



City Research Online

City, University of London Institutional Repository

Citation: Pellet, C. and Tamvakis, M. (2016). Investing in Copper Futures: Evaluation of Absolute Return Strategies Within a Discrete-State Hidden Markov Model. SRIC Foundation.

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/16872/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

---- PUBLICATION RESEARCH PAPER ----

**Investing in Copper Futures: Evaluation of Absolute
Return Strategies Within a Discrete-State Hidden
Markov Model**

By Cédric Pellet & Prof. Michael Tamvakis

1. INTRODUCTION

The commodities market has gain in importance since the 1990s, with more investors, traders and merchants buying futures, hedging positions, speculating and more generally using all kind of complex financial instruments that make up the commodities market. Within this period, many things changed as well. First, commodity prices soared during the first decade of this century during the so-called “commodities boom” or “commodities super cycle”. Second, commodities have become an asset class of its own similar to equities or fixed income and part of the alternative investment world.

While each commodity has a unique profile, base metals and more particularly copper enjoys a special relationship with the world economy. Indeed, the performance of copper is reputed to be an excellent way to gauge economic activity around the world. Copper is said to have an ability to predict turning points in the global economy. This belief comes from copper's widespread uses in most sectors of the economy, from homes and factories to power generation and electronics. Thus, it is copper demand reflected in the copper market price that is viewed as a reliable leading indicator of economic health. Rising copper prices suggest high copper demand and, therefore, a growing global economy while declining copper prices suggest decreasing demand and an imminent economic slowdown.

The objective of this work is to define an absolute return strategy using copper futures by combining the forecast power of a Markov model together with a large set of selected simple trading strategies based on fundamental, economic and financial variables. If the Markov model succeeds in mapping out all states of the copper market, then it becomes easy to find the best performing simple trading strategy in each market regime and combine them in one final strategy. Thus, the Markov trading regimes can be viewed as providing a context for the final trading strategy as well as act as strategy filters. To evaluate the different indicators in a trading situation, backtesting of each simple strategy must be achieved. This process is accomplished by reconstructing, with historical data, trades that would have occurred in the past (Kuepper, 2014). Each indicator has its set of trading rules decided beforehand. The result offers statistics used to evaluate the effectiveness of each strategy. The theory is that any strategy that made money in the past is likely to make money in the future, and equally, any strategy that performed poorly in the past is likely to perform poorly in the future. In the case of commodities, and unlike in the equity markets, researchers usually use returns on commodity futures to backtest a strategy as an investor cannot invest in the spot market. However, participation in futures markets always requires a more active approach because there are no infinitely lived futures contracts.

As financial markets are non-stationary¹, it might not make much sense to consider and model them as one distribution over time. This is where trading regimes are helpful. The theory behind regimes is that non-stationary time series, such as market data, move through different “modes” during which market behavior and dynamics are different. In other words, the assumption is that market returns have been drawn from two or more distinct distributions. If one can recognize these trading regimes, one could easily adapt its strategy to react better to these regime shifts. In short, regime detection will help us decide what simple trading strategy to use in what trading environment.

To address the issue of non-stationary and investigate the market regime switches, a Hidden Markov Model will be used on the COMEX copper front-month futures logarithmic returns from 1985 to 2015. The idea is to characterize the copper price dynamics by a regime-switching model in which parameters change according to an unobserved Markov process. The switching times between the different regimes are not directly observable as in the real markets and defined as an unobservable Markov chain. Commodity prices generally display great and persistent swings, with periods of stable prices and times of high volatility (Chevallier & Ielpo, 2013). The price of copper is no exception, and thus should be successfully modeled as a Markov chain. Usually, studies in traditional financial markets find three relevant states: stable, transition, and unstable. The unstable state clearly shows the lowest return and largest volatility. Previous papers found three existing regimes for the SHFE copper futures market (Cheng, Shao, & Guo, 2013). Lia & Zhanga (2009) also found that three regimes were good enough to describe the LME copper forward market (Lia & Zhanga, 2009). However, using a different technique, Chevallier & Ielpo (2013) found five existing regimes for the COMEX copper futures market (Chevallier & Ielpo, 2013). In their book, they use “tests based on the empirical likelihood surface” to choose the appropriate number of regimes instead of using penalized likelihood tests such as AIC or BIC.

Then, the objective is to see what simple trading strategy performs best in each regime (i.e. posts abnormal returns every time the copper market is in this regime) and to combine these best indicators in one overall absolute return trading strategy. For the strategy to be useful in the future, one will have to assess the current state and anticipate the next one by using the Markov conditional probabilities of migrating to another state, and then apply the corresponding simple trading strategy.

From reading this work, one should come away with a usable tool for deciding which strategies

¹ A non-stationary process has a variable variance and a mean that does not remain near, or returns to a long-run mean over time (i.e. random walk with or without a drift, deterministic trends) (Jordanova, 2007).

to use, given the state of the copper market. This work tries not only to be theoretical, but also aims at being applicable either as a premise of trading strategy that could be further developed or as a general indication for analysts working with the copper market on their everyday job.

2. SIMPLE TRADING STRATEGIES

A. Introduction to the different strategies

The copper market reacts strongly to a number of economic news items. At diverse moments in time and with varying intensities, global economic developments, copper fundamental factors (supply and demand) and financial factors (speculation, use of copper as collateral) exert an influence on the copper price (Burgering, 2014). For example, macroeconomic figures may disrupt the correlation with fundamental factors and may affect price volatility. Economic, financial and external factors have more of an impact on short-term trends in the copper price whereas the long-term trends are more driven by fundamentals.

The indicators used in this work have been selected based on three reasons. First, they are listed in the usual economic news provided by financial information providers such as Bloomberg or Reuters. Second, the empirical academic literature has listed some of them as “the main market movers”, meaning news items that affect the financial markets. See (Balduzzi, 2001), (Guegan, 2009) and (Roache, 2009)). Third, they are all related to the copper market in one-way or another. For the news associated with Europe and the USA, the dataset usually starts in 1985 and ends in 2015. For the news associated with China, the dataset usually begins in 1995 and ends in 2015.

In this chapter, we'll define and backtest decision rules on different indicators applied to COMEX copper futures returns. The rule-based strategies will seek absolute returns via outright long or short directional trades on futures contract. The aim is to deliver positive performance in both bull and bear phases of the copper market (positive return, regardless of market direction) based on different indicators, with the objective of limited losses. Volatility measures are important to consider when developing a trading strategy, especially when the strategy involves futures that can be highly leveraged and that are subject to margin calls.

To avoid over-optimization, the quantitative rules are ones that could apply to all commodities, or a subset of commodities (industrial metals in that case), and are not optimized to the extent that it matches the copper returns perfectly. The rules will be deployed on short-term (every day) or medium-term (every month or every quarter) strategies, depending on the frequency of release of each indicator. Moreover, the position will always be open one day after the indicator has been published, as we do not have intraday data (i.e. every hour) to trade during

the day just after the announcement. The exposure will be 100% short or 100% long at all times. As historical data spanning from 1985 to 2015 are used, it will include different events such as financial crises and other disruptive events. Thus, this will give insight about which and in what kind of trading environment, economic indicators perform best at predicting copper price returns. Moreover, it is often a good idea to backtest over an extended time frame that encompasses several different types of market conditions. Only in this way one can ensure that the strategy has encountered enough variation in the market conditions to stand a reasonable chance of being able to adapt to changing market conditions in the future.

To get the most accurate backtesting results as possible, it is common practice to imitate the broker that will be used when the strategy goes live. For instance, commission amounts, lot sizes, position-sizing rules, margin requirements and much more should be taken into account in the different backtests. However, we will not take it into account in this work, as we often trade only once per month or once per quarter, it is less of importance and greatly simplifies the backtesting work.

B. Choosing the benchmark

Futures are particular as they trade in a succession of short-lived contracts that are only active for a specified number of months. To backtest strategies on historical futures data, one needs to aggregate the futures contract prices to generate a similar continuous stream of prices. The usual rule is always to consider the front-month contract (nearest expiration date) and roll the price into the next front-month contract around expiration time. This process of selection determines the actual contract to consider for any given date. However, gaps can appear between the old contract price and the new one because of contango, backwardation and other factors such as seasonality. Indeed, these factors generate a difference in the price of different expiration date contracts, and these gaps can make the dataset appear fragmented.

