PEER REVIEW HISTORY

BMJ Open publishes all reviews undertaken for accepted manuscripts. Reviewers are asked to complete a checklist review form (**http://bmjopen.bmj.com/site/about/resources/checklist.pdf**) and are provided with free text boxes to elaborate on their assessment. These free text comments are reproduced below.

ARTICLE DETAILS

TITLE (PROVISIONAL)	Disparities in Spread and Control of Influenza in Slums of Delhi: Findings From An Agent-Based Modeling Study
AUTHORS	Adiga, Abhijin; Chu, Shuyu; Eubank, Stephen; Kuhlman, Christopher; Lewis, Bryan; Marathe, Achla; Marathe, Madhav; Nordberg, Eric; Swarup, Samarth; Vullikanti, Anil; Wilson, Mandy

VERSION 1 – REVIEW

REVIEWER	James Koopman Univ. of Michigan, USA
REVIEW RETURNED	22-Apr-2017

GENERAL COMMENTS	Specific Comments 1. When model “agents” are persons rather than infectious agents, using “persons”, especially in an initial overview, communicates more clearly to this reviewer. 2. In the authors minds, colocation is a proxy for airborne transmission. But it would be possible to have large effects of colocation with no airborne transmission. Currently there is no good evidence from any population regarding routes of transmission and co-location effects. Co-location might capture direct physical contact, fomite mediated transmission, or airborne. The controversy regarding the importance of airborne transmission in influenza is far from resolved. So why not say that colocation could be capturing diverse routes of transmission? 3. It is not given that influenza is airborne! There is indirect evidence that it can be airborne. But there is no basis to estimate what fraction of transmissions are airborne in any population. That does not mean that a co-location model is not a good way to capture spatial effects on transmission since all possible routes of transmission will be related to colocation. 4. References 15-18 in no way support the statement that “household size and crowding make it easier to transmit airborne infections”. Reference 16 for example refers to meningococcal disease where it is doubtful that there is any airborne transmission. Measles can be airborne, but again there is no evidence or model analyses that gives us any indication regarding what fraction of transmission is airborne. The famous case of doctor’s office transmission after several hours absence is sometimes quoted as supporting airborne transmission. But any model of air exchange indicates that is in fact evidence against airborne transmission. 5. I think that the major inferences about intervention effects in this paper will be robust to changing transmission parameters to lower values generating lower attack rates. But future work should assess this.

	6. The simplifying assumption in this model whose realistic relaxation seems most likely to affect the intervention effect inferences is the assumption that all transmission is related to temporally specific co-location. This assumption is behind the larger total transmission potential in this model of slum dwellers vs. others. Since we do not know what fraction of influenza transmissions are via large droplet or direct fomite contamination, where co-location at different times might lead to more transmission than airborne would, it would be interesting to examine how expanding contact times at locations would affect these results. But again, that is something that should be done for this publication. 7. The model of vaccine failure should be clarified. A 30 efficacy could be that everyone has a 30% less chance of being infected given that they are exposed or it could mean that 30% are completely protected and the others get no protection. Reality is somewhere in between these unrealistic extremes. This should be made clearer not only in the methods section, but in the discussion section as well, since it could make such a big difference. 8. The comment about interactions between factors (inputs) in a model changing results applies to reality as well as models. Since there are very likely many factors in the real world generating important interactions that are not in the model, this makes prediction based on model analyses problematic. But the first step is understanding of effects in a model, and that is what gives this paper its value. It shows how co-location algorithms that use real world data generate important differences in intervention effects. There is no need for this model to be used for specific predictions for it to have value to both Public Health and to science. 9. The legend for Figure 1 should state that results and summations across 25 simulations. The methodology behind those summations should be made clearer in the methods section. 10. The small gain from the VsSn compared to the other combinations of focused interventions doesn't provide an argument in favor of focusing, given all the social problems that could engender. 11. The discussion might include a segment on how health officials in Delhi could use or benefit from this model and, in turn, how this might stimulate data collection that could be used to improve the validity of inferences made using this model. General Comments 1. The revisions made greatly strengthen this paper over the original. 2. The results are likely to influence both modelers and public health officials (and even more so for the rare modeler who is a public health official) 3. This work is ground breaking and innovative. 4. The results may very well be sensitive to realistic relaxation of many simplifying assumptions. But if this paper has

	the impact that it should, many modelers will hop onto the task of further exploring what affects the inferences made. 5. In order for this work to have the impact it can potentially have, the code used should be made readily available, as should the data from the surveys on which the models are built.

