

Directional Changes: A New Way to Look At Price Dynamics

Edward P K Tsang

**Centre for Computational Finance & Economic Agents (CCFEA)
University of Essex, Colchester CO4 3SQ UK**

5 September 2017

Revised version appeared in Mandal J., Dutta P., Mukhopadhyay S. (eds) Computational Intelligence, Communications, and Business Analytics (CICBA), Communications in Computer and Information Science, Vol. 775, Springer, 2017, 45-55

Abstract

Prices in financial markets are normally summarized by time series, where transaction prices are sampled at fixed time intervals. Directional change is an alternative way of sampling data: transaction price is sampled when a significant change in the price is recorded. In this paper, we explain how directional changes can provide a valuable alternative perspective to price movements. We also describe the frontier of directional change research, which include forecasting, algorithmic trading and market tracking.

1 Introduction

In financial markets, assets are transacted in irregular times. Although all the transaction prices are recorded, they are normally summarized before further research is conducted. The most common way to summarize data is to create time series: from all the transactions, one data point is sampled at each fixed intervals. For example, the final transaction price every day is published as the “daily closing prices”. Most research would start with time series sampled this way. Many statistical methods have been developed over the years to facilitate time series studies.

Guillaume et al. (1997) introduced an alternative way to sample data, based on the concept of “directional change” (DC). The concept has been used by traders, under the name zigzag (Sklarew 1980), though it has not received serious attention by the research community. Details of DC will be introduced in the next section. The basic idea is to sample data at peaks and troughs. Thus the market is partitioned into uptrends and downtrends. The advantage of using DC is that extreme points are always sampled.

The purpose of introducing DC is to add a new perspective to data, which hopefully complements the perspective from time series. Unlike time series, DC research is still in its infancy. Many existing methods applicable to time series are not applicable to data sampled at irregular times. New methods must be invented. This paper summarizes some of the ideas developed so far, and presents the frontier of DC research.

2 Directional Change (DC): an alternative to time series

Unlike time series, which samples data at fixed intervals, DC samples data points at peaks and troughs in the market. What counts as a peak or trough depends on how big a change the observer considers to be significant. Suppose an observer considers 2% to be a significant change in price. When the market is on an uptrend, then whenever the price drops from the latest highest price by 2% or more, the last high will be marked as a peak, and the market is considered to have entered a downtrend. In a downtrend, whenever the

price rises from the latest low by 2% or more, the last low will be marked as a trough, and the market is considered to have entered an uptrend. (The definition above is a recursive one; to start, it does not matter whether one looks for an uptrend or a downtrend.) A formal definition of DC can be found in (Tsang 2010).

Different observers may use different thresholds (2% in the above example). Therefore, different observers may see different peaks and troughs. It is also worth pointing out that peaks and troughs are recognized in hindsight. A peak (trough) is only identified when the price has dropped (risen) by or beyond the threshold. The point at which price has dropped by or beyond a threshold is called a DC Confirmation point. The price change from the last extreme point (peak or trough) to the DC Confirmation point is called a DC Event. The price change from the confirmation point to the next extreme is called an Overshoot (OS) Event.

Figure 1 summarizes the basic concepts of DC. In Figure 1, A and E are troughs whereas C and G are peaks. At point B, price has risen from A by the predetermined threshold. Therefore, it is at point B that A is confirmed to be a trough. Point B is called a DC confirmation point. The market is in an uptrend from B onward, until price drops from the next record high by the predetermined threshold. C is confirmed to be a peak when price drops to D.

