

Predicting epidemics and the impact of interventions in heterogeneous settings: standard SEIR models are too pessimistic

Luc E. Coffeng (MD PhD), Sake J. de Vlas (PhD)

Department of Public Health, Erasmus MC, University Medical Center Rotterdam, Rotterdam, The Netherlands

* Corresponding author: l.coffeng@erasmusmc.nl

Key words: basic reproduction number; infectious disease modelling; estimation; heterogeneity; assortative mixing.

The basic reproduction number (R_0), i.e. the average number of secondary cases produced per primary case in a fully susceptible population, is an established and useful concept to describe the potential for an infectious disease to cause an epidemic, and thus, how intensive interventions must be to effectively prevent an outbreak [1]. When using a deterministic compartmental model to describe transmission of microparasitic diseases, R_0 can be expressed as the product of the average duration of infectiousness times the average number of successful transmission events per time unit (assuming a homogeneous mixing population) or the spectral radius of the next-generation matrix (in case of a structured heterogeneous population) [2,3]. However, deterministic models quickly become too cumbersome when trying to capture many risk strata and/or multiple sources of heterogeneity that are relevant for real-world applications (e.g., individual variation in exposure, disease progression, intervention uptake, age groups, geographical areas, assortative mixing). In such situations, individual-based models (IBMs, sometimes also referred to as microsimulation or agent-based models) are a more convenient and flexible alternative, although they are computationally more expensive. With IBMs, the modeller simulates a finite number of individuals, drawing values of individual characteristics from pre-defined distributions, allowing the inclusion of various heterogeneities. Another benefit of IBMs is that they place no restriction on the probability distribution of sojourn times per disease stage. However, derivation of R_0 for IBMs is not as straightforward as for deterministic models because in IBMs, due to the use of continuous distributions for individuals' characteristics, there are intractably many population states that are challenging to integrate over, especially if multiple sources of heterogeneity are considered. This challenge creates some potential pitfalls when estimating or applying assumptions to model parameters related to the transmission and the effect of interventions with IBMs. We illustrate this by means of a simulation study with an individual-based SEIR model implemented in R package *virsim* (www.gitlab.com/luccoffeng/virsim), which was recently used to explore the impact of a geographically stratified strategy against COVID-19 [4].

We will start by considering a scenario where we have data on the initial exponential growth of an epidemic and where we have external information about the duration of the latency period (average of 4.6 days, with variation between individuals following a Weibull distribution with shape 20) and infectious period (5 days, variation following an exponential distribution), which can be taken to represent COVID-19. These two durations determine the generation interval (average time between onset of infection in infectors and infectees), which together with R_0 determines the initial exponential growth of the epidemic [1]. Let us assume that this initial exponential growth (period between the two black bullets in Figure 1) is well described by a model for a homogeneously mixing population of 10,000 people with a transmission rate of 0.5

day⁻¹, meaning that $R_0 = 5 \times 0.5 = 2.5$ (red line). To reproduce the same initial exponential growth with a model that allows for inter-individual variation in contact rates (assuming a gamma distribution with shape 3.4), the average transmission rate has to be lower: about 0.395 day⁻¹ (green line). Now, $R_0 = 2.5$ as derived from the initial data, is higher than the product of the average transmission rate and the duration of infectiousness. Intuitively, this is obvious as at this early stage of the epidemic, transmission is mostly driven by people with relatively high contact rates. This pattern is slightly more pronounced when we consider assortative mixing (blue line), where people are assumed to interact more with people who have similar contact rates. Here we assume that individual contact rates vary as in the second model variant, but the population is divided in 10 clusters (e.g. villages or neighbourhoods) of 1000 individuals each, and individuals spend 90% of their time in their own cluster and 10% in the population as a whole, and cluster membership is correlated with an individual's contact rate ($\vartheta = 0.26$, in contrast to $\vartheta = 0$ for the first two model variants) [4]. In this case, the same initial exponential growth rate is reproduced with an estimated average transmission rate of 0.386 day⁻¹, which is even (somewhat) lower than that estimated for the model with only inter-individual variation. More noticeable, however, is that the size of the epidemic decreases substantially with increasing assumed levels of heterogeneity, because more high-risk individuals have been "spent" by the epidemic, leaving the rest of the population to support transmission at a slower rate (Figure 1). The slightly drawn-out right tail for assortative mixing does not compensate for the much lower peak, making this the most optimistic scenario regarding overall health impact: 70% cumulatively infected by the end of the epidemic vs. 74% (inter-individual variation only) and 90% (homogeneous population). Clearly, ignoring such individual-level and/or geographic heterogeneities – which do exist in real-world situations – leads to a somewhat pessimistic prediction of the course of the epidemic in the absence of interventions. It also means that when using IBMs, one should be very careful when adopting estimates of transmission rates from other models or when deriving such transmission rates from estimates of R_0 based on other models, which may have relied on other assumptions regarding heterogeneity in transmission.

