Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

A NOVEL APPROACH FOR OUTCOMES ESTIMATION IN HYBRID SIMULATION MODELS OF DISEASE TRANSMISSION AND PROGRESSION

Soham Das[1], Aparna Venkataraman[1,2], and Varun Ramamohan[1]

[1]Dept. of Mechanical Engg., Indian Institute of Technology Delhi, New Delhi, INDIA
[2]Centre for Health Services Research, Faculty of Medicine, The University of Queensland, Brisbane, AUSTRALIA

ABSTRACT

In this study, we consider hybrid simulations consisting of an agent-based simulation (ABS) to model disease transmission in a population, and a discrete-time Markov chain executed as a Monte Carlo simulation to model the heterogeneous progression of the disease in infected agents. In such scenarios, execution of the ABS is stopped at a certain time point. At this point, disease-related outcomes for infected agents are estimated by executing the disease progression Monte Carlo simulation for each infected agent over their lifespans, well beyond the execution horizon of the ABS. This can incur substantial computational expense. We present a novel method to alleviate this computational burden by randomly sampling and allocating disease-related outcomes from a repository of outcomes generated and stored as a one-time exercise prior to execution of the hybrid simulation. We demonstrate the effectiveness of our approach via a stylized hybrid simulation of a hypothetical infectious disease transmission scenario.

1 INTRODUCTION

Infectious disease models developed to inform intervention policies typically incorporate disease transmission dynamics as an important feature, given that doing so captures the second-order benefits of preventing new infections (Kim and Goldie 2008). Decisions on intervention policies are made based on the effects of these interventions on population-level disease-related health and economic outcomes. For example, health outcomes estimated by such models commonly include average life years (the average total time spent alive by an entity of interest) and quality-adjusted life years (an aggregate measure of time spent in each health state multiplied by a utility weight taking values in $[0,1]$, reflecting the health-related quality of life in that state). Economic outcomes commonly include the direct and indirect costs of managing the disease. The health and economic outcomes associated with an intervention (such as treatment, or a disease screening campaign) are often integrated into a single measure of cost-effectiveness of the intervention, which is used to inform decisions around its reimbursement from the perspective of a health payer (such as private insurance providers in the United States or national health systems in the United Kingdom) (Muennig and Bounthavong 2016). Key modeling paradigms used to capture infectious disease transmission dynamics include differential equation based compartmental models and agent-based simulation models (Kim and Goldie 2008).

In this study, we consider a hybrid simulation scenario that arises in this context, involving an agent-based simulation (ABS) of infectious disease transmission dynamics, and a discrete-time Markov chain (DTMC) model of disease progression executed as a Monte Carlo simulation for each infected agent used to estimate the health and economic outcomes for each infected agent (dos Santos et al. 2020; Kar et al. 2022). Such hybrid simulations arise commonly in the context of hepatitis C virus (He et al. 2014; Das et al. 2019) and HIV/AIDS models (Gopalappa et al. 2017). In such models, the stochastic and heterogeneous progression of the disease within an infected individual across their lifespan is captured by the execution of the disease progression DTMC (referred to henceforth as the DP-DTMC) as a single replication of a

Monte Carlo simulation. The outcomes from this Monte Carlo replication - typically life years and time spent in each disease state - are used in conjunction with a background non-disease-related mortality model (referred to henceforth as the non-DRM model) to estimate the total QALYs and costs associated with management of the disease in this individual. The execution of the DP-DTMC (for infected agents) and the non-DRM models begins as soon as an agent enters the model via birth, model initiation, or when the agent experiences a change in infection status (e.g., an uninfected agent acquires the infection). In such hybrid simulations, the ABS itself is executed for a certain time horizon (e.g., for T years of simulation time), at the end of which many infected and uninfected agents are typically present. Lifetime outcomes need to be estimated for each of these agents, which will require continuing the execution of the DP-DTMC and the non-DRM models for the infected agents and the non-DRM model alone for uninfected agents. If the disease in question progresses relatively slowly, such as HCV (wherein average life expectancies are not affected substantially (Aggarwal et al. 2017)) or HIV, and if many agents are present at the end of T, the execution of the DP-DTMC and the non-DRM Monte Carlo replications beyond T can incur substantial computational expense. In this study, we propose and demonstrate a method to reduce the computational expense for such hybrid simulation models of infectious disease transmission dynamics. The hybrid simulation is depicted in Figure 1.

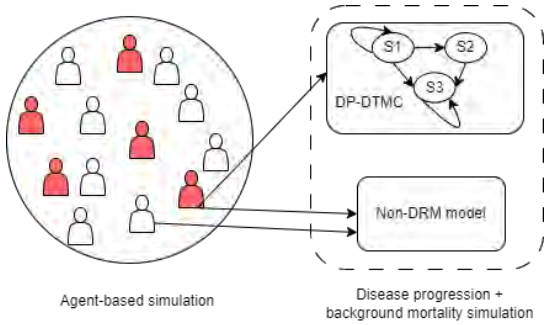

Figure 1: Hybrid disease transmission and progression simulation. DP-DTMC: disease progression - discrete-time Markov chain. Agents in red - infected; agents in white - uninfected. Non-DRM model: non-disease-related mortality model.

