Towards a Comparative Analysis of Policy-making Distribution in Hierarchical Organisations - the Epidemic Management Case Study

Jeanne Jolivet, Julien Raymond, Ada Diaconescu
Telecom Paris, LTCI, IP Paris, Palaiseau, France
name.surname@telecom-paris.fr

Louisa Jane Di Felice
Universitat Autonoma de Barcelona & ICTA Barcelone, Spain
louisajane.df@gmail.com

***Abstract*—Large sustainable systems, e.g. human organisations or socio-technical systems, must manage large-scale crises. This implies prompt responses from policy-makers, including an initial reaction followed by repeated adjustments. Policies must be (re)defined under two opposing difficulties: i) the amount of relevant information to collect and process for effective policies; and ii) the urgency required for their implementation. Most large systems handle this issue via hierarchical organisations, distributing information and control across levels. Still, determining appropriate distributions of policy-making across levels remains a non-trivial problem. This position paper explores the nature of this issue via an illustrative case study: epidemics management. By comparing different policy-making distributions – more or less centralised – it highlights the importance of local details for effective local management and for the ensuing global effects. While still exploratory, this initial study brings to the fore the interplay between local and global policies and the impacts of relevant information for scalable and sustainable management.**

***Index Terms*—hierarchical organisations, adaptive control, policy-making, epidemic management, multi-scale authority**

I. Introduction

Large sustainable systems, including human organisations and technical systems, must manage large-scale crises. Here, policy-makers must provide an initial response followed by repeated adjustments as consequences unfold. A major difficulty in (re)defining appropriate policies stems from two conflicting requirements: i) the necessity to collect and process relevant information for determining effective policies – which takes time; and, ii) the urgency of implementing effective policies.

Most large systems address this issue via hierarchical structures, which meet timing constraints by distributing information processing across levels. Here, top-level policy-makers ensure system-wide coherence based on a global overview; while lower levels provide more adapted local solutions based on narrower, yet locally-relevant information. Importantly, hierarchical organisations of information processing do not necessarily require hierarchical authority, e.g. higher authority can be assigned at the top or bottom, [1]. Notably, this is inline with E.Ostrom's principles for sustainable institutions [2], including (P7) recognition of rights to organise locally; and (P8) nested institutions involving multiple governance layers.

While this general multi-level scheme addresses the aforementioned conflict, determining an appropriate distribution of policy-making across various governance levels remains a non-trivial problem (e.g. [3], [4]). In principle, concentrating authority at the top ensures faster global coherence, while pushing it towards the bottom ensures more effectiveness "on the ground". As local and global effects are highly interrelated (e.g. micro-macro feedbacks [5]) a difficult compromise must be found between fine- and coarse-grained management to ensure system sustainability overall.

In previous work, we proposed a conceptual basis for studying this local-vs-global management interplay [1]. The current paper explores the nature of this issue more concretely via a simulated case study: epidemics management. A crisis situation was selected to illustrate the urgency of policy-making and the need to adjust policies on-the-fly, as their effects are uncertain. The study relies on simulation via an agent-based model (ABM), implemented via NetLogo. The model is inspired from classic epidemic propagation literature [6]. It includes several communities (e.g. cities, villages) with different population densities and resources (e.g. stores, hospitals)– Fig. 1. When an epidemic breaks, policies are defined and adjusted at global (macro) and/or local (micro) levels, depending on the epidemic extent. Policy-making is assessed based on epidemic duration and casualty number.

This study aims to bring to the fore the fact that different governance distributions may significantly influence system behaviour and outcomes, both locally and globally. It does not claim to show how exactly such influence may play out in a particular system. Hence, we do not aim to provide a realistic epidemic model– a vast area on its own. Instead, the presented ABM serves as an experimental tool for exploring multi-scale management dynamics, via an intuitive example. It helps to identify the main generic factors favouring local and global policy-making, irrespective of concrete system instances.

Notably by comparing the results of different policy-making distributions – more centralised (macro) or decentralised (micro) – this study highlights the importance of detailed local information for local management, which in turn impacts global effects. Previous studies also highlighted latency as a key factor in multi-scale system dynamics [7]. Certainly, changing agent behaviour or parameter values in the presented ABM may lead to different outcomes. Still, information access and latency remain determining factors.

