Irving Fisher Committee on
Central Bank Statistics

BANK FOR INTERNATIONAL SETTLEMENTS

IFC Satellite Seminar on "Post-crisis data landscape: micro data for the macro world", co-organised with the Central Bank of Malaysia and the European Central Bank

16 August 2019, Kuala Lumpur, Malaysia

Granular data and stress testing: stepping up to the challenge[1]

Cees Ullersma and Iman van Lelyveld,

Netherlands Bank

[1] This presentation was prepared for the meeting. The views expressed are those of the authors and do not necessarily reflect the views of the BIS, the IFC or the central banks and other institutions represented at the meeting.

Forthcoming in:

Handbook of Financial Stress Testing, editors Doyne Farmer, Til Schuermann, Alissa Kleinnijenhuis, Thom Wetzer. Cambridge University Press

Granular data offer new opportunities for stress testing [1]

Cees Ullersma and Iman van Lelyveld

Main messages

- Data is becoming increasingly more available: more granular, more timely and reported more efficiently. In combination with high quality metadata, this trend allows for connecting and viewing the datasets in new and more informative ways.
- Such developments are changing how policymakers and regulators can operationalise stress testing. It allows stress testing to move towards modelling the complete financial ecosystem. It also enables (macroprudential) stress testing with micro-foundations, taking into account the diversity of market participants. The impact of shocks can be analysed from a highly aggregated macro level all the way down to individual market transactions and participants.

Abstract

Nowadays, more and more granular data is being collected. The 2008-2010 crisis has shown that authorities were missing crucial information for accurately identifying risks in the financial system. This realization has led to a significant increase in the depth and scope of information being reported across the system. At the same time, the cost of reporting has gone down due to further digitalization.

In this paper we will examine how increased granularity allows authorities to put the data to better use and we will focus in particular on stress testing. Stress testing involves postulating severe scenario's that often have yet to materialise. More granular data combined flexibly from different sources allow for more accurate assessment of such adverse hypothetical states of the world. For instance, detailed information on interbank exposures allows for understanding bank to bank contagion much better.

We will briefly discuss the developments in the last decade and discuss what possibilities arise. The new granular data sets allow us to implement stress tests on various levels of aggregations using the very same data sets. Then we will turn to prerequisites in terms of for instance meta data allowing to connect data sets and data governance specifying how to treat the data. Finally, we will discuss remaining challenges, ranging from data quality up to high performance IT infrastructure and a company culture that embraces data.

[1] We are grateful to Roos van den Berg and Jiaqi Zheng for their assistance in constructing the figures and to three anonymous referees whose comments have improved the chapter significantly.

Introduction

Nowadays, more and more granular data is being collected. The 2008-2010 crisis has shown that authorities were missing crucial information for accurately identifying risks in the financial system. For example, due to the over-the-counter (OTC) nature of derivatives markets there is no centralised overview of the market.[2] Participant only observe their own volumes and exposure concentrations. The major US investment banks therefore did not realise that *jointly* they were massively exposed to a single entity, the lightly regulated insurer AIG. In setting their capital buffers and implementing other risk mitigating procedures they were thus ignoring an important yet unobserved concentration risk.

The realization that more information is needed to properly capture the risk in the financial system, has led to a significant increase in the depth and scope of information being reported across the system.[3] At the same time, the cost of reporting has gone down because more and more economic interactions are recorded digitally. Also, as a side effect of internet based activity (i.e., commercial and social networks) new aspects of economic activity have become measurable.[4] Furthermore, the cost of analysis has gone down as well as (open source) tools and methods (e.g., machine learning (ML) and artificial intelligence (AI)) are improving markedly. Not only is more data available for stress testing than ever before, implementing more accurate models is within reach.

The spectacular growth in the granularity of data collected in the last century can be illustrated by comparing the level of detail available in Foster (1922) with that in Ehlers (2018), shown in Figures 1 and 2, respectively. Arguably Foster did not intend to depict each individual transaction contained in the flows but the data was also simply not there. For the flows shown in Figure 2 the disaggregated information is in principle available. Combining this information is another matter which we will discuss in detail below.

[2] See Abad *et al* (2016) for a first effort to shed light on these 'dark markets' and Levels *et al* (2018) for further detail.

[3] See for example the second progress report on the implementation of Phase 2 of the G20 Data Gaps Initiative (FSB and IMF (2017)).

[4] See for example Graff (2018) who shows Brexit related worker migration in and out of the UK in a comprehensive and timely way using LinkedIn profiles.

Figure 1 Granularity 100 years ago: the circuit flow of money

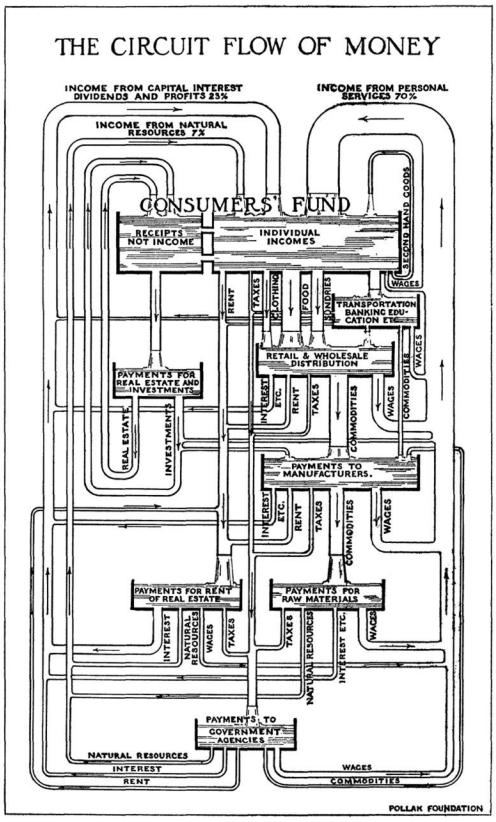

Source: Foster (1922)

Figure 2 Granularity today: the Chinese financial system

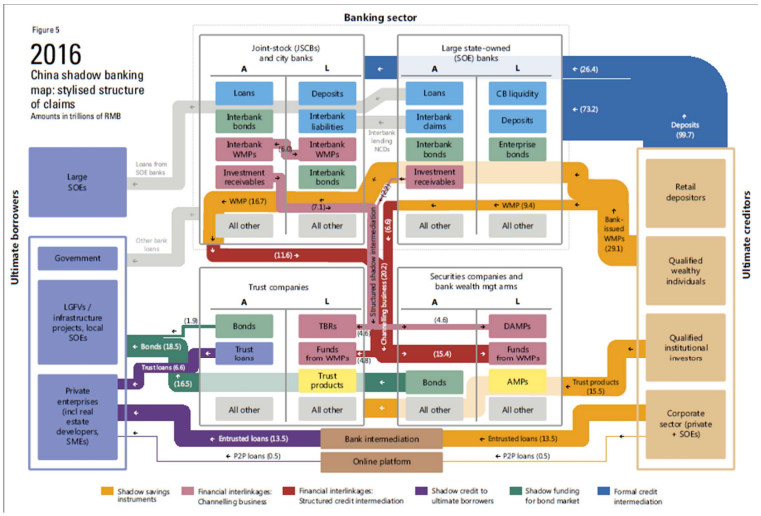

Source: Ehlers *et al* (2018)

Data collection and analysis has been a core competency for most central banks and supervisors for many years. Central banks thus have a deep understanding of how to collect data from – primarily financial – economic agents and are thus well placed to capitalise on the current data trend. At the same time, they are faced with legacy systems that might lock them in inefficient solutions. Some authorities are looking for alternative ways to acquire data. In Austria, for example, direct reporting to the central bank has been replaced by

reporting to a joint venture – owned by the bigger banks – tasked with data handling. This joint venture, in turn, provides the necessary reporting to the authorities (Kienecker (2018)). An alternative is plugging into the systems of private institutions directly. In theory, this would allow authorities to extract the precise data points at precisely the right time without imposing material reporting burden. The practice is more elusive because firms have multiple IT systems with unaligned definitions. For regulatory reporting, significant manual effort is required. Direct access to firms' data will not solve this but will push the necessary manual steps to the authorities.

The breadth of reporting has increased especially in the area of Over-the-Counter (OTC) markets. In OTC markets, such as the interbank market or the derivatives markets, market conditions are more difficult to gauge since contracts are bilateral and generally private information. Recently regulators have introduced new regulation to – at the minimum – introduce exposure reporting (without imposing limits or other constraints). The European Market Infrastructure Regulation (EMIR) is an example of such a regulation covering derivative markets trading. Under this regulation all parties with significant dealings need to report their transactions centrally.[5] These central repositories are for-profit entities that collect the data and disseminate it to the relevant authorities. Another important initiative is AnaCredit which has become operational in the fall of 2018 and will provide a basic version of a European credit register. This will allow authorities to more comprehensively assess credit risk. Furthermore, AnaCredit data can be used to trace out the effects of monetary policy. A final example is the Money Market Statistical Reporting (MMSR) covering large banks and their counterparties in the European money market. It contains deal level information of various short-dated money market contracts (less than a year) and is set up to construct daily interest rate benchmark. An important side effect is that the MMSR data can be used to shed light on many other aspects of money market functioning. Relevant question could, for example, be how systemic or idiosyncratic stress affects the system.

Many data collections have become more granular allowing for made-to-measure aggregations. Whereas traditionally the task of classifying a particular data point was put upon the reporter, a current trend is to request information on a much more granular level and leave the classification to the authorities. For example, reporting firms no longer need to determine the sector of the counterparty but only need to supply the Legal Entity Identifier (LEI). The LEI, which will be discussed in more detail below, is a unique identification number for legal entities. Using the LEI, data can for instance be aggregated to a financial group. The LEI can also be used to generate sector aggregates. By leaving the aggregation process up to the data collector (i.e., the statistical agency, the supervisor or the central bank), consistent handling of the sector allocation is ensured. Furthermore, a change in the sector assignment of a particular firm can be applied consistently across all reported data.

[5] The Dodd-Frank Act in the US has similar stipulations.

A major advantage of data on a granular level is that it allows us to pinpoint the source of stress and the location of the most vulnerable parts of the financial system (Kwast *et al* (2010)). More granularity, however, also implies massive datasets and thus creates computational challenges. Moreover, if we are interested in feedback effects, we need to be able to trace such effects on a sufficiently high frequency. If the time scale is too coarse, we will only be able to see end results once they have already materialised. For example, if overlapping portfolios are an important driver of contagion because they open up a channel to transmit stress, then we need timely information of securities holdings. However, holdings are often, if collected at all on a granular level, only collected on a monthly or quarterly basis (see for instance mutual fund data in the US or the Securities Holdings Statistics in Europe). In the period in between, positions can move in many ways and in coming to stress test we need to make assumptions about how the portfolios change (Wang *et al* (2018)). In sum, examining the possibilities and challenges that new data and tooling bring in the stress test arena is useful, arguably not just for stress testing but for risk identification and management more broadly.

In stress testing many of these new developments come together. Stress tests often require the ad-hoc combination of data in order to be able to flexibly model unlikely scenarios. In the past, data availability dictated which scenarios were feasible to be stressed. For example, interbank exposures were only available as an aggregate without a split by counterpart. It was therefore very difficult to construct linkages from bank to bank. Without an accurate representation of the interbank network, tracing out the transmission of stress through the system is not possible.[6] Stress test were thus restricted to investigate how aggregate risk measures reacted to stress. Furthermore, modelling individual actors was often not possible since available information was not detailed enough to map out utility functions or supply and demand curves. Increased data availability might thus warrant a re-think of more traditional stress testing. A very promising avenue in this respect are heterogeneous agent models with network effects. These models are now computationally feasible and ready to deliver useful insights (Aymanns *et al* (2018)).

To be able to optimally conduct stress testing in this new environment, a number of challenges need to be resolved. First, our data should be of sufficient quality. In particular, the associated metadata should allow us to connect and combine datasets as needed. Second, a conceptual framework on how to handle the data needs to be adopted. Third, we need to develop the capability to tackle complex stress tests. This involves both an adequate IT infrastructure as well as a data-friendly company culture. We will discuss each of these points in turn. Then we will discuss a case study with the Securities Holdings Statistics (SHS) to show how granular data can provide the raw material for micro *and* macro stress tests.

Collecting data: national challenges and opportunities

Typically reporting requirements are developed in a `silo'. These silos are defined by the country specific legal frameworks that govern them. Mandates can be statistical, conduct-of-business (including market conduct), as well as micro- and macroprudential. Statistical

[6] See van Lelyveld and Liedorp (2006) and van Lelyveld *et al* (2011) for early efforts to construct network stress tests for the interbank and reinsurance markets, respectively.

mandates have the longest history. The English King William the Conqueror already ordered a stock-take of productive assets to determine the tax base as long ago as 1085. Market conduct mandates require data with the aim to ensure fair and efficient trading. Microprudential reporting is primarily aimed at collecting data to assess an individual firm while macroprudential supervision takes wider effects into account. For the latter type of supervision we are thus naturally more interested in data that can shed light on the interaction of different financial actors and the effects of financial markets on the wider economy. Each mandate has its own focus and incentives to coordinate across regulatory domains are limited.

