Seasonality of dividend point indexes

Marco Marchioro
www.statpro.com
Version 0.90

December 2016*

Abstract

We analyse a variety of dividend-point indexes, typically used to determine the payoff of dividend derivatives. We show strong evidence of seasonality patterns for all indexes examined. Since dividend-point reset with irregular time intervals, we find the need to use a different seasonality function for the current or any future period. We determine the seasonality function by shifting forward the dividends paid in several past intervals and averaging. Finally we show how to use the seasonality functions to determine the virtual quotes of dividend futures with a constant-maturity expiry from the current date.

Keywords: dividend derivatives, dividend futures, dividend seasonality, dividend term structure, constant-maturity futures

1 Introduction

Investing in value stocks, i.e. in equities that pay-out dividend regularly, has been a well-known strategy for portfolio managers for many decades. However the welcome cash flow stream coming from dividends is typically also associated with the risks intrinsic in holding the underlying equity. From this point of view holding a dividend-paying equity gives exposure to two major risk factors: the stock price, that changes frequently with the financial markets, and the expected dividend stream, that changes because of the variations in the company dividend policies. These two risk factors typically vary at different frequencies: the first one changes fast, while the second one is usually slower. Note that these risk factors are related to two different types of future expected cash flows: the first one is linked to the resale value of the stock, while the second one is related to the future payments of dividends. For a long time investors could only take a position in both factors at the same time. Dividend derivatives are financial instrument that allow investors to take a view on the second risk factor *only*. See reference [3] for an overview of the market for dividend derivatives.

In recent years dividend derivatives became a popular new asset class for institutional investors seeking new sources of alpha and diversification. The growing popularity of dividend swaps and dividend futures has created a new market and there are now specialised desks in handling this type of instruments. Note that these derivatives can be used both for hedging the future dividend cash flows and to take a stance on them, i.e. the investor can either go long or short on the future dividend streams.

There are two major types of dividend derivatives: those based on single equities and those linked to an

*Latest review January 11, 2017

equity index. Single-name dividend derivatives take a position in the future dividend paid by an isolated equity in a future time period. These type of derivatives are traded typically for equities with a large capitalisation. However, the most popular choice for dividend derivatives has a typically payout that depends on the future dividends paid by the stocks making up equity indexes, such as the S&P 500, the Nikkei 225, or the EuroStoxx 50.

In this paper we consider mostly index dividend derivatives because they are the most popular ones. However, the results obtained for index dividend derivatives can be also applied to instruments based on a single stock.

2 Dividend point indexes

In general a derivative is a financial instrument with a payoff depending on an observable financial variable and, in the case of an index-dividend derivative, this variable is given by a *dividend point index*. In figure 1 we plot the graph of the Swiss SMI dividend-point index from December 2008 to December 2015. This index resets every year on the third Friday of December. We note two features of this graph:

- the cumulative sum of the dividends paid in each period is increasing from year to year,
- the index percentage increases happens on similar dates regardless of the year

While the first feature has to do with the bull-market phase in the business cycle, the second one is a stylised fact about dividend indexes.

A dividend index splits the time axis in periods and keeps track of the cumulative dividends paid out in each of these intervals. After the end of a period the index value is typically reset to zero. Then as soon as any equity belonging to the index pays a dividend, which may also happen on the first trading after the peak, the index is updated with an increment proportional to the cash paid out. The dividend point index uses the same scale as the index itself to keep track of the paid dividends. In other words given an equity index and a period in time, the dividend point index is defined as the sum of all dividends paid out since the beginning of the period, normalised in the same way as the index itself. If we think at the index as a fund with a NAV equal to its quote, then the dividend point index represents the cumulative dividends paid out by one share of the fund in the examined period up to a certain date.

To set out the notation used throughout the paper, consider a certain dividend index I so that for any date t we denote with $I(t)$ the observed index value. Suppose that t is in period that ends at date T_{n+1}. Then we have

$$I(t) = \sum_{i=1+T_n}^{t} d_i \qquad \text{for } T_n < t \le T_{n+1}\,, \tag{1}$$

where d_i is proportional to the sum of dividends paid at date i (see reference [5] for more details).

The period over which the index accumulates the dividends is typically of one year, however there are few cases in which the reset is made quarterly. In figure 2 for example, the quarterly index for the S&P 500 is plot for the four periods going from December 2011 to December 2012. Note how the index reset to zero at the end of each quarter.

To lighten the verbosity of the following discussions we use the name *dividend index*, or simply *index*, to denote the dividend point index. On the other hand we use the wording *underlying equity index* or simply *equity index* to refer to the original index of which we observe the dividends of.

2.1 Index reset-date conventions

Dividend indexes track the cumulative amount of dividends paid out in a certain period so that they are naturally periodic. As already stated the periodicity can be either quarterly or yearly. However depending on the specific dividend-point index at hand, there might be different conventions to determine when is the last date before the reset. Most indexes have a period last date that falls on the third Friday of December, however, there are other conventions. For example in the Japanese market the index resets are synchronised with the beginning of the social year, so that the last period

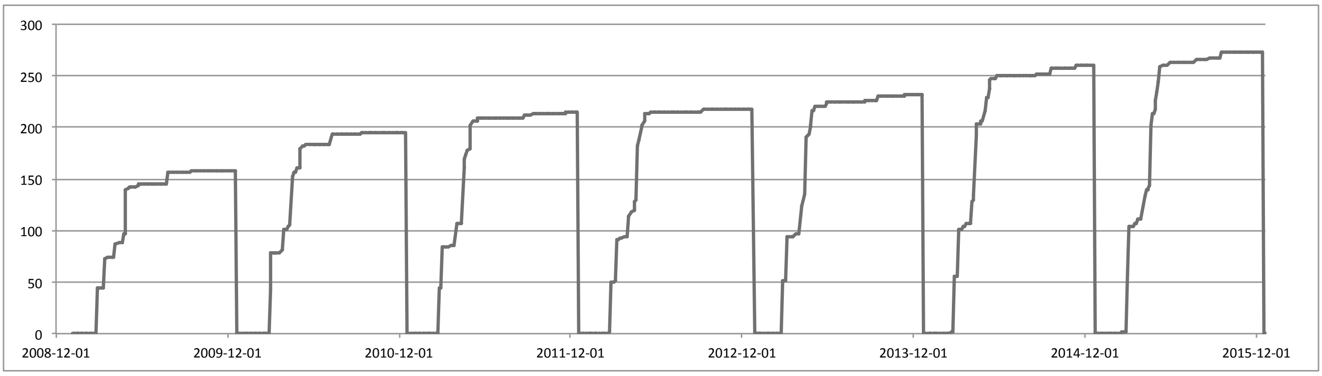

Figure 1: Graph of the Swiss SMI dividend-point index from December 2008 to December 2015. Note that the dividend maximum height increases every year and that there seem to be a periodicity in the timing of the dividend payments.

dates always falls on April 1st. The quarterly indexes usually have the last period date falling on the third Friday of the months of March, June, September, and December. For future reference we define a label for each reset-date convention:

(1) The most common convention that applies to annual indexes is to make the last date before the reset to fall on the third friday of December;

(2) In the Japanese market the reset date is always on the day following April first of every year;

(3) Most quarterly indexes, however, have four reset dates every year and they fall on the third Friday of the months of March, June, September, and December.

In table 1 we list the major dividend-point indexes, their periodicity, and the reset-date convention they follow.

In terms of notation we identify each period by its last date, so that period n ends at T_n. The same period n starts the day after T_{n-1}, a date that we denote with $1+T_{n-1}$. Since it is useful to distinguish among the current period, the periods completely in the future, and the periods completely in the past, we use the following convention:

- T_0 is the last date of the previous period
- T_1 is the last date of the current period
- T_2, T_3, ..., are, orderly, the last dates for all the periods that are completely in the future
- ..., T_{-2}, T_{-1}, are, orderly, the last dates for all periods before the previous one (i.e. that one ending at T_0).

Period lengths and alignements In table 2 we list the start date, the end date, and the number of days for periods ending in the years from 2015 to 2020, for the *Nikkey-225 DVP* and the *S&P 500 annual DVP*. As we can observe from the table, the number of days in each period varies between 364 to 371 days. While, by definition, for the *Nikkey-225 DVP* the period-end date always fall on April 1$^{\text{st}}$. On the other hand, for the *S&P 500 annual DVP*, the intervals end may fall anywhere between December 15$^{\text{th}}$ and December 21$^{\text{st}}$. Periods of different lengths and with different start and end dates pose some challenges when trying to relate dividends across multiple intervals.

