

Which Trend is Your Friend?

Ari Levine and Lasse Heje Pedersen *

Abstract

Managed-futures funds and CTAs trade predominantly on trends. There are several ways of identifying trends, either using heuristics or statistical measures often called “filters.” Two important statistical measures of price trends are time series momentum and moving average crossovers. We show both empirically and theoretically that these trend indicators are closely connected. In fact, they are equivalent representations in their most general forms, and they also capture many other types of filters such as the HP filter, the Kalman filter, and all other linear filters. Further, we show how these filters can be represented through “trend signature plots” showing their dependence on past prices and returns by horizon. Our results unify and broaden a range of trend-following strategies and we discuss the implications for investors.

Keywords: trend-following, momentum, moving average crossover, filtering, managed futures, CTA

JEL codes: G02, G1, G13, G14

First version: Nov., 2010. This version May, 2015

* Ari Levine is at AQR Capital Management, Two Greenwich Plaza, Greenwich, CT 06830, e-mail: ari.levine@aqr.com. Lasse H. Pedersen (corresponding author) is at Copenhagen Business School, New York University, AQR Capital Management, and CEPR; e-mail: lhpedersen@gmail.com; web: www.lhpedersen.com. 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.

1. Introduction

Trend-following investing is the predominant investment style for managed-futures hedge funds, commodity trading advisors (CTAs) and certain macro traders.¹ Trend-following investing can be defined loosely as buying when prices have been rising and (short-)selling when prices have been falling, 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 to one another? These are the questions we seek to address in this paper.

To put our findings in a broader perspective, note that since 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” efficient-market hypothesis, future price moves are completely unpredictable, meaning that trend-following strategies should not work.³ However, price trends may exist if markets are not completely efficient or if risk premia 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. As an 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 the gross domestic product and unemployment rate collected from many sources with errors – try to assess whether the economy is heading into recession or is over-heating. 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 wants to hear. Along the same lines, trend-followers have employed quantitative tools to enhance their signal of the price trend and reduce the noise around it.

In finance, one simple approach to capture price trends is time series momentum (TSMOM) as defined by Moskowitz, Ooi, and Pedersen (2012). The simplest form of a time series momentum signal is the return over some recent time period, e.g., the return over the past 12 months. 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 was negative, the trend is assessed to be downward, and the TSMOM investor short sells gold. Moskowitz, Ooi, and Pedersen (2012) show that investing based

¹ Hurst, Ooi, and Pedersen (2012, 2013) provide a detailed analysis of managed futures strategies and show 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) find strong performance of trend-following strategies, and Park and Irwin provide a survey (2007). Technical trading rules are analyzed more broadly by Lo, Mamaysky and Wang (2000) and Sullivan, Timmermann, and White (2000). Zakamulin (2015) makes an independent analysis of the performance of market timing with moving averages.

³ Fama (1965) provides a detailed summary of the random walk hypothesis for stocks.

on 12-month TSMOM was on average profitable for each of the 58 liquid securities they analyzed over the last 25 years. Their universe included instruments across 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), and Okunev and White (2003)). A CTA using this method will buy when one moving average of recent prices crosses another moving average of prices measured over a longer horizon. The idea is that a “fast” moving average captures the average recent prices, while a “slower” moving average capture 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 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. Hence, any trend signal that can be designed using MACROSS techniques can be similarly designed as a TSMOM strategy. Similarly, TSMOM is a special case of MACROSS signals.

As a different way to state this result, we show how trend filters can be equivalently represented as functions of past prices vs. past returns. We show how to move from prices to returns and vice versa and illustrate these “trend signatures” graphically. We also show how a large class of filtering methods used in science and economics can be viewed as a special case of TSMOM, including the HP filter of Hodrick and Prescott (1997), the Kalman filter, a large set of linear filters, and regression-based methods.

Lastly, we perform an empirical study that compares the performance of common implementations of TSMOM strategies with that of MACROSS strategies. Consistent with our theoretical results, we find that these strategies are closely related empirically and discuss the reasons for the performance differences.

Our results have several implications for trend-following investors. First, both TSMOM and MACROSS are effective in filtering trends. Second, since these methods are equivalent in their most general forms, excellence in trend following does not depend on *which* of these filters is used, but, rather, *how* they are used. In particular, the performance depends more crucially on the choice of trend horizons, diversification across instruments, portfolio construction, risk management, optimally managing transaction costs, and efficient dynamic trading and order routing. Third, our method shows how trend-following strategies can be viewed through “trend signature plots,” showing their dependence on past prices and returns by horizon.

2. Time Series Momentum vs. Moving-Average Crossover

Time Series Momentum

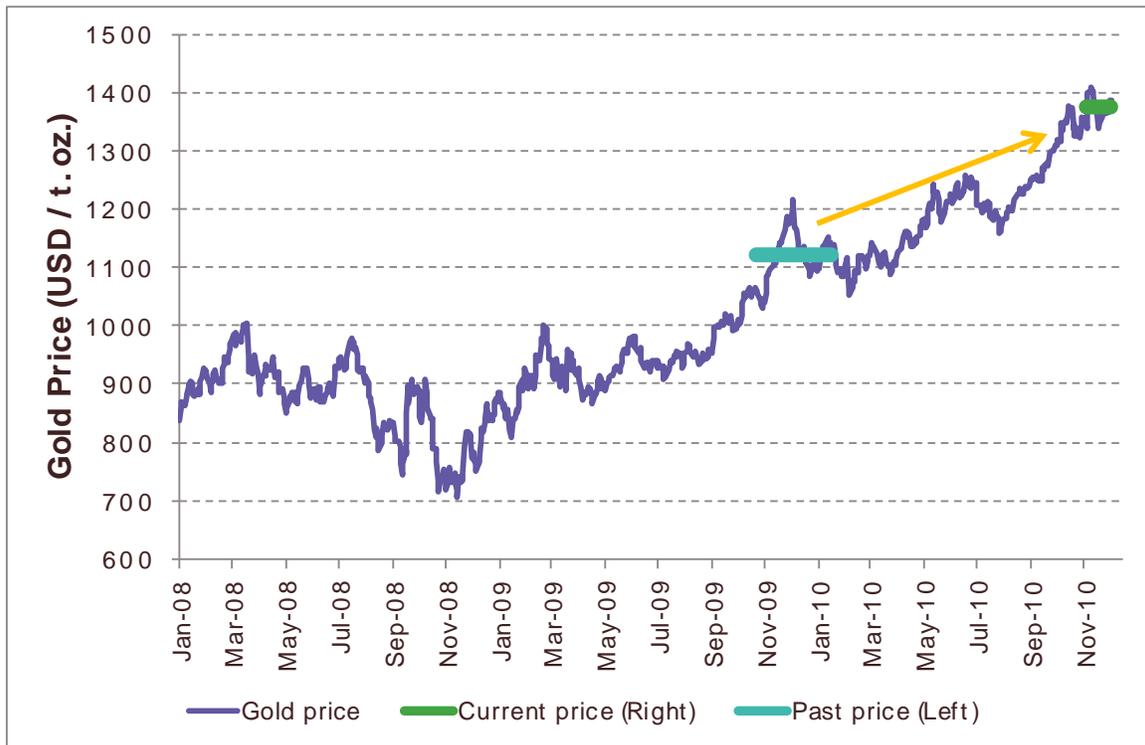
A TSMOM strategy goes long when prices have been moving up, and short when prices have been moving down. The simplest TSMOM signal is the past return over some time period, say m months or days:

$$TSMOM_t^m = return_{t-m,t} \quad (1)$$

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 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, using a return index in the place of prices). In this study, we focus on differences in log prices or index levels for simplicity.

As an example, Figure 1 shows how gold prices have positive TSMOM at the end of 2010 as prices had been trending upward.

Figure 1. Time Series Momentum of Gold Futures. *This figure shows gold prices over a two-year period. The arrow shows the filtered trend from the smoothed past price (left horizontal bar) to the smoothed current price (right horizontal bar).*



There are also more refined TSMOM signals. One way to refine the TSMOM signal is to smooth the prices used to calculate the return:

$$TSMOM_t^m = \text{average}(P_s : s \text{ near current time } t) - \text{average}(P_{s'} : s' \text{ near lagged time } m) \quad (2)$$

Smoothing can be a good idea because it reduces random noise in the data. For instance, focusing on a single past price may be arbitrary and subject to more noise than using an average of multiple past prices (which can be called “back-end smoothing”). Figure 1 uses the average price over the 60-day period indicated by the horizontal line on the left. As an example, Asness, Moskowitz, and Pedersen (2013) use 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 can be suboptimal if recent prices contain important information about the current trend or a trend-reversal, but 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, which we call c :

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

This equation 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. Ooi, Moskowitz, and Pedersen (2012) conduct 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 (MA) of prices, which we call MA^{fast} and MA^{slow} . The fast MA puts more weight on recent prices, whereas 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. In Figure 2 we plot these for gold prices over the same period shown in Figure 1.

