Open, Contingent, Adaptive and Reactive Resilience

– using ABM and other tools to facilitate our collective survival in an uncertain world

Bruce Edmonds[1] [0000-0002-3903-2507] and Vahid Yazdanpanah[2] [0000-0002-4468-6193]

[1] Centre for Policy Modelling, Manchester Metropolitan University, UK
bruce@edmonds.name
[2] Agents, Interaction and Complexity Research Group, University of Southampton, UK
v.yazdanpanah@soton.ac.uk

"The truth is, that we are just not very good at forecasting the character of the risks that we will encounter."
– Oliver Letwin, UK Minister for Preparedness, 2010-16 *[1]*

Abstract. This work puts forward a particular perspective on achieving resilience in the face of deep uncertainty due to complex systems. In principle, various ways of lacking knowledge about the behaviour of the enveloping system and its agents, can make traditional planning approaches ineffective. This calls for resilient strategies that can adapt to achieve their goals, or simply survive, in view of unexpected shocks. We distinguish different types of situations in which different strategies might be helpful and propose requirements for a new strategy suitable for deep uncertainty, what we call *Open, Contingent, Adaptive and Reactive Resilience* (OCARR).

Keywords: policy making, disaster avoidance, adaptation, deep uncertainty, ABM, visualisation, monitoring, risk, adaptative governance, resilience

1 Introduction

After the Ebola outbreak of 2013-16, the UK "Minister for Preparedness" (Oliver Letwin), set up a small unit in the Civil Contingencies Secretariat (CCS) whose purpose was to scan for unanticipated threats to the UK, so that the government could react in a timely manner [1]. For example, it allowed the government to consider whether to ban travel from countries affected by the Zika virus outbreak. Sometime during Theresa May's government this unit was disbanded. In October 2016, the government held an exercise to assess the United Kingdom's preparedness and response to a pandemic influenza outbreak, called "Exercise Cygnus" [2]. Although it may be that some of the recommendations that came from this exercise were taken on board, others (such as the capacity of the social care sector to receive patients from hospital and the need to better organise stocks of resources, such as Personal Protective Equipment) were not. The report after the exercise [2] said: "*the UK's preparedness and response, in terms of its* *plans, policies, and capability, is currently not sufficient to cope with the extreme demands of a severe pandemic that will have a nationwide impact across all sectors.*" It is easy to identify actions that, *in hindsight*, would have been better prepared the UK for the Covid-19 pandemic, however this illustrates a basic problem for governments (and other institutions) — that of how to mitigate the impact of the unexpected.

This paper presents an approach to achieving some level of resilience when in deeply uncertain situations. This approach is called "Open, Contingent, Adaptive and Reactive Resilience" or "OCARR" for short. This approach is not entirely novel in terms of its overall direction nor some of its components, however we do introduce high-level characteristics to support identifying the resilience problem at hand and bring more tools to the table — in particular Agent-Based Modelling (ABM) — and thus contribute towards making it more practical to implement. Such a perspective for problem identification and structuring allows for understanding the type of uncertainty we faced and, applying the appropriate strategies for mitigation.

To motivate OCARR, we first look at some other strategies and discuss their likely scope – that is, the kind of circumstance where they each might be effective. We then present OCARR and then discuss some of the tools available for making this approach work, including ABM. We finish with a discussion of some of the remaining challenges to implementing OCCAR.

2 Traditional Planning – when the effects of policy can be meaningfully predicted

This approach inherits from Engineering, and comprises of the following stages:

1. Identify the goals of the policy;
2. Design ways to measure the extent to which any outcome would meet these goals;
3. Identify the possible strategies one might apply to meet these goals;
4. Predict the costs and benefits of each strategy;
5. Choose the strategy that has the highest predicted benefit and least predicted cost.

In essence, such traditional planning and optimisation is appropriate for decision making in domains with a low level of unpredictability – "surprise-free" domains. In such situations, the outcome of implementing a policy and assessing their potential consequences can be reliably predicted. Here, the decision maker relies on being able to measure and anticipate the costs and benefits of implementing different policies (given a relatively small range of available policies) in order to determine the optimal one. Such an approach suits domains in which the main goal is to achieve the best a system can deliver in view of available resources, e.g., in industrial manufacturing. By definition, this approach is not apt for making decisions in domains characterised by uncertainty or where risks have a high level of unpredictability.

