Online Supplement Material

Cost-Aware Generative Portfolios

Not for Submission

OA.1. Other Results

OA.1.1. Asset Pricing Implications

The main paper establishes that cost-aware training improves net-of-cost performance and alters the composition of basis-portfolios. These findings raise a deeper question: does incorporating transaction costs in Panel Trees impact the pricing information in the generated basis-portfolios? Since each Panel Tree generates a set of tradable leaf basis-portfolios, its maximum-SR allocation defines a tangency portfolio that serves as a pricing-kernel proxy; see Cochrane (2009). Therefore, incorporating transaction costs may change the pricing-kernel proxy implied by the Panel Tree. We address this question in two steps. First, we examine whether the net-of-cost returns generated by the cost-aware and cost-ignorant Panel Trees are spanned by standard asset pricing factor models. Second, we run cross-strategy spanning regressions on the gross and net strategy variants to examine whether the cost-aware Panel Tree shifts the payoff space in economically implementable directions.

OA.1.1.1. Factor Spanning Tests To assess the asset pricing content of the two Panel Tree strategies, we regress each strategy's gross and net returns on a comprehensive set of published factor models and test whether the intercept α in the time-series regression $R_t = \alpha + \boldsymbol{\beta}' \boldsymbol{F_t} + e_t$ is significantly different from zero. A failure to reject a zero alpha is consistent with the factor set spanning the strategy's average return, whereas a significantly positive alpha indicates incremental reward that the factors cannot explain.

Table OA.1 reports the monthly intercepts (%) from time-series regressions of the cost-ignorant (CI) and cost-aware (CA) strategy returns on six factor models. The specifications are CAPM (Sharpe 1964), FF3 (Fama and French 1992), FF5 (Fama and French 2015), FF5 augmented with momentum, short-term reversal, and liquidity (Pástor and Stambaugh 2003), Q5 (Hou et al. 2015), and ALL. The ALL specification combines the augmented FF5 factors with the Q5 factors, for 12 factors in total. Panel A reports results under the maximum-SR objective, and Panel B reports results under MVU with $\delta = 2$. Within each panel, the gross columns report the factor-adjusted returns before transaction costs, whereas the net columns show the factor-adjusted returns after transaction costs. All t-statistics use Newey–West standard errors (Newey and West 1987).

Panel A reveals a sharp gross-to-net contrast under the maximum-SR objective. Before costs, both strategies earn significantly positive alphas in every specification. The CI gross alpha ranges from 0.39% to 0.60% per month, while the CA gross alpha ranges from 0.60% to 0.77%. After realized transaction costs, however, the CI intercept ranges from −0.16% to 0.04%. It is close to zero under CAPM and FF3, turns negative as more sophisticated factors are added, and is significantly negative at the 5% level under Q5 and ALL. By contrast, the CA net alpha remains positive and statistically significant at the 1% level among these specifications. It declines from 0.43% ($t = 5.01$) under CAPM to 0.26% ($t = 3.02$) under the ALL specification, indicating that the broader factor set cannot absorb all of the cost-aware strategy's net return. Thus, a high gross alpha does not by itself imply implementable factor-adjusted performance: only the CA strategy retains a reliably positive intercept after transaction costs.

Panel B shows a similar contrast under MVU ($\delta = 2$). Before costs, the CI and CA alphas range from 2.02% to 3.69% and from 2.53% to 3.51%, respectively, and all are statistically significant at the 1% level. After costs, the CI alpha ranges from −0.17% to 1.38%. It remains significantly positive under CAPM and FF3 but becomes statistically insignificant from zero once more factors are included. The CA net alpha, in contrast, remains between 1.17% and 2.15%. For example, it equals 1.98% ($t = 3.59$) under FF5, 1.69% ($t = 2.98$) under the augmented FF5 model, 1.17% ($t = 2.02$) under Q5, and 1.30% ($t = 2.35$) under the ALL specification.