This method is based on the premise of interchangeability of outcomes across *iid* replications of a simulation, and the fact that the generation of outcomes for each agent via execution of the DP-DTMC and the non-DRM models represents the output of a replication of a Monte Carlo simulation. Given that we are interested in the distribution of outcomes across agents and not in the outcomes associated with individual agents, we propose a method to randomly sample and allocate outcomes of interest (e.g., QALYs) from a repository of outcomes to agents. This repository of outcomes is generated as a one-time exercise prior to the execution of the hybrid simulation model. In subsequent sections, we describe our proposed scheme for generation of the outcomes repository and appropriate sampling and allocation of outcomes to agents in the hybrid simulation. We empirically demonstrate the effectiveness of our 'outcomes sampling and allocation' (OSA) approach, in terms of comparability of aggregate outcomes and reduction in computational expense, in comparison to the status-quo 'without sampling and allocation' (WSA) approach via a stylized hybrid agent-based simulation of the disease transmission and progression dynamics of a hypothetical infectious disease.

We now discuss the literature relevant to outcomes generation for hybrid models of infectious disease transmission dynamics, and our research contributions with respect to the extant literature.

The hybrid simulation scenario, involving an ABS capturing the transmission and individual-level interaction dynamics of some phenomenon and a Monte Carlo simulation used to estimate outcomes related to the heterogeneous, individual-level impact of the phenomenon, most commonly arises in cost-effectiveness models that inform reimbursement decisions for a health intervention. A significant portion of the health

economics and cost-effectiveness modeling literature for infectious diseases is based on static modelling approaches that work with a fixed (closed) cohort of agents without considering individual-level interactions or growth of the cohort. Examples include VanDeusen et al. (2015), Aggarwal et al. (2017) and Taguchi et al. (2020) in the context of HIV/AIDS, HCV infection, and human papillomavirus (HPV) infection, respectively. Dynamic approaches, which advantageously can accommodate individual-level interactions relevant to disease transmission, include differential equation based compartmental models and individual-level agent-based simulations. Compartmental models take an aggregate perspective to modelling disease transmission dynamics, wherein compartments represent different conditions like health and disease states and transitions between compartments are governed by rates applied to differential equations. Examples of such modelling approaches are Schobert et al. (2012) (HPV), Song et al. (2015) (HIV), Lim et al. (2018) and Das et al. (2019) (HCV).

Individual-level models consider individual sample paths, including interactions between individuals, that are relevant to disease spread and progression. Most hybrid simulations in this context arise in this case, as described in the beginning of the previous paragraph. Examples of such hybrid simulations, which include agent-based simulations for modelling disease transmission and disease progression models for estimating outcomes for those infected, include the studies of Olsen and Jepsen (2010) (HPV), He et al. (2016) (HCV), Gopalappa et al. (2017) (HIV) and Zhou et al. (2021) (influenza). The disease progression models used are usually DTMCs, as in Olsen and Jepsen (2010), He et al. (2016) and Ayer et al. (2019), but may also involve other approaches such as individual-level survival curves (Gupta et al. 2021). Most of these studies do not report clearly their approach towards outcomes estimation, and it is likely most either assign average values to individuals (in which case, the stochasticity/heterogeneity of outcomes is not preserved) or execute an approach similar to the WSA approach. None of these models, in our understanding, employ an approach similar to the OSA method for reducing the computational expense of outcomes estimation.

Our research contribution thus involves development of a way for substantially reducing the computational expense incurred by hybrid simulations with ABSs for modelling interactions relevant to spread of a phenomenon and individual-level models for estimating the outcomes resulting from interactions within the ABS, and this forms the key research contribution of this study. Additionally, health economics studies for infectious diseases only adopt a lifetime horizon for outcomes estimation when working with fixed/closed cohorts, which is likely why many studies for slow-moving diseases such as HCV only work with fixed/closed cohorts. This is because of the modeling complexity and computational demands of incorporating lifetime outcomes within agent-based models that work with much smaller time horizons. Our study, via the OSA approach, provides a way to conveniently incorporate lifetime horizons for outcomes estimation within agent-based models. This approach may be extended to such hybrid simulations in other domains as well.

2 METHODOLOGY

We first describe the status-quo WSA approach to provide the necessary context before introducing our proposed OSA approach for outcomes estimation.

2.1 The *Without Sampling & Allocation* Method for Outcomes Estimation

Given that, without loss of generality, we consider estimation of outcomes for uninfected agents as well as infected agents, we assume generation of outcomes begins as soon as an agent enters the model, either via birth or at model initiation. At each time step, relevant outcomes are incrementally accumulated. If the agent is uninfected, then the non-DRM model alone is executed and relevant outcomes (usually life years alone) begin to be recorded and accumulated, and if the agent becomes infected, both the DP-DTMC and non-DRM models are executed and the relevant outcomes (life years, QALYs and costs) are recorded and accumulated in each time step. For example, with daily time steps, the daily disease-related cost corresponding to a particular disease state s is added for every day the agent spends in s. Similarly, for

QALYs, at every time step, the health-related quality of life weight (say q_s) associated with state s is multiplied with the time step unit (e.g., δt, corresponding to a day) to yield the QALY value ($q \times \delta t$) added at that time step. The WSA algorithm is presented in Algorithm 1.

