Quantifying the Effects of Norms on COVID-19 Cases Using an Agent-Based Simulation

Jan de Mooij[1(✉)], Davide Dell'Anna[1], Parantapa Bhattacharya[2],
Mehdi Dastani[1], Brian Logan[1], and Samarth Swarup[2]

[1] Intelligent Systems, Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands
{a.j.demooij,d.dellanna,m.m.dastani,b.s.logan}@uu.nl

[2] Biocomplexity Institute and Initiative, University of Virginia, Charlottesville, USA
{parantapa,swarup}@virginia.edu

Abstract. Modelling social phenomena in large-scale agent-based simulations has long been a challenge due to the computational cost of incorporating agents whose behaviors are determined by reasoning about their internal attitudes and external factors. However, COVID-19 has brought the urgency of doing this to the fore, as, in the absence of viable pharmaceutical interventions, the progression of the pandemic has primarily been driven by behaviors and behavioral interventions. In this paper, we address this problem by developing a large-scale data-driven agent-based simulation model where individual agents reason about their beliefs, objectives, trust in government, and the norms imposed by the government. These internal and external attitudes are based on actual data concerning daily activities of individuals, their political orientation, and norms being enforced in the US state of Virginia. Our model is calibrated using mobility and COVID-19 case data. We show the utility of our model by quantifying the benefits of the various behavioral interventions through counterfactual runs of our calibrated simulation.

Keywords: Large-scale social simulation · Norm reasoning agents · Computational epidemiology

1 Introduction

In social systems in general, and in the science of epidemiology in particular, human behavior has always been recognized to play a crucial role [9]. This is especially true in the COVID-19 pandemic since, prior to the availability of vaccines, efforts at containing the epidemic have emphasized behavioral changes, such as mask wearing, physical distancing (e.g., keeping 6 ft apart), and social distancing (e.g., working from home, schooling from home). Compliance with these recommendations has varied widely, both spatiotemporally and demographically [12]. In most places, these non-pharmaceutical interventions (NPIs) were implemented

starting in March 2020. For example, in the US state of Virginia nine Executive Orders (EOs) were implemented between March and July 2020. Were some of these EOs more effective than others in limiting the spread of COVID-19? More generally, what determines the effectiveness of NPIs? Does their timing and sequence matter? These are all important questions to answer for developing effective mitigation plans for the next major epidemic. In this work, we propose an agent-based simulation approach for these problems, focusing on an analysis of the EOs implemented in Virginia.

Computational models of disease spread, including agent-based simulations, have become quite sophisticated. However, incorporating realistic models of human behavior in these simulations remains a challenge [5,10]. Most models assume a certain level of compliance with a behavioral intervention, and apply it uniformly at random [16]. In reality, however, compliance can be highly non-uniform as it depends on a number of factors, including: demographics, peer influence, political orientation, risk assessments, and beliefs about the efficacy of the behavior [2,4]. To improve epidemic simulations, we therefore need methods for the realistic modeling of behavior.

Belief-Desire-Intention (BDI) models developed in the MAS community, particularly those incorporating normative reasoning, are a natural fit for this problem [15]. However, it has been challenging to find appropriate data to calibrate such behavior models in simulations. Our approach is to use cellphone-based mobility data and a synthetic population [1] to create a data-driven simulation which is sufficiently detailed that the effects of behavioral responses to the EOs can be evaluated. To address the challenges of scaling, we adapt the BDI-based multiagent programming technology, 2APL [6,7], to support discreet time steps and deferral of action execution. We integrate this new library, Sim-2APL, with a new distributed agent-based simulation framework we call PanSim. This aspect of the work is presented in our companion paper [3]. In the current paper, we focus on the simulation design and evaluation. Our main contribution here is a framework that allows detailed investigation of the effects of non-pharmaceutical interventions through the use of multiple sources of data and appropriate behavioral models for agents.