Figure 2. Moving-Average Crossover Indicator for Gold Futures. *This figure shows the gold price, a fast 20-day moving average (MA), and a slower 260-day MA. Given that the fast MA is above the slow one at the end of the sample, the filtered trend is up as indicated by the arrow.*



The MACROSS strategy depends on which MA is higher: the fast one or the slow one. As we can see in the figure, the fast 20-day MA is above the slow 260-day MA at the end of the time period that we plot. This means that recent prices are above past prices, resulting in an upward trend.

Of course, this is just one example of an MACROSS strategy. Other strategies arise by varying the time horizons (here the 20 and 260 day averaging periods). Further, we need not give each day's price an equal weight in the moving average. Another common method is to weight past prices exponentially (called exponentially-weighted moving average as discussed in more detail later). More generally, we can compute the MA's using any weighting scheme, which we will denote by w (where the weights can for instance be equal weights or exponential weights as described in the examples further below). Hence, in general we can write the MA's mathematically as:

$$\begin{aligned}
 MA_t^{fast} &= \sum_{s=1}^{\infty} w_s^{fast} P_{t-s+1} \\
 MA_t^{slow} &= \sum_{s=1}^{\infty} w_s^{slow} P_{t-s+1}
 \end{aligned}
 \tag{4}$$

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

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

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

$$MACROSS_t = MA_t^{fast} - MA_t^{slow} \quad (6)$$

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.

Moving Average Crossover as Time Series Momentum

The *MACROSS* signal is the difference between two MA's and therefore a combination of past prices:

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

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

Specifically, the *MACROSS* equation (6) is equivalent to the *TSMOM* strategy (3) with coefficients on past returns c_s computed as follows⁴

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

These implied coefficients c_s are positive for all *MACROSS* strategies where the fast MA is uniformly faster than the slow MA as given by Equation (5), which is true for the standard *MACROSS* strategies. It is natural that these *TSMOM* coefficients are positive since this means that the strategy is trend following (whereas negative coefficients would have indicated a bet on trend reversal).

Further, the implied return weights c_s approach zero as the number of lags s increases (assuming that the weights w_s^i sum to 1 for each i). While the coefficients c_s have no specific "scale," it is natural to normalize them such that they to sum to one:

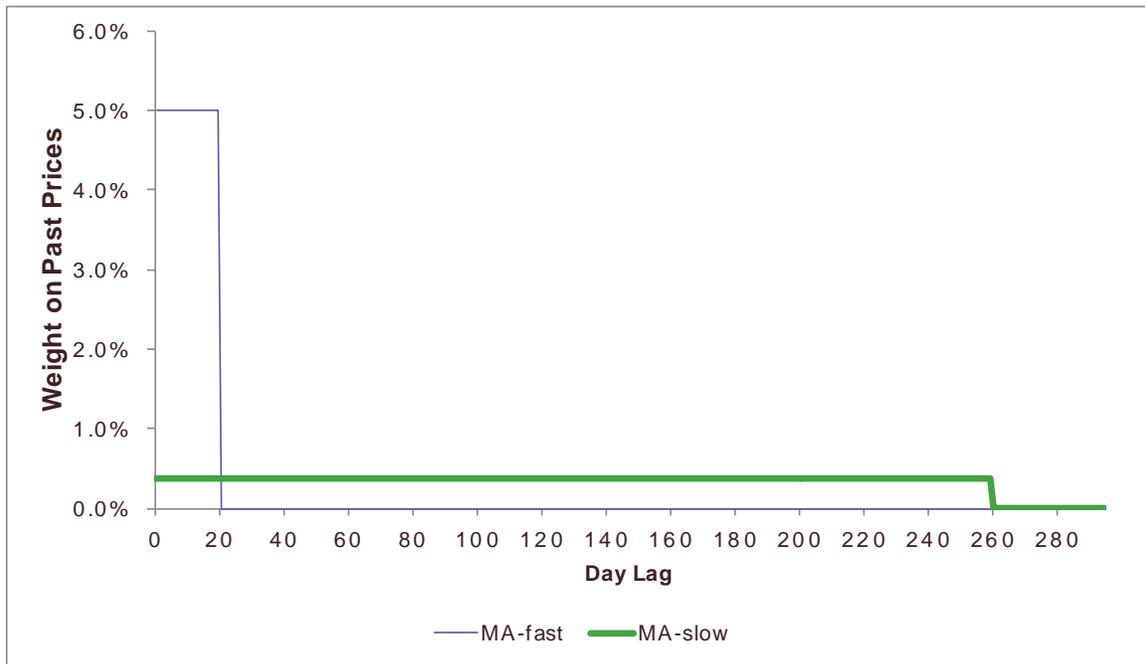
⁴ To see this, note that the coefficients on the price P_{t-s+1} at each time must be equalized in the two different ways of writing the trend signal, i.e., $c_s - c_{s-1} = w_s^{fast} - w_s^{slow}$. This can be iterated to arrive at the expression for c given the initial value $c_1 = w_1^{fast} - w_1^{slow}$.

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

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

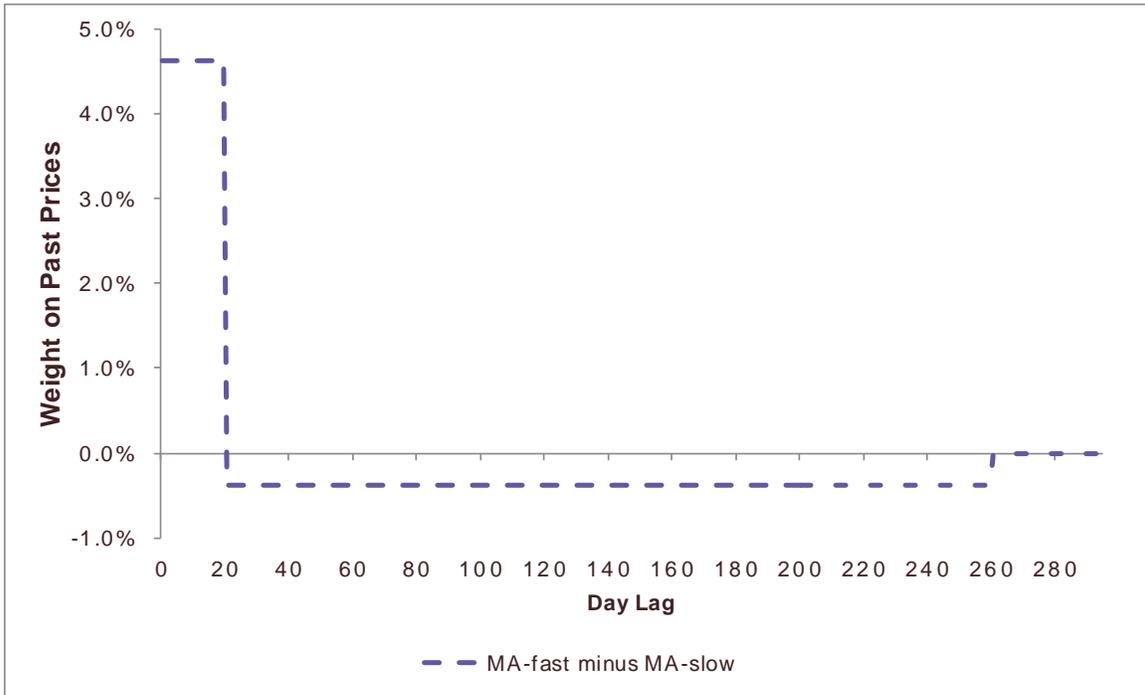
To understand this conversion from MACROSS to TSMOM (i.e., the conversion from “price space” to “return space”), consider the 20-day vs. 260 day equal weighted MA strategy. This strategy compares past prices based on the coefficients illustrated in Figure 3.