We observe that such an approach is encapsulated in the UK government's "Green Book" [3] – a guide for how to appraise and evaluate policies, projects and programmes. As discussed, this strategy only works if the decision problem a decision maker is facing with satisfies a number of conditions. There has to be a manageable number of policies to be tested, the extent that each of these meets the goals has to be measurable and predictable, and the costs and benefits have to be commensurable so one can choose the best option. If conditions do not hold, then other strategies are necessary.

3 The Adaptive Approach – when history repeats itself

The human body's metabolism is arguably one of the most complex systems. However, how such a system behaves in response to the flu virus is known and well-explored. This is the situation when we are interacting with complex but how it responds to available stimuli (e.g., policies and interventions) are predictable, in a sense, the world is complex but repeats similar patterns. The complexity makes the explicit calculation of policy outcomes hard. However, in such settings, it suffices to improve measurable outcomes via adapting our policies in a responsive way. This includes:

1. Trying out various policies either in real-life case studies or trials;
2. Evaluating the efficacy of policies and how the system responded when these policies were tried; and
3. Adapting our policies in a responsive way in the light of those evaluations.

This approach is implicit whenever policies are evaluated – the prediction of consequences is not enough, rather one looks at how they performed *in practice* (e.g. as in [4]). In principle, adaptive approaches in decision making are not aiming to anticipate unknown risks but mainly focus on improving measurable outcomes even when the probability of them occurring is unknown.

Clearly, as proved by how the UK government responded to the Covid-19 pandemic, adaptive approaches with a responsive view are not sufficient for ensuring resilience and building the capacity to handle unknown and unpredictable events. We agree with Oliver Letwin [1] that strategies developed to deal with known threats caused by flu were not enough to deal with the effects of unknown viruses.

4 The Participatory Approach – when one needs to reconcile multiple viewpoints, goals and preferences

Another form of uncertainty in policy making comes from the sheer variety of values, beliefs and viewpoints of the stakeholders involved. Imagine the case that a city council with a limited budget aims to invest in either building a new school, expanding a park in the area, or installing smart traffic management infrastructures in intersections. Such a decision affects a heterogenous set of stakeholders, including the households in the borough as well as public and private sector representatives. Here, the optimality of a policy is not merely of an objective nature but depends on whether (and to what extent) it is aligned with the values and viewpoints of stakeholders. Such situations have sometimes been called "wicked systems" (e.g. in [5]). The main question here is to determine the *collective's preference* based on a fair and justifiable aggregation or negotiation process [6]. This calls for a participatory approach that:

1. Determines a set of policy options with (potentially conflicting) trade-offs and desirability for different stakeholders;
2. Presents the policies to representatives of the stakeholders;
3. Facilitates discussions aimed to improve the understanding of different parties and the decision maker about the distribution of viewpoints and rationales;
4. Applies some method to determine the preferred policy from the collective's point of view (negotiation among parties, preference aggregation etc.)

There is an appendix to the Magenta book [7] that discusses this kind of approach (among others). In our city council scenario, a person might be opposed to improving the traffic infrastructure but then changes their mind after hearing the rationale others have for supporting the idea, or maybe keeps their stance but accepts the final policy as they now have an understanding of the decision-making process in a transparent way and feels consulted by being involved in the decision-making process.

An applicable set of tools here are *problem structuring* methods such as DPSIR (drivers, pressures, state, impact and response model of intervention) [8]. Such tools can help visualize the causal relations and interactions between society and various components that characterize the risk. Specially in the context of environmental risks, DPSIR has been proved to be effective for structuring and understanding different aspects of the risk and for identifying aspects that need (policy) interventions [9].

The participatory approach is not necessarily aiming for an optimal policy from any one stakeholder's perspective, to obtain the most resilient option or to deal with unanticipated risks [10]. The main aim is to improve the acceptability of the policy, to capture uncertainties caused by heterogeneous and potentially conflicting perspectives among stakeholders, and to improve the understanding of those affected by the policy on how the final decision is produced [7].