This initial analysis provides generic insights into the interplay between local (micro) and global (macro) policies. It opens promising research directions for addressing this key management problem in scalable and sustainable systems.

II. Approach Overview

A. Case Study and Qualitative Model Description

The study aims to analyse infection propagation across four adjacent communities (Fig. 1), with different population densities and available resources (stores and hospitals) – e.g. a densely populated 'big city' where inhabitants have access to nearby resources; and a sparsely populated 'countryside' with resources farther away from most inhabitants (cf. Table III). Inhabitants go shopping every day to a store selected based on its distance from home.

The epidemic case study starts with one infected individual and spreads as individuals get in contact while shopping in the same store. After a period of being infected, an individual may either: recover, and be immune for a while; or require hospitalisation. In the latter case, they go to the nearest hospital. If admitted, they recover after a while; otherwise, e.g. no available room, they do not survive (Fig. 3).

To contain the epidemics and limit casualties, policy-makers issue lockdown constraints, in terms of the maximum distance from home (R_{max}) that individuals can move within. This implies that some stores may be within or outside the lockdown area, for each policy. This, in turn, makes a policy more or less 'acceptable', causing individuals to either comply or transgress it. E.g., if too few stores are available within the confined area individuals may decide to go to a 'forbidden' store that provides a desired item, despite the lockdown. Complying agents who wish to access a 'forbidden' store will stay home instead. Based on the observed acceptability of their policy, authorities may adjust the lockdown distance, so as to limit transgressions and disease propagation.

Governance is organised hierarchically, with two main scales. The global scale (macro) represents the entire population over the four communities. The local scales (micro: $\mu_1..\mu_4$) represents the four separate communities. We study the epidemics propagation depending on two policy-making distribution schemes. In both cases, the general lockdown strategy is issued at the global scale, as above. First, a lockdown distance (R_{max}) is set based on the proportion of infections. Then, the policy is repeatedly adjusted (larger or shorter R_{max}) depending on how the epidemics evolves (e.g. new infections or hospital occupancy); and also on the ensuing acceptability rates and transgression rates.

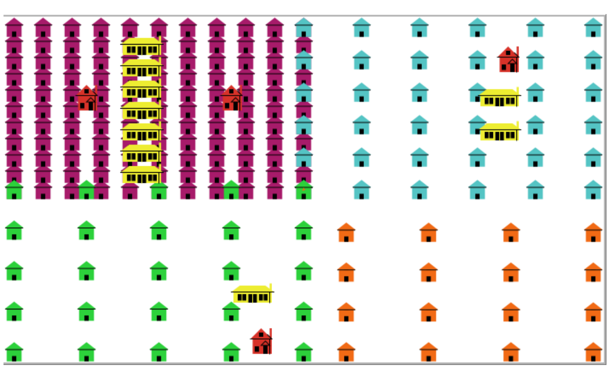

Fig. 1. Community model: 1) big city (purple); 2) city (blue); 3) small city (green); 4) countryside (orange) – with stores (yellow) and hospitals (red); and 1-3 inhabitants/house

The two considered governance variants differ as follows. In a more centralised variant (referred to as 'macro') the global level sets the initial policy and then adjusts it by assessing the entire managed region (i.e. all four communities at once). In a more decentralised variant (micro), local governance resets policies separately for each community, depending on the specific infection and acceptability status. Hence, the two variants only differ in how they assess the current situation (i.e. infection and acceptability) before adjusting the policy. This results in different policy instances (actual R_{max} values across communities), while the general policy-making strategy stays the same. This corresponds to a governance scheme where the top levels set the general strategy, e.g., to ensure global coherence and fairness; and the lower levels reify it to match local constraints. Further differentiation can be explored in future work, by allowing local governance to also adapt the global policies – e.g. customising the adjustment rate of R_{max} depending on local factors.