Figure 3. Moving-Average Crossover Coefficients: Fast and Slow Averages. *The fast moving average (MA) is an equal-weighted average over the past 20 days, while the slow MA is an equal-weighted average of the past 260 days.*



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

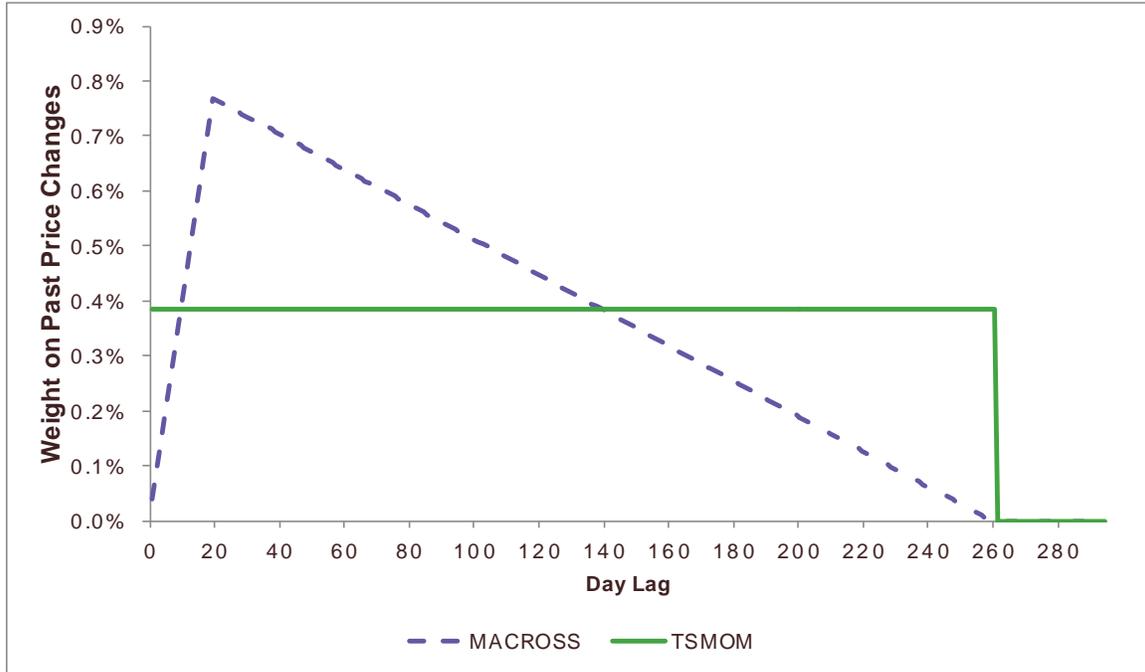
Figure 4. Price Signature Plot: Moving-Average Crossover, Equal-Weighted. *As seen in equation (7), the MACROSS signal is computed as a weighted average of past prices, where the weights are the fast MA(20) minus the slow MA(260). The resulting signal puts positive weights on the most recent 20 days 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 weights on recent prices is positive, the weights on distant prices are negative, and the weights eventually go to zero. This shows that a moving average strategy can be interpreted as a TSMOM strategy where both the front-end price and the back-end price have been smoothed. Indeed, the MACROSS strategy is like a TSMOM strategy where 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 days 21 to 260, and the return is then computed as the difference between these two smoothed prices.

We can also use equation (8) above to translate the MACROSS coefficients for price levels into TSMOM coefficients for price *changes*, i.e., returns. The coefficients in “return space” are seen in Figure 5, where we have normalized the weights to sum to 1 in both cases.

Figure 5. Return Signature Plot: Trend Coefficients in Return Space. *This figure shows how much weight the trend indicator places on each daily return in the past for the simple 260-day time series momentum (TSMOM) and the 20/260 moving-average crossover (MACROSS), respectively, where the weights are normalized to sum to one. That is, a trend indicator can be viewed as an average of daily returns (rather than prices), as the figure shows.*



The graph shows two lines: The solid line shows 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 dotted 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 using price changes instead. The MACROSS assigns the most importance to intermediate price changes, and lesser weight to the most recent price changes and very old price changes.

Time Series Momentum as Moving-Average Crossover

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 P_{t-12M}):

$$TSMOM_t^m = P_t - P_{t-m} \quad (10)$$

This shows that a TSMOM strategy can be viewed as an MACROSS where the recent MA^{fast} is simply the current price P_t (i.e., the weighting scheme puts all the weight in one

price, namely the most recent one : $w_1^{fast} = 1$, $w_s^{fast} = 0$ for $s > 1$). Similarly, the distant MA^{slow} is simply the lagged price (i.e., its weighted scheme puts all its weight on that price: $w_m^{slow} = 1$).

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

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

$$\begin{aligned} w_1^{fast} - w_1^{slow} &= c_1 \\ w_j^{fast} - w_j^{slow} &= c_j - c_{j-1} \end{aligned} \tag{11}$$

This shows how to choose the *difference* between the weights of recent MA^{fast} and the distant MA^{slow} . However, there are many choices of moving averages that produce the same signal, since adding and subtracting the same price has no effect. In contrast, the momentum weights c are unique, as are the weights on past prices $w_s^{fast} - w_s^{slow}$, so these are more fundamental parameters of the filtering process.⁵

Example: Exponentially-Weighted Moving-Average Crossover

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, with an exponential decay of $\theta > 0$, we have

$$EWMA_t = \frac{1}{1-\theta} \sum_{j=0}^{\infty} \theta^j P_{t-j} \tag{12}$$

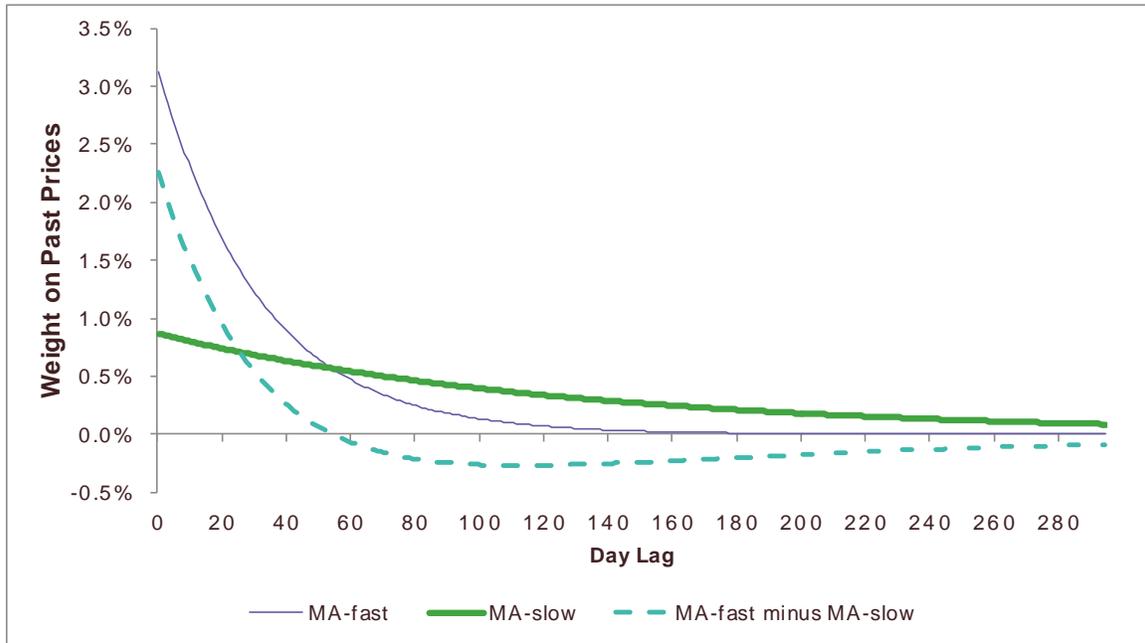
Instead of parameterizing by the decay θ , it is more intuitive to consider the center of mass (COM) of the moving average, defined as

$$COM = \frac{1}{1-\theta} \sum_{j=0}^{\infty} \theta^j j = \frac{\theta}{1-\theta} \tag{13}$$

⁵ One particular choice of MA functions with positive weights that add up to one is as follows. For the fast MA, we can let $w_j^{fast} = c_j / \bar{c}$ for all j , where $\bar{c} = \sum_{j=1}^{\infty} c_j$. For the slow MA, we can let $w_1^{slow} = 0$ and $w_j^{slow} = c_{j-1} / \bar{c}$ for $j \geq 2$.