REVIEWER	Elizabeth Lee Georgetown University, USA
REVIEW RETURNED	30-Jun-2017

GENERAL COMMENTS	Building off Chen et al. 2016, the goal of this paper was to examine the effectiveness of multiple interventions in preventing the epidemic spread of influenza in synthetic networks representing the population of Delhi with and without explicit slum features. This work explores the importance of incorporating demographic heterogeneity and considers the contact pattern features of slum populations that generate distinctly different epidemic dynamics from those of a more homogeneous urban population. Here are my comments: Methods: - From what I understand, Network 1 essentially averages the demographic features of slum and nonslum populations over the entire network, while Network 2 spatially groups slum populations. Is it true that the data informing Network 1 activities includes surveys from slum populations? If not, the two networks are not particularly comparable; Network 2 would simply have greater connectivity than Network 1, so it isn't surprising that the interventions would be less effective in Network 2 across the board. Perhaps it would be more apt to compare the disease risk of Network 1 individuals and non-slum populations in Network 2, as is alluded to in Figures 4 and 5. It would be more interesting to better understand the intertwined disease risk of slum and non-slum populations, and how trickle-down effects between these populations change intervention effectiveness (as examined in Figure 5). - Disease transmission occurs stochastically among co-located individuals engaging in similar activities, and it seems that a single transmissibility value is used for all activity types. Can the authors please clarify why we should assume that household interactions are equally likely as shopping center or workplace interactions to result in disease transmission? There is ample empirical literature to suggest that household interactions and community-based interactions have different risks of transmission. For instance, one can imagine that spending two hours at a shopping mall where there is an infected individual does not carry the same risk as spending two hours with an infected family member at home. At the very least, this should be mentioned as a limitation to the activity-based modeling approach. Results: - Second paragraph of Results and Analysis: I don't understand what "differences" are being tested by the t-test described here. Is it, for instance, that the distribution of epidemic sizes for the 25 simulation replicates are being compared between Network 1 and Network 2? If that is indeed the case, can the

	authors please explain why this is an appropriate test to use? If these replicates are “sample means” that can be compared with the t-test, what are the populations for which the test provides comparison? Certainly I think it makes sense to report error bars on all of the simulation results, but I’m not convinced that reporting the simulation results as “statistically significantly different” is appropriate (or, in fact, needed). - I am surprised that the stay-home intervention was so effective among Network 2 simulations (Figure 3b) and slum populations (Figure 4) since slum populations had an average of over 15 household contacts. Can the authors comment on this result? Discussion: - This paper incorporates detailed demographic features into the population model, yet the modeled interventions span a wide range of compliance and coverage values, some of which are unrealistic. While it’s useful to consider these ranges from a theoretical standpoint, it may be useful for the authors to highlight and discuss the implications for models where the interventions fall within the realistic range. Is there any literature on compliance for non-pharmaceutical interventions or vaccination coverage for slum and non-slum populations, in Delhi and other parts of India, or in other populations with comparable demography? What policies or historical examples exist that have triggered school closures at 1% and 5% cumulative infection rates? - Intuition from models of epidemic spread on networks tells us that increasing heterogeneity in degree should also increase the epidemic size and peak and reduce the time from epidemic onset to peak. The authors may consider linking this theory with their comparison of Networks 1 and 2. - “Despite this being the first work of its kind – to model...” – This may be overstating the novelty of this work. There are other examples of agent-based models in city populations (e.g., Kumar S, Piper K, Galloway DD, Hadler JL, Grefenstette JJ. Is population structure sufficient to generate area-level inequalities in influenza rates? An examination using agent-based models. BMC Public Health. 2015;15(1):947.), and this paper is a straightforward extension of the models in Chen et al. 2016. Minor comments: - Why is “Influenza” capitalized in different parts of the paper