5 Contingency Planning – when there are known threats but where probabilities cannot be assigned to them

For a long time, at least since the Ebola outbreak, the threat of pandemics was known. However, it was not clear to policy makers how probable a pandemic was – on that would require strict lockdowns in almost all countries, and affect socioeconomic systems drastically. It was a known threat but the probability of its occurrence was unknown but estimated to be low. This shows that for such a class of threats, we lack methods to assign accurate probabilities to the risk (and hence the expected cost). This suggests a strategy of puting in place contingency mechanisms to ensure the resourcefulness of those systems that may be affected.

This class of undesirable events with low (estimated) probability but high effect correspond to what Nasim Taleb calls *black swans* [11], that is, situations that a system rarely faces — and are difficult to learn about — but where their occurrence can dramatically affect the performance, e.g. Global warming reaching a tipping point. For such cases, relying on probabilities is not helpful as it is not possible to gather enough information to accurately characterise it sufficiently. This viewpoint corresponds with the presented perspective in the chapter on Economics in Non-Equilibrium Social Science and Policy [12], and further elaborations by Frank Dignum [13] (namely that the assumption that "*agents are able to gather and process substantial amounts of information*" is unrealistic and harmful for making effective decisions in real-life problems). Specifically, for the above-mentioned class of problems, relying on probabilities has proved to be not only ineffective but the cause of other side effects (e.g., damage caused by mishandling a crisis).

This is not only an issue caused by an expected utility approach that one may adopt for making decisions but is also a result of the computational principle that soley aims for *efficiency*. This principle — to aim to allocate as few resources possible for solving a problem — is recently questioned, e.g., by Moshe Vardi [14] in his proposal to learn from the pandemic experience and look for computational contingency planning methods that violate the efficiency principle but incorporate some level of redundancy, and ensure a degree of resilience.

As a way forward, and to have an appropriate response to similar threats that may be caused by climate change, we ideate applying contingency planning by:

1. Identifying known threats with high potential impacts (not computing expected costs or benefits by multiplying with probabilities);
2. Investing in contingency mechanisms to mitigate the threat and avoid its consequences; and
3. Reserving redundant resources to adapt with such highly improbable scenarios in case mitigation mechanisms fail and the threat becomes an inevitability.

This approach is effective when the risk can be easily identified with potentially high (but unknown) severity and/or probability and relatively low contingency cost [15]. Note that contingency planning can be implemented in a participatory manner. To that end, we envisage the applicability of multiagent incentivisation and coordination methods [16]. Such methods can be a base for taking into account how concepts such a social norms and institutions can be used to nudge the behaviour of key decision makers towards the adoption of a policy and for ensuring a level of policy acceptability in a society. The other line of research that can be applied for building consensus and aggregating individual judgements (e.g., on a given policy) is computational social choice and voting theory [17]. Using judgement aggregation techniques, one can obtain an unbiased view on the choice of a collective based their knowledge about individual preferences. This can support the implementation of policies that are in line with what the stakeholders prefer and achieve more widely acceptable interventions.

In this case, ABMs can help the identification process, e.g., by simulating how a risk may pan out over time and gathering insight about its scale, hence categorising it as a high or low impact one. In addition, ABM is also useful for analysing and visualising the effect of policies and providing views of the available data to enable policy actors to effectively "drive" a resilient policy and move to the implementation phase. For example, an ABM might be analysed using Machine Learning techniques to try and identify the different parameter regimes whereby certain risks emerge, and then the ABM might be simplified and experimented on within those regimes to understand broadly how and why those outcomes could emerge there. To that end, resilience governance needs to take advantage of more recent data-driven tools to inform model-driven techniques (such as ABM) in order to develop indicators for developing situations that might require a policy response.

6 OCARR – for living with deep uncertainty

If one is trying to deal with truly uncertain events – that is events which are not predictable, in the sense of being able to assign any meaningful probability to them, that do not repeat (so are not learnable) and are too many to do contingency planning for – then none of the above strategies are very helpful. However, that does not mean that nothing useful can be done. It is this kind of situation that OCARR is designed for.