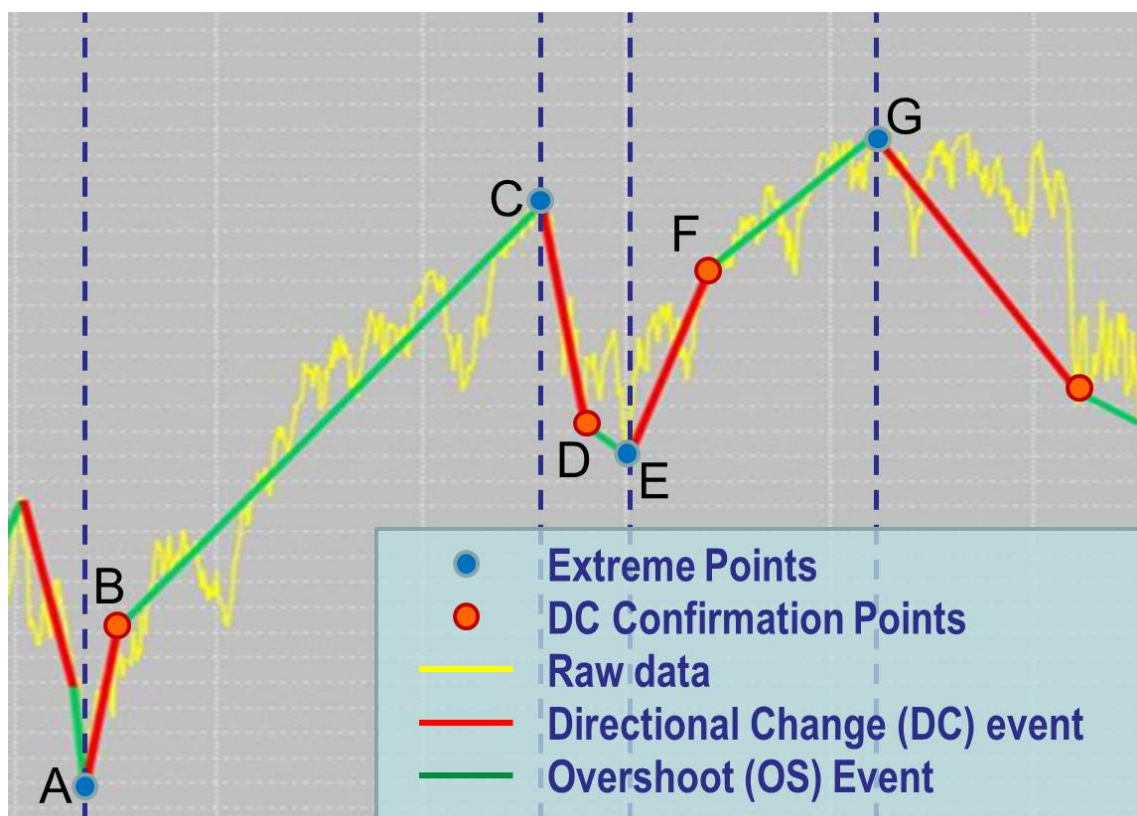


Figure 1 – Data summarized as Directional Changes

Under DC, Price changes are partitioned into alternating uptrends and downtrends: each trend comprises a DC Event and an OS event that follows. In Figure 1, between extreme points A and C is an uptrend which comprises a DC Event (AB) and an OS Event (BC);

between extreme points C and E is a downtrend which comprises a DC Event (CD) and an OS Event (DE).

One important goal in data science is to turn data into information. What counts as information? Information is something that the observer can make sense of. DC partitions the market into uptrends and downtrends, which correspond to the concepts of “bull” (uptrend) and “bear” (downtrend), which all traders understand. So DC is a useful way to extract information from financial data.

3 Describing price movements under DC

By sampling different data points, DC sees price movement from an angle different from time series. Under time series, one fixes time (in the x-axis) and measures changes in price (in the y-axis). Under DC, one fixes the threshold in price change (in the y-axis), and let data determine when to sample the next extreme point, i.e. let data determine the next value on the x-axis. This also determines the time at which the next data point is sampled.

3.1 Basic DC indicators: T, TMV and R

To describe price movements under DC, we need to develop a vocabulary. For each trend, one could record the total price changes from one extreme point to the next. One can also record the time that it takes to complete the trend. For example, see the trend from A to C in Figure 2. The **Time** that it takes to complete the trend is labelled with **T** (at the bottom of the figure). We define **Total Price Movement**, denoted by **TMV**, as the absolute percentage change from A to C, normalized by the threshold. For example, suppose the threshold used is 5%, and price change from A to C is 18%, then $TMV = \frac{|18\%|}{5\%} = 3.6$. The absolute value is used to allow TMVs in uptrends and downtrends to be compared. The value is normalized by the threshold so that changes under different thresholds are comparable with each other. The Annualised **Return** in the trend AC is labelled **R** in Figure 2. We shall focus on the three indicators T, TMV and R in this paper. A wider range of indicators is described in (Tsang 2017).

