Observation-driven Cross-Asset Value*

Florian Ielpo† Julien Royer‡

October 2025

Abstract

We propose a novel approach to generating cross-asset valuation factors based on a trend-cycle decomposition. We compare a model-agnostic econometric approach with more recent deep-learning techniques, to estimate trends in the log price of a vast group of assets and indices. We show how the resulting cyclical component can be used as a value signal. This value signal positively correlates to known value signals used in the industry and in the academia, while showing a robust and positive Sharpe ratio. The resulting value factor negatively correlates to timeseries momentum factors, as usual for value factors.

Keywords: Trend filtering; Value; Factor Investing; Deep Learning;

*This article owes to a preliminary work of Lambert Njo. The first author wishes to thank him. The authors also thank the participants of the Computational and Financial Econometrics conference in 2024, London, UK.

†Lombard Odier Investment Managers, Avenue des Morgines 6, 1211 Petit Lancy, Switzerland and Centre d'Economie de la Sorbonne, 103 Boulevard de l'Hôpital, 75013 Paris, France. Corresponding author : f.ielpo@lombardodier.com.

‡Lombard Odier Investment Managers, Avenue des Morgines 6, 1211 Petit Lancy, Switzerland.

1 Introduction

Value is one of the first premium that has been uncovered in the financial literature. The idea of a value strategy is simple: assets prices can be disconnected from the fundamental value of the assets. If an investor is able to estimate the unobserved fundamental value of an asset, the spread between the asset price and this estimation allows to identify cheap and expensive assets. By buying cheap assets and selling expensive one, an investor can harvest a premium - the value premium. Since the essential paper "Value and Momentum Everywhere" paper of Asness et al. (2013), the investment community's understanding of hedge fund strategies has progressed enormously. What was previously seen as pure alpha suddenly became commoditized systematic beta, as most asset class seem to be exhibiting the same double patterns. First, assets outperforming and underperforming in the past tend to continue doing so in the coming weeks, forming a momentum effect. Second, assets showing elevated (respectively low) prices in comparison with their history tend to generate negative (respectively positive) returns, forming a value effect. Both effects correlate negatively.

If both effects are now well known, forming value strategy is always a source of complexity for a variety of reasons. First, there is a lack of consensus within the financial community on how to gauge the valuation of an asset. Using financial metrics such as rates or dividend yields shows a lack of consistency across asset classes. Second, within an asset class, common valuation metrics can strongly depend on sectors. For instance, "price-to-book" ratios make sense for a majority of sectors, but much less so for financials. Finally, some of these financial metrics are reliant over expectation variables, such as the price-earnings ratio: practitioners have a tendency of using analysts forecasts instead of past realized values, adding a final layer of complexity. An interesting solution to all of these issues has been proposed in Asness et al. (2013), using asset prices (futures in their context and ours) and dividing these prices by their value 5 years ago. By doing so, the value analyst obtains a mean-reverting ratio across all asset classes, at the cost of assuming that value trends last for an about 5-year period.

The link between "long term" past returns and valorisation is not new and was first documented by Bondt and Thaler (1985) on individual stocks. Basu (1977) who showed that deviations from the CAPM could be explained by valuation metrics. Fama and French (1992) showed that Value was not only a premium but also a factor in the cross-section of US stock returns. Fama and French (1996) find that portfolios formed on book-to-market measures are highly correlated with portfolios formed on a 5-year return metric. Black (1986) actually finds a justification of the Value premium as the counterpart of the Momentum premium. The author postulates that market trends are created by relatively informed trades that push asset prices far from the unobserved fundamental value of an asset, until the spread becomes so large that mean reversion will occur. Another behavioral explanation to the value premium, Lakonishok et al. (1994) finds that the value factor appears because of errors in the expectations of investors, attracted by more glittering strategies. Other explanations have been proposed: from a risk perspective, Fama and French (1992) and Berk et al. (1999) claim that value firms are said to be riskier in bad times due to the more mature businesses in which they operate or due to reversibility costs. Lempérière et al. (2017) show that Value, in the case of equity markets is associated with a positive skewness of which he questions the significance. Straying away from equities, Asness et al. (2013) highlight how, although the rationale across asset classes for Momentum is well established, finding satisfying justification for the existence of the Value factor on other asset classes than equities is difficult.

Measuring valuation as long-term reversal thus appears highly appealing at it provides an "economic model"-free approach only based on price data. An open question[1] however remains on the choice of the econometric specification that will deliver such metric. One of the issues with the 5-year window is its "ad hoc" flavor, but not only. A genuine value measure should produce mean-reverting and stationary signals: the 5-year window ratio happens to deliver some of that, but not in a controlled way. In this article we aim at building upon the intuition that the easiest measure to gauging valuation is the price of assets, consistent with the economic theory. Instead of a lookback period of our own choosing, we propose to perform a classical trend-cycle decomposition of the log of asset prices. The "cycle" part of this decomposition should resemble what Asness et al. (2013) are trying to exploit, only this time in a timeseries model driven way. Several methods could be used to extract such a value signal and non-parametric methods such as the Hodrick-Prescott or the Baxter-King filters which would be natural candidate, shouldn't they exhibit instability in their estimation of such a decomposition on the tails of samples. The right tail is of an utmost interest to us, given that is where the investment signal will sit. Recently, Hamilton (2018) proposed an appealing contender to this filtering

[1] An alternative literature is based on the affine model, exploiting such models to connect fundamentals with rates and earnings dynamics, which drive some sort of fundamental value to equities and bonds. Examples of this literature are Werner and Lemke (2009), Lettau and Wachter (2011), Bekaert and Grenadier (1999), Bansal and Yaron (2004) and Ielpo and Kniahin (2020).

approach, namely an econometric filter. That band-pass filter produces cycle components that are stationary by construction. Our experiment shows how this cycle element is also mean-reverting, with a half-life which is observation-driven instead of being explicitly controlled by a model. In this article, we therefore explore the use of Hamilton's trend/cycle decomposition in order to gauge how one can use it to derive value measures later used to build cross-asset value factor using a consistent metrics across all asset classes.

