DO WTI OIL PRICE MOVEMENTS FOLLOW A MARKOV CHAIN?

By Richard R. Conn CMA, MBA, CPA, ABV, ERP

This paper is the first in a two-part series which is, ultimately, dedicated to outlining a methodology for predicting future oil prices using Markov properties. Far from being unique or innovative, there have been many excellent academic works on this subject already written (see References). However, the difficulty with academic papers is that they are typically written for the academic community and not the average practitioner. This paper is squarely aimed at providing the typical financial analyst (who already has a rudimentary understanding of matrix algebra and the basic qualities of Markov probabilities) with an overview of creating a WTI (West Texas Intermediate) Monte Carlo-styled oil price prediction model.

The primary difference between the approach taken here compared with the academic papers is that the focus of this series will be upon forecasting the future states of the outcomes (i.e. directionally the steps involved in getting from the current WTI level to some future prediction of WTI) compared with estimating the future price level of WTI. Most of the academic papers are intent upon estimating whether future oil price will be higher or lower than current prices and, if so, but how much. We, instead, will be interested in predicting the random walk that WTI is likely to take between now and the long-term foreseeable future rather than upon attempting to estimate the size of each of those steps (i.e. the magnitude of the price change at each point in time). Only after we have devised a path that future oil price movements are likely to follow will we turn our attention to the question of estimating how large each of those price changes reasonably can be expected to be. The practicality of this approach will become evident as the methodology is unveiled.

AN OVERVIEW OF THE MECHANICS

While the explanation will be tedious, the details of what the Markov properties are and why they are important to the movements of oil prices is unavoidable. At the same time, as little space as possible will be dedicated towards the underlying theory – this series is dedicated towards coming up with a practical and working model.

Markov Chains are probabilistic models/observations about movements within a state space. Considering the movement of oil prices (and we shall only concern ourselves with the movement of WTI) there are **only three states**. At the end of any given day, for example, the WTI price could have moved up, moved down or stayed the same, relative to the closing price of the previous day. This latter state would be described as a 'loop' – the price looped back to where it was such that the new 'state' is the same as the old.

FIGURE 1

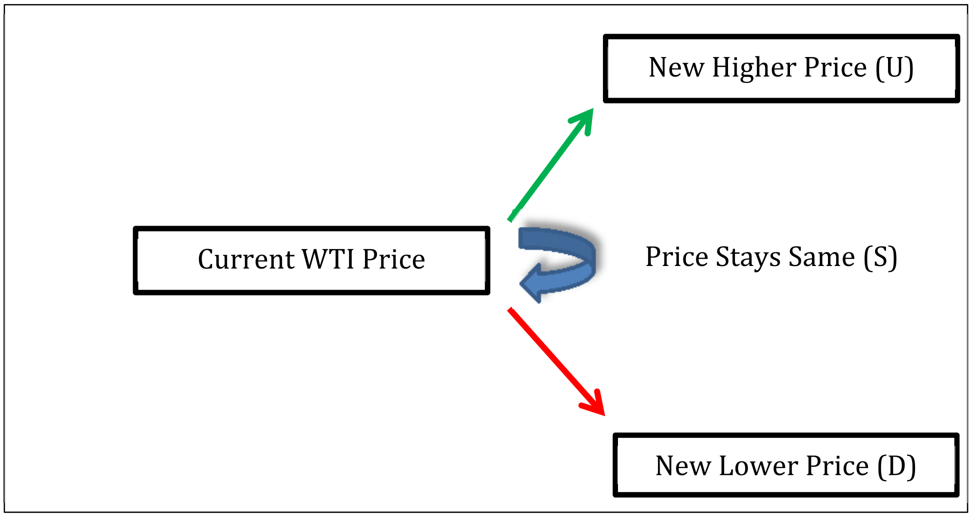

Such an observation should lead the analytically curious to the question: 'With what frequency have WTI prices increased in the past relative to how often they have declined or remained the same?' Being able to answer this question may lead to being able to make better predictions about the future path of WTI prices. If we abbreviate such movements as D,S,U for Down, Same, Up respectively, then we have a shorthand means of quickly symbolizing long strings or chains of price movements. If, for example, we are describing a sequence of price changes over six days (i.e. five movements) as DDSUD, then, that means the Second day ending price was Down from the first; the Third was Down from the second; the Fourth was the Same as the Third; the Fifth was Up from the Fourth; and finally the Six was Down from the Fifth.

The primary characteristic of Markovian movements between the various states is that only the current position within the state space is in any way relevant to the next movement. In the example above, if we are attempting to predict where WTI will be at the end of Day Seven (either D,S, or U relative to Day Six), the only information that is useful in making that prediction is that Day Six ended as a "D" day. The fact that the previous four movements prior to Day Six were DDSU in no way improves our statistical chances of correctly estimating how Day Seven will end. Another way to say this is that XXXXDD (where Day Seven ends as a 'Down' Day, for example) has just as much probability of occurring as DDSUDD or UUUUDD or USDSDD or any permutation where XXXX represents any random string of "D","S","U" (and the mathematically astute will recognize that there are 3^4 = 81 permutations possible here). In other words the history or path that WTI prices took before arriving at the D of Day Six has no bearing upon the direction that prices will take on Day Seven. Another way of saying this is that the price movement on Day Seven is completely independent of the direction of the price movements at any time prior to Day Six.

Why should we care if WTI prices exhibit Markovian characteristics? Because, if WTI prices are independent of the path upon they took to get to the current state AND we can identify or determine a probability for which state they might next move to, then this will quite easily allow us to build a simple probabilistic model for predicting where WTI prices may be going. Conversely, if

WTI prices are path dependent, then the level of complexity involved in calculating where they can next be expected to move to increases exponentially. Imagine, for example, that it is now Day Six which has ended as a D Day and the probability of where price may end tomorrow is dependent upon the path prices took in the previous five days (i.e. the four price movements prior to the "D" of Day Six). That would mean that each of the possible 81 different price paths that may have occurred in the last five days could have its own unique probability in determining where price moves to on Day Seven. And now imagine how complex it may be to predict price movements if tomorrow's price were in some way dependent upon which price path had been taken over the prior 50 or 100 or 1,000 trading days? That would be 3^49 or 3^99 or 3^999 possible permutations respectively – numbers too large to comprehend, much less assign probabilities to.