3 Seasonality of dividend indexes

In order to improve the transparency with investors, companies making decisions on their dividend policies usually try to be consistent over long time periods. In a

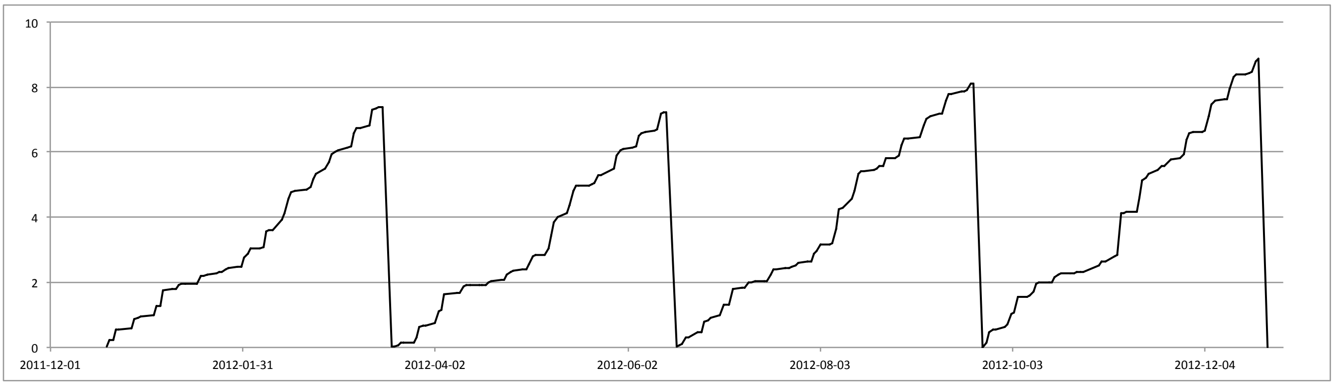

Figure 2: Graph of the S&P 500 quarterly dividend point index from December 2011 to December 2012. Note that the index resets to zero at the end of every quarter.

Index Name	**Periodicity**	**Reset Date**
Copenhagen 20 DVP	Annual	Third Friday in December
Helsinki 25 DVP	Annual	Third Friday in December
Stockholm 30 DVP	Annual	Third Friday in December
Swisse SLIDP	Annual	Third Friday in December
Swisse SMIDVP	Annual	Third Friday in December
Swisse SMIMDP	Annual	Third Friday in December
Nikkey-225 DVP	Annual	April first
S&P 500 quart. DVP	Quarterly	Third Friday Quarterly
S&P 500 annual DVP	Annual	Third Friday in December

Table 1: List of common dividend-point indexes, their periodicity, and their reset date.

fantasy world, for example, a company that sells umbrellas will have excess cash at the end of the rainy season and their dividend policy might be to pay investors at the beginning of the dry season. For this sample company, we may not know the exact amount of dividend paid next year, however, we can reasonably forecast the time of payments. Similarly, because of bookkeeping reasons, a company that receives steady cash flows during the whole year may still decide to pay dividends at the beginning, in the middle, or at the end of every quarter. Therefore for a value company, with a policy to pay dividends at the end of every quarter, we might not now the future dividend values however we can be confident in the payment dates.

For these and possibly other reasons we suspect that there is some periodicity in the dividend payments. In order to confirm this suspicion we may also look at the long term graph of major dividend indicies. For example in figure 1 we can look at the plot of the Swiss SMI dividend index for seven consecutive years. As noted earlier the increment in the index seem to happen more less at the same dates.

The seasonality of dividend payments is an important feature of this type of indexes and being able to measure it should bring benefits to dividend-derivative investors.

Nikkey-225 DVP			S&P 500 annual DVP		
start date	**end date**	**days**	**start date**	**end date**	**days**
2014-04-02	2015-04-01	365	2014-12-20	2015-12-18	364
2015-04-02	2016-04-01	366	2015-12-19	2016-12-16	364
2016-04-02	2017-04-01	365	2016-12-17	2017-12-15	364
2017-04-02	2018-04-01	365	2017-12-16	2018-12-21	371
2018-04-02	2019-04-01	365	2018-12-22	2019-12-20	364
2019-04-02	2020-04-01	366	2019-12-21	2020-12-18	364

Table 2: For the periods ending in the years from 2015 to 2020, we list the start date, the end date, and the number of days for the *Nikkey-225 DVP* (left side) and the *S&P 500 annual DVP*, (right side). The dates are formatted according to the ISO-8601 standard, i.e. according to the `YYYY-MM-DD` convention.

Measure of seasonality There are at least two major motivations to measure the dividend-index seasonality. The first one is purely economic: we want to know in when in the year dividends are paid. The second application is related to the risk-neutral evaluation of dividend derivatives and is the focus on this paper. The quote of a dividend futures, for example, represents the risk-neutral expectations of the total dividends paid in a certain future period (e.g., period n with $n \geq 2$). However we usually need a seasonality function to tell us when, in the future interval, those dividends are expected to be paid.

For the purposes of this paper, not only we are interested in the evaluation of dividend derivatives, we are also interested in their risk simulations using the historical simulation method (see, e.g., reference [4]). A challenge in performing the historical simulations for the dividend term structure is that the dividend futures are typically quoted on fixed dates, while we would like to have available constant-maturity futures. In other words we would like to have dividend future quotes for maturities that are always forward in time of an integral number of years, starting from the current date (or multiple quarters for quarterly indexes).

3.1 Overlapping historical periods on the reference interval

As noted earlier, we are interested in measuring the index seasonalities to split the expected dividends in a reference interval among the different dates. In terms of notation we define the *reference interval* as either the current period, i.e. the period with n=1, or any other period completely in the future (i.e. with $n \geq 2$). We are going to compute some sort of seasonality function averaging the index data from the observed *historical past periods*, i.e. periods with $n \leq 0$. We assume the reference interval to start at T_{r-1} and end at T_r, with $r \geq 1$ When r=1 we are considering the current period: in this case some of the dividends have been already paid. On the other hand for $r > 2$ the start date in the future with no actual dividend paid so far.

In order to compute a seasonality measure we need to overlap historical periods with the reference interval. However in general different periods have a different number of days. Also periods may start, or end, at different dates compered with the reference interval. As a consequence there are many different ways in which we can *cut* the historical periods and overlap them to the reference one.

To better understand the problem visually, we examine a fictional sample dividend index with a past history of three years that we call *Year A*, *Year B*, and *Year C*. We then assume the reference interval to be given by the current period: *Year R*. In figure 3 the different years are schematically drawn as a sequence of pointy segments. Note that the periods are shown to have different lengths (not to scale) to draw attention to the fact that they may contain different number of days. Also typically the periods start at different calendar dates, determined for example by the date of the third Friday in December.

There are at least three ways in wich we can align periods from different years:

1. *date-alignement*, in which the same calendar date, i.e. day and month, from different years are matched;
2. *period-start alignement*, in which we assume that the first dates of the periods are matched, regardless of the calendar date the fall in;
3. *period-end alignement*, in which the last dates of each period are matched.

In figure 4 we visually reppresents the different types of alignements for the same fictional periods as those pictured in figure 3.

Accrual index By definition dividend-point indexes track the *cumulative* dividends paid out by the equities in the index over a certain period. This means that the level of the index at one point depends on *all* the dividends paid out since the period start date. For example when we are using the *date alignement* in trying to match the dividends paid at a certain calendar date for different years, we need to keep track of the index level at each date. It turns out that is easier, instead of dealing with the dividend-point index itself, to go back to the actual paid dividends. Therefore, before matching the data for different periods, we compute the *accrual index* as the index difference from one day to the previous one. Note that in the calculation of the accrual index we need to be careful when computing the difference at the interfaces of two periods, as we do not want to include artificially high jumps.

Dealing with quarterly indexes If we are dealing with quarterly indexes we need to be even more careful when overlapping different periods. For example we do not want to match dates that belong to different seasons. Therefore after we split the index history in different years, we further cut it in different quarters. In figure 5 for example we show how to align periods coming from different dates and different quarters for a fictional quarterly index.