The OCARR strategy is essentially threefold, so we will present this conceptually first, before looking how to make this a practicality by using particular tools and approaches. It can be abstracted to the following steps:

1. Identify as many possible future events and their subsequent trajectories as one can, regardless of their perceived likelihood.
2. Analyse the trajectories that would pose a significant threat – understanding how they might come about and develop.
3. Implement "indicators" or "early warning systems" that give an indication of when these threatening trajectories might be emerging at an early stage, so that the adaptive response can be enacted rapidly.

There are a number of points and caveats to be made about each of these.

In this kind of situation one can never identify *all* of the possible events and trajectories, nor indeed exclude all those that may seem possible but, for reasons that are not apparent, would not actually occur. The aim for this stage is to identify a good number of the non-obvious possibilities, erring on the side of inclusivity, and add them to the obvious ones. Thus, contributions of possibilities from a variety of inputs, models, stakeholders etc. is desirable here.

The judgement of what trajectories pose a significant threat is quite a political one. Questions such as "*who would this be a threat to?*" and "*to which values and goals does this pose a threat?*" need to be answered and answering them will not be a purely technical matter, but are of a sociotechnical nature as they necessarily involve a political process. Values and viewpoints are important here [6].

Given what is identified are mere possibilities, it can be hard to justify implementing contingency measures for all of them if this involves any significant level of cost. Under OCCAR, instead of contingency measures, it is suggested that 'early warning' indicators might be a more practical response. The 'indicators' might be as simple as abstracting some statistic from some existing data, setting up some new data-collection process, or designing a visualisation to show when some more complex pattern is starting to emerge. This is not full-blown contingency planning due to the costs of this. They should be designed so that: (a) how they relate to the data is fairly transparent, so those looking at them know what they mean and (b) they give an easily read indication of the extent to which the threatening trajectories might be arising, so that policy actors can effectively 'drive' policy – as one drives a car, reacting to events to avoid collisions. Each indicator needs to be relatively cheap to implement, because there may need to be a large number of them corresponding to some unlikely outcomes.

7 Reasoning Methods and ABM Tools to Support OCARR

Here, we elaborate on reasoning methods, computational tools, and problem structuring techniques with the potential to support the three phases of OCARR (Figure 1).

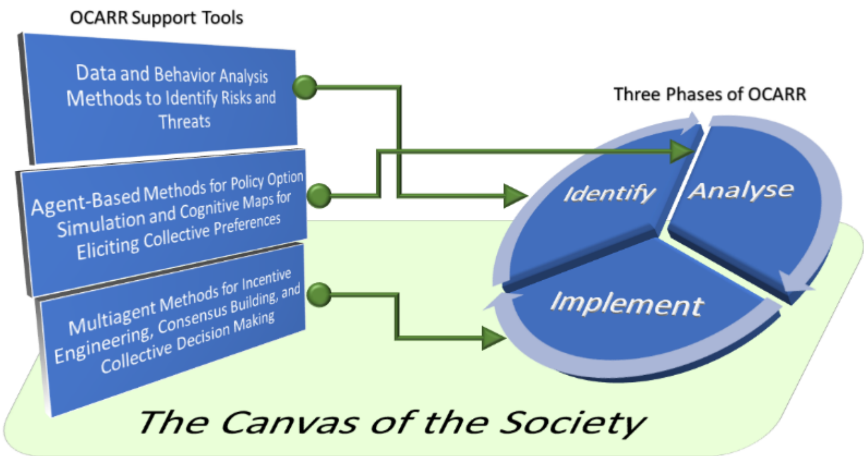

Fig. 1. An illustration of an implementation of OCARR

As discussed, the first phase of OCARR is to identify future possibilities thoroughly and gather information about their features. This identification leads to a better understanding of the category of the risk we are facing and then to use the appropriate tools in the upcoming OCARR phases. Here, Agent-Based Models (ABMs) are applicable as they can expand the set of considered possibilities by including those that are difficult to infer on *a priori* grounds, namely the emergent outcomes. In other words, they can help identify how things might go wrong [18]. Why these can occur might become obvious *after* inspecting the ABMs but not necessarily – the processes that contribute might be so complex that they are not very amenable to human understanding. At this stage of OCARR it is important that as many different kinds of sources to suggest possibilities as possible are included for the maximum possible coverage. ABMs are well suited to this possibilistic approach, due to the large number of aspects of complex systems for which we lack good evidence (e.g., how actors react).

