lucerna partners

Reputational incentives for authorised push payment scams

Publishing performance metrics by bank

August 2021

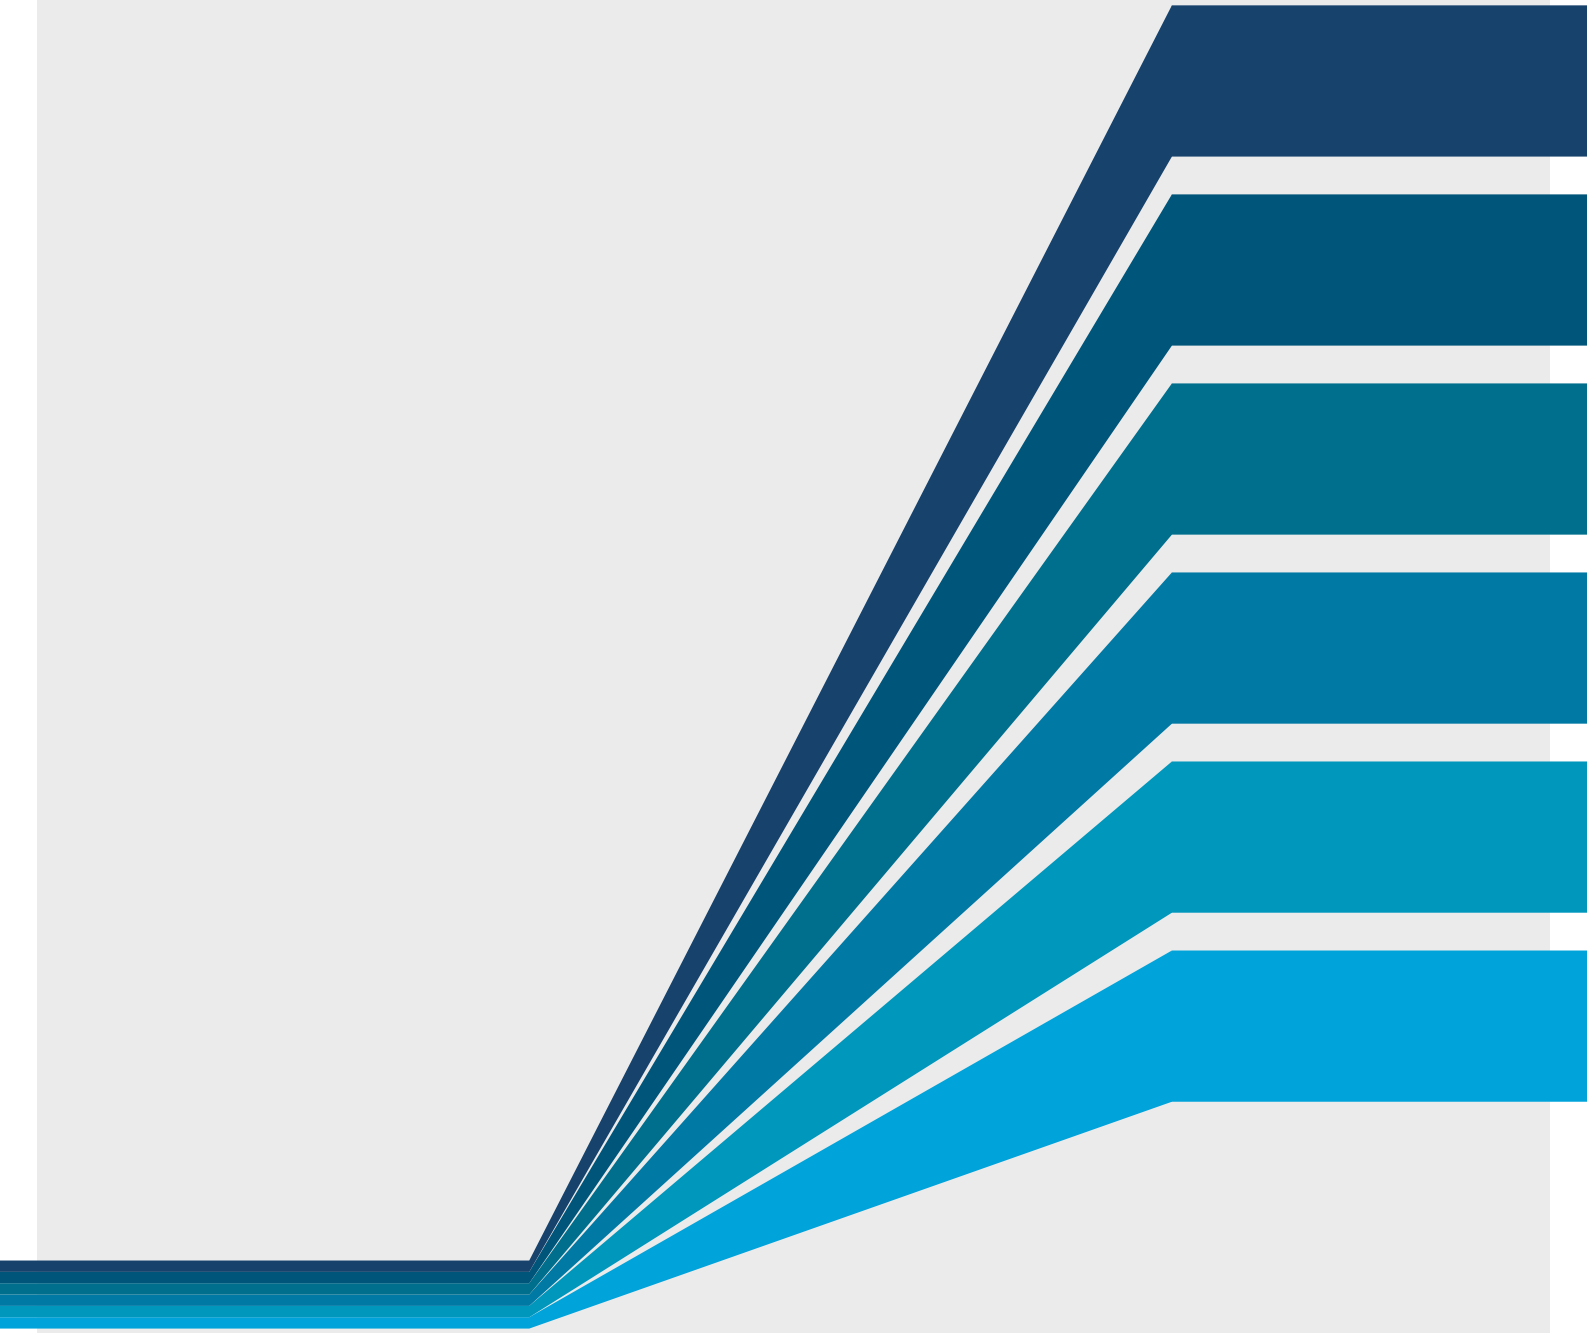

Contents

Executive Summary

The PSR asked Lucerna to carry out independent research to identify and recommend metrics around the publication of APP scam data by Payment Services Provider (PSP). We were tasked with assessing to what extent each identified metric might incentivise PSPs to do more to prevent APP scams taking place and to protect consumers when they did fall victim to scams.

The PSR asked us to consider the pros and cons such as potential benefits, costs, and unintended consequences, but did not ask us to recommend whether the PSR should go ahead and require the publication of the metrics.

Our starting assumption was that an effective reputational incentive would focus on telling consumers, potentially via consumer bodies and the media, what was most relevant to them:

- How likely is my bank to give me my money back if I am a victim?
- How likely is my bank to protect me from APP scams?
- How much does a bank help the scammers?

We identified and tested – through a programme of interviews with banks, industry stakeholders and consumer representative bodies – the following headline metrics:

A: The proportion of APP scammed customers who are left - fully or partially - out of pocket:

- by volume: total APP scam cases[1] where the cost is fully or in part borne by the victim, as a percentage of all the sending bank's APP scam cases; and
- by value: total value of APP scam losses borne by victims, as a percentage of sending bank's total APP scam value.

[1] On a 'cases closed' basis.

B: The bank's APP scam rate as a sending bank

- by volume: total number of APP scam payments as a percentage of total number of push payments by consumers; and
- by value: total value of APP scams, as a percentage of total value of push payments by consumers.

C: The bank's APP scam rate as a receiving bank, taking into account repatriation

- total value of APP scam payments received minus the value repatriated, as a percentage of total value of push payments received from consumer accounts.

Publishing sending bank data by PSP (Metric A and Metric B)

We consider that the relative likelihood that a scammed consumer is left out of pocket (Metric A), and the relative likelihood of falling victim to a scam (Metric B), is highly relevant information for consumers. Publishing data on these metrics is likely to attract comment from consumer groups and the media and therefore is very likely to have a reputational impact on banks.

We think it is very unlikely that publishing Metric A would, because of a moral hazard, skew consumers' behaviour adversely. It is conceivable, if consumers were very confident they would get their money back, they might take a little less care over small transactions.

Given current levels of full reimbursement, we think there is very little risk that publishing this data would result in consumers being very confident that they would not be left out of pocket if they fall victim to a scam.

We heard a range of views on the degree to which publishing sending bank scam rates, in Metric B, risked providing help to scammers by potentially identifying banks with weaknesses on prevention although most stakeholders did not consider publishing overall scam rates would give material aid to scammers.

