CASE STUDY
FEBRUARY 2021

# Determining the Optimal Dynamic Credit Card Limit

**Prepared by**

Juan Licari
Juan.Licari@moodys.com
Managing Director

Olga Loiseau-Aslanidi
Olga.Loiseau-Aslanidi@moodys.com
Head of risk modelling of Economics and Business Analytics APAC

Vera Tolstova
Vera.Tolstova@moodys.com
Assistant Director-Economist

Masood Sadat
Masood.Sadat@moodys.com
Economist

**Contact Us**

Email
help@economy.com

U.S./Canada
+1.866.275.3266

EMEA
+44.20.7772.5454 (London)
+420.224.222.929 (Prague)

Asia/Pacific
+852.3551.3077

All Others
+1.610.235.5299

Web
www.economy.com
www.moodysanalytics.com

In this paper, we use a Markov decision process (MDP) model for credit card profitability to find the optimal dynamic credit limit policy. The state space is represented by customers' behavioural scores and balances to allow for more granular policy setting. We use U.K. credit card data as a case study to build a behavioural scorecard for the scoring bands to segment customers by risk. We then demonstrate the optimal credit limit optimisation and compare it with the actual data outcomes.

# Determining the Optimal Dynamic Credit Card Limit

BY JUAN LICARI, OLGA LOISEAU-ASLANIDI, VERA TOLSTOVA AND MASOOD SADAT

In this paper, we use a Markov decision process (MDP) model for credit card profitability to find the optimal dynamic credit limit policy. The state space is represented by customers' behavioural scores and balances to allow for more granular policy setting. We use U.K. credit card data as a case study to build a behavioural scorecard for the scoring bands to segment customers by risk. We then demonstrate the optimal credit limit optimisation and compare it with the actual data outcomes.

## Introduction

Credit cards have traditionally been one of the most profitable lines of business for lenders. The market has grown rapidly in recent decades, with transactions on major credit card brands reaching more than $450 billion globally in 2020 and expected to grow by more than 40% by 2025.[1]

This boom together with enhancements in computational power has made quantitative analytical methods play a bigger role in risk-adjusted strategy development such as credit card limit management. While some lenders still rely on expert-driven policies or static risk-return approaches without considering the default risk and profitability over time, many are turning to dynamic analytics to make optimal decisions for managing limits to increase profitability and protect from losses.

Various techniques are available for modelling the credit limit as one of the key parameters of credit card profitability. Academic studies and industry practitioners find that credit limit policy affects credit card users' behaviours and profits earned. For instance, Soman and Cheema (2002) argue that a higher credit limit encourages the usage of credit. Trench et al. (2003) use the Markov decision process to optimise a customer's lifetime value by changing either the customer's credit limit or interest rate. So and Thomas (2011) use the MDP model to show that the dynamic credit limit policy based on behavioural scores and dynamic programming optimisation techniques help to account for changing borrower behaviour and lead to higher expected profit.

In this paper, we set up a model for the profitability of credit cards and use it for the optimal dynamic credit limit policy evaluation. This model is a version of the Markov decision process, where the states of the system are based on the borrower's behavioural scores as well as credit card balances. The decisions are made in each time period, affecting the return from the borrowers and the probabilities to which state they move next. This approach is applied to a case study of U.K. credit card portfolios with historical data for more than 4 million accounts obtained from the European Data Warehouse.

To design the states of the system, we build a behavioural scorecard model using the methodology described in Loiseau-Aslanidi, Thiagarajah and Tolstova (2020), and set up a bucketing algorithm to segment the obtained behavioural scores and balances. To test if a higher profit can be achieved, we condition the dynamic optimal credit limit policy on the previous period balance information rather than on credit limit. This allows for a more granular optimal policy setting, as it may lead to higher expected profit by freezing balances on accounts for high-risk borrowers, that is, setting a limit equal to balances. As the credit has already been granted and a limit assigned in the dataset, the results of the optimal dynamic limit modelling are compared to actual data in the portfolio to show that optimal limit policy boosts expected profit through the change in the distribution of accounts across behavioural scores and balance buckets.

## Model specification

### Theoretical settings

The model designed in this paper belongs to the family of models used for credit card policies optimisation with enhancements and simplifying assumptions.[2] Leveraging a classical MDP model, we apply a dynamic programming approach to solve for optimal credit limit policy.

