PEER REVIEW HISTORY

This paper was submitted to a another journal from BMJ but declined for publication following peer review. The authors addressed the reviewers' comments and submitted the revised paper to BMJ Open. The paper was subsequently accepted for publication at BMJ Open.

(This paper received three reviews from its previous journal but only two reviewers agreed to published their review.)

ARTICLE DETAILS

TITLE (PROVISIONAL)	Impact of the Southwark and Lambeth Integrated Care Older People's Programme on hospital utilisation and costs: controlled time series and cost consequence analysis.
AUTHORS	Exley, Josephine; Abel, Gary A; Fernandez, Jose-Luis; Pitchforth, Emma; Mendonca, Silvia; Yang, Miaoqing; Roland, Martin; McGuire, Alistair

VERSION 1 – REVIEW

REVIEWER	Meredith Rosenthal Harvard TH Chan School of Public Health, USA
REVIEW RETURNED	12-Jun-2018

GENERAL COMMENTS	It is a widely held notion in health policy that poor coordination and lack of prospective management of patient care leads to preventable hospital admissions in the elderly, driving up the cost of health care. Health systems have experimented with a variety of strategies to address the needs of high risk patients, including but not limited to older patients: the chronic care model, some versions of patient centered medical homes, high risk care management in primary care, stand-alone "ambulatory ICUs", etc. These models typically share common elements – risk assessment, engagement of a multi-professional team, a care plan, non-visit based care, and robust information systems. Despite the common sense appeal of such initiatives, the case for their effectiveness is mixed and few of these programs have demonstrated cost savings. Exley et al. add to this body of evidence a rigorous evaluation of the "Older People's Programme". The authors could do a better job in the introduction situating the current paper in the context of related studies including those evaluating Wagner's chronic care model, the various Medicare pilot/demonstration programs including the Comprehensive Primary Care Initiative and the CMS high risk care management demonstration that produced the results you cite in FN 15. You may wish to compare and contrast these different approaches, but they are clearly related and your findings have some similarities with what we know from these other efforts. First and foremost, targeting matters. The programs that have demonstrated cost savings almost always are focused on the most complex patients – those with the highest baseline risk of hospitalization. Similarly your finding that HA and ICM increased outpatient use and some

	admissions reminds me of a recent evaluation of a readmission reduction program in the U.S.: A patient navigator program reduced readmissions among frail elderly but increased them in a younger group of high-utilizers who had behavioral health and social complexity. [Balaban, R.B., Galbraith, A.A., Burns, M.E. et al. J GEN INTERN MED (2015) 30: 907. https://doi.org/10.1007/s11606-015-3185-x] Building on these insights from related literature, I can see a couple of implications for this study. First: it would be ideal to stratify results by patient health status, and examine the possibility of a heterogenous response. Is that possible in your data? Second, your “dose” measure for HA and ICM also might capture the practice’s perception of need and/or the extent of targeting by the practice in terms of focus on higher risk patients. Which is not to say that analysis lacks merit, but the interpretation might be a little more complicated than just a dose-response. Particularly given you findings you might think a bit more about how practices (and their patients) with higher percentages of HA and ICM differ from practices with lower percentages. You conclusion seems reasonable – it may be time to stop trying to get blood from a stone – though I continue to wonder whether a better-targeted effort (I can’t tell how well targeted this one was but you would know) would generate different results. You might also note that we should try to find out whether the increased utilization observed for HA and ICM is really an indicator of unmet need (your guess) or low-value care that resulted from increased contact with the health care system. Since you executed your planned evaluation you should report everything of course, but I was puzzled why you looked at length of stay at all – can you say why you thought it could be impacted by the program? Through the discharge planning? A couple of tiny methodological quibbles: 1. I was surprised you used Poisson rather than negative-binomial models given the likely over-dispersion of outcomes. 2. Despite having conducted propensity-score matching in exactly the way you do here many times, I have become convinced that other methods, including coarsened exact matching or propensity score weighting are preferable: https://gking.harvard.edu/publications/why-propensity-scores-should-not-be-used-formatching

REVIEWER	Frances Bunn University of Hertfordshire, UK
REVIEW RETURNED	22-Jun-2018

GENERAL COMMENTS	This is an interesting paper on an important topic. It is well written and I think the conclusion (that we should focus on improving care rather than cost saving) is important. I have a few minor comments. 1. I would have liked a bit more detail about how the holistic assessment was done. For example, was the GP responsible for all the assessment? I doubt that most GPs would be able to provide adequate information on issues such as benefits and housing. 2. Who acted as case manager and what did the case

	management entail? 3. The authors report that uptake of holistic assessments were lower than expected. Are they able to give any reasons for why this might have been? 4. The conclusion of the paper was that interventions should focus on improving care rather than reducing costs. I would have liked to see some care outcomes reported in this paper - we are unable to assess whether the intervention had any impact on outcomes for patients and their family carers. The lack of such data might be flagged up as a limitation.