So we want WTI price movements to possess Markovian characteristics, but we cannot simply assume that they do. Some formal test procedure needs to be applied to empirical WTI historic data in order to assess if those movements have followed a Markov chain (i.e. are path independent)

AN INTUITIVE APPROACH

Even before considering a more rigorous approach that involves the use of matrix algebra and the Chi-square test, there is an intuitive approach that should allow us to set a lower-boundary about the price-path dependency of historic WTI price movements. While this informal method will not allow us to absolutely conclude that WTI price movements do follow a Markov chain, it will alert us in the event that they clearly do not.

If WTI price movements are Markovian and price-path independent, then it is only the current state that has any bearing upon where price will move next. This means that, selecting a sub-sample of, say DD, price movements from a very large sample of WTI price history we would expect to find the same relative frequency of DD price chains as the occurrence of XDD price chains. That is, given that the probability for prices to decline twice in succession (which necessarily would require a 3 day observation period), we would expect to find the same probability for an XDD price chain to occur over a four day observation period (where the "X" could represent either a Up, Down or Same movement). If this were not the case, then one might begin to suspect that price-path had some bearing upon what course the third price movement took.

Expanding this logic, if post hoc WTI price movements are price-path independent, we would expect XXDD to occur with the same relative frequencies as DD chains. And, similarly, XXXXXXXDD ten day observations should occur with the same frequency and so should one-hundred-day chains ending in DD occur in just the same frequency.[1]

[1] The caveat here is that we are speaking about relative frequency. If, for example, our subsample is 1,001 days long, there are 1,000 chains of length two possible. This is because, on the first day, there is no movement, so it takes two days upon which to observe the first movement. Sequentially, then 1,000 two-period chains could be observed over a 1,001 day subsample. If n is the number of days in the subsample, and k is the length of the chain, then, n-(k-1) = number of chains possible. Therefore, 1,001 – (2 – 1) = 1,000. Similarly, if we want to observe the frequency of length three chains, now there would only be 1,001 – (3 – 1) = 999 chains possible. In that case we may be measuring the ratio of XDD chains that occur over 999 possible

EIA EMPIRICAL EVIDENCE: The U.S. Energy Information Administration (EIA) maintains a large database of historical WTI prices (spot prices, measured at Cushing, OK) and makes that available for easy download in either, daily, weekly or month-end form. That database begins on Jan. 2, 1986 and therefore includes approximately 7,100 trading day observations up to present day (this paper was written in early May, 2014). By drawing our conclusions from a 28 year sample, we are implicitly conditioning our results to be long-termed in nature. Any single-event WTI price shock that has occurred since January 1986 will, to some extent, be reflected in the data – but tempered by the long-term status quo. Similarly, any short or mid-term cycles that the WTI has exhibited will be somewhat mitigated in our results. Had we wanted to construct a mid-term estimator of WTI prices, we would have been faced with the difficulty of deciding which portion of the historic data best reflected the expected characteristics of the forthcoming mid-term.

We will, however, examine the data for short and long-term volatility, with the intent of identifying any permanent transitions that may have occurred over time. But we will not do so until after our observations about frequency has been presented.

outcomes. Similarly, a chain of length 100 (i.e. would require 101 observation days) could only possibly occur 1,001 – (100 – 1) = 902 in this subsample and therefore it would be frequency of the number length 100 chains ending in DD relative to the 902 possible outcomes of a chain this length that would provide the important statistic in this case.

TABLE 1: FREQUENCY OF 'XX' ENDING CHAINS IN EIA DATA

Two Day Chains:

		Day n Observation		
		D	S	U
Day (n + 1) Observation	D	22.11%	0.71%	24.63%
	S	0.83%	0.07%	0.84%
	U	24.50%	0.95%	25.37%

Ten Day Chains:

		Day n Observation		
		D	S	U
Day (n + 1) Observation	D	22.10%	0.71%	24.63%
	S	0.83%	0.07%	0.84%
	U	24.50%	0.95%	25.37%

One-Hundred Day Chains:

		Day n Observation		
		D	S	U
Day (n + 1) Observation	D	22.03%	0.71%	24.67%
	S	0.82%	0.07%	0.84%
	U	24.54%	0.95%	25.37%

One-Thousand Day Chains:

		Day n Observation		
		D	S	U
Day (n + 1) Observation	D	22.49%	0.60%	24.67%
	S	0.65%	0.08%	0.72%
	U	24.61%	0.76%	25.42%

Looking first at the 2 Day Chain data, at the intersection of DD, this tells us given all the possible 2 day chains that could have been observed in the EIA data (there were 7,143) 22.11% of the time a Down movement on Day n was followed by a subsequent Down movement (i.e. DD) on Day n + 1. Similarly, 25.37% of the time an Up movement was followed, on the next day, by another Up movement (i.e. UU).

Moving to the data representing the 1,000 day chains we find that, of all the possible 1,000 day chains that could have been observed in that data (of which there were 7,144 – (1000 – 1) = 6,145 unique observations), we see very similar statistics as the in 2 Day observations. There is roughly a 22.5% chance of observing a 1,000 day chain that ends in DD. There is only a sixth-tenths of one-percent chance of observing a 1,000 day chain that ends in SD. And, there is approximately a 25.4% chance of a 1,000 day chain ending in UU. These are very similar numbers to the observations in the 2 day chains which are, in turn, very close to those statistics observed with the 10 day chains and the 100 day chains. Had we observed a dramatically changing statistic as the chains got longer, we might have begun to suspect that the WTI moves in a price-path dependent manner. But the daily outcomes[2] have remained remarkably consistent regardless of the length of the chain, therefore it is appropriate to do some more formal testing.

MARKOV PROBABILITIES IN MATRIX FORM

It is convenient to present Markov single-stage data in Matrix form. For ease of exposition, we will speak of the daily price movements when explaining the methodology, but later the weekly observations will also be presented and compared to the daily. There are two presentation alternatives possible (one is the transposition of the other) and a few words on Matrix protocol are probably in order.

