Trend-Adaptation of Moving Averages (TAMA)

Dr. Patrick Winter
Berlin, Germany
research@patrick-winter.de

Abstract

Largely unnoticed so far, Moving Averages (MAs) exhibit a trend bias: They underestimate (overestimate) the true value of the underlying asset during upward (downward) trends. This can not only result in misleading interpretations and forecasts, it also propagates to downstream indicators and ultimately the whole trading system. In this paper, we propose a robust method – TAMA – to correct the bias, i.e., to make any MA trend-adaptive. We exemplify this for common MAs and evaluate the effects qualitatively and quantitatively. Our findings indicate that trend adaptation greatly improves their performance in typical use cases. Additionally, we outline some new chart analysis tools that could be derived from TAMA: a Harmonic MA, a forecasting procedure, and a bands indicator.

1. Introduction

Moving Averages (MAs) are among the most fundamental tools of technical analysis. Their primary job is to filter out quasi-random noise from a price development – smoothing it thereby –, which is supposed to allow a better insight into the current "true value" of the underlying asset. Analysts then can interpret this value directly, forecast the continuation based on it, or use it as input for downstream indicators (e.g., the MACD or Bollinger Bands) and trading systems.

For all these applications, it is obviously of crucial importance that the MA's value actually represents the asset's true value without distortion, i.e., unbiased. If it does not, interpretations and forecasts could be misleading, and, even worse, downstream indicators and trading systems could produce wrong signals *without anyone noticing*, as the MA input is subsumed by them.

Fortunately, most commonly used MAs are unbiased. This is at least what many traders believe, and for good reasons indeed: When we look at the formula of the Simple MA (SMA), for instance, it is identical to the formula of the arithmetic mean, which is theoretically known to be not only unbiased but even the best unbiased linear estimator under some conditions (BLUE; e.g., Blom 1976).

In the practice of technical analysis, however, cases occur daily like the one exemplarily shown in Figure 1 (for the EUR/USD exchange rate in mid-2022): The employed MAs – regardless of their concrete specifics – seem to systematically underestimate price movements during upward trends and systematically overestimate them during downward trends (red circles) – and thus not to be unbiased at all!

Figure 1: Typical bias of MAs during trends.

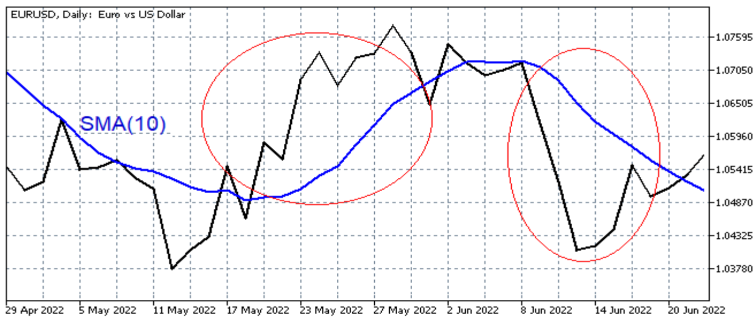

We will investigate where this apparent contradiction comes from in a moment, but let us first stay with the figure and ask another question: Since it can clearly be seen that – and even how far – the MA is from the price, can it then not simply be shifted mathematically just as it can be mentally, formally correcting its distortion? This is in fact possible, but unfortunately only if the price development is viewed retrospectively (as we do it here), i.e., if one already knows when trends have occurred. The usual intention to use an MA is to calculate it during the ongoing price development ("online"), of course, and at those points in time, the data necessary for such a correction are obviously not yet available. This may also be a main reason why MAs, despite their great importance, still have not been trend-adapted.

However, this does not mean that the approach has to be abandoned altogether. Although the trend is not fully foreseeable at the time of MA calculation, methods exist to estimate it, even robustly despite the then very short trend window, and the information obtained thereby can be integrated into the MA. In the following, we will develop a rather simple method to achieve this (Section 2) and show then, in Section 3, how it can be used to make *any* MA trend-adaptive[1], i.e., to remove its bias during trends to a large extent. Our particular method also allows for some unique extensions, which are presented in Section 4. In Section 5, we evaluate by data whether and, if so, by how much MAs can be improved this way regarding their typical use cases. Section 6 finally concludes this work with some outlooks for future research.

2. The TAMA Method

2.1 A Closer Look at the Problem

Let us now return to the aforementioned seeming contradiction: Why does a MA like the SMA, which has a formula that is theoretically unbiased, still exhibit a bias during trends?

The reason for this lies in the fact that the prices p_t are not realizations of the same – especially identically distributed – random variable, as it is assumed in the calculation of the arithmetic mean (i.e., they do not form a sample). Rather, each price p_t should be interpreted as the sole realization of its *own* random variable P_t, and it is by no means guaranteed that this variable must have the same distribution – or even the same location (or expectation) – as the previous one (P_{t-1}). On the contrary: It is precisely the *definition* of a trend that this is not the case.

Figure 2 illustrates this problem with an idealized price development smoothed by a 5-period SMA, examined at each point in time t at the end of a day, without knowledge of the future. First consider $t = 5$. The price has stayed at a value of 52 for five consecutive days here; thus, $SMA_5(5)=1/5 \cdot 5 \cdot 52=52$. This would also be the best possible forecast for $t = 6$; and indeed, that day ends again with a value of 52. Therefore, $SMA_5(6)=52$, and the previous forecast would now also be the best for $t = 7$. But things turn out differently: The price now rises to 54. This information is one that holds significance for analysts. As a counter example, consider $t = 10$: The price here has increased equidistantly for five consecutive days. Under the paradigm that all price values between $t = 6$ and $t = 10$ were drawn from the same distribution, the best estimate of their expected value – and thus the best smoothing for $t = 10$ as well as the best forecast for $t = 11$ – would indeed be that of the SMA, namely $SMA_5(10)=1/5 \cdot (52+54+56+58+60)=56$. Under the paradigm that a trend exists, i.e., $E(P_t)=E(P_{t-1})+m$ for a constant m (here $m = 2$), the corresponding values would, in contrast, lie exactly on the trendline, at 60 for $t = 10$ and at 62 for $t = 11$. The deviation of the SMA from this is therefore irrelevant – or rather *misleading* – information for analysts under the latter paradigm, since it merely results from the mathematical neglect of the trend.

