

Design and Incentives in Operational Risk Data Consortia

Homolya, Daniel*- Kiss, Hubert Janos[†]

October 14, 2009

Abstract

In operational risk management, using external data coming from data consortia belongs to the international best practice. These consortia are formed by individual banks that report anonymously their operational risk events related to processes, people, systems or external events. The validity of the data provided cannot be checked fully by a datapool manager, so in principle the quality of the data is questionable. We study using a repeated-game framework the conditions that ensure truthful reporting to the database. First, we show that in the stage game reporting truthfully is not a unique equilibrium. When other banks are believed to lie, a best response is to lie as well. However, as we go to repeated game banks may have sufficient incentives to report the true data, provided the expected punishment upon detected misreport is high enough. Since the parameters determining the expected punishment are variables set by the consortium, the proper choice of them is crucial for the appropriate functioning of a consortium.

JEL Classification: C7, C73, G21, G32

Keywords: operational risk, risk management, data consortia, repeated games

1 Introduction

Application of external data for decision making is indispensable in case of lack of own experience concerning certain events or problems. Need for applying external data emerges in the case of operational risk of credit institutions as well. Operational risk (henceforth, OR) - defined as "the risk of loss resulting from

*Magyar Nemzeti Bank (central Bank of Hungary), senior analyst, Financial Stability Department - Corvinus University Budapest, Ph.D. candidate, Department of Investments and Corporate Finance, e-mail:homolyad@mnbb.hu. The views expressed in the paper are those of the authors and do not necessarily reflect the official views of Magyar Nemzeti Bank, the central bank of Hungary.

[†]Dpto. Análisis Económico: Teoría e Historia Económica, Facultad de Ciencias Económicas, Universidad Autónoma de Madrid, Ciudad Universitaria de Cantoblanco, 28049 Madrid (Spain). Email: hubert.kiss@uam.es

inadequate or failed internal processes, people and systems or from external events" (BIS [2004])- is one of the main novelty of the Basel II process affecting credit institutions' capital requirements. This kind of risk has dualistic characteristics: high frequency low impact (HFLI) and low frequency high impact (LFHI) events are dominating in the frequency-severity space. LFHI events can be measured with difficulties, due to the short observation period of internal loss databases. However, if a risk event has not realised, it does not mean that there is no risk exposure. To handle these problems (that is, to substitute internal, historical experience) expert judgements, scenario analysis techniques or external loss data are usually applied. In this paper we concentrate on external loss databases. External loss data could stem from so-called public databases and consortium-type databases. In the latter case banks are sharing and pooling data on operational risk events with each other on a voluntary basis. We name just a few examples: ORX (international), GOLD (British Banking Association), DAKOR (German public banks), ORD (German saving cooperations), DIPO (Italian Banking Association) and HunOR (Hungarian Operational Risk Database) (see Voit [2007], Homolya-Szabolcs [2008]). OR management is a relatively new phenomenon in banking. Only a small number of benchmark studies on individual banks' operational risk management practice have been published, but those suggest that using external data that stems from databases belongs to the international best practice (see e.g. case of Deutsche Bank in Aue-Kalbrener [2006]).

Secrecy is an integral part of bank management. Banks do not like to disclose information about their functioning. This is especially true for data on failures. That is why in this kind of data consortia banks are pooling operational risk event data anonymously, without full checking possibility by other members and datapool manager, external auditing. Financial reports of banks, as in case of other corporations are verified by external auditors, prudential data reports for authorities are examined by supervisors and central banks depending on the supervisory structure. In case of a voluntary database a datapool manager ("quality insurer") can check the consistency (but not the full validity) of the data. Thus, banks may misreport as the consortium does not have the authority to conduct an overall check. Nevertheless, the lack of published evidence about banks misreporting their OR events and anecdotal evidence suggest that these data consortia contain data with quite good quality. This could be explained by two reasons:

- Banks could have fears of examinations of supervisors. If misreporting is detected by the supervisor, this could have direct consequence in form of a fine. ("fear of punishment")
- Operations of these data consortia are mainly based on a contract signed by banks. This expresses a commitment of banks, and a compliance with a written agreement is a natural requirement for banks operating in a market economy. ("gentleman agreement")

In this study, we focus on the first point and we show how the possibility

of being detected to misreport and the ensuing punishment (monetary or non-monetary as the loss of reputation) may induce banks to report truthfully their OR events. We use a repeated-game setup in which banks play a simultaneous stage game each period. In each period, banks are hit by an OR-shock and they may report truthfully their OR events or they may lie. Banks send their true or false reports to a center that computes a set of values based on the reports. Banks use this set of values to manage their operational risk. Banks play this simultaneous game without a definite end, therefore the game is infinitely played.

We suppose that banks find using the value computed by the center useful only if they believe that all the other banks have reported truthfully. Hence, the first-best case is when each bank reports truthfully its OR events, the center reports back a true set of values which enables the banks to prepare optimally for next period's OR shock (causing OR event). We will consider two types of values that the banks receive from the center. If a bank is able to calculate the true set of values given that the other banks report truthfully, then we call this type of value recoverable. In the other case, we call it irrecoverable. The distinction of recoverable and irrecoverable characteristics is useful, as banks' behaviour could be influenced by the fact that other banks report truthfully, providing a free lunch for the misreporting bank. We show that in the case of recoverable value in the stage game the unique Nash equilibrium implies banks misreporting. For irrecoverable values, the set of equilibria alters and truthful reporting is also a Nash equilibrium. Nevertheless, the multiplicity of equilibria still contains the possibility that banks lie.