Across the six benchmark factor models, the CA net alpha remains positive and statistically significant under both portfolio objectives. These models include market, size, value, profitability, investment, momentum, short-term reversal, liquidity, return on equity, and expected growth factors. The CA strategy's *annualized* net alphas range from 3.1% to 5.2% under the maximum-SR objective and from 14.0% to 25.8% under MVU. The CI net alphas become negative under richer models in Panel A and remain significantly positive only under CAPM and FF3 in Panel B.

These findings add an implementation dimension to Cong et al. (2025), who show that Panel Tree test assets expand the mean–variance frontier and generate returns unexplained by benchmark pricing models. Before costs, both the CI and CA strategies display the factor-adjusted performance associated with Panel Tree basis-portfolios. Cost-aware training preserves positive factor-adjusted performance after transaction costs across the six benchmark models.

OA.1.1.2. Cross-Strategy Spanning Tests The factor model tests treat observable factors as the benchmark. Next, we ask whether the cost-aware and cost-ignorant strategies span each other. If the two Panel Trees simply recover the same underlying payoff space, with cost awareness only

Table OA.1 Factor-Model Regressions for Panel Tree Returns

Factor model	CI Gross	CA Gross	CI Net	CA Net
Panel A: **Maximum Sharpe Ratio**				
CAPM	0.60***	0.77***	0.04	0.43***
	(9.31)	(8.71)	(0.62)	(5.01)
FF3	0.58***	0.74***	0.02	0.41***
	(9.55)	(8.67)	(0.26)	(4.87)
FF5	0.48***	0.65***	-0.08	0.31***
	(8.32)	(7.91)	(-1.37)	(3.87)
FF5 Augmented	0.45***	0.64***	-0.12*	0.29***
	(7.72)	(7.60)	(-1.95)	(3.61)
Q5	0.39***	0.60***	-0.16**	0.26***
	(5.98)	(6.99)	(-2.31)	(3.14)
ALL	0.39***	0.60***	-0.15**	0.26***
	(6.43)	(6.64)	(-2.40)	(3.02)
Panel B: **Mean–Variance Utility** ($\delta = 2$)				
CAPM	3.69***	3.44***	1.38***	2.07***
	(7.46)	(7.70)	(2.89)	(4.59)
FF3	3.61***	3.51***	1.31***	2.15***
	(7.47)	(7.44)	(2.77)	(4.51)
FF5	2.91***	3.37***	0.60	1.98***
	(6.05)	(6.15)	(1.27)	(3.59)
FF5 Augmented	2.47***	3.08***	0.17	1.69***
	(5.05)	(5.46)	(0.35)	(2.98)
Q5	2.33***	2.53***	0.12	1.17**
	(4.76)	(4.35)	(0.25)	(2.02)
ALL	2.02***	2.65***	-0.17	1.30**
	(4.19)	(4.74)	(-0.36)	(2.35)

Notes. This table reports the monthly alphas (%) from time-series regressions of gross and net Panel Tree strategy returns on benchmark asset pricing factors. CI and CA denote cost-ignorant and cost-aware Panel Trees. Gross and Net denote gross and net strategy returns. Panel A reports results under the maximum Sharpe ratio, and Panel B reports results under mean–variance utility. In each panel, rows 1–5 use, respectively, CAPM; the Fama–French three-factor model (FF3); the Fama–French five-factor model (FF5); FF5 augmented with momentum, short-term reversal, and liquidity; Q5, which contains the q-factor market, size, investment, return-on-equity, and expected-growth factors. The last row ALL includes all 12 factors. Newey–West t-statistics are in parentheses. All specifications use 480 months from January 1984 through December 2023. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

reducing turnover, then each strategy's return should be spanned by the other, and the cross-strategy alphas should be zero. By contrast, if the basis-portfolios generated by the cost-aware Panel Tree expand the mean–variance frontier of the cost-ignorant Panel Tree, the spanning pattern will be asymmetric.