FIGURE 2

		Day n Observations		
		Down	Same	Up
Day (n +1) Observations	Down	DD	SD	UD
	Same	DS	SS	US
	Up	DU	SU	UU
		100%	100%	100%

		Day (n + 1) Observations			
		Down	Same	Up	
Day n Observations	Down	DD	DS	DU	100%
	Same	SD	SS	SU	100%
	Up	UD	US	UU	100%

The difference between the two presentations above deals with whether we are identifying the 'Day n' observations as rows and the 'Day n+1' observations as columns, or the other way around. The convention usually is the former, 'Day n' are rows, 'Day n + 1' are columns and, as a result, the total of each row sums to 100%. The difference between the two is a matter of personal preference, but it is important to keep in mind that the statistic we are interested in is that *given that we are at current state X as at 'Day n'*, we wish to know the probability of moving to state X on 'Day n + 1'. For example, given that we have observed a 'D State' on 'Day n' we are interested in learning the

[2] For reasons of space the weekly statistics will not be presented here by reserved for Appendix 1. Similar to the daily stats, there is a good deal of consistency moving between the lengths of chains. One notable distinction between the Daily and Weekly observations was that it was noticeably more likely for a two-week chain to finish on a UU (about 30% of the time), than it was for a two-day chain to do so (only about 25% of the time).

frequency with which this transitions to a 'DU' chain, for example, on Day n + 1. Therefore, the sum of all the 'DX' chains in that subsample total 100% (which, by definition, they must).

It would be a mistake to obtain the test statistic in the wrong order. For example, to determine the frequency of 'XD'. In words this would be akin to asking ourselves, given the fact that 'Day n' has ended in the 'D State' what is the probability of observing the 'X State' on 'Day n – 1'? Such an observation would not be in accordance with Markovian characteristics.

For reasons of personal preference, this paper will use the unconventional Day n Observations as columns (ergo, the Columns will sum to 100%), and apologizes are offered to those readers more accustomed to the more frequent form where the rows sum to 100%.

MARKOV PROBABILITIES OBSERVED IN WTI PRICE MOVEMENTS SINCE 1986

The following table shows the probability of observing an X price movement on Day n + 1 given the occurrence of an X movement on the preceding day:

TABLE 2: MARKOV PROBABILITIES IN WTI

2 Day Chain[3]

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.61%	41.13%	48.44%
	S	1.74%	4.03%	1.65%
	U	51.65%	54.84%	49.90%
		100.00%	100.00%	100.00%

10 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.60%	41.13%	48.44%
	S	1.74%	4.03%	1.65%
	U	51.66%	54.84%	49.90%
		100.00%	100.00%	100.00%

100 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.48%	40.98%	48.49%
	S	1.74%	4.10%	1.65%
	U	51.78%	54.92%	49.86%
		100.00%	100.00%	100.00%

1000 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	47.10%	41.57%	48.56%
	S	1.36%	5.62%	1.41%
	U	51.53%	52.81%	50.03%
		100.00%	100.00%	100.00%

[3] Some licence is taken with the naming convention here. Recognize that an observation of two periods (e.g. SU) requires three separate days because it can only be known what the 'n' day movement is at the end of the second day. Similarly, the n + 1 day observation can only be completed at end of day three. Therefore, when we speak of a Two Day Chain, we really mean three days have elapsed. Further, for the 10-Day Chain, Day 'n' is the observation taken on Day 10 (the 9th price movement) and Day (n + 1) is the observation taken on Day 11. Using post hoc data looking backwards, we are always interested in the XX movements *at the end of the chain*. This will not be the case when we use Markov probabilities to predict future price movements.

Similar to our observations regarding the simple frequencies table, we note that the Markov probabilities also display a surprising consistency regardless of how long the price movement chain is. It is not as intuitively obvious, in this case whether this consistency is a good thing or not.

However, the statistical interpretation of outcomes is similar. Looking at the 2 Day Chain data, for example, of all the occasions where an occurrence of DX was observed, 46.60% of those times resulted in a DD observation, 1.74% of the time a DS was observed and 51.65% of the times a DU event had occurred.

Before proceeding on with a more in depth description of Markov probabilities and how they might be tested, it will be useful to digress to a comparison of the Daily vs. Weekly Markov Matrices as this information will become relevant later in the discussion. The following is simply a repeat of Table 2 as above, except now the week-end closing price data has been juxtaposed (on the right) alongside the daily data:

TABLE 3: DAILY VS WEEKLY MARKOV PROBABILITIES

Daily WTI Price Chain Data

2 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.61%	41.13%	48.44%
	S	1.74%	4.03%	1.65%
	U	51.65%	54.84%	49.90%
		100.00%	100.00%	100.00%

10 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.60%	41.13%	48.44%
	S	1.74%	4.03%	1.65%
	U	51.66%	54.84%	49.90%
		100.00%	100.00%	100.00%

100 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.48%	40.98%	48.49%
	S	1.74%	4.10%	1.65%
	U	51.78%	54.92%	49.86%
		100.00%	100.00%	100.00%

1000 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	47.10%	41.57%	48.56%
	S	1.36%	5.62%	1.41%
	U	51.53%	52.81%	50.03%
		100.00%	100.00%	100.00%

Weekly (i.e. week-ending) WTI Price Chain Data

2 Week Chain

		Weekly n Observation		
		D	S	U
Weekly (n +1) Observations	D	51.66%	50.00%	42.51%
	S	0.58%	0.00%	0.26%
	U	47.76%	50.00%	57.23%
		100.00%	100.00%	100.00%

10 Week Chain

		Weekly n Observation		
		D	S	U
Weekly (n +1) Observations	D	51.17%	50.00%	42.44%
	S	0.59%	0.00%	0.26%
	U	48.25%	50.00%	57.31%
		100.00%	100.00%	100.00%

100 Week Chain

		Weekly n Observation		
		D	S	U
Weekly (n +1) Observations	D	51.55%	50.00%	42.47%
	S	0.62%	0.00%	0.27%
	U	47.83%	50.00%	57.26%
		100.00%	100.00%	100.00%

1000 Week Chain

		Weekly n Observation		
		D	S	U
Weekly (n +1) Observations	D	51.14%	0.00%	41.54%
	S	0.46%	0.00%	0.00%
	U	48.40%	100.00%	58.46%
		100.00%	100.00%	100.00%

Salient observations in comparing the daily vs. weekly probabilities are:

1. Given that a particular week has ended in a U state, it is about 8% more likely that it will be followed by another U week-end than is it the case that a UU occurrence would be observed in the daily data.
2. The corollary is that, given that a week has ended in a D state, it is marginally less likely that it will be followed by a U (DU's only occur approximately 48% of the time in the weekly

data) than would be the case in the daily (where DU's occur approximately 52% of the time).

