

Which Trend is your Friend?

Levine, Ari; Heje Pedersen, Lasse

Document Version

Final published version

Published in:

Financial Analysts Journal

DOI:

[10.2469/faj.v72.n3.3](https://doi.org/10.2469/faj.v72.n3.3)

Publication date:

2016

License

CC BY

Citation for published version (APA):

Levine, A., & Heje Pedersen, L. (2016). Which Trend is your Friend? *Financial Analysts Journal*, 72(3), 51-66.
<https://doi.org/10.2469/faj.v72.n3.3>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 18. Jun. 2025



Which Trend Is Your Friend?

Ari Levine & Lasse Heje Pedersen

To cite this article: Ari Levine & Lasse Heje Pedersen (2016) Which Trend Is Your Friend?, Financial Analysts Journal, 72:3, 51-66, DOI: [10.2469/faj.v72.n3.3](https://doi.org/10.2469/faj.v72.n3.3)

To link to this article: <https://doi.org/10.2469/faj.v72.n3.3>



© 2016 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 27 Dec 2018.



Submit your article to this journal [↗](#)



Article views: 559



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 20 View citing articles [↗](#)

Which Trend Is Your Friend?

Ari Levine and Lasse Heje Pedersen

Managed futures funds and commodity trading advisers (CTAs) use heuristics or statistical measures often called “filters” to trade on price trends. Two key statistical measures of trends are “time-series momentum” and “moving-average crossovers.” We show, empirically and theoretically, that these trend indicators are closely related. In fact, they are equivalent representations in their most general forms. They also capture many other types of filters, such as the Hodrick–Prescott (HP) filter, the Kalman filter, and all other linear filters. We show how these filters can be represented through “trend signature plots,” demonstrating their dependence on past prices and returns by horizon.

Trend-following investing is the predominant investment style for managed futures hedge funds, commodity trading advisers (CTAs), and certain macro traders.¹ Trend-following investing can be defined loosely as buying when prices have been rising and selling short when prices have been falling. The strategy is based on the idea that these price trends are more likely to continue than not. Several studies have found trend-following investing to be profitable,² but what is the best way to identify a price trend? What methods exist for identifying trends, and how do they compare with one another? These are the questions we seek to address in this article.

To put our findings in a broader perspective, we note that because day-to-day price changes are “noisy,” finding a trend that predicts the *next* day’s price move in any market is never easy. According to the so-called random walk or efficient market hypothesis, future price moves are completely unpredictable, meaning that trend-following strategies should not work.³ Price trends may exist, however, if markets are not completely efficient or if risk premiums change over time.

Finding a price trend among noisy random price moves presents a challenge similar to that of “filtering” information from the noise in many other applications,

such as astronomy, audio, ballistics, image processing, and macroeconomics. For example, engineers who track ballistic missiles based on noisy radar information attempt to filter out noise to determine the missile’s direction. Similarly, macroeconomists and central bankers who receive imperfect economic data—such as estimates of GDP for countries and unemployment rates collected from many sources (with errors)—try to assess whether an economy is heading into a recession or is overheating. Investors trading on trends in financial markets face the similar challenge of assessing the direction that prices are headed by filtering noisy price data. In the world of audio, Ray Dolby developed the Dolby system to reduce noise in music recordings and enhance the “signal” that the listener hears. Along the same lines, trend followers use quantitative tools to enhance the signal of the price trend and reduce the noise around it.

In finance, a simple approach to capture price trends is time-series momentum (TSMOM) as defined by Moskowitz, Ooi, and Pedersen (2012). The simplest form of a TSMOM signal is the return over some recent time period (e.g., the past 12 months). For instance, if investing in gold has resulted in a positive return over the past 12 months, then the trend is assessed to be upward and the TSMOM investor buys gold. If the past return is negative, the trend is assessed to be downward and the TSMOM investor sells gold short. Moskowitz et al. (2012) showed that investing based on 12-month TSMOM was, on average, profitable for each of the 58 liquid securities they analyzed over the past 25 years. Their universe included instruments from among the world’s most liquid commodity futures, equity futures, bond futures, and currency forwards.

Another way to assess trends in financial markets is the moving-average crossover (MACROSS) method (e.g., Brock, Lakonishok, and LeBaron 1992; Okunev

Ari Levine is a managing director at AQR Capital Management, LLC, Greenwich, CT. Lasse Heje Pedersen is a professor at Copenhagen Business School, Frederiksberg, Denmark, and New York University, NY, and a principal at AQR Capital Management, LLC, Greenwich, CT.

Editor’s note: This article was reviewed and accepted by Executive Editor Stephen J. Brown.

Editor’s note: Both authors are affiliated with AQR, a global investment management firm running long-only and alternative investment products, including managed futures funds.

and White 2003). A CTA using this method will buy when one moving average of an asset's recent prices crosses another moving average of the asset's prices measured over a longer horizon. The idea is that a "fast" moving average captures the average recent prices whereas a "slower" moving average captures where prices used to be. If recent prices are above where prices used to be, then the trend is assessed to be upward and the MACROSS investor buys.

Both the TSMOM and MACROSS methods can be refined in various ways (e.g., by relying on different trend horizons). We show that the most general form of MACROSS can be viewed as a special case of the most general TSMOM strategy, and vice versa. A different way of stating this result is that trend filters can be equivalently represented as functions of past prices versus past returns. We illustrate how each trend signal can be represented graphically using "trend signature plots," again based on either past prices or returns. We also show how a large class of filtering methods used in science and economics can be viewed as a special case of TSMOM.

Finally, we report an empirical study that compares the performance of common implementations of TSMOM strategies with that of MACROSS strategies, and we discuss the implications for investors.

Time-Series Momentum vs. Moving-Average Crossover

In this section, we first describe separately the time-series momentum strategy and moving-average

crossover strategy. Then, we analyze how they can substitute for each other, and we illustrate our results through a number of specific examples.

Time-Series Momentum. A TSMOM strategy goes long when prices have been moving up and short when prices have been moving down. The simplest TSMOM signal is the past return over some time period—say, m months or days:


$$\text{TSMOM}_t^m = \text{Return}_{t-m,t} \quad (1)$$

For instance, 12-month momentum considers the return over the past 12 months. The return can be computed as the ratio of prices, P_t / P_{t-m} , or as the ratio of two points of a return index that takes dividends or coupons into account for cash instruments and handles roll yields and implicit financing for futures. Alternatively, it can be computed as the difference of (log) prices, $P_t - P_{t-m}$ (or, again, with a return index used in the place of prices). In this study, for simplicity, we focused on differences in log prices or index levels. As an example, **Figure 1** shows how gold prices had positive TSMOM at the end of 2010 because prices had been trending upward.

More-refined TSMOM signals are also possible. One way to refine the TSMOM signal is to smooth the prices used to calculate the return:

$$\text{TSMOM}_t^m = \text{Average}(P_s : s \text{ near current time } t) - \text{Average}(P_{s'} : s' \text{ near lagged time } m). \quad (2)$$

Figure 1. TSMOM of Gold Futures, January 2008–January 2011



Note: The arrow shows the filtered trend from the smoothed past price (left horizontal bar) to the smoothed current price (right horizontal bar).