1 According to the latest Nilson report.

2 Herasymovych, 2019: "Modeling credit process as an MDP is far from new, however it has been mainly used to optimize credit limits (So, 2009; So and Thomas, 2011), credit prices (Trench et al., 2003), collection policies (Briat, 2005) or as a credit scoring system itself (Malik and Thomas, 2010; Regis and Artes, 2015)."

This framework allows for sequential decision-making considering the evolution of a customer's behaviour over time.

For simplicity, we assume that the dynamics of account movements across states is defined by a stationary first order Markov chain. We identify different states that the account can be in and look at the account movements between different states during a given period. We assume an infinite account lifetime horizon and use a monthly time period.[3]

The application of behavioural score as a state variable was introduced by So and Thomas (2011). In our model, we also add the balance as a second state variable. The behavioural scores may depend on balance, limit and other variables that depend on them, such as utilisation. If utilisation is a behavioural score component, either the previous credit limit or utilisation itself can be used as a third state variable. In our case, we do not expand the state space further as we use the current balance-to-limit-at-origination ratio rather than the ratio based on the optimal limit predicted by our model.[4]

To define the state space in our model, we combine the previous time period behavioural scores and balances into bands or buckets using the bucketing algorithm described below. Hence the state variables are the previous period balance band $b_{t-1}$ and the behavioural scores band $s_{t-1}$. To obtain the behavioural scores, we employ an automated model-building toolbox described in Loiseau-Aslanidi, Thiagarajah and Tolstova (2020) to build a behavioural scorecard model. The toolbox features algorithms leveraging a modified logistic regression with predefined constraints imposed via supervised binning and variable selection.

The MDP model timing is depicted in Chart 1. At the beginning of the current period *t*, a decision maker observes the state variables and takes them as given when deciding on the credit limit $l_t$ which becomes a function of $b_{t-1}$ and $s_{t-1}$. A credit limit $l_t$ affects customer behaviour and defines the balance bucket $b_t$ by the end of period *t*. We take each credit limit level to be one of the possible actions that can be chosen.

By the end of period *t*, each account transits to states $b_t$ and $s_t$ with probability $b_t(b_t,s_t\,|b_{t-1,}s_{t-1,}l_t)$ which is constant over time and depends only on previous period buckets $b_{t-1}$ and $s_{t-1}$, and credit limit $l_t$, such that:

» $l_t \geq b_{t-1,}$ the limit cannot be set below the balance at the beginning of period

» $l_t \geq b_{t,}$ the balance by the end of period cannot exceed the limit

Each account can default with probability $pd(b_t)$ that depends on the behavioural score bucket. If the account defaults, the losses are equal to the value of balance by the end of the period.[5] If the account does not default, the corresponding profit is equal to *r*

## Chart 1: Model Timing Diagram

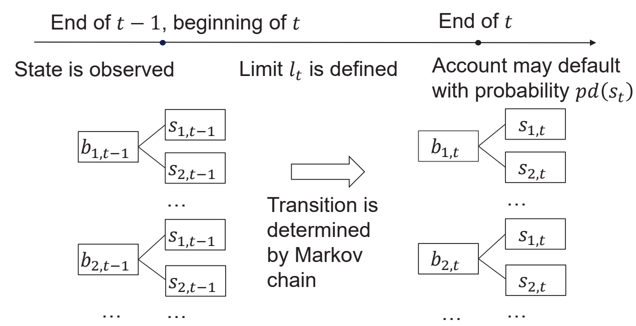


Source: Moody's Analytics

* $b_t$, where *r* is a monthly interest rate for simplicity assumed to be constant.

The expected profit at the end of period t depends on the buckets for, $b_{t-1}$, $s_{t-1}$ and $l_t$ and can be defined as:

The optimal credit limit is a policy function that maximises future

$$\pi(b_{t-1}, s_{t-1}, l_t) = \sum_{b_t, s_t} p(b_t, s_t |\, b_{t-1}, s_{t-1}, l_t) * \left(-pd(s_t) b_t + \left(1 - pd(s_t)\right) r b_t\right)$$

discounted profit obtained from the customer. A dynamic programming problem for the optimal credit limit evaluation is specified as follows:

$$V(b_{t-1}, s_{t-1}) = \max_{l_t(b_{t-1}, s_{t-1})} \{\pi(b_{t-1}, s_{t-1}, l_t) + \beta \textstyle\sum_{b_t, s_t} p(b_t, s_t |\, b_{t-1}, s_{t-1}, l_t) * V(b_t, s_t)\} \quad (1)$$

where V $(b_{t-1}, s_{t-1})$ is a value function that depends on the state variables, $0 < \beta < 1$ i is the parameter capturing the time value of money.

