Wharton
UNIVERSITY of PENNSYLVANIA
JACOBS LEVY EQUITY MANAGEMENT CENTER
for Quantitative Financial Research

Accounting for the Anomaly Zoo: a Trading Cost Perspective

PRESENTER
Andrew Chen, Federal Reserve Board

CO-AUTHORED WITH
Mihail Velikov, Penn State University

So many anomalies, so many questions...

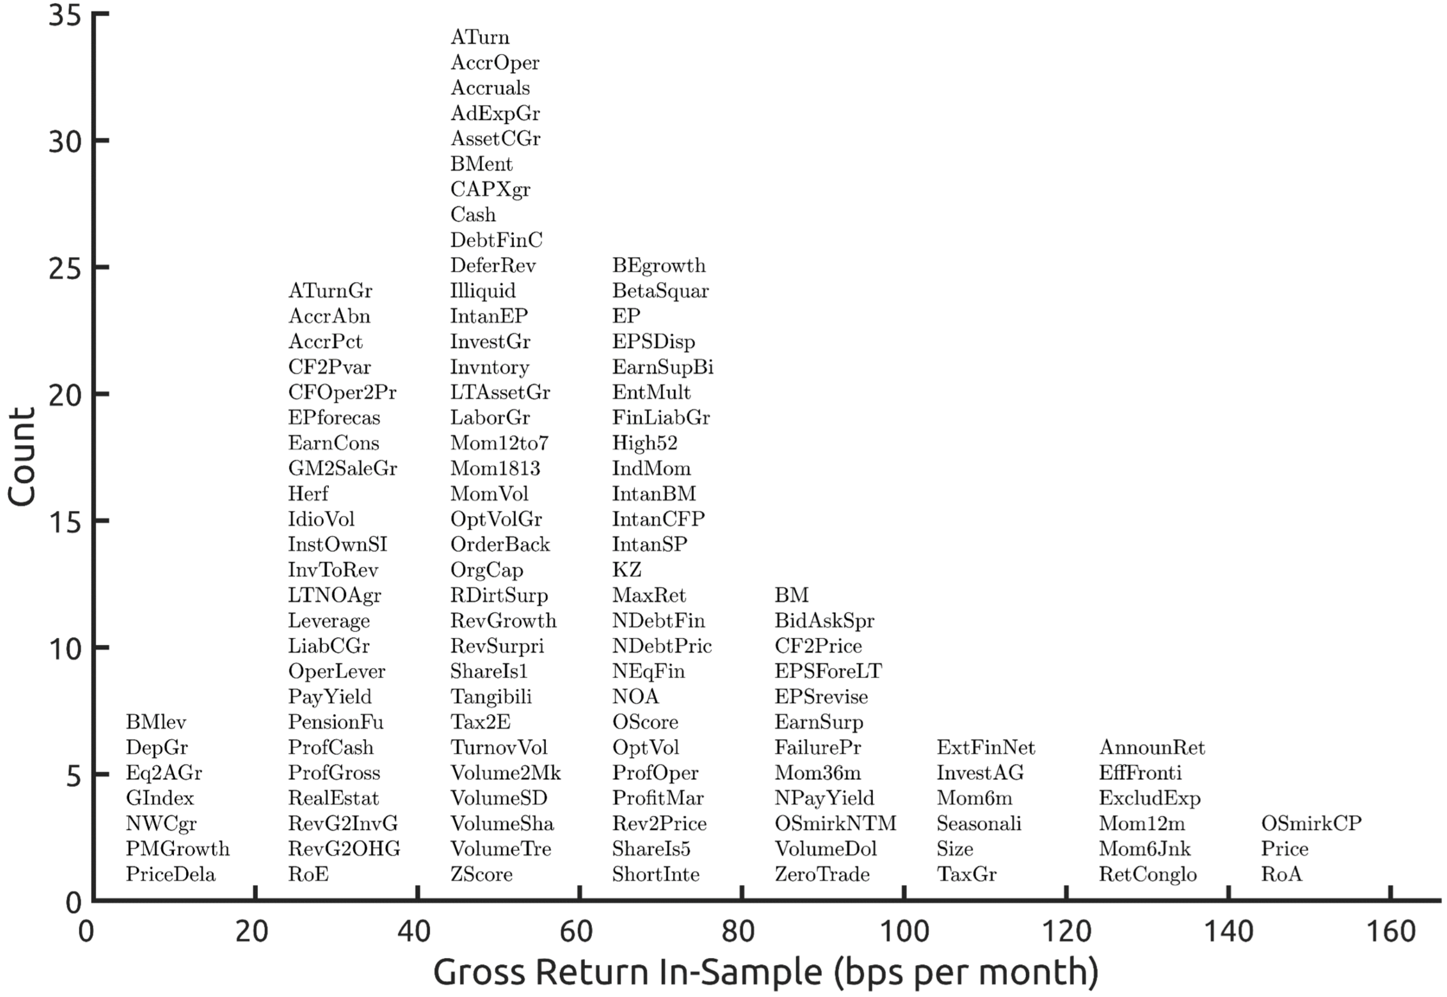

- What kind of factor model can explain this zoo? Can such models be rationalized?
- Which anomalies are redundant? Which have synergies?
- What share of these returns is due to data-mining?

We don't address any of these

Our question is more basic:

How much profit should investors expect (in the future) from investing in anomalies?

(We just want to know the expected return)

Existing literature does not answer the simple question:

The Standard Approach	The Problem
Average returns over decades of history	Data mining bias + investor learning => Can't expect historical returns to persist into the future (McLean and Pontiff 2016)
Measure gross returns (before trading costs)	Gross returns are not profits

This Paper:

We study post-publication returns net of costs for 120 anomalies

Costs = effective bid-ask spreads (TAQ/ISSM)

Post-publication net returns are tiny:

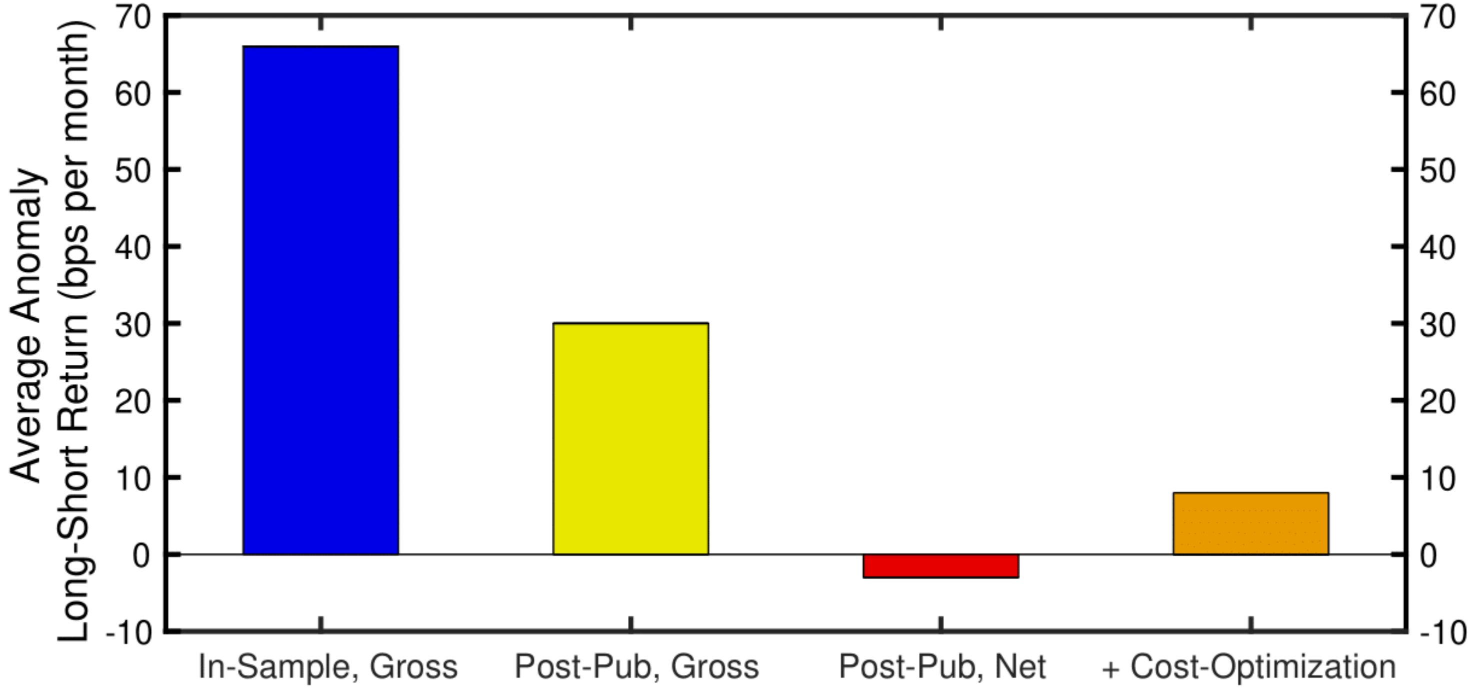

Average investor should expect tiny profits from the average anomaly

Related Literature

Many, many papers study trading costs of anomalies

- Stoll and Whaley (1983); Ball, Kothari, and Shanken (1995); Knez and Ready (1996); Pontiff and Schill (2001); Korajczyk and Sadka (2004); Lesmond, Schill, and Zhou (2004); Hanna and Ready (2005); Frazzini, Israel, Moskowitz (2015); Novy-Marx and Velikov (2016) ...

What's new: by far the most comprehensive set of anomalies (120)

- Allows for inferences regarding short post-publication samples
- Get us much closer to expected profits

Caveats

We do not attempt to study

- Implementation shortfall (Frazzini et al 2015; Briere et al 2019)
- Price impact (Frazzini et al 2015; Briere et al 2019)
- Combining multiple anomalies (DeMiguel et al Forthcoming)

Our goal is a simple benchmark expected return

Our benchmark: uses effective bid-ask spreads for single strategies

- lower bound cost for the average trader, irrespective of portfolio size
- starting point for studying more complex issues

Roadmap

1. Anomalies data and trading cost measures
2. Results
   a) Average published strategy
   b) Average cost-mitigated strategy
   c) Selected cost-mitigated strategies (adjusted for selection bias)

Anomalies data and trading cost measures

Anomalies Data

Begin w/ Chen and Zimmermann's (2018) 156 replicated characteristics

- Remove 34 that are not continuous
  - Need cost mitigation to understand costs, need continuity for cost mitigation
- Remove 2 that are somewhat hard to call anomalies
  - CAPM beta
  - Tail risk beta (Kelly and Jiang 2014)

Remaining: 120 published anomalies

- 50% focus on Compustat accounting variables
- 30% use purely price data
- 20% use analyst forecasts, institutional ownership, volume, etc

Short post-publication samples require a large number of anomalies

Trading Costs: Basics

Procedure:

1. Track portfolio weights over time
2. Whenever position is entered or exited: assume half the effective bid-ask spread is paid

Effective bid-ask spread:

$$[\text{Effective Spread}] = 2[\ \log[\text{Trade Price}] - \log[\text{Quote Midpoint}]\ ]$$

- For buys: trade price > midpoint (pay too much)
- For sells: trade price < midpoint (earn too little)

Interpretation: Lower bound cost to average trader

Lower bound cost

- Omits shorting costs and price impact
- Even the tiny net returns we find are unattainable to many traders

For average trader:

- Technically, a small liquidity demander
- Sophisticated arbitragers may supply liquidity (and bear other costs) (Frazzini, Israel, and Moskowitz 2018; Cont and Kukanov 2017)

Reminder: our goal is a simple benchmark expected return

Trading Costs: Data

Post-publication costs: high-frequency data

- 2003-2016: Daily TAQ (milli/nano-second timestamps)
- 1993-2003: Monthly TAQ (second timestamps)
- 1983-1992: ISSM
  - NASDAQ data starts in 1987

In-sample costs: average 4 low frequency proxies (1926-1982)

- Gibbs (Hasbrouck 2009)
- High-low spread (Corwin and Schultz 2012)
- Volume-over-volatility (Kyle and Obizhaeva 2016)
- Close-high-low (Abdi and Ranaldo 2017)

High-frequency data is important for post-publication samples

Low-Frequency Bias Over Time

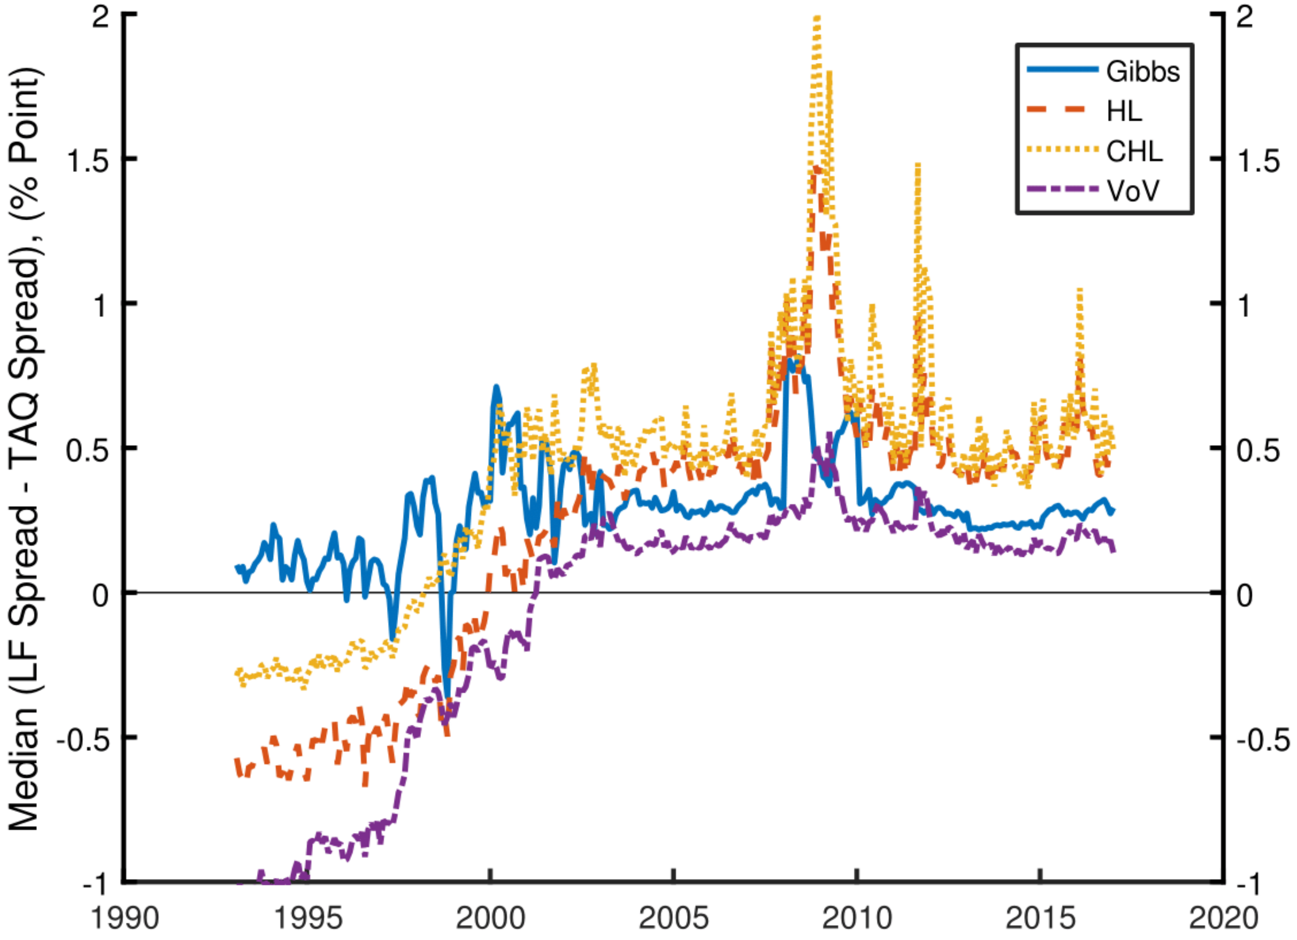

- Low-freq spreads are 25-50 bps upward bias in recent data

Our effective spread over time

- Huge spreads in 1930s-1940s
- Spreads rise in 1970s as NASDAQ enters CRSP
- Spreads plummet in 2000s with electronic trading

