The impact of varying class sizes on epidemic spread in a university population

Alex Best, Prerna Singh, Charlotte Ward, Caterina Vitale, Megan Oliver, Laminu Idris, Alison Poulston

School of Mathematics and Statistics, University of Sheffield, Sheffield, S3 7RH, UK

Abstract

A common non-pharmaceutical intervention (NPI) during the Covid-19 pandemic has been group size limits. Further, educational settings of schools and universities have either fully closed or reduced their class sizes. As countries begin to reopen classrooms, a key question will be how large classes can be while still preventing local outbreaks of disease. Here we develop and analyse a simple, stochastic epidemiological model where individuals (considered as students) live in fixed households and are assigned to a fixed class for daily lessons. We compare key measures of the epidemic - the peak infected, the total infected by day 180 and the calculated R_0 - as the size of class is varied. We find that class sizes of 10 could largely restrict outbreaks and often had overlapping inter-quartile ranges with our most cautious case of classes of 5. However, class sizes of 30 or more often result in large epidemics. Reducing the class size from 40 to 10 can reduce R_0 by over 30%, as well as signficantly reducing the numbers infected. Intermediate class sizes show considerable variation, with the total infected varying by as much as from 10% to 80% for the same class size. We show that additional in-class NPIs can limit the epidemic still further, but that reducing class sizes appears to have a larger effect on the epidemic. We do not specifically tailor our model for Covid-19, but our results stress the importance of small class sizes for preventing large outbreaks of infectious disease.

1. Introduction

The classic Susceptible-Infected-Recovered (SIR) epidemiological model has long been used to model the spread of infectious disease in human, animal and plant populations (Kermack and McKendrick, 1927; Anderson and May, 1979). More recently, its extended SEIR (Susceptible-Exposed-Infected-Recovered) framework has formed a central pillar of much of the modelling of the Covid-19 pandemic, often including highly realistic movement and contact networks (Ferguson et al, 2020; Kucharski et al, 2020; Firth et al, 2020; Kain et al, 2021). A key non-pharamceutical intervention (NPI) for populations across the world during the Covid-19 pandemic has been restricting population mixing through 'lockdowns', with people encouraged to stay at home and avoid mixing with individuals outside their household unless essential. This has often included closing educational settings of universities and schools, with 31 countries enacting full school closures and reduced schooling in a further 48 countries (UNESCO, 2020). As countries move to reopen these settings, an important question is how classes can be organised to minimise further disruption to students' education whilst limiting epidemic spread. There have been some excellent, in-depth modelling studies of infection spread in educational settings, especially universities, with a range of NPIs included, often with a focus on testing and isolation strategies (Bahl et al, 2020; Brook et al, 2020; Cashore et al, 2020; Lopman et al, 2020). Here we focus on the question of how class sizes may impact an epidemic.

An important measure of infectious disease growth and severity is the basic reproductive ratio, R_0 (Anderson and May, 1982; Heesterbeek, 2002). This well-known term defines the average number of new infections from one infected individual in an otherwise disease-free population. In mathematical definitions, R_0 is broadly the product of three quantities: the potential (or probability) of infection upon an infectious contact, the infectious period and the number of disease-free individuals contacted per unit time (Dietz, 1993). An important consequence of this is that the larger the population that can be

contacted by an infectious individual - the effective population size - the greater the potential spread of the epidemic. Thus, if a population can be partitioned in to smaller sub-groups with minimal mixing between them, the spread of disease can be significantly limited. This is a reality populations all over the world have experienced during the Covid-19 pandemic, with 'lockdown' measures aimed at limiting mixing of households or groups. Data suggests that such NPIs - such as closing businesses, closing schools and, of relevance to our study, limiting group gathering sizes - have reduced the R_0 of Covid-19 by as much as 60% (Brauner et al, 2020). While there are many studies examining epidemic spread in heterogeneous environments and/or on networks, to our knowledge few modelling studies have specifically explored how group size limits might impact the extent and severity of an epidemic. Kain et al (2021) found that 'chopping off the tail' of individual infection distributions - in effect preventing large gatherings - could effectively restrict an epidemic in their Covid-19 parameterised model, reducing both mean transmission and the variance in outcomes. In their broader examination of Covid dynamics on a university campus, Brook et al (2020) found that a group size limit of 6 could reduce the effective reproductive ratio from 1.05 to 0.86, but limits of 50 had almost no impact on disease spread. As we seek to reopen universities and schools, some mixing between households will be essential. An important consideration, then, is the degree to which we could allow some mixing whilst still limiting the extent of the epidemic.

