Working Paper Draft

The impact of sample rotation patterns and composite estimation on survey outcomes

Philip A. Bell

July 1998

This Working Paper Series is intended to make the results of current research within the ABS available to other interested parties. The aim is to present accounts of ABS developments and research work or analysis of an experimental nature so as to encourage discussion and comment.

Views expressed in this paper are those of the author(s) and do not necessarily represent those of the Australian Bureau of Statistics. Where quoted or used, they should be attributed clearly to the author(s)

ABS Catalogue No. 1351.0

Contents

INQUIRIES

- *For further information about the contents of this Working Paper, please contact the author:*
  Philip Bell — *telephone:* (02) 6252 7053 *or facsimile:* (02) 6252 8015 *or e-mail:* philip.bell@abs.gov.au
- *For information on the* Working Papers in Econometrics and Applied Statistics, *please contact the Managing Editor,*
  Leanne Johnson — *telephone:* (02) 6252 5355 *or facsimile:* (02) 6252 8015 *or e-mail:* leanne.johnson@abs.gov.au *or write* c/- Analytical Services Section (W42b), Australian Bureau of Statistics, P.O. Box 10, Belconnen, ACT 2616, Australia.

- *For information about other ABS statistics and services* please refer to the back of this publication.

Abstract

The aim of a repeated survey is to allow one or more items to be monitored across time. For survey design purposes this aim has often been simplified to two objectives: good estimates of the item for each period, and good estimates of period to period change. In the Australian Labour Force Survey (LFS) these objectives lead to a design with high overlap between successive monthly samples.

Focusing on good estimates of the "underlying trend" of the series, and how it changes over time, could lead to quite different survey designs. Previous work suggests that a sample rotation pattern with no month to month overlap would provide better trend estimates. Unfortunately such a rotation pattern gives poor estimates of month to month change.

This paper considers an alternative estimator, the linear composite estimator, in combination with various sample rotation patterns. A rotation pattern is presented in which individuals are sampled for two successive months out of every four months, giving a 50 percent overlap of sample between successive months. By using composite estimation this rotation pattern yields improved estimates of trend while maintaining good estimates of month to month change.

1 Introduction

1.1 Survey outcomes and sample design

The aim of a repeated survey is to allow one or more items to be monitored across time. For survey design purposes this aim has often been simplified to two objectives: good estimates of the item for each period, and good estimates of period to period change. In the Australian Labour Force Survey (LFS) these objectives lead to a design with high overlap between successive monthly samples.

This paper suggests that survey designers should take account of objectives related to longer term change across time. For many surveys, successive estimates behave quite erratically from period to period. Users of such data will often be attempting to assess the "underlying direction" of the series, perhaps using some smoothing technique or making an assessment "by eye". In doing so they are incorporating information from a number of periods up to the current period. Survey designs that seek to optimise the survey for such longer term assessments may be quite different to those that are optimal for period to period change.

Tallis (1995) suggested that high overlap between successive surveys for the LFS reduces the ability to detect turning points in the economy. This and work by Sutcliffe and Lee (1995) suggest that a sample rotation pattern with no month to month overlap would provide better estimates of the underlying direction of the series. This paper extends this work by considering an alternative estimator, the linear composite estimator, in combination with various sample rotation patterns. Composite estimation is not currently used in the LFS, though a different form known as the AK composite estimator has been used for many years in the US Current Population Survey (Gurney and Daley (1965)).

Section 2 defines a variety of outcomes for a repeated survey. Besides a number of standard estimates, it introduces a "trend" estimate that attempts to smooth out seasonal effects and local irregularities. This trend is introduced as a surrogate for the various methods of assessing "underlying direction" of the series. Outcomes of interest are measures of the value and direction of the trend at the end of the series, and also how much the trend at a time point is revised as estimates for later times become available. Variance and mean squared error for the various outcomes are defined.

Sections 3 and 4 describe two aspects of survey design that can be changed to alter these survey outcomes. Section 3 describes survey rotation patterns, which control the overlap between the units selected in the survey for different months. The current, high overlap pattern for the LFS is presented, along with two alternative rotation patterns that would lead to lower overlap between successive months.

Section 4 describes different survey estimators. It presents a class of linear composite estimators which make use of data from a number of successive months. These estimators make use of the correlation structure of the survey estimates to produce estimators with lower variance than the simple estimator. How useful these estimators are depends on the correlation structure and hence on the survey rotation pattern.

Section 5 presents the effects of the available rotation patterns and estimators on the various survey outcomes, in the case of the LFS. It is seen that the different designs are good for different outcomes, with the current rotation pattern good for month to month change but inferior to the other patterns for assessing longer term direction of the series.

Section 6 gives conclusions of the paper. While previous studies have presented the impact of rotation pattern on trend, this paper is new in assessing the combined impact of composite estimation and rotation pattern. One of the rotation patterns presented (the "2 in 2 out" pattern), which has performed poorly in previous studies, is seen to be quite effective in combination with composite estimation. The final message to survey designers is the importance of knowing the key outcomes of the survey and using this information in assessing different survey designs.

2 A discussion of survey outcomes

2.1 Level and movement objectives

The basic aim of any survey is to provide estimates of various population characteristics with sufficient accuracy for the uses to which they are put. In a one-off survey this maps to a fairly clear objective - we want to get low bias and low sampling error for one or more key estimates.

In a repeated survey we wish to provide good estimates not just of values at a single time point, but also of how the population is changing over time. These objectives are related, since sufficiently accurate estimates at each time point must result in a good picture of changes over time. Because of this, much sample design work has been focused on obtaining good cross-sectional estimates (or *level estimates*). For this purpose the focus of design work is typically the size and composition of the sample and how to use any available extra data such as population benchmarks.

Designing for good level estimates leaves considerable room for affecting the quality of longitudinal measures. Consider the estimate of change between two months (the lag one *movement estimate*). The sampling error on this estimate depends not just on the sampling error on the level estimates but also on the correlation between estimates from the two months. The best estimates of movement will result from a high correlation - this can often be obtained by retaining a large portion of the sample common to the two months.

The key design parameter affecting the estimates of change is the overlap between successive samples. Maintaining high overlap between repeats of a survey is operationally convenient, since many sampled units have been located and have some experience of the survey. High overlap also improves the estimates of lag one movement in cases where a unit's responses for an item are highly correlated between successive periods.

Many repeated surveys have been designed with estimates of level and lag one movement as the sole design objectives. This leads to survey designs that have high overlap between successive survey periods. The motivation for such a design is easy to express to users of the survey, and is unlikely to raise controversy.

2.2 Objectives related to longer term change

Unfortunately, in many repeated surveys it would be inappropriate for users to respond strongly to the movement from one period to the next. The lag one movement may behave quite erratically. One reason is sampling error on the estimates - the survey may simply not be large enough to detect real period to period movements of the size users wish to respond to. A second reason is that the true sequence of population values is affected by irregularity - short term and transient changes in the population which have little relationship to policy evaluation or prediction of future values.

