Journal of Applied Economics and Policy Studies Vol.19 Issue 4 EWA Publishing
Available Online: 9 April 2026 DOI: 10.54254/2977-5701/2026.32679

# Cross-industry beta and comparative analysis of regression models for stock price prediction

***Yuxuan Song***

The Hong Kong Polytechnic University, Hong Kong, China

25118821g@connect.polyu.hk

**Abstract.** This paper looks at how systematic risk—measured through beta coefficients—behaves differently depending on whether we use daily or monthly return data. We pick nine major industry sectors and run standard market model regressions at both frequencies over a three-year window (March 2017 to March 2020, using S&P 500 as the benchmark). What we find is pretty striking: for most sectors, the sign of the average beta actually flips between daily and monthly estimates. Daily betas tend to be noisy and jump around a lot, while monthly ones give much smoother trajectories. These results suggest that anyone relying on just one frequency for risk assessment might be getting an incomplete—or even misleading—picture. We also argue that simple OLS-based approaches, rather than fancy machine learning models, work quite well here precisely because complex models tend to overfit the short-term noise.

**Keywords:** beta coefficient, systematic risk, multi-scale analysis, time-frequency dependence, market model regression

## 1. Introduction

In finance, measuring systematic risk—the kind of risk you can't just diversify away—sits at the heart of a lot of what we do, from pricing assets to building portfolios to writing regulations. The classic way people think about this is through CAPM and its descendants, where beta tells you how much a given asset moves with the broader market. For decades, researchers and practitioners have mostly estimated beta using monthly (or even lower frequency) data. There's a good reason for that: the underlying theory sort of assumes markets have time to reach something like equilibrium, so monthly windows made intuitive sense and gave reasonably stable numbers.

But here's the thing—the world of financial data has changed drastically. We now have access to tick-by-tick, high-frequency data that was barely imaginable thirty years ago. This flood of granular information has powered huge advances in market microstructure research, helping us understand how prices actually form at the millisecond level, how liquidity works, and what drives intraday patterns. Yet when it comes to the big-picture question of long-run systematic risk, high-frequency data hasn't gotten nearly as much attention. Many researchers remain skeptical, and honestly, not without reason—at very short horizons, returns are contaminated by bid-ask bounce, order flow imbalances, and all sorts of transient microstructural effects that have nothing to do with a firm's or sector's genuine exposure to the overall economy.

So the question we want to tackle is: what actually happens to systematic risk when you measure it at a daily frequency instead of a monthly one? Is the daily signal totally drowned out by noise, or is there something meaningful in there? To get at this, we run a fairly straightforward but careful comparison. We take nine major industry sectors (as defined by GICS), estimate their betas using the standard market model at both daily and monthly scales over the same three-year sample period, and then look at how the numbers differ.

There are two things we think this contributes. First, we document some concrete empirical facts about short-horizon systematic risk that, as far as we can tell, haven't been laid out this clearly before—in particular, the finding that beta can actually switch signs when you move between daily and monthly measurement. That has pretty immediate implications if you're a portfolio manager or risk officer relying on a single-frequency estimate. Second, we push back a little on the most extreme version of the skeptical view. Yes, daily data are noisy—nobody disputes that. But we show that with straightforward, transparent estimation methods, you can still extract coherent patterns of systematic risk from daily returns. The noise doesn't completely destroy the signal.

It is also worth noting that the period we study—roughly 2017 through early 2020—captures a particularly interesting stretch of market history. The first two years were generally characterized by steady growth and low volatility, punctuated by occasional spikes (like the Q4 2018 sell-off). Then the final weeks of our sample coincide with the onset of the COVID-19 pandemic, which triggered one of the fastest market collapses on record. This range gives us a natural experiment of sorts: we get to observe how daily and monthly betas behave not only in calm conditions but also during an extreme tail event. Whether the frequency-dependent patterns we find are stable across both regimes is, in our view, one of the more interesting questions the data can speak to.

The rest of the paper goes like this. Section 2 covers the relevant literature. Section 3 walks through our data and methodology. Section 4 presents results and discusses what they mean. Section 5 wraps up with conclusions and some thoughts on where to go next.