For the identification of features, the class of data-driven methods to mine the gathered data from the behavior of the system can be helpful in specifying agents in a simulation by associating processes with actors, and consequently inform behavioral rules in ABMs [19]. The task of revealing possible outcome trajectories are less than if one wants to know what will probably happen since one can be more inclusive as to the included mechanisms in the model, and worry less about available evidence or noise.

Here, we follow Pearl [20] and argue that data analysis is apt for understanding associations but are not to support an intervention analysis or counterfactual reasoning, both key for risks identification and policy analysis.

Note that some tools have the potential to contribute to more than one phase of OCARR. In the implementation phase of OCARR, we face the challenge of incentivising the choice of indicators – to have a reasonably fair process for building consensus among the society members about this. Without such mechanisms and tools to support these steps, even if the society is well aware of a threat, we cannot ensure our resilience against it because we would not agree on which threats to watch for. In particular, complex ABMs are not suitable for presentation to stakeholders or policy actors since even modellers find them hard to understand. Rather, the understanding of emergent processes that come from them might be used to design transparent indicators.

The tools provided could be as simple as techniques to visualise the data coming in from the indicators. These could be designed so that possible emergent processes, analysed and understood in the analysis phase can be shown in a dramatic but transparent manner. Such visualisations are much easier to understand than complex ABMs and can be a common point of reference for political discussion as how to react. However, care must be taken to keep the channels of information and evidence open, so that there is not an over-emphasis on what is displayed.

8 Open Problems for ABM and Social Simulation Research

As discussed, problems we aim to address using OCARR are (not limited to those caused by complexities in general, e.g., due to the presence of various stakeholders but) mainly caused by failing to identify the *type of uncertainty* one may face in a complex system. Indeed, OCARR's focus on understanding the nature of uncertainties relates to but is distinguishable from the state of art on simulation-based policy making in complex systems [21] [22]. In principle, while they focus on simulating how different policies may affect a society of stakeholders with the aim to find optimal policies, we introduce a problem structuring/identification step that precedes the (simulation-based) policy appraisal phrase. To that end, OCARR guides the process of policy analysis and provides input to modelling steps (e.g., to how the environment should be modelled in view of the type of uncertainty at hand and how to define the agents' behavioural rules).

No process can identify all unexpected threats – there is always a chance that some factor that is completely new or out-of-context appears or becomes significant. In addition, the fact that threats that lead to deleterious outcomes, as well as efforts to mitigate them, are complex in nature and might make possible other undesirable possibilities inevitable. However, the body of work on ABM and social simulation research has the capacity to support various phases of OCARR. In this section, we elaborate on more concrete challenges in OCARR and specifically focus on the class of open problems for the social simulation community.

Embedding Qualitative Aspects: As discussed, OCARR looks at the problem of addressing uncertainties not only as a system *analysis* problem but in combination with the *identification* and *implementation* phase. This is due to the fact that analysing the optimality of a policy, in particular in uncertain situations, is entangled with its implementation and the character of the threat we are faced with. While in studies that merely focus on analysis-oriented use of ABMs, one may focus only on using quantitative data sources, in OCARR one can also incorporate qualitative aspects such as the subjective preference of stakeholders and the tacit knowledge of experts about social, cultural, and political dynamics. To that end, we envisage the development and integration of methods to inform ABM techniques using qualitative evidences [23]**.**

Integrating Live Inputs: ABMs traditionally focus on a dynamic process of analysing the interrelation between agents but using some form of static data. This is to gather data and then start the simulation process. This causes common limitations (1) that there may not be enough data to support the simulation process and cover various aspects of the analysis and (2) that the real system at hand may change during the simulation and analysis phase. Such caveats have a more drastic effect when we are focusing on uncertain and unanticipated threats. A way forward for the ABM community is to invest in the so-called *live simulations* [24] as simulation settings in which the behavioural rules are being updated lively (in contrast to being fixed at start point of the simulation process). This is similar to virtual worlds and game settings in which the inputs gathered during the simulation are fed back to the simulation itself. Such a dynamic social simulation perspective has the potential to capture uncertainties that are unknown at the time of initial data gathering and is an appropriate response to the dynamic nature of problems OCARR aims to address.