Smoothing can be a good idea because it reduces random noise in the data. For instance, focusing on a single past price might be arbitrary and subject to more noise than using an average of multiple past prices (which could be called “back-end smoothing”). Figure 1 uses the average price over the 60-day period, indicated by the dotted horizontal line. Asness, Moskowitz, and Pedersen (2013) used back-end smoothing of cross-sectional momentum signals.

Smoothing recent prices (“front-end smoothing”) also reduces noise but has the potential drawback of delaying the signal. With front-end smoothing, recent price changes are smoothed out and, therefore, only gradually affect the trading signal. This approach can be suboptimal if recent prices contain important information about the current trend or a trend reversal, but it can be helpful in reducing turnover.

Of course, traders may want to define the TSMOM signal over a variety of horizons. In the extreme, one could use a series of daily (or monthly) returns and give each day’s return a separate weight, c :

$$\text{TSMOM}_t^c = \sum_{s=1}^{\infty} c_s (P_{t-s+1} - P_{t-s}). \quad (3)$$

Equation 3 means that a general TSMOM signal can be generated by considering all the past daily price changes and assigning importance to each day based on how long ago it happened. For instance, one might want to rely more on recent price changes in assessing the current price trend. Moskowitz et al. (2012) conducted a detailed analysis of how


returns at various lags predict future returns. A trend-following strategy is characterized by having positive coefficients c_s , whereas negative coefficients correspond to reversal trades.

Moving-Average Crossover. The MACROSS strategy first computes two moving averages (MAs) of prices, which we call MA^{fast} and MA^{slow} . The fast MA puts more weight on recent prices; the slow MA puts more weight on past prices. As an example, we can compute an equal-weighted MA over the past 20 weekdays as a measure of recent prices and a 260-day average as a measure of where prices used to be. **Figure 2** shows the plot of these moving averages for gold prices over the same period shown in Figure 1.

The MACROSS strategy depends on which MA—the fast one or the slow one—is higher. In Figure 2, the fast MA is above the slow MA at the end of the time period; that is, recent prices are above past prices, resulting in an upward trend.

Of course, other MACROSS strategies can be devised by varying the time horizons (here, the 20- and 260-day averaging periods). Furthermore, each day’s price need not have an equal weight in the moving average. Another common method, which is discussed in detail later, is to weight past prices exponentially (called “exponentially weighted moving average”). Generally, we can compute the MAs by using any weighting scheme, w , where the weights can, for instance, be equal weights or exponential weights. Hence, in general, we can write the MAs mathematically as

Figure 2. MACROSS Indicator for Gold Futures, January 2008–January 2011



Note: Because the fast moving average is above the slow moving average at the end of the sample, the filtered trend is up, as indicated by the arrow.

$$\begin{aligned} \text{MA}_t^{\text{fast}} &= \sum_{s=1}^{\infty} w_s^{\text{fast}} P_{t-s+1} \\ \text{MA}_t^{\text{slow}} &= \sum_{s=1}^{\infty} w_s^{\text{slow}} P_{t-s+1}. \end{aligned} \quad (4)$$

The idea that one MA is faster than the other can be captured mathematically by the requirement that the fast MA place more weight on the most recent prices:

$$\sum_{j=1}^s w_j^{\text{fast}} \geq \sum_{j=1}^s w_j^{\text{slow}} \text{ for all } s. \quad (5)$$

The trading signal is then the MACROSS—that is, the difference between these moving averages:

$$\text{MACROSS}_t = \text{MA}_t^{\text{fast}} - \text{MA}_t^{\text{slow}}. \quad (6)$$

Hence, the MACROSS signal tries to measure whether recent prices, as captured by MA^{fast} , are above or below more distant prices, as captured by MA^{slow} . Intuitively, a positive MACROSS means that recent prices are higher than past ones, indicating a rising trend.

MACROSS as TSMOM. The MACROSS signal is the difference between two MAs and, therefore, a combination of past prices:

$$\text{MACROSS}_t = \sum_{s=1}^{\infty} (w_s^{\text{fast}} - w_s^{\text{slow}}) P_{t-s+1}. \quad (7)$$

This equation shows that MACROSS signals, in general, can be viewed as combinations of past price *levels*. Similarly, the general TSMOM equation (Equation 3) shows that TSMOM is a combination of past price *changes*. We can go back and forth between price levels and price changes, however, if we change the coefficients accordingly.

Specifically, the MACROSS equation (Equation 7) is equivalent to the TSMOM strategy (Equation 3) with coefficients on past returns c_s computed as follows:⁴

$$c_s = \sum_{j=1}^s (w_j^{\text{fast}} - w_j^{\text{slow}}). \quad (8)$$

Implied coefficients c_s are positive for all MACROSS strategies in which the fast MA is uniformly faster than the slow MA as given by Equation 5, which is true for the standard MACROSS strategies. Naturally, these TSMOM coefficients are positive because the strategy is trend following (negative coefficients would have indicated a bet on trend reversal).

Furthermore, implied return coefficients c_s approach zero as the number of lags, s , increases (assuming that the weights w_s^i sum to 1 for each i). Although the coefficients c_s have no specific “scale,” it is natural to normalize them in such a way that they sum to 1:

$$\bar{c}_s = \frac{c_s}{\sum_{j=1}^{\infty} c_j}. \quad (9)$$

In this way, the TSMOM signal can be viewed as a weighted average of past returns.

To understand this conversion from MACROSS to TSMOM (i.e., the conversion from “price space” to “return space”), consider the 20-day versus 260-day equal-weighted MA strategy. This strategy compares past prices on the basis of the coefficients illustrated in **Figure 3**.

When we take the MACROSS (i.e., the fast average minus the slow average), then we have the weights on past prices depicted in **Figure 4**.

Figure 3. MACROSS Coefficients: Fast and Slow Averages


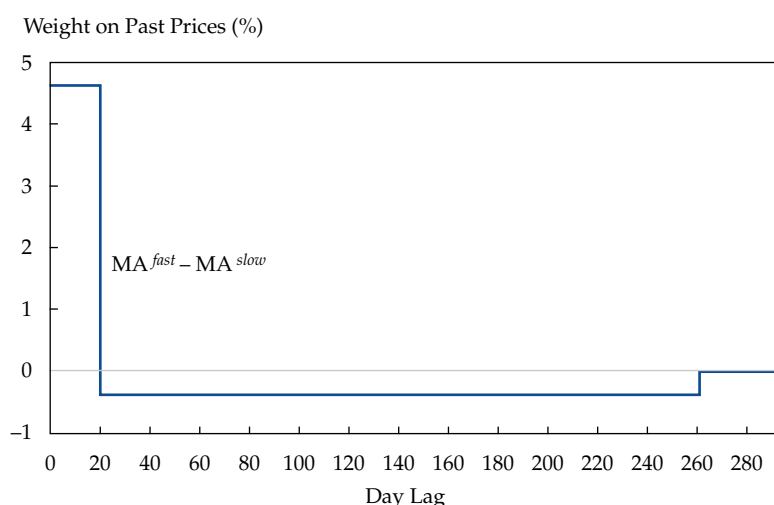



Figure 4. Price Signature Plot: MACROSS, Equal Weighted

Notes: As shown in Equation 7, the MACROSS signal was computed as a weighted average of past prices, where the weights are the fast MA (20 days) minus the slow MA (260 days). The resulting signal puts positive weights on the most recent 20 days and negative weights on the past 21–260 days.