We tend to agree with the arguments we heard that that if a bank's consumers are disproportionately becoming victims of scams then that bank is already being targeted by scammers. We also heard that it should not typically take sending banks a long time to rectify the sorts of weaknesses in their controls that might cause APP scam rates to increase. There is likely to be significant lead time before the first publication and this would allow banks to address any existing vulnerabilities – they have already been on notice for some time that the regulator is seriously considering publication.

We cannot say that publishing Metric B data would be entirely without risk. But we have not seen convincing evidence that any risks are likely to outweigh the benefits of publishing. Formal consultation would provide a further opportunity to explore evidence for any material risk.
We think publication of Metric B data alongside Metric A is likely to give a fuller, and potentially a fairer, picture than publishing Metric A alone.

Publishing receiving bank[2] data (Metric C)

Many, but not all, industry stakeholders thought that consumers would not be interested in receiving bank data. All of the consumer bodies we spoke to disagreed, and argued that consumers care about banks funding scammers, and information about which banks send the most funds to scammers could affect consumers' choice of bank – since consumers do, of course, also bank with some banks that are receiving banks.

Consumer groups told us this would be important information. We believe they are likely to highlight it, so it could affect banks' reputations generally, and this is likely to contribute to the choices consumers make.

We think publishing receiving bank APP scam rates by total value (taking into account the funds that receiving banks intercept and repatriate) would provide for an effective reputational incentive for banks (and for other receiving entities) to tackle scams in their role as a receiving bank.

[2] We use the term 'receiving bank' to cover all receiving entities including PayPal and so on.

We think publishing the data net of repatriation best targets the incentives on the outcome of minimising customer losses from scams.

Most industry stakeholders told us that publishing receiving bank data is more likely to help scammers than publishing sending bank data. At the same time, we also heard that scammers are already very aware which banks are easiest to exploit in terms of receiving payments.

Some stakeholders told us that it could take considerably longer for receiving banks to rectify any weaknesses in their systems and processes, than to address any vulnerabilities in sending bank protections.

Receiving banks would likely have many months from detecting a new vulnerability shown in their data before that data was reported and then published. If banks were able to demonstrate to the PSR that this was not long enough to mitigate a substantial risk of helping scammers, the PSR could consider collecting Metric C data alongside the other metrics but publishing it with a delay.

The delay should be sufficient to give receiving banks time to address any identified vulnerabilities.

This may not be necessary, but it is an option that the PSR may wish to consider and allow individual banks the opportunity to make specific confidential submissions to the PSR about their own circumstances.

Recommendations

Should the PSR be minded to require publication we recommend Metrics A, B and C as appropriate candidates, with Metric C published net of repatriation and Metric A data (by volume and value) published by scam value band as well as for all scams as a whole.

In the detail of our report, we assess Metrics A, B and C – and other suggested metrics and variations – against a framework of: incentive effect; practicality; potential for unintended consequences; and fairness. The table below gives a summary of our findings.

Findings on headline Metrics A, B and C

	Incentive effect on banks	Practicality	Potential for unintended consequences	Fairness of comparisons	Recommendation
A: The proportion of APP scammed customers who are left - fully or partially - out of pocket	High	High	Low	High Assuming B also published	Publish by volume and value
B: The bank's APP scam rate as a sending bank	High	High	Low But some uncertainty	High	Publish by volume and value
C: The bank's APP scam rate as a receiving bank, taking into account repatriation	Medium Consider delay between reporting and publication	High	Low With time to address major vulnerabilities before publication (if necessary)	High	Publish by value - consider delay between reporting and publication

Thank you

A number of banks, industry stakeholders and consumer bodies were extremely generous in helping us with our report. They gave us a considerable amount of their time and shared their expertise and knowledge. We are grateful for this assistance.

Section 1: Introduction

Lucerna was asked to consider the pros and cons of publishing APP scam data as a reputational incentive on banks.

What we were asked to do

In February 2021 the PSR published a call for views outlining proposed measures to help prevent Authorised Push Payment (APP) scams and protect consumers. One proposal was to publish APP scam data by Payment Services Provider (PSP) – referred to as Measure One.

At present, no firm-level information is publicly available on APP scam losses, reimbursement and repatriation levels. The PSR considered that if such information were made publicly available, it may give PSPs stronger incentives to do more to prevent APP scams taking place and to protect consumers when they did fall victim to scams. The PSR asked Lucerna to produce an independent report to identify and recommend metrics that PSPs could be required to publish.

The report would consider the pros and cons of metrics such as potential benefits, costs, and unintended consequences.

Our understanding of Measure One

We understand the high-level objectives for APP scams policy is to reduce APP scam losses to consumers by: preventing scams; repatriating defrauded funds; and appropriately reimbursing[3] consumers who fall victim to scams.

Effective incentives across the system will help to drive both preventative action and protection of consumers. We have been asked to consider one type of incentive – a reputational incentive on banks[4] based on publication of comparative APP scam metrics by individual bank.

[3] In this report 'reimbursement' includes reimbursement from both funds repatriated by receiving banks and funds the sending bank provided, unless it is clear we are making a distinction between the two.

[4] The term 'banks' is used to refer to all PSPs participating in the relevant payment systems over which authorised push payments are sent by consumers.

Consumers might access the comparative information directly, or consumers and others might be become aware of banks' relative performance via consumer bodies' campaigning and media coverage. The reputational incentive could work by:

- directly informing consumers' choice of bank, through providing information on the relative likelihood of being scammed and reimbursed by different banks; and
- affecting a bank's reputation for addressing scams and protecting consumers, with indirect implications for consumers' choices, and for banks' regulatory, political and market risk.

The reputational incentive, if effective, could encourage banks to both prevent APP scams and ensure consumers who fell victim to scams were not left out of pocket.

To the extent the cost to sending banks of reimbursing victims increases, this could further incentivise them to prevent scams, including by taking steps to encourage receiving banks, third party facilitators and consumers themselves to play their part.

The purpose of publishing comparative information by individual bank is not to directly educate consumers on how to better protect themselves nor to incentivise actions by third parties, important though both of these are.

Our approach

We identified the main types of metric suggested in, and in response to, the PSR's call for views. We also identified other key choices to be made in implementing Measure One, such as the range of banks to be covered and who should publish the data. We developed criteria to assess each potential metric and set of metrics.

To gather evidence for our assessment, we reviewed key responses to the PSR's call for views and interviewed a selection of banks including most Contingent Reimbursement Model (CRM) Code banks[5], consumer bodies and other key stakeholders. See the Appendix for a list of interviewees.

Our recommendations are as robust as they can be given this was a short project, but a number of issues should be followed-up in formal consultation on any PSR proposal to take forward Measure One and we set these out in our report.

[5]The CRM Code is a voluntary code that sets out consumer protection standards on banks to help reduce the number of APP scams, and specifies broadly the approaches signatories should take in determining whether consumers are reimbursed for scam losses.

Section 2: Potential metrics

We designed and tested three headline metrics for publication as well as a number of other potential variations.

Our starting assumption was that an effective reputational incentive would focus on telling consumers - potentially via consumer bodies and the media - what was most relevant to them.

Our assumption, supported by interviews with consumer-focused bodies and some but not all banks, was that consumers would want to know:

- How likely is my bank to protect me from APP scams?
- How likely is my bank to give me my money back if I am a victim?
- How much does a bank help the scammers?

We noted that many stakeholders wanted the set of metrics to provide a balanced view covering not just reimbursement levels to consumers but also banks' performance in preventing scams.

Based on this, we identified – and tested with interviewees - the following three potential headline metrics for publication:

A: The proportion of APP scammed customers who are left - fully or partially - out of pocket:

- by volume: total APP scam cases[6] where the cost is fully or in part borne by the victim, as a percentage of all the sending bank's APP scam cases; and
- by value: total value of APP scam losses borne by victims, as a percentage of sending bank's total APP scam value.

B: The bank's APP scam rate as a sending bank

- by volume: total number of APP scam payments as a percentage of total number of push payments by consumers; and

[6]On a 'cases closed' basis.

- by value: total value of APP scams[7], as a percentage of total value of push payments by consumers.

C: The bank's APP scam rate as a receiving bank, taking into account repatriation

- total value of APP scam payments received minus the value repatriated, as a percentage of total value of push payments received from consumer accounts.

We assumed metrics should be expressed as a relevant percentage, rather than an absolute number, because the aim is to compare performance easily between different sized banks.

We also assumed that all the relevant the data would be on a 'cases closed' basis, i.e. based on scams which had been investigated and confirmed, and any reimbursement decided.

More granular versions of the metrics

A number of responses to the call for views argued that data published under Measure One ought to be split by different scam types. That is, purchase scams, investment scams, and so on, because different scam types have quite different characteristics. So we explored proposals to split Measure One data by scam type and by scam value.

Metrics on reimbursement decisions

Under the CRM Code, a bank's decision on whether or not to reimburse a consumer who has become a victim may be determined by a number of factors, including:

- whether the sending and/or receiving bank has been deemed as complying with the Code's protection standards;
- whether a victim has been deemed vulnerable (to scams); and
- whether a victim has been deemed 'at fault', under a number of criteria, including whether they have been deemed to have ignored an 'effective warning' or a warning the payee name did not match under Confirmation of Payee (CoP), or to have acted unreasonably, grossly negligently or fraudulently.