The primary benchmark for copper is quoted on the LME with the 3-month contract. The issue is that this contract is quoting the price of a new 3-month forward every day, thus backtesting on the LME data would not be realistic. Indeed, the day after one buys a forward contract it ceases to be a 3-month forward. Even if the forward ultimately converges with the cash market, the only data we have from the LME are quoted prices, each day, for a new 3-month forward. Therefore, when extreme and/or persistent backwardation or contango happen, relying on the quoted prices of a new 3-month forward can result in dramatic errors in the PnL estimates (Crittenden, 2010).

To prove this point, one can create a synthetic equity curve representing one contract rolled every three months, using the LME 3-month forward and the LME spot market. For instance,

one buys the current 3-month forward, then sells it three months later at spot price, and rebuys one new 3-month forward, and so on and so forth. “Most of the time the synthetic equity curve closely tracks the quoted prices for LME Copper, both 3-month forwards and cash” (Crittenden, 2010). However, there was a massive drawdown in the 1980's that only appears if one rolls the position. The contango was extreme and persistent during that period.

The luxury we have with copper is that not only the LME quotes the red metal but also COMEX. Moreover, COMEX copper quotes normal futures contracts every month. The continuous front-month COMEX copper futures clearly follow the LME synthetic rolling 3-month and show the effect of contango and backwardation. Thus, the chosen benchmark for this work will be the “COMEX Copper Futures, Continuous Rolling Contract #1” from the Quandl website. Quandl's continuous contracts are created using the simplest possible roll algorithm: "end-to-end concatenation". There's no price adjustment, and the roll dates are simply the last trading dates.

C. Data transformation

The simple percentage change is not accurate in the case of market prices because it is not possible to reliably add together the simple percentage change numbers over multiple periods. However, the continuously compounded return can be scaled over a longer time frame. Thus, the benchmark is computed as continuous log returns with a daily, monthly, quarterly or yearly frequency depending on the indicators chosen for the strategy.

Most of the indicators series are monthly growth rates (changes in levels expressed with respect to the previous month) computed as $(S_t - S_{t-1})/S_{t-1}$ where S_t is the variable of interest at time t so that they can be interpreted as percentage changes. MoM (Month-over-Month) measures tend to be more volatile as they are more affected by one-time events (e.g. stock market crash). For strongly trending variable, we use the month-on-month growth rate (better than daily).

Quarter-over-Quarter (QoQ) figures in our case are computed as change compared to the previous similar quarter one year ago, such as comparing Q2-2014 with Q2-2013. Several significant indicators are released on a quarterly basis, such as the Gross Domestic Product (GDP). Quarter-over-Quarter numbers will tend to be more volatile than Year-over-Year figures but less volatile than Month-on-Month numbers.

Year-over-Year (YoY) numbers report the changes in a year's worth of data, in comparison with the previous year. YoY statistics integrate more data and thus can give a better long-term picture of the underlying report figure.

Predictors that were not transformed to growth rates are: the spread between the 10 year U.S. Treasury Bond and the 3 months U.S. Treasury Bill yield, the VIX index, the TED spread and the PMI New Orders/Inventories ratio. Changes in copper demand and inventories are computed as first differences (in metric tons). Changes in net long/short traders' positions are calculated in the same way.

Last, all indicators that weren't on a daily scale were transformed into daily data to integrate daily copper returns into every backtests. This was done to reflect as much as possible trading reality, and especially daily drawdown and the fact that some indicators are published at the beginning of the month and others around the 15th of each month.

D. Backtest of the different simple strategies

All charts and statistics in this chapter were generated with the econometric software R and several financial packages such as PerformanceAnalytics and Quantmode.

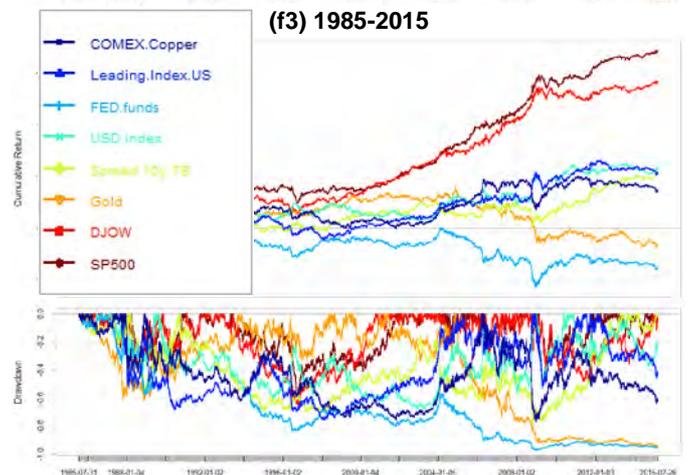
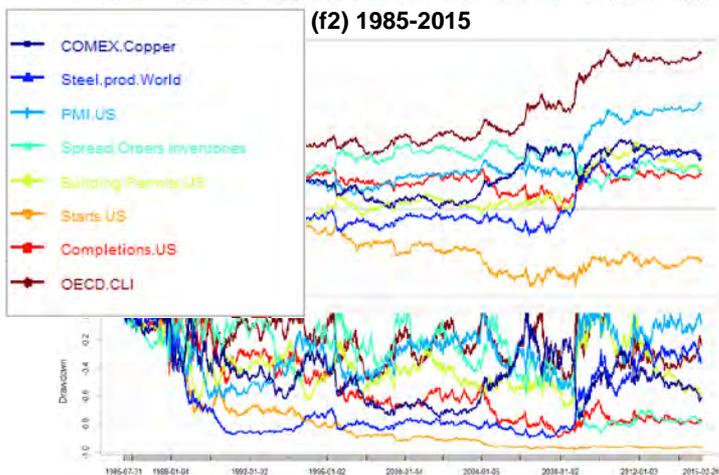
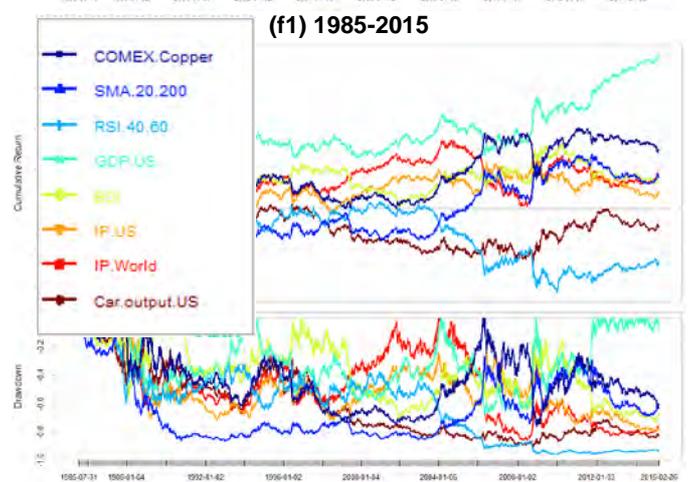
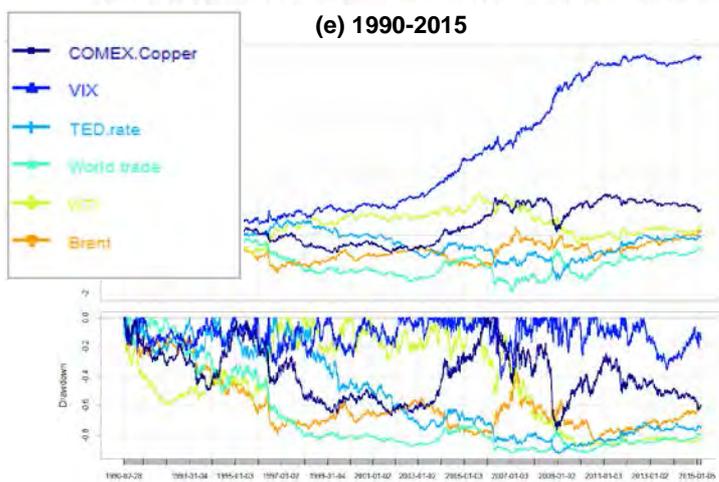
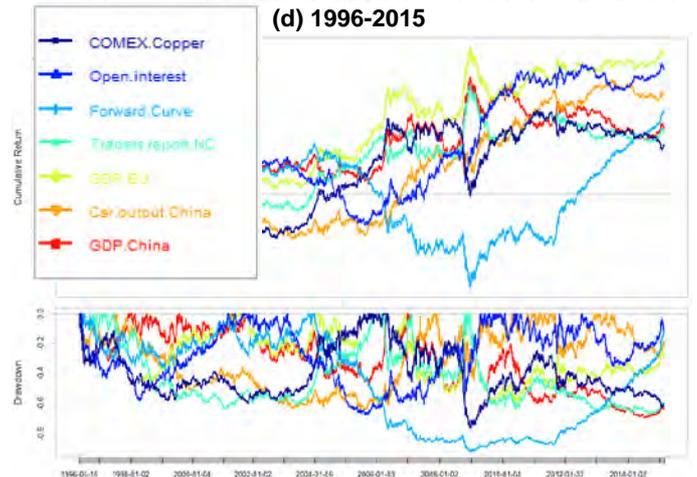
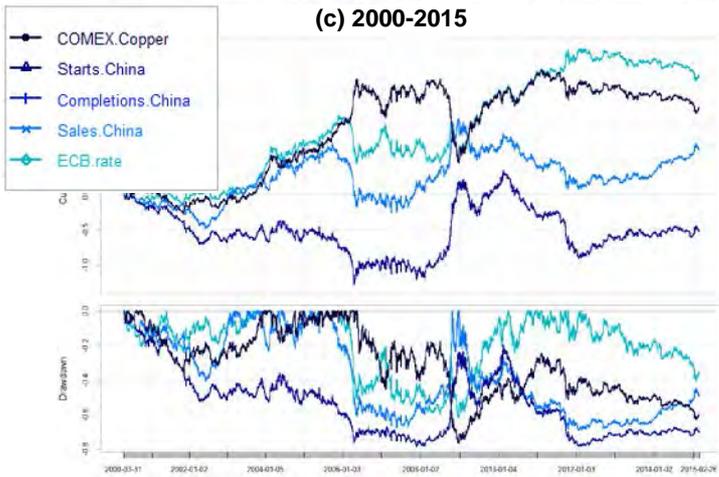
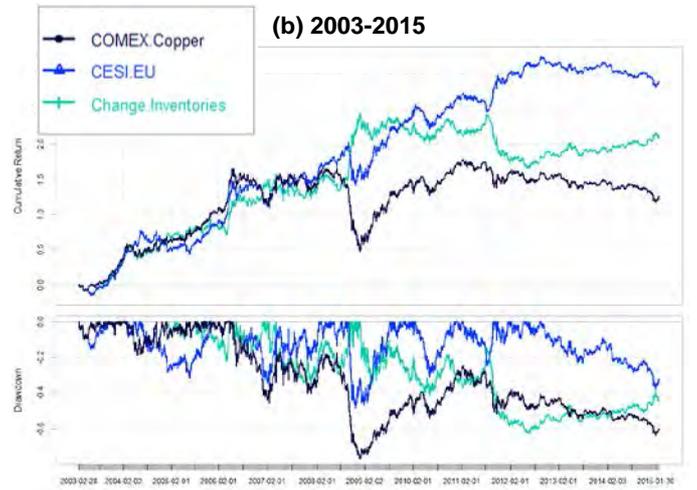
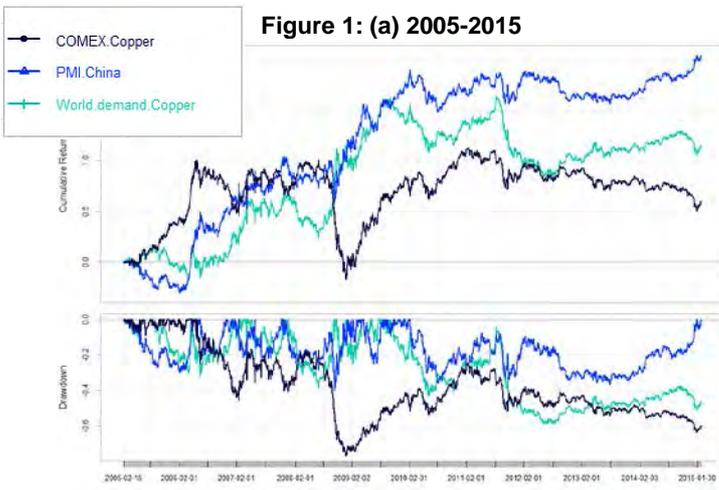
Table 1: Summary of the backtested simple strategies