A second potential pitfall relates to predicting the impact of an intervention with some known or assumed effect on transmission. Let us assume we have estimated the transmission rate given some initial exponential growth (as in Figure 1), and let us now predict the impact of an intervention that reduces all individuals' transmission rates by 50% for the remainder of the epidemic (Figure 2). Because in models with increasing levels of heterogeneity in transmission, high-risk individuals are "spent" earlier during the epidemic, the net effect of an intervention on the epidemic curve will be greater. Where the homogeneous model shows a small peak (red), this peak is lower and narrower for a model with inter-individual variation in contact rates (green), and completely absent in the model with assortative mixing (blue), barring the rise in case numbers around 5 days after implementation of the intervention, which is due to the assumed latency time of infection being 4.6 days. This phenomenon highlights that model predictions for the impact of interventions with some known or assumed effect are also more pessimistic when heterogeneity in transmission is not sufficiently captured. It further means that, again, one should be very careful when adopting estimates of intervention effects from other models, as such estimates are conditioned on assumptions regarding heterogeneity in transmission.

Short communication

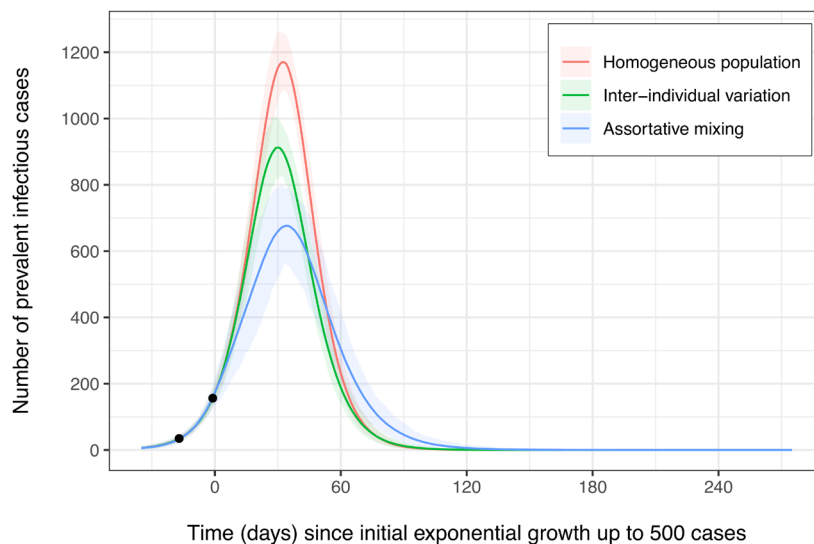


Figure 1. Epidemic curve for five scenarios with regard to heterogeneity in transmission as predicted by a stochastic SEIR model. Each scenario represents an infectious disease with a latency time of 4.6 days and infectious period of five days. Graph lines represent the average of 500 repeated stochastic simulations in a population of 10,000 individuals; shaded bands represent the central 95% percentiles of the stochastic simulation results. For each scenario (coloured lines), the overall transmission rate in the population was calibrated such that the initial exponential growth (from 100 to 500 cumulative cases (i.e., up to 5% of the population), indicated by the black bullets), averaged over repeated simulations, matched that of a simple homogeneously mixing model with $R_0 = 2.5$ (red line). For the scenario with inter-individual variation (green), individual transmission rates were assumed to follow a Gamma distribution with shape 3.4, such that the 2.5th and 97.5th percentiles of the distribution differed by a factor 10. For assortative mixing (blue), we assumed that the population was divided into 10 clusters, where individuals experienced 90% of the force of infection from within their own cluster and 10% from the population as a whole, with individual cluster membership being correlated with an individual's contact rate ($\theta = 0.26$).




Figure 2. Predicted effect of an intervention with a known impact on the transmission rate (50% reduction) after a period of initial exponential growth (black bullets) for three scenarios of heterogeneity. Underlying modelling assumptions and average transmission rates are as in Figure 1. Graph lines represent the average of 500 repeated stochastic simulations in a population of 10,000 individuals; shaded bands represent the central 95% percentiles of the stochastic simulation results.

