Transaction-cost-aware Factors*

Federico Baldi-Lanfranchi†

May 17, 2024

Abstract

I propose transaction-cost-aware (TCA) factors that are optimized to explain the returns investors can earn in practice, net of transaction costs. My methodology targets the trade-off between (i) acquiring exposure to risk factors and (ii) saving on transaction costs incurred in the process. Models that include TCA factors come closer to spanning investors' feasible efficient frontier. When trading is costly, TCA factor models increase net maximum squared Sharpe ratios by up to a factor of 2.5. TCA construction is most beneficial for high-turnover factors, such as momentum, that are otherwise unprofitable net of costs. I therefore suggest that asset pricing tests should focus on TCA factors to draw valid inferences.

JEL Classification: G11, G12, G14.

Keywords: Factor models, transaction costs, optimal trading.

*I thank my advisors Pierre Collin-Dufresne and Andreas Fuster for invaluable guidance. I am also grateful to Lorenzo Bretscher, Francesco Celentano, Kent Daniel, Victor Demiguel, Rüdiger Fahlenbrach, Julien Hugonnier, Semyon Malamud, Mohammad Pourmohammadi, and seminar participants at EPFL, HEC Lausanne, and Frontiers of Factor Investing conference for helpful comments and suggestions. All remaining errors are my own.

†EPFL and Swiss Finance Institute. E-mail: federico.baldilanfranchi@epfl.ch

1 Introduction

Investors demand compensation *net of transaction costs* to take on non-diversifiable risks. Identifying priced risk exposures has been the largest collective effort in asset pricing over the last forty years: Harvey, Liu, and Zhu (2016) review more than 300 candidate risk factors. However, these factors are typically designed and evaluated overlooking transaction costs that investors would incur to trade them.

In practice, asset pricing factors require turnover to preserve the link between expected returns and conditioning information on firm characteristics. When researchers identify characteristics that predict future returns, they construct factor portfolios that exploit this conditioning information. Individual assets receive weights at time t that reflect characteristic realizations at $t-1$. In the next period, new characteristic realizations require investors to revise weights on factor constituents.

Prior work assumes that investors always trade to perfectly realign factors and conditioning information. The resulting factors are transaction-cost-*unaware* (TCU) because they rebalance in full irrespective of how expensive this adjustment is. In this paper, I argue that the TCU approach is only sensible absent transaction costs. Instead, I take the perspective of investors who evaluate the benefits of realigning with conditioning information against the ensuing rebalancing costs. I propose transaction-cost-*aware* (TCA) factors that address this rebalancing trade-off.

When new characteristic realizations become available, TCA factors target the implied weights (to gain characteristic exposure) but rebalance only partially towards these weights at a fixed trading intensity τ (to control transaction costs). Equivalently, a TCA factor can be seen as a weighted average of its previous period allocation and the target factor. For instance, at the end of period $t-1$ an investor holding TCA momentum shifts a share τ of her investment towards the target, that is the TCU momentum factor UMD_t. She instead leaves the remainder invested at her current allocation. The larger the optimal trading intensity, the faster the investor adjusts towards the target factor.[1]

[1] Reducing the *frequency* at which factors reconstitute offers a heuristic alternative to contain transaction costs. However, this cost-mitigation technique is fundamentally different in nature and outcome from trading intensity optimization. Factors that reconstitute at low frequencies revise portfolio weights according to a coarser information set that ignores higher frequency changes in characteristics. Conversely, TCA factors exploit all available information but control the speed at which portfolio weights come to reflect new characteristic realizations. The latter approach is preferable for two main reasons. First, TCA trading directly aligns with Fama's (1991) efficiency argument, according to which investors act on new information to the extent to which the marginal benefits outweigh marginal costs. Second, rebalancing frequencies are chosen ad hoc while trading intensity is the outcome of a maximization process.

Both TCA and TCU factors can be evaluated on their net returns (which investors earn after costs) or their gross returns (which include compensation that accrues to liquidity providers). These two distinctions contextualize four approaches to asset pricing inference.

The standard protocol in the literature tests *gross* TCU factors. Unfortunately, this approach conflates net returns and transaction costs. Gross returns mask the cost of trading as a positive contributor to factor performance. As a result of this, an important contribution by Detzel, Novy-Marx, and Velikov (2023) shows that before-cost inference favors factors with high turnover. Such factors carry a large transaction cost component that is incorrectly accrued as a gain to the investor. Loosely speaking, neglecting the cost of trading transforms the rebalancing trade-off into a rebalancing *incentive*.

Detzel et al. (2023) and Li, DeMiguel, and Martin-Utrera (2023) revise model selection using *net* TCU factors. However, resulting inferences are not necessarily informative on whether the underlying characteristics are priced in the cross-section. I show that the performance of net TCU factors is largely reflective of their construction inefficiencies. When transaction costs are nonzero, these factors rebalance too aggressively. Suboptimal construction compresses net risk-premia, biasing results against high-turnover factors.

The main focus of this paper is on net TCA factors. These factors are meaningful for rational investors who also care about dimensions of risk other than variance. When trading is costly, these investors acquire multi-factor exposure in a cost-efficient way.[2] Optimal trading intensities are factor-specific, and driven by three channels. First, factors that trade on persistent characteristics command lower trading intensities. For instance, size factors require less aggressive rebalancing than profitability factors: Firms that are large today tend to stay large in the future, but current profitability does not guarantee large earnings going forward. Second, characteristics that correlate positively with transaction costs drive down optimal trading intensities. This is the case of momentum since recent underperformers tend to have larger bid-ask spreads. Lastly, factors with slow-decaying risk-premia receive higher trading intensities because each dollar spent in transaction costs buys a longer streak of high returns.

I judge TCA factors on two main criteria. First, models that replace TCU factors with their TCA variations should come closer to spanning the *feasible* efficient frontier. To this end, I use the squared Sharpe ratio criterion (Sh^2) of Barillas and Shanken (2017) as a model comparison tool after correcting factor returns for proportional transaction

[2]In this sense, TCA factors extend and generalize the insights of Gârleanu and Pedersen (2013). They solve the optimal trading rule for an investor with mean-variance preferences that faces quadratic transaction costs.

costs.[3] Their methodology ranks models on the squared Sharpe ratio achieved by a mean-variance-efficient (MVE) combination of their factors. This metric quantifies how closely the factors span the efficient frontier.[4] I focus on the six factor models studied in Detzel et al. (2023), which provide a representative account of low-dimensional specifications used in asset pricing research. Second, each TCA factor should individually explain differences in net average returns equally or better than its TCU counterpart. I examine performance at the factor level through spanning regressions.

TCA factors improve the pricing ability of all models I consider. In terms of net Sh^2 ratios, TCA models perform 28% to 150% better than their TCU counterparts. Net Sh^2 improvements achieved when switching from TCU to TCA factors *within* a candidate model often overshadow net Sh^2 differences *across* competing TCU models. I also document that these reductions in net-of-cost pricing errors are statistically significant and persist out-of-sample. Further, factor models that already employ prominent cost mitigation heuristics still obtain meaningful net Sh^2 improvements when factor trading intensities are also optimized.

Spanning regressions confirm that TCA factors are also individually better suited to explain differences in net asset returns. All eleven TCA factors I consider produce positive net alphas when regressed on their TCU counterparts, and six of these alphas are statistically significant. Conversely, TCU factors leave negative or insignificant intercepts on the TCA versions.

TCA models' success in describing net returns comes largely from improvements in the factors that are most expensive to trade. This is exemplified by the momentum factor under three scenarios: (i) with TCU construction and ignoring transaction costs; (ii) with TCU construction and after costs; and (iii) with TCA construction and after costs. In my sample, TCU momentum earns a gross premium of 0.64% per month, the highest among the factors I consider. This large gross premium overestimates the performance that investors realize in practice. Momentum also incurs the most transaction costs. These expenses are particularly severe when trading the TCU factor, which requires 63 bps per month in trading costs. On a net basis, the premium on TCU momentum drops

[3] Proportional costs offer a conservative estimate of overall transaction costs. Investors experience additional implementation frictions due to fixed costs, short-selling fees, price impact costs, and taxes on dividends and capital gains. However, a more comprehensive gamut of trading frictions makes the assumption of TCU trading relatively more restrictive. Expanding the set of frictions considered would thus result in larger benefits from TCA factor construction. Li et al. (2023) prove that the maximum squared Sharpe ratio criterion remains valid as a model comparison tool when transaction costs have a proportional form.

[4] This approach provides a general model ranking tool, whose validity is not restricted to a specific choice of test assets.

from 0.64% to a negligible 0.01%. However, a large share of this performance drop comes from inefficient rebalancing. Instead, TCA momentum only incurs 30 bps in monthly transaction costs, less than half of the TCU version. Trading momentum at the optimal intensity increases the factor's net return 22-fold and its annualized net Sharpe ratio by 0.19. TCA construction also clarifies momentum's importance in spanning the feasible efficient frontier. Models I review imply MVE portfolios that load marginally on net TCU momentum. Conversely, the weight on the net TCA version consistently exceeds 10%.

Model comparison in the TCU case does not reflect the true importance of high-turnover factors. As a consequence, the Sh^2 criterion ranks models differently within the classes of TCA and TCU factors. To illustrate this, consider first TCU factor models. Out of the six I review, the six-factor model of Barillas and Shanken (2018) (BS6) dominates before costs. It has the highest Sh^2 (2.25) and performs 8.4% better than the second-best model. However, this superior performance is largely illusory. Five out of six of the factors in the model reconstitute at a monthly frequency, and are therefore expensive to trade. After accounting for transaction costs, BS6 ranks second-worst and its Sh^2 drops by 80%. Underperformance on a net basis is primarily due to construction inefficiencies. Moving from TCU to TCA factors almost doubles the model's net Sh^2, positioning it as the third-best performer among the six considered.

Related Literature: This paper contributes to an emerging literature on model comparison with transaction costs.[5] My methodology directly optimizes factor construction for the cost of trading. This is in contrast with prior work which corrects TCU factor returns for transaction costs but restricts investors from explicitly optimizing factor design by taking these costs into account. Chernov, Dahlquist, and Lochstoer (2024) show that an out-of-sample MVE portfolio of G10 and emerging market currencies prices prominent currency factors net of transaction costs. Dickerson and Nozawa (2024); Dickerson, Nozawa, and Robotti (2023) draw net-of-cost comparisons between competing factor models within the space of corporate bonds. This paper is most related to Detzel et al. (2023) and Li et al. (2023) who study how successfully existing TCU factor models price the cross-section of stock returns under different transaction cost functional forms.[6] My

[5] Luttmer (1996) and Korsaye, Quaini, and Trojani (2021), among others, provide a theoretical characterization of stochastic discount factors in the presence of transaction costs.

[6] A directly related literature reviews the performance of anomalies in the cross-section of stock returns after correcting their returns for transaction costs. Novy-Marx and Velikov (2016) document that most TCU anomalies do not survive this correction when traded individually. DeMiguel, Martin-Utrera, Nogales, and Uppal (2020) find that transaction costs increase diversification benefits from trading multiple anomalies jointly.

empirical results also center around equities, but the TCA methodology can in principle be applied to any asset class.

TCA factors advance work in this literature in three ways. First, they are closer to practical implementation, in which investors actively optimize for frictions. Second, my methodology resolves distortions in inference caused by inefficient rebalancing and revives high-turnover factors. Third, TCA construction reinforces the linkage between empirical factor models and theoretical results that motivate them. According to Arbitrage Pricing Theory (Ross, 1976), investment opportunities that survive arbitrage activity must reflect compensation for risk. In practice, rational arbitrageurs only eliminate opportunities that are profitable after costs. Hedge funds employ sophisticated execution algorithms because cost mitigation expands the set of profitable trading opportunities. Therefore, the APT logic applies more closely to strategies that are implemented efficiently in the face of transaction costs, such as TCA factors. Investment opportunities that deliver positive gross alphas are not necessarily in violation of the APT. Such trading strategies are unattractive for arbitrageurs if alphas turn negative after costs, despite cost-aware execution.[7] Conversely, the APT is silent about strategies that earn negative net alphas when traded inefficiently. My methodology recognizes that such investment opportunities may still expand the efficient frontier if they turn profitable with cost-aware implementation.

I also document that factors that update conditioning characteristics infrequently face significant turnover in intermediate months. For example, the five factors of Fama and French (1993, 2015) reconstitute each June but incur between 35.9% and 64% of their yearly transaction costs in the remaining eleven months. This additional turnover is not accounted for in prior work and comes from two channels. First, research focuses on long-short factors that can be conveniently interpreted as traded excess returns. When factor legs earn uneven returns or corporate events occur, factors pick up a net exposure to the risk-free rate and this interpretation breaks down.[8] Therefore, factor investors face transaction costs to maintain dollar neutrality, even absent changes in firm characteristics. Second, researchers often restrict constituent weights better to identify the association between firm characteristics and expected returns. For instance, Fama and French assign equal weights to portfolios of small and large stocks within each leg of their factors. Maintaining this constraint further increases turnover.

[7] Detzel et al. (2023) make a similar argument in the TCU setting. However, once investors are allowed to optimize trading intensity, only those strategies that do not deliver positive net alphas *at any trading intensity* remain consistent with the APT.

[8] Cash dividends decrease the invested amount in the factor leg they originate from. M&A transactions also break dollar neutrality if the stocks of the target and acquirer are in opposite factor legs.

More generally, I connect the factor literature to papers on cost-aware trading that follow Gârleanu and Pedersen (2013). Gârleanu and Pedersen (2013) solve the dynamic optimal portfolio problem for myopic mean-variance investors that face quadratic transaction costs. Collin-Dufresne, Daniel, and Sağlam (2020) extend the framework to accommodate stochastic transaction costs, and show that investors should rebalance more heavily when costs are low. Collin-Dufresne, Daniel, and Saglam (2022) characterize the optimal trading rule of non-myopic investors in a similar setting. Jensen, Kelly, Malamud, and Pedersen (2022) propose a machine learning methodology to evaluate investment strategies against their net returns for each level of risk.