These results in the stage game suggest that to induce truthful reporting we need repeated interactions. Since punishment plays a central role in repeated games, we need misreporting to be observable. As stated before, banks' reports to the database could not be validated fully by the datapool manager. Nevertheless, we suppose that with some probability the supervisory authority has the chance to detect if a bank lies. We provide more details later on how the authority may detect a misreporting. A possible story is that as the supervisory authority carries out a regular on-site examination in a bank, it may get to know that a certain OR event happened in the entity. Supposing that the authority has access to the values reported to the center, it may realize that the bank has misreported. We assume that the detection of misreporting entails a punishment. If the expected punishment is high enough, then any bank will find it optimal to report truthfully. This result is in full consonance with the spirit of repeated games: if the interaction is repeated and the stakes and punishments are high enough, then the optimal result may arise as the unique outcome.

Since banks cannot observe directly if another bank has misreported, banks cannot condition their strategy on the history of previous reports. Therefore, in each period banks decide whether to report truthfully or not based on if they believe that the other banks report truthfully or not. The possibility of punishment makes that a bank expects other banks to report truthfully, and then this bank will report truthfully as well. All banks behaving in this manner leads to the first best.

Regarding the punishment, consequences of detected misreporting are reg-

ulated in multilateral agreements among participants of the data consortium. Obvious legal consequence can be the suspension of database membership for a finite time horizon or forever, which means that the bank will not receive data from the data consortium in the future. Moreover, misreporting banks could receive a stigma, and other banks will not like to cooperate with this bank in case of other data exchange initiatives or other type of banking interactions. Although hard evidence is not available, industry-wide rumours suggest that due to this fear of punishment and reputational risk banks tend to treat reporting with as much effort as possible after the launch of the data consortium.

To our best knowledge, our paper is the first to study the conditions and incentives that ensure that banks report truthfully to a operational risk data consortium. Our results show that the key element in a data consortium is that member banks should have a lot at stake. That is, the expected punishment (monetary or reputational) should be high enough to deter banks from misreporting.

The study is organized as follows. In the next section, we review briefly the literature. In section 3 we present the model with the one-shot game, whereas in section 4 we proceed to the analysis of the repeated game. Section 5 concludes.

1.1 Related literature

The literature on operational risk deals generally with two aspects of operational risk data consortia. On the one hand, the main features of risk consortia are studied from an institutional point of view (conditions of establishment, members, contents of data reporting etc.) (e.g. Voit [2007], Wood [2007]). On the other hand, topics related to the usage of external data (e.g. mixing of internal data and external data) are examined from a statistical point of view (e.g Baud et al. [2002], Na et al. [2005]). However, the topics of these articles do not cover the incentives of data reporting from theoretical point of view, and the optimal processes to be applied in order to reach good quality of reporting.

Our paper is closely related to the emerging literature on multi-party computation games (see, for example Smorodinsky-Tennenholtz [2006]). In a multi-party computation game a set of agents has access to some private information. Each agent may get to know her own information at a cost. The agents would like to compute jointly a function whose inputs are the private information of the agents. Agents are selfish and would like to know the true value of the function without giving away their private information. This setting resembles our problem of truthful reporting. This study differs from that approach in that we rely on punishment to induce the socially efficient behavior, whereas in multi-party computation game the focus is on designing mechanisms that elicit the true private information.

Our study is related to the literature on repeated games (see, for instance chapter 5 in Fudenberg-Tirole [1991]). Our model is somewhat special because banks cannot observe the truthfulness of other banks directly. The authority has the possibility of observing that a bank has lied. Thus, the monitoring is delegated to the authority. For simplicity, we assume away any problems related

to principal-agent problems.

2 Model

There is a set of banks $N = \{1, 2, \dots, n\}$ and a special agent called 'the center'. It is assumed that there exists a private secure communication line between every bank $i \in N$ and the center. Banks are *ex ante* identical. At the beginning of the period each of them is hit by an OR shock. Then, banks choose whether to report truthfully or to lie to the center. After reporting the center calculates some set of values based on the reported shocks and then this information is reported back to the banks. Based on this external information and on internal data the bank manages its OR and prepares for the next period's shock. Banks prepare themselves against operational risk shocks by allocating capital for absorbing operational risk based on calculation of the loss profile. External data are good inputs into risk measurement and management, as banks are able to compare internal data with external data ("benchmarking"), and mix the data from different datasources.

Due to the nature of operational risk the data from the center (e.g. banking associations, special consultant firms) which is based on industry-wide experience helps to build up an appropriate OR management. To have access to it a bank has to provide information about the own OR events. We assume that while each bank values true external data, but no one likes to give away information about internal losses.

