3 July 2020

Exabel

Alpha testing of credit card data

1. Introduction

Alternative data has recently gained much traction in finance due to the unique insights into companies' performance and macroeconomic effects one can derive from them. These data sets may, for example, be used to predict key performance indicators for companies or lead indicators of sector performance. In this paper we discuss another application, namely using alternative data to allocate funds in a portfolio.

Many alternative data sets are very noisy, and at best provide indirect indications of performance. One type of alternative data that stands out, however, is data collected from credit card transactions. This data can be used as a proxy for sales numbers, which is an immediate indicator of company performance, meaning that the path from data to value is shorter than for many other alternative data sets.

Having access to credit card data along with the ability to analyse it properly, allows an actor in the financial market to obtain early high-quality estimates of company performance and makes it possible to identify consumer trends early. To be able to extract as much value as possible from this high-quality data, it is essential to have the capability to perform statistical analyses on the data.

The purpose of this white paper is to demonstrate that there is real value in the transaction data by performing backtesting based on signals extracted from the data. In order to achieve this, we have to handle the marked seasonality inherent in transaction data in order to identify clear signals in the data, and then show that these signals give actionable insights into the performance of the companies in the data set. We do this in the following way: First, we outline the methodology we use for evaluating information content in the data. Secondly, we present a simple transformation of the credit card data, serving as our baseline signal, and show that this baseline signal has value according to that methodology. Thirdly, we analyse simple variations of this signal and show that the signal is robust to these variations, and in some cases improve the performance. The credit card data used in this white paper comes from our alternative data partner, **1010data**, which is one of the leading vendors in this space.

All analyses in this paper are performed using the **Exabel** platform, which makes it easy to perform the required analyses in a way that does not require the user to have a background in statistics.

2. Method

In this section we describe the methodology we use to demonstrate that the credit card transaction data has real value, in that it can be used to produce *actionable* insights which can help with driving real investment decisions. This consists of two steps: We first first analyse the data and extract a clear signal from it, and then use that signal to simulate buy and sell decisions in a way that hopefully produces profit. If the signal appears to be robust and performs well in such a simulated scenario, it is an indication that the signal has real information content and can assist in making investment decisions or building portfolios.

The first challenge we face is that data collected from real-world sources generally has much noise, and we have to peel away the noise to identify signals that can be given a meaningful interpretation. In the case of credit card data such noise can consist of weekly and seasonal fluctuations and random variations in consumer behaviour, influenced by a host of different factors.

This signal is a proxy for sales numbers; however, our objective is not just to estimate the sales numbers, but to use the signal to take advantage of price movements, and thus make profitable investment decisions.

In order to determine if a signal has profit potential, we simulate a trading strategy by evaluating the signal on historical data. We pick a reallocation frequency (e.g. daily, weekly or monthly) and for each date for which we want to reallocate the assets, we perform the following calculations:

1. Use the signal to rank all the companies.
2. Simulate a strategy where we buy or short one dollar worth of assets based on one of the following methods:
    a. Buy the top *x*% and short the bottom *y*% of the companies.
    b. Adjust all the signal evaluations so that the mean is zero, and allocate the assets according to the magnitude of the signal, with negative values indicating shorting.
3. Calculate the returns the given allocation strategy gives.
4. Repeat steps 1-3 for all allocation dates.

The simulations are made using *equal bet size*, meaning that one unit of funds is invested each day allocations are made, without reinvesting the returns from the previous allocation. A strategy generating, e.g., 30% each year will consequently appear as a linear growth in our figures, whereas if one reinvested the returns, one would see exponential growth.

In the remainder of the paper we will identify examples of such signals and demonstrate that the simulated trading strategy yields profit when used on historical data. Note that perceived value from simulating on historical data does not guarantee that the same strategy will work in the future; company returns are highly non-stationary time series and many factors may change their future behaviour.

3. Description of data

The credit card data we are looking at is collected from a panel of consumers from the United States and aggregated to a company level. The aggregated data is available with daily time resolution.

1010data does several steps of preprocessing of the credit card data in order to give the user a consistent view of the data. One example of such a preprocessing step is to adjust the credit card spend for acquisitions, so that the data view is consistent before and after mergers and acquisitions. In this analysis we have access to credit card data for around 300 companies in the US.