Depending on a country's institutional framework, the legal mandates are assigned to different agencies potentially making sharing data difficult. Given the often complex governance, it is no wonder that reporting requirements are a patchwork of over- and underlapping components. Moreover, sharing data among agencies is often difficult because existing laws prohibit sharing confidential data. Even in the US, a country with a long tradition of publishing a wide range of financial information, this seems to be the case. This point was reiterated by Ruth Judson (Federal Reserve Board) at a recent G20 workshop on data sharing (Judson (2017)). She noted that "The U.S. has problems similar to those of other countries with regard to sharing data across agencies and countries".

Collecting data: cross-border challenges

Reporting and sharing data across national jurisdictions is even more complicated than in a specific jurisdiction. For example, the Basel Accords are a comprehensive set of rules for banks aiming to create a level playing field globally. The accords are themselves not legally binding, however. Basel committee member countries have to put the accord in national legislation first. For the current accord, Basel III, this is in Europe first and foremost the so-called CRD-IV package. CRD IV contains reporting requirements for licensed entities, that is banks. There are generally very few reporting requirements for firms outside the prudential regulatory perimeter. Furthermore, there is little coordination of reporting requirements for, for example, insurance undertakings and banks.

In the field of statistics, various bodies aim to coordinate regulation and the collection of (regulatory) statistics across jurisdictions. Standards are set by international institutions such as the UN, OECD, IMF and Worldbank. Agreements reached are then either implemented on a best-efforts basis or cemented in national law.

Collecting data: new opportunities through meta data and code distribution

A consequence of how reporting requirements are typically designed, is that meta data is not well coordinated. Arguably, the precise reporting need to achieve their mandate is best determined by the relevant authorities. For example, microprudential banking supervisors are best placed to determine what information is needed to safeguard a bank. They should thus decide on what data points to request and at what frequency. Macroprudential supervisors might need other information because the scope of their mandate is wider. The same holds for other supervisors (e.g., insurance sector or financial markets supervisors).

Standardization across various reporting frameworks can be beneficial. Statistical agencies and standard setters have agreed on a long list of standards to ensure consistent measurement of social phenomena. For instance, there are International Organisation for Standardization (ISO) standards for country naming conventions. In financial markets, other conventions have been agreed on such as the International Securities Identification Number (ISIN, ISO 6166) and the Committee on Uniform Security Identification Procedures identification number (CUSIP) used in the US and Canada. These conventions do not yet have a complete coverage in terms of products or global acceptance but good progress is being made. For example, if regional breakdowns are reported using standard ISO country codes then we can easily combine information from, say, pension fund exposures to mutual fund investments in particular countries or regions. More complex issues arise if we want to unequivocally define economic activity across sectors: what constitutes a sector? What are the defining characteristics of a loan? If different sectors agree to use the same meta data frameworks, combining and validating data across different sectors can be simplified significantly. The choices in one domain can thus have an external effect and affect the usefulness and efficiency of reporting elsewhere. If for instance banking supervisors agree to universally adopt the LEI, other financial supervisors can follow and then matching exposures of banks with those of, for example, insurers becomes much easier. Berner and Judge (2019) support this line of thought and argue forcefully that developing standards jointly with industry increases social welfare.

An example of a key identifier in metadata is the Legal Entity Identifier (LEI).[7] The advantages of the LEI are discussed in more detail in Bottega and Powell (2011). The LEI is a 20-digit, alpha-numeric code based on a standard (i.e., ISO 17442). It connects to reference information that enables clear and unique identification of legal entities participating in financial transactions. The LEI initiative is a joint effort initiated by the Group of 20, the FSB and many regulators around the world. They have emphasised the need to make the LEI a broad public good. The LEI Regulatory Oversight Committee is at the highest level of LEI governance. Below this is the Global LEI Foundation, which brings together the information curated by local LEI granting organisations. These are typically credit registries and chambers of commerce. The number of LEIs is growing rapidly, fostered by new regulatory requirements, and now exceeds 1 million active identifiers (LEI ROC (2018)). The ECB estimates that the LEI covers at least securities with a total value of €95 trillion worldwide as of November 2017. This is a 25% increase since the end of January 2017. Although much has improved recently, some major challenges remain. For one, as shown in Figure 3, the coverage is far from even nor complete.

[7] Other initiatives in the ESCB to improve data collection are the Banks' Integrated Reporting Dictionary (BIRD), the Integrated Reporting Framework (IReF), the Register of Institutes and Affiliates Data (RIAD), and the Centralised Securities Data Base (CSDB).

Figure 3 Number of Legal Entity Identifiers (LEI)

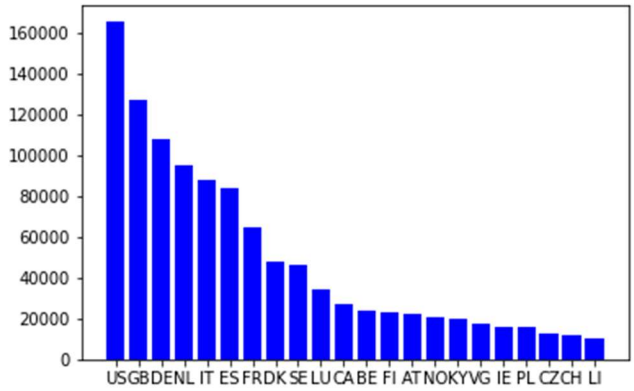

Source: GLEIF (www.gleif.org). Parsing by Martien Lubberink

More technical challenges remain unresolved as well: for example, if a LEI lapses because the fees have not been paid, then a new LEI needs to be applied for. The entity will now have both an old inactive LEI as well as a new LEI. The old LEI might linger in reporting data for a long time if the entity entered into long running commitment such as a long-dated swap.

The LEI makes it possible to start assigning activity and risks exactly to those subsidiaries that are responsible. For example, for trading and holding of securities it is clear to whom to assign the flow and stock of securities. With multiple jurisdictions this allocation is of critical importance (Cerutti and Schmieder (2012), Fang and van Lelyveld (2014)). A poignant example is the default of Lehman where US authorities moved assets from Lehman's British subsidiaries before British authorities could react. Without information on the location of an asset (and how fungible it is) it is difficult to assess the robustness of a (local) institution.

To clarify how such allocation could work, we plot the legal structure of Rabobank in Figure 4. The information is taken from the GLEIF website and shows the structure and location of all the legal entities with LEIs. The different continents are coloured as explained in the legend. It is clear that even a relatively small player like Rabobank already has a very complicated structure spanning multiple jurisdictions. Note that this is only a partial picture since not all the units have a LEI attached to them. Given this structure, authorities in different jurisdictions can allocate activities to these units and assess whether these activities pose a local or global threat to financial stability.

Being able to break down a reporting entity into its constituent parts is very useful in stress testing because it allows for pinpointing where and how stress would affect a firm (Cerutti and Schmieder (2012)). If a bank for instance extends mortgages in two subsidiaries in two distinct geographic areas with quite different business cycles, then being able to stress the mortgage book using two different cycles can be quite valuable. Alternatively, if different resolution schemes are in force for different parts of a bank, claims and liabilities used in the stress test need to be broken down by jurisdiction (cf. Joint Forum (2012), in particular on the role of internal guarantees). Although a unique identifier such as the LEI is a necessary first step, we currently do not yet receive all the necessary information on a

sufficiently disaggregated basis to allow for ad-hoc choice of aggregation. That is, a stress test of a bank with the option that some of the subsidiaries are allowed to default still requires extensive data work. In principle, however, we can construct a reasonably precise picture of a large and complex internationally active bank including a breakdown for the significant subsidiaries.

Figure 4 Legal structure of Rabobank

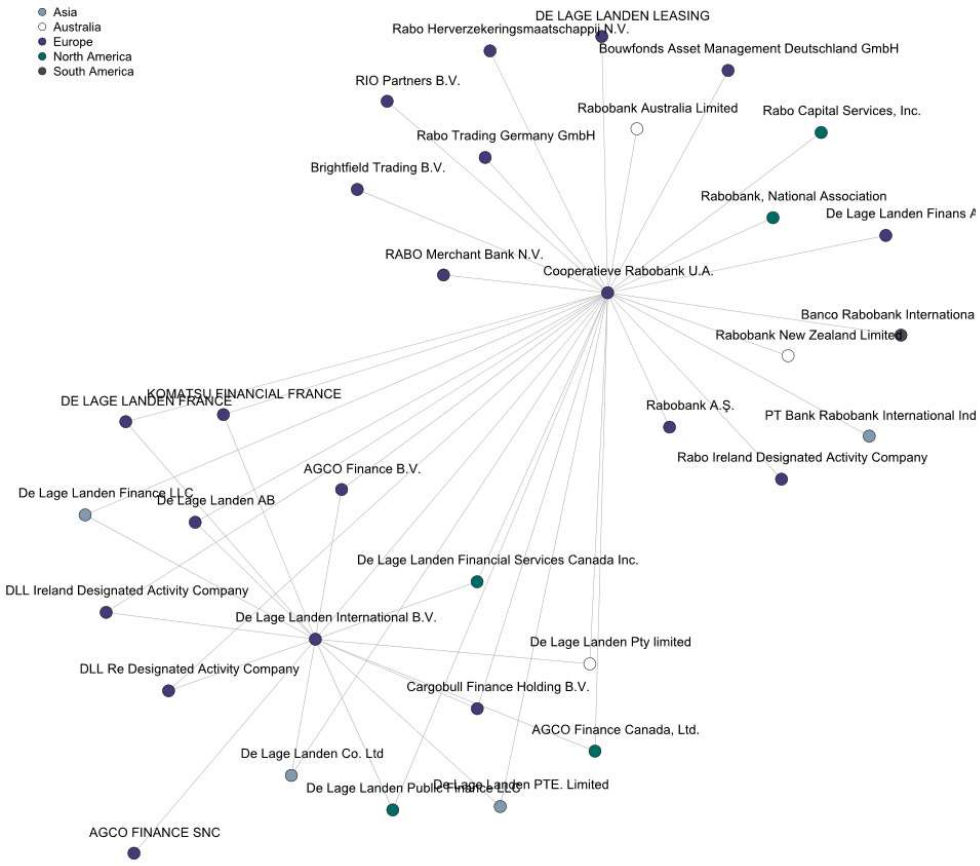

Ensuring a consistent reporting framework is quite a challenge although some progress is made. For example, the European Commission has undertaken an extensive study to assess whether the current reporting requirements are sufficiently harmonised (EC (2019)). Overall, the requirements are effective in meeting the objective of enabling supervisory authorities to fulfil their statutory tasks and mandates. While the EC concludes that EU-level requirements are coherent in a broad sense, they are found to have a range of inconsistencies both across as well as within reporting frameworks. Many of these inconsistencies appear minor and purely technical in nature. They nonetheless put a burden on reporting entities and supervisors. Key inconsistencies identified include non-aligned definitions, different formats of data fields or templates and inconsistent timing of mandated reporting. Note that EC's assessment did not back up the repeated claim by stakeholders that there are a significant number of duplicate requirements between different reporting frameworks. The EC did identify numerous cases of broadly or very similar data being requested. Authorities could thus consider if their information need merits multiple reports.

As an example of the mundane but in practice quite real complications that arise in combining information we can examine the differences for selected fields of the EMIR, MiFIR, REMIT and SFTR requirements as shown in Figure 5.[8] The data fields shown are present in all four frameworks but carry different labels. This could in principle be resolved by making a mapping based on the definitions. However, in some cases the definitions (or, equivalently, the admissible values) differ. For example, the data fields in the second row exhibit a validation inconsistency. Under EMIR, the status of the reported LEI shall be "Issued", "Pending transfer" or "Pending archival" while MiFIR reports are also accepted if the status of the LEI is '"Lapsed". For long-running swaps this might increasingly become an issue because lapsed LEIs would then generate a missing value in the EMIR data – but not in MiFIR. Likewise, there is a partial inconsistency in terms of content of the data fields in rows 2 and 3. Under REMIT, the reporting ID is not required to be an LEI308. Similarly, REMIT accepts a MIC, LEI, ACER code or ' XBIL' (for bilateral trades) as 'organised market place' while EMIR, MiFIR and SFTR always require this field to be populated with a MIC code. Linking REMIT to the other data sets is thus not straightforward.

Figure 5 Comparison of selected fields in EMIR, MiFIR, REMIT, and SFTR

EMIR	MiFIR	REMIT	SFTR
Field Identifier	Field Identifier	Field Identifier	Field Identifier
Currency of price	Price Currency	Price currency	Price currency
Report submitting entity ID	Submitting entity identification code	Reporting entity ID	Report submitting entity
Venue of execution	Venue	Organised market place	Trading venue

Source: EC (2019).

