

Improving the efficiency of individually randomised clinical trials by staggering the introduction of the intervention

Richard Hooper¹, Charles Knowles²

¹ Centre for Primary Care & Public Health, Blizard Institute, Queen Mary University of London, London UK

² Centre for Neuroscience & Trauma, Blizard Institute, Queen Mary University of London, London UK

Address for correspondence:

Dr Richard Hooper,
Centre for Primary Care & Public Health,
Queen Mary University of London,
Yvonne Carter Building,
58 Turner Street,
London E1 2AB

r.l.hooper@qmul.ac.uk

Abstract

In cluster randomised trials the introduction of the intervention can be staggered in different clusters, leading to a stepped wedge design. This strategy can lead to gains in efficiency, which might also translate to the context of individually randomised trials, though this has been relatively unexplored. Here we present one illustrative example. We consider trials in which participants start in a control condition such as routine care, and can cross over at any stage to the active intervention. We make the assumption that the effect of the intervention is the same however long the delay before a participant crosses over to the intervention condition. We consider designs for a trial with three repeated assessments, including a baseline, and show that a three-sequence design with staggered introduction of the intervention in two of the sequences estimates the treatment effect after one period more efficiently than a parallel groups design.

Keywords

Clinical trials, efficiency, individually randomised, stepped wedge trials, waiting list control

1. Introduction

Staggering the introduction of an intervention for different individuals in a clinical trial can be a useful strategy when the trial is cluster randomised, for example at the level of a general practice or a local community, giving rise in this case to stepped wedge designs and other design variants.^{1,2} One motivation for doing this may be to reduce the total number of participants needed for the trial, while another motivation may be practicality.^{1,3,4} It is reasonable to wonder whether the same strategy could similarly improve the efficiency of an individually randomised trial, though this seems not to have been explored in the literature. Here we illustrate one scenario – a trial involving a baseline assessment of outcome and two follow-up assessments, with the first of these being the primary outcome – where an individually randomised trial with a staggered intervention design needs fewer participants to achieve the same precision or statistical power (for the primary outcome, at least) as a more familiar parallel groups design.

In Sections 2 and 3 we set out our assumptions and statistical model, and in Section 4 we review the general solution for the variance of the treatment effect estimate in this framework. Section 5 describes possible designs for a trial featuring three repeated assessments including a baseline, and compares the efficiency of these designs. Section 6 shows how conclusions are affected if we relax some of the assumptions of the model. Finally in Section 7 we discuss whether staggering the intervention is a useful approach in general for individually randomised trials.

2. Context

We focus here on pragmatic trials of healthcare technologies or innovations in which participants who start in the control condition (usually routine care) can cross over at any stage to the active intervention. We assume that participants are assessed at baseline (randomisation) and at regular times thereafter. While stepped wedge trials often randomise clusters simultaneously, we imagine here the more usual situation for individually randomised trials in which participants are randomised sequentially over a period of time. To compare trial designs we must make further assumptions about how the intervention works: here we assume the effect of the intervention is the same however long the delay before a participant crosses over to the intervention condition (at least over the timescale for which we plan to follow participants in the trial). This is, in fact, an implicit assumption in any study where we might consider offering the intervention to control participants at the end of the follow-up period on ethical grounds.

To avoid making too many assumptions we will allow the expected outcomes of participants who remain in the control condition to change over time following randomisation. (We acknowledge a tension between this generalisation and the earlier assumption that the effect of intervention should be the same whenever it is introduced. The latter assumption is likely to become untenable if there is a significant, natural deterioration or improvement in participants.) In cluster randomised stepped wedge trials it is common to assume further that the effect of the intervention is maintained at a constant level once it is introduced. While this might be defensible when the intervention is a policy change instituted at cluster level, and participants at each assessment are new, incident cases, it is an assumption which is almost never made in individually randomised trials where participants are followed longitudinally. Here we will allow the effect of treatment to change over time (for example to be attenuated) following its introduction.