Notation:

- t: time index. Note that t can exceed the simulation time horizon n (i.e., the point at which the execution of the disease transmission ABS is stopped.
- $a_t \in \{1,2,...,A\}$: age at time t, where A is the maximum age limit for an agent in the model; for example, $A = 100$ years.
- SI_t: status of infection at time t; SI_t = 1 if the agent is infected, 0 otherwise.
- $s_t \in \{1,\ldots,m\}$: disease state at time t.
- P: outcome of interest.
- P_t^{ac}: Accumulated value of P until time t.
- δP_t: incremental value of P at time t.
- *Initialize*: function that initializes the age, status of infection and disease state for the agent.
- *Evaluate*: function that evaluates the incremental parameter value δP_t at time t to be added to the accumulated value P_{t-1}^{ac} based on a_t and $SI_t \times s_t$ (if the agent is uninfected, then $SI_t \times s_t = 0$, and if the agent is infected, then $SI_t \times s_t = s_t \in \{1,\ldots,m\}$, the disease state).
- *Update*: function that updates SI_t and s_t based on the transmission and progression modules.

Algorithm 1 The *without sampling and allocation* approach for outcomes generation.

$t \leftarrow 0$
a_0, SI_0, $s_0 \leftarrow$ *Initialize* (a, SI, s)
$P_0^{ac} \leftarrow 0$
while Agent is alive **do** ▷ Updates are assumed to occur at the end of t
  $t \leftarrow t + 1$
  $a_t \leftarrow a_{t-1} + 1$
  $SI_t, s_t \leftarrow$ *Update* (SI, s)
  $\delta P_t \leftarrow$ *Evaluate* $(a_t, SI_t \times s_t)$
  $P_t^{ac} \leftarrow P_{t-1}^{ac} + \delta P_t$
Return P_t^{ac}

Note that Algorithm 1 is developed from the perspective of a single agent - in other words, the value of the simulation clock at $t = 0$ may be different for every agent, except for those present at model initiation. Upon entry of the agent into the model, four variables are initialized - age (a_t), status of infection (SI_t), disease state (s_t, if $SI_t = 1$) and accumulated outcomes (P_t^{ac}) - and indexed by time. Upon entry of the agent, t is set to 0 and the value of P_t^{ac} is also 0. The values of age (a_0), status of infection (SI_0) and disease state (s_0, if $SI_0 = 1$) form the initial conditions for the agent. At every time step after entry of the agent into the model and until its death, the following steps are performed:

- The time step t is incremented by 1 unit.
- The age of the agent is incremented by 1 unit.
- The infection status of the agent is updated, which can change due to disease transmission (if agent is uninfected) or due to cure (if agent is infected).
- The disease state for infected agents is updated in accordance with the DP-DTMC model.
- The incremental value of the outcome, δP_t, is calculated as a function of the infection status, disease state, and age.
- δP_t is added to the variable indicating the outcome value accumulated until time step $t-1$ (P_{t-1}^{ac}), yielding the outcome value accumulated by the end of time step t (P_t^{ac}).

When the agent dies (say at time t_e), the outcome value $P_{t_e}^{ac}$ accumulated until t_e is returned. We assume that the agent exits the model by death if their age reaches A. Importantly, the process depicted in Algorithm 1 is executed until the death of the agent in question - that is, it can extend well beyond the end of the ABS execution time horizon n.

We now describe the OSA approach for outcomes estimation.

2.2 The *Outcomes Sampling & Allocation* Method for Outcomes Estimation

Under the OSA approach, none of the components of the hybrid simulation model need to be executed beyond the ABS execution time horizon n. Allocation of outcomes for each agent is done at model entry and at time points when they experience a change in infection status. Upon a change in infection status, the previous allocation of outcomes is cancelled and a fresh allocation is made based on the then age and disease state. To account for outcomes incurred from the time of model entry to the point at which the most recent change in infection status occurs, the process of incremental accumulation of outcomes is continued during the ABS model execution period alone.

Outcomes are randomly sampled and allocated from a repository of outcomes created as a one-time exercise. The outcome allocated to each agent can be considered as a realization of a random tuple representing the stochastic progression of the disease as well as non-DRM if the agent is infected, and as the realization of a random variable representing stochastic non-DRM if the agent is uninfected. We also recall here that the initial conditions that determine the outcomes for an infected agent are their age and disease state upon model entry and age at model entry alone for uninfected agents. Note that age and/or disease state at model entry become relevant only for those agents present at model initiation (i.e., $t = 0$), and for those infected or cured at $t > 0$, only the age at t is relevant. This is because upon infection, the disease state s_t is always 1 (the initial stage of the disease). Similarly, upon cure, the agent reverts to a 'healthy' uninfected state ($SI_t = 0$). Therefore, the repository consists of vectors of outcome realizations corresponding to each combination of health state (given by $SI_t \times s_t$) and age. Each such vector consists of a large number (in our case 10,000) realizations of the outcome generated via a one-time execution of the DP-DTMC and non-DRM models, as applicable. If there are k outcomes ($k = 3$ in our case: life years, QALYs and costs) of interest, then k such vectors are generated for each combination of agent age a and health state $SI \times s$. Thus if there are A possible ages and $m+1$ possible health states (m disease states and one uninfected state), then k vectors of outcome realizations are generated for each of the $A \times (m+1)$ combinations of agent initial conditions.

Sampling and allocation from the repository, once stratified by the combinations of age and health state, is straightforward. Upon model entry, or a change in health state (infected to uninfected, or vice versa), the vectors of outcomes corresponding to the agent's then age and health state are identified, and outcomes are sampled with replacement and allocated per the OSA approach depicted in Algorithm 2. The algorithm below captures the OSA approach for a single outcome P from the perspective of a single agent who enters the model at time t_0.

