Algorithmic bias and financial services

A report prepared for Finastra International

March 2021

Important notice

This report has been prepared by KPMG LLP ("KPMG") solely for Finastra International Limited ("Finastra" or "Addressee") in accordance with the terms of engagement agreed between Finastra and KPMG, dated 17 February 2021.

KPMG's work for the Addressee was performed to meet specific terms of reference agreed between the Addressee and KPMG and that there were particular features determined for the purposes of the engagement.

The report should not be regarded as suitable to be used or relied on by any other person or any other purpose. It is issued to all parties on the basis that it is for information only and that it will not be copied or disclosed to any third party or otherwise quoted or referred to, in whole or in part, without our prior written consent.

This report is not suitable to be relied on by any party wishing to acquire rights against KPMG (other than Finastra) for any purpose or in any context. Any party other than Finastra that obtains access to this document or a copy and chooses to rely on this document (or any part of it) does so at its own risk. To the fullest extent permitted by law, KPMG does not accept or assume any responsibility to any readers other than Finastra in respect of its work for Finastra, this document, or any judgements, conclusions, opinions, findings or recommendations that KPMG may have formed or made.

KPMG does not assume any responsibility and will not accept any liability in respect of this document to any party other than Finastra.

KPMG does not provide any assurance on the appropriateness or accuracy of sources of information relied upon and KPMG does not accept any responsibility for the underlying data used in this draft report.

The opinions and conclusions expressed in this report are those of KPMG and do not necessarily align with those of Finastra.

Contents

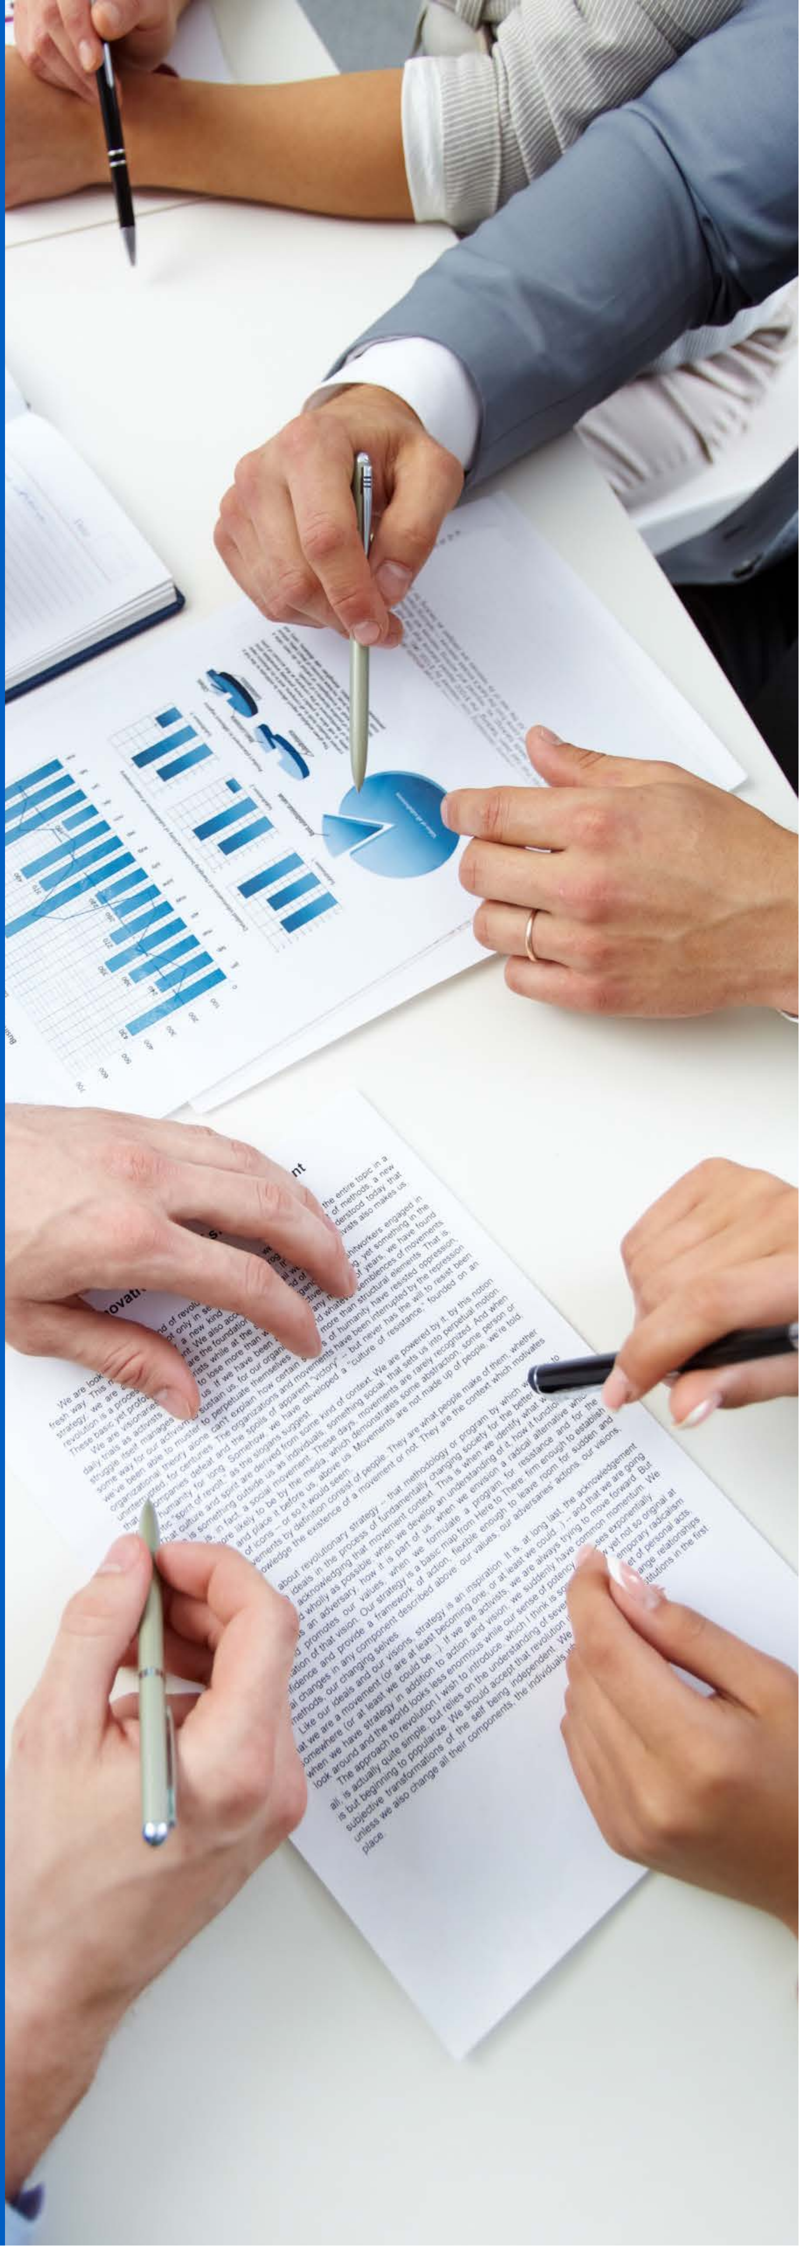

01

Executive summary

The objective of this report is to stimulate and add to the discussion regarding how to mitigate potential algorithmic bias in financial services.

To do so, this report contains four key sections. We explore the fundamentals of algorithms – how they work, what they require, and how they are used in day-to-day business. We set out how bias can be introduced into algorithms. We consider how some algorithms are used in financial services, and the relevance of the biases explored given those uses. We also set out key statistics on the market sizes for a range of products in the sector used by end consumers.

While algorithms themselves can be highly complex, the **fundamentals of algorithms** are straightforward. In simple terms, algorithms are a set of mathematical instructions that enable computers to complete a specific task, for example, pricing an insurance policy. They span a continuum of complexity, from simple rules-based approaches through to algorithms that learn with greater autonomy. In short, algorithms enable organisations – including those in financial services – to utilise a range of data, drawing out patterns and making predictions that a human alone could not calculate either in the same time or with the same level of accuracy.

A layer of complexity to be aware of when designing, refining and implementing an algorithm is **algorithmic bias.** Biases can manifest themselves within data being used in the algorithm, in the way an algorithm is designed, and even in how the algorithm is used (i.e. in interpretation of the results by a human).