	Period	Trading rules	Cumulative return	Annualized return	Annualized Std Dev	Annualized Sharpe	Annualized Info ratio	Modified CVaR	Max drawdown
Benchmarks									
	2005-2015	Buy & Hold	60.10%	6.03%	33.79%	0.17	/	-4.90%	76.67%
	2003-2015	Buy & Hold	125.33%	10.59%	32.95%	0.32	/	-5.20%	76.67%
·	2000-2015	Buy & Hold	122.15%	8.36%	30.77%	0.27	/	-4.60%	76.67%
	1996-2015	Buy & Hold	75.03%	4.08%	30.29%	0.13	/	-4.90%	76.67%
	1990-2015	Buy & Hold	86.87%	3.55%	28.56%	0.12	/	-4.70%	76.67%
	1985-2015	Buy & Hold	149.10%	5.15%	29.14%	0.17	/	-5.10%	76.67%
Quantitative / Technical									
·	1985-2015	(T20)>(T200); 1; -1	77.28%	2.67%	29.14%	0.09	-0.06	-4.70%	86.70%
·	1985-2015	(T40)>(T60); 1; -1	-133.40%	-4.60%	29.14%	-0.15	-0.21	-4.30%	96.18%
Economy / Sentiment									
·	1996-2015	(T)>=(T-1); 1; -1	212.44%	11.55%	30.29%	0.38	0.16	-4.70%	60.32%
·	1996-2015	(T)>=(T-1); 1; -1	100.58%	5.47%	30.29%	0.18	0.03	-2.80%	69.60%
·	1985-2015	(T)>=(T-1); 1; -1	376.28%	13.01%	29.14%	0.44	0.19	-3.70%	67.90%
·	1990-2015	(T)>=(T-1); 1; -1	-50%	-2%	28.56%	-0.07	-0.12	-4.40%	92.82%
·	1985-2015	(T)>=(T-1); 1; -1	82.27%	2.84%	29.14%	0.09	0	-2.70%	73.23%
·	2003-2015	(T)>=(T-1) OR (T)>0; 1; -1	289.15%	24.44%	32.92%	0.74	0.45	-5%	48.23%
Cu Fundamentals									
·	2005-2015	(T)>=(T-1); 1; -1	115.20%	11.55%	33.78%	0.34	0.12	-4.20%	58.89%
·	1985-2015	(T)>=(T-1); 1; -1	41.74%	1.40%	29.14%	0.04	0.08	-4.80%	80.07%
·	1985-2015	(T)>=(T-1); 1; -1	80.65%	2.70%	29.14%	0.09	-0.05	-4.70%	86%
·	2003-2015	(T)>0; -1; 1	208.26%	17.60%	32.94%	0.53	0.15	-4.60%	62.28%
·	1996-2015	(T)>=(T-1); 1; -1	156.64%	8.52%	30.29%	0.28	0.1	-4.20%	64.61%
·	1985-2015	(T)>=(T-1); 1; -1	-48.40%	-1.60%	29.14%	-0.05	-0.15	-4.50%	91.28%
·	1985-2015	(T)>=(T-1); 1; -1	148.99%	5.15%	29.14%	0.17	0	-4.70%	89.89%
·	2005-2015	(T)>=(T-1); 1; -1	203.70%	20.42%	33.77%	0.6	0.32	-4.40%	40.10%
·	1985-2015	(T)>=(T-1); 1; -1	300.54%	10.39%	29.14%	0.35	0.12	-3.50%	61.39%
·	1985-2015	(T)>=(T-1); 1; -1	119.83%	4.14%	29.14%	0.14	-0.02	-3.20%	67.87%
·	1985-2015	(T)>=(T-1); 1; -1	-148.20%	-5.10%	29.14%	-0.17	-0.24	-4.00%	97.80%
·	1985-2015	(T)>=(T-1); 1; -1	95.24%	3.29%	29.14%	0.11	-0.04	-3.50%	89.22%
·	2000-2015	(T)>=(T-1); 1; -1	-52.60%	-3.60%	30.77%	-0.11	-0.26	-2.60%	78.89%
·	2000-2015	(T)>=(T-1); 1; -1	-52.60%	-3.60%	30.77%	-0.11	-0.26	-2.60%	78.89%
·	2000-2015	(T)>=(T-1); 1; -1	62.39%	4.27%	30.77%	0.13	-0.08	-3.50%	69.31%
·	1985-2015	(T)>=(T-1); -1; 1	113.13%	3.91%	29.14%	0.13	-0.03	-3.70%	86.23%
·	1985-2015	(T)>=(T-1); 1; -1	449.12%	15.52%	29.14%	0.53	0.25	-4.20%	51.58%
·	1985-2015	(T)>=(T-1); 1; -1	215.85%	7.46%	29.14%	0.25	0.06	-4.60%	75.99%
Financials									
·	1990-2015	(T)>=(T-1); 1; -1	22.83%	0.90%	28.56%	0.03	-0.06	-3.90%	88.34%
·	1990-2015	(T)>=(T-1); 1; -1	14.41%	0.50%	28.56%	0.02	-0.07	-4.40%	84.47%
·	1985-2015	(T)>=(T-1); 1; -1	-81.60%	-2.80%	29.14%	-0.09	-0.18	-3.30%	94.69%
·	1985-2015	(T)>=(T-1); 1; -1	564.67%	19.52%	29.14%	0.67	0.39	-2.90%	64.52%
·	1985-2015	(T)>=(T-1); 1; -1	684.38%	23.65%	29.14%	0.81	0.51	-2.70%	58.67%
·	1985-2015	(T)>=(T-1); -1; 1	224.54%	7.76%	29.14%	0.26	0.06	-3.20%	67.76%
·	1985-2015	(T)>=(T-1); -1; 1	-152%	-5.20%	29.14%	-0.18	0.26	-4.30%	97.07%
·	1990-2015	(T)>=(T-1); -1; 1	605.58%	24.78%	28.56%	0.86	0.61	-2.70%	41.89%
·	1985-2015	(T)>=(T-1); 1; -1	220.05%	7.61%	29.14%	0.26	0.05	-3.50%	78.83%
·	1990-2015	(T)>=(T-1); -1; 1	-10.30%	-0.40%	28.56%	-0.01	-0.09	-4.10%	92.87%
·	2000-2015	(T)>0; -1; 1	168.12%	11.51%	30.77%	0.37	0.08	-4.40%	61.83%
Market Dynamics									
·	1996-2015	(T)>=(T-1); 1; -1	190.76%	10.37%	30.29%	0.34	0.15	-4%	67.61%
·	1996-2015	(T)>=(T-1); 1; -1	93.45%	5.08%	30.29%	0.16	0.02	-4.70%	67.17%
·	1996-2015	(T)>=0; -1; 1	127.90%	30.29%	22.95%	0.22	0.06	-3.50%	92.39%