B. High-level Analysis of Policy-driven Behaviour

In the studied scenario, authorities aim to limit population infection over the short term (e.g. to avoid hospital overloads). Their strategy is to limit individuals crossing each-other by constraining their movement distance from home. We consider p_t the part of the population who leave their home at time t. This value can be decomposed into two further ones: c_t, the complying population part who leave by respecting the restriction; and x_t, the transgressing population part who leave despite the restriction because its acceptability is low.

$$p_t = c_t + x_t \tag{1}$$

Let i be an agent, with $i = 1..N_a$, and N_a the total number of agents (395 in our case); and let j be a store, with $j = 1..N_s$, N_s the total number of stores (10 in our case). Then, at each time, we consider $\phi(i,j)$ the probability that an agent i wants to access store j, such that $\sum_{j=1}^{N_s} \phi(i,j) = 1$, $\forall i$. $d(i,j)$ is the distance between store j and the home of agent i. $R_{max,t}$ is the lockdown distance imposed by authorities at t and $acc_{i,t}$ is the acceptability of agent i at t. We partition stores into two subsets: $J_{IN,i,t} = \{j \in [1, N_s] | d(i,j) \leq R_{max,t}\}$, the stores within the lockdown area of agent i at t; and $J_{OUT,i,t} = \{j \in [1, N_s] | d(i,j) > R_{max,t}\}$ the stores outside this area. We can then estimate c_t and x_t as follows:

$$\begin{aligned} c_t &= \frac{1}{N_a} \sum_{i=1}^{N_a} \left[\sum_{j \in J_{IN,i,t}} \phi(i,j) \right] \\ x_t &= \frac{1}{N_a} \sum_{i=1}^{N_a} \left[(1 - acc_{i,t}) \sum_{j \in J_{OUT,i,t}} \phi(i,j) \right] \end{aligned} \tag{2}$$

Thus, as $\sum_{j\in J_{IN,i}} \phi(i,j) + \sum_{j\in J_{OUT,i}} \phi(i,j) = 1$, we get:

$$p_t = 1 - \frac{1}{N_a} \sum_{i=1}^{N_a} \left[acc_{i,t} \cdot \sum_{j\in J_{OUT,i,t}} \phi(i,j) \right] \qquad (3)$$

This indicates that the proportion of moving individuals p_t can only be limited by reducing the proportion of frustrated individuals j_{OUT} who decide to transgress. This implies finding a trade-off between the lockdown distance $R_{max,t}$, which lowers moving agents if they comply; and the ensuing $acc_{i,t}$, which increases moving agents who decide to transgress.

On the one hand, the factor $\sum_{j_{OUT}} \phi(i,j)$, representing the movement constraint on individual i, increases as $R_{max,t}$ decreases. Thus, lower R_{max} would increase lockdown and minimise movement p_t. Yet, on the other hand, the second factor $acc_{i,t}$, representing acceptability, decreases when low $R_{max,t-k}$ is experienced in the past (it does not depend on current $R_{max,t}$). Hence, if authorities impose too restrictive constraints (i.e. low $R_{max,t}$) they risk high disappointment and future transgressions – thus rendering constraints ineffective. Furthermore, regaining acceptability takes time, hence necessitating long periods of low lockdown and risking further infections. In short, lockdown must be more restrictive to constrain agent movement, yet not too restrictive to risk transgression. This trade-off is all the more crucial as acceptability features an inertial effect, taking time to recover.

III. Model Specification

To specify the agent-based model, we rely on ODD standard protocol: Overview, Design concept and Detail [8].

A. Overview

1) Purpose: The purpose of the agent-based model (ABM) is to us help analyse how information (i.e. who knows what when), flowing through a multi-scale system, influences the effectiveness of decisions and actions taken at different scales, in terms of the resulting system dynamics and goal achievement. We focus here on the effectiveness of policy-making distribution across different scales, within a social governance system, managing an epidemic situation. Specifically, the ABM allows comparing the effects of more central policies (macro), based on coarser-grained global information, with respect to more local policies (micro), based on finer-grained, regional information.

2) States variables and scales: The model comprises tree hierarchical governance scales: individuals, communities and global. Individuals feature the state variables in Table I:

Communities include a set of individuals, of hospital beds and of stores; and R_{max} for the micro case. The global scale holds an aggregate of community information (e.g. sum or average). There are two levels of policy-making:

- **Micro-scale**: each community decides on its containment measures (R_{max}), based on the number of infected people and the ability to hospitalise them in the community.
  Macro-scale: A global lockdown measures is decided depending on the global number of infected people and the community's ability to accommodate them.

var. name	var. values
my-community	positive integer $\in \{1,2,3,4\}$ in our case
my-house	positive integer: house id
status	enum.: 'susceptible', 'infected', 'immune'
hospitalized	boolean: yes/no
movement	$integer \in \{0,1\}$
target	positive integer: store id or house id
acceptability	real value $\in [0,1]$
transgressing	$integer \in \{0,1\}$