Lastly, this paper relates to the literature on cost mitigation techniques and factor strategies (Arnott, Li, and Linnainmaa, 2024; Detzel et al., 2023; Novy-Marx and Velikov, 2019; Rattray, Granger, Harvey, and Van Hemert, 2019). These approaches apply heuristic adjustments to factor strategies that reduce their implementation costs. I instead mitigate costs by optimizing the trade-off between risk and net returns directly.

The remainder of the paper is organized as follows. Section 2 documents the pitfalls of TCU factor construction. Section 3 illustrates the TCA methodology. Sections 4 to 6 contain empirical results on model selection, spanning regressions, and additional cost mitigation benefits. Section 7 concludes.

2 Distortions in inference with TCU factors

Figure 1 illustrates that the Sh^2 criterion of Barillas and Shanken (2017) can produce misleading results, both before and after costs, when applied to TCU factors. I compare three models. The first two are the CAPM of Sharpe (1964) and Lintner (1975), and the six-factor model of Fama and French (2018). The subscript c denotes that the model uses cashflow profitability in place of accruals-based profitability to construct the profitability factor. The third model is a variation of the FF6Fc specification that I design to have a particularly high cost of trading: I term this high-cost model FF6Fc^{HC}.

FF6Fc^{HC} exploits the same characteristics as FF6F, but adds three adjustments that deliberately inflate transaction costs. First, I reconstitute all factors in the model at the end of each month, rather than each June. This modification amplifies turnover and makes factors more expensive to trade. Second, I replace the size factor with one that invests only in small stocks, which have larger bid-ask spreads than large stocks. Put differently, I replace the FF6 size factor with the excess return on its long leg.[9] Third, I

[9] Fama and French (2018) test the pricing ability of a similarly constructed measure of size.

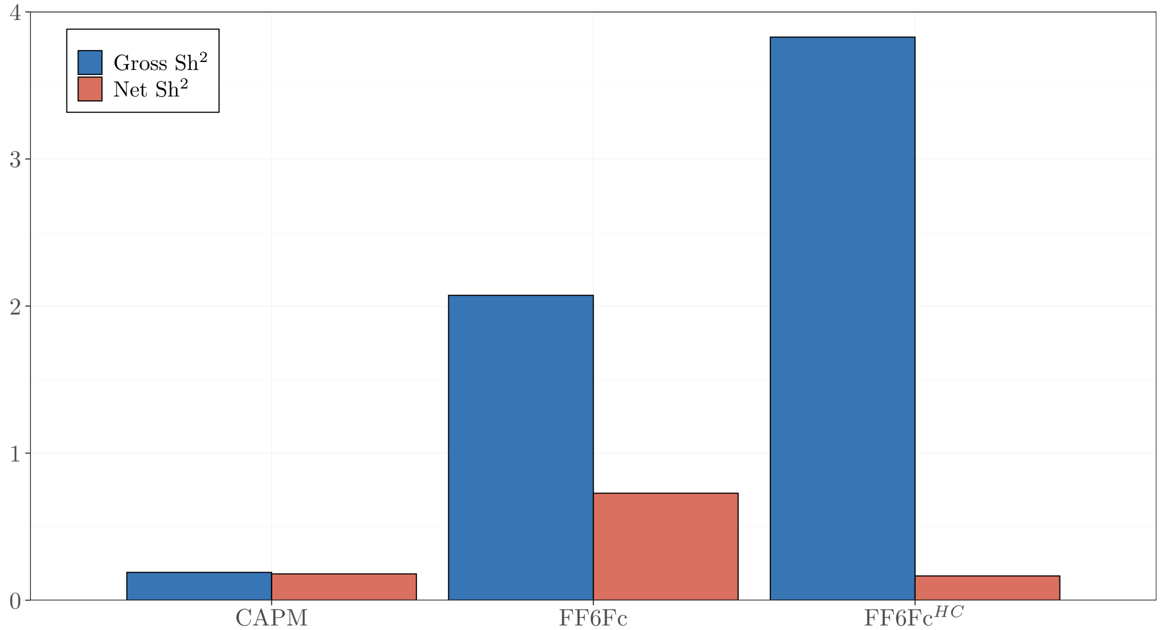

Figure 1: The effects of TCU construction. The above figure compares the performance of three asset pricing models. All three models include traditional, TCU factors. The vertical axis tracks the ex-post squared Sharpe ratio Sh^2 achieved by each model between July 1972 and December 2022. The first model is the classic CAPM of Sharpe (1964) and Lintner (1975). FF6c is the Fama and French (2018) six-factor model, where the subscript c denotes that the profitability factor is constructed on cashflows rather than operating profits. FF6Fc^{HC} is a variation of FF6Fc that requires larger transaction costs. It reconstitutes the FF6Fc factors every month and restricts the asset universe to the 50% of constituents with the highest transaction costs. The left (blue) bars show squared Sharpe ratios before accounting for transaction costs. The right (red) bars measure the Sh^2 using factor returns corrected for transaction costs.

restrict the set of constituents to stocks that, conditionally on qualifying for investment in a particular factor portfolio, are in the upper 50% of the transaction cost distribution. I apply this last adjustment to all factors except the market.[10]

Before costs, FF6Fc improves the ex-post squared Sharpe ratio Sh^2 by roughly 1.7 compared to the CAPM. This superior performance is not surprising. More than fifty years' worth of research in empirical asset pricing separates the two models. Further, the FF6Fc specification nests the CAPM, complementing the market with five additional factors. The CAPM's Sh^2 thus sets a lower bound for the more parameterized model.

[10] For instance, the high-cost value factor, HML^{HC} revises its composition at the end of each month, taking equal long positions in the HS^{HC} and HB^{HC} portfolios. At the same time, the factor shorts an equivalent dollar amount in the LS^{HC} and LB^{HC} portfolios. The HS^{HC} portfolio is invested in assets with high book-to-market and small market capitalization. Specifically, it only loads on the 50% of stocks with the highest transaction costs in the small size and high book-to-market segment (in a value-weighted fashion). I form the HB^{HC}, LS^{HC}, and LB^{HC} portfolios in a similar way.

Moving from the FF6Fc model to its high-cost version is instead exceedingly hard to justify from an economic standpoint. Yet, FF6Fc^{HC} delivers a comparable increase in Sh^2 with respect to its FF6Fc baseline. Accounting for transaction costs quickly resolves this tension: the apparent pricing ability of FF6Fc^{HC} is entirely fictitious. Out of the FF6Fc^{HC}, only the market is still profitable to trade net of costs. Therefore, FF6c^{HC} achieves the same net Sh^2 as the CAPM, 0.18, which is 95% lower than its gross-of-cost estimate. In practice, an investor trading an MVE portfolio of the FF6Fc^{HC} factors would move farther away from the achievable efficient frontier compared to the baseline FF6Fc model. Put differently, FF6Fc^{HC} only delivers a high Sharpe ratio before costs because gross returns on its factors include a large transaction cost component. FF6Fc^{HC} factors are thus informative about transaction costs investors might incur upon trading, but fail to explain the cross-section of returns they can earn in practice. Nonetheless, a researcher employing the Sh^2 criterion would conclude that we should prefer FF6Fc^{HC} as an asset pricing model.

Detzel et al. (2023) show a similar example, where a single factor based on low-volatility and industry relative-reversals (LV-IRR) dominates before costs, but delivers a negative net Sharpe ratio after correcting for transaction costs incurred. The main difference is that high turnover in LV-IRR comes from the economics of the underlying signal. Reversal factors are costly to trade because such signals rapidly revert to the mean. The case at hand shows that distortions in inference can also arise entirely due to construction choice. Figure 1 confirms that the FF6Fc factors do expand the available efficient frontier, even after correcting for the cost of trading, when they are constructed with the original methodology. The underperformance of FF6Fc^{HC} is thus not indicative of scarce promise in the Fama and French (2018) characteristics. Rather, it is an artifact of the construction methodology.[11] This fact is made apparent by the lack of substantive differences between the economic characteristics that drive FF6Fc^{HC} and FF6Fc. High pairwise correlations between factors in the two models solidify this point. For instance, the correlation between the two momentum factors, UMD and UMD^{HC}, is 92% before costs and 89% after costs.

This example shows that even inference based on *net* Sh^2 can produce misleading results when factor construction is not optimized. Discretionary choices can mask the pricing ability of the underlying characteristics, which is ultimately what researchers set out to demonstrate. In particular, the combined effect of net-of-cost inference and

[11]In this example, the construction of the FF6Fc^{HC} is deliberately suboptimal. In practice, academics can achieve similar results indirectly. Factors that rebalance more frequently and use complex sorting methodologies that overweight costly-to-trade assets are exposed to similar problems.

TCU factors sets an unreasonably high bar to clear for characteristics with little persistence. Asset pricing tests in the net TCU setting can suffer from low power and lead to over-rejections. At the same time, discretionary choices of reconstitution frequencies and other factor attributes can translate into large differences in performance, both before and after costs, that are uninformative about the pricing ability of the underlying characteristics. This phenomenon results in a proliferation of factors that exhibit little variation in terms of economic motivation and contributes to Cochrane's (2011) "factor zoo" problem.

3 TCA Factors

3.1 Trading partially toward the target

I consider an economy with N investable assets. Factors are characteristic-sorted portfolios that load on the asset universe, following the Fama and French (1993) blueprint. Each factor k is a fixed dollar portfolio that loads on stock i in month t with weight w^*_{it}, where the factor subscript is suppressed for legibility. Trading is costly and investors trade off the benefits of tracking the underlying characteristic closely against transaction costs they incur upon rebalancing. They do so by choosing an unconditional, factor-specific *trading intensity* $\tau \in (0, 1]$, so that dollar positions in each security satisfy:

$$x_{it}(\tau) = \tau \cdot w^*_{it} + (1 - \tau) \cdot w_{i,t-1}(\tau)(1 + \tilde{r}_{it}) \tag{1}$$

At the end of each month, investors move a share τ of their holdings toward target weights w^*_{it}. They retain the remainder invested at their current allocation $w_{i,t-1}(\tau)(1+\tilde{r}_{it})$, which reflects returns excluding dividends on each stock $\tilde{r}_{it}$.[12] Larger values of τ imply a more aggressive rebalancing schedule. Normalizing to keep investment in each portfolio leg constant yields the weights

$$w_{it}(\tau) = \frac{x_{it}(\tau)}{n_{it}(\tau)}\,, \tag{2}$$

where the normalizer n_{it} is

[12] I correct returns for M&A dividends, which are not included in the standard CRSP field. When M&A transactions are settled in cash, investors receive cash directly in a brokerage account. I assume that investors incur transaction costs when they re-invest such cash proceeds in the market, but not on the cash dividend itself. Sabbatucci (2015) shows that M&A dividends are substantial, and amount to 30% of total shareholder payout over the last 20 years.

$$n_{it} = \sum_{j=1}^{N_t} x_{jt}(\tau)\, \mathbb{1}\{sign(x_{jt}(\tau)) = sign(x_{it}(\tau))\}\,. \tag{3}$$

Rebalancing is costly and investors incur transaction costs $\mathrm{TC}_{it}(\tau)$, which reflect their trading intensity choice:

$$\mathrm{TC}_{it}(\tau) = \Big| w_{it}(\tau) - w_{i,t-1}(\tau)(1 + \tilde{r}_{it}) \Big| c_{it} \tag{4}$$

where c_{it} is the proportional (one-way) cost of trading 1\$ in stock i in month t. I estimate c_{it} from daily CRSP data following the guidelines of Abdi and Ranaldo (2017). If quote data is available, c_{it} is the quoted bid-ask spread scaled by twice the contemporaneous mid-point and averaged over month t. Otherwise, I estimate c_{it} with the CHL spread estimator that Abdi and Ranaldo (2017) propose. I detail the estimation process in Appendix A.

I compute net-of-cost factor returns $\tilde{f}_t$ similar to Detzel et al. (2023)

$$\tilde{f}_t(\tau) = f_t(\tau) - \mathrm{TC}_t(\tau) \tag{5}$$

where the gross return f_t and transaction costs associated with the factor, TC_t, are given as follows:

$$\begin{aligned} f_t(\tau) &= \sum_{i=1}^{N} w_{i,t-1}(\tau)\, r_{it} \\ \mathrm{TC}_t &= \sum_{i=1}^{N} \mathrm{TC}_{it}(\tau) \end{aligned} \tag{6}$$

When $\tau = 1$, TCA weights reduce to target weights, and $w_{it}(1) = w^*_{it}$. In other words, factor portfolios rebalance fully toward target weights in each period. This case recovers the setting in Detzel et al. (2023), in which investors incur transaction costs but cannot adjust their trading accordingly. Further restricting $c_{it} = c = 0$ nests the standard case of frictionless trading, which is the de facto standard in the empirical asset pricing literature. Conversely, as τ approaches zero, TCA factors move closer to "buy and hold" portfolios, where trading only occurs to keep the invested amount constant.

I evaluate competing factor models based on the maximum squared Sharpe ratio criterion of Barillas and Shanken (2017), which has recently gained increasing popularity in the asset pricing literature. Barillas and Shanken (2017) rank models on the squared Sharpe ratio achieved by a mean-variance efficient combination of their factors, $\text{Sh}^2(f)$. The methodology builds on Gibbons, Ross, and Shanken (1989), who show that augmenting a set of factors f with test assets R improves the achievable squared Sharpe ratio by

$$\alpha_R \Sigma^{-1} \alpha_R = \text{Sh}^2(R, f) - \text{Sh}^2(f) \tag{7}$$

where α_R are the pricing errors from regressing R on f. When R includes all possible factors, $\text{Sh}^2(R, f) = \text{Sh}^2(R)$ and minimizing pricing errors becomes equivalent to maximizing $\text{Sh}^2(f)$.

Similar to Detzel et al. (2023), I maximize squared Sharpe ratios of the *net* factors. In other words, I select the model that comes closest to spanning the mean-variance frontier investors can achieve *in practice*, after accounting for transaction costs. However, I substantially deviate from Detzel et al. (2023) in terms of the nature of the factors considered. In this paper, I extend the factor space to all portfolios that can be generated by trading toward a basis set of K target factors with the K-vector of factor-specific trading intensities τ. For each model, I choose τ and factor weights θ to maximize

$$\text{Sh}^2 = \underset{\theta,\tau}{max} \left\{ \frac{\mathbb{E}\Big[\theta' f_t(\tau) - |\theta|' \, \text{TC}_t(\tau)\Big]^2}{\mathbb{V}\Big[\theta' f_t(\tau) - |\theta|' \, \text{TC}_t(\tau)\Big]} \right\} \tag{8}$$

subject to $1'\theta = 1$ and $\tau \in (0, 1]^K$.