#### Bucketing algorithm

We apply an automated bucketing algorithm for the state variables to identify the bands for the behavioural scores and balances in line with So and Thomas (2011). Nevertheless, instead of just splitting each state variable into buckets with different dynamics, our approach also focuses on identifying the bands with different risk. We also ensure enough observations in each bucket for the scores and balances, which is essential for the transition probabilities estimation.

Our algorithm consists of the two key steps below.

1. The sequential splitting of behavioural scores and balance values using our binning algorithm described in Loiseau-Aslanidi, Thiagarajah and Tolstova (2020) that transforms the values of the state variables into several groups, known as buckets or bins. We also specify a maximum number of cut points for each state variable according to specified criteria:
   » The behavioural scores are split to differentiate between buckets with low and high default rates
   » The balance values are split to differentiate between low and high-risk accounts within each score bin.
2. The algorithmic refinement of the balance cut points obtained from the previous step. This aims to verify that the number of

3 Since the average duration for U.K. credit cards in the case study data sample is around 20 years, this assumption is realistic in our case.

4 The remaining attributes include balance growth rates and variables that depend neither on balance nor on limit. The dependence of the behavioural score on the balance and balance growth rates makes two state variables interconnected. This does not pose technical issues, as the association between the state variables is captured by the transition probability matrix of MDP. We do not consider the incompatible balance and behavioural scores buckets.

5 An average loss given default (LGD) for U.K. credit cards estimated for EDW data is equal to around 95%; therefore, this is a realistic assumption for our case study.

## Chart 2: U.K. Credit Card Scorecard Model

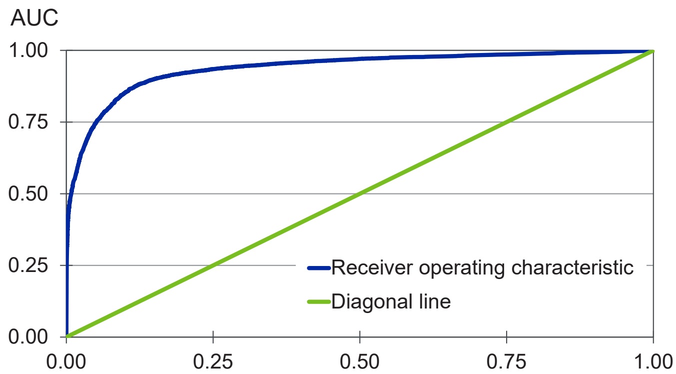


Source: Moody's Analytics

## Chart 3: U.K. Credit Card Interest Rate

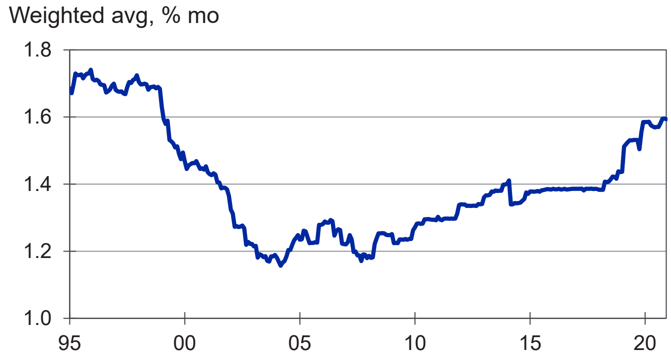


Sources: Bank of England, Moody's Analytics

observations in each combination of behavioural score and balance buckets is at least as high as the minimum requirement:

- Combine balance buckets cut points into one vector
- Iteratively eliminate the cut points for which the requirement of minimum number of observations is not satisfied for at least one combination of behavioural score band and corresponding balance bucket.

