

Identifying Synergistic Interventions to Address COVID-19 Using a Large Scale Agent-Based Model*

Junjiang Li¹^[0000-0002-4144-106X] and Philippe J. Giabbanelli¹^[0000-0001-6816-355X]

Department of Computer Science & Software Engineering, Miami University, Oxford OH 45056, USA <https://www.dachb.com> {lij111,giabbapj}@miamioh.edu

Abstract. There is a range of public health tools and interventions to address the global pandemic of COVID-19. Although it is essential for public health efforts to comprehensively identify *which* interventions have the largest impact on preventing new cases, most of the modeling studies that support such decision-making efforts have only considered a very small set of interventions. In addition, previous studies predominantly considered interventions as independent or examined a single scenario in which every possible intervention was applied. Reality has been more nuanced, as a subset of all possible interventions may be in effect for a given time period, in a given place. In this paper, we use cloud-based simulations and a previously published Agent-Based Model of COVID-19 (Covasim) to measure the individual and interacting contribution of interventions on reducing new infections in the US over 6 months. Simulated interventions include face masks, working remotely, stay-at-home orders, testing, contact tracing, and quarantining. Through a factorial design of experiments, we find that mask wearing together with transitioning to remote work/schooling has the largest impact. Having sufficient capacity to immediately and effectively perform contact tracing has a smaller contribution, primarily via interacting effects.

Keywords: Cloud-based simulations · Factorial Analysis · Large-scale simulations · Synergistic Interventions

1 Introduction

Mathematical models of COVID-19 commonly use ‘compartmental models’, which are systems of coupled differential equations that predict global quantities such as the number of infections at any given time [6]. Agent-based models (ABMs) were later used to capture heterogeneity in populations [17], representing how different individuals (e.g., in age, gender, or socio-economic factors) have different risks, or willingness and abilities to comply with preventative measures [5].

* The authors thank the Microsoft AI for Health program for supporting this research effort through a philanthropic grant.

In the absence of a widely used vaccine, our study and previous ABMs rely on *non-pharmaceutical interventions* including individual-level preventative measures (e.g., stay at home, social distance, using face masks) and subsequent interventions (e.g., contact tracing, quarantining). There are several limitations to these studies. First, previous studies focused on the effects of a *small subset of interventions* commonly adopted by national governments, and not all interventions are applied in conjunction. In a sample of 10 studies [4], we noted that an average of only 2.3 interventions are used simultaneously, which is significantly less than the number of interventions that have been implemented or considered by governments [15]. Second, when several interventions are implemented, there is limited analysis to assess whether synergistic effects are obtained or whether most of the benefits can be attributed to only some of the policies. Identifying the right set of synergistic interventions is an important information for policy-making, particularly as compliance becomes an issue. Finally, we frequently note reporting issues such as an insufficiently motivated number of runs (since stochastic models need replicated runs to achieve a sufficient confidence interval) or a coarse resolution when the target population is large.

To support policymakers in identifying the right set of interventions and provide the necessary confidence in high-stake computations, our paper uses Design of Experiment (DoE) techniques and large-scale cloud-based simulations that measure the individual and interactive effects of interventions at a detailed level. Specifically, our contributions are twofold: ① We measure the impact of *six interventions* on disease prevalence and mortality by accounting for interactive effects. Our interventions include face masks, social distancing, stay-at-home orders, testing, contact tracing, and quarantining. We simulate these interventions at *various levels of adherence*, thus accounting for possible variations in behavioral responses. ② We provide all results at a very accurate population scale of 1 : 500 for the USA and within a confidence interval of at least 95%. This requires a massive number of computationally intensive experiments, which are performed through the cloud via the Microsoft Azure platform.

The remainder of the paper is structured as follows. In section 2, we describe the procedures of our experiments, including the principles of factorial designs of experiments. Our results are discussed in section 3. For an overview on modeling the biology and policies for COVID-19, or potential methodological limitations to this study, we refer the reader to our supplementary online materials [4].

2 Methods

Our simulations are performed **using the Covasim platform**. We ensure the accuracy of the simulations within our application context by using *(i)* a resolution¹ of 1 : 500 and *(ii)* a 95% Confidence Interval² [16, pp. 184–186]. Our

¹ Each simulated agent represents 500 real-world people. Given the US population of 328 million people, we have a simulated agent population of over 650,000.