To formalize this test, we construct four strategy return series: Cost-Ignorant Gross, Cost-Ignorant Net, Cost-Aware Gross, and Cost-Aware Net. Table OA.2 reports cross-strategy spanning regressions across these four variants. Each cell corresponds to a time-series regression $Y_t = \alpha + \beta X_t + e_t$, where the row identifies the test asset Y_t and the column identifies the spanning factor X_t. The failure to reject the null hypothesis of zero α is consistent with the column strategy explaining the expected returns of the row strategy, whereas a significant α rejects the null.

The results have three implications. First, the gross-versus-net comparison of the same strategy quantifies the implementation wedge attributable to transaction costs. Under maximum SR, the intercept is 0.57% ($t = 18.10$) for CI Net explaining CI Gross, and 0.36% ($t = 19.92$) for CA Net explaining CA Gross. Under MVU, the corresponding estimates are 2.32% ($t = 17.29$) and 1.41% ($t = 20.65$). The reverse regressions produce intercepts of similar magnitude with the opposite sign. This pattern is consistent with lower transaction costs for the cost-aware strategies.

Second, the two gross strategies fail to span one another. Under maximum SR, regressing CI Gross on CA Gross yields an intercept of 0.37% ($t = 5.59$), while the reverse regression yields 0.48% ($t = 5.76$). Under MVU, the corresponding intercepts are 2.14% ($t = 3.39$) and 1.94% ($t = 3.19$). The bidirectional rejections indicate that cost awareness changes the composition of the underlying basis-portfolios. This finding corroborates the structural evidence in the Panel Tree Structure section of the main paper: the cost-aware and cost-ignorant Panel Trees select different characteristics and partition the cross section differently, leading to distinct payoff spaces even before transaction costs are deducted.

Third, the net-space comparison reveals an asymmetry. Under maximum SR, CI Net regressed on CA Net produces an intercept of −0.09% ($t = -1.49$), whereas CA Net regressed on CI Net produces 0.42% ($t = 5.41$). Under MVU, the corresponding estimates are 0.49% ($t = 1.00$) and 1.67% ($t = 3.71$). Thus, when CI is treated as a return factor, CA earns a positive and precisely estimated alpha under both objectives; when CA is treated as a return factor, the CI intercept is insignificant from zero. It implies that the CA strategy contains incremental pricing information that is not spanned by the CI strategy.

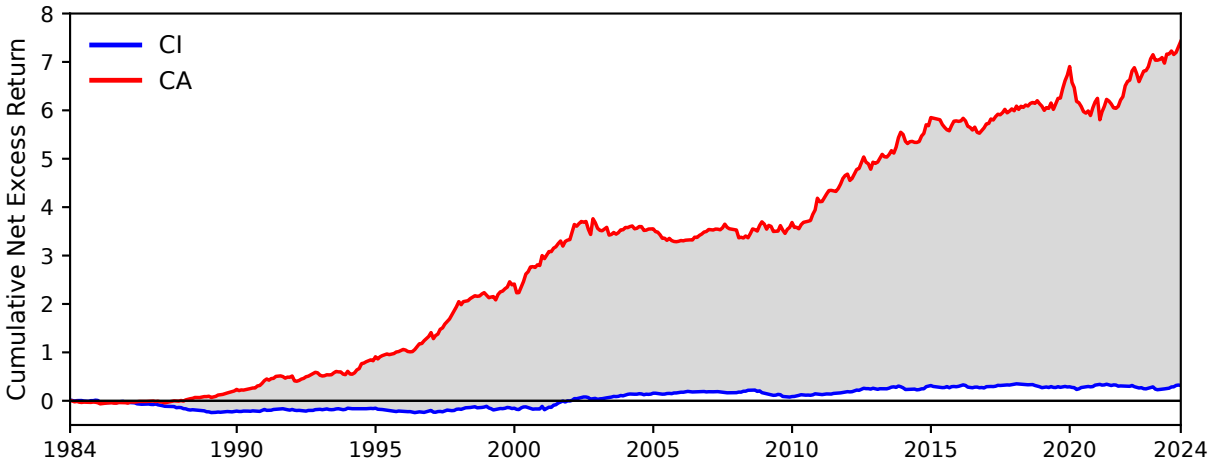
Table OA.2 Cross-Strategy Spanning Tests