Algorithmic bias can be prevalent for several reasons including the use of samples of data that are not representative of the population, overweighting certain data more than others without valid justification, and/or due to the beliefs held by human end users.

Given the current global pandemic, a particularly pertinent bias to be aware of is **emergent bias** (i.e. a bias that is not known at the time of an algorithm's development but arises in the future). As our societal view of what constitutes bias evolves over time, algorithms must also evolve. In the context of COVID-19, algorithms may be using historical data – for example in relation to job types (and job security), or income. Given the structural changes taking place in the global economy, these data may no longer be reflective of the current or future situation. Without continual adjustment and evaluation of algorithms, organisations run the risk of producing results which, potentially unfairly, offer different outcomes to different groups of consumers.

Algorithms have been used in consumer-facing financial services products for many years, with the tools and techniques used advancing substantially over time. The use of algorithms in financial services includes **credit scoring, rate setting and insurance pricing,** in addition to back-office functions such as **credit risk** and **fraud risk** management.

Financial services providers increasingly rely on algorithms to make decisions on how to price a policy or whether to provide credit and at what price, with the level of human involvement in these decisions reducing as a result. **Understanding the potential for algorithmic bias** within the underlying algorithms is therefore key.

Consumer expenditure in the product groups we discuss in this report is significant. For example, **consumer lending was over £440bn, $6,100bn and €280bn in the UK, US and France, respectively, last year**. Both the provision and costs to consumers of this credit (e.g. the interest rates charged) will be informed in many cases by the types of algorithms we discuss.

It is therefore important that algorithms are constructed in ways that mitigate against the potential for bias. Not doing so could run the risk of **financial harm to the end consumer, a loss of trust, legal action and/or regulatory enforcement.**

02

The basics of algorithms

Organisations continually strive to serve their customers better. As technology evolves at these organisations, such as moving key front and back office systems to the cloud,[1] the adoption of Artificial Intelligence (AI) is gaining substantial momentum.

In financial services, the applications and potential benefits of the use of AI are varied. Uses range from algorithms to improve and tailor customer experiences, through to developing sophisticated models that are allowing credit provision to customers that previously would not have been considered.

In this section of our report, we discuss the basics of algorithms. We start with a discussion of algorithm fundamentals, including the main types of algorithms used in AI and how they work. We then set out examples as to how algorithms can be used, from automated decision making through to reducing costs. We then discuss one of the most crucial elements of algorithms: data. Finally, we discuss how an algorithm moves from a concept to being used in day-to-day business as well as some key limitations.

2.1 Fundamentals of an algorithm

In simple terms, algorithms are a set of mathematical instructions that enable computers to complete a specific task. For example, if a company wants to contact all customers over the age of 65, an algorithm can be written to sort through an entire customer database, selecting only those customers with a date of birth that is older than a specified date.

Algorithms span a continuum of complexity, from explicitly programmed step-by-step rules (such as the example provided above) to algorithms that learn without being explicitly pre-programmed. These latter cases – algorithms that learn from data on how to execute a task, as opposed to using step-by-step instructions – are the foundations of AI.

Machine Learning ("ML") is a subfield within AI. These two terms are often used interchangeably, although have some important differences. ML refers to a set of algorithms which are applied to data in order to identify patterns and/or make predictions. These algorithms can perform well in contexts where traditional statistical models may struggle, for example when the number of variables exceed the number of observations in the data or where the patterns in the data are non-linear.

As shown below, ML can be broadly divided into four categories, according to how the algorithm learns from data.

Supervised learning algorithms try to model relationships and dependencies between a target (e.g. a customer defaulting (or not) on a home improvement loan) and the input features (e.g. the customer's income, the industry they work in, where they live, their age). Targets and features can be defined in different ways. Examples include continuous variables (e.g. a

[1] This, in turn, allows organisations to access and store data in easier ways, aiding the development (and deployment) of algorithms.

numeric variable ranging from 1 to 100) and labelled data (e.g. whether a transaction is fraudulent, not fraudulent, or unknown). Algorithms also place different "weights" on these features. For example, it may be that an algorithm finds a customer's income to be a better predictor of defaulting than the industry the customer works in, so income would receive a higher weighting. Common examples of supervised learning algorithms include: Nearest Neighbour; Naïve Bayes; Decision Trees/Ensemble methods; Linear/Logistic Regression; Support Vector Machines; and Neural Networks.

Unsupervised learning algorithms are mainly used in pattern detection (e.g. fraud detection) and descriptive modelling. These algorithms use statistical techniques on the input data to extract rules, detect patterns and summarise and group data points, which in turn help to derive insights and describe the data. In contrast to supervised learning algorithms, there are no targets (e.g. knowing whether a transaction is fraudulent (or not)) in the data based on which the algorithm can try to model relationships. Examples include various clustering methods (e.g. K-Means, Expectation–Maximization clustering) and association rule learning algorithms

Semi-supervised learning algorithms fall in between supervised and unsupervised learning. In many scenarios, the ability to fully label the data (e.g. to know all previous cases of fraud present in a dataset) is infeasible. In these cases (i.e. where some data is labelled, and some data is not) semi-supervised algorithms are used. These methods exploit the idea that even though labels for the whole dataset are not known (i.e. fraudulent cases vs non-fraudulent cases in the data); the algorithm trains itself on a subset of known cases. In other words, uses the labels indicating 'Fraud' and 'non-Fraud' to derive a 'partially trained' algorithm. The partially trained algorithm then applies what it has learned to the unlabelled cases in the data. In other words, it estimates what label (i.e. 'Fraud' vs 'non-Fraud') is most likely.

Reinforcement learning algorithms continuously learn in an iterative fashion. The algorithm explores a range of possible outcomes. Its outcomes are then compared against actual outcomes (i.e. those observed in the real world). The algorithm is told which of its outcomes are closest to those observed in reality, and, as it iterates over time, its performance improves. Algorithmic performance is the term used to describe how well an algorithm is performing its task (e.g. the closeness of predictions to outcomes).

Note that in all the above cases, different machine learning algorithms can and often should be tested in order to provide the best output.

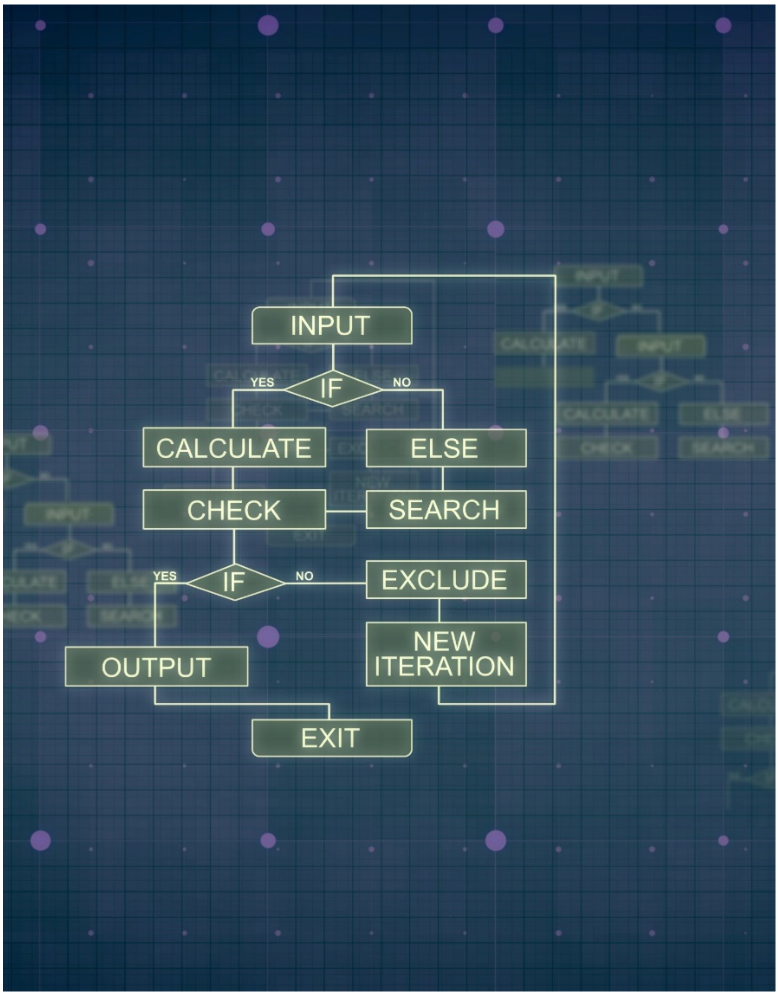

2.2 Use cases of algorithms

With the use of data growing in the last decade, algorithms are increasingly being applied in ways that inform business decisions and automate processes. Below, we explore three overarching categories of their application:

— **Algorithms that improve decisions:** algorithms can be used to process very large quantities of data. As a result, they can provide rich insights, informing decisions in ways which without that processing would not be possible. Applications in financial services include credit scoring, market strategy, pricing and capital optimisation.