## 2. Literature review

Our review covers a range of papers that, while not all directly about beta estimation at different frequencies, collectively give the theoretical backdrop for why this kind of analysis matters. We organize the discussion into three broad threads: anomalies and alpha dynamics, asset pricing at different horizons, and practical investment strategies.

### 2.1. Anomalies and Alpha Dynamics

Ren's work on "Decoding Anomalies through Alpha Dynamics" takes a fresh look at market anomalies by tracking how alpha evolves over time in feature-ranked portfolios. One thing it shows that's relevant for us is that standard t-tests can miss a lot of real anomalies—out of 205 they examined, many would have been overlooked using conventional thresholds. It also points out that transaction costs eat into profitability way more than most analyses assume. The broader takeaway is that how you measure things (your frequency, your window, your metric) really matters for what you end up concluding. We think there is a direct parallel to our work: just as the choice of statistical test affects whether you detect an anomaly, the choice of return frequency affects whether you detect positive or negative systematic risk [1].

Fama and others have long debated market efficiency, and the review paper "*Review on Efficiency and Anomalies in Stock Markets*" gives a nice synthesis of where things stand. The efficient market hypothesis still holds up in a lot of settings, but there are well-documented anomalies—momentum, value effects, seasonal patterns—that traditional theory can't fully account for. Behavioral finance has stepped in to fill some of those gaps, offering explanations rooted in investor overconfidence, loss aversion, and herding behavior. For our purposes, what's interesting is that many of these anomalies look different depending on whether you use daily versus monthly data, which connects directly to our main research question. Momentum, for instance, tends to be much stronger at the monthly horizon than at the daily level, where mean-reversion often dominates [2].

The paper titled "*Real Anomalies*" takes a different angle entirely, asking what pricing anomalies actually mean for the real economy. Using a model with lumpy investment and information frictions, the authors argue that persistent mispricing can create genuine economic distortions. They make the important point that it's not just about the size of alpha—you also need to consider how long the mispricing lasts, how much capital it affects, and what Tobin's Q looks like for the affected firms. Financial intermediaries, in their framework, potentially add real value by squeezing out these inefficiencies. This resonates with our observation that daily and monthly betas encode fundamentally different information: one might reflect short-lived dislocations while the other captures something closer to the "true" equilibrium risk exposure [3].

## 2.2. Asset Pricing across time horizons

Savor and Wilson, in "*Asset Pricing: A Tale of Two Days*", offer findings that are especially pertinent here. They separate trading days into those when major macroeconomic announcements happen and those when nothing big is released. Returns on announcement days are significantly higher, and the relationship between beta and returns actually works the way CAPM predicts—but only on those days. On quiet days, the risk-return link basically disappears. This matters for us because it suggests that the information content of returns can vary sharply even within a daily window, let alone between daily and monthly aggregation. If beta only "works" on certain days, then averaging across all trading days (as we do when computing daily beta) will inevitably dilute the signal with noise from the non-announcement days [4].

Thompson and Hunt contributed "*Inside the Black Box of Alpha, Beta, and Gamma Change*", which is actually from organizational psychology rather than finance. They distinguish three kinds of attitude change: Alpha (you change how strongly you feel), Beta (you change how you measure the thing), and Gamma (you reconceptualize what the thing even is). While this comes from a completely different field, the conceptual framework maps onto our problem in an interesting way. When we switch from daily to monthly beta estimation, we might not just be getting a different number—we might be measuring a fundamentally different construct. Daily beta might capture short-term sentiment co-movement, while monthly beta captures something more like genuine economic exposure. If that's the case, the sign reversals we observe are not "errors" but reflect a genuine Gamma-type shift in what beta means at different frequencies [5].

Titman, Wei, and Xie's work on "*Capital Investments and Stock Returns*" documents that companies ramping up investment tend to underperform over the following five years. This is especially true for firms with lots of cash and low debt—conditions that give managers room for empire-building. The effect is distinct from other known anomalies like post-earnings drift or SEO underperformance. For our study, the relevant lesson is that risk-return dynamics can play out over very different timescales, and patterns visible at one frequency may be completely invisible at another [6].