Banks, whether members of the CRM Code or not, may also apply other policies that determine reimbursement rates, such as automatic reimbursement for certain types of scam losses.

Publishing comparative data on banks' reasons for reimbursement and non-reimbursement could shed light on what lay behind the differences between banks' reimbursement rates. It could, for example, identify where a high reimbursement rate resulted from poor compliance with CRM Code protection standards, or where a bank was more likely to find consumers at fault in particular ways.

[7]Before any repatriation of funds lost.

We explored two potential reasons metrics:

- proportions of decisions to reimburse resulting from, for example, sending bank fault, receiving bank blame, consumer vulnerability, no blame, shared blame with the consumer, automatic reimbursement; and
- proportions of decisions to deem the consumer at fault resulting from, for example, acting unreasonably, ignoring an effective warning or CoP result, gross negligence or acting fraudulently.

Metrics on scam prevention activity

Responses to the call for views suggested a number of metrics that might shed light on specific activities that sending and receiving banks take to try to prevent scams so we explored these kinds of metrics. For example, the numbers of transactions discontinued as a result of warnings or the number of mule account closed by receiving banks.

Metric on facilitation of APP scams by others

Many APP scams are facilitated by scammers' exploitation of vulnerabilities in systems other than the banking system, such as scams advertised on social media and promoted by search engines, imitation telephone calls and scam texts.

Many responses to the call for views suggested the need to include in Measure One a metric on the facilitation of APP scams by others and we explored the case for a metric of this kind.

Metric on complaints

A number of responses to the call for views suggested that Measure One should include a metric on complaints about banks' handling of APP scams, including the number of complaints to the Financial Ombudsman Service and the ombudsman's uphold rate against each bank and so we included this potential metric in our considerations.

Implementation

While not central to what we were asked to do by the PSR, we have also explored a number of other key issues for the implementation of Measure One, including:

- Which banks should the Measure apply to?
- Who should publish the data?
- How should the quality of the data be assured?;
- How often should the data be published? and
- Trialling and testing of the metrics.

Section 3: Interview findings

We tested the potential metrics and variations with a range of key stakeholders.

We asked key stakeholders for their views on how each metric performed against the following characteristics:

Incentive effect - contribution to the reputational incentive on banks, including ease of use and likely impact with consumers;
Practicality - including the feasibility of ensuring adequate quality data;
Potential for unintended consequences - including potential to help scammers, or to adversely skew banks' or consumers' behaviour;
Fairness of comparisons between banks - and therefore effectiveness of overall comparisons between banks.

A number of banking industry stakeholders said that reputational incentives were not needed, with some asking for evidence that published information would drive consumers' choice of bank, and others insisting that protecting their customers is all the incentive they need.

Many said that they disagreed with the proposal to publish information aimed at creating a reputational incentive, and thought it would be a distraction from concentrating on more important measures to prevent scams and educate consumers on the need to protect themselves.

Some banks, however, held the view that reputational incentives would be effective (to varying degrees), and said that league tables do tend get the attention of Boards, and this does drive behaviour and the allocation of resources within banks.

All consumer-focused stakeholders we spoke to were in favour of publishing data, even where they thought moving on to mandating reimbursement for all consumers would be an even more effective way to protect consumers.

A consumer body held the view that, given the state of switching in banking, information on APP scams is perhaps not likely on its own to drive switching.

But it considered that performance on APP scams was important as part of the wider picture of a bank's reputation and over time this could have an impact on consumer choice and switching.

METRIC A: THE PROPORTION OF APP SCAMMED CUSTOMERS WHO ARE LEFT - FULLY OR PARTIALLY - OUT OF POCKET

Incentive effect
A number of stakeholders said the reimbursement rate alone was not enough for the consumer to make an informed choice and a sending bank's APP scam rate was important for context.

One consumer group told us that publishing data on how many people get all their money vs only part of their money back, would be important because they suspected some banks were 'negotiating' partial settlements at the expense of consumers. One bank agreed this was happening across the industry.

> *"In our member's experience, banks are likely to be more resistant to reimbursing high value frauds and to only reimburse part of the amount lost by the customer."*
>
> **Fraud Advisory Panel.**

One bank asked whether, beyond providing a reputational incentive, the PSR intended to use the data to drive particular outcomes, for example CoP phase 2 or better consumer recovery times.

A Code bank preferred expression of metrics by both volume and value, in order to tell the right story because of the range of different scam types and values.

Another Code bank also noted that expressing metrics by value and by volume can give different information.

> *"You do have to be careful; the data can be very misleading - [approximately] 80% of PP scam value will be caused by 15% of volume."*
>
> **A bank.**

Unintended consequences: Skewing consumer behaviour
We heard many arguments from banking industry stakeholders that publication of reimbursement rates could lead to a moral hazard for consumers, where a consumer took less care, particularly for low value payments.

Other stakeholders, including one bank, were deeply sceptical of these arguments, pointing out that being a victim of scams is embarrassing and hurtful in itself, and it also would be unusual for consumers to be completely confident of getting all of their money back.

> *"There are seat belts in cars but how often do you crash your car to test if they work?."*
>
> **A bank.**

METRIC B: THE BANK'S APP SCAM RATE AS A SENDING BANK

Fairness
One bank argued that information on reimbursement rates could not be published without also publishing information on how likely a consumer was to be scammed – because a bank with a low APP scam rate could have a high non-reimbursement rate compared to other banks because it took greater steps to prevent customers making payments to scammers (the implication being that those consumers who did were therefore at fault).

Another bank made the point that an institution with weak controls would be subsidising organised crime even if it was reimbursing consumers.

A consumer body said the answer to the argument that reimbursement comparisons might not be fair, because one bank might be better at prevention than others, was to publish more data. The solution is not to hide the headline numbers.

Some banks said that differences between the consumer bases of banks, differences in particular between small and large banks, could drive differences in scam rates.

Practicality
Two industry stakeholders said that expressing Metric B as the number of scam payments (not cases) as a percentage of total consumer payments was the most appropriate approach.

We heard that scams are recorded at the date of case closure, which is not the same as the date the scam happened or payments were made. We asked about whether this would materially change the data and were told that the majority of cases are resolved within a few weeks though this does vary by bank.

But in any case it was not likely to be material over a period of months, and it would be extremely difficult to record the outcome of a scam complaint case back dated to when a scam payment was made or when the scam happened.

One bank counselled caution in publishing the total value of scams because such information could be market sensitive. It said that the issue did not arise for volume scam data and the problem with value might perhaps be eased by publishing a scam rate as a percentage of total consumer payments value.

Unintended consequences: Helping scammers
Some banks argued that publishing how likely a customer of a bank was to get scammed would help scammers. For example, a scammer might sort data obtained by bank and target the customers of the bank that appeared to have the weakest controls first.

Another bank when presented with the example of scammers sorting data for the most profitable prospects said that could happen but scammers would very quickly make their way to the rest of the list in any event.

Another bank said that while information which revealed a bank had a control weakness could be exploited, whether this was a real problem would depend on the timescale of publication. That is, whether a bank had enough time to identify and fix any weaknesses before the information became public. But one bank said that it can take a long time to put in place defences so it is not necessarily the case that a bank can defend themselves before the data is published. Changes might take six, nine, twelve months or more.

Some banks did not agree with the arguments that revealing a high scam rate made a bank more of a target.

Which? noted that it published information about the security of banks, but it had not seen evidence from banks that it provided a guide to scammers or caused increases in scams.

> *"We work with experts to find weaknesses in banks' security and publish detailed reports on this - the result is the banks tighten up on their security, and none have ever complained they are targeted by fraudsters as a result of our work."*
>
> **Which?**

METRIC C: THE BANK'S APP SCAM RATE AS A RECEIVING BANK, TAKING INTO ACCOUNT REPATRIATION

Incentive effect
A number of industry stakeholders said the PSR should collect information on receiving banks APP scam rate and mule accounts but not publish it. One of them noted they were seeing big differences between banks in terms of which receiving banks were hosting scammers.

One bank questioned the benefit of publishing receiving bank data, arguing that if there were no mule accounts, scammers would just operate the scam so the money went to pre-paid cards or crypto merchants, so targeting the receiving banks may not reduce scams. Another bank said that the information would be of no use to consumers.

One bank pointed out that not all those receiving money from scams are banks, there are other entities.

All consumer-focused stakeholders thought publishing information on receiving banks' rate of APP scams would be important.

A number of consumer-focussed stakeholders recognised that while non-repatriated scam loss was less important in directly informing consumer choice of bank than whether or not a victim would get their money back, consumers do care about defeating scammers and making ethical choices.

Information on receiving banks' performance could indicate to consumers which banks were "the fraudsters' friends" and paid scammers the most money. One noted that You and Yours on Radio 4 frequently covered scams and there were often remarks by consumers about the receiving banks.

> *"If there is one bank full of mule accounts we absolutely need to know that! And yes, consumers would be interested."*
>
> *"A bank is a receiving bank one day and a sending bank the next - it's all one bank."*
>
> **Stakeholder with a consumer focus**

One bank emphasised the importance of publishing information on both reimbursement and repatriation rates. It said a higher repatriation rate reflects on the effectiveness of the receiving bank and not the sending bank so the data shouldn't give credit to a sending bank when a consumer isn't out of pocket if the reason for that is the effectiveness of the receiving bank stopping the money reaching a scammer.

