,2$, and $t=0,1,2$, we use $p_{i,t}$ to denote the price of asset i at time t, which will be determined endogenously in equilibrium. Following the tradition in the micro-structure literature, we use the price change to denote "returns": the return of asset i at time t is defined as $r_{i,t}=p_{i,t}+d_{i,t}-p_{i,t-1}$. The prices of the two assets at $t=2$ are pinned down by the final dividends $d_{1,t}$ and $d_{2,t}$. The Appendix reports the details of the derivation of equilibrium prices at $t=0,1$. Our focus is the equilibrium prices at $t=1$:

(5) $$p_{1,1}=v_1-\frac{A_1}{1+A_1}w_1,$$

(6) $$p_{2,1}=v_1+\frac{A_1}{1+A_1}w_1.$$

The above equations show that due to the hedging demand caused by the endowment shocks, the consumer prefers asset 2 over asset 1. As a result, we have $p_{1,1}<p_{2,1}$ in equilibrium. Moreover, the price difference, $p_{2,1}-p_{1,1}$, is at its maximum in the absence of the arbitrageur (i.e., $A_1=\infty$). It is reduced when the arbitrageur exploits the anomaly (i.e., $A_1<\infty$) and disappears only when the arbitrageur is risk neutral (i.e., $A_1=0$).

From the prices in (5) and (6), we obtain the correlation between the returns of the two assets

(7) $$Corr_0(r_{1,1},r_{2,1})=\frac{\text{var}_0(v_1)-\text{var}(z_1)-\text{var}_0\left(\frac{A_1}{1+A_1}w_1\right)}{\text{var}_0(v_1)+\text{var}(z_1)+\text{var}_0\left(\frac{A_1}{1+A_1}w_1\right)},$$

which will be the focus of our empirical analysis. By contrasting the pre-discovery equilibrium (i.e., $A_1=\infty$) and the post-discovery equilibrium (i.e., $A_1<\infty$), we obtain the discovery effect on $Corr_0(r_{1,1},r_{2,1})$. We first consider the following two extreme cases.

Case 1 (consumer's effect). We shut down the effect from the arbitrageur's risk aversion shock by setting $A_t=A$. It is straightforward to verify from equation (7) that

$$Corr_0(r_{1,1},r_{2,1})|_{A=\infty}<Corr_0(r_{1,1},r_{2,1})|_{A<\infty}.$$

That is, the discovery of the anomaly increases the correlation. Intuitively, a consumer's hedging demand creates a negative correlation between the two assets. An arbitrageur's trading dampens the hedging demand from the consumer and hence increases the correlation.

[6]In the Supplementary Material, we present an alternative model, which explicitly analyzes the arbitrageur's wealth effect.

Case 2 (arbitrageur's effect). We shut down the effect from the consumer's wealth shock by setting $w_t = w$. It is then straightforward to verify from equation (7) that

$$Corr_0(r_{1,1}, r_{2,1})|_{A=\infty} > Corr_0(r_{1,1}, r_{2,1})|_{A<\infty}.$$

That is, the discovery of the anomaly decreases the correlation. Intuitively, the arbitrageur has a long–short position in assets 1 and 2 to exploit the anomaly. Suppose her investment intensity increases (i.e., risk aversion decreases), she will buy asset 1 and sell asset 2. This increases asset 1's return but decreases asset 2's. Similarly, when her investment intensity decreases (i.e., risk aversion increases) she will unwind her long–short positions, that is, sell asset 1 and buy asset 2, which decreases asset 1's return but increases asset 2's. In both cases, the arbitrageur's trading pushes the returns of the two assets in opposite directions, reducing their correlation.

The two cases above show that the overall discovery effect depends on which effect dominates. We argue that the arbitrageur's effect should dominate and hence Case 2 is more empirically relevant for the following three reasons.

First, hedge funds (arguably the most preeminent arbitrageurs in the stock market) routinely employ substantial leverage while typical investors, who derive wealth mostly from labor income, barely have access to leverage. As a result, the value of the hedge fund portfolio tends to be much more volatile than the portfolio value of regular investors.

Second, as is well known in the literature, hedge fund clients have a strong tendency to chase past results (Fung, Hsieh, Naik, and Ramadorai (2008)). A high hedge fund return typically attracts large capital inflows from investors while a low hedge fund return is often followed by withdrawal from its clients. This further amplifies the fluctuations in the AUM.

Third, financial disturbances often alter risk perceptions and margin restrictions. This may interact with the market liquidity, further reinforcing each other and posing substantial risks to highly leveraged arbitrageurs (Brunnermeier and Pedersen (2009), Adrian, Moench, and Shin (2010), and He and Krishnamurthy (2013)).

All three aspects above are vividly illustrated in the 2008 Global Financial Crisis. Hedge funds with high leverage suffered large losses and capital withdrawal, and faced more stringent margin requirements. As a result, the risk-bearing capacity of hedge funds fluctuates substantially, causing large movements in asset prices. Recent empirical evidence is also consistent with the view that arbitrageurs play a prominent role in determining asset prices. For example, Koijen et al. (2022) find that hedge fund transactions are most effective, among all investors, in affecting asset prices. Therefore, it is reasonable to expect the arbitrageur's effect to dominate, leading to the following hypothesis:

Hypothesis 1. The discovery of an anomaly reduces the correlation between the returns of the long and short legs.

The aforementioned intuition also suggests that if the arbitrageur's trading intensity is more volatile, she needs to adjust her long–short positions more drastically, causing a larger reduction in the correlation. In the model, the arbitrageur's trading intensity is determined by her risk aversion, which can be considered a shorthand for the arbitrageur's risk-bearing capacity. In empirical analysis, an arbitrageur's risk-bearing capacity is often represented by her wealth. Hence, we propose the following hypotheses.

Hypothesis 2. The discovery effect in Hypothesis 1 is stronger if the arbitrageur's wealth volatility is larger.

The above two hypotheses focus on the effects on asset prices. The model illustrates that the underlying driving force is the trading by the arbitrageur. On this aspect, we have the following two hypotheses.

Hypothesis 3. After the discovery of an anomaly, the arbitrageur increases her portfolio to exploit the anomaly.

Hypothesis 4. After the discovery of an anomaly, the arbitrageur increases (reverses) her positions in the anomaly when her wealth increases (decreases).

III. Empirical Analysis

In this section, we empirically examine the hypotheses proposed in the previous section. Section III.A describes the data. Section III.B tests the main model prediction that the discovery of an anomaly reduces the correlation between the returns of its long and short legs (Hypothesis 1). To shed light on the underlying mechanism, we directly connect the discovery effect with the aggregate trading by hedge funds. Section III.C connects the discovery effect with the volatility of the aggregate AUM of hedge funds (Hypothesis 2). Sections III.D and III.E examine hedge fund trading intensities by testing Hypotheses 3 and 4, respectively.