3.2 Choosing target weights

TCA factors require a choice of target weights w^*_{it} to trade toward. In an ideal scenario, such weights would be informed by a theoretical model that maps economic characteristics into risk-factor premia. In practice, the search for theoretically motivated linkages between risk-factors and economic characteristics is still an ongoing effort. I thus set target weights w^*_{it} so that all basis factors reconstitute at a monthly frequency. To do so, I run characteristic sorts underlying each factor's construction at the end of each month, irrespective of the original reconstitution frequency. I use market information as of t to construct portfolio weights for the following period. Instead, I update accounting characteristics at a six-month delay, in line with the original Fama and French (1993) methodology. The motivation for choosing monthly reconstituted target weights is threefold.

First, absent theoretical guidance, it is often unclear what lags of economic characteristics are relevant for expected returns. Taking the example of the value effect first illustrated by Basu (1983), what lag of book-to-market is most predictive of returns? TCA factors allow to sidestep this issue. Equation (1) shows that TCA weights are an exponentially-smoothed combination of past target weights. For a given level of transaction costs, the maximization problem (8) will suggest a lower trading intensity τ^*_{HML} if past values of book-to-market are more predictive of expected returns than recent realizations.

A second argument in favor of monthly reconstituted factors centers on the information set available to investors. The choice of target weights I propose always trades on the most recent information available on the underlying characteristic.

A third and more important reason to deviate from established conventions in the literature relates to transaction costs. For the purpose of this discussion, it is useful to distinguish between factor *reconstitution* and factor *rebalancing*. On each *reconstitution* date, the econometrician defines the investable asset universe for each factor, she sorts securities by the chosen characteristic(s) and assigns them to sub-portfolios formed at the intersections of such sorts. I refer to intermediate dates, in which the econometrician observes factor returns but no reconstitution takes place, as *rebalancing* dates.

Absent corporate events, portfolio assignments are only revised on reconstitution dates. Reconstitution is generally the leading source of turnover in factor construction because investors incur transaction costs to adjust their allocation. However, investors also need to engage in costly trading on rebalancing dates. Rebalancing needs arise to ensure that factor portfolios are well-defined excess returns and meet potential equal-weight, value-weight, or rank-weight constraints imposed for identification.

I find that the Fama-French factors, despite reconstituting each June, still experience significant turnover in other months due to rebalancing activity. While such turnover is costly, it acts on stale information, since characteristics entering portfolio sorts are only updated in June. Table 1 shows that non-June turnover and the transaction costs incurred because of it are substantial for the Fama-French factors. An investor holding a 100$ position in the size factor SMB would have incurred 63 cents worth of rebalancing costs each year due to turnover in months other than June. Such expenses would have

Table 1: Non-June rebalancing. The table below quantifies turnover and transaction costs incurred by the Fama-French factors in non-reconstitution months. I first compute transaction costs (TC) and turnover (TO) at the month and factor level. Transaction costs are $\text{TC}_{kt} = \sum_{i=1}^{N_t} |w_{ikt} - w_{i,k,t-1}(1 + \tilde{r}_{ikt})| c_{ikt}$ and turnover is $\text{TC}_{kt} = \sum_{i=1}^{N_t} |w_{ikt} - w_{i,k,t-1}(1 + \tilde{r}_{ikt})|/2$. Columns 3 and 5 respectively show the magnitudes of TC and TO in months other than June, expressed in %. I first sum TC and TO incurred between July and May of each year and report yearly averages. Columns 4 and 6 show shares of TC and TO incurred on rebalancing dates as a % of the yearly total. The sample spans from 1972 to 2022.

		Transaction Costs (TC)		**Turnover (TO)**	
	Characteristic	Non-June Level (%)	Non-June Share (%)	Non-June Level (%)	Non-June Share (%)
SMB	Size	0.63	64.0	51.1	60.4
HML	Value	0.69	46.8	54.1	44.9
RMW	Profitability	0.69	47.7	54.8	45.5
RMWc	Cash Profitability	0.68	39.9	54.4	37.6
CMA	Investment	0.69	35.9	54.8	32.4

amounted to 64% of the yearly transaction costs required to hold the size factor, with the remainder being incurred on reconstitution dates. Point estimates for transaction costs (TC) and turnover (TO) incurred when trading other factors are similar between July and May. Factors instead differ more heavily in the portion of turnover and trading costs originating on rebalancing dates as opposed to reconstitution dates. The investment factor, CMA, experiences the lowest share of non-June transaction costs (turnover), which is 35.9% (32.4%) of the yearly total. Such figures are nevertheless substantial, and suggestive that holding the Fama-French factors is not a passive endeavor, even absent reconstitution concerns.

Non-June rebalancing needs arise to keep the long and short ends of each factor balanced. If either leg of factor k outperforms the other at month-end t, the factor picks up a net exposure to the risk-free rate and loses its interpretation as a tradable excess return over the following period. Each factor leg is in turn an equally-weighted combination of sub-portfolios. For instance, the long leg of the value factor, HML, assigns equal weights to the portfolios of small- and large-value stocks. Similarly, the growth portfolio is an equally weighted mix of the small-growth and large-growth portfolios. Correcting for differences in returns across portfolios in the same factor leg also requires additional trading. Lastly, each of the constituent portfolios is a value-weighted portfolio. Therefore corporate events and dividends affecting any of the constituents also induce a rebalancing need. Taken together, these considerations suggest that reconstituting target weights w_{it}^* at a frequency that matches observed returns may be beneficial, as it reduces turnover that acts on stale information.

4 Model Comparison

In this section, I run horse races between competing asset pricing models. I focus on the six factor models covered in Detzel et al. (2023). These specifications have the benefit of being low-dimensional and have high tenure in the literature. Table 2 summarizes candidate models and factors. FF5 is the ubiquitous six-factor model of Fama and French (2015), to which FF6 adds a momentum factor. I denote with a subscript c the versions of the two models that replace accruals-based profitability with cash profitability. HXZ4 is the q-theory model of Hou, Xue, and Zhang (2015). Barillas and Shanken (2018) show that a combination of FF6 and HXZ4 factors, together with the monthly updated value factor of Asness and Frazzini (2013), achieves the largest Sh^2 before costs. I denote their model BS6.

I investigate the performance of each model (i) gross-of-cost using traditional (TCU factors), (ii) net-of-cost, but still assuming TCU factor construction, and (iii) net-of-cost with optimized trading intensity. I construct factors entering the first two sets of models following the documentation provided on the authors' webpages. I replicate weights in each factor, stock, and month to obtain factor-level excess returns before and after the cost of trading. I instead reconstitute all TCA versions of the factors on a monthly basis, irrespective of the original reconstitution frequency. Appendix B reports the construction methodology of characteristics entering TCA factors and replication statistics for TCU factors.

4.1 Maximum Squared Sharpe Ratios

Figure 2 illustrates the benefits of rebalancing factors conservatively in the presence of transaction costs. I show how models that trade toward TCA target weights fare net of costs for each possible choice of trading intensity. To stack the deck against results, I restrict factors to rebalance at the same intensity within each model. This restriction sets a conservative benchmark: optimal trading intensities for individual factors are likely heterogeneous, due to differences in turnover, return persistence, and the average cost of trading constituents.

The relationship between trading intensity and net Sh^2 is hump-shaped and appears smooth across all models. Net Sh^2 initially increases rapidly in τ, because factors gain exposure to the underlying characteristics. Rebalancing benefits die down when transaction costs become more substantial, and net Sh^2 peaks for values of τ between 20 and

Table 2: Candidate factors. The table below summarises the candidate factors and models that I evaluate in this section. FF5 and FF6 are the factor models of Fama and French (2015,1). The subscript c denotes variations of the above models that replace the operating profitability factor with a cashflow-based one. HXZ4 is the q-theory model of Hou et al. (2015). BS6 is the empirically motivated model of Barillas and Shanken (2018). BS6 replaces the standard value factor with a monthly-reconstituted version, which is due to Asness and Frazzini (2013).

Factor	Characteristic	Reconstitution	FF5	FF5c	FF6	FF6c	HXZ4	BS6
MKT	Market	Monthly	✓	✓	✓	✓	✓	✓
SMB	Size	June	✓	✓	✓	✓		✓
HML	Value	June	✓	✓	✓	✓		
RMW	Profitability	June	✓		✓			
RMWc	Cash Profitability	June		✓		✓		
CMA	Investment	June	✓	✓	✓	✓		
UMD	Momentum	Monthly			✓	✓		✓
ME	Size	Monthly					✓	
IA	Investment	Monthly					✓	✓
ROE	Profitability	Monthly					✓	✓
HMLm	Value	Monthly						✓

30%, depending on the model. Net Sh^2 declines past this level since excessive rebalancing erodes compensation for increased risk exposure.

Strikingly, HXZ4 starts delivering higher net Sh^2 than the baseline when τ is as low as 4.6%. Put differently, when transaction costs are present, it is preferable to retain 95.4% of funds invested at the previous period allocation, rather than fully rebalance the HXZ4 factors. A similar lower bound is consistent across models and all candidates outperform their TCU counterparts at a 6.9% trading intensity.

Two main empirical findings emerge from Figure 2. First, TCA models outperform TCU versions even without increasing model complexity. Since trading intensities are fixed at this stage, net squared Sharpe ratios only optimize over factor weights in the ex-post MVE portfolio. Put differently, TCA and TCU models have the same number of estimated parameters in this setting. Therefore, Figure 2 dismisses potential concerns that TCA factors may overfit in-sample with respect to their TCU baselines. Figure 2 also shows that performance gains from using TCA factors are remarkably robust to trading intensity misspecification. Limiting rebalancing activity delivers benefits over the baseline across all models and for a wide range of trading intensities. Even for the MVE portfolio implied by the TCA FF5c model, which is the closest to its baseline, the tangency portfolio lies above its TCU benchmark for all trading intensities between 6% and 78.9%.

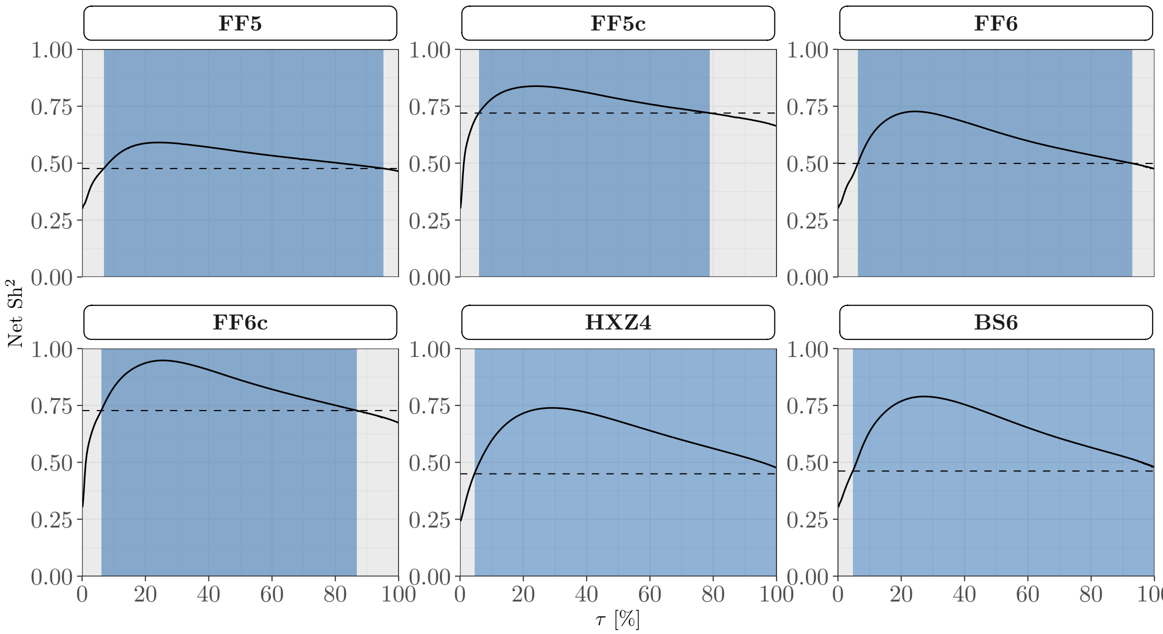

Figure 2: Benefits of conservative rebalancing. The figure graphs the relationship between trading intensity τ and net model performance. Each panel represents one of the candidate models. The continuous lines show the net Sh^2 each model achieves when trading toward its TCA target with intensity τ. Dashed lines represent instead the baseline net Sh^2 which can be achieved through TCU factors. The shaded region highlights the set of trading intensities that deliver higher or equal Sh^2 with respect to the baseline. The sample ranges between July 1972 and December 2022.

In the remainder of the paper, I relax the assumption of a common trading intensity at the model level. Figure 3 quantifies improvements in the ability to span the achievable efficient frontier when I allow each factor and model pair to rebalance at the optimal trading intensity. TCA factors deliver net Sh^2 that are 28% to 84% higher than their TCU counterparts. The FF6c model dominates both with an optimal choice of τ and when $\tau = 1$. However, four out of five of the remaining TCA models have higher Sh^2 than the TCU FF6c. Strikingly, investors would be better off pricing assets with any of these four suboptimal models, but choosing trading intensity optimally, rather than rebalancing naively and pricing assets with the best-performing TCU model. This point underscores the pitfalls of factors that are constructed without recognizing that investors alter their trading decisions when trading is costly.

While all TCA models perform better after costs, models where transaction costs are a larger concern benefit most from an informed choice of trading intensity. HXZ4 is the worst-performing TCU model, with a net Sh^2 of 0.45. The empirical challenges the model faces after fees are unsurprising since HXZ construct their factors using a 2x3x3 sorting methodology that magnifies turnover and amplifies the weight of small stocks. However, a more conservative choice of trading intensity improves the model's performance by 78%, resulting in a net Sh^2 of 0.79.