The other readily apparent observation we can make about the Weekly Data, is that this much less stability in the occurrence of the SX chains. This stands to reason because the number of occasions upon which the first week will end at exactly the same price as the starting observation will be quite rare.

THE MARKOV TRANSITION MATRIX

If we accept that only the current state is an indicator of where the next state is likely to be (and, therefore, all the previous price movements prior to the current state have no bearing in the matter), then there is a single matrix that will encapsulate all that can be known about probable price movements. This is called the Transition Matrix and, in the case of the daily WTI price data is represented by:

TABLE 4: MARKOV TRANSITION MATRIX

2 Day Chain

		Day n Observation		
		D	S	U
Day (n +1) Observations	D	46.61%	41.13%	48.44%
	S	1.74%	4.03%	1.65%
	U	51.65%	54.84%	49.90%
		100.00%	100.00%	100.00%

This tells us that, given that we have already observed a D movement today (on Day n), there is a 46.61% probability of another D movement happening tomorrow (on Day (n + 1)). But does it give us any predictive insight into what is likely to happen the day after tomorrow (on Day (n + 2))?

For example, given that today a D movement has already occurred, what is the probability that another D movement will occur tomorrow followed by another D movement the day after that? Well, today we know there is a 46.61% probability that a subsequent D will occur on Day (n + 1) and then tomorrow, given that a D movement had already happened on that day, the probability of another D (on Day (n +2)) is 46.61% again. So, the probability of a DDD is 46.61%^2 = 21.7249%

But what if we were attempting to predict the more relevant instance where we know today has ended with a D movement (for example), but we wish to predict if Day (n + 2) would end with a D movement? Symbolically, this can be represented by the chain DXD. The probability of observing a D on Day (n + 2) given that a D has occurred on Day n must be the sum of the probabilities: DDD + DSD + DUD. These are the only three paths beginning with the current state of D that will allow us to finish with a D on Day (n + 2). Using the 2 Day probabilities from the Transition Matrix, we can calculate this probability as:

DDD =	(DD) x (DD)	= 46.61% x 46.61%	= 21.7249%
DSD =	(DS) x (SD)	= 1.74% x 41.13%	= 0.7157%
DUD =	(DU) x (UD)	= 51.65% x 48.44%	= 25.0193%
	PROBABILITY OF DXD		= 47.46%[4]

Markov chains have this somewhat counter-intuitive and ironic quality: past movements have no predictive power – knowing the price path upon which today's state came into being is of no practical assistance in estimating what tomorrow's price movement will be – only today's state has any bearing upon tomorrow. And yet, when estimating further forward, say to the day after tomorrow, then the price path does become relevant. This is because tomorrow's state is dependent upon what state we are at today and so too will be the state on Day (n + 2) because we are attempting to predict it prior to knowing with absolute certainty how tomorrow will end. This is the 'chain-like' feature of Markov chains.

While we could continue to calculate all the possible paths remaining as we did above for DXD, namely: DXS, DXU, SXD, SXS, SXU, UXD, UXS and UXU, it will be easier at this point just to note that, because of the marvels of matrix algebra, all these outcomes will be provided simply by raising the Transition Matrix to the power of two:

FIGURE 3 : TRANSITION MATRIX RAISED TO 2ND POWER

	Matrix A					Matrix A^2		
	D	S	U			D	S	U
D	46.61%	41.13%	48.44%	^2		47.46%	47.39%	47.43%
S	1.74%	4.03%	1.65%		=	1.73%	1.78%	1.73%
U	51.65%	54.84%	49.90%			50.81%	50.82%	50.83%

Therefore, if, on Day n we are in State S, and we are attempting to estimate what the probability is that Day (n + 2) will end also in State S, Matrix A^2 tells us there is a 1.78% chance of that occurring. And, it should be readily apparent that estimating the outcome on Day (n + 3) is represented in Matrix A^3 and that the probable outcomes on Day (n + 100) are captured in Matrix A^{100} , etc. Therefore, if we wish to know what the probability is that WTI will close either in an D, S or U state 1,000 days hence, we only need an inexpensive matrix calculator to key in Matrix A as above and raise this to the power of 1,000. Could any approach be simpler? In this case, the answer is YES, because, as we will next discuss, the process very quickly approaches a steady state.

STEADY STATE LIMIT VALUE OF TRANSITION MATRIX

Ignoring all the reasoning as to why this is true[5], it turns out that, if the Transition Matrix is 'regular' (meaning all its elements are positive – and, of course, the columns all sum to 1), then raising the

[4] Those with a greater familiarity of Matrix Algebra will recognize this as the dot product of the first column and the transposition of the first row.

matrix to increasingly larger powers will eventually bring it to a 'steady state' wherein all the elements approach a limiting value (referred to as an asymptotic quality). As it happens, the WTI transition matrix reaches its limit very quickly (in less than 5 days), as can be seen below:

[5] The steady-state vector can be found by first finding the eigenvector resulting from the eigenvalue +1. This eigenvector will be a multiple of the steady-state vector which can be found by scaling its components such that they sum to 1.

TABLE 5: TRANISITION MATRIX RAISED TO INCREASING POWERS

Matrix A

	D	S	U
D	46.61%	41.13%	48.44%
S	1.74%	4.03%	1.65%
U	51.65%	54.84%	49.90%

Matrix A^2

	D	S	U
D	47.46%	47.39%	47.43%
S	1.73%	1.78%	1.73%
U	50.81%	50.82%	50.83%

Matrix A^3

	D	S	U
D	47.44%	47.44%	47.45%
S	1.74%	1.74%	1.74%
U	50.82%	50.82%	50.82%

Matrix A^4

	D	S	U
D	47.44%	47.44%	47.44%
S	1.74%	1.74%	1.74%
U	50.82%	50.82%	50.82%

Matrix A^{10}

	D	S	U
D	47.44%	47.44%	47.44%
S	1.74%	1.74%	1.74%
U	50.82%	50.82%	50.82%

Matrix A^{100}

	D	S	U
D	47.44%	47.44%	47.44%
S	1.74%	1.74%	1.74%
U	50.82%	50.82%	50.82%

Matrix A^{1000}

	D	S	U
D	47.44%	47.44%	47.44%
S	1.74%	1.74%	1.74%
U	50.82%	50.82%	50.82%

It can be seen in the above that the WTI steady state is achieved by Day (n + 4) – the probabilities in Matrix A raised to the 4th power are exactly the same when raised to the 10th, 100th or 1,000th power. Note as well that the probabilities in each of the columns duplicate each other – this means that, by the fourth day, it does not matter what the initial state was – the outcome probabilities are the same regardless of where one started. Therefore, for long-term estimation, one need not take into account the starting state, the probability of experiencing a D, S or U on any given day in the future remains static, namely; D = 47.44%, S = 1.74%, and U = 50.82%[6]