This article unfolds as follow: Section 2 presents the methodology we retain to build the value signal in a general context. Section 3 presents empirical applications showing the interest of the approach. Section 4 concludes. Robustness checks are relegated to the appendix.

2 The construction of the trend-based value factor

2.1 Building the valuation signal the econometrician way

Let P_t be the price of a given asset or the value of an index, total return or not. Let $p_t = \log P_t$, the log of the asset price or of the value of the index. That variable is typically an integrated variable, exhibiting a stochastic trend or a persistence so near random-walk it is not stationary. Filtering out this stochastic trend yields a cyclical component that is mean-reverting and can be interpreted as the variation around an unobservable fair value of the asset. We thus propose to leverage recent trend filtering techniques to derive a valuation factor[2].

Let us denote φ_t a measurable function with regard to the σ-field $\mathcal{F}_t = \{p_u, u \leq t\}$, parametrized by a vector of parameters θ_0. Trend filtering commonly writes

$$p_t = \varphi_t(\theta_0) + \varepsilon_t \tag{1}$$

where ε_t corresponds to the mean-reverting residuals around the trend component $\varphi_t(\theta_0)$. For example, in the well-known HP filter of Hodrick and Prescott (1997), the parameter θ_0 (usually written λ) is fixed and $\varphi_t^{\mathcal{HP}} = \varphi_t^{\mathcal{HP}}(\theta_0) = \varphi_t^{\mathcal{HP}}(\lambda)$ is

[2]See Hodrick (2020) for a discussion on recent trend filtering techniques.

numerically obtained as

$$\varphi_t^{\mathcal{HP}} = \min_{\varphi_t} \sum_t (p_t - \varphi_t)^2 + \lambda \left[(\varphi_{t+1} - \varphi_t) - (\varphi_t - \varphi_{t-1}) \right]^2 .$$

Note that the dependence of φ_t on both $\varphi_t + 1$ and φ_{t-1} - as any two-sided filter - often leads to distortions at the boundary of the sample. More recently, Hamilton (2018) proposed an alternative trend filter, popularized as the H filter, based on the linear regression between p_t and past values of p_t lagged by multiple quarters. The filter writes as

$$\varphi_t^{\mathcal{H}}(\theta_0) = \beta_0 + \sum_{h=1}^{l} \beta_h p_{t-(h+4)Q} + \varepsilon_t, \tag{2}$$

where $t - qQ$ denotes a lag of q quarters and l denotes the maximum number of quarterly lags considered in the regression set to 4 in the following as recommended in Hamilton (2018). The vector of parameter $\theta_0 = (\beta_0, \beta_1, \dots, \beta_l)' \in \mathbb{R}^{l+1}$ can be estimated either via Ordinary Least Squares or Maximum Likelihood. Let us denote $\widehat{\varphi}_t^{\mathcal{H}} = \varphi_t^{\mathcal{H}}(\widehat{\theta}_n)$ the estimated time-varying trend and $\widehat{\varepsilon}_t$ the residuals

$$\widehat{\varepsilon}_t^{\mathcal{H}} = p_t - \widehat{\varphi}_t^{\mathcal{H}}$$

from which we intend on deriving a valuation metric.

2.2 Building the valuation signal the data scientist way

Although the H filter is well suited for isolating the mean-reverting component from the stochastic trend in asset prices, it has certain limitations. Originally developed for low-frequency economic time series—typically characterized by a higher signal-to-noise ratio—it may be less effective when applied to equity prices, which are notoriously noisy. Furthermore, despite its tractability and ease of estimation, the model's linear structure may be too restrictive to capture the complex, nonlinear dynamics underlying price behavior. To address these challenges, we propose an alternative approach that leverages deep learning techniques to extract residual components more effectively.

Machine learning applications in finance have garnered increasing attention in recent years (see Kelly and Xiu (2023) for a recent review). Neural networks, in particular, have become popular as they have theoretical underpinnings as "universal approximators" for any smooth predictive function (Hornik et al. (1989)). They usually rely on a feed-forward structure, that consists in an “input layer” of raw predictors, one or more “hidden layers” with individual neurons that interact and nonlinearly transform the predictors, and an “output layer” that aggregates hidden layers into an ultimate outcome prediction. Formally, as in Gu et al. (2020), we can define a deep feed-forward neural network in a recursive way. Let y a set of predictors, L the number of hidden layers in the neural network, and $K^{(l)}$ the number of neurons in each layer $l = 1, \ldots, L$, and denote $x^{(l)} = (1, x_1^{(l)}, \ldots, x_{K^{(l)}}^{(l)})'$ the vector of output from each neuron k in layer l. Recursively, each neuron k in layer $0 < l < L$ yields

$$x_k^{(l)} = f(x^{(l-1)'}\theta_k^{(l-1)})$$

where f denotes the activation function and the final output is given by the aggregation function

$$g(y, \theta) = x^{(L-1)'}\theta^{(L-1)}$$

where $\theta_k^{(l-1)}$ and $\theta^{(L-1)}$ are vectors of parameters to estimate.

Recently, Triebe et al. (2019) introduced the AR-net, a particular form of a deep feed-forward neural network specifically designed for time series. Indeed, the predictors in the network are lagged values of the variable to predict, as it is the case for autoregressive models in standard econometrics. Interestingly, the H filter can be represented with an AR-net without any hidden layer and the four past value of the log price, lagged by 5 to 8 quarters, as presented in Figure 1, but the AR-net framework allows for a more flexible representation. For example, the addition of hidden layers, as presented in Figure 2 would allow potential non-linearities.

