

ARTICLE TYPE

Investigating the assumptions of the self-controlled case series method

Heather J. Whitaker*^{1,2} | Yonas Ghebremichael-Weldeselassie^{1,3} | Ian J. Douglas² | Liam Smeeth² | C. Paddy Farrington¹

¹School of Mathematics and Statistics, The Open University, Milton Keynes, UK

²Department of Non-Communicable Disease Epidemiology, London School of Hygiene and Tropical Medicine, London, UK

³Warwick Medical School, The University of Warwick, Coventry, UK

Correspondence

*Heather Whitaker, School of Mathematics and Statistics, The Open University, Milton Keynes, MK7 6AA, UK. Email: heather.whitaker@open.ac.uk

Summary

We describe some simple techniques for investigating two key assumptions of the self-controlled case series (SCCS) method, namely that events do not influence subsequent exposures, and that events do not influence the length of observation periods. For each assumption we propose some simple tests based on the standard SCCS model, along with associated graphical displays. The methods also enable the user to investigate the robustness of the results obtained using the standard SCCS model to failure of assumptions. The proposed methods are investigated by simulations, and applied to data on measles, mumps and rubella vaccine, and antipsychotics.

KEYWORDS:

Assumptions, Event dependent exposures, event dependent observation period, Self controlled case series

1 | INTRODUCTION

The self-controlled case series (or SCCS) method is used to investigate the potential association between an exposure and an acute adverse event of interest. (1, 2, 3, 4) In this method, estimation is conditional on the total number of events observed for each individual over a pre-defined period of observation. In consequence, only cases (that is, individuals who have experienced at least one event) are included in the analysis, and all time-invariant confounders that act multiplicatively on the baseline incidence are automatically controlled. The method has proved to be particularly useful for investigating adverse events following vaccination (5) or infection, and is increasingly being used in other areas of pharmacoepidemiology, (6) notably when there is concern about confounding. (7)

Unlike many other more traditional analysis methods, such as cohort and case-control methods, the SCCS method makes use of information after the acute adverse event of interest. In order for the SCCS method to produce unbiased estimates of the association between exposure and event, two assumptions are required: that occurrence of an event does not influence subsequent exposures, and that it does not influence the end of the period of observation. In pharmacoepidemiology, an event may alter subsequent exposure if it is a contra-indication to treatment, so no exposure is possible after the event. Conversely, an event may increase post-event exposure if a medication exposure is prescribed as a direct consequence of the event; for example, such a situation might arise if the events were falls and the exposures were a pain-relief medication that may be prescribed after a painful fall. The assumption that events do not influence the period observed might be violated if the adverse event carries substantial mortality, as in the case of stroke or myocardial infarction, so the event date and observation end date are not always independent. (8) Note that the timing of mortality unrelated to the adverse event is independent of the event time and does not violate SCCS model assumptions.

These two assumptions are required because the SCCS method conditions on the exposure history throughout the observation period for each case, and on that observation period. In particular, this may involve conditioning on exposures and observation time after an event: this feature distinguishes the SCCS method from most other methods used in epidemiology, including the case-crossover method, (7, 9) an alternative case-only method.

The assumption that exposures should not depend on prior events is equivalent to requiring that the exposure should be external, or exogenous (10); this equivalence was previously discussed. (3) The assumption that observation periods should not depend on events is a type of non-informative censoring assumption: if it fails, then the end of observation may carry information about event times. However, unlike the censoring of the outcome event usually encountered in survival analysis, (10) censoring here is of the planned end of observation, after an event has occurred.

These assumptions relate specifically to the event process, the exposure process and the observation process, and are quite separate from issues of confounding by additional variables. Time-invariant multiplicative confounders factor out of the SCCS likelihood, and are thus adjusted automatically. Time-varying confounders, (11) on the other hand, can affect the results obtained in SCCS studies, and need to be included in the model along with age or calendar time. However, this is not the issue considered in the present paper.

Variants of the SCCS method have been proposed for situations in which either of the two assumptions described above is violated. (12, 13, 14) However, it would also be useful to have simple methods of analysis to identify when the assumptions are violated, and whether results obtained using the standard method of analysis are likely to be robust to such violations.

In this paper we propose some simple sensitivity analyses to check these assumptions, and to investigate the robustness of the results to their failure. The methods we propose are all based on the original standard SCCS method, (1, 2) so are straightforward to apply: they can be used prior to trying more complex elaborations of the method designed to cater for circumstances in which the assumptions do not hold. (12, 13, 14) We also highlight some graphical tools to display salient features of the data.

In Section 2 we present the key features of the SCCS method, including the SCCS likelihood. In Section 3 we discuss the assumption that events do not influence subsequent exposures. In Section 4, we discuss the assumption that events do not influence the observation period. In Section 5 we present an evaluation of the proposed methods by simulation. Section 6 contains some applications to real data followed by final remarks in Section 7.

2 | THE SCCS METHOD

Suppose that individual i in a given cohort is followed up during an observation period $(a_i, b_i]$, for the occurrence of an event of interest. Individual i may experience $n_i \geq 0$ events over this observation period. Individual i may also experience exposures: at time $t \in (a_i, b_i]$, the exposure history at t is x_i^t . In the simplest standard SCCS model, (1, 2) x_i^t includes a single categorical exposure function on $(a_i, t]$ taking levels k , $k = 0$ corresponding to the unexposed reference level and $k = 1, 2, \dots, K$ corresponding to risk levels following the onset of exposure(s). This may be represented by a time-varying factor $x_i(t)$. We let $x_i = x_i^{b_i}$ denote the exposure history up to the end of observation, b_i . The event rate is $\lambda_i(t|x_i^t)$.

Over the observation period, individual i 's event rate is modified by age (or calendar time or other measured temporal factors), exposure history, and fixed factors specific to the individual. We assume that these influences are captured by the following multiplicative incidence model:

$$\lambda_i(t|x_i^t) = \phi_i \psi(t) \exp\{x_i(t)^T \beta\}.$$

Here ϕ_i represents the combined effect of individual-specific factors, $\psi(t)$ is the effect of age (or time), assumed to be common to all individuals in the cohort, and $x_i(t)$ is a vector of exposures at age (or time) t derived from x_i^t . The focus of inference is the vector parameter β , which is the log relative incidence associated with exposure.

So far we have defined a cohort model to represent the event rate. The SCCS method bases estimation of β on a likelihood that involves only the cases within that cohort, that is, those individuals that experience one or more events during the observation period. Suppose that there are N such cases. To avoid a proliferation of indices, we re-use i and label the cases $i = 1, \dots, N$. Case i might experience $n_i \geq 1$ events, so the total number of events is $M = \sum_{i=1}^N n_i$. The n_i events for case i occur at times t_{ij} , $j = 1, \dots, n_i$ in $(a_i, b_i]$.