A. Data

We obtain our sample of anomalies by combining those studied in Stambaugh, Yu, and Yuan (2012) and Green, Hand, and Zhang (2017). We restrict our analysis to those where the anomaly portfolios can be constructed based on the CRSP, Compustat, and IBES data. We exclude anomalies where the sorting variables are interactions of multiple characteristics. We further restrict our sample to anomalies where the sorting variables are continuous rather than dummies so that decile portfolios can be formed. Our final sample includes 99 anomalies. The list of these anomalies is provided in the Supplementary Material.

For each month t, we calculate the sorting variable for each anomaly based on the information at the end of month $t-1$, assuming that annual accounting data are available if the firm's fiscal year ended at least 6 months ago, and that quarterly accounting data are available if the fiscal quarter ended at least 4 months ago. We exclude stocks with prices below $5 to avoid microstructure effects in penny stocks. Monthly stock returns are from CRSP and include delisting returns. We follow Kozak, Nagel, and Santosh (2018) to construct anomaly portfolio returns with

TABLE 1

Summary Statistics

Table 1 reports the summary statistics on the anomalies in our sample. For each anomaly, its excess correlation X is the correlation between the returns of its deciles 1 and 10 minus the correlation between the returns of its deciles 5 and 6, as defined in equation (8). The first and second columns are based on equal-weighted and value-weighted portfolios, respectively.

	Equal-Weighted	Value-Weighted
No. of anomalies	99	99
Average correlation among anomaly returns	0.05	0.05
Mean publication year of the anomalies	2000	2000
Median publication year of the anomalies	2001	2001
Percentage of anomalies based on working papers	10%	10%
Mean long–short monthly anomaly return	0.49%	0.38%
Number of anomalies with t-statistic >1.65	78	62
Mean of the excess correlation measure X	−0.10	−0.14
Std. of the excess correlation measure X	0.12	0.15

NYSE breakpoints. Our sample of anomaly returns, when available, start from 1926 and all anomaly returns end in Dec. 2020.

We use the publication date as a proxy for the discovery time. For unpublished anomalies, we use the date of the latest working paper.[7] It is not obvious how to choose the "discovery time" for each anomaly. The choice is necessarily subjective to some extent. Hence, we do not take the literal interpretation that those anomalies were secrets before the publication time and became public information afterward. The essence of discovery time in our model is the time when a large number of arbitrageurs start exploiting the anomaly. It is natural to expect the publicity to attract more attention from arbitrageurs after the publication date, compared to before the publication date. Arbitrageurs are therefore more likely to exploit it. If an anomaly has multiple articles that focus on it, we choose the publication date of the most-cited article.

Table 1 provides the summary statistics of the anomalies in our analysis. The average correlation among anomaly returns is low (0.05), consistent with the estimates in other studies (McLean and Pontiff (2016), Green et al. (2017)). Hence, meta-analysis of a large number of anomalies can potentially gain significantly more insights about the systematic properties across anomalies, as each anomaly provides nearly independent information.

Among the 99 anomalies, for equal-weighted (value-weighted) portfolios, 78 (62) of them have average anomaly returns with t-statistics greater than 1.65. The average long–short return of these anomalies is 49 (38) basis points per month for equal-weighted (value-weighted) portfolios.[8]

To measure hedge funds' trading activities, we utilize the classification in Agarwal, Jiang, Tang, and Yang (2013), which combines the information in the 13F institutional holdings data and hedge fund name information from a union of five

[7]Around 90% of the anomalies are based on published papers. Excluding anomalies based on working papers does not change the results of the paper.

[8]This is consistent with the prior evidence that anomaly profits are weaker for value-weighted portfolios (Hou et al. (2020)). We also construct anomaly portfolios after excluding stocks with market cap below the bottom 20th percentile NYSE market cap threshold. As shown in Tables O3 and O4 in the Supplementary Material, our main results remain qualitatively similar.

major hedge fund databases to identify the hedge funds in 13F. The 13F institutional holdings data cover by far the largest number of institutional investors. All institutional investment managers (including foreign investors) that have investment discretion over $100 million in Section 13(F) securities (mostly publicly traded equity) are required to disclose their quarter-end holdings in these securities. A 13F-filing institution is classified as a hedge fund if its major business is sponsoring/managing hedge funds according to the information revealed from a range of sources, including the institution's own websites, SEC filings, industry directories and publications, and news article searches. A Form 13F is filed at the "management company" level rather than at the "portfolio" or the individual fund level. We identify the hedge fund holdings for the period in which we have the hedge fund AUM data. Our final sample consists of 942 unique hedge funds. The holdings data cover 1981–2020, where 1981 is the first full calendar year when 13F holdings data are available.

We obtain monthly hedge fund AUMs from TASS for the period of 1981–2020. Since we examine anomalies in the U.S. equity market, we focus on U.S. equity hedge funds. We compute the percentage change in AUM for each fund, and aggregate them into the value-weighted average of percentage AUM change of all funds. We denote this measure as the percentage change in hedge fund wealth.

B. Correlation

1. Measurements and Specifications

To test Hypothesis 1, we examine whether the correlation coefficient between the excess returns of deciles 1 and 10 decreases after the discovery of the anomaly. For each anomaly *i*, we compute the excess correlation as

$$X_{i,t} \equiv \rho_{i,t}^{1,10} - \rho_{i,t}^{5,6}, \tag{8}$$

where $\rho_{i,t}^{1,10}$ is the correlation coefficient between the monthly excess returns of deciles 1 and 10 of anomaly i during the 5 years prior to month t, and $\rho_{i,t}^{5,6}$ is similarly defined for deciles 5 and 6. This adjustment controls for a potential time trend for the correlation among stocks. The motivation is the following. To exploit the anomaly, arbitrageurs are more likely to take larger long–short positions in deciles 1 and 10 than in deciles 5 and 6. Hence, the correlation between deciles 5 and 6 may share a common time trend with the correlation between deciles 1 and 10, but should be less subject to the discovery effect. As shown in Table 1, the excess correlation has a mean of −0.10 and standard deviation of 0.12.

We then regress the excess correlation $X_{i,t}$ on DISCOVERY_i with anomaly fixed effects, where DISCOVERY_i is a dummy variable that takes the value of 0 before the discovery of anomaly i and 1 afterward.

We consider a variety of weighting schemes that not only examine the robustness of the results but also shed light on the underlying mechanism. Specifically, we consider two weighting methods at the stage of anomaly portfolio formation, and three weighting methods at the regression stage to give certain anomalies higher weights.