### Application to U.K. credit card data

In this section, we apply the modelling approach described above to the actual data from credit card customers. The credit card data are obtained from the European Data Warehouse for the pool of U.K. credit cards covering a wide range of anonymous clients. The dataset includes nearly 4 million unique clients' accounts and 150 million observations. The historical observation period is from May 2014 to January 2019. The available characteristics include information on current balance, limit, arrears balance, next minimum payment, payment method, customer's employment status, and geographical region.

#### Behavioural scorecard model

We begin by designing the behavioural scorecard probability of default model using an automated model building procedure described in Loiseau-Aslanidi, Thiagarajah and Tolstova (2020). We use our toolbox with the binning and the variable selection algorithms that solve relevant optimisation problems subject to constraints required in credit scoring. The resulting model contains the best set of characteristics for the most comprehensive risk profile while addressing issues that affect model reliability and applicability.

The analytical tractability of the PD model is implemented through the constraints into the automated model building procedure. These constraints guarantee the presence of intuitive trends in the default rates across intervals of binned variables and the intuitive signs of coefficients in the model. For instance, the constraints include a positive/negative relationship between the default rate and balance in arrears/limit at origination.

The behavioural PD scorecard model we build is a weight-of-evidence logistic regression with the key drivers listed below. This strongest set of characteristics is chosen from the available initial characteristics' analysis so that all weak criteria are eliminated. The attributes in the final model are chosen through the forward stepwise regression with constraints on coefficient signs to guarantee an intuitive impact of model drivers on the default rate.

- Utilisation ratio based on limit at origination
- Balance in arrears to original balance
- Next minimum payment to current balance
- 3-month absolute growth in balance
- Months on book
- Payment method
- 12-month absolute growth in balance

The model demonstrates an accurate in-sample model fit and high quality out-of-sample performance measured by a Gini coefficient of 87.6%. The AUC is shown in Chart 2, and the model is interpretable by default.

We apply a mapping to transform the estimated PDs from the scorecard model to scores. Such mapping refers to the range and format of scores in a scorecard and the rate of change in odds for increases in score. There are various mappings to transform the estimated PDs from the scorecard model to the scores options used in the industry. In this paper, we use one of the most common approaches where the relationship between the log odds and the scores can be presented as a linear transformation:

$$Score = \theta + \delta \cdot ln\left(\frac{1-PD}{PD}\right),$$

where and θ and δ are such that the score equals 500 when good/bad odds equals 1 and the score increases by 50 points when good/bad odds double. These assumptions translate into θ = 500 and δ = 50/log (2).

## Chart 4: Default Rate by Score Buckets

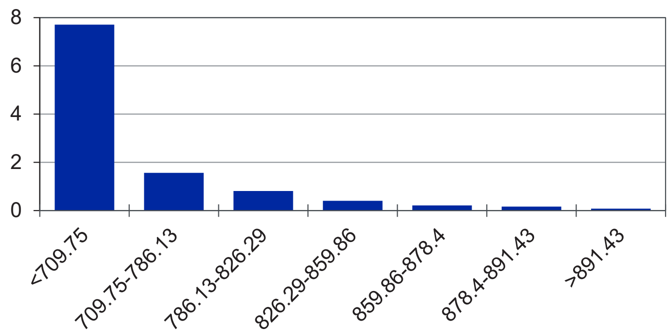


Source: Moody's Analytics

## Chart 5: Optimal Policy Limits by Score Buckets

Lower-balance bucket

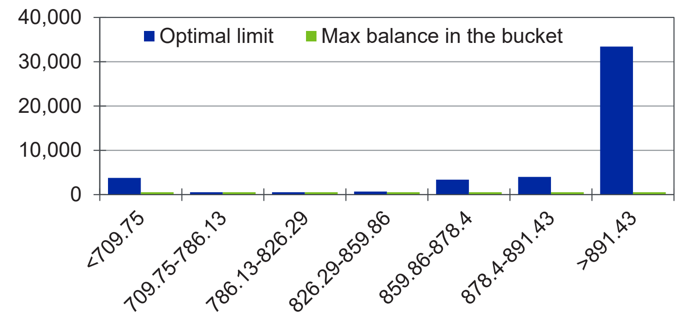


Source: Moody's Analytics

### Credit limit optimisation

The behavioural scores to assess the default risk of the customers obtained above are used in the MDP model. For this exercise, we use 1 million randomly sampled accounts from U.K. EDW data for credit cards to solve for the optimal credit limit using the Eq (1) with the time discount parameter set to 0.995 as in So & Thomas (2011).