The SCCS likelihood is then

$$L = \prod_{i=1}^N \frac{\prod_{j=1}^{n_i} \psi(t_{ij}) \exp\{x_i(t_{ij})^T \beta\}}{\left(\int_{a_i}^{b_i} \psi(t) \exp\{x_i(t)^T \beta\} dt \right)^{n_i}}.$$

Note that the individual-specific term ϕ_i has factored out: the method adjusts automatically for time-invariant random and fixed effects that act multiplicatively on the event rate. The SCCS likelihood may be derived from a model for the underlying cohort by conditioning, for each case, on the observation period $(a_i, b_i]$, the exposure history x_i up to b_i , and the number of events n_i in the observation period. The method is also valid for non-recurrent events (so $n_i = 1$), in the limit where the $\phi_i \rightarrow 0$, which in practice requires events to be uncommon: this is a rare disease assumption. Further details of this derivation have been given previously. (1, 3)

Several modelling approaches are available to represent the age (or time) effect $\psi(t)$. For simplicity, we shall refer to the $\psi(t)$ as age effects from now on, though $\psi(t)$ could alternatively represent calendar time, typically used to capture seasonal variation, or time since first diagnosis. The simplest models age with piecewise constant functions, (1) which we refer to as the standard SCCS model. Alternatives include the semiparametric model where $\psi(t)$ is left unspecified, (3) fractional polynomial methods in which age (or time) effects are piecewise constant but smoothed by fractional polynomials, (15) and fully smoothed spline-based models in which age (or time) is represented by a linear combination of M-splines.

(16, 18) Additionally, alternative methods are available for modelling the exposure effects using splines. (17, 18) However, the methods proposed in this paper are based on the standard SCCS model with piecewise constant functions representing both age and exposure. An example of an individual's observation period split into age and exposure groups as in a standard SCCS set up is shown in figure 1 .




FIGURE 1 Example standard SCCS observation period for one case with two exposures and four age groups.

For the parameter β to retain the same interpretation as in the cohort model, two key assumptions are required, namely that the observation period $(a_i, b_i]$ and the exposure history x_i are unaffected by the occurrence and timing of events. These are non-trivial assumptions, specific to the SCCS method. We consider them in turn in the following sections.

3 | EVENT-DEPENDENT EXPOSURES

The impact of event-dependent exposures on the log relative incidence parameter β is predictable. If occurrence of an event increases the chance of a subsequent exposure, we would expect β to be biased downward. Conversely, if occurrence of an event reduces the chance of a subsequent exposure, β would be biased upwards. (12) Situations like this commonly arise: for example, the adverse event may be a contra-indication to the exposure, occurrence of an adverse event may result in delayed vaccination; or a period of hospitalization may alter the treatment regime.

3.1 | A test for event-dependent exposures

Suppose that the impact of an event on the exposure process occurs in a time interval of duration τ after the event. Then, if events precipitate exposures, we would expect an excess of exposures in the interval $(t, t + \tau]$ where t is the time of event. Equivalently, we would expect an excess of events in the interval $[v - \tau, v)$ where v is the time of the exposure onset. Similarly, if events inhibit exposures in an interval of length τ after the event, we would expect a dearth of exposures in $(t, t + \tau]$, and equivalently a dearth of events in $[v - \tau, v)$.

This can easily be tested by extending the SCCS model to include a pre-exposure 'risk' period of duration τ prior to each exposure. Overlaps with other risk intervals can be dealt with either by giving precedence to the most recently started risk period, or by giving precedence to the pre-exposure risk period, or by treating the overlap as a new risk period with its own parameter. The SCCS model is fitted in the usual way, with an extra parameter θ for the log relative incidence in the pre-exposure period.

In certain circumstances, the parameter θ has a direct interpretation as the log relative incidence of exposure in a time interval of duration τ after an event. This has been elucidated previously. (3) Briefly, when events are rare so that the probability of two events occurring within a time interval of duration τ is small, the SCCS likelihood incorporating the effect of event-dependent exposures is approximately

$$L \simeq \prod_{i=1}^N \frac{\prod_{j=1}^{n_i} \psi^*(t_{ij}) \exp\{x_i(t_{ij})^T \beta + z_i(t_{ij})\theta\}}{\left(\int_{a_i}^{b_i} \psi^*(t) \exp\{x_i(t)^T \beta + z_i(t)\theta\} dt\right)^{n_i}}$$

where $z_i(t)$ is the number of exposures in $(t, t + \tau)$, and $\psi^*(t) = \psi(t)$ when $\theta = 0$. When τ is small, the number of exposures in an interval of length τ is typically 0 or 1. Then the covariate $z_i(t)$ is equivalent to specifying a pre-exposure 'risk' period of length τ . Thus, at least for small values τ , the parameter θ can be interpreted as the log relative incidence of an exposure in an interval τ after an event.

3.2 | Graphical representations

We highlight several graphs to explore event-dependence of exposures, the first of which has been used by several authors. (19, 20) This plots the distribution of the interval $t_{ij} - v_{ik}$ between exposure and event times, for each combination of events j and exposures k within case i , where $j = 1, \dots, n_i, k = 1, \dots, m_i$, and $i = 1, \dots, N$. Thus, this is a histogram of the values $t_{ij} - v_{ik}$, which we refer to as the exposure-centered interval plot. Positive values correspond to events after an exposure, negative values to events before an exposure. Only cases in individuals with at least

one exposure are included in this graph. If multiple exposures need to be differentiated – as may be the case with different doses of a vaccine – this can be achieved by different shadings on the same exposure-centered interval plot, or by using separate plots for the different exposures. If a peak or trough in events prior to exposure is apparent, the presence of short-term event-dependent exposures is indicated.

The observation periods and ages at exposure typically vary between cases. Consequently, the numbers of exposed events under observation can vary substantially in different parts of the interval plot. Therefore a further plot, indicating the total number under observation at each time, may help interpretation. Since an event time for case i can take any value in $(a_i, b_i]$, the interval $t_{ij} - v_{ik}$ can take any value in $(a_i - v_{ik}, b_i - v_{ik}]$. Let $I_{ik}(t)$ be the indicator function for the interval $(a_i - v_{ik}, b_i - v_{ik}]$, taking the value 1 if $t \in (a_i - v_{ik}, b_i - v_{ik}]$ and 0 otherwise. Then let

$$S(t) = \sum_{i=1}^N \sum_{k=1}^{m_i} n_i \times I_{ik}(t).$$

Now plot $S(t)$ against t . For each t , $S(t)$ is the total number of exposed events under observation at time t after exposure, given the constraints imposed by the observation periods and exposures. The shape of the plot represents the expected shape of the exposure-centered interval plot in the absence of any age effect, exposure effect and event-dependent exposure. It may be useful to help interpret whether there is an excess or dearth of events pre-exposure in the exposure-centered interval plot.

Another useful graph is obtained by plotting the estimated parameter θ against τ , for a range of values of τ , $\tau_1 < \dots < \tau_k$. This plot, with confidence limits on the estimates, helps to indicate the extent of possible event-dependence of exposures. To visualise the robustness of the estimates of the parameter of primary interest, β , the estimated values of β obtained with pre-exposure periods using different values of τ can also be plotted.