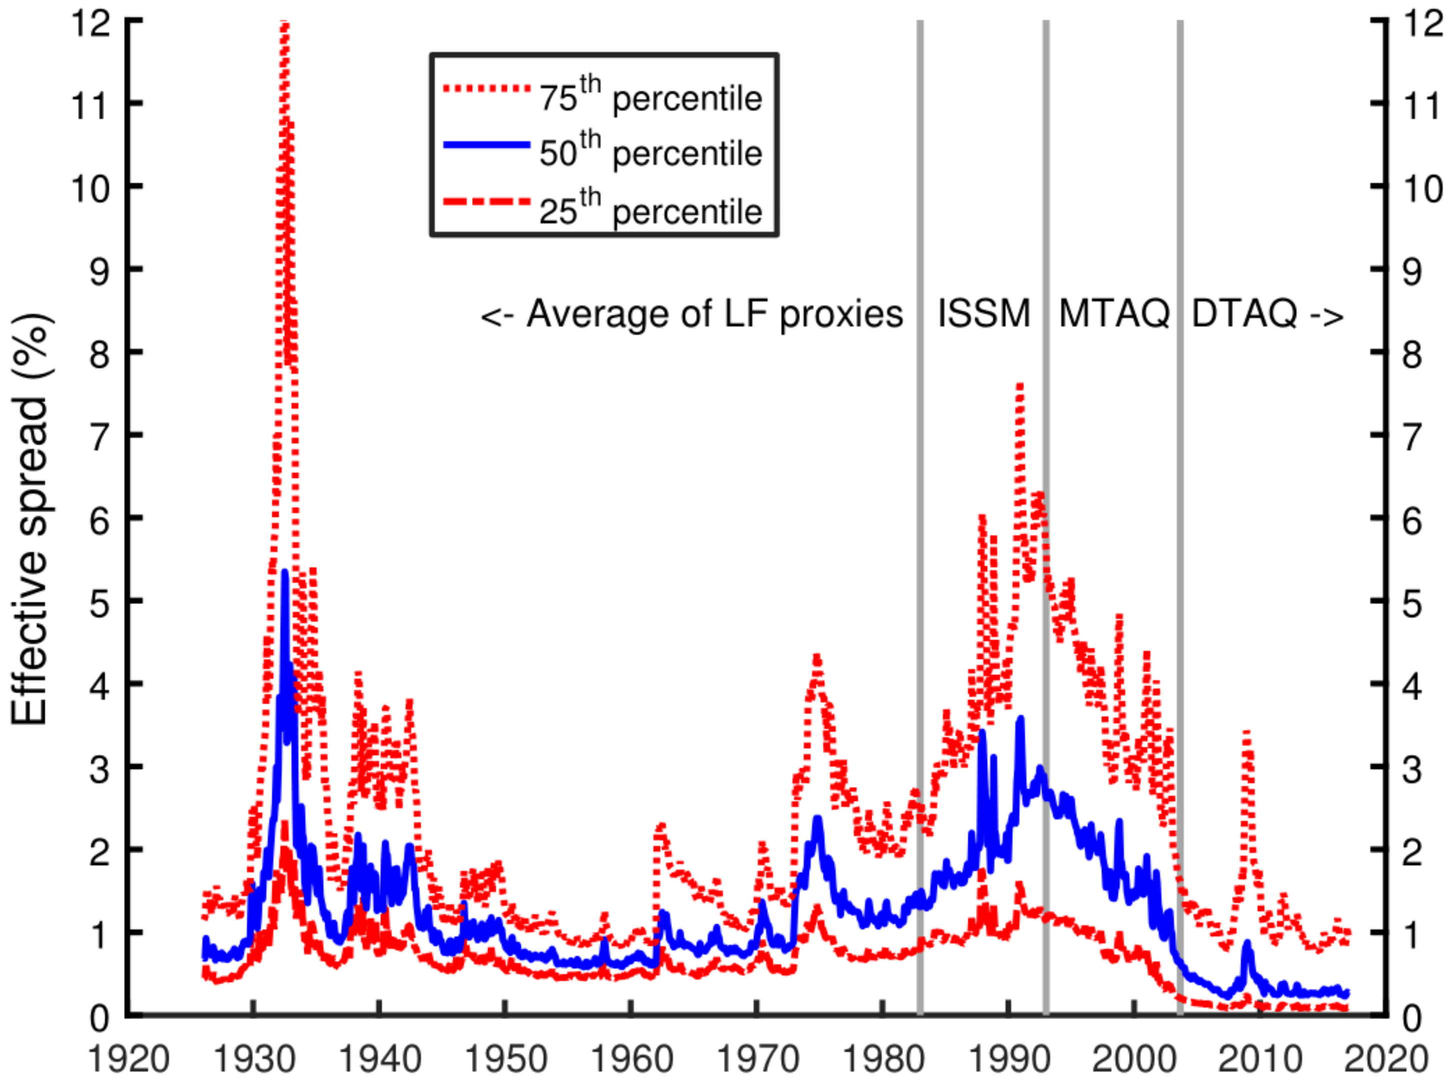

Is the average published strategy profitable?

Published Strategies

Almost all anomaly publications focus on equal-weighting

- (McLean and Pontiff 2016; Chen and Zimmermann 2018)

And use simple strategies:

- Long/short stocks in extreme quantiles
- Rebalance when signal updates

Same approach here: equal-weighted long-short quintiles
+ rebalancing when signal updates

- Quick, simple picture of net returns
- Next: cost-mitigated strategies

Result 1: Average investors should expect no profit from the average published strategy

- Standard errors are small
- Net returns are negligible even in-sample

	Average Across 120 Anomalies (%, Monthly)			
	Gross Return	Turnover (2-sided)	Ave Spread Paid	Net Return
In-Sample	0.66	31	2.19	0.05
	(0.04)	(4)	(0.06)	(0.06)
Post-Publication	0.30	30	1.11	**-0.03**
	(0.04)	(4)	(0.06)	(0.05)

- Decomposition

[Net Return] ≈ [Gross Return] − [Turnover] × [Spread]

= 30 bps − 0.30 × 111 bps = −3 bps per month.

Why are trading costs so large post-decimalization?

Decimalization: spread ≈ \$0.01, price ≈ \$20 ⇒ spread ≈ 5 bps.