***How to read the trading rules:** (T)>=(T-1);-1;-1 : if the most recent indicator (T) is higher or equal to the previous indicator (T-1), then position is long (1), otherwise position is short (-1)

***Sharpe and Information ratio:** As the standard deviation fails to differentiate between upside deviation and downside deviation, all strategies will have the same standard deviation as the benchmark. Thus, the Sharpe ratio that defines risk as standard deviation won't be a good measure of performance. The information ratio should do a better job. It answers two questions: first, did the active strategy outperform the passive benchmark? Second, was the active strategy able to outperform the benchmark consistently? Generally speaking, an information ratio in the 0.40-0.60 range is considered reasonably good.

***Modified Conditional Value at Risk:** As the standard method for calculating VaR is not appropriate for returns distributions that are asymmetric or display fat tails, the Modified Value at Risk (or MVaR) is used here. Indeed, the MVaR adjusts the standard deviation to account for skew and kurtosis in the returns distribution (higher negative skew and higher kurtosis act to increase VaR). However, the tail end of the distribution of loss is not typically evaluated with the MVaR. On the other hand, the Modified Conditional Value at Risk (MCVaR) is more sensitive to the shape of the loss distribution in the tail end of the distribution. The MCVaR at 95% level represents the expected return on the portfolio in the worst 5% of the cases.



Strategies from 2005 to 2015:

The “PMI.China” strategy is the best performing one, with a cumulative return of 2.037 over the considered period, and an annualized return of 20.42%. The PMI index “provides the earliest clues of how the economy has fared during the previous four weeks” (Chevallier & Ielpo, 2013). The manufacturing industry is an early cyclical sector and has a high correlation with overall economic activity. It is a good indicator to identify or interpret implications for economic growth at an early stage. When global industrial activity increases and order books are filling, demand for base metals such as copper will increase (Burgering, 2014). The PMI is setting the tone for the upcoming month and other indicator releases. Moreover, China consumes around 45 percent of the globe’s copper and thus its economic growth is one of the most important drivers of copper’s future demand. Furthermore, both strategies have a lower maximum drawdown than the benchmark. In a more general note, both strategies can forecast the 2008 crisis and post huge gain during the period. They also have a slightly lower MVaR than the benchmark, meaning that a particular loss won’t exceed, at 95% confidence level, 3 percent on any given day.

Strategies from 2003 to 2015:

The “CESI.EU” strategy is the best performing one, with a cumulative return of 2.89 over the considered period, and an annualized return of 24.42%. The “Change.Inventories” strategy best forecasts the 2008 crisis compared to the “CESI.EU” strategy that drops at the beginning of the crisis and then reverses. The Citigroup Economic Surprise Index (CESI) tracks how economic data are comparing relative to the consensus forecasts of market economists. The index increases when economic data beat economists’ consensus estimates and decreases when data are below estimates. For example, when economic data are missing expectations (the index falls), one would expect copper to fall, or at least gain less, and it seems that it is indeed the case. Of course, the indicator is not perfect because of factors such as data’s time lag, the forward-looking nature of the market, and shifting expectations. Moreover, the EU accounts for around 18% of global copper consumption with Germany being the main player in Europe. One of the main components of the CESI index is the German GDP and each GDP revision (every month) is counted in the index calculation. Thus, the CESI should be more precise than the strategy based only on the official quarterly published European GDP. Lastly, this index can be a powerful tool for predicting where the economy is headed for another reason. Generally speaking, economists lag the market. Indeed, economists tend to be overly pessimistic from their recent experience when the economy starts entering a new growth’s phase. As a result, they expect that many measures of the economy will not perform as well as they do in reality. The surprise Index will turn positive. On the opposite case, as the economy enters a new contraction’s phase, economists maintain the optimism that was built up during the "good times". As the real world decline in economic metrics, the surprise Index turns negative.

Strategies from 2000 to 2015:

The “ECB.rate” strategy is the best performing one, with a cumulative return of 1.68 over the considered period, and an annualized return of 11.51%. This strategy and the benchmark returns are negatively skewed, which for investors can mean a greater chance of extremely negative outcomes. All other strategies related to China housing are performing worse than the benchmark. Moreover, “Starts.China” and “Completions.China” are showing the same statistics, meaning both strategies are the same. No strategies particularly stand out from the benchmark in that section.

Strategies from 1996 to 2015:

The “GDP.EU” strategy is the best performing one, with a cumulative return of 2.12 over the considered period, and an annualized return of 11.55%. All strategies perform better than the benchmark. Three strategies (GDP.EU, GDP.China and Traders.report.NC) forecast the 2008 crisis. However, these three strategies also drop when the copper market came back in 2009. The tech bubble of 2000-2002 does not seem to have affected copper returns at that time. The GDP.EU strategy probably performs well for the following reason. Given its widespread use in transportation, manufacturing and construction the copper price is sensitive to any economic slowdown. Any time GDP numbers or growth numbers beat expectations; there is going to be an underlying belief that demand for raw commodities will increase. As the EU accounts for around 18% of global copper consumption, trends in European GDP matter for the copper market.