The third potential pitfall we illustrate pertains to estimating the effect of an intervention from data and projecting that effect forward in time. This time, for “data”, we will consider four weekly data points (open circles in Figure 3) after initiation of the intervention (50% reduction in transmission). These “data” were based on the average of 500 repeated simulation with the homogeneously mixing model with a fixed parameter set as in Figure 2. Let us then recalibrate the transmission rate and intervention effect for each of the three models, using approximate Bayesian computation based on sequential Monte Carlo [5] (<https://gitlab.com/luccoffeng/abcsmc>). The estimated transmission rates were 0.51 (95%-Bayesian credible interval (BCI): 0.45–0.58) for the homogeneous model (red), 0.41 (95%-BCI: 0.35–0.49) with inter-individual variation (green), and 0.43 (95%-BCI: 0.35 – 0.54) for assortative mixing (blue). The estimated effect of the intervention was a 51% reduction (95%-BCI: 39–59%) in the average transmission rate for the homogeneous model (red), or a 49% reduction (95%-BCI: 39–61%) when assuming inter-individual variation (green), and a 46% reduction (95%-BCI: 31–60%) when assuming assortative mixing. This phenomenon highlights that multiple combinations of the average transmission rate, level of heterogeneity in transmission, and the effect of interventions can more or less reproduce the data equally well; more heterogeneity corresponds with a lower estimated impact of the interventions. Now, if we project the impact of these interventions forward in time, assuming that these interventions are continued indefinitely (solid lines in Figure 3), our three alternative assumptions about heterogeneity result in qualitatively rather similar predictions that would have the same implications for policy: this intervention is effective at controlling the epidemic, taking in the order of five to six months.

If we consider the possibility of fully lifting interventions when a target prevalence of, say, 100 prevalent infectious cases has been reached (dashed lines in Figure 3), our three alternative assumptions about heterogeneity result in distinctly different outcomes that have completely different implications for public health: a pronounced second wave that might cause problems for health care (homogeneous model, red); a smaller, more “manageable” second wave (inter-individual variation, green); or no second wave at all but a gently tapering epidemic tail (assortative mixing, blue). Even with additional data on the decline in infection numbers after the first peak, it would be almost impossible to distinguish which of the three alternative assumptions is most plausible. These different outcomes are driven by two factors: (1) higher estimated overall transmission rate for the homogeneous model (red); and (2) selection and depletion of high-risk susceptible individuals during the initial phase of the epidemic in the models with inter-individual variation (green) and assortative mixing (blue), resulting in 14% lower (green) and 19% lower (blue) average contact rates of remaining susceptible individuals at day 60.

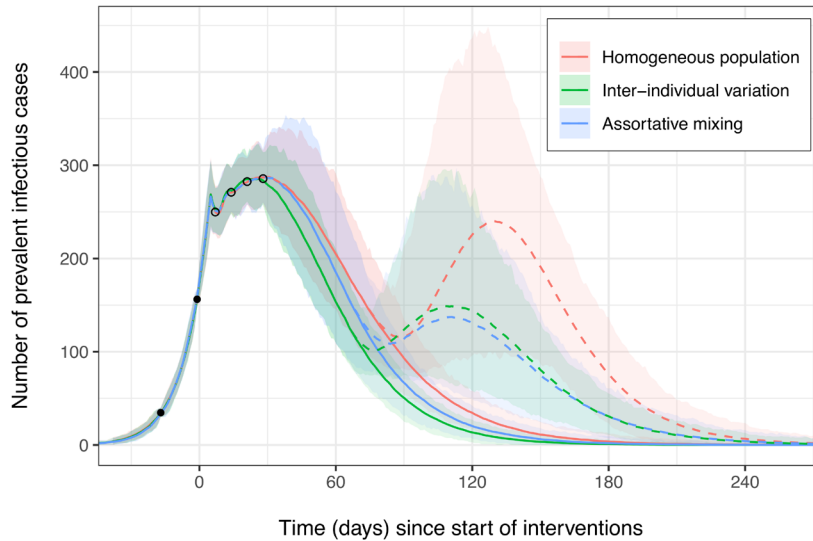


Figure 3. Predicted long-term effect of an intervention with an unknown impact that has to be estimated from data (open circles) after a period of initial exponential growth (black bullets). The solid lines represent a scenario where interventions are continued indefinitely. The dashed lines represent the trajectories if all interventions are suspended when a target of <100 prevalent infectious cases has been reached. Shaded areas represent 95%-Bayesian credible intervals that capture uncertainty about both the transmission rate and effect of interventions.

It has long been recognised that estimates of R_0 and predictions for the impact of interventions can be highly sensitive to model structure and assumptions about unknown or uncertain model parameters [6–8]. Also, individual heterogeneity has been shown to be important for disease emergence and superspreading events, as demonstrated for e.g. SARS [9]. However, to our knowledge, no efforts have been made to investigate the impact of heterogeneity on the later stages of epidemic control. Lloyd previously showed for deterministic SIR and SEIR models that, given the initial exponential epidemic growth rate, estimates of the transmission rate, R_0 , and the associated required effectiveness of interventions are higher when assuming higher variation in individual sojourn times [7]. We add to this with our finding that more highly heterogeneous contact patterns drive estimates of the transmission rate down (i.e., into the opposite direction) due to selection and depletion of high-risk individuals during the initial stages of an epidemic. Lloyd could not have discovered this with deterministic models, as these (unlike IBMs) assume that the distribution of characteristics of remaining susceptible individuals remains the same over the entire course of an outbreak.