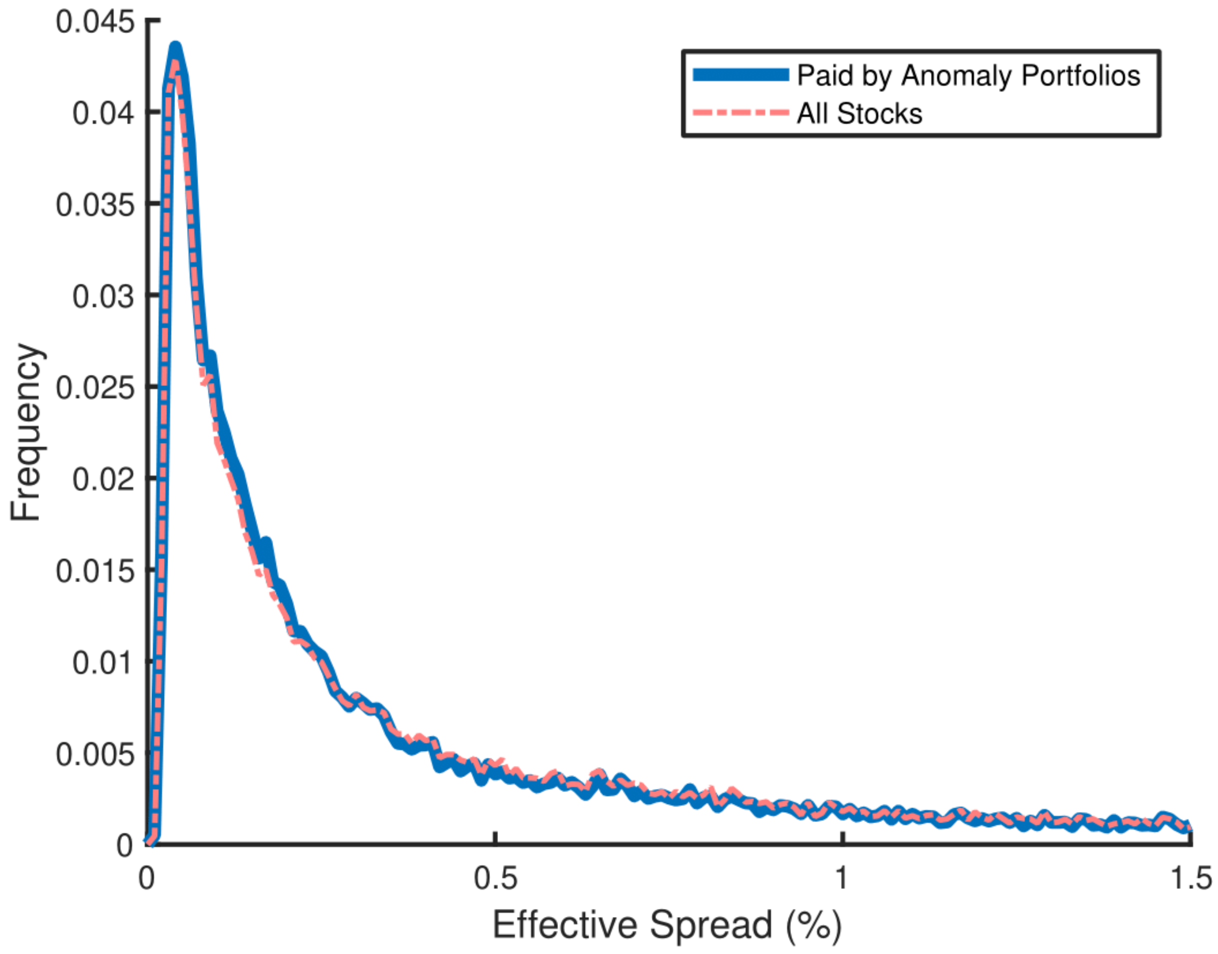

But 5 bps represents the **mode**

- Spreads have an extremely long right tail
- Mean spread = 67 bps
- Published strategies require trading across the entire distribution

Recap: is the average published strategy profitable?

No.

- 30% turnover × 111 bps spread wipes out profits

But these strategies completely ignore costs

Can smarter strategies earn profits?

Is the average cost-mitigated strategy profitable?

Cost Mitigation Overview

We combine two techniques

1. Value-weighting: reduces spreads paid
2. Buy/Hold Spreads: reduces turnover

These two together outperform several other cost mitigations

- (Novy-Marx and Velikov 2016, 2018)

Empirical Exercise

1. Optimize two techniques in-sample
2. Re-examine post-publication net returns

The Buy/Hold Spread: mimics optimal trading under trading costs

(Magill-Constantinides 1976; Brandt, Santa-Clara, Valkanov 2009)

Optimization Overview

Choose weighting and buy/hold spreads to maximize in-sample net returns

More formally:

$$\left\{ \{w_i^*\}_{i=1}^{N},\ \{\mathsf{bhs}^*(q)\}_{q=1}^{4} \right\} = \arg\max\ \text{In-Sample Net Return}\left(w_i, \mathsf{bhs}(q(\tau_i))\right)$$

where

$$\begin{aligned} w_i &\in \{\text{equal-weighted}, \text{value-weighted}\} \\ \mathsf{bhs}(q) &\equiv \text{buy/hold spread as a function of turnover quartile } q \end{aligned}$$

Specification aims to balance performance and robustness

Before cost-mitigation (in-sample)