Informing Policy Implementation and Incentivisation Mechanisms: It is unrealistic to expect that all categories of stakeholders and policy makers follow OCARR given potential conflicts of interest and the political context of a society. A challenge for our community is on the need for intermediary modules to allow the integration of ABM with incentivisation mechanisms. In particular, how different incentivisation mechanisms affect the behaviour of agents in an ABM and how such a behaviour change feeds back and affects the effectivity of the incentives. For instance, giving discounts and subsidies on an insurance policy that supports those living on flood-prone areas may nudge the behaviour of society to invest in insurance (as a climate change adaptation tool). However, this behaviour change may in itself provide a social incentive and motivates others to migrate to live in riverbanks which increases the risk of being flooded. Such dynamics and interconnections between threat analysis and policy implementation necessitates focusing on ABMs that are aware of interconnected and compound risks [25].

9 Concluding Remarks

We presented an approach that focusses on identifying the nature of threats and aims for ensuring resilience against unanticipated possibilities. We deem that our community has the potential to address parts of the problem.

As an immediate future work, we aim to focus on design and development of threat identification tools. Such tools need to be dynamic as the nature of threat may change through time. For instance, a threat for which we once suggested traditional planning may be identified as one that needs contingency planning – e.g., based on new evidences. Clearly situations do not always fall neatly into one of the delineated categories discussed, in which case the OCARR approach needs to work in parallel with other strategies and adapt its implementation phase. Also, instead of a linear process to identify a threat and then merely focus on analysis and implementation, OCARR has to be circular and incorporate re-identification, and re-analysis processes to expand our identification of possibilities and our understanding of them. This opens interesting research directions on how to integrate other approaches, e.g., the participatory approach and contingency planning, in a dynamic way.

As argued by Oliver Letwin [1], the way we analyse policies and reason about the appropriateness needs to reflect the characteristics of the threat in question. In some domains, e.g., when we are facing uncertainty about the nature of a disease and the virus that caused it, we may reason cautiously and prescribe the medicine to cure the disease even if only under one eventuality we can imagine that the patient is suffering from that disease (we say cautiously because only under one eventuality we may be correct). In other cases, for instance in ascribing liability to those who were accountable for addressing a threat — but failed to do so — we may reason more sceptically and apply sanctions to a person, only if under all worlds we are uncertain about she is blameworthy. To deal with real-life threats, policy makers need a range of approaches, including OCARR to deal with them and then operational tools to aid the understanding of the threat, and drive their policy accordingly.

Without such dynamic tools for threat identification, behaviour analysis, and policy implementation, our societies will suffer from the political context of blame avoidance and end up in short-termism. With this proposal, we motivated the need for a shift in approach and positioned OCARR in relation to ongoing work on ABM approaches.

Acknowledgements

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through the Trustworthy Autonomous Systems Hub (EP/V00784X/1).

Bibliography

[1] BBC Radio 4, "Lessons on a crisis, episode 1: preparedness," BBC, 16 March 2021. [Online]. Available: https://www.bbc.co.uk/programmes/m000t6ql. [Accessed 23 May 2021].

[2] Public Health England, "Annex B: Exercise Cygnus Report," 20 10 2016. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/927770/exercise-cygnus-report.pdf. [Accessed 25 5 2021].

[3] HM TREASURY, "The Green Book: appraisal and evaluation in central government," 3 12 2020. [Online]. Available: https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-governent. [Accessed 25 5 2021].

[4] HM Treasury, "Magenta Book - Central Government guidance on evaluation," March 2020. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/879438/HMT_Magenta_Book.pdf. [Accessed 29 5 2021].

[5] H. W. Rittel and M. M. Webber, "Dilemmas in a general theory of planning," *Policy sciences,* vol. 4, pp. 155--169, 1973.