When forming anomaly portfolios, we consider both equal- and value-weighted portfolios. It is often noted that anomaly profits are higher for equal-weighted portfolios which focus more on smaller stocks (Hou et al. (2020)). Hence, tests based on equal-weighted portfolios might have a stronger statistical power. However, the analysis based on value-weighted portfolios is perhaps more relevant for assessing economic significance.

We further consider three different weighting methods in our regressions. First, in the baseline regression, all anomalies are weighted equally. However, one might expect that some anomalies attract more attention from arbitrageurs than others. Hence, we use the citation of the original academic study of an anomaly as a proxy for the attention it receives from arbitrageurs. To measure citations, we obtain Google citation counts, as of Oct. 26, 2016, of the studies that first discovered the anomaly.[9] In the second weighting method, anomalies are weighted by their Google citation counts. That is, following Ang, Hodrick, Xing, and Zhang (2006), we use WLS by weighting anomalies based on their Google citation counts. Note that anomalies discovered earlier have more time to accumulate citations. As a result, these raw citation counts tend to place higher weights on anomalies discovered earlier. Hence, in the third weighting method, we weight anomalies based on the citation counts per year. These two citation measures allow us to focus on alternative sets of prominent anomalies. For convenience, we will refer to these three weighting methods as "NoCite-weighted," "RawCite-weighted," and "CitePerYear-weighted," respectively.

2. Regression Results

The first three columns in Panel A of Table 2 report the results based on equal-weighted portfolios. Column 1 shows that when all anomalies are weighted equally in the regression, the coefficient of DISCOVERY is −0.04 ($t = -7.07$). That is, consistent with our model prediction, the correlation between deciles 1 and 10 is reduced, on average, by 0.04 in the post-discovery sample. This effect is stronger if we weight anomalies by their citation counts. As shown in columns 2 and 3, the coefficient of DISCOVERY is −0.06 and −0.05 for the RawCite- and CitePerYear-weighted regressions, respectively. Both are statistically significant at the 1% level. If one interprets the citation count as a proxy for arbitrageurs' attention to an anomaly, these results lend further support to our hypothesis. The results based on value weighted portfolios, reported in columns 4–6, show the same pattern with larger coefficient estimates (−0.05, −0.1, and −0.08 for NoCite-, RawCite-, and CitePerYear-weighted regressions, respectively).

Since all anomalies are based on the same set of stocks, the anomaly portfolios often share stocks. Hence, when one anomaly is discovered, it may affect the correlation measures of related anomalies. Therefore, the discovery effect in the cross section is likely to be weaker. Nevertheless, we also run the regressions with both anomaly and month fixed effects. The results are reported in Panel B of

[9]Although the Google citation service was launched in 2004, it covers articles published before that and hence does not create an obvious bias. Nevertheless, as a robustness check, we split our sample into two based on discovery time. One covers anomalies discovered before the end of 1998 and one covers those after. The results remain similar to those in Table 2.

TABLE 2

Discovery and Correlation

Table 2 reports the results from the panel regressions of the excess correlation $X_{i,t}$, defined in equation (8), on DISCOVERY$_{i,t}$, which is a dummy variable that is 0 when t is before the anomaly i's discovery time and 1 afterward. The observations in regressions are weighted in three ways: i) equal-weighted; ii) weighted by anomaly's raw citation counts; and iii) weighted by anomaly's citation counts per year. The citation count of an anomaly is its Google citation count as of Oct. 26, 2016. The regressions in Panel A include anomaly fixed effects and those in Panel B include both anomaly fixed effects and month fixed effects. Constant terms are omitted. t-statistics, reported in parentheses, are based on standard errors that are double-clustered by anomaly and time (month). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Anomaly Fixed Effects

	Equal-Weighted Anomaly Portfolios			Value-Weighted Anomaly Portfolios		
	NoCite Weight	RawCite Weight	CitePerYear Weight	NoCite Weight	RawCite Weight	CitePerYear Weight
Dependent Variable: $X_{i,t}$	1	2	3	4	5	6
DISCOVERY$_{i,t}$	−0.04***	−0.06***	−0.05***	−0.05***	−0.10***	−0.08***
	(−7.07)	(−8.23)	(−6.86)	(−5.33)	(−5.07)	(−4.61)
Anomaly fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	No	No	No	No	No	No
No. of obs.	80,309	80,309	80,309	80,309	80,309	80,309
R^2	0.37	0.71	0.64	0.40	0.77	0.65

Panel B. Anomaly and Time Fixed Effects

	Equal-Weighted Anomaly Portfolios			Value-Weighted Anomaly Portfolios		
	NoCite Weight	RawCite Weight	CitePerYear Weight	NoCite Weight	RawCite Weight	CitePerYear Weight
Dependent Variable: $X_{i,t}$	1	2	3	4	5	6
DISCOVERY$_{i,t}$	−0.02**	−0.06***	−0.05***	−0.02**	−0.10***	−0.07***
	(−2.36)	(−8.14)	(−6.42)	(−2.03)	(−5.48)	(−5.27)
Anomaly fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
No. of obs.	80,309	80,309	80,309	80,309	80,309	80,309
R^2	0.52	0.74	0.68	0.51	0.78	0.67

Table 2. The overall evidence remains similar. In specifications where anomalies are weighted equally (columns 1 and 4), the coefficient of DISCOVERY is −0.02 and is statistically significant at the 5% level. In the other four columns, where anomalies are weighted by their citation counts, the estimated discovery effect is close to that in Panel A.

3. Robustness

In the above analysis, an anomaly's discovery time is chosen as the publication year of the main academic article on the anomaly. It is conceivable that the more appropriate discovery time can be earlier or later. We conduct a simulation to examine the effect of the choice of the discovery time. Specifically, we randomly assign a pseudo discovery year to each anomaly from the distribution of anomaly publication years in our sample. We then rerun the regressions in Panel A of Table 2 to obtain the estimate of the coefficient of DISCOVERY. We iterate this procedure 10,000 times to obtain the simulated distribution of the regression coefficient estimates. If the publication time is close to the true discovery time (i.e., the time when arbitrageurs start exploiting the anomaly), the estimate of the regression coefficient based publication time should be stronger than most of the estimates based on randomly assigned pseudo discovery time. This is indeed the case. For

https://doi.org/10.1017/S0022109023000145 Published online by Cambridge University Press

equal-weighted anomaly portfolios, for example, only 4.6% of the estimated effects based on pseudo discovery time are stronger than the estimate based on the publication time.[10]

One potential concern on the citation-weighted analysis is the skewness of citation counts. In our sample, the skewness estimates of Cite and CitePerYear are 2.37 and 2.24, respectively. Hence, this weighting method places significant weights on a small number of anomalies. To further examine the robustness of our results, we use log(1 + Cite) and log(1 + CitePerYear) as two alternative weighting schemes for our regressions. The skewness estimates of log(1 + Cite) and log(1 + CitePerYear) are −0.02 and −0.08, respectively. The estimated discovery effect based on these two alternative weights remain similar.