REVIEWER	Dr Jonathan Stokes Manchester Centre for Health Economics, University of Manchester, United Kingdom
REVIEW RETURNED	29-Jun-2018

GENERAL COMMENTS	Thank you for the opportunity to review a very timely and interesting paper looking at effects of an integrated care intervention on secondary care utilisation and costs. Generally, the analysis appears to be very well-conducted and presentation is excellent. I just have a few minor suggestions for improvements (in detail below). In summary, some clarification/additional discussion would be nice in some areas, and some additional details on the modelling to reassure approach/assumptions. • Abstract (p.2, lines 27-30). This section could be more clearly written, not immediately obvious which costs relate to which outcome o Suggest something like: "…Management were associated with increases in elective activity and costs; £126 increase in outpatient attendance and £936 in elective admission costs per Holistic Assessment carried out, and £576 increase in outpatient and £5,858 in elective admission costs per patient receiving Integrated Care Management" • Introduction (p.3, summary box). Maybe make clearer the flow of what elements directly follow from other here, and what applies to which specific patientgroups, e.g. as far as I understand, the HA, ICM and multi-disciplinary team meetings all seem to follow on from each other and are only for a subset of high-risk patients that need to progress through each stage. Not immediately clear whether the other services apply to all older people or not. This information would help reader make judgement call on potential spillover effects. (Perhaps also expand on "at risk" here – make more clear that risk defined by clinician judgement following the holistic assessment.) • Results. Would be good to see table comparing baseline matching characteristics between groups before and after matching showing indication of reduced bias/how well matched the practices were. • Discussion o Discuss limitation of analysing costs without looking at effectiveness measure. Potential for health/patient experience gains from addressing unmet need or other spillover effects of closer working across sectors? o Discussion of change in effects over years? o More discussion of spillover effect findings (potentially small saving/cost neutral, excluding start-up costs, plus decreased A&E,

	electives and outpatient) versus effects on those directly targeted would be good. Does this suggest potential for prevention to impact more positively on programme aims than focus on immediately high-risk patients alone? o Looks like you can be a bit more specific on your theory of increased cost due to unmet need. This appears to be quite clearly shown through route of utilisation, i.e. planned, elective admissions and outpatient attendance. o (p.9, lines 38-41) “While it is standard to allow for confounders by using external controls, many areas of England were adopting some sort of initiative to better co-ordinate care.” ♣ This is also dealt with effectively with your dose analysis where you restrict to practices within the local area where you know the precise integration activity – add as strength • References o References 6 and 16 are currently the same. o Another paper that might be worth comparing to is this one https://bmjopen.bmj.com/content/6/4/e010468.short where we likewise looked at both direct intervention effects and spillovers effects of a case management integrated care intervention. Similar indications to your results. o Might also be worth referencing this recent large systematic review of integrated care by Baxter et al. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5946491/ Might give you a reference to speculate on potential effectiveness on wider outcomes than utilisation of secondary care and costs? • Appendix o (Appendix p.4, lines 28-31) “In this model a categorical variable is included taking a value of 0 for all observations prior to the start of SLIC and for control practices at all time points. For SLIC practices it takes the value of the number of years since the start of the intervention, i.e. 1 in 2012/13, 2 in 2013/14 etc”. The model appears to show some similarities to a difference-in-differences with gradual introduction (e.g. Propper et al. 2002 https://www.ncbi.nlm.nih.gov/pubmed/11939240), where your categorical variable acts as a more conventional ‘intervention*POST’ term (adding the additional dose of treatment in terms of years since start). Have you tested the intervention/controls for parallel trends in the pre-period? Can you show that this assumption is satisfied (I’d guess it will be if well-matched)/correct for it if not? o Assignment of integrated care models, I would argue, is largely decided at the CCG level (which were mandated to implement some form of integration with the 2012 Health & Social Care Act). Have you clustered standard errors by CCG to account for this design? o (Appendix p.6-11) ‘Further results of the statistical analysis’. The yearly breakdown of results of the overall intervention effects is interesting, maybe worth mentioning a bit more detail in the main text? Arguably, it shows that utilisation is generally decreasing (with the exception of ACSCs) in the latter years, so potentially longer-term effects might show better cost-effectiveness (are you able to break down your net cost estimates in Table A14 by year too to test this)?