Average net return

= 5 bps/month

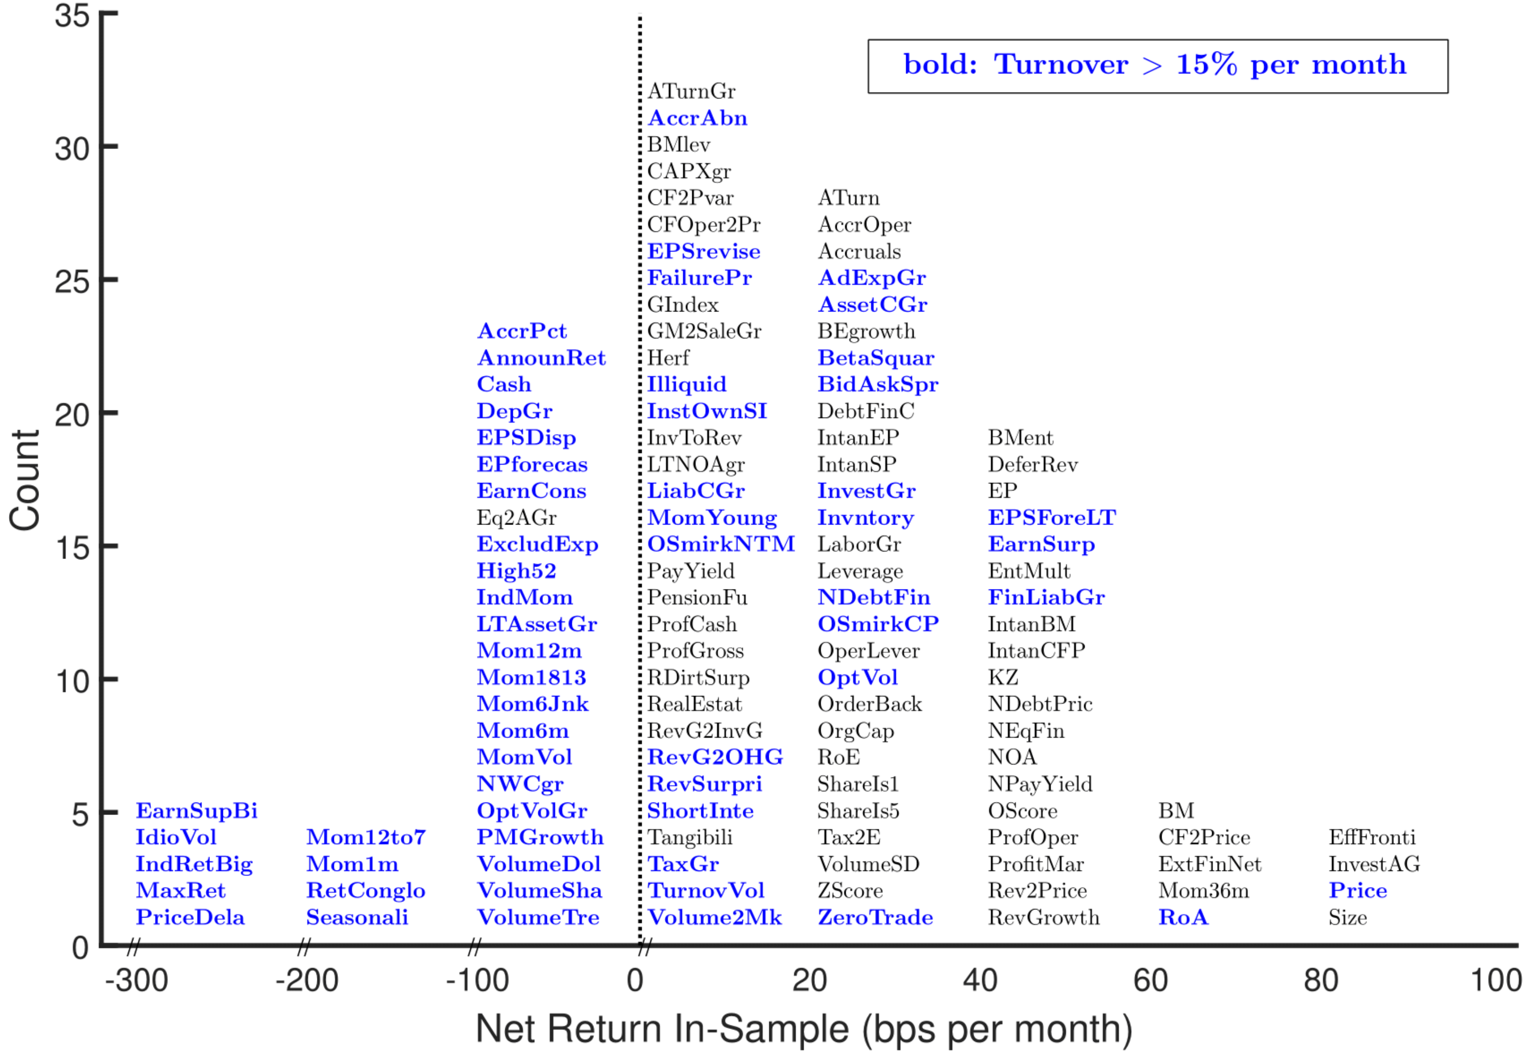

After cost-mitigation (in-sample)

Average net return

= 38 bps/month

Cost-mitigation works well (in-sample)

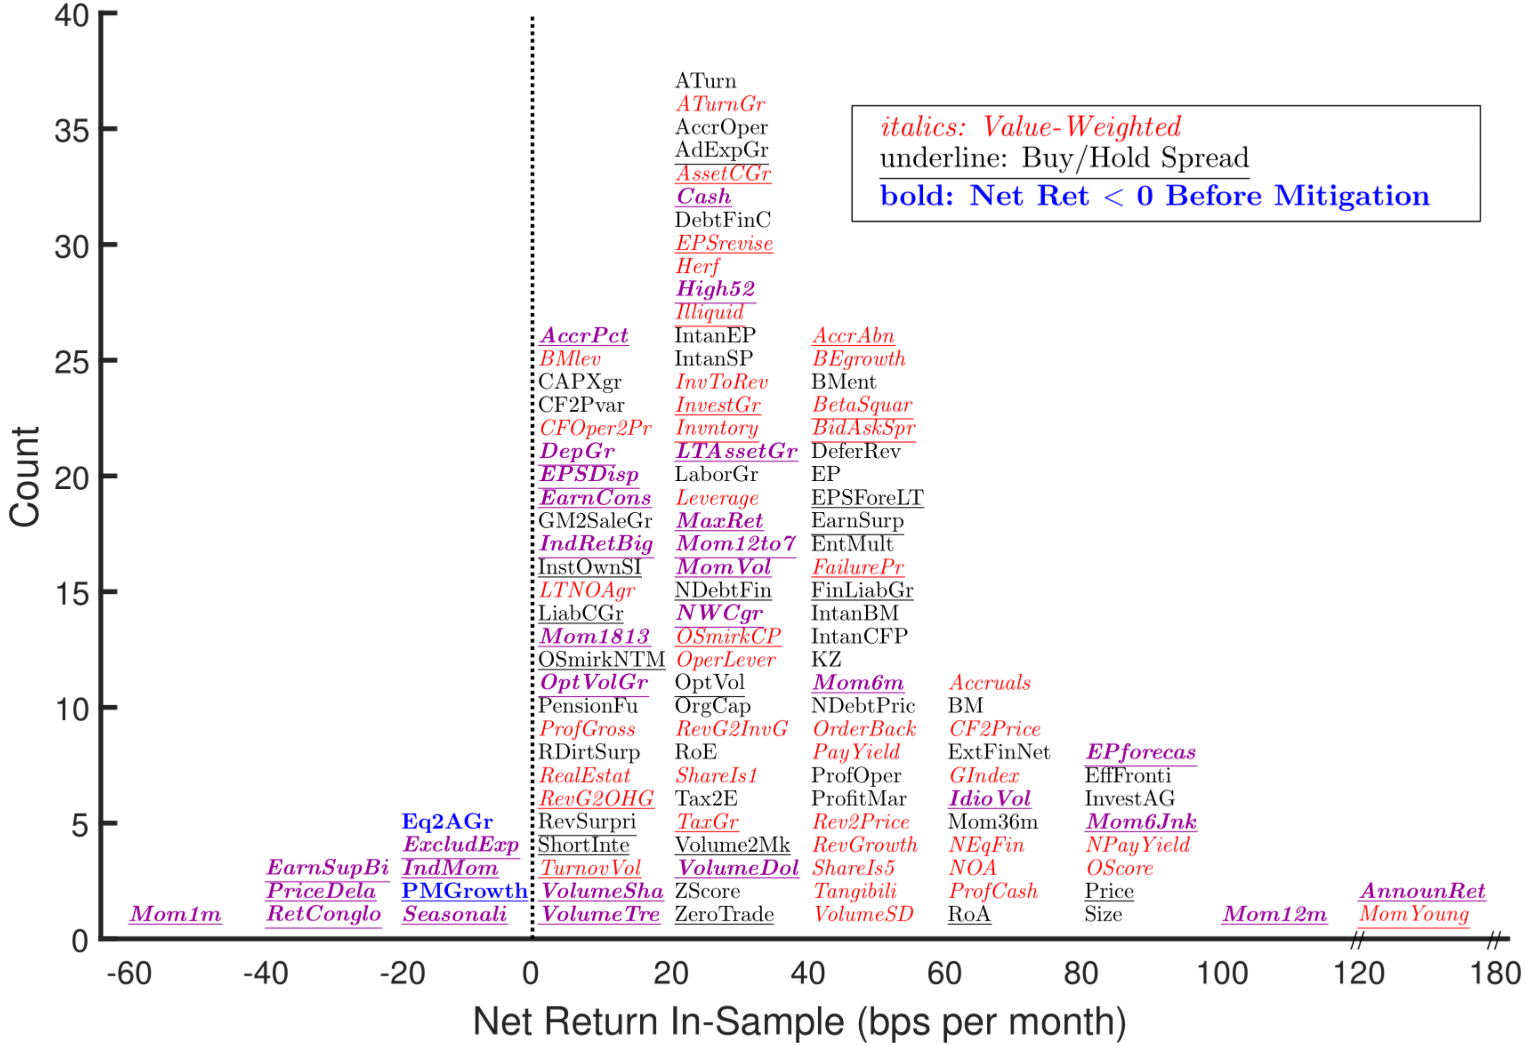

Result 2: Average investors should expect tiny profits from the average cost-mitigated strategy