As an example, consider a trial of bariatric (weight loss) surgery.⁵ Weight loss achieved by this type of surgery, and improvements in obesity-related conditions such as type 2 diabetes and hypertension, can be rapid, but need not be permanently maintained. In the UK, bariatric surgery is available through the National Health Service (NHS) to very obese patients (especially those with

medical complications) who can demonstrate they have tried other weight loss methods without success. “Without success” means we can assume that weight is not decreasing over time in the population waiting for surgery. We might assume further that that weight remains at a constant level, though it is also possible that it will continue to increase. In trials of bariatric surgery vs non-surgical treatments published to date, changes in weight or body mass index have typically been assessed over a period of 12 or 24 months.⁶ Waiting times for bariatric surgery in the NHS, meanwhile, can be longer than a year,⁷ making a waiting list control design appealing. Although bariatric surgery is already available, a Cochrane review has called for more high quality trials to evaluate and compare operative techniques.⁶

We can imagine a waiting list control design for a trial in which a comparator group who receive non-surgical treatment for the first 12 months of their involvement are then offered surgery, with all participants followed for a further 12 months and assessments made at baseline, 12 and 24 months. The assumption is that the effect of surgery is the same in both groups. (If we model change on a log-ratio scale rather than a difference scale this assumption might still be reasonable even if we are concerned that weights in the comparator group are increasing over the first 12 months.)

3. Statistical model

Under the general assumptions set out above, the mean of a continuous outcome measure for an individual follows a trajectory with baseline β (the mean at time 0) and time effects τ_t , $t = 0 \dots T$, with $\tau_0 = 0$ for identifiability. When the individual crosses over to the active intervention there are additional effects of treatment: α_1 in the first period after crossing over, changing to α_2 one period later, then to α_3 , and so on. One (or more) of these treatment effects α_k will generally be the primary target for estimation in the trial.

A trial design consists of two or more randomised ‘sequences’ (we use this term to be consistent with terminology used for cross-over and stepped wedge trials), which differ according to when (or if) the intervention is introduced (Figure 1). For simplicity we assume an equal number of participants, n , in each sequence. We assume a mixed linear regression model for outcome with a random effect of participant. That is, for individual $i = 1 \dots n$ in sequence $j = 1 \dots J$ at time $t = 0 \dots T$ after randomisation the outcome is

$$Y_{ijt} = \beta + \tau_t + A_{jt} + \eta_{ij} + \varepsilon_{ijt} \quad (1)$$

where

$$\begin{aligned} \eta_{ij} &\sim N(0, \sigma_b^2), \\ \varepsilon_{ijt} &\sim N(0, \sigma_w^2), \end{aligned}$$

with η_{ij} , ε_{ijt} all independent, and

$$A_{jt} = \begin{cases} 0 & \text{if sequence } j \text{ is in the control condition at time } t \\ \alpha_k & \text{if sequence } j \text{ crosses over to the intervention between times } t - k \text{ and } t - k + 1, k \geq 1. \end{cases}$$

Under this model assessments of the same participant at different times have an “exchangeable” correlation structure, with the same correlation, r , between outcomes at any two time-points:

$$r = \sigma_b^2 / (\sigma_b^2 + \sigma_w^2).$$

We also define the total variance σ^2 as

$$\sigma^2 = \sigma_b^2 + \sigma_w^2,$$

so that (σ^2, r) becomes an alternative parameterisation of the variance components (σ_b^2, σ_w^2) .

4. Precision of treatment effect estimates

If we write outcomes Y_{ijt} as a single column vector \mathbf{Y} , and write (1) in simpler matrix form as

$$\mathbf{Y} = \mathbf{X}(\beta, \tau_1, \dots, \tau_T, \alpha_1, \alpha_2, \dots)' + \mathbf{e}, \quad \mathbf{e} \sim \mathbf{N}(\mathbf{0}, \mathbf{V})$$

then for known \mathbf{V} the generalised least squares estimators (which are the best linear unbiased estimators) of the parameters are

$$(\hat{\beta}, \hat{\tau}_1, \dots, \hat{\tau}_T, \hat{\alpha}_1, \hat{\alpha}_2, \dots)' = (\mathbf{X}^T \mathbf{V}^{-1} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{V}^{-1} \mathbf{Y}$$

with variance-covariance matrix

$$\text{Var}(\hat{\beta}, \hat{\tau}_1, \dots, \hat{\tau}_T, \hat{\alpha}_1, \hat{\alpha}_2, \dots)' = (\mathbf{X}^T \mathbf{V}^{-1} \mathbf{X})^{-1}.$$

