Financial Conduct Authority

MS14/6.2: Annex 10

Market Study

Credit Card Market Study

Interim Report: Annex 10 - Account Level Data

November 2015

Introduction

1.1 This section describes the data we requested from credit card issuers and a credit reference agency; how we matched these datasets together; and our approach to standardising issuers' internally derived credit risk scores.

Data requests

1.2 To inform our analysis we requested account level information from credit card issuers and a credit reference agency. We worked with firms to design our data request in a way that avoided collecting any personally identifiable information or credit card numbers.

Account level data

1.3 We requested data from eleven issuers that represented the range of business models and types of issuers (monoline providers, banks) that operate in the credit card market.

1.4 For each issuer in our sample, we requested a range of information on all credit card accounts in their portfolio which were active at some point over a five year period (1 January 2010- 31 January 2015). In total, we collected information on 74mn accounts which covered around 80% of all credit card accounts active during our sample period.[1]

1.5 The data requested covered consumer details at the point of origination, monthly account activity and product features. These are described below.

Information collected from consumers at the point of origination (Table 1)

1.6 Here we requested information issuers collect from consumers at the application stage. This included information on the initial terms of their credit card agreement (such as APR and credit limit) and consumer specific information (such as their application score and demographic information).

Statement cycle account level data (Table 2)

1.7 The requested account level data was provided on the basis of the statement cycle[2] for all accounts held over the period 1 January 2010 – 31 January 2015. Some credit card accounts were opened after 1 January 2010. For these accounts, information was reported from the date the account was opened up to 31 January 2015 or the closing date (whichever occurred first). For accounts which were closed in this period, information was reported up to the account closing date.

1.8 We requested information for each account over each statement cycle, such as opening and closing balances, number and value of transactions, fees and interest applied to the account, repayments made, internal risk scores assigned to the account, costs of servicing the account (split by overhead costs, funding costs and operational costs) and arrears data.

[1] CRA data indicated that there were 95mn accounts open over our sample period. Of these, we received account level data for 74mn accounts, which equates to around 80% of the market.

[2] The statement cycle is the timespan over which the information presented in each statement relates.

1.9 For some issuers, costs reported at the account level were not readily available. Where possible, issuers employed a methodology to allocate costs. As such, analysis using these cost variables was done on a firm by firm basis, and took into account the various methodologies and assumptions employed by issuers in providing this information.

Product information (Table 3)

1.10 Here we requested a monthly snapshot of product details and features between 1 January 2010 and 31 January 2015. This allowed us to observe how product features evolved over time.

1.11 We worked closely with issuers to ensure account level data submissions were as complete as possible. However, there were some instances where information could not be provided at all or within the timescales. Often this was the result of issuers not holding the requested information in an easily attainable format or no longer holding the information.

Credit Reference Agency Data

1.12 We requested, from a CRA, an anonymised list of all credit card holders that had an open account at some point over the period 1 January 2010 to 31 January 2015. For each credit card holder, we requested a list of all credit card accounts that they had open for at least part of this sample period, as well as the account opening and closing dates. We used this list of consumers to link credit card accounts in our sample to an individual consumer.

1.13 This data indicated that the whole credit card market includes 38mn consumers that had at least one credit card at some point over our sample period (34mn of these consumers were captured in our account level data request) and 95mn accounts were open at some point over the same period (74mn were captured by our account level data request).[3]

1.14 In this request, we also asked for the CRA credit risk score assigned to consumers at a specific point in time (25 May 2015). This score gives a prediction of the probability of delinquency in the next 12 months.

[3] We understand that CRA data includes accounts sold to debt collection agencies, as these accounts continue to be reported to CRAs.

Data cleaning process

1.15 On receiving the data, we carried out a number of quality checks. This involved on-going engagement with the credit card issuers in our sample. The key areas of checks we completed are outlined below.

On-going queries and engagement with firms

Reconciliation checks on the account level data

The purpose of this was to check our understanding of the data submitted by issuers. Here we reconciled the statement cycle data submitted by issuers, to monthly aggregated data.

Standardisation of variables across issuers

The purpose of this exercise was to ensure that variables submitted by issuers were labelled and formatted consistently.

Frequency checks on all submitted variables

The purpose of this was to identify any significant omissions or variables that exceeded our expected range

Summary data derived from account level data submissions sent to issuers for review

The purpose of this exercise was to check that our summary of issuers' data resonated with their credit card portfolio(s).

