Journal of Financial Economics

Index rebalancing and stock market composition: Do indexes time the market?☆

Marco Sammon [a], John J. Shim [b],*

[a] Harvard Business School, Harvard University, United States of America
[b] Mendoza College of Business, University of Notre Dame, United States of America

ARTICLE INFO

JEL classification:
G11
G23

Keywords:
Stock indexes
Index providers
Index funds
Rebalancing
Issuance
Buybacks

ABSTRACT

Value-weighted indexes must rebalance in response to stock market composition changes, e.g., issuance, buybacks, and IPOs. In doing so, existing index funds implicitly engage in market timing. Index funds' long-short rebalancing portfolios have an annualized return of 4.61% and load negatively on value and profitability factors. We estimate these trades impose a 46–69 bps annual index-level performance drag. We explore alternative value-weighted indexes that rebalance less and delay responding to compositional changes. Despite still closely tracking the market, these indexes improve market timing and lower trading costs, saving 50 bps annually, an order of magnitude greater than index fund fees.

1. Introduction

Early finance theories, such as the Capital Asset Pricing Model (CAPM) and modern portfolio theory, prescribe that holding the market portfolio is optimal. And, in practice, the most common way to "own the market" is to buy index funds. Financial advisors and academics have followed suit, largely recommending diversified index funds as blanket financial advice for would-be stock market investors. However, when investors buy an index *fund*, they are relying on an *index* to operationalize the objective of "owning the market", as index funds nearly perfectly mimic the indexes they track.

Early indexes were launched in the 1950s, two decades before the launch of the first retail index fund in 1976. The most popular broad-based indexes are value-weighted (i.e., weighted by market capitalization), and are defined by a set of rules that determine what stocks to hold, how much to hold, and when to buy and sell shares, i.e., when to rebalance. In this paper, we ask whether existing total-market indexes effectively track the overall market, what causes a value-weighted index fund to rebalance, and whether rebalancing rules ultimately benefit investors.

We develop a simple framework to show the equivalence between holding a value-weighted total stock market portfolio and owning the same fraction of each stock's shares outstanding. We call this fraction the *ownership ratio*. This equivalence implies that changes in shares outstanding, i.e., changes in the composition of the stock market, necessitate rebalancing to reestablish a constant ownership ratio across stocks. Using this framework, we show that the rebalancing rules used by the most popular total market indexes lead to a large tracking error with respect to the daily-rebalanced total market. In addition, the long/short portfolio capturing value-weighted index funds' rebalancing trades has an annualized return of 4.61%, which translates to a 46 bps to 69 bps index-fund-level performance drag. We construct hypothetical

☆ Nikolai Roussanov was the editor for this article. We thank audiences at Colorado Boulder, Four Corners, Geode Capital Management, Harvard Business School, Maryland Junior Faculty Research Conference, Notre Dame, and SFS Cavalcade, and we also thank Malcolm Baker, Rodney Comegys, Mark Egan, Tim Estella, Anna Helmke (discussant), Juhani Linnainmaa (discussant), Tim Loughran, Ananth Madhavan, Ben Matthies, Chris Murray, Lubos Pastor, Jeff Pontiff, Jim Rowley, Shri Santosh, Emil Siriwardane, Erik Stafford, Adi Sunderam, Russ Wermers (discussant), James White (discussant) and Rafa Zambrana for helpful comments and suggestions.

Received 30 May 2025; Received in revised form 17 December 2025; Accepted 27 December 2025

Available online 7 January 2026
0304-405X/

indexes to understand how alternative rebalancing rules can benefit investors while being cognizant of tracking the overall market. We find that an index using a "sleepy" rebalancing approach, characterized by rebalancing every one to two years and instituting a built-in lag of one to two years before incorporating changes in shares outstanding into ownership ratios, significantly improves returns. Specifically, it both boosts returns by approximately 40 to 50 basis points per year and improves return per unit of *market tracking error*, which we define as the tracking error of the index with respect to the daily-rebalanced total market. The majority of the benefits stem from waiting to respond to endogenous market composition changes, which avoids poor market timing.

We now provide more details on our findings. Our paper has three parts: (1) we establish that a value-weighted index has an *ownership ratio* interpretation, which helps identify the set of choices available to an index designer, (2) we document how existing indexes perform from a market tracking error and a rebalancing perspective, and (3) we describe an alternative, implementable value-weighted index that would have delivered superior overall returns and higher returns per unit of market tracking error.

We start by outlining the methodology used by a value-weighted total market index and how an index fund would operate to follow that index with no index tracking error. The methodology highlights that the fund holds the same fraction of each constituent stock's shares outstanding. We call this fraction the *stock ownership ratio*. For a value-weighted total market index fund, each stock ownership ratio is identical and equal to the overall *fund ownership ratio*, defined as the fund's assets under management (AUM) divided by the total market capitalization of the stock market.[1] So, a simple way to think about how a value-weighted index fund determines its holdings is by maintaining a single ownership ratio. That is, it keeps each stock-level ownership ratio and the fund ownership ratio constant and equal to one another.[2]

The ownership ratio framework also helps to understand how changes in the composition of the stock market affect a value-weighted index and necessitate index rebalancing. If a stock increases its shares outstanding through some form of issuance (e.g., a Seasoned Equity Offering (SEO) or stock-based compensation), all else equal, the ownership ratio for that stock would go down, making it misaligned with all other ownership ratios. In order to maintain a value-weighted index, the index (and index fund) must purchase some of the newly issued shares and fund the purchase by proportionally selling all other stocks until the ownership ratios are realigned. The same logic applies when many stocks issue and/or repurchase shares, and when the fund must purchase an Initial Public Offering (IPO) or reinvest proceeds from a take-private transaction.[3]

To prevent excessive turnover, real-world indexes use *index-eligible shares outstanding*, instead of actual shares outstanding, to compute the ownership ratio, and update index-eligible shares outstanding only at prespecified rebalancing dates. Because a value-weighted index is equivalent to the ownership ratio framework, the only discretion a value-weighted index has is when and how to update index-eligible shares outstanding, which would trigger rebalancing to reestablish constant ownership ratios. That is, a value-weighted index designer picks (1) how frequently to rebalance, and (2) the timeliness of the shares outstanding data used to update index-eligible shares (i.e., whether to institute a minimum delay in incorporating shares outstanding changes).

Our framework clarifies that an *index* only adjusts holdings when it updates index-eligible shares outstanding. This contrasts with an *index fund*, which must also handle flows and other cash management considerations. Therefore, an index is defined by rules that determine how it responds to changes in the composition of the market. In fact, such compositional changes are the *only* reason a value-weighted index fund needs to rebalance. These rules affect index fund returns in two ways: trading costs (i.e., what and how much to trade), and the market timing of trades (i.e., if composition changes have predictive power for future returns). For rhetorical ease, we sometimes interchangeably use the terms index and index fund, under the assumption that the index fund has essentially no index tracking error with respect to the underlying index itself (which is true in reality).

The second part of our paper studies existing U.S. total stock market indexes. We define *market tracking error* (MTE) as the standard deviation of the difference between the index return and the actual daily U.S. stock market. In line with the CAPM's prescription of "holding the market", we define the true daily U.S. stock market as all ordinary common shares traded on major exchanges with a valid closing price the previous day. This is distinct from traditional (index) tracking error, which is defined as the standard deviation of index fund returns minus index returns. We find that the major value-weighted indexes all have significant MTE, ranging from around 75 bps per year to over 300 bps per year. For context, the large funds tracking these indexes have *index* tracking error of around 5 bps per year or less.

Our ownership ratio framework implies that if an index has a large MTE, it must be due to that index's rebalancing policy. As a result, we ask whether these rebalancing policies lead to higher or lower returns for investors. That is, do indexes deviate from the market to offer higher average returns for investors? To this end, we decompose value-weighted index funds' holdings into a combination of four distinct portfolios: the portfolio held after the previous rebalancing event, the long-short rebalancing portfolio of existing stocks (i.e., the intensive margin), the long-short rebalancing portfolio of new index additions/deletions (i.e., the extensive margin), and the flow-induced scaling portfolio.

We find that the combined rebalancing portfolios have an annualized return of −4.6% per year and an annualized alpha with respect to the Fama French 5-factor model augmented with momentum, reversal, and issuance factors of approximately −4.0% per year. This is primarily driven by the intensive margin rebalancing portfolios because they are larger in dollar terms, although the extensive margin portfolios have significantly more negative returns. In addition, the combined rebalancing portfolios have negative and statistically significant loadings on value (HML), profitability (RMW), and investment (CMA). This indicates that index rebalancing inadvertently takes on factor exposure, something surprising for investors that buy a total market index fund to specifically avoid non-market factor exposure.

Because rebalancing portfolios constitute only a fraction of an index fund's AUM, we must take into account their size to understand the effect on fund-level returns. Empirically, we find that the dollar value of these rebalancing portfolios ranges from roughly 10% to 15% of AUM. Therefore, a back-of-the-envelope calculation implies that a −4.6% annual return on index funds' rebalancing portfolios manifests at the fund level as a performance drag of 46–69 bps per year. So, while index investors have (rightly) long been focused on expense ratios, the rebalancing rules of the underlying indexes have imposed costs that are an order of magnitude larger.

Given existing index design choices lead to significant MTE *and* lower returns, we explore hypothetical value-weighted total stock market indexes to see if different choices can benefit investors. As our ownership ratio framework points out, value-weighted indexes can only vary via two discretionary choices: the frequency of rebalancing and

[1] With a subset of the market, the denominator would be the total market capitalization of the stocks in that market segment.

[2] With float adjustments, the intuition is the same but all quantities become float-adjusted quantities. That is, the stock level ownership ratio becomes shares held divided by float-adjusted shares outstanding, and the AUM ownership ratio becomes AUM divided by float-adjusted total stock market capitalization.

[3] For the many stock indexes that target particular segments of the market, additions and deletions from the index can be interpreted in the same way as an IPO or stocks going private for the purpose of index rebalancing.

the minimum delay in shares outstanding data used to update index-eligible shares outstanding. We construct our hypothetical indexes to range in rebalancing frequency from once a month to once every two years, and with a minimum shares outstanding delay of 0 days to 4 years. To clarify, the shares outstanding delay imposes a minimum time that must elapse before a change in shares outstanding can be used to update index-eligible shares at the next rebalancing date. Following the ownership ratio framework, we compute exactly how each hypothetical index would rebalance given the history of shares outstanding and stock returns from 1977 to 2023. And, using the holdings of these hypothetical funds, we calculate the returns of each index. However, this does not account for trading costs, which are embedded in the returns of the real-world indexes that we wish to compare to.

In this paper, we take the view that trading costs materialize as price impact costs, and the price impact often occurs in anticipation of rebalancing. Greenwood and Sammon (2025) document that, over the past 20 years, index price effects have increasingly occurred prior to the actual rebalancing date, due to the growing prevalence of front running by investors anticipating index fund demand. Given that the announcement and effective dates no longer exhibit significant price movements, we think of the front running as largely shifting the price impact from index fund trading earlier in time. That is, if there were no front running, price impact costs would be borne directly by the index fund at, say, the closing auction on the index rebalancing effective date. We estimate these costs using a price impact model based on which stocks and how much of each stock our hypothetical index would trade in a given rebalancing event. Finally, in computing the total return for each hypothetical index, we subtract our estimate of trading costs.

We find that our hypothetical index returns in excess of the daily total stock market return increase with longer shares outstanding delays and with less frequent rebalancing. Delays show more consistent improvement in returns than less frequent rebalancing, where a two-year delay yields an excess return of around 50 bps per year, regardless of rebalancing frequency. Although the effect of longer delays on returns is roughly monotonic, the benefits plateau with delays longer than two years.

While longer delays and less frequent rebalancing improve returns, they also increase MTE. This is mechanical, as such delays make the index less responsive to changes in market composition and, as a result, the index portfolio and index return will differ from the market portfolio and market return. To better understand how index rebalancing rules affect returns at the expense of tracking error, we define the *market information ratio* (MIR) as the index return (net of trading costs) minus the daily-rebalanced total market return, all divided by the index's MTE. MIR comes from the original motivation of the paper: if investors' goal is to own the total stock market, there may be some cost of deviating from that benchmark. And, given existing real-world indexes have a large MTE, it seems sensible that, if an index is to incur some MTE, it should at least make choices that maximize returns per unit of MTE. We find that a shares outstanding delay of 1 to 2 years roughly maximizes MIR.

Finally, we identify the source of these benefits. The index design choices of when and how to rebalance directly change the size and composition of the long-short rebalancing portfolios described above. That is, these choices affect returns through trading costs and market timing (i.e., the predictability of returns after market compositional changes).