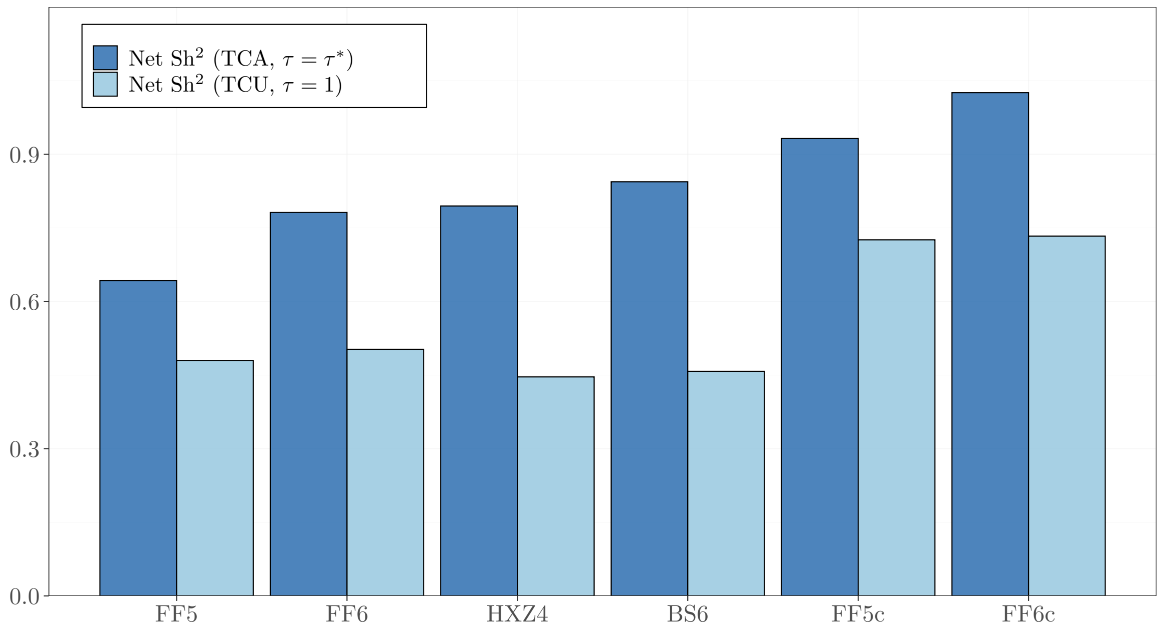

Figure 3: Net-of-cost model comparison: TCA and TCU factors. The above figure plots net model Sh^2. The left bars (dark blue) show the performance of each model with TCA factors. TCA factors trade toward monthly reconstituted target weights with optimized trading intensity τ^*. The right bars (light blue) measure how close the same models come to spanning the achievable efficient frontier when factors are TCU and trade with $\tau = 1$. I sort models by the Sh^2 they achieve in their TCA version. The sample ranges between July 1972 and December 2022.

Additional performance benefits of TCA models stem from the inclusion of momentum, which is especially costly to trade due to its fleeting portfolio composition. Momentum strategies suffer from large turnover because returns are typically less persistent than accounting characteristics: stocks that have done well in recent times may not continue their good runs in the future. Therefore, the BS6 factors, which include momentum, benefit even more than HXZ4 factors from transaction-cost-aware trading, with a performance improvement of approximately 84%. Further, adding momentum to the FF5 model only increases Sh^2 by 0.02 in the TCU case. Performance improvements are even more modest when considering cash profitability versions of the two models.

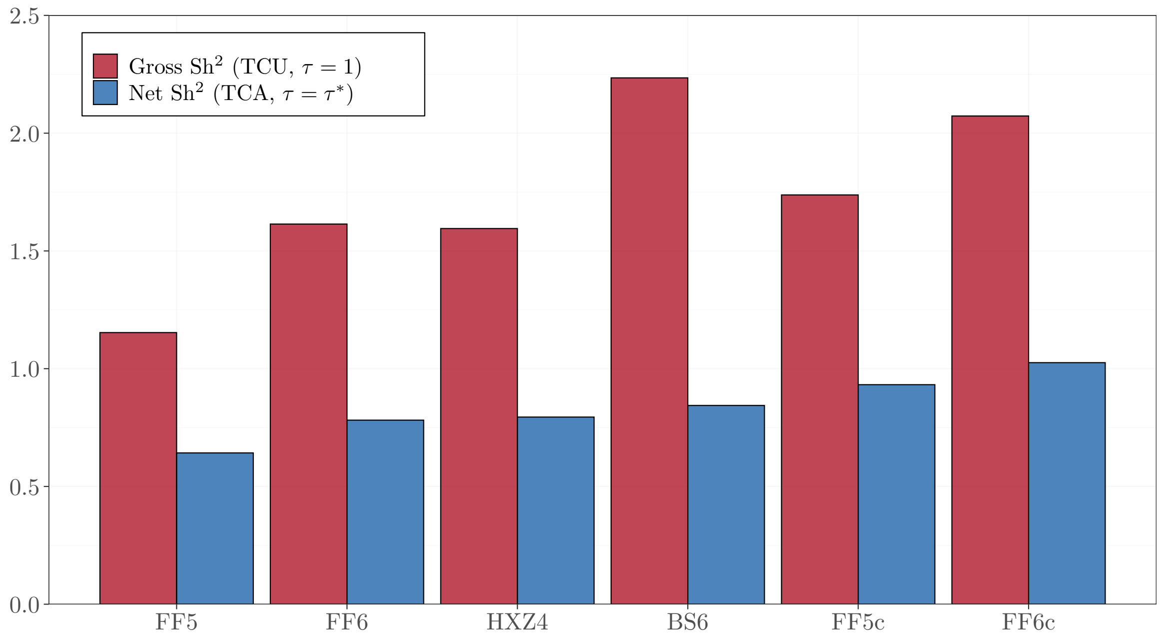

Figure 4: Model comparison: net TCA factors and gross TCU factors. The above figure compares model Sh^2 before costs and net of costs with TCA factors. The left bars (dark red) show the performance of each model before costs, using gross TCU factors. The right bars (dark blue) show the performance of each model with TCA factors. TCA factors trade toward monthly reconstituted target weights with optimized trading intensity τ^*. Models are sorted by their Sh^2 with TCA factors. The sample ranges between July 1972 and December 2022.

However, TCA models do benefit from momentum exposure. The TCA versions of FF6 and FF6c outperform their less parametrized counterparts by 0.13 and 0.10 respectively.

Heterogeneous benefits from transaction-cost-aware trading translate into heterogeneity in the ranking of the models considered. HXZ4 and BS6 outperform FF5 and FF6 in their TCA versions, while TCU models produce the opposite ranking. Similar to Detzel et al. (2023), standard statistical tests of Sh^2 differences cannot be applied to this setting. Asymptotic results on Sh^2 comparison rely on the delta method.[13] However, net Sh^2 ratios are not differentiable in the presence of proportional transaction costs. Equation 8 shows that the Sh^2 depends on the absolute value of factor weights in the MVE portfolio. Nonetheless, factor-level results, which are the object of section 5, show that the HXZ4 and BS6 factors benefit most from transaction-cost-aware trading on an individual basis. This finding is suggestive that differences in model rankings produced by TCA factors are likely robust. Further, FF5 and FF5c are approximately nested in their more parametrized FF6 and FF6c counterparts, provided factor trading intensities

[13] See Barillas and Shanken (2018) and Barillas, Kan, Robotti, and Shanken (2020).

on common factors are close. Investors would thus prefer the larger model in each of the two pairs, given their higher net Sh^2, even absent a rigorous asymptotic theory.

Importantly, the two different rankings in Figure 3 suggest that model comparison efforts after fees can be biased by the effects of discretionary construction choice. Inference on asset pricing models should instead contrast performance after optimizing trading intensity for transaction costs incurred in the process. This insight complements findings in Detzel et al. (2023), who show that net and gross Sh^2 produce different rankings if factors are TCU. Figure 4 highlights that this phenomenon extends to TCA factors. Both TCA Fama-French models incorporating cash profitability outperform BS6 after fees, while the latter model dominates gross-of-cost.

Figure 4 also shows that before-cost model comparison dwarfs differences between the models in the more realistic setting where investors experience transaction costs and construct factors accordingly. Apparent differences in model performance largely manifest due to arbitrary construction choices and when the cost of trading is overlooked.

4.2 Optimal trading intensities

Figure 5 shows how τ^* varies across TCA factors and candidate models. Optimal trading intensities are far below 100%. Trading factors conservatively thus brings investors closer to the achievable efficient frontier in the presence of transaction costs.

It is important to understand through which channels models benefit from optimizing trading intensity. Does the choice of τ^* reflect meaningful economic properties of the underlying characteristics? Or is it an artifact of minor construction details and correlation effects at the model level? Figure 5 can only offer partial insights on this point, due to the narrow set of factors considered. However, results seem to point toward the former interpretation. Estimated values of τ^* differ substantially between factors, but competing models assign similar trading intensities to each individual factor. Optimal trading intensities are also consistent within factor themes. The operating profitability factor, RMW, constitutes the only outlier in this respect. Estimated τ^* appear instead similar between size factors (SMB and ME), investment factors (CMA and IA), and profitability factors (RMWc and ROE).

Factor-level regularities in optimal trading intensities are also desirable from an empirical standpoint. Fixing factor-level trading intensities across models can be appealing, provided that it carries minor implications in terms of performance. This restriction may be particularly convenient when comparing a large set of models, especially if one believes that none of the candidates are correctly specified.

Benefits from reducing trading intensity seem lowest for the operating profitability factor, RMW, which rebalances roughly 70% of its allocation on a monthly basis. The high τ^* for RMW is in contrast with what I find for its cash-based version. The optimal trading intensity for RMWc is close to 25%, less than half of τ^*_{RMW}. This suggests that turnover in the operating profitability factor is inflated by mean reversion in accruals. Sloan (1996) finds that variation in operating profits coming from accruals is less persistent than the cashflow component. The ROE factor of HMZ, which sorts stocks on quarterly ROE, also trades substantially more slowly. More granular information about profitability allows for gradual adjustments in factor composition which can be smoothed over time.

The market and the two size factors, SMB and ME, require the least trading. Since these factors are value-weighted and sorted on market equity, they are close to self-rebalancing. Investors only need to adjust the composition of the market factor due to corporate events and net issuance, which may translate into higher-than-average transaction costs on the affected securities. In line with this intuition, τ^*_{MKT} ranges between 2.1% (FF6) and 3.3% (FF5c). SMB and ME face additional turnover with respect to the market due to *migration*, i.e. the rebalancing need that arises when securities move between the equally-weighted sub-portfolios that constitute each factor. The larger turnover aligns with even lower optimal trading intensities, which range between 1.5% and 2%.

4.3 Ex-Post Optimal MVE Weights

In this section, I investigate how TCA models realize improvements over their TCU counterparts. It is possible that transaction-cost-aware trading delivers similar improvements for all factors, leaving their relative importance unchanged. Conversely, if any of the candidate factors drive larger benefits from more conservative trading, their weights in the MVE portfolios should increase. Table 3 shows that TCA models produce different ex-post efficient MVE portfolio weights with respect to both gross and net TCA factors.

Ignoring transaction costs understates the relevance of the market factor with respect to the TCA case in Panel C. Overinvestment in the market results in higher leverage across all models. TCA models also attach higher weights to the size factor, with the

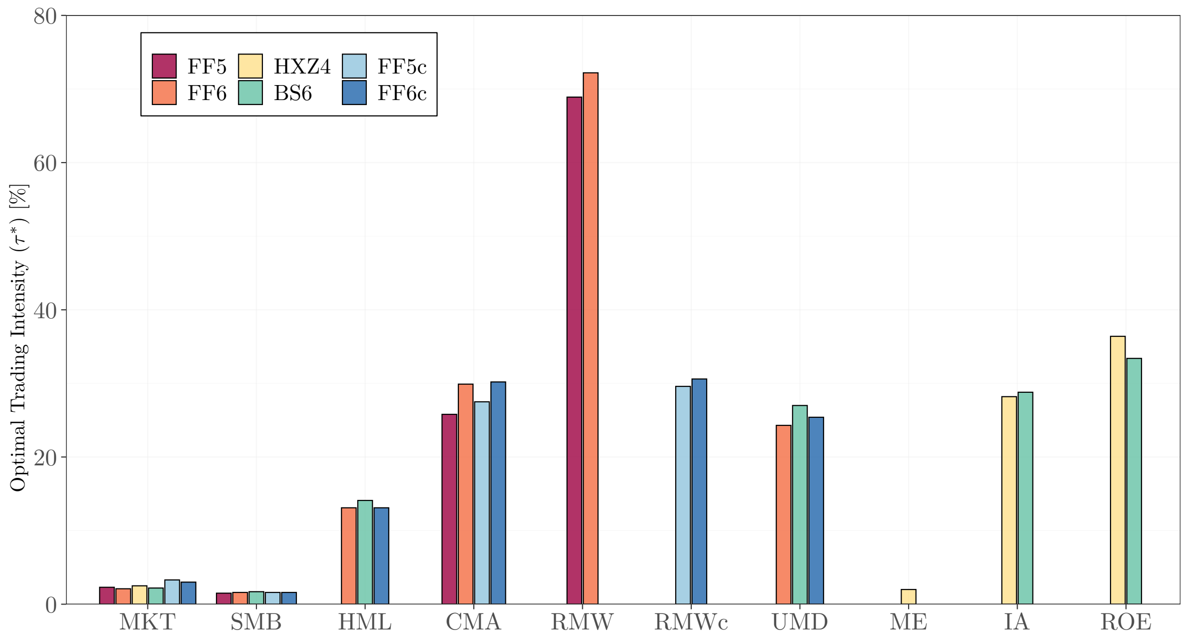

Figure 5: Optimal trading intensities. The above figure shows estimated trading intensities for each factor and model. The sample ranges between July 1972 and December 2022. Optimal trading intensities are undetermined for factors that receive zero weight in a given model. I therefore drop the corresponding model and factor pairs.

exception of ME in HXZ4. Figure 5 shows that market and size are the factors with the slowest optimal trading, which aligns with their gain in relative importance in Panel C of Table 3. Changes in weights are most striking for the BS6 model, which includes 5 monthly reconstituted factors out of 6. The weight on the investment factor increases from 4% to 32% when moving from the optimal factor portfolio before costs to the TCA case. This change comes at the expense of the value, profitability, and momentum factors which are costlier to trade.