Additional Notation:

- P_t^{al}: Allocated value of outcome P at time t
- P^{ad}: Adjusted (in conjunction with incremental outcomes) value of P
- *Allocate*: function that randomly samples and allocates P_t^{al} based on a_t and $SI_t \times s_t$ from the outcomes repository

Algorithm 2 The *outcomes sampling & allocation* approach for outcomes generation.

$t \leftarrow t_0$, $P^{al} \leftarrow 0$ ▷ t_o is the time at which the agent enters the model
a_{t_0}, SI_{t_0}, $s_{t_0} \leftarrow$ *Initialize* (a, SI, s)
$P^{al}_{t_0} \leftarrow$ *Allocate* $(a_{t_0}, SI_{t_0} \times s_{t_0})$
$P^{ac}_{t_0} \leftarrow 0$
while $t \leq n$ **do** ▷ Updates are assumed to occur at the end of t
  $a_t \leftarrow a_{t-1} + 1$
  Update (SI_t, s_t)
  if $| SI_t - SI_{t-1} | > 0$ **then** ▷ Change in infection status occurs
    $P^{al}_t \leftarrow$ *Allocate* $(a_t, SI_t \times s_t)$
    $P^{ad} \leftarrow P^{al}_t + P^{ac}_{t-1}$
  $\delta P_t \leftarrow$ *Evaluate* $(a_t, SI_t \times s_t)$
  $P^{ac}_t \leftarrow P^{ac}_{t-1} + \delta P_t$
  if agent dies **then**
    Return P^{ad}
if Agent is alive **then**
  Return P^{ad}

It can immediately be seen from Algorithm 2 that no part of the hybrid simulation is executed beyond the ABS time horizon n. Under the OSA approach, as mentioned before, sampling and allocation of outcomes occurs upon an agent's entry into the model or upon a change in health state (infection or cure). The final value of the outcome returned under this approach is the most recently allocated and adjusted value prior to the agent's exit from the model, P^{ad}, and not the most recently recorded incrementally accumulated outcome value P^{ac}.

The basic notion underlying the OSA approach is that under the same initial conditions of agent age and health state (the tuple $(a, SI \times s)$), the outcome for each agent can be considered as the output of a Monte Carlo simulation replication. The vectors of outcomes in the repository, for a given tuple $(a, SI \times s)$, are also generated as the outputs of a large number (10,000 in our case) of *iid* replications of the same DP-DTMC and non-DRM (where applicable) Monte Carlo simulation models.

Under the OSA approach, the outcome returned for an agent is based on the most recently allocated and adjusted outcome prior to its exit from the model. If this allocation and adjustment is done at time t, this outcome is the sum of the incrementally accumulated outcome P^{ac}_{t-1} and the allocated outcome P^{al}_t. Under the WSA approach, the outcome returned for an agent can be also be broken down similarly, as the sum of the incrementally accumulated outcome until and excluding t, P^{ac}_{t-1} and the incrementally accumulated outcome from t onwards (including t), which we can denote as P^{ac}_{t+}. P^{ac}_{t+} can be thought of as being estimated by executing the DP-DTMC and the non-DRM models as a Monte Carlo simulation from t onwards.

Now, under the OSA approach as well, P^{al}_t is estimated in the same manner - by executing the DP-DTMC and the non-DRM models as a Monte Carlo simulation assuming that the agent starts their sojourn in the age group a and the health state (disease state if infected) at the time the allocation is done (t). However, in this latter case, the execution is done as a one-time exercise external to the execution of the hybrid simulation, and is done to generate a large number of *iid* outcome realizations that form the outcomes repository. Even in the WSA case, execution of the DP-DTMC model begins only when the agent is infected (if infection occurs at $t > 0$) or enters a particular state (for infected agents present at $t = 0$). Further, the Markovian assumption underpinning the DP-DTMC renders unnecessary the consideration of any time spent in the state prior to its execution. Similarly, we assume that the agent ages or age groups are also determined such that the risk of non-DRM is approximately the same for all agents within an age group regardless of how much time they have spent in the age group. For example, the risk of non-DRM is the same regardless

of whether an agent is of age a or $a+0.5$ as long as the ages a and $a+0.5$ are considered to be within the same age group a. Under these conditions, it is likely that P_t^{al} (from the OSA approach) converges in distribution to that of P_{t+}^{ac} (from the WSA approach) as long as the number of agents in the ABS cohort is sufficiently large and the number of samples corresponding to the $(a, SI \times s)$ is also sufficiently large.

An explanation of why outcomes need to be reallocated each time the infection status changes and why outcomes need to be incrementally accumulated during the ABS model execution period is warranted here. First, we recall that the outcomes repository is stratified by the age and health state that represent the point from which the outcome needs to be estimated - that is, the point at which the agent is assumed to begin their sojourn in said age (group) and health state. We also note that the outcome returned for an agent is the most recently allocated and adjusted outcome prior its exit from the model. The need for reallocation of outcomes is immediately clear from this fact: given that we do not know if and when the next change in health state will occur prior to an agent's exit from the model, reallocation must be done upon initiation (an agent's health state may not change at all prior to its exit) and each time its health state changes.

