



Moving Average: Holy Grail or Fairy Tale - Part 1

By Theodore Wong

June 16, 2009

Advisor Perspectives welcomes guest contributions. The views presented here do not necessarily represent those of Advisor Perspectives.

Buying and holding a diversified portfolio works well during good times, but falls short when supposedly uncorrelated asset classes drop in unison in bear markets. Are there alternative investment strategies that work for all seasons? The 10-month *Moving Average Crossover (MAC)* system is a solid candidate, as it sidestepped two recent bear markets in 2000 and 2008. But did it work in previous bear markets? Is 10 months the optimum length?



Let's examine historical evidence to find out if MAC is the Holy Grail or just a fairy tale.

Background

Electrical engineers routinely use moving average as a *low-pass filter* in analog and digital signal processing. It blocks transient perturbations from the input and only allows the core signals to pass, hence the term *low-pass filter*. Transient perturbation is a fancy name for short-lived popcorn noise that obscures the underlying signal.

Random spikes in an otherwise smooth signal are undesirable. We can reduce the amplitude of these noisy spikes by averaging the values of the data points neighboring on either side of the spike. *Figure 1* shows how the filtered output closely tracks the original signal but the unwanted spikes are attenuated. The degree of noise suppression depends on the number of points used in the averaging. The longer the averaging period, the smoother the output. Because we can't predict when random spikes will appear, we slide the filter block across the entire data stream from start to finish. The term *moving average* literally describes this function.

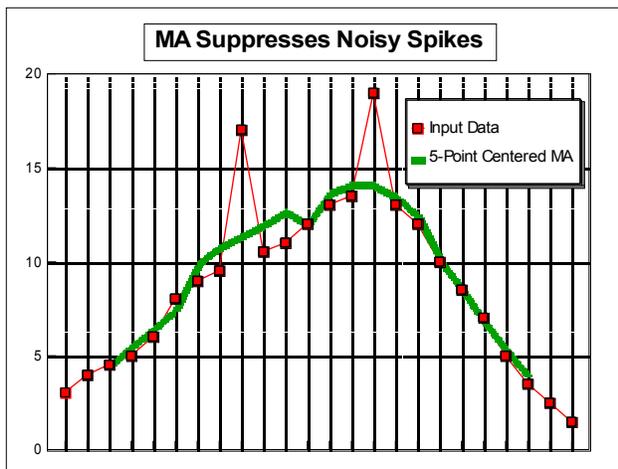


Figure 1

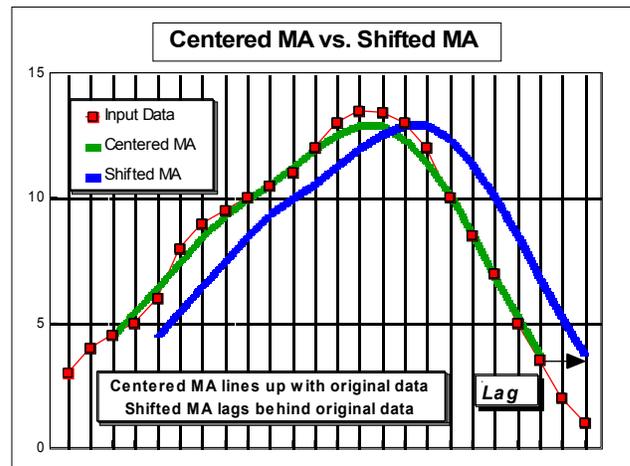


Figure 2

As an engineer, I have always been skeptical of the way stock market technicians plot MA curves. Traditionally, engineers align the midpoint of the MA curve with the center of the original data curve. This way, the MA curve is centered with respect to the original time series as shown by the red curve in *Figure 2*. Technicians, on the other hand, shift the end of the MA curve to match the most recent price point as shown by the blue curve in *Figure 2*. The lag between the original data curve and the shifted MA curve created by this peculiar plotting convention is the core of the MAC system. Without shifting the MA, there is no lag. Without the lag, there is no crossover.

The Moving Average Crossover system

MAC is the simplest and probably the oldest trading system. You buy when the price rises above its moving average, and you sell when it drops below. Although there are several forms of moving averages, I prefer the [Exponential Moving Average \(EMA\)](#) to the [Simple Moving Average \(SMA\)](#) because EMA gives slightly smaller lag.

My assumptions

To keep things simple, I made three assumptions:

1. All proceeds after sales are kept in non-interest-bearing cash.
2. No transaction fees.
3. No taxes.

The first assumption penalizes MAC in favor of buy-and-hold. Parking proceeds in Treasury Bills would obscure the central focus of my study because short-term interest rates varied widely throughout history.

The second assumption has a small positive bias toward MAC. But fees on index funds and ETFs (which I assume, for the purpose of my analysis, have been around since the Civil War) are less than 10 basis points and will not significantly affect my results.