Contrasting Panels B and C of Table 3 shows how ex-post MVE weights differ between TCA and TCU factor models after transaction costs. Forcing $\tau = 1$ results in an overly conservative allocation. As spread factors are rebalanced more aggressively than optimal, models reduce loadings on costlier-to-trade factors to compensate. Weights on the market factor thus increase across all models with respect to the TCA case. In turn, failing to account for transaction costs in factor design can dampen the relative importance of costlier-to-trade factors when performing inference net of transaction costs. This effect is especially apparent with the momentum factor, UMD. TCU models load only marginally on momentum, to the point that BS6 places zero weight on UMD, effectively collapsing into a 5-factor model. This finding is consistent with the minor

Table 3: Ex-post mean-variance efficient weights. The table below reports factor weights in the ex-post efficient mean-variance portfolio. Panel A shows optimal loadings on TCU factors in the standard asset pricing inference setting, in which the cost of trading is ignored. Panel B reports factor weights in the Detzel et al. (2023) setting, in which TCU factors are evaluated net of transaction costs. Panel C shows instead ex-post efficient weights on TCA factors.

Panel A: Gross TCU Factor Models											
Model	MKT	SMB	HML	CMA	RMW	RMWc	UMD	ME	IA	ROE	HMLm
FF5	18	9	-3	46	30						
FF5c	17	13	-3	29		44					
FF6	18	8	4	33	23		14				
FF6c	17	11	2	23		37	10				
HXZ4	15							15	36	34	
BS6	14	10					19		4	27	26
Panel B: Net TCU Factor Models											
Model	MKT	SMB	HML	CMA	RMW	RMWc	UMD	ME	IA	ROE	HMLm
FF5	26	4	7	34	29						
FF5c	22	10	4	18		46					
FF6	25	4	9	29	26		7				
FF6c	22	10	5	16		44	3				
HXZ4	27							6	38	29	
BS6	24	8					0		30	32	6
Panel C: Net TCA Factor Models (τ^*)											
Model	MKT	SMB	HML	CMA	RMW	RMWc	UMD	ME	IA	ROE	HMLm
FF5	22	11	0	42	25						
FF5c	19	15	0	26		40					
FF6	20	9	11	26	18		16				
FF6c	18	13	5	20		34	10				
HXZ4	21							13	43	23	
BS6	20	11					10		32	17	10

differences in Sh^2 observed in Figure 3 between both TCU versions of the FF5 and FF6 models. Momentum plays a much larger role in TCA factor models, in which weights on the factor more than double, and become sometimes larger than in the before-cost benchmark.

TCA models exhibit sparsity with respect to the gross-of-cost benchmark: HML washes out of the TCA versions of FF5 and FF5c. The net-of-cost maximum Sharpe ratio criterion introduces additional sparsity with respect to the TCU case because transaction costs in equation (8) effectively act as a LASSO penalty. The exclusion of HML reflects the extended drawdown that value suffers in the recent sample. On top of its poor recent performance, HML is positively correlated with the investment and profitability factors and provides scarce diversification benefits. However, HML resurfaces in models that also include a momentum factor. Asness and Frazzini (2013) show that trading value and momentum jointly is beneficial in light of their negative correlation.[14]

4.4 Out-of-Sample Results

Standard model comparison tests such as those of Gibbons et al. (1989) and Barillas et al. (2020) cannot be applied with proportional transaction costs. In this section, I turn to the bootstrap approach of Fama and French (2018) and Detzel et al. (2023) to draw a statistical comparison between models that employ TCA and TCU factors.

I obtain 10,000 simulations of in-sample (IS) and out-of-sample (OS) net Sh^2 of each model as follows. First, I partition the 606 available months between July 1972 and December 2022 into consecutive pairs: (1,2), (3,4), ..., (605,606). In each simulation run, I sample 303 pairs with replacement. I then randomly assign one month within each sampled pair to the IS period and its partner month to the OS period. If any given pair is sampled more than once, I keep the same assignment of months within the pair to the IS and OS periods. This ensures that no overlap exists between the IS and OS samples. I then obtain factor weights θ and trading intensity τ that maximize the IS net Sh^2 for each TCA factor model specification. In the OS period, I use the estimated parameters to back out the OS net Sh^2. I repeat the same procedure for TCU models, setting τ to 1 for all factors and optimizing net Sh^2 only with respect to factor weights in the IS period.

Table 4 compares the net TCA and net TCU versions of each model. TCA models outperform in close to 100% of the IS simulations. In line with previous results, the BS6 model benefits most from TCA construction, with HXZ4 closely following. Their IS net Sh^2s improve by 120% and 91%, respectively. However, the absolute magnitudes of IS net Sh^2 should be interpreted with care: Ex-post Sh^2 estimates suffer from a well-known small-sample bias (Barillas et al., 2020; Fama and French, 2018; Jobson and Korkie, 1980). With the benefit of hindsight, models overweight factors whose realized returns

[14] In Detzel et al. (2023), HML drops out from all models after accounting for transaction costs in the TCU setting they consider. This is in contrast with my findings. The higher relative importance of HML in Panel B of Figure 3 with respect to Detzel et al. (2023) reflects the joint effect of (i) differences in the stock-level estimator for c_{it}, which has lower bias and higher correlation with TAQ effective spreads (ii) costs incurred by infrequently reconstituted factors on rebalancing dates, which Detzel et al. (2023) neglect, and (iii) differences in the sample considered, which includes one additional year in this paper.

are large compared to expected returns. This issue is even more pronounced in the IS bootstrap since only half of the sample is available for estimation.

The OS net Sh^2s are substantially lower, reflecting that investors would have only been able to realize part of the paper performance of the models. The IS bias is larger for TCA models in light of their additional flexibility. However, OS results confirm that TCA models robustly outperform their TCU counterparts and that improvements are more substantial for models with high-turnover factors.

Table 4: Bootstrap model comparisons between TCA and TCU factor models. The table below presents net-of-cost bootstrap comparisons between TCA and TCU versions of each candidate model based on 10,000 simulation runs. I first partition the 606 sample months, July 1972 to December 2022, into adjacent pairs: (1,2), (3,4), ..., (605, 606). Each simulation run then draws 303 pairs with replacement and randomly assigns one month within each pair to the in-sample (IS) period and the partner month to the out-of-sample (OS) period. If the same pair is sampled more than once, I keep the same IS/OS split within this pair. The first and second columns report the average net Sh^2 estimated for each model across IS simulation runs, respectively in its TCU and TCA versions. The third column, "TCA Best" presents the share of IS simulation runs in which the TCA model outperforms the TCU version, expressed in %. The last three columns present analogous results for OS simulations.

	In-Sample			Out-of-Sample		
	TCU Sh^2	TCA Sh^2	TCA Best	TCU Sh^2	TCA Sh^2	TCA Best
HXZ4	0.70	1.34	100.0	0.35	0.65	90.1
BS6	0.81	1.78	100.0	0.30	0.66	92.5
FF5	0.70	1.13	98.8	0.31	0.34	64.2
FF5c	0.95	1.64	99.6	0.51	0.67	76.5
FF6	0.78	1.59	100.0	0.30	0.51	84.5
FF6c	1.01	2.03	100.0	0.48	0.77	85.2

Panel B shows that FF6c is still the dominant model OS, confirming previous results. However, it is now only selected in 40.8% of the simulation runs. Conversely, the relative strength of FF5c and HXZ4 improves in the OS bootstrap. These results are suggestive that part of the IS outperformance of the larger FF6c and BS6 models is due to the small-sample bias discussed earlier.[15] In addition, including momentum in OS simulations is less beneficial, due to the large 2009 crash in this strategy (Daniel and Moskowitz, 2016). Simulation runs in which the crash is (is not) in the IS period underinvest (overinvest)

[15] Barillas et al. (2020) show that the small-sample bias in model Sh^2 increases with the number of factors.

in momentum. Both effects help FF5 and HXZ4 to close the distance with factor models that include a momentum factor.

Table 5: Bootstrap model comparisons for TCA factor models. The table below presents net-of-cost bootstrap comparisons between candidate models that use TCA factors. All figures reported are based on 10,000 simulation runs. I first partition the 606 sample months, July 1972 to December 2022, into adjacent pairs: (1,2), (3,4), ..., (605, 606). Each simulation run then draws 303 pairs with replacement and randomly assigns one month within each pair to the in-sample (IS) period and the partner month to the out-of-sample (OS) period. If the same pair is sampled more than once, I keep the same IS/OS split within this pair. The first and second columns report the average net Sh^2 estimated for each model. Columns 2 to 6 report the % of simulation runs in which the row model outperforms the column model. The last column shows the % of simulation runs in which the corresponding model dominates. Panel A presents results for IS simulation, while Panel B for OS simulations.

Panel A: In-Sample TCA Bootstrap Results							
		Probability that the Row Model outperforms the Column Model (%)					
	Average Sh^2	BS6	FF5	FF5c	FF6	FF6c	Best
HXZ4	1.34	2.2	74.1	26.2	26.0	6.1	0.9
BS6	1.78		97.8	63.6	77.0	25.1	23.9
FF5	1.13	2.2		1.3	1.2	0.1	0.0
FF5c	1.64	36.4	98.7		55.1	1.9	1.5
FF6	1.59	23.0	98.8	44.9		2.5	1.3
FF6c	2.03	74.9	99.9	98.1	97.5		72.4
Panel B: Out-of-Sample TCA Bootstrap Results							
		Probability that the Row Model outperforms the Column Model (%)					
	Average Sh^2	BS6	FF5	FF5c	FF6	FF6c	Best
HXZ4	0.65	44.6	87.3	47.7	67.9	36.2	22.8
BS6	0.66		88.6	50.0	79.7	33.0	17.2
FF5	0.34	11.3		3.6	18.4	5.5	0.3
FF5c	0.67	50.0	96.4		72.0	28.8	17.4
FF6	0.51	20.3	81.6	28.0		8.0	1.5
FF6c	0.77	67.0	94.5	71.2	92.0		40.8

Table 5 in Appendix C shows bootstrap simulations comparing TCU models. Similar to Detzel et al. (2023), who focus on model comparison in the TCU setting, I find that

FF6c is the best-performing model IS, but FF5c marginally wins out OS. This suggests that trading momentum is less beneficial still in the OS feasible portfolio of investors who do not optimize trading intensity. This is in part because conservative trading mitigates crash risk in the TCA momentum factor. Setting $\tau = 0.225$, which is close to the optimal trading intensity for UMD in the FF6, FF6c, and BS6 models, results in a kurtosis of about 9.36. The classic UMD factor with $\tau = 1$ has a substantially larger kurtosis of 13.38.

Overall, keeping the set of candidate characteristic-based factors constant, factor models with cost-aware construction substantially outperform TCU ones. Differences in performance *across* TCA factor models are less evident. Therefore, bootstrap simulations seem to confirm that cost-aware construction can dominate OS gains from careful model selection.

5 Rebalancing Trade-off

In earlier sections, I relate improvements in pricing ability delivered by TCA factors to a reduction in transaction costs. The premise is that, if characteristic C is priced in the cross-section of expected returns, investors face a trade-off between securing high exposure to C and containing transaction costs incurred in the process. However, competing channels may also contribute to pushing optimal trading intensities below 100%. In this section, I discuss other factors that may result in conservative trading and present evidence that transaction costs are the main driver of trading intensities.

Target factor weights are often empirically motivated, and likely not efficient even absent transaction costs. Optimizing trading intensity may therefore improve the efficiency of individual factors, both before and after the cost of trading. This is in stark contrast with the literature on dynamic price impact, where the aim portfolio is a weighted average of current and future expected MVE portfolios, which are efficient by construction (Collin-Dufresne et al., 2020; Gârleanu and Pedersen, 2013; Jensen et al., 2022). First, recent characteristic realizations may be noisy measures of their systematic components. Conservative trading attaches larger weights to lagged characteristic values and may produce more efficient factors absent transaction cost considerations. In this vein, Novy-Marx (2012) argues that momentum is largely driven by returns 12 to 7 months before portfolio formation, while more recent performance seems less informative. Further, Daniel, Mota, Rottke, and Santos (2020) show that characteristic sorts also pick up systematic components that contribute to portfolio variance, but do not represent priced variation in before-cost expected returns. Ehsani and Linnainmaa (2022); Fama and French (2020) propose related procedures to isolate priced variation and improve factor efficiency. If lagged characteristic realizations are less correlated with unpriced systematic components, reducing trading intensity may again prove beneficial irrespective of the cost of trading. While desirable, the above effects are unrelated to the transaction cost channel explored in this paper.