4 | EVENT-DEPENDENT OBSERVATION PERIODS

Unlike event-dependent exposures, the direction of bias if observation periods are event-dependent is generally not easily predictable. (14) We consider the situation in which observation may be terminated early, early termination being potentially influenced by the event history. To this end we need some more notation.

Let b_i^* denote the planned end of observation of individual i , as determined by age and calendar time boundaries on case ascertainment. The actual end of observation, b_i , may be less than b_i^* . Thus, $b_i = \min\{b_i, c_i\}$, where c_i denotes the age at which the time line of individual i is censored. If $c_i < b_i^*$, then $b_i = c_i$ and the planned observation period $(a_i, b_i^*]$ is censored. If $c_i \geq b_i^*$, then $b_i = b_i^*$ and the observation period is uncensored. Censoring could be due to death when events are potentially fatal, or to any other event-dependent cause of end of follow-up. Event-dependent observation periods arise when the occurrence of an event increases the probability of censoring sufficiently that $c_i < b_i^*$. This may (or may not) induce bias in the parameter β of primary interest. In contrast, end of follow-up due to causes independent of the event will not bias β .

4.1 | A robustness test for event-dependent observation periods when censoring information is available

Suppose first that, for each individual i , we have information on both the actual end of observation b_i and on whether observation was censored or not, that is, on the indicator $I_i = I(c_i < b_i^*)$, equal to 1 if $c_i < b_i^*$ and to 0 otherwise. Thus cases can be classified in two groups: those observed over the full period $(a_i, b_i^*]$ (we shall refer to these as the uncensored cases) and those censored at $c_i < b_i^*$ (the censored cases). If almost all cases are in the first category, the intensity of event-dependent censoring is low and can be ignored, and a standard SCCS analysis will be valid. This suggests, heuristically, a simple sensitivity analysis: compare the results of a standard analysis on all cases, with the results of a standard SCCS analysis on just the fully observed cases. (4) If the estimates of the exposure-associated parameters β differ, event-dependent observation periods may be a problem.

More formally, this may be investigated by fitting, first, the standard SCCS model to all cases, and then fitting a model including the interaction of the exposure effect with the indicator variable I_i . The interaction between the censoring indicator and age is not included, because censoring may affect age effects even when there is little impact on exposure effects, which are of primary interest. The two models are nested, so a likelihood ratio test can be used to test the interaction. If significant, this suggests that event-dependence of observation periods may affect the estimated exposure effect. More importantly, however, the estimates of the exposure effect in the two groups can be compared to assess robustness to failure of the assumption that observation periods are not event-dependent.

A justification of this procedure is as follows. Let h_i^t denote the event history of individual i to time t in $(a_i, b_i]$. Let $\mu(s|h_i^s)$ denote the censoring hazard at time s , given the event history in $(a_i, b_i]$ to time s . At age b_i the event history for individual i is $h_i^{b_i} = \{t_{i1}, \dots, t_{in_i}\}$. If there are no events by age s , then $h_i^s = \emptyset$ and $\mu(s|h_i^s) = \mu(s)$ does not depend on the event history in $(a_i, b_i]$. The adjusted SCCS likelihood contribution of case i may be written

$$L_i = \frac{\prod_{j=1}^{n_i} [\psi(t_{ij}) \exp\{x_i(t_{ij})^T \beta\}] \mu(b_i|h_i^{b_i})^{I_i} S(b_i|h_i^{b_i})}{\int_{a_i}^{b_i} \int_{s_1}^{b_i} \dots \int_{s_{n_i}}^{b_i} \prod_{j=1}^{n_i} [\psi(s_j) \exp\{x_i(s_j)^T \beta\}] \mu(b_i|s_1, \dots, s_{n_i})^{I_i} S(b_i|s_1, \dots, s_{n_i}) ds_{n_i} \dots ds_2 ds_1},$$

where

$$S(y|h^y) = \exp\left(-\int_{a_i}^y \mu(s|h^s)ds\right),$$

and

$$S(y|s_1, \dots, s_{n_i}) = \exp\left(-\int_{a_i}^y \mu(s|\{s_1, \dots, s_{n_i}\} \cap (a_i, s])ds\right).$$

In general, the likelihood contribution L_i differs according to the value of I_i . For example, if case i is censored, so that $I_i = 1$, then L_i involves the term $\mu(b_i|h_i^{b_i})$ in its numerator, but not if case i is uncensored ($I_i = 0$). This difference is what the proposed test is designed to pick up. On the other hand, if $\mu(s|h_i^s) = \mu(s)$ for all s in $(a_i, b_i]$, then the difference between the contributions for censored and uncensored cases disappears. The denominator of L_i becomes

$$\begin{aligned} \int_{a_i}^{b_i} \int_{s_1}^{b_i} \dots \int_{s_{n_i}}^{b_i} \prod_{j=1}^{n_i} [\psi(s_j) \exp\{x_i(s_j)^T \beta\}] \mu(b_i|s_1, \dots, s_{n_i})^{I_i} S(b_i|s_1, \dots, s_{n_i}) ds_{n_i} \dots ds_2 ds_1 \\ = (n_i!)^{-1} \times \left(\int_{a_i}^{b_i} \psi(s) \exp\{x_i(s)^T \beta\} ds \right)^{n_i} \end{aligned}$$

and the likelihood contribution reduces to that of a standard SCCS likelihood with age effect $\psi(t)$ and exposure effect β .

If the estimates of β differ substantially between censored and uncensored cases, or according to whether censored cases are included or excluded, more detailed investigations and an adjusted SCCS method may be required. (14)

4.2 | A test for the identification of event-dependent observation periods when censoring information is not available

Suppose now that information on whether or not the observation period has been censored is unavailable. In this case, $I(c_i < b_i)$ is not known, and the test described above cannot be done – and nor can an adjusted SCCS method be applied. (14) However, another approach can be used to identify if event-dependent observation periods are present.

If events precipitate the end of the observation period, we would expect to observe a higher than usual rate of occurrence of events in the period just before the end of observation. A simple test for this is to include an additional, terminal risk period, of duration κ say, just before the end of observation. Thus, we include in the model the additional, terminal risk period $(b_i - \kappa, b_i]$. It makes sense to try different values of κ .