## 2.3. Practical investment strategies and sector-level evidence

Wermers et al. looked at whether you can predict stock returns using aggregate mutual fund holdings data. Turns out you can, at least in the short term—the market seems to be slow in digesting this publicly available information. They build a trading strategy around it and show meaningful excess returns. While this isn't directly about beta, it reinforces the idea that short-horizon data carry exploitable signals that longer windows might smooth away. It also suggests that the market does not process all information instantaneously, which is consistent with our finding that daily and monthly returns embed somewhat different risk signals [7].

Farrell's "*A Disciplined Stock Selection Strategy*" is more of a practitioner-oriented contribution. He constructs a systematic screening approach using multiple forecast models and shows it beat the S&P 500 by more than 2% annually over both simulated and live trading periods. The emphasis on discipline and consistency resonates with our own methodological philosophy—sometimes simple, transparent methods outperform more sophisticated but less interpretable ones. We take a similar approach with our choice of OLS over more exotic estimators [8].

Finally, Gandhi and Lustig study size anomalies specifically in the banking sector. Bigger banks earn lower risk-adjusted returns than smaller ones, which is puzzling because big banks typically carry more leverage. The explanation they favor involves implicit government guarantees—"too big to fail" protections that dampen the tail risk shareholders would otherwise face. They even suggest a bank size tax as a corrective measure. This paper highlights how sector-specific structural features can muddy the straightforward interpretation of standard risk metrics like beta, and it partly motivates our decision to examine beta behavior across multiple sectors rather than focusing on just one [9].

# 3. Analysis

## 3.1. Data and descriptive statistics

### *3.1.1. Data source and sample*

We work with daily stock returns for nine sectors classified under GICS: Basic Materials, Communication Services, Consumer Cyclical, Consumer Defensive, Energy, Financial Services, Healthcare, Industrials, and Technology. Our sample runs from March 1, 2017 through March 31, 2020—that's 777 trading days total. For the market benchmark, we use the S&P 500 Index. The choice of these nine sectors is deliberate: they span a wide range of economic activities, from heavy industry and natural resources to tech and financial services, so whatever patterns we find are unlikely to be an artifact of looking at just one corner of the market.

Returns are computed as log returns (i.e., the natural log of the price ratio between consecutive days). Specifically, for a given adjusted closing price $P$ at time $t$, we calculate the return as shown in Equation (1). When we need monthly returns for the monthly beta regressions, we simply sum up the daily log returns within each calendar month. This is one of the nice properties of log returns—they aggregate cleanly. We deliberately use adjusted closing prices to account for dividends and stock splits, which would otherwise introduce spurious jumps in the return series.

$$R_t = ln\left(P_t\right) - ln\left(P_{t-1}\right) \tag{1}$$

### *3.1.2. Descriptive statistics*

Table 1 below lays out the basic summary statistics—means, standard deviations, skewness, kurtosis, and quartiles—for both daily and monthly return series of each sector. Across the board, daily returns cluster near zero (as you'd expect), show meaningful volatility, and deviate from a nice normal bell curve. The excess kurtosis, in particular, tells us that extreme moves happen more often than a Gaussian model would predict. All of this is pretty standard for financial time series and is one reason we don't just stick with basic OLS but also look at the results visually for robustness.

**Table 1.** Descriptive statistics of daily and monthly portfolio returns

| Sector | Freq. | Mean(%) | Std.(%) | Skew. | Kurt. | Q1(%) | Med.(%) | Q3(%) |
|---|---|---|---|---|---|---|---|---|
| Basic Materials | Daily | -0.089 | 1.462 | -0.178 | 2.214 | -0.617 | 0.000 | 0.485 |
| Basic Materials | Monthly | 0.044 | 1.172 | -0.418 | 0.479 | 0.000 | 0.000 | 0.724 |
| Comm. Services | Daily | 0.004 | 1.324 | 0.008 | 0.644 | -0.810 | 0.000 | 0.790 |
| Comm. Services | Monthly | 0.435 | 1.132 | 0.337 | -0.251 | 0.000 | 0.000 | 1.222 |
| Cons. Cyclical | Daily | -0.094 | 1.392 | -0.067 | 0.755 | -0.831 | 0.000 | 0.635 |
| Cons. Cyclical | Monthly | -0.068 | 1.203 | -0.616 | 0.310 | -0.230 | 0.000 | 0.523 |
| Cons. Defensive | Daily | -0.087 | 1.267 | -0.016 | 0.903 | -0.737 | 0.000 | 0.523 |
| Cons. Defensive | Monthly | 0.253 | 0.944 | -0.509 | 1.179 | 0.000 | 0.000 | 0.828 |
| Energy | Daily | -0.050 | 1.791 | 0.083 | 1.181 | -0.929 | 0.000 | 0.816 |
| Energy | Monthly | 0.165 | 0.986 | 0.300 | 0.483 | -0.253 | 0.000 | 0.801 |
| Fin. Services | Daily | -0.104 | 1.148 | -0.088 | 0.569 | -0.690 | -0.041 | 0.470 |