An example of how credit card data looks is given in Figure 1, for Amazon. Here we see clear seasonal patterns and holiday effects, and a growth trend, which is common for this type of data.

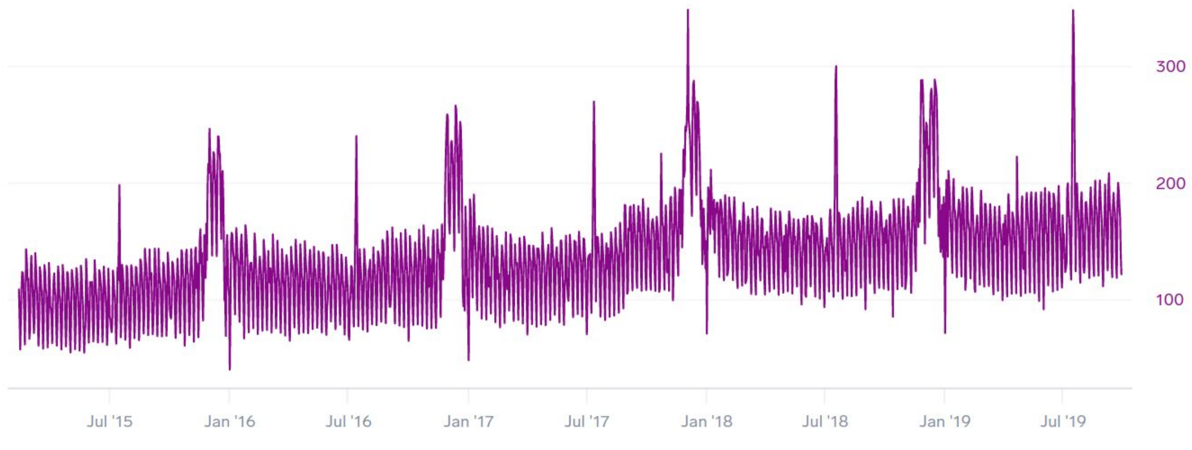

Figure 1: Acquisition-adjusted credit card spend for Amazon.com, Inc.

4. Baseline signal

The first step is to find a baseline signal in the data, which we can use in the initial analysis. Due to the noise and the pronounced periodicity in the data, we cannot expect the raw untransformed data to give us good performance. We therefore have to apply certain transforms to the data to remove the noise and irregularities and extract a clear signal.

Our purpose for this section is to extract a relatively simple signal which illustrates the power of the transaction data. It turns out that the minimum we have to do to get a workable signal is to handle the seasonal variability in the data. In this section we therefore create a signal which transforms the raw data in a way that, to a large extent, cancels the seasonal effects. In the next section we apply the backtesting algorithm of Section 2 to the signal developed here to demonstrate that the signal can yield real value. However, seasonality is not the only issue we face when analysing transaction data. We therefore suggest other transforms in Section 6, designed to tackle other issues, seeking to improve this relatively simple baseline signal.

In credit card data we encounter different types of periodicity, which depend on consumer behaviour. The first one is weekly effects, as consumers' behaviour varies a lot depending on the weekday. A simple way to counter this is to aggregate data on a weekly basis, typically by using seven-day moving averages.