In terms of trading costs, we find that less frequent rebalancing yields lower trading costs by accessing a flatter part of the price impact curve, which is increasing and concave in quantity traded. In terms of magnitudes, we find that trading cost savings are around 3 bps to 9 bps per year, similar in magnitude to index fund expense ratios. However, we find that market timing benefits are an order of magnitude greater than the effects of reducing trading costs. A minimum two-year delay in updating shares outstanding yields greater returns of 50 bps per year. This magnitude aligns with the fund-level cost implied by our long-short rebalancing portfolio estimates for real-world value-weighted index funds. In other words, indexes designed to rapidly respond to compositional changes in the market may inadvertently load on asset pricing factors – such as issuance or other factors related to net issuance – that are associated with negative returns.

As a final test, we examine whether the results of our delayed rebalancing strategy can be replicated by dynamically trading the market factor and an issuance factor. After all, the rebalancing rules are essentially modifying the way an index responds to changes in shares outstanding. While an ex-ante version of this dynamic trading strategy achieves a slightly higher Sharpe ratio than our delayed rebalancing indexes, it does so by taking on significantly greater market tracking error and delivering lower average returns. Therefore, this strategy yields a lower market information ratio and would require leverage to match the market's average return. We view this as further evidence that our delayed rebalancing portfolios more practically benefit investors, and highlight ways to avoid the unique issuance-related costs that value-weighted market investors are exposed to.

Contributions. We make three sets of contributions to the literature. The first contribution is to identify the ownership-ratio approach to understanding an index and index fund. A value-weighted index fund owns a percentage of the market or market segment and owns the same percentage of each constituent stock's shares outstanding. We think of the ownership ratio as a simple way to understand how an index fund operates and how it must trade in response to both cash considerations and to changes in the composition of the stock market. Further, in the Appendix, we provide empirical evidence that the ownership-ratio logic explains how real-world index funds actually trade.

The second contribution is to provide an estimate of the implicit costs incurred by passive investors that attempt to track the total U.S. stock market with index funds. Because most index funds rebalance quarterly with no delay in shares outstanding used to update index-eligible shares, they are exposed to market timing by firms. Said differently, we identify the issuance effect that index fund investors are implicitly exposed to and show that this issuance effect is both related to existing issuance factors but also has a component that cannot be explained by known factors. This type of issuance exposure leads to performance drag that is an order of magnitude larger than index fund expense ratios.

The third contribution is to offer an alternative "sleepy" index design which minimizes the market timing and transaction costs paid by index fund investors per unit of market tracking error. While media and academic scrutiny have focused on the cost of expense ratios and management fees to investors, we show the costs of index construction and rebalancing policies to investors are about 10 to 20 times as large. This finding emphasizes that, while index providers only have a few degrees of freedom in designing an index, the choice of rebalancing horizon and frequency has a substantial effect on how those indexes perform. At least in our sample of almost 50 years of data, slightly adjusting index design could have rebated these costs (and more) back to index fund investors.

1.1. Related literature

The most closely related paper to ours is Pedersen (2018), which also points out that index funds must trade as the composition of the stock market changes. However, Pedersen (2018) focuses on the role of active managers in this context. Because passive investors are price takers, active managers can justify high fees by pricing compositional changes and providing liquidity to non-passive investors and firms that enter and exit the market (on the intensive and extensive margin). Our focus is similar, but we show that, in many respects, index funds are not truly "passive" with respect to these changes, but rather active responders through their choice of rebalancing policy, and ultimately that the speed of their response creates costs for passive investors.

Our paper is also related to Harvey et al. (2025), who show evidence of measurable price pressure following mechanical stock-bond rebalancing, and, like us, show that this materializes as a cost to investors. We highlight that our results, which focus on longer horizon returns, are more likely to be driven by poor market timing than predictable price pressure. In particular, our finding that flow-based scaling does not lead to predictable returns at the quarterly frequency suggests that, if there is price pressure, both the pressure and subsequent reversal are accounted for within the quarterly horizon that we examine. This is further confirmed by our analysis that isolates market timing in our hypothetical index portfolios by avoiding the "crowded" trades around real-world index rebalancing events.

Our research contributes to several other strands of literature, the first of which focuses on market timing by firms. Within this area, there is an extensive literature that studies IPOs, firm issuance, and buybacks in the cross section of equities. IPOs and SEOs exhibit long-run underperformance, as shown in Ritter (1991), Loughran and Ritter (1995), and Spiess and Affleck-Graves (1995). In addition, acquisitions using stock (Loughran and Vijh, 1997) also forecast lower future stock returns for the acquiring company. Ikenberry et al. (1995) study share repurchases (i.e., buybacks) and find that stock returns are higher following repurchase announcements, especially for value stocks.

Several studies also find that the cross-section of net issuance has predictive power for stock returns. Daniel and Titman (2006) find that 5-year composite issuance predicts future stock returns, with the strongest predictive power over the 1- to 3-year horizon. Pontiff and Woodgate (2008) find that a 1-year and 5-year measure of past net issuance has a negative relation with future returns, with the 1-year measure serving as the somewhat stronger predictor. Fama and French (2008a,b) also examine the relationship between share issuance and expected returns and find similar results. More recent work by Baba-Yara et al. (2024) shows that issuance effects are among the most persistent of firm characteristics, with net issuance predicting returns for up to 4 years after the issuance itself. Finally, within the context of firm issuance and index funds, Sammon and Shim (2024) show that firms take the other side of passive buying and tend to do so when valuations are high.

There has also been important work that studies *aggregate* issuance and buybacks. For example, Baker and Wurgler (2000) document that aggregate equity issuance has predictive power for future overall stock market returns. Dichev (2007) argues that actual investor dollar-weighted returns are lower than buy-and-hold returns, precisely because of the timing of equity capital flows (i.e., issuance and repurchases). Dichev (2007) finds that the difference between these two return measures is about 1% per year in the U.S., and is positive and significant for almost all international stock markets as well. There are additional findings on equity offerings and the aggregate stock market from the IPO literature. IPOs tend to occur in "waves", often coinciding with periods of high market valuations (Pástor and Veronesi, 2005; He, 2007; Chemmanur and He, 2011). This behavior aligns with broader evidence of market timing by firms. Empirically, this pattern is associated with subsequent underperformance, as stocks that go public during IPO waves tend to generate lower returns in the periods that follow.

Finally, our paper contributes to the growing literature on index fund design. Sammon and Murray (2024) show that the mechanical buying of stocks immediately post-IPO by index funds leads to underperformance and that obtaining allocations during IPOs would yield better outcomes. Similarly, Arnott et al. (2023) demonstrate that trading with index rebalancers (i.e., investors trading in the same direction as index funds' demand shocks) often exposes these traders to post-trade reversion, and that waiting to, e.g., buy stocks added to major indexes results in better performance.

More broadly, there is an ongoing debate about where an index's boundaries should lie (Sauter, 2022). These practical conversations about index funds and rebalancing tend to focus on trading costs, which are especially salient to index fund managers if they lead to index tracking error. Our results suggest that the effect of poor market timing by index funds is an order of magnitude larger than estimated index fund trading costs and index fund fees. We believe, therefore, that it may be beneficial to broaden the set of index fund design considerations to include market tracking error, implicit trading costs, and market timing.

2. Data

Fund data. We employ an extensive data filtering and cleaning approach. This is necessary because our empirical analysis uses changes in stock-level holdings for each index fund. Working with changes is very sensitive to data issues, largely because the errors in the underlying data are in *levels*. Because changes are small relative to levels, erroneously inferring position-level errors as position-level changes results in noise that is an order of magnitude larger than most actual changes. To reduce the effect of these position-level errors, we follow the approach in Sammon and Shim (2023).[4]

We start with all stock-fund-quarter Thomson Reuters S12 mutual fund holdings data between 1980 and 2023. We merge this data with the CRSP stock data using historical CUSIP. We correct for splits and other share adjustments by multiplying raw shares held by CRSP's cumulative factor to adjust shares, and divide prices by CRSP's cumulative factor to adjust prices.

We then use the CRSP MF Links database to link Thomson Reuters S12 fund numbers ($fundno$) to CRSP fund numbers ($crsp\_fundno$). When referring to funds throughout the paper, we use the S12 $fundno$, which groups all share classes associated with individual $crsp_fundno$s into a single S12 $fundno$. We identify passive funds following the procedure in Appel et al. (2016), although results are similar if we identify passive funds only by the index fund flag in the CRSP mutual fund database.

We apply several filters at the fund-quarter level which are standard in the literature, designed to account for incubation bias and funds with significant non-equity holdings. We also remove all fund-quarter observations affected by two data timing issues discussed in the S12 documentation: stale data and reporting delays. This, again, is important because of our use of holdings changes. Finally, we apply several data-quality filters at the *stock-quarter* level to reduce the effect of outliers. See Appendix A for a detailed explanation of our data cleaning process.

Index fund and index data. The main sample focuses on index funds tracking the three major index families that are the primary focus of this paper: the S&P family of indexes (500, 400, and 600), the Russell family of indexes (1000, Mid-Cap, 2000, and 3000) and the CRSP market-capitalization-based indexes and the CRSP total market index. Restricting to funds tracking these indexes leaves us with 2344 fund-quarter observations across 82 index funds. We extend the analysis to all index funds in Appendix D.1 and D.2. This sample is significantly larger, with 41,153 fund-quarters across 1376 passive funds. We obtain data on indexes' total returns from Thomson Reuters Datastream and S&P Capital IQ.

Stock data. We define the universe of stocks considered throughout the paper as all ordinary common shares (CRSP share codes 10 and 11) traded on major U.S. exchanges (CRSP exchange codes 1–3). We also use the CRSP data for shares outstanding, returns, and prices. Finally, we use CRSP daily stock data to construct the daily value-weighted total stock market portfolio.

[4] See the data appendix in Sammon and Shim (2023) for more details. The replication code can be found on the authors' websites.

3. Conceptual framework for a value-weighted stock index

In this section, we provide a conceptual framework to better understand that a value-weighted total stock market index – and corresponding index fund – achieves value weights by maintaining a constant ownership ratio across stocks. We then discuss the minor real-world tweaks that the most popular stock indexes add, and show how this slightly changes the intuition from our simple example. Lastly, we precisely explain how index rebalancing works according to the ownership ratio framework.

Our goal in this section is to define what it means to be a value-weighted index, i.e., to make explicit that the ownership ratio logic is the underlying principle behind the design of *all* value-weighted indexes. This framework will serve as the foundation for the hypothetical indexes and index funds we study in Section 5, ensuring that they are grounded in the same logic that governs real-world value-weighted indexes.

We will also often interchangeably use the term index and index fund, even though they are distinct entities. We do this because in this section and in Section 5, we assume that there exists an index fund that perfectly tracks the underlying index with zero index tracking error. This assumption is mostly for rhetorical purposes — an *index* does not literally hold positions, and an *index fund* does not update index-eligible shares outstanding. But, because index funds mimic their index up to a scaling factor, it is convenient to sometimes use the term index while considering, say, holdings or flows.

3.1. Ownership ratio framework

Our framework starts with a value-weighted total stock portfolio. For simplicity, we start with an index that does not apply any stock-level filters or float adjustments to shares outstanding. Therefore, this index includes every stock in this investment universe, and the weight of each stock is that stock's market capitalization as a fraction of total market capitalization.

Define the *fund ownership ratio*, Ω_t, as

$$\Omega_t = \frac{aum_t}{totalme_t}, \tag{1}$$

where aum_t is the assets under management of the fund that perfectly tracks our index, and $totalme_t$ is the total market capitalization of the universe of stocks we consider. In addition, define the *stock ownership ratio*, $\Omega_{i,t}$, as

$$\Omega_{i,t} = \frac{sharesheld_{i,t}}{shrout_{i,t}}, \tag{2}$$

where $sharesheld_{i,t}$ is the number of shares held by the fund in stock i at time t, and $shrout_{i,t}$ is the number of shares outstanding for stock i at time t.

An index fund that holds a value-weighted portfolio holds positions such that the fund ownership ratio and each stock ownership ratio are equal, or

$$\Omega_t = \Omega_{i,t} \quad \forall i. \tag{3}$$

This expression provides another way to view how value-weighted index funds work. It says that owning a value-weighted total stock market portfolio is identical to owning a constant fraction of the shares outstanding of every stock, and that fraction is given by what fraction of the overall market the fund owns. That is, if a fund tracking a value-weighted index has AUM equal to 1% of the total market capitalization, it exactly achieves this by owning 1% of every stock's shares outstanding. For complete details on the equivalence between Eq. (3) and a value-weighted index/index fund, see Appendix B.1.

It is intuitive to think of a value-weighted portfolio as being determined by the weights (and, as a result, prices) of each stock. However, the weight-centric view obscures a key insight: despite daily price fluctuations that change weights, value-weighted index funds require no trading activity to maintain the proper value-weighted holdings. The ownership ratio reveals why — by holding a fixed percentage of each stock's shares outstanding, the portfolio automatically maintains correct value weights.

