Spectral Risk Factors and the Limits of Spanability

Daniel Batista*[1,2]

[1]Université de Genève
[2]Swiss Finance Institute

March 23, 2026

Preliminary draft – please do not circulate without permission.

Abstract

Standard asset-pricing factors are usually treated as single return series, yet the risks they summarize need not be homogeneous across horizons. I decompose the market, momentum, and value factors into orthogonal spectral components using a downsampled Haar discrete wavelet transform and estimate sparse mimicking portfolios for each band with elastic net and band-specific complexity controls. The results show clear heterogeneity across factors. Market risk is broadly distributed across horizons and remains economically trackable even as performance weakens at longer horizons. Momentum is largely a short-horizon object: its high-frequency components are tracked reasonably well, but spanability deteriorates quickly as the horizon lengthens. Value looks different from both. Its more informative variation shifts toward medium and lower frequencies, but that variation is also less evenly represented by traded returns. The clearest boundary case is HML D5, corresponding roughly to 32–64 trading days, for which the full-sample procedure does not recover a stable mimicking portfolio. Evidence of spanability is also stronger on the component's natural wavelet grid than after translation back to daily frequency.

*Email: daniel.batistadasilva@unige.ch

*Replication files and code available upon request.

1 Introduction

Standard factor models compress systematic risk into a small number of return series (Fama and French, 1993, 2015). That compression is useful, but it also masks an economic distinction that is central to how risk is transmitted through markets. Some shocks arrive and dissipate quickly, reflecting liquidity conditions, trading imbalances, or rapid information diffusion. Others evolve more slowly, through changes in expectations, discount rates, or broader macro-financial conditions. Once these sources of variation are aggregated into a single factor return, it becomes difficult to tell which part of the factor is tied to short-horizon dynamics, which part reflects slower-moving risk, and whether markets are equally able to represent both.

This paper studies how much of the horizon-specific variation in market, momentum, and value can be represented by traded returns. I decompose each factor into orthogonal components using a downsampled Haar discrete wavelet transform and then estimate a mimicking portfolio for each band. Throughout the paper, D1 through D6 denote dyadic bands ranging from roughly 2–4 trading days to 64–128 trading days. The objective is not to propose a new stochastic discount factor, introduce a new family of priced factors, or evaluate welfare gains from horizon-specific trading strategies. It is to recover the temporal structure of familiar risks and to identify where traded representations become less stable.

In this setting, spanability refers to empirical accessibility. A component is more spanable when a stable traded proxy can track it closely and when standard spanning restrictions are difficult to reject at the relevant horizon. This distinction matters because a factor may look spanable in the aggregate while containing components that are only imperfectly represented once the factor is separated by horizon. Looking band by band therefore allows the paper to distinguish between the economic existence of a systematic risk and the market's ability to intermediate that risk.

To implement that idea, I estimate a sparse mimicking portfolio for each spectral band from a broad return universe that includes equities, bonds, and macro-financial series. The construction uses elastic net regularization (Zou and Hastie, 2005) with band-specific complexity controls, so that the problem remains commensurate with the sharp decline in effective sample size at slower horizons. All estimation is first conducted on the dates at which the component itself is defined—its natural wavelet grid—and then revisited after the corresponding mimicking portfolios are expressed at the daily frequency. This comparison is important because economic representability at a component's own horizon need not coincide with clean implementation in calendar time.

The evidence reveals a clear economic pattern. The market factor is the least specialized object in the sample. Its spectral variation is spread across horizons, and tracking performance declines gradually rather than collapsing at long horizons. In pseudo-out-of-sample terms, the correlation falls from about 0.70 in D1 to about 0.25 in D6, which still leaves non-trivial long-horizon representation. This is what one would expect if market risk aggregates shocks arriving at different speeds.

Momentum looks very different. Its economically accessible content is concentrated at short horizons. The pseudo-out-of-sample correlation is about 0.41 in D1 and drifts toward zero by D6. More importantly, the short-horizon components are also the ones for which traded returns provide the clearest representation of momentum. The factor becomes progressively harder to span as one moves away from the continuation window that gives it its usual interpretation.

