agresearch
āta mātai, mātai whetū

White Paper

Detecting systemic risk earlier through artificial intelligence

Mark Wever, Munir Shah and Niall O'Leary

May 2020

Acknowledgement: this project has received financial support from "New Zealand Bioeconomy in the Digital Age", a science program funded by the Ministry of Business, Innovation and Employment

Pitch

Systemic risk, such as M.bovis or COVID-19, do great harm to the New Zealand economy. Artificial intelligence (AI) can identify such risks earlier and facilitate quicker responses and better outcomes. We present a novel, state-of-the-art method for systemic risk detection based on AI. Development and deployment of this method could save the New Zealand economy hundreds of millions of dollars over the next decade.

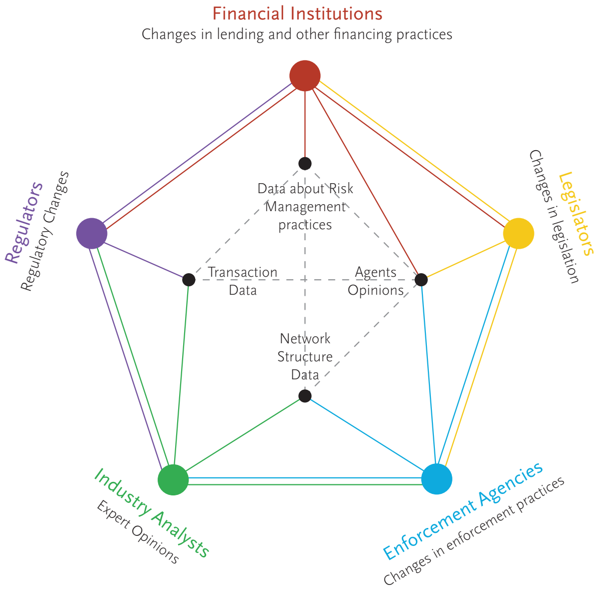

Foreword

Systemic risks are potential trigger events or developments that could undermine the viability of entire networks or systems. Examples include contagious food diseases and the bankruptcy of keystone companies. Systemic risks are likely to happen more frequently going forward, as markets (e.g. financial) are becoming increasingly complex and opaque, technological change is accelerating at an unprecedent rate and the global political situation is becoming more unstable.

These same fundamental changes to the way the world operates mean it is becoming more difficult to anticipate and deal with risks when they do occur. This is in large part because the systems we need to protect are becoming more complex, and so their fragility/resilience to shocks are becoming harder to gauge.

The tools available to researchers, regulators and companies for detecting and governing systemic risk have not kept pace with these challenges. Traditional methods are generally periodic assessments. These are too intermittent, too slow and too narrow in focus for timely systemic risk detection.

We present a framework for detecting systemic risk using artificial intelligence (AI) in this paper to address these shortcomings. An AI based system could assess systemic risk in an automated, continuous, and comprehensive manner with greater vigilance and reliability than current methods. This would provide regulators and companies in production networks with earlier warning signals of a wider range of systemic risks on the one hand, and more up-to-date measures of the fragility/resilience of the system against these risks on the other hand.

To illustrate the value, we discuss how such an framework would work when it is applied to biosecurity related threats and risks, such as mycoplasma bovis (M.bovis). The spread of the M. bovis bacterium has resulted in an ongoing, established and major biosecurity incursion affecting the New Zealand dairy industry. A framework such as the one proposed here could assist with the early detection of an incursion like M. bovis, or help to mitigate the damage.

While we focus in the latter part of the paper one specific type of risk (biosecurity hazards) in one specific type of production network (agrifood), the fundamentals and principles on which our framework is based are applicable to a wide range of risks and contexts (e.g., the detection of systemic risk within the financial services industry).

What is systemic risk?

Systemic risks are potential trigger events or developments that could severely disrupt or even do unrepairable damage to a network [5, 14]. This network can be an individual industry, a sector, or even the economy as a whole. Furthermore, systemic risks are not limited to the economic domain, as their occurrence may impact the environment, the political system and other spheres of society [8, 13].

Systemic risks are often characterized by failures at critical nodes in a network [22]. In networks that are vulnerable to systemic risk, such failures can cause rippling effects across the network [6, 24]. In a financial context, for example, if inter-bank loans don't get repaid and investors fear other bankruptcies, the bankruptcy of a large and well-connected bank may cause other banks to go bankrupt as well [1, 2].

Systemic risks tend to occur more frequently, and with more severe consequences, in networks that are complex and dynamic [3]. Many social and economic networks fall into this category (e.g., cities, markets), as do, for example, many biological networks (e.g. reef ecosystems) [4].

In such networks, the relationships and interactions between the different elements of the network are often unclear and unstable [5, 6]. This makes it difficult for observes to understand what internally generated systemic risks may be lurking within the network, also difficult to understand how the network will respond to externally systemic risk."

Why is systemic risk so hard to deal with?

Risk management complications

Risk-managers usually have a host of difficulties in complex and dynamic networks, including problems with deciding:

- What techniques and models to use to assess the risks to which a network is exposed
- What scenarios to develop and test
- What data and data sources to use as input for model and scenario building
- How to generate or access this data in a timely fashion.

Difficulties arises for various reasons [5, 7, 8]:

- Key information about the network and its constituent components may be missing or out-of-date (as it describes the network in previous states, while it has already shifted to a new state)
- The nature of the systemic risks to which the network is exposed may be unknown
- The conditions, or state of the network, that increase the likelihood of a systemic risk occurring are not well understood
- The impact of a systemic event is generally very difficult to estimate, even in the case of well-known risks (e.g., it is difficult to forecast all possible ways in which climate change will impact economic performance).
- The effect of proposed or planned interventions in the network are difficult to predict.

Companies, policy-makers, and citizen/voters face various disincentives and biases against investing in systemic risk management activities [9, 10, 11]. As a result, societies also tend to commit insufficient resources to undertaking them [12, 13].