Filling data gaps When performing the alignements of the historical periods with respect to the reference interval, see figure 4, we sometimes need to trim the past period data and some other times we might need to fill some gaps. While trimming the excess data is usually not a problem we need to define how to fill the gaps of missing accruals. We do this by *borrowing* the missing data from the previous historical period when the gap is at the beginning of the reference period and from the next period when it is at the end.

Comparing different period alignements We compare the alignement results of the accrual index according to the different alignement techniques.

In figure 6 we show the first four weeks for the numerical results of the alignement of the S&P 500 quarterly accrual index in the first quarter. Since for this index the last period date always falls on a Friday (in this case the third Friday of March), when we use the *date-alignement*, so that the periods are synchronised on the same calendar date, the start and the end dates typically fall on different weekdays. On the other hand, when using the *period-start* or the *period-end* alignement, the intervals are matched according to the weekday.

For this specific case we note that the companies forming the index are more likely to pay dividends at different years on the same weekday rather than on the same date. As a consequence, as seen in figure 6, the *date alignement* seems less suitable as it scatters dividend payments in a diagonal pattern. For this index it seems that the `period-start` and the `period-end` alignement provide very similar results.

3.2 Reconstructed dividend indexes

The computation of the accrual index is only necessary to stick together dividend-index data coming from

Figure 3: Visual representation of historical periods for a fictional dividend index spanning three years in the past (pointy segments coloured in light blue). The reference interval, in this case *year R*, is represented by a pointy segment coloured in light red.

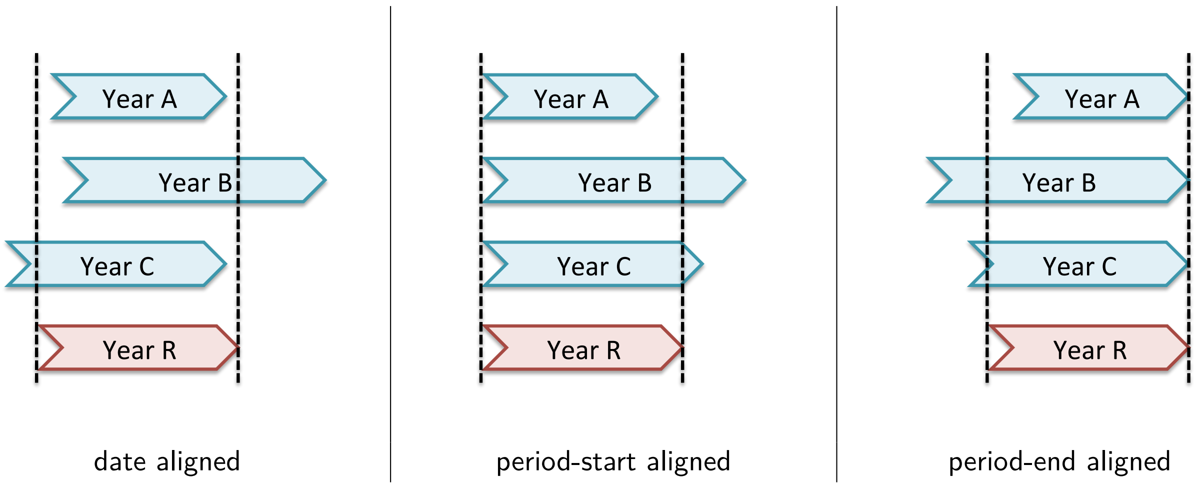

Figure 4: Three examples of possible alignement for the periods of figure 3: on the left-hand side we represent the *date* alignement, in the center pane we have the *period-start* alignement, and on the right-hand side the *period-end* alignement.

different periods. After we are done with the period overlapping we can reconstruct the artificial index with respect to the reference period by computing the cumulative sum of the accruals. In figure 7, for example, we show the reconstructed index for the ten periods from Q1-2007 to Q1-2016 with the respect to the Q1-2017 reference period. We notice that,

1. although it is not perfectly clear, there seems to emerge a seasonality feature common to all periods
2. and that the reconstructed index is not smooth because of the presence of weekly jumps.

Weekly smoothing of accrual indexes As seen in figure 7 the weekly jumps prevent us from clearly see a seasonality pattern in the reconstructed dividend index. In order to obtain a smooth result, for example, one could simply perform a running average of some sort on the reconstructed index, or use some other kind of smoothing technique. However such a procedure would inevitably alter the total amount of dividend paid in a certain period. We propose instead a smoothing procedure that does not alter the total amount of dividends paid in each week.

The procedure is very simple: starting form the accrual index we split the reference period in weeks and

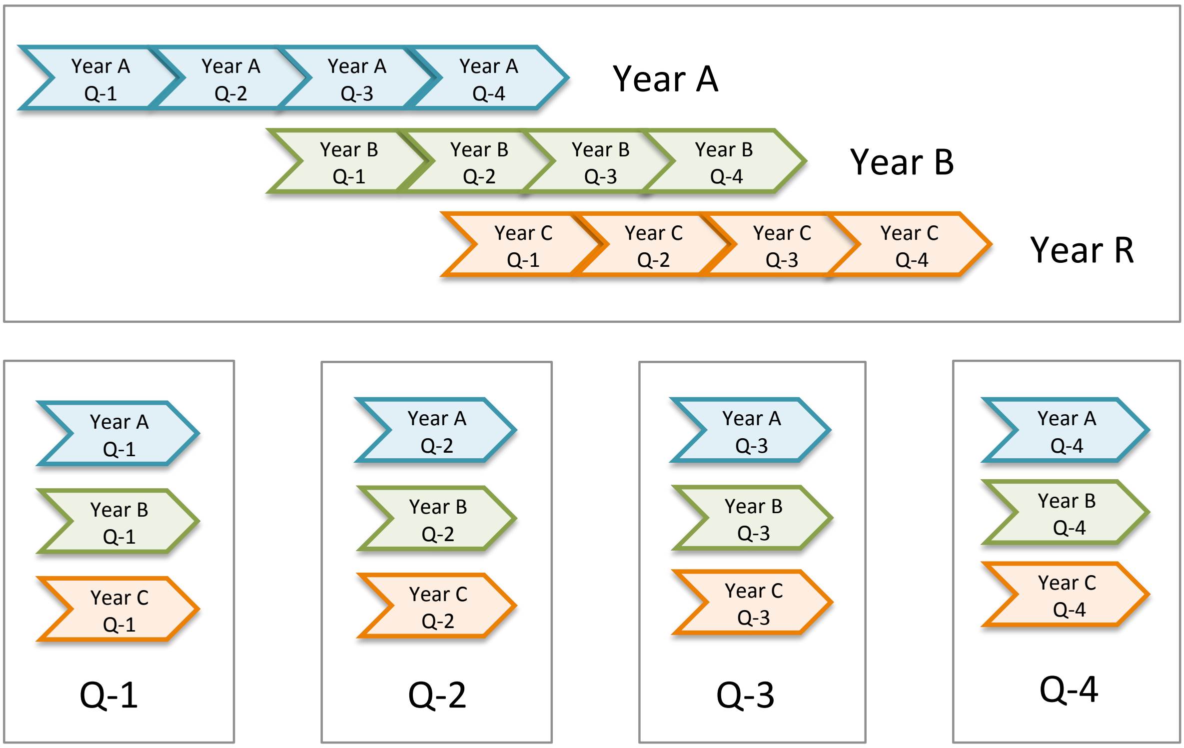

Figure 5: Split of the history of a fictional quarterly dividend index in years and seasons. The alignements are then performed preserving the quarters.

compute the total amount of dividend paid in each week. Then we create a new accrual-index time series, the *smooth accrual index*, in which the weekly dividends are equally spread among the days that belong to the originating week. By construction at the end of the period the smooth and the original accrual index reach the same level.

In figure 8 we show the smooth reconstructed dividend index for the same data as the graph of figure 7. Notice how the smooth dividend index no longer contains the weekly wiggles and that we did not alter the total level of dividends paid during the period.

In summary, starting from the original dividend-point index we compute the accruals, then we align the period data with respect to the reference period, we apply the weekly smoothing and end up with a reconstructed index as that shown in figure 8.

Normalised period results Observing figure 8, the seasonality feature is now more apparent but it is still not easily recognisable and therefore not very useful. The problem is that the different period results end with very different amounts of total dividends. Therefore, in order to be able to obtain an average seasonality function that we can apply to the reference interval, we normalise the period results so that they all have a value of 100% at the period end.