3.2. When does an index fund trade?

There are two types of events that cause an index fund to trade.[5] The first event is a cash-driven "scaling" event. These events largely occur at the index fund level instead of the index level. These events include fund inflows or outflows, payment of dividends, and receipt of delisting proceeds. Inflows and outflows change the AUM of the fund, or the numerator in the *fund* ownership ratio, Ω_t. In order to reestablish a constant ownership ratio, the fund must proportionally scale up or scale down the holdings of each stock. Delisting events work similarly to inflows, as again the fund will use delisting proceeds to scale up existing positions.[6]

The second event that causes an index fund to trade is a change in stock market composition. This can come from one or more firms that change shares outstanding via a seasoned equity offering (SEO), a share repurchase, or new shares awarded via employee compensation programs (and thus added to shares outstanding). When the fund or stock-specific ownership ratios change because of changes in shares outstanding, the index fund no longer holds a value-weighted market portfolio and must trade to reestablish constant ownership ratios as presented in Eq. (3). For example, issuance via an SEO in stock i will cause $\Omega_{i,t}$ and Ω_t to both decline, but all $\Omega_{j,t}$ for $j \neq i$ will be unchanged. This means that the index must adjust positions by increasing its position in i (to raise $\Omega_{i,t}$), and fund that by decreasing positions in all other stocks j until they match the new fund ownership ratio Ω_t. In the (typical) event where many stocks adjust shares outstanding, the index will increase positions of stocks that issued relatively more and decrease those that issued relatively less/repurchased relatively more.

These adjustments show that changes in shares outstanding lead to heterogeneous trading across positions (in contrast to cash-driven scaling described above), which is widely referred to as index *rebalancing*. New additions – e.g., IPOs for total market indexes – have a similar effect as issuance by existing firms. In the case of an IPO, the total market capitalization increases, lowering Ω_t. As a result, the index must scale down all existing positions to raise cash and buy shares of the new IPO. For existing (non-IPO) stocks, cross-sectional differences in net issuance are the *only* determinant of rebalancing-motivated trading by index funds.

3.3. Index-eligible shares outstanding and index design

In our ownership ratio framework, any change in shares outstanding would trigger the need to trade to reestablish a constant ownership ratio. This would cause an index to potentially rebalance daily because shares outstanding can change frequently, leading to excessive trading by funds that track the index. The solution that many indexes

[5] Many index funds now operate as an ETF instead of a mutual fund structure. The same analysis applies to arbitrage activity for the purpose of creation/redemption by market makers and authorized participants (APs). That is, this activity is quite literally the flows into and out of an ETF. The fact that most large equity ETFs have premiums close to zero suggests that the flows are priced close to NAV.

[6] Delistings are a special case among cash events, in that they change the denominator of the fund ownership ratio and set the delisted stock's ownership ratio to zero. Consider a delisting where a stock pays a liquidating dividend to holders of the stock. This delisting will cause the size of the total stock market to decline. This will cause the *fund* ownership ratio to increase, and the fund reestablishes a constant *stock* ownership ratio by increasing holdings in each of the non-delisted stocks.

adopt is to use *index-eligible shares outstanding* instead of actual shares outstanding when equating ownership ratios. This approach updates index-eligible shares outstanding at regular intervals, but holds index-eligible shares outstanding fixed until the next rebalancing event. As a result, the index does not need to immediately rebalance when market composition changes. This also means that index providers' policies on the frequency and methodology for updating index-eligible shares outstanding will dictate *how much* index funds trade at rebalancing events (which determines trading costs) and *when* index funds will make the trades (which determines the index's "market timing").

For example, say an index rebalances today (and establishes a constant ownership ratio), and will rebalance next one month from now. If a stock issues shares the day after the rebalance, and another stock issues shares the day after that, these changes will only be incorporated into the various ownership ratio denominators at the next rebalance in about a month. At that rebalance, all stocks' index-eligible shares outstanding will be updated simultaneously, typically with the most recent shares outstanding amounts. New additions will also be incorporated at this time into the index.[7] Because of these updates, the index fund will buy or sell all stocks in the portfolio. For brevity, we will sometimes use the term "shares outstanding" instead of the full term "index-eligible shares outstanding" when describing the index choices that determine how and when to rebalance.

Our framework, coupled with the practice of updating index-eligible shares outstanding only at regular intervals, makes clear that a value-weighted index is defined by two choices: how and when to update shares outstanding. Because a value-weighted index is beholden to maintaining a constant ownership ratio, the shares held in each stock are simply a function of the fund's AUM, the total stock market capitalization, and index-eligible shares outstanding. For an index, the only degree of freedom is index-eligible shares outstanding, which can be determined via "how" to update shares outstanding and "when" to update shares outstanding.

For the "how", an index may, for example, choose to use a more stale value of shares outstanding instead of the most recent at the time of rebalancing. For the "when", an index may rebalance every year instead of every quarter. Both of these choices will affect the composition and timing of the rebalancing portfolio, i.e., the long-short portfolio that is used to adjust existing positions to reestablish a value-weighted index/index fund. These choices, as we will see in Section 4, will affect how an index differs from the actual total stock market portfolio and the way in which an index implicitly times the market by responding to firm issuance and repurchase activity. In Section 5, we will explore a wide range of choices on how and when to update index-eligible shares outstanding, which will impact both index returns and how well an index tracks the market.

3.4. Rebalancing

To precisely describe how the index/index fund adjusts in response to changes that could affect either the fund ownership ratio or the individual stock ownership ratios, we must make some assumptions about the timing of these changes.

We assume the following events occur sequentially between rebalancing periods (e.g., over a quarter), and all except for returns directly affect the fund ownership ratio or one of the stock-specific ownership ratios:

1. **Returns.** Each stock has some return. Together, these returns also generate a value-weighted market return. Returns do not directly affect the ownership ratio because it changes the numerator and the denominator of the fund ownership ratio by exactly the same factor.
2. **Flows.** The fund may have some inflow or outflow. This changes the fund ownership ratio by changing the fund AUM.
3. **Issuance/Buybacks.** Firms can repurchase or issue shares of their own stock. Buybacks or issuance in a given stock affect both that stock's ownership ratio and the fund ownership ratio by changing the denominators.
4. **Additions/Deletions.** New stocks enter the market via an IPO or existing stocks delist. This changes the fund ownership ratio through total market capitalization, and, for IPOs, introduces new stock ownership ratios that must satisfy Eq. (3).

Rebalancing in the face of the combination of compositional changes amounts to a fixed-point problem, in the sense that a fund must rebalance in a way that satisfies the stock-level *and* fund ownership ratio rules. The resulting rebalancing for each stock i that did not delist or was not an IPO is

$$\frac{shares_{i,t}}{shares_{i,t-1}} = \frac{\left(1 + r_{mkt,t} + flow_t\right)\left(1 + issue_{i,t}\right)}{\left(1 + r_{mkt,t} + issue_{mkt,t} + adddrop_{mkt,t}\right)}, \tag{4}$$

where $r_{mkt,t}$ is the value-weighted market return and $issue_{i,t}$ is the net issuance of stock i over the rebalancing window t, and is measured as a fraction of last period's shares outstanding (with negative values indicating buybacks). The two key variables in the denominator are aggregate variables: $issue_{mkt,t}$ is the aggregate dollar net issuance between rebalancing events as a fraction of last period's total market cap, and $adddrop_{mkt,t}$ is the dollar value of IPOs minus delisted firm values scaled by last period's total market cap. Complete details on deriving the necessary rebalancing conditions are presented in Appendix B.2.

The rebalancing expressed in Eq. (4) shows that a fund will be a net buyer of stock i if the flow into the fund and issuance by the firm are larger than the growth in total market cap. For example, if i's issuance is large relative to overall market issuance and IPOs, the fund will need to buy shares in stock i when it rebalances to reestablish holding the value-weighted market portfolio. As an extreme example, if all firms issued the same amount of equity as a percentage of shares outstanding, the fund would not need to rebalance.

A seemingly counterintuitive feature of Eq. (4) is that it includes $r_{mkt,t}$, even though the ownership ratio framework tells us that value-weighted index funds do not need to trade in the face of returns.[8] Including $r_{mkt,t}$ is necessary in Eq. (4) for two reasons. The algebraic reason is that $shares_{i,t}$ in Eq. (4) comes from the ownership ratio at t relative to the ownership ratio at $t-1$, and returns capture one of the ways in which the two components of the ownership ratio – AUM and total market cap – each evolve from $t-1$ to t. The second is to ensure the scale of rebalancing is correct in the sense that, e.g., a flow will have a smaller effect on rebalancing if the flow is unconditionally large but small relative to returns. So, while returns do not directly cause any rebalancing, they affect how much rebalancing must be done from other sources.

As a check, we take Eq. (4) to the data to see if rebalancing in real world index funds follows the rebalancing predicted by our ownership ratio framework. While we have made several strong assumptions that have allowed us to simplify our framework, we still find strong support that the ownership ratio logic can, rather precisely, be used to understand how real-world index funds rebalance. We describe the empirical tests and results in Appendix B.3.

In addition, there are several real-world considerations that we have abstracted from. None of these materially change how to *conceptually* think about a value-weighted index or index fund. We provide details and adaptations to incorporate these real-world details into our

[7] Some indexes will selectively update shares outstanding outside of the regular rebalance based on large single compositional changes — e.g., a large SEO or a large IPO.

[8] The exact expression for rebalancing has the fund return in the numerator of Eq. (4) instead of the market return. Given that in this example, our index fund perfectly tracks the index, the fund return and market return are equivalent. We use this simplification in the empirical verification below in Appendix B.3, given that most index funds have minimal index tracking error.

framework in Appendix B.4, including float adjustments and dividends. We also outline how our framework can help clarify some common misconceptions about value-weighted indexes in Appendix B.5.

4. Do existing indexes track the market?

In this section, we assess whether total market indexes actually track the market. To do so, we define a new measure that we call *market tracking error* (MTE), which quantifies the magnitude of deviations between an index's returns and the return of the value-weighted market as a whole. We show that real-world indexes have large MTE, and we argue that these deviations are a function of how and when indexes choose to rebalance. We also define a new measure called the *market information ratio* (MIR) to evaluate the trade-off between the returns to deviating from the market and MTE. Finally, we show that for real-world index funds, these rebalancing trades lead to poor returns, both in terms of alpha and negative loadings on asset-pricing factors known to predict high returns (e.g., value and profitability). In this section, we focus on index and index fund returns between 1990 and 2023, although the exact time period is different for each index based on data availability.

4.1. Market tracking error and market information ratio

To test whether total market indexes actually track the market, we define a new measure: *market tracking error* (MTE). Whereas index tracking error reflects how closely an *index fund* replicates its benchmark index (and is typically small for most large index funds), MTE captures deviations between an *index* and the market itself. Specifically, we define

$$\text{MTE} = \sigma(r_{\text{idx},t} - r_{\text{mkt},t}), \tag{5}$$

where $r_{\text{idx},t}$ is the return of the index and $r_{\text{mkt},t}$ is the return to the market.

In Eq. (5), MTE crucially depends on how "the market" is defined. We define a tradable version of the market as the return to a value-weighted portfolio of all ordinary common shares traded on major U.S. exchanges in CRSP with a valid closing price the previous trading day. This means we treat a recent IPO as eligible only starting from the close of its first trading day.[9] For context, our daily market portfolio is very similar to the Fama–French market factor, which is commonly used throughout the literature and also only includes ordinary common shares traded on U.S. exchanges. The return correlation between our market factor and the Fama–French market factor is 0.9998.

Unlike in standard index tracking error, trading costs and front running are *not* forgiven in MTE. Index tracking error evaluates index fund performance relative to the price at which the stock was added to the index. Because index funds often trade in the closing auction on the actual day the stocks are added to or dropped from indexes (at the closing auction price), their trades appear to have no price impact and create no tracking error. That is, if their trades have price impact at the close, it will affect both the index and index fund in exactly the same way. However, this ignores the cost of being front run by intermediaries, who move prices in anticipation of index trading (see Greenwood and Sammon (2025) for details). In contrast, MTE captures the effect of securities that are part of the market but not yet in the index, and the return drag from trading at prices affected by anticipatory demand.

By construction, deviating from the market leads an index to earn different returns. To capture the trade-off between potentially earning higher returns from deviating from the market and taking on additional MTE, we define the *market information ratio (MIR)* as

$$\text{MIR} = \frac{\mathbb{E}[r_{\text{idx},t} - r_{\text{mkt},t}]}{\text{MTE}}, \tag{6}$$

where $r_{\text{idx},t}$ is the return to the focal index and $r_{\text{mkt},t}$ is the return to the market (as before, defined as the value-weighted return on all ordinary common shares with valid closing prices on major U.S. exchanges). Eq. (6) is nearly identical to the standard Treynor-Black information ratio (IR) formula (Treynor and Black, 1973), with one main change: the residual return and volatility are not based on regressing returns on the market factor. Rather, it is the difference between the portfolio return and the daily-rebalanced market return. In other words, we implicitly set the market beta to 1 and estimate the IR just as in the Treynor-Black IR.