  – As an example, traditional motor insurance pricing models were built on a relatively small number of variables (e.g. a vehicle's price, its age, and the customer's age). Now, ML algorithms can incorporate many factors (e.g. previously observed driving behaviour, weather, traffic in a customer's local area), resulting in more accurate predictions of what the optimal price for a policy should be.

— **Algorithms that reduce human effort:** algorithms can be used to reduce the human effort required for specific tasks – typically those which have significant repetition. In turn, this can lead to decreased costs.

  – As an example, many financial services organisations use AI chatbots to communicate with customers. These chatbots use natural language processing ("NLP") and simulate human conversations for customer enquiries, only passing customers to humans for further interactions in more complex cases.

— **Algorithms that solve complex issues:** algorithms can be applied to solve complex problems which have previously been considered as either too hard or too complex for humans.

  – For example, in financial services, fraud detection has long been a challenging and time-consuming task for humans, especially considering the ever-growing number of datasets that, in combination, need to be considered. However, ML-based anomaly detection algorithms are increasingly being used to process real-time data, from a vast array of sources, and discover hidden correlations between user activities and possible fraud.

2.3 Data requirements of algorithms

2.3.1 Why is data important

To train an ML algorithm correctly, it is essential to have the right data. It needs to be accurate, and there typically needs to be enough variation in the data for the algorithm to find patterns. If there are issues with the data going into an algorithm, there will be issues with the output. Three aspects of data that affect algorithmic performance are:

— **Data sample size**

  Sample size is the number of individual pieces of data. Datasets that are too small may lead to an algorithm failing (e.g. being unable to identify patterns) or lead to bias in results. Some categories of algorithms require significantly higher volumes of data relative to others.

  A critical step in training an algorithm, for example using an algorithm to understand relationships between a target (such as a loan default) and a set of features (including, for example, credit score) is to split the dataset into (at least) two parts. Typically, algorithms first learn on a "training" dataset, then the results are applied to a "test" dataset to check their validity. This approach helps prevent several issues that can occur when building models, such as overfitting, which limits its generalisability in real-world settings. In the context of sample size, it is therefore important to have a dataset that is large enough to create valid training and testing datasets.

Approaches are increasingly being refined and explored that attempt to minimise the impact of small sample sizes in machine learning, such as bootstrapping and cross validation. We do not cover these here.

— **Data dimensions**

  Data dimensions are the data fields, known as features or variables (e.g. age, gender, income). For the supervised learning models discussed above, these should be the features that affect what is being predicted (i.e. the target). Without enough of the relevant features, algorithmic performance will suffer. For example, if the target is to predict whether a customer will default on a car loan, a range of potential features (e.g. customer historical behaviour, income, levels of other debt, job type) could be covered in the data fields. If certain features were missing (e.g. income or other debts) the predictive accuracy may be reduced, potentially significantly.

  At the same time, irrelevant features may reduce model performance. For example, including whether a customer supports a certain football team would be unlikely to add to its predictive accuracy; it creates 'noise' in the training step. These variables should be filtered out of the data using feature selection approaches.

  In addition to features present in raw data, using domain knowledge can help create new features that add to predictive accuracy further. For example, using the postcodes in the raw data could generate new features, such as 'closest city names', or 'distance to the city centre'. This process is referred to as feature engineering.

— **Data quality**

  Inaccurate data will impact model performance. Errors can include typos or mislabelled information, missing values in some data fields, and data in inconsistent forms (e.g. a number being stored as text). Appropriate data cleaning and data imputation approaches should be conducted to prepare data ahead of modelling.

2.3.2 Sources of data

Data sources can be divided into two main categories:

— **Internal data** are generated and collected within the business. Internal data are held at both the customer and organisation level. It includes:

  – Back-office data such as human resource, operations maintenance, finance and governance data.

  – Middle-office data, such as customer support or other business support (e.g. data collected at contact centres).

- front-office, such as any data generated directly from customers (e.g. applications for financial products or how a customer has responded to a marketing promotion).

These data can be both structured (e.g. data tables saved in databases in a row/column format) and semi-structured or unstructured data (e.g. paper documents or call records). Organisations increasingly use sophisticated data environments and hold data strategies covering data collection, exploitation and storage in order to best use internal data.

— **External data** are generated outside of the business. It includes:

  – Open data, which are data that anyone can freely access, use and share. Examples of these data include economic information from governments, weather and geographic data. These data often have a licence to permit usage, including transforming, modifying, reusing and redistributing the data, even commercially. Fees may be charged for access to cover the cost of creating, maintaining and publishing usable data.[2]

  – Paid data, which are available from third-party data providers at a cost.

  – Trust-based shared data, which is data shared between companies. For example, ten large pharmaceutical companies, including Johnson & Johnson, AstraZeneca and GSK, agreed on a data collaboration strategy in June 2019, so that they can train their machine learning algorithms on each other's data to accelerate drug-discovery.[3]

  – Social data is collected from social media platforms. It contains user posts and relevant metadata, such as the number of shares/likes, hashtags, comments, as well as post location and time. Social data are often free to access but may require licences for use subject to copyrights.

Prior to using data, it is important to understand relevant data laws. For example, the General Data Protection Regulation (GDPR) covers data protection and privacy in the EU. It came into force in May 2018, and applies to all organisations operating within the EU as well as organisations outside the EU that offer goods or services to individuals in the EU.[4] The Data Protection Act 2018 (DPA 2018) is the UK's implementation of the General Data Protection Regulation (GDPR), which controls how personal information is used by organisations, businesses or the government.[5]

2.4 Using algorithms in the day-to-day

In this section, we set out the process of moving from an idea as to how an algorithm may support business decision making through to using it in reality – known as being "in production".

2.4.1 Proof of concept model

The first stage is to build a **proof of concept ("PoC") model**. It includes three stages which are not mutually exclusive – organisations will often go through multiple iterations of each step prior to moving on to deployment.

An essential first step when building the PoC is to identify a clear **use case**. Doing so typically requires consulting a range of stakeholders across an organisation – from those with business understanding to shape commercial value through to technical modelling expertise to assess feasibility. In some cases, business owners or end users may lack a clear understanding of data availability and over-estimate their data quantity and quality, and hence raise unrealistic use cases.

Data preparation includes obtaining the data from the sources described above. As part of this, organisations typically define how frequently these data will need to be updated going forward, which will in turn inform how often the model itself is updated and deployed at the later stage. Not all data is necessarily obtained during a PoC phase.

[2] https://www.europeandataportal.eu/elearning/en/#/id/co-01

[3] Accessed via Financial Times, https://www.ft.com/content/ef7be832-86d0-11e9-a028-86cea8523dc2

[4] Information Commissioner's Office, 22 March 2018, Guide to the General Data Protection Regulation (GDPR)

[5] https://www.gov.uk/data-protection

Initial **modelling and testing** is then completed across three phases:

— Training: selecting an algorithmic model and using training data to train the model.

— Validation: "tuning"[6] to obtain the best performing model.

— Testing: using a separate testing dataset in order to provide an unbiased evaluation of the model's performance.

2.4.2 Model deployment

If the PoC is successful, then the second stage is to move to **deployment**.

There are several aspects to be considered in deployment. These include choosing:

— A **user interface.** This should maximise the usability and the end user experience after deployment.

— A suitable **deployment model**. Organisations increasingly use Cloud technology to deploy algorithms. They must choose between three commonly used deployment models: on-premises (sometimes called private clouds); cloud (sometimes called public clouds); or a hybrid approach which mixes cloud-based and on-premises resources.