In figure 9 we show a plot of the normalised reconstructed smooth period results. We notice that other than the graph for the period Q1-2009 (which was a pretty special period for the stock markets, as a multi-year bottom was reached) all the other plots tend to

Aligned-dates

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016
2016-12-17	0	0.04	0.04	0.05	0.01	0	0.02	0.02	0.1	0.31
2016-12-18	0	0.01	0.04	0.01	0	0	0.03	0.06	0.31	0
2016-12-19	0.06	0.04	0.01	0	0	0.01	0.05	0.26	0.01	0
2016-12-20	0.05	0.39	0	0	0.03	0.2	0.3	0	0	0
2016-12-21	0.32	0.22	0	0.02	0.18	0.02	0.1	0	0	0.24
2016-12-22	0.2	0	0.08	0.01	0.18	0.32	0	0	0.29	0.21
2016-12-23	0	0	0.18	0.25	0.09	0.01	0	0.28	0.1	0.16
2016-12-24	0	0.01	0.4	0	0	0	0.18	0.2	0.08	0.01
2016-12-25	0	0	0	0	0	0	0	0	0	0
2016-12-26	0.02	0.03	0	0	0	0	0.02	0	0	0
2016-12-27	0.3	0.36	0	0	0.01	0.02	0.26	0.34	0	0
2016-12-28	0.06	0.03	0	0.15	0.01	0.28	0.14	0	0	0
2016-12-29	0.02	0	0.28	0.19	0.22	0.04	0	0	0.36	0.4
2016-12-30	0	0	0.18	0.11	0.03	0.06	0	0.14	0.25	0.22
2016-12-31	0	0.03	0.02	0	0.01	0	0.01	0.12	0	0.01
2017-01-01	0	0	0	0	0	0	0	0	0	0
2017-01-02	0	0.29	0.16	0	0	0	0.27	0.26	0.28	0
2017-01-03	0.26	0	0	0	0.05	0.04	0	0.06	0	0
2017-01-04	0	0	0	0.06	0.02	0.27	0	0	0	0.46
2017-01-05	0	0	0.03	0	0.16	0	0	0	0.14	0
2017-01-06	0	0	0	0.5	0.49	0.48	0	0.03	0	0.71
2017-01-07	0	0.01	0.52	0.06	0	0	0.01	0.02	0.63	0.06
2017-01-08	0.44	0.48	0.04	0	0	0	0.5	0.54	0.04	0
2017-01-09	0	0.1	0.01	0	0	0.03	0.02	0	0.01	0
2017-01-10	0.17	0.03	0	0	0	0.02	0	0	0	0
2017-01-11	0.02	0.08	0	0	0	0.13	0.16	0	0	0
2017-01-12	0	0	0	0	0.14	0.03	0	0	0	0.01
2017-01-13	0	0	0.06	0.11	0	0	0	0.2	0.21	0.22
2017-01-14	0	0	0.08	0.03	0.04	0	0	0.01	0.06	0.03

Period-start

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016
2016-12-17	0	0	0	0	0	0	0	0	0	0
2016-12-18	0	0	0	0	0	0	0	0	0	0
2016-12-19	0	0.01	0.08	0.02	0.03	0.01	0.18	0.28	0.29	0.24
2016-12-20	0.06	0	0.18	0.01	0.18	0.2	0	0.2	0.1	0.21
2016-12-21	0.05	0.03	0.4	0.25	0.18	0.02	0.02	0	0.08	0.16
2016-12-22	0.32	0.36	0	0	0.09	0.32	0.26	0	0	0.01
2016-12-23	0.2	0.03	0	0	0	0.01	0.14	0.34	0	0
2016-12-24	0	0	0	0	0	0	0	0	0	0
2016-12-25	0	0	0	0	0	0	0	0	0	0
2016-12-26	0	0.03	0.28	0.15	0.01	0	0.01	0.14	0.36	0
2016-12-27	0.02	0	0.18	0.19	0.01	0.02	0	0.12	0.25	0.4
2016-12-28	0.3	0.29	0.02	0.11	0.22	0.28	0.27	0	0	0.22
2016-12-29	0.06	0	0	0	0.03	0.04	0	0.26	0	0.01
2016-12-30	0.02	0	0.16	0	0.01	0.06	0	0.06	0.28	0
2016-12-31	0	0	0	0	0	0	0	0	0	0
2017-01-01	0	0	0	0	0	0	0	0	0	0
2017-01-02	0	0.01	0.03	0.06	0.05	0	0.01	0.03	0.14	0.46
2017-01-03	0	0.48	0	0	0.02	0.04	0.5	0.02	0	0
2017-01-04	0.26	0.1	0.52	0.5	0.16	0.27	0.02	0.54	0.63	0.71
2017-01-05	0	0.03	0.04	0.06	0.49	0	0	0	0.04	0.06
2017-01-06	0	0.08	0.01	0	0	0.48	0.16	0	0.01	0
2017-01-07	0	0	0	0	0	0	0	0	0	0
2017-01-08	0	0	0	0	0	0	0	0	0	0
2017-01-09	0.44	0	0	0	0	0.03	0	0.2	0	0
2017-01-10	0	0.02	0.06	0	0	0.02	0	0.01	0.21	0.01
2017-01-11	0.17	0.15	0.08	0.11	0.14	0.13	0.18	0.05	0.06	0.22
2017-01-12	0.02	0.06	0.03	0.03	0	0.03	0.01	0.04	0.05	0.03
2017-01-13	0	0.03	0	0.04	0.04	0	0.03	0.04	0.03	0.08
2017-01-14	0	0	0	0	0	0	0	0	0	0

Period-end

	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016
2016-12-17	0	0.22	0	0	0	0.03	0	0	0	0
2016-12-18	0	0	0	0	0	0	0	0	0	0
2016-12-19	0	0	0.08	0.02	0.03	0	0.02	0.28	0.29	0
2016-12-20	0.06	0.01	0.18	0.01	0.18	0.01	0.03	0.2	0.1	0.24
2016-12-21	0.05	0	0.4	0.25	0.18	0.2	0.05	0	0.08	0.21
2016-12-22	0.32	0.03	0	0	0.09	0.02	0.3	0	0	0.16
2016-12-23	0.2	0.36	0	0	0	0.32	0.1	0.34	0	0.01
2016-12-24	0	0.03	0	0	0	0.01	0	0	0	0
2016-12-25	0	0	0	0	0	0	0	0	0	0
2016-12-26	0	0	0.28	0.15	0.01	0	0.18	0.14	0.36	0
2016-12-27	0.02	0.03	0.18	0.19	0.01	0	0	0.12	0.25	0
2016-12-28	0.3	0	0.02	0.11	0.22	0.02	0.02	0	0	0.4
2016-12-29	0.06	0.29	0	0	0.03	0.28	0.26	0.26	0	0.22
2016-12-30	0.02	0	0.16	0	0.01	0.04	0.14	0.06	0.28	0.01
2016-12-31	0	0	0	0	0	0.06	0	0	0	0
2017-01-01	0	0	0	0	0	0	0	0	0	0
2017-01-02	0	0	0.03	0.06	0.05	0	0.01	0.03	0.14	0
2017-01-03	0	0.01	0	0	0.02	0	0	0.02	0	0.46
2017-01-04	0.26	0.48	0.52	0.5	0.16	0.04	0.27	0.54	0.63	0
2017-01-05	0	0.1	0.04	0.06	0.49	0.27	0	0	0.04	0.71
2017-01-06	0	0.03	0.01	0	0	0	0	0	0.01	0.06
2017-01-07	0	0.08	0	0	0	0.48	0	0	0	0
2017-01-08	0	0	0	0	0	0	0	0	0	0
2017-01-09	0.44	0	0	0	0	0	0.01	0.2	0	0
2017-01-10	0	0	0.06	0	0	0.03	0.5	0.01	0.21	0
2017-01-11	0.17	0.02	0.08	0.11	0.14	0.02	0.02	0.05	0.06	0.01
2017-01-12	0.02	0.15	0.03	0.03	0	0.13	0	0.04	0.05	0.22
2017-01-13	0	0.06	0	0.04	0.04	0.03	0.16	0.04	0.03	0.03
2017-01-14	0	0.03	0	0	0	0	0	0	0	0.08

Figure 6: Tables with the three different type of alignements, for the S&P 500 quarterly accrual index, for ten consecutive years from 2007 to 2016, with the first quarter of 2017 as the reference period. The cells are conditionally formatted so that the larger accruals have darker shades of green.