The MA weights $w_s^{fast} - w_s^{slow}$ usually have this shape as a function of the time lag s . The weight on recent prices is positive, the weights on distant prices are negative, and the weights eventually go to zero. Thus, a moving-average strategy can be interpreted as a TSMOM strategy in which both the front-end price and the back-end price have been smoothed. Indeed, the MACROSS strategy is like a TSMOM strategy in which the current price is computed as the average of the past 20 days' prices, the past price is computed as the average of the prices from day 21 to day 260, and the return is then computed as the difference between these two smoothed prices.

We can also use Equation 8 to translate the MACROSS coefficients for price levels into TSMOM coefficients for price *changes* (i.e., returns). The coefficients in return space are shown in **Figure 5**, where we have normalized the weights to sum to 1 in both cases. The graph shows two lines: The dotted line is the simplest TSMOM signal. It gives equal weight to the price change (or return) on each of the past 260 days; that is, it assesses the direction of the trend based on the average return. The solid line plots the MACROSS coefficients derived from Equation 8. It shows that, even though MACROSS is defined as a moving average of price levels, it can be computed instead from price changes. The MACROSS assigns

Figure 5. Return Signature Plot: Trend Coefficients in Return Space

Note: This figure shows how much weight the trend indicator places on each daily return in the past for the simple 260-day TSMOM and the 20/260 MACROSS strategies.

the most importance to intermediate price changes and less weight to the most recent price changes and very old price changes.

TSMOM as MACROSS. The simple TSMOM signal can be computed as the difference between the current price (or log price or return index), P_t , and the lagged price, P_{t-m} (e.g., the price 12 months ago, or $P_{t-12 \text{ months}}$):

$$\text{TSMOM}_t^m = P_t - P_{t-m}. \quad (10)$$

Equation 10 shows that a TSMOM strategy can be viewed as a MACROSS in which the recent moving average, MA^{fast} , is simply the current price, P_t ; that is, the weighting scheme puts all the weight on one price—namely, the most recent one: $w_1^{\text{fast}} = 1, w_s^{\text{fast}} = 0$ for $s > 1$. Similarly, the distant moving average, MA^{slow} , is simply the lagged price; that is, its weight scheme puts all the weight on that price: $w_m^{\text{slow}} = 1$.

More-refined TSMOM signals can also be viewed as MACROSS. If one uses front-end smoothing, then the MA^{fast} becomes a (possibly weighted) average of recent prices. Similarly, if back-end smoothing is used, then MA^{slow} becomes an average of lagged prices.

If many momentum horizons are used simultaneously, with coefficients c , as discussed previously, then the MACROSS weights can be computed as follows:

$$\begin{aligned} w_1^{\text{fast}} - w_1^{\text{slow}} &= c_1 \\ w_j^{\text{fast}} - w_j^{\text{slow}} &= c_j - c_{j-1}. \end{aligned} \quad (11)$$

Equation 11 shows how to determine the *difference* between the weights of MA^{fast} and MA^{slow} . Many choices of moving averages produce the same signal, however, because adding and subtracting the same price has no effect. In contrast, momentum weights c are unique, as are the weights on past prices, $w_s^{\text{fast}} - w_s^{\text{slow}}$, so these variables are more fundamental parameters of the filtering process.⁵

Example: Exponentially Weighted MACROSS.

An exponentially weighted moving-average (EWMA) crossover is similar to a simple MACROSS, but the fast and slow moving averages are exponentially weighted instead of equal weighted. Specifically, an exponential decay of $\theta > 0$ produces

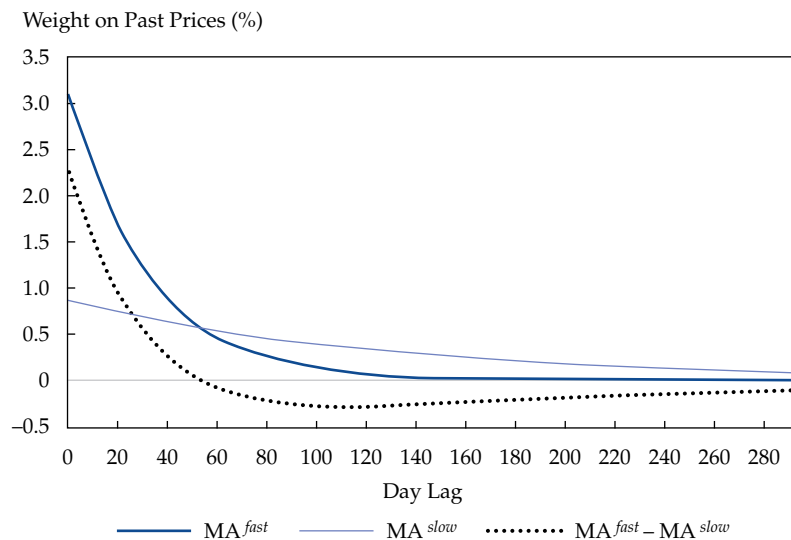
$$\text{EWMA}_t = \frac{1}{1-\theta} \sum_{j=0}^{\infty} \theta^j P_{t-j}. \quad (12)$$

A more intuitive approach than parameterizing by decay θ is to consider the center of mass (COM) of the moving average, defined as

$$\begin{aligned} \text{COM} &= \frac{1}{1-\theta} \sum_{j=0}^{\infty} \theta^j j \\ &= \frac{\theta}{1-\theta}. \end{aligned} \quad (13)$$

The COM can be useful for forming an intuition about the effective length of the moving average. For an exponentially weighted MACROSS, the price weights, $w_j = \theta^j / (1-\theta)$, look similar to those of a simple MACROSS but are smoother, as we show in Figure 6.

Figure 6. Price Signature Plot: MACROSS with Exponential Weights



Notes: This figure shows how much weight an EWMA crossover puts on past prices. The fast EWMA uses a COM of 32, and the slow EWMA uses a COM of 128. The effective weights of the EWMA crossover are also shown.

An exponentially weighted MACROSS also implies *return* weights c that are similar to the equal-weighted case but smoother, as shown in **Figure 7**.

The Equivalence of All Other Linear Filters

The literature on signal processing includes many types of linear filters for applications in science and engineering. In fact, a large set of linear filters of prices can be viewed as TSMOM and MACROSS signals if we allow any weights c and w .⁶ Specifically, linear filters corresponding to positive return weights c have a natural interpretation as a TSMOM signal. Hence, TSMOM and MACROSS trend indicators represent many classic filtering techniques.