[6] L. n'i Aodha and B. Edmonds, "Some pitfalls to beware when applying models to issues of policy relevance," in *Simulating Social Complexity*, 2017, pp. 801--822.

[7] CECAN – the Centre for Evaluating Complexity Across the Nexus, "Magenta Book 2020 - Supplementary Guide: Handling Complexity in Policy Evaluation," HM Treasury, 2020.

[8] European Environmental Agency, "Environmental indicators: Typology and overview," European Environmental Agency, 1999.

[9] S. Mehryar, R. Sliuzas, A. Sharifi, D. Reckien and M. van Maarseveen, "A structured participatory method to support policy option analysis in a social-ecological system," *Journal of environmental management,* vol. 197, pp. 360--372, 2017.

[10] N. A. Jones, P. Perez, T. G. Measham, G. J. Kelly, P. d'Aquino, K. A. Daniell, A. Dray and N. Ferrand, "Evaluating participatory modeling: developing a framework for cross-case analysis," *Environmental management,* vol. 44, p. 1180, 2009.

[11] N. N. Taleb, The black swan: The impact of the highly improbable, Random house, 2007.

[12] P. Ormerod, "Economics," in *Non-Equilibrium Social Science and Policy: Introduction and Essays on New and Changing Paradigms in Socio-*

Economic Thinking, Springer, 2017, pp. 19-44.

[13] F. Dignum, "Review of Non-Equilibrium Social Science and Policy: Introduction and Essays on New and Changing Paradigms in Socio-Economic Thinking (Understanding Complex Systems)," Journal of Artificial Societies and Social Simulation, 2018.

[14] M. Y. Vardi, "Efficiency vs. resilience: What COVID-19 teaches computing," *COMMUNICATIONS OF THE ACM,* 2020.

[15] C. R. Sunstein, Averting Catastrophe: Decision Theory for COVID-19, Climate Change, and Potential Disasters of All Kinds, NYU press, 2021.

[16] V. Yazdanpanah, S. Mehryar, N. R. Jennings, S. Surminski, M. J. Siegert and J. van Hillegersberg, "Multiagent Climate Change Research," in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 2020.

[17] F. Brandt, V. Conitzer, U. Endriss, J. Lang and A. D. Procaccia, Handbook of computational social choice, Cambridge University Press, 2016.

[18] B. Edmonds and L. Aodha, "Using agent-based modelling to inform policy–what could possibly go wrong?," in *Multi-Agent Systems and Agent-Based Simulation (MABS-2018)*, 2018.

[19] F. Squazzoni, J. G. Polhill, B. Edmonds, P. Ahrweiler, P. Antosz, G. Scholz, É. Chappin, M. Borit, H. Verhagen, F. Giardini and N. Gilbert, "Computational Models That Matter During a Global Pandemic Outbreak: A Call to Action," *Journal of Artificial Societies and Social Simulation,* vol. 23, p. 10, 2020.

[20] J. Pearl, "Theoretical impediments to machine learning with seven sparks from the causal revolution," *arXiv preprint arXiv:1801.04016,* 2018.

[21] P. Davidsson, F. Klügl and H. Verhagen, "Simulation of complex systems," in *Springer Handbook of Model-Based Science* , Cham, Springer,, 2017, pp. 783-797.

[22] N. Gilbert, "Agent-based social simulation: dealing with complexity," in *The Complex Systems Network of Excellence*, 2004, pp. 1-14.

[23] B. Edmonds, "Using qualitative evidence to inform the specification of agent-based models," *Journal of Artificial Societies and Social Simulation,* vol. 18, p. 18, 2015.

[24] S. Swarup and H. S. Mortveit, "Live Simulations," in *19th International Conference on Autonomous Agents and Multiagent Systems*, 2020.

[25] C. A. Phillips, A. Caldas, R. Cleetus, K. A. Dahl, J. Declet-Barreto, R. Licker, L. D. Merner, J. P. Ortiz-Partida, A. L. Phelan, E. Spanger-Siegfried and others, "Compound climate risks in the COVID-19 pandemic," *Nature Climate Change,* pp. 586--588, 2020.