VERSION 1 – AUTHOR RESPONSE

Reviewer's comment	Authors' response
Reviewer 1	
The authors could do a better job in the introduction situating the current paper in the context of related studies including those evaluating Wagner's chronic care model, the various Medicare pilot/demonstration programs including the Comprehensive Primary Care Initiative and the CMS high risk care management demonstration that produced the results you cite in FN 15. You may wish to compare and contrast these different approaches, but they are clearly related and your findings have some similarities with what we know from these other efforts. First and foremost, targeting matters. The programs that have demonstrated cost savings almost always are focused on the most complex patients – those with the highest baseline risk of hospitalization. Similarly your finding that HA and ICM increased outpatient use and some admissions reminds me of a recent evaluation of a readmission reduction program in the U.S.: A patient navigator program reduced readmissions among frail elderly but increased them in a younger group of high-utilizers who had behavioral health and social complexity. [Balaban, R.B., Galbraith, A.A., Burns, M.E. et al. J GEN INTERN MED (2015) 30: 907. https://doi.org/10.1007/s11606-015-3185-x]	We have expanded the introduction with an additional paragraph. In this we cite the Wagner model and have included the recent systematic review by Baxter suggested by referee 3 and two other systematic reviews of case management as our main additional sources. In this new paragraph we also discuss the high/medium risk approach, citing literature which shows that while the potential for reduced utilisation is greatest in the highest risk groups (even after accounting for regression to the mean), a much wider group (I.e. moderate risk) needs to be targeted if the aim is to reduce overall healthcare costs as, in absolute numbers, most unscheduled admissions are from moderate rather than high risk patients. The new paragraph (paragraph 2 of the introduction) reads as follows: *The conceptual basis behind case management interventions lies in the Chronic Care Model (Wagner et al 1996) which includes using clinical information systems to plan patient care and redesigning the delivery of care to meet the needs of patients with chronic illness. Payers have focused on these elements of the model to identify patients with high healthcare costs, hoping that better targeted and coordinated care will improve care and reduce costs, though often with a focus on costs as the primary outcome. Uncontrolled studies of healthcare utilisation in this group often show reduction in utilisation which may simply result from regression to the mean (Roland and Abel 2012) and systematic reviews of rigorous evaluations of case management and interventions to integrate or coordinate care have on the whole shown much smaller effects (Stokes et al 2015, Baker et al 2018, Baxter et al 2018). One problem is whether case management interventions should target this highest risk group who are likely to show the greatest impact on individuals but unlikely to show much impact on healthcare costs overall (Roland and Abel 2012). The South London Integrated Care (SLIC) reported in this paper took a wide population approach, intending originally to carry out health assessments on 50% of all residents 65 and over combined with a range of primary and secondary care interventions and targeted case management for those identified as at risk.* • *Wagner E, Austin B, von Korff M. Organising care for patients with chronic illness. Milbank Quarterly 1996; 74: 511-544* • *Stokes J, Panagioti M, Alam R, Checkland K, CheragoSohi S, Bower P. Effectiveness of case management for 'at risk' patients in primary care: a systematic review and meta-analysis. PLoS One 2015; 10: e0132340* • *Baxter S, Johnson M, Chambers D, Sutton A, Goyder E, Booth A. The effect of integrated care: a systematic review of UK and international evidence. BMC Health Services Research 2018; 18: 350*