What is not immediately captured by TSMOM and MACROSS filters is nonlinear effects, such as whether the signs of the returns have been consistent for a time period or, conversely, whether returns have been accelerating recently. However, variations of TSMOM can also account for such effects. For example, we have thus far limited ourselves to strictly positive weights c . If we loosen this restriction to include sets of weights the *sum* of which is positive but which include some negative weights, we can create trend measures that implicitly include differences in returns. These metrics can be interpreted as including acceleration/deceleration measures. For example, if we put positive weights on more recent returns and negative weights on more distant returns, we are measuring whether returns have been stronger recently than in the past; in other words, we are measuring acceleration.

The HP Filter. TSMOM and MACROSS techniques also capture, as a special case, the so-called HP (Hodrick–Prescott) filter. Hodrick and Prescott (1997) reported that the method (also called the Whittaker–Henderson Type A method; see Whittaker 1923 and Henderson 1924) has been used in actuarial sciences to smooth mortality rates; in astronomy (e.g., by Giovanni Schiaparelli in 1867); and in ballistics (e.g., by John von Neumann in the 1940s). The method is widely applied in macroeconomics, where it is used to filter out the business-cycle trends from noisy data on GDP growth.


The HP filter is based on the idea that prices (or GDP) have a growth component g and a cyclical component z (i.e., $P_t = g_t + z_t$). The trend is the change of the smooth growth component, $\text{trend}_t = g_t - g_{t-1}$. The growth component is filtered from the price data by finding a growth path that implies small trend variations, $\text{trend}_t - \text{trend}_{t-1}$ (stable trend), and small noise terms, z_t (good fit):

$$\min_{g_1, \dots, g_T} \sum_t (z_t)^2 + \lambda \sum_t (\text{trend}_t - \text{trend}_{t-1})^2, \quad (14)$$


using a parameter that determines how stable the filtered trend is.

We show in Appendix A that growth component g_t is a moving average of past prices. Hence, the $\text{trend}_t = g_t - g_{t-1}$ is a difference between two moving averages (the MA at time t and the MA at time $t-1$). Therefore, the HP trend is a MACROSS signal. As an example, **Figure 8** shows the weights on prices for the two growth components, g_t and g_{t-1} , as well as for the trend_t , based on an HP filter with parameter $\lambda = 10^4$.

Figure 7. Return Signature Plot: MACROSS Using Exponential Weights



Notes: This figure shows the weights that an EWMA crossover puts on past returns. The COMs used for the EWMA are, respectively, 32 and 128 for the fast and slow EWMA. For comparison, a 260-day TSMOM signal is also shown. Weights have been normalized to sum to 1 in both cases.

Figure 8. Price Signature Plot: HP Filter


Notes: This figure shows the weights on prices for the two growth components, g_t and g_{t-1} , and for $trend_t$ (their difference) for $\lambda = 10^4$. Note that, although g_{t-1} depends on P_t , the final trend does not depend on any future information.

Furthermore, because a MACROSS signal is also a TSMOM signal, the HP filter is also a TSMOM signal. That is, the HP trend can be written as a weighted average of past price changes, as we depict in Figure 9.

What is special about the HP filter is that it implies a particular shape for the MA weights and for the momentum return weights. The shape of the weights is similar to an exponentially weighted MA. The weights on returns are not strictly positive,

however, because the filter has a small amount of negative weight. Thus, the filter can be thought of as a combination of a simple TSMOM signal plus a small amount of acceleration.

The Kalman Filter. The Kalman filter (Kalman 1960) can be used to optimally estimate hidden variables of dynamic linear systems with noisy observations. A full treatment of Kalman filtering is beyond our scope here, but in the context of trend detection,

Figure 9. Return Signature Plot: HP Filter

Note: This figure shows the weights on past returns for $trend_t$, for $\lambda = 10^4$.

the Kalman filter can be applied to estimate the underlying (and hidden) trend variable driving returns.

The particular application of the Kalman filter will depend on the model used for the underlying data. If we know more about the underlying dynamics of the system, we can put more structure around the model, which may help in estimating the parameters of a Kalman filter. For price data, however, even simple random walk models tend to capture most of the important dynamics of the price series. So, it is not clear that adding more structure to the underlying data-generating process would be productive. If the goal is simplicity, the special case of the “local trend” model may be a good choice. The local trend model treats prices as a random walk with a trend, where the trend itself is a random walk whose value is not directly observable.⁷ The Kalman filter can then be used to estimate this underlying trend.

Harvey (1984) showed that in applying the Kalman filter to a local trend model, *the resulting optimal trend estimate is simply an exponentially weighted moving average of past returns*. In other words, in the absence of more model structure than the linear trend model, using the Kalman filter to estimate a trend results in strictly positive weights on past returns (i.e., a TSMOM-type signal). The COM parameter of the exponentially weighted moving average is determined by the parameters of the underlying model (which may themselves be estimated from the data).

Unlike the EWMA crossover discussed earlier, which consists of two exponentially weighted moving averages on past prices (one fast and one

slow), the Kalman filter results in an exponentially weighted moving average on *returns*. This pattern is different from the simple TSMOM or EWMA crossovers, as shown in **Figure 10**.

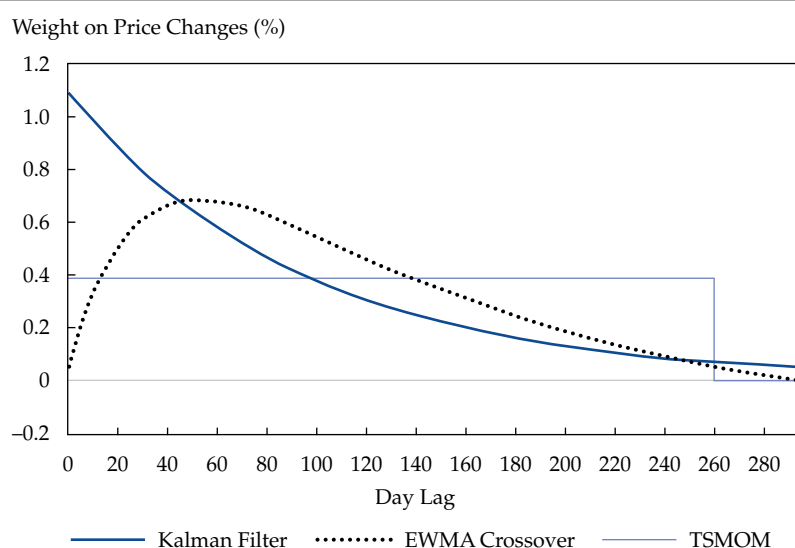
Trend Estimation Using Ordinary Least-Squares Trend Regression. Another intuitive trend measurement method that turns out to be equivalent to a generalized TSMOM signal is a regression-based trend estimate on prices. To estimate the trend in a price series over a certain period of time, we can estimate an ordinary least-squares (OLS) best-fit straight line through the price series. **Figure 11** shows how this approach works for a one-year trend estimate of gold prices.