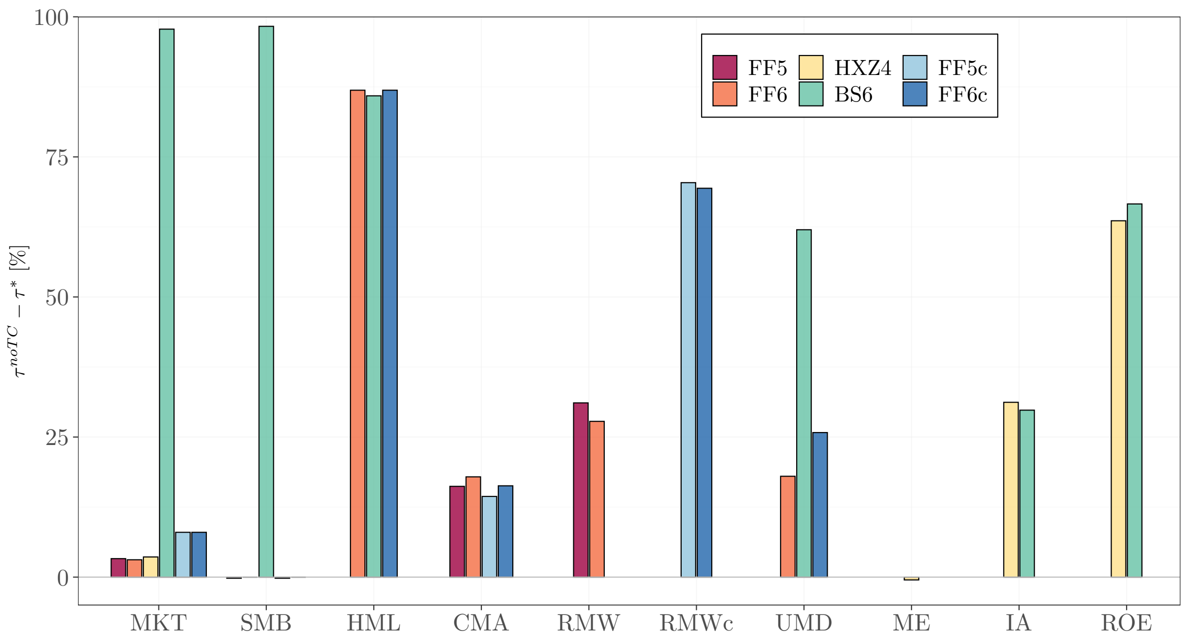

Figure 6: Optimal trading intensities with and without transaction costs. The above figure shows how transaction costs affect optimal trading intensity for each factor and model. I denote τ^{noTC} the optimal trading intensity absent transaction costs, i.e. when $c_{it} = c = 0$ (when trading is costly). The optimal trading intensity when trading is costly, which is displayed in figure 5, is instead τ^*. The vertical axis shows the differences between τ^{noTC} and τ^*. The sample ranges between July 1972 and December 2022. Optimal trading intensities are undetermined for factors that receive 0 weight in a given model. I therefore drop the corresponding model and factor pairs.

Figure 6 disentangles before-cost efficiency and the effect of transaction costs. I compute the difference in optimal trading intensity with and without transaction costs for each factor and model. When the cost of trading is set to zero, investors rebalance factors more aggressively, and optimal trading intensities approach 100% for the wide majority of models and factors. These findings substantiate that transaction costs are the leading driver of conservative trading.

The market and size factors represent the only partial exceptions. Turnover in such factors is strongly reflective of net issuance, as discussed in section 4.3. Conservative trading in MKT, SMB, and ME aligns with a large literature on equity issuance and stock returns. Loughran and Ritter (1995); Ritter (1991); Spiess and Affleck-Graves (1995); Stigler (1963), among others, show that firms tend to underperform following both seasoned and initial equity offerings. Under this view, delaying rebalancing can reduce the net issuance exposure of the market and size factors, and result in higher excess returns gross and net of the cost of trading. Daniel and Titman (2006,1) show that net issuance is a strong negative predictor of stock returns, suggesting that priced systematic variation rather than noise or unpriced variation may contribute to slower trading in market and size factors.[16]

Table 6: Individual factor premia. The table below zeroes in on individual factors. I report average monthly premia, μ, in % points, the associated t-statistics t, and the average transaction costs incurred per month, TC. Factor premia that are significant at the 5% level are in bold. I compare factors under three scenarios. Columns 1 and 4 refer to TCU factors evaluated before accounting for the cost of trading. Columns 2 and 5 relate to the same set of factors, after correcting excess returns for the cost of trading shown in column 7. Columns 3, 6, and 8 present matching results for TCA factors. Lastly, column 9 shows the difference in annualized net Sh between TCA factors and traditional TCU factors.

	μ (%)			t			TC (%)		
	TCU Gross	TCU Net	TCA Net	TCU Gross	TCU Net	TCA Net	TCU	TCA	Net ΔSh
MKT	**0.58**	**0.56**	**0.60**	3.09	3.01	3.46	0.02	0.01	0.07
SMB	0.14	0.05	0.14	1.12	0.43	1.44	0.09	0.03	0.14
HML	**0.37**	0.24	0.24	2.96	1.90	1.91	0.13	0.06	0.00
RMW	**0.32**	0.18	0.19	3.29	1.88	1.88	0.13	0.16	0.00
RMWc	**0.41**	**0.25**	**0.25**	4.99	3.01	2.99	0.16	0.16	0.00
CMA	**0.30**	0.13	**0.20**	3.72	1.59	2.41	0.17	0.14	0.12
UMD	**0.64**	0.01	0.23	3.66	0.07	1.40	0.63	0.30	0.19
ME	0.2	0.08	0.16	1.94	0.60	1.32	0.17	0.05	0.11
IA	**0.36**	0.14	**0.27**	4.14	1.58	3.12	0.22	0.14	0.22
ROE	**0.56**	**0.21**	**0.24**	5.24	1.97	2.19	0.35	0.25	0.03
HMLm	**0.38**	0.10	0.24	2.43	0.64	1.91	0.28	0.06	0.18

Trading partially toward target weights could also translate into higher pricing ability because it optimizes before-cost diversification between the factors. Table 6 presents evidence suggesting that this channel is not the main driver of my results. For each candidate factor, I compare the individual performance of its TCA and TCU versions. Individual comparisons shut down the diversification channel and quantify cost savings that TCA construction can achieve when trading factors in isolation.

All TCA factors have higher or equal premia with respect to the net TCU case. Only three of the eleven factors are 5% significant when $\tau = 1$. After optimizing trading

[16] Baker and Wurgler (2000) demonstrate that the equity share in new issues negatively predicts market returns more specifically.

Table 7: Spanning regressions. The table below reports coefficient estimates and associated t-statistics for spanning regressions. Statistically significant coefficients at the 5% level are shown in bold. For each factor candidate, I regress the net returns on the TCA version on the net returns earned by its TCU counterpart, and vice versa. Columns 2 and 3 investigate whether TCU factors span TCA factors. The dependent and independent variables are reversed in columns 4 and 5. Regression intercepts (α) are in basis points per month.

		Net TCA on Net TCU		Net TCU on Net TCA	
	Characteristic	α (bps)	β	α (bps)	β
MKT	Market	**0.09**	**0.91**	-0.07	**1.06**
		(2.58)	(124.06)	(-1.94)	(124.06)
SMB	Size	0.11	**0.6**	-0.08	**0.96**
		(1.71)	(28.67)	(-1.02)	(28.67)
HML	Value	0.01	**0.93**	0.01	**0.96**
		(0.33)	(71.68)	(0.31)	(71.68)
RMW	Profitability	0.00	**1.00**	0.00	**0.95**
		(0.21)	(112.36)	(0.19)	(112.36)
RMWc	Cash Profitability	0.01	**0.97**	0.01	**0.95**
		(0.37)	(81.32)	(0.51)	(81.32)
CMA	Investment	**0.08**	**0.92**	-0.05	**0.95**
		(2.6)	(64.96)	(-1.86)	(64.96)
UMD	Momentum	**0.22**	**0.88**	**−0.23**	**1.05**
		(4.86)	(85.74)	(-4.65)	(85.74)
ME	Size (monthly)	**0.09**	**0.98**	**−0.09**	**0.98**
		(3.53)	(116.93)	(-3.33)	(116.93)
IA	Investment (monthly)	**0.14**	**0.92**	**−0.12**	**0.96**
		(4.73)	(66.62)	(-3.87)	(66.62)
ROE	Profitability (monthly)	0.03	**1.00**	-0.01	**0.92**
		(1.04)	(80.00)	(-0.42)	(80.00)
HMLm	Value (monthly)	**0.17**	**0.72**	**−0.17**	**1.12**
		(3.02)	(50.07)	(-2.43)	(50.07)

intensities, t-statistics on the monthly premia generally increase, and both investment factors become significant at the 5% level. TCA factors are also substantially cheaper to trade, except the profitability factors.[17] Lastly, net (annualized) Sh are also larger or

[17] Figure 5 shows that the optimal trading intensity for the TCA RMW is close to 70%. Additional turnover due to monthly reconstitution of the target weights thus overstates the reduction in transaction costs from less aggressive trading.

equal for all TCA factors. These findings suggest that TCA factors deliver improvements over TCU factors through cost reduction.

Table 6 allows for qualitative comparisons, but offers limited insights in terms of inference. Differences in net Sh are not differentiable due to transaction costs, and cannot be tested directly. To address this problem, I run spanning regressions of TCA factors against their TCU counterparts. Table 7 shows that TCA factors deliver positive alphas over traditional TCU factors. These alphas are statistically significant at the 1% level for six out of the eleven factors, and particularly large for those that reconstitute at a monthly frequency in their original formulation. Conversely, TCU versions of the UMD, ME, IA, and HMLm factors do not span their TCA counterparts, and deliver significant negative alpha at the 5% or 1% level. This reinforces that optimizing trading intensity delivers significant benefits when transaction costs are present, over and above the effects of diversification. All betas are large and statistically significant, suggesting that the risk exposures of the two sets of factors are similar in nature.

6 Additional Cost Mitigation

Previous results highlight how TCA construction can substantially reduce transaction costs incurred when trading academic factors, and how cost savings translate into improvements in the factor models that investors can construct in practice. In this section, I evaluate how optimizing trading intensity fares when compared to two prominent cost mitigation approaches in the literature: "banding" and "netting". The banding strategy of Novy-Marx and Velikov (2019) reduces turnover by combining factor portfolios with buy-and-hold spreads. Factors that employ banding require more extreme characteristic realizations to enter active positions in stocks than to unwind these positions. Netting arises naturally for sophisticated investors, who can trade factor constituents directly. Such investors can net out offsetting trades across long and short legs of different factors. DeMiguel et al. (2020) term "trading diversification" the reduction in transaction costs that arises when netting is possible. Detzel et al. (2023) find that both of these strategies help mitigate transaction costs.

There are two important caveats when comparing TCA construction with netting and banding. First, the three cost mitigation approaches are complementary. TCA factors can accommodate additional cost mitigation by applying netting and banding to their target portfolios. Second, while in principle the size of the banding buy-and-hold spread can be optimized, Detzel et al. (2023) set an ex-ante a 20% buy-and-hold spread for all factors. Conversely, TCA factors are the outcome of the maximization problem (8), which optimizes factors directly while accounting for their transaction costs. This approach aligns closely with the market efficiency argument of Fama (1991). Under this revised version of the efficient market hypothesis, investors exploit available information to the full extent to which this is beneficial after accounting for the trading costs they incur. Accordingly, my methodology incorporates information on characteristics embedded in the TCA target while optimally accounting for transaction costs. Different characteristics command different optimal trading intensities which in turn can be related to characteristics' cross-sectional and time-series properties.

6.1 Banding

I construct cost-mitigated factors that employ banding following Detzel et al. (2023). I start with replications of the original TCU factors and apply a 20% buy-and-hold spread to portfolio cutoffs. For instance, standard academic factors are long stocks in the top 30% of the underlying characteristic's distribution and short those that land in the bottom 30%. Conversely, the cost-mitigated factors only enter long (short) positions when characteristics cross the top (bottom) 20% and maintain these positions until characteristics fall out of the top (bottom) 40%.

I also combine banding and TCA construction to evaluate the joint benefits of the two methodologies. To this end, I first apply banding to the monthly reconstituted TCA targets. I then solve problem (8) again using the updated targets.[18]

Figure 7 compares the net Sh^2 achieved by candidate models in four scenarios: (i) with their original TCU design, (ii) with TCU design and banding, (iii) with TCA construction but without banding and, (iv) with both TCA construction and banding. An individual comparison between banding and TCA construction shows that the latter methodology delivers larger improvements in net Sh^2, with the only exception of BS6. In addition, incremental benefits materialize when the two methodologies are applied jointly. Synergies are particularly apparent in the case of the FF5 and FF5c models. All factors in these models reconstitute at a yearly frequency and benefit little from banding in isolation. However, TCA construction allows the investor to reconstitute these factors on a monthly basis, while keeping transaction costs in check. This also increases the scope for banding, as turnover in the target portfolio increases.

[18] Detzel et al. (2023) apply banding only to monthly reconstituted factors. Since all TCA factors reconstitute at a monthly frequency, I also implement buy-and-hold spreads on the TCU factors that reconstitute less frequently to set a higher hurdle for my methodology.

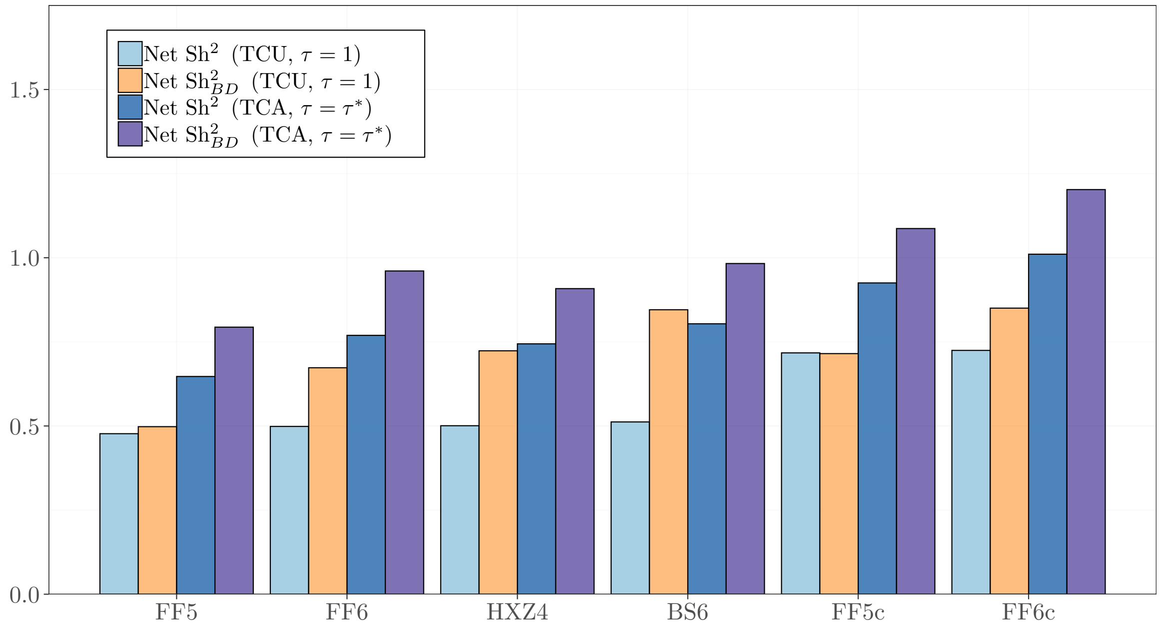

Figure 7: Robustness to banding. The figure compares annualized net Sh^2 that candidate models achieve under four different scenarios. Starting from the left, the first set of bars (light blue) shows the pricing ability of models including traditional TCU factors, which rebalance fully in each period. The second set of bars (orange) adds banding. The remaining sets of bars show the net Sh^2 of TCA factor models, respectively without and with banding (dark blue and purple). The sample ranges between July 1972 and December 2022.