To overcome restrictions on sharing data, which could be used for more comprehensive stress tests, different approaches have been suggested – and in some cases executed as well. One way is to distribute code and only share the outcomes. Although the granular underlying data might be sensitive, the model outcomes – such as for instance regression coefficients – generally do not reveal the information encoded in the individual observations. In a project undertaken by the Basel Committee on Banking Supervision Working Group on Liquidity Stress Testing, participating researchers ran the code on their respective financial network data and then fed back the results to a shared repository. This resulted in an overview of network characteristics and network vulnerabilities for a wide range of networks (Anand *et al* (2018)).[9] A similar approach of code and result sharing is

[8] EMIR is a regulation on OTC derivatives, central counterparties and trade repositories, MiFIR covers reporting of transactions in financial instruments which are traded on a trading venues, REMIT deals with wholesale energy market integrity and transparency, and, the SFTR provides securities financing transactions reporting. See EC (2019) and the references and links therein for further details.

[9] The code for this project is available at https://github.com/imanvl/RTF_NTW_Horse. This approach is taken much further in the biomedical sciences where similar security and storage concerns are at play. See for instance the Global Alliance for Genomics and Health (https://www.ga4gh.org).

followed by the International Banking Research Network.[10] An alternative is to assign a trusted party who handles data collection and analysis.[11] For example, supervisors of the largest banks in the world have drawn up a legal framework that governs the exchange of very granular, bank specific information. This data is now being collected and analysed in the International Data Hub (IDH) hosted by the Bank for International Settlements.[12] The detailed reports are then sent back to the data contributors. Over the years the IDH has become more engaged in other regulatory workgroups, contributing to for instance analysis on the importance of CCPS, but external publication is still out of the question.

Data collection can have external effects. Although reporting requirements are almost always drafted with a single purpose in mind, the data collected can have multiple uses. For example, loan-level information on bank loans used in banks' internal credit risk models is generally primarily collected for microprudential supervision. With such data, supervisors can assess the validity and robustness of a bank's credit risk model. The data can also be used for other micro prudential purposes such as stress testing an individual bank. However, if such data is collected consistently *across* banks, then such data becomes potentially useful for macroprudential supervisors as well. The very same data can then also be used as input for an aggregate stress test. The aggregate stress test could in this case be micro-founded with attention for the distribution in credit risk and move beyond just examining credit aggregates. Alternatively, credit registry data – primarily collected for market transparency, statistical and monetary policy purposes – could be used for micro- and macroprudential stress tests. At DNB we are now investigating to what extent the information in the granular loan tapes, which we request ad-hoc in on-site microprudential examinations, overlaps with information in AnaCredit, the ESCB's credit registry. The advantage of AnaCredit is timeliness and cost. For AnaCredit, banks have put in place a regular reporting schedule which reduces the need for manual steps and hence costs. The disadvantage is that the definitions in place are not aligned with the supervisory definitions needed to assess compliance with, for example, Basel rules. At the very least, the AnaCredit information can provide a timely proxy for bank's credit risk.

To 'recycle' data in other domains might dictate additional requirements that do not seem immediately useful to the reporting parties. For example, the definition of what precisely constitutes a `bad' loan might currently be left up to a reporting bank (as the definition ties into internal control procedures) and the definition of when a loan is past due is then firm specific. This firm-specific choice might be determined by accounting rules, regulation or legacy IT systems. For understanding a bank's credit risk, the precise definition of past due is less important than that this definition is consistent over time and that the data is recorded accurately. To make the data useful for macroprudential analysis, however, it is important to use a single definition applied consistently across firms. Furthermore, if data is to be used for cross-country analyses then these definitions need to be agreed upon internationally as

[10] See https://www.newyorkfed.org/ibrn for more information.
[11] Yet another alternative is to encrypt the data in such a way that relevant information can be shared without revealing individual reporter's data. See Flood *et al* (2013).
[12] The IDH is part of the Financial Stability Board's Data Gaps initiative. See FSB (2009) for more details.

well.[13]

The ECB has undertaken a pilot to see how integration of highly granular data can be accomplished (Lauro en Traverso (2018)). The use case of this pilot was the analysis of the impact of the QE programme on banks' balance sheets on both the asset as well as the liability side. A key finding is that integrating two separate sets requires even higher data quality standards of the contributing sets than in the case of separate reporting; the data needs to be not just internally consistent but also match across reporting frameworks. Authorities thus need to be "prepared for integration" by putting in place integrated data management policies and adhering to best practices. Barbic *et al* (2017) provide detail on how, based on microprudential information, the ESCB has developed a data set with consolidated banking data which can be used for macroprudential analysis.

Data governance

To reap the benefits associated with combining different datasets on a regular basis we first need a conceptual framework of how we handle data. One possible starting point is the Four Quadrant Model we employ at DNB (Figure 6). In this model, the quadrants are delineated on two dimensions. First, on the y-axis we plot whether data is produced on demand (pull) or as part of a regular production cycle (push). An example of the former is the construction of a dataset for a dashboard built by a business analyst. The analyst could for instance be interested in a holistic view of the risk of a bank and would like to see risk metrics coming from different reports all presented in a single view. In this case, the analyst is actively asking for the data. An example of the pushed data is the regular prudential reporting process. The collection of such data is mandated by law and such data will be collected without the analyst actively asking for it. This is not to say the analyst is not interested, merely that he or she does not have to instigate the process. The second dimension, delineated on the x-axis, is whether the data collection is a structured, systematic process or an ad-hoc one. For instance, for producing macro-economic statistics, well-developed processes are in place which have been honed over the years. For ad-hoc research projects, in contrast, entirely new data sets or even methods to collect data need to be developed.

The combination of these two delineating dimensions results in four quadrants for which we require different data governance. To ensure that data in Quadrant I remains highly structured, we need to apply clear definitions and maintain high data quality. Generally, we define a Logical Data Model (LDM). An LDM is an exact definition of a data set. It defines aspects such as the number of variables, the storage type(s) and a description of what the variable supposedly measures. It also describes the relationships between variables. If reported data does not pass the data quality tests dictated by the LDM then one or more resubmissions are required until the data is of sufficient quality to be admitted to the production database. Quadrant 1 also requires an audit trail of data point changes otherwise the integrity of the data cannot be guaranteed. For example, analysist could – unintentionally or otherwise – change important benchmark interest rates or inflation rates.

[13] Initiatives such as the development of the Financial Industry Business Ontology (FIBO) are very helpful. FIBO is an open source business conceptual ontology of how financial instruments, business entities and financial processes work across the global financial industry (See https://edmcouncil.org/page/fiboproductsaccessre).

Quadrant II holds those (intermediate) datasets that users can, for instance, visualise in a dashboard. As discussed above, the analyst actively ask for data and are therefore also more involved in ensuring the quality of the data. Jumping ahead, Quadrant III covers those sets that require no centralised governance because they are not meant to be put onto the DNB data platform. Datasets in this quadrant are public or are local ad-hoc data collections. The ownership, and hence the responsibility of properly maintaining the data, lies elsewhere (for public data) or with the local user.

For stress testing, Quadrant IV is the most important quadrant. For the most part stress tests are bespoke exercises focusing on particular risks or scenarios. Although there might be a legal basis requiring periodic stress tests, the details are typically not spelled out. This implies the lower half of the y-axis of Figure 6. In some countries stress tests are becoming an annual fixture and hence the data requests might fit better in a Quadrant I process. However, to our knowledge, this has not happened anywhere. In terms of data governance, the best place for stress test analyses is Quadrant IV. Here the analyst can bring together various data sources usually required in stress testing. Since stress test scenario's change regularly, the input required will also differ each time.

Figure 6 The Four Quadrant Model for data governance

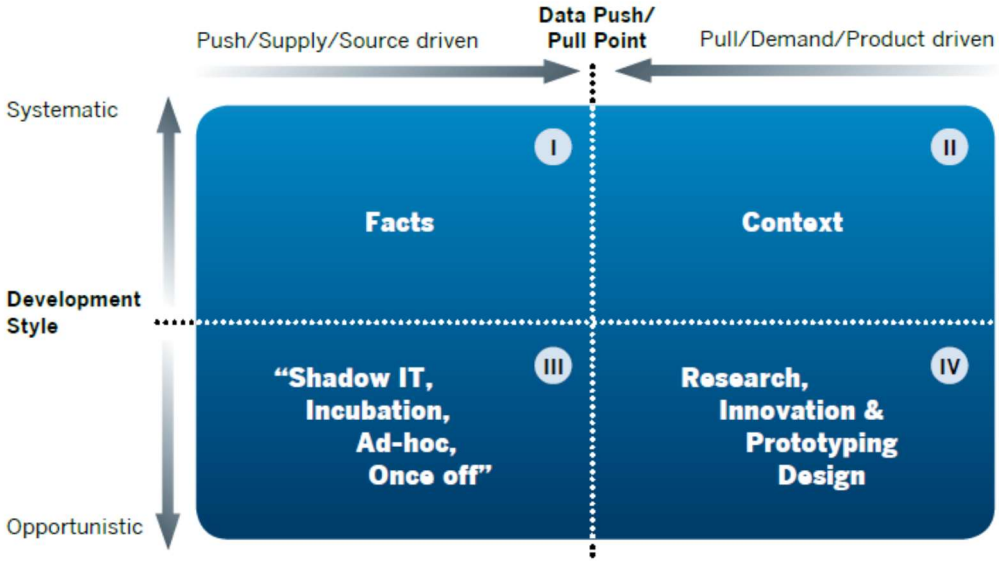

The choice of where to place a particular data set in the quadrant model depends on various factors. For example, the *importance* of the analysis using the data determines how much effort should be put into getting the collection process to the very highest level. With new data sources this presents a problem because the universe of possible useful analyses is not immediately apparent to management while the costs are real and immediate. *Repeated usage* might merit further investment in turning the reporting into a process. Another key factors is *available budget*. Finally, there is some value in the *flexibility required*. Tying yourself to the mast with a very prescriptive data intake process also implies that deviations from the process are costly. A deviation could for instance be requesting an accelerated delivery with lower data quality standards because a crisis is unfolding and authorities value noisy but timely information over perfect but later information. In such data files, some checks or required resubmissions might have been foregone. The

information will be available faster but will also be noisier.

A major advantage of promoting a data set to Quadrant I is that data quality is unambiguous. This makes, for instance, combining data straightforward; since identifying keys are clear, no time is lost in matching data points from different sources. Furthermore, issues with missing or tainted should have been solved. It is up to the data owner to weigh the cost of promoting the data set to Quadrant I/II against the benefits of ease of use.

The challenge is to let the governance model create value in stress testing. As noted, stress testing often combines highly structured data (e.g., prudential CALL reports in the US or FINREP and COREP in the EU) with ad-hoc data sources (e.g., banks' proprietary model parameters). Moreover, if the stress test covers multiple sectors then information from different types of institutions (e.g. banks, insurers, and pension funds) needs to be shared by supervisors with different mandates. From Quadrant IV, the access to the high-quality structured data, which resides in Quadrant I, should be without impediments for those with the right access-rights. Given that the structure is clear, direct SQL-querying of databases can be implemented without impinging on confidentiality. Such direct access will reduce data wrangling costs, reduce errors and potentially make it possible to deliver stress tests results in a more timely fashion.

Note that the Quadrant model does not dictate the physical set up but primarily describes the governance of data. Traditionally, the physical setup (i.e., the hardware and the servers deployed on it) was tightly linked to the governance. Some servers were maintained in production with very tight governance while others were research and development servers with a much looser regime. With the appropriate access right procedures and rights maintenance software this tight coupling can become looser – applications and data can be anywhere because their use can be controlled.
The questions to ask about data governance can thus be decoupled from the physical set-up: What can we expect of the quality of the data? How important is traceability of data mutations? Who is allowed access? In practice, the Quadrant model still does map to different types of workspaces. At DNB, we use primarily the tooling for our large databases (SAS), a structured Business Intelligence (BI) process and strict governance for Quadrant I and II. For Quadrant IV we have a Research Area Network and a high-performance cloud environment where analysts have the freedom to deploy a much wider range of (open source) tooling. With this comes the responsibility to maintain this tooling and vouch for the solidity of the analysis.

In an ideal world all data would come with a completely fleshed out Logical Data Model (LDM). An LDM defines the structure expected in the data. It will stipulate data-types (e.g., string or float) and relations between data-points (for example, total assets should be equal to total liabilities). If the data complies with the LDM a complete description of the data is ensured. Such a description can for instance be used to ensure consistency across jurisdictions and communicate with reporting entities. It also points out exceptions but does not in itself provide the `business rules' that need to be applied to find practical solutions to reporting errors. For example, if the LDM prescribes using two-digit ISO country codes then a sensible business rule might be to also allow valid three-digit ISO codes even though strictly speaking these are a violation of the LDM. In the cleaning process a mapping from

the two-digit to the three-digit ISO codes could be applied.

In practice, we define LDMs for a limited set of reports. Defining an LDM is costly because it entails fleshing out all the possible admissible permutations in a reporting framework. This is not trivial since many relevant regulations are ambiguous in their definitions (since more often than not they are the result of negotiations and the compromise has to reflect multiple views). Especially with ad-hoc requests there is insufficient time to work out a fully-fledged LDM. Moreover, it might not be cost-efficient to define a complete LDM for a single-shot data request.