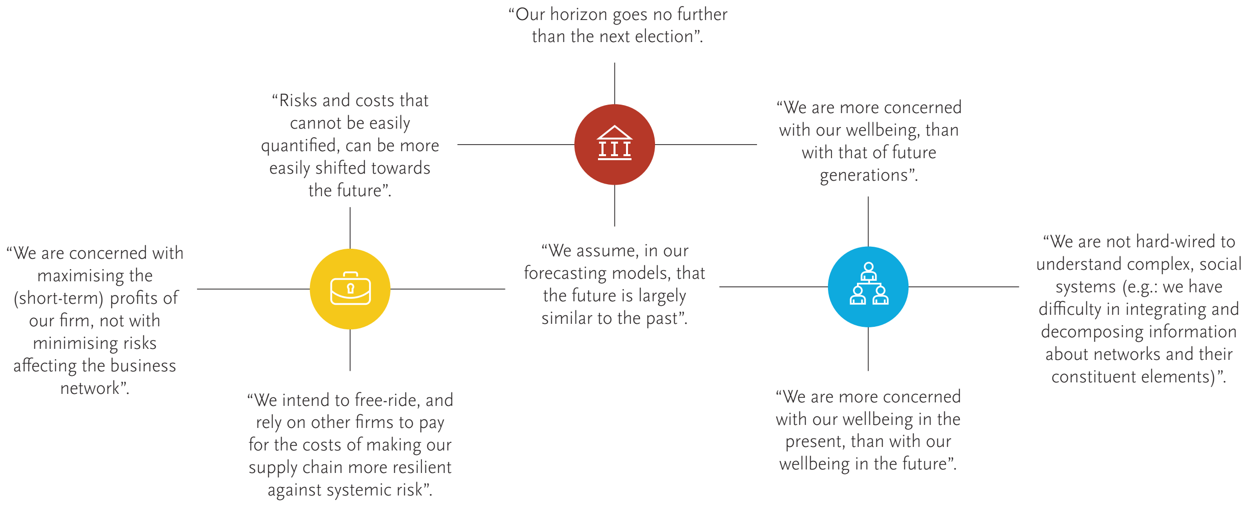
"Our horizon goes no further than the next election".
"Risks and costs that cannot be easily quantified, can be more easily shifted towards the future".
"We are more concerned with our wellbeing, than with that of future generations".
"We are concerned with maximising the (short-term) profits of our firm, not with minimising risks affecting the business network".
"We assume, in our forecasting models, that the future is largely similar to the past".
"We are not hard-wired to understand complex, social systems (e.g.: we have difficulty in integrating and decomposing information about networks and their constituent elements)".
"We intend to free-ride, and rely on other firms to pay for the costs of making our supply chain more resilient against systemic risk".
"We are more concerned with our wellbeing in the present, than with our wellbeing in the future".

Can Artificial Intelligence help?

Artificial intelligence (AI) can play an important role in helping human decision makers, such as risk managers/analysts, deal with systemic risk earlier and more effectively [14, 15]. How? By helping them overcome some of the difficulties they have in timely becoming aware of and understanding changes in complex systems.

Benefits for detection and management of systemic risk

- Wider range of early warning signals
- Prompt analysis of warning signals ▹ more time to respond
- Continuous assessment of more system failure points / system resilience
- Continuously updated assessments of potential systemic risk impacts
- Risk and resilience monitoring and impact assessment ▹ action prioritisation recommendations

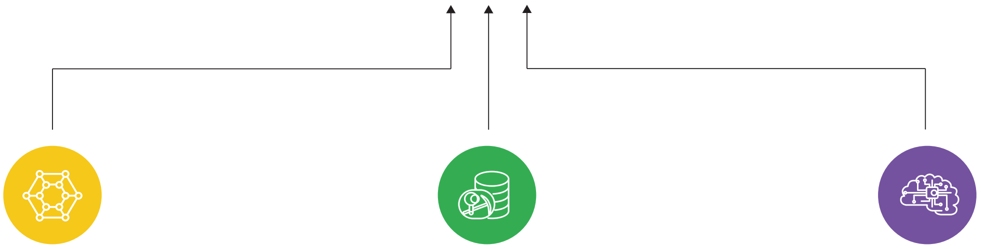

Natural language processing, video and imaging processing, speech to voice algorithms, and automatic data collection and curation, enables analysts to efficiently dissect vast and more varied data sets [16 ,17].

Machine learning principles and data mining algorithms enhances the ability of analysts to identify subtle patterns and relationships in the system [18].

Combination of automatic data collection and machine learning speeds up the process by which models of the system can be developed, tested and modified. This increases system modelling capacity dramatically.

What will we do differently?

- AI is more central to our approach than in existing approaches [14]. This would allow us to deal with far more diverse and larger data sets, which help us in turn to detect a wider range of potential early warning signals of systemic risk.

  While we could deal with far larger data sets, we would nonetheless be better able to differentiate warning signals of risk from noise. This is because we use and combine a wider set of complementary methods and tools to analyse that data [19, 20]. "If you search for a needle in haystack with your bare hands, a single haystack looks large. When you use a metal-detector and magnifying glasses, you could find the needle in 10 haystacks".
- Our approach would offer a more continuous assessment of the state of the system or network in question. While most approaches acknowledge the dynamic nature of the networks in which systemic risks arise, they nonetheless tend to offer largely periodic assessments [19, 21, 22]. In our approach, which would be more automated and in which data would be assessed from a much wider range of sources, our algorithms would continuously scan for patterns or changes in either the network, or its external environment.
- Our approach relies neither on ad-hoc selected indicators nor on a large unified theory to assess the network, but on various small-scale models working together as a "system of models" [3, 6, 20]. This helps us to give more specific and practical instructions to help our algorithms identify key agents, processes and interactions at different levels of the network.
- An important limitation of many approaches to the study of systemic risk is that they often rely on a limited range of techniques and data sets [3, 14]. However, to be able to understand complex systems, a wide range of techniques and data sets is necessary, and a lot of thinking is required in order to make these different techniques work well together. This, in order to be able to take a look at the system from many different angles and to develop a varied of scenarios about how the system may evolve [7, 19]."

  Stand-alone techniques may have their uses for generating a model about the system, for anticipating a risk or making a prediction, or for helping us understand a part of the system. However, with complex and dynamic systems, we need multiple models, anticipate various risks, and understand the system as a whole. With such systems, often a multitude of futures is equally plausible [5, 7].

  To be able to anticipate these futures, also a multitude of different methods and techniques is required. For that reason, we will use in our approach a range of different types of methods to identify early warning signals of risks at the one hand, and the fragility of the network against these risks at the other hand.

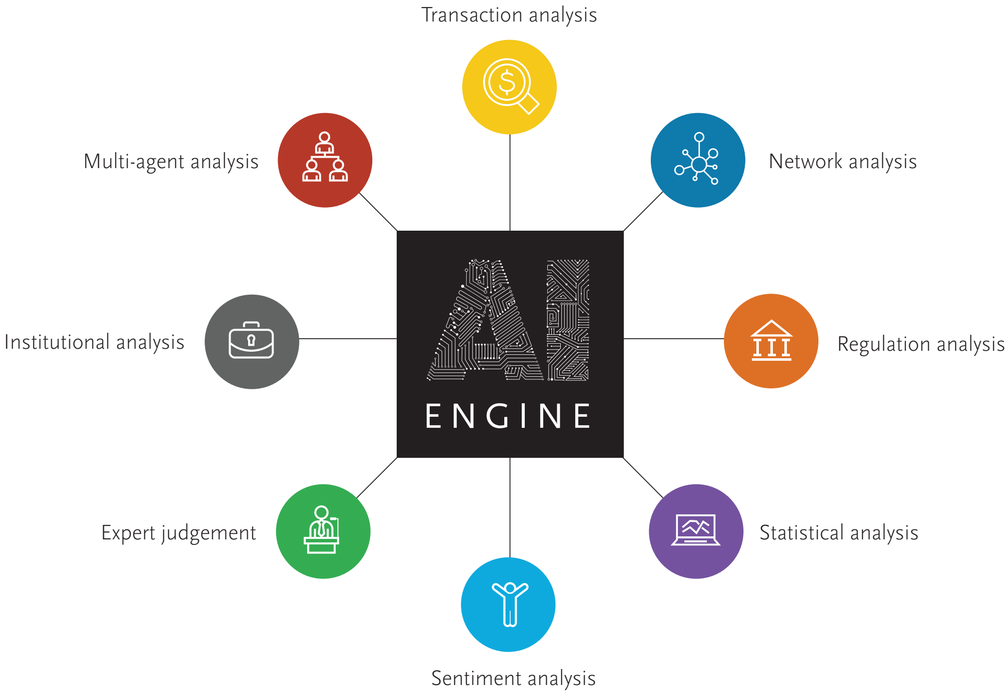

Our approach

Our approach in six steps:

1. Conceptualize the system as a set of partially overlapping production networks
2. Identify the industries that are central to the functioning of the system
3. Identify key agents, institutions and processes within these industries
4. Search for indicators within these industries that could indicate the onset of known systemic risks
5. Undertake horizon scans to identify early warning signals of black swans ("unimaginable" systemic risks)
6. Assess the ability of the industries to withstand, recover and adapt to these risks. [4, 5, 23]

(For the sake of simplicity, the six steps are outlined in a sequential manner. However, they are in fact largely iterative and parallel processes, where also the "later" processes inform the "earlier" ones. For example, the extent to which a production network is able to withstand an bankruptcy of a keystone company (step 6), partially helps to determine the amount of resources that should be allocated to searching for (additional) indicators of this risk (step 4).)

Steps 1-3: Identify systemically important industries within the larger economic network, as well as keystone companies those industries.

Step 1

Conceptualise an economy or sector as a set of overlapping industries

Step 2

Does the industry pose a systemic risk to the economy in case it fails or goes into decline?

Step 3

If this company fails or goes into decline, can the other companies in the industry still thrive or at least survive?

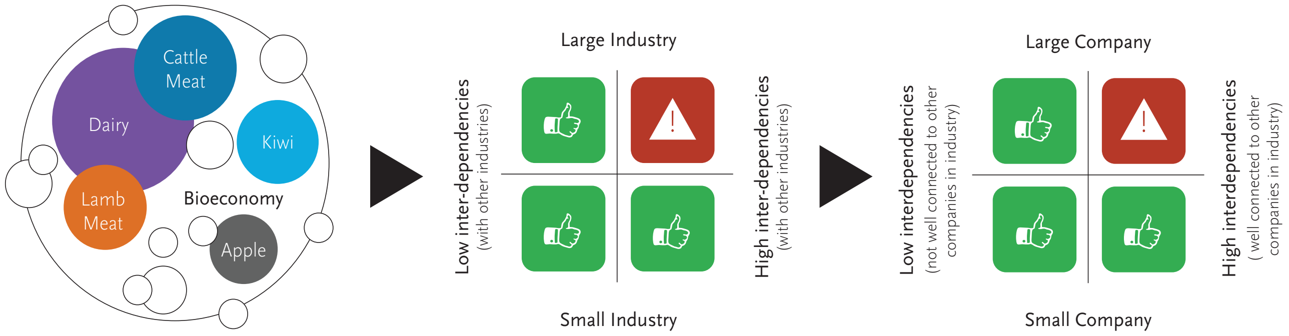
Cattle Meat
Dairy
Kiwi
Lamb Meat
Bioeconomy
Apple
Large Industry
Low inter-dependencies (with other industries)
High inter-dependencies (with other industries)
Small Industry
Large Company
Low interdependencies (not well connected to other companies in industry)
High interdependencies (well connected to other companies in industry)
Small Company

Step 4: Detect early warning signals of known systemic risks affecting keystone industries and companies

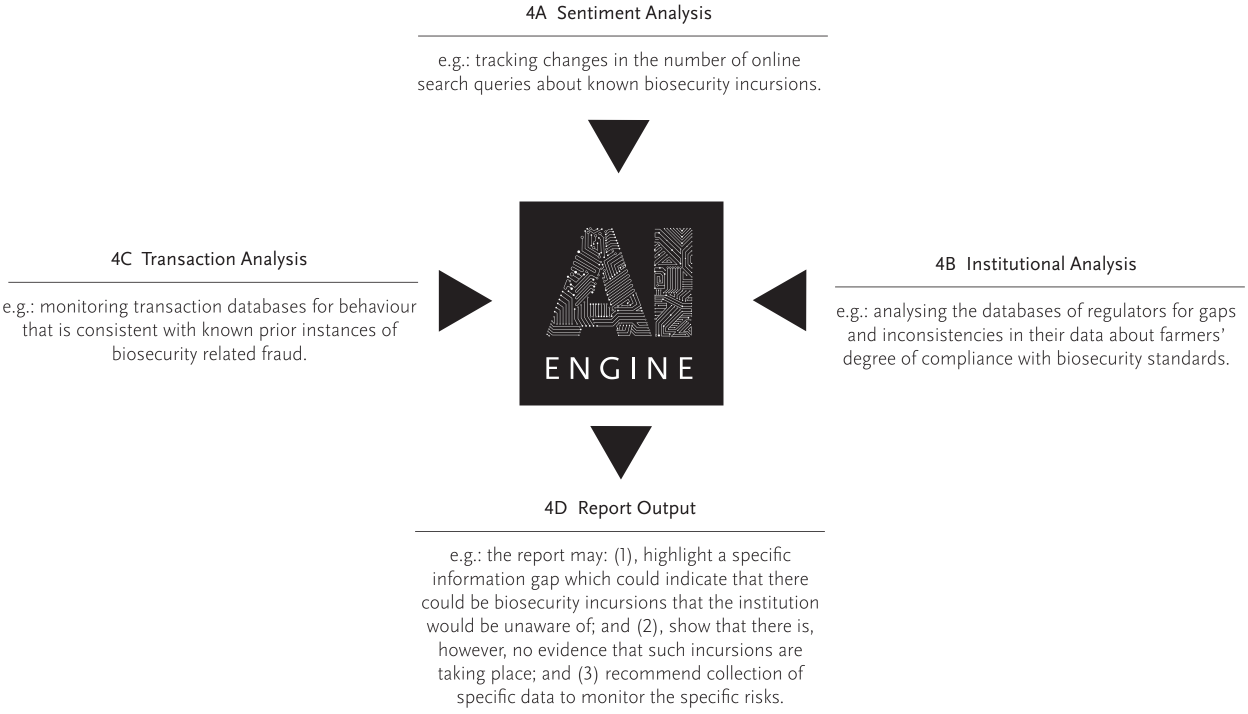

4A Sentiment Analysis

e.g.: tracking changes in the number of online search queries about known biosecurity incursions.

4C Transaction Analysis

e.g.: monitoring transaction databases for behaviour that is consistent with known prior instances of biosecurity related fraud.

4B Institutional Analysis

e.g.: analysing the databases of regulators for gaps and inconsistencies in their data about farmers' degree of compliance with biosecurity standards.

