Difficult decisions in uncertain times: AI and automation in commercial lending

Received (in revised form): 28th March, 2021

Sean Hunter
Chief Information Officer, OakNorth, UK

Sean Hunter is chief information officer at OakNorth. Prior to this, Sean was one of the first commercial engineers at Palantir Technologies in Europe, where he led trader oversight partnerships with large financial institutions, particularly Credit Suisse, which led to being co-head of the joint venture known as Signac. Before Palantir, Sean was a strategist at Goldman Sachs for eight years, working in a number of areas including equities, fixed income and algorithmic trading. Prior to Goldman Sachs he was IT director at a dot com start-up, writing its initial systems and managing a growing development team through two initial public offerings (IPOs). Sean is an adviser to entrepreneurs at Antler, the social impact start-up Kamayi, and is on the editorial board of the *Journal of Digital Banking*. Sean contributed a chapter to *The RegTech Book* (Wiley Press) and is co-author on two patents in the field of anomaly detection.

OakNorth, 57 Broadwick Street, London, W1F 9QS, UK
Tel: +44 (0)797 691 6762; E-mail: sean.hunter@oaknorth.com

Onur Güzey
Head of Artificial Intelligence, OakNorth, UK

Onur Güzey serves as head of artificial intelligence at OakNorth. Prior to this, he was an assistant professor of computer engineering at Istanbul Sehir University. Onur received his BSc in computer engineering from Istanbul Technical University. He earned his MSc and PhD degrees in electrical and computer engineering from University of California, Santa Barbara, and received his MBA from MIT Sloan School of Management. His previous work experience includes positions at Mentor Graphics, Intel and McKinsey & Co.

OakNorth, 57 Broadwick Street, London, W1F 9QS, UK
Tel: +44 (0)755 542 5033; E-mail: onur.guzey@oaknorth.com

Abstract Progress in artificial intelligence (AI) and automation has improved many parts of financial services. These techniques have struggled, however, to make inroads in many areas of commercial lending, largely because of the relative unavailability of sufficient data. Traditional techniques of extrapolation from historical data are also inadequate in times of significant disruption (such as the COVID-19 pandemic). In this paper we discuss these challenges and present techniques such as driver analysis, nowcasting and the use of AI to enable granular subsector classification and forecasting. These allow greater use of data-driven AI-augmented decision making even where full decision automation is not necessarily possible or desirable. Finally, we examine the case study of OakNorth Bank in the UK, which has used these techniques to achieve very promising results since its launch in 2015.

KEYWORDS: nowcasting, commercial lending, driver analysis, small to medium enterprises (SME)

INTRODUCTION

Automation has been used extensively in banking to streamline many arduous processes and improve efficiency. This has had a particular impact in retail banking where, for example, low-friction, highly automated customer offerings have been identified as one of the most important emerging trends in retail banking for 2021.[1] As many loans to businesses are very complex, however, most banks still use a decision-making process that is largely manual. Fintech companies such as Kabbage and Iwoca have been successful in providing an automated application and decision-making process for smaller and more commoditised loans such as simple working capital facilities. As loans become larger and more bespoke, however, it becomes harder to make informed credit decisions at speed and at scale, in a cost-effective manner. As such, banks today still use processes that have remained largely unchanged for the last 20 years. In order to improve these, the focus has primarily been on digitising an existing process to move away from paper rather than radically altering the process to allow for more efficiency.

One of the challenges facing any attempt to improve this is the availability and quality of data. While public data is generally available for large borrowers who have listed equity or debt, for smaller borrowers information is typically hard to come by. Even if that was not the case, any attempt to build an automated decision process for commercial lending along the lines that have been so successful in the credit card industry, for instance, would rapidly hit the limits of available data, given how narrow and specific are the decision-making criteria for commercial lending.

To build decision models for high-volume retail loans, a typical approach is to score each loan using a standardised credit score (eg FICO) and then build a Markov state transition model, which divides loans into states of default (from 'current', through '30', '60' and '90 days delinquent' to 'defaulted', for example) and then estimates the probability of a loan to a borrower in a given credit score band moving from one category to another (see Figure 1).