Value is different again. Relative to market and momentum, its more interpretable content shifts toward slower bands, but that shift comes with weaker and less stable tradability. Short-horizon HML components are only modestly tracked. Some slower bands show stronger in-sample alignment, yet the out-of-sample evidence is uneven rather than monotone. The most revealing result is HML D5, corresponding to roughly 32–64 trading days. In the full sample, the estimation procedure does not deliver a stable mimicking portfolio for this component. I read this as evidence that spanability becomes fragile at that horizon: the component is visible in the data, but its traded representation is not robust in the full sample.

A second central result concerns implementation. When the spanning tests are run on the component's natural wavelet grid—that is, with both the left-hand-side component and the traded return series sampled on the dates associated with that band—the joint evidence is often permissive. Several components look broadly consistent with spanning at their own horizon. The picture becomes less favorable when the same mimicking portfolios are taken back to the daily frequency. Economic spanability at a component's own horizon is therefore not the same as clean daily implementability.

The contribution of the paper is mainly economic. It shows that familiar factors bundle together risks that live at different horizons, and that the market's ability to represent those risks is itself horizon-dependent. Spectral decomposition is useful here because it helps separate the parts of factor variation that are easier to represent with traded returns from those that are harder to represent.

The remainder of the paper proceeds as follows. Section 2 places the paper in the literatures on frequency-domain asset pricing, multi-horizon risk, and mimicking portfolios.

Section 3 describes the decomposition, portfolio construction, and data. Section 4 presents the empirical results. Section 5 concludes.

2 Related Literature

This paper sits at the intersection of three literatures. The first is frequency-domain asset pricing. Dew-Becker and Giglio (2016), Bandi et al. (2021), and Neuhierl and Varneskov (2021a,b) show that exposures and risk compensation can differ across frequencies, so the timing of fluctuations is part of the economic content of risk rather than a nuisance detail. This literature mainly studies how asset prices reflect frequency-specific risk. The present paper instead asks how much of that horizon-specific variation can be represented by traded returns.

A related literature studies persistence and the economic distinction between transitory and slow-moving shocks. Ortu et al. (2020) and Ortu et al. (2023) show that components with different persistence can be associated with different mechanisms and different asset-pricing implications. Their results reinforce the idea that aggregating everything into a single time-domain object can conceal economically meaningful heterogeneity. The contribution here is not another persistence decomposition per se, but an empirical map from that heterogeneity to tradable proxies.

Wavelet-based studies of financial data provide a natural bridge between those ideas and the present design. Wavelets are well suited to this setting because they localize variation by horizon while preserving the time ordering of the series. Existing applications use them to study volatility, co-movement, and horizon-specific return behavior, including work on the value premium by Kang et al. (2017). Here they serve a practical role: they separate horizon-specific components of familiar factors and make it possible to study whether those components admit stable traded proxies.

The paper also connects to the literature on mimicking portfolios and tradable proxies. Fama and French (1993) showed long ago that useful factor representations often come from portfolios built to capture an underlying economic exposure. More recent work emphasizes that such proxies should be sparse, stable, and robust to correlation in the underlying asset universe; see, for example, Brodie et al. (2009) and Kelly et al. (2021). The present paper brings that logic to the frequency domain by asking whether the horizon-specific components of familiar factors can be turned into stable traded objects.

3 Empirical Framework and Data

Let F_t denote a daily factor return, such as the market, momentum, or value factor. I decompose F_t into orthogonal detail components and a low-frequency residual using a downsampled Haar discrete wavelet transform,

$$F_t = \sum_{j=1}^{J} D_j(t) + A_J(t).$$

Each $D_j(t)$ isolates fluctuations in a dyadic band. In the present setup, D1 captures movements at roughly 2–4 trading days, D2 captures 4–8 days, and so on, with D5 corresponding to approximately 32–64 days and D6 to 64–128 days.

The choice of a downsampled orthogonal transform is important for the interpretation of the results. Orthogonality allows variance to be assigned cleanly across horizons, so the variance decomposition has a direct economic reading. Downsampling avoids the overlapping structure of redundant transforms, which would otherwise create mechanical serial correlation and make tracking look better than it really is. Because the transform is downsampled, each $D_j(t)$ is observed only every 2^j dates. All band-level estimation and inference are therefore conducted on the natural grid of the component itself rather than on an artificially filled daily series.