Matrices \mathbf{X} and \mathbf{V} have their simplest form when n and σ^2 are both 1. We refer to these as $\mathbf{X}_{(1)}$ and $\mathbf{V}_{(1)}$. $\mathbf{X}_{(1)}$ has a single row for each repeated assessment in each randomised sequence (see the examples in Figure 1). $\mathbf{V}_{(1)}$ is the Kronecker product $\mathbf{I}_J \otimes \mathbf{R}$, where \mathbf{I}_J is the identity matrix of order J and \mathbf{R} is the matrix of correlations between J repeated assessments of the same individual. In the examples in Figure 1, where $J = 3$,

$$\mathbf{R} = \begin{pmatrix} 1 & r & r \\ r & 1 & r \\ r & r & 1 \end{pmatrix}.$$

Then for general n , σ^2 and r the variance-covariance matrix can be written

$$\text{Var}(\hat{\beta}, \hat{\tau}_1, \dots, \hat{\tau}_T, \hat{\alpha}_1, \hat{\alpha}_2, \dots)' = \frac{\sigma^2}{n} (\mathbf{X}_{(1)}^T \mathbf{V}_{(1)}^{-1} \mathbf{X}_{(1)})^{-1}$$

or, writing the total sample size as $N = Jn$,

$$\text{Var}(\hat{\beta}, \hat{\tau}_1, \dots, \hat{\tau}_T, \hat{\alpha}_1, \hat{\alpha}_2, \dots)' = \frac{\sigma^2}{N} J (\mathbf{X}_{(1)}^T \mathbf{V}_{(1)}^{-1} \mathbf{X}_{(1)})^{-1}. \quad (2)$$

5. Designs for trials with three repeated assessments including a baseline

5.1 Parallel groups designs with baseline assessments

All the designs we consider in this article include a baseline assessment at time 0, at the point of randomisation. Such baseline assessments are common in trials: they need not delay the overall

timetable of the study, and they always improve the precision of the treatment effect estimate.⁸

We use the term “parallel groups design with one or more baseline assessments” to refer to a two-sequence design in which one sequence crosses over to the intervention condition during the trial while the other remains in the control condition throughout. We use designs of this type as comparators when discussing the relative efficiency of alternatives which stagger the intervention in different sequences. The parallel groups design illustrated in Figure 1a, for example, allows us to estimate the two treatment effects α_1 and α_2 , with (see Appendix)

$$\text{Var}(\hat{\alpha}_1) = \text{Var}(\hat{\alpha}_2) = \frac{\sigma^2}{N} 4(1 - r^2).$$

If we have no interest in α_2 but still want to take advantage of three repeated assessments of each individual then we might prefer the design illustrated in Figure 1b, which effectively includes two baseline assessments. In this case (see Appendix)

$$\text{Var}(\hat{\alpha}_1) = \frac{\sigma^2}{N} 4 \left(1 - \frac{2r^2}{1+r} \right).$$

5.2 A two-sequence, waiting list control design

The design illustrated in Figure 1c could be termed a waiting list control design: here the sequence which acts as the control in the first follow-up period is then given the intervention and both sequences include a second follow-up. In practice, trials described as having a waiting list control design would not always bother with this second follow-up, but (as with the parallel groups design) it allows us to estimate both treatment effects α_1 and α_2 , this time with (see Appendix)

$$\text{Var}(\hat{\alpha}_1) = \frac{\sigma^2}{N} 4(1 - r^2)$$

and

$$\text{Var}(\hat{\alpha}_2) = \frac{\sigma^2}{N} 8(1 + r - 2r^2).$$

5.3 A three-sequence design with staggered introduction of the intervention in two sequences

Now consider the three-sequence design in Figure 1d, in which the intervention is introduced in a staggered fashion in two of the sequences, with a third sequence remaining in the control condition throughout. For this design we used numerical matrix inversion to evaluate the expression in (2) for different r .

5.4 Relative efficiency of different designs

Because the variance of a treatment effect estimate from a given design is proportional to σ^2/N , the respective constants of proportionality (more specifically their inverse ratio) determine the asymptotic relative efficiency of two competing designs – the ratio of sample sizes needed by both to achieve the same precision.

Figure 2 shows the asymptotic relative efficiency of the waiting list control design (Figure 1c) and the design with staggered introduction of the intervention (Figure 1d) relative to a parallel groups design

with one or two baseline assessments (Figure 1a or Figure 1b) for different r . Observe that the waiting list control design is no more efficient for estimating α_1 than a parallel groups design, and is between 50% and 67% less efficient for estimating α_2 . The staggered intervention design, on the other hand, is between 56% and 77% more efficient for estimating α_1 than the parallel groups design with a single baseline measurement, though it is also between 22% and 29% less efficient for estimating α_2 . If we use a parallel groups design with two baseline assessments as the comparator instead, the staggered intervention design is still between 34% and 56% more efficient for estimating α_1 .