In summary, the above analysis shows a strong and robust discovery effect on the correlation between the long- and short-leg returns, especially for anomalies that are more likely to attract arbitrageurs' attention. The correlation between deciles 1 and 10 of an anomaly decreases by 4% to 10% after its discovery. This is a sizable reduction and represents 33% to 83% of the standard deviation of the correlation measure.

C. The Role of Arbitrageurs

In our model, the discovery effect operates through arbitrageurs: their trading activity reduces the correlation between deciles 1 and 10. Hence, a direct test of this view is to examine whether this correlation is indeed related to arbitrageurs' trading activities.

Hypothesis 2 implies that the post-discovery correlation between deciles 1 and 10 of an anomaly is decreasing in the volatility of arbitrageurs' wealth. To test this hypothesis, we need a proxy for the volatility of arbitrageurs' wealth. It is impossible to directly observe the wealth of arbitrageurs. As a compromise, we measure the aggregate AUM of hedge funds, which are often considered to be the archetypal arbitrageurs in financial markets. The implicit assumption is that the volatility of the aggregate AUM of all hedge funds is positively correlated with the volatility of the wealth of arbitrageurs. For each month during 1986–2020, hedge fund wealth volatility is calculated as the standard deviation of the percentage change in hedge fund wealth during the previous 5 years, excluding the current month t. Since this volatility measure is persistent and has a time trend, we follow Chen, Da, and Huang (2019) and Lo and Wang (2000) to detrend it. Specifically, we take the difference between wealth volatility and its 12-month moving average, and normalize this difference by the 12-month moving average. We denote this detrended volatility measure as WEALTH_VOL_t.

We then regress the excess correlation measure $X_{i,t}$ on the interaction between the discovery dummy and the hedge fund wealth volatility $\text{DISCOVERY}_i \times \text{WEALTH\_VOL}_t$. Our model implies that the discovery effect should be stronger when the hedge fund wealth volatility is higher, that is, the interaction coefficient should be negative.

[10]Details on this simulation are reported in Table O2 in the Supplementary Material.

Consistent with this prediction, column 2 of Table 3 shows that the estimate of the coefficient for the interaction term is −0.03 ($t = -3.05$). As shown in columns 4 and 6, the estimated interaction coefficient is larger in RawCite- and CitePerYear-weighted regressions. This finding is consistent with the interpretation that the discovery effect is stronger for anomalies that are more robust and attract more attention from arbitrageurs.

It is interesting to contrast these results with the conventional intuition that hedge fund wealth volatility tends to be higher when the stock market is more volatile (e.g., in a financial crisis). Since individual stocks tend to be more correlated when the market is more volatile, this conventional intuition implies a *positive* relation between hedge fund wealth volatility and the correlation between deciles 1 and 10 (the opposite of our model prediction). To examine this prediction, we regress the excess correlation measure on the hedge fund wealth volatility. As shown in columns 1, 3, and 5 of Table 3, the relation between the excess correlation measure and the hedge fund wealth volatility is indeed significantly positive. Hence, the results in columns 2, 4, and 6 suggest that the post-discovery arbitrageur's effect dominates the effect from the conventional intuition, turning the correlation between $X_{i,t}$ and hedge fund wealth volatility negative.

We repeat our analysis based on value-weighted portfolios, and the results remain similar. For example, as shown in columns 8, 10, and 12 of Table 3, a 1-standard-deviation increase in wealth volatility will result in a 3% to 10% greater reduction in correlation post discovery than before discovery, representing 25% to 83% of the standard deviation of the correlation measure.

D. Discovery and Arbitrage Trading

In this section, we test Hypothesis 3. That is, we directly examine if the discovery of an anomaly is followed by more hedge fund trading that exploits the anomaly. For each stock, we construct the aggregate hedge fund holdings in this stock as the number of shares held by all hedge funds at the end of the quarter divided by the total number of shares outstanding. We then follow Chen et al. (2019) to measure the arbitrage trading in this stock as the current-quarter hedge fund holdings minus the average hedge fund holdings in the stock over the past four quarters.[11] Finally, for each anomaly, following Puckett and Yan (2011) and Dong, Feng, and Sadka (2019), we construct its hedge fund trading intensity measure as the average arbitrage trading in the decile-1 stocks of the anomaly minus that of the decile-10 stocks, weighted by the value of the stocks traded.

Specifically, for each anomaly, we construct its hedge fund trading intensity as

$$A_{it} = L_{it} - S_{it}, \tag{9}$$

where L_{it} and S_{it} are the average arbitrage trading in the long- and short-leg stocks, respectively. A positive value of A_{it} suggests that hedge funds buy more (or sell less) decile-1 stocks than decile-10 stocks. That is, hedge funds appear to be trading in

[11] Some studies have used short positions to infer arbitrage trading. This approach is less effective for our purposes since many short positions are reported to be missing in the Compustat database, especially before 2000 (Calluzzo, Moneta, and Topaloglu (2019), Chen et al. (2019)).

TABLE 3

Discovery, Wealth Volatility, and Correlation

Table 3 reports the results from the panel regressions of the excess correlation $X_{i,t}$, defined in equation (8), on the dummy variable $\text{DISCOVERY}_{i,t}$, the wealth volatility WEALTH_VOL_t (scaled by its standard deviation), and their interaction term. $\text{DISCOVERY}_{i,t}$ is 0 when t is before the anomaly i's discovery time and 1 afterward. WEALTH_VOL_t is the detrended aggregate AUM volatility of U.S.-equity-focused hedge funds. The observations in the regressions are weighted in three ways: i) equal-weighted; ii) weighted by anomaly's raw citation counts; and iii) weighted by anomaly's citation counts per year. The citation count of an anomaly is its Google citation count as of Oct. 26, 2016. Constant terms are omitted. t-statistics, reported in parentheses, are based on standard errors that are double-clustered by anomaly and time (month). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	Equal-Weighted Anomaly Portfolios						Value-Weighted Anomaly Portfolios					
	NoCite Weight	NoCite Weight	RawCite Weight	RawCite Weight	Cite PerYear Weight	Cite PerYear Weight	NoCite Weight	NoCite Weight	RawCite Weight	RawCite Weight	Cite PerYear Weight	Cite PerYear Weight
Dependent Variable: $X_{i,t}$	1	2	3	4	5	6	7	8	9	10	11	12
WEALTH_VOL_t	0.005**	0.02***	0.01**	0.03***	0.03**	0.08***	0.005*	0.02***	0.01**	0.04***	0.03**	0.09***
	(2.01)	(3.80)	(2.47)	(3.99)	(2.47)	(2.79)	(1.84)	(4.55)	(2.54)	(3.85)	(2.53)	(2.85)
$\text{DISCOVERY}_{i,t} \times \text{WEALTH\_VOL}_t$		−0.03***		−0.04***		−0.1**		−0.03***		−0.05***		−0.1**
		(−3.05)		(−2.98)		(−2.31)		(−4.10)		(−2.97)		(−2.35)
$\text{DISCOVERY}_{i,t}$		−0.01*		−0.08***		−0.07***		−0.01		−0.08***		−0.06***
		(−1.75)		(−7.96)		(−4.95)		(−1.56)		(−6.88)		(−4.31)
Anomaly FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of obs.	37,409	37,409	37,409	37,409	37,409	37,409	37,409	37,409	37,409	37,409	37,409	37,409
R^2	0.45	0.45	0.72	0.73	0.70	0.72	0.48	0.48	0.74	0.75	0.73	0.74