Ultimately, our study investigates what happens to an epidemic when a population partitioned into households is mixed for a short time period each day into fixed groups. The situation loosely in our minds is of a university cohort living in accommodation who attend a class each day. Similarly, we might consider a school population attending classes, or a local community forming interaction bubbles. However, we stress that our study is a relatively simple, theoretical study of the impact of mixing, and we make no strong claims about the precise values or predictions our model makes. For clarity

and transparency we state here a few of our key assumptions: (i) all classes occur simultaneously with immediate transitions between class and home, (ii) we do not include testing and isolation, (iii) there is no further mixing than in classes and homes, either within the population or externally, (iv) we do not include 'superspreaders' but instead assume all individuals experience the same transmission rates. We discuss the possible impact of each of these assumptions in the Discussion. We also do not attempt to parameterise or structure our model specifically for Covid-19. Rather, we seek to identify the general patterns that result from mixing partititioned populations into different sized groups.

2. Methods

We develop and run stochastic simulations of an epidemiological model using python (code is available from GitHub, https://github.com/abestshef/classsizeSEIR). The underlying epidemiological model is an SEIR (Susceptible-Exposed-Infected-Recovered) framework where, within a setting ('home' or 'class'), the dynamics would be given by the following ordinary differential equations,

$$\frac{dS}{dt} = -\beta SI \tag{1}$$

$$\frac{dE}{dt} = \beta SI - \omega E \tag{2}$$

$$\frac{dI}{dt} = \omega E - \gamma I \tag{3}$$

$$\frac{dR}{dt} = \gamma I, \tag{4}$$

where β is the transmission coefficient (with βI the 'force of infection'), ω is the rate of progression from exposed to fully infected and γ is the recovery rate. We note that our state variables are defined as densities, not proportions, and that the total population size is $N = S + E + I + R$ rather than being normalised to 1. The stochastic simulations use a Gillespie algorithm (Gillespie, 1977) to calculate waiting times between events. The possible events are initial infection ($S \to E$), progression to full infection

($E \to I$) and recovery ($I \to R$). Which event occurs at a chosen time point depends on their relative probabilities at that point and in the relevant setting. Each day is divided in to fixed time periods where all students are in each setting, either home or class. We assume immediate movement between settings, with classes occuring during $t \in [\text{day} + 0.4, \text{day} + 0.5]$, roughly equivalent to a 2.5 hour period. The event probabilities will be different in each setting; thus when a transition time is reached, the waiting time is stopped and recalculated from the transition point. Besides these transitions, infection is assumed to occur evenly throughout the 24 hour period.

The total population size is $N = 1000$ and students are randomly divided in to houses of size n_h and classes of (target) size n_c. We take two household sizes, $n_h = 10$ (100 households) and $n_h = 5$ (200 households; arguably reasonable averages for university halls and private housing respectively). Every household has exactly n_h individuals. Target class sizes vary in steps of 5 from 5 to 50. Where N/n_c is an integer, exactly that many classes are created, all containing exactly n_c individuals. Where this is not the case, the classes contain either n_c or $n_c - 1$ individuals (for example, for a target class size of 15, there are 62 classes of 15 students and 5 classes of 14 students). Both the house and class composition is fixed in each simulation. We randomly choose 25 individuals to be infected at the beginning of each simulation, and all other individuals are susceptible (we do not investigate the effect of varying this initial number). Our key investigation will be to vary average class sizes and explore the impact on the epidemic. We also compare results where transmission is high ($\beta = 0.5\gamma$) and low ($\beta = 0.2\gamma$). These values are chosen partly to give reasonably realistic estimates for R_0 and also because the lower value represents a 60% reduction of the higher value, which we use to represent a reduction due to NPIs (see below). While these values would appear to produce very high values of the basic reproductive ratio, R_0, in the mean-field model ($\beta = 0.5\gamma, N = 1000 \implies R_0 = \beta N/\gamma = 500$), it is well known that the actual R_0 is considerably lower in individual-based models,

especially when interactions networks are small (Keeling and Grenfell, 2000). We directly calculate R_0 from our simulations (see below) and found across all the results presented here the median R_0 fell in the range [0.66,3.68]. We also additionally examine the case where NPIs in the class (e.g. masks, ventilation, distancing) reduce transmission from the high to the low value (a reduction of 60%). We additionally assume $\gamma = 1/14$ and $\omega = 1/7$ in all simulations, giving a latent period of 7 days and infectious period of 14 days.