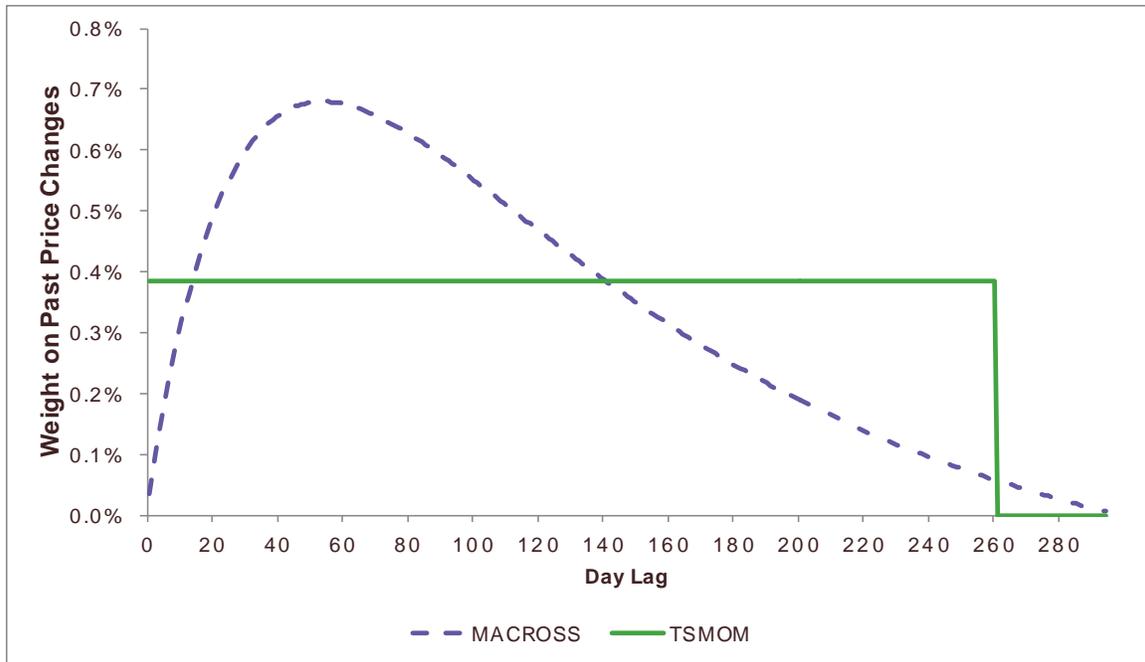
The center of mass can be useful for forming an intuition about the effective length of the moving average. For an exponentially weighted MACROSS, the price weights $w_j = \frac{\theta^j}{1-\theta}$ look similar to those on a simple MACROSS, but smoother, as we show in Figure 6.

Figure 6. Price Signature Plot: Moving-Average Crossover with Exponential Weights. *This figure shows how much weight an exponentially weighted moving average (EWMA) puts on past prices. The fast EWMA uses a center of mass COM=32, and the slow EWMA uses COM=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.

Figure 7. Return Signature Plot: Moving-Average Crossover using Exponential Weights. *This figure shows the weights that an exponentially weighted moving average (EWMA) crossover puts on past returns. The centers of mass used for the EWMA are 32 and 128 for the fast and slow EWMA, respectively. For comparison, a 260 day time series momentum (TSMOM) signal is also shown. Weights have been normalized to sum to one in both cases.*



3. On the Equivalence of All Other Linear Filters

The literature on signal processing includes many types of linear filters across 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 . More 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 are non-linear 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 account for such effects as well. For example, we have thus far limited ourselves to strictly positive weights c . If we loosen this restriction to include sets of weights whose *sum* 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 weight on more recent returns, but 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.

⁶ Specifically, any filter $f(\cdot)$ on data series P that is a causal (i.e., only depends on the past), linear (i.e., $f(aX+Y)=af(X)+f(Y)$), and time-invariant (i.e., function does not depend on time) can be represented as a sum of weighted past values of P .

The Hodrick Prescott Filter

TSMOM and MACROSS techniques also capture as a special case the so-called Hodrick and Prescott (HP) filter. Hodrick and Prescott (1997) report that the method is also called the Whittaker-Henderson Type A method (Whittaker (1923)) and has been used in actuarial sciences to smooth mortality rates; in astronomy (e.g., by Schiaparelli in 1867); and in ballistics (e.g., by von Neuman 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.

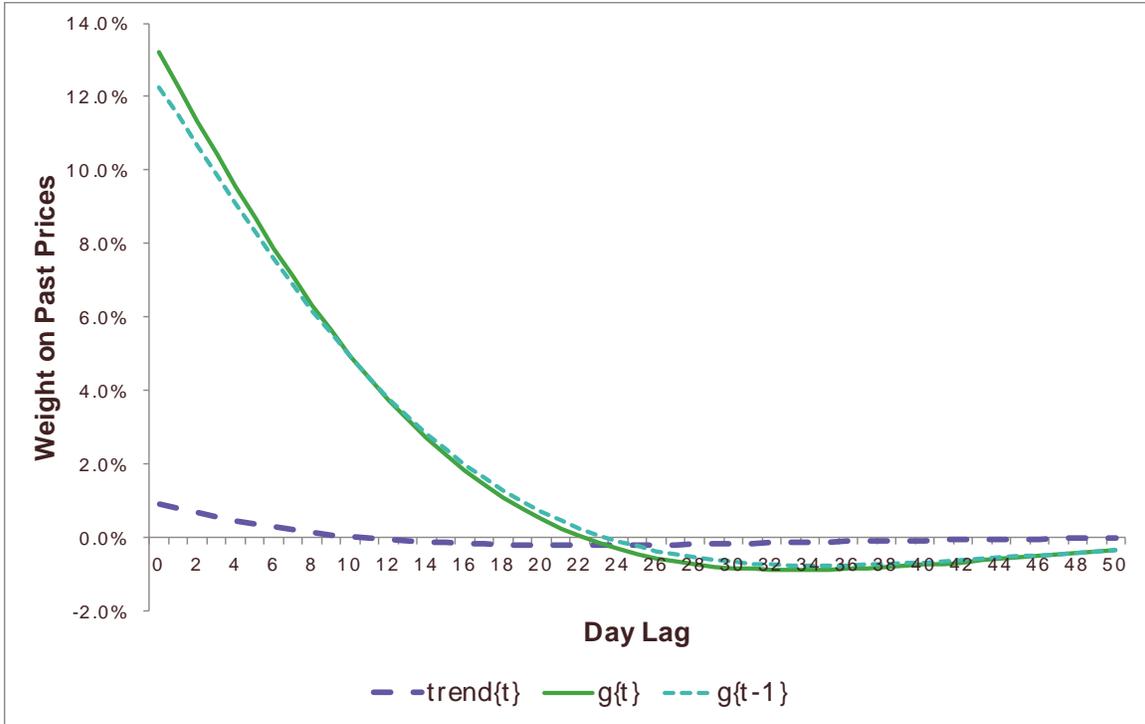
The HP filter is based on an idea that prices (or GDP) have a growth component g and a cyclical component z , i.e., $P_t = g_t + z_t$. The trend is the change of the smooth growth component, $trend_t = g_t - g_{t-1}$. The growth component is filtered from the price data by finding a growth path that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

$$\min_{g_1, \dots, g_T} \sum_t (z_t)^2 + \lambda \sum_t (trend_t - trend_{t-1})^2 \quad (14)$$

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

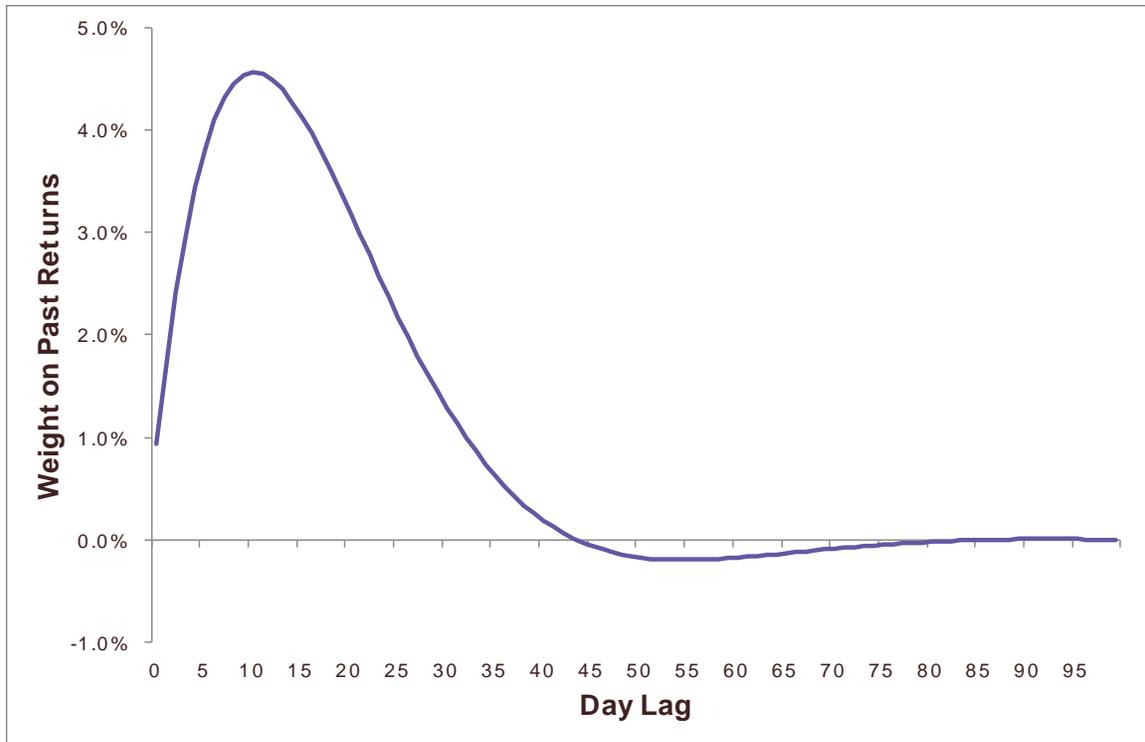
We show in Appendix A that the growth component g_t is a moving average of past prices. Hence, the $trend_t = g_t - g_{t-1}$ is a difference between two moving averages (the MA at time t and the MA at time $t-1$), and, therefore, it is an MACROSS signal.