TABLE I
Individual state variables

3) Process overview and scheduling: Time is simulated in discrete steps and follows a daily cycle. Each cycle is split into 'day' and 'night', each consisting of an equal number of steps (`midday`). Each day, each individual picks a random time to (Fig. 2):

- Select a target to move to (Fig. 2)
- Update acceptability: if the target is outside the lockdown zone, it decreases, otherwise it increases, by a `step-acceptability` amount;
- Move to target, unless already on it;
- Infect someone else (*infectiousness* probability, Fig.3);
- Return home (unless already home).

Each day, governance levels (micro or macro):

- Collect status information (Algorithm 1) and update the containment measure R_{max} (Algorithm 2).

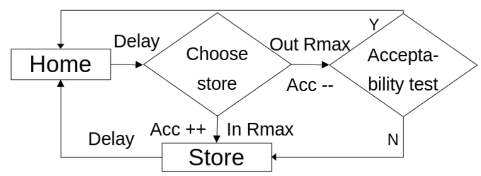

Fig. 2. Individual target selection. Acceptability (Acc) decreases when "choose store" is "Out of R_{max}"

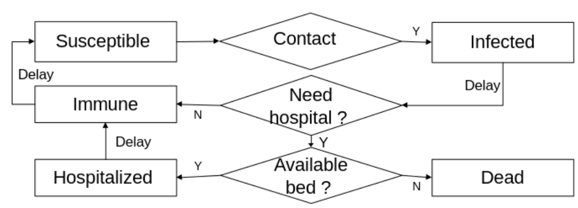

Fig. 3. Individual state diagram.

Algorithm 1: collect_Status_Information

Require: $Rmax, infected, population, acceptability$
  if Authority = macro-scale **then**
    inf ←sum (infected)
    pop ←sum (population)
    acc ←avg (acceptability)
    R_{max} ← Update_containment(inf, pop, acc)
  end if
  if Authority = micro-scale **then**
    for k in communities **do**
      inf ←infected[k]

$pop \leftarrow$ population[k]
$acc \leftarrow$ acceptability[k]
$R_{max} \leftarrow$Update_containment(inf, pop, acc)
end for
end if

Algorithm 2: update_Containment

Require: $inf, pop, acc, Rmax$
if $inf > limit\_infection \times pop$ **and** $acc > limit\_acceptability$ **then**
$R_{max} \leftarrow R_{max} - step\_down\_containment$
else
$R_{max} \leftarrow R_{max} + step\_up\_containment$
end if

B. Design Concept

Adaptation. Authorities (community or global) adapt lockdown policies according to the number of infected individuals and acceptability status. They aim to keep infected agents below 5% and acceptability over 60%. R_{max} is adjusted accordingly, at each step, by adding or removing $sigma = 35$.

Interaction Infected individuals may infect others who are in 'susceptible' states, when they cross paths (e.g. at the store or at home), given an `infectiousness` probability.

Stochasticity. Disease spreading follows several stochastic rules. Individuals randomly choose the store they want to go to each day, following a Gaussian distribution with standard deviation `sigma` with respect to the distance between the store and their home. Individuals go to the target store at a random time during the day and stay there for a random duration (`store-duration`). The duration of infection, of hospitalization and of immunity follow exponential distributions with a corresponding expected value parameter: `immunity-duration`, `infection-duration` and `hospital-duration`, respectively. We also tested Gaussian distributions for these durations, with similar results. When an individual wants to go to a 'forbidden' store, it respects the constraint with a probability equal to its acceptability; or it goes to the store. The individual infected initially is fixed and always at the same location.

Observation. To compare the simulations, the following data are collected daily, for each community:

- lockdown policy R_{max}: maximum travel distance that authorities impose to individuals;
- Number of individuals infected;
- Number of deceased individuals;
- Average acceptability: values closer to 1 imply good acceptance of lockdown policies; and values closer to 0 imply low acceptance and transgression likelihood;
- Average transgression: if an individual wants to go to a store that is outside the lockdown area and their acceptability is low then they may transgress and go to the store anyway (with `acceptability` probability);
- Average movement: if an individual moved during the day, their 'movement' state is set to 1; otherwise to 0.