Network models are an increasingly popular technique to capture heterogeneity in contact patterns [8, 10], explicitly capturing variation in contact frequencies between pairs of persons. However, most network models assume that the quality and thus transmission potential is the same for all contacts, and thus still underestimate the level of heterogeneity in transmission [10]. To capture the selection process that we illustrate here, a network model would have to also capture variation in quality of contacts via weighted edges. This would lead to network nodes with highly weighted edges to be more likely to be selected and depleted during the initial stage of an epidemic, leading to a smaller epidemic size. Of course, quantifying the distributions of edge weights comes with a major data collection challenge; typically, we can only recognise a limited number of different contacts types (e.g., family, neighbour, colleague, classmate), and even these may vary considerably between individuals in terms of transmission potential.

In summary, with a simple simulation exercise we illustrate the consequences of not adequately modelling heterogeneity when predicting the impact of epidemics and the effects of interventions. If heterogeneity is ignored (i.e., no or too little inter-individual variation or assortative mixing), then that model will to some degree overestimate the transmission rate and

the potential course of the epidemic. For instance, this may happen when basing the average transmission rate in a heterogeneous model on estimates of R_0 that are derived from case numbers and the generation interval via a standard procedure [11]. This phenomenon fits the more general notion that estimates of R_0 are specific to the model type and structure used, as well as assumptions about parameters that may be hard to estimate [12]. We further show that when models do not capture enough heterogeneity, predictions for the impact of interventions on case numbers are relatively pessimistic. If the effect of an intervention on the transmission rate needs to be estimated, this is overestimated when the model captures less heterogeneity than present in reality. This will have few consequences for the predicted long-term impact if such an intervention is continued over an extended period of time. However, should such an intervention be suspended, say, after reaching a target, the potential for a second epidemic wave will depend strongly on assumptions about heterogeneity, with more heterogeneity resulting in lower remaining epidemic potential, due to selection and depletion of high-risk individuals during the early stages of the epidemic. This phenomenon has likely also affected current model predictions regarding COVID-19, as most transmission models assume homogeneous mixing or at most employ a simple age-stratification. The lack of accounting for a sufficient degree of inter-individual variation and/or assortative mixing may have led to over-cautious predictions of durations of lock-downs and required vaccine coverage levels.

Acknowledgements

The authors thank Nico Nagelkerke for critically reading the text.

Funding

The authors acknowledge funding from ZonMw for the Dutch COVID-19 Monitoring Consortium (www.zonmw.nl, grant 10430022010001). LEC acknowledges funding from the Dutch Research Council (NOW, <https://www.nwo.nl/en>, grant 016.Veni.178.023).

References

1. Diekmann O, Heesterbeek H, Britton T (2013) *Mathematical Tools for Understanding Infectious Disease Dynamics*. New Jersey: Princeton University Press.
2. Diekmann O, Heesterbeek J a P, Roberts MG (2010) The construction of next-generation matrices for compartmental epidemic models. *J R Soc Interface* **7**: 873–885.
3. Inaba H (2012) On a new perspective of the basic reproduction number in heterogeneous environments. *J Math Biol* **65**: 309–348.
4. de Vlas SJ, Coffeng LE (2021) Achieving herd immunity against COVID-19 at the country level by the exit strategy of a phased lift of control. *Sci Rep* **11**: 4445.
5. Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MPH (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J R Soc Interface* **6**: 187–202.
6. Anderson RM, May RM (1991) *Infectious Diseases of Humans: Dynamics and Control*. Oxford & New York: Oxford University Press.
7. Lloyd AL (2009) Sensitivity of Model-Based Epidemiological Parameter Estimation to Model Assumptions. *Mathematical and Statistical Estimation Approaches in Epidemiology*. Dordrecht: Springer Netherlands. pp. 123–141.

Short communication

8. Bansal S, Grenfell BT, Meyers LA (2007) When individual behaviour matters: homogeneous and network models in epidemiology. *J R Soc Interface* **4**: 879–891.
9. Lloyd-Smith JO, Schreiber SJ, Kopp PE, Getz WM (2005) Superspreading and the effect of individual variation on disease emergence. *Nature* **438**: 355–359.
10. Keeling MJ, Eames KTD (2005) Networks and epidemic models. *J R Soc Interface* **2**: 295–307.
11. Wallinga J, Lipsitch M (2007) How generation intervals shape the relationship between growth rates and reproductive numbers. *Proc Biol Sci* **274**: 599–604.
12. Delamater PL, Street EJ, Leslie TF, Yang YT, Jacobsen KH (2019) Complexity of the basic reproduction number (R_0). *Emerg Infect Dis* **25**: 1–4.