Another bank agreed that since repatriation depends on the receiving bank, data on this should be published by receiving bank separately from sending bank data. Other banks stressed the importance of incentivising receiving banks to improve their performance.

Unintended consequences: Helping scammers
Many industry stakeholders warned against publishing information on receiving banks scam rates or mule accounts, which could tell scammers which banks were the most vulnerable in terms of them accessing bank accounts to receive money.

We asked why publishing receiving bank data was more sensitive to this form of unintended consequence than publishing sending bank data.

Most banking industry stakeholders considered publishing receiving bank data was more risky - we should stress that in many cases without agreeing publishing sending bank data was a good idea - and thought it would be more complicated and time consuming for a receiving bank to close weaknesses. Reasons given included that a bank's exposure might depend on:

- the profile of its customer base (we heard scammers might attempt to buy unused UK bank accounts from people now living in other countries);
- a bank's decisions on how accessible to make bank accounts to those who might struggle to provide a permanent address or certain forms of identification, for example, homeless people or people living in very temporary accommodation; or
- the growing numbers of unwitting mules as scammers respond to banks stepping up controls - inbound receiving bank payment profiling can be considered a new frontier in detection and prevention.

In response to the question of whether the scammers already know about banks' weaknesses, one bank said that whilst it was true that some scammers would know of weaknesses, publishing the information would mean all scammers would know.

MORE GRANULAR VERSIONS OF THE METRICS
A number of banking industry stakeholders argued for data to be split by scam type because the values and volumes of scams differed hugely between scam types and so the proposed metrics ran the risk of consumers' not understanding the origin and nature of different scams.

We heard support from some in the banking industry for splitting reimbursement rate specifically by scam value, for example high vs low value scams.

This was because many banks pay-out on low value scams but it is very important to see the approach of these banks when the scam involved life-changing amounts.

One bank argued against splitting the data by value band because it made a judgment about what amounts matter to consumers; even small amounts might be life changing for some people. Another industry stakeholder agreed that the bands might be seen to place less emphasis on small amounts that still might be very important to people. One bank thought splitting the data to distinguish between purchase scams (a band of scams less than, for example, £300) could be helpful, as the Code was never really aimed at preventing all purchase scams.

A consumer body said that the biggest concern was reimbursement of life-changing loss amounts, so the published data needs to show if any bank is only reimbursing smaller amounts.

METRICS ON REIMBURSEMENT DECISIONS

Incentive effect
A number of banking industry stakeholders did not consider that consumers would care about the reasons for reimbursement so there would be limited value in publication. One bank did not think the data would be informative because all the reasons for blaming the consumer are the same based on reasonable belief or the consumer didn't act on an effective warning.

One consumer body considered it was important to distinguish between a bank that was good at stopping scams vs a bank that was bad at prevention and was just paying out reimbursement. Another consumer body didn't think that Code reasons would be a useful indicator from a consumer perspective.

Practicality
One bank considered that, for publication of reimbursement and consumer blame reasons, much clearer and consistent mandatory definitions would need to be imposed. Currently banks' application of reasons in the CRM code is inconsistent. The receiving bank at fault category, in particular, was not reliable. Another stakeholder agreed that the reasons in the Code were not used in a consistent way.

Another bank said it did not look at the individual consumer blame reasons but instead they looked at the whole case and so do not collect the data by Code reason. This bank considered there were really only two reasons for blame - no reasonable basis or gross negligence – and no bank should be refusing to pay because a consumer ignored an effective warning. Another bank agreed cases often involved, or could be interpreted as involving, more than one reason.

Unintended consequences: Skewing consumers' behaviour
One bank feared publishing reasons could show consumers what "excuses" to use when claiming reimbursement.

METRICS ON SCAM PREVENTION ACTIVITY

Incentive effect
Many banks stressed that they thought the focus should be on preventing scams but none put forward any firm proposals for a metric that would accurately show how good a bank might be at preventing scams. We heard that data on specific prevention metrics would be difficult to gather. For example, "did a consumer abort the payment because they got a warning, or because the doorbell rang?"

One banking industry stakeholder thought that although it would be really difficult to accurately capture a metric that gave a good picture on how good banks were at preventing scams, the banks would be willing to work on this if the PSR decided to publish Measure One data because it would be very important to show the prevented fraud rate for a bank.

A consumer body said consumers do not care about data on prevention, beyond the proposed core metric on how likely they were to be scammed.

Unintended consequences: Helping scammers
A bank argued that identifying that a bank had specific control issues, such as ineffective warnings, could be a "gift" to scammers.

METRICS ON FACILITATION OF APP SCAMS BY OTHERS

Most banks and industry stakeholders that we spoke to thought that data on how others – social media companies and telecom firms in particular – were involved in facilitating scam activity was very important.

This is clearly a very important point to the banking industry, with many expressing the sentiment that others needed to be on the hook for the part they played and pressured to identify solutions, just as the banks were being pressured to seek solutions. Many pointed out that the origin of scams is completely out of banks' control and there is clearly a degree of frustration that others are not also being compelled to publish data.

A consumer body noted that no-one disagreed that third parties also had part of the responsibility for addressing APP scams. But in a world where scams happen, publishing data by individual bank is about informing consumers' decisions on who to bank with. Another consumer body said that while it accepts the banks aren't completely responsible, as a whole, the banking industry is responsible for the money flowing from consumers to scammers.

> *"The banks have just got to do a much better job - scammers are inside the system, they are getting bank accounts, they are using the faster payment system - it's down to banks to get them out."*
>
> **Stakeholder with a consumer focus**

We heard that the industry was beginning to systematically collect data on the origin of scams and were seeing that many were enabled online by fake advertisements. A sample of this data has already been published. It was collected by manually looking at cases and establishing the origins of the scam. This process might be automated by 2022 and good quality data available.

A number of banks argued that it would be helpful for a regulator to collect and publish the banks' data on third party enablers of APP scams. This would carry more credibility than the banks themselves simply publishing the data to argue that APP scams were not their fault ('they would say that").

METRICS ON COMPLAINTS

Few interviewees expressed any view on whether it would be useful to incorporate data from the Financial Ombudsman Service on uphold rates. One bank said that this could be tricky because the ombudsman is not accepting arguments that a consumer received an effective warning and there is a fundamental difference between some banks and the ombudsman about the interpretation of the Code reasons.

> *"We across the industry have [around] a 99% uphold rate at the Financial Ombudsman, so there is a fundamental difference about what the Code reasons mean."*
>
> **A bank.**

The Financial Ombudsman Service told us that it would release its data on APP scam cases on request – it has a policy of releasing any data that could legitimately be a subject of a freedom of information request. So a consumer body, for example, could simply ask for it.

The ombudsman confirmed that its data was different from banks' data. In particular ombudsman uphold rate data is available when cases have become somewhat historic depending on the time taken to resolve second stage complaint cases. It also told us that the second stage data has also been affected by 'waves' of different reasons for high uphold rates as the Code was introduced and changed.

IMPLEMENTATION

Which banks?

Many industry stakeholders argued that requiring only Code banks to publish could discourage banks from joining the Code, with one arguing that it could also dissuade banks signing up to future voluntary arrangements if they feared doing so could lead them into future mandatory requirements that other banks avoided by not stepping up and doing the right thing.

A number of banking industry stakeholders told us that scams are moving to non-Code banks, and there must be level playing field between Code and non-Code banks. One argued that, if only Code banks published, this could be exploited by other banks to enhance their reputation.

A consumer body noted, however, that the potential impact on the reputation of banks from publication cut both ways. Publication by Code banks only (at first) could enable them to argue they were being transparent and have nothing to hide. If they are good performers, this might enhance their reputation.

Who should publish the data?
Many stakeholders did not have strong views on who published. One was not in favour of just individual banks publishing as they could could cherry-pick the way they expressed the data. One expressed a preference for the regulator publishing rather than UK Finance.

We heard that it would not be appropriate for UK Finance to publish the data because to do so would be incompatible with its role as a trade body representing industry players.

Consumer groups were generally in favour of the PSR publishing the data, as this would mean it had greater credibility.

Quality of the data
We heard that banks did not measure reimbursement in the same way, against the same time-periods, and so on.

One bank argued that there was a need for clearer definitions for published reporting to work including on what is a scam, what is fraud and how different banks report cases.

A consumer body said that it had seen examples of banks telling victims that a scam was their fault and not telling them about the Code - batting them off – and in some cases, frontline bank staff were not aware of the Code. If this is widespread, it could cast further doubt on data quality in that many scams may not be making it into the system to even be recorded as a scam case.

We heard that UK Finance could collect, clean and ensure data consistency (but not publish). More than one bank said UK Finance currently did a good job around auditing data definitions and collection, and processes around collecting the data worked well and should be used if the PSR decided to go ahead with publishing data for Measure One.

In response to a question from us about how the PSR could be assured the data had been collected properly by banks, one bank said if the PSR made a formal information request the banks would have a binding duty to make sure the data was accurate.

How often should the data be published?
There was no strong disagreement with a 6-monthly publication timetable. Some consumer bodies and banks might favour moving towards quarterly, some in the industry considered annual publication could be appropriate.