2   Simulation Design

In this section, we describe our COVID-19 simulation, the key components of which are illustrated in Fig. 1. We start with a synthetic population of the US state of Virginia, where agents have realistic demographics, weekly activity schedules, and activity locations drawn from real location data. In our simulation, each individual in the synthetic population is represented by a norm-aware Sim-2APL agent (Sect. 2.2) which reasons about whether to comply with the various EOs that were implemented in Virginia (Sect. 2.3). The agents interact via a disease model implemented in the novel PanSim distributed environment (Sect. 2.4). In Sect. 3 we show how we calibrate the parameters of our simulation with real-world data, while in Sect. 4 we evaluate our simulation by comparing the disease progression when different norm interventions are put in place.
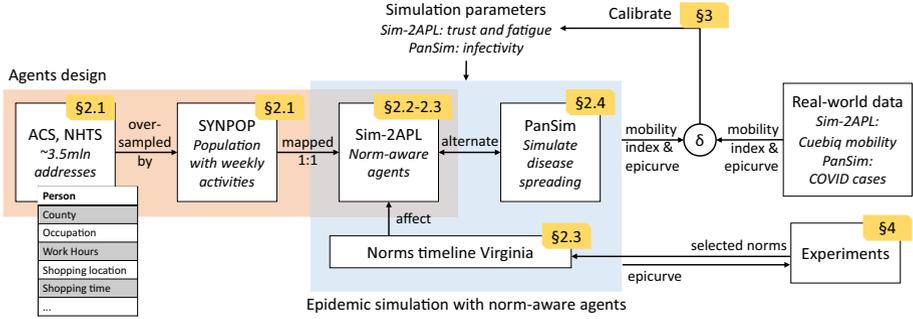
Fig. 1. COVID-19 simulation setting.

2.1   Data Sets Used in the Simulation

We use four data sets in this work, as described briefly below.

Synthetic Population of Virginia, USA: Agents in our simulation are drawn from a synthetic population of the state of Virginia, USA. This synthetic population has been constructed from multiple data sources including the American Community Survey (ACS), the National Household Travel Survey (NHTS), and various location and building data sets, as described in [1]. This gives us a very detailed representation of the region we are studying (multiple counties within Virginia). Agents are assigned demographic variables drawn from the ACS, such as age, sex, race, household income, and political orientation. In each county c, we label each household as Democratic with probability equal to the percentage of Democratic voters in the 2016 U.S. presidential elections in county c, and Republican otherwise. Agents are also assigned appropriate typical weekly activity patterns by integrating data from the NHTS. For each activity, each agent is assigned an appropriate location, using data about the built environment from multiple sources, including HERE, the Microsoft Building Database, and the National Center for Education Statistics (for school locations).

Mobility Data: In order to model the changes in mobility due to various Executive Orders (EOs) implemented between March and July 2020, we use anonymized and privacy-enhanced cellphone-based mobility data provided by Cuebiq. This data set contains location pings generated from the cellphones of a large number of anonymous and opted-in users throughout the USA. Cuebiq collects data with informed consent, anonymized all records and further enhanced privacy by replacing pings corresponding to home and work locations with the centroids of the corresponding Census blockgroups. We aggregate the data to the county level as follows. First we calculate the average *radius of gyration* for cellphone users in the county. The radius of gyration is given by $r = \sum_l d(l, l_c)/k$, where l is the location (latitude and longitude) of the user, l_c is the centroid of all the locations visited by the user on that day, k is the number of locations visited by the user on that day, and d is the Haversine distance. We then calculate a

mobility index as the percentage change in average r over all users in a given region on a given day compared with the average for the same day of the week in the same region during January and February of 2020, i.e., before any EOs were issued. For example, the mobility index for a specific Monday in May 2020 is the percentage change in the average r on that day compared to the average over all Mondays in January and February 2020.

COVID-19 Case Data: We use county-level COVID-19 case data from USA Facts to calibrate the disease model in our simulation. A caveat is that the number of confirmed cases probably under-counted the number of actual cases substantially, especially early in the epidemic, due to limited testing. We compensate for this in the simulation calibration by choosing a scale factor of 30, i.e., we assume that the actual number of cases was $30\times$ the reported number of cases. This arbitrary choice can straightforwardly be changed without affecting the methodology in our work.

Executive Orders in Virginia: We use a data set on Executive Orders that were implemented in each state in the USA [14] from the Johns Hopkins Coronavirus Resource Center [11]. From this we extract the ones that were implemented in Virginia in the period between March 1st and June 30th, 2020. In the simulation, EOs are represented by *norms* that agents may obey or violate, as described in Sect. 2.3. We quantify the benefits of these EOs through counterfactual runs of our calibrated simulation in Sect. 4.