Sometimes the LDM approach might seem overbearing and over-engineering for a single report. For instance, if we are designing an application to capture the contact information at supervised institutions (i.e., a mail address or phone number) it might seem overkill to work with an LDM. However, this contact-information might feed into a larger system that in the end is meant to be shared through an internal or a public register. If the data source is not captured unequivocally, then the register data will be tainted. In practice users across the institution will start to keep track of contact information leading to inconsistent and incomplete information. The costs materialise in the relatively simple address application while the benefits only materialise later for the organisation as a whole. The same holds for more important information such as the legal structure of supervised firms or the identifiers of financial instruments. Fairly attributing the costs of a local solution that is 'too much' at the local level but has institution-wide benefits is a significant management challenge.

The challenge is to find a balance between LDM fetishism and `anything goes' reporting. Although adhering to an LDM ensures that the data we admit to the final database meets a specified standard, there can be good reasons to diverge. For instance, if in a crisis situation a new datasource needs to be available on short notice, then defining an LDM will take too long. Alternatively, for small, one-off projects the costs of definition might create too much overhead. Nevertheless, the concept of an LDM is worthy to be applied more widely. The adoption of LDMs is however hampered because amongst policymakers – mostly economists and lawyers by training – the concept is relatively unknown. Unfortunately, LDMs are thus generally an afterthought to reporting requirements that have been the subject of long running negotiations. The structure and the content of the templates is then already set. Ambiguities that a rigorous approach such as an LDM definition process bring to the surface are then difficult to address. It would therefore be helpful if LDM-like thinking would be embraced from very early on in the reporting design. Furthermore, the external benefits which materialise only later and elsewhere in the organisation should be incorporated in weighing the pros and cons. One way to accomplish this is subsidise drawing up the LDM.

Acceptance of LDM-like thinking is hampered by, on the one hand, something we could term LDM-fetishism and, on the other hand, the typical way of how LDMs are presented to the non-initiated. LDM-fetishism is sticking to the model no matter what. This is not helpful as policy makers and researchers generally prefer noisy information over no information at all. In some, if not most cases, a greatly simplified LDM is more than adequate. The second reason LDMs are not widely embraced is because of the way they are presented. The typical end user is used to thinking in terms of a template or an Excel sheet. The LDM, however,

reduces these Excel sheets to their essence and strips common elements to be defined in a separate place. The logical structure is then presented graphically. These wall sized graphics, printed in tiny fonts and to the uninitiated unintelligible symbols, are difficult to digest for most supervisors or policy officers. Modellers should put more effort in presenting the LDM in ways that appeal to the end user and that make clear how they add value.

Developing a stress testing capability

Once all these governance challenges, technical difficulties and conceptual issues have been resolved, we will be able to paint a complete picture of the financial ecosystem from all angles.[14] This requires a more or less comprehensive coverage of all relevant financial sectors. For some sectors, further detail is required to be able to identify systemically important entities. For example, data on lightly or non-regulated sectors is generally more sparse and if financial activity in these sectors picks up and starts to become a factor influencing systemic risk, then new ways of unearthing this information need to be found. For instance, although currently fintech firms' activity is only a fraction of the established financial sector (Frost (2020)), this might change in the future. For other sectors, a coarser aggregate can be sufficient. For instance, if contracts are linear then stress test effects are less susceptible to tail scenarios and we thus have less need to the full distribution of exposures. It also requires that we collect stocks of all relevant balance sheet positions (including off balance sheet items and derivatives exposures). Furthermore, to detect fast moving risk shifting we will need to collect financial market transaction data (i.e., flow data).

Ideally, the validity of the data is beyond reproach. In practical stress testing, significant effort is spent on discussions about whether the often patchy data quality warrants the conclusion. However, discussions about the interpretation of stress test results should ideally not be about the validity of the bare facts; these discussions should focus on the conclusions and their policy implications. In practice, some of the data might indeed be noisy and – as some argue – this is inevitable. In stress test we often have no choice, however. We should be pragmatic and get the data up acceptable levels with sensible business rules.

With data on the entire ecosystem, we will be able to deliver the raw material for both micro- and macroprudential stress tests. In a sense, we can seamlessly zoom from a highly aggregate macro level all the way down to individual transactions of individual market participants. Moreover, we can also connect markets and see what market participants or sectors are doing in different instruments. Depending on the mandate, regulators and supervisors are interested in different scales of aggregation. For instance, a conduct-of-business supervisor will be interested in individuals trading with insider knowledge. A microprudential supervisor will be looking at the solvency of a particular firm, while a macroprudential authority will look for vulnerabilities at the macro level. With granular data, views for each of these can be generated from a single data source. To be able to zoom into different time scales we need the most granular level to start with.

To illustrate the notion that collecting granular data allows for flexible "zooming", we show data on securities holdings in the Euro area. For compiling national accounts

[14] See Bijlsma *et al* (2018) for an example.

statistics, ESCB member banks collect securities holdings of Euro area residents combined in the Securities Holding Statistics (SHS). Each central bank collects individual securities holdings from relevant entities in their jurisdiction. In the Netherlands, the reporting sample includes around 800 reporters. In principle, we can thus move from country aggregates all the way down to individual firms holding a particular instrument. We visualise this process in Figure 7 in three panes.[15] In each pane we use the same data but focus on different levels of aggregation. For clarity, we collapse smaller countries to a 'rest of the world' (RoW) category and show holder countries in the bottom half of the circle while issuers are shown in the top half.

In the left pane (Pane A), we show the total exposures of country holders to various issuing countries. This could be called the `macro' view. Here we see that France is the largest holder (8.3 Tr euro) and that its holdings have a significant home bias: 60% of its holdings are issued domestically with 500 Bn euro (6%) issued in Germany a distant second. It is clear that home bias is an entrenched stylised fact. This level of aggregation is useful for macro-economic analyses of international capital stock and flows and is most useful for macro stress tests.

Next, if we are interested in a particular sector, we can zoom in as in Pane B. We term this the `meso' view. Here, we show the exposures of the different sectors in the Netherlands to different other issuing countries. The largest holdings are concentrated in the pension fund sector, with the Monetary Financial Institutions (MFI) or banking sector a close second. We see that for most sector the largest exposure is to Dutch firms (i.e., a self-loop).

Figure 7 Securities data at different levels of granularity

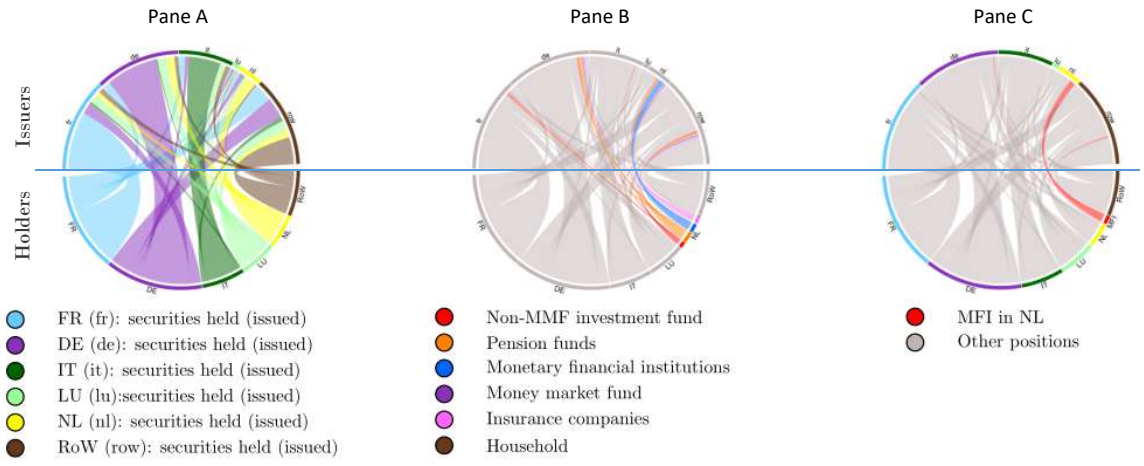

The rightmost pane (Pane C) highlights the exposures of a single sector, in this case the banking sector. Showing all sectors, as in Pane B, can sometimes conceal important exposures if we are interested in just a single sector. Alternatively, we could use the colour to denote the issuing country. Such a view would allow us to focus on the source of foreign shocks.

Zooming in even further – not shown due to data confidentiality – we get to the `micro'

[15] See Sigl-Gloeckner (2018) for a survey of network visualization methods, in particular for the SHS data.

view where we could show to what country-sector one particular pension fund is exposed. This would be of interest to the pension fund's supervisor when assessing a fund's country risk. Finally, we can discern the 'nano' view. Here we could highlight the exposure of a single fund to a single instrument (identified by ISIN). As an additional flourish, we can show all related instruments issued by the same corporate group or, in the same vein, to a particular group of instruments (cf Boermans and van Wijnbergen (2018)).

Having the ability to zoom is potentially extremely valuable for the efficiency of stress testing. The same source data can be used for stress testing from different micro and macro angles. For example, if we are worried about country risk exposures of a particular sector, we take a `meso' view and stress the valuation of the instruments held. This is generally a financial stability concern. Conversely, if we are interested in a micro-prudential approach, we can just select the individual institution and apply the stress parameters to its portfolio items. The same data can thus serve different constituencies. Incidentally, having granular data at our disposal also means that macroprudential stress tests can have micro-foundations, meaning that they can model/simulate shock amplification with much higher degrees of verisimilitude, potentially capturing otherwise-hidden risks.

With the underlying data of the SHS, national authorities can also move to the individual firm, be it a bank, a pension fund, or an insurance firm. The portfolios of the larger banking firms are available in full detail. On this level, we can currently already add a lot of other firm-level data. For example, we can add public information from commercial data providers such as SNL and ORBIS. For research purposes, we can add prudential information from *inter alia* FINREP and COREP. With such information we can for instance study how a bank suffering liquidity stress – as measured in the regulatory liquidity coverage ratio (LCR) reports – reacts by engaging in precautionary hoarding of liquid securities (Acharya and Merrouche (2013)) or is forced to fire sale (Diamond *et al* (2011)).

Going forward, it should also become much easier to zoom using other cross-country data collections. For example, the European Market Infrastructure Regulation (EMIR), Securities Financing Transaction Regulation (SFTR) and Money Market Statistical Reporting (MMSR) are all granular data collections that can be analysed on different aggregation levels.

Consistent data across different markets also allows for stress testing of connected markets. For example, a recent trend is that exposures in derivative market are ever more mitigated by the exchange of collateral. This is driven by regulation mandating central clearing but also by firm's reduced appetite for counterparty credit risk. By exchanging collateral, counterparty credit risk is reduced significantly. However, since collateral requirements are very sensitive to market volatility, the reduction in credit risk is accompanied by an increase in liquidity risk. Clearly, sudden spikes in market volatility translate directly into collateral calls. And funding such calls can be particularly expensive or difficult at precisely these same times. To properly understand either the derivative or funding markets such as the repo market, we should thus be able to analyse the two markets jointly (Adrian *et al* (2013)).

Collecting consistent information for an entire market might allow for better risk assessment than what financial markets can currently achieve. Financial markets are

generally very efficient in digesting information. However, a financial market cannot price a risk it cannot observe. In many OTC markets, only a few core players can form a reasonably coherent view on all participants' positions. Such market-making parties can capitalise on the information embedded in order flow or – accumulated over time – knowledge of counterparties' positions (Duffie and Manso (2007)). In a dispersed market, however, such centralised knowledge is not available. The market can thus sometimes miss pockets of risks due to, for example, concentration risk. A well-known example is AIG, the American insurer that turned out to be a key player connecting the financial markets in the 2008 crisis. Information on the network as a whole can potentially improve on fragmented information. For example, Squartini, Garlaschelli and Van Lelyveld (2010) show that network information can show fragmentation in interbank network connections significantly earlier than market signals. Such market fragmentation can potentially hamper a market's ability to redistribute risks.

Note that stress tests are often kept simple on purpose in order to be able to tell causal stories (Drehmann (2009)). Some data analytics methods do not fit well with such a strategy. The results from machine learning algorithms, for instance, are difficult to trace back to the actions of individual actors. This makes explaining the outcomes to mostly uninitiated users of the results (i.e., management or the general public) rather difficult. Without a comprehensible narrative, defending stress test outcomes becomes very difficult. One way forward would be to try to model the complexities in the tradition of Herbert Simon (Simon (1962)). The understanding of the underlying processes would allow for the translation into causal stories. Another way forward would be to tell a broader story about how shocks to the system, depending on their nature and magnitude, can either be absorbed or are destabilizing (Wiersema *et al* (2020)). Finally, in some cases contributions to particular outcomes can be attributed to particular agents or risk drivers using the concept of Shapley values (Joseph (2019)).

Challenges in a new environment

The new possibilities for stress testing discussed so far also come with several challenges as well in terms of a) governance, b) the IT capacity to handle the data volume, and c) the skill sets needed. We will discuss each of these in turn.

Governance. Collecting data on a more granular level implies more, often sensitive, detail. Moreover, the granular level sometimes contains information on natural persons. For example, loan level information might contain natural persons' names which are in principle not needed for stress testing (beyond being the primary key to link other information (e.g., tax records). Handling such data demands additional care. In this case, it might be useful to anonymise the loan level data while preserving the ability to combine the data with other sources. In Europe, the General Data Protection Regulation (GDPR), which is being enforced as of May 2018, now comes with much tighter regulation of what firms and authorities can and cannot do with natural persons' data. In collecting data that contains or might contain natural persons data, adequate thought needs to be given to ensure fair use while still adhering to the GDPR.