	Explanatory Strategy (factor)			
Dependent strategy	CI Gross	CI Net	CA Gross	CA Net
Panel A: **Maximum Sharpe Ratio**				
CI Gross	–	0.57***	0.37***	0.48***
		(18.10)	(5.59)	(8.15)
CI Net	-0.56***	–	-0.20***	-0.09
	(-16.70)		(-2.81)	(-1.49)
CA Gross	0.48***	0.76***	–	0.36***
	(5.76)	(9.38)		(19.92)
CA Net	0.13	0.42***	-0.35***	–
	(1.59)	(5.41)	(-18.44)	
Panel B: **Mean–Variance Utility**				
CI Gross	–	2.32***	2.14***	2.81***
		(17.29)	(3.39)	(5.53)
CI Net	-2.27***	–	-0.16	0.49
	(-17.27)		(-0.25)	(1.00)
CA Gross	1.94***	3.07***	–	1.41***
	(3.19)	(6.84)		(20.65)
CA Net	0.53	1.67***	-1.40***	–
	(0.87)	(3.71)	(-20.62)	

Notes. This table reports pairwise spanning regressions across four strategy return series: gross and net returns of the cost-ignorant Panel Tree (CI Gross and CI Net) and gross and net returns of the cost-aware Panel Tree (CA Gross and CA Net). Panel A reports results under the maximum Sharpe ratio, and Panel B reports results under mean–variance utility. In each regression, the row strategy is the dependent variable and the column strategy is the explanatory return factor. Each off-diagonal entry reports the monthly intercept (in percent), followed by its Newey–West t-statistic in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

OA.1.1.3. Summary Across the six benchmark factor models, the CI net alpha is not robust: it becomes negative under sophisticated factor models for maximum SR and remains significantly positive only under CAPM and FF3 for MVU. By contrast, the CA net alpha remains positive and statistically significant across all six models under both portfolio objectives. The cross-strategy regressions show a similar pattern: the CA net return earns a positive intercept when the CI net return serves as the explanatory factor, while the reverse intercept is statistically insignificant from zero. Together, these tests indicate that the CA strategy contains incremental pricing information not captured by either the benchmark factor models or the CI strategy returns.