The need for adjustment of the allocated outcome with the incrementally accumulated outcome also emerges from the need for reallocation of outcomes each time an agent's health state changes. Consider an agent B that enters the model in an infected state at time t_1 and is cured at t_2. This implies that the outcome for this agent - for example, QALYs - is the sum of the QALYs incurred between t_1 and t_2 and the QALYs incurred from t_2 onwards. Given that the QALYs that are allocated are predicated on the agent's age and health state at t_2, and do not consider the agent's trajectory prior to t_2, it is clear that they represent the QALYs incurred from t_2 onwards. Therefore, adding the incrementally accumulated outcomes until t_2 to the allocated outcomes at t_2 is required to estimate the outcome in its entirety.

We now describe the computational illustration of the OSA approach.

3 COMPUTATIONAL IMPLEMENTATION OF THE OSA APPROACH

3.1 Hybrid Simulation Development

All computation relevant to this study was conducted on the MATLAB scientific computing platform, and empirical experiments were executed on a workstation with a 12$^{\text{th}}$ generation Intel Core *i9* processor, with a a base clock speed of 3.2 gigaHertz and 32 gigabytes of memory.

In order to illustrate the OSA approach, we create a stylized hybrid simulation consisting of four modules: (i) demographics, (ii) disease transmission, (iii) disease progression, and (iv) treatment. This stylized model was based on the hybrid simulation developed for modeling the transmission and progression of the hepatitis C virus in the Indian context (Das et al. 2019).

The model consists of a dynamic cohort of agents, initialized with 10,000 agents (with 400 infected agents) and with non-DRM mortality operating in the background. At 100 years of age, an agent dies automatically. The distribution of the ages of the 10,000 agents was determined based on data from the 2011 Indian census (Ministry of Home Affairs, Government of India 2011). New agents enter the model through births, incorporated through an hourly probability of birth. We executed the model using daily time steps. The model execution time horizon was 10 years. The OSA approach was executed only for these 10 years (3600 time steps). However, the WSA approach was executed until the deaths of all agents who entered the model until the end of the 10-year model execution time horizon. However, non-DRM and the DP-DTMC models were executed after the intervention period.

Disease transmission is modeled via an environment within the model created to facilitate agent interactions. Any agent could visit this environment during a time step with a certain probability. We assume that infection transmission takes place in groups of size 10, and the composition of these groups changes daily. The presence of one infected agent can contaminate the group with a certain probability, and once the group becomes contaminated, any uninfected agent within an interaction group could get infected.

With regard to disease progression, the DP-DTMC model comprises four disease states other than disease-related death (DRD). The DP-DTMC model is depicted in Figure 2. Quality of life and disease state costs vary according to the disease state. Treatment takes place at the end of every year (including once in the beginning), hence there are 11 treatment 'camps' over the 10-year period. We explored three uptake rates (30%, 50% and 70%) for treatment. The uptake rate represents the proportion of infected agents who undertake treatment every year. By treatment-naive patients, we imply patients who had not failed treatment earlier. The cure rates of treatment are assigned such that they decrease and the treatment costs increase as the disease stage (state) becomes more advanced. Experiments using each uptake rate are carried out for ten replications. Both undiscounted and discounted values of life years, QALYs and costs are calculated as averages over the cohort of interest. The parameterization of the stylized hybrid

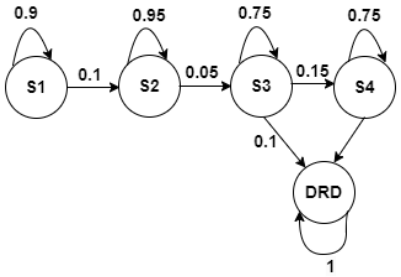

Figure 2: The disease progression discrete-time Markov chain. *Notes.* DRD = disease-related death.

simulation is provided in Table 1. The same parameterization was used for generating the repository as well as for carrying out simulation experiments using the WSA and the OSA approaches. Also, while the ABS and the DP-DTMC and the non-DRM models were executed using daily time steps under both the OSA and the WSA approaches, the generation of the repository was carried out by Monte Carlo simulation replications of the DP-DTMC and the non-DRM models using annual time steps. This is because a large number of agent age and health state combinations needed to be executed for the generation of the outcomes repository and using daily time steps would have substantially increased the computational runtime of the exercise.

We now describe the results of our computational experiments.

3.2 Computational Experiments: Results

Computational experiments were conducted comparing the outcomes generated by the hybrid simulation run under both the OSA and the WSA approaches. Our hypothesis is that the distributions of the outcomes from the hybrid simulation under both approaches will be similar (no statistically significant differences) with a substantial reduction in average runtimes for the OSA approach. First, in Table 2, we provide outcomes (life years, QALYs and costs, both discounted and undiscounted) and runtimes from 10 replications for the OSA and the WSA approaches for a single treatment uptake rate (50%) to illustrate in detail the results from the hybrid simulation model. Then, in Table 2, we provide summary statistics (averages and standard deviations) of the outcomes and runtimes from both approaches for three uptake rates: 30%, 50% and 70%.

Table 1: Parameterization of the stylized hybrid simulation. *Notes*: DRD = disease-related death.