Assume that each bank is hit in period t by a Bernoulli shock with unknown parameter p_t . Suppose that banks' shocks are independent and the probability is uniform across banks. Therefore, with probability p_t bank i suffers a shock of magnitude 1, and with the complementary probability it suffers a shock of magnitude 0. The resulting loss is represented by $x_i \in \{0, 1\}$. We denote the vector of the true shocks by $x = (x_1, \dots, x_n)$. Let \hat{x}_i denote the report made by bank i , whereas the vector of reports is represented by \hat{x} .

The strategy in any period (that is, in the stage game) of bank i is to be truthful (denoted as tr) or to lie (denoted as l). Thus, $s_i = \{tr, l\}$ for $i \in N$. Since the game is repeated, banks have to specify their strategies for each period, so $\mathbf{s}_i = s_{i,t}$ for $t = 1, 2, 3, \dots$ where $s_{i,t}$ is bank i 's strategy in period t . Being truthful means that the report made by the bank i (\hat{x}_i) is equal to the realization of the shock ($\hat{x}_i = x_i$). Lying implies that $\hat{x}_i \neq x_i$.

Suppose that the function $f(x)$ is the best predictor for the next period's probability of being hit by a shock of magnitude 1. For simplicity, assume that banks find external data useful only if they believe them to be truthful. Therefore, from a bank's perspective the main question is if all the other banks will report truthfully or not; and in face of this belief it should decide how to report. This assumption can be seen as if the bank believes that the values given by the center will be useful for its OR-management if the reports have a sufficiently good quality. Based on the banks' reports, the center calculates $f(\hat{x})$.

In this paper, we disregard potential principal-agent problems between the center and the banks, the center cannot be captured or influenced by the individual banks.

2.1 Stage game

First, we analyse the stage game. We define utilities, strategies and study two types of data.

2.1.1 Utility

For the bank, participating in a data consortium does not affect other parts of the bank, so we assume that the utility of it relates in an additive manner to the total utility of the bank. Hence, banks will report truthfully if the benefits outweigh its costs ($c_i \geq 0$). Costs include tangible efforts (e.g. enhanced documentation costs) and intangible elements (like disutility from reporting weaknesses) as well. If a bank does not participate in the consortium, then its utility related to the consortium is zero.

Assume for simplicity that the benefits and costs are uniform (1 and c , respectively) across banks.

As stated already, banks value only external information ($f(\hat{x})$) which is believed to be truthful. Without loss of generality we normalize the benefits of knowing the truthful external data and managing properly next period's operational risks to 1. We suppose that bank i 's utility if it reports truthfully has the following form:

$$u_i(f(\hat{x}), c \mid \hat{x}_i = x_i) = \begin{cases} 1 - c & \text{if } f(\hat{x}) = f(x) \\ -c & \text{otherwise.} \end{cases} \quad (1)$$

Note that the OR shock does not appear in the utility function. The utility function takes into account if OR is properly managed with the help of external data (benefit of 1) or if it is not (zero benefit).

If bank i lies, then the utility will depend on whether it can somehow calculate \bar{x} or not, as we will consider later. The immediate costs of lying are zero.¹ Formally:

$$u_i(f(\hat{x}), c \mid \hat{x}_i \neq x_i) = \begin{cases} 1 & \text{if } f(\hat{x}) = f(x) \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The outside option, that is not participating in the consortium is supposed to yield zero utility. Note that lying when $f(\hat{x}) \neq f(x)$ and the outside option yield the same payoff. Hence, we suppose that staying out of the consortium or not being serious as a consortium member are the same from the perspective of the bank.

These payoffs capture the main problem involved in anonymous reporting. They imply that a bank is willing to participate in the consortium and report

¹The cost of lying that is not immediate is the expected punishment.

truthfully if there is some mechanism which ensures that reports are truthful. Otherwise, staying out is a better option.

A necessary condition of participating in the consortium is that $1 > c$, which is supposed to hold.

2.1.2 Strategy

The optimal action of any individual bank relies on what it believes about the truthfulness of the other banks' report. To make these beliefs more concrete, we suppose that the banks hit by the zero shock report truthfully. On the contrary, banks hit severely (shock 1) have incentives to hide this fact. Therefore, an individual bank may believe that the banks hit severely by the OR shock either report truthfully or they lie. Since, except for the shock, banks are identical, these two beliefs cover all the possibilities.

We have a simultaneous-move game. Banks report their type without knowing what the other banks report. It implies that banks cannot condition their report on observable reports made by other banks. As a consequence, bank i 's strategy in the stage game is $s_i = \{tr, l\}$, as stated earlier.

We focus on pure strategy equilibria. As we will argue later, mixed strategies do not change substantially the results.

2.1.3 Reporting

Let us briefly consider how individual banks report to the operational risk data consortia. Generally, banks are reporting with full identity to a center, who checks the consistency of data, but could not validate fully them and then publishes the data anonymously. There are also exceptions. In the case of the Hungarian operational risk data consortium (HunOR), banks report their OR events with anonymity to the center.² The practice of the Hungarian data consortium reflects cultural and confidence aspects. Hungarian banks are reluctant to report in absence of total anonymity.

The center generally reports back individual events' data with anonymity or scaled data. Scaling means transforming one bank's loss into other bank's loss based on relationship between size indicators of institutions and loss event data. For instance, OR data occurred in the smallest bank are scaled up by a factor representing the difference of size between the smallest bank and the bank that wants to access to the OR data.