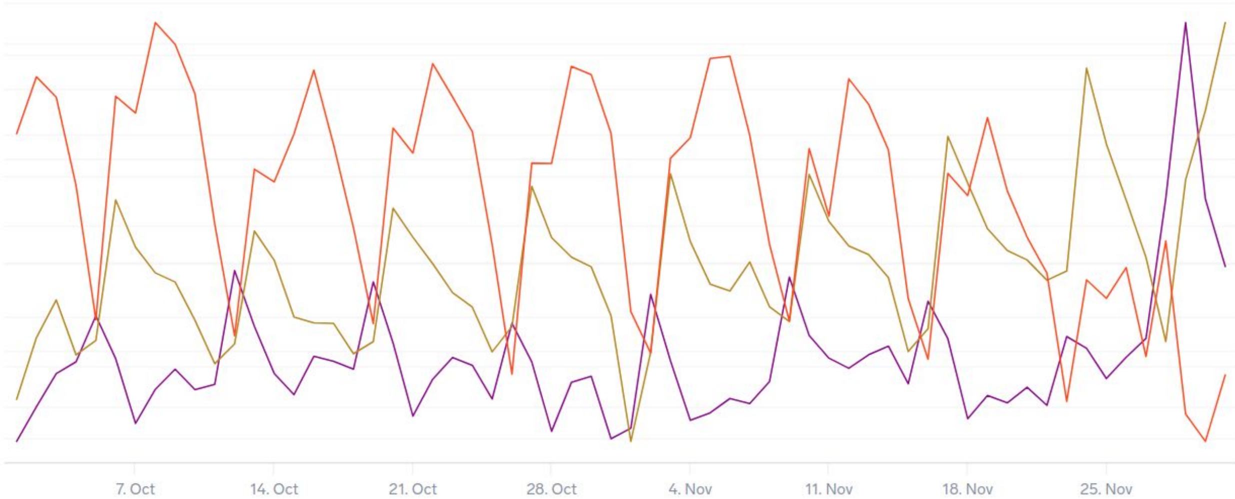

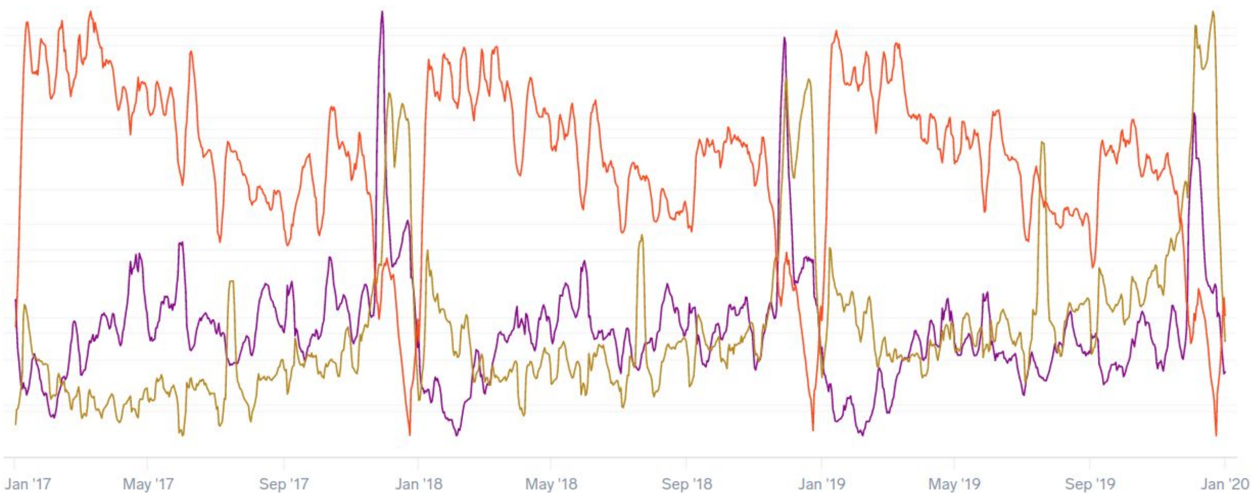

Figure 2: Top: Illustration of weekly effects for the apparel retail sector (purple), consumer non-cyclicals (yellow) and passenger transportation (red). Bottom: Illustration of seasonal effects for the apparel retail sector (purple), consumer non-cyclicals (yellow) and passenger transportation (red).

The second kind of periodicity is the seasonal effect varying with the time of year, which is more insidious as there are several different factors at play. One important thing to notice is that seasonal effects in credit card data are sector specific: for example, the department store sector and other sectors selling goods to consumers, generally have spending peaks near the Christmas holiday, while the travel sector has a spending drop during this period. See the bottom part of Figure 2 for an illustration of these effects in three selected sectors. Because of this it is easier to compare companies within the same sector.

A natural way to counter this kind of periodicity is to use a year-over-year transformation of the raw data. Combining these two solutions (weekly and annual seasonality), we take our baseline signal to be a year-over-year comparison of seven-day moving averages. In formulas:

E1.
$$y_i^k = \sum_{j=0}^{k} x_{i-j}$$

E2.
$$z_{i,k} = \frac{y_i^k - y_{i-365}^k}{y_{i-365}^k}$$

where x_i denotes the credit card spend for a specific company at time *i*. A plot of this signal for Amazon and Netflix is given in Figure 3.