AR-net are however limited in their ability to model near unit root processes. Additionally, more advanced neural networks such as recurrent neural networks (Lecun et al. (1998)) or long-short term memory neural network (Hochreiter and Schmidhuber (1997)) have been introduced to model time series, but they also struggle to capture long-range dependence as noted by Chien et al. (2021) - a particular drawback as we are interested in a long term trend-cycle decomposition. Recently, Triebe et al. (2021) introduced a deep learning model designed to handle non-stationary time series. Their *Neural Prophet* model can be seen as a hybrid method building upon the classical seasonal autoregressive integrated moving average (SARIMA). In

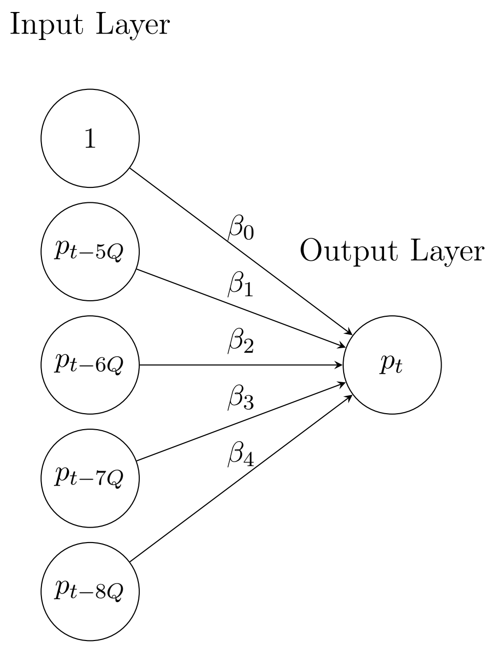

Figure 1: AR-net representation of the H filter of Hamilton (2018)

this model, the log price time series is decomposed into three components[3]

$$p_t = T(t) + A(t) + L(t) + \varepsilon_t$$

where $T(t)$ is a deterministic trend with potential breaks, $A(t)$ is a AR-net capturing the autoregressive nature of p_t, and $L(t)$ is another AR-net measuring the effect of lagged exogenous variables on the price of the asset. We can thus derive deep-learning based residuals by specifying the trend filtering as

$$\varphi^{\mathcal{NP}} = T(t) + A(t) + L(t).$$

2.3 Building the cross-asset value factor

We now detail how we obtain the composition of the proposed cross-asset value factor building on Asness et al. (2013). Let $\hat{\varepsilon}_{it}$ be the estimated residuals for asset i at time t in a set of I assets $(1, 2, 3, \ldots, I)$. As in our model, $\hat{\varepsilon}_{it}$ can be interpreted as the deviation from the "long term" value of the price series, a positive $\hat{\varepsilon}_{it}$ means

[3] In the original Neural Prophet model, the authors consider three additional components: $S(t)$ and $E(t)$ being the seasonal and holiday components that we do not include as they should be irrelevant for financial asset prices, and $F(t)$ a "future-known exogenous variables" that we do not include as it would make our filter non-causal.

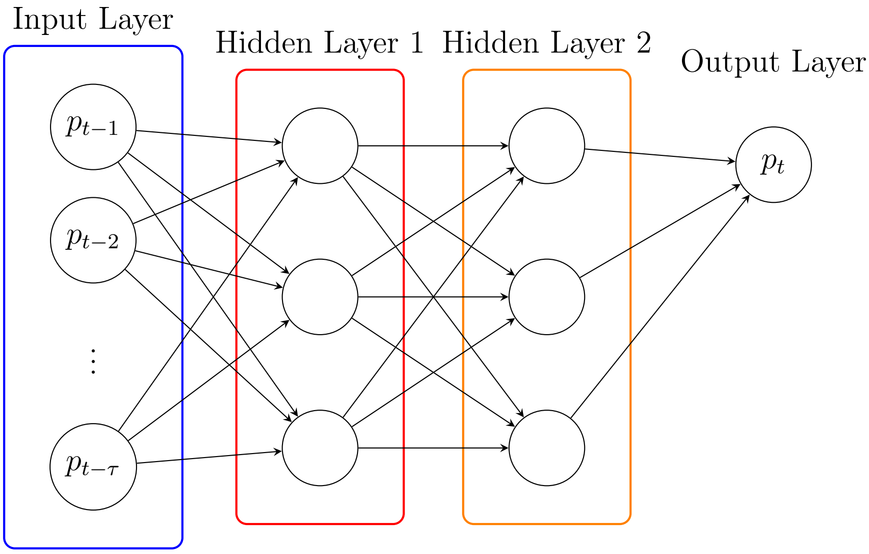

Figure 2: Example of an AR-net with τ lags and two hidden layers with three neurons

an expensive asset, while a negative one means a cheap asset. We thus propose to construct our valuation metrics as

$$s_{it} = -\frac{\hat{\varepsilon}_{it}}{\sigma(\hat{\varepsilon}_{it})} \tag{3}$$

where $\sigma(\hat{\varepsilon}_{it})$ denotes the volatility of the residuals. The negative sign allows to have a positive signal for value asset and a negative one for expensive assets, while the volatility rescaling is necessary to have comparable metrics across different assets and asset classes. Let ω_t^i be the portfolio weight for asset i at time t. It is obtained from the below formulaic expression:

$$\omega_{it} = -c_t \left(\text{rank}(s_{it}) - \frac{1}{I} \sum_{k=i}^{I} \text{rank}(s_{kt}) \right), \tag{4}$$

where rank(.) is the rank function. The weights across assets within the investment universe made of the I assets sum to zero, representing a Dollar-neutral long-short portfolio. c_t is a scaling factor such that the overall portfolio is scaled to one dollar long and one dollar short, as in Asness et al. (2013).

Let r_{it} be the return on asset i at time t. The factor portfolio performance will compute as

$$r_t = \sum_{k=i}^{I} \omega_{it-1} r_{it}. \tag{5}$$

As noted in Asness et al. (2013), using such a ranking function instead of creating simple spread portfolio let the investor benefit from more diversified exposures and less extreme weights.