https://doi.org/10.1017/S0022109023000145 Published online by Cambridge University Press

the "right" direction to exploit the anomaly. Likewise, a negative value can be interpreted as hedge funds trading in the "wrong" direction.

Panel A of Table 4 reports the mean of these trading intensity measures. The average trading intensity A_{it} across all anomalies is 0.05% in the pre-discovery sample. It increases more than 12-fold to 0.62% in the post-discovery sample. This simple comparison is consistent with the hypothesis that hedge funds start trading in the "right" direction after the discovery of an anomaly. The average of L_{it} is 0.06% and 0.32% in the pre- and post-discovery samples, respectively. That is, after an anomaly is discovered, hedge funds appear to buy more of the stocks in the long leg of the anomaly. The average of S_{it} is 0.01% and −0.31% in the pre- and post-discovery samples, respectively. After the discovery of an anomaly, hedge funds appear to reduce their holdings of the stocks in the short leg of the anomaly.

We regress the trading intensity measure A_{it} on DISCOVERY$_{it}$, with A_{it-1} and A_{it-2} as controls. The regression results are reported in columns 1–3 of Panel B of Table 4. Consistent with the model prediction, the coefficients of DISCOVERY are 0.3%, 0.8%, and 0.8% ($t = 2.92$, 2.53, and 2.37) for the NoCite-, RawCite-, and CitePerYear-weighted regressions, respectively. Columns 2 and 3 show that the estimated effects in RawCite- or CitePerYear-weighted regressions are close to three times as large as the effect in the NoCite-weighted regression in column 1. These results suggest that following the discovery of an anomaly, hedge funds exploit the anomaly more intensively. This effect is especially strong for anomalies that are more likely to be robust and attract more attention from arbitrageurs.

These estimated effects are economically significant. They imply that, on average, the discovery of an anomaly is accompanied by an increase in the quarterly arbitrage trading intensity of 0.3% to 0.8% of the total shares outstanding of the traded stocks. This is equivalent to $3 to $9 billion of additional trading volume per quarter.

We also examine the long and short legs separately. We first regress L_{it} on DISCOVERY$_{it}$. As shown in columns 4–6 of Table 4, the estimates of the coefficient of DISCOVERY$_{it}$ are significantly positive. That is, following the discovery, hedge funds buy more stocks of the long leg of the anomaly. We then regress S_{it} on DISCOVERY$_{it}$. As shown in columns 7–9, the estimates of the coefficient of DISCOVERY$_{it}$ are significantly negative. That is, following the discovery, hedge funds appear to reduce their holdings in the stocks of the short leg of the anomaly.

These results add to the debate on whether institutional investors contribute to correcting or exacerbating anomalies by trading in the "right" or "wrong" direction, or whether they simply play no roles, unconditionally or after discovery. On the one hand, Edelen, Ince, and Kadlec (2016) show that financial institutions trade in the wrong direction of anomalies and appear to cause anomalies. Other studies suggest that institutions are merely noise traders in the presence of anomalous returns (Ali, Chen, Yao, and Yu (2008)), play no role in anomalies (Lewellen (2011)) or even drive sentiment-induced mispricing (DeVault, Sias, and Starks (2019)).[12] On the

[12]Based on meta-analysis of a large number of anomalies in all countries around the world except the United States, Jacobs (2016) provides suggestive evidence that unconditionally, arbitrageurs may trade on the wrong side of anomalies; Jacobs and Muller (2020) further show that even after publication, arbitrageurs do not exploit anomalies, possibly due to limits to arbitrage.

TABLE 4

Discovery and Arbitrage Trading: Hedge Funds Evidence

Panel A of Table 4 reports the mean of hedge funds' anomaly trading intensities for the pre- and post-discovery sample periods. $L_{i,q}$ and $S_{i,q}$ are hedge funds' trading intensities in the long and short legs of anomaly i in quarter q, respectively. $A_{i,q}$ is $L_{i,q}$ minus $S_{i,q}$. Panel B reports regression results. The dependent variables are $A_{i,q}$, $L_{i,q}$, and $S_{i,q}$. The independent variables include the dummy variable DISCOVERY$_{i,q}$, $A_{i,q-1}$, and $A_{i,q-2}$. DISCOVERY$_{i,q}$ is 0 when q is before the anomaly i's discovery time and 1 afterward. The observations in regressions are weighted in three ways: i) equal-weighted; ii) weighted by anomaly's raw citation counts; and iii) weighted by anomaly's citation counts per year. The citation counts of an anomaly are its Google citation counts as of Oct. 26, 2016. Constant terms are omitted. t-statistics, reported in parentheses, are based on standard errors that are double clustered by anomaly and time (quarter). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A. Mean Anomaly Trading Intensity

	$A_{i,q}$ (%)	$L_{i,q}$ (%)	$S_{i,q}$ (%)
Before discovery	0.05	0.06	0.01
After discovery	0.62	0.32	−0.31

Panel B. Discovery Effect on Arbitrage Anomaly Trading Intensity

Dependent Variable	$A_{i,q}$			$L_{i,q}$			$S_{i,q}$		
	NoCite Weight	RawCite Weight	Cite PerYear Weight	NoCite Weight	RawCite Weight	Cite PerYear Weight	NoCite Weight	RawCite Weight	Cite PerYear Weight
	1	2	3	4	5	6	7	8	9
DISCOVERY$_{i,q}$	0.3***	0.8**	0.8**	0.2**	0.4**	0.4**	−0.2**	−0.5**	−0.5**
	(2.92)	(2.53)	(2.37)	(2.38)	(2.05)	(2.05)	(−2.14)	(−2.16)	(−2.00)
$A_{i,q-1}$	0.3***	0.4***	0.4***	0.3***	0.3***	0.3***	0.2***	0.3***	0.3***
	(8.53)	(6.17)	(6.84)	(10.65)	(5.42)	(5.73)	(5.52)	(14.34)	(11.91)
$A_{i,q-2}$	−0.02	0.02	0.02	0.003	0.03	0.03	−0.01	−0.04	−0.03
	(−0.84)	(0.67)	(0.83)	(0.16)	(1.03)	(1.29)	(−0.61)	(−1.06)	(−0.90)
Anomaly FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of obs.	14,586	14,586	14,586	14,586	14,586	14,586	14,586	14,586	14,586
R^2	0.15	0.35	0.32	0.14	0.21	0.20	0.12	0.21	0.20

other hand, Calluzzo et al. (2019) analyze 14 anomalies and provide evidence that financial institutions appear to increase their positions that exploit the anomalies after they are broadly publicized. These studies are typically based on a small number of anomalies, while our results are based on a broad set of anomalies. Moreover, our analysis is guided by a theoretical framework, which allows us to exploit its novel implication on the behavior of the aggregate arbitrage capital.