We solve for optimal limit policies for various values of monthly interest rates, including 1.5%, 1.6%, 1.75%, 2% and 2.5%. The results of optimal credit policy are similar for all cases except for interest rates of 2% and above. The latter case involves slightly higher limits for the intermediate behavioural score buckets. For the more detailed analysis presented here, we consider 1.5% as it corresponds to the average interest value observed in the last three years (see Chart 3).

In our exercise, the number of buckets for scores and balances determined by the bucketing algorithm are 7 and 20, respectively.[6] The resulting score buckets are demonstrated in Chart 4. The higher-risk bucket with scores below or equal to 709.75 has the highest default rate, while the lower-risk bucket with scores above 891.43 has the lowest default rate. The picked balance buckets are arranged by size, being more granular in the first five buckets and gradually becoming less granular in the remaining buckets.

Charts 5 and 6 demonstrate the optimal credit policy example for a lower (450;540] and a higher balance bucket (1483;1566], respectively, generated by the model for $r = 1.5\%$.[7] The optimal credit limit policy recommends increasing the credit limit for customers with the highest behavioural scores in both high- and low-balance buckets. Moreover, the policy recommends increasing the limit for relatively inactive or low-balance accounts even in cases when the account is relatively risky with lower behavioural scores. Such a policy can encourage borrowers to use their credit cards, while it is unlikely that those accounts moving to the bucket with a high balance could result in high losses in case of default. Hence the expected losses are low even if the limit is increased.

## Chart 6: Optimal Policy Limits by Score Buckets

Higher-balance bucket

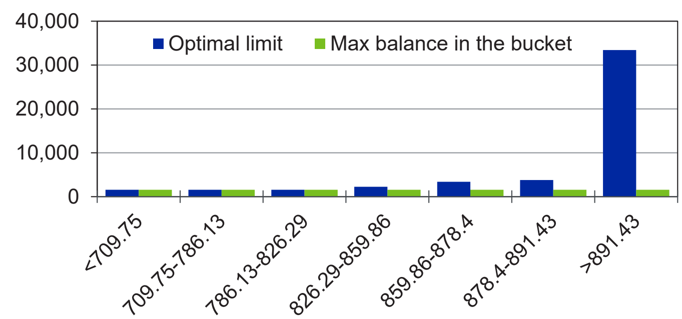


Source: Moody's Analytics

To compare the MDP model and the actual policy, we consider the following four cases for policy, profit and transition for each interest rate:

1. Optimal policy, model profit, model transition
2. Optimal policy, model profit, real transition
3. Real policy, model profit, real transition
4. Real policy, real profit, real transition

The optimal policy is the optimal credit limit policy estimated by the MDP model, while the real policy is the actual limit observed in the data.

## Table 1: Total Discounted Profit Per Account

| | Optimal policy | | Real policy | |
|---|---|---|---|---|
| Monthly interest rate, % | Model profit Model transition | Model profit Real transition | Model profit Real transition | Real profit Real transition |
| 1.50 | 61.92 | -154.71 | -180.14 | -153.44 |
| 1.60 | 90.76 | -111.17 | -135.35 | -108.40 |
| 1.75 | 116.02 | -45.73 | -68.15 | -40.83 |
| 2.00 | 200.44 | 63.90 | 43.84 | 71.78 |
| 2.50 | 362.56 | 283.99 | 267.82 | 297.00 |

*Source: Moody's Analytics*

6 When running the bucketing algorithm for the state variables, we set the minimum number of observations in each bucket to 10,000, while the maximum number of cut points for scores and balances binning steps are set to 20.

7 Similar conclusions hold for other interest rates.

The model profit is the expected profit determined by the default probability of corresponding behavioural score buckets, while the real profit is the actual profit calculated based on the observed defaulted and non-defaulted states in the data.

The model transition means that the distribution of accounts across states evolves according to the Markov chain with transition matrix estimated based on the historical data. The real transition means that the actual distribution of accounts observed in the data is used.

The profit computed for each case defined above is shown in Table 1. We calculate the total discounted profit per account starting with the initial distribution of accounts in 2014 and summing up the monthly profits weighted by the time discount parameter until the last observation period in 2018.