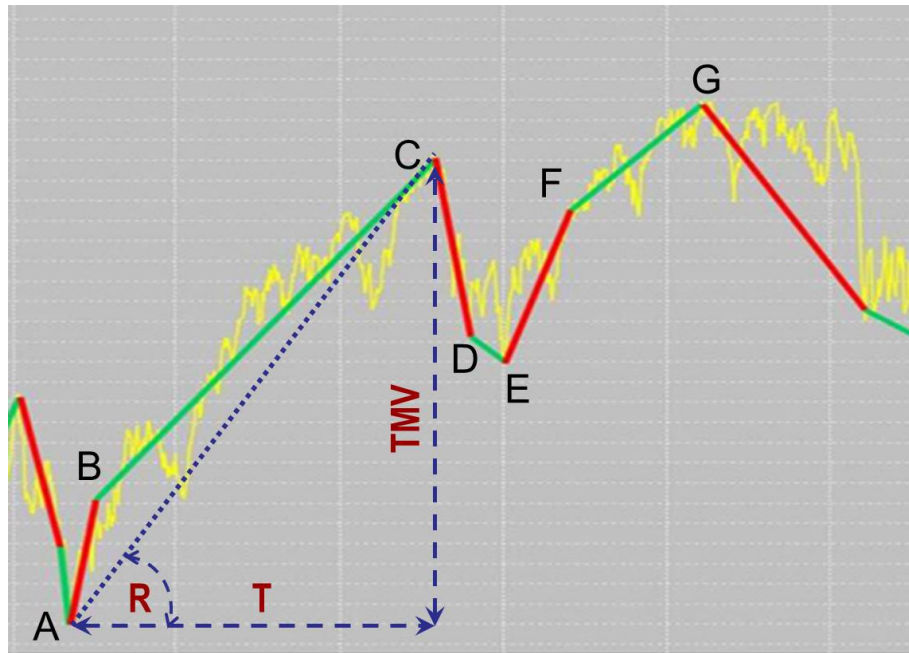


Figure 2 – Basic indicators under DC

3.2 Measuring Volatility under DC: frequency and magnitude

Under time series, volatility can be measured by the standard deviation of the returns over a period of time (e.g. the 7-days volatility measured by the standard deviation of the last 7 daily returns). Under DC, each trend may take a different amount of time (T) to complete. The smaller T is, the more frequently the market has changed between uptrend and downtrend. In other words, $\frac{1}{T}$ is a measure of volatility in terms of trend-switching **frequency**. One may use $\frac{1}{Median T}$ to measure the volatility over a period, where Median T is the median T value of all the trends in that period.

Frequency is only one way to measure volatility under DC. TMV measures the **magnitude** of price change in each trend. The higher the magnitude, the more volatile the market is. Given that theta is fixed within an observation, high TMV indicates high overshoot.

Frequency and magnitude are orthogonal with each other. Figure 3 shows hypothetical frequencies and magnitudes in three markets A, B and C. Market A is more volatile than Market B because while they share the same TMV, Market A takes only half of the time that Market B takes to complete each trend. Market B is more volatile than Market C, because although they take the same time to complete each trend, the trends in Market B have bigger TMV values than those in Market C.




Figure 3 – The hypothetical volumes and frequencies of three markets; Market A is more volatile than Market B (which is lower in frequency), which is more volatile than Market C (which is lower in volume)

4 Research frontier

Having described the basic concepts in DC, we shall describe the frontier of DC research. This includes researches that help building trading strategies and early warning systems, which could benefit traders, market makers and regulators.