Table 1 Panel A reports annualized market tracking error (MTE), measured in basis points, for three major U.S. total market indexes, as well as the widely followed S&P 500, across different subperiods. The S&P Total Market and S&P 500 exhibit the largest MTEs overall, with average deviations of 309 and 234 bps per year over the full sample of available data. The next largest market tracking error is for the Russell 3000, at 145 basis points per year. This, again, is perhaps unsurprising, as the Russell indexes rebalance less frequently, and Russell applies float adjustments to many stocks in their investment universe. Finally, the CRSP Total Market Index, which more closely approximates the full investable market by including micro-cap stocks and rapidly incorporating large IPOs and SEOs, exhibits the lowest MTE.

Panel B of Table 1 reports the MIR for each index. The broad takeaway from this panel is that, for these indexes, the MIR is generally economically small. The first row shows that although the MIR for the S&P 500 has been negative over the full sample, it was positive over the last few years when large-cap stocks outperformed small-cap stocks. Similarly, the Russell 3000, which includes more small-cap stocks, has a positive overall MIR, but has a negative MIR in recent years. Finally, the CRSP total market index has a seemingly large negative MIR, but this is primarily a reflection of its very low MTE.

4.2. Returns from market tracking error

The ownership ratio logic established in Section 3 describes how a value-weighted index must be rebalanced to account for compositional changes in the stock market. Most real-world indexes respond relatively quickly – typically within 1 calendar quarter – to such compositional changes.

At a high level, Eq. (4) shows that the fund will rebalance towards stocks that had net issuance greater than the value-weighted average, and vice versa for stocks with net issuance less than the value-weighted average. Given the long literature connecting cross-sectional differences in issuance/buybacks with future returns (e.g., Daniel and Titman (2006), Pontiff and Woodgate (2008) and Fama and French (2008a), among many others), these index rebalancing policies may affect returns by implicitly timing the market. By market timing, we mean that indexes may, even inadvertently, time the buying of stocks when firms themselves think prices are high and future returns will be low (and therefore issue shares), and vice versa for selling stocks.

We take a portfolio-based approach to quantify the relationship between index fund rebalancing and future returns. This is built on the ownership ratio logic of Section 3, which implies that one can construct the overall portfolio ultimately held by a given index fund at the end of a quarter from four conceptually distinct portfolios: (1) the long-only portfolio held at the end of last quarter, (2) the long-short portfolio from index rebalancing for existing positions from changes in market composition on the intensive margin, (3) the long-short portfolio of

[9] Importantly, our definition of $r_{\text{mkt},t}$ differs from CRSP's VWRETD (Value-Weighted Return Including Dividends) series, which fixes portfolio weights within a quarter, i.e., it ignores within-quarter issuance, buybacks and some IPOs, but, more notably, includes non-ordinary-common share securities e.g., REITs and even ETFs and ETNs. The latter does make a small but noticeable difference in returns.

Table 1
Market tracking error and the market information ratio for major indexes.

	Panel A: Market tracking error (Basis points)				
	All data	1990–1999	2000–2009	2010–2019	2020–2023
S&P 500 (1990–2023)	233.9	259.0	281.1	134.8	233.2
S&P total market (1995–2023)	309.0	173.8	218.5	444.4	204.0
Russell 3000 (2005–2023)	144.9		100.0	173.4	106.9
CRSP total market (2011–2023)	74.6			59.9	99.6
	Panel B: Market information ratio				
	All data	1990–1999	2000–2009	2010–2019	2020–2023
S&P 500 (1990–2023)	−0.058	0.045	−0.198	−0.109	0.130
S&P total market (1995–2023)	−0.159	−0.179	0.099	−0.345	0.064
Russell 3000 (2005–2023)	0.009		0.143	−0.011	−0.057
CRSP total market (2011–2023)	−0.326			−0.509	−0.109

Notes. This table reports the annualized market tracking error (MTE) and market information ratio (MIR) for several popular total market indexes and the S&P 500. Data for the S&P 500 starts in 1990. Data for the S&P Total Market starts in 1995. Data for the Russell 3000 starts in 2005. Data for the CRSP Total Market starts in 2011. MTE is computed as the standard deviation of the difference between daily index returns and daily-rebalanced total market returns, multiplied by the square root of 252 (i.e., the average number of trading days in a year). The units of MTE are basis points. MIR is the index return minus the daily market return, all divided by the index's market tracking error (MTE).

additions to and deletions from the index, and (4) the long or short portfolio representing the proportional scaling of past positions based on flows and other cash-driven events. The scaling portfolio is different from the long-only portfolio held at the end of the last quarter in that it is directional, i.e., if there are outflows from a fund, the expected scaling will be represented as a short portfolio.

We call the second portfolio, based on deviations of existing positions from perfect scaling, the "rebalancing portfolio", and we use it to study whether intensive margin index fund rebalancing has predictive power for future returns. We also use the expected scaling portfolio as a placebo test, because this portfolio is less likely to be subject to poor (or beneficial) market timing, given that flows to index funds are likely uninformative about future returns. In addition, to the extent price pressure from the index fund's own trading affects returns, it should show up as reversal for the scaling portfolio in the subsequent quarter. Finally, even though index additions and deletions have been studied extensively and are not necessarily the focus of this paper (see, e.g., Shleifer (1986) and Harris and Gurel (1986), among many others), we will construct a portfolio analogue in this setting (identified as the third portfolio in the list above) so it is comparable to our intensive margin rebalancing portfolio.

To construct the long-short intensive margin rebalancing portfolio for each index fund and each quarter in our data, we compute the long leg representing all of the stocks that the fund rebalanced towards (i.e., purchased proportionally more or sold proportionally less) and the short leg representing all of the stocks the fund rebalanced away from (i.e., sold proportionally more or bought proportionally less). Our construction proceeds as follows. First, among all stocks which are held both in quarter t and in quarter $t-1$, we compute the fund's change in position after accounting for the expected change attributable to scaling due to flows and funding needs/proceeds from the imbalance of new fund additions and deletions, or

$$\Delta shares^{rebal}_{i,j,t} = shares_{i,j,t} - (1 + flow_{j,t} - adddrop_{j,t}) \times shares_{i,j,t-1}. \quad (7)$$

In words, this is the difference between the number of shares of stock i held by fund j in quarter t, and the predicted number of shares given perfect proportional scaling of past positions in the face of flows ($flow_{j,t}$) and add/drop imbalances ($adddrop_{j,t}$). Note that Eq. (7) is a direct application of the ownership ratio logic in Equation A.15. For each fund-quarter, the long leg of the portfolio is composed of stocks with $\Delta shares^{rebal}_{i,j,t} > 0$, while the short leg is composed of stocks with $\Delta shares^{rebal}_{i,j,t} < 0$.

We then compute portfolio weights separately for the long and short portfolios to be proportional to $|\Delta shares^{rebal}_{i,j,t}| \times p_{i,t}$, where $p_{i,t}$ is the end-of-quarter price. In words, for each fund-quarter, we construct a value-weighted portfolio that is long stocks with an increase in shares held relative to what was predicted by perfect scaling and is short stocks with a decrease in shares held relative to scaling. The weights in each stock and in each leg of the portfolio are given by the expected dollar trade that the fund made to deviate from perfect scaling divided by the total expected dollar trading in that leg. An underlying assumption of this exercise is that we construct weights assuming all trading takes place at the end of the quarter (i.e., the trade price corresponds to the final price of the quarter). Finally, for all rebalancing portfolios, if a stock delists over the following quarter, we assume any delisting proceeds are invested in the market (defined for this exercise as MktRF + RF from Ken French's data library) for the remainder of the quarter.

We construct the scaling portfolio following the exact same procedure, but using $shares_{i,j,t-1} \times (flow_{j,t} - adddrop_{j,t})$ instead of $\Delta shares^{rebal}_{i,j,t}$. As mentioned above, this portfolio captures *expected* trading by the fund if it were to perfectly scale up/down existing positions given flows and the add-drop imbalance. This acts as a placebo test, as signed scaling is unlikely to be affected by market timing with respect to firm decisions unless investor flows time the market or if scaling is large enough to lead to persistent price pressure. As with the intensive margin rebalancing portfolio, stocks are weighted by the expected dollar trade in that stock divided by the total expected dollar trade in that fund and quarter. For each fund-quarter, there is only one leg of this portfolio, as the fund can only either have net positive or net negative expected scaling activity. To mirror the long-short construction for the rebalancing portfolios, we assign the market return to the missing leg for each fund-quarter.

We construct the "additions and deletions portfolio" (i.e., the extensive margin rebalancing portfolio) using the same procedure, where the long leg is composed of index additions, and the short leg is composed of index deletions. Specifically, the quantity is defined as the number of shares held at the end of quarter t for new additions, and is defined as the number of shares held at the end of quarter $t-1$ for deletions. As with the other portfolios, positions are weighted by the dollars of trading in each stock (again assuming that all trade occurs at the end of quarter price) relative to the total dollars of trading for each leg of the portfolio. If a fund drops a stock because the stock delisted, we assume that the total dollars of trading is equal to shares held at the end of $t-1$ times the delisting amount per share (all split-adjusted). If this portfolio is missing the long or short leg, we assign the missing leg the return of the market. For example, if the fund had several drops but did not have any adds, we create a short leg that has the market return.[10]

Finally, we create a total rebalancing portfolio, which combines the intensive and extensive margin rebalancing portfolios. The long leg of

[10] Results are similar if we instead assign the missing portfolio legs the risk-free rate of return.

the total rebalancing portfolio combines the long leg of the intensive and extensive margin rebalancing portfolios, but re-weights each stock based on the expected trade divided by the total amount traded in the combined long leg. The short leg is constructed analogously. Because new additions and deletions are typically small relative to the total amount of intensive margin trading, this portfolio generally has returns closer to the intensive margin rebalancing portfolio than the extensive margin rebalancing portfolio.

We construct these portfolios for a set of funds that are most likely to be value-weighted index funds. As described in Section 2, to be included in the sample, we require that (1) the fund is classified as passive in the CRSP Mutual Fund database, and the fund either has (2A) a median active share (Cremers and Petajisto, 2009) with respect to the major value-weighted benchmarks (S&P 500, 600, 400, Russell 1000, 2000, 3000) of less than 5%, or (2B) the fund is one of the large Vanguard funds benchmarked to the value-weighted CRSP indexes, e.g., VTI. The sample runs from 1996 to 2023 due to data availability for the set of funds that pass our filters.[11] In Appendix D.2, we replicate these results for all funds identified as passive in the CRSP mutual fund data.

For each portfolio constructed for each fund each quarter, we consider the returns to that portfolio in the quarter *following* the trades. For example, if an index fund adds 10 stocks in 2020 Q4, the long leg of that fund's extensive margin rebalancing portfolio considers the returns to those stocks in 2021 Q1. These long-short portfolios at the fund-quarter level are the unit of observation in our analysis.

Table 2 presents the results of regressing these portfolio returns on a set of asset pricing factors. In column 1, the constant term is the mean return of our intensive margin rebalancing portfolio and shows that the portfolio has an average annualized return of −4.15%. This is consistent with stocks that index funds rebalance towards underperforming the stocks they rebalance away from. And, this is consistent with indexes engaging in poor market timing, as stocks issuing relatively more equity are in the long leg of the portfolio and those issuing relatively less (or buying back relatively more) are in the short leg.

One might be concerned that the results in column 1 are driven entirely by price pressure. Specifically, the index fund rebalancing itself might affect prices, which in turn would affect future returns. We believe this is not the case for several reasons. First, with these intensive margin rebalancing portfolios, we are removing the portion of index fund trading explained by flows, which is significant and has been shown to predict returns due to price pressure (see, e.g., Frazzini and Lamont (2008) and Lou (2012)). Second, an index fund may actually be a net *seller* of some of the stocks in the "long" leg of the intensive margin rebalancing portfolio (or a buyer in the "short" leg). This could happen if, e.g., the fund received an outflow of 5%, but the fund only reduced its position in the stock by 2% due to a contemporaneous issuance in that stock. In this case, the fund is "long" the stock from the perspective of the rebalancing portfolio, even though it sold the stock in practice, i.e., it sold less than it otherwise would have if only responding to flows.

In column 2, we add the market factor and find the alpha does not change significantly from the mean return in column 1, and the market beta estimate is economically small. Finally, in column 3, we control for the Fama–French 5 factors, as well as momentum, short-term reversal, and an issuance factor.[12] Even after controlling for a host of factors, the alpha is still negative and statistically significant at −3.2% per year, albeit with a somewhat smaller magnitude.