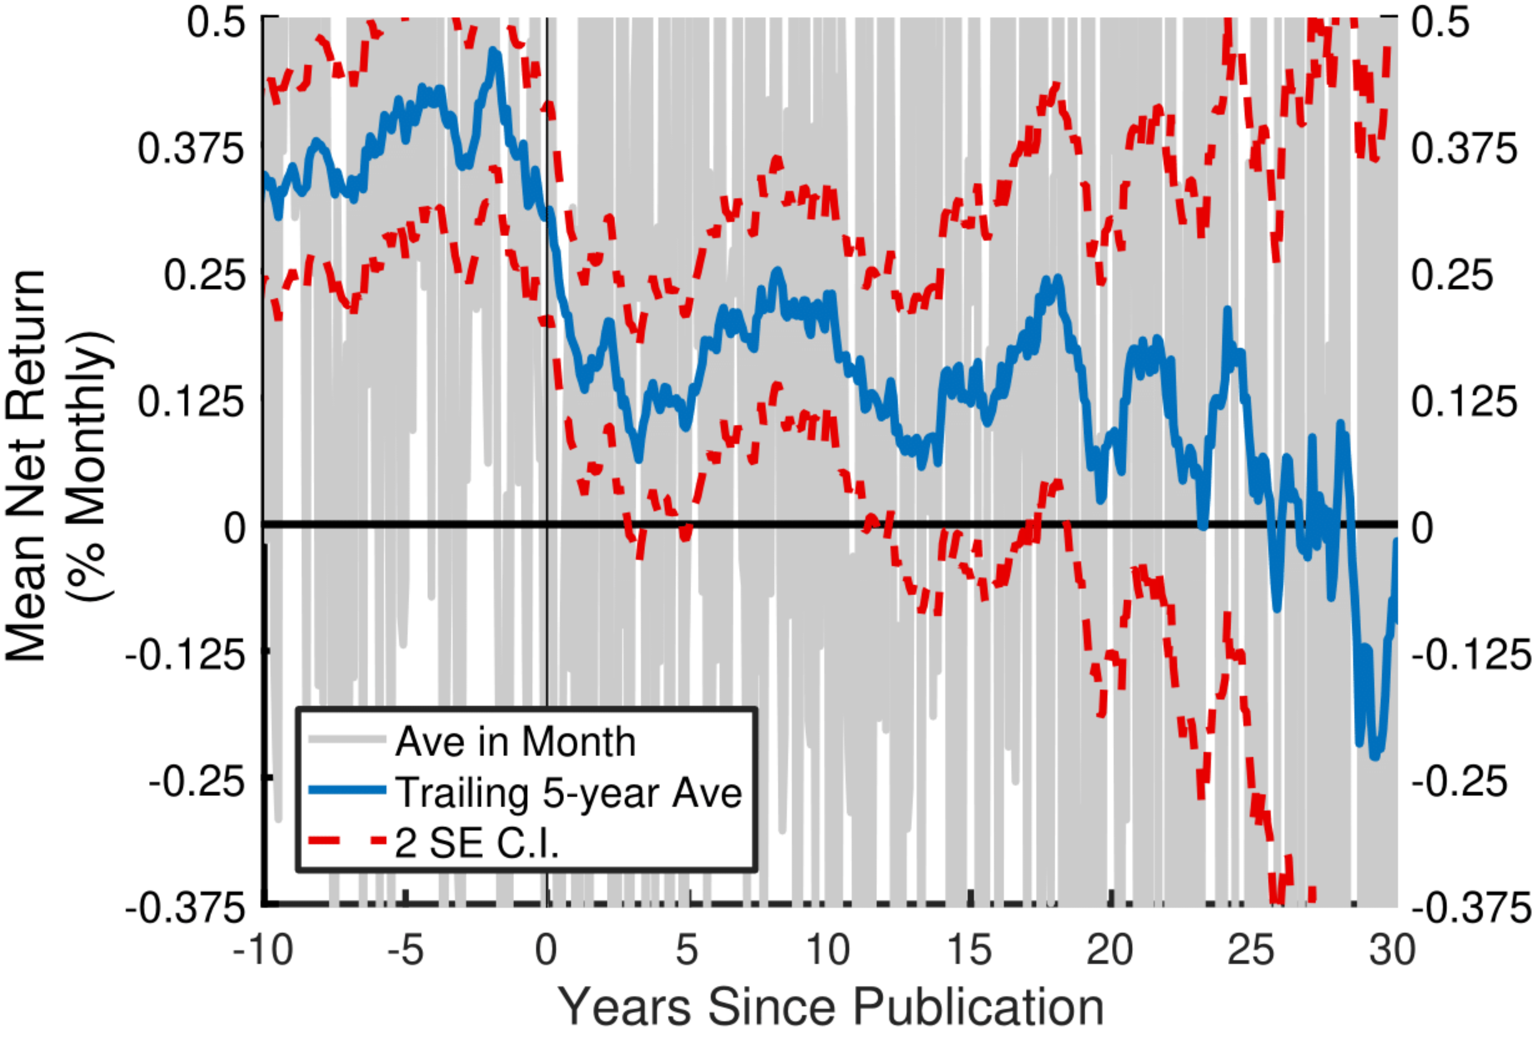

- Sizable in-sample net returns plummet around publication
- Average **4-13 bps/month** after publication, depending on how you take the average

Selected Cost-Mitigated Strategies

Size, B/M, and momentum are among the better performers

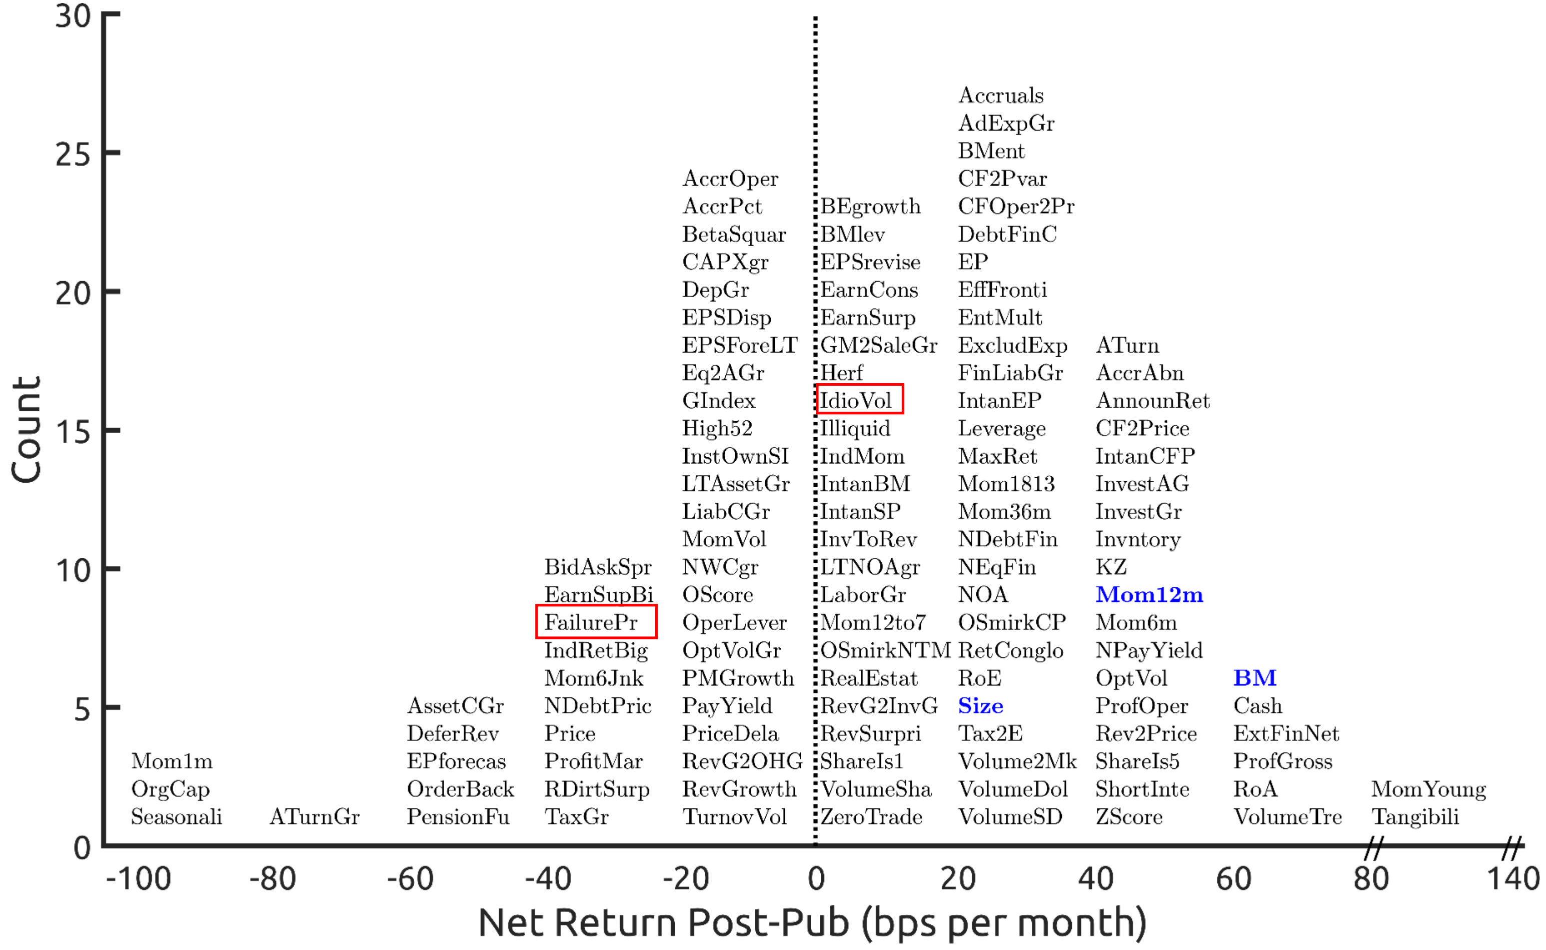

- Consistent with recent papers that measure implementation shortfall
  - Frazzini et al (2015)
  - Briere et al (2019)
- Are size, value, and momentum special?
- Or are they lucky?
- What about idiovol or distress (FailurePr)?