Strategies from 1990 to 2015:

The “VIX” strategy is the best performing one, with a cumulative return of 6.05 over the considered period, and an annualized return of 24.78%. In a more general note, only two strategies (VIX and World.trade) forecast the 2008 crisis. The VIX measures the volatility implied by option prices on the S&P500 over the coming month. A higher value in the VIX indicates that market participants expects an overall negative economic or financial outlook and hence an increased aversion to risk (suggesting there is a demand for protection). The COMEX copper seems connected in some way or another to the US stock market; thus it is obviously linked to the VIX. Another reason the strategy is working so well is probably because the VIX is said to have predictive power in the foreign exchange market, and thus should contain predictive information about copper. Indeed, Brunnermeier et al. (2009) have shown that the VIX predict higher returns in carry trade strategies used in the foreign exchange market (Brunnermeier, Nagel, & Pederson, 2009). They found that carry trades lose money on average in times of rising VIX because the carry trades are unwound at that moment. Moreover, Rabobank found out that the VIX could help understanding currency moves in commodity producer’s country. The theory is as follows: an increase in the USD contributes to the sell-off in commodities even if demand has looked quite strong, in making US commodities less competitive against those priced in depreciating currencies. And in turn, “the VIX, as a measure of market volatility and investor panic,

provides a useful measure of investors' enthusiasm for investing into the dollar, viewed as a safe haven in times of market panic. If fear remains elevated, this will place continued pressure on non-US dollar currencies" (Rabobank, 2011). This is also linked to what Brunnermeier et al. (2009) have found. Indeed, in the case of an USD increasing while other currencies decreasing, carry trades would most probably suffer, as the USD - the low-interest rate currency - would increase, and commodity currencies such as the Australian dollar – the high-interest rate currency - would decrease, reducing the overall spread between the two and thus resulting in carry trade being less attractive.

Strategies from 1985 to 2015:

The "SP500" strategy is the best performing one, with a cumulative return of 6.84 over the considered period, and an annualized return of 23.65%. However, it's the "OECD.CLI" strategy that has the lowest maximum drawdown but with a cumulative return of "only" 4.49.

Like the GDP.EU strategy, the GDP.US strategy probably performs well for the same reason. Given its widespread use in transportation, manufacturing and construction the copper price is sensitive to any economic slowdown. Any time GDP numbers or growth numbers beat expectations; there's going to be an underlying belief that demand for raw commodities will increase. As the US accounts for around 8% of global copper consumption, trends in US GDP matter for the copper market. Furthermore, the COMEX copper is quoted in the USA, and will probably react more significantly to the US data such as the American GDP numbers.

The PMI.US strategy also performed quite well, probably for the same reasons that were mentioned regarding the PMI.China strategy. Even though the U.S. manufacturing "is not the huge component of total GDP that it once was, this industry is still where recessions tend to begin and end" (Barnes R. , Economic Indicators: Purchasing Managers Index (PMI), 2014). Thus, it is a good indicator to identify or interpret implications for economic growth at an early stage and, as we've seen, copper is reacting to GDP from the US and the EU.

The OECD composite leading indicators (CLI) is a measure of world economic strength. The system is centered "on the growth cycle approach, where business cycles and turning points are measured and identified in the deviation-from-trend series" (OECD, 2015). The aim is to provide early signals of turning points in economic activity (business cycles). The OECD computes leading indicators for 30 OECD countries and major emerging economies. Then, the CLI indicator is released monthly, reflecting conditions two months prior. But as the turning points of the CLI consistently precede those of the business cycle between 6 to 9 months, it is still considered as a good leading indicator. As industrial metals are pro-cyclical - their returns vary with economic activity - means that copper is affected by the business cycle. Indeed, copper consumption will vary depending on the industrial activity. In short, copper returns should be higher during expansions than during recessions.

Therefore, the strategy just longs copper when cyclical assets should perform better and shorts it when defensive assets (their returns are relatively insensitive to economic activity) should perform better.

It was recently documented that returns in the U.S. stock market are predictive for returns in various other global equity markets. The explanation is that the U.S. stocks market “is a news or information hub, where news about economic and financial developments are most efficiently absorbed and reacted to” (Buncic & Moretto, 2014). In that sense, copper quoted in the US market can also follow this pattern, which would explain why the DJOW and SP500 strategy are performing so well. Another reason is that rising equity prices can be a forward-looking indicator of the strength in economic activity. Thus, everyday participants in the copper market would adjust their positions according to what happened in the stock market the day before. Following that thought, it would be interesting to see if there exist any relation between the S&P 500 futures and the copper futures market. Furthermore, the S&P and copper forward curves would be interesting to compare as well.

3. MARKOV TRADING STRATEGY

This chapter investigates, by constructing a hidden Markov model, the state transition behaviors in the price fluctuation of copper from July 1985 to February 2015. Indeed, the Markov model is a method applied to investigating the structural changes of nonlinear time series (Cheng, Shao, & Guo, 2013). In finance, hidden Markov models are generally used to modeling equity or bond prices with business cycles (Bae, What are the regimes in the commodity market?, 2012). However, there were only a few studies for the regime analysis of the commodity market in general. See the Introduction for the list of the few papers that talk about regimes in the different copper futures markets. The model will use continuous front-month COMEX copper futures data, including a total of 7291 daily data from July 1985 to February 2015. The trading prices were selected from the daily closing prices of the market and then transformed to log yield rates of closing prices.

All the results obtained in this chapter were computed with the statistical software R and more specifically with the package DepmixS4.

A. Selection of the number of states and learning algorithm

Two assumptions are essential when constructing a hidden Markov model (Bae, What are the regimes in the commodity market?, 2012):

- 1) The conditional probability of the current observation given the current state is independent of previous states and observations: $P_{Y(t)|S_{1,t}, Y_{1,t-1}} = P_{Y(t)|S(t)}$
- 2) The state process has the Markov property. In other words, the conditional probability

distribution of futures states given previous states and observations depends only on the present state: $P_{S(t+1)|S_{1,t},Y_{1,t}} = P_{S(t+1)|S(t)}$

Under these basic assumptions, one has to decide on two parameters of the hidden Markov model:

- 1) The first thing to decide is the probability distribution/density function of the observations given the present state. Indeed, to infer the hidden states from the returns, one has to specify a distribution model. In this work, we suppose that the daily returns of the COMEX copper market follow a multivariate normal distribution that depends on the market regimes (Donninger, 2012). However, we know that assets returns are not normal as they have fat tails. Mulvey & Zhao (2010) argue that hidden Markov models used with the normal distribution can address relatively accurately fat-tailed distributions (Mulvey & Zhao, 2010). Other more complicated time series model such as ARCH or GARCH could be used.
- 2) Another parameter that should be determined is the number of hidden states. We cannot directly observe the regimes/ market states. The only observable information is the daily returns of the COMEX copper. No one knows how many states are in the market. In this work, we assume that the number of hidden states is discrete and finite. To decide on the right number to use, we will use both Bayesian information criterion and Akaike's information criterion.

We assume that the daily returns of the COMEX copper market are the observations of the model and the returns at time t depend on the regime at that time. The model by using Fraser (2008) can be written as : $Y(t)|S(t) = y(t)|s(t) \sim N(\mu_{s(t)}, \Sigma_{s(t)})$, where $\mu_{s(t)}$ denotes the mean of the daily return series at the end of time t , $\Sigma_{s(t)}$ is the covariance matrix of daily returns under state $s(t)$. The assumption of normality of returns and finite number of states allow the Baum–Welch algorithm to estimate model parameters. The Baum–Welch algorithm is an expectation–maximization (EM) algorithm for applications of HMM. Given the initial set of parameters θ and the realized series of observations $y_{1,T}$, the Baum-Welch algorithm iterates forward-backward estimation and gives a solution that always converges to a local maximum of the likelihood function $P_{\theta}(y_{1,T})$.

Last, we need to determine the proper number of states. A higher number of states give a more precise model but at the same time increase computational burden because of exponentially increasing number of parameters that need to be estimated (Bae, Kim, & Mulvey, Dynamic asset allocation for varied financial markets under regime switching framework, 2013). Moreover, it becomes difficult to interpret the meaning of estimated parameters for each state when the existence of many states is assumed. Two criterias for determining the appropriate number of states are often used.