Figure OA.1 Cumulative Net Excess Returns

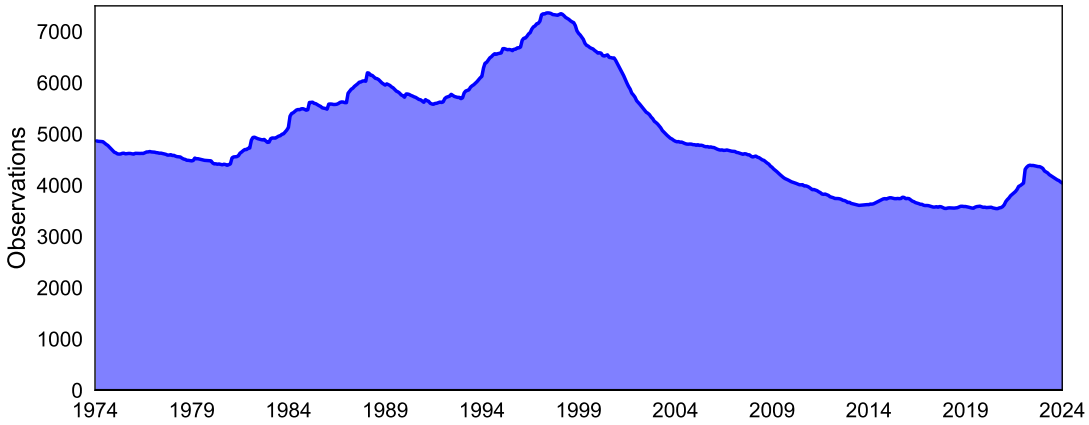


Note. This figure reports the out-of-sample cumulative net excess returns of two strategies. The blue line is for the cost-ignorant (CI) Panel Tree, while the red line is for the cost-aware (CA) Panel Tree. The out-of-sample period spans from 1984 to 2023. The gray shaded area highlights the difference in performance between the two strategies over time, illustrating the value added by internalizing transaction costs into the objective function.

OA.1.2. Cumulative Wealth

Figure OA.1 reports the cumulative net excess returns of the cost-ignorant and cost-aware Panel Trees under the maximum-Sharpe objective over the 1984–2023 out-of-sample period. The shaded region highlights the cumulative performance gap between the two strategies and shows how the monthly net-performance differences accumulate over time. The cumulative net excess return of the cost-aware strategy begins to separate from that of the cost-ignorant strategy in the late 1980s. The gap widens further through the 2000s and 2010s, and remains substantial at the end of the sample. By 2023, the cost-aware Panel Tree reaches 742.3% in cumulative net excess return, compared with 31.8% for the cost-ignorant Panel Tree.

Figure OA.2 The Number of Stocks



Note. This figure shows the number of common stocks in our monthly sample from January 1974 to December 2023. The sample is constructed following Jensen et al. (2023) using CRSP–Compustat data and is restricted to U.S.-traded common stocks with non-missing lagged market equity and returns, and positive book equity and lagged market equity.

OA.2. Stock Data

We use the publicly available dataset from Jensen et al. (2023), which combines return data from CRSP with accounting data from Compustat. We restrict our sample to common stocks (shrcd: 10 and 11) traded on USA exchanges (`excntry = USA`). To be included in our sample, stocks must have non-missing data on lagged market equity and returns, as well as positive book equity and lagged market equity.

Our monthly sample starts in January 1974 and ends in December 2023. The number of stocks in each month is shown in Figure OA.2. The average and median numbers of monthly stock observations are 5,551 and 5,597 for the first 25 years, and 4,405 and 4,103 for the latter 25 years. We use the 153 characteristics studied in Jensen et al. (2023), and each characteristic is standardized cross-sectionally to the $[0, 1]$ interval. For missing characteristic values, we first impute using the industry median based on SIC, NAICS, and GICS classifications; if a value remains missing, we then impute using the cross-sectional median. We apply cross-sectional winsorization at the 1% and 99% levels to mitigate the impact of outliers.

OA.3. Tree Structure under MVU

This section examines whether the reallocation of split capacity documented under the maximum Sharpe-ratio objective also appears under mean–variance utility (MVU) with $\delta = 2$. The empirical implementation allows at most nine splits, but an MVU tree stops earlier when no admissible split improves the objective. Across the 40 rolling windows, the cost-ignorant (CI) and cost-aware (CA)

Table OA.3 Reallocation of MVU Split Capacity across Rolling Windows

Panel A: **Aggregate Split Allocation by Meta-category**

	Avg. frequency		Frequency share (%)		Avg. score		Score share (%)	
Meta-category	CI	CA	CI	CA	CI	CA	CI	CA
Accounting-fundamental	4.80	5.80	58.72	70.30	22.38	28.63	51.59	65.54
Price-volume	3.38	2.45	41.28	29.70	21.00	15.05	48.41	34.46