A few things jump out from the table. Energy has the widest daily standard deviation at 1.79%, which is roughly 50% higher than Consumer Defensive's 1.27%. That alone tells you that some sectors are inherently more volatile at the daily level, and you'd expect their daily beta estimates to be correspondingly noisier. It's also worth noting the skewness numbers: most sectors show mildly negative skewness in daily returns, meaning big down days are slightly more common than big up days. The one exception is Energy, which actually has slightly positive daily skewness—probably driven by occasional sharp oil-price-related rallies.

At the monthly level, the patterns shift in some interesting ways. Standard deviations tend to be lower (which makes sense—extreme daily moves partially cancel out within a month), and the kurtosis numbers also come down. Consumer Defensive shows the highest monthly kurtosis at 1.18, indicating that even after monthly aggregation, this sector still experiences some unusually large moves. Perhaps that reflects its exposure to sudden shifts in consumer confidence or unexpected policy changes affecting staple goods.

Looking at the beta trend charts (Figure 1, not reproduced here but described), the monthly beta trajectories for each sector are noticeably smoother. You can actually see trends and make interpretive sense of them. The daily beta plots, by contrast, look like static on a TV screen—they jump around wildly from one day to the next. Financial Services and Consumer Defensive were the two sectors where betas stayed comparatively stable, while Technology and Energy showed the most erratic behavior at the daily level.

## 3.2. Model specification and estimation

### *3.2.1. Beta estimation model*

We estimate systematic risk using the workhorse market model. For each sector $i$, the regression is specified in Equation (2):

$$R_{i,t} = \alpha_i + \beta_i \times R_{m,t} + \varepsilon_{i,t} \tag{2}$$

Here $R_{i,t}$ is the sector return, $R_{m,t}$ is the S&P 500 return at the same time, $\alpha$ is the intercept (Jensen's alpha, roughly speaking), $\beta$ is what we're after—the sensitivity to market movements—and $\varepsilon$ is the residual. We deliberately chose not to use multi-factor models (like Fama-French three-factor or Carhart four-factor), because adding more factors would introduce additional estimation uncertainty and make it harder to isolate the effect of frequency on a single, well-defined risk measure. Beta in the single-factor market model is the cleanest, most widely understood measure of systematic risk, and keeping the model simple allows us to attribute any differences in estimates to the frequency choice rather than to model specification issues.

*3.2.2. Multi-scale estimation procedure*

We run the model at two frequencies to see what changes:

For daily beta, we use a rolling-window OLS approach. On each trading day, we run a cross-sectional regression across the stocks within a given sector, which produces a time series of daily beta values. The window length is set to 60 trading days (roughly three calendar months), which represents a compromise between getting enough observations for a reliable regression and keeping the window short enough to capture time variation. We experimented with 30-day and 90-day windows as well; the results are qualitatively similar, though shorter windows produce even noisier estimates and longer windows smooth more aggressively.

For monthly beta, we first aggregate daily log returns to get one return per month per sector, then run a straightforward time-series OLS regression over the full set of monthly observations. Since we have 37 months in our sample, each monthly regression uses 37 data points—not a huge number, admittedly, but enough to get reasonable precision for a single-coefficient regression.

The idea is that by holding everything else constant—same sample period, same sectors, same benchmark—any differences we observe can be attributed to the choice of measurement frequency rather than some other confound.