For any given group of loans, it is possible to estimate the future outcomes using something like Monte Carlo simulation. This process relies heavily on large quantities of loans that all have similar parameters in order for the internal state transition probabilities to be accurate. For credit card loans, pools of millions of loans would be used to fit such a model.

For commercial lending, however, even a large lender is likely to have a few hundred thousand loans at most, and once you divide these by sector, size and vintage of loan, revenue size of company and region, you very quickly get to a number of loans in any given bucket that is way too small to estimate the probability matrix. As a simple example, given a typical average tenor of 2.5–3.5 years, it is highly likely that for many sectors, a bank will have no defaults for loans which are one or two years old.

	To				
From	**current**	**d30**	**d60**	**d90**	**defaulted**
current	0.99	0.01	0.00	0.00	0.00
d30	0.30	0.10	0.60	0.00	0.00
d60	0.10	0.00	0.10	0.70	0.10
d90	0.05	0.00	0.00	0.00	0.95

Figure 1: Hypothetical state transition table giving the probability of loans moving between states of default from one interest period to the next[2]

How, then, can it estimate a probability of default for these loans? In addition, if you make an estimation error when fitting a model for small loans, the resulting losses are by definition also small. As the loan sizes increase, the potential losses increase too, and rapidly approach a point where an unexpected event affecting several borrowers could have a meaningful impact on the capital of a small bank.

Regulators have therefore cautioned against naïve attempts at automated decision making for these types of loan,[3] and banks have instead relied on experienced credit officers to make decisions, backed by a fundamental financial analysis of the borrower and a peer analysis showing the borrower's financial and operating performance in comparison with a peer set of comparable companies. They also will usually ask for assets to provide additional security on any loan. This approach suits mature sectors in which a bank has existing exposures that can be used as a benchmark for any new prospective borrowers, but it is difficult to assess new or disruptive businesses or ones that are in sectors in which a bank has less experience. It also makes it harder for borrowers who operate asset-light businesses, as even a good business may well not have collateral to pledge against a potential loan. This tends to exclude service businesses and many newer business models. Because this process is time-consuming, banks will tend to competitively bid for borrowers that allow for an easy decision, leading to a race to the bottom on yield and generally unprofitable business. Meanwhile, good businesses can find it difficult to obtain finance if the loan size is too small, data is hard to come by, or other circumstances add additional complexity to the decision-making process.

The impact of COVID-19 has worsened this picture. Both the severity of the economic impact and the scale of intervention by governments and policymakers to counter its economic impact have led to a situation which is unprecedented; therefore historical data is unhelpful in making predictions about the prospects for any business borrower applying for a loan at present.

> 'We're in this economy where everybody bases their models prediting the future on the past and of course we've never been in a situation where [we] effectively have been forced to shut down the economy with this much fiscal stimulus.' William Demchak, Chairman, President and CEO, ONC[4]

The approach outlined above is particularly unsuitable because banks may be comparing a new borrower's latest financials (which reflect the crisis) with those of existing peer borrowers in their portfolio (which may be dated and therefore not show the effect). Risk rating models that fit to historical data will also not accurately predict default probabilities under such conditions. Data sources giving a broader macroeconomic picture will typically have a lag and therefore, while accurate, will only show the state of conditions at some time in the past rather than what is happening now. All of this creates a climate of uncertainty, which can lead to banks pulling back and not giving borrowers the finance they need if they are to restart the economy.

There are a number of ways in which AI and automation can help with these challenges. First, we can use statistical techniques to create forward-looking analysis. For example, if we are considering a loan in a given sector, we can look at macroeconomic variables and use those in an attempt to predict the possible future revenues, costs and market size in that sector. This in turn can be used to make intelligent sector-specific scenarios that can be used to predict borrower performance and understand the borrower's future financial needs and risks.

We begin by looking at all the potential macroeconomic variables we have available.