C. Details : Initialization

The ABM was implemented using NetLogo multi-agent platform and configured as listed in Table II.

Parameter	Value
`midday`	100 steps
`infectiousness`	30%
`hospitalisation-chance`	15%
`immunity-duration`	5 * `midday` steps
`infection-duration`	3 * `midday` steps
`hospital-duration`	3 * `midday` steps
`store-duration`	`midday` / 5 steps
`sigma`	35
`step-up-containment`	`sigma`
`step-down-containment`	`sigma`
`limit-infection`	5%
`limit-acceptability`	60%
`step-acceptability`	3%

TABLE II
SYSTEM CONFIGURATIONS

The governed system contains four communities with different housing densities (Table III). Stores and hospitals were included to ensure approximately 0.03 stores and 0.01 hospital beds per inhabitant[1]. Each house hosts 1, 2 or 3 individuals – preset in an alternate manner across adjacent houses. The epidemic starts by infecting the individual with $id = 34$, in the South-East of community 1 (big city). Infectiousness is set to 30% as this value enables lockdown to influence epidemic spreading – i.e. higher values would spread the epidemic to the entire population and lower values would not spread at all.

community	density	population	houses	hospitals	stores
1. big city	2.2	242	121	2	7
2. city	0.6	72	36	1	2
3. small city	0.4	50	25	1	1
4. countryside	0.3	31	16	0	0

TABLE III
COMMUNITY SETTINGS

IV. EXPERIMENTS & RESULTS

We ran three types of experiments, with 100 repetitions each: macro policy-making (in orange); micro policy-making (in blue); and without lockdown, as reference (in green). Each simulation runs for 300 days (60,000 steps). We observe different infection spreading dynamics in the three experiment types. To compare results, we show average measurements over 100 simulations, for visibility reasons – cf. Figs. 4, 5, 6, 7; and complement these with more detailed graphs showing the discrepancies between communities – cf. Figs. 8, 9, 10. We provide more detailed data, both global and per each community, in Table IV.

Simulations begin with an infection peak – cf. first days in Fig. 4. The macro model reacts, on average, with a stricter lockdown than the micro model (cf lower R_{max} in Fig. 5). This leads to fewer infections in the macro model during the first 25 days. However, this also highly reduces overall acceptability from day 25 (cf Fig. 6) leading to transgressions

[1]https://fr.statista.com

Observed	autho-rity	global	big city	city	small city	country-side
R_{max}	micro	89	73	91	90	100
	macro	93	93	93	93	93
	no	130	130	130	130	130
% infected	micro	4	4	4	3	3
	macro	3	3	3	3	3
	no	3	3	4	3	3
% deceased	micro	32	34	31	29	25
	macro	39	41	38	34	32
	no	49	50	51	47	46
% Acceptability	micro	86	87	80	88	88
	macro	81	94	74	89	67
	no	100	100	100	100	100
% Transgression	micro	7	7	10	5	6
	macro	10	2	14	5	18
	no	0	0	0	0	0
% Movement	micro	82	80	82	84	82
	macro	88	89	87	86	89
	no	99	99	99	99	99

TABLE IV
DATA AVERAGED OVER TIME AND OVER 100 SIMULATIONS, FOR EACH KIND OF AUTHORITY AND EACH COMMUNITY

and risk of increased infections. Hence, the macro governance loosens restrictions (increases R_{max}, Fig.5), leading to a second peak of infections (Fig. 4) and deaths (Fig. 7). This results in higher number of total deaths for the macro strategy compared to the micro approach.

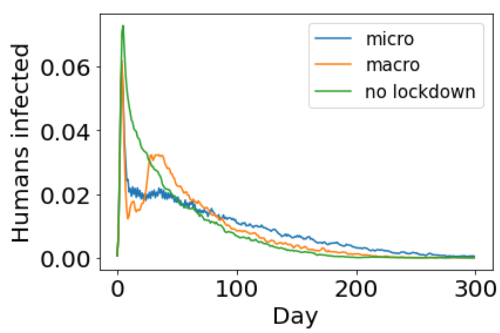

Fig. 4. Total infected individuals

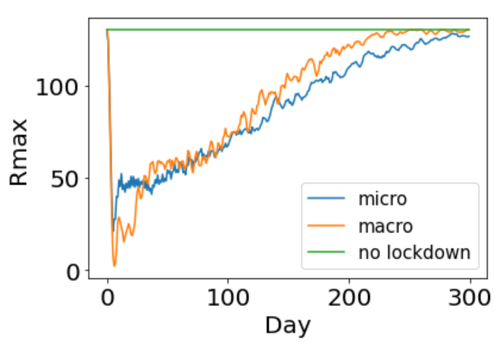

Fig. 5. Lockdown distance R_{max}. The plot for the micro case depicts the average R_{max} of all communities