Volume. To collect, digest and store larger volumes of data, authorities might need to explore unfamiliar terrain, for example storing data outside their own servers. Although

generally the costs of commercially available cloud services (i.e., computing and storage), have come down drastically, many government agencies are not yet ready to take advantage of this since this would mean outsourcing IT capacity in the cloud. The decision to outsource is an important strategic choice that also comes with operational, security and legal risks. Arguably, security risks have economies of scale and could thus possibly be better handled by larger and more specialised outside parties. Note that in some cases, the data is not allowed to leave the jurisdiction and in smaller jurisdictions this limits the number of providers that are able to provide the services demanded.

To get a better understanding of the challenges, possibilities and limitations of using cloud services, DNB has started an initiative in early 2018 to bring confidential data to the cloud. Here, we can more easily apply advanced and computationally expensive techniques to datasets of several billion observations. Such volumes are already too big for most central banks and supervisors and thus potentially useful insights are foregone. Moreover, one might even argue that the mandate to collect transaction-level data is rendered pointless if analysts do not have the IT infrastructure needed to study the data. One of the pilot projects undertaken to build on our capacity to analyse at scale is a computationally expensive stress test of margining requirements in the OTC interest rate derivatives market. The stress tests looks into the liquidity risk of margin procyclicality which has already attracted some attention in the academic literature (cf Glasserman and Wu (2018)).

Skill set. If an organisation is to become truly data driven, then data should be embedded centrally in the business process. This requires that those involved should not be 'data challenged'. Some would argue that regulators are more likely to be 'data challenged' than the regulated, and that this creates its own pitfalls. Traditionally, market conduct supervisors are trained lawyers while central banks rely on economists. Both lawyers and economists generally have attractive outside options in the private sector. Authorities cannot remunerate on the same level as the private sector but try to compete by providing a better work-life balance and an appeal to civic duty. In addition, working in government for some time allows professionals to build a network and get a much better understanding of the regulatory process. This is very valuable for a career in the private sector and the wages foregone can thus be seen as an investment in higher future income in the private sector. However, Luca Enriques has argued that such a trade-off might not be relevant for IT engineers and Data Scientist in Regtech or Suptech.[16] Developing IT applications for regulatory compliance or supervisors is much less dependent on soft skills and localised knowledge. It will therefore be more difficult for government agencies to attract top-level IT engineers. Given the ample outside opportunities for people with these skills, it becomes even more important to offer interesting and rewarding jobs. Hopefully, the very rich data sets central banks and supervisors can use and the appeal of a meaningful job that offers rewards to society at large is sufficiently convincing to data scientists and IT engineers to accept a pay cut.

It is crucially important that data scientists interact with supervisors and policy officers. In doing so, data scientists will be inspired by real world problems, whereas their colleagues in policy and supervision will get a better understanding of where central banking and

16

supervision can benefit in practice from the data revolution. New techniques such as artificial intelligence and machine learning will not drive out professional judgement, which in terms of analysing shocks and their impact on the financial sector will remain essential for figuring out which shocks are feasible and for thinking outside the box more broadly. In terms of policy making on the basis of stress testing, judgement calls will remain essential, for instance in case of dilemmas where a choice is needed between two options which each affect a different group in society.

Discussion and conclusion

To really make new opportunities in data pay off for stress testing requires hard work. The new and enticing possibilities are not a free lunch. To achieve them means to adhere to data governance principles, for instance by using meta data consistently and on a large scale, collecting and cleaning new data sets, applying LEI and UPI codes etc. For a single stress testing project the costs may seem disproportional. But usually there are substantial external benefits. As we go further down this road, new opportunities are emerging for an all-encompassing risk management approach on the basis of modelling, in a comprehensive way, the financial ecosystem.

In this chapter we have explored the opportunities that more granular data and closer to real time reporting bring. Granular – even transaction level - data allows us to aggregate the data in many ways. This is ideal for stress testing as here flexibility is key; the strength of stress testing is that it is possible to entertain unlikely scenarios that have never materialised. Often these involve new markets or players for which we don't regularly collect data. These opportunities also come with challenges. Data collection is often organised in silos. For instance, insurance supervisors focus on insurance firms (e.g., Solvency II) while reporting for monetary policy operations (e.g., MMSR) covers primarily banks. In looking at the impact of interest rate setting on the wider economy it might, however, be interesting to analyse how money market funding conditions are affected by insurance firms' repo transactions – and vice versa. Such analysis is hampered by inadequate coordination on reporting standards (frequency, metadata, etc.).

Having more and more granular data in-house also puts more weight on the governance of that data. For instance, who is the owner of the data? Dispersed ownership is an invitation to free riding and cost shifting. At DNB, the owner is therefore always a single person, although he or she is free to involve others for input or funding. Given that more granular data is often also useful more broadly, finding an organisational form that takes into account external effects is becoming more important. Furthermore, the owner should aim for maximally distributing the data to be able to realise its full potential. This implies a 'share, unless …'-policy which should be accompanied with well implemented access rights on the ground. Such a data policy should hold within organisation as well as for the interaction between organisations involved – that is the macroprudential supervisor (ESRB and OFR/FSOC in the EU and the US, respectively) and supervisory authorities and central banks.

The flood of new data potentially also has reputational risk for authorities. For some parts of the system authorities knew very little and were thus blind sighted for the build-up of risks. In the aftermath of the crisis, a massive increase in data reported ensued. The

challenge authorities now face is to effectively use all the data and turn it into information and actionable analysis. Failing to incorporate the new data into analyses runs the risks of instigating a large public backlash the next time a crisis materialises. Some might conclude that authorities did have crucial information reported but failed to act on it.

In sum, new data and tooling have the potential to significantly improve stress testing but challenges remain. The sheer volume of data requires an investment in data processing capacity. More importantly, this requires a change in the way we use the data. Traditionally, those involved in stress testing have a supervisory or a financial stability background. Knowledge of database management and optimizing database queries is thus often insufficient to handle large volumes efficiently. Another challenge is how to communicate results to senior management and lay people. Scrolling through the data in an Excel sheet to get a feel for the data is no longer feasible. In presenting results, more care should thus be given to supply appropriate dashboarding as well. Characterizing high dimensional data with new concepts – such as the eigenvector centrality of a bank in a market – needs to be accompanied with sufficient explanation to instil a sense of comfort in decision makers.

References

Abad, J., Aldasoro, I., Aymanns, C., D'Errico, M., Fache Rousová, L., Hoffmann, P., Langfield, S., Neychev, M., Roukny, T., 2016. Shedding Light on Dark Markets: First Insights from the New EU-wide OTC Derivatives Dataset. ESRB Occasional Paper 11.

Acharya, V. V., Merrouche, O., 2013. Precautionary Hoarding of Liquidity and Inter-Bank Markets: Evidence from the Sub-prime Crisis. Review of Finance 17, 107-160.

Adrian, T., Begalle, B., Copeland, A., Martin, A., 2013. Repo and Securities Lending. Federal Reserve Bank of New York Staff Reports 529.

Anand, K., van Lelyveld, I., Banai, A., Christiano Silva, T., Friedrich, S., Garratt, R., Halaj, G., Hansen, I., Howell, B., Lee, H., Martínez Jaramillo, S., Molina-Borboa, J. L., Nobili, S., Rajan, S., Rubens Stancato de Souza, S., Salakhova, D., Silvestri, L., 2018. The Missing Links: A Global Study on Uncovering Financial Network Structure from Partial Data. Journal of Financial Stability 35, 107-119.

Aymanns, C., Georg, C.-P., Golub, B., 2018. Illiquidity Spirals in Coupled Over-the-Counter Markets.

Barbic, G., Borgioli, S., Klacso, J., 2017. The journey from micro supervisory data to aggregate macroprudential statistics. ECB Statistics Paper Series 20.

Berner, R., Judge, K., 2019. The Data Standardization Challenge. ECGI Working Paper Series in Law 598.

Bijlsma, M., Castro Campos, M., Chaudron, R., Jansen, D.-J., 2018. Building a Multilayer Macro-Network for the Netherlands: A New Way of Looking at Financial Accounts and International Investment Position Data. IFC Working Papers.

Boermans, M. A., van Wijnbergen, S., 2018. Contingent convertible bonds: Who invests in European CoCos? Applied Economics Letters 25, 234-238.

Bottega, J. A., Powell, L. F., 2011. Creating a Linchpin for Financial Data: Toward a Universal Legal Entity Identifier. Finance and Economics Discussion Series 7.

Cerutti, E., Schmieder, C., 2012. The Need for "Un-consolidating" Consolidated Banks' Stress Tests. IMF Working Paper WP/12/288.

Diamond, D. W., Rajan, R. G., 2011. Fear of Fire Sales, Illiquidity Seeking, and Credit Freezes. The Quarterly Journal of Economics 126, 557-591.

Drehmann, M., 2009. Macroeconomic stress-testing banks: a survey of methodologies. In: Quagliariello, M. (ed.), Stress-testing the Banking System: Methodologies and Applications, Cambridge University Press.

Duffie, D., Manso, G., 2007. Information Percolation in Large Markets. American Economic Review 97, 203-209.

Ehlers, T., Kong, S., Zhu, F., 2018. Mapping Shadow Banking in China: Structure and Dynamics. BIS Working Paper 701.

European Commission, 2019. Fitness Check of EU Supervisory Reporting Requirements. Commission Staff Working Document SWD.

Fang, Y., van Lelyveld, I., 2014. Geographic Diversification in Banking. Journal of Financial Stability 15, 172-181.

Financial Stability Board, International Monetary Fund, 2009. The Financial Crisis and Information Gaps. Report to the G-20 Finance Ministers and Central Bank Governors.

Financial Stability Board, International Monetary Fund, 2017. Second Phase of the G-20 Data Gaps Initiative Second Progress Report. Report to the G-20 Finance Ministers and Central Bank Governors.

Flood, M., Katz, J., Ong, S., Smith, A., 2013. Cryptography and the Economics of Supervisory Information: Balancing Transparency and Confidentiality. OFR Working Paper 11.

Foster, W. T., 1922. The Circuit Flow of Money. American Economic Review 12, 460-473.

Frost, J., 2020. The Economic Forces Driving FinTech Adoption across Countries. DNB Working Paper 663.

Glasserman, P., Wu, Q., 2018. Procyclicality in Sensitivity-Based Margin Requirements. Working Paper.

Joseph, A., 2019. Shapley regressions: A framework for statistical inference on machine learning models. XArchiv Working Paper.

Joshua Graff, 2018. The real shape of Britains post-Brexit talent brand. Linkedin Pulse.

Judson, R., 2017. Enhancing Data Availability: Recent U.S. Experience with Banking Data. In: Workshop on Data Sharing, Deutsche Bundesbank.

Kienecker, K., Sedlacek, G., Turner, J., 2018. Managing the processing chain from banks source data to statistical and regulatory reports in Austria. OeNB Statistiken 3.

Kwast, M. L., Holden, S., Jurcevic, D., Van Lelyveld, I., 2010. Norges Bank Stress Testing of Credit Risks - Report of an External Review Panel.

Lauro, B., Traverso, R., 2018. Data Fitness for Integration. Mimeo.

Legal Entity Identifier Regulatory Oversight Committee (LEI ROC), 2018. The Global LEI System and regulatory uses of the LEI. Progress Report.

Levels, A., De Sousa Van Stralen, R., Petrescu, S. K., Van Lelyveld, I., 2018. CDS market structure and risk flows: the Dutch case. DNB Working Paper 592.

Sigl-Gloeckner, P., 2018. Visualising Financial Systems. Thesis Imperial College London.

Simon, H. A., 1962. The Architecture of Complexity. Proceedings of the American Philosophical Society 106.

Squartini, T., van Lelyveld, I., Garlaschelli, D., 2013. Early-Warning Signals of Topological Collapse in Interbank Networks. Nature Scientific Reports 3, 3357.

The Joint Forum, 2012. Report on intra-group support measures.

van Lelyveld, I., Liedorp, F., 2006. Interbank Contagion in the Dutch Banking Sector: A Sensitivity Analysis. International Journal of Central Banking 31, 99-133.

van Lelyveld, I., Liedorp, F., Kampman, M., 2011. An empirical assessment of reinsurance risk. Journal of Financial Stability 7, 191–203.

Wang, D., van Lelyveld, I., Schaumburg, J., 2018. Do Information Contagion and Business Model Similarities Explain Bank Credit Risk Commonalities? DNB Working Paper 619.