Figure 2: Idealized price development with an SMA(5).

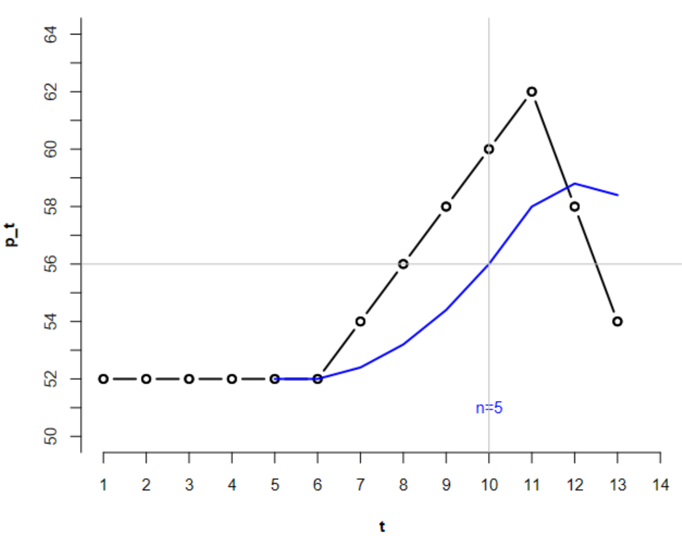

2.2 Idea and Approach

But how can one know which paradigm is the correct one in each case? Absolute certainty only comes in hindsight; while the price is evolving, estimations are necessary. Fortunately, the situation without a trend can easily be modelled as a special case of the situation with a trend, and then it can be evaluated based on the past price how likely this special case is.

The most intuitive approach that comes to mind may be to create a simple linear regression model of the price development with the time index as the only explanatory variable and to train it then with the data from the current MA-window (so for an n-period MA at time t, using the prices p_{t-n+1} to p_t). The estimated value $\hat{p}_t$ for p_t (even though this price is already known!) that results from the regression line obtained this way could then represent the value of the trend-adaptive MA at t.

In fact, this approach corresponds mathematically to the trend-adaptive version of the SMA (the TASMA), as we will see later.[2] However, it has some significant weaknesses (also when compared to other TAMA variants); in particular, it is not robust, e.g., regarding the choice of n. Figure 3 illustrates this problem and its solution based on the price development from Figure 2.

Figure 3: Various regression lines with trend on the price development from Figure 2.

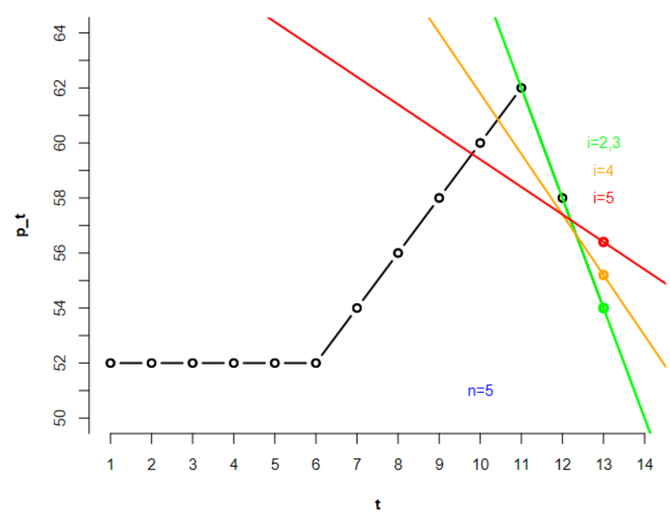

Let us first consider only the red line. This is the regression line that would result from the process described above for $n = 5$. Its calculation has included the entire downtrend present at the current time $t = 13$, i.e., the points $t = 13$, $t = 12$, and $t = 11$, but also the points $t = 10$ and $t = 9$, which – though only clearly recognizable in this illustrative example – do not belong to it. If a different period length had been chosen instead, another line would have resulted, e.g. the green one for $n = 3$, which hits the trend much better. So the choice is not irrelevant, and the trend will not always have started just by chance at the edge of the period window. Introducing a parameter $i \leq n$ that lets the user specify a different start time would also not help much, as this parameter is not known, and certainly not the same for all windows.

The idea for "robustifying" the regression is therefore to perform it not only for *one* point in time of the window, but for *all* points in time $t - i + 1$ within it (cf. lines in Figure 3), i.e., multiple times. Then, a trend that starts only after the left edge of the window, as it is the case in the illustration, is still found. Additionally, while points that do not belong to it still have an influence on some estimates, only the "more leftward" ones are affected by this anymore (in the example for $i \in [4;5]$). The (overall) estimate $\hat{p}_t$ for p_t is then obtained as an average of the individual estimates $\hat{p}_{t_i}$ (cf. coloured dots).

2.3 Formalization

We will now formalize this idea, proceeding “backwards” from the last step to the first. This last step is the calculation of an average of the price estimates $\hat{p}_t$ at time t:

$$TAGMA_n(t) = \frac{1}{n} \cdot \sum_{i=1}^{n} w_i \cdot \widehat{p}_{t_i} + r(t). \qquad (1)$$

To take into account that there are multiple ways to calculate “an average”, we have introduced here weights w_i and a residual constant $r(t)$ that may depend on t but not on the price estimates. This flexibility will play a crucial role later, and it will also become clear then why we have denoted the result as $TAGMA_n(t)$.