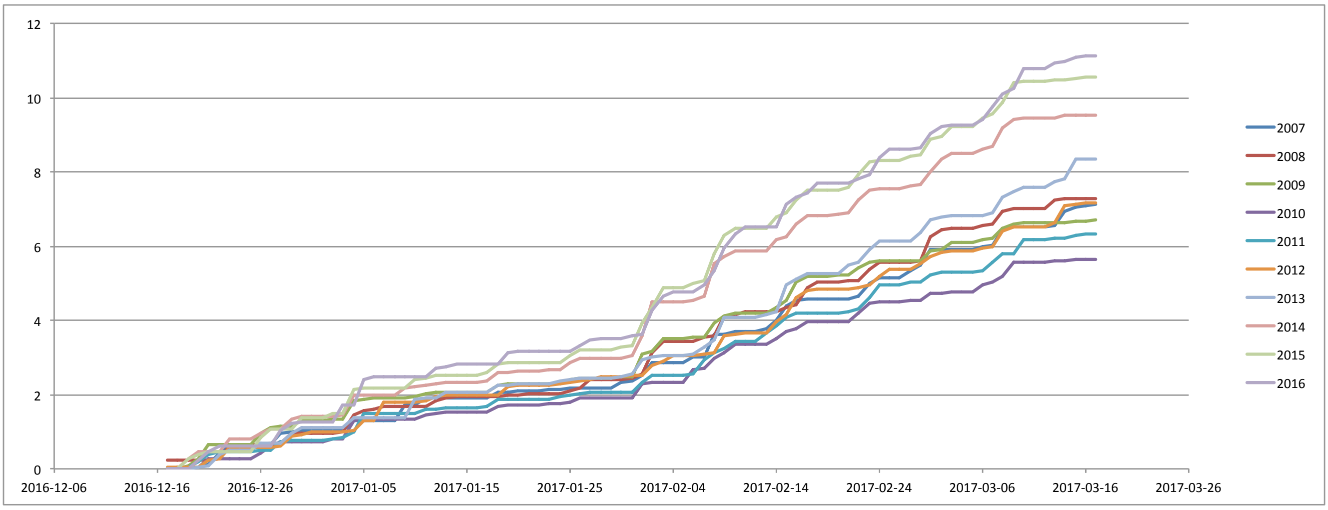

Figure 7: Graph of the ten reconstructed indexes, for the S&P 500 quarterly index, for the years from 2007 (Q1) to 2016 (Q1) with the respect to the Q1 of 2017 reference period.

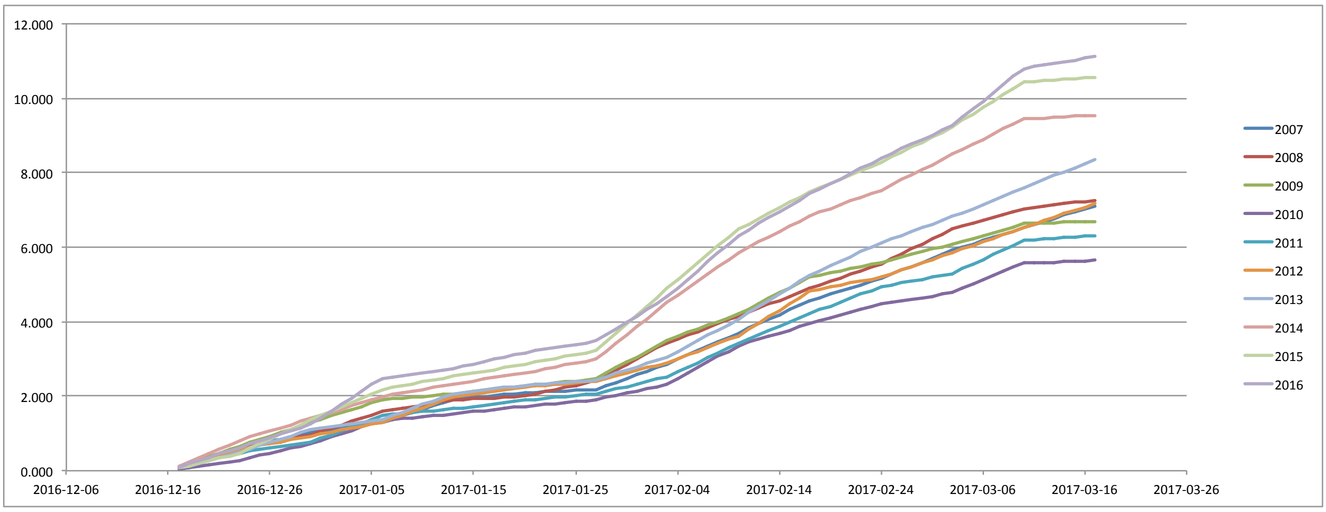

Figure 8: Graph of the reconstructed *smooth* index, for the ten periods from 2007 (Q1) to 2016 (Q1), with respect to the Q1 of 2017, for the same data set as figure 7.

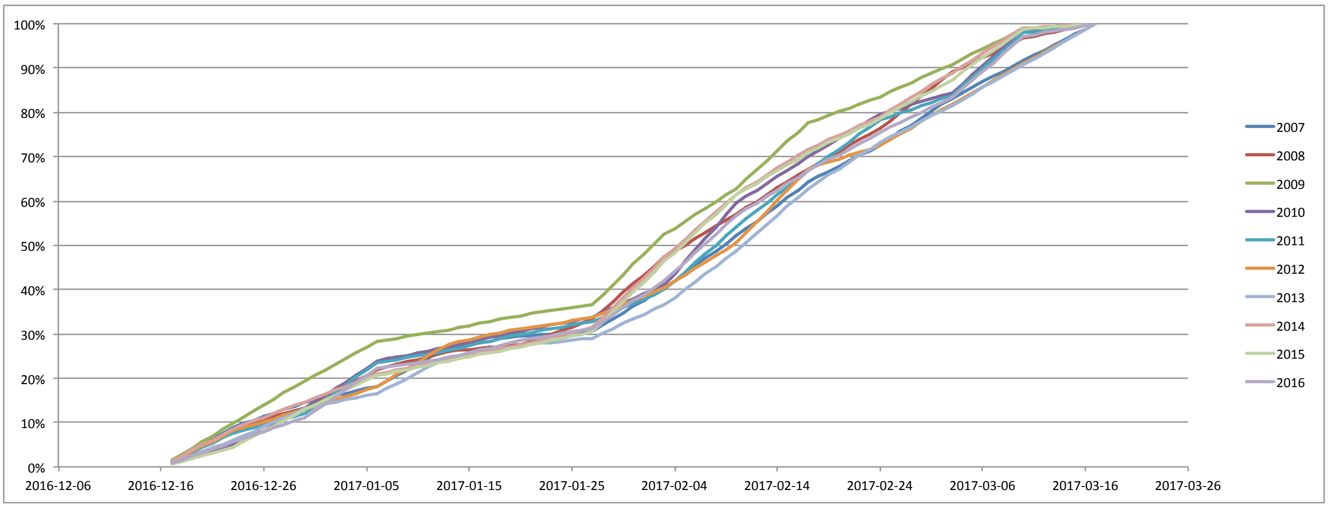

Figure 9: Normalised reconstructed smooth period results, for the S&P 500 quarterly index, for the ten periods from Q1-2007 to Q1-2016 with respect of the reference period Q1-2017.

overlap. We consider this overlapping evidence of the seasonality of the considered dividend index.

3.3 The dividend seasonality function

The normalisation of the reconstructed dividend index by its value at the end of the period allows us to compare results from different intervals (see figure 9 for an example). At this point it seems natural to take the average over all the periods to extract the *dividend seasonality function*, or for brevity the *seasonality function*. Therefore we define the seasonality function for the reference period r as

$$S_r^{[I]}(t) = \left\langle \frac{I'(t)}{I'(T_r)} \right\rangle \qquad \text{for } T_{r-1} < t \leq T_r \,,$$

where the symbol $\langle . \rangle$ denotes the average over all historical past periods and I' denotes the smooth reconstructed dividend index.