One Code bank noted that the schedule may not be that important – it might only need to be repeated a few times to have the necessary impact, noting that the CMA consumer service data was a big prod to executive teams of banks who were shown to be poor performers and performance across the industry improved quickly.

Section 4: Assessment of headline metrics

We used the interview feedback, the responses to the PSR's call for views and our analysis to assess the key metrics.

We have not been asked to recommend whether Measure One should in principle be taken forward but only to identify the pros and cons of different metrics and to make recommendations on the metrics that could be included were Measure One to proceed.

Based on interviews, the responses to the PSR's call for views, and our own analysis, we summarise below our assessment of each potential metric against our criteria set out earlier.

METRIC A: THE PROPORTION OF APP SCAMMED CUSTOMERS WHO ARE LEFT - FULLY OR PARTIALLY - OUT OF POCKET

- by volume: total APP scam cases[8] where the cost is fully or in part borne by the victim, as a percentage of all the sending bank's APP scam cases; and
- by value: total value of APP scam losses borne by victims, as a percentage of sending bank's total APP scam value.

[8] On a 'cases closed' basis.

Incentive effect

The relative likelihood that a scammed consumer is left out of pocket, even if they have been reimbursed in part, is highly relevant information for consumers, likely to attract comment from consumer groups and the media, and therefore very likely to have a reputational impact on banks.

In part reimbursement arises: where a bank has failed to meet CRM Code protection standards and where the consumer is also deemed at fault; where a receiving bank has succeeded in repatriating some of the lost funds; or where (we were told) a bank has negotiated with a consumer to provide an in part rather than full reimbursement.

We do not think that publishing the number of in part reimbursed cases is appropriate. This is because we think consumers would principally care whether or not a scam will leave them out of pocket and splitting the metric in this way would not significantly improve the incentive effect.

We heard some arguments that a distinction should be made between reimbursement and repatriation but we do not think consumers would care a great deal whether their bank gives them repatriated funds or simply reimburses them, as long as they are not left out of pocket. In any case the repatriation rate is more relevant to receiving banks' performance (see headline Metric C).

Practicality
CRM Code banks already report the relevant data, by cases closed within the reporting period and by total scam value. The CRM Code sets out relevant definitions.

Some cases take longer than others to resolve. Repatriation can in some cases take a significant time to be finalised. Cases can be marked as closed, but repatriation is later actioned. We consider that, if the reporting period is 6 months or longer, the effect on published data comparisons of any late repatriation would be small.

Potential for unintended consequences
APP scams may involve very different sums, from £10s to £10,000s. There is a risk that publication of data on the volume of cases not fully reimbursed alone could skew banks towards reimbursing many small, but few large, value scams. It is important therefore that Metric A is published by both total value as well as by volume of cases. See also our discussion later of publishing data by scam 'value band'.

We think it is very unlikely that publishing Metric A would, because of a moral hazard, skew consumers' behaviour adversely.

It is conceivable, if consumers were very confident they would get their money back, they might take a little less care over small transactions. But given current levels of full reimbursement, we think there is very little risk that publishing this data would result in consumers being very confident that they would not be left out of pocket if they fall victim to a scam.

Fairness
We recognise the arguments made by many banks that publishing data on reimbursement alone, in the absence of data about banks' performance in preventing APP scams, could give an incomplete and potentially unfair picture. (See discussion below of headline Metric B).

We explored whether it would be fairer to have a de minimis threshold so that, for example, consumers left less than 10% out of pocket would count as fully reimbursed.

But on the one hand a consumer left 10% out of pocket could still bear a loss of thousands of pounds, while on the other hand if less than 10% was a small sum it is not clear why a bank would choose to leave a such consumer out of pocket at all.

We therefore don't recommend a de minimis threshold. We consider publishing Metric A by value as well as volume helps paint a fair picture of bank's performance on reimbursement.

METRIC B: THE BANK'S APP SCAM RATE AS A SENDING BANK

- by volume: total number of APP scam payments as a percentage of total number of push payments by consumers; and
- by value: total value of APP scams, as a percentage of total value of push payments by consumers.

Incentive effect

We consider the relative likelihood of falling victim to a scam is highly relevant information for consumers, likely to attract comment from consumer groups and the media, and therefore very likely to have a reputational impact on banks.

We think that expressing this relative likelihood both as a proportion of the value of consumer payments made and of the number of payments conveys a rounded picture about the bank's performance in preventing both numbers of scams and the largest scams involving life-changing sums.

Practicality

CRM Code banks already report the number and value of APP scam payments. The Code sets out relevant definitions.

The Code covers push payments by micro-enterprises (with less than 10 employees) and consumers. We understand that typically banks can broadly distinguish payments made from consumer or personal accounts from payments made from business or corporate accounts.

Banks may not distinguish these categories in precisely the same ways, and in particular accounts used by micro-enterprises might fall into either category. We expect it would nevertheless be possible for banks to report total consumer payments in a way that is consistent enough to use as the denominator in Metric B[9].

There is likely to be a mismatch between the timing of recording confirmed scam payments and recording total push payments.

We consider the effect of this mismatch is likely to be limited, particularly over a six- or twelve-month reporting period, given limited volatility of total payment numbers over longer time periods.

We heard that data on the value of APP scams could be market sensitive. We are confident this can be dealt with by ensuring appropriate procedures for publication of the data, noting that regulators often need to publish market sensitive information.

Potential for unintended consequences

We heard a range of views on the degree to which publishing sending bank scam rates risked providing help to scammers by potentially identifying banks with weaknesses on prevention though most stakeholders did not consider publishing overall scam rates would give material aid to scammers.

We tend to agree with the arguments that if a bank's consumers are disproportionately becoming victims of scams then that bank is already being targeted by scammers.

We also heard that it should not typically take sending banks a long time to rectify the sorts of weaknesses in their controls that might cause APP scam rates to increase. There is likely to be significant lead time before the first Measure One publication and this would allow banks to address any existing vulnerabilities – they have already been on notice for some time that the regulator is seriously considering publication.

A sending bank that suffered a specific new vulnerability would typically have a period of months in which to rectify any weaknesses before these showed up clearly in published Measure One data, assuming publication every 6 or 12 months (and of course in dramatically exceptional circumstances could make representations to the PSR).

[9]The total number of UK businesses with less than 10 employees is not large (a little over a million) compared to the total number of UK consumers (and micro-enterprise and charity scam payments comprise only around 4% of APP scams), so we don't think any differences between banks in how total micro-enterprise payments were counted would raise material concerns.

We cannot say that publishing Metric B data would be entirely without risk. But we have not seen convincing evidence that any risks are likely to outweigh the benefits of publishing. Formal consultation would provide a further opportunity to explore evidence for any material risk.

Fairness
While we can understand there could theoretically be a trade-off between a bank's performance on preventing APP scams and its reimbursement rate, the data available to us hasn't allowed us to verify any trade-off.

Regardless, we think publication of Metric B data alongside Metric A is likely to give a fuller, and potentially a fairer, picture than publishing Metric A alone.

We explored expressing the APP scam rate, by volume, in different ways, including expressing it as:

- scam cases as a percentage of the total number of consumer accounts;
- scam cases as a percentage of the total number of consumer payments; and
- scam payments as a percentage of the total number of consumer payments.

We think, on balance, scam payments as a percentage of the total number of consumer payments makes most sense. This is clear, and directly communicates the likelihood of a payment being a scam payment. Unlike using cases, it has the potential to skew comparisons between banks because banks may be subject to different scams with different average numbers of payments per scam.

But we think any skewing would be limited over a six- or twelve-month reporting period. Using consumer accounts as the denominator could skew comparisons between banks with different types of consumer base who use their accounts in different ways.

There is a potential for differences between banks' consumer bases to drive differences in sending bank scam rate and this seems more likely if smaller banks were included in Measure One (see our later discussion on implementation and which banks should be included).

If smaller banks are included, the data could be reviewed before publication and, if it appears that there are legitimate explanations for different outcomes depending on size of bank, the data could be grouped into, for example, large banks, digital challenger banks and so on.

If any consistent differences were identified between groups in ways that could be driven by differences in consumer base rather than banks' actions, then it would be possible to present the published data by type of bank group.

METRIC C: THE BANK'S APP SCAM RATE AS A RECEIVING BANK, TAKING INTO ACCOUNT REPATRIATION

- total value of APP scam payments received minus the value repatriated, as a percentage of total value of push payments received from consumer accounts.

We also considered whether the value of app scam payments to receiving banks should in addition be published 'gross' without netting-off repatriation.

We recognise that not all push payment-receiving entities are banks. Payments may also be received, for example, by crypto exchanges, PayPal and so on. We use the term 'receiving bank' to cover all receiving entities.

Incentive effect
We think publishing receiving bank APP scam rates by total value (taking into account the funds that receiving banks intercept and repatriate) would provide an effective reputational incentive for banks (and for other receiving entities) to tackle scams in their role as a receiving bank.

We think publishing the data net of repatriation best targets the incentives on the outcome of minimising customer losses from scams.

We heard that consumers care about banks funding scammers, and information about which banks send the most funds to scammers could affect consumers' choice of bank – since consumers do, of course, also bank with some banks that are receiving banks.