One of the reasons the alpha shrinks is because the intensive margin rebalancing portfolio loads negatively and significantly on the value factor (HML), the profitability factor (RMW), the investment factor (CMA), and the issuance factor. These loadings indicate that the intensive margin rebalancing portfolio has a tilt towards growth, unprofitable, aggressive investment firms, as well as firms which have recently issued equity. This is not surprising for two reasons. First, our ownership ratio framework shows that cross-sectional differences in issuance are the sole driver of intensive margin rebalancing. Second, growth stocks, unprofitable stocks, and firms with aggressive investment are precisely the types of firms most likely to issue new equity and least likely to repurchase shares. Interestingly, including the issuance factor in the regression does not reduce the alpha of the intensive margin portfolio to zero. And, we find this is true even if we use issuance factors built with different definitions or based on different horizons (Daniel and Titman, 2006; Pontiff and Woodgate, 2008; Fama and French, 2008a). We believe this is because, while these measures do capture return predictability associated with issuance, they do not fully capture the issuance effect specifically relevant for value-weighted market investors. We return to this point in Section 5.2.3.

Collectively, the results in columns 1–3 point to the intensive margin rebalancing portfolio as systematically timing its trades in a way that leads to low risk-adjusted returns as well as tilting away from factors with high average returns. That is, the stocks in the portfolio and/or the timing of these trades are a cost to index investors.

Columns 4–6 replicate the results in columns 1–3, but with our index scaling portfolio. Here, the constant terms are insignificant and near zero. This is additional evidence that a mechanical price pressure story is unlikely to be driving our intensive margin rebalancing results. The logic is that any buying by passive funds might push up prices today (i.e., transitory price pressure) and therefore lead to low returns the following quarter (and vice versa for selling by passive funds). This price-pressure effect, however, should be true for any type of fund trading, whether it is flow-based or rebalancing. And, a zero alpha, seen in columns 4–6, is inconsistent with a pure price pressure story. The lack of evidence for the price pressure story in Table 2 is also consistent with our estimates of returns from the hypothetical indexes we form in Section 5. There, we find that trading costs, measured as expected transitory price impact, are an order of magnitude smaller than return predictability from issuance-related market timing.

For completeness, columns 7–9 replicate the results in columns 1–3, but with our index additions and deletion portfolio (i.e., the extensive margin). Here, all of the constant terms are negative and significant, which is consistent with results in, e.g., Wurgler and Zhuravskaya (2002) that returns are low after index additions and high after index deletions.[13]

Finally, columns 10–12 report the returns to the total rebalancing portfolio, which combines both the intensive and extensive margin rebalancing portfolios. The estimated constant term in column 10 is −4.6%. After adding in all the asset pricing factors, the alpha becomes −4.0%, which is the headline figure we present in the abstract. Note that these estimates are effectively a weighted average of the results in columns 1–3 and 7–9. Because the extensive margin portfolio represents a smaller dollar share of rebalancing activity, the total rebalancing portfolio estimates are closer to the intensive margin alphas in columns 1–3.

[11] Initially, the sample starting in 1996 seems late, as the well-known index ETF SPY was launched in 1993, and the Vanguard 500 index fund existed well before that. As discussed in Appendix B.3, these pre-1996 observations are removed by the data quality filters described in Section 2.

[12] In this table, we use an issuance factor from Chen and Zimmermann (2022) based on forming quintiles of net issuance over the past year following Pontiff and Woodgate (2008). Appendix D.6 shows that results are similar using issuance over the past 5 years following Daniel and Titman (2006) or over the past year with a one year delay following Fama and French (2008a).

[13] Note that, for funds like the CRSP total market fund, index deletions are rare for stocks that do not delist. Such deletions typically occur only when a stock no longer meets eligibility criteria, such as changes in its country of domicile following a tax inversion. Also, index funds have no additions or deletions in some quarters. This is why the number of observations in columns 7–9 are smaller than the number of observations in all other columns.

Table 2
Rebalancing portfolio returns.

	Intensive margin rebalancing			Scaling			Extensive margin rebalancing			Total rebalancing		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Market		0.0781***	0.119***		0.0058	−0.0197		0.0568	−0.0226		0.0695***	0.038
		(0.021)	(0.028)		(0.010)	(0.012)		(0.038)	(0.046)		(0.022)	(0.027)
Size			−0.161***			0.0603**			0.382***			0.131**
			(0.046)			(0.030)			(0.096)			(0.058)
Value			−0.116***			−0.0171			−0.119			−0.111**
			(0.043)			(0.021)			(0.073)			(0.043)
Profitability			−0.169***			0.0499*			−0.0245			−0.188***
			(0.057)			(0.027)			(0.099)			(0.059)
Investment			−0.117*			0.0151			−0.415***			−0.192***
			(0.060)			(0.024)			(0.098)			(0.056)
Momentum			0.0252			−0.0357**			0.171***			0.0552*
			(0.030)			(0.017)			(0.051)			(0.030)
Reversal			−0.154***			0.0057			0.0769			0.0016
			(0.039)			(0.020)			(0.070)			(0.042)
Issuance			−0.146***			0.0066			0.0385			0.0413
			(0.052)			(0.025)			(0.096)			(0.061)
Alpha	−4.149***	−4.950***	−3.202***	0.363	0.304	0.295	−5.635***	−6.217***	−4.889***	−4.612***	−5.326***	−3.974***
	(0.638)	(0.671)	(0.701)	(0.300)	(0.313)	(0.340)	(1.123)	(1.165)	(1.299)	(0.642)	(0.645)	(0.669)
Observations	2339	2339	2339	2339	2339	2339	2301	2301	2301	2344	2344	2344
R-Squared	0.000	0.008	0.101	0.000	0.000	0.012	0.000	0.001	0.048	0.000	0.006	0.066

Notes. This table presents regressions of our rebalancing and scaling portfolios' returns on asset pricing factors for value-weighted index funds tracking the S&P 500, 400, and 600, the Russell 1000, 2000, and 3000, and the CRSP market-capitalization based indexes. The intensive margin captures rebalancing activity for existing positions. The scaling portfolio is based on how the index fund is predicted to scale existing holdings up and down based on flows. The extensive margin portfolio is long/short stocks added/dropped by the index fund. The total rebalancing portfolio is composed of all intensive and extensive margin rebalancing trades. Portfolios are constructed in quarter t to examine returns in $t+1$. We use quarterly index fund holdings data from 1996 Q1 to 2023 Q3 to construct portfolios. The unit of observation is at the fund-quarter level. All returns are in percent and annualized. Robust standard errors are in parenthesis.

The results in Table 2 support the idea that the rule-driven value-weighted index fund rebalancing implicitly leads to timing the market, largely in response to firm changes in shares outstanding. In Table A.1 in Appendix D.2 shows that these results apply more generally with a sample of all index funds, some of which we know are not value-weighted. Specifically, we find that in the expanded sample of all index funds, the intensive margin rebalancing portfolios still have negative average returns and alphas.

Although the rebalancing returns and alphas are economically large, they represent the returns to only the rebalanced portion of the fund. The degree to which these returns affect fund-level returns is a function of how large the rebalancing portfolio is relative to AUM. We do a back-of-the-envelope calculation on the approximate size of the rebalancing portfolio by examining the expected rebalancing trade size from our ownership ratio framework.

For the Vanguard Total Market Index Fund (Ticker: VTI), the index fund which most closely resembles the total-market index fund described by our ownership ratio framework, the median total size of the rebalancing portfolios is between 10% and 20% of AUM. This figure has been slowly declining in recent years to around 10% to 15% of AUM. In contrast, most of the other index funds in our sample have larger rebalancing trades as a fraction of AUM. If we take a conservative range of the size of the rebalancing portfolio to be around 10% to 15% of AUM annually, the 4.61% average total rebalancing portfolio return implies a performance drag of 46 bps to 69 bps at the *fund* level. These figures align with the market timing benefit shown in Section 5, which stems from our hypothetical indexes designed to eliminate the costly rebalancing we observe in real-world index funds. As a sense of scale, the often scrutinized expense ratio of these index funds is mostly between 3 bps and 9 bps per year. That is, our analysis suggests that index design choices on how and when to update index-eligible shares outstanding amount to index performance drag that is an order of magnitude greater than expense ratios.

5. Designing a better index

As outlined in Sections 3.1 and 3.3, the only discretion an index has is how and when to adjust shares outstanding (by updating index-eligible shares). Designing a value-weighted index is defined by two choices: when to rebalance (i.e., rebalancing frequency), and how long to delay incorporating shares outstanding to update index-eligible shares (i.e., how stale the shares outstanding used is). The first choice, *rebalancing frequency*, determines how often the index adds and drops stocks and updates index-eligible shares outstanding. More frequent rebalancing leads to smaller trades per event and less market tracking error, but may increase trading costs and cause the fund to more quickly take the opposite side of firm issuance.

The second choice, *shares outstanding staleness*, governs the minimum delay between a change in a firm's shares outstanding and when that change is reflected in index-eligible shares. Using the most recent data is appealing from the perspective of minimizing market tracking error. But, this leads index funds to take the other side of endogenous compositional changes by firms that are known to forecast returns, even over long horizons (Daniel and Titman, 2006; Pontiff and Woodgate, 2008; Fama and French, 2008a; Baba-Yara et al., 2024). Introducing a delay, therefore, can reduce exposure to firm-driven market timing. In addition, if price impact is correlated with firm activity, then the choice of shares outstanding staleness will also impact trading costs.

In this section, we analyze these choices by constructing hypothetical indexes using a variety of different rebalancing frequencies and shares outstanding staleness. We then decompose the benefits of better-performing indexes into the contribution from economizing on trading costs and better market timing.

5.1. Index design and performance

In this subsection, we describe the construction of our hypothetical value-weighted total market indexes, how we measure the performance of an index (and the index fund that tracks it), and provide performance results.

5.1.1. Index design

Our hypothetical indexes will vary based on how frequently they rebalance and the staleness or delay of the shares outstanding used to update index-eligible shares. Given the framework in Section 3, these are the two key decisions for *every* value-weighted index fund.

Rebalancing frequency governs when rebalancing trades by the hypothetical index funds occur, e.g., at the end of each month, quarter, year, etc. The shares outstanding delay determines what information is used at each rebalance, e.g., shares outstanding data that is stale by at least a day, a month, a year, etc. before it can be incorporated into index-eligible shares. Appendix C provides several examples designed to illustrate how rebalancing frequency and shares outstanding delay are distinct dimensions of index design. In this subsection, we show how systematically varying these two parameters affects index fund performance. Later, in Section 5.2, we decompose these additional returns into reduced trading costs and market timing benefits.

5.1.2. Measuring performance

As outlined in Section 4.1, the market information ratio (MIR) is designed to capture the tradeoffs between better returns (from improved market timing and lower trading costs) and deviations from the market portfolio. In this section, we modify Eq. (6) to account for estimated trading costs. That is, $r_{\text{idx},t}$ is the return of the index fund *net of expected trading costs*. Importantly, the "benchmark return", i.e., $r_{\text{mkt},t}$ does not include trading costs, so any trading costs incurred by our index funds will directly negatively affect MIR.

Another measure of performance we use is simply the net-of-trading cost index return in excess of the market, which captures how index design affects both trading costs and market timing. However, this measure does not capture the possible costs of deviating from the total stock market portfolio. That is, an index may have high returns but, despite still operating as a value-weighted index by following the ownership ratio framework, may deviate substantially from the total stock market portfolio.

Trading costs. While most real-world index funds appear to trade at zero cost, this is largely an artifact of how index-linked trades are executed. Index funds typically rebalance by trading in the closing auction, where standard measures of price impact (such as the effective or realized bid–ask spread) do not apply (Chinco and Sammon, 2024). Furthermore, focusing only on price changes around the closing auction ignores the anticipatory trading that occurs days, weeks, or even months before the index rebalancing event. Intermediaries typically pre-position themselves to provide liquidity to index funds ahead of the rebalance, especially when the expected trade is predictable (Greenwood and Sammon, 2025; Sammon and Murray, 2024). For example, in the case of a large index addition, this type of "front running" would push up prices *ahead* of the announcement and the inclusion date, effectively leading the index fund to buy at a high price.

These front-running costs are not captured by standard index tracking error, because the index itself adds the stock at the closing auction price, i.e., after the price impact has already occurred. However, front running still affects the returns earned by the fund relative to the market as a whole because this price impact is transitory. We view this anticipatory price movement as a type of transaction cost and, in this subsection, we use two standard models of price impact to quantify these trading costs as a function of rebalancing frequency.