4.1 Stylised facts in the market

One of the most significant findings in DC is the discovery of power laws in the market (Glattfelder et al 2011). By observing tick-to-tick data in foreign exchange markets over a long period of time, Glattfelder et al (2011) found stunning statistics in the market. The most significant two are shown in Figure 4 and explained below:

- i. The average percentage change in an overshoot (OS) is approximately the same as the threshold used (θ) to observe DC. In other words, TMV approximately equals to 2.
- ii. The average time that it takes to complete an overshoot event is approximately twice the time that it takes to complete a DC event.

The above power laws were purely based on observations. It is up to researchers to offer explanations. In general, being able to see regularities that others do not see, even if no explanation could be offered, gives the observer an edge in competitions. The above observations have been used in the design of trading algorithms (Golub et al 2017), which will be explained in Section 4.3.




Figure 4 – Two important power law observations in the foreign exchange market (Glattfelder et al 2011): (i) the average overshoot (OS) size is approximately equal to the threshold used (θ) to observe DC; (ii) the average time that it takes to complete an OS event is approximately twice as long as the time to complete a DC event.

4.2 Forecasting under DC

DC summarizes price movements in the past. Extreme points are only recognized in hindsight. As it is the case in time series, one could attempt to forecast under DC. There are many ways to formulate forecasting problems under DC. Three formulations are illustrated in Figure 5 and explained below.

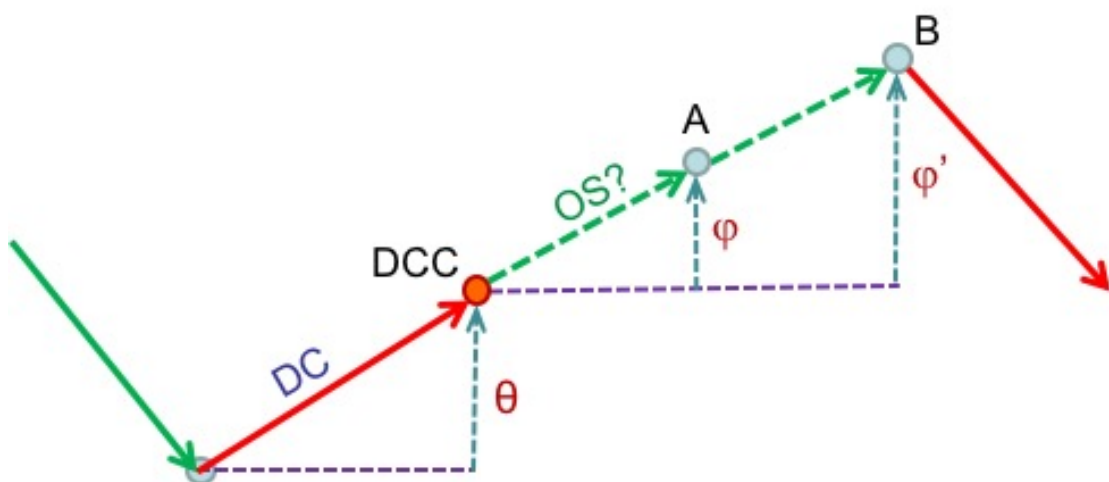


Figure 5 – A forecasting problem under DC: At DC confirmation point (DCC), to forecast whether overshoot would reach or exceed φ (point A); or, at point A, to forecast whether overshoot would reach or exceed φ' (point B)

- Formulation 1: At DC confirmation point (DCC) in Figure 5, one may attempt to forecast the magnitude of the overshoot. One could ask: how far would price go above DCC before we see the next DC? In other words, one may attempt to forecast the TMV value, which is a real number.
- Formulation 2: One could ask at DCC: would the price rise above DCC by a percentage ϕ before we see the next DC? Here the threshold θ could be used as a reference. For example, would $\phi \geq 0.5\theta$? Would $\phi \geq \theta$? In this case, one attempts to forecast a Boolean value.
- Formulation 3: One could also ask the above question at any time: for example, at point A, one may ask: would the price rise above DCC by a percentage ϕ' before we see the next DC?

Machine learning methods can be used to tackle the above forecasting problems. The key to any forecasting problem is to find the “independent variables”, which are variables that could provide information about the variable to be forecasted.