VERSION 1 – AUTHOR RESPONSE

Reviewer 1

Building off Chen et al. 2016, the goal of this paper was to examine the effectiveness of multiple interventions in preventing the epidemic spread of influenza in synthetic networks representing the population of Delhi with and without explicit slum features. This work explores the importance of incorporating demographic heterogeneity and considers the contact pattern features of slum populations that generate distinctly different epidemic dynamics from those of a more homogeneous urban population.
Here are my comments:

Methods:
- From what I understand, Network 1 essentially averages the demographic features of slum and nonslum populations over the entire network, while Network 2 spatially groups slum populations. Is it true that the data informing Network 1 activities includes surveys from slum populations? If not, the two networks are not particularly comparable; Network 2 would simply have greater connectivity than Network 1, so it isn't surprising that the interventions would be less effective in Network 2 across the board. Perhaps it would be more apt to compare the disease risk of Network 1 individuals and non-slum populations in Network 2, as is alluded to in Figures 4 and 5. It would be more interesting to better understand the intertwined disease risk of slum and non-slum populations, and how trickle-down effects between these populations change intervention effectiveness (as examined in Figure 5).

Response:
The data in Network 1 does not include surveys from the slum populations. Everyone in Network 1 has the non-slum activities whereas in Network 2, people in the slum regions have the slum activities as taken from the slum survey. Given the higher connectivity among the slum population in Network 2, one can intuitively conclude, as you suggest, that interventions would be less effective under Network 2 compared to Network 1 but the magnitude of the difference is harder to assess, especially for different epidemic outcomes such as the difference in peak, time to peak and attack rate. We agree that it would be interesting to compare the disease risk of Network 1 individuals with non-slum populations in Network 2. This discussion has been added to the "Results and Analysis" section just after Figure 4. Also Figure 5 has been updated to include an additional column on Network 1 and more discussion regarding this has been added to the "Constrained resource allocation among slum and non-slum areas" part of the "Results and Analysis" section.

- Disease transmission occurs stochastically among co-located individuals engaging in similar activities, and it seems that a single transmissibility value is used for all activity types. Can the authors please clarify why we should assume that household interactions are equally likely as shopping center or workplace interactions to result in disease transmission? There is ample empirical literature to suggest that household interactions and community-based interactions have different risks of transmission. For instance, one can imagine that spending two hours at a shopping mall where there is an infected individual does not carry the same risk as spending two hours with an infected family member at home. At the very least, this should be mentioned as a limitation to the activity-based modeling approach.

Response:
In our model the disease risk depends on the contact time but not on the activity type as you correctly observed. This is a limitation of our study and has now been mentioned in the "Strengths and limitations" and "Discussion" section of the paper.

Results:
- Second paragraph of Results and Analysis: I don't understand what "differences" are being tested by the t-test described here. Is it, for instance, that the distribution of epidemic sizes for the 25 simulation replicates are being compared between Network 1 and Network 2? If that is indeed the case, can the authors please explain why this is an appropriate test to use? If these replicates are "sample means" that can be compared with the t-test, what are the populations for which the test provides comparison? Certainly I think it makes sense to report error bars on all of the simulation results, but I'm not convinced that reporting the simulation results as "statistically significantly different" is appropriate (or, in fact, needed).

Response:
For each of the 25 replicates, the index cases are selected with a different random seed so the mean of outcome variable addresses the stochastic variation. For example, in the case of the epidemic size, the numbers are the exact values of replicates' epidemic size and not the 'sample means'. The two-sample t-test is used to compare the mean epidemic size in two networks. The p-value of the t-test determines if the epidemic sizes under the two networks are statistically significantly different or not. We did not put error bars since the errors across 25 replicates are quite small. This is also consistent with the small p-values reported.

- I am surprised that the stay-home intervention was so effective among Network 2 simulations (Figure 3b) and slum populations (Figure 4) since slum populations had an average of over 15 household contacts. Can the authors comment on this result?

Response:
The stay-home intervention (SHO) is effective because the model assumes that as soon as a person becomes infectious, that person is confined to home. Thus, a person that becomes infectious can infect their family members, as you say, but if these other members become infectious, then they, too, will be confined to home. Thus, home-bound people can infect their family members, but no one beyond their family (for 100% compliance). As compliance rate increases, this effect approaches, roughly, a "family-based" isolation intervention (similar to ISO), consistent with the results in the two figures that you cite, for ISO and SHO. This has been added to the description of Figure 3.