The first is a model where percent price impact depends on trade size relative to average daily volume (ADV) scaled by volatility (Almgren and Chriss, 2001; Frazzini et al., 2018):

$$PI_{i,t}^{ADV} = \beta \cdot \sigma_{i,t} \cdot \left(\frac{\text{Trade Size}_{i,t}}{\text{ADV}_{i,t}} \right)^{0.5} \tag{8}$$

where β is a price impact coefficient. Empirically, we calculate $\sigma_{i,t}$ as the standard deviation of stock i's daily return over the 22 trading days before the index rebalancing event, and we calculate $\text{ADV}_{i,t}$ as the average daily trading volume in stock i over the same period. To reduce the influence of outliers driven by stocks with very low trading volume, we cap $PI_{i,t}^{ADV}$ at 50% (which is well over the 99.5^{th} percentile of the distribution of estimated $PI_{i,t}^{ADV}$ in our sample). β is set to 0.2 based on the trading cost model estimates from Frazzini et al. (2018).[14]

The second is a square-root model, in the spirit of Almgren et al. (2005) and Gabaix et al. (2006), where percent price impact is proportional to the square root of trade size in dollars:

$$PI_{i,t}^{sqrt} = \lambda_{i,t} \cdot \sqrt{\text{DollarVolume}_{i,t}} \tag{9}$$

where $\lambda_{i,t}$ is a stock-time specific coefficient of price impact. Empirically, we use the WRDS millisecond intraday indicator suite's "Lambda (Price Impact Coefficient)" estimate and calculate a moving average over the 22 trading days before the index rebalancing event. We use $PI_{i,t}^{ADV}$ as our main measure because we can get estimates for our full sample. We provide estimates separately for $PI_{i,t}^{sqrt}$ in Appendix D.4, which are limited to only when TAQ data is available. We find largely similar results from both models.

Applying these models to the expected trades of index funds is motivated by the idea that, in a world without pre-positioning, we would observe, at minimum, the same price impact on the day of the rebalance as we currently see (cumulatively) in the pre-rebalancing period because intermediaries are "pre-positioning" the shares specifically for index funds. This assumption is particularly relevant for the hypothetical index funds we study, which are simulated and, for the most part, not subject to any realized price impact resulting from pre-positioning.[15]

5.1.3. Index performance

Excess index returns. We compute returns for each hypothetical index that represents each combination of rebalancing frequency and shares outstanding delay. That is, we calculate exactly how each hypothetical index would rebalance and what portfolio it would end up holding given how shares outstanding evolved up to that point. To get the final index return, we subtract trading costs estimated from the price impact models above.

When we use the price impact models to calculate price impact, we use the ownership ratio framework to predict the size of the exact rebalancing trades of each hypothetical fund at each rebalancing event. We treat this as the trade size and use the price impact model with historical liquidity parameters from the rebalancing date to estimate price impact for each trade.[16] We combine the estimated price impact for each stock based on its contribution to the portfolio and subtract the weighted sum from the hypothetical index return to get the final, net-of-trading-cost index return that we use in the numerator of the MIR from Eq. (6).

Importantly, in Eqs. (8) and (9), the size of the trade – which is a function of the size of the fund – matters for price impact. We assume that each hypothetical index fund owns 10% of total market capitalization, a large but plausible number for the passive ownership

[14] Before 1982 (and for most of 1982 itself), daily trading volume is missing in CRSP for Nasdaq-listed stocks. This creates a problem for applying Eq. (8), as it requires calculating average historical trading volume for each stock. For Nasdaq stocks before 1983, we first calculate the median trading cost for NYSE stocks using Eq. (8) and apply that to all Nasdaq listings.

[15] Note, however, that our hypothetical index funds typically trade in the same direction as real-world index funds, e.g., buying after issuance and selling after buybacks. As a result, if we were to rebalance them on a similar schedule as real-world index funds (e.g., at the end of each quarter), some front-running effects might still be present in the observed prices. Because we subtract expected price impact costs from returns for every hypothetical index, we may double count price impact costs for our hypothetical indexes that coincide with when actual index funds trade (i.e., quarterly rebalanced indexes with no shares outstanding delay).

[16] Importantly, we exclude expected trading costs associated with cash events, e.g., inflows, outflows, investment and payment of dividends and reinvestment of delisting proceeds. Instead, we focus on quantifying the trading costs associated purely with index funds' *rebalancing* activities.

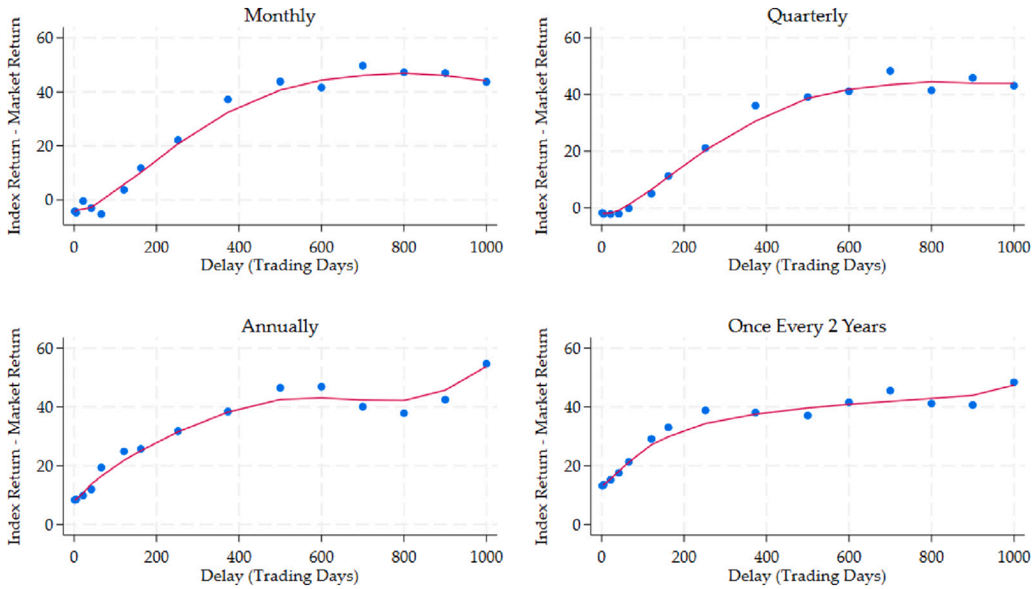

Fig. 1. Excess index return.

Notes. This figure plots the return of each hypothetical index fund in excess of the daily total stock market return as a function of index design choices. Those choices are the rebalancing frequency and shares outstanding staleness (or delay). The index return is net of trading costs. Trading costs are computed using rebalancing trades that would have occurred at each rebalance and a price impact model. The hypothetical index funds use shares outstanding and return data for each stock from 1977 to 2023.

industry as a whole. This may overstate trading costs in the earlier part of our sample and understate them in the most recent years. However, assuming the fund is a constant size relative to total index capitalization prevents mechanically introducing a time series trend in trading costs. In addition, it avoids the effect of cumulative benefits compounding and skewing the results over time.

The first measure of performance we examine is index returns in excess of the daily market return (i.e., the numerator in the MIR). Fig. 1 plots this excess return for each hypothetical index fund from 1977 to 2023. Each of the four panels holds fixed the rebalancing frequency (monthly, quarterly, annually, and biennially), and the shares outstanding delay of the hypothetical index fund is given on the x-axis. There are several notable patterns. First, longer shares outstanding delays, regardless of rebalancing frequency, increase returns, with a magnitude of about 40 bps to 50 bps per year for a two-year delay (~500 trading days). This magnitude is similar to the back-of-the-envelope fund-level performance drag estimated at the end of Section 4.2. We also note that the excess return is slightly negative for the monthly and quarterly rebalancing with little to no shares outstanding delay. This comes from the fact that price impact costs are concave — more frequent rebalancing leads to landing on a relatively steep part of the price impact curve. The addition of a delay with these higher-frequency rebalancing policies offsets some of these trading costs, likely with benefits from better market timing.

Market tracking error and market information ratio. Our second measure of index performance uses MIR, or index excess returns per unit of market tracking error. Before examining MIR as a function of index design choices, we show how these choices impact market tracking error. Fig. 2 shows the results. Unsurprisingly, both less frequent rebalancing and more delayed shares outstanding increase market tracking error. This is natural — rebalancing less often and using more stale shares outstanding means that an index is less responsive to compositional changes in the stock market, and the delay in responsiveness means the index's portfolio looks different from the market portfolio.

While using more stale shares outstanding improves index excess returns, it also leads to greater MTE. If an investor, motivated by long-standing financial advice and academic research, wants to track the market portfolio, it seems sensible to consider a measure of performance such as MIR, which measures the return benefits of index design against the costs of deviating from the market. Fig. 3 presents MIR as a function of rebalancing frequency and shares outstanding delay. In essence, MIR takes the returns from Fig. 1 and divides by the MTE of Fig. 2. The MIR results show that there is an interior maximum: some shares outstanding delay improves MIR but, eventually, a long enough delay causes the MIR to decline. This is true for every rebalancing frequency, and interior maximums are roughly found with longer delays when the rebalancing frequency is shorter. For example, the monthly rebalancing has a maximum MIR around a two-year shares outstanding delay, while the biennial rebalancing has a maximum MIR at around a one-year delay. This is because less frequent rebalancing automatically builds in a longer delay — even if an index has little to no shares outstanding delay, it will still, on average, wait a long time for the next rebalancing, automatically introducing some staleness.

5.2. Decomposing excess index returns

To better understand why both lower rebalancing frequency and/or greater shares outstanding staleness improve excess index returns and the market information ratio, it is helpful to understand how much of the improvement comes from economizing on trading costs vs. better market timing, and the degree to which they are positively or negatively correlated. We first examine the trading cost component of our hypothetical index returns, then study the contribution from market timing. Finally, we examine if the index return can be dynamically replicated by holding a combination of a total market index fund and an issuance factor.

5.2.1. Trading costs

We separately report the price impact costs for each hypothetical index fund, as outlined above in Section 5.1.2. Fig. 4 reports our estimates of trading costs as a fraction of AUM across all of the hypothetical index funds we examine above. The figure uses estimates from the ADV-based model from Eq. (8), which allows us to compute trading costs over the full sample period from 1977 to 2023. We present parallel estimates using the model based on estimated price impact coefficients in Eq.

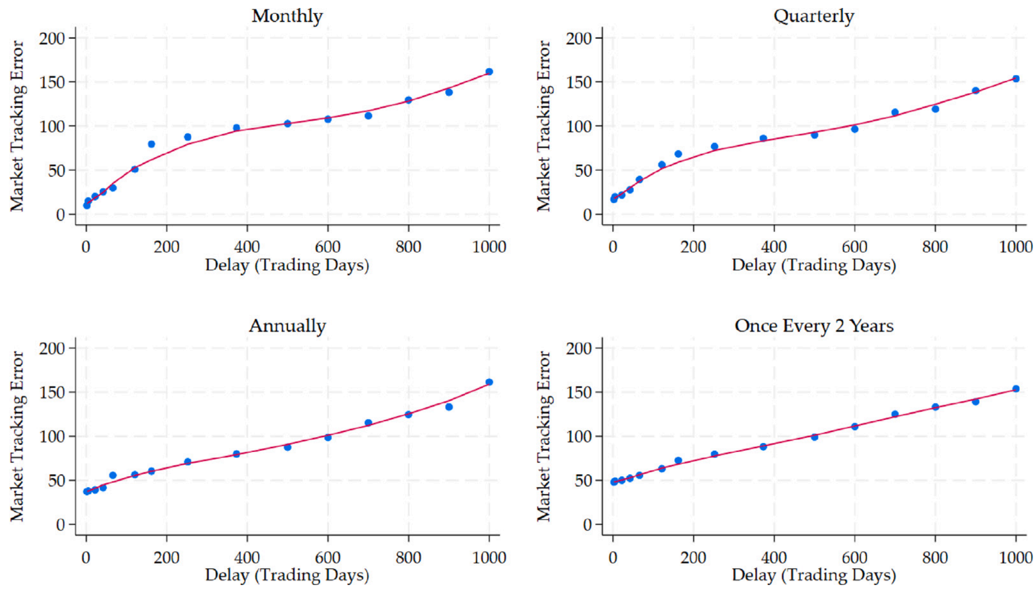

Fig. 2. Market tracking error.
Notes. This figure plots our measure of market tracking error (MTE) of each hypothetical index fund as a function of index design choices. Those choices are the rebalancing frequency and shares outstanding delay. MTE is the standard deviation of the hypothetical index return (net of trading costs) minus the daily total market return. The hypothetical index funds are constructed based on shares outstanding and return data for each stock from 1977 to 2023.