² We identify and perform a sufficient number of replications in each scenario such that the average results are within a 95% Confidence Interval (95% CI).

simulation scenarios correspond to combinations of interventions at various levels (detailed in subsection 2.2). To quantify the effects of stipulated interventions, we report the total infections after 180 days. It is important to note that our simulated time period starts in September 2020 rather than at the beginning of the pandemic, when most of the population was susceptible and almost none had recovered. By choosing a more current starting time, our simulated dynamics take into account the presence of recovered agents as well as individuals who are already at various stages of the infection. We made two modifications to `Covasim` accordingly: randomly selected agents are set to recovered based on the amount estimated by the CDC and scaled to our population resolution; and a separate random subset of agents are set to infected, with their date of infection chosen among the 14 days preceding the start of the simulation (based on CDC data) such that agents are realistically set to different stages of infection.

As `Covasim` already contains a large number of population parameters calibrated for the US, we used the same default values in our work. There were three exceptions, as the knowledge base upon which `Covasim` was built has since evolved and needed to reflect our current understanding of real-life disease dynamics. As detailed in our supplementary online materials [4], we updated the model regarding the distribution of incubation period, proportion of symptomatic cases, and testing delays.

We focused on **four categories of interventions**: mask wearing (realized as direct reductions of the susceptibility of simulated agents), lockdowns, testing, and contact tracing. Several parameters are required for each category in order to precisely characterize how the intervention will unfold. For instance, testing is a matter of *how many* tests are available on a daily basis, *when* to test individuals who're entering quarantine as they were exposed to the virus, and the extent to which tests are *reliable* (i.e. test sensitivity). We list the parameters for each category of intervention in Table 1 together with the range of values that *could* be used and the specific subset that we do use. Our choice is motivated by the references provided in the table and is often limited to a binary due to the experimental set-up explained in the next subsection. Every intervention is applied for the entire duration of the simulation. Each intervention takes place in certain networks, thus we account for the possibility that agents may wear masks in one network but not in another.

Our goal is to measure the impact of interventions on disease prevalence. We use a 2^k **factorial design** Design of Experiments (DoE) in which each parameter is set to two values, designated as 'low' and 'high' [12, pp. 233–303]. Our four categories of interventions result in 9 parameters (Figure 1), which are listed as X_1, \dots, X_9 in Table 1 together with their low and high values. A factorial design serves to investigate the synergistic effects of these parameters, by measuring the response y (average number of infections at the end of the simulation) for each simulated combination. As mentioned in section 3.1, y is obtained over a number of replications necessary to fit the 95% confidence interval to within 5% of the average. To determine the contribution of each parameters (individually as well as in groups), we compute the variances in response contributed by each

Table 1. Covasim interventions used here, with the following shorthands: $w \rightarrow$ work network, $s \rightarrow$ school network, $h \rightarrow$ home network, and $c \rightarrow$ community network.

Intervention	Parameters	Possible Values	Values Chosen	Var.	Ref.
Mask Wearing	Fractional reduction in transmissibility	$[0, 1]$	80%		[1]
	Applied networks	$\mathcal{P}(\{w, s, h, c\})$	$\{\{w, s, h, c\}, \{w, s\}, \{c\}\}$	X_1	
Lockdowns	Fractional reduction in contacts for $\{w, s\}$	$[0, 1]$	{5%, 30%}	X_2	[7]
	Fractional reduction in contacts for c	$[0, 1]$	{10%, 30%}	X_3	[7]
	Applied networks	$\mathcal{P}(\{w, s, h, c\})$	$X_1 \setminus h$		
Testing	Number of tests per day	Any integer	$\{6, 11.1\} \times 10^5$	X_4	[8]
	When to test quarantined individuals	{start, end, both, daily}	{start, end, both}	X_5	
	Test sensitivity	$[0, 1]$	{55%, 100%}	X_6	[10]
Contact Tracing	% of contacts of a tested person that can be traced	$[0, 1]$	{0.2, 1}	X_7	[9]
	Delay in contact tracing	Any integer	{0, 7}	X_8	[9]
	Start contract tracing without waiting for results	{True, False}	{True, False}	X_9	

combination and perform an F -test. In other words, the variance is decomposed over individual parameters (X_1, \dots, X_9) as well as interactions of parameters. We speak of 2nd order interaction when examining the joint effect of two parameters ($X_1X_2, X_1X_3, \dots, X_8X_9$), 3rd order interaction for three parameters, and so on. We include up to 3rd order interactions to scan for possible effects that would only happen when three interventions are applied jointly.