Figure 8. Price Signature Plot: Hodrick and Prescott Filter. *This figure shows the weights on prices for the two growth components, g_t and g_{t-1} , as well as for $trend_t$, their difference, for $\lambda=10^4$. Note that while g_{t-1} does depend on P_t , the final trend does not depend on any future information.*



Further, since MACROSS's are TSMOM signals, the HP filter is also a TSMOM signal. That is, the HP trend can be written as a weighted average of past price changes.

Figure 9. Return Signature Plot: Hodrick and Prescott Filter. *This figure shows the weights on past returns for $trend_t$, for $\lambda=10^4$.*



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. However, the weights on returns are not strictly positive – there is a small amount of negative weight in the filter. Thus, the filter can be thought of as a combination of simple TSMOM plus a small amount of acceleration.

The Kalman Filter

The Kalman Filter (Kalman (1960)) is a technique which can be used to optimally estimate hidden variables of dynamic linear systems with noisy observations. A full treatment of the Kalman Filtering technique is beyond our scope here. But in the context of trend detection, the Kalman Filter can be applied to estimate the underlying (and hidden) trend variable driving returns.

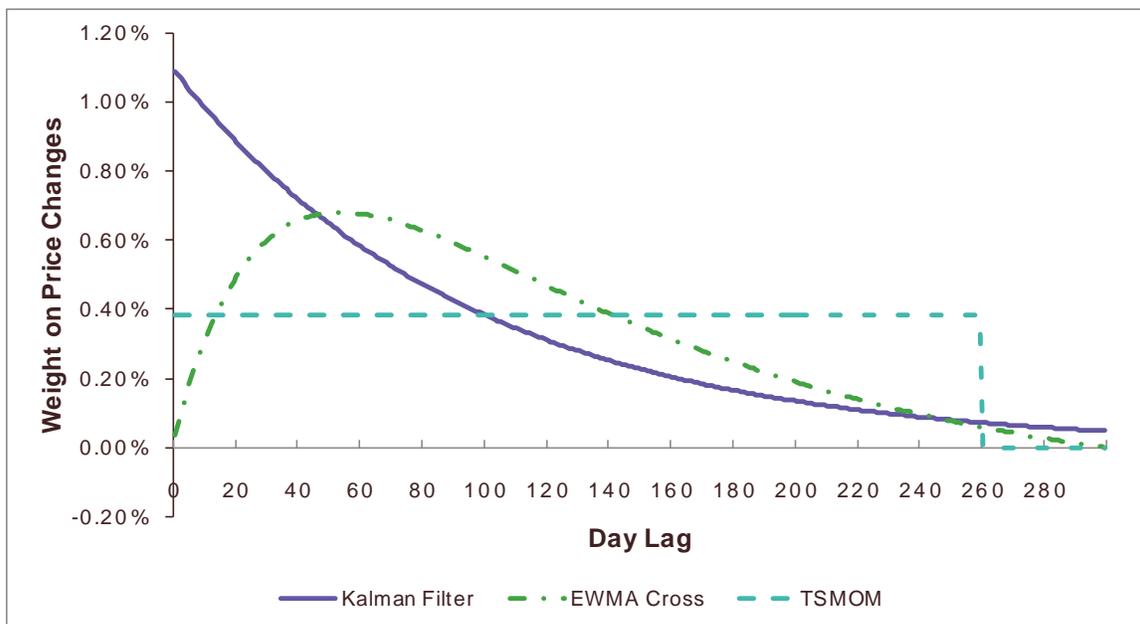
The particular application of the Kalman Filter will depend on the model used for the underlying data. If we know more about the underlying dynamics of the system, we can put more structure around the model, which may help in estimating the parameters of a Kalman Filter. But for price data, even simple random walk models tend to capture most of the important dynamics of price series. So it is not clear that it would be productive to add more structure to the underlying data generating process. If we are striving for 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 a random

walk whose value is not directly observable.⁷ The Kalman Filter can then be used to estimate this underlying trend.

Harvey (1984) shows 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, absent 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 center-of-mass 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 is different than the simple TSMOM or EWMA crossovers, as shown in Figure 10:

Figure 10. Return Signature Plot: Kalman Filter. *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 center of mass of 96 is used to form the exponential weights. Also shown are weights for an exponentially weighted moving average (EWMA) crossover with centers of mass of 32 and 128 for the fast and slow EWMA, respectively, as well as weights for a simple 260 day time series momentum (TSMOM) signal.*

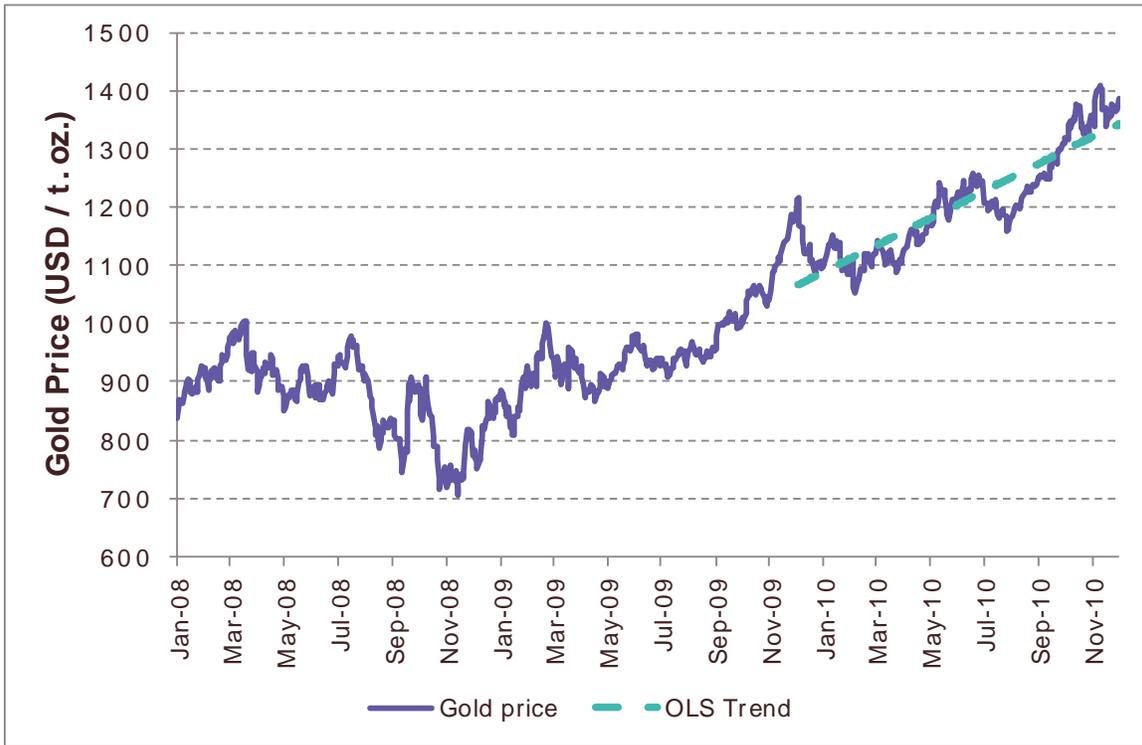


⁷ Slightly more formally, $return_t = trend_t + \varepsilon_t$, and $trend_t = trend_{t-1} + \eta_t$, where η_t and ε_t are both iid, normally distributed with mean 0 and constant variance, and independent of one another.