5.5 Sample size calculation

Suppose we want to design a trial in the bariatric surgery example so that we have 80% power at the 5% significance level to detect a difference in body mass index at 12 months (α_1) of 2kg/m^2 , assuming a standard deviation of 5kg/m^2 , and a correlation between any two repeated measurements of body mass index of 0.8.

Using the expressions in Section 5.1, the variance of the treatment effect estimate from a parallel groups design with one or with two baseline assessments is $1.44\sigma^2/N$ or $1.1556\sigma^2/N$, respectively. The variance of the treatment effect estimate from the waiting list control design in Section 5.2 is $1.44\sigma^2/N$. Evaluating expression (2) in the case of the three-sequence staggered intervention design shows that the variance of the treatment effect estimate is $0.8374\sigma^2/N$.

The ‘fundamental equation’ of sample size tells us that if the variance of the treatment effect estimate is $w\sigma^2/N$ then the sample size needed to demonstrate an effect of at least δ with 80% power at the two-sided 5% significance level is

$$N \geq (z_{0.975} + z_{0.80})^2 \frac{w\sigma^2}{\delta^2},$$

where z_p is the $100p^{\text{th}}$ centile of the standard Normal distribution.⁹

Thus we would need 72 participants for a parallel groups design with a single baseline assessment (or for a waiting list control design), 58 participants for a parallel groups design with two baseline assessments, but only 42 participants for the staggered intervention design (sample sizes have been rounded up to a multiple of the number of sequences – 2 in the case of parallel groups designs, and 3 in the case of the staggered intervention design).

6. Relaxing the exchangeability assumption

We could relax the assumption of exchangeability by considering an alternative correlation matrix, \mathbf{R} . The following correlation matrix for three repeated assessments assumes there is one correlation, r_1 , between outcomes assessed at consecutive time-points, and another correlation, r_2 , between outcomes assessed at the first and last time-points:

$$\mathbf{R} = \begin{pmatrix} 1 & r_1 & r_2 \\ r_1 & 1 & r_1 \\ r_2 & r_1 & 1 \end{pmatrix}.$$

Figure 3 plots the relative efficiency of the design with staggered intervention relative to a parallel groups design with one or two baseline assessments, using the above correlation matrix. The Figure shows the relationship between relative efficiency and r_1 for fixed values of r_2 (note that in practice we would not expect r_2 to exceed r_1). For any fixed value of r_2 the design with staggered

intervention is always at least 56% more efficient for estimating α_1 than a parallel groups design with a single baseline measurement. The design with staggered intervention is also at worst 33% less efficient for estimating α_2 than a parallel groups design with a single baseline measurement (the curves for fixed r_2 tend asymptotically to a constant 67% as r_2 approaches 1). If we use a parallel groups design with two baseline assessments as the comparator instead, the picture is less one-sided: for some combinations of r_1 and r_2 with $r_1 > r_2$ the staggered intervention design is then less efficient for estimating α_1 .

7. Discussion

The two-sequence waiting list control design has the appealing quality, at least, that all participants get the intervention in the end, though whether this meets ethical objections to a parallel groups design has been debated.¹⁰ It does not offer any advantage in terms of efficiency over a parallel groups design, at least in the situation modelled here. A three-sequence design with staggered introduction of the intervention in two of the sequences, on the other hand, offers a more efficient way (in terms of the number of participants we need to recruit, and the time taken to recruit them) to estimate the treatment effect after one period than a parallel groups design with the same length of follow-up. If our primary hypothesis relates to the treatment effect after one period, with a secondary hypothesis relating to the treatment effect after two periods, we could address both with the staggered intervention design – the primary hypothesis much more efficiently and the secondary hypothesis only slightly less efficiently than with a parallel groups design. In our bariatric surgery trial with body mass index assessed at baseline, 12 and 24 months, for example, the treatment effect at 12 months is estimated more efficiently with the three-sequence design than with a parallel groups design, and the effect at 24 months slightly less efficiently. We are not aware of any published or ongoing trials using the design in Figure 1d.