There are two special cases when applying the 2^k factorial design to COVID-19 interventions. First, we cannot assume which values of X_1 (networks for mask wearing) and X_5 (whether to test at the start, end, or both times in a quarantine) have the highest or lowest impact on disease incidence. Consequently, we performed simulations for each of the three possible values of X_1 and X_5 . We identified the maximum and minimum number of infected individuals across these simulations, and thus set the high and low values accordingly (Table 2). For example, consider that we have 10 sick individuals when testing at the start of quarantine, 100 when testing at the end, and 1000 when testing at both times. In this case, ‘start’ is the intervention with highest impact ($X_5^{high} = start$) while ‘both’ does the least to control the spread of the virus ($X_5^{low} = both$). Second, not all $2^9 = 512$ combinations of the high and low levels are valid because of the

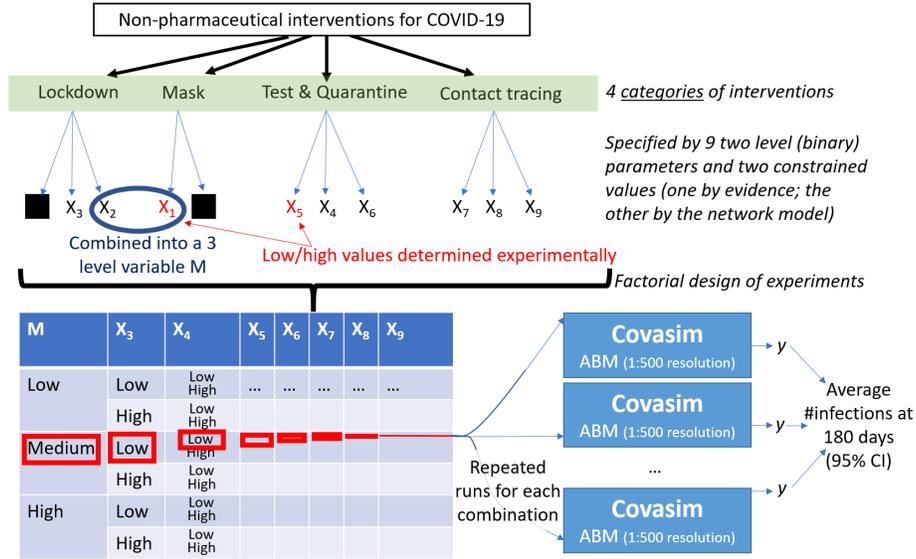


Fig. 1. Process flowchart from the identification of interventions (top) and their operationalization in the simulation to the factorial Design of Experiments (DoE) and repeated runs of a stochastic model targeting a 95% Confidence Interval.

interdependency between the networks for mask wearing and lockdown. If an intervention applies to work and school (e.g., working and learning remotely) then the fractional reduction of contact in the *community* should be 0 rather than $X_3^{low} = 0.1$, since the community network is not concerned by the intervention.

To resolve this issue, we first note that the highest impact intervention is the one that targets all networks, hence $X_1^{high} = \{w, s, h, c\}$ as seen in Table 2. The core question is thus the *low* value, which is either $\{w, s\}$ or c . This choice sets either X_2 or X_3 (resp. the fractional reduction in contacts for c and $\{w, s\}$) to 0%. Consequently, $2^7 = 128$ combinations are invalid and $384 = 3 \times 2^7$ remain. To avoid the complications of an incomplete parameter space, we combine two parameters into a new variable M with 3 levels (Table 3) and keep the rest unchanged. We can interpret this variable to be the level of “social distancing” in the broadest sense that is enforced. Moreover, the problem reduces to an ordinary factorial design analysis with one 3-level factor, which is analyzed using established methods [12, pp. 412–414].

3 Results and Discussion

Mask wearing together with transitioning to remote work/schooling has the largest impact (Table 4). It interacts with the constructs having the next largest impacts: ability to perform contact tracing and whether to start contact tracing without waiting for test results. Our finding about the importance of masks

Table 2. Maximum responses caused by each choice of variables X_1 and X_5 .

Variable	Value	Max Response	Level
X_1	$\{w, s, c, h\}$	678.93	High
	$\{w, s\}$	278571.55	—
	c	359937.49	Low
X_5	start	359862.78	—
	end	359606.72	High
	both	359937.49	Low

Table 3. Correspondence between levels of M and combinations of levels of X_1 and X_2 , given c as the low level for X_1 .

M	X_1	X_2
High	$\{w, s, c, h\}$	30%
Med	$\{w, s, c, h\}$	5%
Low	c	0%

and remote work/school differs from early works (in the first half of 2020) regarding COVID-19. Early systematic reviews released as preprints considered that “the evidence is not sufficiently strong to support widespread use of face-masks as a protective measure against COVID19” and that “masks alone have no significant effect in interrupting spread of [influenza-like illnesses ...] or influenza” (cited in [11]). More recently, a systematic review from December 2020 concluded that only four out of seventeen studies ‘supported use of face masks [and] a meta-analysis of all 17 studies found no association between face mask intervention and respiratory infections’ [13]. *However, once results are adjusted* for factors such as age and sex, the meta-analysis “suggests protective effect of the face mask intervention” [13]. Similarly, the most recent analyses and commentaries agree that reducing social network interactions in settings such as universities [3] is needed to avoid large outbreaks. Consequently, our result regarding the large effect of masks and remote work/schooling contributes to the more recent evidence base on interventions regarding COVID-19.