I exclude tax effects for several reasons. First, tax rates vary with income levels. Second, top marginal tax rates changed dramatically in the past 138 years, from below ten percent before 1910 to above ninety percent in the 1950s. Third, buy-and-holders are not exempt from tax; tax payments are merely deferred. When they eventually sell their holdings, their entire cumulative gains will be taxed. Ignoring taxes is a balanced compromise, and does not give the MAC system an unfair advantage.

I will revisit the fee and tax assumptions after presenting my results.

Let the contest begin: MAC versus buy-and-hold

To compare performance between MAC and buy-and-hold, I used Compound Annual Growth Rates (CAGRs) and 138 years of monthly data for the S&P500 total return index (with dividend reinvestment) from 1871 to 2009. I examined a wide range of MA lengths, from two to twenty-three months.

The buy-and-hold benchmark returned 8.6 percent over this period, and is represented by the red bar in *Figure 2*. The green bars represent the CAGRs for different moving average lengths. CAGRs below 11 months consistently beat buy-and-hold. Above that, they reach diminishing returns. The quasi bell-shaped curve suggests that the distribution is not random.

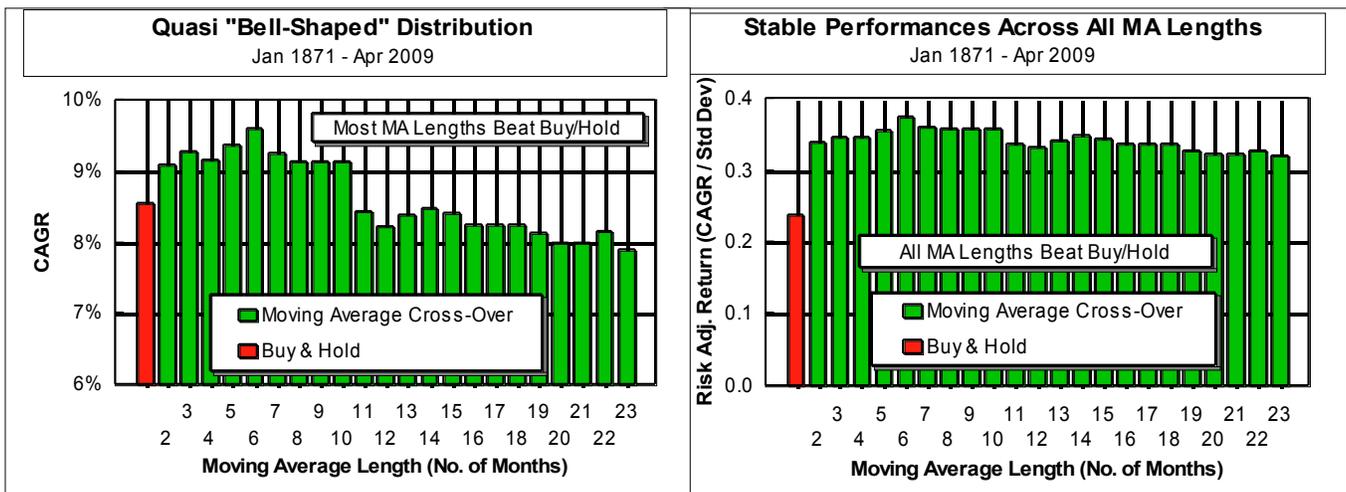


Figure 3

Figure 4

Figure 4 provides even more compelling support for the MAC system. I calculated risk-adjusted returns using the ratio of CAGR to its standard deviation, measured monthly from January 1871 to April 2009. Standard deviation of returns is a generally accepted measure of risk. By this definition, the MAC system beats buy-and-hold hands down across all MA lengths. The stability in risk adjusted return performance and their insensitivity to the MA length show that MAC is a robust system.

Standard deviation treats both up and down volatility as risk. Judging from the “[missing out the best days](#)” argument buy-and-holders embrace, I presume that they don’t consider upside volatility as risk - only downside volatility. A more relevant measure of downside risk is equity drawdown. Drawdown is the percentage decline from the most recent equity peak. There are two ways to evaluate drawdown: average drawdown and maximum drawdown. Figures 5 and 6 respectively show the results of the two methods.