To make sensible decisions in such a series users need a longer term view of changes in the population. This requires comparing data over longer periods. A movement over three or four periods may be used, or some smoothing of the data over time. For a monthly survey, users may take quarterly averages as a way of smoothing the data - these can then be compared across time, being a more stable series.

Sophisticated users of a repeated survey recognise the danger of responding to the lag one movement in its own right. This is evidenced by the widespread use of methods aimed at providing a more reliable long term picture. However, survey designers have rarely recognised estimation of longer term change as a survey objective.

It turns out that the best survey designs for estimating longer term change may be very different to those that are best for estimating lag one movement. In particular, a low overlap between periods may lead to improved estimates of longer term change.

2.3 Introducing the "trend"

It is difficult to define what we mean by longer term change, which makes it hard to measure this aspect of the performance of a survey design. One approach to this is to produce a variety of measures, such as movements at longer lags or, for a monthly survey, quarterly averages and their movements. We will follow this approach in some of the evaluation, demonstrating that the same survey designs are appropriate for improving a variety of measures.

In addition, we will introduce the "trend" of the series. The term "trend" here refers to a smoothing of the series that attempts to remove seasonal variation as well as short term irregular variation. The trend results from a time series decomposition of the series into trend, seasonal and irregular components (and other components such as trading day effects).

Many statistical agencies use methods of time series decomposition based around the X11 program (Shiskin *et. al.* 1967). For the purposes of this paper we choose a method that was derived as a linear approximation to the X11 method by Dagum *et. al.* (1996). This method is used to represent the sort of trend outcome obtained by time series decomposition in most agencies. Because it uses a linear transformation of the survey estimates, it is straightforward to analytically derive measures of accuracy of estimates under this trending procedure.

While the formulae and results presented in this paper are specific to the particular trend used, they should give a good indication of what users are achieving with their various smoothing techniques. In this sense the trend given is presented as a surrogate for what users and agencies are currently doing to determine the direction of the underlying series.

2.4 Estimates and variances for outcomes

The variance matrix of the survey estimates

Let Y_t be the true population value of the item of interest at time t, and let y_t be the survey estimate for time t. Write $Y = \{Y_t\}$and $y = \{y_t\}$as column vectors containing this data for times t=1,...,N.

The survey estimates are assumed to be unbiased, and standard methods can be used to calculate their variances and covariances using the survey data. The variance-covariance matrix (or more simply, variance matrix) of the survey estimates is given by

$$V \quad = \mathrm{E}(y - Y)(y - Y)',$$

for E( ) indicating expectation across possible samples.

In practice it is often appropriate to smooth estimates of variance and covariance across time to obtain the estimate of V. This requires making some assumptions about the stationarity of the sampling error e.g. $\mathrm{var}(y_t) = \sigma^2$ and $\mathrm{cov}(y_t, y_{t-k}) = \sigma^2\rho_k$.

Variance of derived estimates

Let $\alpha = \{\alpha_t\}$ be a vector of parameters that define a linear combination $\alpha' y = \Sigma_t \alpha_t y_t$ of the survey estimates. The variance of such a linear combination is given by

$$\operatorname{var}(\alpha' y) = E(\alpha' y - \alpha' Y)(\alpha' y - \alpha' Y)' = \alpha' V \alpha$$

This formula can be used to obtain estimates of movements at various lags, or other derived estimates such as quarterly averages. For example, the lag 1 movement uses $\alpha' = (0\ 0\ \ldots\ 0\ \text{-}1\ 1)$. The movement between two quarterly averages would use $\alpha' = (0\ 0\ \ldots\ 0\ \text{-}\frac{1}{3}\ \text{-}\frac{1}{3}\ \text{-}\frac{1}{3}\ \frac{1}{3}\ \frac{1}{3}\ \frac{1}{3})$.

Under the simple stationarity assumptions given above the lag *k* movement has variance given by $\operatorname{var}(y_t - y_{t-k}) = 2\sigma^2(1-\rho_k)$. It is clear that this variance is minimised by a survey design which gives large correlation at lag *k*.

Outcomes related to trend estimates

Under the linear approximation to X11, the trend for any time point is a linear combination of values at a number of time points. Assume the number of time points available N is large. Let M= $\{t : N - M < t \le M\}$ be a set of time points defining the middle of the series - points far enough from the beginning and end of the series that adding more estimates would not appreciably affect the trend for time points in M.

Write T_{M} as the matrix which gives trend values for time points in M based on the N data points, so that the estimated trend is $T'_{\mathrm{M}} y$. We call $T'_{\mathrm{M}} y$ the *mid trend*.

The true trend is defined to be $T'_{\mathrm{M}} Y$. That is, the true trend is the result of applying our trending method if we knew the series of true population values for a sufficiently large number of times before and after the period of interest.

Write $T'_{\mathrm{E}} y$ for the trend for time points in M estimated from data in M only - we call this the *end trend*. This is not unbiased for the true trend, since its expectation is $T'_{\mathrm{E}} Y \neq T'_{\mathrm{M}} Y$.

Outcomes of interest are given in the form $\alpha' T'_{\mathrm{E}} y$ (end estimates) or $\alpha' T'_{\mathrm{M}} y$ (mid estimates). We define three outcomes that appear critical: *level* of trend uses $\alpha = (0\ 0\ \ldots\ 0\ 1)$, *movement* of trend (at lag 1) uses $\alpha = (0\ 0\ \ldots\ 0\ \text{-}1\ 1)$ and *curvature* of trend uses $\alpha = (0\ 0\ \ldots\ 0\ 1\ \text{-}2\ 1)$.

Movement of the trend may be more important to users than its level. Users are often interested in turning points, where the trend changes from increasing to decreasing. This clearly is related to trend movement. Curvature of the trend is the second difference of the trend, and it is concerned with changes in the trend direction. Such changes are also of key interest to users, and it seems clear that a good estimate of turning point requires a small sampling error on the change in trend movement between successive time points, ie. on the curvature.

Finally, for any trend outcome the value at the end of the series is modified as estimates for later months become available. The *trend revision* for a given outcome will be defined as the difference between its value at the end of the series (based only on data to time M) and its value in the middle of the series (ie. after all revisions). The revision is thus given by $\alpha' T'_{\mathrm{M}} y - \alpha' T'_{\mathrm{E}} y$.