As now classic in the literature, diagnosing the existence of a cross-asset value factor based on the proposed model requires finding a positive Sharpe ratio, while such a strategy should display a negative correlation to a momentum strategy (Dollar neutral or not). Also note that the long-short setup we have adopted for literature-consistency reasons can easily be modified to obtain a directional value factor by using weights computed as:

$$\omega_{it} = -c_t \times \text{rank}(\hat{\varepsilon}_{it}), \tag{6}$$

with c_t now acting as a control variable for risk and the overall leverage that the factor can reach. In the next section we present empirical applications showcasing how this value factor construction can be used in different settings.

An interesting property of this valuation approach is that we can select l such that p_t can be I(1) and $\hat{\varepsilon}_t$ can be I(0), while displaying some form of persistence, similarly to most mispricing measures. Fitting an Autoregressive process to the estimated $\hat{\varepsilon}_t$ will allow to explicitly measure the persistence and half-life of the value anomaly. This offers an interesting by-product to our approach, coming alongside stationarity.

3 Empirical Application

In this section, we propose different applications of the proposed methodology with the intention of gauging its ability to create value strategy simply from the price or index level of financial assets.

Please note this section does not yet include results based on the deep-learning approach presented in Subsection 2.2. All Value factors are currently computed from the H filter presented in Subsection 2.1. and we are working towards a new version where we compare the two approaches.

3.1 A first illustration

Before commenting the output table from the different value factors we have composed using this methodology, we propose here a short subsection comparing a classic measure of valuation –the dividend yield ratio[4]– to our approach in the pure case of equities. The success of the dividend yield in assessing the valuation of equities owes to two factors: first, its simplicity of computation; second, the fact that it can be compared across most sectors and with long-term yields for relative valuation assessment. It also comes with limitations, as its difficulties of application for companies pertaining to the growth style. Also, when willing to obtain a cross-asset value factor, a similar kind of metric needs to be built for assets such as commodities, which happens to be a non-consensual case in the literature.

Figure 1 below compares the price-to-dividend ratio of the S&P500 and the valuation measure we introduced in Section 2. It has been obtained from 5-year rolling regressions with $i = (1, 2, 3, 4)$, therefore using four different quarters as recommended in Hamilton (2018). This comparison is not quantitative but visual, as a first endeavor to draw a line between what is widely accepted as a valuation measure and what we propose in this article, thus echoing with the 5-year return and book-to-market relationship (see Gerakos and Linnainmaa (2017)). As clear from Figure 3, both time series show a strong covariation. A linear regression to assess how much of the dividend yield variation gets explained by our measure shows an R^2 of 78%, when our measure is purely based on price information.

3.2 Investigating the cross-asset value factor across samples

Next, we created four different datasets composed only with liquid futures, similarly to Moskowitz et al. (2012). These futures relate to four different markets: bonds, equities, currencies, and commodities and are listed in the Appendix. Our database starts on January 2nd, 2000 and ends on January 31st, 2024, encompassing slightly more than 24 years of daily prices. Their list is provided in a table in the appendices. Our value measure is computed using 5 years of rolling overlapping daily data and using 4 quarterly lags, that is $i = (1, 2, 3, 4)$.

Our empirical strategy unfolds as follows: first, we created value portfolios at a single asset class level in order to assess if our approach could generate a value factor

[4] see for example Basu (1983)

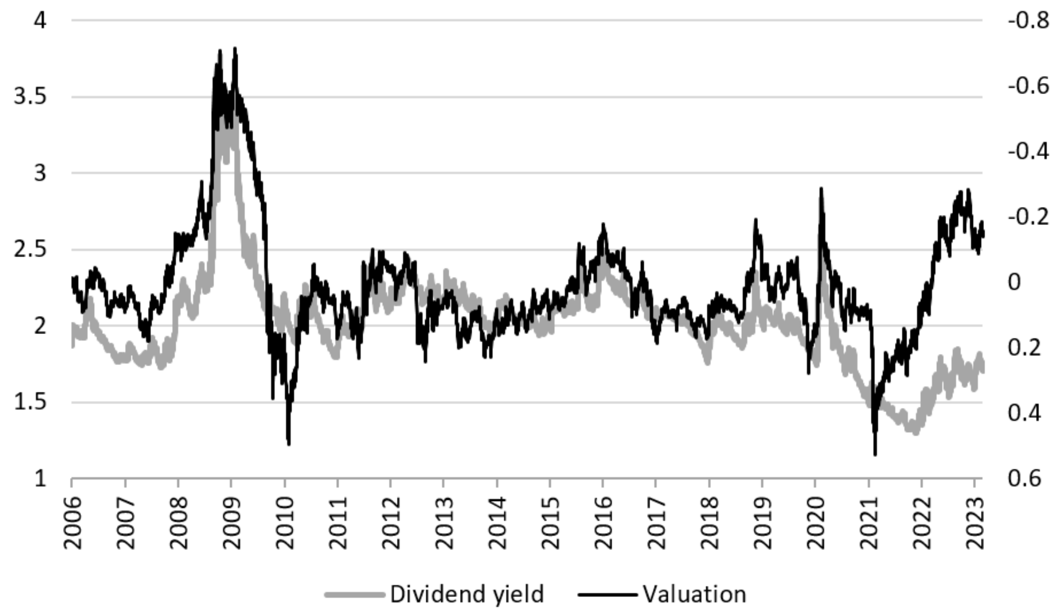

Figure 3: Price/Dividend ratio of the MSCI World vs. the residual from the Hamilton's regression

within asset classes. Then we created value portfolios by progressively combining asset classes together, starting with bonds and equities, then adding commodities and finally currencies. Also, we ran two different kinds of factor composition: one targeting an ex-ante volatility of 15% using an exponentially moving average with parameters 0.94 and one not targeting volatility, for which the 1 dollar long 1 dollar short legs drive the risk of the portfolio.

With these settings, we obtained the results shown in Table 1 and the indexed performances shown in Figure 4. From a single asset class perspective, each value factor sees a positive performance, be it with the volatility targeted version or not. Sharpe ratios fluctuate between 0.11 in the case of the forex market for the volatility targeted version and 1.12 in the case of equities without volatility targeting. The skewness of these value returns has a tendency to be positive: large positive returns are more frequent than large negative returns. This is consistent with what one can expect from a value factor: unless they are value traps, value trades have a tendency of delivering small negative returns (given their short carry nature) and large positives when the value anomaly rapidly gets repriced.