— An appropriate **deployment pattern**. These are effectively the frequency with which the algorithm needs to be used. At a high level, there are three deployment patterns:

  – Real-time: instant use of the model is required (e.g. chatbots);

  – Streaming: the model performs every time an input changes; and

  – Batch inferencing: the model performs periodically, typically processing a set of historical data at once (e.g. a weekly credit risk prediction).

— The process for **monitoring, feedback and iteration**. It is important to monitor and get feedback on model performance. This can be used to adjust the model as needed. This should be a continual, iterative process supported via a Continuous Integration/ Continuous Deployment (CI/CD) pipeline. Evaluations should be conducted frequently – for example to explore whether the system is consistently performing as expected, or whether outcomes show any bias.

[6] Tuning refers to a process whereby certain inputs (for example, the number of trees used in a tree-based model) into the algorithm are changed, and the algorithmic performance then compared to alternative inputs.

2.5 Key limitations

The discussion above shows the clear range of potential benefits of using algorithms. However, there are several potential challenges to implementing them. These include:

— **Data limitations**: algorithms make decisions based on historical (or input) data. As discussed earlier in this section, these data can be wide ranging and potentially costly to gather. Input data also directly affects the performance of the algorithm; if data is lacking either in quantity or quality, the output from the algorithm itself may not be robust.

— **Explainability and transparency**: another challenge in using algorithms is that they can be difficult to explain and may not be transparent in their decision-making process for end users. Some algorithms, such as generalised linear regression, are easier to explain – they produce numerical outputs which can be easily linked back to input data. On the other hand, some algorithms can be harder to explain, both in terms of the process followed as well as the outputs reached (e.g. Neural Networks).

— **Algorithmic bias**: if not properly controlled for, the output of on algorithm can contain bias. In these cases, the algorithm may suggest outcomes which are systematically less favourable to, for example, individuals within a group, where there is no relevant difference between groups that justifies such a difference. Algorithmic bias is discussed further in the next section.

03

Algorithmic bias

As the adoption of AI continues to accelerate, potential issues and unintended consequences of its use in business are increasingly being explored. This section of our report discusses one such issue; algorithmic bias.

We start with some background on algorithmic bias, including how it can be defined. Definitions of bias can vary. For example, different countries have different laws to help counteract discrimination or bias, the academic literature on what constitutes a potential bias continues to evolve, and regulators are taking varying approaches in how biases are considered. In this section, we present a range of algorithmic biases, drawing on examples which have been discussed in the academic literature and by regulators, along with potential mitigants. This is not intended to be an exhaustive list, nor do we suggest that all biases can – or should – be fully mitigated. Similarly, while we discuss certain legal and regulatory approaches that have been introduced by governments and regulators in order to reduce the negative impact of algorithmic bias, we do not intend this section to provide an exhaustive summary of relevant legal and regulatory issues.

3.1 Background

Algorithmic bias can be defined in a variety of ways and can also mean different things in different contexts. In the broad sense, algorithmic bias *"relates to outcomes which are systematically less favourable to individuals within a particular group and where there is no relevant difference between groups that justifies such harms"*.[7]

As algorithms become more broadly used, bias could have significant harmful impacts on the users of products and/or services that rely on algorithms. This could, in turn, also lead to difficulties for the developers and users of the algorithms. It is important for the developers and users of the algorithms to understand how bias can manifest, and thus follow appropriate mitigations to ensure outcomes are as fair as possible.

Legally protected characteristics

In the UK, it is against the law to discriminate against someone based on several 'protected characteristics' which have been defined by the Equality Act 2010[8]. These are:

[7] https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/
[8] https://www.equalityhumanrights.com/en/equality-act/protected-characteristics

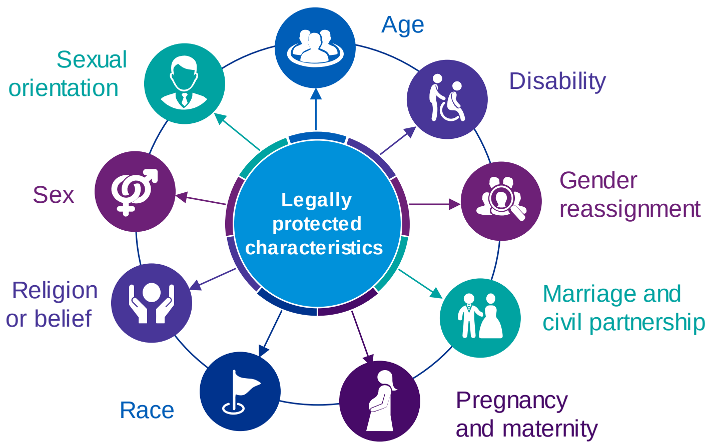

Similar laws are also in place in many other jurisdictions. These include those jurisdictions which Finastra have asked us to estimate market sizes of certain products for in this report, including the United States, France and Hong Kong. Although Singapore does not have any legislation which expressly prohibits discrimination on the grounds of race, ethnicity, religion, gender, disability or sexual orientation, the constitution outlines that all persons are entitled to equal protection of the law and that there shall be no discrimination based on religion, race, descent or place of birth.[9,10]

3.2 Factors that can lead to algorithmic bias

There are several factors that can lead to algorithmic bias. The main factors include:

— Biases in data;

— Biases in algorithmic design; and/or

— Biases in human use.

We now discuss each of these factors in turn.

3.2.1 Biases in data and mitigation strategies

Algorithmic bias can arise from the use of underlying data which also contains bias. Bias in data can come from a variety of sources. These include, but are not limited to:

— **Historical bias** may exist in the data due to historical factors. For example, it may be the case that residents in a specific area have historically had relatively higher loan default rates than the national average. This could therefore feed into loan decisions for that group and may result in a higher proportion of declined loan applications.

Such biases may also be reinforced over time. For example, the prevalence of the higher loan default rates in historical data may further reduce the applicants' credit scores, thus increasing the likelihood of the future loan applications being declined, which is a data point that may get used in future iterations of the algorithm. This is known as a 'feedback loop'.

— **Sample bias** may exist in the data either because the sample size is small and/or it is skewed towards different groups that are not representative of the entire population. For example, if an

[9] DLA Piper, Employment Discrimination, https://www.dlapiperintelligence.com/goingglobal/employment/index.html?t=09-discrimination

[10] "While the Constitution provides that all persons are entitled to the equal protection of the law and that there shall be no discrimination based on religion, race, descent or place of birth, successful challenges on constitutional grounds are rare." (DLA Piper, Employment Discrimination)

insurance algorithm is trained using a small dataset in which most claimants happen to be from a certain area, the algorithm may increase the premiums for people from this area.

— **Human bias** is often included in the data that algorithms learn from, which is especially the case in supervised machine learning. These algorithms essentially learn what they have been taught by humans. They can therefore carry all the conscious or unconscious biases from human judgement. For example, if an algorithm is using the decisions made by an insurance underwriter as labels to assess the risk in insuring a car or home, it may replicate any biases prevalent in the underwriter's decisions.

— **Measurement bias** arises due to how data was measured, collected and/or stored. This can range from how decimal numbers are rounded up or down in systems, through to systematic data recording errors, due to faulty equipment and/or operation. Measurement bias usually applies systematically across all relevant data. However, in some cases, it may incorrectly impact a particular group more than other groups (e.g. if data collection methods are poorer in one area compared to another).

— **Proxy bias** can arise where variables used within an algorithm are correlated – either on a standalone basis, or once combined with other variables – with potentially sensitive attributes (e.g. use of a geographical variable, such as a neighbourhood, which is correlated with particular ethnicities). In most cases, it is challenging to determine if a variable is correlated with protected characteristics and if it should be included in algorithm training. Proxy bias can mean that simply excluding the protected attributes of individuals in the data may not eliminate or even mitigate the potential impact of bias on particular groups. In fact, protected attributes are often included in data so that adequate measures of bias and the effectiveness of mitigation methods (which we explore below) can be gathered.

— **Emergent bias** is more difficult to anticipate and in turn control for. This refers to a bias which emerges some period after an algorithm is used, often as a result of changing societal knowledge, population, or cultural values.[11] As a result, the data selected and used within the algorithm may not reflect the emergence of the new knowledge or societal values.