The spectral components are statistical objects and are not directly tradable. For each band j, I therefore estimate a mimicking portfolio,

$$\widehat{D}_j(t) = R_t' w^{(j)},$$

where R_t collects returns from a broad traded-return universe and $w^{(j)}$ is chosen to approximate $D_j(t)$ on the dates where that component is defined. The weights are estimated with an elastic net objective (Zou and Hastie, 2005),

$$\min_{w^{(j)}} \left\| D_j - R w^{(j)} \right\|_2^2 + \lambda_j \left[\alpha \left\| w^{(j)} \right\|_1 + (1-\alpha) \left\| w^{(j)} \right\|_2^2 \right],$$

with $\alpha = 0.5$ in the baseline implementation. This is a natural compromise for the present problem. The return universe is high dimensional and strongly correlated, so pure subset selection is unstable, especially in slower bands where the effective sample size is small.

Two further design choices matter. First, I retain assets with up to 50% missing observations within a band and impute missing returns with the cross-sectional mean on that date. This keeps the design matrix rectangular without manufacturing asset-specific signal in the missing entries. Second, I choose λ_j by block cross-validation and average weights across repeated block partitions. The point of the exercise is not to find the single best-fitting in-sample portfolio, but to recover a stable representation of the component.

Model complexity is tied explicitly to the effective sample size n_j of each band. As j increases, the number of usable observations falls sharply. To keep the problem commensurate with the amount of information in the data, I restrict the maximum number of active positions to

$$dfmax_j = \lfloor \sqrt{n_j} \rfloor$$

and pre-screen the candidate set to the $2dfmax_j$ assets with the strongest marginal correlation with $D_j(t)$. These controls become tighter in slower bands by construction. Economically, this matters because low-frequency spanability should not be declared on the basis of extremely flexible portfolios that are only possible because the estimator is allowed to overfit a small number of observations.

The resulting portfolios are evaluated along two dimensions. The first is tracking performance, measured primarily by the correlation between $D_j(t)$ and $\widehat{D}_j(t)$ on the component's natural grid, together with block pseudo-out-of-sample correlations. These are the central objects in the paper because they speak directly to whether a horizon-specific risk is economically accessible. Hedging regressions are reported as a secondary consistency check: if a component is well tracked, the residual variance left after projecting it on the mimicking portfolio should also be lower.

The second dimension is spanability. I run band-by-band intercept tests and joint GRS tests using standard traded factors on the right-hand side. These tests are implemented both on the wavelet grid, where the regressions are run on the dates associated with each band, and after translating the mimicking portfolios back to daily frequency. The comparison is deliberate. Wavelet-grid tests ask whether the component is representable at its own horizon. Daily tests ask whether that representation survives in calendar time.

The return universe combines CRSP common-stock returns with bond and macro-financial return series, including corporate bond indexes, discount-bond returns, and macro-financial instruments that are informative about slower-moving risks. Because the orthogonal wavelet transform requires dyadic sample lengths, each factor is truncated to the largest power-of-two window available in the data. In the current implementation, that yields 2^{14} daily observations for market and momentum and 2^{13} observations for HML. The effective estimation window therefore differs across factors even though the underlying raw data are broader.

Because the return universe is assembled to maximize coverage across bands, the exercise is best interpreted as evidence on representability across horizons rather than as a literal implementation backtest.

4 Results

The evidence is easiest to read as an economic map of where factor risk lives and how much of that risk can be represented by traded returns. The full set of tables is reported in the Appendix. The main text focuses on the pattern that survives across them.

Table 1 shows that, variance-wise, the three factors look more similar than different: most of their variation sits in the shorter bands. The later differences in tracking and spanability therefore do not come from a trivial case in which one factor is mechanically short-horizon and another is mechanically long-horizon. They emerge when one asks how much of that variation can be turned into a stable traded proxy.

The market factor is the least specialized object in the sample. Its high-frequency components are tracked well, but the decline with horizon is gradual rather than abrupt. Table 2 shows pseudo-out-of-sample correlations of about 0.70 in D1, 0.49 in D2, 0.39 in D3, and still around 0.25 in D5 and D6. The market therefore looks like a genuinely multi-horizon source of systematic risk. Some parts of it are naturally easier to span than others, but no single horizon defines the factor.