Matching rates between credit reference agency and account level data

Here we checked the matching rates between credit reference agency data and account level data, for each issuer.

Approach to standardising issuers' internal credit risk score

1.16 As part of our data request to issuers, we requested they provide us with the internal credit risk scores they use to inform their decision making for each account and each statement cycle. Below is a description of how issuers use credit risk scores and our approach to standardising them so they are broadly comparable across issuers, as well as over time and across products for a particular issuer.

Credit risk scores

1.17 Credit risk scorecards (risk scores) are an important tool for credit card issuers to assess the creditworthiness (probability of default) of consumers at the point of acquisition as well as over the life of a credit card account. Risk scores are used to rank individual consumers in terms of creditworthiness and are estimated from issuers' risk models that take into account various attributes of an individual (for example, their debt payment history, age of oldest account). Issuers use risk scores to inform their lending decisions as well as the terms they offer. This allows them to manage their exposure to credit risk.

Rationale for standardising issuers' credit risk scores

1.18 For our analysis, we needed the measure of credit risk that was used by issuers to inform their lending decisions. Also, in some cases, we needed the risk score at different points in time. This was the case when assessing the performance of accounts (or consumer outcomes) by different risk profiles (i.e. the risk assigned to the account at the beginning of the period of analysis). As such, we asked issuers to provide internal credit risk scores at the account level for each statement cycle.

1.19 In some instances we used the CRA risk scores and sense checked the results using issuers' internal risk scores. As the CRA risk score was reported on 25 May 2015 (i.e. after the end of our sample period), it reflected the performance of accounts at that point of time (i.e. it is an ex-post measure of the credit risk of consumers in our sample).

1.20 Issuers use different methodologies to assess the creditworthiness of consumers which leads to different credit risk score variables being used by each issuer. These different risk scores are also recorded differently (e.g. using different scales), such that even if two firms judged the same consumer to be equally risky they may assign different credit score numbers.

1.21 Issuers also have different risk appetites and target markets, which means they use different risk score thresholds to determine whether they will accept a credit card application (for example, for some products issuers may only accept applications from lower risk consumers, for others they may accept mid-range risk consumers).

1.22 In addition, issuers evaluate the efficacy of their risk scores over time and adopt new, or modify existing, scores to cater for changing market needs or to improve their predictive capacity. They may also use different risk scores for separate products.

1.23 The implications of these factors are that risk scores are not comparable between issuers and in some cases are not comparable across time or the products of an individual issuer. As such, it was necessary to standardise them. Our approach to standardising risk scores is set out below, first on a step-by-step basis and then with a more detailed description of each step.

Approach to standardising issuers' internal credit risk scores

1.24 In standardising issuers' internal credit risk scores, we followed the steps summarised below:

- **Step One: Identifying credit risk scorecard changes**

  Here we identified the various risk scores used by an individual issuer. For example, as mentioned earlier issuers may change their risk scores over time or use different risk scores for different product ranges. These needed to be identified and standardised.

- **Step Two: Constructed a proxy for the probability of default**

  The next step was to construct a proxy for the probability of default, which we refer to as the 'bad rate'.

- **Step Three: Defined fifteen risk categories**

  For each issuer and using a single month as a starting point, we identified the range of internal credit risk scores that corresponds to a particular 'bad rate' (as an illustrative example, a 'bad rate' of 5% may correspond to those accounts falling within the range of risk scores 500-800 for Issuer A, but 1200-1600 for Issuer B), and then created fifteen risk categories from this (from low risk (1) to high risk (15)).

- **Step Four: Replicating for additional months and issuers**

  For each issuer we replicated step three for additional months, refining our risk category thresholds, where necessary. Where an issuer used multiple risk scorecards (for example, changed their scorecards over time or for different product ranges), these were also standardised using this approach.

- **Step Five: Assign the relevant risk category to each account**

  Finally, we assigned to each account in our sample one of the fifteen risk categories for each statement cycle.

1.25 The below table provides an illustrative example of the outcome of this standardisation process. For instance, if we were to define 'Risk Category One' as having a 'bad rate' of between 0 to 1%, this corresponds to accounts with internal risk scores of between 800 and 850 for Issuer A. In the case of Issuer B, it relates to accounts with the underlying risk scores of between 1200 and 1400. The effect of the standardisation process is to treat both these sets of accounts as equally risky.