2.2 Agents Activities and Deliberations

Each agent in the synthetic population is characterized by its weekly activity schedule, a set of typical daily activities over the course of one week. The schedule defines the location, start time and duration of all agents' activities as one of 7 distinct high level *activity types*: HOME, stay at or work from home; WORK, go to work or take a work-related trip; SHOP, buy goods (e.g., groceries, clothes, appliances); SCHOOL, attend school as a student; COLLEGE, attend college as a student; RELIGIOUS, religious or other community activities; and OTHER, any other activity, including recreational activities, exercise, dining at a restaurant, etc. For example, one activity in an agent's schedule could state "SHOP at location l between 7 p.m. and 8 p.m." These activity types categorize a larger number of low-level activity types, including but not limited to those describing the categories above. The high level activity types are what the agents use for reasoning, while the lower level activity types – which we do not use for reasoning because they are not guaranteed to have been sampled accurately during the creation of the synthetic populations – are only used to assign the location and activity time and duration according to the activity schedule. Each simulation step corresponds to one day, and at each simulation step, each agent retrieves and performs the activities from its activity schedule for the day of the week corresponding to that step.

We interpret each activity in an agent's daily schedule as a (to-do) goal for the corresponding Sim-2APL agent. For each activity (i.e. goal) in its daily schedule,

the agent generates a plan based on its goal, identifies any norms applicable to the activity, and decides whether it will obey or violate the norm(s) (See Sect. 2.3 below). If there are no applicable norms, or if the agent decides not to obey the norm, the agent uses the default plan for the to-do goal, i.e., the planned daily activity. However, if the agent decides to obey an applicable norm, the default plan for the daily activity is transformed into a norm-aligned plan. For example, if a norm specifies a mask should be worn in public places, the SHOP activity in the example above will be transformed into a SHOP activity with a "wearing a mask" modality.

2.3 Reasoning with Norms

We consider 11 norms representing a subset of the Executive Orders implemented in the state of Virginia (US). We distinguish regimented norms (R) that cannot be violated by agents, from non-regimented norms (NR) where agents may autonomously decide whether to comply with the norm or not. In addition, some norms have parameters that further specify the applicability of a particular instance of the norm to the activity itself or to the agent considering that activity. For example, the *type* parameter of the *BusinessClosed* norm specifies the type of business to which the norm instance applies (e.g., an instance may specify that only Non-Essential Business (*NEB*) should close), while the *size* and *type* parameters of the *SmallGroups* norm specify the maximum size of groups permitted in a context of a particular type (e.g., no more than 10 people are allowed in a *public* space). The *type* parameter of *SchoolsClosed*, finally, specifies the grade levels that are closed (e.g., *K-12* specifies all K-12 level schools are closed, i.e. the norm applies only to activities of type SCHOOL when the agent performing the SCHOOL activity is attending K-12 level education). The norms are summarized in Table 1 and briefly explained in Table 2. Figure 2 shows the date on which each norm came into force.

Factors Influencing Agent Decisions. If a regimented norm applies to an activity of an agent, the agent simply obeys the norm. If a non-regimented norm applies, the agent's decision whether to obey or violate the norm is influenced by a number of factors determined by the agent's beliefs and preferences regarding the activity. For example, in deciding whether to maintain physical distancing in a particular shop (i.e., to obey a *MaintainDistance* norm during a SHOP activity in a particular shop), agents take into account how many other agents they have observed maintaining physical distancing (*dist*) in the shop in the past, and their *trust* in the government[1]. Note that a norm may not be applicable to (relevant for) certain activities or agents, e.g., the norm *WearMaskPublInd* is not applicable to WORK or SCHOOL. Each factor is represented by a real value in the interval [0, 1], and the factors are summarized in last five columns of Table 1.

---

[1] Our choice of the factors influencing the agents' decisions, as well as of the norms mentioned above, should be considered as a 'proof of concept' to illustrate our framework. In more realistic simulations, elicitation of the most relevant factors in a well-designed study would be paramount. This is left for future work.

