A Machine Learning Approach to Spinoff Investment Optimization

Mehdi Chouiten and Romain Ekert

***Abstract*—A spinoff is an event that consists of the creation of a new company based on an existing division of a mother company. Often, this event is due to a new strategic vision of the mother company or to unsatisfying financial performance of the spun off division. In our study, we try to capitalize on several features to measure their impact on spinoffs success or failure and thus, build a predictive model that allows us to select best spinoffs to invest in. Our method aims to predict the stock price performance over different time horizons: 6, 12,18, and 24 months. Allowing profitable exits to investors (either stock traders or option traders). Using mainly Bloomberg platform data, we compared several machine learning algorithms (SVM, Random Forest, Gradient Boosting) and different methodologic approaches (Binary classification, time-series classification, multi-class clustering) to build an efficient, yet improvable, model to reach our goal.**

***Index Terms*—Spinoff, investment, machine learning, modeling, option trading, time-series classification, stock prediction.**

I. Introduction

In this study, we looked at financial events called spinoffs, which consists of the creation of an independent company through the division of a parent company. Spinoffs are mostly performed by businesses wishing to streamline their operations by selling less productive or unrelated subsidiary.

There is little consensus as to whether firms that find themselves spun off from other companies – either as new, standalone companies, or under the stewardship of new parent companies – perform better or worse than they did before. It is linked to the reason of the spin-off and past performance evolution [1], [2]. In 1992, Woo, Willard, and Daellenbach [3] found that, on average, the performance of divested units after the spin-off does not improve compared with the three years before divestment. But in 1999 Desai and Jain [4] found that long-run performance of both the former parent company and the divested unit is strongly positive, provided that the spin-off increases the company's focus. A 2010 meta-analysis from Lee and Madhavan [5] detailed many of the different issues that make divestiture so hard to evaluate consistently.

We assumed that it is possible to find patterns in the post-spinoff evolution of the companies' stocks prices and that some fundamental characteristics from the companies involved will allow us to predict this evolution. We aimed to use the research to design a spinoff oriented investment strategy based on both financial research (like in [6] and [7]) and machine learning.

II. Data

Our data consist of historical spinoffs that occurred between 1990 and 2016.

We retrieved those data as an excel file from the Bloomberg software, selecting only spinoffs events that have actually been completed. When available, we also enriched or updated the data from other data sources.

For each spinoff in our database, we have information on its operating industry, financial ratios and indebtedness for both the parent and the new company (NewCo) resulting from the spinoff. We also have the NewCo's monthly share price history during the first four years following the spinoff completion date.

Examples of the features taken in consideration are EBIT Margin, Profit Margin, Total Debt, Debt to Shareholders Equity, Spinoff Date, Horizon of prediction and more than 50 financial ratios.

III. Approach

Our goal is to predict whether a spinoff is likely to succeed or not, measurement of success being the NewCo's stock price evolution over the horizon of 6, 12, 18 and 24 months.

First, we intended to predict if the price had gone up (or not) after 6 months (12 months, 18 months, …) considering all information available at completion date, resulting in a binary classification problem with the probability of belonging to each class as our output.

But we soon realized that this output could not be used as an investment strategy. Indeed, if we managed to predict an increase in the share price after six months, we have no guarantee that it will stay this way during the following months.

We then intended to predict the NewCo's share price behavior over a specific time period following the spinoff.

To do so, we first need to perform a clustering of the post-spinoff stock price time series, focusing on their shapes and more specifically on the months where their extrema are reached.

Manuscript received May 23, 2018; revised July 20, 2018.
Mehdi Chouiten is with Datategy SAS, France (e-mail: mehdi.chouiten@datategy.net).

Once we achieved the clustering task, each spinoff is assigned to one of our clusters (defined here as a type of evolution) so that we moved to the classification task where we used our features about the parent and the NewCo at spinoff completion date in order to predict in which cluster a NewCo's share price evolution belongs to.