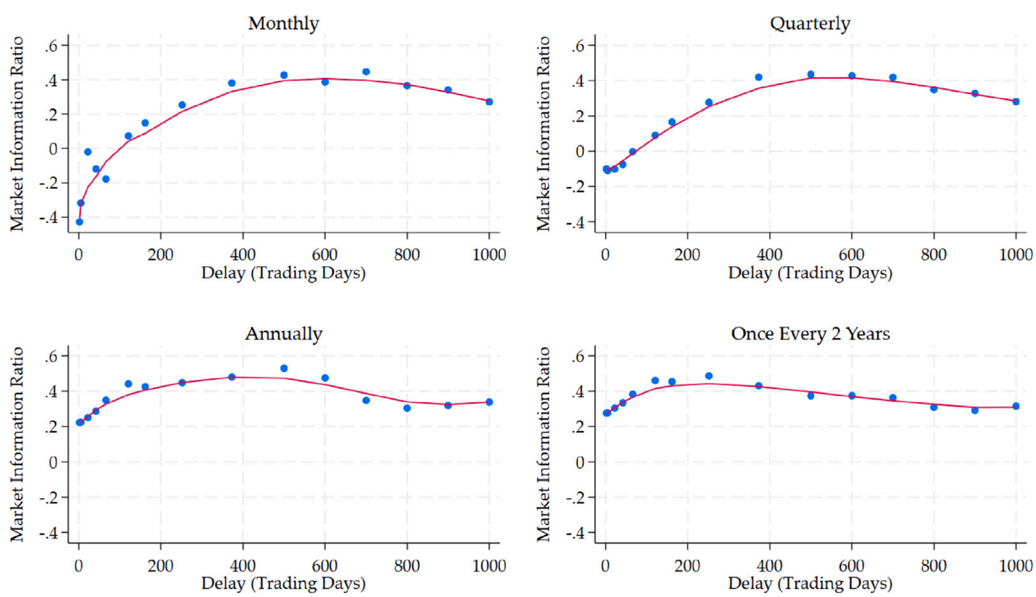

Fig. 3. Market information ratio.
Notes. This figure plots our measure of the market information ratio (MIR) of each hypothetical index fund as a function of index design choices. Those choices are the rebalancing frequency and shares outstanding delay. MIR is the ratio between the hypothetical index fund excess return and its market tracking error (MTE). The index fund excess return is the hypothetical fund return net of trading costs minus the daily total market return. MTE is the standard deviation of the index fund excess return. The hypothetical index funds are constructed based on shares outstanding and return data for each stock from 1977 to 2023.

(9) in Appendix D.4, and find all our main conclusions are essentially unchanged.[17]

There are three main takeaways from the figure. First, trading costs are generally modest when scaled by AUM, ranging from 4 to 9 basis points per year. This is consistent with the findings from Madhavan et al. (2022). To be clear, trading costs are significant as a dollar amount, but are modest only relative to AUM. As a reminder, here we set AUM to be 10% of the overall stock market capitalization. This value is still the same order of magnitude as index fund expense ratios, which suggests that carefully selecting rebalancing rules only for the purpose of economizing on trading costs is at least as important as the expense ratios that investors closely track. Second, less frequent rebalancing (e.g., moving from monthly to annually) reduces trading costs, even though total trading volume per year is similar across strategies.[18] This

[17] We are limited to a shorter sample period for the second price impact model, which is why we use the ADV model for the main results.

[18] Delaying rebalancing does not reduce total trading over the long run, except in cases where a firm is listed and subsequently delisted before it would have entered the index.

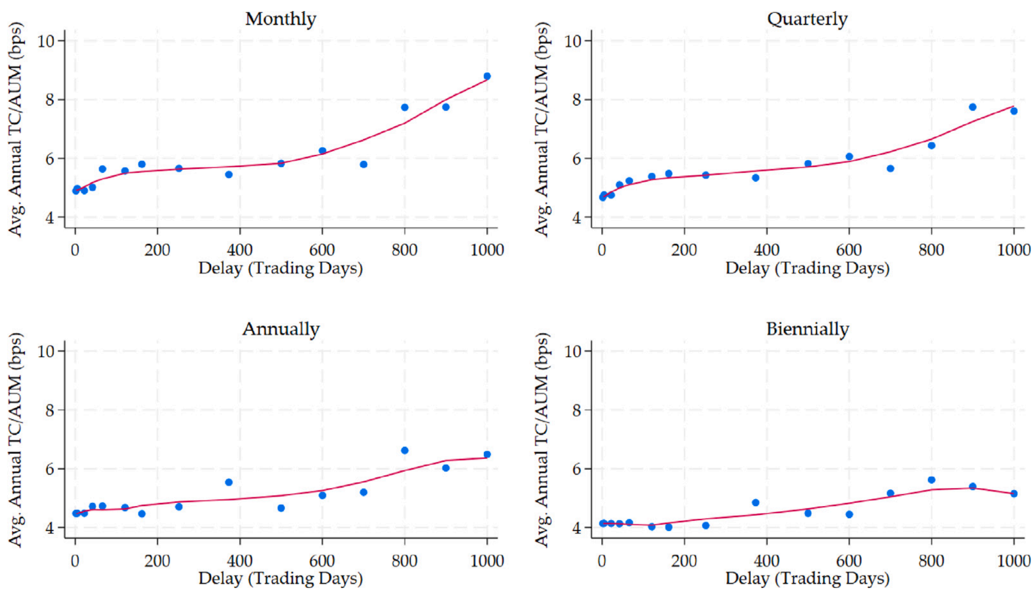

Fig. 4. Estimated trading costs by rebalancing frequency and shares outstanding delay.
Notes. This figure reports the average annual trading costs as a fraction of AUM for hypothetical index funds under different rebalancing frequencies and shares outstanding staleness. Trading costs are estimated using the ADV-based model in Eq. (8), applied over the full sample from 1977 to 2023.

is primarily because trading costs in Eq. (8) are concave in trade size, so trading less frequently in larger quantities can reduce the overall total cost of execution. Third, a larger shares outstanding delay can increase trading costs. One explanation is that, around issuance and repurchase events, trading volume may be elevated, and therefore the trade size as a fraction of ADV is lower.

5.2.2. Market timing

As shown in Section 3, a value-weighted index must trade in response to changes in the composition of the stock market. The choice of *shares outstanding staleness* determines how quickly to respond to firm activity like IPOs, equity issuance, and buybacks.

Most real-world index funds rebalance quarterly and update index-eligible shares outstanding with the most recently available data, i.e., they use shares outstanding with no staleness. As documented in Section 4.2, this relatively rapid response to compositional changes leads to poor returns. But, because these trades by index funds are also predictable, it is somewhat difficult to fully separate market timing from price impact costs. In other words, front-running induced price pressure ahead of predictable rebalancing in response to compositional changes could explain part of the poor performance of the index fund rebalancing portfolios, rather than market timing per se.

To isolate the market timing component, we again construct hypothetical index funds following the ownership ratio logic, except we instead allow funds to rebalance at a *daily* frequency while delaying the incorporation of shares outstanding into index-eligible shares by some fixed amount of time. Fig. A.4 in Appendix C provides an illustrative example. The combination of daily rebalancing and a constant delay in shares outstanding avoids, for the most part, forcing our hypothetical index funds to trade at the exact same time as real-world index funds (e.g., the end of a quarter), which helps to mitigate the effects of front running by avoiding the "crowded" times of index-induced trading and isolate the market-timing component of returns. Further, by imposing a fixed lag (as opposed to a calendar-based lag), it allows us to hold constant the time horizon over which the market can adjust prices in the face of firm activity affecting shares outstanding.[19]

We find that hypothetical index fund returns are monotonically increasing in the shares outstanding delay. Fig. 5 plots the average market timing benefit, measured as the hypothetical index return in excess of the daily stock market return. This index return, because it is meant to isolate only the market timing benefit, does *not* incorporate trading costs. The relationship between shares outstanding staleness and the market timing benefit is also concave: the benefits grow steadily from 0 to 600 trading days (a bit more than two calendar years) and begin to plateau at approximately 70 bps per year beyond that. In Appendix D.3, we show that the market timing benefits within each decade still roughly plateau at around a 600 day delay, although the level at which the benefits plateau varies substantially, ranging from 30 bps up to 140 bps.

5.2.3. Replicating index returns

The ownership ratio framework in Section 3 highlights that the only reason a value-weighted index fund needs to rebalance is to account for compositional changes through issuance, buybacks, and IPOs. Therefore, one way to interpret our hypothetical index funds is that they are some combination of a total market index fund and an issuance factor, where the issuance factor is long stocks with below average net issuance (or have done net buybacks) and short stocks with above average net issuance over the horizon of the rebalancing delay. And, having a positive weight in this issuance factor allows an investor to offset the market portfolio's tendency to have greater exposure to recently issued equity and less exposure to stocks that experienced buybacks.

A natural approach to mimic the delayed rebalancing hypothetical index return is to create an optimal combination of the market and an issuance factor. After all, our delayed rebalancing strategy is directly related to cross-sectional differences in issuance, but still largely holds something similar to the market. We do this in two ways. First, we identify the ex-post optimal combination of the market and issuance factors. Second, we construct an implementable ex-ante trading strategy that computes the weight in the market and issuance factors for each year based on optimizing over the past five years.

For the ex-post optimal portfolio, we construct the optimal without a risk-free asset. The reason for this is that the two-risky-asset tangency portfolio can be very sensitive to means and covariances, which can lead to extreme portfolio weights. We also assume a risk aversion of 5 for our exercise. This assumption has the benefit of reducing extreme

[19] If, for example, we used shares outstanding from the end of the last quarter, the staleness could vary across stocks depending on when their shares outstanding changed within the previous quarter.

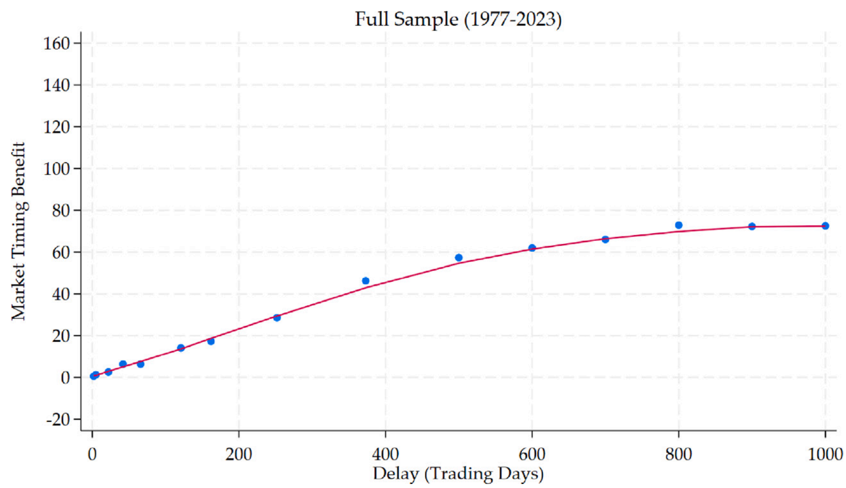

Fig. 5. Market timing.
Notes. This figure shows the market-timing benefit of delaying the response to changes in shares outstanding. The market timing benefit is defined as the hypothetical index return minus the daily total market return. The hypothetical indexes considered here rebalance daily but vary the delay in responding to changes in shares outstanding. The sample spans 1977–2023.

values in the portfolio weights, which is especially useful when moving to the dynamic, ex-ante trading strategy. Just as above, we use the Fama–French market factor (which is rebalanced quarterly) and the 1-year issuance factor of Pontiff and Woodgate (2008).

We find that the optimal in-sample portfolio has weights of 0.21 on the issuance factor and 0.79 on the market. This yields a Sharpe ratio of 0.83. In comparison, our 500-day delay index with quarterly rebalancing has a Sharpe ratio of 0.76. Of course, the optimal issuance/market combination portfolio is formed on an ex-post basis while our 500-day delay index is based on a simple tweak to value-weighted index rebalancing — the optimal portfolio *should* have a higher Sharpe ratio.

We also estimate the optimal ex-ante portfolio on a rolling basis, similar to a trading strategy that an investor could actually implement in real time. We construct weights in the market and issuance factor for each year by computing the optimal portfolio from the past 5 years of data. We find that this strategy has a Sharpe ratio of 0.785, only slightly higher than our 500-day delay and quarterly rebalance index's Sharpe ratio of 0.76.

While the Sharpe ratios are quite similar, the trading strategy achieves its returns in a fundamentally different way from our delayed rebalancing index. First, the average return is much lower with the trading strategy — 10.35% per year vs. 13.20% for our hypothetical index. This makes the barriers to entry for many value-weighted index fund investors somewhat high. That is, to achieve the same average return as a value-weighted market investor, the trading strategy investor must lever up the optimal issuance-market factor portfolio. Given that there is a large retail segment of value-weighted index fund investors, the costs and complexity of implementing such a strategy may be infeasible.