In Section 3.1, we introduced the indicator TMV, which is the normalized price change in the current trend. For convenience, we introduce the variable Overshoot Value (OSV) here. Let OSV^θ be the percentage price change from the DC confirmation point, normalised by the threshold θ . Some reflection should convince the readers that $OSV^\theta = 1 - TMV$.

Bakhach et al (2016a) attempted to forecast under Formulation 2 above. Bakhach et al (2016a) show that the OSV at point DCC under the DC observation with 2θ , i.e. $OSV^{2\theta}$, is a good indicator on whether $\phi \geq \theta$ will be true. In laymen terms: the bigger picture could tell us a lot about the smaller picture. This is illustrated in Figure 6.

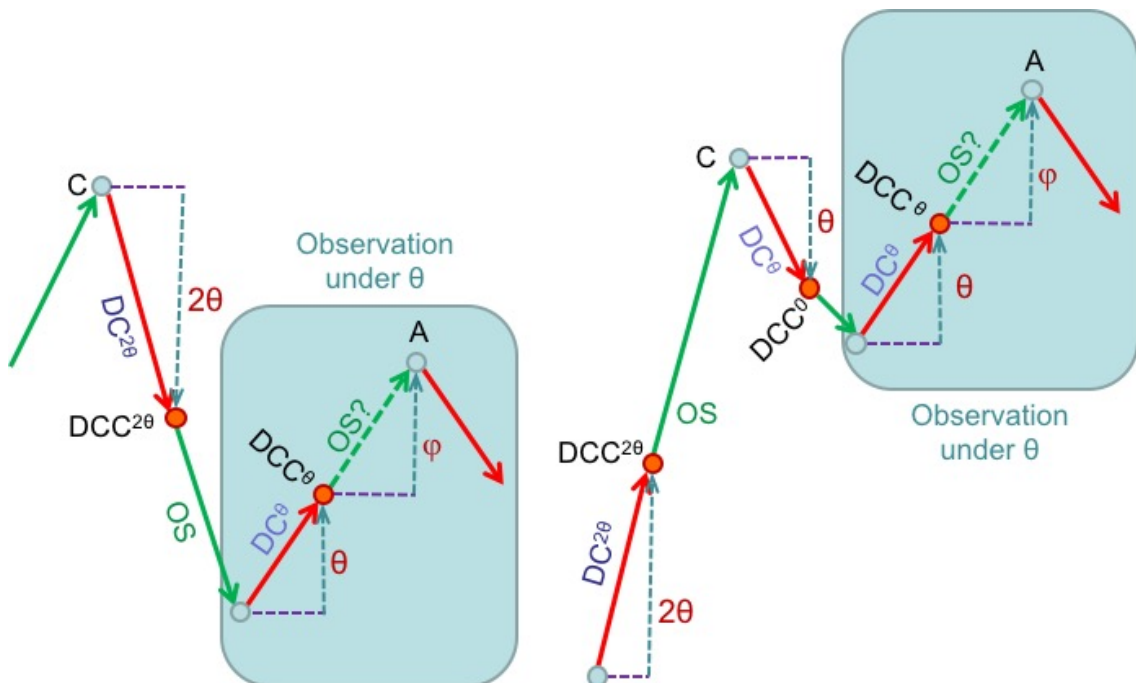


Figure 6 – Observation from a bigger picture (2θ) could benefit observation in a smaller picture (θ , shown in the shaded areas); (a) On the left of the figure, the DCC^0 point is still part of the downtrend after $DCC^{2\theta}$ was observed by DC using 2θ ; (b) on the right of the figure, the DCC^0 point is part of the uptrend when observed by DC using 2θ

On the left of Figure 6, the DCC^θ point observed under threshold θ is part of the downtrend after $DCC^{2\theta}$ was observed under threshold 2θ . In other words, the DCC^θ point confirms an uptrend (as observed under θ) within a bigger-scaled downtrend (as observed under 2θ). On the right of figure 6, the DCC^θ point is part of the uptrend when observed by DC using 2θ . In other words, the DCC^θ point confirms an uptrend (as observed under θ) within a bigger-scaled uptrend (which was confirmed after $DCC^{2\theta}$ was observed under 2θ). Results in Bakhach et al (2016a) suggest that the $OSV^{2\theta}$ value at the DCC^θ points in the two figures could help us to forecast whether price will reach point A or not.