Recent work has highlighted the difficulties in representing outcomes from stochastic epidemic models (Juul et al, 2020). First, to visualise the 'typical' time courses, we follow the methods of Juul et al (2020) to present the 'most central' 50% of simulation runs. 100 simulations are run, discretised and stored. We then repeatedly sample subsets of these stored runs (100 samples of 20 curves) and increase the 'score' of any run that falls entirely within the bounds of the sampled curves between time-points 10 and 150. Secondly we present three key measures of the epidemic - the peak number infected, the total number infected by day 180 and the calculated R_0 (see below) - from 100 simulation runs for each class size using box and whisker plots. These highlight the median values, the inter-quartile range (IQR; 25%-75%), the maximum/minimum (or 1.5×IQR if smaller) and any outliers (values greater than 1.5×IQR). Alongside these, we compare the IQR of the class of 5 (the 'most cautious' approach) with all other class sizes, noting where the IQRs do and do not overlap using shading of the boxplots. This allows us to explore whether class sizes can be raised above this cautious level without causing large changes to the outcome.

A brief note on the basic reproductive ratio, R_0; in our simple SEIR structure we would have $R_0 = \beta\hat{N}/\gamma$, which depends on the effective disease-free population size $\hat{N}$. However, the interpretation of $\hat{N}$ will vary depending on the degree of mixing. Here we make a direct calculation of R_0 in each simulation by recording the number of infections caused by the 25 initially infected individuals, an intuitive measure of R_0 as might be estimated during

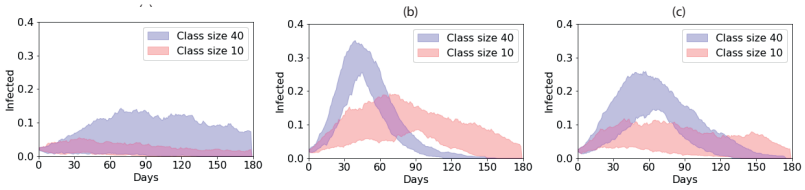

Figure 1: Bounds of the 'most central' 50% of 100 simulation runs for large average households for average class sizes of 40 (blue) and 10 (red) where infection rates are (a) low, $\beta = 0.2\gamma$, (b) high, $\beta = 0.5\gamma$ and (c) high at home but low in class.

a real epidemic.

3. Results

3.1. Large households

Taking an average household size of 10 and comparing the most central 50% of runs for average class sizes of 10 and 40 (figure 1a-c), it is clear that smaller class sizes substantially restrict the epidemic. When infection rates are low ($\beta = 0.2\gamma$, figure 1a), with a class size of 40 there is considerable variability, with some of the central curves showing minimal spread but others reaching peaks above 15% infected. Reducing the class size to 10 clearly restricts the central epidemics, with few curves peaking above 5% infected and in some cases the epidemic completely finishing by day 60. For greater infection rates ($\beta = 0.5\gamma$, figure 1b) there is a clear epidemic in all of the central runs for any class size, but is clearly more severe with the larger groups, with the peak of the central runs increasing from never more than 20% for a class size of 10 to always more than 26% for a class size of 40. Finally we investigate the impact of having simple NPIs in place in classes such that the infection is reduced (from $\beta = 0.5\gamma$ to $\beta = 0.2\gamma$) while in class but not at home. Compared to the previous case we do see reductions in the epidemic, with the peaks lowered by around 10%. Noticeably, however, solely reducing the class size from 40 to 10 (figure 1b blue v red) causes a greater reduction in the epidemic than solely instituting the in-class NPIs in a class of 40 (figure 1b red v figure 1c blue).

Looking in more detail for varying class sizes using the boxplots, with low infection rates (figure 2a-c) we again clearly see that greater class sizes lead to larger epidemics in terms of all three measures. The colouring highlights that only a class size of 10 has an overlapping IQR for both peak and total infections, while class sizes up to 25 have an overlapping IQR for peak infections only, meaning sizes of 30 or above have clearly different outcomes to a class of 5. Moreover, for a class size of 10 the top of the inter-quartile range (IQR) is 19% total infecteds, but for a class size of 40 the bottom of the IQR is 58%, emphasising the large effect of different sizes. While all class sizes have overlapping R_0 IQRs with the class of 5, since this is essentially a logarithmic quantity of epidemic growth it is not unexpected, and the median value is reduced from 1.68 for a class size of 40 to 1.40 for a class size of 10.

We see considerable variation for intermediate class sizes, with the minimum and maximum total infected for classes of 35 and 40 extending from below 10% to above 80%, suggesting different locations could experience very different epidemics purely due to stochastic variation. We also note that we assumed the same number of initial infections in all simulation runs. Different institutions would likely start with different number of infections, which would further increase the heterogeneity of outcomes.