Note however that the presence of event-dependent censoring of observation periods does not imply that the exposure parameter of interest, β , is biased. Thus, suppose that

$$\mu(t|h_i^t) = \mu(t) + \|h_i^t\|\theta,$$

where $\|h_i^t\|$ is the number of events for individual i that have occurred prior to t in $(a_i, b_i]$. In other words, the censoring hazard is $\mu(t)$ before the first event, and $\mu(t) + k\theta$ between the k th and $(k+1)$ th events. It then follows that

$$\begin{aligned} \mu(b_i|h_i^{b_i}) &= \mu(b_i|s_1, \dots, s_{n_i}) = \mu_i(b_i) + n_i\theta, \\ S(b_i|h_i^{b_i}) &= e^{-n_i b_i \theta} \times \prod_{j=1}^{n_i} e^{\theta t_{ij}} \times \exp\left(-\int_{a_i}^{b_i} \mu(s)ds\right), \end{aligned}$$

and

$$S(b_i|s_1, \dots, s_{n_i}) = e^{-n_i b_i \theta} \times \prod_{j=1}^{n_i} e^{\theta s_j} \times \exp\left(-\int_{a_i}^{b_i} \mu(s)ds\right).$$

The adjusted SCCS likelihood contribution for case i is then identical whether the case is censored or uncensored, and equal to

$$L_i = \frac{\prod_{j=1}^{n_i} \psi^*(t_{ij}) \exp\{x_i(t_{ij})^T \beta\}}{\left(\int_{a_i}^{b_i} \psi^*(s) \exp\{x_i(s)^T \beta\} ds\right)^{n_i}},$$

where $\psi^*(t) = \psi(t) \exp(\theta t)$. Thus L_i is a standard SCCS likelihood contribution and will yield unbiased estimates of β , though the estimate of the age effect is likely to be seriously upwardly biased (if $\theta > 0$), especially at older ages.

This example serves to emphasize the key point that failure of the assumption that observation periods are event-dependent does not imply that the exposure effect is biased. The test proposed in this section is therefore likely to be useful only for ruling out the presence of bias due to event-dependent observation periods.

4.3 | Graphical representations

A key graphical display when it is suspected that event-dependent censoring of observation periods is present is the histogram of the times from event to end of observation. Ideally, two histograms can be produced according to whether observation times are censored. Thus, a first histogram could include just the values $c_i - t_i$ for individuals with censored time of observation (that is, $c_i < b_i$), and a second could include just the values $b_i^* - t_i$ for individuals with $c_i > b_i^*$. If a spike is apparent close to zero in the censored data histogram, the presence of event-dependent observation periods is indicated. Alternatively, in the absence of censoring information, the histogram could include all values $b_i - t_i$, though this is more difficult to interpret. Such histograms indicate nothing about the robustness of estimates.

5 | SIMULATIONS

We investigate the various methods proposed to test whether exposures depend on events and whether events affect the end of observation periods.

5.1 | Event-dependent exposures

5.1.1 | Method

To investigate the performance of the tests proposed to identify existence of event dependent exposures, we began by fixing the number of cases to 100 and 200, with each case having a fixed start and end of observation period at 0 and 300 days. Ages at start of exposure were then simulated from a uniform distribution. All exposure-related risk periods were assigned a length of 10 days. The numbers of exposures per case considered were 1 and 5 (generated so as not to overlap).

Each individual's observation period was divided into intervals based on the one or five exposure risk periods and age group cut points at 100 and 200 days. Then on the basis of the given true age and exposure related relative incidence values the interval in which the event occurred was determined using a multinomial distribution. Three different true exposure-related relative incidence values of 1, 2, and 5 were investigated. Age effects in the three age groups had relative incidence values of 1, 2 and 3. An age at event within the interval where the event occurred was simulated from a uniform distribution. Once a complete data set was obtained this way, it was then modified in three ways, to delay, remove and add exposures after the event day, each with probability 0.5 and 0.9. For the delay scenario, exposures beginning within an interval of $\tau = 15$ and $\tau = 30$ after the event were delayed by τ days. Also a delay scenario that we label $\tau = \text{permanent}$ was defined, where all exposures beginning between the event day and the end of the observation period were delayed by 15 days. For the remove exposures scenario, exposures beginning within an interval τ after the event day were removed. Values of $\tau = 15$, $\tau = 30$, were assigned and also $\tau = \text{permanent}$, where any exposure beginning between the event day and the end of observation period were removed. For the add exposures scenario, the complete data were modified by adding exactly one extra exposure in an interval of length τ from age at event. The new exposure start times were generated from a uniform distribution in the interval between age at event and age at event + τ , where $\tau = 15$, $\tau = 30$ and $\tau = \text{permanent}$ (i.e the interval between age at event and age at end of observation period).

Both the complete and modified data were then analyzed by SCCS models with and without pre-exposure periods of length 15, 30 and 60 days. The three ways of handling overlapping risk periods explained in 3.1 were applied. To recap, these were to give the exposure risk period precedence, to give the pre-exposure period precedence and to create a new period where overlaps occur.

5.1.2 | Results and evaluation

The bias in the log exposure related relative incidences β due to event-dependent exposures was calculated. The bias was defined as the difference between the median($\hat{\beta}$) obtained when a model without pre-exposure risk periods was fitted to the modified data (where events affect subsequent exposures) and median($\hat{\beta}$) for the complete unmodified data (the unbiased gold standard), where ($\hat{\beta}$) is the log of the estimated exposure related relative incidence. The performances of the proposed tests to identify event-dependent exposures were evaluated by calculating the power of the tests. The power was calculated as the proportion of significant pre-exposure relative incidence estimates at the 5% level of significance. The bias due to event dependent exposures and the power of the tests proposed in Section 3.1 are presented in Table 1 for $n = 200$, 5 exposures and a true relative incidence of 2 (similar results were obtained for $RI = 1, 5$, and slightly reduced power for $n = 100$ as expected).

As shown in Table 1, for the complete unmodified data (labeled Prob=0), the power values, as expected, are close to 5% for the methods with exposure taking precedence and with a new overlap period. However, for the method where pre-exposure takes precedence over exposure risk period the power values are higher than expected, hence this method cannot be recommended for use. There is little difference between power

TABLE 1 Simulation results for true exposure relative incidence = 2 and 5 initial exposures per case. The first column gives the true value of τ simulated. The second column 'Prob' gives the probability that exposures were modified following an event (Prob=0 indicates unmodified data). The third column 'Bias' gives the bias of the simulated data. The final 9 columns give the power of the three tests with different ways of handling overlaps: exposure period takes precedence, pre-exposure period takes precedence, or new overlap period, with different values of τ fitted.