Mean squared error and revision for trend outcomes

The variance of a mid trend estimate $\alpha' T'_{\mathrm{M}} y$ is given by

$$\begin{aligned}\mathrm{var}(\alpha' T'_{\mathrm{M}} y) &= \mathrm{E}(\alpha' T'_{\mathrm{M}} y - \alpha' T'_{\mathrm{M}} Y)(\alpha' T'_{\mathrm{M}} y - \alpha' T'_{\mathrm{M}} Y)' \\ &= \alpha' T'_{\mathrm{M}} V T_{\mathrm{M}} \alpha\end{aligned}$$

The mean squared error of the corresponding end trend estimate $\alpha' T'_{\mathrm{E}} y$ is

$$\begin{aligned}\mathrm{mse}(\alpha' T'_{\mathrm{E}} y) &= \mathrm{E}(\alpha' T'_{\mathrm{E}} y - \alpha' T'_{\mathrm{E}} Y)(\alpha' T'_{\mathrm{E}} y - \alpha' T'_{\mathrm{E}} Y)' \\ &= \alpha' T'_{\mathrm{E}} V T_{\mathrm{E}} \alpha + \alpha' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) Y Y' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) \alpha\end{aligned}$$

The first term here is the variance matrix $\mathrm{var}(\alpha' T'_{\mathrm{E}} y)$ of the end trend estimate, while the second term is the squared revision that would occur given the true data. It due to the bias which arises because the end trend does not predict the true trend perfectly. This second term is independent of the sample design.

The mean squared revision matrix for this outcome is given by

$$\begin{aligned}&\mathrm{E}(\alpha' T'_{\mathrm{E}} y - \alpha' T'_{\mathrm{M}} y)(\alpha' T'_{\mathrm{E}} y - \alpha' T'_{\mathrm{M}} y)' \\ &\quad = \alpha' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) V (T_{\mathrm{E}} - T_{\mathrm{M}}) \alpha + \alpha' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) Y Y' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) \alpha\end{aligned}$$

Both the mean squared error at the end and the mean squared revision contain a component that does not depend on survey design. Since we are focused on the effect of sample design it is appropriate to exclude this component from our measurements. So for estimates of $\alpha' T'_{\mathrm{M}} Y$ the key measures to calculate are the variance of the trend estimates, $\mathrm{var}(\alpha' T'_{\mathrm{M}} y)$and $\mathrm{var}(\alpha' T'_{\mathrm{E}} y)$, and the variance of the revisions, given by

$$\mathrm{var}(\alpha' T'_{\mathrm{E}} y - \alpha' T'_{\mathrm{M}} y) = \alpha' (T'_{\mathrm{E}} - T'_{\mathrm{M}}) V (T_{\mathrm{E}} - T_{\mathrm{M}}) \alpha .$$

3 Impact of survey rotation pattern

3.1 Rotation pattern in the LFS

Methods for controlling overlap between successive surveys will depend on the nature of the repeated survey. We will describe overlap control that uses a fixed survey rotation pattern. The details will be from the LFS, a monthly household survey that controls overlap by using a rotation pattern. Much of this description will apply straightforwardly to similar household surveys.

The LFS is a survey of the civilian population of Australia aged 15 years or older. Dwellings are selected first by selecting geographic areas, and then by choosing a cluster of dwellings from each area. Data is collected for all in-scope individuals in these dwellings.

The initial stage of this multi-stage selection process is to select geographic areas. These are divided into eight "rotation groups" which are used to control rotation of dwellings into and out of the survey.

The current "rotation pattern" in the LFS consists of sampling the same dwellings from a rotation group each month for eight months. In the next month new dwellings from the same geographic areas are selected, and they will be sampled for eight more months. The month at which new dwellings are selected is different for each rotation group.

This rotation pattern ensures that there is an overlap between sampled dwellings in seven eighths of the geographic areas between any two successive months. This gives high correlations between successive estimates from the same rotation group.

3.2 Alternative rotation patterns

The current LFS rotation pattern is referred to as "8 in", since a new set of dwellings remains in sample for 8 months. This paper focuses on two alternative patterns which result in reduced correlation between successive months.

The first alternative will be referred to as the "1 in 2 out" pattern. In this rotation pattern each dwelling is sampled once a quarter up to a total of eight times. In the other months of the quarter, different dwellings from the same geographic regions would be sampled. This rotation pattern would produce no sample overlap from month to month.

The second alternative will be called the "2 in 2 out" pattern. In this rotation pattern each dwelling is sampled two months in a row out of every four months, for a total of eight times in sample. Different dwellings from the same geographic regions would be sampled on the other two months of the four. With this rotation pattern half of the sample would be common to consecutive months.

These patterns can be varied by reducing or increasing the number of times each dwelling is sampled. The specific patterns compared in this paper sample each dwelling eight times, and for the same sample size they require the same number of geographic areas. So the methods have a similar cost to maintain, and the sample at any time point will be equally clustered under each of the rotation patterns.

Other statistical agencies use different rotation patterns for their labour force surveys. Statistics Canada uses a 6 in pattern. The U.S. Current Population Survey uses a 4 in 8 out pattern, while Japan uses a 2 in 10 out pattern. These last two patterns allow considerable overlap between samples a year apart, with the objective of improving estimates of year to year movement. Both the alternative patterns presented here also allow overlap a year apart.

4 Impact of composite estimation

4.1 Simple estimates

Let $\tilde{y}_{r,t}$ be an estimate of Y_t based on data from the rth out of R rotation groups. Define the series of simple estimates $\hat{y}^{S} = \{\hat{y}_t^{S}\}$ in which the estimate for a given time point is the mean of the rotation group estimates for that time point (i.e. $\hat{y}_t^{S} = \frac{1}{R}\sum_{r=1,\ldots,R} \tilde{y}_{r,t}$).

This simple estimate may differ somewhat from the standard survey estimate, since the survey estimates are typically not calculated as the mean of the rotation group estimates. The simple estimates are used in this paper as proxies for the standard survey estimates.

4.2 Linear composite estimates

The simple estimates at a time point depend only on survey values obtained at that time point. By using values obtained at nearby time points it is possible to improve on these simple estimates by taking advantage of the autocorrelations between estimates at the rotation group level.

Let $\tilde{y}_W = \{\tilde{y}_{r,t}\}_{r=1,\ldots,R;\, t\in W}$ be a column vector of the rotation group estimates based on a set of times W (known as the window). Define a linear composite estimator as a linear combination $\beta'\tilde{y}_W$ of the rotation group estimates which is unbiased for the value of interest.

Let C_W be the matrix such that $\mathrm{E}(\tilde{y}_W) = C_W Y_W$, for Y_W the true population values for times in the window W. Then the expected value of a linear combination $\beta'\tilde{y}_W$ is given by

$$E(\beta'\tilde{y}_W) = \beta'\mathrm{E}(\tilde{y}_W) = \beta' C_W Y_W \ .$$

To obtain an unbiased estimator of an outcome $\alpha' Y_W$ requires imposing the constraints $C'_W\beta = \alpha$.