Panel B: **Largest Category-level Changes (CA minus CI)**

Largest increases	Δ Freq.	Δ Score	Largest decreases	Δ Freq.	Δ Score
Investment	+0.55	+2.58	Momentum	-0.30	-2.23
Quality	+0.23	+2.48	Short-Term Reversal	-0.38	-2.10
Value	+0.65	+2.10	Profit Growth	-0.25	-1.20
Low Leverage	+0.08	+1.00	Size	-0.23	-1.03

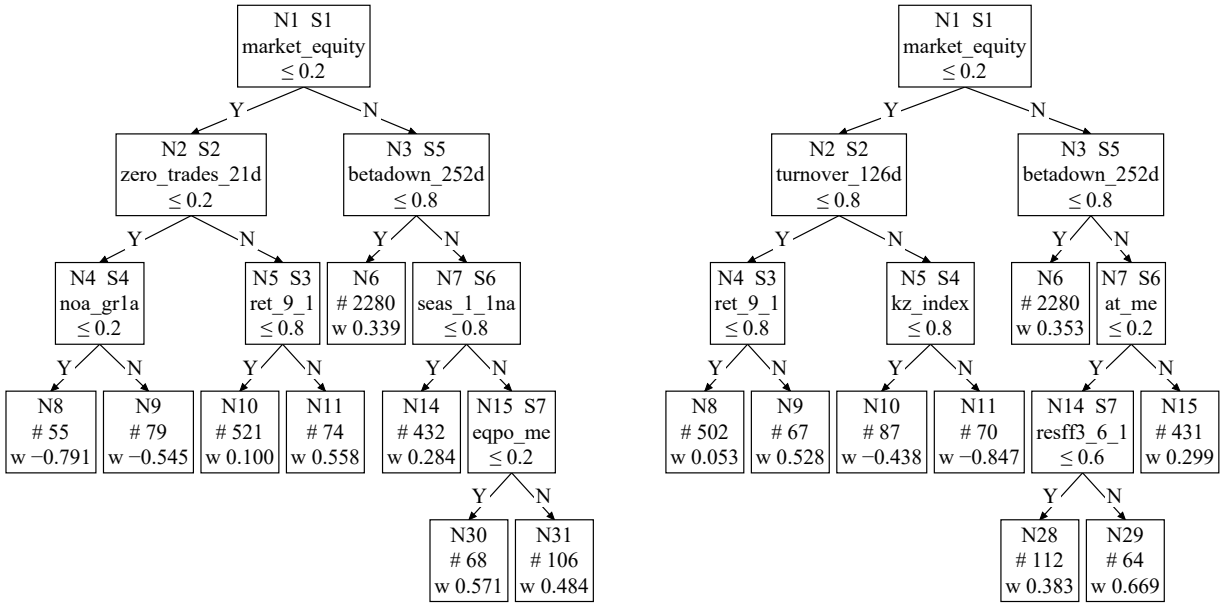
Notes. This table summarizes realized split allocation under the cost-ignorant (CI) and cost-aware (CA) MVU trees, averaged over 40 rolling windows. Panel A aggregates splits into two meta-categories: accounting-fundamental and price-volume. Frequency shares divide category-level average frequency by each strategy's average realized number of splits. Within each rolling window, the split score is 9 for the first split, 8 for the second, and so on down to 1 for the ninth split; absent splits receive zero. Panel B reports the four largest increases and decreases in average score (CA minus CI), together with the corresponding changes in average split frequency.

trees average 8.18 and 8.25 realized splits, respectively. We therefore measure frequency shares relative to each strategy's average realized number of splits and assign a score of zero to unfilled split positions.

Table OA.3 shows that cost-aware training tilts realized split capacity toward accounting-fundamental signals. Panel A reports that their frequency share rises from 58.72% under CI to 70.30% under CA, while their score share rises from 51.59% to 65.54%. Because the score measure places greater weight on earlier splits, the 13.95 percentage-point increase in score share, compared with the 11.58 percentage-point increase in frequency share, indicates that the shift occurs near the top of the tree as well as in the number of selections.