True τ	Prob	Bias	Expo takes precedence			Pre-expo takes precedence			New overlap period		
			$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$
Delay exposures											
	0	0	0.05	0.05	0.05	0.08	0.11	0.12	0.05	0.04	0.04
15	0.5	0.003	0.77	0.05	0.05	0.72	0.08	0.11	0.79	0.05	0.04
	0.9	0.007	0.97	0.06	0.05	0.97	0.06	0.11	0.97	0.06	0.04
30	0.5	0.019	0.75	0.86	0.05	0.7	0.81	0.1	0.77	0.87	0.05
	0.9	0.036	0.97	0.97	0.06	0.97	0.97	0.08	0.97	0.97	0.06
Perm (shift 15)	0.5	0.049	0.74	0.57	0.32	0.69	0.44	0.17	0.76	0.62	0.38
	0.9	0.063	0.97	0.95	0.68	0.97	0.92	0.45	0.97	0.98	0.78
Remove exposures											
	0	0	0.04	0.06	0.05	0.07	0.11	0.12	0.05	0.05	0.05
15	0.5	0.029	0.72	0.6	0.36	0.68	0.44	0.19	0.72	0.59	0.36
	0.9	0.062	0.93	1	0.83	0.94	0.99	0.56	0.93	1	0.83
30	0.5	0.063	0.7	0.83	0.64	0.65	0.76	0.53	0.7	0.83	0.64
	0.9	0.127	0.93	0.93	1	0.94	0.94	1	0.93	0.93	1
Perm	0.5	0.334	0.39	0.58	0.63	0.32	0.47	0.54	0.39	0.58	0.65
	0.9	0.866	0.93	0.93	0.93	0.93	0.94	0.94	0.93	0.93	0.93
Add exposure											
	0	0	0.04	0.04	0.04	0.07	0.11	0.11	0.04	0.04	0.05
15	0.5	-0.107	1	1	0.99	1	1	1	1	1	1
	0.9	-0.196	1	1	0.99	1	1	0.99	1	1	1
30	0.5	-0.103	0.99	1	0.99	1	1	1	1	1	1
	0.9	-0.189	1	1	0.99	1	1	0.99	1	1	1
Perm	0.5	-0.116	0.53	0.57	0.59	0.81	0.87	0.87	0.71	0.77	0.82
	0.9	-0.216	0.94	0.97	0.99	0.99	1	1	0.99	1	1

achieved when exposure takes precedence and when a new overlap period is given, except when an exposure is added at any time after the event (labeled perm), where the method with a new overlap period has higher power. Hence, we favor the method with a new overlap period.

The absolute bias is very small when exposures are delayed by true $\tau = 15, 30$ days (within a period of τ days) and the test has only the nominal 5% power when the fitted τ is longer than the true τ (as the total number of exposures will be exactly the same, only timings are different). A permanent delay represents a shift of any post-event exposure by 15 days, hence it is unsurprising that power drops when fitted $\tau = 60$.

When exposures are removed the absolute bias increases noticeably as true τ increases, as more exposures are removed (multiple exposures per case can be removed in this scenario). Power is consistently high when exposures are removed with probability 0.9. When exposures are removed with probability 0.5, power achieved is highest when the fitted τ is well matched to the true τ .

When exposures are added the test has very high power to detect a difference. Power is below 1 when an exposure is added at any time (perm), and increases with increasing τ fitted.

To demonstrate how bias is affected by including a pre-exposure risk period in an SCCS model, we also calculated the bias as a difference between the median($\hat{\beta}$) estimate obtained with a pre-exposure risk period included in the model for the modified data and the median($\hat{\beta}$) estimate for the complete unmodified data (with no pre-exposure period). These results are presented in Table 2 .

From Table 2 it can be seen that, as expected, when events reduce the probability of being exposed (delay, remove) the exposure relative incidence is overestimated. Similarly, when events increase the chance of being exposed the direction of bias, as expected, is downwards. Inclusion of a pre-exposure risk periods usually adjusts the bias in the correct direction, but there are exceptions and sometimes the bias is over-corrected. For example, when exposures are added with probability 0.9, almost all events land in the pre-exposure period and the way overlaps are dealt with makes a considerable difference. When exposure takes precedence, estimates can become massively over-corrected, and when pre-exposure takes precedence, bias can become considerably worse. A better balance is struck when a new period is assigned. Hence, overall, we favor the method where a new period is given for overlaps.

When exposures are added or removed at any time after the event ('Perm'), it may not, in theory, be possible to achieve complete bias correction (an alternative SCCS method should be used). The results for the new overlap period method are consistent with this, achieving only partial bias correction when exposures are added or removed permanently (true $\tau = \text{Perm}$). Bias correction ought to be optimized when the fitted τ matches the period affected well. The total period affected is $\tau + 10$, as the length of the exposure risk period is always 10. Since these exact values were not fitted this is a little unclear from the results, and how the slight bias created in age effects by the addition, removal or shift of exposures is

TABLE 2 Simulation results for true exposure relative incidence = 2 and 5 initial exposures per case. The third column labelled 'Nopreexp' gives the bias without any pre-exposure period fitted. The final 9 columns give the bias in estimating exposure relative incidence with pre-exposure risk periods for the three different ways of handling overlaps and values of τ fitted.

True τ	Prob	Nopreexp	Expo takes precedence			Pre-expo takes precedence			New overlap period		
			$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$
	0	0	0.006	0.007	0.012	0.011	0.019	0.018	0.005	0.004	0.004
Delay exposures											
15	0.5	0.003	-0.147	-0.012	-0.004	-0.078	0	0.014	-0.08	-0.003	0.001
	0.9	0.007	-0.252	-0.019	-0.002	-0.127	-0.003	0.016	-0.135	-0.002	0.004
30	0.5	0.019	-0.125	-0.29	0.004	-0.051	-0.116	0.014	-0.055	-0.104	0.019
	0.9	0.036	-0.209	-0.491	0.009	-0.08	-0.181	0.021	-0.09	-0.189	0.033
Perm (Shift 15)	0.5	0.049	-0.096	-0.104	-0.107	-0.032	-0.008	0.006	-0.029	-0.012	0.015
	0.9	0.063	-0.185	-0.209	-0.207	-0.066	-0.012	-0.002	-0.07	-0.039	0.004
Remove exposures											
15	0.5	0.029	-0.107	-0.127	-0.124	-0.042	-0.042	-0.04	-0.044	-0.048	-0.051
	0.9	0.062	-0.2	-0.237	-0.235	-0.067	-0.082	-0.081	-0.067	-0.086	-0.093
30	0.5	0.063	-0.066	-0.207	-0.212	-0.002	-0.03	-0.02	-0.003	-0.034	-0.031
	0.9	0.127	-0.119	-0.386	-0.397	0.01	-0.056	-0.05	0.011	-0.056	-0.057
Perm	0.5	0.334	0.268	0.189	0.067	0.301	0.297	0.291	0.301	0.292	0.285
	0.9	0.866	0.677	0.463	0.084	0.799	0.78	0.717	0.799	0.78	0.717
Add exposure											
15	0.5	-0.107	0.588	0.586	0.592	-0.056	-0.036	-0.009	0.16	0.042	-0.027
	0.9	-0.196	2.114	2.11	2.149	-0.16	-0.089	-0.039	0.245	0.032	-0.087
30	0.5	-0.103	0.171	0.586	0.586	-0.071	-0.037	-0.011	0.018	0.048	-0.023
	0.9	-0.189	0.376	2.106	2.133	-0.151	-0.084	-0.049	0.003	0.039	-0.081
Perm	0.5	-0.116	0.003	0.086	0.226	-0.099	-0.089	-0.101	-0.06	-0.06	-0.061
	0.9	-0.216	0.015	0.195	0.625	-0.2	-0.203	-0.199	-0.124	-0.118	-0.118

also unclear. Considering the results with the new overlap period for the add and delay scenarios, fitting τ too long (e.g. fitting $\tau = 60$ when true $\tau = 15, 30$) always under-corrects the bias, while fitting τ too short always over-corrects the bias. For the remove scenario, biases tend to be over-corrected. Over-correction is less of a concern for the delay and remove scenarios, than for the add scenario where estimates are biased upwards.