Momentum is much more concentrated in short-horizon tradable variation. The pseudo-out-of-sample correlations fall from about 0.41 in D1 to 0.27 in D2, then continue to fade until they are essentially zero in D6. That pattern fits the usual economic interpretation of momentum better than a generic statement about "high-frequency risk." More importantly, the short-horizon components are also the ones that traded returns capture most clearly.

Value is slower-moving, but the right summary is "slower and unevenly spanable," not simply "low-frequency and well captured." Its short-horizon components are only modestly tracked. Some slower bands line up more strongly in-sample, yet the out-of-sample evidence is uneven. In Table 2, D4 posts a relatively strong in-sample correlation but a much weaker pseudo-out-of-sample correlation, while D6 shows the clearest slower-band alignment but also sizeable uncertainty across folds. The most informative result is D5. In the full sample, the estimation procedure does not select a stable mimicking portfolio for HML D5, and the corresponding hedging entries are absent in Table 3. At the same time, cross-validation still produces a weak average correlation in that band. I therefore interpret D5 as evidence of fragile or incomplete spanability rather than as proof that the component is economically irrelevant. This is precisely the kind of boundary case the paper is designed to identify.

Table 3, which reports hedging regressions, tells the same story and is best read as a consistency check rather than as a separate contribution. Hedge effectiveness decays gradually with horizon for market, fades more quickly for momentum, and is uneven for HML, with no stable full-sample hedge for D5.

Tracking and spanability are related but distinct. A component may be tracked well enough to be economically meaningful without satisfying exact spanning restrictions, and a component may appear spanable on its natural grid without yielding a convincing daily implementation. The spanning tests help organize that distinction.

On the wavelet grid, the joint GRS results in Table 6 are broadly permissive. For the true components and for the mimicking portfolios, the joint null of zero intercepts is generally not rejected. That does not mean the grid-based evidence is perfectly clean: the band-by-band intercepts in Tables 4 and 5 still reject in a number of cases, especially for market-related objects. Interpreted cautiously, the joint evidence suggests that when the regressions are run on the dates associated with each band, a sizeable part of the variation can be represented within the existing traded-factor space.

The picture is less favorable once the same mimicking portfolios are evaluated at the daily frequency. Table 6 shows that the daily GRS rejects for market, while the reverse daily tests reject for all three factor families. The broader point is that a component may look spanable when evaluated on its own grid and still be harder to reproduce as a daily return series.

This distinction is central to the paper's interpretation of limits to spanability. The wavelet grid is the right place to evaluate whether a horizon-specific risk is representable in the traded-return span in principle. The daily grid is the right place to evaluate whether investors can actually carry that risk as a practical object. The gap between the two is not a technical nuisance. It is the empirical signature of implementation frictions.

HML D5 is again the most revealing example. It is a meaningful spectral object, but the data do not support a stable full-sample traded proxy for it. In that sense, D5 is not simply a weaker version of market D6 or momentum D5. It is a case in which the limit to spanability becomes visible in the estimation itself.

Taken together, the results point to a simple interpretation of standard factors. The market factor aggregates systematic risk arriving at many speeds. Momentum is primarily a short-horizon phenomenon whose economic content weakens rapidly outside that window. Value contains slower-moving variation, but precisely those slower components are less completely intermediated by traded returns. Spectral decomposition is useful here not because it merely reorganizes time-series variation, but because it separates the part of factor variation that markets can readily intermediate from the part that remains only imperfectly represented. The limits of spanability are therefore not random estimation failures. They vary systematically with both the factor under study and the horizon at which the risk materializes.

5 Conclusion

This paper studies the extent to which horizon-specific components of standard asset-pricing factors can be represented by traded returns. Using a downsampled orthogonal wavelet decomposition and sparse mimicking portfolios, it shows that the answer differs sharply across factors and across horizons.

Market risk is broadly distributed and remains economically trackable even at slower horizons, although tracking weakens as the horizon lengthens. Momentum is much more short-horizon, both in its spectral structure and in the part of that structure that can be represented by traded returns. Value is slower-moving, but that slower-moving content is also the least fully intermediated. The clearest manifestation is HML D5, a medium-horizon component for which the full-sample procedure does not deliver a stable traded proxy.