The optimum choice of β minimises the variance of the composite estimator (ie. $\mathrm{var}(\beta'\tilde{y}_W) = \beta'\mathrm{var}(\tilde{y}_W)\beta$) under these constraints. The matrix $\mathrm{var}(\tilde{y}_W)$ is the variance matrix of the rotation group estimates, which depends on the rotation pattern being used.

Using standard results for minimisation of a quadratic form under linear constraints (see for example Rao (1973) p. 65) the optimal β is given by $\beta^W(\alpha) = \mathrm{var}(\tilde{y}_W)^{-1} C_W Q^- \alpha$, for Q^- any generalised inverse of $(C'_W \mathrm{var}(\tilde{y}_W)^{-1} C_W)$. Writing $\beta^W = \mathrm{var}(\tilde{y}_W)^{-1} C_W Q^-$ this reduces to $\beta^W(\alpha) = \beta^W \alpha$.

Thus $\hat{y}^W = \beta^{W\,\prime}\tilde{y}_W$ is the linear composite estimator based on the window W that is unbiased for Y_W and has minimum variance. The optimal linear composite estimator of an outcome $\alpha' Y_W$ is $\alpha'\hat{y}^W$.

The dependence of the composite estimators on the window W is important, as different windows will give different estimators. Note that all estimates based on the same window will agree, in the sense that the estimates of $\alpha'_1 Y_W$ and $\alpha'_2 Y_W$ add to the estimate of $(\alpha_1 + \alpha_2)' Y_W$.

4.3 Composite estimators and revisions

In a repeating survey the first composite estimate available for a time point M will be based on a window of points $W = \{t : M - L < t \le M\}$. It is possible to update previous estimates to be based on this same window. This will improve the estimates for those time points, and ensures that other estimates such as movement estimates will be optimal. Unfortunately it will result in revisions of the survey estimates as new data arrives.

A sensible approach is to use a fixed size of window L for composite estimation, and to allow a fixed number R of revisions. When a new month of data arrives, the window is moved and optimal composite estimates are computed for the final time point and the previous R time points. Estimates from earlier times are left fixed at their last computed value.

With a large window and a number of revisions, the composite estimates from this approach will be nearly optimal for any linear combination of the population characteristics. They will, for example, be nearly optimal for estimating trend and items such as movement of trend. With no revisions the only estimate that is optimal is the end level estimate. Nevertheless, a strategy with no revisions is attractive to users, and some evaluation of this option will be presented.

Looking at trend revision is complicated by composite estimation with revision. Suppose we write y^{E} as the vector of composite estimates available at time M, and y^{M} as the vector of composite estimates available at a later time N when M is in the middle of the series. The trend revision on a composite estimate of $\alpha' T'_{\mathrm{M}} Y$ becomes

$$\operatorname{var}(\alpha' T'_{\mathrm{E}} y^{\mathrm{E}} - \alpha' T'_{\mathrm{M}} y^{\mathrm{M}})$$

The elements of y^{E} and y^{M} are linear combinations of the rotation group estimates $\tilde{y}_{r,t}$, and so the variance of this trend revision can be calculated based on the variance matrix of these estimates.

5 Outcomes for various survey designs

5.1 Details of the LFS situation

For the calculations in this paper, monthly estimates of persons by labour force status were obtained for each rotation group, categorised by month, sex, age (grouped as 15–19, 20–24, ... , 50–54, 55–64, 65+) and part-of-state (14 geographic regions covering Australia). Within these categories, the estimates for each rotation group were pro-rated to match known population benchmarks.

The autocorrelation structure of these rotation group estimates has been discussed in previous papers - Bell and Carolan (1998) and Bell (1998). The following model for the autocorrelations is assumed:

$$\begin{aligned} \operatorname{Corr}(\tilde{y}_{r,t}, \tilde{y}_{r,t-k}) &= \rho_{\mathrm{W}k} && \text{for estimates from the same set of dwellings} \\ &= \rho_{\mathrm{B}k} && \text{for estimates from the same rotation group} \\ & && \text{but from different sets of dwellings.} \end{aligned}$$

This model assumes that the sampling error autocorrelation in a rotation group depends only on the lag and on whether the rotation group has a common sample of dwellings between the two time points. The values $\rho_{\mathrm{W}k}$ and $\rho_{\mathrm{B}k}$ will decrease as lag k increases, with $\rho_{\mathrm{W}k} \geq \rho_{\mathrm{B}k}$. In the case of the LFS, the following four parameter model fits the autocorrelations well on data up to lag 7:

$$\rho_{Wk} = (1 - r_U^2)(\theta_P{}^k r_P^2 + \theta_B{}^k (1 - r_P^2)) \text{ and} \tag{10}$$

$$\rho_{Bk} = (1 - r_U^2)\theta_B{}^k (1 - r_P^2) \text{ .} \tag{11}$$

The current rotation pattern does not allow rotation groups to have common dwellings at lags over 7 months, so the model was used to extrapolate for these longer lags. It appears that the results are not very sensitive to this extrapolation. For discussion of the model, including interpretation of the parameters r_U, r_P, θ_P and θ_B, please refer to the previous papers. The correlations assumed for this paper at various lags are shown in table 1. They assume the fitted parameter values $\theta_P = 0.87697$, $\theta_B = 0.94$, $r_U = 0.3101$ and $r_P = 0.90456$ for proportion employed and $\theta_P = 0.81164$, $\theta_B = 0.94$, $r_U = 0.50038$ and $r_P = 0.91713$ for proportion unemployed. Standard errors σ_E assumed for the simple estimates are 0.21 percentage points for proportion employed and 0.11 percentage points for proportion unemployed.

Using these autocorrelations the variance matrix var($\tilde{y}_W$) can now be produced for any given rotation pattern and window W.

Table 1: Estimated autocorrelations of rotation group estimates

Item	Lag k:	1	2	3	4	8	12	18
Proportion employed	ρ_{Wk}	0.80	0.71	0.64	0.57	0.36	0.23	0.12
	ρ_{Bk}	0.15	0.15	0.14	0.13	0.10	0.08	0.05
Proportion unemployed	ρ_{Wk}	0.62	0.52	0.44	0.37	0.19	0.11	0.05
	ρ_{Bk}	0.11	0.11	0.10	0.09	0.07	0.06	0.04

5.2 Estimators being compared

To specify linear composite estimate requires defining a window size L, and number of revisions R and the item for which the estimator is to be optimised. In the tables and graphs presented, the composite estimator being used will be denoted by the notation C L,R and the simple estimator by the letter S. For most comparisons, the composite used will be C11,5. This composite uses a window of 11 months of data, and allows estimates to be revised five times.

The composites will be optimised for the item proportion unemployed. One reason is that the correlations assumed for proportion unemployed may be more typical of other variables than the higher correlations assumed for proportion employed. An estimator optimised for proportion unemployed achieves as much as is possible with the lower correlations, while still achieving good results for proportion employed.