Driver analysis then uses statistical methods to identify relevant drivers for a target time series among all these potential macroeconomic variables. Figure 2 illustrates an easy-to-evaluate correlation graph that shows the relationship between sector revenue and some candidate drivers. Such methods can be used to automatically filter irrelevant macroeconomic variables or identify drivers that can be utilised to create forward-looking analysis for a particular sector. Combined with more advanced forecasting methods and frequent data updates, driver analysis can provide an easily explainable and powerful method for enabling forward-looking analysis.

As additional data sources become available, we can use similar analysis to determine whether any relevant drivers are present in these new sources and whether or not the new sources would be additive.

A related technique is nowcasting, which uses the techniques of prediction to fill in the gaps between the last reliable historical data points and the present time, allowing a bank to effectively eliminate the problem of data source lag and understand the current state of its portfolio and borrowers. Utilising high-frequency datasets, nowcasting predicts the current value of key sector-level indicators such as sector revenue. These predicted indicators can serve as an early indicator, bringing attention to the potential risks developing in a portfolio. The statistical methods used — which include back testing and regression — provide confidence intervals that can help evaluate the accuracy of the predictions. Figure 3 illustrates nowcasting applied to sector-level revenue published by US Census. In this particular example, utilising sector-specific high-frequency data around foot traffic, macroeconomic indicators, mobility and credit card transactions, nowcasting can provide current predictions for data that is otherwise released with more than two months' lag.

There are also more mundane ways in which AI can help to automate tasks and make this decision process easier. For example, data sources frequently have an incomplete or inaccurate sector code for a given company. Even where a sector code is filled in by the borrower or relationship manager at a bank, it may not be the most accurate description of the borrower's activities. This is problematic if a bank is attempting to use sector-specific analysis in making the loan decision. Nevertheless, it would be possible, for example, to train