4D Report Output

e.g.: the report may: (1), highlight a specific information gap which could indicate that there could be biosecurity incursions that the institution would be unaware of; and (2), show that there is, however, no evidence that such incursions are taking place; and (3) recommend collection of specific data to monitor the specific risks.

Step 5: Undertake horizon scans to detect 'black swans' earlier

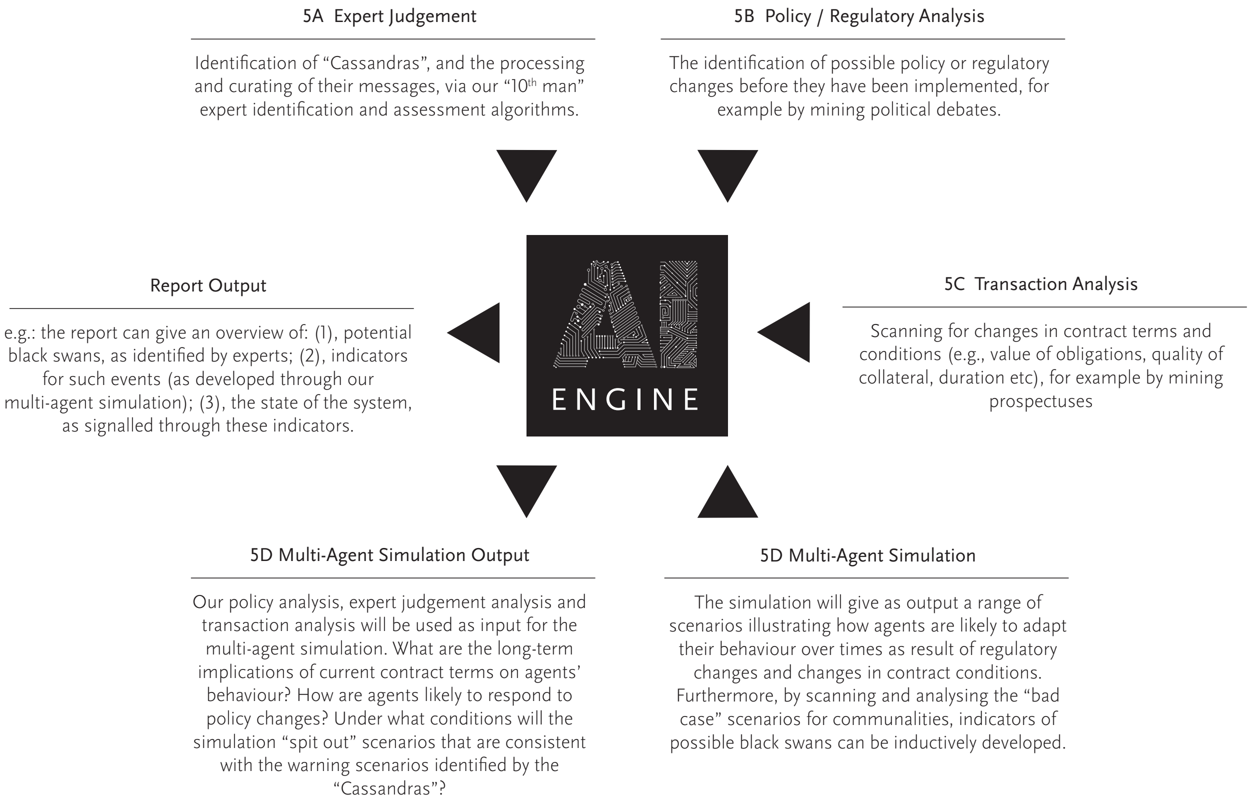

5A Expert Judgement

Identification of "Cassandras", and the processing and curating of their messages, via our "10th man" expert identification and assessment algorithms.

5B Policy / Regulatory Analysis

The identification of possible policy or regulatory changes before they have been implemented, for example by mining political debates.

Report Output

e.g.: the report can give an overview of: (1), potential black swans, as identified by experts; (2), indicators for such events (as developed through our multi-agent simulation); (3), the state of the system, as signalled through these indicators.

5C Transaction Analysis

Scanning for changes in contract terms and conditions (e.g., value of obligations, quality of collateral, duration etc), for example by mining prospectuses

5D Multi-Agent Simulation Output

Our policy analysis, expert judgement analysis and transaction analysis will be used as input for the multi-agent simulation. What are the long-term implications of current contract terms on agents' behaviour? How are agents likely to respond to policy changes? Under what conditions will the simulation "spit out" scenarios that are consistent with the warning scenarios identified by the "Cassandras"?

5D Multi-Agent Simulation

The simulation will give as output a range of scenarios illustrating how agents are likely to adapt their behaviour over times as result of regulatory changes and changes in contract conditions. Furthermore, by scanning and analysing the "bad case" scenarios for communalities, indicators of possible black swans can be inductively developed.

Step 6: Assess the robustness, resilience and adaptive capabilities of keystone industries and companies

Would industry be able to withstand the known systemic risks and black swans identified in the earlier steps?

How quickly would the industry recover if one (or more) of these events were to occur?

Could the industry adapt to new circumstances after the event (if necessary)?

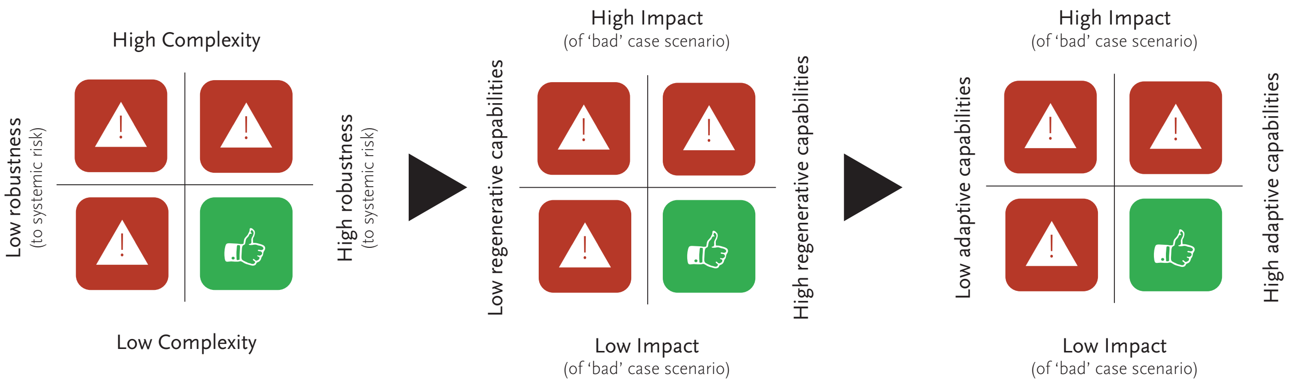
High Complexity
Low robustness (to systemic risk)
High robustness (to systemic risk)
Low Complexity
High Impact (of 'bad' case scenario)
Low regenerative capabilities
High regenerative capabilities
Low Impact (of 'bad' case scenario)
High Impact (of 'bad' case scenario)
Low adaptive capabilities
High adaptive capabilities
Low Impact (of 'bad' case scenario)