	• *Baker J, Grant R, Gopalan A. A systematic review of care management interventions targeting multimorbidity and high care utilization. BMC Health Services Research 2018; 18: 65* *Roland M, Abel G. Reducing emergency admissions: are we on the right track? BMJ 2012; 345: e6017*
Building on these insights from related literature, I can see a couple of implications for this study. First: it would be ideal to stratify results by patient health status, and examine the possibility of a heterogenous response. Is that possible in your data?	We do not have data on health status of individual patients and so are not in a position to do this. In theory we could get diagnoses from previous hospital admissions from the dataset but we do not currently have permission to access these data at individual patient level for data governance reasons.
Second, your "dose" measure for HA and ICM also might capture the practice's perception of need and/or the extent of targeting by the practice in terms of focus on higher risk patients. Which is not to say that analysis lacks merit, but the interpretation might be a little more complicated than just a doseresponse. Particularly given you findings you might think a bit more about how practices (and their patients) with higher percentages of HA and ICM differ from practices with lower percentages.	Although it is possible that higher numbers of patients in a practice receiving HAs and ICM might have been driven by patient need, the fact that our analysis considers within practice changes would imply that if this were the case there were accelerating numbers of admissions in practices with the greatest need and that these were not stemmed by the intervention. We have added this as a clarification in the discussion.
You conclusion seems reasonable – it may be time to stop trying to get blood from a stone – though I continue to wonder whether a better-targeted effort (I can't tell how well targeted this one was but you would know) would generate different results. You might also note that we should try to find out whether the increased utilization observed for HA and ICM is really an indicator of unmet need (your guess) or low-value care that resulted from increased contact with the health care system	We have now added this as a recommendation for further research in the discussion section as follows: ***"***By accounting for within practice changes, our analysis implies that if the number of HAs and ICM conducted within a practice was being driven by patient need then the number of planned admissions was accelerating in practices with the greatest need. *If this is true it is an important message; programmes aimed at integration may not always be cost reducing. In particular, it would be of value to understand whether the increased elective care was to receive interventions that are likely to have a major effect on quality of life or whether these were likely to be for interventions of lower value."*

Since you executed your planned evaluation you should report everything of course, but I was puzzled why you looked at length of stay at all – can you say why you thought it could be impacted by the program? Through the discharge planning?	Simplified discharging was intended to improve the discharge process for patients returning home or to a care home and anticipated to reduce length of stay. We have now expanded the description of the key elements of the intervention in the Box.
I was surprised you used Poisson rather than negative-binomial models given the likely overdispersion of outcomes	We agree with the reviewer that we expect the data to be over dispersed. However, rather than dealing with this through a negative binomial model we have chosen to address this by stipulating random effects for practices. In this way we address the source of the over dispersion directly.
Despite having conducted propensity-score matching in exactly the way you do here many times, I have become convinced that other methods, including coarsened exact matching or propensity score weighting are preferable: https://gking.harvard.edu/publications/why-propensity-scores-shouldnot-be-used-formatching	We recognise that there are issues with propensity score matching, however, this is not what we have done. We used the so-called genetic matching algorithm which is a computer intensive method which aims to achieve maximum balance across the range of matched practices for the variables included in the matching. However, as we point out, due to the issues with closing practices the effectiveness of the matching in the analysis was far from perfect and so the exact method used is somewhat immaterial. For this reason we chose to adjust for variables used in the matching (excluding outcome variables). We have now specified that we used the genetic matching algorithm and provide a reference in the main body of the text, as this information was previously only presented in the appendix.
Reviewer 2	
I would have liked a bit more detail about how the holistic assessment was done. For example, was the GP responsible for all the assessment? I doubt that most GPs would be able to provide adequate information on issues such as benefits and housing Who acted as case manager and what did the case management entail?	We have added additional detail to the introduction to clarify what an HA is and who was involved in undertaking ICM.
The authors report that uptake of holistic assessments were lower than expected. Are they able to give any reasons for why this might have been?	The challenges associated with implementation were examined in a process evaluation conducted by the King's Fund and reported separately. We have highlighted that there were numerous reasons that the programme failed to achieve its planned activity targets and referenced the King's Fund report.