The i-th price estimate should be the result of its own linear regression, thus lie on the corresponding regression line. Denoting $\hat{\alpha}_i$ as its intercept and $\hat{\beta}_i$ as its slope, then $\hat{p}_{t_i}$ can be formalized as

$$\widehat{p}_{t_i} = \hat{\alpha}_{t,i} + \hat{\beta}_{t,i} \cdot i, \qquad (2)$$

since what is sought is the i-th point on the line in each case. To illustrate, it is helpful to look at Figure 3 again, for example for i=3 (green line): This is estimated by the points p_{11}, p_{12}, and p_{13}, and if we number these from left to right, the estimation time $t = 13$ corresponds exactly to the i-th (and last) point.

(2) refers to simple linear regressions, and for such, general formulas for the straightforward calculation of $\hat{\alpha}_i$ and $\hat{\beta}_i$ are known.[3] Adapted to our case, the following relationships are derived (see Appendix):

$$\hat{\alpha}_{t,i} = \bar{p}_{t,i} - \hat{\beta}_{t,i} \cdot \frac{i+1}{2} \quad \text{and} \qquad (3a)$$

$$\hat{\beta}_{t,i} = \begin{cases} \frac{12}{i^3 - i} \cdot \sum_{j=1}^{i} \left(j - \frac{i+1}{2}\right) \cdot \left(p_{t-i+j} - \bar{p}_{t,i}\right), & \text{if } i > 1 \\ 0, & \text{else} \end{cases} \qquad (3b)$$

where $\overline{p}_{t,i}$ denotes the arithmetic mean of the prices from p_{t-i+1} to p_t.

3. From MA to TAMA

3.1 General Procedure

Now that we have formalized our idea, we could call the TAGMA a new indicator, choose any weights w_i and any residual $r(t)$ for it, and include it in our trading system to replace the classical MAs.

However, our actual goal was different: Instead of replacing the classical MAs by something new, we just wanted to trend-correct them. This means that what remains is to find a relationship between any MA – call it “Generic MA” (GMA) – and its trend-adaptive (TA) version, the TAGMA (hence the name).

To do so, we start with the formula that (almost)[4] every concrete MA is a special case of:

$$GMA_n(t) = \sum_{i=1}^{n} u_i \cdot p_{t-i+1} + s(t), \qquad (4)$$

with weights u_i (which can, of course, also be 0) and a price-constant residual $s(t)$, both factors to characterize the concrete MA. This almost looks like the definition of the TAGMA in (1) already, but there is an important difference: The GMA is based on different actual prices p_{t-i+1}, while the TAGMA is based on different estimates $\hat{p}_{t_i}$ of the *same* price p_t.[5]

To make both approaches comparable, an essential insight is required: If the TAGMA is the TA-version of the GMA, it must match the latter *for all arbitrary price values and points in time* when trend adaptation is “turned off”. We will later introduce that capability as an extension to the TAMA method in more detail; for now, it should suffice to say that TA enters (2) through the dynamic slope coefficient $\hat{\beta}_{t,i}$ and thus can be deactivated by setting $\hat{\beta}_{t,i} = 0$ instead of its value from (3b). Using (2) and (3a) with this, we can demand the following equality independent of specific price values:

$$TAGMA_{n,TA=0}(t) = \frac{1}{n} \cdot \sum_{i=1}^{n} w_i \cdot \bar{p}_{t,i} + r(t) \overset{!}{\triangleq} \sum_{i=1}^{n} u_i \cdot p_{t-i+1} + s(t) = GMA_n(t). \qquad (5)$$

One can show (see Appendix) that this is satisfied if and only if the concretization parameters of the TAGMA are chosen as follows:

$$w_i = \begin{cases} n \cdot i \cdot (u_i - u_{i+1}), & \text{if } i < n \\ n^2 \cdot u_n, & \text{else (if } i = n) \end{cases} \quad \text{and} \qquad (6a)$$

$$r(t) = s(t). \qquad (6b)$$

Thus, a fairly simple relationship has been found that allows the calculation of the weights for the corresponding TAMA with deactivated trend adaptation from the knowledge of the MA’s formula. With the "reactivation" of TA by re-setting $\hat{\beta}_{t,i}$ to its value from (3b), its construction is complete.

3.2 TA-versions of Common MAs

We can now apply our newly found relationship to any MA, and we will do so in the following for the three arguably most common MAs. Our approaches will be illustrative for other MAs as well.

Simple Moving Average (SMA)

The SMA is characterized by $u_i = \frac{1}{n}$ and *s(t)* = 0. Thus, we have from (6a):

$$w_i^{SMA} = \begin{cases} n \cdot i \cdot \left(\frac{1}{n} - \frac{1}{n}\right), & \text{if } i < n \\ n^2 \cdot \frac{1}{n}, & \text{else} \end{cases} = \begin{cases} 0, & \text{if } i < n \\ n, & \text{else (if } i = n) \end{cases} \tag{7}$$

This means that all regressions except the "largest one" (the one taking into account all *n* points up to the beginning of the period window) are excluded from the calculation! Thus, the resulting TASMA is based only on this single regression, making it generally fragile and sensitive to the choice of *n*. This issue has been discussed in more detail in Section 2.1. As a consequence of it, the TASMA is usually more of theoretical interest and not recommended to be used in practice.

Weighted Moving Average (WMA)

The Weighted MA (WMA) is based on the idea of assigning the weight $u_i = \frac{n-i+1}{N}$ to the price p_{t-i+1}, which is *i* - 1 days away from the current price p_t, so that older influences linearly decrease in importance compared to newer ones. Here, $N = n + (n-1) + \ldots + 1 = \frac{n^2+n}{2}$ is a normalization factor. According to (6a), the weights of the TAWMA are determined as follows:

$$w_i^{WMA} = \begin{cases} n \cdot i \cdot \left(\frac{n-i+1}{N} - \frac{n-(i+1)+1}{N}\right), & \text{if } i < n \\ n^2 \cdot \frac{1}{N}, & \text{else (if } i = n) \end{cases} = \frac{2 \cdot i}{n+1}. \tag{8}$$

Similar to the SMA, the WMA represents an interesting special case in a sense, as the otherwise necessary distinction in (6a) is not needed due to the arithmetic progression of its weights.