This function grows from 0% at the beginning of the reference period, to 100% at its end. It represents the expected percentage of dividends to be paid by date t during the reference interval. The increment of the seasonality function provides the average percentage amount of dividends paid at a certain date. Note that for every index I, at least in principle, we have a countable infinity of seasonality functions label by $n{=}1, 2, 3, \ldots$. In figure 10 we plot, as a red line, the dividend seasonality function S[Q1] for the S&P 500 quarterly index for the period Q1 of year 2017.

What should we compare the seasonal normalised reconstructed index with? To answer this question we need to consider the case in which there is no seasonality in the index. In this case the dividend are paid uniformly during the period. Hence, we consider a fictional dividend index, the *uniform dividend index* denoted by U, that pays every day the exact same dividend amount. Even thought such an index is very boring, it can be used as a benchmark for all other dividend indexes. The *normalised uniform index* $U_{\%}$ is just a straight line that grows from 0% to 100%. In figure 10 we plot, as a blue line, the normalised uniform index for the Q1 (year 2017) reference period.

Given a certain index and a seasonality function S, when S is below $U_{\%}$, i.e. in the plot when the red line is below the blue one, it means that the index has not kept up dividend payments with respect to the uniform index. In this case we expect more dividend payments later on. On the other hand when the seasonality function is above $U_{\%}$, i.e. in the plot when the red line is above the blue one, more dividends have been paid with respect to the uniform index. Since both indexes need to reach 100% at the end of the period we expect less dividend payments afterwards.

Given a certain dividend index I, in order to obtain a quantitative measure of the deviation from the uniform-dividend line we define the *rebased seasonality function* (or just *rebased function*) $Z[I]$ as the difference between the seasonality function and $U_{\%}$. In other words

$$Z(t) = S(t) - U_{\%}(t) \,.$$

In figure 10 we plot, as a green line, the rebased seasonality function $Z[Q1]$. By definition this function always starts and ends at zero. The function might be positive, indicating an overpayment of dividends, or negative denoting an underpayment of dividends with respect to the uniform index. Finally note that if there is no seasonality in the original dividend index, the averaged excess dividend payments from one year would cancel out defect payments from another year, therefore in this case the rebased seasonality function is close to zero. A non-null rebased seasonality function indicates the time non-uniformity in the dividend payments that we call seasonality. Note that the convenience of computing the rebased seasonality function is just in providing a graphical intuition about the dividend payments; in the actual quantitative computations the dividend seasonality function is more useful.

Steps leading to the dividend seasonality function Before proceeding in illustrating the numerical results for the the indexes of table 1, we summarise the steps leading to the computation of the dividend seasonality function.

1. We start by collecting the historical values of a

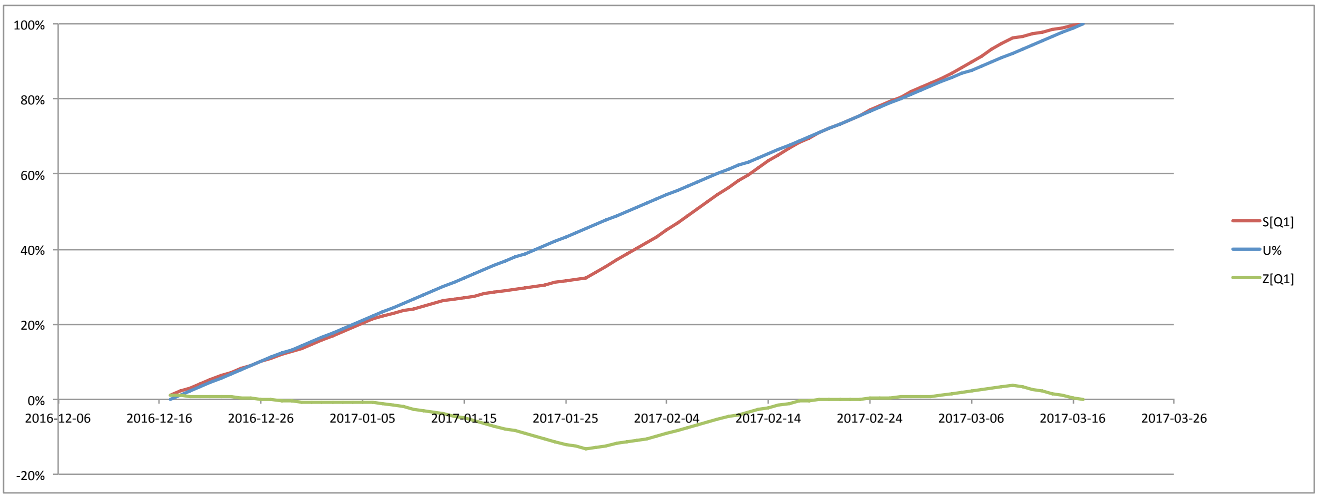

Figure 10: Plot of the dividend seasonality function $S[\text{Q1}]$, for the S&P 500 quarterly index for the period Q1 of year 2017 (red line), together with the normalised uniform dividend index $U_{\%}$ (shown as a blue straight line). For reference, as a green line, we also plot the rebased seasonality function $Z[\text{Q1}]$. (The plot uses the same data as figure 9.)

dividend-point index and keep in mind its reset-date convention.

2. We only consider the past period for which we have complete historical data (i.e. we do not consider period with partial data).

3. We determine the accrual index by computing the index increments and carefully handle the period boundaries.

4. We choose an alignment type and overlap the historical period accruals to the reference interval, filling the gaps or trimming the data appropriately.

5. If needed, we apply a smoothing method, such as the weekly smoothing technique explained earlier.

6. For each past period, we compute the reconstructed dividend index by taking the cumulative sum of the accrual index.

7. We normalise the reconstructed index by dividing it by its value at the end of the reference interval.

8. We compute the dividend seasonality function as the average of the normalised reconstructed index for each historical period.

Extra We compute the rebased seasonality function by subtracting the value of the uniform index to the dividend seasonality function.

3.4 Seasonality functions for major indexes

We show in this subsection the numerical results for the rebased seasonality function of some selected dividend-point indexes.

The Nikkey 225 dividend point index In figure 11 we plot the numerical results for the rebased seasonality function for the Nikkey 225 dividend point index. For the Nikkey 225 index since the rebased function oscillates from -15% to +25% we observe a very strong seasonality. Observing the rebased-seasonality-function shape we notice that the dividend payments are mostly concentrated in June and November. Also there seems to be a *last-minute* dividend payment right before the

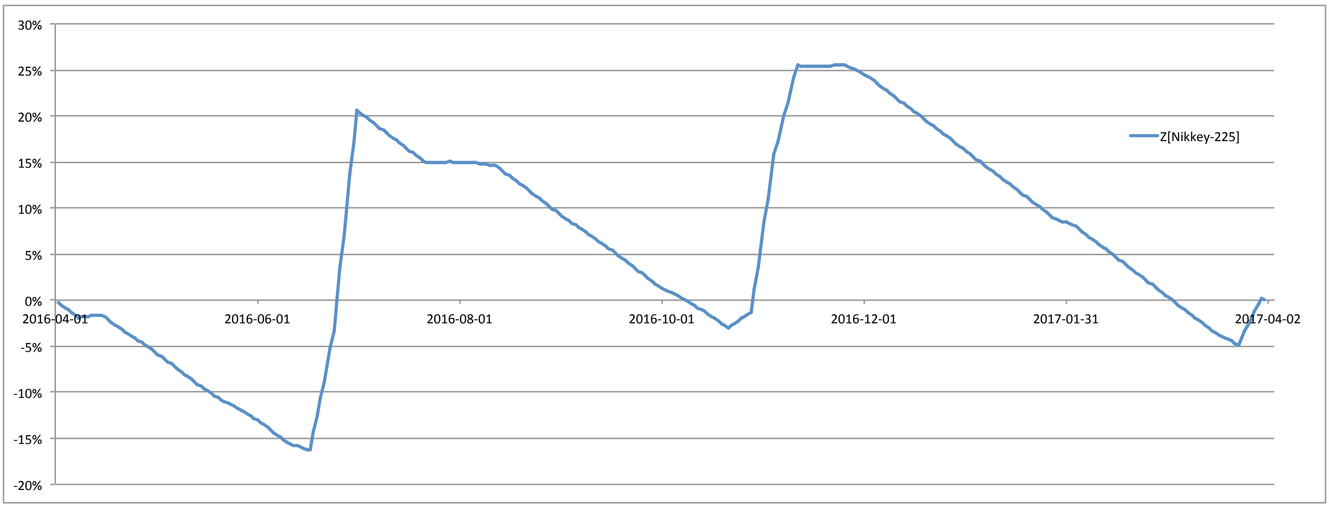

Figure 11: Rebased seasonality function, for the Nikkey-225 index, for the reference interval ending on April 1st, 2017.

end of the period at the end of March.