The conclusion of the paper was that interventions should focus on improving care rather than reducing costs. I would have liked to see some care outcomes reported in this paper - we are unable to assess whether the intervention had any impact on outcomes for patients and their family carers. The lack of such data might be flagged up as a limitation	We have now highlighted in the limitation section that we were not able to examine these outcomes. The section reads as follows: *"In line with the specified aims of the Programme, the analysis focused on measures of secondary care use. However, the intervention might have had an effect that went beyond these outcomes such as improved patient outcomes or experience of care. Future studies should seek to look at whether interventions that aimed to improve coordination of care resulted in gains beyond direct costs of hospital use and whether this differed for patients with different health status."*
Reviewer 3	
Abstract (p.2, lines 27-30). This section could be more clearly written, not immediately obvious which costs relate to which outcome o Suggest something like: "…Management were associated with increases in elective activity and costs; £126 increase in outpatient attendance and £936 in elective admission costs per Holistic Assessment carried out, and £576 increase in outpatient and £5,858 in elective admission costs per patient receiving Integrated Care Management"	Thank you for the suggestion we have corrected the text.
Introduction (p.3, summary box). Maybe make clearer the flow of what elements directly follow from other here, and what applies to which specific patient groups, e.g. as far as I understand, the HA, ICM and multi-disciplinary team meetings all seem to follow on from each other and are only for a subset of high-risk patients that need to progress through each stage. Not immediately clear whether the other services apply to all older people or not. This information would help reader make judgement call on potential spillover effects. (Perhaps also expand on "at risk" here – make more clear that risk defined by clinician judgement following the holistic assessment.)	To improve the clarity we have introduced a new paragraph in the introduction, which highlights that the SLIC Programme targets older people. We have also rewritten the summary box and included more detail on how the HA, ICM and CMDT interventions are connected in the main body of the text.

Results. Would be good to see table comparing baseline matching characteristics between groups before and after matching showing indication of reduced bias/how well matched the practices were	We have introduced the results of the matching in the opening section of the results, highlighting the variables which show the greatest differences in baseline matching. We have included the seventeen graphs illustrating these analyses in the appendix rather than the main paper as to do the latter seemed rather to overbalance the paper. The relevant section in the results section of the main paper now reads as follows: *We examined the extent to which the intervention and control practices were matched for baseline characteristics, the results of which are shown graphically in figures A1 to A17 in the appendix. The variables showing substantial departures from the national profile are the percentage of patients who are over 65 (figure A10), over 85 (figure A11), the practice deprivation score (figure A13) and white (figure A14). In particular we see that intervention practices tend to have fewer old patients compared to England and are on average located in more deprived areas (i.e. their IMD score is higher). In general, the matching has done a reasonable job of reproducing the distribution of matching variables in the intervention practices, even for those variables where substantial departures are seen from the national distribution. However, some small, and statistically significant, deviations remain. As described above, we further adjusted for practice characteristics in the analysis to isolate so far as possible the effect of the intervention.*
Discuss limitation of analysing costs without looking at effectiveness measure. Potential for health/patient experience gains from addressing unmet need or other spillover effects of closer working across sectors?	As above, we have now highlighted that this is a limitation with the analysis however, we consider that it is too speculative to discuss spillover effects, not least because spillover is precisely what this type of multi-disciplinary working is intended to achieve. We have added further discussion about unmet need (see response to referee 1 and the possibility that the additional elective care was for lower value interventions). See response below regarding spillover.
Discussion of change in effects over years?	We already include in the discussion a comment that some of the benefits of the programme were only seen towards the end of the four year programme, citing a US case management intervention where costs initially rose with reductions in costs only seen after a period. We have added the following sentences to the end of this paragraph: *"It may be that some of the negative evaluations of previous case management programmes result from insufficient time for the evaluations to bed in. These may take more time than payers anticipate, especially where changes to working practices and culture are required" (ref Ling T, Brereton L, Conklin A, Newbould J, Roland M. Barriers and facilitators to integrating care: experiences from the English Integrated Care Pilots International Journal of Integrated Care 2012. 12; 24th July)*