The comparisons here are based upon a series of N=90 months, with the middle of the series defined to be all but the first 12 and last 12 months (i.e. M=78 is used). These values are sufficient to give results near those of the ideal situation, which would have N very large and M considerably smaller.

5.3 Results for various rotation patterns and estimates

Broad comparison

Table 2 presents the standard errors (SEs) achieved at the end of the series for various outcome measures, for four rotation patterns and with simple and composite estimation. The fourth rotation pattern is the 4 in 8 out pattern used in the U.S. Current Population Survey. The standard errors are given as a percentage of the standard error of a simple estimate of level.

Table 2: Standard errors of simple and composite estimates at the end of the series, proportion employed
(as % of SE for simple level estimate)

Pattern and estimator		original		quarterly average		end trend	
		level	movement	level	movement	level	movement
8 in	S	100	75	88	80	97	19
"	C(11,5)	94	66	82	67	89	16
1 in 2 out	S	100	130	66	55	72	15
"	C(11,5)	98	127	65	53	71	14
2 in 2 out	S	100	102	76	71	81	17
"	C(11,5)	89	78	72	60	78	15
4 in 8 out	S	100	85	84	91	94	22
"	C(11,5)	91	70	77	70	84	17

The current 8 in pattern achieves the best standard errors for lag 1 movement - this is expected, since this design has the greatest overlap at lag 1. It does not perform particularly well for the longer term measures (level and movement of quarterly average and level and movement of trend). For all outcomes the composite estimates give lower standard errors than the simple estimates.

The 1 in 2 out pattern gives very poor standard errors for lag 1 movement, but is very good for the longer term indicators in this table. Composite estimation achieves relatively little for the 1 in 2 out rotation pattern. It seems unlikely that composite estimation would be used with this rotation pattern, given the extra complexity involved. For this reason only the simple estimator will be presented for the 1 in 2 out pattern in later results.

The 2 in 2 out pattern appears as something of a compromise between good long term estimates and good lag 1 movement estimates. Standard errors under the 2 in 2 out pattern are greatly improved by composite estimation, especially for the lag 1 movement estimate. In fact, composite estimation transforms this rotation pattern - with composite estimation the standard errors compare well with those achieved under other designs given her. Only the composite estimator will be presented for the 2 in 2 out pattern in later results.

Finally, the 4 in 4 out estimator is shown as an example of what is achieved under other rotation patterns. This rotation pattern is much improved by composite estimation, and in fact the U.S. survey uses a composite estimator (though not of the form described in this paper).

Comparison to results from simple estimator for current pattern

The remaining comparisons will be between four designs:

8 in S	current rotation pattern, simple estimator
8 in C	current rotation pattern, composite estimator
2 in 2 out C	2 in 2 out with composite estimator, and
1 in 2 out S	1 in 2 out with simple estimator

Graph 1 presents a bar chart giving standard errors for the same outcomes as table 2, but expressed as a percentage of the standard error achieved under the "8 in S" design. Graph 2 is the same but for proportion unemployed.

Graph 1: Standard error, proportion employed (relative to 8 in S)

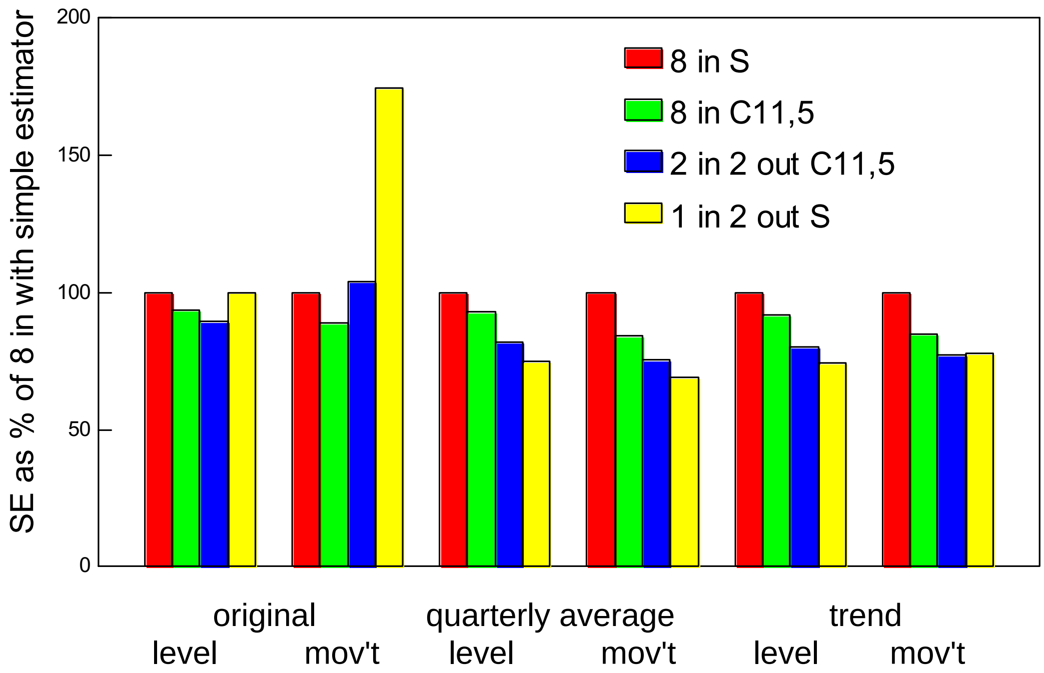

Graph 2: Standard error, proportion unemployed (relative to 8 in S)