Indexes from different European markets In figure 12 we plot the rebased seasonality functions for the first six indexes listed in table 1. Since all these indexes reset on the third Friday od December their reference intervals coincide.

Note how strong is the seasonality feature for these indexes and that the seasonality function ranges from -20% to about +50%. Differently from the Japanese index, for these European indexes the bulk of dividend payments seems to be concentrated between the months of March and May. Note that when the rebased function has a gradient of -1 there are no dividend payments, hence for most indexes we can observe that there are virtually no payments after the month of July until the February of the next year. This stylised fact can be explained by the fact that, in these European markets, the dividend payments are based on the accounting results of the previous calendar year.

The S&P 500 quarterly index For the quarterly S&P 500 index the seasonality analysis is complicated by the fact that the index resets more than once a year. In figure 13 we show a plot of the four rebased functions, one for each quarter, for the S&P 500 dividend-point index. Since there are only about thirteen weeks in each period, we expect smaller oscillations of the rebased seasonality function. However the index seasonality can still be observed clearly for all four quarters.

Notice how the rebased functions for the different seasons resemble each other. They all seem to drop in the first few weeks, just to bounce back up soon afterwards, but then to drop again in the middle of the quarter. All functions reach their minimum value between -15% and -12%.

In order to properly investigate the similarity of the graphs for the seasonality function for the different quarters we overlap them on the same graph. In figure 14 we plot all four rebased functions on the same graph using $U_{\%}$ as the abscissa. This graph makes the similarity between the different quarters very clear. This similarity has the simple explanation that, for the North American market, differently from the other markets previously examined, the dividend cycle falls on a quarterly basis (and this is probably the reason for which we have a quoted quarterly index). Also, observing the same graph, we notice that most dividend payments for the companies

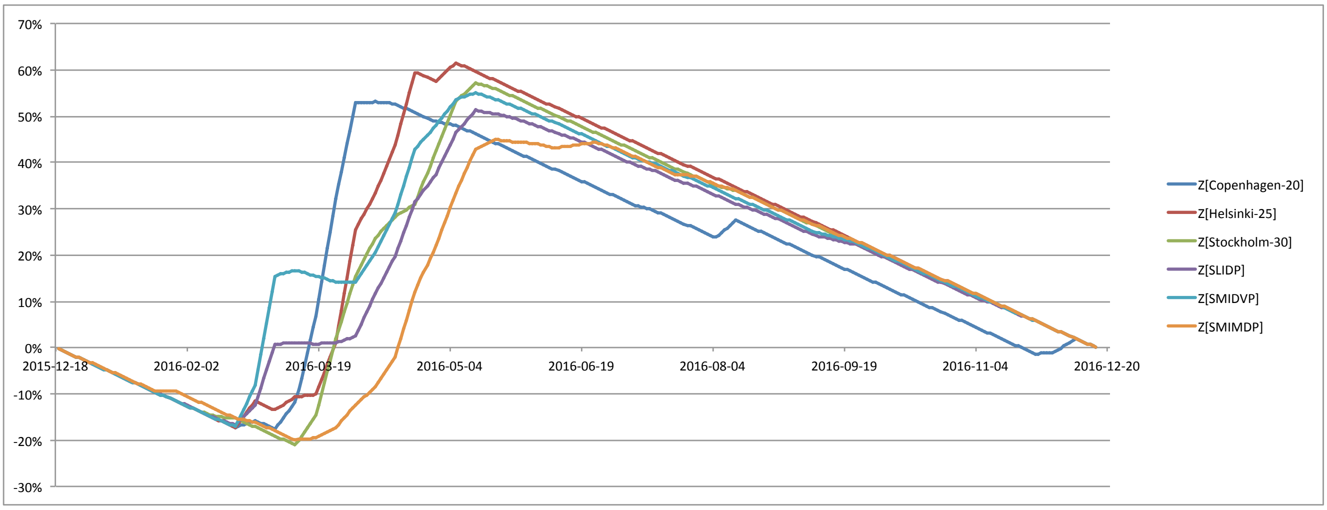

Figure 12: Rebased seasonality functions for the first six indexes listed in table 1. The reference interval is the year ending in December 2016.

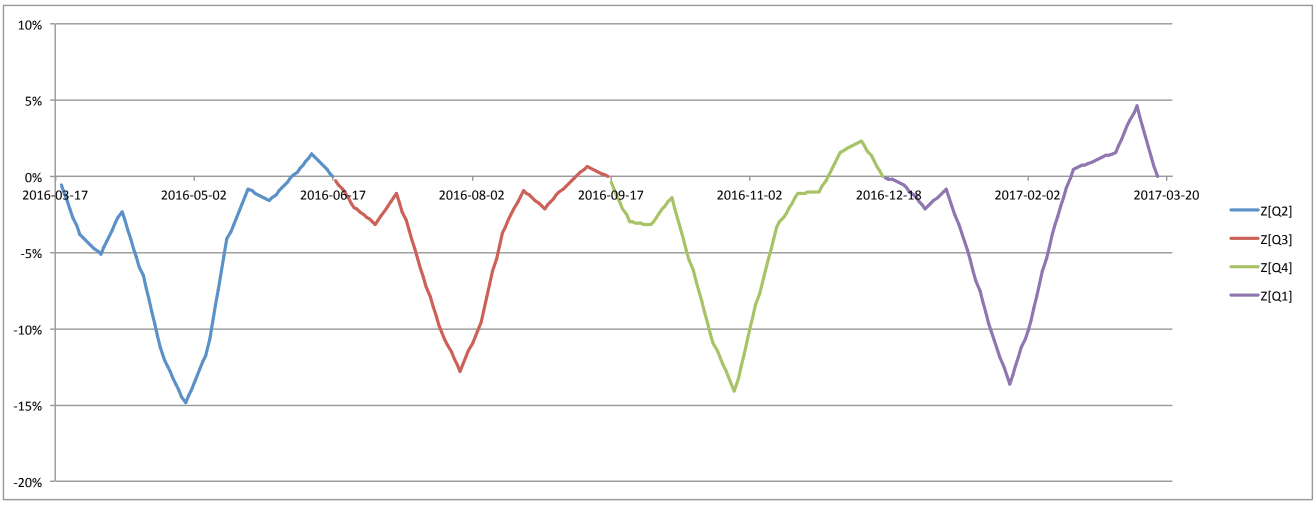

Figure 13: Rebased seasonality functions, for the quarterly S&P 500 dividend index, for the four quarters Q2, Q3, Q4, and Q1.

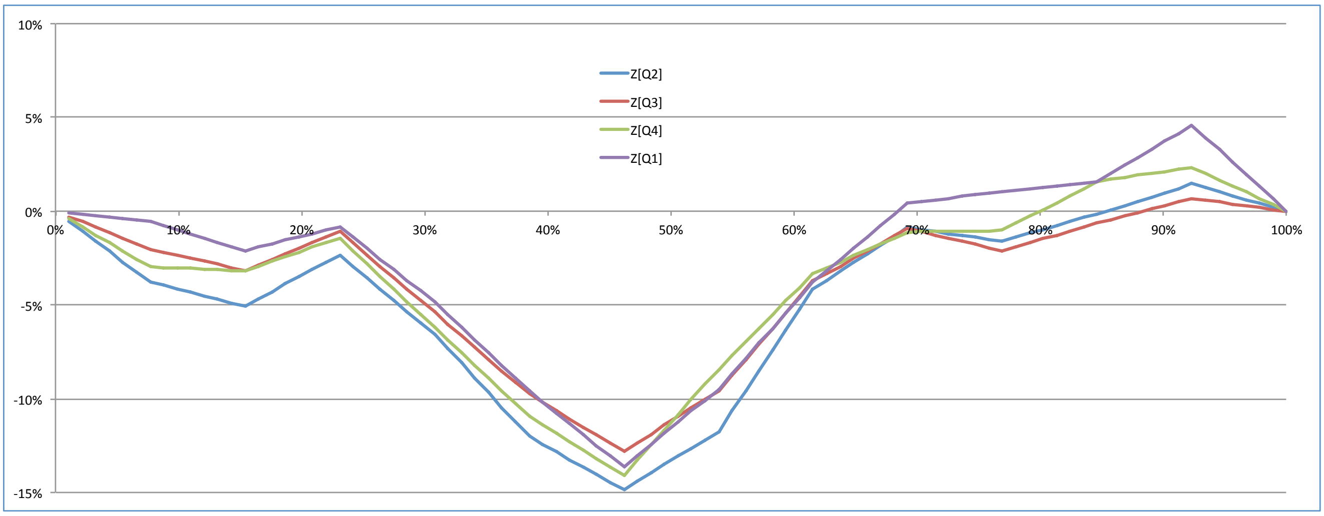

Figure 14: Overlapping of the rebased seasonality functions, for the quarterly S&P 500 dividend index, for the four quarters Q2, Q3, Q4, and Q1, plotted against the uniform index $U_{\%}$. Note that the results were computed for different reference intervals however all of them with a length of thirteen weeks.

in the S&P 500 index are made in the second half of each quarter.