Consumer groups told us this would be important information and we believe they are likely to highlight it, so it could affect banks' reputations generally, and this is likely to contribute to the choices consumers make even if the information does not directly drive switching on its own.

Practicality
Data on receiving banks is not currently reported, so there would need to be appropriate preparation and a dry run.

We consider this data should be collected and reported by sending banks, because:

- receiving banks will not have access to all the data, in particular whether a payment is a scam and the total number of payments they receive from, specifically, consumer accounts; and
- even while not all banks are subject to Measure One, comparative data could be collected on all receiving banks.

We would envisage each sending bank subject to Measure One reporting their APP scam value, repatriated funds received, and total consumer payments sent, split by individual receiving banks.

The data from each sending bank would then be aggregated centrally into combined data on each receiving bank. Receiving banks should be given an opportunity to review and to corroborate or challenge their own data before it was published.

One industry stakeholder pointed out that the set of receiving banks of particular interest might change on a regular basis, and also that "on-us" transactions should be included for a fair picture.

We understand some sending banks may need to invest in changes to their case management systems in order to report repatriation received by receiving bank and the timescales for this should be explored by the PSR in any consultation it issues.

Potential for unintended consequences
Based on the evidence we heard, we consider publishing data on scam rates by receiving bank might help scammers exploit receiving banks with relative vulnerabilities in systems or processes.

For example, vulnerabilities in their onboarding processes resulting in easier establishment of mule accounts, or in their detection of unusual patterns of payments received into accounts that have been taken over by scammers.

At the same time, we heard that scammers are already very aware which banks are easiest to exploit in terms of receiving payments.

We also heard that it could take considerably longer for receiving banks to rectify any weaknesses in their systems and processes, than to address any vulnerabilities in sending bank protections.

We heard timescales ranging from six to eighteen months once a vulnerability has been identified, depending on the issue and bank.

We note that receiving banks would likely have many months from detecting a new vulnerability in their data before that data was reported and then published under Measure One. But if banks were able to demonstrate that this was not long enough to mitigate the risk of helping scammers, the PSR could consider collecting Metric C data alongside the other metrics, but publishing it with a delay sufficient to give receiving banks time to address any identified vulnerabilities.

We think that a significant reputational incentive would still be maintained, because banks would know the data would still become public, even if delayed.

We note that the point we made previously about the banks being on notice of the intention to publish also applies. That is, it is known that the regulator is seriously considering the matter of publication, and the timescale before publication would be at least the time it takes the PSR to consult, reach a decision, arrange to collect the data, run any necessary trials and checks and publish the data.

Whether a further delay to publication is needed, and the precise length of any delay, should be explored further in formal consultation and perhaps individual banks may wish to make confidential submissions to the PSR about their particular circumstances.

We think that publishing the Metric C data before netting-off repatriation could potentially skew receiving banks' actions to tackle the harm from scams. Banks should be focused on whatever set of actions minimise the net impact on customers.

Fairness

The data reported on each receiving bank would not be complete, since not all sending banks are likely to be reporting under Measure One at first (see later discussion) – although we expect that the six largest banks would account for a substantial proportion of total payments.

But the data ought to be a broadly fair and representative of each receiving bank's performance and therefore suitable for making fair comparisons.

Publishing a bank's performance both as a receiving bank and a sending bank gives a fuller picture of their overall performance on APP scams.

For example, a bank that performed relatively well in preventing scams on its own customers might be hosting a disproportionate number of mule accounts involved in scams on other banks' consumers.

But unlike between Metrics A and B there is no obvious potential trade-off between a bank's performance on Metric C and on other proposed metrics. We do not consider it unfair to publish Metric C alone for a bank.

RECOMMENDATIONS

We recommend Metrics A, B and C are published under Measure One, with Metric C published net of repatriation.

As part of formal consultation, the PSR should give consideration – based on evidence from banks - to whether there is any need to delay publication of data under Metric C sufficiently to allow receiving banks to address vulnerabilities that the data might reveal to scammers and, if so, for how long.

Section 5: Assessment of other metrics

Using the same sources, we assessed some additional metrics and variations of metrics against our criteria.

MORE GRANULAR VERSIONS OF THE METRICS

Proposals to split published data by:

- scam type, for example, purchase scam, investment scam; and
- the amount of money involved in the scam by value band, for example, up to £1000, £1000-10,000 and so on.

Incentive effect

In principle splitting data by scam type could inform consumers of their separate likelihoods of being scammed and their separate likelihoods of being reimbursed for each scam type (romance scam, purchase scam) by bank. But it's not clear to us that consumers could make good use of such data. Consumers do not typically know in advance what type of scam they might encounter.

Publishing data by scam type could also add considerable complexity to Measure One, and increase scope for distracting attention from the headline metrics which may diminish the strength of the reputational incentive.

At the same time, we recognise that scam types covered by the CRM Code have different origins and modes of operation and this should of course be reflected in other initiatives dealing directly with consumer education and bank prevention activities.

We think how banks treat high-value, life-changing scam losses, from whatever type of scam, is highly relevant to consumers. This is where there is potentially most harm to consumers and where the incentive for banks not to ensure full reimbursement may also be greatest.

Publishing Metric One, the consumers left out of pocket, both on a case volume and total value basis should shed some light on whether a bank is tending disproportionately to reimburse low value scams.

But we think publishing data on consumers left out of pocket (volume and value) split by scam 'value band' is also justified, despite the extra complexity, to inform consumers of their likelihood of getting their money back for scams involving different sums.

Practicality
Data is already reported by Code banks by scam amount in value bands of less than £1000, £1000-£10,000 and greater than £10,000.

We have considered whether these current band thresholds are the most appropriate. We think there is a case for reducing the upper threshold from £10,000 to £5,000 to capture more fully what can be life-changing sums for some people.

We heard a suggestion for an additional upper band of £100,000+, though we think comparisons between banks might be skewed by the smaller number of scams of this size. The location of value band thresholds could be considered further in formal consultation.

Potential for unintended consequences
Publication of Metric One by value band does more to help address the potential risk, identified earlier, that publishing data on the volume of cases not fully reimbursed alone could skew banks towards reimbursing many small but few large money scams.

We heard that publishing Metric One by value band might be perceived as downplaying the importance of reimbursing low value losses - and relatively low value losses can be highly significant, even life-changing, for some consumers. We have proposed that the upper value band threshold could capture more of the large, potentially life-changing sums than existing reporting. We do not think simply reporting data disaggregated by value band would reduce focus on reimbursement of low value losses. Indeed reimbursement rates for these would be more transparent.

Fairness
For smaller banks, splitting Metric One by scam value is likely to increase the risk of comparisons between banks being skewed by a handful of cases, particularly in the upper value band. If smaller banks were included in Measure One, the data could be reviewed to determine whether splitting data by value band was appropriate for smaller banks.

Recommendations
We recommend that Metric One data (by volume and by value) is published by scam value band as well as for all scams as a whole. Appropriate value bands might be less than £1,000, £1,000-5,000 and greater than £5,000, but the precise thresholds should be determined by further consultation.

METRICS ON REIMBURSEMENT DECISIONS
Proposals for additional metrics on the:

- proportion of decisions to reimburse resulting from, for example, sending bank fault, receiving bank blame, consumer vulnerability, shared blame with the consumer, no blame, automatic reimbursement; and
- proportion of decisions to deem the consumer at fault resulting from, for example, acting unreasonably, ignoring effective warning or CoP result, gross negligence or acting fraudulently.

Incentive effect
There is a potential risk that some banks could gain credit for apparently good performance that in fact resulted under CRM Code rules from their poor adherence to CRM Code standards. Publishing reasons for reimbursement decisions might help reveal if this was the case. Publishing reasons for non-reimbursement decisions (that is, consumer blame) might provide incentives for banks to justify their use of reasons for blaming consumers.

However, CRM Code reasons are not relevant to banks not in the CRM Code, and seem unlikely to be understood by consumers. Consumers are principally interested in whether or not they are likely to be out of pocket, and are unlikely to care much about the range of CRM Code reasons why this might be the case. Including complicated reasons metrics would increase the complexity of the published metric suite and could distract attention from the headline metrics that consumers are most interested in.

Practicality
We have been told there is a lack of consistency between banks in how they apply reasons for decisions on reimbursement, and that banks don't necessarily record consumer blame decisions by individual CRM Code reason.

If reasons metrics were to be included in Measure One, considerable preparation work would be needed by banks, UK Finance and the PSR before the data could be ready for publication, including:

- improving consistency of application and reporting of reasons for reimbursement decisions;
- developing a new reporting process for consumer blame reasons (which are not currently reported); and
- developing an approach for how non-CRM Code banks would report on the reasons for their decisions.

There could be significant scope for discretion by banks as to how they classified their reasons for decisions in order to put themselves in the best possible light. Even costly independent audit processes might be unlikely to be able to prevent this.

Fairness
Given evidence of a lack of consistency between banks in their application of CRM Code reasons, and the scope for discretion in deploying reason categories, it seems unlikely that comparisons between banks would be particularly meaningful or fair.

To the extent that a bank's relative performance on reimbursement under Metric A was flattered unfairly by poor adherence to CRM Code prevention standards, this ought to be mitigated by the effect of poor prevention standards on its APP scam rate under Metric B.

Recommendations
We do not recommend including reasons metrics in Measure One.