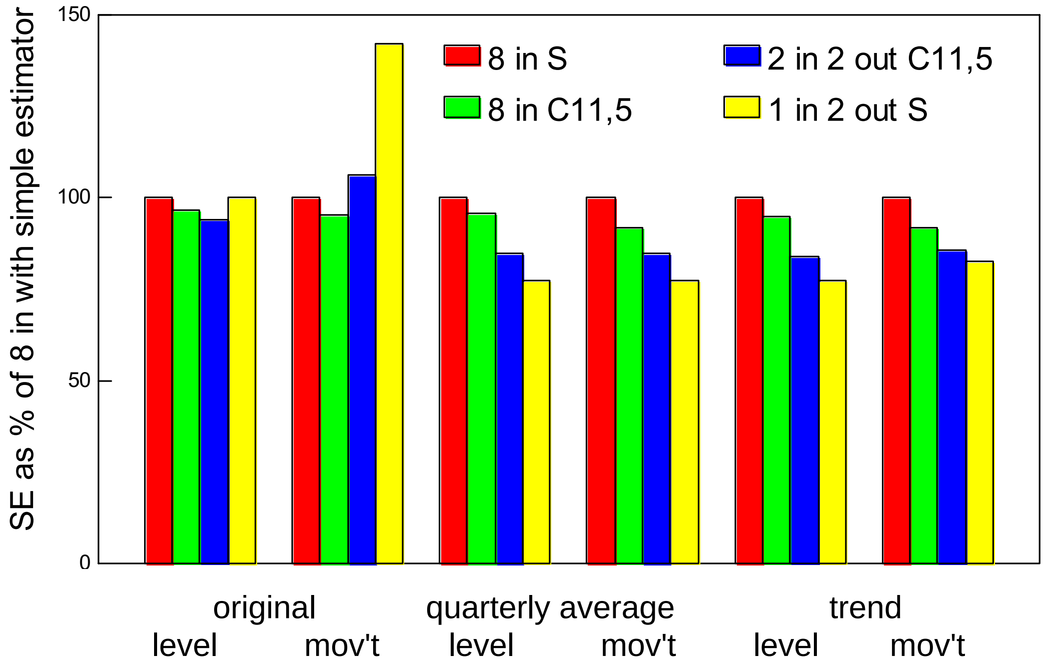

Graph 3 presents the standard errors for movement of proportion employed at various lags. "2 in 2 out C" performs well for movement at lag 3 or more, and is not too bad for lag 1 or 2. "1 in 2 out S" is good at some specific lags, but very poor at lags 1 and 2.

Graph 3: Standard error (movements), proportion employed (relative to 8 in S)

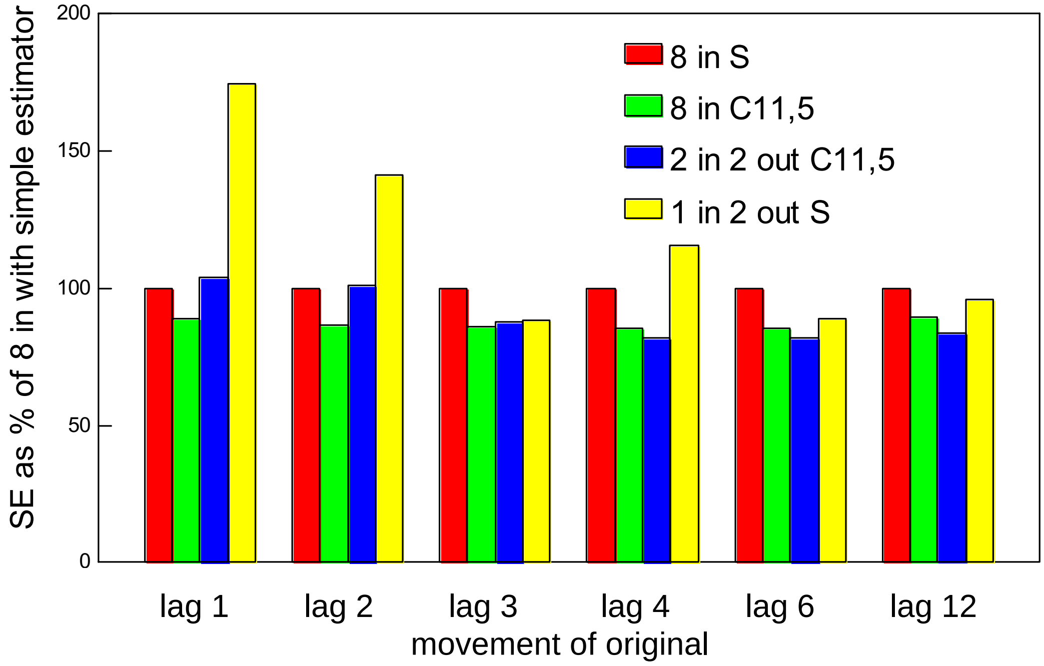

Graph 4 presents outcomes related to the end trend. Standard errors are given for the trend level, the trend movement at lag 1 and lag 3, the trend curvature, the revision of the trend level and the revision of lag 1 trend movement. "1 in 2 out S" has the lowest standard errors for most of these measure, but has the highest standard error for trend curvature. The curvature of the trend at the end apparently is affected by the poor behaviour of the lag 1 movement under this design. "2 in 2 out C" performs consistently well for these trend outcomes.

Graph 4: Standard error (trends), proportion employed (relative to 8 in S)

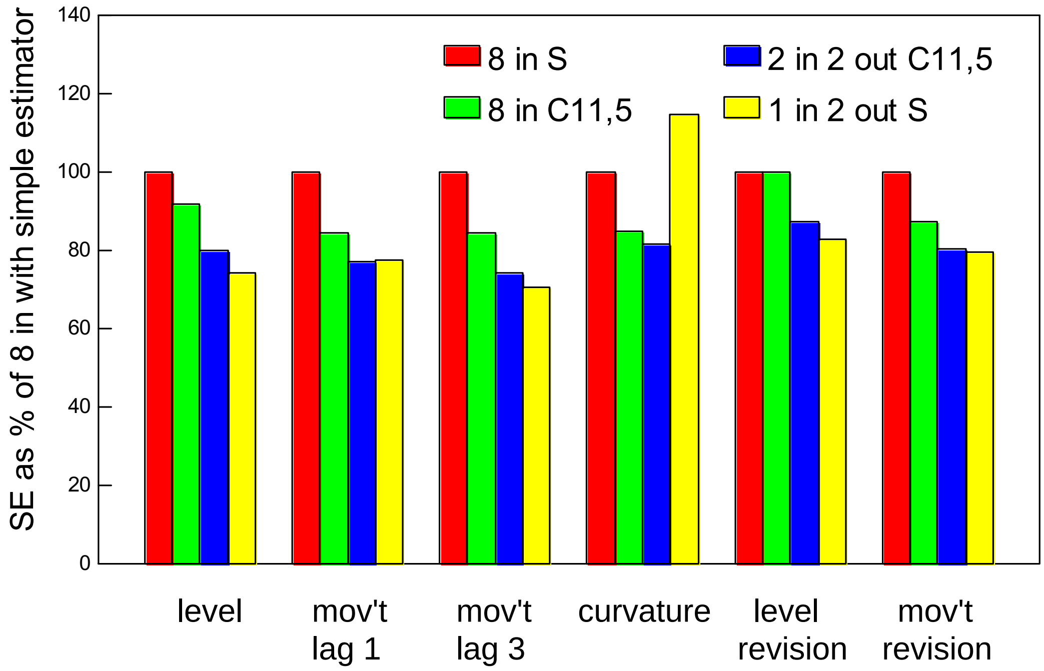

Comparison between composite estimators

The composite estimators presented above used 11 months of data and assumed 5 revisions. It may be desirable to use a smaller window, and to use fewer or no revisions. The negative to this is that with a small window or few revisions the estimators will be less optimal, particularly for the longer term outcomes.

Comparisons of six different estimators are given in graph 5 for the "8 in" rotation pattern and in Graph 6 for the "2 in 2 out" pattern.