As an example, COVID-19 has significantly changed the way that societies around the globe are operating. In the case of financial services, algorithms previously using income will now need to consider how to control for, for example, furlough payments. Insurance algorithms will need to start accounting for the relative risk profile of different occupations. Without appropriately considering these new factors in the data, and in an algorithm's design, algorithms may generate biased results – for example, to favour groups who were least affected by COVID-19. Furthermore, even with the same set of variables in an algorithm(s), the pattern in the data may change over time (known as 'data drift'), which raises the necessity of regularly updating/re-training the model to mitigate this.

Mitigation strategies can be used to tackle the bias in **data**. Several strategies can be considered by the developers and users of algorithms in different steps/stages**.**

— **Data defining:** ensuring careful, methodical selection of features/variables that can support the defined business problem.

— **Data gathering:** checking whether the dataset is representative (or not) for defined groups. If not, collecting more data for underrepresented group(s) may be required (if possible). It will also be important to minimise any errors in data collection, measurement and storage process.

— **Data labelling:** setting guidelines and rules for mitigating bias in data labelling. This can be done by using labelling groups and checking processes to mitigate subjective bias from individuals.

— **Data pre-processing:** for a given dataset, a variety of **pre-processing** methods can be used to transform the original data before inputting it into the algorithm in a bid to mitigate potential

[11] Batya Friedman et al., 1996, Bias in Computer Systems

biases. These include reweighting the training samples in each group[12] and removing disparate impact from the features so that the group membership can no longer be inferred.[13] More recent approaches have been to use a form of machine learning – conditional Generative Adversarial Networks (cGANs) – to generate synthetic, fair data with selective properties from the original data.[14]

3.2.2 Biases in algorithmic design and mitigation strategies

In addition to bias in the data itself, bias can also be introduced during the design of an algorithm. The outcomes of algorithms can be different even when the same input data has been used. Biases in algorithmic design include:

— **Objective bias** can arise from the aim of the algorithm itself. For example, an algorithm used for assessing the credit default risk of mortgage borrowers may aim to maximise the *overall* success rate of predictions. However, over or under recording of particular groups could mean the algorithm is more accurate for some groups than others – impacting their outcomes – even though the overall accuracy is maximised across all of the groups.

— **Weighting bias** reflects the fact that weights are often applied to features in an algorithm, and if these are not applied correctly, outcomes can be impacted.

For example, an algorithm used to estimate car insurance premia may use features such as a person's age and where they live, but these may not carry equal weight. Studies have found that algorithms used by many car insurers in the US were relying on credit scores more heavily than driving records. This meant that a single driver who only had a 'good' credit score paid $68 to $526 more per year, on average, than similar drivers with the best scores, depending on his or her home state.[15]

— **Evaluation bias** arises from the methods used to evaluate algorithms. As set out in Section 2 on algorithm fundamentals, it is common practice to divide data into the training and test datasets to check algorithmic performance. Inappropriately split data, however, can lead to the sample bias described earlier in this section, creating a subsample of data which has gone from having no bias, to having bias introduced.

Mitigation strategies can be used as part of algorithmic design:

— **In-processing** mechanisms can be adopted, which modify the algorithms directly to account for fairness. These methods include incorporating fairness into the objective function or imposing a constraint in the model training process.[16,17] Several fairness metric definitions are available such as *Equalised Odds*, *Equal Opportunity*, *Demographic Parity*, and *Counterfactual Fairness*.[18]

— **Post-processing** mechanisms can be performed after training by accessing a dataset which was not used in the algorithm's training (i.e. **Post-processing** mechanisms can be performed after training by accessing a dataset which was not used in the algorithm's training (i.e. test dataset). These methods include reassigning the labels predicted by the algorithm.[19]

[12] Faisal Kamiran et al., 2011, Data Pre-Processing Techniques for Classification without Discrimination
[13] Michael Feldman et al., 2014, Certifying and Removing Disparate Impact
[14] http://ecai2020.eu/papers/348_paper.pdf
[15] The Truth about Car Insurance, https://www.consumerreports.org/cro/car-insurance/auto-insurance-special-report/index.htm
[16] Toshihiro Kamishima et al., 2012, Fairness-Aware Classifier with Prejudice Remover Regularizer.
[17] Elisa Celis et al., 2018, Classification with Fairness Constraints: A Meta-Algorithm with Provable Guarantees.
[18] Ninareh Mehrabi et al., 2019, A Survey on Bias and Fairness in Machine Learning
[18] Rachel KE Bellamy et al., 2018, AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias

— Finally, algorithms can be tested using **bias detection** techniques from comparing outcomes for different groups to computer simulations. This can be done before the algorithm is put into production. Researchers have recently developed a three-level rating system which can determine the relative fairness of an algorithm.[20]

3.2.3 Bias in human use

Algorithmic outcomes are often interpreted and used by humans. As a result, the use of the outcomes can also be subject to bias. These biases can include decision bias, belief bias and interpretation bias. We do not discuss these human biases further in this report. However, it is important for organisations to be aware of the potential for bias in the human use of algorithm outputs. **Mitigation strategies** can also be used to help prevent this. It is important for algorithm users to understand the functions and limitations of the algorithm under different contexts and, as part of that, algorithm developers should provide clear handover instruction. Training can also be provided to users on the types of biases discussed above.

UK Competition and Markets Authority work on algorithmic harms

— There has been an increasing focus on the use of algorithms in recent years including any potential unintended consequences associated with their use.

— In the UK, the Competition and Markets Authority (CMA) has recently launched a programme of work on analysing algorithms, with an aim to better identify and address harms caused by algorithms.[21]

— The CMA published a paper in January 2021 describing harms which might arise from algorithms, many of which involve personalisation. The paper notes that personalisation can be harmful as it might not be transparent, might target vulnerable consumers or have unfair distributive effect.

— The CMA is likely to take further actions in this space, by providing guidance to businesses around algorithms as well as identifying and remedying existing harms.

[20] Biplav Srivastava et al., 2018, Towards Composable Bias Rating of AI Services

[21] CMA, Algorithms, competition and consumer harm: call for information, https://www.gov.uk/government/consultations/algorithms-competition-and-consumer-harm-call-for-information

04

Algorithms in financial services

Financial services organisations are becoming increasingly aware of the significant benefits that using algorithms can deliver, from improving the customer experience and organisational productivity through to enhancing human interaction with the financial service ecosystem. Algorithms are enabling the development of new products and in turn demand that would not have been possible using previous technologies.

Given the above, organisations across the sector are undergoing a fundamental shift in the way they operate and interact with their customers. In short, organisations have had to transform into data and insight driven businesses.

A combination of external factors has fuelled this shift. Key factors include:

- Rapidly changing consumer preferences – informed both by what consumers observe from competitors within the sector, and the experiences they have outside of it – is requiring organisations to offer a personalised customer experience across a range of services and products
- Increasing regulatory pressure is driving organisations to establish better control and transparency across their business operations
- Growth in eCommerce is driving organisations to deploy counter measures to safeguard data and customers from cyber-attacks and fraudulent activities
- Uncertain economic conditions and physical location shutdowns during the Covid-19 pandemic has driven enforced digitalisation across the sector
- A reduction in margins, coupled with competition from new, disruptive businesses, have driven a focus on cost reduction, which can be facilitated by technology
- Continued maturity of AI algorithms and models supported by a substantial increase in the range of available data. In certain jurisdictions, this includes the ability to share data through “open banking”

Organisations continue to transform into the data and insight driven businesses described above. Namely, they are:

— leveraging transactional and behavioural data on customers using algorithms to drive vast improvements in customer experience; and

— embedding automation across core businesses and a range of decision-making processes, driving out cost.

In the remainder of this section we explore what the above means in practice.

To do so, we set out the various “use cases” of algorithms in financial services, which includes discussions of algorithms that lead to direct impacts to customers compared to those which are indirect- (i.e. that impact the customer, but relate to parts of a business that the customer is not directly exposed to). The purpose of this section is to give an overview as to the types of cases available to organisations in the sector. We then discuss the potential impact of algorithmic bias given the scope of these use cases.