The post spinoff time period over which we will predict the stock price evolution type is set as a parameter of our model, depending on how far in time we want to predict and also because our historical data have different post spinoff learning time period available (spinoffs that happened in 2016 have a maximum of 12 months learning period).

This approach allowed us to use our output as an investment strategy. For example, spinoffs with their minimum share price being at completion date, 6 months or 1 year after induce strategies where one should invest respectively at completion date, after 6 months or 1 year; conversely spinoffs with their maximum at completion date should not be included in our portfolio.

The output being, for each spinoff, the probability of belonging to a cluster, each cluster corresponding to a specific trading strategy.

IV. Binary Classification

To do rapid prototyping, we first simplified our approach, skipped the clustering task, and moved to a binary classification problem. Class 1 being where NewCo's share price is minimum at completion date over the post spinoff period (representing spinoffs that we are more willing to invest in at completion) and class 0 being all other types of evolutions.

To do so, we compared several algorithms used for classification, such as SVM (Support Vector Machines), Random Forest and Gradient Boosting, using a GridSearch to find the optimal set of parameters, one of them being the weights assigned to each class as spinoffs belonging to class 1 are under-represented in our data. This is a classical manner to handle unbalanced classes in datasets.

We measured our performance using the ratio of well predicted spinoffs belonging to class 1 over the total of predicted spinoffs belonging to class 1 since our investment strategies will only be based on spinoffs belonging to this class.

The selected model was the Random Forest with the following set of parameters: a forest composed of 1000 trees, a minimum of 2 samples per leaf, the squared root of the total number of features as the maximum number of features to consider when looking for the best split at each node of a tree, and a weight of 2 on samples belonging to class 1.

In a horizon of 36 months, our model has an average performance of 0.81. It predicted on average that 1% of all spinoffs belonged to class 1 (instead of the actual 15%) but with a performance of 0.81.

To be more clear, if our test set is composed of 100 spinoffs, 15 of them actually belong to class 1 and our model will predict that only 1 spinoff belong to class 1 but with a probability of 0.81 to actually belong to this class.

In comparison, the SVM algorithm managed to predict the 15 spinoffs belonging to class 1 but it also predicted 50 more spinoffs that actually belong to class 0, resulting in a performance of 0.3.

To improve our performance, we used for each sample the probability of belonging to each class, and kept only the highest one as a valid prediction, fixing the threshold as the minimum between the number of well predicted class 1 spinoff and half of the total number of predicted class 1 spinoffs. This resulted in a performance of 1.0 but lowered the number of predicted spinoff to 0.3 in average which is less than 1, therefore this algorithm will rarely predict that a spinoff belong to class one but when it does, you can be sure that it is true.

V. Multi-Class Classification

Then we moved to the multi-class problem where each class is defined as a type of evolution resulting from an unsupervised clustering of the share prices time series.

A. *Time Series Clustering*

Time series clustering is to partition time series (as shown in Fig. 1) data into groups (as shown in Fig. 2) based on similarity or distance, so that time series in the same cluster are similar. The aim is to identify categories regrouping time series with a common “behavior”.

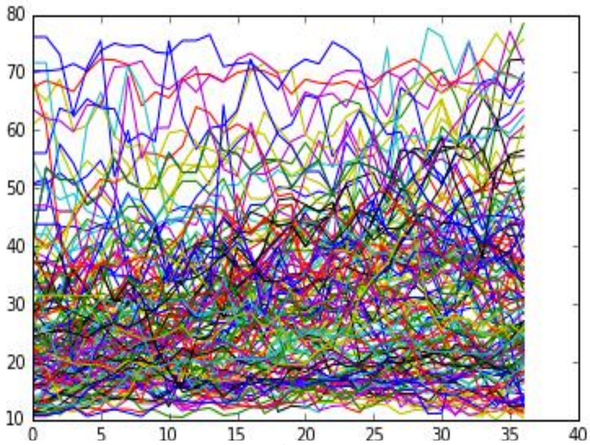

Fig. 1. Plot of all time-series.