Discussion:
- This paper incorporates detailed demographic features into the population model, yet the modeled interventions span a wide range of compliance and coverage values, some of which are unrealistic. While it's useful to consider these ranges from a theoretical standpoint, it may be useful for the authors to highlight and discuss the implications for models where the interventions fall within the realistic range. Is there any literature on compliance for non-pharmaceutical interventions or vaccination coverage for slum and non-slum populations, in Delhi and other parts of India, or in other populations with comparable demography? What policies or historical examples exist that have triggered school closures at 1% and 5% cumulative infection rates?

Response:
Several citations are provided in a paragraph under "Interventions" within the "Methods" section.

- Intuition from models of epidemic spread on networks tells us that increasing heterogeneity in degree should also increase the epidemic size and peak and reduce the time from epidemic onset to peak. The authors may consider linking this theory with their comparison of Networks 1 and 2.

Response:
This is a complicated issue, and one that is useful in exploring from a graph theoretic perspective. First, the maximum degree in Network 1 is 170, and that in Network 2 is 180. Thus, there is not a big

change in maximum degree, nor in the range of degrees. Second, the average degrees in the two networks is about 30 to 33 as shown in the revised Table S1 of the supplemental information. Based on increasing the number of in-home edges by 10 (say from 4 to 14 family members), this will not make a huge change to the average degree nor the maximum degree of nodes in a network, as observed above. We have added these comments to the Methods section.
Some more general comments. The particular graph measures to employ are not straight-forward, nor is it straight-forward to correlate these changes with dynamics. There are effects of epidemic thresholds, as used by Faloutsos at CMU, among others, which accounts for graph structure and the virus transmission model. This measure is related to the epidemic size. The speed of contagion (e.g., for time-to-peak) is related to the graph diameter and degree distribution. Both epidemic size and speed are also affected by duration of interactions (which is not a graph measure). For example, one path between two nodes may be dominant, even if only this single path exists, if the contact durations among pairs of nodes along this path are sufficiently large. Alternatively, multiple disjoint paths between two nodes may result in tenuous transmission between them if the durations of interaction are sufficiently small.

- “Despite this being the first work of its kind – to model...” – This may be overstating the novelty of this work. There are other examples of agent-based models in city populations (e.g., Kumar S, Piper K, Galloway DD, Hadler JL, Grefenstette JJ. Is population structure sufficient to generate area-level inequalities in influenza rates? An examination using agent-based models. BMC Public Health. 2015;15(1):947.), and this paper is a straightforward extension of the models in Chen et al. 2016.

Response:
The novelty statement has been removed.

Minor comments:
- Why is “Influenza” capitalized in different parts of the paper?
Response:
All instances have been changed to “influenza” in the manuscript.

Reviewer 2

BMJ Open Article Resubmission Review
Specific Comments
1. When model “agents” are persons rather than infectious agents, using “persons”, especially in an initial overview, communicates more clearly to this reviewer.

Response:
This has been changed per the reviewer’s comment.

2. In the authors minds, colocation is a proxy for airborne transmission. But it would be possible to have large effects of colocation with no airborne transmission. Currently there is no good evidence from any population regarding routes of transmission and colocation effects. Co-location might capture direct physical contact, fomite mediated transmission, or airborne. The controversy regarding the importance of airborne transmission in influenza is far from resolved. So why not say that colocation could be capturing diverse routes of transmission?

Response:
This has been changed per the reviewer’s comment. The statement is under Methods, under disease model.

3. It is not given that influenza is airborne! There is indirect evidence that it can be airborne. But there is no basis to estimate what fraction of transmissions are airborne in any population. That does not mean that a co-location model is not a good way to capture spatial effects on transmission since all possible routes of transmission will be related to colocation.

Response:
We have modified the text throughout the paper and used the phrase “environmentally-mediated” instead of “airborne”, to allow for “airborne” as one of the possibilities.