Parameter	**Estimate**
Demographics	
Initial population	10,000 (initially infected: 400)
Maximum age limit (A)	100 years
Birth rate	0.2 / hour
Background mortality	1.6×10^{-5} / agent / day
Disease Transmission	
Probability of an agent visiting the environment on a day (representing an interaction with a non-zero risk of infection)	0.05
Interaction group size (group changes every day)	10
Probability that an infected agent contaminates the group sub-environment	0.1
Probability of an uninfected agent getting infected from a contaminated group sub-environment	0.02
Disease Progression	
Annual Disease Progression probabilities	S1 to S2: 0.10 S2 to S3: 0.05 S3 to S4: 0.15 S3 to DRD: 0.10 S4 to DRD: 0.25
Disease state distribution among the pool of infected agents present at model initiation	S1: 200, S2: 150, S3: 40, S4: 10
Health-related quality of life utility weights	S1: 0.80, S2: 0.70, S3: 0.55, S4: 0.40
Disease state costs (annual)	S1: 1,000, S2: 4,000, S3: 10,000, S4: 50,000
Treatment	
Percentage of infected patients treated every year (uptake rate)	50%
Cure rates for patients	S1: 0.95, S2: 0.90, S3: 0.85, S4: 0.80
Treatment costs (INR)	S1, S2: 10,000; S3, S4: 20,000

Table 2: Model outcomes at 50% uptake rate for the WSA and the OSA approaches for 10 replications. *Notes.* Runtimes are in seconds. WSA: Without sampling & allocation approach, OSA: Outcomes sampling & allocation approach, u/d: undiscounted, d/c: discounted.

WSA Approach										
Outcome of interest	**Replication**									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Runtime	334	330	319	540	323	623	319	324	320	325
Life years (u/d)	60.39	60.33	60.34	60.48	60.57	60.71	60.78	60.2	60.39	60.44
QALYs (u/d)	57.89	57.86	57.8	57.98	58.1	58.2	58.31	57.66	57.89	57.91
Costs (u/d)	60,791	60,236	61,530	60,951	60,590	60,879	60,198	61,186	61,049	62,516
Life years (d/c)	27.63	27.53	27.64	27.64	27.67	27.72	27.71	27.57	27.6	27.67
QALYs (d/c)	26.13	26.04	26.12	26.14	26.18	26.21	26.24	26.05	26.1	26.16
Costs (d/c)	35651	35,346	35,768	35,590	35,390	35,576	35,147	35,736	35,575	36,262
Outcome sampling outcomes										
Outcome of interest	**Replication**									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Runtime	164	199	167	200	165	200	258	167	198	167
Life years (u/d)	61.97	61.92	61.84	61.85	61.56	61.51	61.7	61.64	61.92	61.43
QALYs (u/d)	59.32	59.33	59.2	59.24	58.99	58.89	59.11	59.03	59.34	58.84
Costs (u/d)	63502	62770	64323	62691	62094	63141	62826	63261	61826	62963
Life years (d/c)	28.08	27.94	27.99	27.99	27.95	27.92	27.94	27.95	27.98	27.88
QALYs (d/c)	26.52	26.41	26.43	26.44	26.41	26.26	26.41	26.41	26.44	26.33
Costs (d/c)	36373	35,980	36,884	36,318	36,001	36,309	36,077	36,258	36,022	36,316

The results in Tables 2 and 3 indicate that the OSA approach yields comparable estimates of the average life years, QALYs and costs in comparison to those from the WSA approach, in both undiscounted and discounted terms. However, the outcomes yielded by the OSA approach marginally overestimate those from the WSA approach. To explain this, let us assume that an agent is of age *a* when a change in status of infection takes place. Under the OSA approach, outcomes are randomly sampled from the repository and allocated based on age *a* and the new health state. In doing so, the fact that agent has most likely already spent some time in the age *a* prior to the change in their infection status is not taken into consideration. On the other hand, the initial conditions during the Monte Carlo simulation replications executed to generate the repository represented the point at which an agent started their sojourn in age *a*. This slight mismatch in initial conditions most likely is the cause for the marginal overestimation of outcomes by the OSA approach.

Table 3: Summary statistics of model outcomes at three treatment uptake rates under the WSA and OSA approaches for outcomes estimation. *Notes.* Runtimes are in seconds, standard deviations are in parentheses. WSA: Without sampling & allocation approach, OSA: Outcomes sampling & allocation approach, u/d: undiscounted, d/c: discounted.

Uptake rate	**30%**		**50%**		**70%**	
Outcomes of interest	**WSA**	**OSA**	**WSA**	**OSA**	**WSA**	**OSA**
Runtime	311	158.2	375.7	188.5	489.8	314.7
Life years (u/d)	55.80 (0.14)	57.21 (0.22)	60.71 (0.22)	61.86 (0.13)	63.32 (0.15)	64.61 (0.13)
QALYs (u/d)	52.08 (0.15)	53.43 (0.21)	58.21 (0.23)	59.26 (0.12)	61.61 (0.16)	62.85 (0.14)
Costs (u/d)	83857 (593)	84507 (711)	61006 (671)	63096 (613)	46772 (621)	47599 (546)
Life years (d/c)	26.38 (0.04)	26.78 (0.05)	27.65 (0.06)	27.97 (0.05)	28.32 (0.04)	28.62 (0.04)
QALYs (d/c)	24.20 (0.05)	24.58 (0.05)	26.15 (0.07)	26.43 (0.04)	27.28 (0.04)	27.55 (0.04)
Costs (d/c)	45133 (285)	45214 (366)	35605 (299)	36254 (268)	29691 (309)	29955 (270)

Table 3 explores how the OSA approach performs at multiple uptake rates. Consistent comparability in the magnitudes of the outcomes can be seen across uptake rates, accompanied by significant reductions in runtimes. The average runtime reduction was 49.13% for a 30% uptake rate, 49.83% for a 50% uptake rate, and 35.75% for a 70% uptake rate, with an overall average reduction of 44.90%.