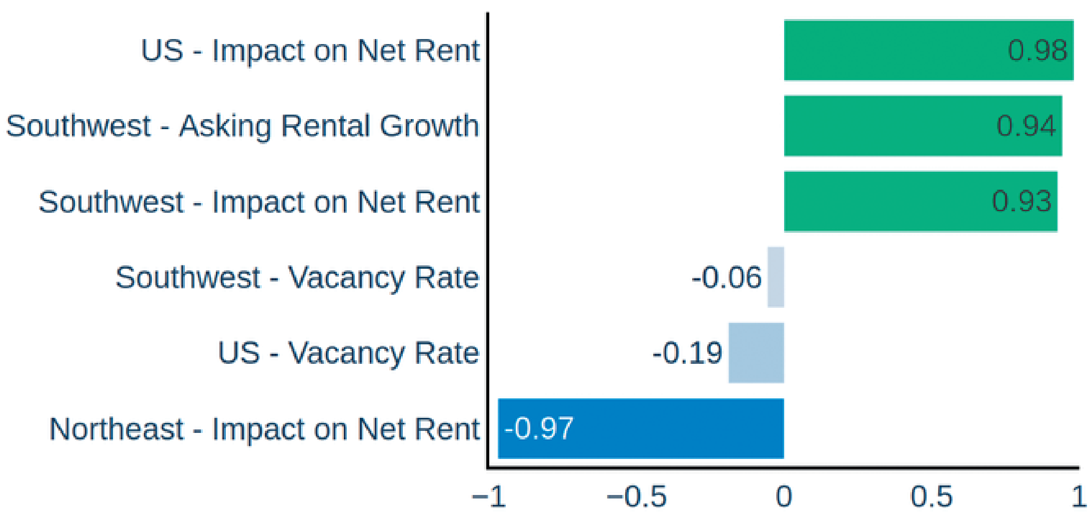

Figure 2: Driver analysis based on correlation of drivers and sector revenue

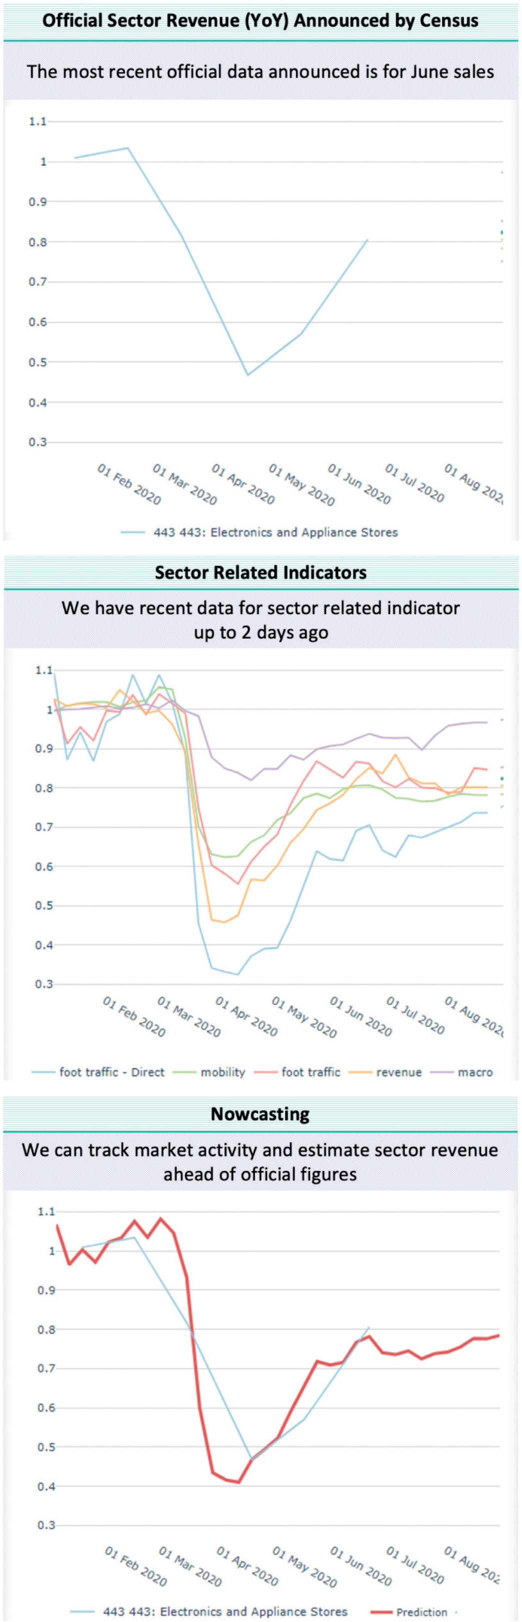

Figure 3: Nowcasting

a classifier on companies where sector data is present and trustworthy, and then use that to estimate likely sector codes for other businesses. Entity resolution is a very common problem affecting companies attempting to build a holistic picture from multiple data sources. For example, because data in the real world is messy, it can be difficult to deduce that a company in one source is the same company in another source. After all, there are frequently many companies with similar names, some of which form groups with common ownership while others may share a name by pure coincidence. It is possible, however, to determine linkages using probabilistic methods, for example,[5] and thereby deduce which companies are most likely to be linked. This is, of course, not a perfect solution, but is a significant improvement on no linkage at all — provided that a suitable confidence cut-off can be chosen to exclude matches that are too likely to be dubious.

Having a better understanding of the sector in which its borrowers operate and the current and likely future conditions in those sectors can allow a bank to perform forward-looking stress testing on its portfolio to see how borrowers are likely to perform under a given set of stress scenarios. For example, OakNorth has created over 262 sector-specific COVID stress scenarios and has used these to rate how vulnerable borrowers are, based on liquidity, debt capacity and profitability through those scenarios. This in turn can be used to produce a blended COVID vulnerability rating and allow a bank to target specific actions at the borrowers who are most likely to benefit. For example, if a borrower has vulnerable liquidity but their future profitability outlook is strong, then if they have additional debt capacity a bank may consider lending them additional funds, assuming the borrower can afford to pay back the loan as long as they get help to survive in the short term. This is a positive outcome for both bank and borrower.

Applying granular stress scenarios across an entire portfolio allows a much more robust and targeted response to the current crisis than merely pulling back from sectors entirely, simply because some businesses in those sectors are suffering.