Now looking at the cross-asset setup, the results happen to be relatively similar: the Sharpe ratios are in the region of 0.2 to 0.4, that is to say in line with the

Not Volatility Targeted							
	Sovereign	**Equities**	**Commodities**	**FX**	**B-E**	**B-E-C**	**B-E-C-F**
CAGR	3.82%	16.18%	11.14%	1.53%	3.35%	4.97%	4.21%
Ann. Return	3.72%	15.43%	11.99%	1.75%	4.54%	5.90%	4.87%
Ann. Vol	4.58%	13.73%	18.90%	7.50%	16.47%	15.54%	13.33%
Sharpe	0.81	1.12	0.63	0.23	0.28	0.38	0.37
Skewness	0.43	1.21	-0.16	0.08	-0.01	0.06	0.14
maxDD	-0.09	-0.17	-0.47	-0.26	-0.38	-0.36	-0.32
Volatility Targeted							
	Sovereign	**Equities**	**Commodities**	**FX**	**B-E**	**B-E-C**	**B-E-C-F**
CAGR	15.23%	17.91%	7.61%	0.41%	0.98%	3.05%	2.43%
Ann. Return	14.97%	17.16%	8.31%	1.69%	2.23%	4.13%	3.56%
Ann. Vol	16.00%	15.78%	15.63%	16.05%	16.03%	15.69%	15.75%
Sharpe	0.94	1.09	0.53	0.11	0.14	0.26	0.23
Skewness	0.46	0.32	-0.15	-0.33	0.33	0.07	0.05
maxDD	-0.26	-0.23	-0.44	-0.64	-0.41	-0.55	-0.57

Table 1: Performance statistics of the single asset class and cross asset class value factors. Reading note: "B" stands for Bonds, "E" for equities, "C" commodities and "F" forex.

historical Sharpe ratio of a standard risk premia (see Ilmanen et al. (2014)). That number diminishes in the case of a volatility targeted portfolio to 0.14 to 0.26. There, out of the six skewness showed in Table 1, five are positive. The proposed approach thus seems to deliver appealing results. However, within the proposed framework, one can also analyze the statistical properties of our value factor, testing for the non-stationarity of prices, the stationarity of the value signal and finally investigating its half-life.

3.3 Correlation to momentum and existing cross-asset value

In order to further explore the characteristics of our approach to creating value factors, we present in this subsection a correlation analysis between the factor performances showed in the previous subsection and the performance of (1) a momentum strategy and (2) the Asness et al. (2013) value factor. The momentum strategy is built after the methodology presented in Moskowitz et al. (2012) and uses past 12 months of returns to measure trends across assets.

Figure 3 shows all of these correlations in one place. The factors obtained with our approach for all samples are negatively correlated with their momentum equivalent. These correlations range from -18% for Equities to -64% for Forex in the case of single asset datasets. In the case of cross-asset samples, that correlation is

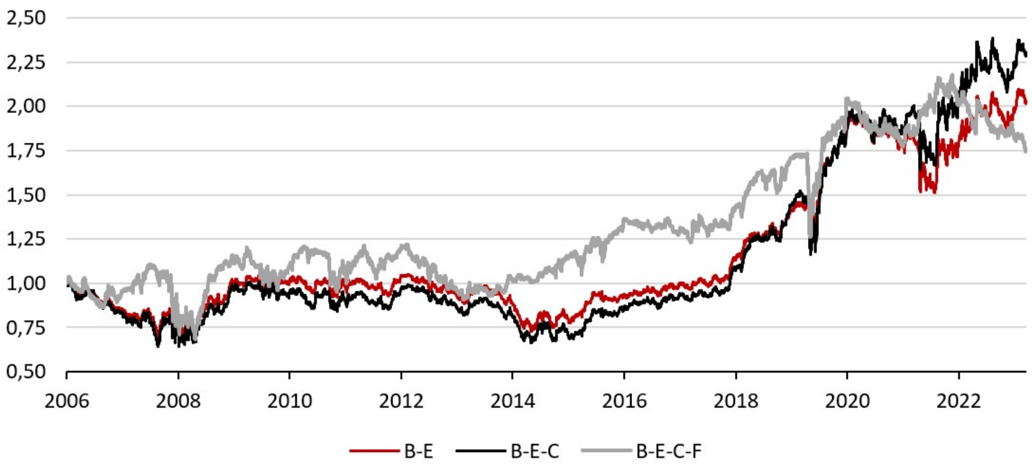

Figure 4: Indexed performance across datasets

globally about −60%. Hence, the correlations obtained with cross-asset universes are stronger than in the case of single asset classes. The correlations between our value factors and the one from Asness et al. (2013) are all positive. In the cross-asset case, they are also more important in scale than in the case of single asset classes. Both empirical results are therefore consistent with the intuition when it comes a value factor.

3.4 On Value's skewness

As mentioned earlier, value is often viewed as the counterpart to momentum; thus, one might expect Value to exhibit a skewness that is the opposite of that observed in a typical momentum strategy. Following this logic, as momentum strategies are expected to demonstrate downside risk and hence negative skewness, one would expect, according to Lempérière et al. (2017), a positive skewness in the performance of the various return streams obtained from each sub-sample. Figure 6 offers another perspective on this topic by displaying sorted returns from the lowest to the highest. A strategy with positive skewness should exhibit a thicker right tail compared to its left tail, and Figure 6 illustrates how, for each dataset considered here, the positive returns are significantly larger than the negative returns. Thus, our strategy contributes to the literature by identifying a positive asymmetry in Value strategies and factors.