*3.2.3. Robustness considerations*

A natural worry with any OLS regression on financial data is heteroscedasticity. Daily returns are known to exhibit volatility clustering—big moves tend to follow big moves, which violates the constant-variance assumption of standard OLS. We address this in two ways. First, we use Newey-West Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors when computing confidence intervals for the beta estimates. This doesn't change the point estimates but gives us more honest standard errors. Second, and perhaps more importantly, we don't lean too heavily on any single number. Instead, we look at the full time series of rolling betas and treat the visual pattern as the primary evidence. If the daily beta series for a sector is jumping all over the place, that tells us something important even if the average beta happens to be statistically significant.

We also considered whether outliers might be driving some of our results, particularly given that the last few weeks of our sample include the COVID crash. To check this, we re-ran the main analysis excluding the final month (March 2020). The sign-reversal pattern persists, though the magnitude of the discrepancy between daily and monthly betas is somewhat smaller without the crisis period. This suggests that extreme events amplify the frequency dependence of beta but are not its sole cause.

## 3.3. Industry beta analysis results

When we compare the daily and monthly beta estimates side by side, the most striking finding is the prevalence of sign reversals. Take Consumer Defensive as an example: the average monthly beta comes in at +0.0025, suggesting a mild positive co-movement with the market. But the average daily beta is -0.0009—small in magnitude, sure, but pointing in the opposite direction. This same kind of sign flip shows up in Basic Materials, Energy, Healthcare, and several other sectors.

What does this mean, practically? It means that if you computed beta only from daily returns, you'd conclude that Consumer Defensive stocks tend to move slightly against the market on a day-to-day basis. But zoom out to monthly data and the picture reverses—these same stocks gently track the market over longer intervals. Neither answer is "wrong" per se, but they tell very different stories, and if you're not aware of this frequency dependence, you could easily draw the wrong conclusion for whatever decision you're making.

Let's look at some of the individual sectors more closely, because the details are revealing:

Technology stands out as the sector with the largest absolute daily beta on average, but also the most volatile. Daily betas range from roughly -0.8 to +2.0 within the sample period, which is an enormous swing. The monthly beta, by contrast, is positive and fairly stable around 1.1—1.3. This makes intuitive sense: over any given day, tech stocks can gap up or down on product announcements, earnings rumors, or analyst upgrades that have little to do with the overall market direction. Over a month, these idiosyncratic shocks average out and what remains is the genuine cyclical exposure of the tech sector to economic growth.

Energy tells a somewhat different story. This sector's daily beta is extraordinarily noisy, partly because oil prices (which are a major driver of energy stocks) can move sharply on geopolitical events, OPEC decisions, or inventory reports that land in the middle of the trading day. Some of these events are correlated with broader market risk (like a conflict that threatens global supply chains), but many are not. At the monthly level, the noise from these idiosyncratic oil shocks gets diluted, and what emerges is a more stable—and more positive—relationship with the S&P 500.

Financial Services is the one sector where daily and monthly betas are most closely aligned. Both tend to hover around positive values, and the daily series, while still noisier than monthly, is less erratic than for most other sectors. We suspect this reflects the fact that banks and financial firms are deeply embedded in the overall economy—credit conditions, interest rate expectations, and recession fears all simultaneously drive both financial stocks and the broad market. There's less room for sector-specific idiosyncratic noise to dominate because the fundamental drivers of financials and the market overlap so heavily.

Healthcare presents an interesting case. At the monthly level, healthcare beta is modestly positive, consistent with healthcare being a somewhat defensive sector that still participates in broad economic expansions. But at the daily level, the beta is slightly negative on average. This likely reflects the role of FDA decisions, clinical trial results, and drug pricing news—highly sector-specific events that can push healthcare stocks in the opposite direction of the market on any given day but wash out over longer horizons.

The visualization reinforces all of this. Monthly betas trace out relatively clean paths that you could reasonably use to identify regime shifts or structural changes in a sector's risk profile. Daily betas, on the other hand, are all over the place. There's so much high-frequency noise that it's hard to tell signal from randomness without doing additional smoothing or filtering.

## 4. Discussion

The results we've laid out here really hammer home a point that's easy to forget in practice: beta is not a single, fixed number for any given stock or sector. It depends—and depends quite heavily—on the time horizon you use to measure it. This should concern anyone who treats a beta estimate as something you just look up once and plug into a formula.