Data consortia do not report scaled and unscaled data together. If banks know the scaling mechanism, then - from scaled and unscaled data - they would be able to recover the size indicators of the source bank of the loss event, which would breach anonymity. Besides individual loss events' data, the data consortia are reporting statistical measures (such as mean, median, standard deviation, skewness, kurtosis, etc.).

²In the background, all the events are marked with an identifier for the bank. However, it can be reached only by system administrators in case of empowerment by members of the data consortium.

The previous short description of reporting practices reveals an inherent trade-off. Banks can use external OR data most efficiently if they get a detailed information about what has happened exactly. Nevertheless, detailed information may reveal the identity of the bank where the OR event occurred which goes against anonymity. Hence, a consortium has to strike the right balance in this issue. We try to capture a similar idea by distinguishing between recoverable and irrecoverable data. In our distinction, we do not go as far as to identify a bank based on the data it provided. We are just interested whether the value computed by the center gives incentives to report truthfully or not. If the true value can be computed by an individual bank given that the other banks report truthfully, then we call the value recoverable. Otherwise, the value is irrecoverable. We show that the latter gives more incentives to banks to be truthful.

2.1.4 The average function - a case of recoverable values

In this subsection, we suppose that the average of shocks is the best predictor for the next period's probability of being hit by a shock of magnitude 1. Therefore, banks want to know the average of shocks (denoted by \bar{x}) which is computed by the center. Based on the reports the center computes

$$f(\hat{x}) = \frac{1}{n} \sum_{i=1}^n \hat{x}_i.$$

Notice that believing that all other banks report truthfully does not imply that bank i will do it as well. If bank i believes that all the other banks report truthfully, then for any $c > 0$, a bank has no incentive to incur the costs, since it is easy to calculate the true average of shocks, \bar{x} . Given truthful reports of all the banks except i , the average can be calculated as $\bar{x} = f(\hat{x}) + \frac{(x_i - \hat{x}_i)}{n}$. By recovering \bar{x} and without incurring the cost bank i 's utility is 1, whereas being truthful implies the cost c .

This result holds for any function $f(\cdot)$ the banks might be interested in for which given the truthfulness of the others a bank can calculate $f(x)$ for any reported \hat{x}_i . Let \hat{x}_{-i} denote the vector of reports of all the banks except for bank i .

Definition 1 *We call the set of values $f(\cdot)$ calculated by the center recoverable if given the truthful reporting of all the banks except i for any \hat{x}_i and any \hat{x}_{-i} bank i can calculate $f(x)$.*

Now we are ready to state the first result.

Proposition 1 *In the stage game, lying is the dominant strategy if $f(\cdot)$ is recoverable.*

Proof. If $f(\cdot)$ is recoverable, then even if all the other banks report truthfully, bank i would not report truthfully, because it is a profitable unilateral deviation.

If bank i reports truthfully, then its utility is $1 - c$. By lying bank i does not incur the cost, but benefits from knowing $f(x)$, so its utility is 1. Clearly, the latter utility is higher, than the former one, so lying pays off. If bank i believes that not all the other banks report truthfully, then bank i has no incentives to incur the costs. If bank i lies, then its utility is 0, whereas truthful report costs c without any reward. Therefore, independently of what the other banks do, lying is the dominant strategy. ■

The idea is conveyed naturally by the following payoff matrix:

bank $i \setminus$ bank $-i$	tr	l
tr	$1 - c, 1 - c$	$-c, 1$
l	$1, -c$	$0, 0$

Bank i is the row player, and all other banks (bank $-i$) are the column players. The strategy tr of all the banks except i stands for the case when all of these banks report truthfully, while lying means that at least one of these banks misreports. This game is a prisoner dilemma with the unique pure-strategy Nash equilibria: (l, l) .

This result shows that if banks played just a one-shot reporting game, the calculated value by the center would be a false one. This false value could not be used for efficient OR management.

2.1.5 Irrecoverable value

We show briefly that even if the set of values calculated by the center were not recoverable, truthful reporting would not be guaranteed either.

Definition 2 *We call the set of values $f(\cdot)$ calculated by the center irrecoverable if given the truthful reporting of all the banks except i for any \hat{x}_i and any \hat{x}_{-i} bank i cannot calculate $f(x)$.*

An example of irrecoverable data is when banks are hit by a continuous shock, that is $x_i \in [0, 1]$ and report on the whole zero-one interval. Suppose that they are interested in the average and the variance. Hence, the center now reports back a vector of two elements, $f(\hat{x}) = \frac{1}{n} \left[\sum_{i=1}^n \hat{x}_i, \sum_{i=1}^n (\hat{x}_i - \bar{x})^2 \right]$. From the data on the average, bank i is able to calculate the true \bar{x} even if she misreports it, but she cannot recover the variance. She can calculate the true $(x_i - \bar{x})^2$, but she cannot do it for the other banks, that is $\sum_{j \neq i}^n (x_j - \bar{x})^2$, because she does not know x_{-i} .