Trend Estimation Using Ordinary Least Squares Trend Regression

Another intuitive trend-measurement method which turns out to be equivalent to a generalized TSMOM signal is a regression-based trend estimate on prices. In order 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. For example, Figure 11 shows how this works for a one-year trend estimate in gold prices.

Figure 11. OLS One-Year Trend Estimate of Gold Prices. *This figure shows the price of gold in USD / t. oz., and an OLS trendline estimated using the last year of data shown.*



This trend estimation methodology turns out to be 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*.

This set of weights in Figure 12 below is parallel to an MACROSS – recent prices get positive weight, and more distant past prices get negative weight. For the sake of comparison, the simple MACROSS weights are shown as well. This can also be translated into weights on returns using equation (8) as seen in Figure 13.

Figure 12. Price Signature Plot: OLS Trend. *This figure shows the weights on prices for the OLS trend estimator which estimates a best fit trend-line through the past 260 days of prices, compared to the simple 20/260 moving average crossover (MACROSS) signal.*

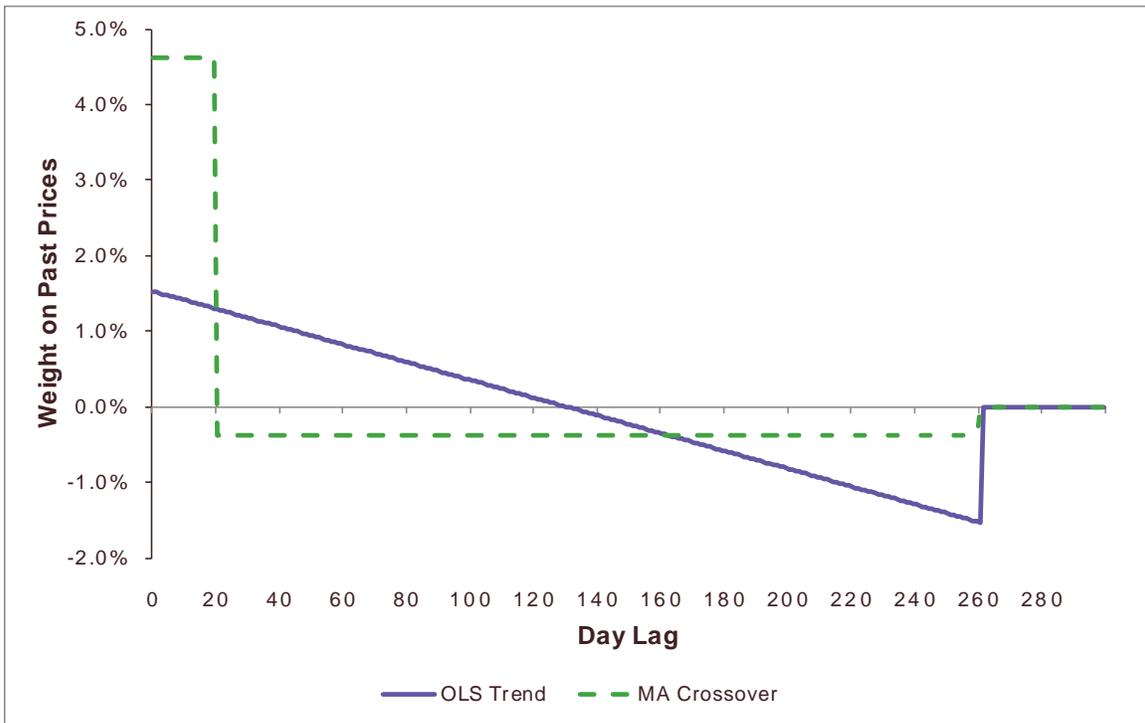
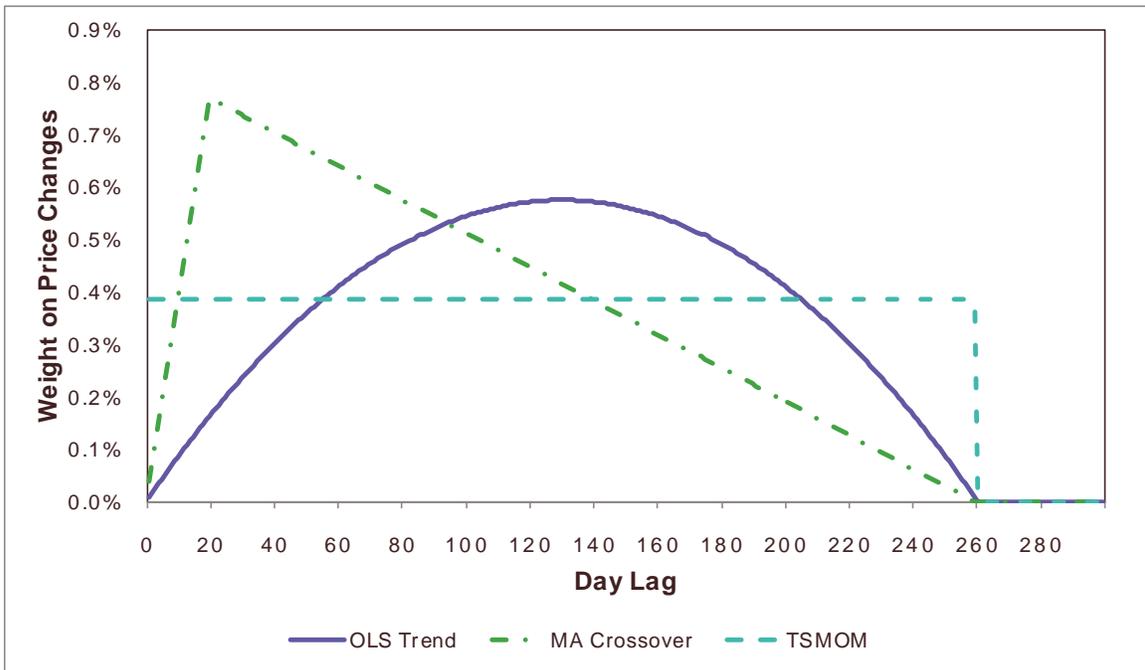


Figure 13. Return Signature Plot: OLS Trend. This figure shows the weights on past price changes, or returns, for the OLS trend estimator, compared to the simple 20/260 MACROSS signal and the simple 260-day TSMOM signal. Each of the three sets of weights is normalized to sum to one.



We see that the OLS regression method gives the most weight to returns at the center of the window, while returns that are at the extremes of the window (either most recent or most distant) are down-weighted by this scheme.

There are many other operations which ultimately are versions of weighting past returns, similar to TSMOM. For example, using the multi-resolution approach, wavelets can be used to extract trends at various resolutions from a price series.⁸

4. 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 focus our empirical study on these trend indicators.

Despite this equivalence between the most general classes of TSMOM and MACROSS, there is room for differences to emerge between the common implementations of those signals as well as from non-linear transformations done as part of portfolio construction. Therefore, it is interesting to empirically study how simple TSMOM strategies compare to simple MACROSS strategies.

Data

We use prices from 24 commodity futures, 13 developed government bond futures, 12 currency pairs from 9 underlying currencies, and 9 developed equity indices. These 58 instruments were chosen for their liquidity by Moskowitz, Ooi, and Pedersen (2012), and we extend their dataset so that our data covers prices from January 1985 through April 2015. Signals are calculated from a return index (rather than from prices directly) that is formed by rolling futures and forward prices, and which therefore implicitly incorporates 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. Since futures and forwards have implicit financing, these return indices are naturally excess of cash. The list of instruments and their sources are detailed in Appendix C.

Methodology

We seek to construct three standard TSMOM strategies and three standard MACROSS strategies that are relatively comparable. For the TSMOM strategies, we consider 1-

⁸ For a full treatment of using wavelet filtering in financial time series, see Gencay, Selcuk, and Whitcher (2001)

month, 3-month, and 12-month trends, following the methodology of Hurst, Ooi, and Pedersen (2013). Specifically, the TSMOM signals are parameterized by the number of look-back days such that $TSMOM(n)$ is calculated as the log return index today minus the log return index n days ago.

$$signal_t^{TSMOM(n)} = P_t - P_{t-n} \quad (15)$$

We consider values of n equal to 22 trading days (approximately 1 month), 66 trading days (3 months), and 260 trading days (12 months). Every day, we form a new portfolio that is held for one day. The results are similar for one-month holding periods, except that the Sharpe ratios are naturally lower. We use daily rebalancing to focus on the connections between the different trend strategies with minimal noise due to infrequent rebalancing.