Our findings depend on the assumption that the effect of the intervention is independent of the time when a participant crosses over from the control to the intervention, but this may be reasonable in trials conducted in healthy or chronically ill populations whose health status is relatively stable. We have only focused on designs with unidirectional cross-over (such as occurs with surgery), and we have not made any assumptions about what happens if and when a participant crosses from the intervention condition back to the control. The most efficient design of all may be a two-way cross-over design,¹¹ but this requires additional assumptions and is clearly not appropriate to every intervention.

We have compared the efficiency of designs that are ‘complete’ in that they include an assessment of each individual at each time-point. All the designs we have considered are run over three time-points, so their relative efficiency in terms of the total number of individuals needed is the same as their relative efficiency in terms of the total number of assessments. If we want a design that minimises the number of assessments rather than the number of individuals then we may want to consider ‘incomplete’ designs, some of which have been investigated elsewhere.¹²

The staggered intervention design presented here should not be considered a ‘magic bullet’ for improving the efficiency of any pragmatic clinical trial – there are still challenges. Extending the length of follow-up opens up the possibility of differential attrition in different sequences, leading to bias. We need sufficient time available for the trial to allow us to schedule two follow-up assessments, and the added burden of the second follow-up must be acceptable to participants. The extended follow-up could mean some participants remain in the control condition for longer than they would have to wait for treatment outside the trial – something they would need to consider before consenting.

The superiority of the staggered intervention design also depends on our being more interested in the effect of the intervention after just one time period than in the effect at a later time. This will not apply to every context – in fact in evidence-based medicine the main focus is more usually on the long-term maintenance of an effect. We do not want to encourage short-termism in general in the selection of the primary endpoint.

In spite of their limitations, however, we think that trial designs with staggered introduction of the intervention will repay further study in the individually randomised case, just as they have stimulated renewed methodological interest in the field of cluster randomised trials.

Appendix

Variances of treatment effect estimators from the designs considered in this article can be derived by analogy with parallel groups designs where a sustained treatment effect α is estimated over v follow-up assessments adjusting for u baseline assessments⁹: in this case the variance of $\hat{\alpha}$ is

$$\text{Var}(\hat{\alpha}) = \frac{4\sigma^2}{N} \left(\frac{1 + (v-1)r}{v} - \frac{ur^2}{1 + (u-1)r} \right). \quad (3)$$

In the parallel groups design in Figure 2a, for example, the assessments of outcome at the second follow-up do not contribute to the estimate of α_1 . This is easily demonstrated: if the estimator for α_1 is a linear combination of the means in different sequences at different times,

$$\hat{\alpha}_1 = \sum \lambda_{jt} \bar{Y}_{jt},$$

then it has expectation

$$E(\hat{\alpha}_1) = \beta \sum \lambda_{jt} + \tau_1(\lambda_{11} + \lambda_{21}) + \tau_2(\lambda_{12} + \lambda_{22}) + \alpha_1\lambda_{11} + \alpha_2\lambda_{12},$$

and since $\hat{\alpha}_1$ is unbiased the coefficients of α_2 and τ_2 must be zero, hence $\lambda_{12} = \lambda_{22} = 0$. This means we estimate α_1 with a single follow-up assessment adjusting for a single baseline, so we can use equation (3) with $(u, v) = (1, 1)$. Similarly for estimating α_2 . The parallel groups design in Figure 2b, meanwhile, has $(u, v) = (2, 1)$.

In the case of the waiting list control design in Figure 2c, a similar argument shows we can use equation (3) for the variance of $\hat{\alpha}_1$, with $(u, v) = (1, 1)$. If we write $\hat{\alpha}_2$ as a linear combination of the means in different sequences at different times,

$$\hat{\alpha}_2 = \sum \lambda_{jt} \bar{Y}_{jt},$$

then it has expectation

$$E(\hat{\alpha}_2) = \beta \sum \lambda_{jt} + \tau_1(\lambda_{11} + \lambda_{21}) + \tau_2(\lambda_{12} + \lambda_{22}) + \alpha_1(\lambda_{11} + \lambda_{22}) + \alpha_2\lambda_{12},$$

and in order for this to be unbiased we need

$$\lambda_{12} = 1, \quad \lambda_{22} = -1, \quad \lambda_{11} = 1, \quad \lambda_{21} = -1, \quad -\lambda_{10} = \lambda_{20} = \lambda,$$

i.e.