When infection rates are larger (figure 2d-f), there are very large epidemics no matter the class size, especially in terms of the total number infected. There are clearly no class sizes where the IQRs overlap with the smallest class for the peak and total infected, while only class sizes of 30 or smaller have overlapping IQRs for R_0. A class size of 15 or above results in more than 94% of the population infected in every single simulation run. Smaller class sizes do lead to noticeabley lower peaks, however - the top of the IQR for class sizes of 15 is 22% and for a class size of 10 it falls to 17%. Reducing the class size from 40 to 10 also leads to a drop in the median R_0 from 3.40 to 2.64 (a 22% reduction).

In-class NPIs lead to a modest reduction in the severity of the epidemic at all

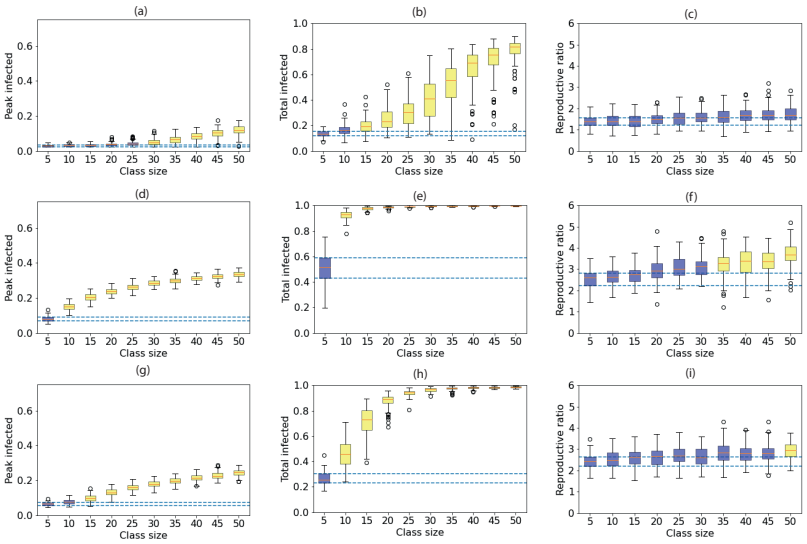

Figure 2: Measures of epidemic severity for household size of 10 for different (target) class sizes from 100 simulation runs at each class size. Boxplots of the peak (a,d,g) and total infected by day 180 (b,e,h) and the calculated R_0 (c,f,i), where infection is low, $\beta = 0.2\gamma$ (a,b,c), high, $\beta = 0.5\gamma$ (d,e,f) and high at home but low in class (g,h,i). The orange lines denote the median, the boxes the 25th and 75th centiles, the whiskers to 1.5 the interquartile range and circles any outliers. The dashed lines mark the inter-quartile range (IQR) for the class of 5, and colouring of the boxes whether the IQRs of each class size do (blue) or do not (yellow) overlap the class of 5's IQR.

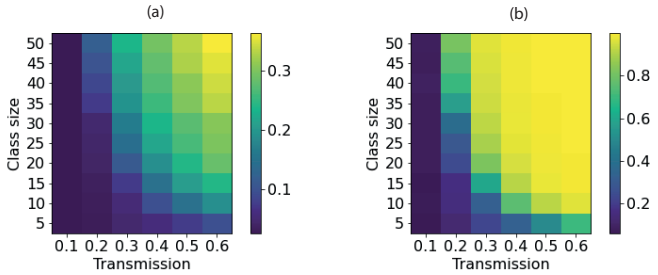

Figure 3: Heatmaps of median peak and total infections for varying (target) class sizes and transmission (where β is the figure given multiplied by γ. The household size was 10 and there were 50 simulation runs for each combination of class size and tarnsmission rate.

class sizes (figure 2g-i), though the epdiemic remains signficiantly larger for larger class sizes. Interestingly, a class size of 10 with no NPIs (peak IQR 14%-17%, total IQR 91%-95%) generally results in smaller epidemics than a class size of 40 with NPIs (peak IQR 21%-23%, total IQR 97%-99%). Thus group size limits in themselves may lead to better outcomes than many other mitigation measures (given our assumptions). The combination of small class sizes and the in-class NPIs can dramatically reduce the severity of the epidemic. Comparing a class size of 40 without NPIs to a class size of 10 with NPIs, the median peak is reduced from 32% to 8%, the median total from 99% to 64% and the median R_0 from 3.40 to 2.48. The class size of 10 has an overlapping IQR with the class of 5 for the peak infected, but no classes overlap for total infected.