Finally, in Figure 3, we provide the cumulative distribution functions (CDFs) for each outcome - discounted and undiscounted - for an uptake rate of 50% for one replication under both approaches. It is evident from the plots that the CDFs for all outcomes generated by both approaches are nearly identical, reinforcing the effectiveness of the OSA approach with respect to the WSA approach.

4 DISCUSSION

In this paper, we present a novel 'outcomes sampling and allocation' approach for estimating outcomes from individual-level hybrid simulation models of infectious disease transmission dynamics. We consider a specific hybrid simulation scenario wherein the infectious disease transmission dynamics are captured by an agent-based simulation and the progression of the disease and background non-disease-related mortality are captured by Monte Carlo simulation models. We demonstrate the effectiveness of our proposed approach in terms of the similarity of distributions of outcomes and substantial reductions in computational runtimes in comparison to the status quo 'without sampling and allocation' approach.

The OSA approach yields average reductions in computational runtimes ranging from 35% - 50%, with reductions in runtimes going up to approximately 70% for individual replications. This can alleviate the computational burden of hybrid simulations involving component agent-based models that may already incur a substantial computational overhead, especially if the cohort sizes considered in the ABSs are large.

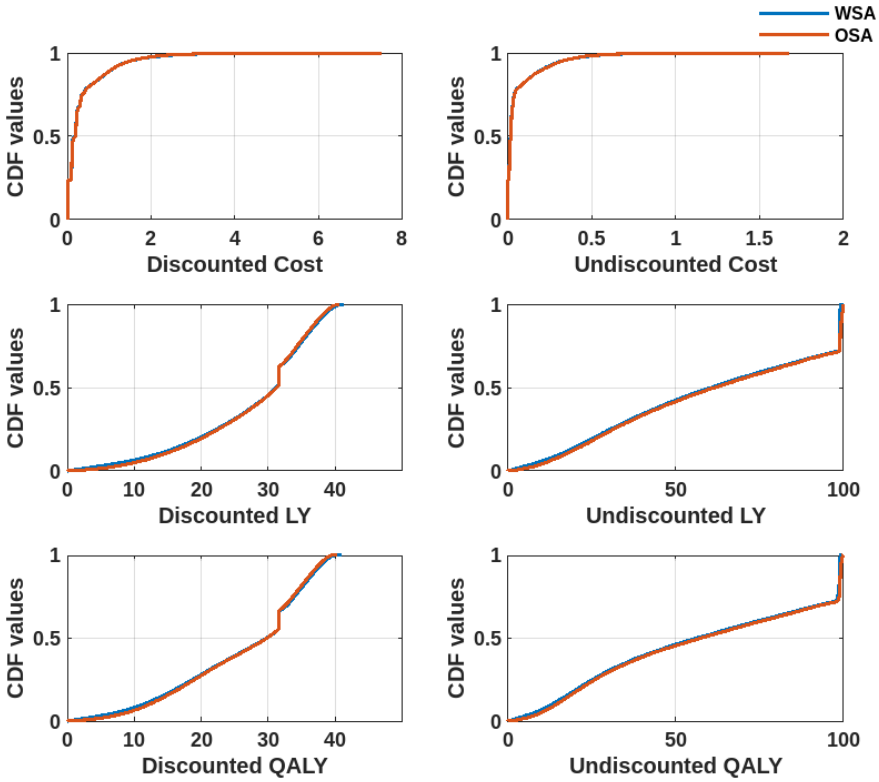

Figure 3: Cumulative distribution functions (CDF) for model outcomes from a single replication, with 50% treatment uptake rate. *Notes.* WS: without sampling & allocation approach, OS: outcomes sampling & allocation approach, LY: Life Years, QALY: quality-adjusted life-year.

The computational experiments described above use a repository that is large enough to yield outcomes that are identically distributed to those from the WSA approach. However, when generating the repository is more computationally intensive, estimating the minimum size of the repository becomes important. This may be done by using power calculations around the hypothesis tests used to determine whether the outcomes from the OSA and WSA approaches exhibit statistically significant differences.

Several future avenues of research can be pursued with regard to this approach. Our approach currently requires incremental accumulation of outcomes for each agent during the execution of the component ABS that are then used to adjust the allocated outcomes to yield the final outcome for an agent. An outcomes repository generation, sampling and allocation scheme that completely obviates the need for incremental accumulation of outcomes may lead to considerable increases in computational runtime reductions. In addition, the impact of intervention uptake rate, which directly influences the number of agents remaining at the end of the model execution time horizon, may also be explored in more detail. Finally, the extension of this approach to hybrid models involving compartmental models of the dynamics of transmission of some phenomenon (including infectious diseases) that may require heterogeneous estimates of outcomes at the individual entity level within certain model compartments may also be investigated.

REFERENCES

Aggarwal, R., Q. Chen, A. Goel, N. Seguy, R. Pendse, T. Ayer *et al*. 2017. "Cost-Effectiveness of Hepatitis C Treatment Using Generic Direct-Acting Antivirals Available in India". *PloS One* 12(5):e0176503.

Ayer, T., C. Zhang, A. Bonifonte, A. C. Spaulding and J. Chhatwal. 2019. "Prioritizing Hepatitis C Treatment in US Prisons". *Operations Research* 67(3):853–873.