Applying our approach - selected case studies

At present, we are further developing and applying our approach in a specific context: the detection and assessment of potential biosecurity incursions affecting the New Zealand agri-food sector. This includes:

- An assessment of the structural robustness/fragility of the production networks that make up the New Zealand agri-food sector against biosecurity incursions.
- An evaluation of data, systems and approaches used by key agents (e.g., large companies, well-connected companies, as well as the suppliers of such companies) and institutions (e.g., regulators) to prevent such incursions.
- The detection of anomalies in the behavior of the companies within the sector that may indicate the onset or outbreak of an incursion affecting human, animal or plant health. The focus here lies predominantly on identifying anomalies within the social aspects of the system rather than the biophysical aspects (for example, by monitoring for suspicious behavior of companies, by scanning for experts that give unheeded warning of possible incursions, etc.). That is, the focus lies mainly on identifying anomalies that suggest companies are taking advantage from structural fragility in the network or from monitoring failures at the institutional level.

Initially, we are developing the AI engine specifically for the detection of biosecurity hazards within the context of the pastoral industries. Subsequently, we will start the process of adapting and "training" the algorithms for the detection of biosecurity hazards within other parts of the agri-food sector as well.

More broadly, our approach, which is predominantly concerned with identifying frictions and cracks in the system that are caused by human and institutional failures, should also be applicable to the detection and assessment of other types of risks, and within other types of contexts, than biosecurity hazards within the agri-food sector.

The principles-based nature of our approach will allow us to, over time and with effort, also adapt the algorithms to a more diverse set of contexts (e.g., the detection of systemic risk within the financial services industry).

Assessing network fragility to biosecurity incursions by means of Artificial Intelligence

The AI engine will monitor for changes in the structure of food production networks, especially looking for: (1), changes in the structure that make it more difficult to detect biosecurity incursions; and (2), changes in the structure that will increase the rate or speed of transmission of an incursion if one were to occur. In general, these are changes that make the network supporting the physical flows of products, services or interactions amongst the agents more densely connected, more complex and more homogenous [6, 24], and changes that make the network supporting the flows of information amongst the agents less connected and more opaque [25].

For example, a denser physical network means that an incursion will spread more easily across agents. Particular concerns in this context include: an increase in the size of potential "super-spreaders" (larger nodes at central points in the network, such as larger traders and other intermediaries), better connected "super-spreaders" (e.g., traders that cover a larger geographical area, that connect a wider range of farmers to a wider range of processing companies or locations, etc.), more linkages across previously isolated hubs (e.g., more trading between intermediaries, for example for inventory management purposes), etc.

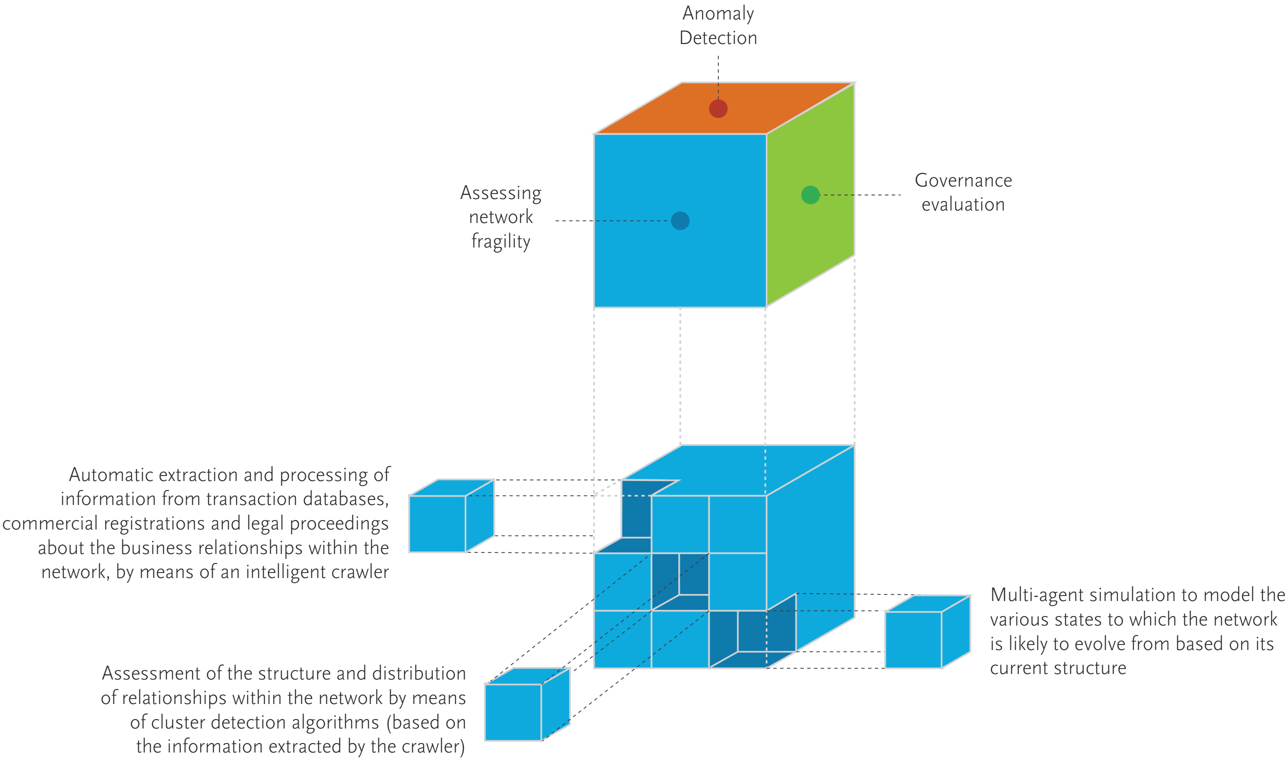
Anomaly Detection
Assessing network fragility
Governance evaluation
Automatic extraction and processing of information from transaction databases, commercial registrations and legal proceedings about the business relationships within the network, by means of an intelligent crawler
Multi-agent simulation to model the various states to which the network is likely to evolve from based on its current structure
Assessment of the structure and distribution of relationships within the network by means of cluster detection algorithms (based on the information extracted by the crawler)

Detecting biosecurity related fraud through Artificial Intelligence

* For example, we will try to detect anomalies in the behavior of potential super-spreaders that are suggestive of biosecurity incursions, as well as in the behavior of the agents that form part of the hub surrounding such key nodes.

We will use pattern recognition algorithms to determine whether any anomalies that have been detected in such hubs by the lower-level algorithms are just random noise, or a potential indication of a biosecurity incursion that should be taken seriously and investigated further.

Because we will pool and analyze data from a wide range of data sources, both structured and unstructured, we will be able to detect patterns that would be missed by more traditional approaches.