As shown in Table 2, the convention used to handle overlaps between risk periods and pre-exposure periods can have a substantial impact on the results, the best being achieved when an additional parameter is used for such overlaps. In many applications, there is at most one exposure, and so the issue of overlaps is much less important – in fact it only arises when exposures are added after an event. Table 3 shows the bias in such circumstances. As expected, if each case experiences at most one exposure, the convention used for overlaps is immaterial when post-event exposures are removed or delayed.

5.2 | Event-dependent observation periods

5.2.1 | Method

To explore the tests for event-dependent observation periods simulations were set up with 100 and 200 cases, all with observation periods of length 300 days and age group cut offs at 100 and 200 days. Only one exposure risk period was generated. Firstly the timings of the exposure risk periods were simulated, five scenarios were considered: (1) a uniform distribution of exposure start times over the observation period with exposure risk periods lasting 30 days (except where they overlap the ends), (2) a linear increase in exposure start times over the observation period with exposure risk periods lasting 30 days, (3) a linear decrease in exposure start times over the observation period with exposure risk periods lasting 30 days, (4) a linear increase in exposure start times with indefinite exposure risk periods that end with the end of observation, and (5) all exposure start times set at 0 with a linear decrease in exposure end times over the observation period. We understand that bias from event-dependent observation periods will tend to be in one direction when trends in the timing of exposures over the observation period are present and the last two scenarios where all exposures are loaded at the end or at the beginning of the observation period were aimed at generating considerable bias. (Conversely exposure risk periods evenly spread over the observation period should give rise to minimal bias, this will be especially true if there are many intermittent exposures hence we did not consider multiple exposures for these simulations.)

Event days were simulated as for event-dependent exposures with exposure-related relative incidences $\exp(\beta) = 1, 2$ and 5, and age-related relative incidences $\exp(\alpha) = 3, 2$ and 1, in the first, second and last age groups respectively. This decreasing age effect was chosen to create more bias; if events tend to occur toward the beginning of the study period then there is a greater proportion of observation period to lose if observation periods are event-dependent.

TABLE 3 Simulation results for true exposure relative incidence = 2 and 1 initial exposure per case. The third column labelled 'Nopreexp' gives the bias without any pre-exposure period fitted. The final 9 columns give the bias in estimating exposure relative incidence with pre-exposure risk periods for the three different ways of handling overlaps and values of τ fitted.

True τ	Prob	Nopreexp	Expo takes precedence			Pre-expo takes precedence			New overlap period		
			$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$	$\tau=15$	$\tau=30$	$\tau=60$
	0	0	-0.006	-0.009	-0.014	-0.006	-0.009	-0.014	-0.006	-0.009	-0.014
Delay exposure											
15	0.5	-0.001	-0.046	-0.007	-0.008	-0.046	-0.007	-0.008	-0.046	-0.007	-0.008
	0.9	-0.002	-0.057	-0.011	-0.01	-0.057	-0.011	-0.01	-0.057	-0.011	-0.01
30	0.5	0	-0.041	-0.069	-0.01	-0.041	-0.069	-0.01	-0.041	-0.069	-0.01
	0.9	-0.001	-0.055	-0.114	-0.014	-0.055	-0.114	-0.014	-0.055	-0.114	-0.014
Perm (Shift 15)	0.5	-0.014	-0.064	-0.057	-0.051	-0.064	-0.057	-0.051	-0.064	-0.057	-0.051
	0.9	-0.021	-0.077	-0.083	-0.084	-0.077	-0.083	-0.084	-0.077	-0.083	-0.084
Remove exposure											
15	0.5	0.011	-0.002	-0.007	-0.008	-0.002	-0.007	-0.008	-0.002	-0.007	-0.008
	0.9	0.05	-0.002	-0.007	-0.012	-0.002	-0.007	-0.012	-0.002	-0.007	-0.012
30	0.5	0.045	0.015	-0.006	-0.01	0.015	-0.006	-0.01	0.015	-0.006	-0.01
	0.9	0.098	0.05	-0.006	-0.016	0.05	-0.006	-0.016	0.05	-0.006	-0.016
Perm	0.5	0.295	0.268	0.245	0.188	0.268	0.245	0.188	0.268	0.245	0.188
	0.9	0.574	0.527	0.468	0.355	0.527	0.468	0.355	0.527	0.468	0.355
Add exposure											
15	0.5	-0.459	0.362	0.348	0.347	-0.313	-0.289	-0.231	-0.012	-0.033	-0.057
	0.9	-0.722	1.766	1.776	1.903	-0.609	-0.522	-0.32	-0.137	-0.231	-0.314
30	0.5	-0.446	-0.124	0.358	0.344	-0.386	-0.281	-0.234	-0.248	-0.012	-0.043
	0.9	-0.705	-0.096	1.761	1.827	-0.643	-0.542	-0.387	-0.436	-0.165	-0.265
Perm	0.5	-0.502	-0.367	-0.29	-0.179	-0.436	-0.428	-0.418	-0.408	-0.36	-0.294
	0.9	-0.781	-0.53	-0.361	-0.068	-0.677	-0.657	-0.672	-0.617	-0.533	-0.42

SCCS models were fit to this complete data, estimates for $\hat{\beta}$ gained may be considered the gold standard, allowing for random variation. Then a 'censoring day' (or early end of observation period) was simulated for each individual. Censoring days were generated under two scenarios: uniform between the event day and the planned end of observation, or at a linearly increasing rate after the event day up to a maximum of 30 days. The observation period ended at the censoring day for fixed proportions of randomly selected individuals: 80%, 50%, 20%, 10% or 5%. SCCS models were fit to these five partially censored data sets. The test for robustness described in subsection 4.1 was applied. The test for identification of event-dependent observation periods in 4.2 was also applied, with κ set at 15, 30, 60 and 90 days.

5.2.2 | Results and evaluation

To evaluate the robustness test, we compared the bias, defined as the median($\hat{\beta}$) for the partially censored data - median($\hat{\beta}$) for the complete data, with the power, defined as the percentage of test rejections at the 5% level of significance. Given the many simulation scenarios (60), results were summarized graphically in Figure 2. The first thing to note on Figure 2 is the differing x-axis scales, which shows the increasing extent of the bias as the proportion of censored individuals increases. The results for 5% of observation periods ending early are not shown, but as expected, the bias was minimal. The model including the interaction term sometimes ran into computational problems when only 5 cases (5% censored, $n = 100$) belonged to the censored group, for example when none of the 5 cases experienced an exposure due to censoring. There are close to the expected 5% test rejections (at the 5% level of significance) when the bias is very close to 0, and the percentage of test rejections rises as the observed bias moves further from 0. A null simulation, for which cases were labelled as censored when in fact observation periods were not censored, also produced around 5% power (results not shown). The bias and number of test rejections depended heavily on the scenario simulated, for example greater power was achieved with a sample size of $n = 200$, in comparison with $n = 100$. Overall, the robustness test appears to work as intended.