Final question: Can we expect selected strategies to be profitable?

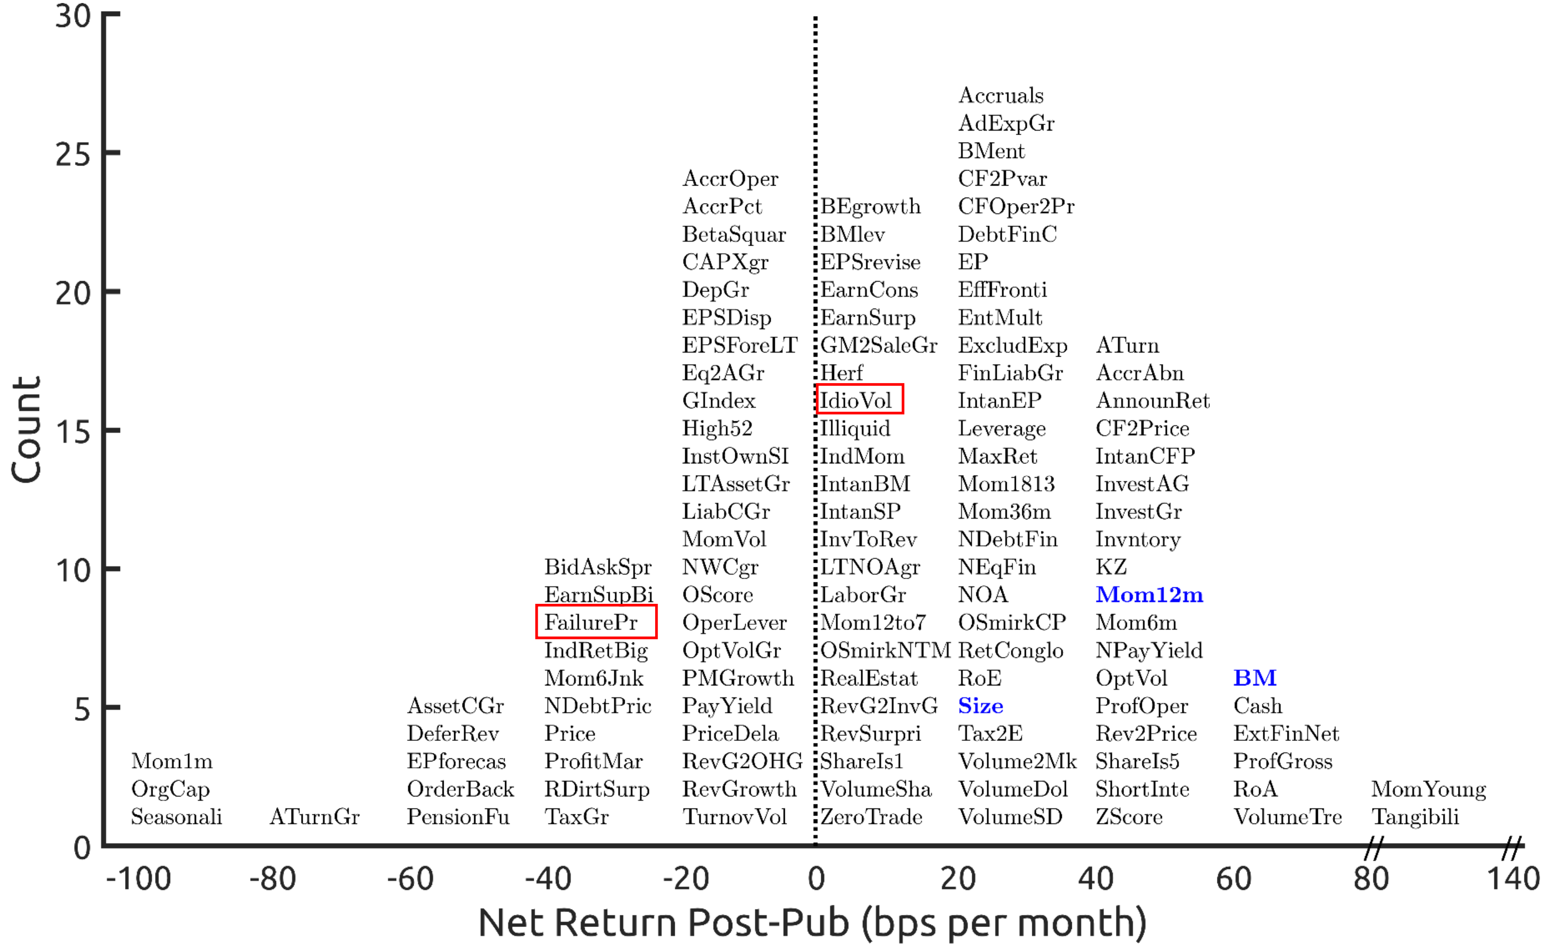

Tricky question: need to adjust for selection (hindsight) bias

We use two adjustments

- Forecast post-pub returns using in-sample information
- Empirical-Bayes adjustment

Bias adjustment 1: Forecasting post-pub net returns

Exercise:

1. Sort anomalies on in-sample turnover or net return
2. Examine mean post-publication net returns

	Post-Pub Net Returns (% monthly)			
In-Sample	Predictor Quartile			
Predictor	1 (Worst)	2	3	4 (Best)
Turnover	-0.18	-0.01	0.18	**0.21**
	(0.06)	(0.05)	(0.05)	(0.05)
Net Return	0.00	0.10	0.06	0.13
	(0.04)	(0.05)	(0.06)	(0.05)

Even the best predictors provide only ≈ 20 bps/month

- Excludes shorting costs, price impact
- Shorting costs average 10-20 bps (Cohen et al 2007)

Bias adjustment 2: Empirical Bayes adjustment

Uses empirical Bayes / "big-data" methods (Efron 2010; Azevedo et al 2019; Liu et al Forthcoming)

1. Model unobserved expected return μ_i

$$\bar{r}_i = N(\mu_i, \mathsf{SE}_i)$$
$$\mu_i \sim N(\mu_\mu, \sigma_\mu)$$

2. Estimate μ_μ σ_μ by method of moments

3. Bayes formula gives bias adjusted expected return

$$\mathbb{E}(\mu_i | r_i, \mathsf{SE}_i) = s_i \mu_\mu + (1 - s_i)\bar{r}_i$$
$$s_i \equiv \frac{\mathsf{SE}_i^2}{\sigma_\mu^2 + \mathsf{SE}_i^2}$$

Bias adjustment 2: Empirical Bayes adjustment

	Bias-Adjusted Net Return Percentiles (%, monthly)			
	50 pct	75 pct	90 pct	95 pct
Including EW	0.13	0.16	0.19	**0.21**
VW only	0.07	0.07	0.07	0.07

Once again:

- Even the best predictors provide only ≈ 20 bps/month
- Restricted to value-weighting => 7 bps

Result 3: average investors should expect only tiny profits from selected, cost-mitigated anomaly strategies.

Intuition: Why is selection bias so large?