Graph 5: Standard error, proportion employed (relative to 8 in S)

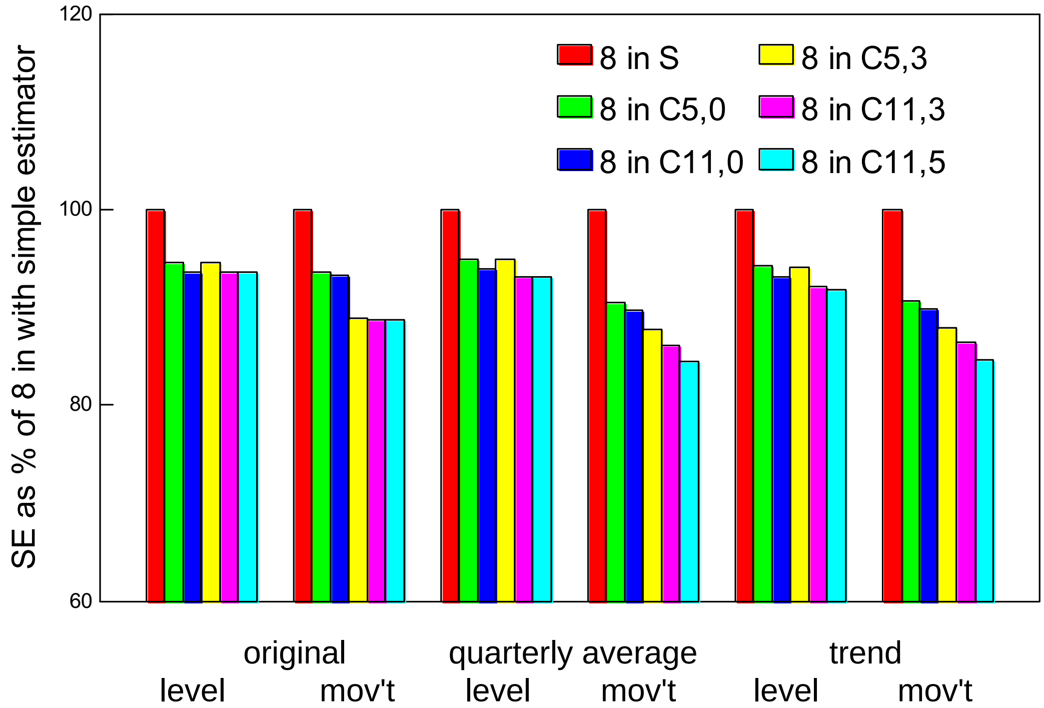

Graph 6: Standard error, proportion employed (relative to 8 in S)

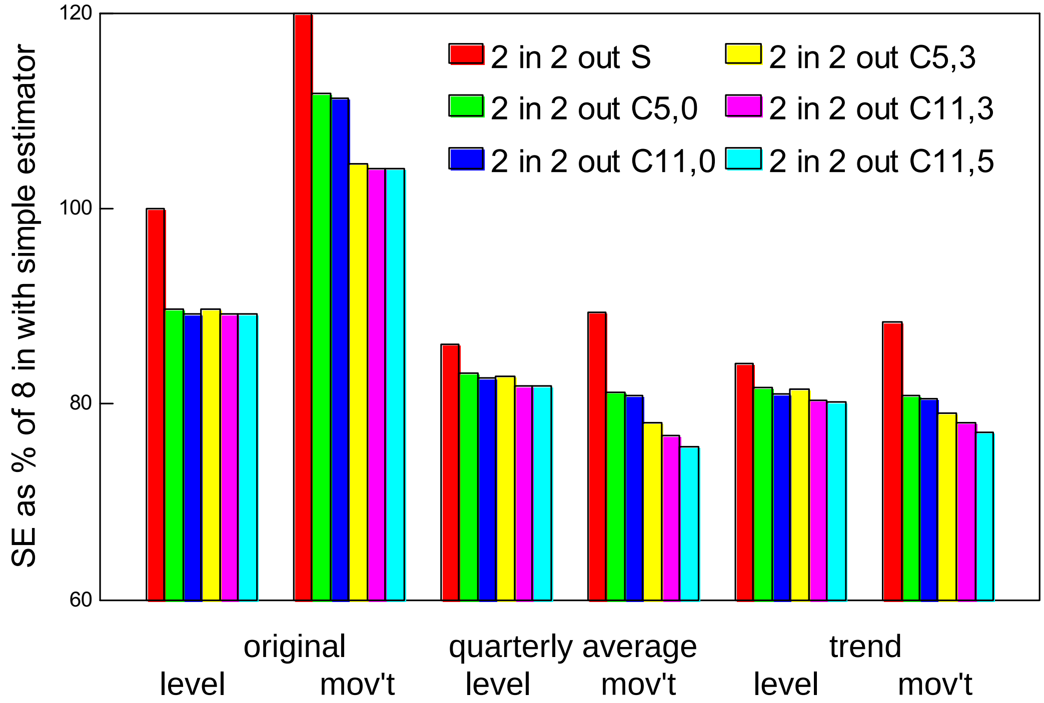

The general picture is similar for both patterns, with standard errors improving as window size and number of revisions increase. Revisions are particularly important to achieve the best lag 1 movement estimates. Longer windows always reduce the standard errors, but have the greatest effect for long term indicators, particularly movement of quarterly average, and movement of trend.

5.4 Simulating a series of survey errors

It is useful to get a feel for the effect of the different designs on the series of survey estimates. To do this, sampling error was simulated by drawing from the multivariate normal distribution with mean 0 and variance matrix var($\tilde{y}_W$). To aid comparing across designs, the same random numbers were used to simulate from each design. This effectively produces a simulated set of rotation group estimates for each rotation pattern under the assumption that the true population values were 0 for all time points. These can then be used to produce a series of simple and composite estimates.

Graphs 7 and 8 show simulated series of sampling errors for the "8 in S" and "1 in 2 out S" designs respectively. Superimposed are two trends applied to these sampling errors, one based on 78 points only (the end trend) and the second based on 90 points (the mid trend, including data from 12 months beyond the last point shown).