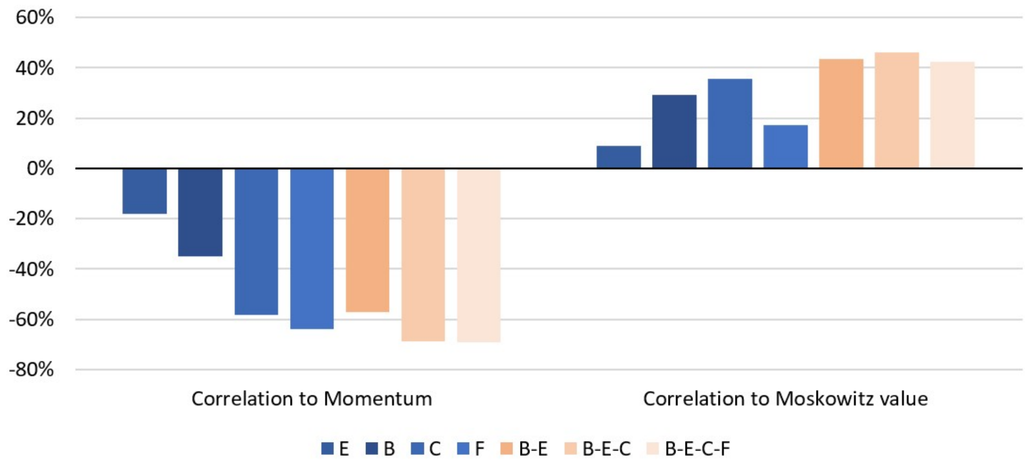

Figure 5: Correlation of the cross-asset factors with the momentum and Moskowitz value factor

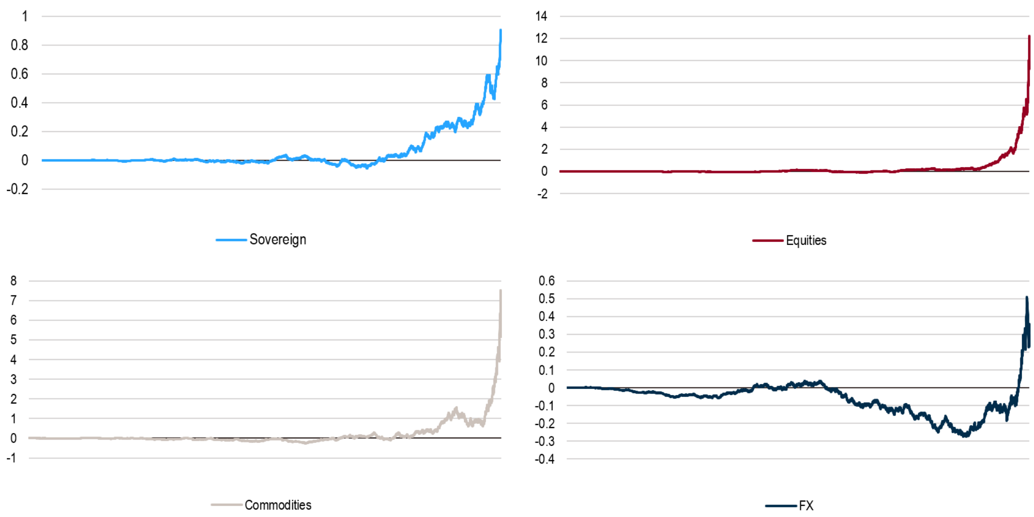

Figure 6: Sorted returns for the Value strategy per dataset

3.5 Value as a US Equities Factor

To complement the illustration presented in Subsection 3.1, we investigate the empirical behavior of an equity Value factor built on US stocks, in the spirit of Fama and French (1992). We consider the S&P 500 universe, which corresponds to the 500 largest US companies. We rebalance our universe monthly, from January 1990 to December 2024, resulting in a dataset of 1287 stocks. For each stock prices series, we compute the valuation signal (3) and every last day of the month, we build a long-short portfolio based on (4) from the stocks that are members of the S&P 500 index at that date[5]. Figure 5 presents the cumulative performance of the strategy compared to Fama-French factors obtained from Kenneth French website.

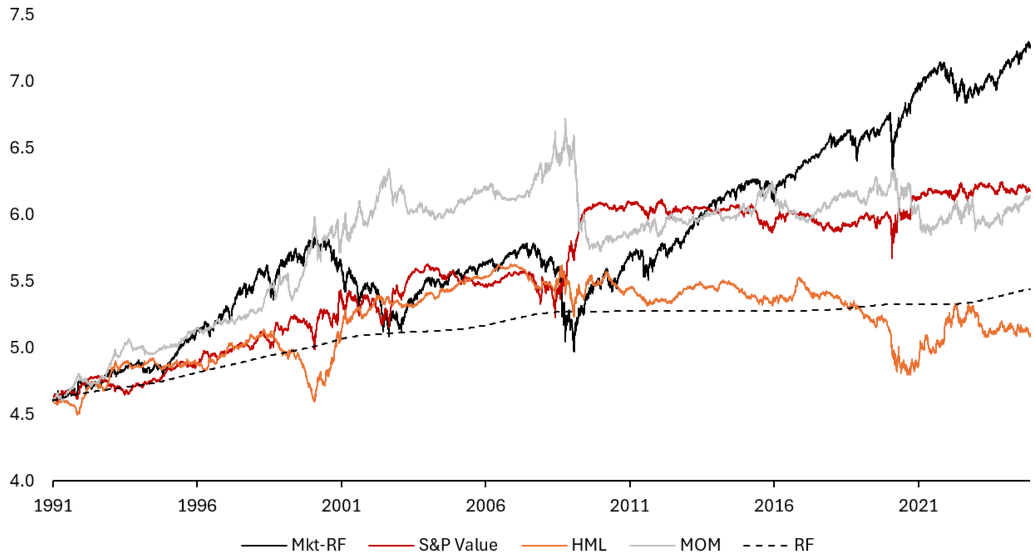

Figure 7: Performance (in log) of the Value factor compared to the Market portfolio in excess of the risk free rate (Mkt-RF), the Momentum factor (MOM), Fama-French Value factor (HML) and the risk free rate (RF).