Second, our delayed index still follows our ownership ratio framework, and therefore can be viewed as a (slightly modified) value-weighted market index. Something that makes this clear is the correlation of each strategy with the market. The optimal issuance-market portfolio has a return correlation of 0.9504 with our daily-rebalanced market. The 500-day delay index that is rebalanced quarterly has a return correlation of 0.9986 with the market. This tighter correlation with the market also implies that the 500-day delay index has a higher market information ratio via a low market tracking error. So, if the goal is to take a more passive stance and invest in a value-weighted total market index fund, just as academics and financial advisors often recommend, the delayed rebalancing fund adheres more closely to that guidance. Said differently, our hypothetical indexes can nearly mimic the performance of a more sophisticated factor strategy with a simple tweak to a value-weighted index fund's rebalancing policy.

Finally, we also document factor exposures and alphas of our delayed rebalancing portfolio returns in Appendix D.5. Figure A.8 shows that longer shares outstanding delays lead to nearly monotonically increasing loadings on factors known to predict returns, including profitability, investment, and issuance. However, controlling for these factors does not fully eliminate the delayed rebalancing portfolios' alpha — the one-year and two-year delay portfolios still achieve a fund alpha of around 40 bps to 50 bps per year. We view this as additional evidence that combining the market with a dynamic factor trading strategy cannot fully replicate the benefits of delayed rebalancing.

5.3. Discussion

A natural critique of our results is that less frequent rebalancing constitutes a form of "active" management. We believe this concern is misplaced, as our hypothetical index funds follow the same ownership ratio framework that defines what it means to be a value-weighted index fund. To clarify this point, we start by defining what we mean by "passive" management.

We define a passive strategy as being governed by a transparent and systematic set of rules for security selection and rebalancing. For example, the Russell U.S. Indexes follow a clearly articulated methodology, detailed in a publicly available document. While this definition includes value-weighted indexes, it is also broad enough to capture more unusual indexes that, e.g., do not implement value-weighting schemes. An example is COWZ, which is an ETF that tracks an index of the top 100 companies in the Russell 1000 in terms of free cash flow, and weights these 100 stocks based on trailing 12-month free cash flow yield (see more details on the ETF website).

While active funds could, in principle, adopt transparent and rule-based rebalancing strategies, nearly all active funds do not make the details of their strategy public. And further, active strategies often incorporate proprietary signals or discretionary elements. We consider any fund or index that does not use a transparent, rule-based rebalancing policy as actively managed.

Based on this definition, our hypothetical rebalancing strategies would squarely fall into the passive category: they follow a systematic set of rules (the ownership ratio framework) which can be transparently implemented with public information (data on firms' shares outstanding). That is, they operate just like existing value-weighted indexes except they use a stale value of shares outstanding instead of the most current value.

Another concern is that, even though our sleepy rebalancing strategy could be implemented in real time with no look-ahead bias, our results are all from in-sample tests. That is, we cannot say that the benefits of delayed rebalancing will continue going forward or that the benefits of delayed rebalancing would persist if the mass of index investors grows considerably. However, we have two main reasons to believe that the patterns we document could continue to hold out of sample.

First, there are compelling reasons why firms prefer to issue more shares when prices are relatively high and buy back shares when prices are low (Myers and Majluf, 1984). Said differently, we believe that it is unlikely that firms will suddenly start issuing shares when prices are relatively *low*, leading to higher returns in the future (and vice versa for high prices and buybacks). This implies that, while it is possible that the relationship between issuance and future returns goes to zero (and the additional returns from our delayed rebalancing strategy also diminish), it would be surprising if they went negative.

The second reason is that the relation between issuance and future returns has been widely known since the mid-1990s (Loughran and Ritter, 1995), and yet the issuance effect has persisted to this day. So, while the concern for any empirical pattern is that it will not hold out of sample (McLean et al., 2020), we now have had more than two decades of data that supports the idea that these returns may come from something more persistent that, for whatever reason, does not get arbitraged away.

Related to this, one might wonder whether the effects we document would persist if our delayed-rebalancing strategy were widely adopted. Currently, index funds own about 15% to 20% of the US stock market. And, if all of these index funds switched to a delayed rebalancing policy after issuance, other (presumably more elastic) investors would have to temporarily absorb an additional 20% of each issuance, lowering prices and discouraging issuance in the first place. In the extreme case where all investors held only index funds which delay responses to compositional changes, newly issued shares would lack natural buyers, underscoring the important role of non-index investors in pricing new securities (Pedersen, 2018). While this is not the focus of our paper, it does suggest a natural limit: if widespread adoption of delayed rebalancing discouraged issuance, the effect of sleepy rebalancing on returns would diminish, as value-weighted index funds *only* rebalance in response to cross-sectional differences in issuance.

6. Conclusion

In this paper, we develop a simple ownership ratio framework to explain how a value-weighted index fund works: it owns the same fraction of each constituent's shares outstanding, and this fraction is equal to the ratio of fund AUM to the total market capitalization of the constituent stocks. This perspective clarifies that value-weighted index funds rebalance solely due to changes in the composition of the market — stocks entering or exiting the public market, or changes in shares outstanding for existing companies. Our ownership ratio framework also highlights that an index provider has limited flexibility in designing a value-weighted index, which can only vary along two dimensions: how often the index rebalances and the timeliness of the shares outstanding used during each rebalance (i.e., shares outstanding staleness or delay).

The ownership ratio framework also shows that index funds mechanically rebalance towards stocks that issued relatively more shares and away from stocks that bought back more shares. Because firms tend to be informed "traders" in their own shares (Loughran and Ritter, 1995; Ikenberry et al., 1995; Daniel and Titman, 2006; Pontiff and Woodgate, 2008; Fama and French, 2008a) and the return predictability associated with this informed trading is long-lived (Daniel and Titman, 2006; Baba-Yara et al., 2024), index funds that relatively quickly take the other side of firm issuance and repurchase activity have low returns from poor market timing.

We estimate that the long-short rebalancing portfolios of real-world value-weighted index funds have an annual alpha of −4.0%. In addition, the rebalancing portfolios' returns load negatively on several asset pricing factors known to predict returns, including value, profitability, and investment. In other words, the trades required to track the value-weighted market portfolio implicitly subject index fund investors to factor exposure. Importantly, although the *only* reason index funds rebalance is due to cross-sectional differences in issuance, conditioning on existing issuance factors does not fully eliminate the negative alpha. Thus, we believe these rebalancing portfolios identify the specific issuance effect value-weighted index fund investors are exposed to.

We use these insights to explore alternative index designs that can avoid poor performance but still respect the ownership ratio framework. Specifically, we use a set of hypothetical indexes to explore the costs and benefits of various combinations of rebalancing frequency and shares outstanding staleness (or delay). We find that hypothetical index returns increase with less frequent rebalancing and more stale shares outstanding, with benefits as high as 50 bps per year in excess of the daily rebalanced total market return. An index with quarterly rebalancing and a delay of around two years maximizes the market information ratio, or the ratio of excess returns to tracking error with the daily market. Importantly, these benefits are not based on an ad hoc rule — rebalancing frequency and shares outstanding delay are parameters that our ownership ratio framework shows *must* be specified by any value-weighted index.

Our paper demonstrates that index fund rebalancing is driven by market compositional changes and implicitly subjects value-weighted index fund investors to a market timing strategy. That is, index investors are exposed to a specific issuance-related factor that is not completely explained by known issuance factors as well as other asset pricing factors. And, a simple modification to index design – delaying how quickly an index responds to compositional changes – would have produced savings to index investors an order of magnitude greater than index fund expense ratios.

CRediT authorship contribution statement

Marco Sammon: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **John J. Shim:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors have nothing to disclose.

Appendix A. Supplementary data

The online appendix can be found on the authors' websites.

Data availability

Replication Package for "Index Rebalancing and Stock Market Composition: Do Indexes Time the Market?" (Original data) (Mendeley Data)

References

Almgren, R., Chriss, N., 2001. Optimal execution of portfolio transactions. J. Risk 3, 5–40.

Almgren, R., Thum, C., Hauptmann, E., Li, H., 2005. Direct estimation of equity market impact. Risk 18 (7), 58–62.

Appel, I.R., Gormley, T.A., Keim, D.B., 2016. Passive investors, not passive owners. J. Financial Econ. 121 (1), 111–141.

Arnott, R.D., Brightman, C., Kalesnik, V., Wu, L., 2023. Earning alpha by avoiding the index rebalancing crowd. Financ. Anal. J. 79 (2), 76–97.

Baba-Yara, F., Boons, M., Tamoni, A., 2024. Persistent and transitory components of firm characteristics: Implications for asset pricing. J. Financial Econ. 154, 103808.
Baker, M., Wurgler, J., 2000. The equity share in new issues and aggregate stock returns. J. Finance 55 (5), 2219–2257.
Chemmanur, T.J., He, J., 2011. IPO waves, product market competition, and the going public decision: Theory and evidence. J. Financial Econ. 101 (2), 382–412.
Chen, A.Y., Zimmermann, T., 2022. Open source cross-sectional asset pricing. Crit. Finance Rev. 27 (2), 207–264.
Chinco, A., Sammon, M., 2024. The passive ownership share is double what you think it is. J. Financial Econ. 157, 103860.
Cremers, K.M., Petajisto, A., 2009. How active is your fund manager? A new measure that predicts performance. Rev. Financial Stud. 22 (9), 3329–3365.
Daniel, K., Titman, S., 2006. Market reactions to tangible and intangible information. J. Finance 61 (4), 1605–1643.
Dichev, I.D., 2007. What are stock investors' actual historical returns? Evidence from dollar-weighted returns. Am. Econ. Rev. 97 (1), 386–401.
Fama, E.F., French, K.R., 2008a. Average returns, B/M, and share issues. J. Finance 63 (6), 2971–2995.
Fama, E.F., French, K.R., 2008b. Dissecting anomalies. J. Finance 63 (4), 1653–1678.
Frazzini, A., Israel, R., Moskowitz, T.J., 2018. Trading Costs, vol. 3229719, SSRN.
Frazzini, A., Lamont, O.A., 2008. Dumb money: Mutual fund flows and the cross-section of stock returns. J. Financial Econ. 88 (2), 299–322.
Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H.E., 2006. Institutional investors and stock market volatility. Q. J. Econ. 121 (2), 461–504.
Greenwood, R., Sammon, M., 2025. The disappearing index effect. J. Finance 80 (2), 657–698.
Harris, L., Gurel, E., 1986. Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. J. Finance 41 (4), 815–829.
Harvey, C.R., Mazzoleni, M.G., Melone, A., 2025. The unintended consequences of rebalancing.
He, P., 2007. A theory of IPO waves. Rev. Financial Stud. 20 (4), 983–1020.
Ikenberry, D., Lakonishok, J., Vermaelen, T., 1995. Market underreaction to open market share repurchases. J. Financial Econ. 39 (2–3), 181–208.
Lou, D., 2012. A flow-based explanation for return predictability. Rev. Financial Stud. 25 (12), 3457–3489.
Loughran, T., Ritter, J.R., 1995. The new issues puzzle. J. Finance 50 (1), 23–51.
Loughran, T., Vijh, A.M., 1997. Do long-term shareholders benefit from corporate acquisitions? J. Finance 52 (5), 1765–1790.
Madhavan, A., Ribando, J., Udevbulu, N., 2022. Demystifying index rebalancing: An analysis of the costs of liquidity provision. J. Portf. Manag. 48 (6).
McLean, R.D., Pontiff, J., Reilly, C., 2020. Taking sides on return predictability. Available At SSRN 3637649.
Myers, S.C., Majluf, N.S., 1984. Corporate financing and investment decisions when firms have information that investors do not have. J. Financial Econ. 13 (2), 187–221.
Pástor, L., Veronesi, P., 2005. Rational IPO waves. J. Finance 60 (4), 1713–1757.
Pedersen, L.H., 2018. Sharpening the arithmetic of active management. Financ. Anal. J. 74 (1), 21–36.
Pontiff, J., Woodgate, A., 2008. Share issuance and cross-sectional returns. J. Finance 63 (2), 921–945.
Ritter, J.R., 1991. The long-run performance of initial public offerings. J. Finance 46 (1), 3–27.
Sammon, M., Murray, C., 2024. Primary capital market transactions and index funds. Available At SSRN.
Sammon, M., Shim, J.J., 2023. Do active funds do better in what they trade? Available At SSRN.
Sammon, M., Shim, J.J., 2024. Who clears the market when passive investors trade? Available At SSRN.
Sauter, G., 2022. The ideal index construction. J. Indexes (2).
Shleifer, A., 1986. Do demand curves for stocks slope down? J. Finance 41 (3), 579–590.
Spiess, D.K., Affleck-Graves, J., 1995. Underperformance in long-run stock returns following seasoned equity offerings. J. Financial Econ. 38 (3), 243–267.
Treynor, J.L., Black, F., 1973. How to use security analysis to improve portfolio selection. J. Bus. 46 (1), 66–86.
Wurgler, J., Zhuravskaya, E., 2002. Does arbitrage flatten demand curves for stocks? J. Bus. 75 (4), 583–608.