Wiersema, G., Kleinnijenhuis, A. M., Wetzer, T., Farmer, J. D., 2020. Inherent Instability: Scenario-Free Analysis of Financial Systems with Interacting Contagion Channels. Institute for New Economic Thinking 10

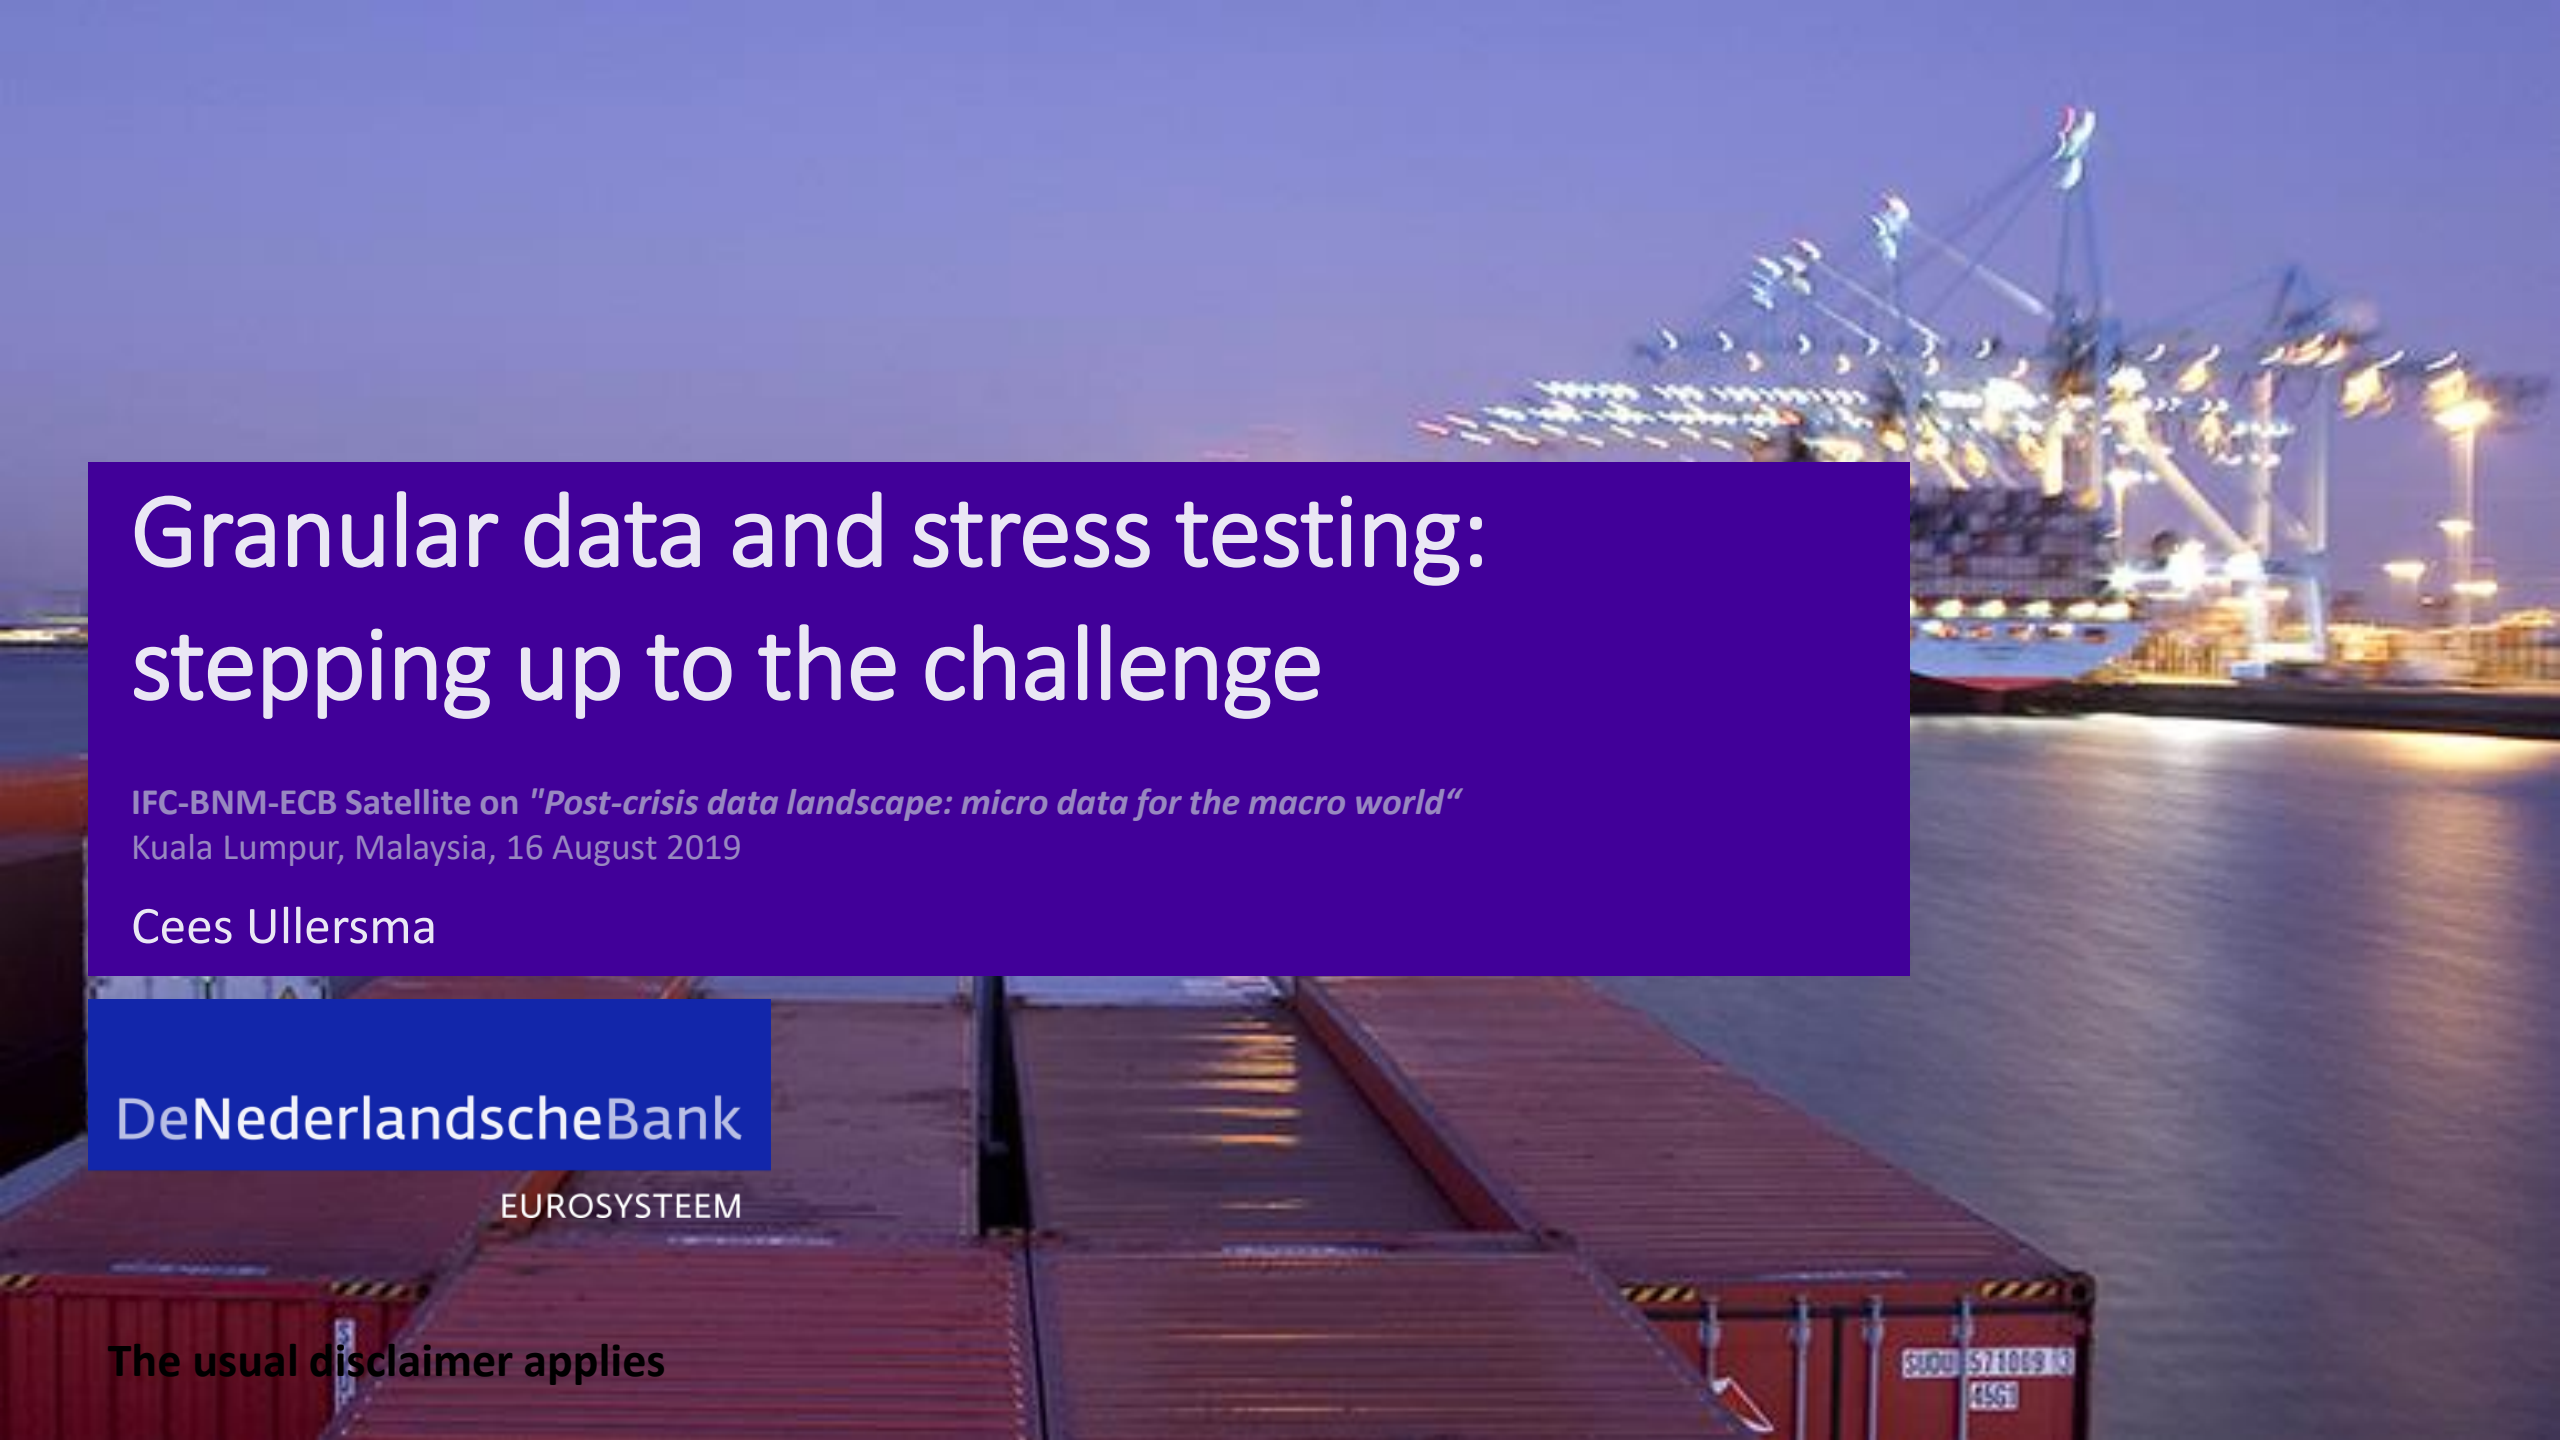

Granular data and stress testing: stepping up to the challenge

IFC-BNM-ECB Satellite on *"Post-crisis data landscape: micro data for the macro world"*
Kuala Lumpur, Malaysia, 16 August 2019