4.1 Algorithm driven financial services use cases

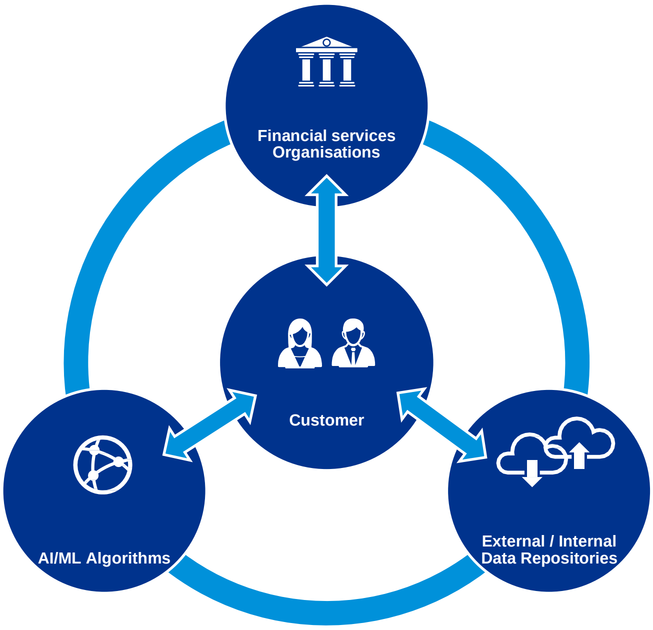

As discussed above, financial services organisations are increasingly using algorithms. Their impact on customers can be classified into two categories – direct impacts through either decisions or interaction, and indirect impacts. We discuss each of these in turn, including examples of relevant use cases.

4.1.1 Direct impacts

Impacts on decisions

Algorithms are being used to make direct decisions on customers based on their data, in turn impacting customers’ choices and behaviour. Below, we explore two key use cases – credit scoring and decisioning (relevant in consumer finance), and automated pricing (relevant in insurance). We provide further detail on how algorithms are used in consumer credit more generally, for example in improving customer experience, in Appendix 2.

Automated credit scoring and decisioning: simplistic, rules-based algorithms have been around and used for years for credit scoring and decisioning. What is new is the shift towards using more advanced techniques, such as the machine learning methods described above, to assess an individual’s or organisation’s creditworthiness using substantial volumes of internal and external data.

At a high level, a scoring approach is used to evaluate data provided by a potential customer – often merged with existing data – in real time. An underlying algorithm then decides to approve, refuse or offer personalised options for a credit application based on the resultant credit score. These methods allow lenders, for example, to distinguish between high default risk applicants and those who are credit-worthy but lack an extensive credit history. At the same time, successful applicants get personalised offers that are aligned to their personal circumstances (e.g. financial status).

Depending on thresholds set by an individual organisation, certain credit decisions are now fully automated. Automated and semi-automated credit scoring and decisioning is widely used by credit providers – including major retail banks – across a range of products, including credit cards and personal loans.

Automated pricing: similar to consumer finance, in the insurance sector, rules-based algorithms have been historically used to determine the risk profile of customers and in turn inform the pricing for a policy. Now, more complex algorithms, including many of the techniques set out in Section 2 above, are being used to assess this risk in real-time based on significant volumes of data. These data include inputs provided by applicants, such as health data and historical claims data, merged with internal datasets to more accurately predict individual risk. In turn, insurers can better identify the optimal prices for premia.

In terms of operations, this shift has helped insurance organisations to reduce the volume of manual work required by employees, reducing costs. In parallel, customers increasingly use price comparison websites to compare providers and identify suitable options in real-time based on the product features and pricing offered.

Impacts on interactions

In addition to making decisions that directly impact the customer, algorithms are being used in day-to-day interactions with financial services providers. Below, we explore a key use – personalised customer self-service. This is relevant both in financial services and across many other sectors.

Personalised customer self-service: Customers are increasingly used to experiencing personalised interaction with companies – regardless of the sector. Part of this includes an ambition to interact with companies in what is close to a real-time way (e.g. avoiding the long wait queues that can be generated in traditional call centres).

To meet this demand in financial services, automated self-serve assistants (e.g. chatbots) use algorithms to generate personalised financial advice based on interactions with customers through an online platform. Natural language processing algorithms are increasingly able to provide instant, self-help customer service, in areas that include customer support, budgeting, setting up savings goals and tracking expenses. These tools are enabling financial services providers to reduce operational support costs whilst providing relevant services to their customer base on a 24/7 basis. Personalised customer self-service is widely used by new-generation challenger banks and insurance providers.

4.1.2 Indirect impacts

Above, we described instances where algorithms are being used to make decisions that directly impact customers. Algorithms are also used, however, in ways that the customer is less aware of. For example, in broader eCommerce, this includes algorithms used for personalised marketing and to alter the choice architecture provided to individual customers. There are also several use cases specific to financial services. Here we focus on two such cases: risk monitoring; and management, and fraud detection.

Predictive risk monitoring and management: as set out above, organisations use algorithmic approaches to determine both credit risk and insurance risk for customers at the point of application, in turn using these to make decisions. Algorithms are also used to develop risk detection models for internal use within financial services organisations. These aim to detect the probability of adverse events occurring (e.g. a customer defaulting on credit) and in turn estimating the associated cost. These models help risk and data experts to discover trends, detect risks, conserve resources and provide better information for robust management of risk.

These algorithms can, on occasion, be used in ways that eventually have “direct” impacts on the consumer. For example, they can enable credit lenders to predict potential defaults ahead of time, and in turn estimate suitable payment plans with financially distressed customers to reduce future default risk, benefitting both parties. Predictive risk monitoring is widely used by credit providers, including major retail banks, across a range of products, including credit cards and personal loans

Real-time fraud detection: growth within eCommerce and online transactions has led to a rise in the number of fraud cases across consumer finance products, including credit cards and online banking transactions. Algorithms increasingly play a crucial role in helping to prevent and identify this type of fraud. For example, fraud detection algorithms analyse customer’s previous purchasing behaviour, locations and transaction values to establish usual behaviours. Where unusual behaviours are then identified by the algorithm, financial services providers are able to generate alerts and/or block/validate transaction before they can be processed. Most of these checks are therefore invisible to the customer.

Methods are continually being refined in order to increase the accuracy of interventions, reducing friction for customers whose transactions are erroneously flagged. Real-time credit fraud detection is widely used by global card payment organisations and major retail banks.

The use cases above are a subset of how algorithms are applied across the financial services sector. They demonstrate the agility, efficiency and reliability that algorithm-based systems can offer, providing new approaches to meet the growing demands of customers.

4.2 Potential for bias

In Section 3, we set out a non-exhaustive list of 12 biases for organisations to be aware of when designing, building and putting algorithms into production. These included biases specific to **data, algorithmic design, and human involvement.**

Above, we set out five cases where algorithms are being used in the financial services sector. These share some common themes, including:

a) To varying extents, they all require significant volumes of **data** to be effective, and these data are likely to change on a frequent basis;

b) The choice of **algorithmic design** is important, along with continuous evaluation, for the results to be meaningful; and,

c) The level of **human involvement** in eventual decision making varies, depending on the preferences of the organisation.

When deploying algorithms for the use cases set out above, it is therefore important that financial services organisations are aware of the potential for algorithmic bias. Not doing so carries risks that could include legal action, regulatory enforcement, and a loss of trust with the consumer. Whilst examples of relevant mitigation strategies have been set out in this report, these strategies continue to evolve. As a result, it is also important that these mitigation strategies, as well as algorithms, are continually reviewed.

05

Market sizing

As outlined in Sections 3 and 4, algorithmic bias can lead to the users of financial products and/or services which rely on algorithms experiencing a degree of harm. The quantum of this harm will depend on the existence, prevalence and magnitude of bias, the context in which the algorithm is applied and the size of the market in which the algorithm is used.

In this section, we provide estimates of the market size of various consumer finance products across five jurisdictions.

Given the differences in reporting across jurisdictions, the market size of each product is not always easily comparable across geographies on a like-for-like basis. This is discussed further in the remainder of this section and in Appendix 1.

5.1 Overarching assumptions

This section details the observed annual consumer expenditure on a range of financial products across five jurisdictions. The data used is summarised in Table 1 below.