E. Wealth Change and Arbitrage Trading

Hypothesis 4 suggests that after the discovery of an anomaly, arbitrageurs will increase (unwind) their positions that exploit the anomaly, when their wealth increases (decreases). To test this prediction, we regress A_{it} on the interaction term of the discovery dummy and the percentage change in hedge fund wealth (scaled by the standard deviation of the change).

The results are reported in columns 1–3 of Table 5. Consistent with our model prediction, they show that the relation between the change in hedge fund wealth and the arbitrage trading intensity in an anomaly becomes significantly more positive after the discovery of the anomaly. Specifically, the coefficient of the interaction term is 0.5% ($t = 2.46$), 0.8% ($t = 2.54$), and 0.7% ($t = 2.54$) for the NoCite-, RawCite-, and CitePerYear-weighted regressions, respectively. That is, after the discovery of an anomaly, the measure of arbitrage trading intensity for the anomaly becomes more correlated with the fluctuation in the AUM of hedge

https://doi.org/10.1017/S0022109023000145 Published online by Cambridge University Press

TABLE 5

Discovery, Wealth Change, and Arbitrage Trading: Hedge Funds Evidence

Table 5 reports regression results. The dependent variables are arbitrageurs' anomaly trading intensity $A_{i,q}$, their trading intensity in the long and short legs $L_{i,q}$ and $S_{i,q}$. The independent variables include the dummy variable DISCOVERY$_{i,q}$, $A_{i,q-1}$, $A_{i,q-2}$, and ΔW_q. DISCOVERY$_{i,q}$ is 0 when q is before the anomaly i's discovery time and 1 afterward. ΔW_q is the change in the aggregate AUM of U.S.-equity-focused hedge funds in quarter q, scaled by its standard deviation. The observations in regressions are weighted in three ways: i) equal-weighted; ii) weighted by anomaly's raw citation counts; and iii) weighted by anomaly's citation counts per year. The citation counts of an anomaly are its Google citation counts as of Oct. 26, 2016. Constant terms are omitted. t-statistics, reported in parentheses, are based on standard errors that are double-clustered by anomaly and time (quarter). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	$A_{i,q}$			$L_{i,q}$			$S_{i,q}$		
	NoCite Weight	RawCite Weight	CitePerYear Weight	NoCite Weight	RawCite Weight	CitePerYear Weight	NoCite Weight	RawCite Weight	CitePerYear Weight
	1	2	3	4	5	6	7	8	9
ΔW_q	0.04	0.1*	0.1*	0.02	0.03	−0.010	−0.03	−0.0008	−0.1
	(0.69)	(1.86)	(1.76)	(0.41)	(0.49)	(−0.15)	(−0.41)	(−1.38)	(−1.63)
DISCOVERY$_{i,q} \times \Delta W_q$	0.5**	0.8**	0.7**	0.3**	0.4***	0.4***	−0.2*	−0.4**	−0.3**
	(2.46)	(2.54)	(2.54)	(2.05)	(2.69)	(2.64)	(−1.89)	(−2.23)	(−2.01)
DISCOVERY$_{i,q}$	0.4**	0.7**	0.6**	0.08	0.1	0.1	−0.2**	−0.3*	−0.5
	(2.58)	(2.00)	(2.04)	(1.01)	(1.02)	(1.02)	(−2.22)	(−1.85)	(−1.56)
$A_{i,q-1}$	0.2***	0.4***	0.4***	0.4***	0.4***	0.4***	0.3***	0.4***	0.4***
	(8.18)	(5.66)	(5.88)	(12.70)	(9.80)	(10.38)	(6.06)	(7.19)	(7.19)
$A_{i,q-2}$	−0.02	0.02	0.03	0.04	−0.04	−0.002	0.03	0.05	0.05
	(−0.76)	(0.61)	(0.91)	(1.63)	(−1.24)	(−0.05)	(1.07)	(1.40)	(1.32)
Anomaly fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. of obs.	14,400	14,400	14,400	14,400	14,400	14,400	14,400	14,400	14,400
R^2	0.13	0.36	0.33	0.29	0.32	0.32	0.21	0.47	0.44

funds. It is interesting to note that, once again, the results based on RawCite- and CitePerYear-weighted regressions are much stronger than those based on NoCite-weighted regressions. This finding is consistent with the interpretation that the effect is stronger for anomalies that are more robust and attract more attention from arbitrageurs.

To examine the long and short legs separately, we regress L_{it} and S_{it} on the interaction term between DISCOVERY$_{it}$ and the percentage change in hedge fund wealth. The results are reported in columns 4–9 of Table 5. Consistent with the model prediction, the interaction coefficient is significantly positive for the long leg (columns 4–6) and significantly negative for the short leg (columns 7–9). These results suggest that after the discovery of an anomaly, an increase of the aggregate hedge fund AUM is followed by more purchase of the stocks in the long leg of the anomaly compared to those in the short leg.

IV. Conclusion

We analyze a stylized model of anomaly discovery, which has implications for both asset prices and arbitrageurs' trading. Our model shows that the discovery of an anomaly reduces the correlation between the returns of its long- and short-leg portfolios, and that this effect is stronger when arbitrageurs' wealth is more volatile. We empirically test these predictions based on 99 anomalies and find clear evidence consistent with our model. Moreover, we also find evidence that after the discovery of an anomaly, hedge funds increase their positions that exploit the anomaly. They also increase (unwind) such trades when their wealth increases (decreases), further supporting the view that the discovery effects work through arbitrage trading.

Finally, all these above results are stronger for anomalies discovered by studies with more citations. This is consistent with the interpretation that the effect is stronger for anomalies that are more robust and attract more attention from arbitrageurs.