Exponential Moving Average (EMA)

The Exponential MA (EMA) extends the idea of the WMA; it also weights prices less that are further back in time, but by an exponential progression:

$$\begin{aligned} EMA_{n,\alpha}(t) &= \alpha \cdot p_t + (1-\alpha) \cdot EMA_n(t-1) = \alpha \cdot \sum_{i=1}^{\infty} (1-\alpha)^{i-1} \cdot p_{t-i+1} \\ &= \sum_{i=1}^{n} \alpha \cdot (1-\alpha)^{i-1} \cdot p_{t-i+1} + \underbrace{\sum_{i=n+1}^{\infty} \alpha \cdot (1-\alpha)^{i-1} \cdot p_{t-i+1}}_{r(t)} \end{aligned} \tag{9}$$

with a smoothing factor 0 < α < 1, which is typically but not necessarily chosen as $\alpha = \frac{2}{n+1}$. Consequently, the weights become progressively smaller, but they never reach the value of 0 and thus never disappear in theory.[6] At first glance, this might seem problematic, as only a weighted sum of precisely *n* prices within the current period window has been allowed so far.

However, if one decomposes the calculation formula as shown in (9), the issue becomes easily resolvable: All other prices for $i \geq n$, including their coefficients, do not enter the weights w_i but are instead represented by the residual term *r*(*t*); this term must be constant with respect to the prices in the window, but not necessarily with regard to the other ones. Thus, the weights w_i can be determined as usual from (6a):

$$\begin{aligned} w_i^{EMA} &= \begin{cases} n \cdot i \cdot (\alpha \cdot (1-\alpha)^{i-1} - \alpha \cdot (1-\alpha)^i), & \text{if } i < n \\ n^2 \cdot \alpha \cdot (1-\alpha)^{i-1}, & \text{else (if } i = n) \end{cases} \\ &= n \cdot i \cdot \alpha \cdot (1-\alpha)^{i-1} \cdot \left(\begin{cases} \alpha, & \text{if } i < n \\ 1, & \text{else (if } i = n) \end{cases}\right). \end{aligned} \tag{10}$$

3.3 The "Simplest" TAMA: The TAHMA

One might have expected that the simplest concretization of the TAGMA in (1), which can be constructed by choosing equal weights of w_i=1 for all *i* and a residual of *r*(*t*)=0, making the calculated average the arithmetic mean, were the TA-version of the SMA, since the latter also uses equal weights and no residual. However, by (7) we have found that this is not the case. Instead, it can be shown (cf. derivation of (6a) and (6b)) that the MA that corresponds to the "simplest" TAMA — the TAHMA — looks like follows:

$$\begin{aligned} HMA_n(t) &= \frac{1}{n} \cdot \left(\frac{1}{1} + \cdots + \frac{1}{n}\right) \cdot p_t + \frac{1}{n} \cdot \left(\frac{1}{2} + \cdots + \frac{1}{n}\right) \cdot p_{t-1} + \cdots + \frac{1}{n} \cdot \left(\frac{1}{n}\right) \cdot p_{t-n+1} \\ &= \frac{1}{n} \cdot \sum_{i=1}^{n} (H_n - H_{i-1}) \cdot p_{t-i+1}, \end{aligned} \tag{11}$$

where $H_i = \sum_{j=1}^{i} \frac{1}{j}$

denotes the *i*-th so-called harmonic number, that is, the i-th partial sum of the harmonic series (H_0 = 0 is defined for notational reasons). We therefore call this MA the Harmonic MA (HMA).[7]

This relationship may appear surprising at first glance, but upon closer examination, it turns out to be very natural: The TAMA is by construction an average of regression estimates, those are also averages themselves, and the basis of the regressions are exactly *i* = 1, 2, ..., *n* points. Now a *consecutive average of averages* corresponds precisely to the essence of the harmonic series.

This inherent connection to the method, its simplicity, and the weaknesses of the TASMA make the TAHMA an especially good candidate for practical use. In fact, we will see later that it also performs well; however, the difference between various TAMAs is generally not very pronounced.

4. Extensions

Now that we have found how the common TAMAs look like, we could proceed directly to their practical application. However, the particular method we have derived for trend adaptation allows for some special extensions that will proof useful there, so that we will introduce them beforehand.

4.1 Trend-Adaptation Switch TA

The first extension has already been mentioned and used to obtain (5): a "switch" that allows deactivating the trend adaptation if desired. This enables a direct evaluation of its effect.

The switch is formally a Boolean parameter, denoted here as *TA*, which is "true" (i.e., 1) by default and must be set to "false" (i.e., 0) to deactivate it. TA can be incorporated by modifying (3b) as follows:

$$\hat{\beta}_{t,i} = \begin{cases} 0, & \text{if } TA = 0 \\ \hat{\beta}_{t,i} \text{ from (3b)}, & \text{else (if } TA = 1) \end{cases} \qquad (12)$$

The effect is immediately clear: Setting the slope of the line β $\hat{\beta}_{t,i}$ to 0 means nothing other than to *rule out* a trend – precisely as MAs implicitly do. This becomes evident in (2) in conjunction with (3a); since the terms associated with $\hat{\beta}_{t,i}$ = 0 vanish, it simply holds that

$$\widehat{p_{t}}_i = \hat{\alpha}_{t,i} = \overline{p}_{t,i} = SMA_i(t).$$

4.2 Forecasting Period P

Many forecasting models for the future price development are based on MAs (see Nau 2014 for an introduction). Of course, TAMAs can also be used as inputs for such a model, and since they take into account not just averages but also their changes during trends, it would be unsurprising to obtain better predictions thereby alone. However, due to the particular calculation of TAMAs – after all, they are based on regression models – it is additionally possible to use them for forecasting *independently*.