Although the preventative approach of using masks and shifting into remote work/school plays the largest role in reducing the likelihood of transmission (by lowering both the number of contacts and the virus transmissibility per contact), we do observe interacting effects with other intervention parameters. In particular, contact tracing is important to mitigate the pandemic, as demonstrated by the case of South Korea [14]. Our study contributes to understanding the specific parameters underlying contact tracing, as results stress the merits of having sufficient capacity to immediately and effectively perform contact tracing.

Limitations such as model validation are detailed in our supplementary online materials [4]. Our simulations used a factorial analysis in the artificial environment afforded by a model: these results could be contrasted to real data, as different US states and countries have implemented different combinations of interventions at various points in time. In addition, our results are reported using the confidence interval method to handle the stochastic nature of the simulations, but parameters may have different levels of uncertainty [2]. Finally, the scale of 1 : 500 (resulting in over 650,000 agents) was the maximum simulation size that we could perform given the number of repeats and hardware memory limitations. Accuracy may thus be improved by simulating a full population.

Table 4. Factors responsible for $\geq 1\%$ of variance in the number of new infection cases

Variable(s)	Meaning	Contribution (%)
M	Masks + remote work/school	66.638
M and X_7	Masks + remote work/school, and ability to trace contacts of tested individuals	13.759
X_7	Ability to trace contacts of tested individuals	8.953
M and X_9	Masks + remote work/school, and starting contact tracing without waiting for test results	2.804
X_9	Contact tracing starts without waiting for test results	2.033
M and X_8	Masks + remote work/school + contact tracing delay	1.077

References

1. Chu, D.K., et al.: Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: A systematic review and meta-analysis. *Lancet* **395**(10242), 1973–1987 (2020)
2. Edeling, W., et al.: Model uncertainty and decision making: Predicting the impact of covid-19 using the covidsim epidemiological code (2020)
3. Ghaffarzadegan, N., et al.: Diverse computer simulation models provide unified lessons on university operation during a pandemic. *BioScience* (2020)
4. Giabbanelli, P.J., Li, J.: Identifying synergistic interventions to address covid-19 using a large scale agent-based model. *medRxiv* (2020)
5. Harper, C.A., et al.: Functional Fear Predicts Public Health Compliance in the COVID-19 Pandemic. *Int J Ment Health Addiction* (2020)
6. IHME COVID-19 Forecasting Team: Modeling COVID-19 scenarios for the United States. *Nature Medicine* (2020)
7. Jacobsen, G.D., Jacobsen, K.H.: Statewide covid-19 stay-at-home orders and population mobility in the united states. *World Med. Health Policy* **12**(4), 347–356 (2020)
8. John Hopkins University: Daily State-by-State Testing Trends. <https://coronavirus.jhu.edu/testing/individual-states> (2020)
9. Kretzschmar, M.E., et al.: Impact of delays on effectiveness of contact tracing strategies for COVID-19: A modelling study. *The Lancet Public Health* **5**(8), e452–e459 (2020)
10. Lisboa Bastos, M., et al.: Diagnostic accuracy of serological tests for covid-19: Systematic review and meta-analysis. *BMJ* p. m2516 (2020)
11. Martin, G., Hanna, E., Dingwall, R.: Face masks for the public during covid-19: an appeal for caution in policy (2020)
12. Montgomery, D.C.: *Design and Analysis of Experiment*. Wiley, 8 edn. (2012)
13. Ollila, H.M., et al.: Face masks prevent transmission of respiratory diseases: a meta-analysis of randomized controlled trials. *medRxiv* (2020)
14. Park, Y.J., et al.: Contact tracing during coronavirus disease outbreak, south korea, 2020. *Emerging infectious diseases* **26**(10), 2465–2468 (2020)
15. Petherick, A., et al.: Variations in Government Responses to COVID-19 (2020)
16. Robinson, S.: *Simulation: The practice of model development and use* (2014)
17. Yang, J., et al.: Prevalence of comorbidities and its effects in patients infected with SARS-CoV-2: A systematic review and meta-analysis. *Int J Infect Dis* **94**, 91–95 (2020)