Proposition 2 *In the stage game, there are multiple equilibria if $f(\cdot)$ is irrecoverable.*

Proof. Contrary to the recoverable case, now even if all the other banks are truthful, bank i cannot calculate $f(x)$ and without the true external data it

cannot manage properly its OR. Given the utilities of truthful reporting (see (1)) if bank i believes that all the other banks report truthfully, then the optimal action is to report truthfully. Nevertheless, if bank i believes that other banks will not report truthfully, then the optimal action is to lie. There are two Nash equilibria. If banks believe that all the other banks will be truthful, then each bank will be truthful and the first best obtains. The corresponding equilibrium strategy profile is $\mathbf{s} = \{s_i = tr\}_{i \in N}$. If banks believe that there will be banks that will lie, then each bank will lie. The corresponding equilibrium strategy profile is $\mathbf{s} = \{s_i = l\}_{i \in N}$. ■

The following payoff matrix represents concisely the previous ideas:

bank $i \setminus$ bank $-i$	tr	l
tr	$1 - c, 1 - c$	$-c, 0$
l	$0, -c$	$0, 0$

This game is a stag hunt with two equilibria: (tr, tr) and (l, l) . Multiplicity of equilibria shows that irrecoverability mitigates the problem of misreporting, because it allows for the possibility that all banks report truthfully. Nevertheless, it does not guarantee the favourable outcome, misreporting remains an equilibrium if others choose to do so as well.

Even though we have shown the possibility that the first best arises in the irrecoverable case, besides the multiplicity of equilibria there are several problems. Suppose that all banks plan to be truthful, but there is nothing that would make the other banks to believe it. Beliefs are not coordinated, so even if all banks would like to be truthful, all of them may end up lying if all of them believe that the others will lie. Therefore, misreporting may become a self-fulfilling prophecy.

Repeated games help to overcome these problems. From repeated games we know that although in a one-shot setup some games result in unfavorable outcome, when it is played infinitely, then the favorable outcome emerges in equilibrium. In repeated interactions, optimal punishments guarantee the uniqueness of the favourable equilibrium.

3 Repeated game

Data-sharing is not defined over a fixed horizon, so it is natural to apply an infinitely repeated setup. In this section first we explain how punishment takes shape in our game. Then, we briefly analyze the benchmark case with full information and after we proceed to study the stochastic detection case. To illustrate our results, we also develop a small numerical example.

Since data consortia are voluntary associations and operational risk events of individual institutions could not be explored fully by external point of view, there could be no full checks on the truthfulness of data. Thus, it is not easy to detect deviation from the optimal play (that is, reporting truthfully) which makes punishments difficult. There are several potential ways to detect that some banks have lied. We could assume that although other banks' actions

are not directly observed, but through the payoffs banks may realize that some other banks have lied. Low utilities related to the database would indicate that there are banks that are misreporting. The problem with that approach is that the fact that some banks lied does not reveal which banks did behave so. As a consequence, it is impossible to impose punishment in an efficient way onto misreporting banks. Notice that it is not efficient that banks punish misreporting banks by being untruthful themselves. Lying supposes a punishment on all other banks, not just the misreporting one.

Therefore, we take another approach that is more efficient. We assume that the financial supervisory authority is able to detect with some probability that a bank or a set of banks have lied. Hence, the authority can identify exactly the misreporting bank which makes punishing efficient. The punishment should deter banks from lying in every period, and this incentive should make the banks believe that all the other banks report truthfully. If a bank is caught misreporting, the punishment is imposed by the authority and not by the subsequent actions of the other banks.

How can the financial supervisory authority detect deviations? Banks want to know about all of the loss events that occur within the organization. After recording and analyzing the event, banks try to fix the problem to avoid future losses related to it. Actually, it is not just profit-maximizing self-interest that prompts banks to do so, but it is also supervisory requirement. Consider Basel II Accord (BIS [2004]) and its legal implementation into the national legislations (eg. Capital Requirement Directive compulsory for the all the banks in the European Union). Under Basel II banks have to allocate capital for operational risk purposes. In case of simpler methods (Basic Indicator Approach or The Standardised Approach) the basis for capital requirements' calculation is gross income, meanwhile in case of the more sophisticated method (i.e Advanced Measurement Approach (AMA) the capital requirement expresses a value-at-risk measure with parameters of one-year time-horizon and 99.9 per cent confidence. Within this legal framework, the banks using AMA should apply accurate external data in their operational risk measurement scheme and supervisors should examine the validity of the data and even method of risk modelling. On the other hand, as supervisory authorities of financial institutions have the right to examine all the processes of banks, thus they may study these internal databases. If the same authorities have some kind of access to the data reported to the center, then the possibility of detecting a misreport exists.

Suppose that a bank can lie as long as it gets detected and suppose also that banks use a common discount factor, δ .

Defining strategies in the repeated setup is relatively easy, since banks cannot condition their action on the observed history, because other banks' actions are not observed. Note that it does not matter whether banks observe that the supervisor has caught a bank lying, because that bank will not play in the subsequent rounds. A bank's optimal action depends in each round on what it believes about the other banks' truthfulness in that round. Thus, there is no need to specify different strategies for any possible history. Strategies just specify whether to report truthfully or not in every period. Therefore, bank i 's

strategy is just $s_{i,t} = \{tr, l\}_t$ for $t = 0, 1, 2, \dots, \infty$.