4.3 Algorithmic trading under DC

Based on stylised facts observed in the market (as described in Section 4.1), Golub et al (2017) developed an algorithmic trading strategy which is called the Alpha Engine. Given that the average overshoot ends when price goes up by the threshold θ used to observe DC (see Figure 4), the Alpha Engine opens a contrarian position when TMV reaches 2; for example, a contrarian will be opened at Point C in Figure 4. It takes profit when price reverts to the next DC confirmation level. If price goes against the trader's position by another θ , the Alpha Engine will increase its position (i.e. short-sell more in the above example). The Alpha Engine has other tactics to improve its odds. For example, it adjusts its threshold dynamically. In an uptrend, upward moves are expected to be longer than the downward moves. Interested readers should refer to (Golub et al 2017).

The Alpha Engine was extensively back-tested on historical data comprised of 23 exchange rates over eight years, from the beginning of 2006 until the beginning of 2014. The trading model suffered from a drawdown of 7.08% while yielding an average profit of 10.05% in the final four years. This is a very promising performance, because in order to test the robustness of the Engine, the authors did not fine-tune the parameters to maximize performance. The Engine serves as a proof of concept, with plenty of room for improvement.

As a contrarian algorithm which buys when the market goes up, and sells when the market goes down, the Alpha Engine provides liquidity and stability to the market. The philosophy behind the Alpha Engine, which also explains why the algorithm is published, is that the profit that it makes is a reward for providing a service (namely, providing liquidity) to the market.

Independent of the Alpha Engine, Bakhach et al (2016b) reported a trading algorithm, which is called the Backlash Algorithm, which is also based on the power law observation described in Section 4.1. Like the Alpha Engine, the Backlash Algorithm is a contrarian algorithm. The algorithm is simpler than the Alpha Engine. It opens a contrarian position when the overshoot reaches a certain threshold, α , where the value of α is tuned by the immediate past history. Unlike the Alpha Engine, which opens a contrarian position when TMV reaches 2, the Backlash algorithm opens its position when TMV reaches $1+\alpha$. The Backlash algorithm takes profits at the next DC confirmation point. This means the profit is limited to θ (like the Alpha Engine). It could lose money if the trend goes against the opened position for another θ or more. Backlash was demonstrated to be profitable in preliminary tests.

The above trading algorithms demonstrate an important point: by exploiting the regularities observed in the market, profitable trading algorithms can be designed. When

the regularities are widely used in trading algorithms, they cease to support profitable trading. Traders will have to work harder to find less obvious regularities. To succeed, a trader has to find new hidden regularities faster than its competitors.

4.4 Pattern recognition

One of the on-going researches in DC is to recognize patterns in the market. Head and Shoulder is a well-known pattern in technical analysis (Chang & Osler 1999). One of the complications in applying Head and Shoulder to trading is to identify the “heads” and “shoulders”. What constitutes to a “head” or a “shoulder” could sometimes be ambiguous. This is where DC can help: once the threshold is specified, extreme points are well defined under DC. Our current research is to use DC to determine the extreme points, and use those extreme points to determine whether the Head and Shoulder pattern has emerged. This eliminates any subjective interpretation of what is and what is not a “head” or “shoulder” point. This research will be reported in (Li & Tsang, forthcoming).

4.5 Market profiling

Under time series, how can we describe the market over a period of time? We can talk about trends and volatility and many other things. What could we say under the DC framework? What are useful descriptive statistics under DC?

Tsang et al (2016) introduces a number of indicators for profiling markets under DC. Above we have already introduced TMV, T, R of a trend. To describe the market over a period, we use the median values of TMV, T and R plus other indicators (including indicators on extreme values, which will not be elaborated here). Median values are used instead of simple average because (a) each trend takes different time to complete, which means the simple average could be misleading; and (b) extreme values are common, based on our observations.