On the contrary, the micro governance starts on average with a less strict lockdown, resulting in slightly more infections and deaths compared to the macro model, for about the first 25 days. Yet, this less strict lockdown prevents a rebound effect (observed in the macro case). After day 25, the micro approach does not need to lower lockdown level, hence preventing any new infection peaks and deaths. We note that both macro and micro strategies are more effective overall than lack of governance (green line in Fig. 7).

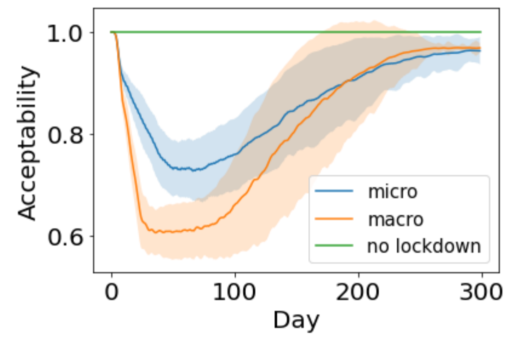

Fig. 6. Average acceptability (with standard deviation)

Moreover, we note that micro-governance selects different lockdown measures for different communities (cf. Fig. 8 and Table IV) to accommodate local specificity and infection dynamics. Large cities impose rather strict lockdown due to higher population density and infections, while rural areas set less stringent lockdown. Still, this leads to a similar acceptability in both urban and rural areas in the micro case. Indeed, as the city has many stores, stricter lockdown has limited effect on acceptability.

Ignoring such local specificity, as in the macro case, leads to transgression and overall ineffectiveness. This can be observed in the transgression dynamics in Figs. 9 and 10. In the micro model, the big city and the countryside have a similar transgression levels (max 2%). However, in the macro case, the countryside have a much higher transgression compared to the micro case (almost double at peak value). This results in increased movement and more contamination.

Overall, the blindness of macro-scale governance that lead to the second infection peak (days 25-50 in Fig. 4) is due to managing distinct communities without local views. Indeed, the big city needs a more restrictive lockdown than the other communities and it represents a large majority of the global population ($> 61\%$). Hence, the macro-scale favors the big city population without noticing that the average acceptability of the smaller communities becomes significantly low, leading to a steep increase in transgressions (Fig. 10). As the macro-scale notices its mistake, it re-adapts its lockdown, following a similar dynamics to that of the micro-scale from then on. Still, this adjustment arrives after many patients had been lost (days 25-50 in Fig. 7).

V. DISCUSSION AND FUTURE PERSPECTIVES

The provided model and simulations highlighted several key aspects for policy-distribution in multi-scale systems: i) the need to access relevant information for defining effective

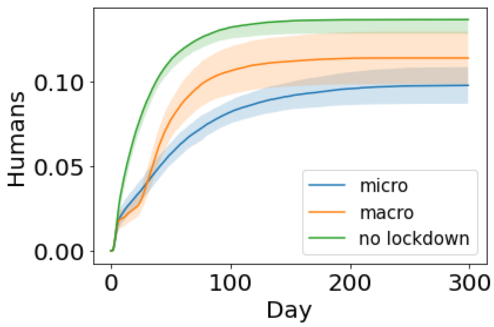

Fig. 7. Cumulative deaths % (with standard deviation)