METRICS ON SCAM PREVENTION ACTIVITY
Proposals for metrics that might shed light on specific activities that sending and receiving banks take to try to prevent scams, for example: numbers of transactions discontinued as a result of warnings; and numbers of mule account closures by receiving banks.

Incentive effect
Banks' overall performance on preventing scams is captured within the proposed APP scam rate metrics (Metrics B and C). Unlike the simple APP scam rates, we do not think consumers would easily know how to interpret data on specific prevention activities. Indeed, it is not clear that data of this kind would be open to easy interpretation. For example: transactions may be abandoned for reasons other than the effectiveness of firms' warnings or interventions; and high numbers of mule account closures could signal poor initial due diligence processes.

Potentially complex additional metrics may also divert attention away from the reputational impact of the headline metrics. For example, focusing attention on the number of transactions discontinued by a bank's customers could divert attention from a headline metric that nevertheless showed the bank maintained a high APP scam rate.

We therefore don't think additional metrics of this kind would contribute to an effective reputational incentive.

Practicality
It is likely that metrics on specific prevention activities could be complex and potentially onerous to design and report on consistently across banks.

For example, it would need to be determined what constituted an abandoned transaction for the purpose of the metric, including at which stage or stages of the transaction process an abandoned transaction 'counted', taking into account the differences in transaction processes between banks and platforms. Banks would all need to be able to consistently report against the agreed definition.

Potential for unintended consequences
Given the practical difficulties of the data under this potential metric, there is a risk that a metric would not give a particularly good picture of the best activities that reduce scams. Banks' efforts might be skewed into those activities that are measured and published even if these are not key to reducing scams.

Fairness
We noted earlier that we had not been able to verify any trade-off between the performance on reimbursement and APP scam rate. In the case of at least one specific prevention activity, there is evidence that there is no trade-off with non-reimbursement rates. The LSB's Code review data analysis found that those banks achieving low overall numbers of discontinued transactions were disproportionately more likely to decide consumers alone were to blame.

It is therefore far from clear that publishing metrics on specific prevention activities would contribute to overall fairness of Measure One comparisons between banks.

Recommendations
We do not recommend including metrics on scam prevention activity in Measure One.

METRICS ON FACILITATION OF APPP SCAMS BY OTHERS
Proposal for an additional metric on the numbers of scams originating with different third parties, such as social media platforms, search engines, telephone companies, etc.

Incentive effect
It could highly be beneficial for data on other facilitators of APP scams to be published. This would help build the case for action by those social media, telecoms and other firms who could play a key role in tackling origination of APP scams.

This data does not however need to be published by individual bank, as the purpose would not be to incentivise banks. Therefore Measure One is not the appropriate place to take this forward.

Practicality
The banking industry is beginning systematically to collect this data at industry level. We think this is a very positive development. This data is likely to be valuable and of great help in the broader efforts to tackle scams.

We think there is a good case for financial and telecoms regulators to be involved in explicitly supporting the collection and publication of banking industry-level data on this issue, and we would be surprised if any regulator able to play a part in reducing scams is not already alive to the issue and ready to take action.

Potential for unintended consequences
Publishing data on the numbers of scams originating from individual third parties, for example social media platforms, could in theory provide scammers with information on how and where to target victims. But in practice scammers are already likely to be well aware of the availability and relative effectiveness of the different third-party channels available to them.

There is potential for an unproductive debate about who carries responsibility for addressing consumer harm from APP scams – banks or third-party facilitators. This would be a mistake and a distraction from incentivising all parties to play their part. Publishing data on third parties separately from Measure One could avoid fuelling this risk.

Fairness
The data does not need to be published by individual bank, so the issue of fairness between banks does not arise.

Recommendations
Publication of banking industry-level data on third parties would clearly help the broader effort to tackle scams, and we think that regulators should do whatever they can to facilitate and encourage this, but we do not recommend including this metric as part of Measure One.

METRICS ON COMPLAINTS

Proposal for an additional metric on complaints about banks' handling of APP scams, including the number of complaints to the Financial Ombudsman Service and the ombudsman's uphold rate against each bank.

Incentive effect
Publication of data on numbers of complaints to the ombudsman and its uphold rate for each bank is important for incentivising banks to respond appropriately to complaints in the first place.

The outcomes of complaints resolved by banks themselves are already be captured in their published Metric A reimbursement rates. Banks' performance on the process of complaint handling is a general consumer service issue that ought to be part of oversight of banks' service performance rather than Measure One.

Practicality
The ombudsman's comparative data on the APP scam complaints it receives and upholds can already be published by the service. There is no need for the PSR to direct banks to provide it for publication under Measure One.

Potential for unintended consequences
We do not think the Financial Ombudsman Service publishing its data aids scammers or adversely skews banks or consumers' actions.

Fairness
The ombudsman's complaints data could in principle be presented alongside Measure One data. However, the uphold rates are available, following investigation and decision-making, on quite a different timescale from the rest of the Measure One data reported by banks on scam and reimbursement rates. So it may not aid – and could distract from – fair comparisons between banks on their latest performance on the Measure One metrics.

Recommendations
The Financial Ombudsman is able to publish its data on APP scam complaints and uphold rates. It is not necessary or appropriate to include a metric on complaints upheld at the ombudsman in Measure One.

Findings on headline Metrics A, B and C

	Incentive effect on banks	Practicality	Potential for unintended consequences	Fairness of comparisons	Recommendation
A: The proportion of APP scammed customers who are left - fully or partially - out of pocket	High	High	Low	High Assuming B also published	Publish by volume and value
B: The bank's APP scam rate as a sending bank	High	High	Low But some uncertainty	High	Publish by volume and value
C: The bank's APP scam rate as a receiving bank, taking into account repatriation	Medium Consider delay between reporting and publication	High	Low With time to address major vulnerabilities before publication (if necessary)	High	Publish by value - consider delay between reporting and publication

Findings on other proposed metrics and variations

	Incentive effect on banks	Practicality	Potential for unintended consequences	Fairness of comparisons	Recommendation
Split metrics data by scam value band	High	High	Low	High Assuming data volatility limited by reporting period / bank size	Split Metric A by value bands
Split metrics data by scam type	Low	-	-	-	Don't include in Measure One
Metrics on reasons for reimbursement and non-reimbursement	Low	Low	Medium	Medium	Don't include in Measure One
Metrics on specific prevention activity	Low	Low	High	Low	Don't include in Measure One
Metric on third party facilitation of scams	Low Data not relevant to individual bank	High	Low	N/A	Don't include in Measure One but support industry-level data publication
Metric on complaints	Medium Data is necessarily historic	Medium Ombudsman can already publish	Low	High	Don't include in Measure One

Section 6: Implementation

We considered some key questions about how any new requirement could be implemented.

WHICH BANKS?

The more banks are covered by Measure One reputational incentives, the greater the likely impact on outcomes for consumers and so in theory, publishing data on all banks is desirable.

Of course, the regulator must take into account regulatory burdens involved in reporting data, which might not be proportionate for the smallest banks, so it may be appropriate to implement a de minimis threshold. The regulator will also consider the practical difficulties of a wide-ranging implementation that could risk delay and losing sight of the urgency of the need to bring about improvements for consumers.

Implementing data reporting to the standard necessary for publication would take more time, and be much more difficult, for non-Code banks than for the Code banks who already report much of the data.

So in the first phase of implementation, we think Measure One reporting requirements should apply to the largest banks, for example the CoP six all of whom are Code members.

Taking this approach - based on bank size rather than Code membership – would avoid the risk of discouraging banks from joining the Code or indeed from joining other future voluntary arrangements. Smaller Code and non-Code banks would be treated the same.

There are precedents for similar performance reporting requirements applying only to the most significant market players, such as the CMA's requirement for publication of banking service quality data.

In subsequent implementation phases, the requirement could be expanded to smaller banks, with the next phase covering perhaps the largest 15 to 20 banks.