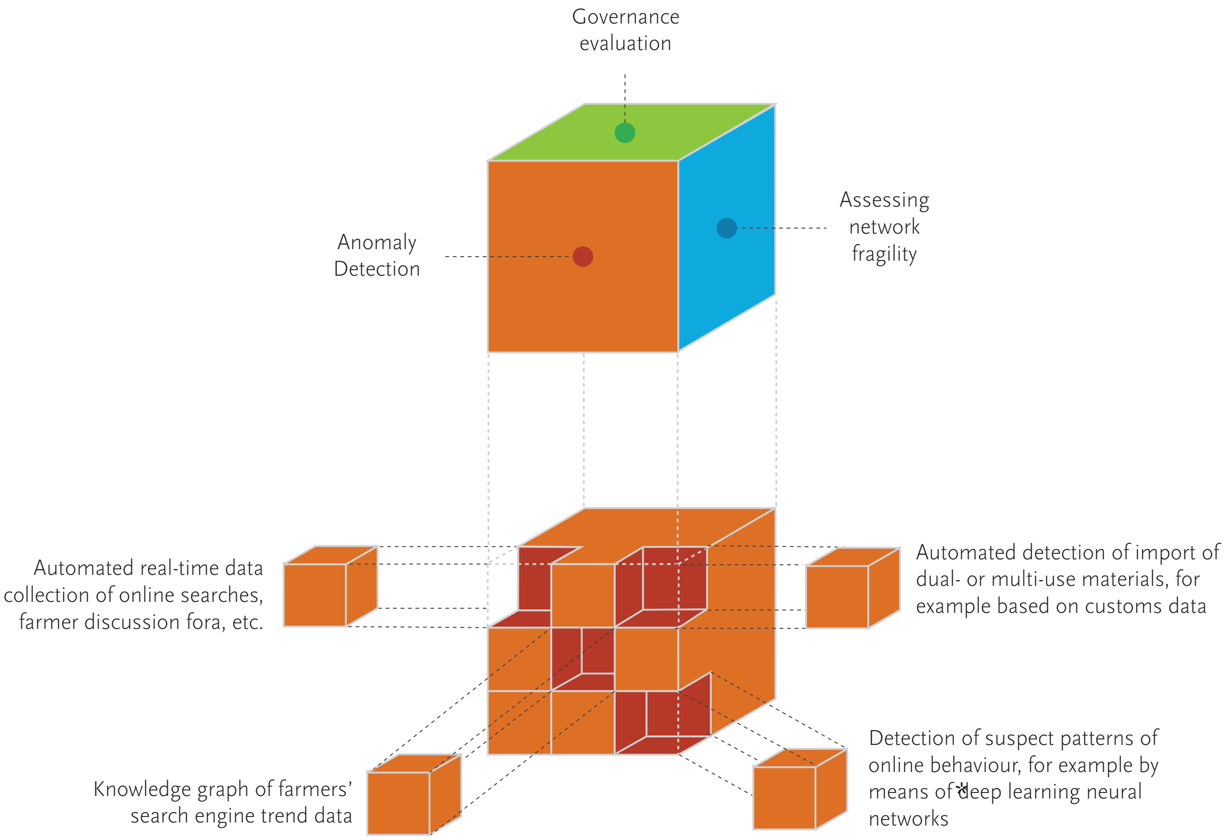
Governance evaluation
Anomaly Detection
Assessing network fragility
Automated real-time data collection of online searches, farmer discussion fora, etc.
Automated detection of import of dual- or multi-use materials, for example based on customs data
Knowledge graph of farmers' search engine trend data
Detection of suspect patterns of online behaviour, for example by means of *deep learning neural networks

When events or patterns are detected that pass a certain predefined trigger or threshold, the users of our system will receive an automatic alert. This alert will be in the form of a multi-layered report, where the user can drill-down to obtain more details about the event that triggered the alert.