4. References 15-18 in no way support the statement that “household size and crowding make it easier to transmit airborne infections”. Reference 16 for example refers to meningococcal disease where it is doubtful that there is any airborne transmission. Measles can be airborne, but again there is no evidence or model analyses that gives us any indication regarding what fraction of transmission is airborne. The famous case of doctor’s office transmission after several hours absence is sometimes quoted as supporting airborne transmission. But any model of air exchange indicates that is in fact evidence against airborne transmission.

Response:
As stated in the previous comment, text has been modified and “airborne” has been replaced with “environmentally-mediated”.

5. I think that the major inferences about intervention effects in this paper will be robust to changing transmission parameters to lower values generating lower attack rates. But future work should assess this.

Response:
This has been added to the future work in the “Discussion” section.

6. The simplifying assumption in this model whose realistic relaxation seems most likely to affect the intervention effect inferences is the assumption that all transmission is related to temporally specific co-location. This assumption is behind the larger total transmission potential in this model of slum dwellers vs. others. Since we do not know what fraction of influenza transmissions are via large droplet or direct fomite contamination, where co-location at different times might lead to more transmission than airborne would, it would be interesting to examine how expanding contact times at locations would affect these results. But again, that is something that should be done for this publication.

Response:
We will address this in the future work and have added it to the future work paragraph in the “Discussion” section.

7. The model of vaccine failure should be clarified. A 30 efficacy could be that everyone has a 30% less chance of being infected given that they are exposed or it could mean that 30% are completely protected and the others get no protection. Reality is somewhere in between these unrealistic extremes. This should be made clearer not only in the methods section, but in the discussion section as well, since it could make such a big difference.

Response:

In our simulation, when a susceptible node is vaccinated, its probability of getting infected by an infectious node is scaled down by the efficacy. If it becomes infectious, its probability of infecting

susceptible nodes is also scaled down by the efficacy. In other words, both incoming and outgoing infection probabilities are reduced by the efficacy. We have clarified this in the manuscript in the "Methods" section under "Interventions".

8. The comment about interactions between factors (inputs) in a model changing results applies to reality as well as models. Since there are very likely many factors in the real world generating important interactions that are not in the model, this makes prediction based on model analyses problematic. But the first step is understanding of effects in a model, and that is what gives this paper its value. It shows how co-location algorithms that use real world data generate important differences in intervention effects. There is no need for this model to be used for specific predictions for it to have value to both Public Health and to science.

Response:
We added a comment at the end of our "limitations" paragraph within the Discussion section.

9. The legend for Figure 1 should state that results and summations across 25 simulations. The methodology behind those summations should be made clearer in the methods section.

Response:
The results in the figure are the averages over 25 replicates at each time point. Thus, they are time-point wise averages. For example, the mean infection rate at day 100 is calculated by taking the average of 25 infection rates that occurred on day 100 of each replicate. It has been clarified in the manuscript in the "Methods" section and in the caption of Figure 1.

10. The small gain from the VsSn compared to the other combinations of focused interventions doesn't provide an argument in favor of focusing, given all the social problems that could engender.

Response:
Yes, we agree that overall, that there are other confounding factors. We also agree that the differences in total infection rates across the four intervention strategies is well within the randomness inherent in our methodology. Our main points are that these strategies produce the largest change in infection rates among a well-defined group that may be used to help slum populations; and also that there is a non-intuitive outcome in these results. See changes made at the end of the Results section in the manuscript.

11. The discussion might include a segment on how health officials in Delhi could use or benefit from this model and, in turn, how this might stimulate data collection that could be used to improve the validity of inferences made using this model.

Response:
A discussion has been added to end of the "Discussion" section under "Public health implications".

VERSION 2 – REVIEW

REVIEWER	James S. Koopman Retired USA
REVIEW RETURNED	21-Aug-2017

GENERAL COMMENTS	The first reviewer made a good point that it would have been better to have kept the number of daily activities and the number of edges more comparable between Network 1 and Network 2. Then it might have been more possible to infer the mechanisms by