Each agent's initial *trust* in the government is determined by sampling a beta distribution $Beta_v(\alpha_v, \beta_v)$ (with $v = R$ for Republican and $v = D$ for Democrat), where $\alpha_v = \mu_v \cdot \kappa$, $\beta_v = (1 - \mu_v) \cdot \kappa$. The means μ_R and μ_D are determined by calibration (explained in Sect. 3); $\kappa = \alpha_v + \beta_v = 100$ characterizes the spread of the distribution, and, for simplicity, is fixed for both distributions. To simulate the decreasing compliance with measures that in reality manifested over time, agents in our simulation decrease their trust in the government by a constant factor f per simulation step after t_f simulation steps (days). Both t_f and f are fixed for all agents and are determined through calibration. The factor *acc* specifies the probability that an agent can be accommodated to work from home (in our simulation $acc = 0.45$ [8], and is the same for all agents). The factors *mask*, *dist* and *symp* specify the fraction of other agents encountered at a certain location who were wearing a mask, maintaining physical distancing, and who were (visibly) symptomatic, respectively. Symptoms are only visible if an agent is actually infected (determined by the disease model PanSim, see Sect. 2.4), but not all infected agents are symptomatic. The factor *all* specifies the number of agents encountered at a given location in excess of the maximum number of agents allowed by the norms currently in force.[2]

Table 1. Which activities are affected by regimented (R) and non-regimented (NR) norms, and the factors influencing the decision to comply with each norm.

Norm	Id	Param	Type	Activity types transformations					Influencing agents believes					
				WORK	SHOP	OTHER	SCHOOL	RELIGIOUS	trust	symp	acc	mask	dist	all
AllowWearMask	n_1	–	NR	mod	mod	mod	mod	mod	x			x		
BusinessClosed	n_2	type	R	del										
EmplWearMask	n_3	–	R	mod										
EncourTelework	n_4	–	NR	del					x	x	x			
MaintainDistance	n_5	–	NR	mod	mod	mod	mod	mod	x				x	
RedBusinessCapac	n_6	perc	R			del								
SchoolsClosed	n_7	type	R				del							
SmallGroups	n_8	size, type	NR	del	del	del	del	del	x	x				x
StayHome	n_9	appl	NR	del	del	del	del	del	x					
TakeawayOnly	n_{10}	–	R			short								
WearMaskPublInd	n_{11}	–	NR		mod	mod		mod	x			x		

Violating or Obeying a Norm. To determine whether to obey or violate a norm n when performing an activity act, the agent calculates a probability $p(n, act)$ of obeying n at the current simulation step, given by:

$$p(n, act) = \frac{1}{1 + e^{(-k \cdot (x - x_0))}} \tag{1}$$

---

[2] Due to space limitations, we refer to the code repository for the specific details of the factors: https://bitbucket.org/goldenagents/sim2apl-episimpledemics.

where x represents the evidence for complying with n computed as the *average* value of the factors (excluding the *trust* factor) that support the compliance with n when performing *act*, $x_0 = 1 - trust$ represents the agent's distrust in the institution that issued n, and k is the logistic growth rate or steepness of the curve ($k = 10$ in our simulation). Note that when the trust in the institution is extreme (e.g., x_0 is close to 0 or 1), the decision to comply with the norm becomes more "resistant" to evidence supporting norm compliance. For example, if the agent has no trust in the institution (i.e., $x_0 = 1$), the probability of complying with a norm that depends only on the factor *mask* is 0.5 when 100% of other agents do wear a mask, but drops off steeply as the value of *mask* declines (when 75% of other agents do wear a mask, the probability to comply with the norm drops to approximately 0.07 for $k = 10$). However, if the trust value is more balanced (e.g., $x_0 = 0.5$), the decision to comply with the norm relies more on the supporting evidence.

Table 2. A brief explanation of the norms enforced in our simulation and of their parameters.

Id	Interpretation	Parameters
n_1	Mask wearing is allowed and encouraged	–
n_2	Businesses of type *type* are closed	*type* $\in \{NEB\}$: the type of business, $NEB = $ Non Essential Business
n_3	Employees working in retail must wear a mask during work activities	–
n_4	Telework is encouraged	–
n_5	Physical distance of 1.5 m should be maintained	–
n_6	Capacity of business should be reduced to *perc*	*perc*: percentage of business capacity
n_7	Schools of type *type* are closed	*type* $\in \{K12, HE, K12 \ or \ HE\}$: the type of school, $K12 = $ primary and secondary education $HE = $ Higher Education (HE)
n_8	The maximum allowed size of groups of type *type* is *size*	*type* $\in \{public, private, all\}$: the target settings, either public, private or both (*all*); *size* $\in \mathbb{N}$: maximum size of groups
n_9	Stay at home if belong to category *appl*	*appl* $\in \{sick \ or \ age \geq 65, all\}$: the group of agents to which the norm applies, either people sick or older than 65 (*sick or age* ≥ 65), or everyone (*all*)
n_{10}	Only take away allowed for restaurants	–
n_{11}	A mask must be worn in public indoor settings	–