Similarly, for the Weekly Data, the Transition Matrix can be seen in Table 3 above (the Two-Week matrix) and this matrix is also quickly asymptotic, reaching a steady-state after being raised to the power of seven.

TABLE 6: WEEKLY STEADY STATE PROBABILITIES

Weekly Data - Steady State	D	S	U
D	46.827%	46.827%	46.827%
S	0.406%	0.406%	0.406%
U	52.767%	52.767%	52.767%
	100.00%	100.00%	100.00%

CHI-SQUARE TESTS

At last we have sufficient background to attend to the initial question: "Do WTI oil price movements follow a Markov Chain?" The Chi-Square statistic is a goodness-of-fit test that we can use to ask the question; 'How likely is it, based upon random variation alone, that the *observed* WTI outcomes would differ as much as they do from the *expected* outcomes – assuming that WTI price does move in a Markov chain?'

The Chi-square statistic is:

$$\chi^2 = \sum\left[\frac{(observed\ count - expected\ count)^2}{expected\ count}\right]$$

However, one cannot use 'scaled' or proportionate inputs for the chi-square calculation – the percentage data that has been so convenient to use in matrix form up to this point will give

[6] Some may point out that these steady-state probabilities derived from the Transition Matrix are simply the sum of the rows in the very first 2-day frequency statistics from Table 1. D = 47.45% = (22.11% + 0.71% + 24.63%); S = 1.74% = (0.83% + 0.07% + 0.84%); U = 50.82% = (24.50% + 0.95% + 25.37%). This is often the case with more sophisticated measurement functions – they only serve to confirm what we intuitively would have suspected to be the truth but without statistical proof. With only the relative frequencies, for example, we did not know that the Markov probabilities would converge to a steady-state limit or how quickly this may happen.

erroneous chi-square statistics.[7] We must actually use the count data, which, for example, for the ten day-chain, was:

TABLE 7

10 Day-Chain Observations

	D	S	U
D	1,577	51	1,757
S	59	5	60
U	1,748	68	1,810
	3,384	124	3,627

The ten-day expected matrix data is a function of the A^{10} matrix in Table 5 – that is, if the WTI prices do follow a Markov chain, we would expect the possibility of a D state on day 10 to be 47.44% (regardless of the starting state), an S outcome would occur 1.74% of the time and, finally an U outcome has the probability of 50.82%. This information, combined with the fact that we know there were 7,135 10-day chains in the historic data (simply the sum of the column totals from Table 6) leads us to the conclusion that the 10-Day *Expected* Matrix would be:

TABLE 8

10-Day Expected Outcome

	D	S	U
D	1,606	59	1,721
S	59	2	63
U	1,720	63	1,843
	3,385	124	3,627

The resulting chi-square statistic from the comparison of these two matrices is 8.44. Applicable 'Degrees of Freedom 'for this test is 8.[8] We will assume a .05 or 5% alpha significance. Our null hypothesis is that the WTI prices *do* follow a Markov chain and therefore any variance we find between the actual observed counts and the hypothetically expected are just the result of random variation (i.e. chance). The chi-square statistic calculates how probable it would be to encounter the calculated variance just based upon random chance alone. If we find that there is less than a 5% chance that such a chi-square outcome could be observed as a result of random chance, then we must reject the null hypothesis and conclude that WTI prices *do not* follow a Markov chain.

Consulting a Chi-square table quickly shows that, at 8 df and $\chi^2 = 8.44$ the statistic is well above the 25% probability mark (any reliable statistics software package will report that the actual probability is 39.1%). Therefore, based upon the 10-day chain alone, we would not reject the conclusion that WTI prices follow a Markov chain.

[7] Specifically, all the *expected* counts must be > than one – and by definition a probability % will not qualify.
[8] There is often considerable confusion about Degrees of Freedom (df) for the chi-square tests. The issue depends upon whether a single-sample test of independence is being performed (in which case, the relevant df would be (rows – 1)(columns – 1) and would amount to 4 for the above. However, a goodness-of-fit test is a comparison between two samples: the actual observed and the hypothetically expected, therefore the df is (k – 1) where k is the number of cells in the two-way table.

Similarly, the χ^2 was calculated for all the 10-Day, 100-Day and 1,000-Day and 10-Week, 100-Week and 1,000-Week chains. In no case was there reason to reject the null hypothesis. While we can reasonably conclude that WTI prices do follow a Markov chain[9], we are more interested in the results of the longer 1,000-period chains as we are attempting to model long-term, rather than short term WTI prices.

[9] To be precise, we have no evidence to reject the assumption that WTI follows a Markov chain. This is decidedly different than having proof positive that WTI *does* follow a Markov chain.

TABLE 9: CHI-SQUARE STATISTICS

Chi-square results (8 df)

CHAIN LENGTH	χ2	α Prob.	Reject Ho?
10-Day Chains	8.45	39.1%	No
100-Day Chains	8.89	35.2%	No
1000-Day Chains	11.00	20.2%	No
10-Week Chains	12.68	12.3%	No
100-Week Chains	12.90	11.5%	No
1000-Week Chains	6.78	56.1%	No

TIME HOMOGENEITY

One of the requirements of Markov probabilities is the assumption of time homogeneity. That is, the probability of moving from a given state to the next state is assumed to remain constant over time. In the case of WTI price movements, we can speculate that this is a requirement that is not strictly adhered to. We know for certain, for example, that WTI prices exhibits periods of greater and lesser volatility (see Figure 4). These periods may also be associated with varying probabilities between D, S, and U states. We might also speculate that, during times of economic boon, the frequency of U movements increases and, correspondingly, the D frequency may increase during periods of economic recession. The pertinent question becomes: 'How might this infraction of the time homogeneity requirement impair our ability to use a Markov probability model to predict future price steps?