	which slum populations exert their effects. A 10% increase in the number of edges and average degree, or an 18% increase in daily activities, could have a huge effect on population level transmission. Ideally in this paper one should be able to distinguish general increases in contact from heterogeneity effects. Would it take that much more work to rerun Network 1 simulations with total daily activities and number of edges corresponding to those in Network 2? But even given this limitation, I still find the analyses presented to be eye opening. Seeing such a big effect should at least raise interest in resolving remaining issues and assessing where the key new knowledge that will benefit mankind lies. None-the-less, now that I am noting the size of total contact differences, I think that the value added by including new runs of network 1 with contact totals the are more comparable to network 2 would add greatly to the value of this paper. I do not find the "percentage increase" argument convincing. Those percentage differences could have arisen just from changing the total contacts in Network 1. Since the Chen et al. paper has already established the potential effects of slums, one would hope that this paper would make the distinction as to whether that effect is due just to the effects on total number of contacts or on a slum focus of intense contacts. I also agree that p-values for mean differences are of no value for this kind of study. Variation between runs rather variation in means is the only entity of interest. I am unconvinced that an individual based network model as a sole method for examining the issues raised about slums and focusing interventions on slums or non-slums is the best approach. A series of different type models that make different assumptions would be more informative. The very small number of runs examined in this paper is certainly an argument against using this method as the only method. I am afraid that on rereading the paper with its minor changes, I was convinced by comments from the first reviewer that my prior review was inadequate and that a comparison of network2 with a modified Network1 that has similar total contacts as network2 is needed. Otherwise the results are too hard to interpret.

REVIEWER	Elizabeth Lee Georgetown University, USA
REVIEW RETURNED	18-Aug-2017

GENERAL COMMENTS	The authors have addressed all of my comments.

VERSION 2 – AUTHOR RESPONSE

Reviewer 1

-Reviewer comment
The first reviewer made a good point that it would have been better to have kept the number of daily activities and the number of edges more comparable between Network 1 and Network 2. Then it might have been more possible to infer the mechanisms by which slum populations exert their effects. A 10% increase in the number of edges and average degree, or an 18% increase in daily activities,

could have a huge effect on population level transmission. Ideally in this paper one should be able to distinguish general increases in contact from heterogeneity eff ects. Would it take that much more work to rerun Network 1 simulations with total daily activities and number of edges corresponding to those in Network 2? But even given this limitation, I still find the analyses presented to be eye opening. Seeing such a big effect should at least raise interest in resolving remaining issues and assessing where the key new knowledge that will benefit mankind lies. None-the-less, now that I am noting the size of total contact differences, I think that the value added by including new runs of network 1 with contact totals the are more comparable to network 2 would add greatly to the value of this paper. I do not find the “percentage increase” argument convincing. Those percentage differences could have arisen just from changing the total contacts in Network 1.

-Response
The networks characteristics arise due to the network generation process. That is, we are not specifying as input, or controlling, the numbers of edges of different types. The social networks are outputs of the synthetic population generation process. Inputs to this process are data such as age and gender distributions, household sizes, home and business locations, and activities. Activities are based on survey data. Edges in the network are formed when people are co-located in the same room at the same time. Our goal here is to use the most realistic populations and derived social networks that the synthetic population construction process enables, to produce highly resolved networks. Our goal is not to develop two networks that have the same number of edges. Our goal is to refine the subnetworks of Network 1, which represent slum regions, in order to more faithfully represent the attributes of the slum dwellers. This refinement results in the construction of Network 2. The survey data from slums is used to determine the activities of the slum dwellers, which results in the network edges (interactions) in the slum regions in Network 2. The extra edges in Network 2 have not been randomly added. They arise due to the activities performed by the slum residents. We hope this clarifies the fact that edges are not the input to the model, but the output of the model.

-Reviewer comment
Since the Chen et al. paper has already established the potential effects of slums, one would hope that this paper would make the distinction as to whether that effect is due just to the effects on total number of contacts or on a slum focus of intense contacts.

-Response
The effect of slum arises from both i.e. a higher number of total contacts as well as more contacts within slums. Work in Chen et al. establishes the effect of slums on epidemics but not the effect of slums on interventions. The focus of this paper is on understanding the effect of slums on interventions.

-Reviewer comment
I also agree that p-values for mean differences are of no value for this kind of study. Variation between runs rather variation in means is the only entity of interest.