Likewise, we consider three MACROSS strategies at similar horizons. We use exponentially-weighted MACROSS signals since these are perhaps the most common in investment practice. The MACROSS signals are parameterized by the center of mass of the fast and slow moving averages, such that an $MACROSS(m, M)$ signal has an m -day center of mass for its fast moving average and an M -day center of mass for its slow moving average:

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

where the weights are $w_s^m = \frac{\theta^s}{1-\theta}$ with $\theta = \frac{m}{1+m}$ as in equations (12)-(13). We choose m and M such that the MACROSS strategies correspond to the TSMOM signals by ensuring that they have similar trend horizons. Specifically, we let M take the values 12, 32, 128 and set m to one quarter of these values. These values are chosen such that M is close to $n/2$ (and divisible by four) for the corresponding $TSMOM(n)$ signals. These choices of M are natural since a $TSMOM(n)$ signal has a center of mass equal to $n/2$ because it gives equal weight to the past n returns.

To put the trading signals on an equal footing, we use the same portfolio construction methodology for TSMOM and MACROSS signals, following Moskowitz, Ooi, and Pedersen (2012). Specifically, for each strategy, our position in an asset i at time t is calculated as follows:

$$position_t^i = 0.65\% * \frac{sign[signal_t^i]}{\sigma_t^i} \quad (17)$$

where $signal_t^i$ is the relevant TSMOM or MACROSS signal and σ_t^i is the volatility of asset i at time t , estimated using an exponentially weighted volatility with a center of mass of 60 days, again following Moskowitz, Ooi, and Pedersen (2012). We multiply by 0.65% in order to target an annualized volatility of 0.65% in each asset. When aggregated, this level of asset target volatility results in an annualized portfolio volatility of approximately 10% for each of the six strategies. We take the sign of the signal for

simplicity, though there are many other possible transformations of the signal that could be used in practice.

Empirical Results

The performance of each of these six strategies is reported in Table 1. We see that TSMOM and MACROSS strategies perform similarly across horizons, both delivering impressive risk-adjusted returns with Sharpe ratios above 1 before transaction costs.

Table 1. Performance of Simple TSMOM and MACROSS Strategies. *This table shows the performance statistics of the six signal portfolios. Excess returns and volatility are annualized, and the Sharpe ratio is the ratio of the two.*

Signal Name	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

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

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

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

Table 2. Regressions of TSMOM on MACROSS and Vice Versa. *Panel A shows the regression results from regressing each MACROSS factor portfolio's daily returns on the three TSMOM factor portfolio returns. T-statistics are shown in parenthesis. The intercept is multiplied by 260 to annualize the daily returns. Panel B shows the regression results from regressing each TSMOM factor portfolio return series on the three MACROSS factor portfolio returns.*

Panel A: Regression of MACROSS on TSMOM						
		Independent Variables				
		TSMOM(22)	TSMOM(66)	TSMOM(260)	Intercept (Annualized)	R ²
Dependent Variable	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%
Panel B: Regression of TSMOM on MACROSS						
		Independent Variables				
		MA(3,12)	MA(8,32)	MA(32,128)	Intercept (Annualized)	R ²
Dependent Variable	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%

For each of the six regressions, the R squared is above 80%. Such a high R-squared shows that these trend signals are closely related to each other and yield strategies that are quite correlated.

In addition, for each of the three MACROSS strategies, we see no significant alpha over the TSMOM signals. In other words, for each of the three MACROSS signals considered here, we do not see any significant performance benefit over a combination of TSMOM signals. This is intuitive given our theoretical results above. Somewhat surprisingly, in the case of *MACROSS(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 do see positive significance in the alphas of some of the TSMOM signals when they are 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, while TSMOM signals, because of their shape, are more easily able to fit an arbitrary MACROSS signal. This can be seen to some extent in the R squares, 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 able to well approximate the *MACROSS(8,32)* signal by combining the three TSMOM signals with relevant weights. In contrast, Panel B shows

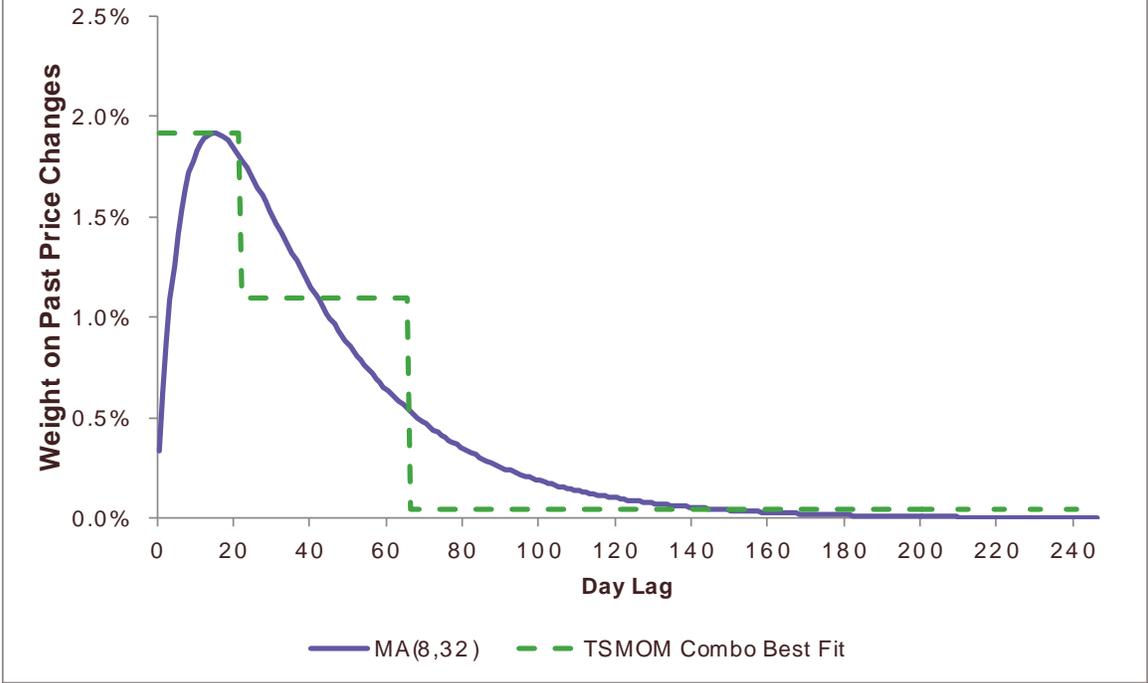
that the MACROSS signals are not able to approximate the TSMOM weight as effectively.⁹

Further, the alphas that we discover in Table 2 may reflect the relative performance of the trend horizons that are over- or under-weighted by the best fit as seen in Figure 14. For instance, as seen in the figure, the MACROSS(8,32) signal gives more weight past returns 60-120 days ago, while the best-fit TSMOM portfolio gives more weight to returns 40-60 days ago, and the latter may predict returns more strongly, leading to a negative alpha. This doesn't necessarily mean that TSMOM is a better way to invest in trend following, as these issues can perhaps be addressed by changing the parameters of the MACROSS signals and including a wider array of MACROSS signals.