The answer to this question is beyond the scope of this paper and is certainly deserving of a paper all on its own. Suffice it to say that, on a year-over-year basis, there is considerable variation amongst the state space probabilities. Complete details of all 28 individual years are provided in Appendix 2. To highlight the daily extremes, however, in some years the DD chain represented as much as 54.5% of all DX movements whereas in others it dropped to only 36.8%. In peak years the UU chain occurred 61.7% of all UX movements and this dropped to as low as 38.8% in other years.

In spite of demonstrating lower overall price volatility, the weekly chain data demonstrated an even greater variance. In some years the DD chain represented 71% of all possible DX outcomes whereas in others it dropped to 23.5%. At the other extreme, UU chains ranged between a high of 71% of all UX outcomes to a low of 36.8%.

If the relevant period upon which to measure time homogeneity is one-year in length, then we can definitely conclude that the WTI price movements are compromised in this characteristic. WTI price movements do not adhere to a homogenous probability matrix through time.

FIGURE 4[10]

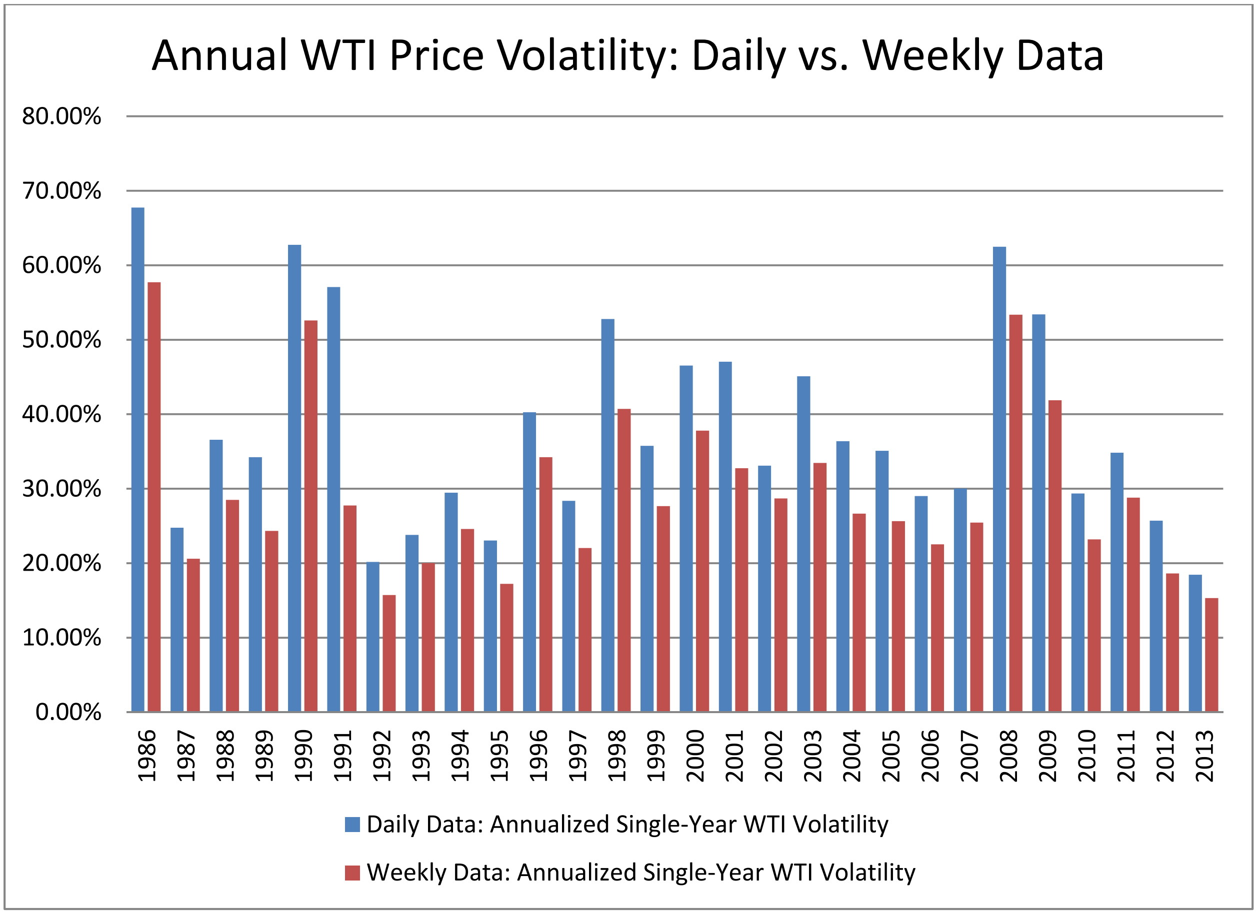

CONCLUSIONS

This paper has discussed the primary characteristics of Markov chains and how they may apply to the movement of WTI prices between price "states" (Down, Same and Up). The most important Markov property is that only the current state has any predictive bearing upon the next state – all previous price movements before the current state are irrelevant in the prediction of where WTI price may go next. We tested for the existence of a Markov chain using simple statistics as well as via a formal Chi-square approach and found, in all cases, that the WTI price movements do strongly exhibit this quality of a Markov chain. Further, we identified the transition matrix for both the daily and weekly Markov probabilities and found that, in both cases, the asymptotic limits are quickly achieved and a steady state is realized only after a few future periods. A confounding complexity, however, is that the post hoc WTI historical data does not demonstrate time homogeneity. In fact, there is considerable variation in the single year Markov probabilities over time and it remains to be seen just how such a defect will impair our ability to reasonably predict future WTI states.

[10] The weekly volatility parallels the ebbs and flows of daily data, but is expectedly less in each year. This is because the daily fluctuations adds 'noise' to the volatility calculation.

We have not yet begun to consider how large the size of each step between the various states should be (i.e. the dollar amount of a U or D move). This topic is left to be investigated in the following paper.