A third kind of seasonality effects are public holidays and other special days, some of which are observed on the same day every year, and some of which move from year to year. Holidays with fixed dates, such as Christmas Day and 4 July are adequately handled by the above transformations. Holidays and other special days that are observed on different dates every year, such as Easter and Black Friday, can still cause spikes in the data, and to perfectly counter these effects, more clever transformations must be done.

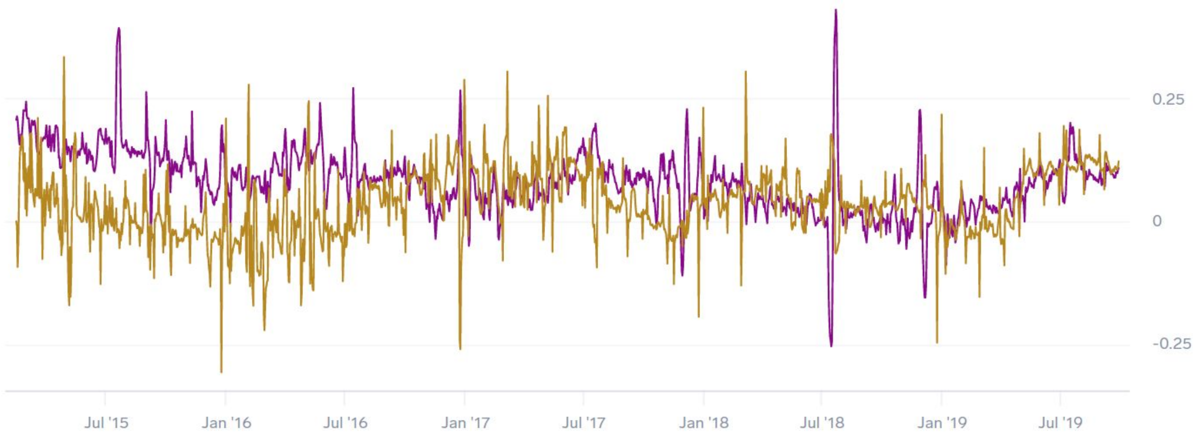

Figure 3: Year-over-year relative change in credit card spend for Amazon (purple) and Netflix (yellow), using a moving average window of 7 days.

The above simple transform is, however, sufficient for demonstrating the value of credit card data, so we will not delve into the technicalities necessary to handle these special calendar effects. Nevertheless, when attempting to maximize the profit reaped from transaction data, such effects should also be taken into account and handled appropriately.

5. Evaluating the performance of the baseline signal

Having identified a baseline signal, we use the backtesting methodology outlined in Section 2 to demonstrate that it has real information content.

We perform the analysis both for the entire universe of companies and for separate sectors. The reason for analysing individual sectors is twofold. First, we expect the transaction data to have varying predictive value in different sectors. Secondly, if the performance is comparable across many sectors, it suggests that the results are more robust, giving us more confidence that they will hold up when applied to unseen data, such as when making future predictions.

In each scenario we calculate the return using both strategies described in Section 2: (a) a long–short strategy where we buy the top x% and short the bottom y% of the companies, and (b) a proportional strategy where the signal evaluations are centred and assets are allocated according to the magnitude of the signal.

All companies

Figure 4 shows the returns for the baseline signal on all the companies combined, using the long–short strategy. The return path here clearly shows a strong performance, indicating that the signal possessed significant value during the backtesting period.

Figure 4: 10% long, 10% short seven-day moving average, daily re-allocation, all companies available. Sharpe ratio: 2.77

In addition to looking at the simulated returns, we can visualise the relation between credit card spend and returns by looking at the accumulated returns relative to the rank. That is, we rank the companies according to the signal values, and plot them together with the returns. If the signal has predictive value, then this curve should be noticeably decreasing. In Figure 5, this plot is shown for year-over-year credit card data. As we can see, there is a clear decreasing trend, with an almost linear relation between the rank and the average returns. The variability in the ranks can, most likely, be attributed to volatility in the returns and the signals.

Figure 5: Accumulated returns by rank.