For completeness, we also ran simulations with a finer resolution of transmission coefficient, β, and recorded the median peak and total infected (figure 3). At the lowest transmission ($\beta = 0.1\gamma$) there are no outbreaks for any class size. However, as the transmission increases we clearly see the increasing important of smaller class sizes, particularly in terms of total numbers infected.

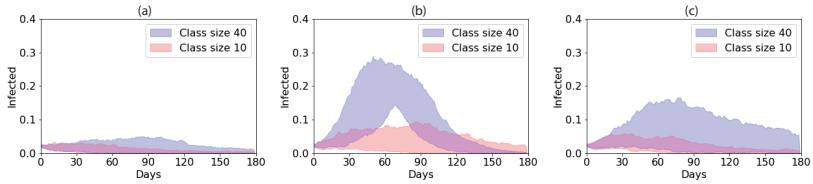

Figure 4: Bounds of the 'most central' 50% of 100 simulation runs for small average households for average class sizes of 40 (blue) and 10 (red) where infection rates are (a) low, $\beta = 0.2\gamma$, (b) high, $\beta = 0.5\gamma$ and (c) high at home but low in class.

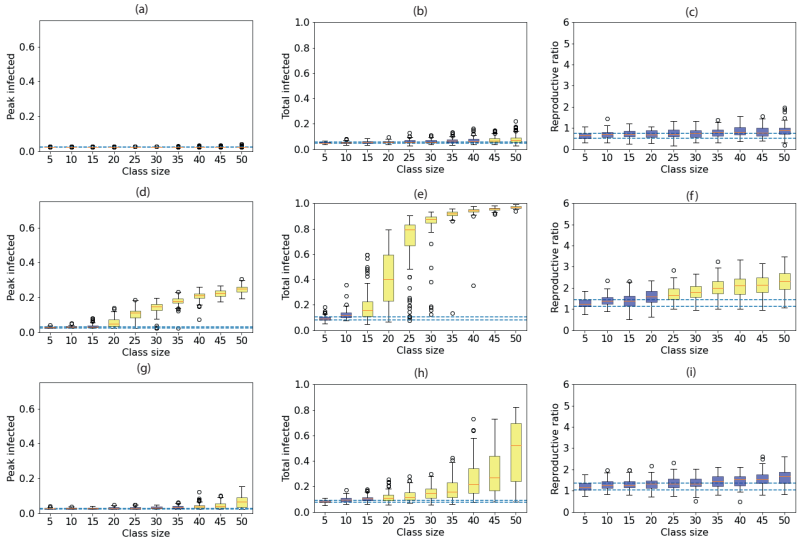

Figure 5: Measures of epidemic severity for household size of 5 for different (target) class sizes from 100 simulation runs at each class size. Boxplots of (left) the peak and (middle) total infected by day 180 and (right) the calculated R_0, where infection is (top) low, $\beta = 0.2\gamma$, (middle) high, $\beta = 0.5\gamma$ and (bottom) high at home but low in class. The orange lines denote the median, the boxes the 25th and 75th centiles, the whiskers to 1.5 the interquartile range and circles any outliers. The dashed lines mark the inter-quartile range (IQR) for the class of 5, and colouring whether the IQRs of each class size do (blue) or do not (yellow) overlap the class of 5's IQR.

3.2. Small households

When the average houshold is reduced to 5, comparing figures 1 and 4 shows that the size of the epidemic is reduced in all cases, since there is naturally less mixing between individuals. When transmission is low ($\beta = 0.2\gamma$, figure 4a, 5a-c) there are no significant outbreaks for any class size. For all class sizes considered the median R_0 is less than 1, and the peak is lower than 5% in every simulation run. All class sizes' peak IQRs overlap with the class of 5's IQR, and classes of 40 or smaller have overlapping IQRs for total infected.

We see dramatic impacts of reducing the class size when transmission is higher ($\beta = 0.5\gamma$, figure 4b, 5d-f). Reducing the class size from 40 to 10 reduced the median R_0 from 2.12 to 1.40 (a 33% dcrease). Even classes of 25 lead to significant outbreaks with the bottom of the IQR being 67% for the total and 8% for the peak, whereas for a class size of 10 the top of the IQR is 14% for total infected and 3% for the peak. Compared to the class of 5, only classes of 10 and 15 have an overlapping IQR for the peak infected and only classes of 10 overlap for total infected. We again see considerable variation in outcomes for fixed class sizes, with total infected in class sizes of 20 and 25 stretching from a minimum of below 10% to a maximum of above 80%.