Appendix. Proofs

We derive the equilibrium prices by backward induction. Let $x_{i,t}$ ($x^a_{i,t}$) be consumer's (arbitrageur's) holding in asset i at $t=0,1$. In the second period, the time 2's consumption of the consumer and the arbitrageur are

$$\begin{aligned} c_2 &= w_1 + (v_1 + z_2 - p_{1,1})x_{1,1} + (v_1 - z_2 - p_{2,1})x_{2,1} + w_1 z_2, \\ c^a_2 &= w^a_1 + (v_1 + z_2 - p_{1,1})x^a_{1,1} + (v_1 - z_2 - p_{2,1})x^a_{2,1}. \end{aligned}$$

At $t=1$, the consumer's first order conditions (FOCs) with respect to $x_{1,1}$ and $x_{2,1}$ lead to the following demand functions:

$$\begin{aligned} p_{1,1} &= v_1 - w_1 - x_{1,1} + x_{2,1}, \\ p_{2,1} &= v_1 + w_1 + x_{1,1} - x_{2,1}. \end{aligned}$$

Similarly, the arbitrageur's FOCs combined with market clearing $x_{i,1} + x^a_{i,1} = 0$ ($i=1,2$) imply the following supply functions:

$$p_{1,1} = v_1 + A_1(x_{1,1} - x_{2,1}), \tag{A-1}$$

$$p_{2,1} = v_1 + A_1(x_{2,1} - x_{1,1}). \tag{A-2}$$

From the above supply and demand functions, we obtain the equilibrium prices (5) and (6).

The consumer's consumption at $t=1$ is

$$\begin{aligned} c_1 &= w_0 + (p_{1,1} + d_{1,1} - p_{1,0})x_{1,0} + (p_{2,1} + d_{2,1} - p_{2,0})x_{2,0} + w_0 z_1 \\ &= w_0 + \left(v_1 - \frac{A_1}{1+A_1}w_1 + v_0 + z_1 - p_{1,0}\right)x_{1,0} \\ &\quad + \left(v_1 + \frac{A_1}{1+A_1}w_1 + v_0 - z_1 - p_{2,0}\right)x_{2,0} + w_0 z_1. \end{aligned}$$

The consumer's FOCs with respect to $x_{1,0}$ and $x_{2,0}$ imply

$$\begin{aligned} &E_0[v_1] - E_0\left[\frac{A_1}{1+A_1}w_1\right] + v_0 - p_{1,0} \\ &= \text{var}(v_1)(x_{1,0} + x_{2,0}) - \text{var}\left(\frac{A_1}{1+A_1}w_1\right)(x_{2,0} - x_{1,0}) + (x_{1,0} - x_{2,0} + w_0), \\ &E_0[v_1] + E_0\left[\frac{A_1}{1+A_1}w_1\right] + v_0 - p_{2,0} \\ &= \text{var}(v_1)(x_{1,0} + x_{2,0}) + \text{var}\left(\frac{A_1}{1+A_1}w_1\right)(x_{2,0} - x_{1,0}) - (x_{1,0} - x_{2,0} + w_0). \end{aligned}$$

Similarly, we obtain the arbitrageur's FOCs with respect to $x^a_{1,0}$ and $x^a_{2,0}$:

https://doi.org/10.1017/S0022109023000145 Published online by Cambridge University Press

$$
\begin{aligned}
&E_0[v_1] - E_0\left[\frac{A_1}{1+A_1} w_1\right] + v_0 - p_{1,0} \\
&= A_0\left(\operatorname{var}(v_1)\left(x^a_{1,0} + x^a_{2,0}\right) - \operatorname{var}\left(\frac{A_1}{1+A_1} w_1\right)\left(x^a_{2,0} - x^a_{1,0}\right) + \left(x^a_{1,0} - x^a_{2,0}\right)\right), \\
&E_0[v_1] + E_0\left[\frac{A_1}{1+A_1} w_1\right] + v_0 - p_{2,0} \\
&= A_0\left(\operatorname{var}(v_1)\left(x^a_{1,0} + x^a_{2,0}\right) + \operatorname{var}\left(\frac{A_1}{1+A_1} w_1\right)\left(x^a_{2,0} - x^a_{1,0}\right) - \left(x^a_{1,0} - x^a_{2,0}\right)\right).
\end{aligned}
$$

From the market clearing condition $x_{i,0} + x^a_{i,0} = 0$ $(i = 1,2)$, we obtain:

$$
\begin{aligned}
p_{1,0} &= v_0 + E_0[v_1] - E_0\left[\frac{A_1}{1+A_1} w_1\right] - \frac{A_0}{1+A_0} w_0, \\
p_{2,0} &= v_0 + E_0[v_1] + E_0\left[\frac{A_1}{1+A_1} w_1\right] + \frac{A_0}{1+A_0} w_0.
\end{aligned}
$$

Equation (7) shows that the post-discovery correlation is higher in Case 1 and lower in Case 2. Hence, this implies that when the arbitrageur's effect dominates, the post-discovery correlation is lower, leading to Hypothesis 1.

As noted in the paper, the arbitrageur's "wealth effect" is captured by the fluctuations in the risk aversion. A larger fluctuation in the arbitrageur's wealth manifests itself as a more volatile risk aversion. In Case 2, equation (7) becomes

(A-3)
$$
Corr_0(r_{1,1}, r_{2,1}) = \frac{\operatorname{var}_0(v_1) - \operatorname{var}(z_1) - w\operatorname{var}_0\left(\frac{A_1}{1+A_1}\right)}{\operatorname{var}_0(v_1) + \operatorname{var}(z_1) + w\operatorname{var}_0\left(\frac{A_1}{1+A_1}\right)},
$$

Hence, the magnitude of the discovery effect in increasing in $\operatorname{var}_0\left(\frac{A_1}{1+A_1}\right)$. Through Taylor expansion at 0, we have $\operatorname{var}_0\left(\frac{A_1}{1+A_1}\right) \approx \operatorname{var}_0(A_1)$. This leads to Hypothesis 2 that the discovery effect is stronger if the arbitrageur's wealth is more volatile.

Equations (A-1) and (A-2) imply that $v_1 - p_{1,1} = A_1\left(x^a_{1,1} - x^a_{2,1}\right)$. Substituting (5) into this equation, we obtain $x^a_{1,1} - x^a_{2,1} = \frac{1}{1+A_1} w_1$. The arbitrageur's positions in the long–short portfolio increase from 0 to a positive value after the discovery, leading to Hypothesis 3. Moreover, as the arbitrageur's risk aversion A_1 decreases (i.e., as her risk-bearing capacity increases), she will increase her positions in the portfolio, leading to Hypothesis 4.

Supplementary Material

To view supplementary material for this article, please visit http://doi.org/10.1017/S0022109023000145.