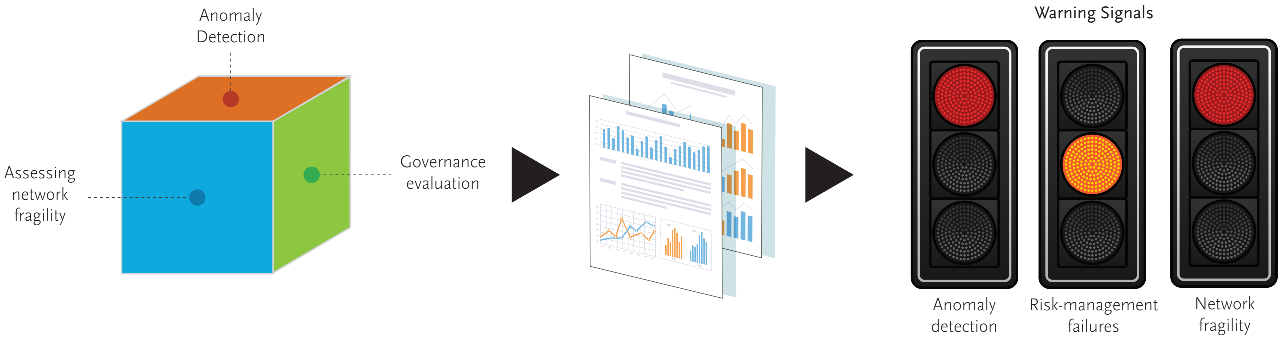

Identifying biosecurity hazards earlier than through machine learning

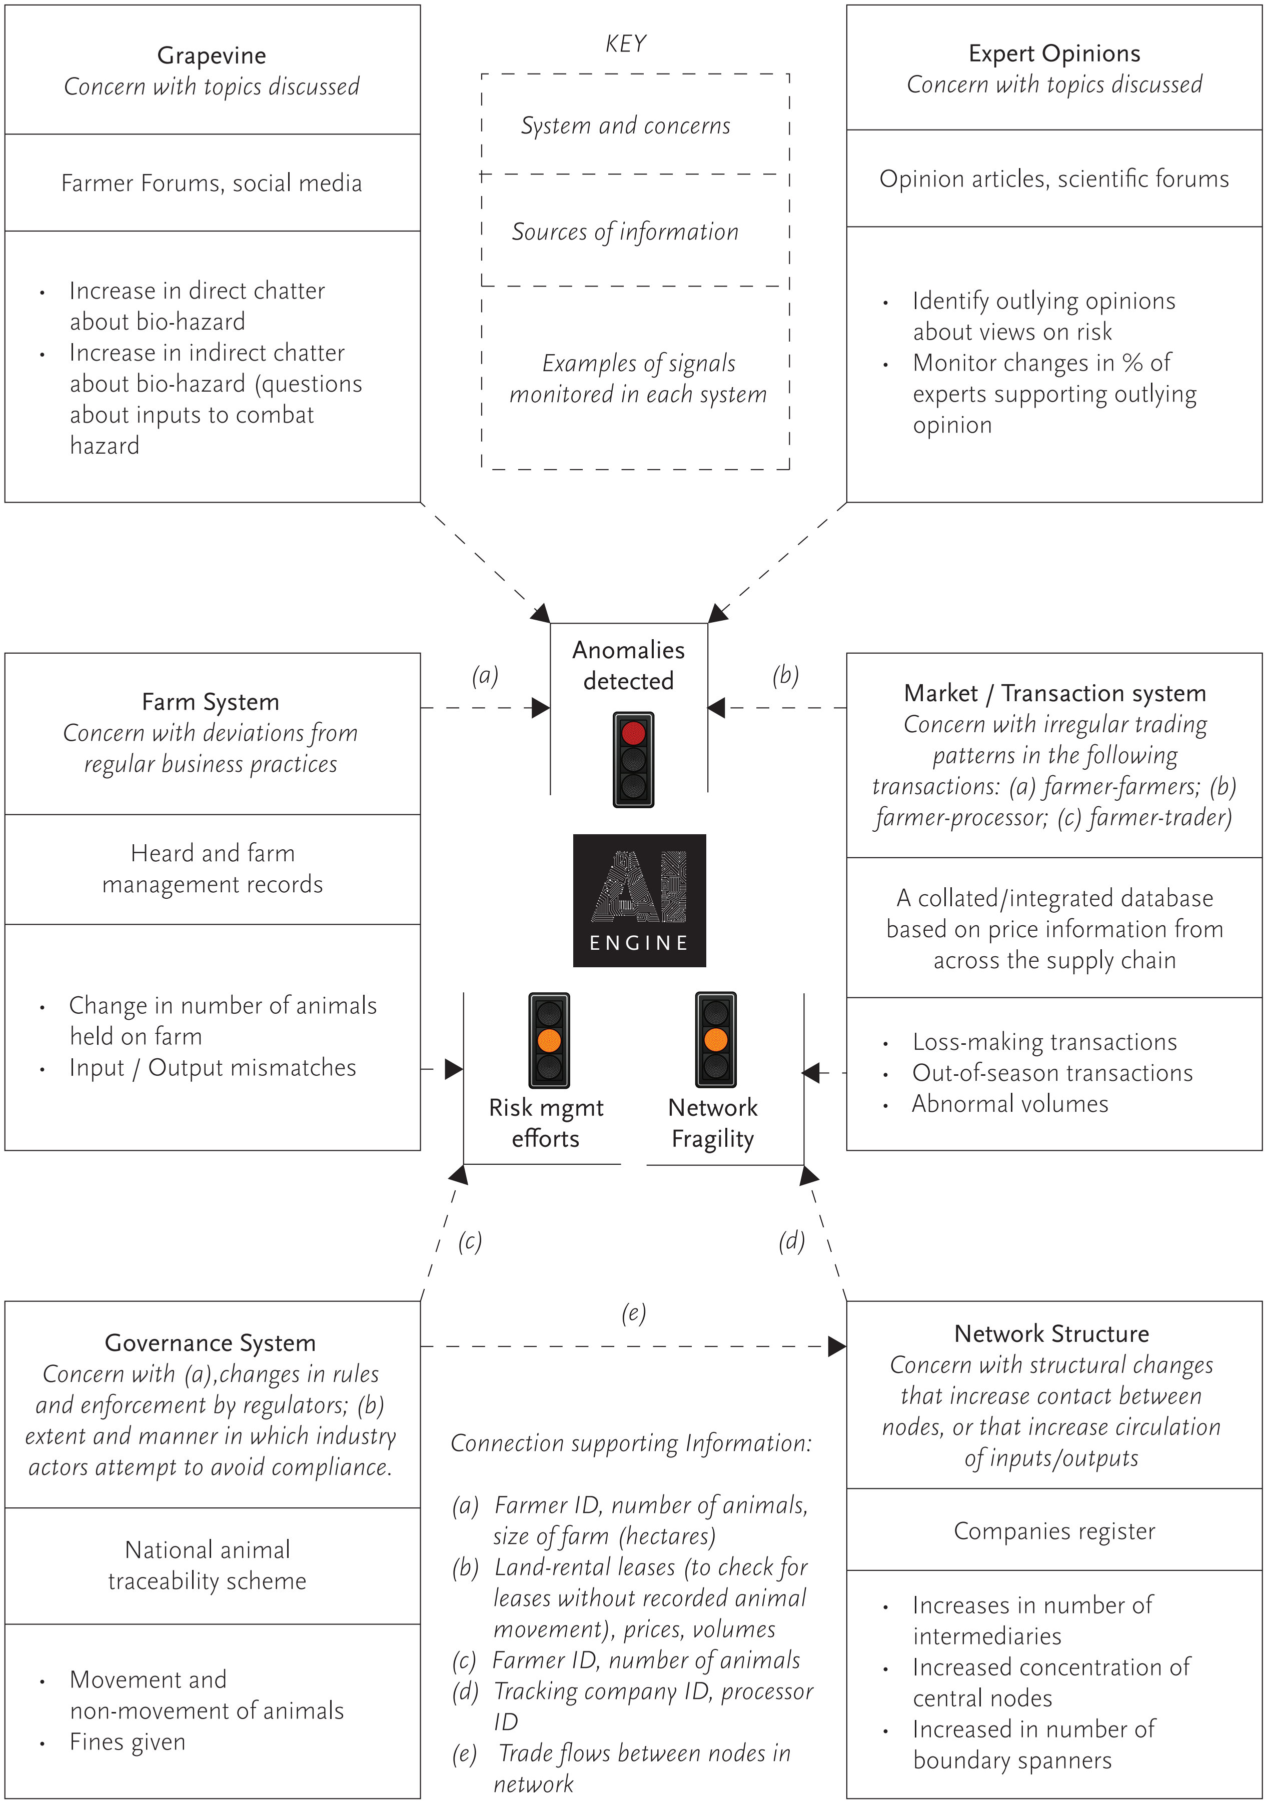

KEY

System and concerns

Sources of information

Examples of signals monitored in each system

Grapevine
Concern with topics discussed

Farmer Forums, social media

- Increase in direct chatter about bio-hazard
- Increase in indirect chatter about bio-hazard (questions about inputs to combat hazard

Expert Opinions
Concern with topics discussed

Opinion articles, scientific forums

- Identify outlying opinions about views on risk
- Monitor changes in % of experts supporting outlying opinion

Farm System
Concern with deviations from regular business practices

Heard and farm management records

- Change in number of animals held on farm
- Input / Output mismatches

Market / Transaction system
Concern with irregular trading patterns in the following transactions: (a) farmer-farmers; (b) farmer-processor; (c) farmer-trader)

A collated/integrated database based on price information from across the supply chain

- Loss-making transactions
- Out-of-season transactions
- Abnormal volumes

Governance System
Concern with (a),changes in rules and enforcement by regulators; (b) extent and manner in which industry actors attempt to avoid compliance.