When an agent violates a norm with respect to a scheduled activity, the norm is ignored for that activity and the agent adopts the default plan for the activity (the to-do goal of the agent). When an agent obeys a norm with respect to an activity, the activity is subject to a transformation. We distinguish three types of transformations of activities:

- *mod*: the modality (in our model either wearing a mask or practicing physical distancing) of the activity is changed. For example, when the norm *Wear-MaskPublInd* is obeyed for the *SHOP* activity, the agent performs that activity while wearing a mask. In the code, the modality is a flag that is interpreted by PanSim and affects the susceptibility or infectivity of an agent (see Sect. 2.4).
- *del*: the activity is cancelled. When an activity is cancelled, it is transformed into a HOME activity, unless the agent can shift the next scheduled activity. For example, if an agent is scheduled to go to WORK, but its working place is closed, the agent will stay HOME, unless in its daily schedule there is a consequent activity (e.g., a SHOP activity) that can be performed earlier.
- *short*: the activity is shortened. For example, when obeying a *TakeawayOnly* norm, the agent will spend less time at the restaurant.

The *Activity Types Transformations* shown in Table 1 specify how the norms affect each activity type. If no transformation is indicated in Table 1 for a pair ⟨norm n, activity type at⟩, the norm n does not apply to activities of type at.
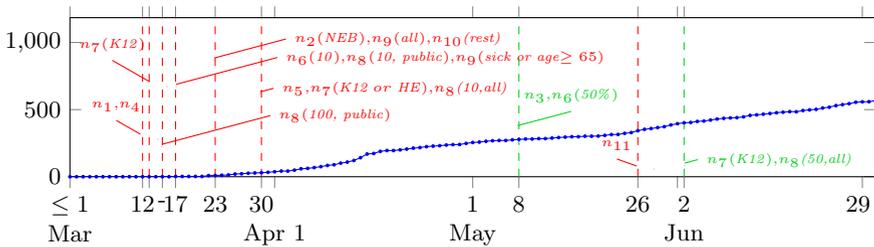

Fig. 2. Cumulative number of combined recorded cases in the counties of Goochland, Fluvanna, and Charlottesville (blue line). Red and green lines are introduction of new restrictions and relaxations of previous ones. (Color figure online)

2.4   Environment Design

To model the spread of COVID-19, we implemented a novel distributed agent-based epidemic simulation platform, which we call PanSim. In PanSim, a simulation progresses in discrete timesteps. When a Sim-2APL agent decides to visit a location, it interacts with other agents visiting that location, and observes the visible attributes exhibited by these agents such as: coughing, wearing mask, social distancing, etc., allowing it to modify its behavior based on its observations

at subsequent timesteps. The probability of symptomatic and asymptomatic agents transmitting or becoming infected per unit time (5 minutes) under different action modalities such as the wearing of a mask or physical distancing, is given by the probabilistic addition of all individual interactions of that day. To simulate cases being introduced from outside, we artificially expose 5 agents during the first 5 days of the simulation, and 3 more agents each simulation day after.

The novelty of PanSim lies in the fact that, unlike previous epidemic simulation frameworks, PanSim has explicit support for modeling human behavior, increasing the number and type of social phenomena that can be modeled, and allowing disease progression to be driven by explicit colocation rather than statistical likelihood of contact between agents. The colocation in turn is the result of locations and times that individual agents – implemented in any agent programming language – can explicitly choose for their activities. PanSim further allows scaling up the number and complexity of agents and visits by distributing the simulation across multiple compute nodes, where each node simulates a distinct set of agents and locations. PanSim synchronizes its instances across compute nodes by sharing only the data relating to agents visiting a location simulated on another node, ensuring all its instances remain synchronized throughout the simulation with minimal communication. Both the framework and experiments showing the scalability are described in detail in the companion paper [3].