When NPIs are included in the class setting (figure 4c, 5g-i), for a class size of 40 the median peak is reduced from 21% without the NPI to 4% with it, and the median R_0 from 2.12 to 1.52. These figures make it roughly equivalent to a class size of 15 without NPIs. We again see that the epidemic is more severe for a class size of 40 with NPIs than for a class size of 10 without (respective median peaks: 4% vs 3%, median totals: 22% vs 12%, median R_0: 1.52 vs 1.40). Combining both smaller classes and in-class NPIs can have substantial impacts: comparing a class size of 40 without NPIs to a class size of 10 with NPIs, we see the median peak reduced from 21% to 3%, the median total infected from 94% to 10% and median R_0 from 2.12 to 1.26. Compared to the class of 5, class sizes of 35 and below have overlapping IQRs for the peak, but only class sizes of 10 and 15 for the total infected. We also see significant

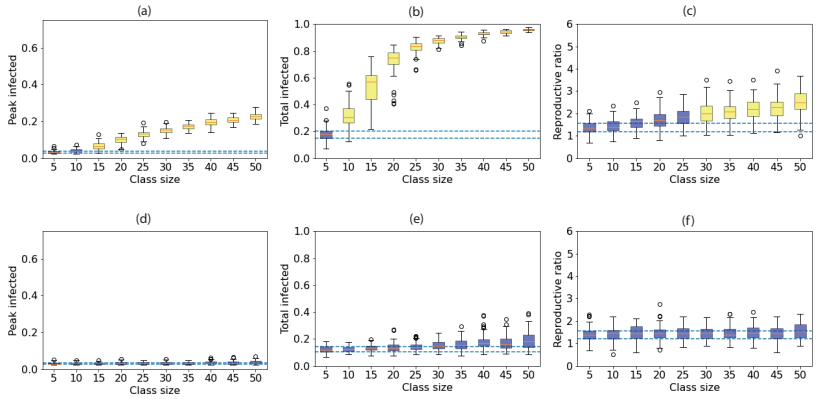

Figure 6: Box plots for peak infected, total infected and R_0 whenthe class length is doubled (a)-(c) or halved (d)-(f). $\omega = 1/7$, $\gamma = 1/14$. $\beta = 0.2\gamma$ in both settings. See earlier figures for description.

variation in outcomes for larger class sizes.

3.3. Time in class

We additionally consider what happens when we alter the time in class to be double ($t \in [\text{day} + 0.4, \text{day} + 0.6]$) or halved ($t \in [\text{day} + 0.4, \text{day} + 0.45]$) compared to above, assuming large housesholds ($n_h = 100$) and low infection rates ($\beta = 0.2\gamma$). Predictably, increasing the class length leads to larger epidemics and decreasing it leads to smaller ones (figure 6). For the shorter classes, all class sizes' peak and total IQRs overlap with the class of 5. Interestingly R_0 is found to be similar for any class size, the median varying only from 1.40 to 1.54. In contrast, no class sizes have overlapping IQRs for total infected and only classes of 10 for peak infected when the class length is doubled.

4. Discussion

As might be expected, a clear result from our model is that the smaller the class size, the lesser the severity of the epidemic. For our paramater sets, reducing the class size from 40 to 10 showed reductions in the median R_0 of up

to 30%. Within that, our results suggest that 'optimal' class sizes across all measures of epidemic severity will rarely exist. Instead, decisions may vary depending on the aim, for example whether to ensure a low peak (to limit pressure on health services) or a low number of total infections (to protect as many individuals as possible from infection). Broadly speaking, small increases in class size from small groups initially lead to only modest increases in the peak of infections but rapid increases in total infections. For example, for high infection rates and small households, increasing the class size just from 10 to 15 led the median peak to increase from 5% to 9% but the median total infected from a modest 27% to a substantial 66%. The average household size and time in class also impacted the severity of outbreak, with larger households and longer classes predictably increasing the potential for large epidemics. In these cases small class sizes were even more essential. Institutional decisions are therefore likely to depend on the desired outcome and specific local conditions, and transparency of decision making will be crucial.