Das, S., D. Sen, V. Ramamohan, and A. Sood. 2019. "An Agent-Based Model of Hepatitis C Virus Transmission Dynamics in India". In *2019 Winter Simulation Conference (WSC)*, 984–995.

dos Santos, V. H., K. Kotiadis, and M. P. Scaparra. 2020. "A Review of Hybrid Simulation in Healthcare". In *2020 Winter Simulation Conference (WSC)*, 1004–1015.

Gopalappa, C., S. L. Sansom, P. G. Farnham, and Y.-H. Chen. 2017. "Combinations of Interventions to Achieve a National HIV Incidence Reduction Goal: Insights from an Agent-Based Model". *AIDS* 31(18):2533–2539.

Gupta, S., C. Mittal, S. Das, S. Shriyam, V. Ramamohan and A. Batra. 2021. "A Simulation Model of Breast Cancer Incidence, Progression, Diagnosis and Survival in India". In *2021 Winter Simulation Conference (WSC)*, 1–12.

He, T., K. Li, M. S. Roberts, A. C. Spaulding, T. Ayer, J. J. Grefenstette *et al*. 2016. "Prevention of Hepatitis C by Screening and Treatment in US Prisons". *Annals of Internal Medicine* 164(2):84–92.

He, T., M. Roberts, J. Grefenstette, and J. Chhatwal. 2014. "Cost-Effectiveness of Hepatitis C Screening in United States Prisons: An Agent-Based Approach". *Value in Health* 17(3):A37.

Kar, E., T. Eldabi, and M. Fakhimi. 2022. "Hybrid Simulation in Healthcare: A Review of the Literature". In *2022 Winter Simulation Conference (WSC)*, 1211–1222.

Kim, S.-Y. and S. J. Goldie. 2008. "Cost-Effectiveness Analyses of Vaccination Programmes: A Focused Review of Modelling Approaches". *Pharmacoeconomics* 26:191–215.

Lim, A. G., H. Qureshi, H. Mahmood, S. Hamid, C. F. Davies, A. Trickey *et al*. 2018. "Curbing the Hepatitis C Virus Epidemic in Pakistan: The Impact of Scaling Up Treatment and Prevention for Achieving Elimination". *International Journal of Epidemiology* 47(2):550–560.

Ministry of Home Affairs, Government of India 2011. "C-13: Single Year Age Returns by Residence and Sex, Census of India 2011". https://censusindia.gov.in/census.website/data/census-tables#, accessed 6th May 2024.

Muennig, P. and M. Bounthavong. 2016. *Cost-Effectiveness Analysis in Health: A Practical Approach*. John Wiley & Sons.

Olsen, J. and M. R. Jepsen. 2010. "Human Papillomavirus Transmission and Cost-Effectiveness of Introducing Quadrivalent HPV Vaccination in Denmark". *International Journal of Technology Assessment in Health Care* 26(2):183–191.

Schobert, D., V. Remy, and O. Schoeffski. 2012. "Cost-Effectiveness of Vaccination with a Quadrivalent HPV Vaccine in Germany Using a Dynamic Transmission Model". *Health Economics Review* 2:1–10.

Song, D. L., F. L. Altice, M. M. Copenhaver, and E. F. Long. 2015. "Cost-Effectiveness Analysis of Brief and Expanded Evidence-Based Risk Reduction Interventions for HIV-Infected People Who Inject Drugs in the United States". *PloS One* 10(2):e0116694.

Taguchi, A., K. Hara, J. Tomio, K. Kawana, T. Tanaka, S. Baba *et al*. 2020. "Multistate Markov Model to Predict the Prognosis of High-Risk Human Papillomavirus-Related Cervical Lesions". *Cancers* 12(2):270.

VanDeusen, A., E. Paintsil, T. Agyarko-Poku, and E. F. Long. 2015. "Cost Effectiveness of Option B Plus for Prevention of Mother-to-Child Transmission of HIV in Resource-Limited Countries: Evidence from Kumasi, Ghana". *BMC Infectious Diseases* 15(1):1–10.

Zhou, Y., A. Nikolaev, L. Bian, L. Lin and L. Li. 2021. "Investigating Transmission Dynamics of Influenza in a Public Indoor Venue: An Agent-Based Modeling Approach". *Computers & Industrial Engineering* 157:107327.

AUTHOR BIOGRAPHIES

SOHAM DAS is a Ph.D. student in Industrial Engineering with the Department of Mechanical Engineering at the Indian Institute of Technology Delhi. He holds a B.E. in Mechanical Engineering from Jorhat Engineering College, Assam, India and an M.Tech. in Industrial Engineering from the Indian Institute of Technology Delhi. His e-mail address is mez198426@mech.iitd.ac.in.

APARNA VENKATARAMAN is a UQ-IITD joint Ph.D. student and a Prime Minister's Research Fellow affiliated with the Faculty of Medicine in the Centre for Health Services Research at The University of Queensland (UQ) and the Department of Mechanical Engineering at the Indian Institute of Technology Delhi. Her e-mail address is aprna.v@uqidar.iitd.ac.in.

VARUN RAMAMOHAN is an Associate Professor with the Department of Mechanical Engineering at Indian Institute of Technology Delhi. He holds a a Ph.D. in Industrial Engineering from Purdue University and a B.Tech. from the National Institute of Technology Tiruchirappalli, India. His e-mail address is varunr@mech.iitd.ac.in.