More discussion of spillover effect findings (potentially small saving/cost neutral, excluding start-up costs, plus decreased A&E, electives and outpatient) versus effects on those directly targeted would be good. Does this suggest potential for prevention to impact more positively on programme aims than focus on immediately high-risk patients alone?	As above, we have introduced some wording into the discussion to highlight that the impact of the intervention might extend beyond secondary care use.
Looks like you can be a bit more specific on your theory of increased cost due to unmet need. This appears to be quite clearly shown through route of utilisation, i.e. planned, elective admissions and outpatient attendance.	We have added a line of text to more explicitly highlight that the increase in costs (see also response to referee 1 on the possibility of the increased utilisation being related to interventions of lower value).
(p.9, lines 38-41) "While it is standard to allow for confounders by using external controls, many areas of England were adopting some sort of initiative to better co-ordinate care." This is also dealt with effectively with your dose analysis where you restrict to practices within the local area where you know the precise integration activity – add as strength	We have added some wording into the discussion to highlight this point. The additional text reads: "*Further, our dose-response, based on the intensity of HAs and ICM, restricted the analysis to practices within Southwark and Lambeth where the level of HA and ICM activity was known. This analysis effectively treats all Southwark and Lambeth practices as controls for each other and overcomes some of the limitations associated with identifying controls."*
References 6 and 16 are currently the same	We thank the reviewer for spotting this error. This has now been corrected.
Another paper that might be worth comparing to is this one https://bmjopen.bmj.com/content/6/4/e010468.short where we likewise looked at both direct intervention effects and spillovers effects of a case management integrated care intervention. Similar indications to your results. Might also be worth referencing this recent large systematic review of integrated care by Baxter et al. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5946491/ Might give you a reference to speculate on potential effectiveness on wider outcomes than utilisation of secondary care and costs?	The issue of spillover is quite complicated for this intervention which started by aiming to screen 50% of the elderly population (though they didn't actually achieve this) – so it's rather different to some interventions which use risk prediction models to identify those at highest risk. We discuss this issue in the expanded introduction (see response to referee 1) and have included the Baxter review in this additional text. Given the geographic dispersal of controls, we think it unlikely that spillover effects would have occurred from SLIC: a greater problem is the presence of similar interventions in other parts of the country, a limitation which we highlight in the paper. We also comment in the discussion on possible wider benefits over and above secondary care utilisation and this is also implicit in the final paragraph of the paper in which we urge payers introducing such programmes to focus on improving care rather than just reducing costs.

(Appendix p.4, lines 28-31) “In this model a categorical variable is included taking a value of 0 for all observations prior to the start of SLIC and for control practices at all time points. For SLIC practices it takes the value of the number of years since the start of the intervention, i.e. 1 in 2012/13, 2 in 2013/14 etc”. The model appears to show some similarities to a difference-in-differences with gradual introduction (e.g. Propper et al. 2002 https://www.ncbi.nlm.nih.gov/pubmed/11939240), where your categorical variable acts as a more conventional ‘intervention*POST’ term (adding the additional dose of treatment in terms of years since start). Have you tested the intervention/controls for parallel trends in the pre-period? Can you show that this assumption is satisfied (I’d guess it will be if wellmatched)/correct for it if not?	As outlined in the second paragraph of the appendix, the models contain a random slope for time clustered at the practice level. This allows for different background trends for all practices, including different trends between intervention and control practices. As such the assumption of parallel trends has not been made.
Assignment of integrated care models, I would argue, is largely decided at the CCG level (which were mandated to implement some form of integration with the 2012 Health & Social Care Act). Have you clustered standard errors by CCG to account for this design?	We have not accounted for clustering by CCG. We have adjusted for clustering by practice to allow for overdispersion and this should, to a large extent also account for CCG level clustering. However, we also note that our control practices have been selected from across the country and so the degree of clustering should be low.
(Appendix p.6-11) ‘Further results of the statistical analysis’. The yearly breakdown of results of the overall intervention effects is interesting, maybe worth mentioning a bit more detail in the main text? Arguably, it shows that utilisation is generally decreasing (with the exception of ACSCs) in the latter years, so potentially longer-term effects might show better cost-effectiveness (are you able to break down your net cost estimates in Table A14 by year too to test this)?	We think it would be unwise to extrapolate trends in the effectiveness of the SLIC programme beyond the bounds of our data given the limited number of years and that decreasing utilisation happens only over two or three timepoints. Any statistical analysis of the rate of change (i.e. slope) would have wide confidence intervals and as it would be based on only two or three points, we prefer not to carry out the analysis suggested here.

VERSION 2 – REVIEW

REVIEWER	Meredith Rosenthal Harvard TH Chan School of Public Health, USA
REVIEW RETURNED	23-Sep-2018

GENERAL COMMENTS	The authors addressed all of the reviewer comments to my satisfaction and the paper is indeed stronger as a result.

REVIEWER	Jonathan Stokes University of Manchester, UK
REVIEW RETURNED	04-Sep-2018

GENERAL COMMENTS	The authors have adequately addressed my previous comments and, in my opinion, the comments of the other reviewers.