If we were to trust our values here as being representative of a real university (which we would caution should be in the light of the many assumptions underlying the model), we would suggest that to ensure the best chance of a restricted epidemic then classes should be limited to 10. In many cases this would have an overlapping inter-quartile range with - and thus not be clearly different to - the most cautious approach of classes of 5 for the peak and total infected. Slightly larger classes - up to 25 - may prevent the peak from rising too high, but will likely result in large numbers of total infections. We would strongly recommend against larger class sizes than 25 based on our parameters and assumptions as both peak and total infected were then consistently clearly different to the most cautious case of a class of 5. In their more detailed study applied to a specific institution, Brook et al (2020) similarly found that small group size limits were the key NPI to reducing infection, showing that a limit of 6 could reduce the effective reproductive ratio from 1.05 to 0.86, but limits

of 50 had almost no impact on disease spread. Similarly, Kain et al (2021) found that 'chopping off the tail' of individual transmission distributions - effectively preventing the grouping of large numbers of individuals - could be a key control measure. Moreover, based on data, Brauner et al (2020) found that restricting gatherings in all settings to 10 people or fewer was one of the most succesful measures at reducing R_0 during Covid-19. It thus appears a consistent result that limiting group sizes to around 10 can be a succesful NPI for slowing or even stopping epidemics.

The stochastic simulations reveal considerable variation in the epidemic time courses, particularly for intermediate class sizes. In some cases, the maximum and minimum of total infecteds in the 100 simulations spread from less than 20% to more than 80%. Thus while methods exist for approximating individual-based models with deterministic systems of ordinary differential equations (Matsuda et al, 1992; Keeling, 1999; Sharkey, 2008), we highlight the importance of using stochastic simulations to appreciate the variety of possible outcomes. In practice, this demonstrates that institutions may make the same decisions about class sizes but experience very different epidemic time courses. As such, institutions may need to consider their 'risk appetite' for organising logistically easier bigger class sizes at the risk of a large epidemic. Again, though, we emphasise that limiting class sizes to 10 or fewer largely prevented significant epidemics.

We investigated a simple case of employing non-pharmaceutical interventions (NPIs) in the class setting that would reduce the infection rate by 60%. We note that this would be a rather strong effect compared to estimates of NPI impacts from data (Brauner et al, 2020; Chu et al, 2020). This reduction was assumed to be due to simple NPIs such as social distancing but the exact method was not explicitly modelled. We found that, as would be expected, this led to smaller epidemics than if no NPIs were present, reducing R_0 by up to 25% for large class sizes. However, we consistently found that the epidemic was smaller in a class size of 10 without NPIs than a class size of 40 with

NPIs. Given that the 60% reduction due to the NPI is already rather strong, this would suggest that reducing the class size may be the most efficient control measure. Both models (Brook et al, 2020) and data analysis (Brauner et al, 2020) have similarly found that group size limitations was one of the most effective NPIs to prevent spread of Covid-19. Of course, both reducing the class size and implementing the NPI could reduce the epidemic considerably, in some cases reducing the median total infected from 96% to just 14% in our model.

As we have stated, we have relatively modest ambitions in this study of exploring how the size of mixing groups (considered as classes here) impact the time course of epidemics in partitioned populations. Other studies have provided highly detailed analyses of models with many realistic assumptions, contact networks and NPIs included (Ferguson et al, 2020; Kucharski et al, 2020; Firth et al, 2020; Kain et al, 2021) including in university settings (Bahl et al, 2020; Brook et al, 2020; Cashore et al, 2020; Lopman et al, 2020). If our study were to be applied to more realistic scenarios or to form the basis of decision making, some key further additions would be necessary. There are some key additions we would highlight. Firstly, we have assumed that all classes occur simultaneously and that transitions are immediate. In reality classes would likely be spread throughout the day and there would be unavoidable mixing during transitions. Spreading the classes out would lower the effective size of the household population for much of the day, potentially reducing infection, while increased mixing during transitions would act oppositely. Secondly we should consider additional NPIs, most importantly isolation of symptomatic (and possibly asymptomatic) individuals, as have been included in other models (Bahl et al, 2020; Brook et al, 2020; Cashore et al, 2020; Lopman et al, 2020). Given that any sort of isolation - whether due to infection or simply imposed - will reduce the degree of mixing and effective population size, such an approach would clearly be expected to further limit the epidemic. Also, we have assumed a closed population with no

mixing outside of households or classes and full adherence by the population. We should account both for further external contacts and for additional mixing between individuals due to socialising, which we would expect to increase the potential for infection. We also assumed that all individuals have the same transmission rate when infected, but there is likely to be individual variation leading to 'super-spreaders' (Lloyd-Smith et al, 2005) and estimates from Covid-19 data suggest a high degree of super-spreading occurring (Wang et al, 2020). This is likely to lead to greater variation in outcomes, with super-spreading often leading to rarer but larger outbreaks (Lloyd-Smith et al, 2005). Further, we assume that all individuals within each household attend exactly one class with the same cohort of classmates. We do not account for the fact that (i) students may live in households with individuals who do not attend classes (family or other friends, for example) or (ii) that students may attend multiple classes. The latter in particular will lead to increased mixing of the population and thus increased transmission. While these additions would undoubtedly change the quantitative values found here, we would expect the fundamental findings - that smaller class sizes lead to smaller epidemics with less variation and that the patterns will vary according to the target measure - will remain.