The payoffs are the same as in the stage game. Bank i 's average discounted value of the payoff stream is:

$$U_i = \sum_{t=0}^{\infty} \delta^t u_i^t.$$

The best situation is obtained when each bank in every period reports truthfully its OR events. This is the case because although one misreporting bank may gain c when $f(\cdot)$ is recoverable, but - due to our assumption that the database is useful only if all reports are truthful - other banks' aggregate net loss is $(n-1)$ that is strictly greater than c given that $1 > c$.³

Definition 3 *The first best obtains if $s_{i,t} = \{tr\}$ for all $i \in N$ and $t = 0, 1, 2, \dots, \infty$.*

In the first best case $u_i^t = 1 - c$ in every period, so $U_i = \frac{1-c}{1-\delta}$.

3.1 The punishment

For simplicity, assume that in each period with probability q the authority has chance to detect if a bank reports truthfully or not. Therefore, q can be viewed as the probability that the supervisory authority examines the validity of the reported OR data. If a bank is found to have lied, then it gets a punishment. There are several ways of punishing a bank: a monetary fine, an exclusion from the consortium for z periods or forever, making it public that the bank lied causing a loss of reputation. In the tradition of repeated games, we consider the case of excluding the bank from participating in the consortium for z periods. We assume that the discounted present value of the foregone utility when lying is detected is P . Or in other words, the punishment of the bank is that it cannot benefit from having good data that enhances its OR management. Formally,

$$P = \sum_{t=0}^z \delta^t (1 - c) = \frac{(1 - \delta^z)(1 - c)}{1 - \delta}.$$

Given δ and c the punishment can be read as: how many periods of exclusion are necessary so that the discounted present value of the punishment equals P . Therefore, P and z are mutually determined.

3.2 Full information case - truthfulness is observable

As a benchmark case, first we consider what happens when $q = 1$, that is each bank's truthfulness is checked in each period. In this case, the payoff matrix in the stage game is the following:

³The net loss of an individual bank that is reporting truthfully if there is a misreporting bank is 1. It is the case, because instead of $1 - c$ the bank receives a payoff of $-c$. Since we suppose that all but one bank reports truthfully, the aggregate net loss amounts to $n - 1$.

bank $i \setminus$ bank $-i$	tr	l
tr	$1 - c, 1 - c$	$-c, 1 - P$
l	$1 - P, -c$	$-P, -P$

Our first result says that if the punishment is harsh enough, then in the unique equilibrium each bank reports truthfully.

Proposition 3 *When the truthfulness of banks' report is checked in each period, then all banks will be truthful if $c < P$.*

Proof. Clearly, if $P > c$, then the dominant strategy of bank i is to be truthful, independently of what the other banks do. Hence, being truthful is a dominant strategy for bank i . Note that we have the same situation from each bank's point of view. Thus, we can say that when punishment is high enough, each bank reports truthfully, so truthful reporting is the unique equilibrium. ■

It is easy to see that if the same condition holds, then each bank will report truthfully in our example of irrecoverable data. We have the following payoff matrix:

bank $i \setminus$ bank $-i$	tr	l
tr	$1 - c, 1 - c$	$-c, -P$
l	$-P, -c$	$-P, -P$

The same logic applies as in the previous case.

These results show that in the case of perfect monitoring a sufficiently high punishment is sufficient to make banks report truthfully. The question is whether the threat of punishment can achieve the same result when the probability of detection is less than 1.

3.3 Stochastic detection case

The previous case does not take into account, that if the probability of detection is small, then misreporting banks may accumulate more utility by lying which makes the discount factor an important part of the analysis. In this subsection, suppose that $q < 1$.

We take the perspective on an individual bank and consider when this bank finds optimal to misreport supposing that all the other banks report truthfully.

If the value computed by the center is recoverable, then given that the others report truthfully, a bank gains 1 instead of $1 - c$ by deviating. Since truthfulness is not observable, a bank can go on with lying until it gets detected. The expected number of periods between being checked is $\frac{1}{q}$, therefore the expected gain from lying is

$$L = \sum_{t=0}^{\frac{1}{q}} \delta^t = \frac{(1 - \delta^{\frac{1}{q}})}{1 - \delta}.$$

After being caught, a bank is excluded for z periods, which supposes a utility loss of P . After these z periods the bank benefits again from the external

data and if the consortium functions optimally, then the discounted utility of participating in it is $U = \frac{1-c}{1-\delta}$ from the date of readmission to the consortium.

A bank will deviate from the optimal equilibrium path (that is, it will lie) at the current period if and only if

$$L - \delta^{\frac{1}{q}}P + \delta^{(\frac{1}{q}+z)}U > U.$$

The expression can be simplified by collecting terms to

$$\frac{L - \delta^{\frac{1}{q}}P}{1 - \delta^{(\frac{1}{q}+z)}} > U.$$

The expression says that if the gains of deviation are high enough even after taking into account the utility loss and discounting, then banks may choose to deviate. The expression also shows that if the discount factor is high enough and/or the punishment is severe enough (or in other words, the length of exclusion is high enough), then banks will not have incentives to deviate.