The purpose of automation in this uncertain and difficult environment is not to replace humans in the decision-making process, but rather to augment the capabilities of the people and allow them to make better decisions and to operate with greater efficiency. If systems can automate sector, peer and financial analysis, produce intelligent insights and highlight borrowers who may have difficulties in the future, the experienced credit analyst is freed up to exercise their judgement and focus on the most important aspects of the decision process. This in turn should lead to better decisions and better credit outcomes.

PREPARING FOR THE PANDEMIC: A CASE STUDY

OakNorth Bank plc, UK

When COVID-19 began to emerge in January 2020, OakNorth Bank's first course of action was to run stress scenarios related to potential supply chain disruption from China on its loan book. By mid-February, before the bank was even sure that COVID-19 had reached British shores, it had already run six full portfolio stress tests across its entire loan book, on a granular, loan-by-loan basis.

From 13th March (ten days before the UK went into its first lockdown), OakNorth Bank began hosting daily ExCo and Credit Committee meetings to discuss how it could best support its borrowers, maintain credit quality, support its deposit clients and operationalise remote working to ensure the safety of its teams. Within 36 hours of the first UK lockdown coming into effect, the bank had run a reverse stress test on its entire loan book, determining how long each business could survive if operating expenses were to stay constant but revenues dropped to zero. This ensured the bank was able to prioritise and in what order.

Over the next two weeks, the bank continually refined the granularity of its stress tests — increasing the severity of the situation based on the UK lockdown lasting three months, then six months, then nine months, and so on. The COVID Vulnerability Rating Framework enabled OakNorth bank to:

- Classify its loan book into granular subsectors and determine the impact of COVID-19 using the sector-specific domain models and forward-looking scenarios;
- Assess each sector through the three stages of the crisis: the initial impact from COVID-19, additional waves with short-term reboots in between, and the new normal;
- Explore a range of outcomes based on structural changes in consumer behaviour, the regulatory impact, government fiscal stimulus and the impact of increased digital usage;
- Utilise this subsectoral analysis to stress test the entire portfolio simultaneously on a loan-by-loan basis and flag which individual obligors might need closer analysis and support;
- Re-underwrite loans to vulnerable businesses and institute closer monitoring while helping management teams understand the stress scenarios;
- Build trust in the scenarios through regular efficacy tests incorporating concepts of 'nowcasting' and 'back-testing'.

This unique approach enabled OakNorth Bank to continue achieving market-leading results and growing its loan book throughout 2020 while maintaining its track record of no credit losses since inception.

CONCLUSION

Commercial lending has always been a challenge, but it is particularly important for banks to make good decisions not only to avoid unnecessary losses but also so they can support borrowers and economic growth. The use of predictive analytics, nowcasting and other similar techniques can be used to provide essential additional context. These are all ways in which AI and automation can be employed to supplement human decision makers within the credit process and lead to better outcomes for both commercial lenders and borrowers.

References

1. Capgemini (2020), ‘Retail Banking: Top Trends 2021’, available at https://www.capgemini.com/wp-content/uploads/2020/11/Retail-Banking-Trends_2021_Web.pdf (accessed 15th June, 2021).
2. For background on the method, see for example Jones, M. T. (2005), ‘Estimating Markov Transition Matrices Using Proportions Data: An Application to Credit Risk’, IMF Working Paper, available at https://www.imf.org/external/pubs/ft/wp/2005/wp05219.pdf (accessed 15th June, 2021).
3. For instance, EBA, ‘Guidelines on loan origination and monitoring’, available at https://www.eba.europa.eu/regulation-and-policy/credit-risk/guidelines-on-loan-origination-and-monitoring (accessed 15th June, 2021).
4. Noonan, L. and Armstrong, R. (May 2020), ‘PNC says fears for US economy promted sale of BlackRock stake’, *Financial Times*, available at https://www.ft.com/content/da9798a6-0d53-491d-b3f7-cccfcb3f975f (accessed 15th June, 2021).
5. Fellegi, I., and Sunter, A. (1969), ‘A Theory for Record Linkage’, *Journal of American Statistical Association*, Vol. 64, No. 328, pp. 1183-1210, available at https://www.tandfonline.com/doi/abs/10.1080/01621459.1969.10501049 (accessed 28th June, 2021).