Cees Ullersma

DeNederlandscheBank

EUROSYSTEEM

The usual disclaimer applies

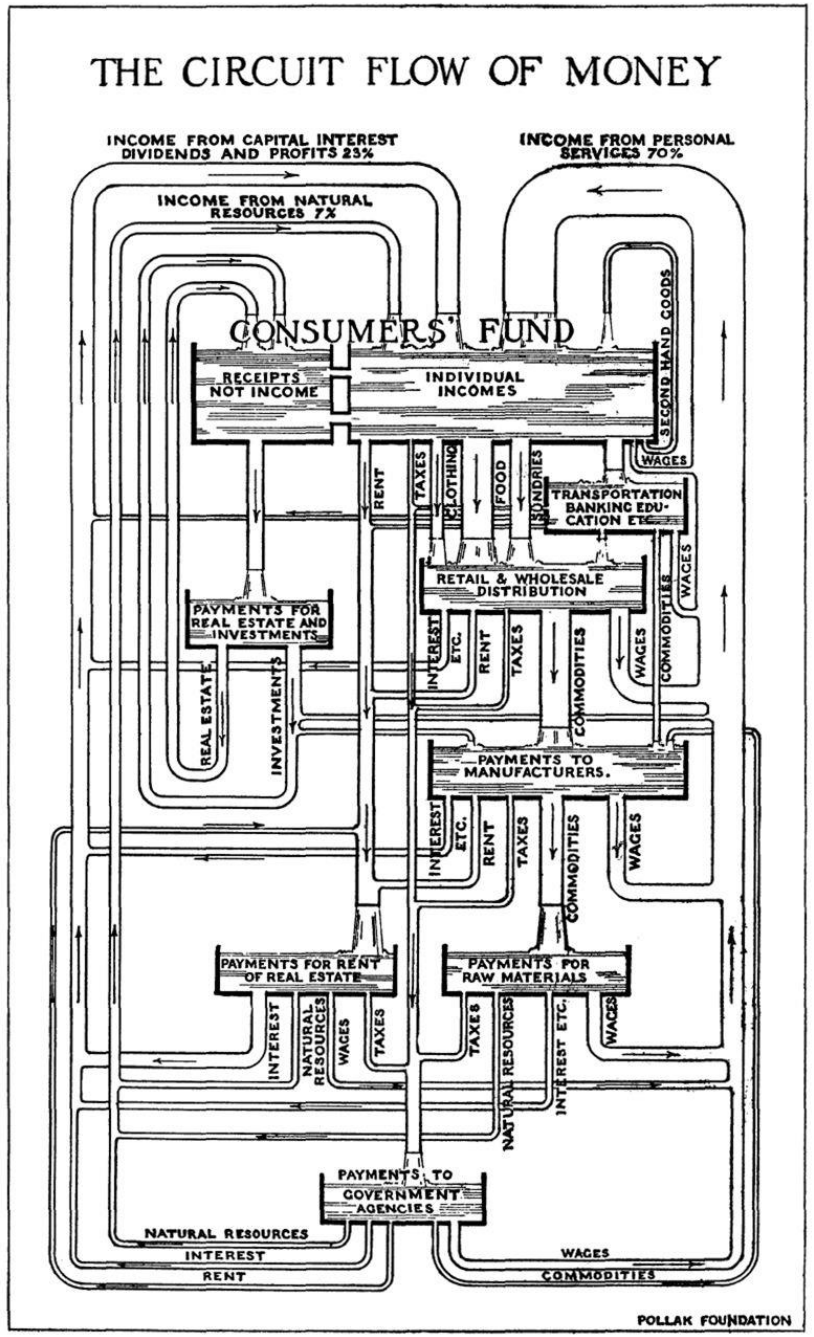

Source: Foster (1922), “The Circuit Flow of Money”, AER

DeNederlandscheBank
EUROSYSTEEM

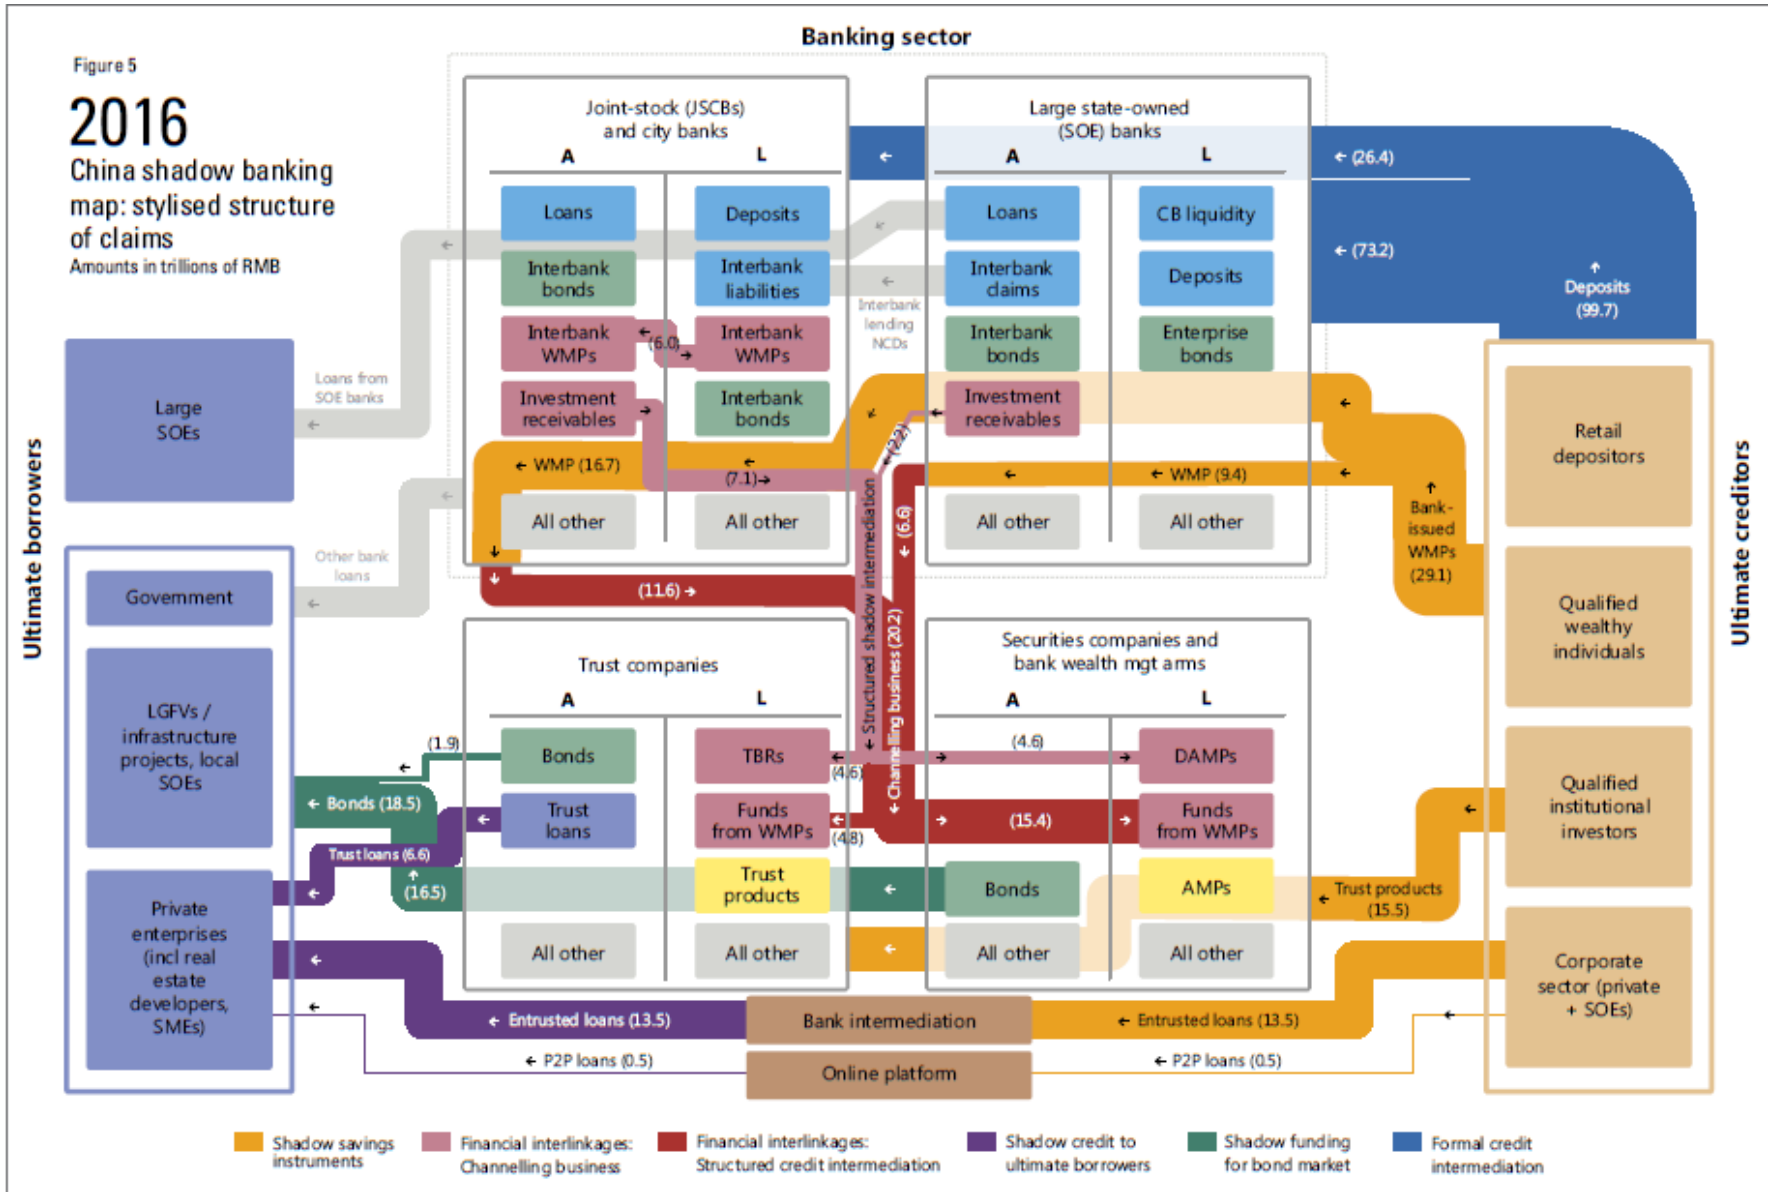

Source: Ehlers et al (2018), "Mapping shadow banking in China: structure and dynamics", BIS WP

Main message:

- Granular data: allows for a mapping of financial eco system
- Data on stocks (balance sheets) and flows (transactions) available more widely
- New opportunities for analysis of shocks
  - Identifying troubled firm, sectors, and markets
  - Trace propagation of shocks through system
- Very promising for stress testing

Excellent starting position

- Availability of granular data
  - Securities Holding Statistics (SHS)
  - Data Gaps → International Data Hub at the BIS
- New OTC reporting
  - Derivatives (EMIR)
  - Money markets (MMSR)
  - Repo (SFT)
- Key competence of central banks and supervisors

Challenges and the way forward

- Data sources diverse
  - Structured & unstructured
  - Numeric & textual
  - Inside & outside
- Obstacles to data sharing
  - Legal restrictions within and across jurisdictions
- Adequate data governance
  - Data ownership and access should be well anchored
- Have the basics in place
  - data catalog, meta data, LEI/UPI
- Open attitude to alternative ways of cooperating
  - Sharing files (eg. mailing code, Sharepoint)
  - GIT: most widely used version control system
  - Container technology

The importance of the Legal Entity Identifier (LEI)