Further algebraic manipulation yields that unilateral deviation is profitable if

$$c > \frac{2\delta^{\frac{1}{q}}(1 - \delta^z)}{1 + \delta^{\frac{1}{q}}(1 - 2\delta^z)}. \quad (3)$$

This expression reveals that if given δ and z the costs of truthful reporting are too high relative to the gains, then banks have incentives to misreport.

If we reverse the condition, then we have the following result.

Proposition 4 *If $c < \frac{2\delta^{\frac{1}{q}}(1-\delta^z)}{1+\delta^{\frac{1}{q}}(1-2\delta^z)}$ holds, then the first best obtains.*

Proof. The condition (3) ensures, that given the truthful reporting of other banks an individual bank does not find it optimal to misreport. Since banks are identical (except for the shock), all of them will choose to report truthfully. Therefore, the first best obtains. ■

Note that the proposition speaks about the punishment (that is P), which entails conditions on δ and z . The intuition of the proposition is straightforward and very much in the spirit of repeated games. Banks do not have incentives to misreport in face of harsh punishments.

These results could not be confirmed by robust empirical evidence at the moment. Nevertheless, anecdotal evidence concerning DIPO (Italian operational risk database) illustrates our result well. There, supervisors might already audited the quality of the data (crosschecking of data recorded in internal databases and data transmitted into the data consortium), and this clearly provides incentives for banks not to misreport loss event data.

For irrecoverable values, L is negative, so banks cannot gain from lying if the other banks report truthfully. As a consequence, the above condition is sufficient for that case as well.

3.4 A numerical example

Although real-life data on costs and discount factors are not available yet, we consider that a numerical example may give more intuition about the result we have obtained. We carry out an analysis that studies how the cost that equalizes both sides of (3) changes as we change other parameters (δ , q or z) of the model. That is, we want to calculate for a given set of parameters the maximum cost of reporting that makes a bank indifferent between being truthful or not. A big increase in this threshold cost can be interpreted as a big gain in utility which the bank derives from a well-functioning consortium. This is the case, since the threshold cost tells how big the cost (relative to the gains that we normalized to 1) should be to eliminate the benefits from reporting truthfully.

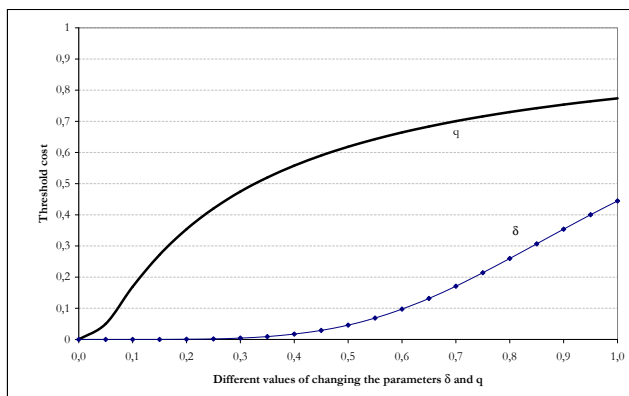


Chart 1: Threshold cost as δ and q changes

Chart 1 shows how the threshold changes when the discount factor and the probability of being examined change.⁴ The increase in the discount factor has a double effect. On the one hand, it increases the discounted value of lying. On the other hand, the punishment becomes also more severe and the opportunity cost of being truthful (U) is also higher. When banks do not value future payoffs (that is, for low values of δ), an increase in δ does not change the threshold cost. As banks assign higher weights to future payoffs, the benefits of reporting truthfully become greater than the benefits of lying. This leads to the result that banks will misreport only if the costs of reporting truthfully are too high.

As the probability of being examined by the supervisory authority increases, the expected benefit of lying decreases. Therefore, banks find optimal to misreport only if the costs of reporting are so high that the benefits of being truthful becomes too small.

The increases in the threshold cost shown in Chart 1 reveal that regarding q the biggest changes in the threshold cost (and consequently in gains from the

⁴When studying the effect of δ on the threshold cost we fix $q = 0.2$. When studying the effect of q on the threshold cost we hold δ constant at 0.9. In both cases we use $z = 2$.

consortium) accrue for relatively low values. The slope of the graph is the highest in that region. An interpretation is that to foster the proper performance of a data consortium the important thing is to give the impression that misreporting banks can be detected. It is not necessary that each bank be audited each year. Banks should have the impression that there is some positive probability of being detected if they misreport. Concerning the discount factor, the change in gains from the consortium are highest for high values. Thus, consortia work better where banks value sufficiently future payoffs. Intuitively, this happens when banks operate in a stable environment, interacting with the same players regularly.