Panel B identifies the categories underlying this reallocation. The largest score increases occur in Investment (+2.58), Quality (+2.48), Value (+2.10), and Low Leverage (+1.00), while the largest declines occur in Momentum (−2.23), Short-Term Reversal (−2.10), Profit Growth (−1.20), and Size (−1.03). Figure OA.4 reports the corresponding category-level frequencies and shares, and Figure OA.5 traces the CA-minus-CI score differences across rolling windows. The direction of the reallocation therefore parallels the maximum-SR evidence, although the category ranking varies with the portfolio objective.

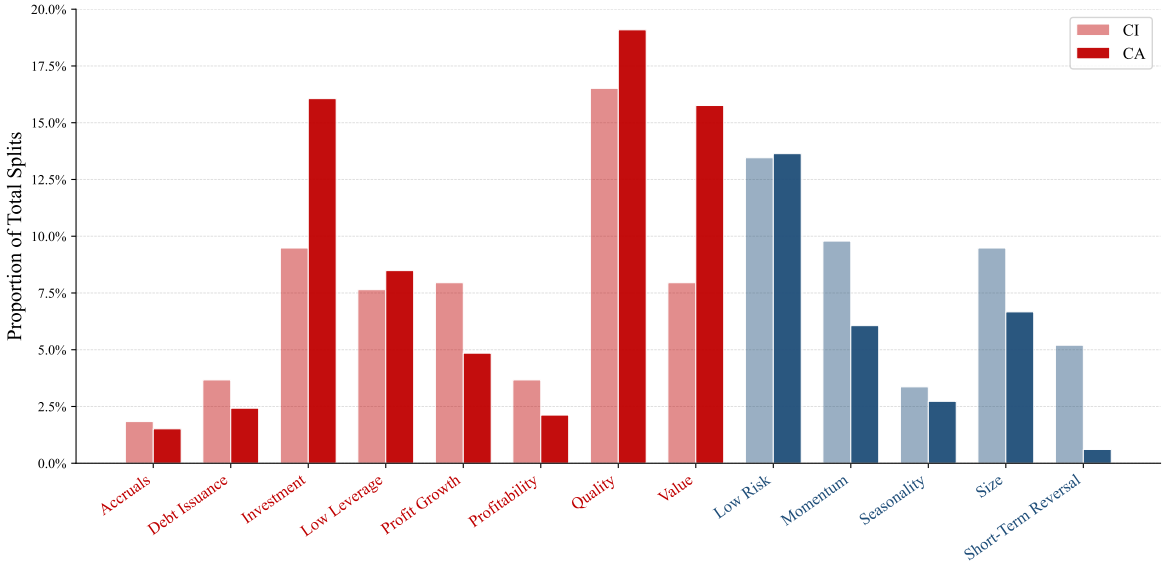
Figure OA.3 Panel Tree Diagrams under the Mean–Variance Objective