The paper also highlights a distinction that is easy to miss in time-domain analyses. Spanability on a component's natural wavelet grid is not the same thing as spanability in daily implementation. A risk can appear economically representable at the horizon where it lives and still become difficult to carry as a daily object. That gap is part of the economic result, not merely a statistical inconvenience.

Because the asset universe is assembled to study representability across horizons, the results are best read as evidence on the boundaries of spanability rather than as a literal trading backtest. Within that scope, the paper shows that standard factors are internally heterogeneous and that their tradability varies systematically with horizon. Future research can study how these boundaries change under alternative real-time investable universes and implementation protocols.

References

Bandi, F.M., Chaudhuri, S.E., Lo, A.W., Tamoni, A., 2021. Spectral factor models. Journal of Financial Economics 142, 214–238.

Brodie, J., Daubechies, I., De Mol, C., Giannone, D., Loris, I., 2009. Sparse and stable markowitz portfolios. Proceedings of the National Academy of Sciences 106, 12267–12272.

Dew-Becker, I., Giglio, S., 2016. Asset pricing in the frequency domain: Theory and empirics. Review of Financial Studies 29, 2029–2068.

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

Fama, E.F., French, K.R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116, 1–22.

Kang, J.W., In, F., Kim, T.H., 2017. Horizon-specific analysis of the value premium: A wavelet approach. Working Paper .

Kelly, B., Pruitt, S., Su, Y., 2021. The three and a half factor model. Working Paper.

Neuhierl, A., Varneskov, R.T., 2021a. Asset pricing with spectral factors. Journal of Financial Economics 140, 495–518.

Neuhierl, A., Varneskov, R.T., 2021b. Frequency-dependent risk. Journal of Financial Economics 140, 644–675.

Ortu, F., Reggiani, P., Severino, F., 2023. Persistence-based capital allocation along the fomc cycle. Working Paper .

Ortu, F., Severino, F., Tamoni, A., Tebaldi, C., 2020. A persistence-based wold-type decomposition for stationary time series. Quantitative Economics 11, 203–230.

Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society Series B 67, 301–320.

Appendix

Tables and Figures

Table 1: Variance Decomposition of Risk Factors by Spectral Band

	Share of Total Variance (%)
Panel A: Market (MKT)	
D1	50.2
D2	23.9
D3	13.6
D4	6.3
D5	2.7
D6	1.8
Residual	1.5
Panel B: Momentum (MOM)	
D1	41.2
D2	27.4
D3	14.1
D4	8.1
D5	3.9
D6	3.1
Residual	2.1
Panel C: Value (HML)	
D1	47.6
D2	26.7
D3	12.5
D4	6.5
D5	2.8
D6	1.8
Residual	2.1

Notes: This table reports the share of total variance of each risk factor accounted for by its spectral components D_j and the residual. Variance shares are computed using the orthogonal wavelet decomposition and therefore sum to approximately 100% for each factor.

Table 2: Tracking Performance of Spectral Components

	In-sample Corr.	OOS Corr.	OOS SD
Panel A: Market (MKT)			
D1	0.716	0.696	0.037
D2	0.524	0.489	0.084
D3	0.463	0.391	0.071
D4	0.345	0.217	0.098
D5	0.385	0.251	0.112
D6	0.466	0.248	0.193
Panel B: Momentum (MOM)			
D1	0.447	0.405	0.051
D2	0.331	0.266	0.070
D3	0.371	0.193	0.072
D4	0.389	0.238	0.084
D5	0.300	0.173	0.101
D6	0.276	-0.027	0.137
Panel C: Value (HML)			
D1	0.385	0.345	0.036
D2	0.314	0.248	0.077
D3	0.335	0.004	0.035
D4	0.588	0.131	0.185
D5	—	0.114	0.165
D6	0.713	0.450	0.333

Notes: This table reports the correlation between each spectral component D_j of the target factor (Market, Momentum, or Value) and its corresponding mimicking portfolio $\hat{D}_j$. In-sample correlations are computed using all available dates for each component. Pseudo-out-of-sample (pseudo-OOS) correlations are obtained via block cross-validation to account for time dependence. Reported standard deviations reflect variability across validation folds. Missing entries indicate components for which no mimicking portfolio is selected.