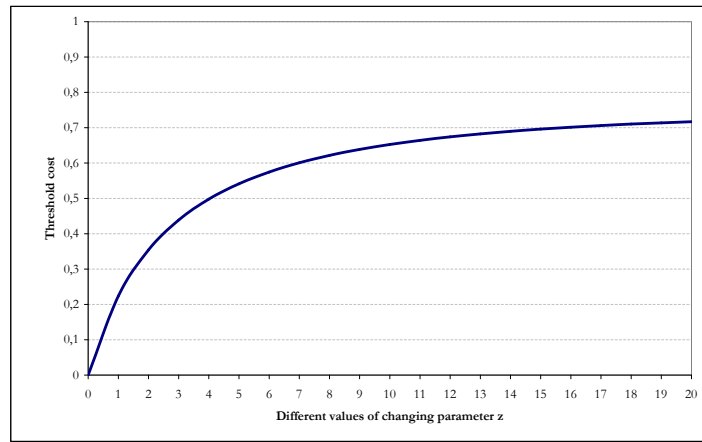


Chart 2: Threshold cost as z changes

Chart 2 illustrates the changes in the threshold cost as the number of periods that a misreporting bank is excluded from the consortium increases.⁵ As expected, a higher z dissuades banks from lying. This is the case, as higher z means a longer exclusion from the benefits of the consortium. In other words, a larger z means a higher punishment. Hence, banks find optimal to misreport only if the costs of reporting truthfully are relatively high.

⁵When studying the effect of z on the threshold cost, we fix $q = 0.2$ and $\delta = 0.9$.




Chart 3: Threshold cost as z and q change jointly

If we analyze jointly the increase in the length of exclusion (z) and the probability of examination (q), then we find that the joint effect is even larger, as depicted in Chart 3.⁶ Note that z and q are the parameters that can be set by the consortium and/or the supervisory authority. The appropriate choice of these parameters is crucial to achieve that banks provide true loss event data. The graph suggests that already for relatively low values of q a sufficiently long exclusion drives up considerably the threshold cost. Hence, if the consortium is designed in a way that reflects these features, then it can be expected to function properly.

3.5 Further issues

We supposed that banks follow pure strategies. It is natural to ask what if bank i follows a mixed strategy prescribing to report truthfully with probability π and lie with probability $1 - \pi$? The previous reasoning goes through, but this new assumption implies that the probability of being caught is less ($q(1 - \pi)$). As a consequence, the first best obtains if punishment is sufficiently severe.

Severe punishment may seem inefficient, because if exclusion is too long, other banks do not receive data on the excluded bank. Less data entail less efficient OR management for the member banks. Our analysis shows that it is not a problem. A severe and credible punishment avoids misreporting in our model. Hence, no bank will be punished, and there will be no loss of valuable data due to exclusion.

⁶In this chart we hold δ constant at 0.9.

4 Conclusion

In the last decade, operational risk management has become an important area in banking, as evidenced by the Basel II Accord. Paralelly, data consortia aiding the adequate OR management have grown in importance. To our best knowledge, we are the first to study from theoretical point of view the conditions under which the anonymous loss event reporting to the consortium is truthful. Since appropriate data quality is instrumental to the good risk management, the problem is highly relevant.

Viewing reporting as a one-shot game reveals the difficulties of arriving at the first best. A repeated-game setup provides a framework that identifies parameters that are crucial to the right design of a consortium. These parameters are the probability that the validity of loss events of a bank is checked and the punishment imposed in case of misreporting. Jointly, they determine the expected punishment upon caught lying. Since these parameters are chosen by the consortium, the appropriate choice can ensure the adequate performance of the consortium.

The main message of our analysis is that banks should perceive a high enough expected punishment if they do not comply with truthful reporting. A numerical example suggests that the combination of a relatively low probability of examination and a long exclusion is enough to induce truthful reporting as the unique equilibrium.

5 References

AUE, F.– KALBRENER, M. [2006]: LDA at work: Deutsche Bank’s approach to quantifying operational risk, *Journal of Operational Risk*, Vol. 1. No. 4. Winter, 2006/2007.

BAUD, Nicolas - FRACHOT, Antoine - RONCALLI, Thierry [2002]: Internal data, external data and consortium data - How to mix them for measuring operational risk, GRO Credit Lyonnais, France, link: <http://www.thierry-roncalli.com/download/oprisk-data-light-version.pdf>

BIS [2004]: International Convergence of Capital Measurement and Capital Standards: a Revised Framework, 26th June 2004

FUDENBERG, Drew - TIROLE, Jean [1991]: *Game Theory*, MIT Press

HOMOLYA, DANIEL - SZABOLCS, GERGELY [2008]: Operational risk data consortia and their applicability - HunOR, the opportunity of Hungarian banks (published in Hungarian), *Hitelintézeti Szemle* (Journal of the Hungarian Banking Association) 2008/1, link:

<http://www.bankszovetseg.hu/anyag/feltoltott/homolya0801.pdf>

HOMOLYA, DANIEL [2009]: The impact of the capital requirements for operational risk in the Hungarian banking system, *Magyar Nemzeti Bank* (Central Bank of Hungary) Bulletin, July 2009

NA, H.S. - MIRANDA, L. Couto - VAN DEN BERG, J. - LEIPOLDT, M.
[2005]: Data Scaling for Operational Risk Modelling, ERIM Report,
link: <http://repub.eur.nl/publications/index/471363749/>.
SMORODONSKY, Rann - TENNENHOLTZ, Moshe [2006]: Overcoming
free riding in multi-party computations - The anonymous case, Games and
Economic Behavior 55 (385-406)
VOIT, JOHANNES [2007]: How to create value from loss data pooling,in
Operational Risk 2.0, Riskbooks, 2007, Incisive Media
WOOD, DUNCAN [2007]: Strength in numbers - Oprisk