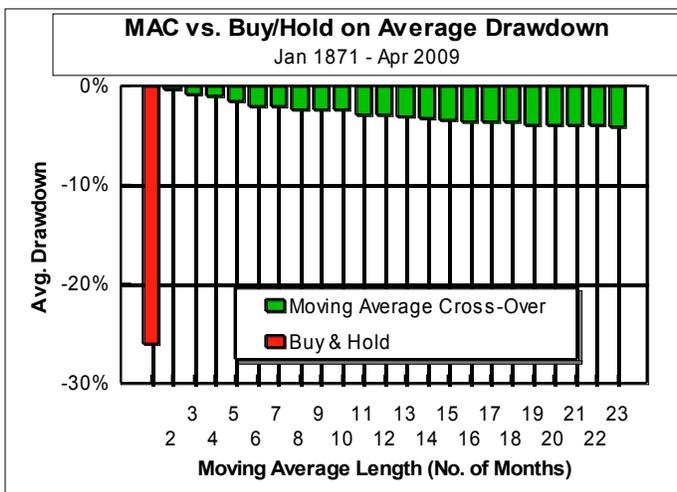


Figure 5

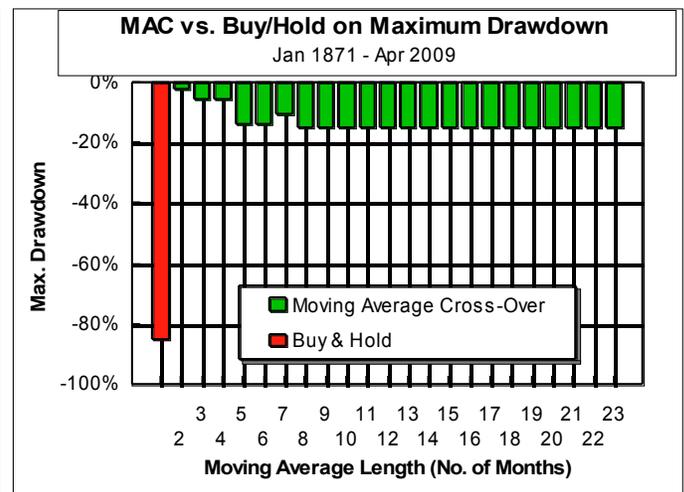


Figure 6

If you don’t view price surges as hazardous and consider only price plunges as risky, then you surely won’t care for the buy-and-hold approach. Buy-and-hold delivers a whopping *negative* 85 percent maximum drawdown, courtesy of the Super Crash from the 1929 peak to the 1932 trough. Even the average drawdown is a painful *negative* 26 percent. In comparison, the maximum drawdown for MAC is only *negative* 15 percent and the average drawdown is no worse than *negative* 4 percent.

I ignored both transaction costs and taxes, so now let's check on these assumptions. *Figure 7* shows the number of round-trip trades (from buy to sell) for the different MA lengths. The average is 0.38 trades per year, or one round-trip every 2.6 years. Even with the 2-month MA, MAC generates only 0.9 round-trip per year, or a holding period of 1.1 years. The low trading frequency of MAC not only keeps transaction costs low, but lowers the tax rates from ordinary income rates to long-term capital gain rates.

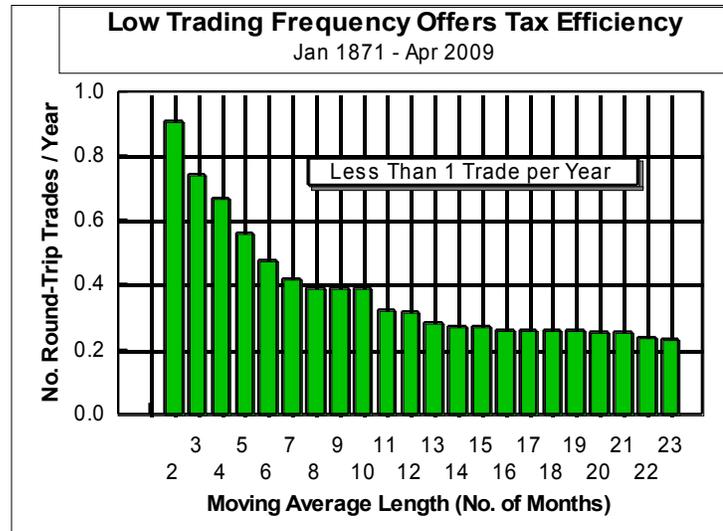


Figure 7

Have we found the Holy Grail?

Based on aggregate performance over the entire 138-year period, the MAC system beats buy-and-hold in both absolute performance and risk-adjusted return. Have we indeed found the Holy Grail that works for all seasons? To find out, stay tuned for Part 2, in which I examine MAC and buy-and-hold on a monthly basis and by decade to see how they compare in all bull and bear markets since 1871.

Theodore Wong graduated from MIT with a BSEE and MSEE degree. He was General Manager of several Fortune-500 companies designing sensors for satellite and military applications. He started a hi-tech company via a LBO in partnership with a private equity firm. He now consults on management and investment best practices. While studying for his MBA, he discovered his true passion was in investment research. He combines engineering analytics with econometrics modeling to enhance quantitative investment analysis. He seeks absolute returns by active risk management in both up and down markets. He can be reached at tswong.advisory@gmail.com.



www.advisorperspectives.com

For a free subscription to the Advisor Perspectives newsletter, visit:
<http://www.advisorperspectives.com/subscribers/subscribe.php>