Table 3: Hedging Effectiveness of Spectral Risk Components

	β	t_{HAC}	R^2 (HE IS)	HE (OOS)
Panel A: Market (MKT)				
D1	0.93	14.2	0.51	0.48
D2	0.88	9.7	0.27	0.23
D3	0.81	7.8	0.21	0.16
D4	0.69	5.1	0.12	0.08
D5	0.61	3.9	0.15	0.10
D6	0.54	2.6	0.22	0.14
Panel B: Momentum (MOM)				
D1	0.82	8.9	0.20	0.17
D2	0.77	6.1	0.11	0.09
D3	0.74	5.4	0.14	0.10
D4	0.70	4.7	0.15	0.11
D5	0.66	3.2	0.09	0.06
D6	0.58	2.1	0.08	0.04
Panel C: Value (HML)				
D1	0.79	7.4	0.15	0.13
D2	0.75	5.9	0.10	0.08
D3	0.71	4.8	0.11	0.02
D4	0.67	3.9	0.35	0.12
D5	—	—	—	—
D6	0.62	2.7	0.51	0.28

Notes: This table reports hedging regressions of each spectral component D_j on its corresponding mimicking portfolio $\hat{D}_j$. The hedge ratio β is estimated by OLS, with t-statistics computed using Newey–West HAC standard errors. R^2 corresponds to in-sample hedging effectiveness (HE IS), defined as one minus the ratio of residual variance to the variance of D_j. Out-of-sample hedging effectiveness (HE OOS) is computed via block cross-validation. Missing entries indicate components for which no stable mimicking portfolio is selected in the full sample.

Table 4: Spanning of True Spectral Components (Grid)

	Market (MKT)	Momentum (MOM)	Value (HML)
D1	3.57*	-0.71	-3.48**
	(1.87)	(-0.53)	(-2.09)
D2	2.17	-0.67	-0.71
	(1.08)	(-0.46)	(-0.38)
D3	0.86	-1.16	-3.36
	(0.40)	(-0.65)	(-1.48)
D4	-3.99*	0.30	-1.16
	(-1.78)	(0.16)	(-0.53)
D5	2.40	-0.05	-2.01
	(1.27)	(-0.03)	(-0.96)
D6	-3.53*	4.32*	0.01
	(-1.77)	(1.72)	(0.06)

Notes: Entries report annualized intercepts (in percent) from spanning regressions of D_j on traded factors, evaluated on the wavelet grid. The right-hand-side factor set is the Fama–French six-factor set excluding the target factor. t-statistics (Newey–West HAC) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Spanning of Mimicking Portfolios (Grid)

	Market (MKT)	Momentum (MOM)	Value (HML)
D1	7.91***	0.24	-0.66*
	(6.08)	(1.29)	(-1.74)
D2	4.69***	0.03	-0.16
	(5.48)	(0.18)	(-0.55)
D3	1.84**	-0.06	0.32
	(2.19)	(-0.40)	(0.69)
D4	0.78	0.07	-0.57
	(1.50)	(0.41)	(-0.92)
D5	0.35	0.05	—
	(0.80)	(0.41)	
D6	-0.14	0.15	0.05
	(-0.61)	(0.31)	(0.08)

Notes: Entries report annualized intercepts (in percent) from spanning regressions of $\hat{D}_j$ on traded factors, evaluated on the wavelet grid. The right-hand-side factor set is the Fama–French six-factor set excluding the target factor. t-statistics (Newey–West HAC) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Joint GRS Tests of Spanning

	True Components (Grid)	Mimicking (Grid)	Mimicking (Daily)	Reverse (Daily)
Market (MKT)	Fail to Reject	Fail to Reject	Reject***	Reject***
Momentum (MOM)	Fail to Reject	Fail to Reject	Fail to Reject	Reject***
Value (HML)	Fail to Reject	Fail to Reject	Fail to Reject	Reject***

Notes: This table reports joint Gibbons–Ross–Shanken (GRS) tests for spanning. The tests evaluate whether the intercepts from regressions of spectral components (true or mimicking) on traded factors are jointly zero. The right-hand-side factor set is the Fama–French six-factor set excluding the target factor. Reverse tests regress the target factor on the set of mimicking portfolios. ***, **, and * denote rejection at the 1%, 5%, and 10% significance levels, respectively.