Table 1: Data used to estimate consumer expenditure

Finance category	Product	US	UK	France	Singapore	Hong Kong
Consumer lending	Credit cards	The data collated for the three consumer lending products is compiled on the same basis by Euromonitor across all jurisdictions. However, these data are stated as either gross lending or transaction value. This is discussed further in Appendix 1.				
	Mortgages					
	Other forms of consumer credit and lending[22]					
Insurance	General insurance	Data is sourced from the Insurance Information Institute (III). This is based on net direct premiums.	Data is sourced from Mintel. This is based on gross premiums.	Data is sourced from Fédération Française de l'Assurance (FFA). This is based on gross premiums.	Data is sourced from the Monetary Authority of Singapore (MAS). This is based on gross premiums.	Data is sourced from the Hong Kong Insurance Authority (HKIA). This is based on gross premiums.
	Life insurance					
	Annuities					

As set out in Table 1, data for the three types of insurance products have been collated from multiple sources, each of which makes its estimations under different assumptions. It is therefore

[22] As defined by Euromonitor. See Appendix 1 for details.

important to consider each jurisdiction independently according to the observations set out in Sections 5.2 to 5.6 and Appendix 1, which contains the assumptions underpinning the numbers presented, rather than comparing data between jurisdictions.

5.2 USA

5.2.1 Consumer lending

The estimated size of the consumer lending market in the US is set out in Table 2 below.

Table 2: Consumer lending in the US

Product	Gross Lending/Transaction Value ($bn)			
	2017	2018	2019	2020
Credit cards	2,257.6	2,485.2	2,688.6	2,382.6
Other consumer credit/lending	1,998.8	2,030.2	2,042.9	1,869.4
Mortgages/housing	1,714.1	1,623.8	2,104.6	1,855.8
Total consumer lending	**5,970.5**	**6,139.2**	**6,836.1**	**6,107.8**

Source: Euromonitor

5.2.2 Insurance

The estimated size of the personal insurance market in the US is set out in Table 3 below.

Table 3: Consumer expenditure on insurance in the US

Product	Net direct premiums ($bn)		
	2017	2018	2019
Health			757.4
Property and casualty (P/C)	297.0	328.7	340.7
General insurance			1,098.1
Accident and health	64.6	62.8	64.7
Life	144.3	143.1	149.8
Life insurance	208.9	205.9	214.5
Annuities	181.8	207.8	217.5
Total insurance			**1,530.1**

Source: III and KPMG analysis

5.3 UK

5.3.1 Consumer lending

The estimated size of the consumer lending market in the UK is set out in Table 4 below.

Table 4: Consumer lending in the UK

Product	Gross Lending/Transaction Value (£bn)			
	2017	2018	2019	2020
Credit cards	119.1	124.8	126.1	94.2
Other consumer credit/lending	139.3	146.2	151.3	112.5
Mortgages/housing	261.2	269.3	268.0	234.9
Total consumer lending	**519.6**	**540.3**	**545.4**	**441.6**

Source: Euromonitor

5.3.2 Insurance

The estimated size of the personal insurance market in the UK is set out in Table 5 below.

Table 5. Consumer expenditure on insurance in the UK[23]

	Premiums (£bn)			
Product	**2017**	**2018**	**2019**	**2020 (est.)**
Motor insurance	13.0	13.1	13.0	13.3
Home insurance	5.9	5.9	6.1	6.1
Accident and health	3.5	3.7	5.1	5.0
Miscellaneous	4.2	4.2	5.0	5.0
General insurance	26.6	26.9	29.2	29.4
Annuities	4.4	4.4	4.3	3.6
Drawdowns	21.4	23.1	24.4	21.5
Annuities/Drawdowns	25.8	27.5	28.7	25.1
Total consumer insurance	**52.4**	**54.4**	**57.9**	**54.5**

5.4 France

5.4.1 Consumer lending

The estimated size of the consumer lending market in France is set out in Table 6 below.

[23] An estimate for life insurance premiums in the UK is not available from Mintel on a like for like basis when compared to other insurance products. As such it has not been included in Table 5.

Table 6: Consumer lending in France

Product	Gross Lending/Transaction Value (€bn)			
	2017	2018	2019	2020
Credit cards	42.3	45.3	47.7	42.2
Other consumer credit/lending	50.5	51.7	53.7	48.7
Mortgages/housing	160.1	168.1	193.5	189.8
Total consumer lending	**252.9**	**265.1**	**294.9**	**280.7**

Source: Euromonitor

5.4.2 Insurance

The estimated size of the personal insurance market in France is set out in Table 7 below.

Table 7: Consumer expenditure on insurance in France

Product	Premiums (€bn)		
	2017	2018	2019
Health and accident	22.5	23.7	24.8
Auto	21.4	22.1	22.8
Property	10.5	10.8	11.3
General insurance	54.4	56.6	58.9
Life Insurance	11.7	12.2	12.7
Annuities	122.9	127.5	131.8
Total insurance	**191.5**	**198.8**	**206.2**

Source: FFA and KPMG analysis

5.5 Singapore

5.5.1 Consumer lending

The estimated size of the consumer lending market in Singapore is set out in Table 8 below.

Table 8: Consumer lending in Singapore

Product	Gross Lending/Transaction Value (SG$bn)			
	2017	2018	2019	2020
Credit cards	38.9	39.9	43.2	39.3
Other consumer credit/lending	45.3	44.6	45.0	42.6
Mortgages/housing	27.7	28.5	28.4	27.8
Total consumer lending	**111.9**	**113.0**	**116.6**	**109.7**

Source: Euromonitor

5.5.2 Insurance

The estimated size of the personal insurance market in Singapore is set out in Table 9 below.

Table 9: Consumer expenditure on insurance in Singapore

Product	Premiums (SG$bn)		
	2017	2018	2019
Singapore Insurance Fund (SIF)			2.7
Offshore Insurance Fund (OIF)			9.9
General insurance			12.6
Non-linked	13.4	13.2	13.2
Linked	4.3	4.2	3.0
Life Insurance	17.7	17.4	16.2
Annuities	0.3	0.4	0.6
Total consumer insurance			**29.4**

Source: MAS and KPMG analysis

5.6 Hong Kong

5.6.1 Consumer lending

The estimated size of the consumer lending market in Hong Kong is set out in Table 10 below.

Table 10: Consumer lending in Hong Kong

Product	Gross Lending/Transaction Value (HK$bn)			
	2017	2018	2019	2020
Credit cards	581.8	653.1	677.2	604.4
Other consumer credit/lending	238.4	253.7	283.2	282.0
Mortgages/housing	396.4	437.5	441.6	386.2
Total consumer lending	**1,216.6**	**1,344.3**	**1,402**	**1,272.6**

Source: Euromonitor

5.6.2 Insurance

The estimated size of the personal insurance market in Hong Kong is set out in Table 11 below.

Table 11: Consumer expenditure on insurance in Hong Kong

Product	Premiums (HK$bn)		
	2017	2018	2019
Accident and health	11.1	12.4	13.8
Motor vehicles	3.3	3.5	3.8
Property damage	2.2	2.1	2.3
General insurance	16.6	18.0	19.9
Non-linked	137.9	133.2	139.4
Linked	12.7	17.4	11.8
Life Insurance	150.6	150.6	151.2
Annuities	7.7	10.8	20.9
Total consumer insurance	**174.9**	**179.4**	**192.0**

Source: HKIA and KPMG analysis

Appendix 1

Market size assumptions

This appendix details the assumptions underpinning the market size estimates set out in Section 5 above.

Consumer lending

In the consumer lending tables in Section 5 above, credit cards is the total value of personal (as opposed to corporate) credit card transactions. Other consumer credit/lending is comprised of auto lending, card lending (less personal credit card transactions), durables lending, education lending, home lending and other personal lending. Mortgages/housing is the total value of mortgage lending across the US.

These figures refer to gross lending or transaction values, and therefore not consumer spending on the cost of credit for these products. The market for consumer credit and lending is diverse and the cost of credit varies significantly across each of the types of lending it covers. For example, the average interest rate charged on credit cards (the annual percentage rate or "APR") is likely to be significantly different to the rate charged on a mortgage. Even within each product category, the interest rate charged is likely to vary substantially. For example, the APR for credit cards will vary depending on the credit score of the customer and the type of card used. Consumer attitudes towards credit are also variable, with some customers paying off balances in a timely manner and incurring minimal or no fees, while others use cards as a credit facility and, willingly or not, incur charges. Similarly, there are a range of mortgage types available and the cost of credit for the consumer will depend on whether rates are fixed or variable, whether payments are allocated to interest or principal and the rates themselves.

As a result, it is difficult to calculate a typical cost of credit from which consumer expenditure on credit cards can be extrapolated. Of course, even a small percentage applied to the market sizes set out in the table above remain significant numbers in absolute terms.