We show in Appendix B that this trend estimation methodology is equivalent to a generalized TSMOM signal. In other words, it can be expressed as a linear combination of weighted past prices and, therefore, as a weighted combination of past price changes.


The set of weights in **Figure 12** is parallel to a MACROSS: Recent prices get positive weight, and more distant past prices get negative weight. For the sake of comparison, the simple MACROSS weights are also shown. The coefficients can also be translated, by using Equation 8, into weights on returns, as seen in **Figure 13**.

The OLS regression method gives the most weight to returns at the center of the time window. Returns at the extremes of the window (either most recent or most distant) are underweighted by this approach.

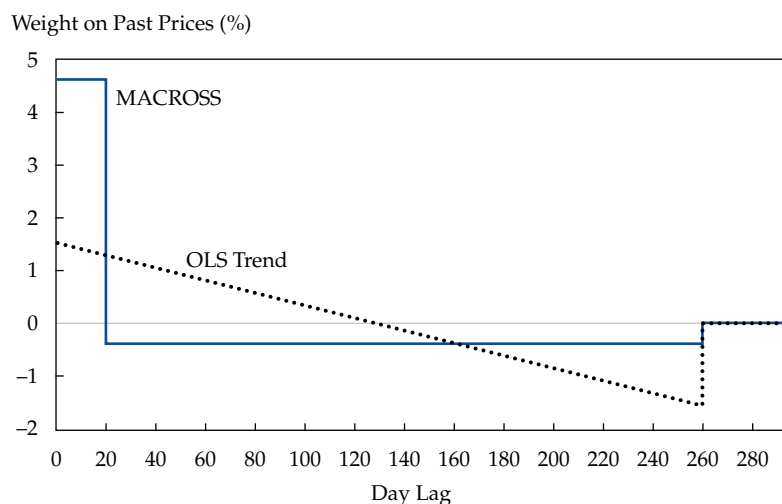
Figure 10. Return Signature Plot: Kalman Filter



Notes: This figure shows the weights on past price changes, or returns, for a Kalman filter applied to a local linear trend model. In this case, a COM of 96 was used to form the exponential weights. Also shown are weights for an EWMA crossover with COMs of, respectively, 32 and 128 for the fast and slow EWMA, as well as weights for a simple 260-day TSMOM signal.

Figure 11. OLS One-Year Trend Estimate of Gold Prices, January 2008–January 2011

Note: The last year of data shown was used.

Figure 12. Price Signature Plot: OLS Trend

Note: This figure shows the weights on prices for the OLS trend estimator, which estimates a best-fit trendline through the past 260 days of prices; it can be compared with the simple 20/260 MACROSS signal.


Many other operations are ultimately versions of weighting past returns that are similar to TSMOM. For example, using the multiresolution approach, “wavelets” can be used to extract trends at various resolutions from a price series.⁸

Empirical Analysis

We have shown theoretically how the most general forms of TSMOM and MACROSS are equivalent and capture all other linear filters. Given that TSMOM

and MACROSS capture all of the other filters and feature prominently in applications, we focused our empirical study on these trend indicators.

In practice, the common implementations of these trend signals are *not* exactly equivalent, so it is interesting to study their empirical similarities and differences. Indeed, differences in the performance of these signals can emerge from the specific parameter choices, from nonlinear transformations completed as part of portfolio construction, and from other practical effects.

Figure 13. Return Signature Plot: OLS Trend

Notes: This figure shows the weights on past price changes, or returns, for the OLS trend estimator, compared with the simple 20/260 MACROSS signal and the simple 260-day TSMOM signal. Each of the three sets of weights was normalized to sum to 1.

Data. We used prices from 24 commodity futures, 13 developed country government bond futures, 12 currency pairs from 9 underlying currencies, and 9 developed country equity indexes. These 58 instruments were chosen for their liquidity by Moskowitz et al. (2012), and we extended their dataset so that our data would cover prices from January 1985 through April 2015. Signals were calculated from a return index (rather than from prices directly) that was formed by rolling futures and forward prices; therefore, the index implicitly incorporated financing cost and “carry” or “rolldown.” The index reflects the actual returns from holding a rolled futures or forward position in an instrument with no cash outlay. Because futures and forwards have implicit financing, these return indexes are naturally excess of cash. A list of instruments and their sources is given in Appendix C.

Methodology. We wished to construct three standard TSMOM strategies and three standard MACROSS strategies that were relatively comparable. Following the methodology of Hurst et al. (2013), for the TSMOM strategies, we considered 1-month, 3-month, and 12-month trends. Specifically, the TSMOM signals were parameterized by the number of look-back days so that $\text{TSMOM}(n)$ was calculated as the log return index today minus the log return index n days ago:

$$\text{Signal}_t^{\text{TSMOM}(n)} = P_t - P_{t-n}. \quad (15)$$

We considered values of n equal to 22 trading days (approximately 1 month), 66 trading days (3 months), and 260 trading days (12 months). For

every day, we formed a new portfolio that was held for one day. The results were similar for one-month holding periods, except that, naturally, the Sharpe ratios were lower. We used daily rebalancing to focus on the connections between the various trend strategies, with minimal noise because of infrequent rebalancing.

We considered three MACROSS strategies at similar horizons. We used exponentially weighted MACROSS signals because they are perhaps the most common in investment practice. The MACROSS signals were parameterized by the centers of mass of the fast and slow moving averages so that a $\text{MACROSS}(m, M)$ signal had an m -day COM for its fast moving average and an M -day COM for its slow moving average:

$$\text{Signal}_t^{\text{MACROSS}(m, M)} = \sum_{s=1}^m w_s^m P_{t-s+1} - \sum_{s=1}^M w_s^M P_{t-s+1}, \quad (16)$$

where the weights are $w_s^m = \theta^s / (1 - \theta)$, with $\theta = m / (1 + m)$ as in Equation 12 and Equation 13. We chose m and M in such a way that the MACROSS strategies would correspond to the TSMOM signals by ensuring that they had similar trend horizons. Specifically, we let M take the values 12, 32, and 128 and set m to one-quarter of these values. The values were chosen so that M would be close to $n/2$ (and divisible by 4) for the corresponding $\text{TSMOM}(n)$ signals. These choices of M are natural because a $\text{TSMOM}(n)$ signal has a COM equal to $n/2$ (in returns) because it gives equal weight to the past n returns.

To put the trading signals on an equal footing, we followed Moskowitz et al. (2012) and used the same portfolio construction methodology for TSMOM and MACROSS signals. Specifically, for each strategy, our position in an asset i at time t was calculated as follows:

$$\text{Position}_t^i = 0.65\% \frac{\text{Sign}(\text{Signal}_t^i)}{\sigma_t^i}, \quad (17)$$

where Signal_t^i is the relevant TSMOM or MACROSS signal and σ_t^i is the volatility of asset i at time t . Again following Moskowitz et al. (2012), we used an exponentially weighted volatility with a COM of 60 days. We multiplied by 0.65% to target an annualized volatility of 0.65% in each asset. When aggregated, this level of asset volatility resulted in an annualized portfolio volatility of approximately 10% for each of the six strategies. We took the sign of the signal for simplicity, although many other transformations of the signal could be used in practice.