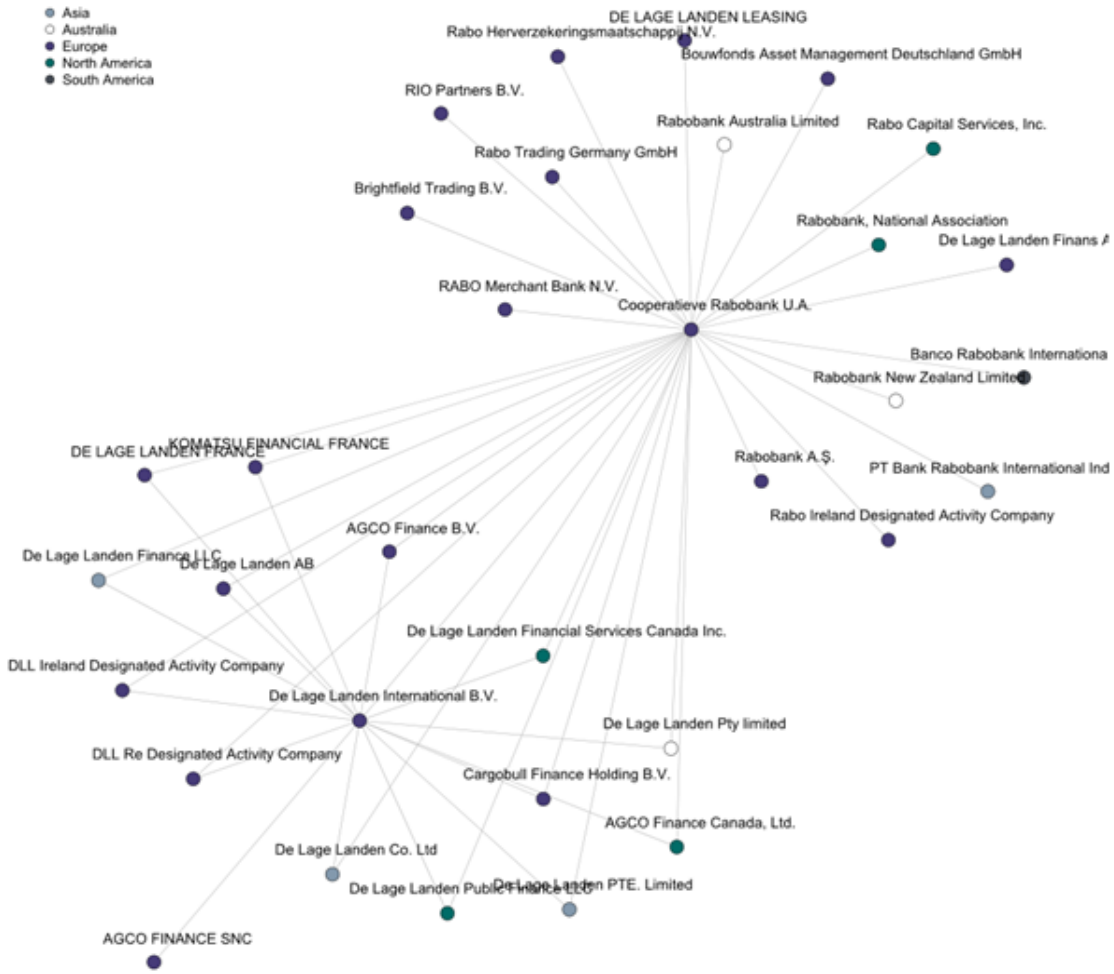

DeNederlandscheBank
EUROSYSTEEM

The same granular data can serve multiple goals

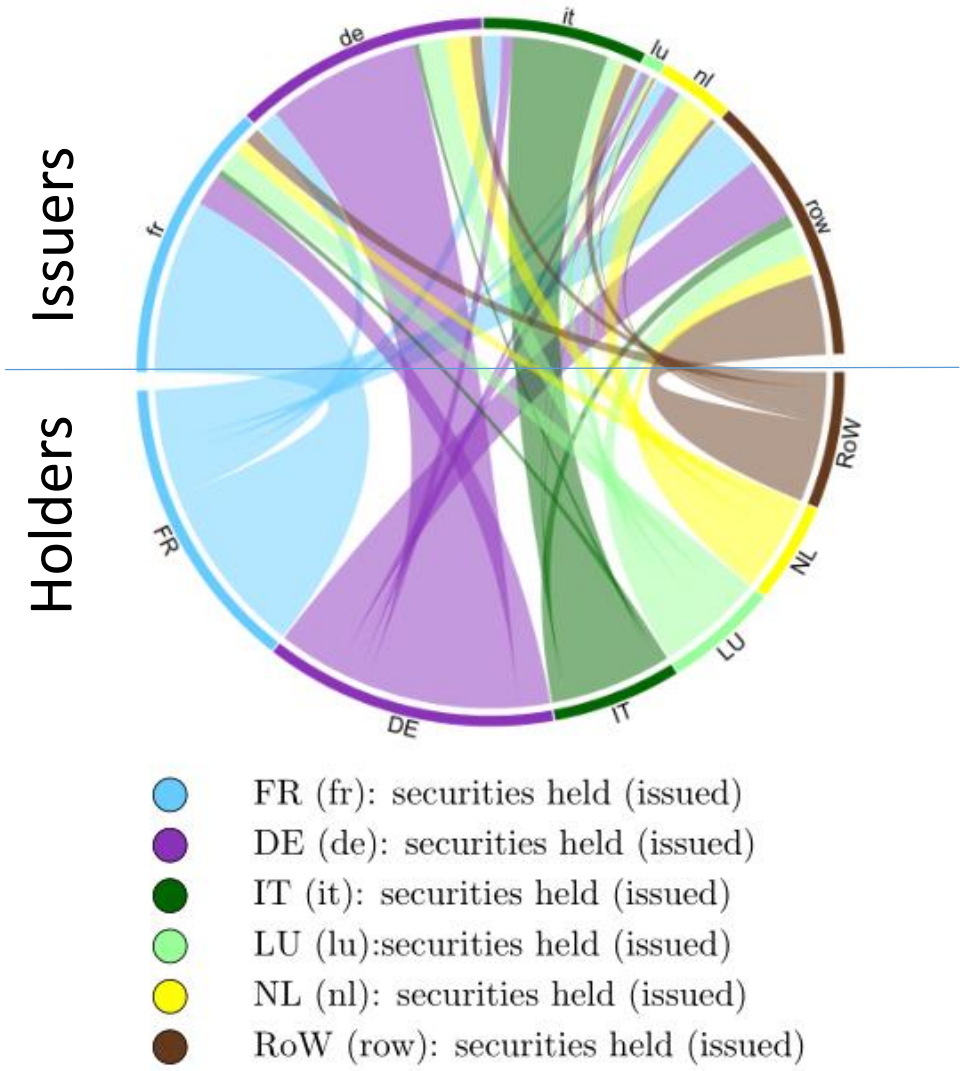

- Data such as SHS has ISIN level exposures on a sector or firm level
- Allows us to see sector-to-sector linkage
- Given security meta data views can extend to focus on:
  - Different securities (eg equity, bonds, …)
  - Maturity
  - Green footprint --> carbon stress test
  - Bond type (eg. CoCo bonds)
- Linkage with country and sector level information
- Macro prudential stress tests

The same granular data can serve multiple goals

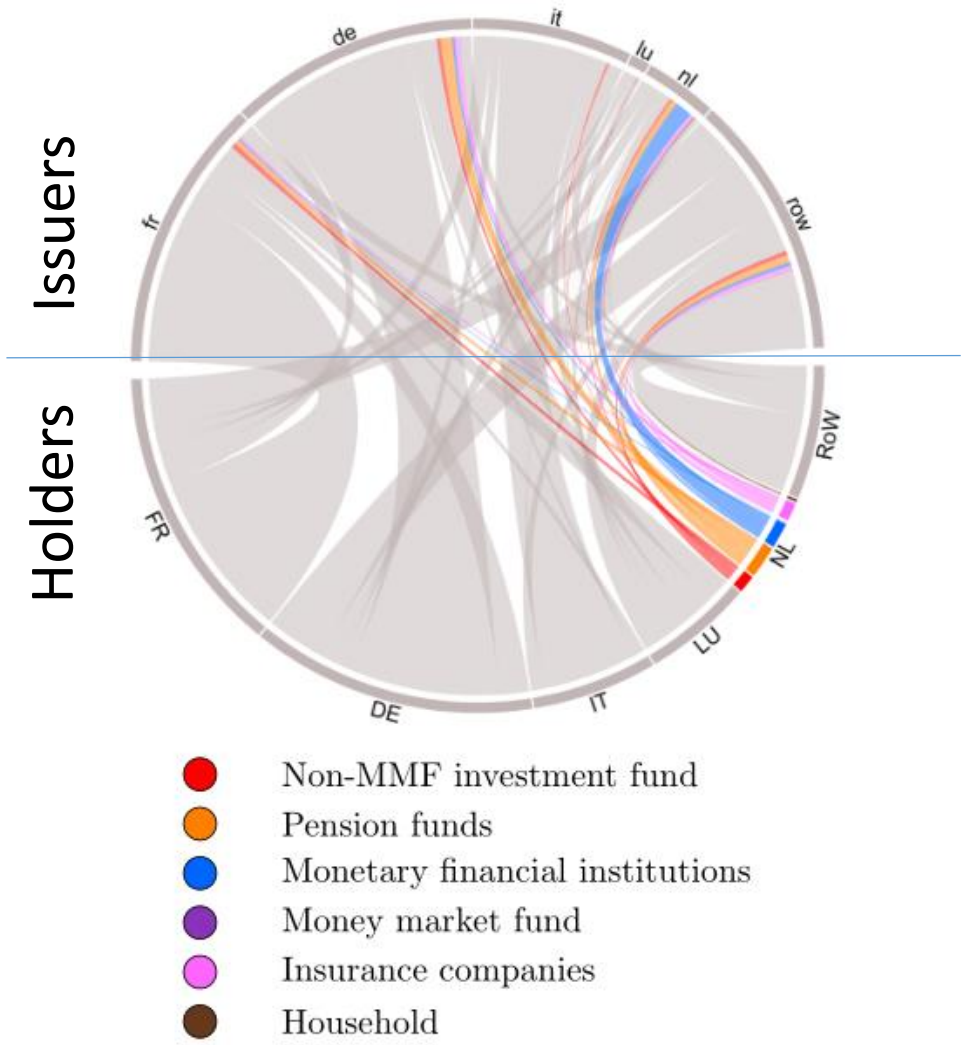

- Same data, but now identifying granular sectors
- Allows to analyse the relative risk for different sectors
- Meso stress test

The same granular data can serve multiple goals

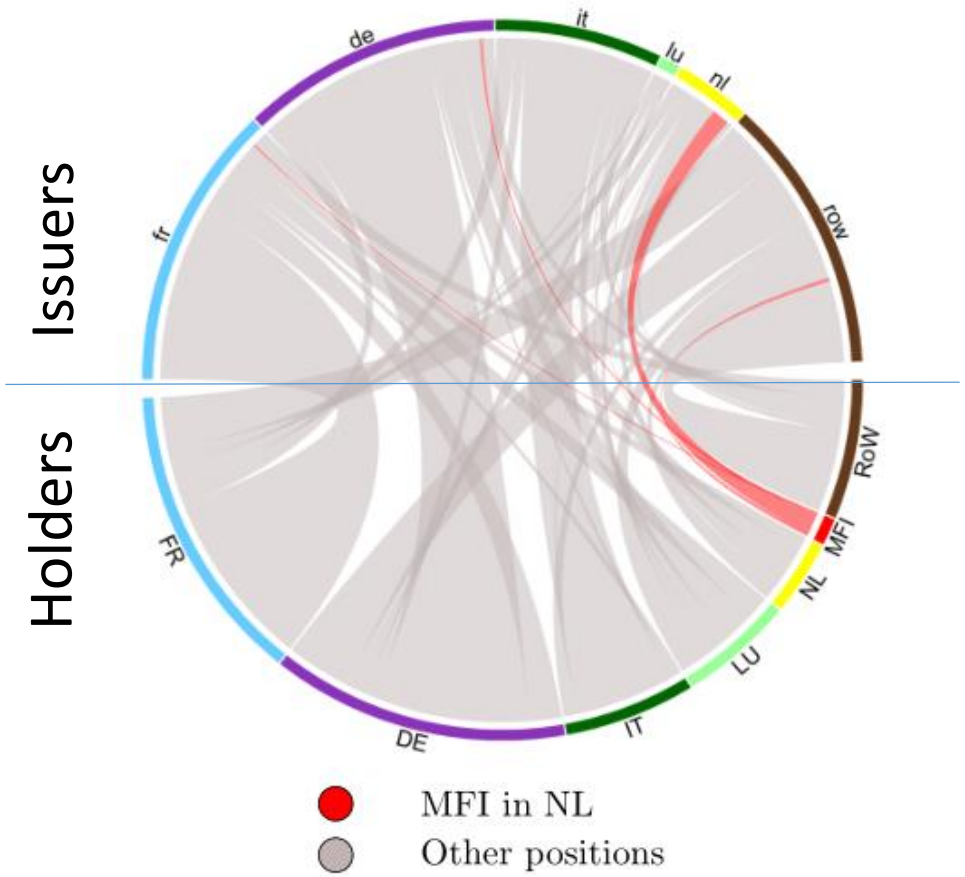

- Drilling further into one sector
- Granularity available but confidentiality precludes showing this detail
- Easy linking to prudential information such as balance sheet or risk information
- Here we could do micro prudential stress tests

Considerations

- The need for causal stories
  - We need to open the black box and have actionable stories
- Find your way in the cloud
  - DNB Data Science Hub has put confidential data in the cloud
- Attract new staff and train existing staff
  - DNB Data Science Hub
  - Traineeship “Data and Technology”
  - Learning trajectory: *Become a “Datapreneur”*
- Let data scientists talk to supervisors and policy officers
  - Joint projects in a hub-and-spokes model

WORD DATAPRENEUR!

- Leertraject
- 6 maanden, 1 dag per week
- Los je eigen data-vraag op!

Data-entrepreneurs in de dop gezocht!

Wil jij zelf aan de slag met data om vraagstukken uit de praktijk op te lossen? Geef je dan op voor een nieuw leertraject: Word Datapreneur. Wie weet mag je wel jouw eigen data-oplossing gaan uitwerken... Iman van Lelyveld en Jasmijn van Loo van de Data Science Hub leggen uit.

Wat is 'Word Datapreneur'?
"Word Datapreneur is een leertraject waarbij DNB'ers een businessprobleem kunnen aandragen dat ze zelf willen aanpakken met behulp van data. Uit de aanmeldingen selecteren we een aantal deelnemers, die vervolgens samen aan de slag gaan om hun eigen en elkaars ideeën te onderzoeken en oplossingen te bouwen. Dit onder begeleiding van collega's van de Data Science Hub."

Waarom dit leertraject?
"Met 'Word Datapreneur' willen we twee dingen bereiken. Ten eerste werken we aan oplossingen voor bestaande problemen. Ten tweede leer je als deelnemer spelenderwijs coderen, om te gaan met data en samen te werken. Je ontwikkelt je tot data-entrepreneur: iemand die ondernemend bezig is met data, ideeën heeft over hoe data een oplossing kunnen bieden voor beleidsvragen, en die ideeën weet te verkopen. Het geeft handen en voeten aan onze ambitie om digitaal vaardig te worden."

Wat houdt het programma concreet in?
"We selecteren zes tot tien groepjes die aan hun businessprobleem gaan werken. Zij beginnen half mei en gaan een halfjaar lang één dag per week aan de slag. Het

Thank you for your attention

Questions?