The results indicate an improvement in profitability when the optimal policy is used. The total discounted profit for the optimal limit policy, model profit and model transition is substantially higher than for other cases when the transition is real. When the model transition is considered along with optimal policy, the monthly profit per account converges to substantially higher values than otherwise, as shown in Chart 7.

The positive impact of optimal policy on expected profit is mostly driven by the adjustments in the distribution of accounts across states with the transition matrix adjusted by the optimal limit. The policy to freeze the relatively high-balance and high-risk accounts restricts their transitions to states with an even higher balance and, potentially, higher losses in the case of default.

## Chart 7: Monthly Profit Per Account

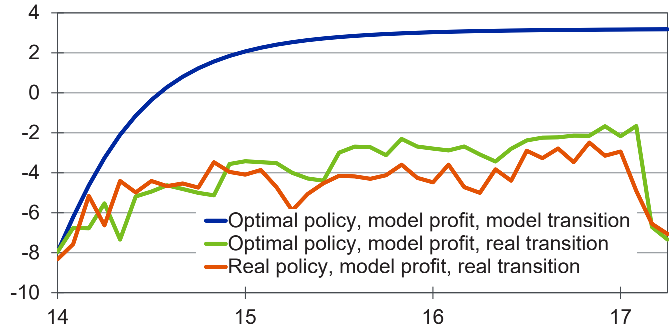


Source: Moody's Analytics

### Concluding remarks

This paper uses the Markov decision process model to generate a dynamic credit limit policy. We link the profitability of a borrower to their default risk by building a behavioural scorecard model to obtain the behavioural scores, which are typically used by lenders to measure risk. In our MDP model, we use both the behavioural score and the balance as state variables and employ the bucketing algorithm to identify the bands for the state variables with different risk. The effectiveness of this modelling approach is shown by its application to a real credit card dataset.

# References


Briat, P. (2005). A Markovian approach to the analysis and optimization of a portfolio of credit card accounts. (Master thesis, National University of Singapore). Retrieved from https://core.ac.uk/download/pdf/48629023.pdf

Herasymovych, M. (2018). Optimizing acceptance threshold in credit scoring using reinforcement learning. (Master thesis, University of Tartu). Retrieved from https://core.ac.uk/download/pdf/159135256.pdf

Loiseau-Aslanidi, O, Thiagarajah, N. S., & Tolstova, V. (2020). Automating interpretable machine learning scorecards. Moody's Analytics white paper.

Malik, M., & Thomas, L. C. (2012). Transition matrix models of consumer credit ratings. International Journal of Forecasting, 28(1), 261-272.

Régis, D. E., & Artes, R. (2016). Using multi-state Markov models to identify credit card risk. Production, 26(2), 330-344.

So, M. C. (2009). Optimizing credit limit policy by Markov Decision Process Models (Doctoral dissertation, University of Southampton).

So, M. M., & Thomas, L. C. (2011). Modelling the profitability of credit cards by Markov decision processes. European Journal of Operational Research, 212(1), 123-130.

Soman, D., & Cheema, A. (2002). The effect of credit on spending decisions: The role of the credit limit and credibility. Marketing Science, 21(1), 32-53.

Trench, M. S., Pederson, S. P., Lau, E. T., Ma, L., Wang, H., & Nair, S. K. (2003). Managing credit lines and prices for bank one credit cards. Interfaces, 33(5), 4-21.

## About the Authors

Dr. Juan M. Licari is a managing director at Moody's Analytics in the London office. He is the global head of the Business Analytics team consisting of risk modellers, economists, and statisticians in the U.K., the U.S., China, UAE, the Czech Republic and Singapore. Dr. Licari's team provides consulting support to major industry players, builds econometric tools to model credit phenomena, and implements several stress-testing platforms to quantify portfolio risk exposure. His team is an industry leader in developing and implementing credit solutions that explicitly connect credit data to the underlying economic cycle, allowing portfolio managers to plan for alternative macroeconomic scenarios. Dr. Licari has extensive hands-on experience as a project lead with respect to development, validation, calibration and monitoring of internal ratings-based models, IFRS 9 and stress-testing credit risk models especially for U.K. banks and financial institutions, for both retail and corporate portfolios. Dr. Licari is actively involved in communicating the team's research and methodologies to the market, including senior management and board members. He often speaks at credit events and economic conferences worldwide. Dr. Licari holds a PhD and an MA in economics from the University of Pennsylvania and graduated summa cum laude from the National University of Cordoba in Argentina.