Table 1: Illustration of outcome of the standardisation process

'bad rate'	**Issuer A** (risk score range)	**Issuer B** (risk score range)	**Risk category**
0-1%	800-850	1200-1400	1
1-3%	700-800	1000-1200	2
3-6%	600-700	800-1000	3
6-10%	450-600	600-800	4

Step One: Identifying credit risk scorecard changes

1.26 To identify the various risk scorecards used by an issuer, we ran descriptive statistics on their internal risk scores for each month to determine if their risk score methodology changed and to determine how long a particular risk score range was applicable. We did this for each of the issuer's different products, to check whether the same risk score scale was being used.

Step Two: Constructed a proxy for the probability of default

1.27 We then constructed a measure of the probability of default (what we refer to as the 'bad rate') that is consistently reported across issuers and which (by definition) is closely correlated with issuers' internal risk scores. A particular 'bad rate', if defined uniformly across issuers, will always represent the same level of risk across issuers, time and products. We calculate and define the 'bad rate', using the available data.

1.28 We considered a number of measures for the 'bad rate', from the percentage of accounts issuers reported as a number of payments in arrears to the percentage that issuers reported as charged-off[4]. We shortlisted those measures that we found to be defined, and reported, consistently across issuers.

1.29 We found three cycles over-due was the most consistently reported variable, and a good proxy for the probability of default. As such, we defined the 'bad rate'[5] as the percentage of accounts that became three cycles over-due over a 12 month period.

Step Three: Defining fifteen risk categories

1.30 To translate the continuous risk scores for each firm into discrete categories, we needed to decide how many discrete categories to use. Defining categories too narrowly (i.e. defining many risk categories) presents the risk that consumers with similar risk profiles fall within different risk categories. In addition, it makes the standardisation very challenging. If defined too broadly (fewer risk categories), there is a risk that consumers with very different risk profiles are captured within the same risk category.

1.31 In order to make our analysis manageable, fifteen categories of credit risk were chosen as a starting point (with one being the lowest risk category, fifteen being the highest risk category). This number was considered sufficiently granular to meaningfully inform our analysis.

1.32 To define the risk categories, we started by plotting the 'bad rate' against the distribution of risk scores, for an individual issuer in a given month. See figure one below for an illustrative example. This example demonstrates how the 'bad rate' and distribution of risk scores are related. Here 'Risk Category Seven' is defined as a 'bad rate' of 7.5% to 9.5% and includes 4% of accounts (for that issuer in a given month) which is the 28th to 32nd percentile of accounts. The corresponding issuers' risk score can then be specified.

[4] For example, charged-off accounts identified by issuers were not consistently reported. Issuers charge-off accounts at different points in time, depending on their internal policies.

[5] The 'bad rate' is defined as: number of accounts that became three cycles over-due over a 12 month period divided by the number of accounts that were less than three cycles over-due at the beginning of the 12 month period.