National animal traceability scheme

- Movement and non-movement of animals
- Fines given

Network Structure
Concern with structural changes that increase contact between nodes, or that increase circulation of inputs/outputs

Companies register

- Increases in number of intermediaries
- Increased concentration of central nodes
- Increased in number of boundary spanners

Connection supporting Information:

(a) Farmer ID, number of animals, size of farm (hectares)
(b) Land-rental leases (to check for leases without recorded animal movement), prices, volumes
(c) Farmer ID, number of animals
(d) Tracking company ID, processor ID
(e) Trade flows between nodes in network

Final thoughts

All of the major systemic risks in recent times, whether it is the 2007-2008 sub-prime crisis or COVID-19, gave off early warning signals. Those signals were not picked up by most regulators and keystone companies in a timely manner, often leading to delayed and insufficient responses. This, in turn, has cost the world economy hundreds of billions of dollars.

An early warning system powered by artificial intelligence (AI), such as the one we have presented in this paper, can help to greatly mitigate systemic risks. How? An AI based system could assist regulators and other types of stakeholders in scanning for risk in a more continuous, automated and comprehensive manner, amongst others by:

- Making it time- and cost-efficient for regulators to gather and process a wide-range of data points about social and other types of systems, through natural language processing algorithms, computer vision and voice recognition. For example, in the context of financial risk, this would allow regulators to rapidly analyze thousands of prospectuses and contracts in order to spot changes in the quality of the securities that are floated in the primary market.
- Greatly enhancing the ability of regulators to identify non-obvious patterns and relationships in such systems, by means of machine learning principles and algorithms more generally. For example, with regard to biosecurity fraud, this would help regulators to better identify suspicious transactions, such as an unexpected increase in chemical purchases in a certain agricultural region.
- Drastically speeding-up the process by which new models about the functioning of such systems can be developed, tested and adjusted (through a combination of automated data collection and machine learning). For example, in the context of the outbreak of a contagious virus, this would help health boards to quickly test the accuracy of existing epidemiological models against the latest data about the spread of the virus.

The assistance that our system could provide would help regulators to:

1. Monitor a much wider range of potential early warning signals of systemic risk;
2. Identify actual warning signals significantly earlier;
3. Provide more up-to-date measures of the fragility of the system against these risks.

In more plain English, it would give regulators an opportunity at an early stage to nip potential systemic risks in the bud or to limit or mitigate effects if preventive action is not (fully) successful.

The costs of developing an early warning system are negligible compared to the potential cost savings that could be obtained by being better prepared against large scale, harmful events such as M.bovis or COVID-19. These types of events are unlikely to stop occurring over the coming decade.

If anything, we will likely see more large-scale crises going forward, as countries are becoming ever more connected and depended on each other. As a result, disruptions in supply chains are more likely to have cross-border effects, biosecurity threats and contagious diseases are more likely to spread across continents, and problems in a domestic financial market are more likely have implications for the global economy.

Faced with this uncertain future, investing in the development of a more sophisticated early warning system is about as safe a bet as they come.

References

I. Haldane, A. and R. May (2011). Systemic risk in banking ecosystems. Nature, 20 (469), 351-355.

II. Kane, E. (2010). Redefining and containing systemic risk. Atlantic Economic Journal, vol. 38, 107-120.

III. Colchester, J. and Filia, F. (2018). Financial complexity and nonlinear dynamics. Complexity Labs and Fasanara Capital.

IV. Meadows, D. (2008). Thinking in systems. Chelsea Green Publishing: Vermont (Canada).

V. IRGC (2018). Guidelines for the governance of systemic risks. Lausanne: International Risk Governance Center.

VI. Haldane, A. (2009). Rethinking the financial network. Speech delivered at the Financial Student Association, Amsterdam, in April 2009.

VII. Walker, W., Marchau, V. and D. Swanson (2010). Addressing deep uncertainty using adaptive policies. Technological Forecasting and Social change, 77, 917-923.

VIII. Taleb, N. (2007). The black swan: The impact of the highly improbable. Random House: New York (NY).

IX. Boston, J. (2016). Anticipatory governance: How well is New Zealand safeguarding the future?. Policy Quarterly, 12 (3), 11-24.

X. Wilensky, U. and M. Resnick (1999). Thinking in levels: A dynamic systems approach to making sense of the world. J. of Science Education and Tech., 8 (1), 1-19.

XI. Ackoff, R. and J. Pourdehnad (2001). On misdirected systems. Systems Research and Behavioral science, 18, 199-205.

XII. Raworth, K. (2018). Doughnut economics: Seven ways to think like a 21st-Century Economist. Chelsea Green Publishing: London (England).

XIII. Boston, J. (2017). Safeguarding the Future: Governing in an Uncertain World. BWB Texts: Wellington (New Zealand).

XIV. Kou, G., Chao, X., Peng, Y., Alsaadi,. F. and E. Herrera-Viedma (2019). Machine learning methods for systemic risk analysis in financial sectors. Technological and Economic development of Economy, 25 (5), 716-742.

XV. Leo, M., Sharma, S. and K. Maddulety (2019). Machine learning in banking risk management: A literature review. Risks, 2019, 7, 1-22.

XVI. Bauguess, S. (2017). The role of Big Data, Machine Learning and AI in assessing risks: A regulatory perspective. SEC Key Note address: OpRisk North America (New York, NY).

XVII. CFA Institute (2019). AI Pioneers in investment management. An examination of trends and use case of AI and big data technologies in investments. CFA Institute: Charlottesville (VA).

XVIII. Rasekhschaffe, K. and R. Jones. (2019). Machine learning for stock selection. Financial Analysts Journal, 74 (3), 70-88.

XIX. Guerra, S., Silva, T., Tabak, B., Penaloza, R. and R. de Mrianda. (2016). Systemic risk measures. Physica A, 442, 329-342.

XX. Bisias, D., Flood, M., Lo, A. and S. Valavanis (2012). A survey of systemic risk analytics. Annual Rev. of Finan. Econ. 4, 255-296.

XXI. Li, S., Wang, M., and J. He (2013). Prediction of banking systemic risk based on support vector machine. Mathematical problems in engineering, Article ID: 136030.

XXII. Cont, R., Moussa, A. and E. Santos (2016). Network structure and systemic risk in banking systems. In J. Fouque and J. Langsam (Eds.): Handbook on systemic risk, pp. 327-368. Cambridge University Press: Cambridge (England).

XXIII. Gramlich, D., Miller, G., Oet, M. and S. Ong (2010). Early warning system for systemic banking risk. Banks and Bank systems, 5 (2), 199-211.

XXIV. Jackson, M. and D. Lopez-Pintado (2019). Diffusion and contagion in networks with heterogenous agents and homophily. Network science 1(1), 49-67.

XXV. Wever, M., Wognum, N., Trienekens, J. and O. Omta (2012b). Managing transaction risks interdependent supply chains. Journal on Chain and Network Science, 12 (3), 243-260.

XXVI. Williams, D. 2018. Why NAIT failed – and what's being done to fix it. Newsroom: www.newsroom.co.nz/2018/05/15/109635/why-nait-failed-and-whats-being-done-to-fix-it# (accessed on 21-10-2019).

XXVII. Pavez, I. Codron, J., Lubello, P. and M. Florencio (2019). Biosecurity institutions and the choice of contracts in international fruit supply chains. Agricultural systems, 176.