Figure 2 also demonstrates how trends in exposure over the observation period influence the direction of the bias: if individuals tend to experience exposure more toward the end of the observation period, relative incidences will be biased upward; whereas if cases tend to experience exposure more toward the beginning of the observation period, relative incidences ($\text{exp } \beta$) will be biased toward 0. The most extreme biases resulted when exposure risk periods always ran to the end of observation (biased upward) or always ran from 0 (biased downward).

The test for identification of event-dependent observation periods was evaluated by plotting the power against the proportion of censored cases, defined as the percentage of test rejections at the 5% level of significance. Results for the scenarios where the censoring day was uniformly distributed between the event day and 300 (the planned end of observation) are given in Figure 3. The proportion of test rejections rises as the proportion censored rises, thus the test does appear to reasonably identify when event-dependent observation periods are present, even for this




FIGURE 2 Test for robustness: power plotted against bias. The four panels show results for different proportions of individuals censored.

scenario, where the timing of censoring is random. For this scenario, the test has greatest power when $\kappa = 15$ or $\kappa = 30$. Note that the uncensored observation period length was 300, so these values of κ correspond to 5% and 10% of the length of the original observation period.

For the scenarios where observation periods were censored within 30 days of the event (results not shown), unsurprisingly, the test had greatest power when $\kappa = 30$. The power was 100% (or very close) when the proportion of cases censored was 20% or more. Even with 5% censoring, the power ranged from 16% to 50%.

We also looked into whether the test for identification of event-dependent observation periods had any power to detect bias, evaluating results as for the test for robustness. We concluded that this test should not be used to detect bias and its use should be limited to identifying whether or not event-dependent observation periods are present in the absence of this information from the data source.

6 | EXAMPLES

We present two examples. The first relates to potential event-dependent exposures, the second to event-dependent observation periods.

6.1 | Measles, mumps and rubella vaccine and blood disorders

These data relate to the potential association between receipt of measles, mumps and rubella (MMR) vaccine and the occurrence of idiopathic thrombocytopenic purpura (ITP), a bleeding disorder, in children. The study was previously reported with a slightly different data set from that used here. (19) The events are admissions to hospital for ITP.

The observation period included all days between 1st October 1991 and 30th September 1994, and between days 366 and 730 of age, inclusive. The data comprise 44 events in 35 children: five children were admitted twice and one child was admitted five times for ITP within the observation




FIGURE 3 Test for identification of event-dependent observation periods: proportion censored plotted against power. The four panels show results for different lengths of κ . Only results with censoring days uniformly distributed between the event day and 300 are shown.

period. Time is measured in discrete days. To allow for variation in the incidence of ITP with age, the analysis uses six age groups: 366-426 days, 427-487 days, 488-548 days, 549-609 days, 610-670 days and 671-730 days of age. In the present analysis, we shall use a single post-MMR risk period: 0-42 days after MMR; thus the day of vaccination (day 0 after MMR) is included in the risk period. This SCCS model gives a relative incidence of 3.23, with 95% confidence interval (1.53, 6.79). Thus there is evidence of a positive association between MMR and ITP in the six weeks following vaccination.

We now investigate the possibility that a hospital admission for ITP might affect subsequent MMR vaccination. While there is little reason to suspect that MMR vaccination would never take place after an ITP admission, it is possible – indeed likely – that such an admission would delay subsequent vaccination. The first step is to plot the intervals between MMR vaccination and ITP onset. This is shown in Figure 4. The panel on the left is drawn using 25-day intervals. There is a clear peak (14 events) in the 50-day period just after 0; in no other 25-day period does the count exceed 4. The panel on the right shows the number of exposed cases under observation. Note that a case with k events is counted k times so that at interval 0, the number under observation is always the number of events in exposed cases; 35 events in 31 exposed cases here (9 events arose in 4 unexposed cases). Taken together, the two graphs do not indicate an evident shortage of events before MMR, which is what would be expected if occurrence of an event were to substantially reduce the likelihood of MMR vaccination. However, there are no events in the 25-day period immediately before MMR, which might indicate a short-term delay in MMR vaccination following an event. On the other hand, such gaps occur at other intervals as well.

This may be investigated further by fitting SCCS models with pre-exposure periods of length τ , for various values of τ . We tried $\tau = 14, 30, 61, 122, 183$ days – that is, 2 weeks, and 1, 2, 4 and 6 months. The results are shown in Figure 5. The panel on the left shows that the relative incidence associated with the pre-exposure ‘risk’ intervals are less than 1, indicating that occurrence of an event inhibits subsequent MMR vaccination during that interval. For $\tau = 14$ and 30 days, the relative incidence is zero as there are no events in those intervals. For $\tau = 61$ it is 0.497, with




FIGURE 4 Left: number of events in exposed cases by time since MMR. Right: number of exposed cases under observation by time since MMR.

95% CI (0.133, 1.86). Similarly, all subsequently confidence intervals include 1. The panel on the right of Figure 5 show the relative incidence associated with MMR vaccination (in the risk period 0–42 days post MMR), when pre-MMR ‘risk’ intervals are included. This suggests that the results are robust to inclusion of a pre-MMR ‘risk’ interval. When $\tau = 0$ we have the original relative incidence of 3.23. As τ increases, the relative incidence declines slightly. For $\tau = 61$ days, it is 2.72, 95% CI (1.22, 6.09). The confidence intervals all lie wholly above 1.

These findings provide some evidence that hospital admission for ITP delays subsequent MMR vaccination for some weeks. The pre-exposure effect is statistically significant for $\tau = 30$ days (likelihood ratio test statistic 6.11 on 1 degree of freedom, $p = 0.013$) but not for the other values of τ . Most importantly, it does not have a major impact on the results, or on the conclusion that MMR is positively associated with ITP.

6.2 | Antipsychotics and stroke

These data relate to the potential association between receipt of antipsychotics and stroke. The original study carried out used a standard self-controlled case series analysis. (21) The original analysis found a significant association between antipsychotics and stroke in both patients with prior dementia (relative incidence 3.50, 95% CI 2.97 to 4.12), and patients without prior dementia (RI 1.41, 95% CI 1.29 to 1.55). The data were later reanalysed to take account of potential event-dependent observation periods owing to the high short-term mortality associated with stroke. (14) The reanalysis confirmed the strong positive association in patients with dementia, but found little evidence of an association in patients without dementia.