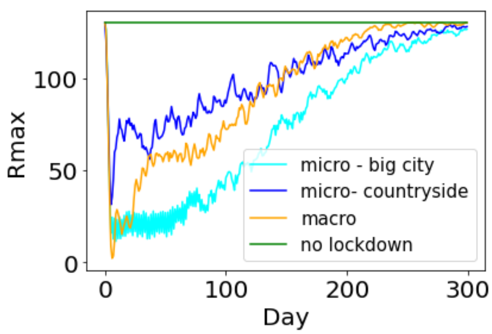

Fig. 8. R_{max} detail, with micro for countryside and big city

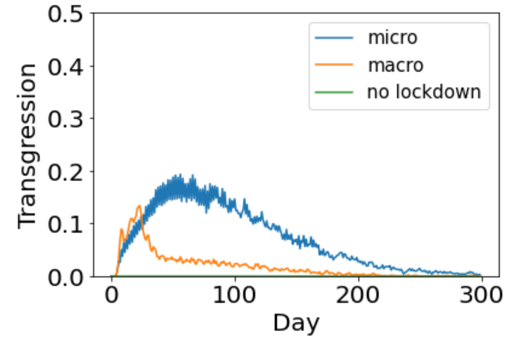

Fig. 9. Average transgression in the big city

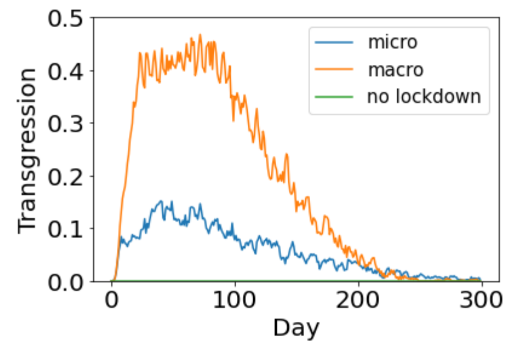

Fig. 10. Average transgression in the countryside

policies; ii) the need for customised policies to accommodate diversity in managed sub-systems and ensure local effectiveness; iii) the potential of local mismanagement to overflow into global ineffectiveness, via sub-system interconnections. At this generic level, we believe that considering these factors is critical for ensuring sustainability and scalability in social and/or technical systems. However, the concrete instances and specific impacts of these factors will differ case-by-case.

Importantly, we do *not* claim that either micro- or macro-governance is better than the other. Indeed, modelling additional factors, with different configurations and experimental settings may lead to different results (e.g. advantages of macro approach). Our main contribution was to illustrate the fact that changing the policy-making distribution across governance levels, with different information visibility, may significantly influence both local and global system behaviour.

That said, the analysed example showed the potential advantages of more decentralised governance in large, heterogeneous systems. These advantages occur whenever access to local details and customised policies become essential to management effectiveness. In the simulated scenarios, macro-governance averaged out differences among diverse communities, leading to general policies that were inappropriate locally.

In contrast, other scenarios may highlight the necessity of global views and uniform policies to ensure system-wide coherence. Our study showed how macro and micro approaches could be combined – i.e. a macro-defined global strategy was micro-adjusted to local conditions. In principle, the micro-scale also had indirect access to some form of global information, by observing the infections dynamics within its community. As individuals cross community boundaries, the epidemic follows similar trajectories within all communities, within certain delays. Hence, local observations may become an effective proxy for inducing global infections. Moreover, the micro-scale could access global information and keep local authority for policy customisation.

In future work we aim to develop a richer model, with more management scales and authority distributions, to better study the effects of various information flows and policy adaptations on local and global system behaviours.

References

[1] Di Felice L. J. and Diaconescu A. Hierarchy beyond top-down control: the architecture of self-organised social systems. In *IEEE Int. Cnf. ACSOS-C / SASSO*, pages 80–85, 2023.
[2] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Canto Classics. Cambridge University Press, 2015.
[3] Görg C. and Rauschmayer F. Multi-level governance and the politics of scale. In *Environmental Governance*, 2009 March.
[4] Diaconescu A. and Pitt J. Holonic institutions for multi-scale polycentric self-governance. In *Int.Wrksp COIN@AAMAS, Paris,2014*, pages 19–35.
[5] Diaconescu A., Di Felice L. J., and Mellodge P. Multi-scale feedbacks for large-scale coordination in self-systems. In *IEEE Int.Cnf. SASO*, pages 137–142, 2019.
[6] Cooper I. and Mondal A. A SIR model assumption for the spread of covid-19 in different communities. In *Chaos Solitons Fractals*, 2020.
[7] Diaconescu A., Di Felice L.J., and Mellodge P. Exogenous coordination in multi-scale systems. *FGCS*, 114:403–426, 2021.
[8] Grimm V. et al. A standard protocol for describing individual-based and agent-based models. *Ecological modelling*, 198(1-2):115–126, 2006.