While it is no the objective of this subsection to discuss the realistic performance of these paper portfolios[6], some interesting conclusions can be drawn. First, our proposed Value factor correlates positively (+0.2) with the High-Minus-Low (HML) factor built on boo-to-price valuation metrics. Although this number may seems

[5]As our methodology relies on long-term mean reversion to capture fundamental value, some stocks may be members of the S&P 500 but not part of the portfolio if the signal cannot yet be computed. We have on average 468 stocks in our Value portfolio.

[6]It is now well documented that equity factors are notoriously difficult to trade and the performance of the long-short portfolios is often challenged in realistic trading conditions (see for example Novy-Marx and Velikov (2016)).

low, it is worth noting that the universe on which the HML is built is different from the S&P 500 universe. In particular, the HML portfolio is computed as the performance of two double-sorted portfolios (based on market capitalization and book-to-price ratio) to capture the Value factor across small and large stocks. Of course, all members of the S&P 500 are large stocks which may explain part of the difference between Fama-French Value factor and ours. Additionally, our Value factor correlates negatively to the Profitability (Robust-Minus-Weak) factor introduced by Fama and French (2015). This negative correlation between Quality and Value is a well-known empirical facts of factor investing, which confirms that our observation-driven valuation method behaves appropriately. Finally, our Value factor presents a staggering -0.7 correlation with the Momentum factor, exhibiting upside convexity when Momentum crashes.

4 Conclusion

This paper introduces a novel observation-driven approach to building a cross-asset Value factor. This simple strategy solves the difficulty that analysts are faced with when comparing the valuation of different asset classes, proposing a unified way to compute such a value measure. We showcase how across single and multi-asset investment universe our approach consistently yields a factor that (1) delivers a positive Sharpe ratio, (2) is negatively correlated to a momentum factor while (3) being positively correlated to known value measures. Moreover, our Value factor presents an appealing convex payoff allowing to hedge momentum tail risk reaffirming the relationship between Value investing and mean-reversion.

References

C. S. Asness, T. J. Moskowitz, and L. H. Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013.

R. Bansal and A. Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

S. Basu. Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis. *The Journal of Finance*, 32(3): 663–682, 1977.

S. Basu. The relationship between earnings' yield, market value and return for NYSE common stocks: Further evidence. *Journal of Financial Economics*, 12(1): 129–156, 1983.

G. Bekaert and S. R. Grenadier. Stock and bond pricing in an affine economy. Working Paper 7346, National Bureau of Economic Research, September 1999.

J. B. Berk, R. C. Green, and V. Naik. Optimal investment, growth options, and security returns. *The Journal of Finance*, 54(5):1553–1607, 1999.

F. Black. Noise. *The Journal of Finance*, 41(3):528–543, 1986.

W. F. M. D. Bondt and R. Thaler. Does the stock market overreact? *The Journal of Finance*, 40(3):793–805, 1985.

H.-Y. S. Chien, J. S. Turek, N. Beckage, V. A. Vo, C. J. Honey, and T. L. Willke. Slower is better: revisiting the forgetting mechanism in lstm for slower information decay. *arXiv preprint arXiv:2105.05944*, 2021.

E. F. Fama and K. R. French. The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465, 1992.

E. F. Fama and K. R. French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):55–84, 1996.

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

J. Gerakos and J. T. Linnainmaa. Decomposing value. *The Review of Financial Studies*, 31(5):1825–1854, 10 2017.

S. Gu, B. Kelly, and D. Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273, 02 2020.

J. D. Hamilton. Why you should never use the hodrick-prescott filter. *The Review of Economics and Statistics*, 100(5):831–843, 12 2018.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 11 1997.

R. J. Hodrick. An exploration of trend-cycle decomposition methodologies in simulated data. Working Paper 26750, National Bureau of Economic Research, February 2020.

R. J. Hodrick and E. C. Prescott. Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1):1–16, 1997.

K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

F. Ielpo and M. Kniahin. Fundamental bubbles in equity markets. *Soft Computing*, 24(18):13769–13796, 2020.

A. Ilmanen, T. Maloney, and A. Ross. Exploring macroeconomic sensitivities: How investments respond to different economic environments. *The Journal of Portfolio Management*, 40(3):87–99, 2014.

B. Kelly and D. Xiu. Financial machine learning. *Foundations and Trends in Finance*, 13(3-4):205–363, 2023.

J. Lakonishok, A. Shleifer, and R. W. Vishny. Contrarian investment, extrapolation, and risk. *The Journal of Finance*, 49(5):1541–1578, 1994.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Y. Lempérière, C. Deremble, T. T. Nguyen, P. Seager, M. Potters, and J. P. Bouchaud. Risk premia: asymmetric tail risks and excess returns. *Quantitative Finance*, 17(1):1–14, 2017.

M. Lettau and J. A. Wachter. The term structures of equity and interest rates. *Journal of Financial Economics*, 101(1):90–113, 2011.

T. J. Moskowitz, Y. H. Ooi, and L. H. Pedersen. Time series momentum. *Journal of Financial Economics*, 104(2):228–250, 2012.

R. Novy-Marx and M. Velikov. A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, 29(1):104–147, 2016.

O. Triebe, N. Laptev, and R. Rajagopal. AR-net: A simple auto-regressive neural network for time-series. *arXiv preprint arXiv:1911.12436*, 2019.

O. Triebe, H. Hewamalage, P. Pilyugina, N. Laptev, C. Bergmeir, and R. Rajagopal. Neuralprophet: Explainable forecasting at scale. *arXiv preprint arXiv:2111.15397*, 2021.

T. Werner and W. Lemke. The term structure of equity premia in an affine arbitrage-free model of bond and stock market dynamics. Working Paper Series 1045, European Central Bank, Apr. 2009.