6.2 Trading Diversification

TCA factors introduced in section 3 are appropriate to represent the opportunity set of investors that trade factors individually. For instance, small investors may be unable to trade factor constituents directly but can gain exposure to individual factors through a combination of ETFs and active factor funds. TCA factors capture the returns such investors can achieve if funds optimize execution and fully pass down the cost of trading, either in the form of investment fees or tracking error.

The above argument suggests that TCA factors presented so far may not be appropriate to capture the cost of capital of sophisticated investors. Such investors would trade factor constituents directly, realize the proceeds from trading diversification, and potentially act as investment intermediaries for less sophisticated agents.

Similar to Detzel et al. (2023), I characterize transaction costs incurred when trading K factors jointly when investors can benefit from trading diversification. The main difference with their approach is that I also allow these investors to optimize rebalancing intensities at the factor level. In this setting, investors have further incentives to adjust trading intensities so that rebalancing trades in a given factor line up more closely with offsetting trades in the remaining ones. This intuition suggests that cost-saving synergies should also emerge when combining netting and TCA construction. Transaction costs in this case take the form:

$$\text{TC}_t^{TD}(\tau, \theta) = \sum_{k=1}^{K} \sum_{i=1}^{N_t} \Big| \theta_k \left[w_{ikt}(\tau_k) - w_{i,k,t-1}(\tau_k) \left(1 + r_{it} - d_{it}\right) \right] \Big| c_{it}. \tag{9}$$

By Jensen's inequality, $\text{TC}_t^{TD}(\tau, \theta)$ sets a lower bound to the cost of trading TCA factors individually.

$$\begin{aligned} \text{TC}_t^{TD}(\tau, \theta) &= \sum_{k=1}^{K} \sum_{i=1}^{N_t} \Big| \theta_k \left[w_{ikt}(\tau_k) - w_{i,k,t-1}(\tau_k) \left(1 + r_{it} - d_{it}\right) \right] \Big| c_{it} \\ &\leq \sum_{k=1}^{K} |\theta_k| \sum_{i=1}^{N_t} \Big| w_{ikt}(\tau_k) - w_{i,k,t-1}(\tau_k) \left(1 + r_{it} - d_{it}\right) \Big| c_{it} \\ &= |\theta|' TC_t(\tau) \end{aligned} \tag{10}$$

I solve again for optimal trading intensities and weights in the ex-post mean-variance efficient portfolio under the assumption that investors can benefit from trading diversification.

$$\text{Sh}^2_{TD} = \underset{\theta, \tau}{max} \left\{ \frac{\mathbb{E}\Big[\theta' f_t(\tau) - \text{TC}_t^{TD}(\tau, \theta)\Big]^2}{\mathbb{V}\Big[\theta' f_t(\tau) - \text{TC}_t^{TD}(\tau, \theta)\Big]} \right\} \tag{11}$$

Figure 8 compares the net Sh^2 candidate models achieve in four scenarios: (i) with TCU factors and without trading diversification, (ii) with TCU factors and trading diversification, (ii) with TCA factors but without trading diversification, and (iii) with both trading diversification and TCA factors. In the HXZ4 and BS6 models, which include factors that are more expensive to trade, the benefits of transaction-cost-aware trading overstate the effects of trading diversification. In the remaining models, the individual impact of TCA construction is comparable to the cost savings from trading diversification. In particular, the FF6c model, which remains the best performing in all four cases, has a virtually equivalent Sh^2 for investors that are restricted from either transaction-cost-aware trading or trading diversification. This observation helps put into perspective the advantages that asset managers can deliver to less sophisticated investors, who may be unable to trade the entire set of factor constituents on the margin. Such advantages may come in the form of cost reduction, rather than through risk-adjusted gross returns - the channel that is typically the object of interest in the mutual fund literature.

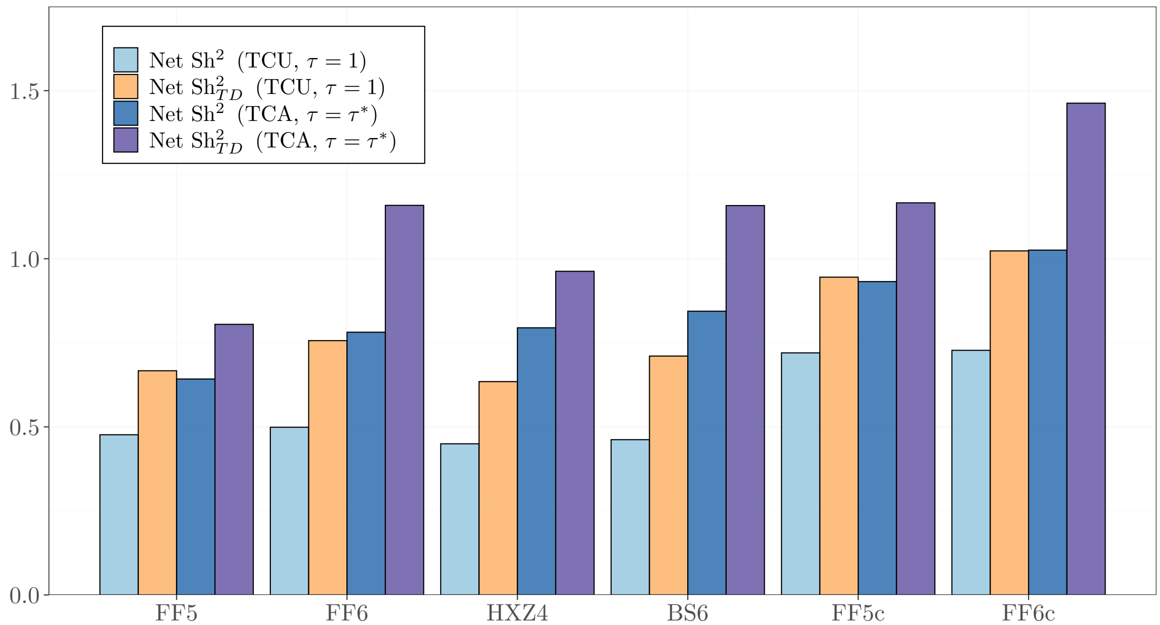

Figure 8: Robustness to trading diversification. The figure compares annualized net Sh^2 that candidate models achieve under four different scenarios. Starting from the left, the first set of bars (light blue) shows the pricing ability of models including traditional TCU factors, which rebalance fully in each period. The second set of bars (orange) adds trading diversification. The remaining sets of bars show the net Sh^2 of TCA factor models, respectively without and with trading diversification (dark blue and purple). The sample ranges between July 1972 and December 2022.

The joint effect of trading diversification and transaction-cost-aware trading further improves the efficient frontier investors can achieve after costs. The dominant model, FF6c, undergoes a 101% Sh^2 with respect to the baseline without transaction-cost-aware trading and trading diversification and performs 42.6% better than the TCA version without TD. The ranking between models also varies from the TCA case without trading diversification. Netting out rebalancing trades across factors naturally favors more parametrized models, since the additional factors introduce additional and potentially offsetting trading motives in the set of constituents. While the FF6c model still dominates the other five candidates, the relative performance of factor models depends on the cost mitigation solutions available to investors. When considering a broader set of models, the tangency portfolio more sophisticated investors can achieve may not only lie higher in the mean-variance plane but may also comprise of a different set of risk-factors. In a similar vein, Li et al. (2023) show that investors with different levels of risk-aversion should benchmark against different factor models when price impact is a concern. Recognizing the effects of transaction costs questions the adequacy of "one-size-fits-all" approaches to factor models.

To qualify asymmetries in relative performance, the FF6 model now outperforms the HXZ4 specification and has a Sh^2 of 1.16, which is equivalent to the Sh^2 of the BS6 model. The FF5c model still outperforms FF6, but only marginally: the distance in Sh^2 between the two shrinks from 0.15 to a mere 0.01. The three models that include the momentum factor - FF6, FF6c, and BS6 - benefit most from trading diversification, as UMD is negatively correlated with the value factor. In my sample, the correlation between momentum and the monthly reconstituted value factor of Asness and Frazzini (2013) is -63%. Overall, the BS6 model is again the one that sees the largest overall performance gains. Its Sh^2 increases by a factor of 2.5 when investors optimize trading intensity and can net out offsetting trades.

7 Conclusion

I show that traditional asset pricing factors are suboptimal if investors incur proportional transaction costs. The cost of trading alters the opportunity set in a fundamental fashion, because it introduces a trade-off between securing risk-factor exposures and controlling rebalancing costs. Factors that are designed while overlooking transaction costs fail to recognize this trade-off, and are unlikely to span the achievable efficient frontier. I instead propose that factors should be constructed in a transaction-cost-aware fashion, evaluating their risk-premia against the necessary cost of trading. I term TCA factors the class of factors incorporating these insights and show that TCA factor models can better characterize the achievable tangency portfolio. Given target weights that provide exposure to a particular characteristic, TCA factors rebalance at the optimal intensity to capture its potential premium, while containing the cost of trading.

TCA factors showcase that factor design is a first-order concern when trading is costly, and meaningful construction can trump the benefits of adopting potentially more parametrized asset pricing models. Out of the set of factor models considered, TCA versions deliver up to 150% larger net squared Sharpe ratios compared to the TCU benchmark.

More importantly, I suggest that discretionary construction choices can bias asset pricing inference. After recognizing the cost of trading, models differ in their relative performance depending on whether trading intensity is optimized or not. This is because, in turn, factors differ in turnover, return persistence, and average cost of constituents. When rebalancing is too aggressive, transaction costs can mask factor premia and dilute the efficiency gains such factors deliver when they are included in asset pricing models. The effect is particularly apparent with the momentum factor. Due to its high cost of trading, momentum plays a marginal role in the ex-post efficient mean-variance portfolio when it is constructed in a TCU way. I find instead that a more conservative rebalancing schedule attributes far greater importance to the momentum factor.

This paper offers a general cautionary note against neglecting frictions in empirical asset pricing research. Investment decisions that are optimal absent frictions may significantly underperform after considering implementation concerns. Consequently, investors modify their optimal allocations to account for friction-induced distortions. Efforts to characterize the opportunity set that either ignore frictions entirely, or restrict investors from optimizing accordingly, may produce misleading results.

References

Abdi, F. and A. Ranaldo (2017): "A simple estimation of bid-ask spreads from daily close, high, and low prices," *The Review of Financial Studies*, 30, 4437–4480.

Arnott, R., F. Li, and J. Linnainmaa (2024): "Smart Rebalancing," *Financial Analysts Journal*, 80, 26–51.

Asness, C. and A. Frazzini (2013): "The devil in HML's details," *Journal of Portfolio Management*, 39, 49–68.

Baker, M. and J. Wurgler (2000): "The equity share in new issues and aggregate stock returns," *the Journal of Finance*, 55, 2219–2257.

Barillas, F., R. Kan, C. Robotti, and J. Shanken (2020): "Model comparison with Sharpe ratios," *Journal of Financial and Quantitative Analysis*, 55, 1840–1874.

Barillas, F. and J. Shanken (2017): "Which alpha?" *The Review of Financial Studies*, 30, 1316–1338.

——— (2018): "Comparing asset pricing models," *The Journal of Finance*, 73, 715–754.

Basu, S. (1983): "The relationship between earnings' yield, market value and return for NYSE common stocks: Further evidence," *Journal of Financial Economics*, 12, 129–156.

Chernov, M., M. Dahlquist, and L. A. Lochstoer (2024): "An anatomy of currency strategies: The role of emerging markets," Tech. rep., Unpublished Working Paper.

Chung, K. H. and H. Zhang (2014): "A simple approximation of intraday spreads using daily data," *Journal of Financial Markets*, 17, 94–120.

Cochrane, J. H. (2011): "Presidential address: Discount rates," *The Journal of Finance*, 66, 1047–1108.

Collin-Dufresne, P., K. Daniel, and M. Sağlam (2020): "Liquidity regimes and optimal dynamic asset allocation," *Journal of Financial Economics*, 136, 379–406.

Collin-Dufresne, P., K. D. Daniel, and M. Saglam (2022): "Optimal Dynamic Asset Allocation with Transaction Costs: The Role of Hedging Demands," *Swiss Finance Institute Research Paper*, 23.

Daniel, K. and T. J. Moskowitz (2016): "Momentum crashes," *Journal of Financial economics*, 122, 221–247.

Daniel, K., L. Mota, S. Rottke, and T. Santos (2020): "The cross-section of risk and returns," *The Review of Financial Studies*, 33, 1927–1979.

Daniel, K. and S. Titman (2006): "Market reactions to tangible and intangible information," *The Journal of Finance*, 61, 1605–1643.

——— (2016): "Another look at market responses to tangible and intangible information," *Critical Finance Review*, 5, 165–175.

DeMiguel, V., A. Martin-Utrera, F. J. Nogales, and R. Uppal (2020): "A transaction-cost perspective on the multitude of firm characteristics," *The Review of Financial Studies*, 33, 2180–2222.

Detzel, A., R. Novy-Marx, and M. Velikov (2023): "Model Comparison with Transaction Costs," *Journal of Finance*, 78, 1743–1775.

Dickerson, A. and Y. Nozawa (2024): "Pricing Corporate Bonds with Credit Risk Primitives," .

Dickerson, A., Y. Nozawa, and C. Robotti (2023): "The Low Frequency Trading Arms Race: Machines Versus Delays," , Unpublished Working Paper.

Ehsani, S. and J. T. Linnainmaa (2022): "Time-series efficient factors," , Unpublished Working Paper, Tuck School of Business.