-Response
The issue of p-values came up because it was not clear to the other Reviewer what ``differences'' were being tested by the t-test. We are not testing the sample-means. From each replicate we get an exact value of the epidemic size for each network; we do 25 replicates and then calculate a mean value across 25 replicates for each network.
The t-test is used to compare the mean epidemic size in

Network 1 and Network 2. The p-value of the t-test determines if the epidemic sizes under the two networks are statistically significantly different or not.

-Reviewer comment
I am unconvinced that an individual based network model as a sole method for examining the issues raised about slums and focusing interventions on slums or non-slums is the best approach. A series of different type models that make different assumptions would be more informative. The very small number of runs examined in this paper is certainly an argument against using this method as the only method.

-Response
We agree that an individual based network model may not be the sole method for studying these problems. We stated that individual models have attractive features such as a high level of granularity.
We agree that other types of models can be profitably employed; we have added a sentence stating this in the Discussion.

-Reviewer comment
I am afraid that on rereading the paper with its minor changes, I was convinced by comments from the first reviewer that my prior review was inadequate and that a comparison of network2 with a modified Network1 that has similar total contacts as network2 is needed. Otherwise the results are too hard to interpret.

-Response
The response to this comment is the same as the response to the first comment made earlier by this reviewer. We refer to the response there.

Reviewer 2

- Reviewer Comment
The authors have previously addressed all of this reviewer's concerns.

-Response
Thank you.

VERSION 3 – REVIEW

REVIEWER	James S. Koopman University of Michigan, USA
REVIEW RETURNED	12-Oct-2017

GENERAL COMMENTS	Achla Marathe and her Network Dynamics and Simulation Sciences Laboratory Biocomplexity Institute team have performed a detailed and informative analysis of a discrete individual stochastic model with detailed characterization of transmission in slums. Their work illustrates how it might be possible for slums to play a crucial role in transmission dynamics that might justify targeting control efforts to them. They claim that (1) “Ignoring the unique attributes of slums in a population overestimates the benefits of the interventions”. (2) “Interventions are more effective in Network 1 than Network 2 for all types of interventions: vaccination, closing schools, staying home, and isolation.” The

	write up implies that Network 1 and Network 2 are comparable except for the latter accounting for slums more specifically. But that is not the case. The total number of contacts is greater for Network 2. Making these unfounded claims detracts from the great value in their work of detailing how slum structure affects transmission dynamics and decisions about what type of interventions will work best in slum and non-slum areas given that immunity is the same in slum and non-slum areas. Of course, this assumption of equal immunity would not hold for seasonal epidemics. In seasonal epidemics, the slums will be protected by their greater immunity. The number of cases where true risk factors appear to be protective factors in infectious diseases is great because this immunity effect is ignored. That deserves a little more emphasis in the discussion. The focus of this research group is on inferring the number and structure of contacts from data about human situations such as family structure, residence, and movement to other social structures. This extremely valuable research focus, however, does not provide a practical basis for population intervention decision processes, as this paper implies it does. To guide interventions, evidence is needed to support the inference that a control action will have greater benefits than some other control action. That evidence must ultimately come from fitting models to data. Projecting behaviors of arbitrarily defined systems that are not shown to fit observed outcomes, as these authors do, is valuable. It enhances understanding of what determines transmission dynamics. But it cannot assess what determinants are actually acting to affect an infection control decision. For that, what is needed is fitting of models to observed outcomes in a manner that specifies the identifiability of parameters or parameter sets or parameter spaces and shows that a particular model fitted to a particular set of data is consistent with one action or another being better. But because one paradigm is being used in a paper that a reviewer like me does not accept is not a reason to withhold publication. It would be nice, however, if the authors would detail the logic leading to claim 1 above. Then if they could go on to detail how the findings in their paper should be used in combination with what data to make a decision, there would be a basis for a better dialogue.

VERSION 3 – AUTHOR RESPONSE

Comment: Achla Marathe and her Network Dynamics and Simulation Sciences Laboratory Biocomplexity Institute team have performed a detailed and informative analysis of a discrete individual stochastic model with detailed characterization of transmission in slums. Their work illustrates how it might be possible for slums to play a crucial role in transmission dynamics that might justify targeting control efforts to them. They claim that (1) “Ignoring the unique attributes of slums in a population overestimates the benefits of the interventions”. (2) “Interventions are more effective in Network 1 than Network 2 for all types of interventions: vaccination, closing schools, staying home, and isolation.” The write up implies that Network 1 and Network 2 are comparable except for the latter

accounting for slums more specifically. But that is not the case. The total number of contacts is greater for Network 2.