The first criterion is the Bayesian information criterion (BIC). BIC gives the adjusted value of likelihood by penalizing the number of parameters and observations. Given any two estimated models, the model with the lower BIC value is the one to be chosen. The second method to measure the goodness of the model is the Akaike's information criterion (AIC), which only uses the number of parameters to penalize the fitness of models. Table 11 shows the different BIC and AIC values for two- to four-state models. The results exhibit that the lowest BIC value is found in the 3-states model. However, the lowest AIC value is found in the 4-states model. The difference is not significant enough to ignore the 3-state model that have less number of states and the improvement from two-state models to three-state models is the biggest. Thus, this work will use the three-state model for easier interpretation and calibration to estimate parameters.

Table 2: Log likelihood, AIC and BIC values

	2 States	3 States	4 States
Log likelihood value	19745.83	19938.9	19974.05
AIC	-39477.66	-39849.79	-39902.1
BIC	-39429.4	-39753.27	-39743.53

Table 3: Transition probability matrices

	State 1	State 2	State 3	State 4
2-state model				
State 1	98.10%	1.80%		
State 2	4.80%	95.10%		
3-state model				
State 1	98.70%	0.00%	1.20%	
State 2	0.00%	90.60%	9.30%	
State 3	0.90%	1.40%	97.50%	
4-state model				
State 1	97.70%	16.30%	0.00%	0.60%
State 2	9.70%	90.20%	0.00%	0.00%
State 3	0.00%	0.00%	76.00%	23.90%
State 4	2.80%	0.00%	72.90%	24.20%

B. Interpretation of the identified regimes

There are 2 main persistent trends in comex copper: one is a positive trend with an annualized expected return of 20.66% (for a volatility of 27.51%), and one is negative with an annualized expected return of -60.48% (for an annualized volatility of 65.46%).

The values of parameters μ_1 , μ_2 , and μ_3 were estimated to be -0.00013, -0.00204 and 0.00082, showing the fluctuation behaviors of COMEX copper prices are mainly transited under three states.

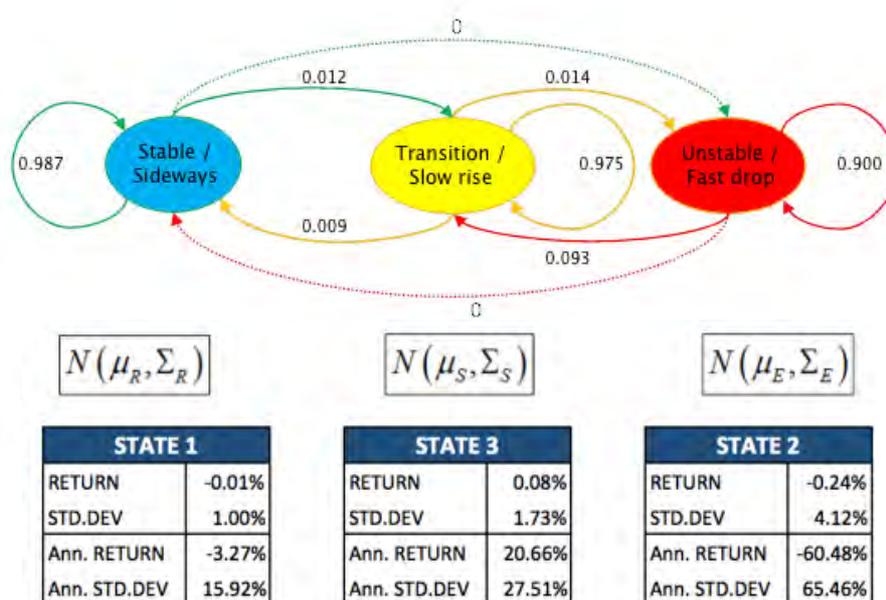
State 1 can be considered as stable or sideways market, with the lowest volatility (around 1%). The average yield rate of COMEX copper prices was -0.24 percent in State 2. This state can be considered as a fast dropping market (volatile bear) with the highest volatility of all three regimes. The average return rate of the prices was 0.08 percent in State 3, and the yield rate was growing in a somewhat slow speed. The market was rising slowly.

Results show that state 1 and state 2 have near-zero or negative returns. This makes sense because “the commodity price is determined by the supply and demand of the market. The lack of supply brings a rise in commodity price, and an effort is made to increase the supply. However, there could be extra supply even after all demand is fulfilled because of the time lag in transporting the goods and this will trigger a sudden collapse in price” (Bae, Kim, & Mulvey, Dynamic asset allocation for varied financial markets under regime switching framework, 2013). The period of supply shortage can be represented as state 3 and the fall of price might occur in state 2.

The main points that we can conclude by looking at the probabilities are:

- 1) The returns of the copper market are ordered as slow growth > sideways > fast drop. Moreover, the sharp drawdown periods are mostly included in the crash regimes.
- 2) Furthermore, the probability of today’s regime being the same as that of yesterday is very high, exceeding 95% for regime 1 and 3.
- 3) If yesterday’s regime was sideways, the probability of today’s regime being crash is practically 0. Likewise, if yesterday’s regime was crash, today’s regime never becomes sideways. The slow growth regime could be used as the “warning sign” for the crash regime being ahead.
- 4) $P_{31}=0.009 < P_{32}=0.014$, which means the probability of a sharp decline of COMEX copper prices after a small rise is greater than the probability of stable/sideways market, but the probabilities of the two mentioned cases are significantly lower than the market staying in State 3 (slow rise) (Cheng, Shao, & Guo, 2013).
- 5) $P_{22}=0.9$, which means the probability that COMEX copper prices continue to steeply decline after great fall is 90% (Cheng, Shao, & Guo, 2013).

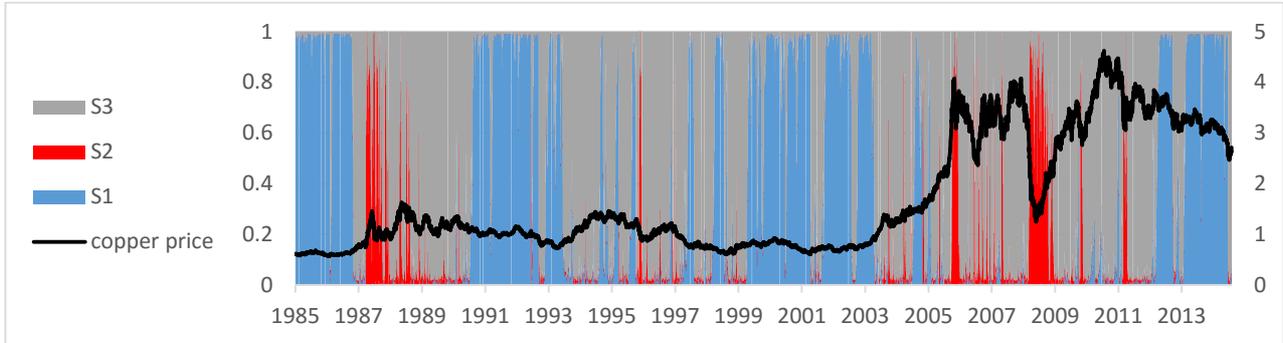
Figure 2: Three states model and transitions probabilities



C. What indicator for what regime?

The different behavior of the copper price in the different state is the primary point of the model. The idea is to find the best simple strategy in each of the three states showed in Figure 3 below.

Figure 3: States transitions real probabilities



The continuity of each state, and what kind of state will mostly appear at each time can be more visibly seen through the observation of smoothed probabilities:

- if the smoothed rate exceeds 0.5 when S=1, then the copper prices are stable/sideways;
- if the smoothed rate tops 0.5 when S=3, then the copper prices are rising slowly;
- if the smoothed rate is more than 0.5 when S=2, then the copper prices are dropping fast.