Tsang et al (2017) explain how these indicators can be used to assess stocks in the UK FTSE 100 index. They help us to see aspects of stock volatilities that cannot be seen under time series. Tsang et al (2017) also show that the combination of measures in DC and time series helps us assess the risks and potential of individual stocks.

The descriptive statistics under DC opens many research areas. Some of our current researches are described below.

4.5.1 Market positioning

Each period in a market can be described by a profile, which comprises a number of indicators. Each profile occupies a position in a profile space. Figure 7 shows a 2-dimensional profile space, where the x-axis measures Returns (R) and the y-axis measures TMV. The full profile described in (Tsang et al 2016) will occupy a position in a high-dimensional space which is difficult to visualize, but could be reasoned with computer programs.

With market positioning, we can find out which market-periods are similar to each other. One application of such similarity observations is in algorithmic trading. For example, the same algorithm may exhibit similar performances in similar markets. A trading strategy that is successful to one market should be tested carefully before it is applied to a market-period which is far away in the positioning space.

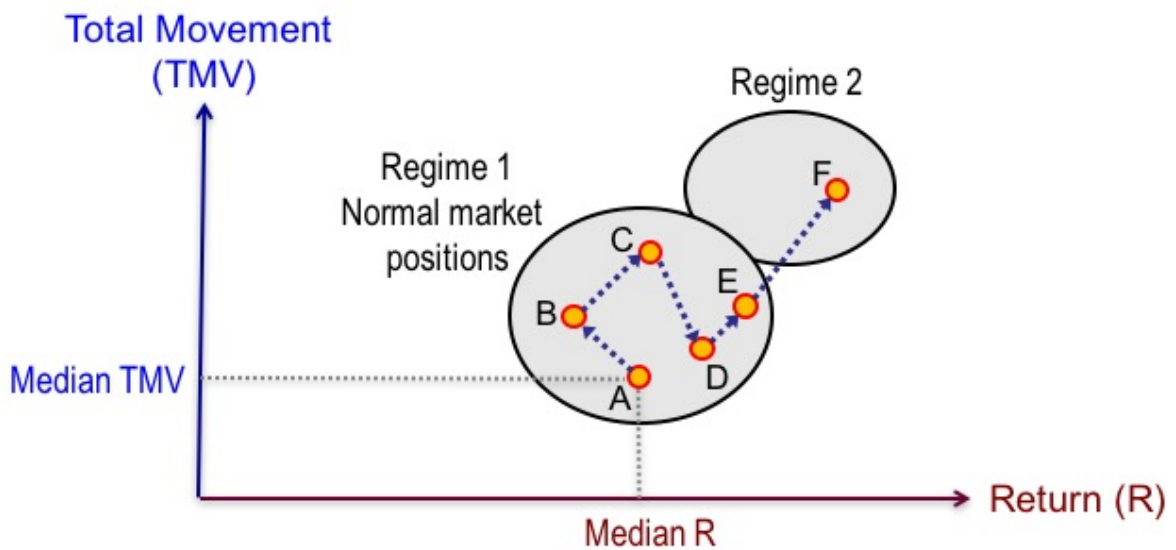


Figure 7 – Market positioning

4.5.2 Regime changes

By observing enough number of markets, with input from financial experts, one should be able to statistically recognize positions where “normal markets” occupy – here financial expertise is required to define what is “normal”. One could also attempt to classify different market regimes. Regime change recognition has been studied under time series. It is one of our on-going researches under DC.

Two hypothetical regimes are shown in Figure 7: profiles A to E fall into the region of “normal positions” (as indicated by financial experts), which is labelled Regime 1. Profile F falls into an abnormal regime, which is labelled Regime 2. Regimes recognized under time series and DC can be used together to enhance our understanding of the markets.

4.5.3 Market position tracking

Another on-going research is market tracking. The idea is to profile the same market over a rolling time window, and track the change of market positions over time. The trajectory of the market positions could help us to recognize significant changes in the market, which could help us to maintain our positions (i.e. buying or selling) or build early warning systems.