3   Calibration

We calibrate the behavior and disease parameters independently from each other in two distinct processes. For this reason, the best parameters for either model were not yet available when calibrating the other. In each process, the parameters for the model not being calibrated were fixed to our best estimations (based on results of earlier trial runs) of the values. In other words, the parameters for the disease model were fixed in the process in which we calibrated the behavior model, and vice versa. Both calibration processes are performed by means of Nelder-Mead (NM) minimization [13]. NM iteratively refines an initial configuration of parameters until it finds a local optimum that minimizes a given objective function, in this case the Root Mean Square Error (RMSE) between observations in the simulation and the real world. Calibration was performed using data from the counties of Charlottesville (41119 unique agents in the synthetic population, 83.25% of which voted Democratic, 16.75% Republican), Fluvanna (24109 unique agents, 45.35% Democratic, 54.65% Republican), and Goochland (20922 unique agents, 37.55% Democratic, $62,45\%$ Republican) for a total of 86150 agents, 61.55% Democratic, $33,45\%$ Republican. These counties have been selected for their proximity, number of agents, and variation in voting preference. For each set of parameters selected by NM, we run 5 different simulations in order to account for non-determinism in the simulation.

Agent Parameters. We calibrate the four parameters of the agent model introduced in Sect. 2.3, i.e., the *means* μ_D and μ_R of the two beta distributions

(a) Mobility index   (b) Confirmed cases
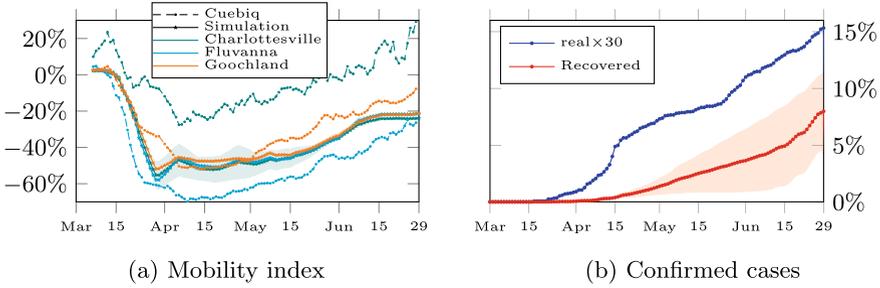
Fig. 3. The mobility index observed in the simulation plotted against that recorded by Cuebiq in each simulated county (a), and percentage confirmed cases ($\times 30$) of the population plotted against that of the recovered agents in the simulation (b).

from which we sample the trust attitudes of Democratic and Republican agents, respectively, the fatigue factor f, and the time step t_f in the simulation at which the fatigue becomes active. We calculate the RMSE between the mobility index in our simulation and in the real-world Cuebiq data (calculated as per Sect. 2.1). We apply a smoothing to the mobility index of each day by averaging it with the mobility index of the 6 preceding days in order to smooth out the intrinsic difference in the weekly repeated mobility trends between the synthetic population and Cuebiq data. We perform these simulations with the disease model parameters fixed to $inf_s = 0.00045$ and $inf_a = 0.0003375$ (best estimate).

Disease Model Parameters. The two parameters of the disease model that are calibrated are the infectivity of symptomatic (inf_s) and asymptomatic (inf_a) agents. We calculate the RMSE between the cumulative infection case count in the three simulated counties and the number of recovered agents in our simulation. The agent parameters are fixed to $\mu_D = 0.776816$, $\mu_R = 0.106955$, $f = 0.0125$, and $t_f = 60$ (best estimate).

Calibration Results. For both calibration processes, we run NM until 10 consecutive configurations of parameters did not improve the objective function. The final parameters determined by our calibration are: $\mu_D = 0.704621$, $\mu_R = 0.004685$, $f = 0.0125$ and $t_f = 60$ for the agent model (RMSE: 17.6574), and $inf_s = 0.0000481$ and $inf_a = 0.0000241$ for the disease model (RMSE: 2052.0222). Figure 3 compares the mobility (Fig. 3a) and the number of recovered agents (Fig. 3b) resulting from these parameters with the real data. The agent parameter calibration found a relatively good fit for the decrease in mobility, including the increase in mobility after the first few weeks. However, the large differences between the different counties could not be reproduced by our simulation. The disease model calibration resulted in a slightly less aggressive spread of the disease than the (scaled) recorded case count in the first few months of the COVID-19 outbreak.