Insurance

USA

Data in Table 3 is from the Insurance Information Institute (III). In the US, general insurance can be split into health insurance and property and casualty (P/C) insurance. The III data covers life insurance and P/C insurance, but only references health insurance in its 2019 report. We have therefore only provided health insurance data for 2019.

Additionally, the split between commercial and personal lines for P/C insurance is only available for 2019. To estimate the personal expenditure for P/C insurance in 2017 and 2018, we have applied the same split to overall P/C expenditure. Furthermore, some health insurance premiums are covered by P/C and life insurance. The data in Table 3 classifies health insurance premiums according to the III's categorisation.[24]

The consumer expenditure recorded across all insurance categories is net direct premiums, that is gross written premiums less expenses. Direct premiums are those that do not include reinsurance premiums.

UK

Data in Table 5 comes from Mintel, which provides actual figures from 2017 to 2019 and estimates for 2020. General insurance is split by motor insurance, home insurance, accident and health and miscellaneous, which incorporates other types of insurance such as travel and pet. General insurance premiums are gross (as opposed to net) and direct.

Annuities and drawdowns are listed separately in Mintel's data and have been presented as such here.

An estimate for life insurance premiums in the UK is not available from Mintel on a like for like basis when compared to other insurance products. As such it has not been included in Table 5.

The data recorded across all insurance categories is gross premiums.

France

All of the data in Table 7 is from the Fédération Française de l'Assurance (FFA). FFA categorises its data into "personal insurance" (*assurance de personne*) and "property and liability insurance" (*assurance de biens et de responsabilit*é). Personal insurance includes endowment insurance (*assurance en cas de vie*) and endowment contracts (*contrat de capitalisation*), both of which have been categorised as annuities in Table 7.

Personal insurance also includes life insurance (*assurance en cas de décès*) and health and accident insurance (*assurance en cas de maladie et d'accidents corporels*), which is included under general insurance in Table 7.

The remaining components of general insurance are all included in FFA's property and liability insurance data under the categories of "auto" (*automobile*) and "property" (*biens et particuliers*). General insurance premiums are gross (as opposed to net) and earned (as opposed to written), meaning that they are net of premiums that were cancelled with no penalty.

Other forms of property and liability insurance listed in FFA's data have not been included in Table 7 as we have deemed them to be commercial in nature and unlikely to be issued to consumers.

Singapore

All of the data in Table 9 is from the Monetary Authority of Singapore (MAS). MAS categorises its general insurance data into two funds: "Singapore Insurance Fund" (SIF) and "Offshore Insurance Fund" (OIF). Data is not split by commercial and personal lines for either SIF or OIF and therefore we have selected property, motor, personal accident and health insurance as the types of insurance most likely to include personal insurance rather than commercial.

[24] The III's breakdown of life insurance also includes group and industrial policies. These have been excluded from the data in Table 2.

MAS does not provide data for motor insurance in OIF or property insurance in OIF prior to 2019. Data for general insurance in these years has been omitted in order to maintain consistency. General insurance premiums are gross (as opposed to net) and direct.

MAS provides data for linked and non-linked life insurance in SIF. The sums assured for linked life insurance are based on investment performance, whereas non-linked life insurance guarantees a fixed sum. Within the linked and non-linked life insurance, data from both annual and singular payments towards premiums is included in Table 9.

Annuities is the total premiums raised from new annuity issuances according to MAS data.

Hong Kong

All of the data in Table 11 is from the Hong Kong Insurance Authority (HKIA). HKIA provides data on the Hong Kong general insurance market but does not break it down by personal and commercial lines. Table 11 contains accident and health, motor vehicle and property damage insurance, as they are the most likely to be personal lines. Other data provided have not been included as they are more likely to be commercial in nature. General insurance premiums are gross (as opposed to net) and direct.

Life insurance is split by linked and non-linked insurance and includes new premiums written for whole life and term assurance.

Annuities includes new endowment premiums and new annuity premiums.

Appendix 2

An exploration of the consumer credit market

Consumer credit is the branch of consumer finance that provides credit or debt to an individual customer for the purchase of goods or services. The provision of consumer credit globally relies on three pillars: a) consumer behaviour; b) credit score data repositories; and c) credit risk models.

In the main body of this report, we focussed on the credit risk models which are used to help organisations decide whether to offer credit to a customer as well as the interest rate that should be charged. To develop the customer's risk profile, these models analyse a customer's past behaviour, credit data from credit check agencies, and their current financial standing.

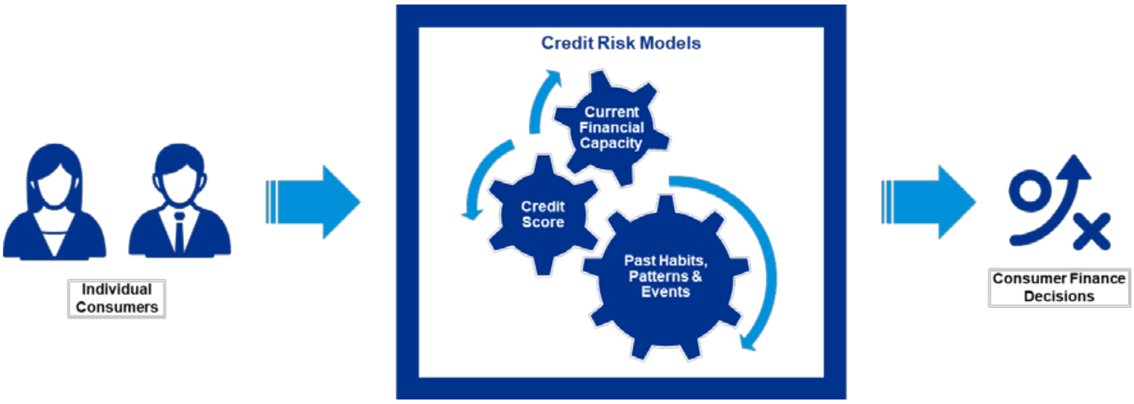

Traditionally, credit risk models used a limited number of variables to develop risk profiles for customers. More recently, advanced algorithms are being used to build up a more accurate risk profile for customers in real time. These algorithms allow for the processing of vast amounts of data which would not be possible for humans or traditional regression-based models to compute.

In the remainder of this appendix, however, we look in more detail at the use of algorithms that are designed to improve customer experience. We discuss two consumer finance products: mortgages and credit cards.

Mortgages

As noted above, algorithms are used in the provision of credit itself – i.e. to assess a customer's creditworthiness.

However, algorithms are increasingly being used to streamline mortgage application processes. Currently, these processes take between 20 and 40 days, with limited ability for customers to track cases in "real-time". This lending process also requires significant manual intervention including the collation and validation of the relevant documents.

By using AI, providers are simplifying these processes. AI-driven optical character recognition (OCR) and computer vision is increasingly being used to automate the capture, review, and validation of data during application processes. Machine learning is being used to streamline document classification and validation using external and internal data sources. Overall, this is reducing the time it takes to process applications, and in turn the friction experienced by customers throughout the mortgage process.

When using these techniques however, developers and users still need to be cognisant of potential algorithmic bias. Many of the potential contributors to bias set out in Section 3, for example the human biases that can be introduced during interpretation, remain relevant.

Credit and payment cards

As discussed in the main body of this report, advanced algorithms are increasingly used by credit card issuers to assess the risk of providing customers with credit.

When it comes to customer experience, the application of algorithms across credit cards can be broadly categorised into two stages: a) customer acquisition; and b) customer engagement and retention.

During the customer acquisition stage, card issuers are increasingly using deep learning capabilities to gain a holistic picture of their target audience. Multiple data sets including historical customer data, previous campaign data, device data and the activities of potential customers (e.g. browsing history) are analysed to inform the acquisitions strategies that are optimal. In short, the insights gathered in these approaches help card issuers to target customers at the right time, on the right device, that will most likely achieve the best response.

Algorithms are also used to engage and retain customers. As noted in Sections 2 and 3, card issuers are using algorithms to detect and prevent fraud by analysing a customer's purchasing behaviour, location and transaction value to generate an alert and/or block a transaction when an unusual transaction is detected – bolstering the customer relationship. Algorithms are also being used to offer increased personalisation to customers. This includes the use of AI based self-serve assistants, which can generate personalised financial advice, including basic support or expenses tracking, using natural language processing. Algorithms can also be used to offer personalised promotions to customers – for example, cashback offers for certain retailers.

www.kpmg.com\uk