In Figure 6 we perform the simulation, using the proportional allocation strategy rather than the long–short strategy. The simulation for this strategy shows higher volatility than for long–short. Thus, when regarding all the companies together, the long–short strategy appears to be the better one. However, we expect that allocating funds proportionally to signal strength will work best for companies that are comparable to each other. Since the relation between credit card transactions and returns varies widely among sectors, we expect this strategy to perform better when used on an individual sector, which we will explore below.

Another interesting point is that the return paths here also indicate that there may be a marked change in the second half of 2018 in how the market reacts to the credit card data. This is, however, difficult to interpret, as all companies here are treated as "equal" and no breakdowns into sectors or other interpretable components are available. Another possible interpretation is that the panel for certain subsets of companies are much more representative for total credit card spend, and for other subsets there is more noise.

Figure 6: Comparison of proportional allocation to 10% long, 10% short allocation, 7 days moving average, daily re-allocation, all companies (259 in this case). The purple line shows proportional allocation. The Sharpe ratio is 2.12 for proportional allocation and 3.06 for long–short

Individual sectors

We proceed by investigating how the baseline signal performs on subsets of the company universe. In particular, we select subsets where we expect the companies within the subsets to be more comparable to each other than when looking at the entire company universe. We choose two RBICS sectors for this investigation.

The first example is the *apparel retail* sector, where the proportion of revenue coming from credit card transactions is comparatively large. For a long-short simulation of this sector given the baseline signal, see Figure 7.

Figure 7: 10% long, 10% short, 7 days moving average, daily re-allocation, 28 companies in the apparel retail sector. Sharpe ratio: 3.35

The simulation results for the proportional strategy yield approximately the same total returns, albeit with lower variance, which results in a larger Sharpe ratio (see Figure 8). Compare this with Figure 6 above where we performed the same simulation for all companies for which we have credit card data.

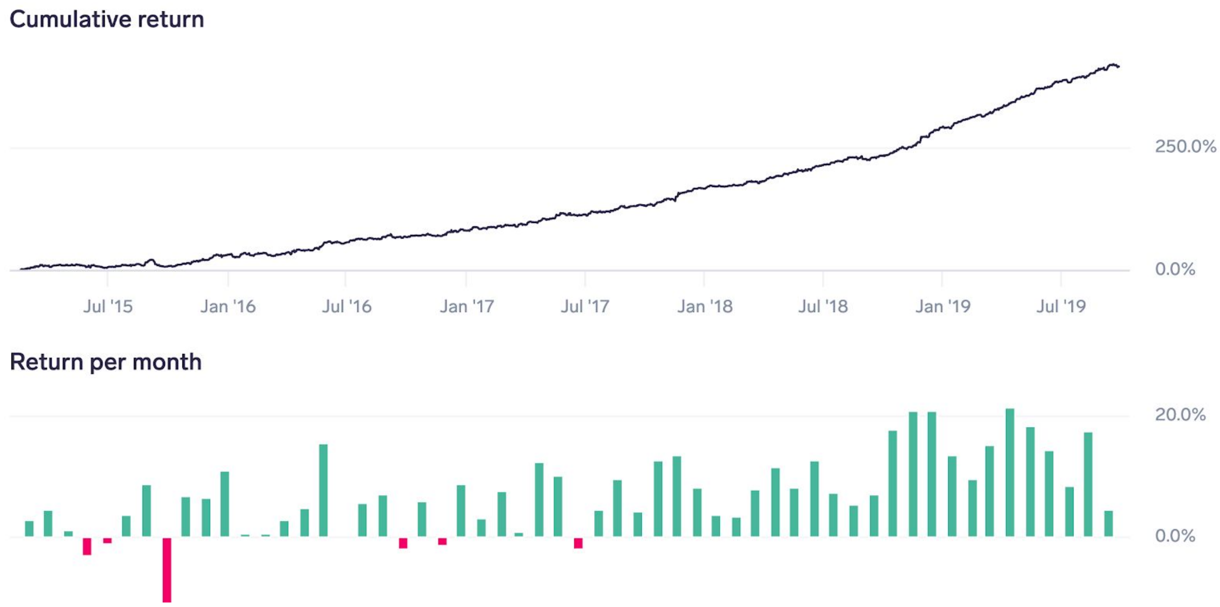

Figure 8: Signal proportional allocation, 7 days moving average, daily re-allocation, the apparel retail sector (28 companies). Sharpe ratio: 3.94

In Figure 9 we see the return paths for three simulations performed for the *business services* sector. In this sector we would expect the credit card transaction data to have less value, because the revenue for these companies cannot easily be traced back to credit card transactions. In addition, the revenue streams may not be as comparable across companies as they are, for example, in the retail sector.