A further extension is introduced for this: In (2), the price at time *t*, p_t, was estimated based on the corresponding point on the i-th regression line, which was found at index i. This was done solely for the purpose of smoothing (averaging), as p_t is already known at (or at the end of) *t*. The price at a later point in time $t + P$, on the other hand, is still unknown by then; of course, however, nothing prevents us from looking for it on the regression lines, too, just at index $i + P$ instead of *i*.

Formally, this requires introducing another parameter, here designated as *P* (for "prognosis" or "period"), which defaults to the value 0. (2) is then to be modified as follows:

$$\widehat{p_{t}}_i = \hat{\alpha}_{t-P,i} + \hat{\beta}_{t-P,i} \cdot (i + P). \qquad (13)$$

(13) means that the i-th estimate $\widehat{p_{t}}_i$ for p_t is calculated based on the data available at time $t - P$, and then extrapolated *P* units into the future from there. Accordingly, a TAMA with this extension provides values not only up to *t* but up to $t + P$, as data are available exactly up to $t = (t + P) - P$.

P has been modelled here as any integer for maximum flexibility.[8] In practice, however, it is generally not advisable to attempt extrapolating the price by more than one unit of time; additionally, the estimation error naturally increases with the distance from the data available at *t*. *P* will therefore usually either have its default value of 0, for which (13) and (2) coincide, or be 1, so it effectively acts as a "forecasting switch".

4.3 Generic Aggregation Function AF

In (1), the *n* estimates $\widehat{p_{t}}_i$ for the price p_t at time *t* were aggregated to an overall estimate through averaging. Though this is the most intuitive method, it is not the only possible one. The right side of (1) can thus be replaced by a more generic aggregation function AF:

$$TAGMA_n(t) = AF(\widehat{p_{t}}_1, \dots, \widehat{p_{t}}_n). \qquad (14)$$

By default, *AF* should still correspond to a (linear) weighted average – in which case (14) is (1) –, but there are applications for which other aggregation functions are even more useful.

A prominent example for such an application are band indicators. Again, TAMAs could simply be used as a replacement for their conventional counterparts here and would likely perform better than these, since the generated bands are then laid around a mean value that more closely reflects the truth, and the interpretation does not need to depend anymore on whether there is a trend or not.

The TAMA method also entails a specific band indicator, however, when appropriate aggregation functions are used. For instance, a price development could be described and analysed using the 20%-quantile band, the 80%-quantile band, and possibly the median (i.e., the 50%-quantile band) of the estimated values $\widehat{p_{t}}_i$. This can even be used to build a trading system out of TAMA. We will demonstrate this idea later on.

5. Practical Evaluation

Having specified the TAMA method and its extensions completely, it is time to return to Figure 1, which has demonstrated the trend-problem of conventional MAs. Continuing the same example, we will now examine qualitatively whether the problem can be solved by TAMA and how this affects the typical use cases of MAs mentioned in the introduction: smoothing, forecasting, and downstream usage. We will also perform a small quantitative study to investigate their performance in the latter two areas.

5.1 Smoothing

Figure 4 reiterates Figure 1, but this time with a TAHMA added to it that has the same period duration (of 10) as the SMA we have looked at earlier, once with and once without trend adaptation.

Figure 4: Figure 1 extended by a TAHMA(10) with and without trend adaptation (TA).

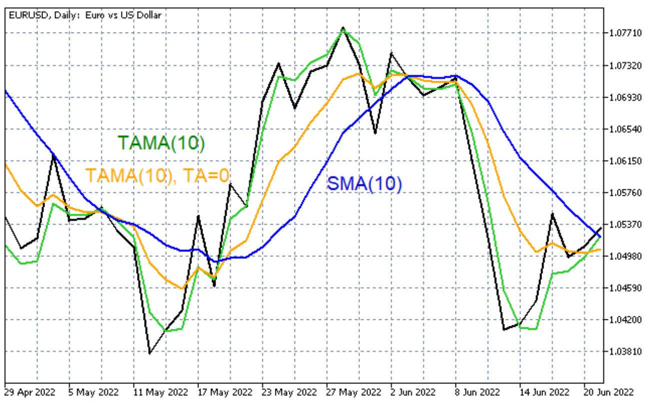

Figure 5: Exchange rate forecast using a TAHMA(10) with *P* = 1, with deviations from the actual price.

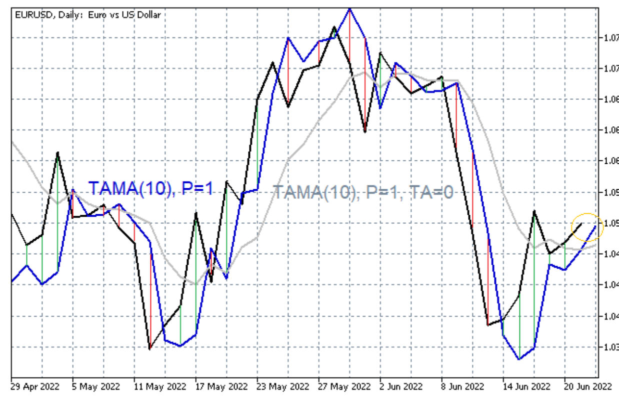

Let us first focus on the version without trend adaptation (orange line), which is equivalent to a HMA(10) from (11). While it is at the same "level" as the SMA, it already provides a smoother picture of the exchange rate: Not only is it practically closer to it at every point in time, it is also significantly less delayed. This results from the HMA, like an EMA, putting more emphasis on more recent prices. However, it can clearly be seen that the trend-problem (cf. red circles in Figure 1) is not solved yet.

That is only achieved when trend-adaptation is switched on, resulting in the actual TAHMA(10) (green line): Both during the upward and downward trends, it not only stays very close to the exchange rate, as it should, but also is no longer systematically below or above it — it is indeed unbiased even during trends. Of course, this also means that with the same period, it smooths slightly less than conventional MAs; it is therefore a plausible recommendation to use somewhat longer period durations for TAMAs in general.