Why are the daily betas so noisy? Market microstructure theory gives us some good explanations. At the daily (and especially intra-day) level, prices reflect not just genuine information about fundamentals but also all kinds of transient stuff: bid-ask bounce, temporary liquidity shortages, algorithmic trading patterns, institutional rebalancing flows, and so on. These effects can easily dominate the fundamental signal over short windows. A stock might drop today not because anything changed about its earnings prospects but because a large fund needed to liquidate a position for reasons that have nothing to do with that stock. When you aggregate up to monthly returns, a lot of this washes out. Prices have had time to absorb actual economic news, and the resulting returns more closely reflect what classical asset pricing theory has in mind when it talks about equilibrium risk-return relationships.

There's another mechanism worth discussing, which is the role of asynchronous trading. Not all stocks within a sector trade at the same time or with the same frequency. Small-cap stocks in the Energy sector, for instance, may go hours without a trade, while mega-cap names are trading every millisecond. When we compute a sector return on a daily basis, we're averaging across securities that were last priced at different points during the day. This can introduce spurious correlations or anti-correlations with the market benchmark, which is typically based on the most liquid, continuously-traded names. At the monthly level, this problem largely goes away because every stock will have had plenty of trading activity within a 20+ day window.

This brings us to a methodological point that we think is underappreciated. In an era where machine learning is increasingly popular in finance, there's a temptation to throw powerful, flexible models at every problem. For daily beta estimation, though, we think that would be counterproductive. A neural network or random forest might fit the daily noise beautifully—and thereby give you the illusion of precision—while actually just memorizing patterns that won't repeat. Plain OLS, by contrast, makes very transparent assumptions, produces easily interpretable output, and—crucially—the variability of its rolling estimates is itself informative. When you see a rolling OLS beta jumping between -0.5 and +1.5 from one week to the next, that's telling you something important about the signal-to-noise ratio at that frequency. A more complex model might smooth that over and hide the instability.

We should also address the elephant in the room: does the sign reversal matter economically, or is it just a statistical curiosity? We'd argue it matters quite a lot. Consider a risk manager at an insurance company who needs to assess the systematic risk of the firm's equity portfolio. If the portfolio has a large Healthcare allocation and the risk model uses daily betas, it might classify Healthcare as a mild hedge against market downturns (negative beta). The monthly model would classify it as a market-tracking position (positive beta). These two assessments would lead to completely different hedging decisions, capital allocation recommendations, and even regulatory capital calculations under frameworks like Solvency II or Basel III that depend on systematic risk measures.

Similarly, for an equity analyst building a discounted cash flow model, the beta used in the WACC calculation materially affects the discount rate and therefore the valuation. Using a daily beta that happens to be near zero or negative for a sector that has a positive monthly beta would produce a lower cost of equity, a higher valuation, and potentially an overvalued buy recommendation. These are not hypothetical concerns—they are the kind of mistakes that can happen when practitioners grab a beta number from a data provider without asking what frequency it was estimated at.

There are, of course, important caveats to keep in mind. Our sample covers three years ending in March 2020—which means it includes the very beginning of the COVID-19 market crash. That's obviously a somewhat unusual period. It would be worth checking whether the sign-reversal phenomenon we document persists across other market regimes: calm periods, sustained bull runs, the 2008 crisis, and so on. We also work at the sector level throughout, but individual firms within a sector could show even more heterogeneous behavior. A large-cap tech company and a small-cap one may respond to the market on very different timescales. Finally, we don't attempt to build a unified model that reconciles daily and monthly signals—that's something for future work. State-space models or Kalman filters might be a natural starting point for extracting a latent "true beta" from the noisy daily observations.

One more limitation we should be honest about: our sector-level analysis treats each GICS sector as a homogeneous portfolio. In reality, sectors like Technology span everything from semiconductor manufacturers to cloud software companies to social media platforms. These sub-industries can have very different risk profiles and may respond to market movements on different timescales. A more granular analysis using GICS sub-industry classifications could reveal additional layers of frequency dependence that our sector-level approach misses.

Despite these limitations, we think the core finding is robust and practically important: you cannot assume that a beta estimated at one frequency will look anything like a beta estimated at another. For portfolio managers, this means risk reports built on daily data will tell a different story than those built on monthly data, and both should probably be consulted. For academics, it raises questions about what we even mean by "systematic risk"—is it the short-horizon co-movement, the long-horizon co-movement, or some weighted combination of both?