Graph 7: Simulated sampling error, 8 in S design, proportion employed

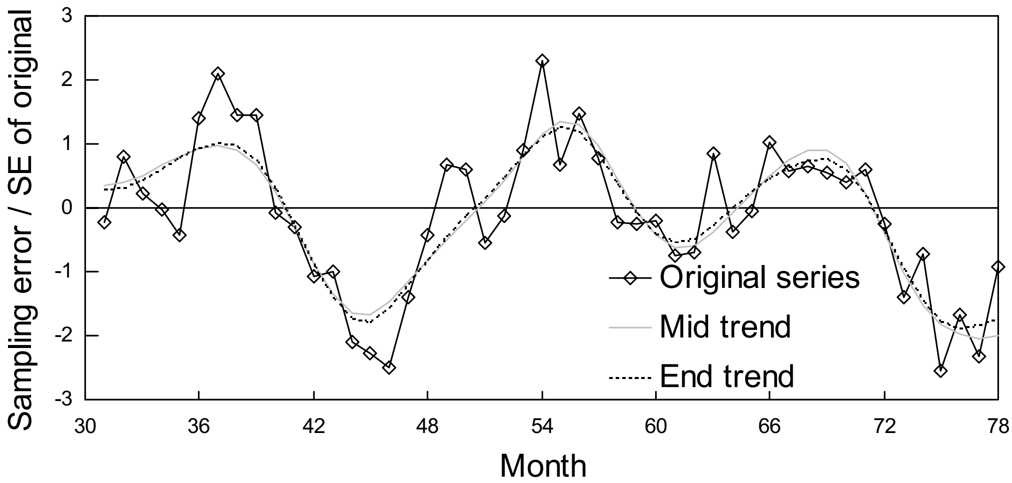

Graph 8: Simulated sampling error, 1 in 2 out S design, proportion employed

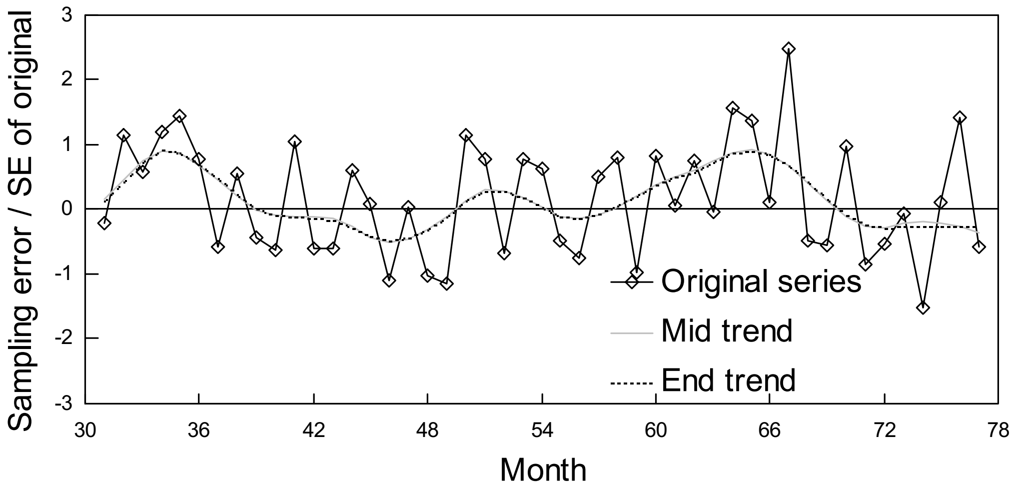

The comparison between "8 in S" and "1 in 2 out S" is quite instructive. The "8 in S" design gives superficially quite well behaved estimates, with small movements between successive points. The problem is that there is a clear longer term movement in the underlying series, induced by the correlation between successive estimates. This apparent trend is spurious, since effectively the true population values are all zero for this simulation.

The "1 in 2 out S" design, in contrast, has large movements at lag 1, and shows an obvious autocorrelation at lag 3. It is harder to discern any great movement in the trend over the period shown - this is good, since the true trend was 0.

The simulations shown are typical of a number of simulations that were run for these two designs. Simulations for the "8 in C" design are similar to the "8 in S" design but with slightly reduced variability. The "2 in 2 out C" design displays behaviour between the extremes represented by the "8 in S" design and the "1 in 2 out S" design. The simulations are relevant because if the "1 in 2 out S" design was adopted, users would be faced with data that looks very different to the current series, and considerably more volatile. For the current estimates quite a lot of sampling error is passing into the trend series - this would be reduced under the "1 in 2 out S" design but at the cost of apparently more irregular estimates.

6 Conclusions

The survey designer is faced with the task of providing estimates that are as useful as possible for the purposes to which they are put. This needs to be tempered with knowledge of the purposes that the data is suitable for. In the LFS case most users would state their key interest as lag 1 movements, yet the data is not suitable for detecting such short term changes in the population. This paper suggests that many survey designs should be aimed at achieving good estimates of longer term change.

The paper suggested a number of outcomes that are of interest to users and that could be assessed in designing a repeating survey. It introduced a trend as a surrogate for the sorts of analysis that users do to determine the longer term behaviour of a series, and specifies outcomes related to the trend. It also examined the effects on the various outcomes of changing two aspects of the survey design - the rotation pattern and the estimator.

In the LFS example there were three main alternatives to the current rotation pattern and estimator. The first was to add composite estimation - this improves standard errors across all outcomes. The second was to move to the "2 in 2 out" pattern with composite estimation - this improved the longer term outcomes further, but was not quite as good for lag 1 movement. The third alternative was to move to a "1 in 2 out" pattern with simple estimation - this could achieve further improvements to the standard error of most longer term outcomes, but was very poor for lag 1 movement. This is very noticeable in the simulated series of sampling errors, where the "1 in 2 out" series appears quite irregular.

There is no magic answer that is best for every possible use of the data. Any design is a trade off - between monthly movement and longer term outcomes, between complexity and simplicity, between cost and accuracy. Designers of repeated surveys should keep in mind the uses made of the data and allow that to influence design choices.

References

Bell, P. A. 1998, 'Using state space models and composite estimation to measure the effects of telephone interviewing on labour force estimates', in *Working papers in Econometrics and Applied Statistics*, ABS, Cat. no. 1351.0, no. 98/2, Canberra.

Bell, P. A. and Carolan, A. 1998, 'Trend estimation for small areas from a continuing survey with controlled sample overlap', in *Working papers in Econometrics and Applied Statistics*, ABS, Cat. no. 1351.0, no. 98/1, Canberra.

Dagum, E.B., Chhab, N. and Chiu, K. 1996, 'Derivation and Properties of the X11ARIMA and Census X11 Linear Filters', *Journal of Official Statistics*, vol. 12, pp. 329-347.

Gurney, M. and Daley, J.F. 1965, 'A Multivariate Approach to Estimation in Periodic Sample Surveys', *Proceedings of the Social Statistics Section of the American Statistical Association*, pp. 242-257.

Rao, C.R 1973, *Linear statistical Inference and its applications, Second edition*, John Wiley and sons, New York

Shiskin, J., Young, A. & Musgrave, J. 1967, *The X-11 variant of census method II seasonal adjustment*, Bureau of the Census, U.S. Department of Commerce, Technical Paper 15.

Sutcliffe, A. and Lee, G. 1995, 'Seasonal analysis and sample design', presented at the Conference of Survey Measurement and Process Quality, Bristol.

Tallis, K. 1995, 'Survey Design and Signal Extraction - Meeting Changing User Expectations', *Proceedings of the Fiftieth Session of the International Statistical Institute, Beijing.*