Empirical Results. The performance of each of the six strategies is reported in **Table 1**. The TSMOM and MACROSS strategies performed similarly for all horizons. Both delivered impressive risk-adjusted returns, with Sharpe ratios above 1 before transaction costs.

Consider next the central empirical question of our project—namely, a comparative study of these two approaches to trend-following investing. For this study, we regressed the return $r_t^{MACROSS(m,M)}$ of each MACROSS factor on all three TSMOM factors:

Table 1. Performance of Simple TSMOM and MACROSS Strategies

Signal	Annual Returns (excess of cash)	Annualized Volatility	Sharpe Ratio
MACROSS(3,12)	10.3%	10.2%	1.01
MACROSS(8,32)	10.9	10.3	1.06
MACROSS(32,128)	12.8	9.7	1.33
TSMOM(22)	9.8	10.1	0.97
TSMOM(66)	12.1	10.1	1.20
TSMOM(260)	14.2	9.8	1.45

Notes: This table shows the performance statistics of six different trend-following strategies. Excess returns and volatility were annualized, and the Sharpe ratio is the ratio of the two.

$$r_t^{MACROSS(m,M)} = \alpha + \beta_1 r_t^{TSMOM(22)} + \beta_2 r_t^{TSMOM(66)} + \beta_3 r_t^{TSMOM(260)}. \quad (18)$$

We also ran the regression with each of the TSMOM factors on the left-hand side: regressing on the three MACROSS factors. In other words, we performed six OLS regressions: regressing each TSMOM (MACROSS) factor portfolio return on the three MACROSS (TSMOM) factor portfolio returns. The results are summarized in **Table 2**.

For each of the six regressions, the R^2 is above 80%. Such a high R^2 shows that these trend signals

Table 2. Regressions of TSMOM on MACROSS and Vice Versa (t -statistics in parentheses)

Dependent Variable	Independent Variables				
A. Regression of MACROSS on TSMOM					
	TSMOM(22)	TSMOM(66)	TSMOM(260)	Intercept	R ²
MA(3,12)	0.76 (30.20)	0.25 (10.85)	0.01 (0.53)	−0.29% (−0.40)	84%
MA(8,32)	0.19 (17.20)	0.73 (65.42)	0.13 (14.29)	−1.78% (−2.41)	86
MA(32,128)	−0.13 (−12.06)	0.18 (16.48)	0.83 (85.76)	0.18% (0.23)	83
B. Regression of TSMOM on MACROSS					
	MA(3,12)	MA(8,32)	MA(32,128)	Intercept	R ²
TSMOM(22)	0.91 (49.02)	−0.01 (−0.84)	−0.04 (−2.53)	1.19% (1.50)	81%
TSMOM(66)	0.03 (2.08)	0.85 (59.46)	0.04 (3.44)	2.06% (2.54)	82
TSMOM(260)	0.14 (10.58)	−0.05 (−3.76)	0.90 (93.09)	1.84% (2.35)	82

Notes: Panel A shows the regression of each MACROSS factor on the three TSMOM factors. Panel B shows the regression of each TSMOM factor on the three MACROSS factors. The intercept was multiplied by 260 to annualize the daily returns.

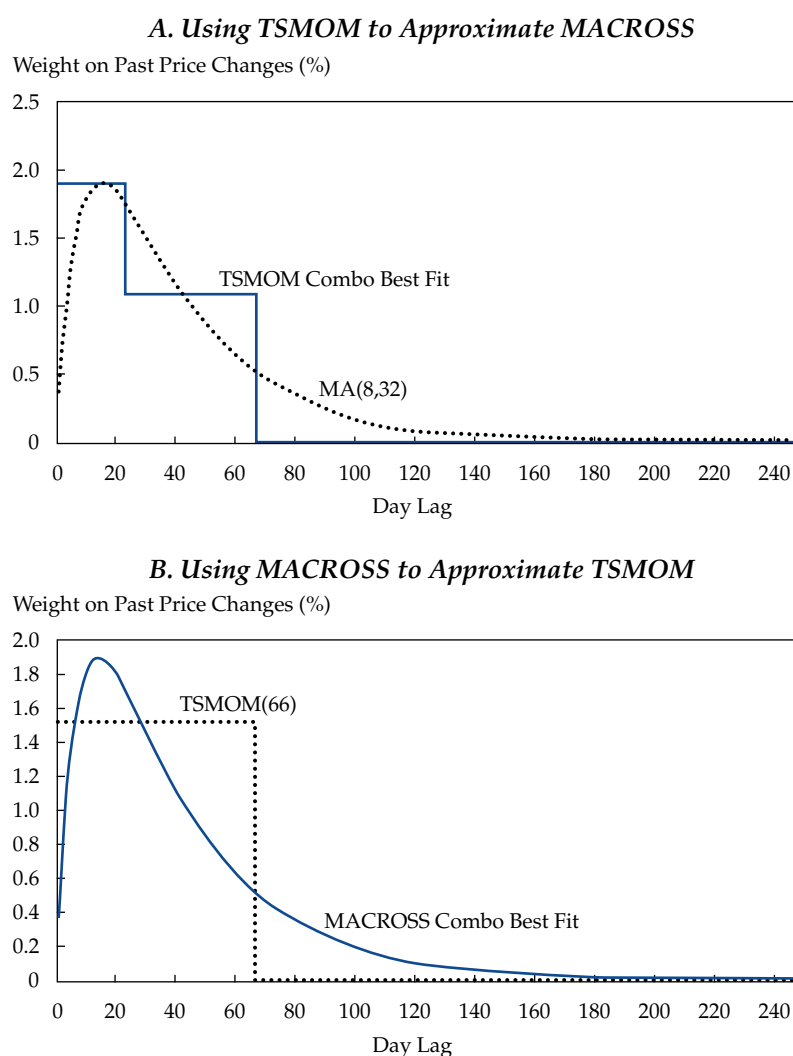
are closely related and that the yield strategies are quite correlated.

In addition, for each of the three MACROSS strategies, we see no significant alpha over the TSMOM signals. In other words, none of the three MACROSS signals considered here provided any significant performance benefit over a combination of TSMOM signals. This result is intuitive in light of our theoretical results. Somewhat surprisingly, however, in the case of MA (8,32), we see a significant negative alpha, meaning that this MACROSS factor would detract if added to the best-fit TSMOM portfolio.