4   Quantifying the Effects of Normative Interventions

We perform an experiment with the calibrated models to understand the relative impact of the measures instigated by the institutions in Virginia on the behavior of its residents. Given the list of $n = 9$ normative interventions that took place in Virginia as per Fig. 2, we run 10 different experiments: in experiment Ei, for $0 \leq i \leq n$, we enact only the first i executive orders. For example, in experiment $E0$, no norm is enforced, i.e., we simulate a scenario where no behavioral intervention takes place; in experiment $E1$, we enact only the first EO, i.e., norms $\{n_1, n_4\}$ starting from March 12th; in experiment $E2$ we enact the first two EOs, i.e., $\{n_1, n_4\}$ starting from March 12th and also $\{n_7(K12)\}$ starting from March 13th, etc. In each experiment we compute the total number of agents that has been infected at the end of the simulation. This time, we include the county of Louisa in the simulation, for a total of 119087 agents. We run each experiment 5 times to account for non-determinism in the simulation.
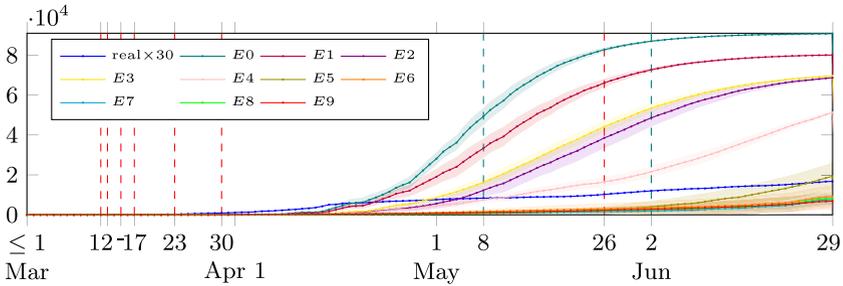

Fig. 4. Cumulative cases in E1-10, and in the real-world ($\times 30$, blue line). (Color figure online)

Figure 4 shows the number of recovered agents at each time step in the simulations (SIR plots available in the code repository), with the standard deviation between the 5 runs shown as the confidence interval. $E0$ shows that if no measures had been taken, the spread of COVID-19 would have been several times more rapid. The higher curves do not show exponential growth until the end of the simulation, since our simulation contained only 119087 agents. After a sufficiently large portion of the population has been infected it becomes increasingly hard for the disease to encounter susceptible agents, slowing the spread.

Table 3 shows the total number of agents that have been infected at the end of the simulation (including those not yet recovered). The experiment $E9$, where all the norms were enforced, shows the lowest number of total infections, with a reduction of 27% in cases compared the $E8$ – in which the maximum group size was completely lifted instead of relaxed from 10 to 50 people – and an 83% reduction compared to the experiment where no norms were enforced.

The largest decrease was from $E5$ to $E6$, closely follows by $E6$ to $E7$. In the last EO in $E6$ the maximum group size of 10 was also applied to private

Table 3. The average number of cumulative cases in each experiment

Exp.	Cumulative cases	Diff. w.r.t. E_{i-1}	Diff. w.r.t. E_0
$E0$	90983.6 ± 41.802		0.0%
$E1$	80711.0 ± 343.627	-11.29%	-11.29%
$E2$	71616.2 ± 241.017	-11.27%	-21.29%
$E3$	71760.6 ± 297.292	$+0.20\%$	-21.13%
$E4$	64674.4 ± 445.714	-9.870%	-28.92%
$E5$	49118.0 ± 4162.414	-24.05%	-46.01%
$E6$	30505.0 ± 10202.892	-37.89%	-66.47%
$E7$	20599.0 ± 4423.636	-32.47%	-77.36%
$E8$	21195.6 ± 4044.473	$+2.90\%$	-76.70%
$E9$	15569.6 ± 5144.708	-26.54%	-82.89%

gatherings, in addition to the already closed *K-12* schools higher education was closed, and physical distancing was declared compulsory. In the last EO in $E7$, the earlier reduction of business capacity to 10 was relaxed to 50% capacity, but offset by requiring all employees to wear masks. Given the large uncertainty in $E6$, we cannot conclusively declare it more effective than $E7$, but rank both as similarly effective. This means that, from the norms considered in this work, restricting the group size in private settings, making physical distancing compulsory, and requiring employees to wear masks were the most effective in reducing the spread of COVID-19.

It should be noted that for the purpose of this work, various simplifications have been applied to the actual norms enforced. Moreover, in practice the EOs (including relaxations) have been issued in response to the actual spread of COVID-19 at that time, while in our simulation they were fixed to their original dates. Nevertheless, these results show that behavioral responses of individual agents to normative interventions, and not just the effect of an assumed level of compliance, can be studied through our proposed simulation framework.