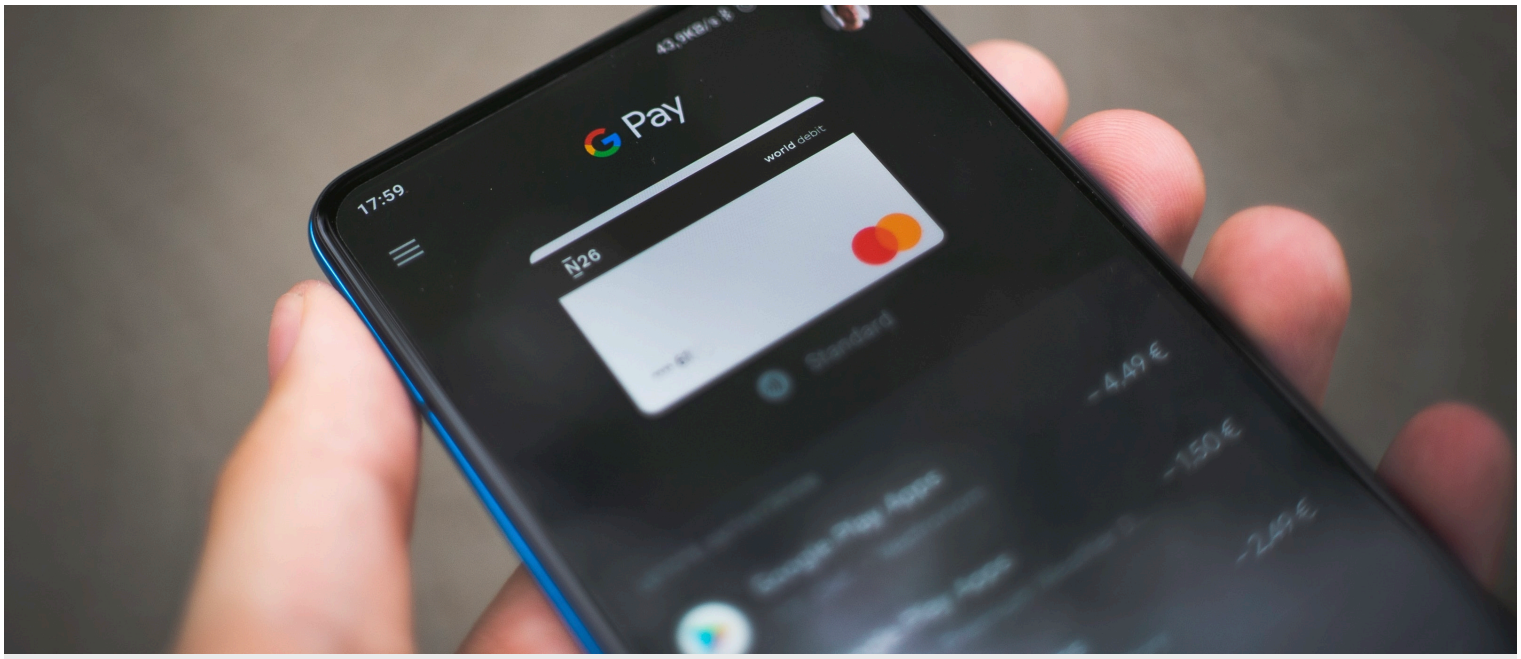

A commitment from the PSR to such expansion could encourage more of these banks to join the Code.

We were told that scammers may be migrating to smaller receiving banks outside the Code. Our proposal that Measure One banks report, as sending banks, on all receiving banks would ensure that all receiving banks were covered by the reputational incentives, even if they were not themselves reporting Measure One data initially.

Recommendations
The requirement to report publishable Measure One data should apply, in the first phase, to the largest banks, with subsequent phases including other smaller banks.

All receiving banks would in any case be subject to reputational incentives from Metric C from the first phase.

WHO SHOULD PUBLISH THE DATA?

Most stakeholders appear to favour the Measure One comparative data, if published, being published by the PSR, and we agree. The data for all banks needs to be easily available and comparable in one place, and the PSR would provide the data with appropriate prominence and authority.

There do not appear to be suitable and available alternatives. Leaving publication to individual banks would make it harder for users to easily access and compare the data. The LSB's locus is over Code signatories, and UK Finance, as a trade body, would not appear to be an appropriate forum for discharging a mandatory regulatory requirement on banks.

Publishing comparisons would of course place an operational requirement on the regulator, and the regulator would want assurance that the data it published was sufficiently robust. We address below how the operational difficulties may be minimised and assurance gained.

In addition to publishing the data comparisons, the PSR could also require the banks subject to Measure One to publish the same comparisons prominently on their own websites. This would be similar to CMA's requirement on banks to publish comparative customer service data - although that data is not always as prominent as it might be on some banks' websites. Banks would also, of course, have discretion to publish greater detail on, and explanation of, their own APP scams data.

Recommendations
The PSR should publish the Measure One comparisons between banks, and require banks to republish the data prominently on their own websites.

QUALITY OF THE DATA

UK Finance currently plays a key role in collecting, collating, cleaning and checking data reported by Code banks. It runs a range of consistency and sense checks, putting back queries to banks. We think UK Finance could continue to play this role at first. This would limit the immediate operational requirements on the PSR, which UK Finance with its existing experience is better placed to undertake.

The PSR should agree the guidance on Measure One reporting that UK Finance provides to banks and the data-checking processes UK Finance undertakes. As part of this, the PSR could consider whether guidance needs to clarify, for the purpose of Measure One reporting, any of the relevant definitions under the Code and ensure that cases are recorded consistently at the point a decision is made.

We have not explored this in detail, but have been told of potential inconsistencies in how, for example, purchase scam payments are categorised.

UK Finance is in a position to discharge this role while Measure One requirements are limited to its members. We think this would remain the case for at least the largest 15-20 banks (depending of course how the market develops).

If and when Measure One reporting requirements were expanded beyond the membership of UK Finance, the PSR is likely to need take on the data collection role. Implementation in phases allows the PSR to prepare itself to do so by shadowing and learning from UK Finance through the initial stages of publication.

Of course UK Finance is not in a good position to assure the robustness and accuracy of banks' internal systems and processes for producing the data. There is currently no self-certification regime for APP scams data reported to UK Finance.

The way forward depends on the PSR's overall approach to assurance, but we consider that the PSR could propose that it asks each bank subject to Measure One to self-certify the quality of its data.

That is, for an appropriate very senior executive (one industry stakeholder recommended this should be the CEO or senior executive at Board level) in each bank to sign-off the data, assuring the PSR of its accuracy.

Recommendations

UK Finance should collect and check data reported by Measure One banks, building on its existing role in relation to data reporting by Code banks.

The PSR should prepare itself to take on this role for when publication requirements extend further than UK Finance membership, by shadowing and learning from UK Finance through the initial stages of publication.

The PSR should require a very senior executive in each bank to assure it that the data provided is accurate and robust.

HOW OFTEN SHOULD THE DATA BE PUBLISHED?

UK Finance currently publishes aggregated data from Code banks at 6-monthly intervals, and this frequency seems broadly appropriate for publication of Measure One data.

While publishing bank-level data more frequently, for example, monthly or quarterly, would be feasible and could give a more up-to-date picture, it:

- could potentially increase the risk of aiding scammers, if banks had less chance to address control failures they identified during the course of a period before data was published;
- could increase volatility in the data, especially for smaller banks, reducing effectiveness and fairness of comparisons;
- would be unlikely to lead to significant additional media coverage; and
- would be more onerous to banks, UK Finance and the PSR.

Publishing annually would give banks even more time to address vulnerabilities they identified, but would reduce the timeliness of data and reduce the number of opportunities for reputational impact.

Recommendations
Measure One data should be published every 6 months.

TRIALLING AND TESTING
We recommended earlier that Measure One is implemented in phases, with the first phase covering the largest banks. This approach would help ensure implementation issues were ironed-out before the requirement was expanded to a large number of banks.

We think an unpublished dry run of reporting, collecting, collating, checking and presenting the data could be undertaken based on the first 3 months Measure One data. This ought to give time for implementation issues to be identified and addressed before publication of the first 6-month period of data.

As part of the dry run, UK Finance could allow each bank to see how they compared with anonymised data for other banks, giving them the chance to challenge the comparisons or review their own data if they considered the comparisons contained errors. The dry run would also provide an opportunity to iron out any inconsistencies between banks in interpreting definitions.

Any delay between reporting and publication of Metric C data on receiving banks would give additional time to dry run and address any implementation issues, including the processes for aggregating the data by receiving bank and putting the results to receiving banks for review and comment.

Like any regulatory measure, the effectiveness of Measure One ought to be evaluated by the PSR following its introduction.

The PSR is in a relatively good position to do this. It already has baseline data on sending bank rates of APP scams and reimbursement by Code bank. If Measure One proves effective as a reputational incentive, an improvement ought to be seen in the reported Measure One data against this baseline.

Of course any improvement seen might have occurred without Measure One. But it may be possible for the PSR to assess the effect of Measure One by comparing any improvement in performance by the banks subject to Measure One with performance of other Code banks and, for scam rates, with the industry as a whole.

Recommendations
We recommend an unpublished dry run of implementation of Measure One based on the first 3 months of data.

Appendix: Interviewees

We are grateful to our interviewees who gave generously of their time and expertise.

A number of banks, industry stakeholders and consumer bodies were extremely generous with their time in helping us in our work.

They gave up their time for interviews, written questions, follow up questions and, in several cases, agreed to give even more time in second interviews so we could check facts and explore arguments we had heard in more detail. We are very grateful for this assistance.

We are very grateful for the time taken by interviewees to engage with us and share their evidence, insight and expertise.

Barclays Bank
Financial Ombudsman Service
Fraud Advisory Panel
HSBC
Lloyds Banking Group
Lending Standards Board
NatWest
Santander
Starling
The Money Charity
TSB
UK Finance
Victim Support
Which?

About Lucerna Partners

Lucerna Partners is a consultancy specialising in regulation and public policy. We advise on strategy, competition, consumer and public policy, and on regulation – whether that's rethinking the regulation of a sector or helping with an immediate and urgent matter. We help firms identify their best regulatory strategies, and present their best case. We help regulators secure the best value for consumers and taxpayers. And we help consumer groups get their voice heard, and policy makers make the right decisions.

Lucerna has worked on payments issues for many years, advising the Financial Conduct Authority (FCA), the Payment Systems Regulator (PSR), payments firms and consumer bodies. Lucerna initiated and delivered for Which? the 2016 super-complaint on Authorised Push Payment (APP) scams, and has provided support to the PSR on the issue subsequently.

+44 (0)207 1935912
enquiries@lucernapartners.com
www.lucernapartners.com