Figure 4: States transitions smoothed probabilities

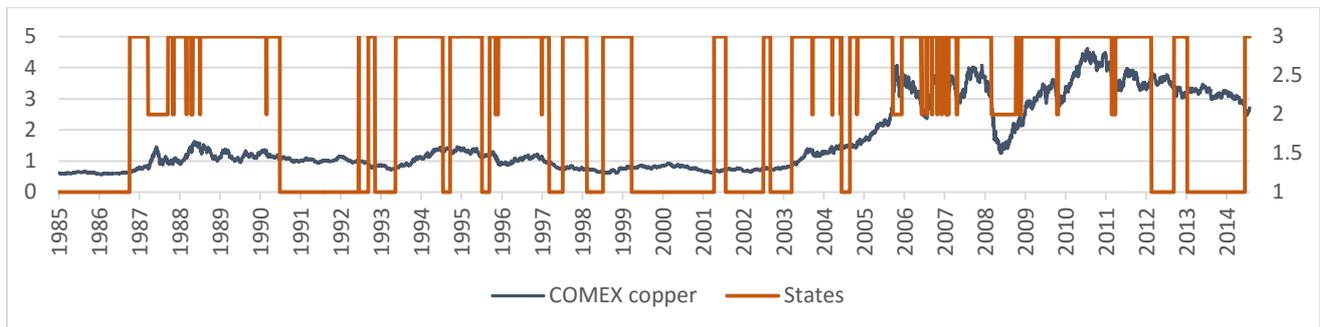


Table 4: B&H benchmark 1985-2015 (3-state model)

	State 1	State 3	State 2
Cumulative return (0)	-72%	292%	-70%
Number of days	2747	4075	469
Number of days (%)	37.67%	55.89%	6.43%

Table 5: Cumulated returns by states from different selected simple trading strategies 1985-2015

	State 1	State 3	State 2
GDP.US	96%	223%	56%
DJOW	217%	160%	186%
SP500	166%	334%	183%
BDI	-24%	-6%	113%
PMI.US	58%	48%	193%
Building.Permits.US	5%	-4%	119%
OECD.CLI	11%	295%	141%
USD.index	-14%	126%	112%
Spread.10y.TB	36%	69%	114%
Leading.Index.US	40%	231%	-56%
VIX (from 1990)	63%	359%	183%

Regime Detection comes helpful when you are trying to decide which strategy to deploy. In regime 1, one should use the DJOW indicator. In regime 3, the SP500 indicator is best suited. Lastly, in regime 2, the PMI.US indicator performs best.

Theoretically, if we backtest it as it is, the combined strategy (from 1985) would have a return of 7.44. In comparison, the best simple trading strategy (SP500) returned 6.84 during the same period. Thus, a 60% return increase in relative term but with an increased MVaR (2.58% vs. 2.4%) and MCVaR (3.55% vs. 2.7%). However, the annualized standard deviation is a bit lower in the Markov strategy (29.10% vs. 29.14%).

If we use a strategy from 1990, the VIX indicator is best compared to SP500 for regime 3. In that case, the Markov strategy returns 7.00 and the best simple strategy from 1990 (VIX) returns 6.05. Again, the MVaR and MCVaR are higher with the Markov strategy (2.52% vs. 2.4%) and (3.50% vs. 2.7%). However, the annualized standard deviation is a bit lower in the Markov strategy (28.50% vs. 28.56%).

If we compare the two Markov strategies, the second one is superior both in terms of annualized return and in terms of risk. Indeed, the Markov90 strategy benefits from a lower annualized standard deviation, a lower maximum drawdown, a lower MVaR and a lower MCVaR. This is underlined by a higher Ann Sharpe ratio and a higher Information ratio. Specifically, the Markov90 helps reduce drawdown in 1996 and 2010.

Figure 5: Markov85 (DJOW+PMI.US+SP500), Markov90 (DJOW+PMI.US+VIX)

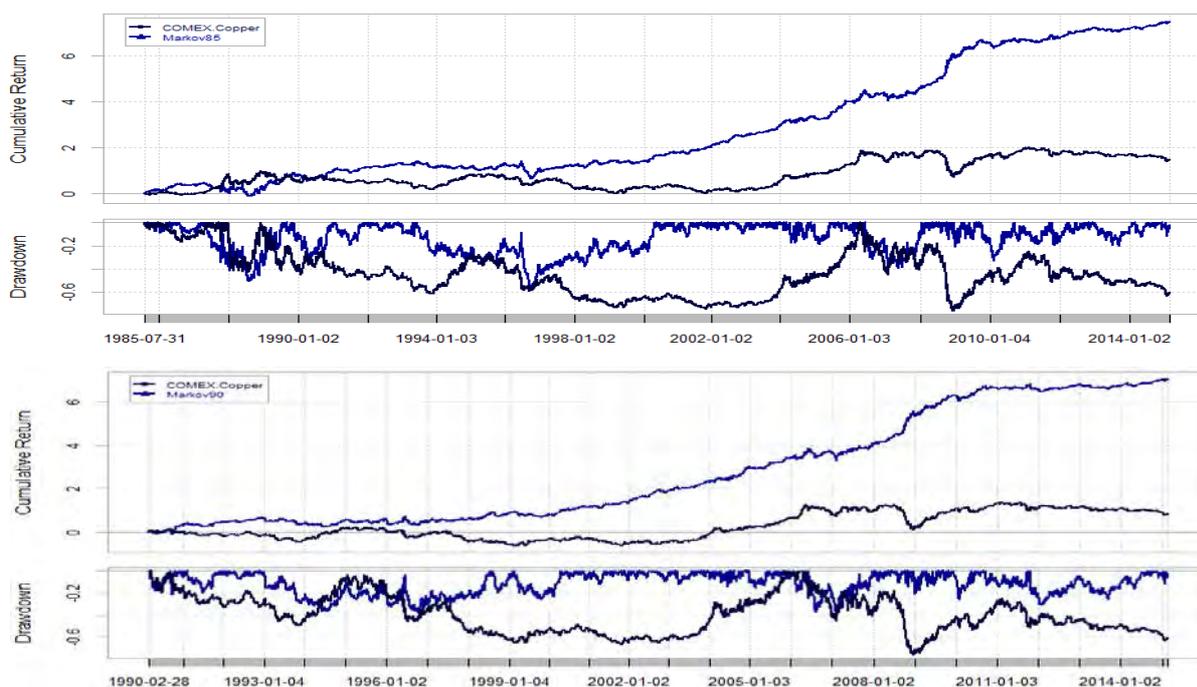


Table 6: Benchmark vs Markov Strategies Statistics

	CMX.Copper85	Markov85	CMX.Copper90	Markov90
Cumulative return (0)	149%	744%	86%	700%
Ann return	5.15%	25.74%	3.55%	28.68%
Ann stdev	29.14%	29.10%	28.56%	28.50%
Max drawdown	76.00%	58.00%	76.00%	45.00%
Modified VaR (95%)	-2.88%	-2.58%	-2.76%	-2.52%
Modified CVaR (95%)	-5.17%	-3.55%	-4.74%	-3.50%
Ann Sharpe ratio (Rf=0%)	0.17	0.88	0.12	1.00
Information ratio	/	0.54	/	0.68

Naturally, a backtest doesn't really make sense as the states were assessed with data going until 2015. This was only done in order to show the theoretical relationship of the simple trading strategies with each market regime. The main issue is that the Viterbi² algorithm uses all data to compute the most likely sequence of states. For instance, the model used the returns for all dates to the Viterbi algorithm. Hence, the Viterbi algorithm might use data in 2010 to determine states in 2009.

The observations from 1985 to 2015 are only used to calibrate this model. Because of the high number of data, we assume that all existing regimes in the copper market already happened at least once and thus are accounted for in the calibration. The probability matrix should be accurate enough to detect the next regimes in the future. Therefore, the two proposed strategies (Markov85 and Markov90) aimed to be used in the future. Moreover, the regimes can be detected using the probability matrix, which should be precise enough as stated before, or can be detected independently by a market analyst. In both cases, depending on the regime, the best simple indicator should be used, as it was properly backtested.

In short, the regimes should be adequately described as enough data are taken into account, but of course it wasn't proved that the Markov probability could detect at the right time the next regime.

In order to address this issue and to backtest the detection ability of the Markov model, we could use different methods:

1. Calibrate the model with 10 years of returns data (1985-1995) and then try predicting the regimes from 1995 to 2015 (out of sample test)
2. Calibrate the model with an expanding window: we start in 1985 and then include every new data in the model
3. Calibrate the model with a rolling window: we always use the last x data to calibrate the model

² When you have found the transition probabilities for the hidden part of the model with the Baum Welch algorithm, then the Viterbi algorithm is used to compute the most likely complete sequence of hidden states conditional on both the returns and the model specification (i.e. probabilities).

The accuracy of the predictions and strategy returns would probably go down as the states should change more often using the expanding or rolling window approach. One solution would be to add a filter that allows the state change only if it is confirmed by a certain number of observations (a sort of delayed state change), but this is for another work.