5   Conclusion

We presented a novel distributed agent-based simulation framework for large-scale multi-agent simulations of norm-governed behaviors in epidemics, and applied it to the case of COVID-19. We modeled a population of agents representing individuals from the state of Virginia, whose daily behavior was determined from multiple data sources, including the American Community Survey. We calibrated and validated the behavior exhibited by the agents, affected by the norms enforced in the state of Virginia (such as school and business closures, mask-wearing and physical distance interventions) using Cuebiq mobility data and the COVID-19 infection data. We used the model to compare the sensitivity of the COVID-19 outbreak size to the different normative interventions. In future

work, we intend to evaluate the scalability of our framework, to introduce more complex agents dynamics, such as inter-agent communications, and to evaluate a number of different hypothesis about the COVID-19 pandemic.

Future work also includes improving the simulation calibration. We believe that reducing the arbitrary scaling of the observed number of cases from 30 to a smaller factor will result in better calibration. Improving the mobility calibration to reflect the variations in mobility index from one county to another may require further refinement of the behavior model. We are also working on scaling up to larger populations, such as all the 133 counties and independent cities in the state of Virginia, which add up to over 7.6 million agents, and evaluating more complex experiment designs.

More broadly, we believe that effective intervention to mitigate novel epidemics requires methods to evaluate the effects of normative interventions in detail, which in turn requires being able to model human behavioral choices and responses. Through the use of substantial real-world data, BDI models of agent reasoning, and a scalable simulation platform, we can come closer to this goal.

References

1. Adiga, A., et al.: Generating a synthetic population of the United States. Technical report, NDSSL 15-009, Network Dynamics and Simulation Science Laboratory (2015)
2. Becher, M., Stegmueller, D., Brouard, S., Kerrouche, E.: Comparative experimental evidence on compliance with social distancing during the COVID-19 pandemic. medRxiv (2020)
3. Bhattacharya, P., de Mooij, J., Dell'Anna, D., Dastani, M., Logan, B., Swarup, S.: PanSim + Sim-2APL: a framework for large-scale distributed simulation with complex agents. In: International Workshop on Engineering Multi-Agent Systems (2021)
4. Chan, D.K.C., Zhang, C.Q., Weman-Josefsson, K.: Why people failed to adhere to COVID-19 preventive behaviors? Perspectives from an integrated behavior change model. Infect. Control Hosp. Epidemiol. **42**(3), 375–376 (2021)
5. Chen, J., Lewis, B., Marathe, A., Marathe, M.V., Swarup, S., Vullikanti, A.K.S.: Individual and collective behavior in public health epidemiology. In: Disease Modelling and Public Health, Part A, vol. 36, pp. 329–368 (2017). Chapter 12
6. Dastani, M.: 2APL: a practical agent programming language. Auton. Agents Multi-Agent Syst. **16**, 214–248 (2008). https://doi.org/10.1007/s10458-008-9036-y

7. Dastani, M., Testerink, B.: Design patterns for multi-agent programming. Int. J. Agent-Oriented Softw. Eng. **5**(2/3), 167–202 (2016)

8. Dey, M., Frazis, H., Loewenstein, M.A., Sun, H.: Ability to work from home. Mon. Labor Rev. 1–19 (2020)

9. Ferguson, N.: Capturing human behavior. Nature **446**, 733 (2007)

10. Funk, S., et al.: Nine challenges in incorporating the dynamics of behaviour in infectious disease models. Epidemics **10**, 21–25 (2015)

11. Johns Hopkins Coronavirus Resource Center: Impact of opening and closing decisions in Virginia, new cases - Johns Hopkins. https://coronavirus.jhu.edu/data/state-timeline/new-confirmed-cases/virginia/. Accessed 07 Oct 2020

12. Katz, J., Sanger-Katz, M., Quealy, K.: A detailed map of who is wearing masks in the U.S. https://www.nytimes.com/interactive/2020/07/17/upshot/coronavirus-face-mask-map.html. Accessed 08 Oct 2020

13. Nelder, J.A., Mead, R.: A simplex method for function minimization. Comput. J. **7**(4), 308–313 (1965)

14. Northam, R.S.: Virginia Governor Ralph S. Northam - executive actions. https://www.governor.virginia.gov/executive-actions/. Accessed 07 Oct 2020

15. Swarup, S., Eubank, S., Marathe, M.: Computational epidemiology as a challenge domain for multiagent systems. In: Proceedings of the Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2014)

16. Verelst, F., Willem, L., Beutels, P.: Behavioural change models for infectious disease transmission: a systematic review. J. R. Soc. Interface **13**, 20160820 (2016)