Olga Loiseau-Aslanidi is the head of risk modelling of Economics and Business Analytics APAC, based in the Singapore office. She manages a team of economists and risk modelers in Prague, Shanghai and Sydney who design models for forecasting and simulation, with an emphasis on stress-testing for three key areas: macroeconomic models, market risk, and credit portfolio risk. During her time at Moody's based in Europe and now in Asia, Olga has led consulting projects with major banks and other financial institutions worldwide focusing on stress-testing, including CCAR, EBA, PRA as well as IFRS 9, IRRBB and IRB model design and implementation. She is directly involved in the research and implementation of Moody's Analytics risk management solutions for market risk and retail credit risk modelling, and often speaks at credit events and economic conferences worldwide, communicating the team's research and methodologies to the market. Before joining Moody's Analytics, she worked for the Academy of Sciences of the Czech Republic and at a consultancy firm focused on macroeconomic forecasting and analysis of emerging market economies. She has published several academic articles and has been teaching graduate courses in economics, statistics and econometrics. She holds a PhD and an MA in economics and econometrics from Charles University (CERGE-EI), following studies for an MSc in mathematics and a BSc with honors in applied mathematics.

Vera Tolstova is an economist at Moody's Analytics in the Prague office. Vera is involved in data analysis, model development and preparation of documentation for consulting projects with banks and international financial institutions. She is also working on developing methodologies for market and credit risk models including PD, LGD and EAD, and their quantitative implementation in statistical software. Prior to joining Moody's, Vera was working as an analyst at a leading Russian pharmaceutical distributor and as a junior researcher at the Center of Economic Research and Graduate Education (CERGE-EI) in Prague. She received her MA in economics from CERGE-EI in 2014, after earning master's and bachelor's degrees from Novosibirsk State University in Russia. She is working on her dissertation thesis with a focus on dynamic macroeconomics with heterogeneous agents, public policy and family economics.

Masood Sadat is an economist at Moody's Analytics in Prague. Masood contributes to projects focused on IFRS 9 stress-testing and forecasting, with a focus on data preparation, modeling and documentation. He received his MA in applied economics from CERGE-EI in Prague, after earning a master's degree from OSCE Academy in Bishkek, Kyrgyzstan and a bachelor's degree from the University of Lucknow in India.

## About Moody's Analytics

Moody's Analytics provides financial intelligence and analytical tools supporting our clients' growth, efficiency and risk management objectives. The combination of our unparalleled expertise in risk, expansive information resources, and innovative application of technology helps today's business leaders confidently navigate an evolving marketplace. We are recognized for our industry-leading solutions, comprising research, data, software and professional services, assembled to deliver a seamless customer experience. Thousands of organizations worldwide have made us their trusted partner because of our uncompromising commitment to quality, client service, and integrity.

Concise and timely economic research by Moody's Analytics supports firms and policymakers in strategic planning, product and sales forecasting, credit risk and sensitivity management, and investment research. Our economic research publications provide in-depth analysis of the global economy, including the U.S. and all of its state and metropolitan areas, all European countries and their subnational areas, Asia, and the Americas. We track and forecast economic growth and cover specialized topics such as labor markets, housing, consumer spending and credit, output and income, mortgage activity, demographics, central bank behavior, and prices. We also provide real-time monitoring of macroeconomic indicators and analysis on timely topics such as monetary policy and sovereign risk. Our clients include multinational corporations, governments at all levels, central banks, financial regulators, retailers, mutual funds, financial institutions, utilities, residential and commercial real estate firms, insurance companies, and professional investors.

Moody's Analytics added the economic forecasting firm Economy.com to its portfolio in 2005. This unit is based in West Chester PA, a suburb of Philadelphia, with offices in London, Prague and Sydney. More information is available at www.economy.com.

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). Further information is available at www.moodysanalytics.com.



## About Moody's Corporation

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). MCO reported revenue of $4.8 billion in 2019, employs more than 11,000 people worldwide and maintains a presence in more than 40 countries. Further information about Moody's Analytics is available at www.moodysanalytics.com.