CONCLUSION

This work uses a selected set of 41 variables to create simple rule-based trading strategies based on COMEX copper futures returns. The role of copper has evolved over time from being a mere commodity used as a primary input in the production process of final goods, to a financial asset that is held and traded for speculative purposes. Therefore, it makes it necessary to use a large and diverse set of variables that control not only for standard demand, inventory and convenience yield factors, but also for the impact that financialization has had on copper.

Covering a sample period from 1985 to 2015, this work shows in the empirical evaluation that some of these simple trading strategies significantly outperform the copper benchmark both in terms of returns and in terms of limit of risk. The best performing simple strategy overall is the « VIX » strategy that shorts the copper if the volatility increased between the day before and the day before last (long in the opposite case). The VIX strategy brings a cumulative return of 6.05 between 1990 and 2015 with a maximum drawdown of 41.89% that happened in 2006 just before the financial crisis. This strategy also has the highest information ratio of 0.61, which means that the strategy consistently outperformed the passive benchmark. The biggest improvement in the performance compared to the benchmark model is realized over the September 2008 to beginning of 2009 period, which seems to have been the worst phase of the 2008 financial crisis. Two other simple strategies are also worth mentioning, the “DJOW” and “SP500” strategies which performed very well with respectively a cumulative return of 5.64 (information ratio of 0.39) and 6.84 (information ratio of 0.51). However, both strategies have much higher maximum drawdowns. The DJOW suffered a maximum drawdown of 64.52%, and the SP500 suffered a maximum drawdown of 58.67%. It seems that the US copper market is linked to the two major US financial market indices. Indeed, the VIX (representing the SP500 expected volatility), the SP500, and the DJOW strategies performed best. All three strategies are taking trading decisions the day after the variable (i.e. VIX) changed. Thus, the copper market is reacting to the change that happened the day before in the US financial markets. The conclusion is that these economic factors are leading (1-day) the US copper futures market.

In the second part of this work, we identify regimes for COMEX copper futures market and apply this information to optimize the best simple trading strategies found in part one. There have only been a few studies investigating regimes in the copper market and only one recent book examining the

COMEX futures with another method. In this study, we find several characteristics of regimes in the copper market by using hidden Markov models under the assumption that the return series of the COMEX Copper futures are generated from a multivariate normal distribution. The parameters of this model reasonably represent the market; the states are clearly separated, and each state has its distinct property. The result showed that the fluctuation behaviors of COMEX copper prices are mainly transited under three states: stable/sideways, unstable/fast drop, and slow rise. The stable state had almost zero return and the lowest volatility, and the slow rise (transition state in our case) state had 20.66% annualized expected return and moderate volatility. Growth of the COMEX copper returns is predominant in the regime defined as the transition state but there is a risk of falling to the unstable/fast dropping state. The unstable state clearly showed the lowest annualized return (-60.48%) and the highest volatility. Smoothed probability and wealth paths show these properties graphically.

Finally, we presented evidence showing that the regime-switching asset allocation strategy significantly improves performance compared with an unconditional static alternative (simple strategies) or compared to the simple buy and hold copper benchmark. For data spanning from 1985 to 2015, a combination of the DJOW indicator (regime 1- stable), the PMI.US indicator (regime 2- fast drop) and the SP500 indicator (regime 3- slow rise) performs best. Theoretically, the Markov strategy would return 7.44, thus 60 percent more in the considered period than the best simple strategy (SP500). However, the max drawdown stays the same in both strategies. The 1985 Markov strategy benefits from an information ratio of 0.54. For data spanning from 1990 to 2015, a combination of the DJOW indicator (regime 1- stable), the PMI.US indicator (regime 2- fast drop) and the VIX indicator (regime 3- slow rise) performs best. Theoretically, the Markov strategy would return 7.00, thus 95 percent more in the considered period than the best simple strategy (VIX). However, the max drawdown is higher with the Markov strategy than with the VIX simple strategy. The 1990 Markov strategy benefits from an information ratio of 0.68.

If we compare the two Markov strategies, the second one is superior both in terms of annualized return and in terms of risk. Indeed, the Markov90 strategy benefits from a lower annualized standard deviation, a lower maximum drawdown, a lower MVaR and a lower MCVaR. This is underlined by a higher annualized Sharpe ratio and a higher Information ratio. Overall, these two strategies outperform the copper benchmark and other simple strategies, and both do especially outstanding during the crash period in 2008. However, both loose during the 1993-1996 period as well as during the year 2006. We can conclude that the regime information helps strategies avoid risk during left-tail events.

Lastly, this regime approach results in trading a collection of “more specialized” strategies (read: less robust), switching between them as the market switches between its various regimes (Liberty, 2010).

Thinking that a regime approach can be more successful than a “one-size-fits-all”, more robust trading strategy (simple trading strategy) makes a strong assumption about regime persistence, specifically that the regime stays “longer than the time it takes to identify the shift to the new regime” (Liberty, 2010). In essence, if the regime switch identification lag were too long, one would probably end up using the incorrect specialized trading strategy to the present regime and suffer from under-performance.

As a further research, it would be useful to investigate whether the Markov probability can detect at the right time the next regime. In order to address this issue and to backtest the detection ability of the Markov model, different methods can be used, including using 10 years of data to calibrate the model and then try to predict the next regime, calibrate the model with an expanding window and calibrate the model with a rolling window. Another point that could be further researched is the evaluation of combinations of different simple strategies. Maybe a combination of 2 simple strategies for each regime of the Markov model could produce better results. Moreover, the maximum drawdown would be probably reduced as well.

BIBLIOGRAPHY

- Bae, G. I. (2012). What are the regimes in the commodity market?
- Bae, G. I., Kim, W. C., & Mulvey, J. M. (2013). Dynamic asset allocation for varied financial markets under regime switching framework.
- Balduzzi, P. E. (2001). Economic News and Bond Prices: Evidence from the U.S. Treasury Market. *The Journal of Financial and Quantitative Analysis*, 36 (4).
- Barnes, R. (2014). *Economic Indicators: Purchasing Managers Index (PMI)*.
- Brunnermeier, M., Nagel, S., & Pederson, L. (2009). Carry Trades and Currency Crashes. *NBER Macroeconomics Annual*, 23, 313-47.
- Buncic, D., & Moretto, C. (2014). *Forecasting Copper Prices with Dynamic Averaging and Selection Models*.
- Burgering, C. (2014). *Copper price & economy*.
- Cheng, H., Shao, L., & Guo, Y. (2013). *State Transition Behaviors of SHFE Copper Prices Based on Markov-switching Model*.
- Chevallier, J., & Ielpo, F. (2013). *The Economics of Commodity Markets*. (J. W. Sons, Éd.) Wiley Finance.
- Crittenden, E. (2010). *London Metals Exchange data & backtesting*.
- Donninger, C. (2012). *The Identification of Market-Regimes with a Hidden-Markov Model*.
- Fraser, A. (2008). *Hidden Markov models and dynamical systems*. Society for Industrial and Applied Mathematics.
- Guegan, D. I. (2009). Further Evidence on the Impact of Economic News on Interest Rates. *Frontiers in Finance and Economics*, 6 (2).
- Kuepper, J. (2014). *Backtesting: Interpreting The Past*. Récupéré sur Investopedia: <http://www.investopedia.com/articles/trading/05/030205.asp>
- Lia, X., & Zhanga, B. (2009). *Price discovery for copper futures in informationally linked markets*.
- Liberty, J. (2010). *Trading Regimes as Strategy Filters*.
- Mulvey, J., & Zhao, Y. (2010). *An Investment Model Via Regime-Switching Economic Indicators*.
- OECD. (2015). *OECD Composite Leading Indicators: Reference Turning Points and Component Series*. Récupéré sur <http://www.oecd.org/std/leading-indicators/oecdcompositeleadingindicatorsreferenceturningpointsandcomponentseries.htm>
- Rabobank. (2011). *Rabobank flags Vix-aid as crop sell-off continues*. Récupéré sur <http://www.agrimoney.com/news/rabobank-flags-vix-aid-as-crop-sell-off-continues--3639.html>
- Roache, S. R. (2009). The Effects of Economic News on Commodity Prices: Is Gold Just Another Commodity? . *IMF Working Papers*.