We would like to demonstrate the superiority of TAMAs for smoothing also quantitatively. However, that were challenging given that entire studies can be dedicated to the question alone how "smoothness" should be captured mathematically and weighed against the accuracy of price representation (e.g., Raudys and Pabarskaite 2016). However, it obviously is always preferable when the latter, given a certain level of the former, does not depend on whether the price "by chance" is in a trend or not, and since TAMAs do not exhibit this weakness anymore, they should dominate MAs in any quantitative investigation on smoothing.

5.2 Forecasting

Figure 5 again shows a TAHMA(10) for the USD exchange rate example (blue line, cf. green line in Figure 4), but this time with $P = 1$ instead of $P = 0$ — i.e., the value at t now only relies on data up to t-1. Consequently, it can now be plotted up to $t + 1$, forecasting the exchange rate for the next day (circle).

How accurate is this forecast over the whole window (blue line)? Examining it on its own at first, it can visually be deemed appropriate: It is usually not too far from the actual value (which is not known until the following day, of course); on some days, it is nearly exact (e.g., 09/05), on others (especially those with significant changes), the deviation is larger (e.g., 12/05). Over- and underestimations (red and green lines, resp.) generally balance each other out, although the former naturally occur more frequently during downward trends, while the latter occur more frequently during upward trends due to the MA basis.[9]

However, we are more interested here in the added value of trend adaptation. To evaluate this, let us first consider what it would mean to forecast a price using conventional MAs, i.e., without accounting for trends: The best possible prediction would then be the trivial one, i.e., assuming the price at $t + 1$ to have the value of the MA at t — in other words, simply shifting the MA one time unit to the right! This is illustrated by the grey line, for which trend adaptation was deactivated using the TA=0 switch: It corresponds exactly to the orange line from Figure 4, just shifted. The value added by the TAMA method can thus be directly determined by comparing the blue to the grey line: The former generally lies much closer to the actual value, especially during trends, corresponding to a significantly better forecast. This is not all that surprising, as dedicated approaches often also rely on regressions.

To quantify the extent of this improvement in general, we conducted a similar analysis as shown here for the 40 DAX members over a 1-year period (01/07/2021-30/06/2022) as example. Forecasting accuracy was measured using the Root Mean Square Error (RMSE), which (in principle) looks at the average size of an estimation residual (cf. red and green lines in Figure 5), independent of its direction. The predictions were based on an EMA with a period of 20 days vs. a TAEMA with the same period. Table 1 shows the results (columns 2-4). These are basically as expected: The estimation error was lower *for each* DAX member when trend adaptation was activated, as this makes additional information — specifically, information about trends — usable. However, the magnitude of this effect may be surprising: The RMSE for the TAEMA was on average lower by nearly 1/3 (31.97%) than for the EMA!

Despite these appealing characteristics, TAMAs likely should not be used as a sole forecasting model. While they can accurately capture averages and their changes, they lack a component that represents the variance around these values. This gap can be bridged, however; either through suitable combinations with "spread indicators" like the Stochastic Oscillator, or by forming and using bands, i.e., interval rather than point estimators. The latter approach is a special case of what will be discussed next.

5.3 Downstream Usage

Finally, we will evaluate whether trend adaptation improves not only MAs themselves but also their downstream usage, i.e., other indicators or whole trading systems that are based on them. As a well-known example for such a downstream indicator, we consider Bollinger Bands; these can also be regarded an example for the interval estimators mentioned above. They are constructed by placing a symmetric channel of size $2 \cdot 2 \cdot \sigma$ around a 20-period SMA in the standard setting, where σ denotes the standard deviation of prices in the current window (for details, see Bollinger 2001). Many investors believe that the price at the next time unit will then lie within this channel with a probability of about 95%. This is, of course, not true in general because none of the assumptions behind this rationale is actually given (incidentally, this is related to the problem discussed in Section 2.1). Nevertheless, there *is* a probability that the expected will happen; it just does not necessarily amount to 95%. Its actual value can be well estimated through backtesting; in our study, it averages around 82% (column 5 of Table 1). But this is only the case for the conventional SMA; if its trend-adaptive version, the TASMA, is used instead, the probability rises to over 90% (column 6), and with a TAEMA, it would still be higher — all for the *same channel size*. This again shows how trend adaptation can improve existing indicators.

Table 1: Comparison of forecasts and bands with and without trend adaptation.

Asset 1/7/21-30/6/22	rmse_ta0	rmse_ta1	-%	boll_ta0	boll_ta1	boll_sig	mima	mima_sig
adidas	130.12	88.78	31.77%	81.36%	93.22%	2	76.17%	0.77
Airbus	58.39	44.30	24.14%	84.75%	87.71%	2	73.19%	0.92
Allianz	90.35	61.39	32.05%	81.78%	91.95%	2	75.74%	0.78
BASF	32.71	20.30	37.94%	81.36%	89.41%	2	77.45%	0.82
Bayer	27.75	16.34	41.11%	80.08%	91.95%	2	71.06%	0.72
Beiersdorf	34.29	26.03	24.10%	86.86%	86.44%	2	74.89%	0.93
BMW	45.34	31.16	31.27%	84.75%	91.95%	2	74.89%	0.83
Brenntag	30.42	22.14	27.23%	84.32%	89.41%	2	73.62%	0.90
Continental	65.30	42.88	34.33%	78.81%	89.83%	2	78.30%	0.79
Covestro	29.49	19.75	33.02%	78.81%	89.83%	2	75.32%	0.83
Daimler Truck	16.69	10.15	39.15%	80.00%	88.33%	2	79.83%	0.77
Deutsche Bank	9.66	6.12	36.64%	84.32%	87.71%	2	75.74%	0.80
Deutsche Börse	55.99	37.12	33.71%	80.08%	92.37%	2	75.32%	0.79
Deutsche Post	25.02	16.61	33.63%	80.93%	88.98%	2	74.89%	0.79
Deutsche Telekom	6.47	4.47	31.00%	82.63%	91.10%	2	75.32%	0.79
E.ON	4.40	2.95	32.80%	80.08%	92.80%	2	74.89%	0.81
Fresenius	19.15	11.28	41.09%	79.24%	94.07%	2	75.74%	0.69
Fresenius Medical Care	25.81	17.44	32.45%	82.20%	91.10%	2	74.89%	0.80
Hannover Rück	65.95	49.66	24.70%	80.51%	90.68%	2	72.34%	0.88
	31.61	20.92	33.83%	83.90%	91.95%	2	79.15%	0.77
HelloFresh	64.32	41.13	36.05%	79.24%	90.68%	2	76.17%	0.77
Henkel vz.	31.30	21.77	30.43%	79.66%	87.71%	2	75.74%	0.80
Infineon	21.25	13.82	34.99%	80.93%	90.25%	2	75.74%	0.77
Linde	109.46	74.28	32.14%	81.36%	90.68%	2	72.34%	0.79
Mercedes-Benz	51.03	34.84	31.72%	87.29%	87.71%	2	77.87%	0.79
Merck	102.64	63.70	37.94%	77.97%	90.68%	2	74.04%	0.73
MTU Aero Engines	104.40	76.85	26.38%	84.32%	88.98%	2	76.17%	0.83
Münchener Rück	107.14	74.11	30.83%	80.93%	94.07%	2	73.19%	0.79