Fama, E. F. (1991): "Efficient capital markets: II," *The Journal of Finance*, 46, 1575–1617.

Fama, E. F. and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33, 3–56.

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116, 1–22.

——— (2018): "Choosing factors," *Journal of Financial Economics*, 128, 234–252.

——— (2020): "Comparing cross-section and time-series factor models," *The Review of Financial Studies*, 33, 1891–1926.

Gârleanu, N. and L. H. Pedersen (2013): "Dynamic trading with predictable returns and transaction costs," *The Journal of Finance*, 68, 2309–2340.

Gibbons, M. R., S. A. Ross, and J. Shanken (1989): "A test of the efficiency of a given portfolio," *Econometrica: Journal of the Econometric Society*, 1121–1152.

Harvey, C. R., Y. Liu, and H. Zhu (2016): ". . . and the cross-section of expected returns," *The Review of Financial Studies*, 29, 5–68.

Hasbrouck, J. (2009): "Trading costs and returns for US equities: Estimating effective costs from daily data," *The Journal of Finance*, 64, 1445–1477.

Hou, K., C. Xue, and L. Zhang (2015): "Digesting anomalies: An investment approach," *The Review of Financial Studies*, 28, 650–705.

Jensen, T. I., B. T. Kelly, S. Malamud, and L. H. Pedersen (2022): "Machine learning and the implementable efficient frontier," *Swiss Finance Institute Research Paper*, 22.

Jobson, J. D. and B. Korkie (1980): "Estimation for Markowitz efficient portfolios," *Journal of the American Statistical Association*, 75, 544–554.

Korsaye, S. A., A. Quaini, and F. Trojani (2021): "Smart stochastic discount factors," *Swiss Finance Institute Research Paper*.

Li, S., V. DeMiguel, and A. Martin-Utrera (2023): "Comparing Factor Models with Price-Impact Costs," , Unpublished Working Paper, London Business School.

Lintner, J. (1975): "The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets," in *Stochastic optimization models in finance*, Elsevier, 131–155.

Loughran, T. and J. R. Ritter (1995): "The new issues puzzle," *The Journal of Finance*, 50, 23–51.

Luttmer, E. G. (1996): "Asset pricing in economies with frictions," *Econometrica: Journal of the Econometric Society*, 1439–1467.

Novy-Marx, R. (2012): "Is momentum really momentum?" *Journal of Financial Economics*, 103, 429–453.

Novy-Marx, R. and M. Velikov (2016): "A taxonomy of anomalies and their trading costs," *The Review of Financial Studies*, 29, 104–147.

——— (2019): "Comparing cost-mitigation techniques," *Financial Analysts Journal*, 75, 85–102.

Rattray, S., N. Granger, C. R. Harvey, and O. Van Hemert (2019): "Strategic rebalancing," , Unpublished Working Paper.

Ritter, J. R. (1991): "The long-run performance of initial public offerings," *The Journal of Finance*, 46, 3–27.

Ross, S. A. (1976): "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, 13, 341–360.

Sabbatucci, R. (2015): "Are dividends and stock returns predictable? New evidence using M&A cash flows," Tech. rep., Unpublished Working Paper, University of California.

Sharpe, W. F. (1964): "Capital asset prices: A theory of market equilibrium under conditions of risk," *The Journal of Finance*, 19, 425–442.

Sloan, R. G. (1996): "Do stock prices fully reflect information in accruals and cash flows about future earnings?" *Accounting Review*, 289–315.

Spiess, D. K. and J. Affleck-Graves (1995): "Underperformance in long-run stock returns following seasoned equity offerings," *Journal of Financial Economics*, 38, 243–267.

Stigler, G. J. (1963): "Public regulation of the securities markets," *Bus. Law.*, 19, 721.

Online Appendix for “Transaction-cost-aware Factors”

A Stock level transaction costs

Measuring the cost of trading factors requires proportional cost estimates at the stock and month level. Chung and Zhang (2014) suggest that daily quoted spreads provide reliable estimates of high-frequency effective spreads. Abdi and Ranaldo (2017) show that CRSP quoted spreads outperform other more sophisticated estimators and recommend adopting the Chung and Zhang (2014) estimator when quote data is available.[1]

I estimate c_{it} from CRSP, using daily quoted bid-ask spreads, if available. In the absence of valid quotes, I employ the CHL estimator of Abdi and Ranaldo (2017). I then fill c_{it} for stock and months that still have missing values based on the methodology proposed in Novy-Marx and Velikov (2016).

A.1 Quoted Spreads

I construct quoted spread estimates following Chung and Zhang (2014). I discard days with non-positive close, bid, or ask prices. I further ensure that bid-ask spreads are non-negative for each observation. The relative bid-ask half-spreads c_{itd} are:

$$c^Q_{itd} = \frac{A_{itd} - B_{itd}}{2M_{itd}} \tag{12}$$

where A_{itd} and B_{itd} are the closing ask and bid prices quoted on day d of month t for stock i. I denote $M_{itd} = (A_{itd} + B_{itd})/2$ the prevailing end-of-day mid-quote. Following Chung and Zhang (2014), I then take c^Q_{it} as the average of c^Q_{itd} estimates over month t, after discarding days with half-spreads exceeding 25% of the mid-quote.

[1] Abdi and Ranaldo (2017) find that the monthly CRSP quoted spread estimator achieves a 96% correlation with TAQ effective spreads and the same mean (0.82%) between October 2003 and December 2015. For comparison, the Gibbs estimator of Hasbrouck (2009), which has seen frequent application in the literature, delivers a correlation of only 40% with the TAQ effective spread, and overestimates its mean by 1.31%.

A.2 CHL Estimator

I compute a second set of effective spread estimates, c_{it}^{CHL}, using the methodology proposed by Abdi and Ranaldo (2017). I use the 2-day corrected version of the estimator, as per the authors' recommendations. I discard observations with non-positive close, high, or low prices, and stock-months with less than 12 valid observations. The proportional cost estimator is then

$$c_{it}^{CHL} = \frac{1}{2D_t} \sum_{d=1}^{D_t} \sqrt{\max\{(p_{itd} - \eta_{itd})(p_{itd} - \eta_{i,t,d+1}), 0\}} \tag{13}$$

where p_{itd} and η_{itd} are respectively the log closing price and the log mid-range $\eta_{itd} = (\log(H_{itd}) + \log(L_{itd}))/2$ on day d. If the leading midrange $\eta_{i,t,d+1}$ is missing, I use the prevailing log midpoint instead, as proposed by Abdi and Ranaldo (2017).

A.3 Imputation

I set c_{it} to c_{it}^{Q}, if available, and use the CHL estimator c_{it}^{CHL} otherwise. This procedure still leaves missing observations for 2.9% of stock months.[2] I fill these observations with the non-parametric methodology proposed in Novy-Marx and Velikov (2016). I impute missing c_{it} with the c_{it}^{Q} of stock j that minimizes the distance

$$\sqrt{(\text{rankME}_{it} - \text{rankME}_{jt})^2 + (\text{rankFF3IVOL}_{it} - \text{rankFF3IVOL}_{jt})^2} \tag{14}$$

where ME is market equity and FF3IVOL is the idiosyncratic volatility from FF3 time-series regressions estimated over three months of daily data.[3]

[2] In my main sample, which runs from July 1972 to December 2022, 84.6% of observations have valid c_{it}^{Q}. Further, an additional 12.4% of observations have a missing quoted spread estimate, but a valid CHL estimate is instead available.

[3] Stock j must be a common stock and must be trading regularly on NYSE, NASDAQ or AMEX.

B Factor construction and replication

B.1 TCU factors

I replicate before-cost returns on TCU factors according to the instructions available on the authors' webpages. Price and market equity data are from CRSP, while accounting signals are available on the annual and quarterly Compustat releases.

Table A.1: Replication Quality. The table below reports replication statistics. The sample ranges from July 1972 to December 2022. Columns 2 and 3 show the average monthly premium μ on the original factor and the replicated estimate μ^r, in percentage points. Column 4 reports the correlation between the two time-series. Column 5 shows the R^2 from time-series regressions of the original factors on the replicated ones. I report t-statistics in brackets.

	μ (%)	μ^r (%)	ρ	R^2
MKT	0.57	0.57	1	1
	(3.06)	(3.07)	(4281.1)	
SMB	0.16	0.14	1	0.99
	(1.33)	(1.12)	(261.69)	
HML	0.33	0.37	0.99	0.99
	(2.62)	(2.94)	(220.4)	
RMW	0.3	0.32	0.99	0.98
	(3.23)	(3.31)	(180.9)	
CMA	0.33	0.3	0.98	0.96
	(4.04)	(3.71)	(125.92)	
UMD	0.63	0.65	1	0.99
	(3.53)	(3.68)	(317.82)	
ME	0.24	0.24	0.98	0.96
	(1.91)	(1.9)	(122.48)	
IA	0.39	0.37	0.97	0.93
	(4.68)	(4.18)	(91.48)	
ROE	0.53	0.56	0.98	0.95
	(4.99)	(5.28)	(111.55)	
HMLm	0.35	0.38	0.96	0.93
	(2.3)	(2.42)	(89.07)	

The replication methodology for the Fama-French factors follows Fama and French (2018) for RMWc and the documentation on Kenneth French's website (https://

mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html) for the remaining factors. I instead follow the notes on Lu Zhang's web page (http://global-q.org/) for the HXZ4 factors and Asness and Frazzini (2013) for HMLm.[4] Table A.1 reports replication statistics.

B.2 Characteristic signals in cost-aware factors

TCA factors target characteristic-sorted portfolios that reconstitute every month. Restrictions on the available asset universe, the sorting methodology, and the characteristics entering each sort match the original TCU factors. However, I revise the computation of characteristics that do not update at a monthly frequency in the original papers. Sorts in month t use contemporaneous market data. I instead update annual accounting characteristics at a six-month lag. Stocks with valid characteristics and fiscal year end at $t-6$ enter the asset universe at the end of month t, and stocks without valid data for months between $t-18$ and $t-6$ drop out. For characteristics based on quarterly accounting data, I use information as of the most recent public quarterly earnings announcement date, as in Hou et al. (2015).

- *Market equity (ME)* - Price times share outstanding, summed across all firm securities. Market equity must be positive to be considered nonmissing. In the sort for month t, size is the contemporaneous ME.
- *Book equity (BE)* - I compute book equity following Fama and French. Book equity is stockholder equity, minus the book value of preferred stock, plus balance sheet deferred taxes (if available), minus investment tax credit (if available). Stockholder equity is the first available value out of (i) shareholder equity, (ii) common equity plus the book value of preferred stocks, and (iii) total assets minus total liabilities. The book value of preferred stock is the redemption, liquidation, or par value, in this order of preference. Investment tax credit is deferred taxes and investment tax credit, or deferred taxes plus investment tax credit, in this order of preference. Investment tax credit only enters the book value computation up to the 1992 fiscal year. Book equity must be positive to be considered nonmissing.
- *Book-to-market (BM)* - In the sort for month t, book-to-market is the ratio of BE at the last available fiscal year end between $t-18$ and $t-6$ and ME at t.
- *Operating profitability (OP)* - Operating profitability is operating profits divided by BE plus minority interest (if available). Operating profits are the difference between total revenue and the sum of cost of goods sold, interest expenses, and

[4] Before-cost HMLm returns are available at https://www.aqr.com/Insights/Datasets.

selling, general, and administrative expenses. In the sort for month t, I take OP at the latest available fiscal year end between $t-18$ and $t-6$. I annualize OP in cases where firms alter their fiscal year ends, and discard firm-years in which the gap between a fiscal year end and the following exceeds 24 months.

- *Investment (INV and I/A)* - Investment is the growth rate of total assets. In the sort for month t, I take INV at the latest available fiscal year end between $t-18$ and $t-6$. I annualize INV in cases where firms alter their fiscal year ends and discard firm-years if the gap between a fiscal year end and the following exceeds 24 months. I/A is the negative of INV.
- *Return on equity (ROE)* - Return on equity is quarterly income before extraordinary items over BE lagged one quarter. In the sort for month t, quarterly income is considered nonmissing if the relative fiscal quarter end is within six months of t.

C Additional Results

Table A.2: Bootstrap model comparisons for TCU factor models. The table below presents net-of-cost bootstrap comparisons between candidate models that use TCU factors. All figures reported are based on 10,000 simulation runs. I first partition the 606 sample months, July 1972 to December 2022, into adjacent pairs: (1,2), (3,4), ..., (605, 606). Each simulation run then draws 303 pairs with replacement and randomly assigns one month within each pair to the in-sample (IS) period and the partner month to the out-of-sample (OS) period. If the same pair is sampled more than once, I keep the same IS/OS split within this pair. The first and second columns report the average net Sh^2 estimated for each model. Columns 2 to 6 report the % of simulation runs in which the row model outperforms the column model. The last column shows the % of simulation runs in which the corresponding model dominates. Panel A presents results for IS simulation, while Panel B for OS simulations.

Panel A: In-Sample TCU Bootstrap Results

		Probability that the Row Model outperforms the Column Model (%)					
	Average Sh^2	BS6	FF5	FF5c	FF6	FF6c	Best
HXZ4	0.70	4.3	50.1	22.3	39.5	16.1	1.6
BS6	0.81		64.5	31.6	53.9	23.6	21.9
FF5	0.70	35.5		2.5	16.3	1.4	0.4
FF5c	0.95	68.4	97.9		81.7	21.8	16.4
FF6	0.78	46.1	83.7	18.3		2.6	1.3
FF6c	1.01	76.4	98.6	78.2	97.5		58.4

Panel B: Out-of-Sample TCU Bootstrap Results

		Probability that the Row Model outperforms the Column Model (%)					
	Average Sh^2	BS6	FF5	FF5c	FF6	FF6c	Best
HXZ4	0.35	63.8	59.6	27.1	60.6	29.4	20.2
BS6	0.30		50.9	18.5	51.5	18.4	5.6
FF5	0.31	49.1		5.7	46.6	11.2	2.1
FF5c	0.51	81.5	94.3		91.2	53.0	38.4
FF6	0.30	48.5	53.4	8.8		5.6	2.0
FF6c	0.48	81.6	88.8	47.0	94.4		31.7