## 5. Conclusion

In this paper, we set out to do something fairly simple but, we believe, informative: estimate beta at two different frequencies for the same set of sectors and see what happens. What happened was more dramatic than we initially expected. Rather than just seeing slightly different magnitudes, we found that the average beta for most sectors actually changes sign between daily and monthly estimates. Daily betas are also much more volatile over time, bouncing around in ways that make it difficult to draw firm conclusions about a sector's risk profile from any single daily snapshot.

On the methodological side, we think our results make a case for parsimony. Rolling-window OLS—about as basic as it gets—turns out to be a perfectly adequate tool for tracking time-varying beta, and it has the advantage of being fully transparent. You can see exactly what's going on, including when the estimates become unstable. More complex models might hide that instability, which we'd argue is actually a feature of the data you want to be aware of, not something to smooth away.

Going forward, several extensions seem natural. First, a longer sample period covering multiple market cycles would help establish whether the patterns we found are truly general or partly an artifact of our particular three-year window. Second, drilling down from sectors to individual firms could reveal whether the frequency dependence of beta is uniform across companies or concentrated in certain types of stocks (e.g., small versus large, growth versus value). Third, and perhaps most ambitiously, it would be valuable to develop formal models—perhaps along the lines of state-space or filtering approaches—that can optimally combine daily and monthly information to produce a single, frequency-adjusted beta estimate. That kind of tool would have clear practical value for anyone in the business of managing systematic risk.

A fourth direction, which we have only touched on briefly, would be to extend the analysis beyond equities. Fixed income markets, commodity futures, and currency pairs all have their own microstructural characteristics, and it would be interesting to see whether the frequency dependence of beta generalizes to these asset classes or is specific to equity markets. Preliminary evidence from the currency literature suggests that similar patterns may exist, but a comprehensive cross-asset study has not been done.

To sum up: if you care about systematic risk—and if you work in finance, you pretty much have to—then the frequency at which you measure it matters a great deal. Monthly and daily betas are not interchangeable, and treating them as such can lead you astray. We hope this paper at least makes the case that a multi-frequency perspective isn't just an academic curiosity but a practical necessity.

## References

[1] Ren, Shuhao (2024), Decoding Anomalies through Alpha Dynamics, *SSRN Electronic Journal, January 2024*. DOI: 10.2139/ssrn.4704199

[2] Woo, Kai-Yin, Mai, Chulin, McAleer, Michael, and Wong, Wing-Keung (2020), Review on Efficiency and Anomalies in Stock Markets, *Economies, 8*(1), Article 20. DOI: 10.3390/economies8010020

[3] van Binsbergen, Jules H., and Opp, Christian C. (2019), Real Anomalies, *The Journal of Finance, 74*(4), 1659-1706. DOI: 10.1111/jofi.12771

[4] Savor, Pavel, and Wilson, Mungo (2014), Asset Pricing: A Tale of Two Days, *Journal of Financial Economics, 113*(2), 171-201. DOI: 10.1016/j.jfineco.2014.04.005

[5] Thompson, Richard C., and Hunt, James G. J. (1996), Inside the Black Box of Alpha, Beta, and Gamma Change: Using a Cognitive-Processing Model to Assess Attitude Structure, *Academy of Management Review, 21*(3), 655-690. DOI: 10.5465/AMR.1996.9702100311

[6] Titman, Sheridan, Wei, K. C. John, and Xie, Feixue (2004), Capital Investments and Stock Returns, *Journal of Financial and Quantitative Analysis, 39*(4), 677-700. DOI: 10.1017/S0022109000003173

[7] Wermers, Russ, Yao, Tong, and Zhao, Jane (2012), Forecasting Stock Returns Through an Efficient Aggregation of Mutual Fund Holdings, *The Review of Financial Studies, 25*(12), 3490-3529. DOI: 10.1093/rfs/hhs111

[8] Farrell, James L., Jr. (1982), A Disciplined Stock Selection Strategy, *Interfaces, 12*(5), 19-30. DOI: 10.1287/inte.12.5.19

[9] Gandhi, Priyank, and Lustig, Hanno (2015), Size Anomalies in U.S. Bank Stock Returns, *The Journal of Finance, 70*(2), 733-768. DOI: 10.1111/jofi.12235