Acknowledgements

Many thanks to Alexander Fletcher for input while developing the stochastic simulation code.

References

Anderson R, May R (1982) Population biology of infectious diseases. Springer-Verlag Berlin

Anderson RM, May RM (1979) Population biology of infectious diseases: Part I. Nature 280:361–367

Bahl R, Eikmeier N, Fraser A, Junge M, Keesing F, Nakahata K, Wang LZ (2020) Modeling covid-19 spread in small colleges. arXiv

Brauner JM, Mindermann S, Sharma M, Johnston D, Salvatier J, Gavenčiak T, Stephenson AB, Leech G, Altman G, Mikulik V, Norman AJ, Monrad JT, Besiroglu T, Ge H, Hartwick MA, Teh YW, Chindelevitch L, Gal Y, Kulveit J (2020) Inferring the effectiveness of government interventions against covid-19. Science URL https://science.sciencemag.org/content/early/2020/12/15/science.abd9338

Brook C, Northrup G, Ehrenburg A, Doudna J, Boots M (2020) Optimizing covid-19 control with asymptomatic surveillance testing in a university environment. medrxiv

Cashore JM, Duan N, Janmohamed A, Wan J, Zhang Y, Henderson S, Shmoys D, Frazier P (2020) Covid-19 mathematical modeling for cornells fall semester. Tech. rep., Cornell

Chu D, Aki E, Duda S, Solo K, Yaacoub S, Schuenemann H (2020) 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. The Lancet 395:1973–1987

Dietz K (1993) The estimation of the basic reproduction number for infectious diseases. Statistical Methods in Medical Research 2:23–41

Ferguson N, Laydon D, Nedjati-Gilani, et al (2020) Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid-19 mortality and healthcare demand. Tech. rep., Imperial College London

Firth J, Hellewell J, Klepac P, Kissler S, Jucharski A, Spurgin L (2020) Using a real-world network to model localized covid-19 control strategies. Nature Medicine

Gillespie DT (1977) A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. Journal of Computational Physics 22:403–434

Heesterbeek J (2002) A brief history of r 0 and a recipe for its calculation. Acta Biotheoretica 50:189–204

Juul J, Græsbøll K, Christiansen LE, Lehmann S (2020) Fixed-time descriptive statistics underestimate extremes of epidemic curve ensembles. Nature physics 17:5–8

Kain M, Childs M, Becker A, Mordecai E (2021) Chopping the tail: How preventing superspreading can help to maintain covid-19 control. Epidemics 34:100,430

Keeling MJ (1999) The effects of local spatial structure on epidemiological invasions. Proceedings of the Royal Society of London B: Biological Sciences 266:859–867

Keeling MJ, Grenfell BT (2000) Individual-based perspectives on r0. Journal of Theoretical Biology 203:51–61, URL https://www.sciencedirect.com/science/article/pii/S0022519399910640

Kermack WO, McKendrick AG (1927) Contributions to the mathematical theory of epidemics - 1. Proceedings of the Royal Society of London B: Biological Sciences 115A:700–721

Kucharski A, Russell T, Diamon C, Liu Y, Edmunds J, Funk S, Eggo R (2020) Early dynamics of transmission and control of covid-19: a mathematical modelling study. Lancet: Infectious Diseases 20:553–558

Lloyd-Smith J, Schreiber S, Kopp P, Getz W (2005) Superspreading and the effect of individual variation on disease emergence. Nature 438:355–359

Lopman B, Liu CY, Guillou AL, Handel A, Lash TL, Isakov AP, Jenness SM (2020) A model of covid-19 transmission and control on university

campuses. medRxiv URL https://www.medrxiv.org/content/early/2020/07/24/2020.06.23.20138677

Matsuda H, Ogita N, Sasaki A, Sato K (1992) Statistical mechanics of population: the lattice Lotka-Volterra model. Progress of Theoretical Physics 88(6):1035–1044

Sharkey K (2008) Deterministic epidemiological models at the individual level. Journal of Mathematical Biology 57:311–331

UNESCO (2020) Covid-19 education recovery, URL https://en.unesco.org/covid19/educationresponse, accessed: 11/02/2020

Wang L, Didelot X, Yang J, Wong G, Shi Y, Liu W, Gao G, Bi Y (2020) Inference of person-to-person transmission of covid-19 reveals hidden super-spreading events during the early outbreak phase. Nature Communications 11:5006