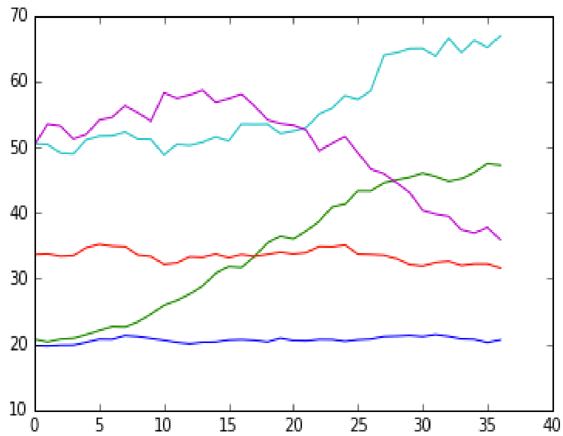

Fig. 2. Plot of clustered time-series.

The first step is to work out an appropriate distance/similarity metric, and then, at the second step, use existing clustering techniques, such as k-means or hierarchical clustering, to find clustering structures [8].

1) Raw-data-based approach

The Dynamic Time Warping (DTW) is a popular similarity measure between time series [9], [10] even if it fails to satisfy the trianle inequality and its computation requires quadratic time. We can avoid most DTW computations with an inexpensive lower bound (LB Keogh).

Unfortunately, using the DTW distance in our clustering task will group time series that are similar in shape but regardless of variations in the time dimension, which is not acceptable in our particular case, therefore the Euclidian distance is preferred since it is very sensitive to distortion in time axis [11].

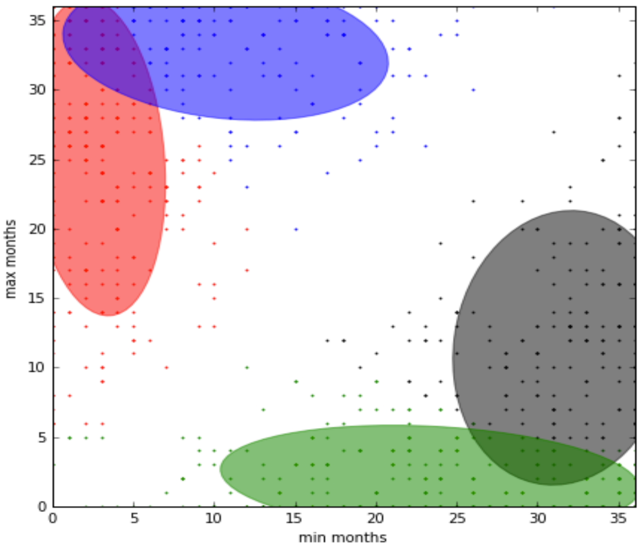

Fig. 3. a. K-Means illustration (case of 4 clusters).

Then we used the K-means (Fig. 3.a. and Fig. 3.b.) algorithm to perform the clustering task. Since the number of clusters is required, we used the silhouette score to choose the optimal number of clusters [12].

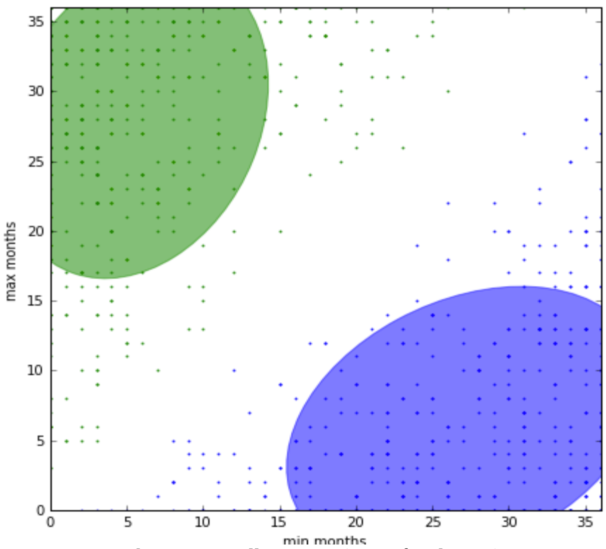

Fig. 3.b. K-Means illustration (case of 2 clusters).