Figure 9: Return paths for 30% long-short (purple), 20% long-short (red) and proportional allocation (yellow). Sharpe 1.23/1.25/1.04 correspondingly.

6. Modification of the baseline signal to handle difficulties in the data

In Section 4 we developed a relatively simple baseline signal, and the analysis in Section 5 showed that the signal is a strong candidate for devising a trading strategy. However, it is worth exploring variations and modifications of this baseline signal, to see if it can be improved upon and to show that it is robust to relatively small changes.

In this section we explore some of the possible modifications to illustrate some of the issues one has to take into account when dealing with alternative data. Note that a final trading strategy may employ a variety of such modifications, and the performance may be different for the various sectors or over time.

Lag

Alternative data sets are seldom provided in real time, so there is almost always a delay between the events the data describes and the time when they are available to use in analyses. This may be due to the data collection procedures, or data post-processing, which may be required to ensure high data quality. In some cases the data provider may offer data with different delays for different prices.

We naturally expect that the fresher the data is, the more valuable it is. It is a good idea for the consumer of alternative data to evaluate the value of this delay, either to …

In Figures 9 and 10 we show simulated returns using data with different delays, which illustrate that longer delays yield smaller returns.

With a lag of 14 days, it seems like that the data may not be fresh enough to give a huge advantage. With a delay of 3 days, the picture is less clear. Here, the sharpe ratios are a bit lower than with no delay, but the annualised returns are higher.

Figure 9: Return path for z_{i7} with delay 0 days (purple), 3 days (yellow), 6 days (red) and 14 days (teal). Subset: Consumer Cyclicals. Sharpe ratios 3.43/3.34/2.59/1.11 respectively.

Figure 10: Return path for z_{i7} with delay 0 days (yellow), 3 days (purple), 6 days (red) and 14 days (teal), all companies. Sharpe ratios 3.54/3.44/2.64/1.66 respectively.

Moving averages

Another parameter that can be varied is the number of days to use in the moving average. As we recall, one important reason for using a moving average in the baseline signal is to cancel the weekly seasonality, which is essential for credit card transaction data. Thus, in this case, while we may vary the number of days, this is a strong argument that the period should always contain a whole number of weeks. For other alternative data sets, this may or may not be an important consideration.

Apart from dealing with seasonality, there are other factors that call for a shorter or longer moving average period. As we discussed above, the value of the alternative data typically decreases the older it gets, which calls for using a shorter period.

On the other hand, using a shorter window makes the analysis more sensitive to noise. This is illustrated in Figure 11, where we see that using a longer window for the moving average makes the curve significantly smoother. Consequently, there is a larger chance that a movement is caused by random fluctuation when using a shorter window. When using a larger window, on the other hand, we risk taking too long to detect actual changes in consumer behaviour, which may lead us to miss out on an opportunity to profit from the data. In Figures 12 and 13 we show that changing the moving average window has a clear effect on the return paths.

Figure 11: The effect of increasing the size of moving averages. In E1, k=7 (yellow) and 56 (purple). Example company: Apple, Inc.

Figure 12: Return paths when increasing the moving average window, all companies. In E1, k=7 (purple), 14 (yellow), 28 (red), 56 (teal). Sharpe ratios: 3.54/3.60/3.00/2.55 respectively.

Figure 13: Return paths for $z_{i,k}$ with k=7 (purple), 14 (yellow), 28 (red) and 56 (teal). Subset: Consumer cyclicals. Sharpe ratios 3.54/3.57/3.08/2.54 respectively.

In these examples, a 14-day moving average seems to be slightly better than using 7 days, but it is not clear if this difference is significant. An interesting point to note is that there appears to be a change in the returns in the late 2018, after which 7 days seems to be a better choice. An explanation for this may be that the market has become more reactive as credit card data has become more widely available.