Figure One: Hypothetical illustration of our risk standardisation

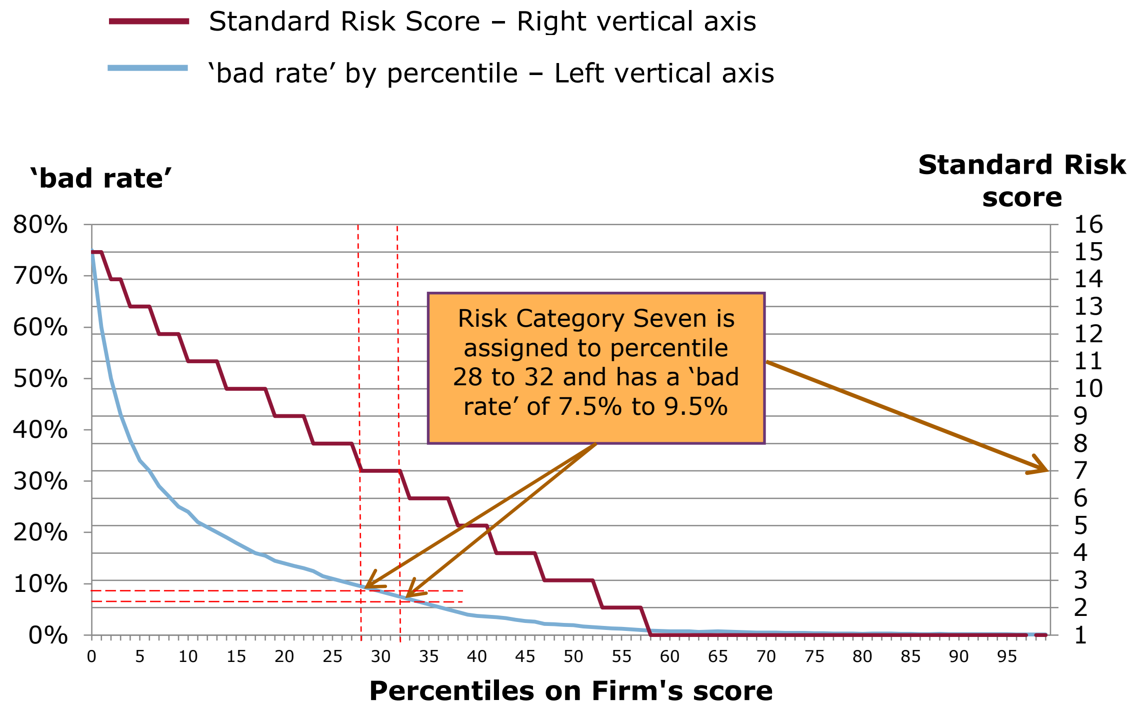

1.33 In terms of process, we started with a given month, and ordered accounts from most to least risky based on an individual issuer's risk score. The blue line in figure one above shows the 'average 'bad rate'' for each percentile of accounts. For example, the first percentile shows that 1% of accounts with the highest risk score have an average 'bad rate' (as calculated in the next 12 months) of around 70% (i.e. on average, 70% of accounts in this percentile fall into more than three months arrears in the next 12 months). As expected, for lower risk accounts, the average 'bad rate' for those accounts is much lower.

1.34 The next step was to translate the distribution of 'bad rates' (i.e. blue line) into the fifteen discrete risk categories (i.e. the red line). To do this, we initially defined a 'bad rate' range for each of the fifteen risk categories using empirical judgement. In doing this, we aimed to capture consumers with a similar risk profile in the same risk category. We also wanted to ensure that the different risk categories captured distinctly different consumers in terms of credit risk.

Step four: Replicating for additional months and issuers

1.35 Once we defined initial 'bad rate' ranges for each risk category, for an individual issuer and for a specific month, we then replicated this process for a second month, ensuring that for an individual risk category the average 'bad rate' was within the 'bad rate' range specified for that category.

1.36 To illustrate, if for month one we defined a 'bad rate' of 0 to 1% for 'Risk Category One' and find that the average 'bad rate' of accounts assigned to that category falls within this 'bad rate' range, we did not amend the 'bad rate' range. However, if when we replicated this for the second month we found the average 'bad rate' lies outside of the range 0 to 1%, we would amend the 'bad rate' range accordingly.

1.37 We followed this process for additional months, and for other issuers, amending the 'bad rate' ranges for each risk category until we had completed the process.

Step five: Assign the relevant risk category to each account

1.38 Once we had defined the 'bad rate' ranges and corresponding internal risk scores thresholds, the risk categories were assigned to each account, for each statement cycle. We made sure to correctly apply the different threshold of each type of risk score.

Sensitivity checks and caveats

1.39 We carried out a number of checks and sensitivities on the risk categories. For example, we checked how closely the risk score categories were correlated to internal risk scores in any particular month. The correlation across issuers ranged from 80% to 97%.

1.40 We recognise that standardising risk scores is a complex process and we have therefore interpreted the results with caution. In particular we were mindful that:

- The fifteen standardised risk categories do not provide an exact measure of the probability of default, but rather rank accounts into particular ranges (with consumers that fall into credit risk category one, being less likely on average to default than those in category two).
- In some cases issuers were not able to provide an internal credit risk score for each account and for each statement cycle. In these cases, we removed these observations from our standardisation process. There is a risk that this may have biased our sample. However, given the small proportion of observations we removed, we do not expect that this had a material effect on our results.
- The reliability of the risk categories rely on the accuracy of the internal credit risk scores provided to us by issuers.

Financial Conduct Authority

25 The North Colonnade Canary Wharf
London E14 5HS
Telephone: +44 (0)20 7066 1000
Website: www.fca.org.uk