Also surprisingly, we see positive significance in the alphas of some of the TSMOM signals when they were regressed on the MACROSS signals. The

fact that we see these significant alphas does not necessarily mean TSMOM specifications are superior to MACROSS specifications or vice versa. It may simply mean that MACROSS signals have a harder time mimicking a TSMOM signal whereas TSMOM signals, because of their shape, more easily fit an arbitrary MACROSS signal. This trait can be seen to some extent in the R^2 s, which are higher when TSMOM signals are the independent variables (although the difference is only a few percentage points). **Figure 14** illustrates this point graphically. Panel A shows how the TSMOM signals are well able to approximate the MA(8,32) signal by combining the three TSMOM signals with relevant weights. Panel B shows that the MACROSS signals

Figure 14. Approximating Moving-Average Weights with Time-Series Momentum and Vice Versa



Notes: Panel A shows the weight of the moving-average MACROSS(8,32) signal on past returns and the combined weights from the three time-series momentum (TSMOM) signals, weighted in proportion to the betas in the regression of MACROSS(8,32) on TSMOM in Table 2. Panel B shows the weight of the TSMOM(66) signal on past returns and the combined weights from the three MACROSS signals, weighted in proportion to the betas in the regression of TSMOM(66) on MACROSS in Table 2.

are not able to approximate the TSMOM weight as effectively.⁹

Furthermore, the alphas shown in Table 2 may reflect the relative performance of the trend horizons that are over- or underweighted by the best fit as seen in Figure 14. For instance, Figure 14 indicates that the MA(8,32) signal gives more weight to past returns 60–120 days ago, whereas the best-fit TSMOM portfolio gives more weight to returns 40–60 days ago. The 40–60-day returns may predict future returns more strongly, which would lead to a negative alpha. This result does not necessarily mean that TSMOM is the better way to invest because these issues may be addressed by changing the parameters of the MACROSS signals and including a wider array of MACROSS signals.

Conclusion

The academic literature and real-world investors have put forth a host of strategies that, on the surface, appear unique but that are all related to trend following at a high level. We sought to unify many of these seemingly disparate strategies in a simple, robust, and intuitive framework. We showed that trends can be filtered out from prices or returns by using a variety of methods, including time-series momentum, moving-average crossovers, and other popular filters. We proved that generalized forms of many trend-based investment strategies are equivalent, and we provided intuition for how the various approaches to trend following vary from strategy to strategy. Furthermore, we showed how each trend signal can be characterized by its “trend signature plots,” which illustrate the trend indicator’s dependence on past prices and returns.

Our results help demystify trend-following investing and put these strategies in a useful perspective for investors. Because each of these signals can be expressed in a unified framework, the conclusion is that the filtering methodology may matter less than the horizons chosen, the portfolio construction, risk management, and other factors that may be useful in identifying the quality of a trend. Our results suggest that investors and managers, rather than looking exclusively at which specific filter to start from, should focus on the robustness and quality of implementation—including optimally managing transaction costs,¹⁰ dynamic trading, diversification, position sizing, portfolio construction, and risk management.

We thank Cliff Asness, Brian Hurst, Ronen Israel, Toby Moskowitz, Yao Hua Ooi, and Dan Villalon for helpful comments and discussions and Jusvin Dhillon for excellent research assistance.

Appendix A. The HP Filter as TSMOM or MACROSS

To find the HP filter, we need to minimize the objective function:

$$\min_g \sum_{t=1}^T (P_t - g_t)^2 + \lambda \sum_{t=3}^T [g_t - g_{t-1} - (g_{t-1} - g_{t-2})]^2. \quad (\text{A1})$$

The objective function can be written in vector form:

$$\min_g (P - g)'(P - g) + \lambda g' \mathbf{K}' \mathbf{K} g, \quad (\text{A2})$$

where the matrix \mathbf{K} is of dimension $(T - 2)$ -by- T , defined as

$$\mathbf{K} = \begin{pmatrix} 1 & -2 & 1 & & 0 \\ & 1 & -2 & 1 & \\ & & \ddots & \ddots & \ddots \\ 0 & & & 1 & -2 & 1 \end{pmatrix}. \quad (\text{A3})$$

To solve this filtering problem, one differentiates the objective function and considers the first-order condition:

$$0 = -(P - g) + \lambda \mathbf{K}' \mathbf{K} g. \quad (\text{A4})$$

Hence, the solution for the growth component is

$$g = (I + \lambda \mathbf{K}' \mathbf{K})^{-1} P. \quad (\text{A5})$$

Equation A5 shows that the growth component is a linear combination of past prices. Indeed, the last row of the matrix $(I + \lambda \mathbf{K}' \mathbf{K})^{-1}$ contains the weights on past prices that give rise to the most recent growth component, g_T . Similarly, the second-to-last row gives the weights for the second most recent growth component, g_{T-1} . Finally, the difference between these components is the current trend, $\text{trend}_T = g_T - g_{T-1}$, which, therefore, is a MACROSS signal or, equivalently, a TSMOM signal.

Appendix B. The OLS Best-Fit Trend as TSMOM or MACROSS

The OLS trend regression estimates the parameters of the following model over some time window that is N periods long:

$$P_t = \alpha + \beta t + \varepsilon_t. \quad (\text{B1})$$

The relevant trend parameter is the estimated slope parameter, $\hat{\beta}$. A positive (negative) value of $\hat{\beta}$ indicates a positive (negative) time trend.

The OLS estimate of $\hat{\beta}$ as a function of the prices P , time t , and window length N is

$$\hat{\beta}_t = \frac{\sum_{s=1}^N (P_{t-s+1} - \bar{P}_{t,N}) \left(\frac{N+1}{2} - s \right)}{\sum_{m=1}^N \left(\frac{N+1}{2} - m \right)^2}, \quad (\text{B2})$$

where $\left(\frac{N+1}{2} - s \right)$ is the de-meaned time index series and $\bar{P}_{t,N}$ is the average price over the window:

$$\bar{P}_{t,N} = \frac{1}{N} \sum_{m=1}^N P_{t-m+1}. \quad (\text{B3})$$

The expression for the estimated slope can be rearranged as follows:

$$\begin{aligned} \hat{\beta}_t &= \sum_{s=1}^N P_{t-s+1} \frac{\left[(N+1)/2 - s \right]}{\sum_{m=1}^N \left[(N+1)/2 - m \right]^2} \\ &=: \sum_{s=1}^N w_s P_{t-s+1}, \end{aligned} \quad (\text{B4})$$

which is clearly a weighted sum of past prices. One can then use Equation 8 to express them as weights c of past price changes (i.e., asset returns).

Appendix C. Data Sources

Item	Source
<i>Commodity futures</i>	
Aluminum	London Metal Exchange
Copper	
Nickel	
Zinc	
Brent crude	
Gas oil	Intercontinental Exchange
Coffee	
Cocoa	
Cotton	
Sugar	
Corn	Chicago Board of Trade
Soybeans	
Soybean oil	
Soybean meal	
Wheat	
Lean hogs	Chicago Mercantile Exchange
Live cattle	
WTI crude oil	New York Mercantile Exchange
RBOB gasoline	
(spliced with unleaded)	
Heating oil	
Natural gas	
Gold	Commodity Exchange, Inc.
Silver	
Platinum	Tokyo Commodity Exchange
<i>Bonds</i>	
Australia 3-year bond	Datastream was used for futures returns, and JP Morgan bond index returns were used before futures returns were available.
Australia 10-year bond	
Euro Schatz	
Euro Bobl	
Euro Bund	
Euro Buxl	
Canada 10-year bond	
Japan 10-year bond	
Long gilt	
US 2-year note	
US 5-year note	
US 10-year note	
US long bond	