(a) Cost-ignorant (b) Cost-aware

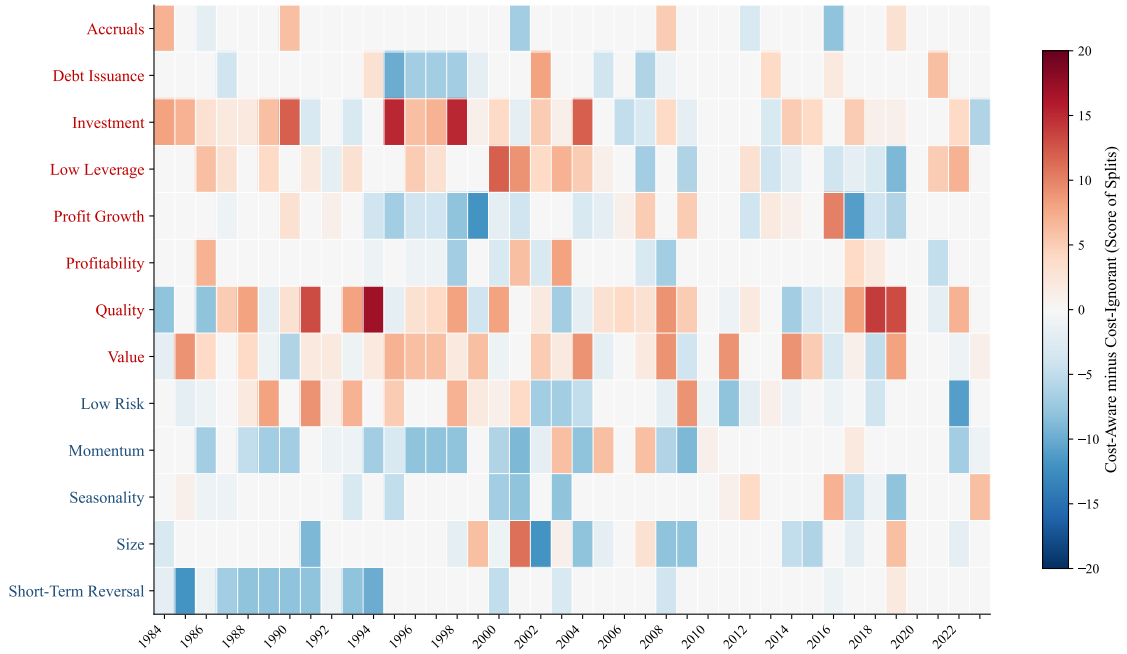
Note. This figure compares the cost-ignorant (Panel a) and cost-aware (Panel b) Panel Trees under the mean–variance objective with $\delta = 2$, with no more than 10 leaves, estimated from January 2013 to December 2022. Nodes display the selected characteristics, cutoff values, node indices (N#), and split order (S#). Terminal leaves display the median monthly number of assets (#) and portfolio weights (w). Characteristic definitions and categories follow Jensen et al. (2023).

Figure OA.4 MVU Split Frequency and Shares by Characteristic Category



Note. This figure reports the average proportion of total splits assigned to each characteristic category under the cost-aware (CA) and cost-ignorant (CI) Panel Trees. The averages are computed across 40 rolling windows, and each tree is allowed at most nine splits per window. The figure provides an intuitive category-level view of how fixed split capacity is reallocated once transaction costs are internalized.

Figure OA.5 MVU Difference in Score-Weighted Split Importance: CA Minus CI



Note. This figure reports, for each rolling window and characteristic category, the difference in split scores between the cost-aware and cost-ignorant Panel Trees. Within each rolling window, the split score is 9 for the first split, 8 for the second, and so on down to 1 for the ninth split; absent splits receive zero. Positive values indicate that the category becomes more important under cost-aware training. Because score is a split-order-weighted measure of split importance, the figure captures not only whether a category is selected but also whether it is used higher in the tree.

References

Cochrane JH (2009) *Asset pricing: Revised edition* (Princeton University Press).

Cong LW, Feng G, He J, He X (2025) Growing the efficient frontier on panel trees. *Journal of Financial Economics* 167:104024.

Fama EF, French KR (1992) The cross-section of expected stock returns. *The Journal of Finance* 47(2):427–465.

Fama EF, French KR (2015) A five-factor asset pricing model. *Journal of Financial Economics* 116(1):1–22.

Hou K, Xue C, Zhang L (2015) Digesting anomalies: An investment approach. *Review of Financial Studies* 28(3):650–705.

Jensen TI, Kelly B, Pedersen LH (2023) Is there a replication crisis in finance? *The Journal of Finance* 78(5):2465–2518.

Newey WK, West KD (1987) A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55(3):703–708.

Pástor L, Stambaugh RF (2003) Liquidity risk and expected stock returns. *Journal of Political Economy* 111(3):642–685.

Sharpe WF (1964) Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19(3):425–442.