If the market has indeed become more reactive, this can be an argument for using a window which is even shorter than seven days. Doing this means reintroducing the problem of weekly periodicity, in which case dealing with this periodicity must be done in a different manner, for example by explicitly modelling weekday effects.

Volatility adjustments

The volatility of the year-over-year credit card data varies between companies, which necessarily also varies the variability in the returns for the various companies. To reduce the risk of our trading strategy, it can therefore be a good idea to downweight companies for which there is higher volatility in the credit card data.

A simple way to do this is to adjust the z_i,k by dividing it with its rolling standard deviation. (That is, the standard deviation of z_i,k in a period immediately preceding the current time point.)

We present one return path for this configuration using a proportional allocation, showing that such adjustments may be necessary to produce a satisfactory return path.

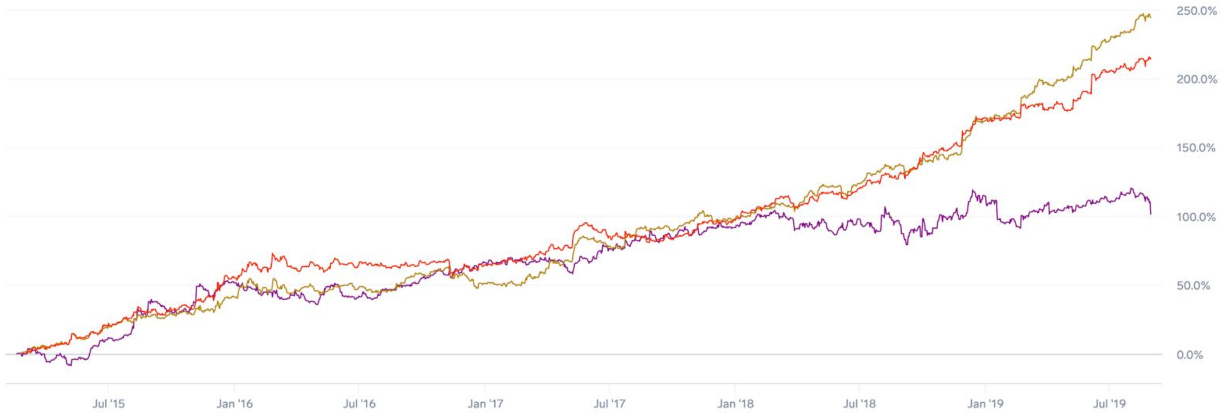

Figure 15. Return paths for consumer cyclicals (proportional allocation). Moving average 14, with no volatility adjustments (purple), 7 days moving standard deviation (yellow) and 14 days moving standard deviation (red). Sharpe ratios: 0.84/2.98/2.65

Rebalancing intervals

So far in the analysis we have simulated a trading strategy with daily rebalancing. While this allows us to always utilise the most recent data in our trading decisions, it may not be practical to rebalance a portfolio every day, and because there are always transaction costs involved in rebalancing, this may drain any profit we stand to make.

We therefore also simulate trading strategies with different rebalancing intervals. In Figure 15 three different rebalancing frequencies are shown: daily (red), weekly (yellow) and monthly (purple) with 1143, 234 and 52 rebalances respectively. The figure shows that even with monthly rebalancing, the data has reasonable information content, but the sweet spot from a practical perspective is probably closer to weekly rebalancing. Figure 17 supports this claim, where the same signal as for Figure 16 is used for simulating the return path, the only difference being that each trade is assumed to have a 10bps cost.