4 Constant-maturity dividend futures

As shown in the previous section, there is hard evidence that the payment of dividends exhibits a seasonality pattern. This seasonality should be taken into account when building a term structure for the expected value of the dividend index. As mentioned in section 3 for pricing and risk-measurement purposes, given a number of futures quotes at fixed maturities we want to compute the equivalent quotes for constant-maturity futures. For simplicity, in the following derivation we assume a yearly dividend index. However, with little modifications, it is easy to adapt the derivation to handle the case of quarterly indexes.

4.1 The seasonality-function approximation

Consider a certain dividend index I. On the current date T_c we can observe the present and the past values of the dividend-point index, i.e. we know $I(t)$ for $t \leq T_c$. On the other hand, when the date t is in the future, the value of $I(t)$ is unknown and should be treated as stochastic. We assume that we can compute the dividend-index expectations with respect to the *risk-neutral measure* (see [1] for more details on the risk-neutral measure).

Suppose also that on the current date T_c we observe on the market the quotes of m dividend futures F_n, with $n = 1, \ldots, m$, each of which expires at T_n (with $T_1 \leq \ldots \leq T_n$). The dates $T_1, \ldots, T_m$ are loosely a year apart from each other (or three month apart for quarterly indexes). Depending on the current date, T_1 could be very close to T_c, or as far as one year in the future (three months for quarterly indexes). Also recall that, by definition, since the current period ends at T_1 we have

$$T_0 < T_c \leq T_1 .$$

Under reasonable assumptions, it can be shown that the quotes F_n are the expected value of the dividend index at the dates T_n, i.e. we can write

$$F_n = E[I(T_n)] = E\left[\sum_{i=1+T_{n-1}}^{T_n} d_i\right] \quad \text{for } n = 1, \ldots, m\,,$$

where the expectation is taken according to the risk-neutral measure. We want to compute the equivalent quotes on $m-1$ constant-maturity futures with expirations at $T'_1, \ldots, T'_{m-1}$ where T'_n is exactly n years after T_c. For convenience of notation we also set $T'_0{=}T_c$.

Let us denote with $F'_1, \ldots, F'_{m-1}$ the values of the (virtual) constant-maturity futures with maturities $T'_1, \ldots, T'_{m-1}$, defined as

$$F'_n = E\left[\sum_{i=1+T'_{n-1}}^{T'_n} d_i\right] \qquad \text{for } n = 1, \ldots, m-1\,. \tag{2}$$

Since we always have

$$T_n < T'_n \leq T_{n+1}\,,$$

we can split the sum in equation (2) in two parts:

$$F'_n = E\left[\sum_{i=1+T'_{n-1}}^{T_n} d_i\right] + E\left[\sum_{i=1+T_n}^{T'_n} d_i\right]$$

that, using definition (1), can be written as

$$F'_n = E[I(T_n) - I(T'_{n-1})] + E[I(T'_n)]\ . \tag{3}$$

The first term on the right-hand side of this expression is inside period n and the second term is inside period $n{+}1$.

Consider now the second term on the right-hand-side of equation (3), it can be written as

$$E[I(T'_n)] = \frac{E[I(T'_n)]}{E[I(T_{n+1})]} E[I(T_{n+1})] = X_{n+1} \cdot F_{n+1}\,,$$

with

$$X_{n+1} = \frac{E[I(T'_n)]}{E[I(T_{n+1})]}\,.$$

The fraction X_{n+1} denotes the expected value of total dividend paid by date T'_n with respect to the total amount of dividends paid in period T_{n+1}. Since period $n+1$ is completely in the future it is reasonable to use the dividend seasonality function to approximate X_{n+1}, so that

$$X_{n+1} \simeq \left\langle \frac{I'(T'_n)}{I'(T_{n+1})} \right\rangle = S_{n+1}(T'_n)\,. \tag{4}$$

In the *seasonality-function approximation* described by equation (4) we are computing a forward quantity X_{n+1} by using an historical average. We believe that this approximation is justified since X_{n+1} is completely into the future.

For the first term on the right-hand-side of equation (3) we need to distinguish two different cases: $n{=}1$ and $n \geq 2$. In case $n{=}1$ we simply have

$$E[I(T_1) - I(T'_0)] = F_1 - I(T_c)\,,$$

because of the additivity of the expectation, since $I(T'_0){=}I(T_c)$ is deterministic. In order to determine the first term on the right-hand-side of equation (3) when $n \geq 2$, we consider the following identities

$$\begin{aligned} E[I(T_n) - I(T'_{n-1})] &= E[I(T_n)] - E[I(T'_{n-1})] \\ &= F_n - X_n \cdot F_n \\ &= (1 - X_n) \cdot F_n \\ &\simeq [1 - S_n(T'_{n-1})] \cdot F_n\,, \end{aligned}$$

where X_n has been defined earlier and in the last line we have made use of the seasonality-function approximation (4). Gathering together all terms we can write the virtual quotes for the constant-maturity futures as

$$F'_1 = F_1 - I(T_c) + S_2(T'_1) \cdot F_2\,, \tag{5}$$

and

$$F'_n = [1 - S_n(T'_{n-1})] \cdot F_n + S_{n+1}(T'_n) \cdot F_{n+1}\,, \tag{6}$$

for $n \geq 2$.

Equations (5) and (6), together with the determination of the seasonality functions S_n, constitute the main results of this paper. Note that in equation (6), the weight given to F_n and that to F_{n+1} may not sum up to 1, because of the different expected seasonality functions of periods n and $n+1$, and because of the possibly different number of days in the two periods.

4.2 Conclusions

In this paper we examine the seasonality of dividend-point indexes for the current and future index periods. Looking at the numerical results of subsection 3.4, we observe that, for all indexes examined, the seasonality of dividends is not negligible. We find that for each reference period we should use a different seasonality function that can be obtained following the steps outlined at the end of subsection 3.3.

One limitation of the approach described in this paper is that, at least in principle, one should have available several years of history for the dividend index. As dividend point indexes currently have limited historical data, we expect to be able to compute seasonality functions with greater accuracy in the next few years.

As shown in section 4, the main practical application of the seasonality functions is the computation of a quote for virtual futures defined with constant maturity. In the work described by reference [2] (see specifically Appendix A), the authors obtain results similar to ours, with few differences. Firstly they use the same seasonality factor for all periods. Secondly they simply average historical data without carefully going through the procedure outlined in section 3. As a consequence we believe that the approach described in this paper provides a better framework for the pricing and risk management of dividend futures.

Finally, a future research work should show how to evaluate dividend derivatives, and possibly how to compute their risk, from the term-structure obtained from the virtual constant-maturity futures.

References

[1] John Hull. *Options, Futures, and Other Derivatives*. Pearson Prentice Hall, 9th edition, 2014. 15

[2] K. Jac, D.J. Frank, and D. Joost. The dividend term structure. http://ssrn.com/abstract=2542578, April 2015. 17

[3] Richard Manley and Christian Mueller-Glissmann. The market for dividends and related investment strategies. *Financial Analysts Journal*, 64, 2008. 1

[4] Marco Marchioro and Dario Cintioli. Foundations of the statpro simulation model. In *The Simulation Model*, Quantitative Research Series. StatPro website, October 2007. Permanent link marchioro.org/papers/foundation-simulation-model/. 5

[5] Scott Mixon and Esen Onur. Dividend swaps and dividend futures: State of play. http://ssrn.com/abstract=2579353, December 2015. 2