Appendix A List of data considered

Bonds	Equities	Commodities		Forex
Treasury note 2y	SP500	Gold	Soybean oil	EUR
Treasury note 10y	Nasdaq	Brent	Gasoil	GBP
Treasury bond	Russell 2000	WTI	Sugar	JPY
Canadian bond	TSX	Cattle	Coffee	AUD
Schatz	Eurostoxx	Nat Gas	Wheat	NZD
Bund	DAX	Soybean	Nickel	CAD
Buxl	CAC40	Corn	Heating oil	CHF
Gilt	AEX	Copper	Zinc	NOK
JGB	FTSE	Silver	Cotton	SEK
Australia 10y	SMI	Aluminum	HRW Wheat	
	Nikkei	Lean hogs	Lead	
	Toppix	Soybean meal		

Appendix B Stationarity of the residuals in Hamilton's regression

One of the interests of the proposed approach comes from its ties to econometrics. Hamilton (2018)) approach is close in spirit to cointegration, as the price series are non-stationary variables while the residual of the regression of price series on their past lags is. Furthermore, that cyclical component should exhibit persistency that can be measured through an AR(1) model, once stationarity conditions have been validated. That estimation will provide us with information regarding the half-life of the value anomaly, an information that could be further used to fine-tune the strategy.

To test for the stationarity of the value signal, given the suspected autocorrelation of it, an Augmented Dickey-Fuller test is appropriate. We do not reproduce the results of the tests we conducted on both in- and out-of-sample value signal, as they all showed a rejection of the null hypothesis of non-stationarity. All our value signals are found to be stationary. We ran a similar test on Asness et al. (2013)'s value measure and failed to reject the null for all of them.

Second, once we made sure our residuals were stationary, we estimated on each of them an AR(1) model from which we derived a half-life. The results are presented in Figure 8 below. The average half-life is 153 days, which is about 7 months. The half-life of commodities (160 days), currencies (168 days) and equities (154 days) is

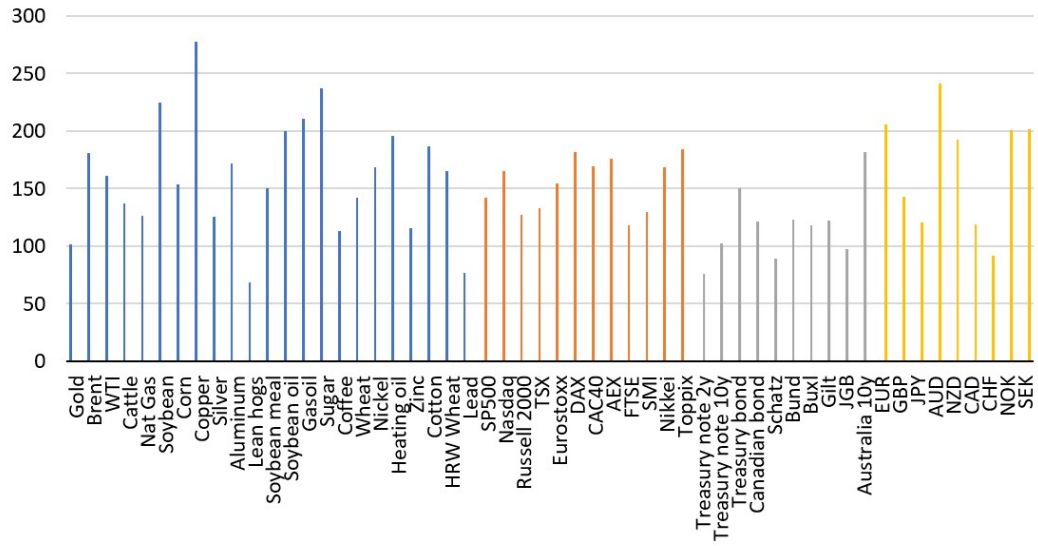

Figure 8: Half-life of the value signal per asset class

marginally higher than the ones for bonds (118 days). This is an advantage of our approach when comparing it to other valuation methodologies.

Appendix C Robustness checks

This section presents a set of robustness checks aimed at providing additional information on the characteristics of the value factor we propose. On Figure 5, we present the breakeven transaction costs (t-cost) for each of the value factors, that is the average transaction costs necessary to make the performance of each factor equal to zero. This breakeven t-cost ranges from 3 basis points (bps) in the case of Forex to 104 bps in the case of equities. Individual asset classes therefore show a large disparity in terms of these transaction costs, while cross-asset factors show a more stable level for these of around 10 basis points. Given the liquidity of the instruments we use here, these transaction cost numbers are higher than usual market conventions, preserving the interest of such a value strategy.

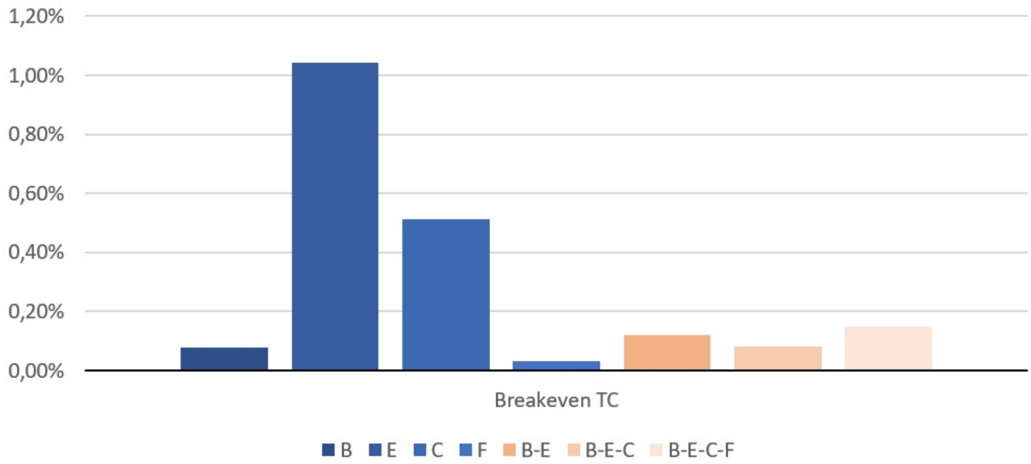

Figure 9: Breakeven t-costs per type of investment universe