For example, in Figure 7, when the market position traverses from A to D, the market is within the normal regime. But if the market moves away from the centre of the normal regime (as the market moves to D and then E in Figure 7), market participants could be alerted. If the market moves into a recognized abnormal market regime, such as position F in Figure 7, then alarm could be raised.

5 Conclusion

Directional change (DC) is a new way to summarize price changes in a financial market, which let data decide when to record a transaction price. In this paper, we have explained how DC could help us to study price dynamics.

To use DC, we need to develop a vocabulary to describe data. In this paper, we have introduced TMV, T and DC Return (R); more can be found in Tsang et al (2016). We have demonstrated how to use this vocabulary to describe volatility in the market. We have also described the frontier of DC research, which involves discovering regularities, forecasting, algorithmic trading, pattern recognition, market profiling, regime change recognition and market tracking.

It is important to point out that time series and DC are not competing ways to study price dynamics; they complement each other. For example, volatility observed under time series may be used alongside observed frequency and TMV values under DC – they all tell part of the story. DC provides an alternative way to look at the market. It opens many doors for future research. We believe that two eyes are better than one. A stereo perspective of the market potentially allows us to see things that we cannot see from one angle.

Acknowledgements

This is an overview of a large project, which has benefited from input by Richard Olsen, Alex Dupuis, Anton Golub, Raju Chinthalapati, Antoaneta Serguieva, Han Ao, Amer Bakhach, Ran Tao, Shuai Ma, James Chen, Shengnan Li, Wing Lon Ng, Hamid Jalalian, Sara Lupino Lane and students from University of Essex.

References

- [1] Bakhach, A., Tsang, E.P.K. & Jalalian, H., Forecasting Directional Changes in FX Markets, IEEE Symposium on Computational Intelligence for Financial Engineering & Economics (IEEE CIFE'16), Athens, Greece, 6-9 December 2016
- [2] Bakhach, A., Tsang, E.P.K., Ng, W.L. & Raju Chinthalapati, V.L., Backlash Agent: A Trading Strategy Based On Directional change, IEEE Symposium on Computational Intelligence for Financial Engineering & Economics (IEEE CIFE'16), Athens, Greece, 6-9 December 2016
- [3] Chang, P.H.K. & Osler, C.L., Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts, *The Economic Journal*, Vol. 109, No. 458, 1999, 636-661
- [4] Guillaume, D., Dacorogna, M., Davé, R., Müller, U., Olsen, R., and Pictet, O., From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets. *Finance Stochastics*, Vol.1, Issue 2, 1997, 95-129
- [5] Glattfelder, J., Dupuis, A., and Olsen, R.B., Patterns in high-frequency FX data: discovery of 12 empirical scaling laws. *Quantitative Finance*, 11(4), 2011, 599-614
- [6] Golub, A., Glattfelder, J. & Olsen, R.B., The Alpha Engine: Designing an Automated Trading Algorithm, *High Performance Computing in Finance*, Chapman & Hall/CRC Series in Mathematical Finance, 2017; SSRN: <https://ssrn.com/abstract=2951348>

- [7] Li, S. & Tsang, E.P.K., Using Directional Changes for Searching Head and Shoulders Bottom Pattern, Working paper WP079-17, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, forthcoming
- [8] C. L. Osler, Identifying noise trader: the head-and shoulders pattern in U.S. equities, tech. report, FRB of New York, FRB of New York Staff Report No. 42, 1998
- [9] A. Sklarew, Techniques of a Professional Commodity Chart Analyst. New York: Commodity Research Bureau, 1980.
- [10] E.P.K. Tsang, Directional Changes, Definitions, Working Paper WP050-10, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, 2010
- [11] Edward P K Tsang, Ran Tao and Shuai Ma, Profiling Financial Market Dynamics under Directional Changes, Working Paper WP074-15, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, June 2015; revised 9 March 2016
- [12] E.P.K. Tsang, R. Tao, A. Serguieva & S. Ma, Profiling High Frequency Equity Price Movements in Directional Changes, Quantitative Finance, Vol.17, Issue 2, 2017, 217-225