Asset (1.7.21-30.6.22)	rmse_ta0	rmse_ta1	-%	boll_ta0	boll_ta1	boll_sig	mima	mima_sig
Porsche	55.81	42.22	24.36%	87.29%	88.98%	2	75.74%	0.85
PUMA	53.21	34.56	35.04%	80.51%	94.07%	2	77.02%	0.78
QIAGEN	18.29	13.62	25.56%	81.78%	89.41%	2	74.47%	0.88
RWE	15.68	12.48	20.42%	79.66%	91.10%	2	74.47%	0.90
SAP	43.96	29.07	33.87%	81.78%	92.37%	2	74.04%	0.76
Sartorius vz.	477.57	309.75	35.14%	82.63%	92.80%	2	76.17%	0.72
Siemens	72.59	51.51	29.04%	81.78%	87.71%	2	74.47%	0.85
Siemens Healthineers	29.67	21.00	29.21%	82.20%	89.83%	2	76.60%	0.86
Symrise	50.83	33.44	34.20%	80.93%	90.68%	2	75.32%	0.77
Volkswagen	100.70	74.26	26.26%	85.17%	88.14%	2	74.89%	0.89
Vonovia	21.35	14.77	30.81%	86.02%	93.64%	2	77.87%	0.83
Zalando	52.13	32.14	38.35%	80.08%	92.37%	2	74.89%	0.75
Average	**59.71**	**40.48**	**31.97%**	**81.96%**	**90.57%**	**2**	**75.40%**	**0.81**

In addition, as already has been mentioned in Section 4.3, the TAMA method also entails its own band indicator, and that will be the object of our last analysis. For this purpose, let us once again look at the example of the EUR/USD exchange rate with a TAHMA(10), but this time with the aggregations functions $AF = \min_i \hat{p}_{t_i}$ and $AF = \max_i \hat{p}_{t_i}$, respectively (Figure 6). Among all quantile-based functions, these are of particular interest because they create the extreme bands, which envelop the price and do never intersect. This may not sound too exciting at first, but it inspires two new evaluation possibilities:

Figure 6: min-max-Bands (AF) of TAHMA(10) around the exchange rate.

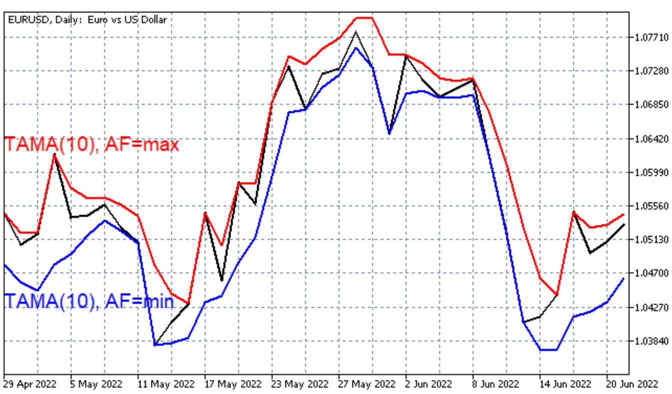

The first is the observation that the price is almost always closer to one band than to the other, often — but not always (e.g., 09/05, 30/05) — hitting the former. It can therefore be assumed that its distance to the bands represents information that can be used to determine the magnitude of the current trend.

The second idea may be even more interesting: If the minimum and maximum of the estimates for the price completely enclose it when p_t is known, what is the meaning of crossing this envelope when p_t is instead forecasted? After all, such a crossing means that the price then falls below or exceeds the smallest or largest possible extension of its current course; this should be a strong signal that something will change!

Figure 7 shows, to verify this, the content of Figure 6 but on a forecast basis, i.e., with $P = 1$. Indeed, there are several places where the envelope formed by the min-max-bands is crossed. Most of these crossings are sequential in the sense that the same band is pierced repeatedly one after another. This could be interpreted as a changing or (more likely) not yet stabilized (estimated) slope. However, more significant is the view of the respective *first* crossing points of a band after a breach point of the other (shown in green and orange in the figure). If these were interpreted as buy or sell signals in the example of the figure, one would have traded well, except for the very first case where an apparent uptrend did not materialize: The large uptrend in mid-May is recognized very promptly and exited exactly when a plateau sets in; another buy signal then occurs only in mid-June after the large downtrend is over.

Figure 7: Figure 6 on forecasting basis ($P = 1$), with respective first crossing points.