We have discussed the conceptual issues surrounding Markov probabilities and matrix algebra sufficiently that we should now be able to apply the steady-state probabilities in an actual Monte Carlo model designed to predict future states of the WTI price movements. This will be the subject of the next and concluding paper in this series.

APPENDIX 1

WEEKLY DATA: Frequency of 'XX' Ending Chains in EIA Data

Two-Week Chains

		Week n Observation		
		D	S	U
Week (n + 1) Observation	D	24.15%	0.20%	22.46%
	S	0.27%	0.00%	0.14%
	U	22.33%	0.20%	30.24%

100% *

Ten-Week Chains

		Week n Observation		
		D	S	U
Week (n + 1) Observation	D	23.81%	0.20%	22.52%
	S	0.27%	0.00%	0.14%
	U	22.45%	0.20%	30.41%

100% *

Ten-Week Chains

		Week n Observation		
		D	S	U
Week (n + 1) Observation	D	24.06%	0.22%	22.46%
	S	0.29%	0.00%	0.14%
	U	22.32%	0.22%	30.29%

100% *

Ten-Week Chains

		Week n Observation		
		D	S	U
Week (n + 1) Observation	D	23.33%	0.00%	22.50%
	S	0.21%	0.00%	0.00%
	U	22.08%	0.21%	31.67%

100% *

* The sum of all 9 observations total 100% as expected

APPENDIX 2

APPENDIX 2: ANNUAL MARKOV PROBABILITIES 1986 THROUGH 2013

DAILY DATA / **WEEKLY DATA**

YEAR				YEAR			
1986	D	S	U	1986	D	S	U
D	49.57%	33.33%	45.16%	D	56.00%	0.00%	42.31%
S	3.42%	0.00%	4.03%	S	0.00%	0.00%	0.00%
U	47.01%	66.67%	50.81%	U	44.00%	0.00%	57.69%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1987	D	S	U	1987	D	S	U
D	50.00%	50.00%	44.26%	D	53.85%	0.00%	46.15%
S	5.00%	0.00%	4.92%	S	0.00%	0.00%	0.00%
U	45.00%	50.00%	50.82%	U	46.15%	0.00%	53.85%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1988	D	S	U	1988	D	S	U
D	36.84%	50.00%	50.37%	D	45.83%	0.00%	44.83%
S	4.39%	0.00%	0.73%	S	0.00%	0.00%	0.00%
U	58.77%	50.00%	48.91%	U	54.17%	0.00%	55.17%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1989	D	S	U	1989	D	S	U
D	38.26%	25.00%	51.49%	D	42.86%	0.00%	38.71%
S	3.48%	0.00%	2.99%	S	0.00%	0.00%	0.00%
U	58.26%	75.00%	45.52%	U	57.14%	0.00%	61.29%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1990	D	S	U	1990	D	S	U
D	50.39%	40.00%	51.22%	D	46.15%	0.00%	57.69%
S	2.33%	0.00%	1.63%	S	0.00%	0.00%	0.00%
U	47.29%	60.00%	47.15%	U	53.85%	0.00%	42.31%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1991	D	S	U	1991	D	S	U
D	43.44%	50.00%	50.78%	D	53.85%	0.00%	42.31%
S	3.28%	0.00%	1.56%	S	0.00%	0.00%	0.00%
U	53.28%	50.00%	47.66%	U	46.15%	0.00%	57.69%

APPENDIX 2: ANNUAL MARKOV PROBABILITIES 1986 THROUGH 2013

	DAILY DATA				WEEKLY DATA		
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1992	D	S	U	**1992**	D	S	U
D	45.24%	54.55%	53.33%	D	52.17%	50.00%	40.74%
S	2.38%	0.00%	6.67%	S	4.35%	0.00%	3.70%
U	52.38%	45.46%	40.00%	U	43.48%	50.00%	55.56%
	100.00%	100.00%	100.00%		100.00%	100.00%	100.00%
1993	D	S	U	**1993**	D	S	U
D	50.00%	14.29%	61.68%	D	61.77%	0.00%	63.16%
S	2.94%	14.29%	1.87%	S	0.00%	0.00%	0.00%
U	47.06%	71.43%	36.45%	U	38.24%	0.00%	36.84%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1994	D	S	U	**1994**	D	S	U
D	51.24%	0.00%	46.51%	D	31.58%	0.00%	42.42%
S	0.83%	0.00%	0.78%	S	0.00%	0.00%	0.00%
U	47.93%	100.00%	52.71%	U	68.42%	0.00%	57.58%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1995	D	S	U	**1995**	D	S	U
D	44.14%	33.33%	44.03%	D	45.83%	0.00%	42.86%
S	1.80%	0.00%	2.99%	S	0.00%	0.00%	0.00%
U	54.05%	66.67%	52.99%	U	54.17%	0.00%	57.14%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1996	D	S	U	**1996**	D	S	U
D	42.59%	33.33%	43.57%	D	42.86%	0.00%	38.71%
S	1.85%	16.67%	2.14%	S	0.00%	0.00%	0.00%
U	55.56%	50.00%	54.29%	U	57.14%	0.00%	61.29%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%
1997	D	S	U	**1997**	D	S	U
D	51.88%	46.15%	54.72%	D	70.97%	50.00%	47.37%
S	3.76%	15.39%	5.66%	S	6.45%	0.00%	0.00%
U	44.36%	38.46%	39.62%	U	22.58%	50.00%	52.63%
	100.00%	100.00%	100.00%		100.00%	100.00%	100.00%

APPENDIX 2: ANNUAL MARKOV PROBABILITIES 1986 THROUGH 2013

DAILY DATA / **WEEKLY DATA**

Daily data:

1998	D	S	U
D	54.48%	66.67%	50.88%
S	1.49%	0.00%	0.88%
U	44.03%	33.33%	48.25%
	100.00%	100.00%	100.00%

Weekly data:

1998	D	S	U
D	67.65%	0.00%	55.56%
S	0.00%	0.00%	0.00%
U	32.35%	0.00%	44.44%
	100.00%	0.00%	100.00%

Daily data:

1999	D	S	U
D	47.66%	60.00%	38.85%
S	0.94%	0.00%	2.88%
U	51.40%	40.00%	58.27%
	100.00%	100.00%	100.00%

Weekly data:

1999	D	S	U
D	36.84%	0.00%	38.24%
S	0.00%	0.00%	0.00%
U	63.16%	0.00%	61.77%
	100.00%	0.00%	100.00%

Daily data:

2000	D	S	U
D	44.14%	50.00%	43.80%
S	0.00%	0.00%	1.46%
U	55.86%	50.00%	54.75%
	100.00%	100.00%	100.00%

Weekly data:

2000	D	S	U
D	45.46%	0.00%	36.67%
S	0.00%	0.00%	0.00%
U	54.55%	0.00%	63.33%
	100.00%	0.00%	100.00%

Daily data:

2001	D	S	U
D	53.49%	0.00%	50.85%
S	1.55%	0.00%	0.85%
U	44.96%	100.00%	48.31%
	100.00%	100.00%	100.00%

Weekly data:

2001	D	S	U
D	53.57%	0.00%	58.33%
S	0.00%	0.00%	0.00%
U	46.43%	0.00%	41.67%
	100.00%	0.00%	100.00%

Daily data:

2002	D	S	U
D	44.44%	14.29%	43.70%
S	2.78%	14.29%	2.22%
U	52.78%	71.43%	54.07%
	100.00%	100.00%	100.00%

Weekly data:

2002	D	S	U
D	45.46%	0.00%	40.00%
S	0.00%	0.00%	0.00%
U	54.55%	0.00%	60.00%
	100.00%	0.00%	100.00%

Daily data:

2003	D	S	U
D	41.38%	50.00%	50.76%
S	1.72%	0.00%	0.00%
U	56.90%	50.00%	49.24%
	100.00%	100.00%	100.00%

Weekly data:

2003	D	S	U
D	47.83%	100.00%	35.71%
S	0.00%	0.00%	3.57%
U	52.17%	0.00%	60.71%
	100.00%	100.00%	100.00%

Daily data:

2004	D	S	U

Weekly data:

2004	D	S	U

APPENDIX 2: ANNUAL MARKOV PROBABILITIES 1986 THROUGH 2013

DAILY DATA | **WEEKLY DATA**

D	41.18%	0.00%	55.04%	D	59.09%	0.00%	29.03%
S	0.00%	0.00%	0.78%	S	0.00%	0.00%	0.00%
U	58.82%	100.00%	44.19%	U	40.91%	0.00%	70.97%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%

2005	D	S	U	2005	D	S	U
D	40.74%	75.00%	43.17%	D	54.17%	0.00%	39.29%
S	2.78%	0.00%	0.72%	S	0.00%	0.00%	0.00%
U	56.48%	25.00%	56.12%	U	45.83%	0.00%	60.71%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%

2006	D	S	U	2006	D	S	U
D	46.28%	0.00%	51.56%	D	45.83%	0.00%	50.00%
S	0.00%	0.00%	0.00%	S	0.00%	0.00%	0.00%
U	53.72%	0.00%	48.44%	U	54.17%	0.00%	50.00%
	100.00%	0.00%	100.00%		100.00%	0.00%	100.00%

2007	D	S	U	2007	D	S	U
D	43.70%	66.67%	49.23%	D	33.33%	0.00%	32.35%
S	1.68%	0.00%	0.77%	S	0.00%	0.00%	0.00%
U	54.62%	33.33%	50.00%	U	66.67%	0.00%	67.65%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%

2008	D	S	U	2008	D	S	U
D	53.49%	0.00%	48.78%	D	68.75%	0.00%	50.00%
S	0.00%	0.00%	0.81%	S	0.00%	0.00%	0.00%
U	46.51%	100.00%	50.41%	U	31.25%	0.00%	50.00%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%

2009	D	S	U	2009	D	S	U
D	50.41%	0.00%	45.80%	D	23.53%	0.00%	37.14%
S	0.00%	0.00%	0.00%	S	0.00%	0.00%	0.00%
U	49.59%	0.00%	54.20%	U	76.47%	0.00%	62.86%
	100.00%	0.00%	100.00%		100.00%	0.00%	100.00%

2010	D	S	U	2010	D	S	U
D	52.42%	100.00%	45.67%	D	47.62%	0.00%	37.50%
S	0.00%	0.00%	0.79%	S	0.00%	0.00%	0.00%

APPENDIX 2: ANNUAL MARKOV PROBABILITIES 1986 THROUGH 2013

DAILY DATA				WEEKLY DATA			
U	47.58%	0.00%	53.54%	U	52.38%	0.00%	62.50%
	100.00%	100.00%	100.00%		100.00%	0.00%	100.00%

2011	D	S	U	2011	D	S	U
D	42.02%	0.00%	51.88%	D	41.67%	0.00%	46.43%
S	0.00%	0.00%	0.00%	S	0.00%	0.00%	0.00%
U	57.98%	0.00%	48.12%	U	58.33%	0.00%	53.57%
	100.00%	0.00%	100.00%		100.00%	0.00%	100.00%

2012	D	S	U	2012	D	S	U
D	47.86%	100.00%	44.78%	D	65.52%	0.00%	45.46%
S	0.86%	0.00%	0.00%	S	3.45%	0.00%	0.00%
U	51.28%	0.00%	55.22%	U	31.03%	100.00%	54.55%
	100.00%	100.00%	100.00%		100.00%	100.00%	100.00%

2013	D	S	U	2013	D	S	U
D	45.38%	0.00%	49.62%	D	60.00%	0.00%	40.74%
S	0.00%	0.00%	0.00%	S	0.00%	0.00%	0.00%
U	54.62%	0.00%	50.38%	U	40.00%	0.00%	59.26%
	100.00%	0.00%	100.00%		100.00%	0.00%	100.00%

REFERENCES

Xu, H., Zhang, Z.X., *A Trend Deduction Model of Fluctuation Oil Prices*, (revised 2011), Institute of International Studies, Fudan University, Shanghai

Mostafaei, H. et al., *Modelling the fluctuations of Brent oil prices by a probabilistic Markov chain*, (2011), Journal of Computations & Modelling, vol. 1, no. 2, 17 – 26

Abazi, A., *Stochastic Volatility in the Crude Oil Prices: A Markov Chain Monte Carlo Approach*, (2003), unpublished thesis

Sirdari, M.Z. et al., *A sationarity test on Markov chain models based on marginal distribution*, School of Mathematical Sciences, Universiti Sains Malaysia