2) Feature-based approach

A different approach to this clustering problem is to represent each time series by a feature vector of lower dimension, in our case the months where the minimum and maximum share prices are reached over the post-spinoff period, in order to perform classical 2-dimensional clustering using Gaussian Mixtures and choosing the number of clusters that minimize the BIC score.

B. *Multi-class Classification*

Once we performed the clustering task, each spinoff is assigned to one of our clusters (defined here as a type of evolution). We now moved to the classification task where we used our features about the parent and the NewCo at spinoff completion date in order to predict in which cluster a NewCo's share price evolution belongs to.

The first approach to multi-class classification relies on extending binary classification problems to handle the multi-class case directly. The second approach decomposes the problem into several binary classification tasks. Various methods are used for this decomposition such as one-versus-all, all-versus-all and error-correcting output coding [14].

VI. Conclusion

This study confirmed that a spinoff induces specific patterns in the evolution of the spun off company's share price and that it is possible to predict this evolution based on information available at the spinoff completion date.

Our focus was mainly on the spun off company's share price evolution but it is very likely that the share price of the parent company also present predictable evolution patterns. Furthermore, our model did not include macroeconomic factors such as general economic conditions, performance of financial markets, competitive factors, changes in laws and regulation as well as other firm specific characteristics such as team management and ownership. Those factors should be of great impact on the predictive power of our model and represent a perspective of improvement of our model.

References

[1] J. J. McConnell, and A. V. Ovtchinnikov, "Predictability of long-term spin-off returns," *Journal of Investment Management*, vol. 2, no. 3, 2004.

[2] P. Cusatis, J. Miles, and J. Woolridge, "Restructuring through spinoffs: The stock market evidence," *Journal of Financial Economics*, vol. 33, 293-311, 1993.

[3] C. Y. Woo, G. E. Willard, and U. S. Daellenbach, "Spin-off performance: A case of overstated expectations?" *Strat. Mgmt. J.*, vol. 13, pp. 433-447, 1992.

[4] H. Desai and P. Jain, "Firm performance and focus: Long-run stock market performance following spinoffs," *Journal of Financial Economics*, vol. 54, pp. 75-101, 1999.

[5] D. Lee and R. Madhavan, "Divestiture and firm performance: A meta-analysis," *Journal of Management*, vol. 36, no. 6, pp. 1345-1371, 2010.

[6] L. Daley, V. Mehrotra, and R. Sivakumar, "Corporate focus and value creation: Evidence from spinoffs," *Journal of Financial Economics*, vol. 45, pp. 257-281, 1997.

[7] J. McConnell, M. Ozbilgin, and S. Wahal, "Spinoffs, ex ante," *Journal of Business*, vol. 74, pp. 245-280, 2001.

[8] T. W. Liao, "Clustering of time series data — a survey," *Pattern Recognition*, vol. 38, no. 11, pp. 1857-1874, 2005.

[9] Y. P. Chen, E. Keogh, B. Hu *et al.*, *The UCR Time Series Classification Archive*, 2015.

[10] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh. "Querying and mining of time series data: Experimental comparison of representations and distance measures," *PVLDB*, vol. 1, no. 2, pp. 1542-1552, 2008.

[11] Keogh and Ratanamahatana, "Making time-series classification more accurate using learned constraints," SDM, 2004.
[12] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of computational and Applied Mathematics*, vol. 20, pp. 53-65, 1987.
[13] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed., Springer, 2009.
[14] M. Aly, S*urvey on Multiclass Classification Methods*, 2005.

Mehdi Chouiten is the CEO of DATATEGY SAS, former lead data scientist in several companies (including finance industry). He contributed to scoring methods and fraud detection algorithms. He holds a PhD in France (CNRS Laboratory). His main focus is on data science having impact on business strategy and operations.