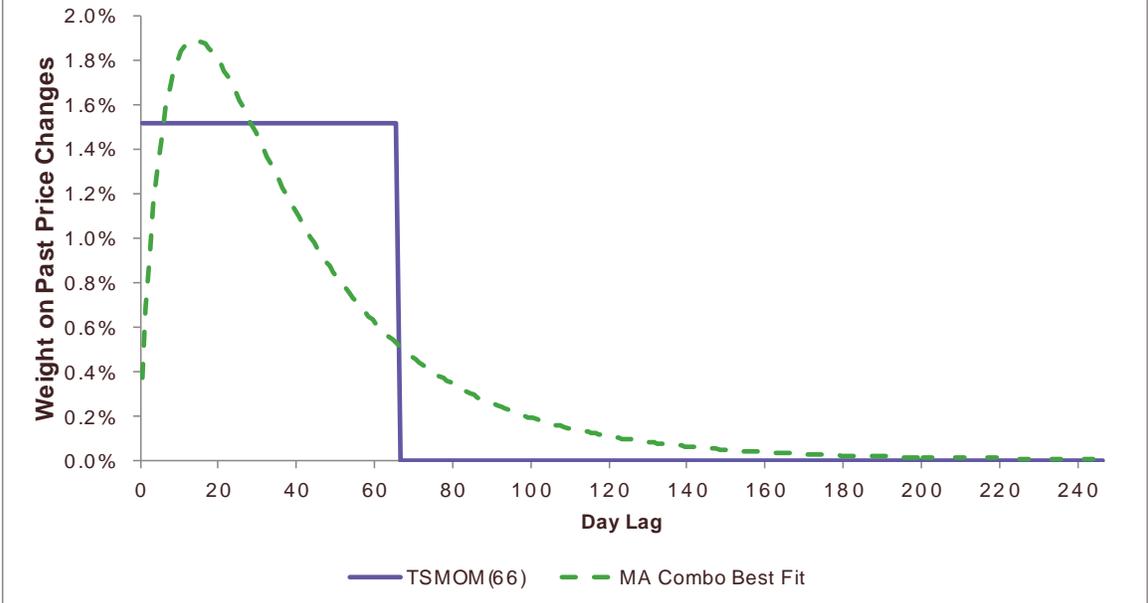
Figure 14. Approximating Moving-Average weights with Time Series Momentum and Vice Versa. *Panel A shows the weight of the moving average MACROSS(8,32) signal on past returns, and the combined weights from the three time series momentum (TSMOM) signals, weighted in proportion to the betas in the regression 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 moving-average crossover (MACROSS) signals, weighted in proportion to the betas in the regression of TSMOM(66) on MACROSS in Table 2.*

⁹ Each plot 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. This is meant to be stylized – it does not perfectly represent the regression since the regression analysis is performed on strategy returns, not the underlying signals. The relation between signal construction and return correlations is far from perfect due to the non-linear portfolio construction, but the stylized results are informative nonetheless.

Panel A: Using TSMOM to Approximate MA



Panel B: Using MA to Approximate TSMOM



5. Conclusion

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

Our results thus further 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, it becomes clear that the filtering methodology may matter less than the horizons chosen, portfolio construction, risk management, and other signals which may be useful in identifying the quality of a trend. Our results suggest that investors and managers focus on the robustness and quality of implementation – including optimally managing transaction costs,¹⁰ dynamic trading, diversification, position sizing, portfolio construction, and risk management – rather than looking exclusively at which specific filter to start from.

¹⁰ See Garleanu and Pedersen (2013)

6. References

- Asness, Cliff, Tobias Moskowitz, and Lasse H. Pedersen (2013), “Value and Momentum Everywhere,” *The Journal of Finance* 68(3), 929-985.
- Brock, W., J. Lakonishok, and B. LeBaron (1992), “Simple Technical Trading Rules and the Stochastic Properties of Stock Returns,” *The Journal of Finance*, 47(5), 1731-1764.
- Erb, C. and Harvey, C. (2006), “The Tactical and Strategic Value of Commodity Futures,” *Financial Analysts Journal*, vol. 62, no. 2, 69-97.
- Fama, E. (1965), “The Behavior of Stock Market Prices,” *The Journal of Business*, vol. 38, no.1, 34-105
- Garleanu, N. and L.H. Pedersen (2013), “Dynamic Trading with Predictable Returns and Transaction Costs,” *The Journal of Finance* 68(6), 2309-2340.
- Gencay, R., F. Selcuk, and B. 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, issue 3, 245-275
- Hodrick, Robert J., and Prescott, Edward C. (1997), “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit, and Banking*, 29, 1–16.
- Hurst, Brian, Yao Hua Ooi, and Lasse H. Pedersen (2013), “Demystifying Managed Futures,” *Journal of Investment Management* 11(3), 42-58.
- Hurst, Brian, Yao Hua Ooi, and Lasse H. Pedersen (2012), “A Century of Evidence on Trend-Following Investing,” *AQR Capital Management*.
- Kalman, R. E. (1960), “A New Approach to Linear Filtering and Prediction Problems”, *Transactions of the ASME, Journal of Basic Engineering*, vol. 82, Series D, 35-45.
- Leser, C. E. V. (1961), “A Simple Method of Trend Construction,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 23, 91–107.
- Lo, A. W., H. Mamaysky, and J. Wang (2000), “Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation,” *The Journal of Finance*, 55(4), 1705-1770.
- Moskowitz, Tobias, Yao Hua Ooi, and Lasse H. Pedersen (2012), “Time Series Momentum,” *Journal of Financial Economics*, 104(2), 228-250.

Okunev, J. and D. White (2003), "Do Momentum-Based Strategies Still Work in Foreign Currency Markets?," *Journal of Financial and Quantitative Analysis*, 38(2), 425-447.

Park, Cheol-Ho, and S. H. Irwin (2007), "What Do We Know About the Profitability of Technical Analysis?" *Journal of Economic Surveys*, 21(4), 786-826.

Silber, William L. (1994), "Technical trading: when it works and when it doesn't," *The Journal of Derivatives* 1, 39-44.

Sullivan, R., A. Timmermann, and H. White (1999), "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap," *The Journal of Finance*, 54(5), 1647-1691.

Zakamulin, V. (2015), "Market Timing with Moving Averages: Anatomy and Performance of Trading Rules," working paper, University of Agder.

Appendix A: The HP Filter as TSMOM or MACROSS

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

$$\min_g \sum_{t=1}^T (P_t - g_t)^2 + \lambda \sum_{t=3}^T (g_t - g_{t-1} - (g_{t-1} - g_{t-2}))^2$$

The objective function can be written in vector form as:

$$\min_g (P - g)'(P - g) + \lambda \cdot g' K' K g$$

where the matrix K is of dimension $(T-2)$ -by- T , defined as:

$$K = \begin{pmatrix} 1 & -2 & 1 & & & 0 \\ & 1 & -2 & 1 & & \\ & & \ddots & \ddots & \ddots & \\ 0 & & & 1 & -2 & 1 \end{pmatrix}$$

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

$$0 = -(P - g) + \lambda \cdot K' K g$$

Hence, the solution for the growth component is:

$$g = (I + \lambda K' K)^{-1} P$$

This shows that the growth component is a linear combination of past prices. Indeed, the last row of the matrix $(I + \lambda K' K)^{-1}$ 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} . Lastly, the difference between these is the current trend, $trend_T = g_T - g_{T-1}$, which therefore is an 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:

$$P_t = \alpha + \beta t + \varepsilon_t$$

over some time window that is N periods long. 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:

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

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:

$$\bar{P}_{t,N} = \frac{1}{N} \sum_{m=1}^N P_{t-m+1}$$

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

$$\hat{\beta}_t = \sum_{s=1}^N P_{t-s+1} \frac{\left(\frac{N+1}{2} - s\right)}{\sum_{m=1}^N \left(\frac{N+1}{2} - m\right)^2} =: \sum_{s=1}^N w_s P_{t-s+1}$$

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

Appendix C: Data Sources

Commodity Futures

Item	Source
Aluminum	LME (London Metal Exchange)
Copper	
Nickel	
Zinc	
Brent Oil	ICE (Intercontinental Exchange)
Gas Oil	
Coffee	
Cocoa	
Cotton	
Sugar	
Corn	CBOT (Chicago Board of Trade)
Soybeans	
Soybean Oil	
Soybean Meal	
Wheat	
Lean Hogs	CME (Chicago Mercantile Exchange)
Live Cattle	
WTI Crude Oil	NYMEX (New York Mercantile Exchange)
RBOB Gasoline (spliced with Unleaded)	
Heating Oil	
Natural Gas	
Gold	COMEX (Commodities Exchange)
Silver	
Platinum	TOCOM (Tokyo Commodity Exchange)

Bonds

Item	Source
Australia 3-year Bond	Datastream is used for futures returns, and JP Morgan bond index returns are used before futures returns are 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	

Currencies

Item	Source
Australia	Spot exchange rates and forward interest rates from Citigroup are used to form return series after 1989. Prior to 1989, spot exchange rates from Datastream are combined with the Interbank Offered Rates (IBOR) from Bloomberg.
UK	
Germany (spliced with euro)	
Japan	
US	
Norway	
Sweden	
Switzerland	
Canada	As above, with 1992 as the switchover point.
New Zealand	As above, with 1996 as the switchover point.

Equity Indices

Item	Source
Australia (SPI 200)	Datastream is used for futures returns, and MSCI country index returns are used before futures returns are available.
France (CAC 40)	
Germany (DAX)	
Italy (FTSE/MIB)	
Japan (Topix)	
Netherlands (AEX)	
Spain (IBEX 35)	
UK (FTSE 100)	
US (S&P 500)	