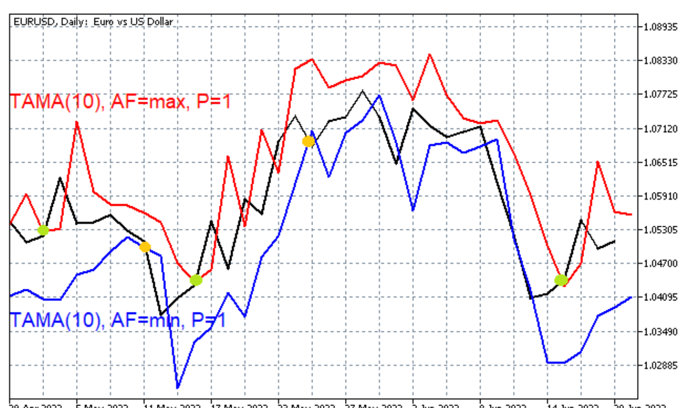

For a quantitative evaluation, the TAMA Bands were studied the same way as the Bollinger Bands. Their hit rate seems lower at first glance (column 8 of Table 1) – on average, 75% prices lie within their channel –, but this channel also is significantly narrower: It requires only about 0.8 standard deviations (column 9) instead of 2 (column 7). Therefore, the TAMA bands could be a good alternative for some applications.

6. Conclusion

In this paper, we have developed a regression-based method to correct the bias MAs exhibit during trends. The method includes a "switch" extension by which trend adaptation can always simply be turned on or off, enabling traders to evaluate its effects easily.

In our example study of 40 DAX members, we have found that turning it on greatly improves an MA's performance in smoothing, forecasting, and downstream usage: Forecasting errors were reduced by roughly 1/3 and Bollinger Bands became more consistent by roughly 10%. Future research can continue to evaluate such effects, e.g. for other MA-based indicators such as the MACD.

The particular variants of TAMAs generally do not differ much regarding performance from our experience, at least no more than the underlying MAs. An exception to this, however, is the TASMA, since its calculation effectively uses only a single regression, which is against the basic idea of TAMA of averaging multiple estimations for robustness. In contrast, the TAHMA we have introduced aligns most naturally with this idea, and thus it is the variant we recommend in the trend adaptation context.

Finally, we have outlined some ideas on what else could be derived from the TAMA method, namely a forecasting procedure and a bands indicator. Both tools also showed promising results in our practical evaluation and are therefore worthy of further exploration in future research.

Implementation for MetaTrader 5

The TAMA method has been implemented for the popular trading program MetaTrader 5 (for the MAs SMA, WMA, EMA, and HMA and the aggregation functions arithmetic mean, min, max, and quantile). It can be downloaded from the MQL5 marketplace (https://www.mql5.com/en/market).

Appendix

The appendix contains the mathematical derivations of (3a) and (3b) for one thing and of (6a) and (6b) for another thing. It has been omitted here for brevity but is available from the author upon request.

Footnotes

[1]"Adaptive" refers to the MAs, as will become clear later, dynamically adapting to trends. This has nothing to do with other "Adaptive MA" indicators, such as the one by Kaufman (1995, pp. 129-153).

[2]The TASMA described here (without the forthcoming extensions) is essentially identical to the widely discussed Linear Regression Slope indicator (e.g., Chande and Kroll 1994, p. 20ff.).

[3] It should be noted that we have no direct interest in these two coefficients but only, indirectly, in the resulting estimated value $\widehat{p_t}_i$. In contrast to the former, the latter is reliable even when the regression is based on very few data points; in the extreme case (i=1) even just one – undoubtedly, p_t itself provides an adequate basis for $\widehat{p_t}_i$ (since a straight line cannot be formed from a single point, $\hat{\beta}_{t,i}$ must be manually set to 0 in (3b) for this case).

[4]The exception are "exotic" MAs that are non-linear *in the prices*, such as the Geometric Moving Average $GeoMA_n(t) = \sqrt[n]{p_t \cdot p_{t-1} \cdots p_{t-n+1}}$. However, many members of this group can be linearized.

[5]There also are two notation-only differences: First, we do not factor out $1/n$ as a normalization constant here because that would hardly make sense in general. Second, MAs are calculated from right to left while the regression lines in (1) are formed from left to right. However, this all is mathematically equivalent.

[6]In practice, there are two different approaches that trading programs can use to handle this. The approach described in the following involves calculating the EMA up to the very first available data point. Alternatively, the calculation can be stopped after exactly n steps. In that case, the coefficient of the n-th price $p_{(t-n+1)}$ must first be corrected by the factor $\frac{1}{\alpha}$. The procedure can then continue analogously to the WMA.

[7]One can sometimes find MAs under this designation that are harmonic in the prices rather than in the coefficients (and thus "exotic") (e.g., gorx1 2020); the HMA presented here should not be confused with those.

[8]For some applications, even a negative P (i.e., a "retrospective") might make sense.

[9]This observation prompts the question of whether one could also apply trend adaptation to the forecast value (instead of the actual value), effectively meaning a "second-order" kind of TAMA, similar to applying an MA to itself. This is presumably not feasible or sensible since the crucial value p_t is missing or already estimated in the forecasting scenario, but this can be a subject for further research.

References

Blom, G. (1976): When is the Arithmetic Mean Blue? The American Statistician, Vol. 30, No. 1, pp. 40–42.

Bollinger, J. (2001): Bollinger on Bollinger Bands. McGraw-Hill Education.

Chande, T.S. and Kroll, S. (1994): The New Technical Trader. John Wiley & Sons.

gorx1 (2020): Harmonic Moving Average, https://www.tradingview.com/script/mpWRYb8F-Harmonic-Moving-Average.

Kaufman, P. (1995): Smarter Trading – Improving Performance in Changing Markets. McGraw-Hill.

Nau, R. (2014): Forecasting with Moving Averages. Fuqua School of Business, Duke University in Durham, NC, USA. https://people.duke.edu/~rnau/Notes_on_forecasting_with_moving_averages--Robert_Nau.pdf.

Raudys, A. and Pabarskaite, Z. (2016): Optimising The Smoothness And Accuracy Of Moving Average For Stock Price Data. Technological and Economic Development of Economy, Vol. 24, No. 3, pp. 984–1003.