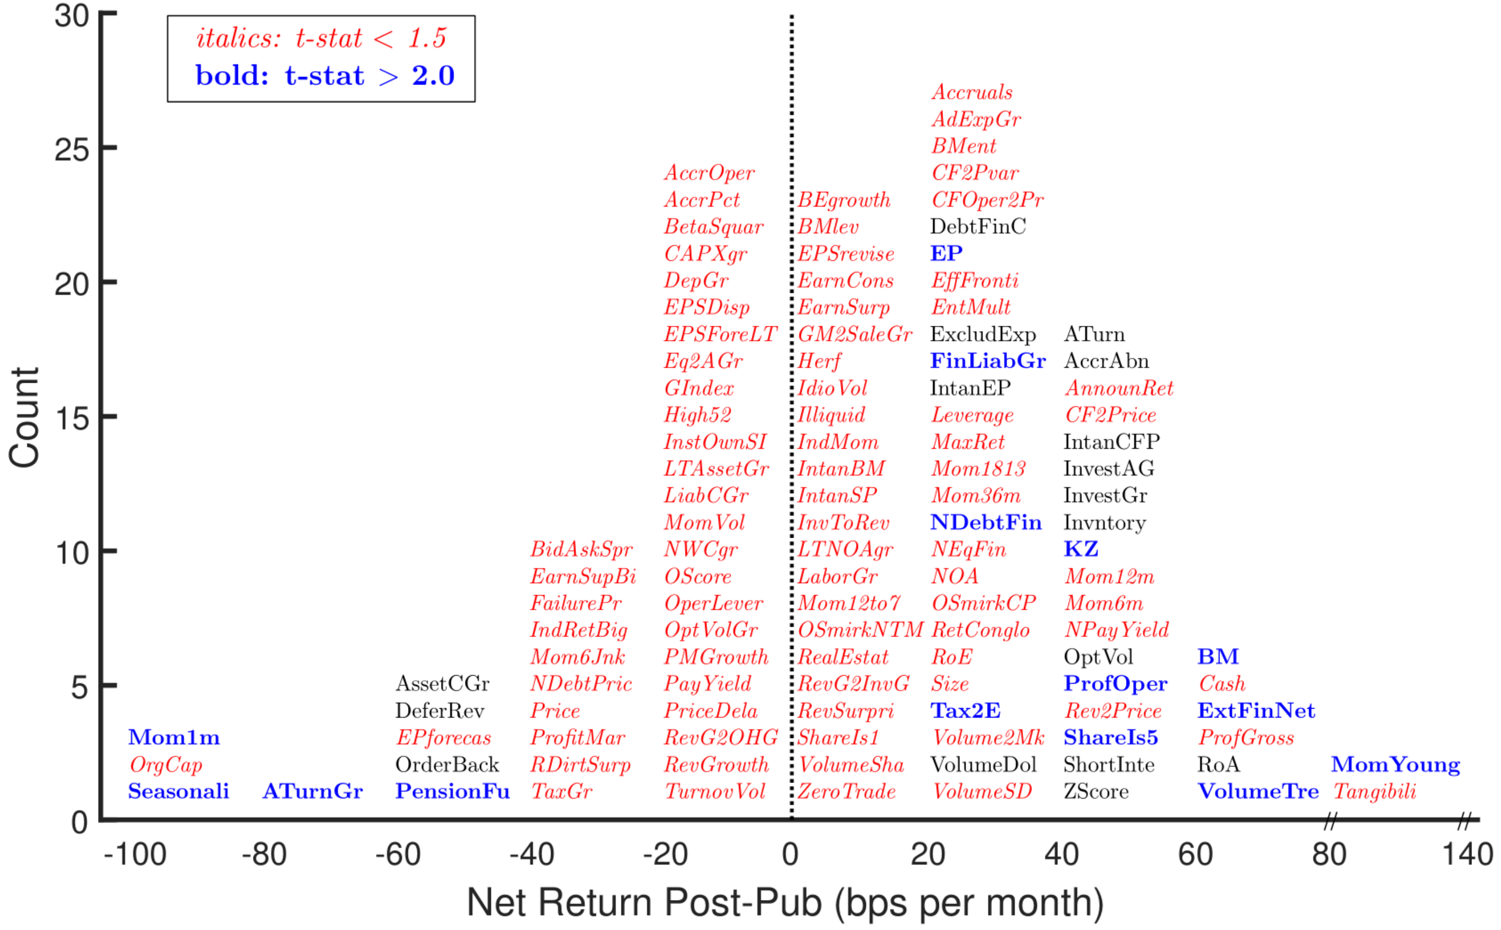

Distribution is close to the null of no predictability

- # |t-stats| > 2.0 = 13%
- No predictability => 5%

Most of the heterogeneity can be explained by noise / luck

Conclusion

We study post-publication returns net of costs for 120 anomalies

Post-publication net returns are tiny

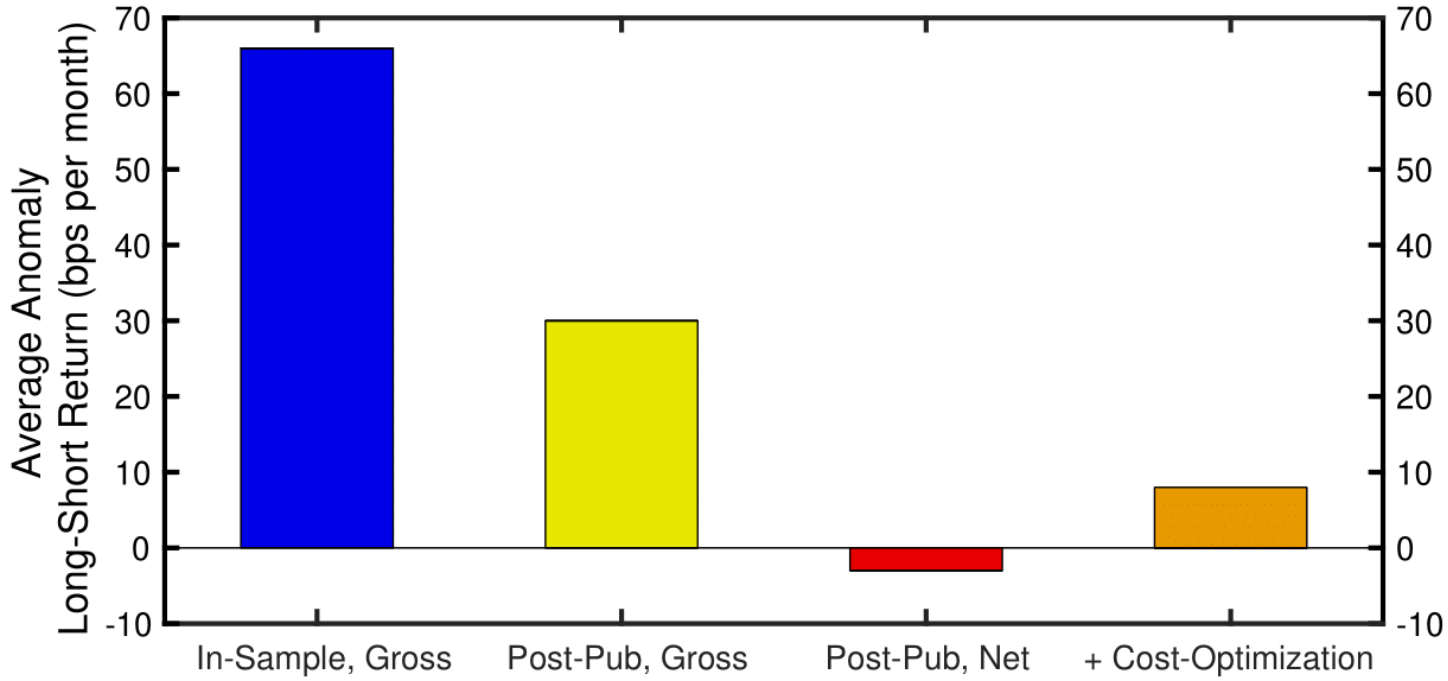

Average investor should expect tiny profits from average anomaly

Even the best anomalies provide only tiny net returns