Item	Source
<i>Currencies</i>	
Australia	Spot exchange rates and forward interest rates from Citigroup were used to form return series after 1989. Prior to 1989, spot exchange rates from Datastream were combined with the Interbank Offered Rates from Bloomberg.
United Kingdom	
Germany (spliced with euro)	
Japan	
United States	
Norway	
Sweden	
Switzerland	
Canada	As above, with 1992 as the switchover point.
New Zealand	As above, with 1996 as the switchover point.
<i>Equity indexes</i>	
Australia (SPI 200)	Datastream was used for futures returns, and MSCI country index returns were used before futures returns were available.
France (CAC 40)	
Germany (DAX)	
Italy (FTSE MIB)	
Japan (TOPIX)	
Netherlands (AEX)	
Spain (IBEX 35)	
UK (FTSE 100)	
US (S&P 500)	

Notes

1. Hurst, Ooi, and Pedersen (2012, 2013) provided a detailed analysis of managed futures strategies and showed that the returns to the strategy can be largely explained by time-series momentum.
2. Silber (1994), Erb and Harvey (2006), and Moskowitz, Ooi, and Pedersen (2012) found strong performance of trend-following strategies, and Park and Irwin provided a survey (2007). Technical trading rules were analyzed broadly by Lo, Mamaysky, and Wang (2000) and Sullivan, Timmermann, and White (1999). Zakamulin (2015) made an independent analysis of the performance of market timing with moving averages.
3. Fama (1965) provided a detailed summary of the random walk hypothesis for stocks.
4. To see this equivalency, note that the coefficients on the price at each time, P_{t-s+1} , must be equalized in the two different ways of writing the trend signal (i.e., $c_s - c_{s-1} = w_s^{fast} - w_s^{slow}$). This equation can be iterated to arrive at the expression for c given the initial value $c_1 = w_1^{fast} - w_1^{slow}$.
5. One particular choice of MA functions with positive weights that add up to 1 is as follows. For the fast MA, let $w_j^{fast} = c_j / \bar{c}$ for all j , where $\bar{c} = \sum_{j=1}^{\infty} c_j$. For the slow MA, let $w_1^{slow} = 0$ and $w_j^{slow} = c_{j-1} / \bar{c}$ for $j \geq 2$.

6. Specifically, any filter $f(\cdot)$ on data series P that is causal (i.e., only depends on the past), linear [i.e., $f(aX + Y) = af(X) + f(Y)$], and time invariant (i.e., the function does not depend on time) can be represented as a sum of weighted past values of P .
7. Slightly more formally, $\text{return}_t = \text{trend}_t + \varepsilon_t$, and $\text{trend}_t = \text{trend}_{t-1} + \eta_t$, where η_t and ε_t are both independent and identically distributed variables, normally distributed with mean 0 and constant variance.
8. A wavelet is a wave-like oscillation with an amplitude that begins at zero, increases, and then decreases back to zero. It is a popular filtering technique because it can extract both time and frequency information from a series. For a full treatment of using wavelet filtering in financial time series, see Gençay, Selçuk, and Whitcher (2001).
9. Each plot in Figure 14 shows a signal's effective weight on past asset returns. It also shows the weighted average of the weighting schemes corresponding to the four explanatory variables weighted in proportion to the betas in the regression. It is meant to be stylized; that is, it does not perfectly represent the regression because the regression analysis was performed on strategy returns, not the underlying signals. The relationship between signal construction and return correlations is far from perfect because of the nonlinear portfolio construction, but the stylized results are, nonetheless, informative.
10. See Gârleanu and Pedersen (2013).

References

- Asness, Clifford S., Tobias J. Moskowitz, and Lasse H. Pedersen. 2013. "Value and Momentum Everywhere." *Journal of Finance*, vol. 68, no. 3 (June): 929–985.
- Brock, William, Josef Lakonishok, and Blake LeBaron. 1992. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance*, vol. 47, no. 5 (December): 1731–1764.
- Erb, Claude, and Campbell Harvey. 2006. "The Strategic and Tactical Value of Commodity Futures." *Financial Analysts Journal*, vol. 62, no. 2 (March/April): 69–97.
- Fama, Eugene F. 1965. "The Behavior of Stock-Market Prices." *Journal of Business*, vol. 38, no. 1 (January): 34–105.
- Gârleanu, Nicolae, and Lasse H. Pedersen. 2013. "Dynamic Trading with Predictable Returns and Transaction Costs." *Journal of Finance*, vol. 68, no. 6 (December): 2309–2340.
- Gençay, Ramazan, Faruk Selçuk, and Brandon Whitcher. 2001. *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. San Diego: Academic Press.
- Harvey, A.C. 1984. "A Unified View of Statistical Forecasting Procedures." *Journal of Forecasting*, vol. 3, no. 3 (July/September): 245–275.
- Henderson, R. 1924. "A New Method of Graduation." *Transactions of the Actuarial Society of America*, vol. 25: 29–40.
- Hodrick, Robert J., and Edward C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, vol. 29, no. 1 (February): 1–16.
- Hurst, Brian K., Yao Hua Ooi, and Lasse H. Pedersen. 2012. "A Century of Evidence on Trend-Following Investing." Working paper, AQR Capital Management.
- . 2013. "Demystifying Managed Futures." *Journal of Investment Management*, vol. 11, no. 3 (Third Quarter): 42–58.
- Kalman, R.E. 1960. "A New Approach to Linear Filtering and Prediction Problems." *Journal of Basic Engineering*, vol. 82, no. 1 (March): 35–45.
- Lo, Andrew W., Harry Mamaysky, and Jiang Wang. 2000. "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation." *Journal of Finance*, vol. 55, no. 4 (August): 1705–1770.
- Moskowitz, Tobias J., Yao Hua Ooi, and Lasse H. Pedersen. 2012. "Time Series Momentum." *Journal of Financial Economics*, vol. 104, no. 2 (May): 228–250.
- Okunev, John, and Derek White. 2003. "Do Momentum-Based Strategies Still Work in Foreign Currency Markets?" *Journal of Financial and Quantitative Analysis*, vol. 38, no. 2 (June): 425–447.
- Park, Cheol-Ho, and Scott H. Irwin. 2007. "What Do We Know about the Profitability of Technical Analysis?" *Journal of Economic Surveys*, vol. 21, no. 4 (September): 786–826.
- Silber, William L. 1994. "Technical Trading: When It Works and When It Doesn't." *Journal of Derivatives*, vol. 1, no. 3 (Spring): 39–44.
- Sullivan, Ryan, Allan Timmermann, and Halbert White. 1999. "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap." *Journal of Finance*, vol. 54, no. 5 (October): 1647–1691.
- Whittaker, E.T. 1923. "On a New Method of Graduations." *Proceedings of the Edinburgh Mathematical Society*, vol. 41: 63–75.
- Zakamulin, Valeriy. 2015. "Market Timing with Moving Averages: Anatomy and Performance of Trading Rules." Working paper, University of Agder (26 August).

Open Access: This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.