$$\hat{\alpha}_2 = (\bar{Y}_{.11} + \bar{Y}_{.12} - \lambda\bar{Y}_{.10}) - (\bar{Y}_{.21} + \bar{Y}_{.22} - \lambda\bar{Y}_{.20}).$$

This is reminiscent of the situation where a sustained treatment effect is estimated over two follow-up assessments adjusting for one baseline assessment, except that this gives an estimator of the form

$$\hat{\alpha} = \frac{1}{2}(\bar{Y}_{.11} + \bar{Y}_{.12} - \lambda\bar{Y}_{.10}) - \frac{1}{2}(\bar{Y}_{.21} + \bar{Y}_{.22} - \lambda\bar{Y}_{.20}).$$

Hence the variance of $\hat{\alpha}_2$ is four times the variance in (3) with $(u, v) = (1, 2)$.

References

1. Hemming K, Haines TP, Chilton PT, Girling AJ, Lilford RJ. The stepped wedge cluster randomised trial: rationale, design, analysis, and reporting. *BMJ*. 2015;350:h391.
2. Hooper R, Bourke L. Cluster randomised trials with repeated cross sections: alternatives to parallel group designs. *BMJ*. 2015;350:h2925.
3. Lawrie J, Carlin JB, Forbes AB. Optimal stepped wedge designs. *Stat Probabil Lett*. 2015;99:210-214.
4. Girling AJ, Hemming K. Statistical efficiency and optimal design for stepped cluster studies under linear mixed effects models. *Stat Med*. 2016;35(13):2149-2166.
5. Vidal J, Corcelles R, Jiménez A, Flores L, Lacy AM. Metabolic and Bariatric Surgery for Obesity. *Gastroenterology*. 2017;152(7):1780-1790.
6. Colquitt JL, Pickett K, Loveman E, Frampton GK. Surgery for weight loss in adults. *Cochrane Database of Systematic Reviews*. 2014;8:CD003641.
7. Bariatric surgery waiting times: A Freedom of Information request to Derby Teaching Hospitals NHS Foundation Trust. WhatDoTheyKnow. https://www.whatdotheyknow.com/request/bariatric_surgery_waiting_times. Accessed 18 July 2018.
8. Frison L, Pocock SJ. Repeated measures in clinical trials: analysis using mean summary statistics and its implications for design. *Stat Med*. 1992;11(13):1685-1704.
9. Machin D, Campbell MJ, Tan SB, Tan SH. *Sample size tables for clinical studies*. 3rd ed. Chichester, UK: Wiley-Blackwell; 2009.
10. Prost A, Binik A, Abubakar I, et al. Logistic, ethical, and political dimensions of stepped wedge trials: critical review and case studies. *Trials*. 2015;16:351.
11. Jones B, Kenward MG. *Design and analysis of cross-over trials*. 3rd ed. London, UK: CRC Press; 2015.
12. Hooper R, Bourke L. The dog-leg: an alternative to a cross-over design for pragmatic clinical trials in relatively stable populations. *Int J Epidemiol*. 2014;43(3):930-936

Figure legends

Figure 1. Designs for trials with three repeated assessments including a baseline, with the corresponding matrix $\mathbf{A} = (A_{jt})$ from equation (1), and matrices $\mathbf{X}_{(1)}$ and $\mathbf{V}_{(1)}$ from equation (2). In each case the schematic on the left shows sequences (rows) against time (columns), with shaded areas showing when the intervention is introduced: (a) parallel groups design with a single baseline assessment; (b) parallel groups design with two baseline assessments; (c) two-sequence waiting list control design; (d) three-sequence design with staggered introduction of the intervention in two sequences.

Figure 2. Relative efficiency with which different designs estimate α_1 (solid line) and α_2 (dashed line) for different correlations r , assuming exchangeability. Efficiency of the waiting list control design (first row) and staggered intervention design (second row) is shown relative to a parallel groups design with a single baseline assessment (left-hand column) and with two baseline assessments (right-hand column).

Figure 3. Relative efficiency with which different designs estimate α_1 and α_2 for different correlations r_1 between consecutive time-points, and r_2 between first and last time-points. Efficiency of the staggered intervention design is shown relative to a parallel groups design with a single baseline assessment (left) and with two baseline assessments (right). Where the matrix inversion routine concluded that a matrix was not positive definite, the graph is continued to zero or to infinity.