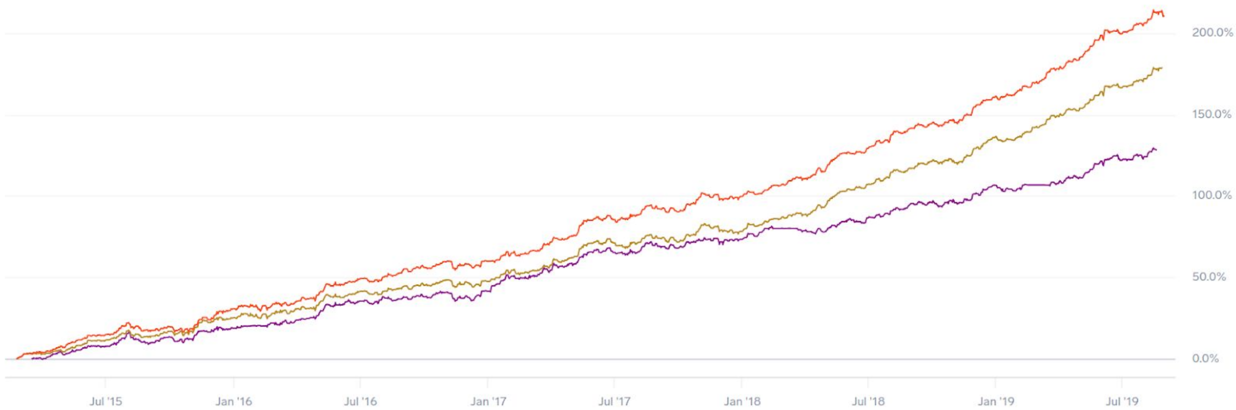

*Figure 16. Return paths for all companies long-short 20%. Moving average 14, with 7 days volatility adjustments. Daily- (red), weekly- (yellow) and monthly rebalancing. Sharpe ratios:*3.82/3.51/2.68

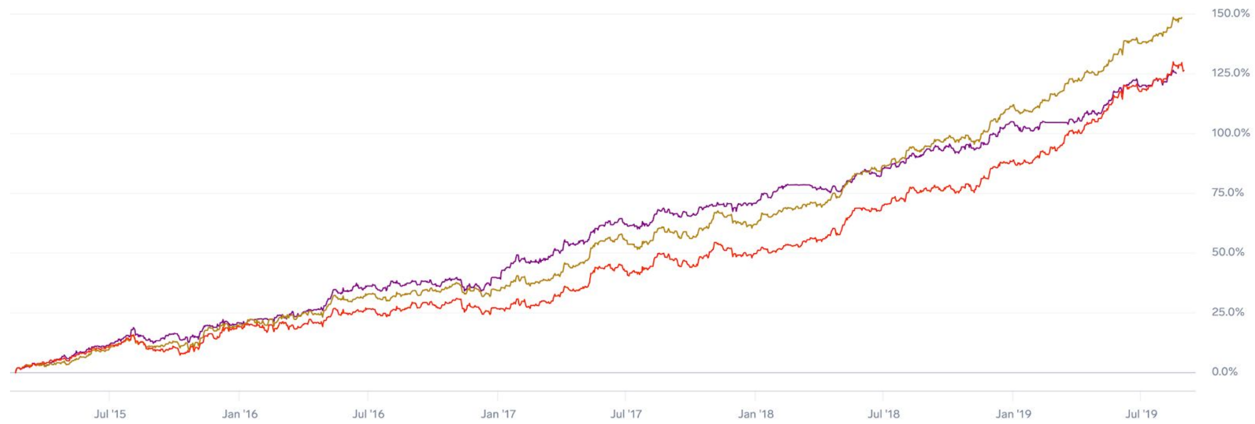

*Figure 17. Return paths for all companies long-short 20%. Moving average 14, with 7 days volatility adjustments, and 10bps cost per trade. Daily- (red), weekly- (yellow) and monthly rebalancing. Sharpe ratios:*2.43/2.89/2.66

7. Conclusion and future work

In this paper we have illustrated that there is clear value in credit card transaction data, and that using a simple transformation of the raw data is adequate for providing a baseline signal that can aid in making investment decisions. Furthermore, we have shown that the baseline signal is robust to changes, and that using the signals on specific subsets of data may be a better option than regarding all companies as equal. We note that the transformations we have used are simple and in their simplicity also intuitive. In practice, one would adjust these signals to give even better performance or to address certain investment aspects.

It is possible to resort to machine learning techniques to improve on the transforms we use for testing information content. Here, we used simple heuristics to define the time series we want to use to rank and allocate on, while it is natural to consider using machine learning techniques to estimate expected returns (alternatively relative expected returns) given credit card spend.

The backtesting strategies we use here are also simple heuristics to test for information content in a signal. If one were to use the signals for real trading, other considerations would apply, and a more sophisticated portfolio building or screening strategies would be necessary. Fortunately both these tools are being integrated into the Exabel platform, facilitating easy use for investment professionals.