Response: We agree that Network 2 has more contacts. In fact Table S1 shows the exact number of nodes and edges in both networks. Both networks have the same number of nodes i.e. 13.8 million but Network 2 has 20 million more edges than Network 1. The extra edges in Network 2 arise due to larger family size in slums and a larger number of activities performed by them. See Table S1. Note that we are not specifying the extra number of edges in Network 2 as an input---the network edges are a consequence of the population generation process which uses survey data. The two networks are comparable because they have the same number of nodes. The difference in edges in network 1 vs. network 2 arises because network 2 accounts for slum attributes and network 1 does not.

Comment: Making these unfounded claims detracts from the great value in their work of detailing how slum structure affects transmission dynamics and decisions about what type of interventions will work best in slum and non-slum areas given that immunity is the same in slum and non-slum areas. Of course, this assumption of equal immunity would not hold for seasonal epidemics. In seasonal epidemics, the slums will be protected by their greater immunity. The number of cases where true risk factors appear to be protective factors in infectious diseases is great because this immunity effect is ignored. That deserves a little more emphasis in the discussion.

Response: The discussion has been updated to include the following.

Despite the detailed modeling effort, there are limitations of this work and areas for improvement in the future. For example, this model assumes that both slum and non-slum individuals have the same level of immunity. This may not be true for seasonal infections. Previous researchers have argued that individuals who live in smaller family sizes, who have access to household amenities and maintain a high level of personal cleanliness, face declining microbial exposure which can modify their immune response and reduce their level of tolerance to respiratory infections [Bjorksten 2004]. Slum households characterized by larger family size and overcrowding, are likely to encounter much higher microbial exposure and therefore may be protected by their greater immunity [Baker et al. 2000, Burstrom et al. 1999].

Björkstén, B., 2004, February. Effects of intestinal microflora and the environment on the development of asthma and allergy. In Springer seminars in immunopathology (Vol. 25, No. 3-4, pp. 257-270). Springer-Verlag.

Baker M, McNicholas A, Garrett N, Jones N, Stewart J, Koberstein V, Lennon D. Household crowding a major risk factor for epidemic meningococcal disease in Auckland children. The Pediatric Infectious Disease Journal 2000 Oct 1;19(10):983-90.

Burström B, Diderichsen F, Smedman L. Child mortality in Stockholm during 1885–1910: the impact of household size and number of children in the family on the risk of death from measles. American Journal of Epidemiology 1999 Jun 15;149(12):1134-41.

Comment: The focus of this research group is on inferring the number and structure of contacts from data about human situations such as family structure, residence, and movement to other social structures. This extremely valuable research focus, however, does not provide a practical basis for population intervention decision processes, as this paper implies it does. To guide interventions, evidence is needed to support the inference that a control action will have greater benefits than some other control action. That evidence must ultimately come from fitting models to data. Projecting behaviors of arbitrarily defined systems that are not shown to fit observed outcomes, as these authors

do, is valuable. It enhances understanding of what determines transmission dynamics. But it cannot assess what determinants are actually acting to affect an infection control decision.

For that, what is needed is fitting of models to observed outcomes in a manner that specifies the identifiability of parameters or parameter sets or parameter spaces and shows that a particular model fitted to a particular set of data is consistent with one action or another being better.

But because one paradigm is being used in a paper that a reviewer like me does not accept is not a reason to withhold publication. It would be nice, however, if the authors would detail the logic leading to claim 1 above. Then if they could go on to detail how the findings in their paper should be used in combination with what data to make a decision, there would be a basis for a better dialogue.

Response: We are not finding determinants that are actually acting to affect an infection control decision. The contribution of this work is in taking a hypothetical epidemic scenario and then considering the impact of interventions on two types of networks; one of which provides a more accurate representation of the ground truth. The goal is to show that the missing edges in network 1 overestimate the impact of interventions and underestimate the epidemic outcomes.