References

Adrian, T.; E. Moench; and H. S. Shin. "Financial Intermediation, Asset Prices, and Macroeconomic Dynamics." *Staff Report*, 422 (2010).

Agarwal, V.; W. Jiang; Y. Tang; and B. Yang. "Uncovering Hedge Fund Skill from the Portfolio Holdings They Hide." *Journal of Finance*, 68 (2013), 739–783.

Ali, A.; X. Chen; T. Yao; and T. Yu. "Do Mutual Funds Profit from the Accruals Anomaly?" *Journal of Accounting Research*, 46 (2008), 1–26.

Ang, A.; R. J. Hodrick; Y. Xing; and X. Zhang. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, 61 (2006), 259–299.

Basu, S. "The Relationship Between Earnings' Yield, Market Value and Return for NYSE Common Stocks: Further Evidence." *Journal of Financial Economics*, 12 (1983), 129–156.

Brunnermeier, M., and L. Pedersen. "Market Liquidity and Funding Liquidity." *Review of Financial Studies*, 22 (2009), 2201–2238.

Calluzzo, P.; F. Moneta; and S. Topaloglu. "When Anomalies Are Publicized Broadly, Do Institutions Trade Accordingly?" *Management Science*, 65 (2019), 4555–4574.

Chen, Y.; Z. Da; and D. Huang. "Arbitrage Trading: The Long and the Short of It." *Review of Financial Studies*, 32 (2019), 1608–1646.

Cho, T. "Turning Alphas into Betas: Arbitrage and Endogeneous Risk." *Journal of Financial Economics*, 137 (2020), 550–570.

DeVault, L.; R. Sias; and L. Starks. "Sentiment Metrics and Investor Demand." *Journal of Finance*, 74 (2019), 985–1024.

Dong, X.; S. Feng; and R. Sadka. "Liquidity Risk and Mutual Fund Performance." *Management Science*, 65 (2019), 1020–1041.

Dong, X., N. Kang, and J. Peress. "Fast and Slow Arbitrage: Fund Flows and Mispricing in the Frequency Domain." Working Paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3688189 (2020).

Dong, X.; Y. Li; D. E. Rapach; and G. Zhou. "Anomalies and the Expected Market Return." *Journal of Finance*, 77 (2022), 639–681.

Dow, J., and G. Gorton. "Arbitrage Chains." *Journal of Finance*, 49 (1994), 819–849.

Edelen, R. M.; O. S. Ince; and G. B. Kadlec. "Institutional Investors and Stock Return Anomalies." *Journal of Financial Economics*, 119 (2016), 472–488.

Engelberg, J.; R. D. McLean; and J. Pontiff. "Anomalies and News." *Journal of Finance*, 73 (2018), 1971–2001.

Feng, G.; S. Giglio; and D. Xiu. "Taming the Factor Zoo: A Test of New Factors." *Journal of Finance*, 75 (2020), 1327–1370.

Fung, W.; D. A. Hsieh; N. Y. Naik; and T. Ramadorai. "Hedge Funds: Performance, Risk, and Capital Formation." *Journal of Finance*, 63 (2008), 1777–1803.

Green, J.; J. Hand; and X. F. Zhang. "The Characteristics that Provide Independent Information About Average U.S. Monthly Stock Returns." *Review of Financial Studies*, 30 (2017), 4389–4436.

Gromb, D., and D. Vayanos. "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs." *Journal of Financial Economics*, 66 (2002), 361–407.

Gu, S.; B. Kelly; and D. Xiu. "Empirical Asset Pricing via Machine Learning." *Review of Financial Studies*, 33 (2020), 2223–2273.

Guo, L.; F. W. Li; and K. J. Wei. "Security Analysts and Capital Market Anomalies." *Journal of Financial Economics*, 137 (2020), 204–230.

Hanson, S. G., and A. Sunderam. "The Growth and Limits of Arbitrage: Evidence from Short Interest." *Review of Financial Studies*, 27 (2014), 1238–1286.

Harvey, C.; Y. Liu; and H. Zhu. "… and the Cross-Section of Expected Returns." *Review of Financial Studies*, 29 (2016), 5–68.

He, Z., and A. Krishnamurthy. "Intermediary Asset Pricing." *American Economic Review*, 103 (2013), 732–770.

Hou, K.; C. Xue; and L. Zhang. "Replicating Anomalies." *Review of Financial Studies*, 33 (2020), 2019–2133.

Jacobs, H. "Market Maturity and Mispricing." *Journal of Financial Economics*, 122 (2016), 270–287.

Jacobs, H., and S. Muller. "Anomalies Across the Globe: Once Public, No Longer Existent?" *Journal of Financial Economics*, 135 (2020), 213–230.

Karolyi, G. A., and S. Van Nieuwerburgh. "New Methods for the Cross-Section of Returns." *Review of Financial Studies*, 33 (2020), 1879–1890.

Koijen, R. S., R. J. Richmond, and M. Yogo. "Which Investors Matter for Equity Valuations and Expected Returns?" NBER No. w27402. (2022).

Kozak, S.; S. Nagel; and S. Santosh. "Interpreting Factor Models." *Journal of Finance*, 73 (2018), 1183–1223.

Kyle, A., and W. Xiong. "Contagion as a Wealth Effect of Financial Intermediaries." *Journal of Finance*, 56 (2001), 1401–1440.

Lewellen, J. "Institutional Investors and the Limits of Arbitrage." *Journal of Financial Economics*, 102 (2011), 62–82.

Liao, G. Y. "Credit Migration and Covered Interest Rate Parity." *Journal of Financial Economics*, 138 (2020), 504–525.

Lo, A. W., and J. Wang. "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory." *Review of Financial Studies*, 13 (2000), 257–300.

Lou, D., and C. Polk. "Comomentum: Inferring Arbitrage Activity from Return Correlations." *Review of Financial Studies*, 35 (2022), 3272–3302.

McLean, R. D., and J. Pontiff. "Does Academic Research Destroy Stock Return Predictability?" *Journal of Finance*, 71 (2016), 5–32.

Pontiff, J. "Costly Arbitrage: Evidence from Closed-End Funds." *Quarterly Journal of Economics*, 111 (1996), 1135–1151.

Pontiff, J. "Costly Arbitrage and the Myth of Idiosyncratic Risk." *Journal of Accounting and Economics*, 42 (2006), 35–52.

Puckett, A., and S. Yan. "The Interim Trading Skill of Institutional Investors." *Journal of Finance*, 66 (2011), 601–633.

Shleifer, A., and R. Vishny. "The Limits of Arbitrage." *Journal of Finance*, 52 (1997), 35–55.

Stambaugh, R.; J. Yu; and Y. Yuan. "The Short of It: Investor Sentiment and Anomalies." *Journal of Financial Economics*, 104 (2012), 288–302.