The adjusted SCCS method is relatively complex, (14) so it is desirable to apply a suitable sensitivity analysis to decide whether it is required. We applied the methods of the present paper, using the same exposure categories: time on antipsychotics as the main exposure risk period, followed by two washout periods of 91 days each (the original paper used five 35-day washout periods). The original paper used 5-year age bands; a finer age categorization was used with the adjusted SCCS method. (14) Here we also used 5-year age bands, but with grouped ages below 45 and above 95 to avoid sparseness.

Of the 6789 cases aged over 20 years used in the analysis, 5218 were censored and 1571 were observed right up until the nominal end of observation. Figure 6 shows the histograms of the time from event to actual end of observation in the two groups. The two histograms differ




FIGURE 5 Left: relative incidence associated with the pre-exposure 'risk' interval. Dashed line is at $RI = 1$. Right: relative incidence associated with MMR, when pre-exposure 'risk' intervals of specified duration are included. Dashed line is at RI with no pre-exposure period.




FIGURE 6 Left: intervals from event to censoring in censored cases. Right: intervals from event to end of observation in uncensored cases.

markedly. Thus, the histogram on the left shows a very sharp mode for intervals less than 0.5 years. These are likely to be due to stroke-induced deaths, and suggest that event-dependent observation periods may well be a problem in these data.

We therefore classified cases according to whether the observation periods were censored or not, and fitted the standard SCCS model and a model with an interaction between the censoring indicator and the exposure variable (the interaction model). In all models we assumed separate age and exposure effects for cases with and without prior dementia, but (for simplicity) combined all types of antipsychotics. Let E and A denote the exposure and age factors, D the 2-level factor for prior dementia and C the 2-level factor for presence of censoring. The standard model, expressed as a model formula using the usual notation, (22) includes the terms

$$D.E + A + A.D$$

and the interaction model includes the terms

$$D.E + C.E + D.C.E + A + A.D.$$

The main effects of C and D are not included as they are time-invariant and drop out of the likelihood. In R code, the model formulae are entered as $D/E + A * D$ and $(D * C)/E + A * D$, respectively. In addition, the model includes an offset term for the length of the interval, and an individual factor term.

The likelihood ratio test for the interaction model compared to the standard model gave a test statistic of 37.20 on 6 degrees of freedom, $p < 0.0001$. This suggests that there is significant event-dependence of observation periods. More importantly, the exposure effects for censored and uncensored cases are very different in cases with no prior dementia, as shown in Table 4 .

TABLE 4 Relative incidence (RI) and 95% confidence interval (CI) in exposed period by dementia group and model

	No dementia		With dementia	
	RI	95% CI	RI	95% CI
Standard model	1.42	(1.30, 1.56)	3.42	(2.90, 4.03)
Interaction model				
Censored	1.63	(1.47, 1.81)	3.47	(2.91, 4.12)
Uncensored	0.91	(0.75, 1.11)	3.11	(1.92, 5.04)
Model with terminal 'risk' period	1.16	(1.06, 1.27)	2.97	(2.53, 3.50)

The results from the uncensored group suggest that the significant effect of antipsychotics on stroke in cases without dementia obtained with the standard model may be spurious, as the significance of the effect disappears in this group. The association in cases with dementia, however, appears genuine: the estimates are virtually the same in censored and uncensored cases. These observations accord with the full analysis undertaken using a more complex model. (14)

We also did separate analyses in patients with and without dementia. The likelihood ratio test for the interaction versus standard model in cases with no prior dementia gave a test statistic of 32.39 on 3 degrees of freedom, $p < 0.0001$. In cases with prior dementia, the test statistic was 4.95 on 3 degrees of freedom, $p = 0.175$. Thus, in the group with dementia, there is only marginal evidence of event-dependent observation periods, with little impact on the exposure estimates, unlike the group without dementia.

Finally, we also fitted a terminal 'risk' period of $\kappa = 183$ days. This choice was based on the sharp peak in intervals from event to end of observation under 6 months in Figure 6 . The parameter θ was estimated to be 3.23, 95% CI (3.04, 3.44), and unsurprisingly the effect is highly significant. Table 4 also shows the values of the relative incidences associated with antipsychotics in patients with and without dementia in this analysis; similar results were obtained with $\kappa = 365$ days. They suggest that including a terminal 'risk' period reduces the bias, but does not eliminate it.

7 | FINAL REMARKS

We have presented some simple methods of analysis to test the validity of key assumptions of the self-controlled case series model, along with some related graphics. These methods can be implemented within the standard SCCS framework, without recourse to more complex adjusted SCCS methods. (12, 13, 14) The methods can be thought of as sensitivity analyses. It is important to note that, even if the assumptions are found to be questionable, the estimates of the parameters of interest may exhibit very little bias. Thus, our methods enable the researcher to check the likely robustness of the results obtained by standard SCCS analyses. When estimates are found not to be robust to failure of assumptions, more complex analyses may be required. (12, 13, 14)

The proposed test for event-dependence of exposures generally has good power in the moderate sample sizes investigated. Fitting a pre-exposure risk period generally reduces the bias caused by event-dependence of exposures, provided that overlaps between risk periods and the pre-exposure risk periods are allocated additional parameters. When occurrence of an event leads to additional exposures, other methods of handling overlaps can seriously destabilize the estimates. Our results suggest that when events delay subsequent exposures by a short time, the correction methods described here may not be worth applying as they may overcorrect the bias. Such scenarios, in particular, are not likely to be so simple in practice and require further investigation. For example, the situation where vaccination is delayed due to the adverse event is one such scenario, but delays will be of varying length and longer delays may be better represented by the 'remove exposures' scenario for which the correction reduces bias.

When events have a lasting impact on subsequent exposures, as when the event is a contra-indication to treatment, fitting a pre-exposure risk period will not correct the bias. To do so, it is necessary to use an adjusted SCCS method, (12, 13) or consider using an alternative self-controlled design. (7) However, it is usually known whether the event inhibits or precipitates further exposures, and so the direction of the bias is known.

We investigated two tests for event-dependence of observation periods. Both tests were found to have acceptable power with the moderate sample sizes investigated. The robustness test displayed a very clear relationship between power and the amount of bias resulting from event-dependence, and is therefore much more useful than the second test proposed, in which only the presence of event-dependence, rather than the degree of bias it may induce, may be detected. In circumstances where event-dependence of observation periods may arise, it is important to obtain the information on censoring required by the robustness test.

As with all sensitivity analyses, it is difficult to provide general guidance as to when results are to be deemed insufficiently robust to failure of assumptions. This will usually depend on context: if the substantive conclusions are unchanged, then the results may be deemed to be robust.

Acknowledgement

This research was supported by an MRC methodology grant MR/L009005/1

References

- [1] Farrington CP. Relative incidence estimation from case series for vaccine safety evaluation. *Biometrics* 1995;51:228–235.
- [2] Whitaker HJ, Farrington CP, Spiessens B, Musonda P. Tutorial in biostatistics: The self-controlled case series method. *Stat Med* 2006;25:1768–1797.
- [3] Farrington CP, Whitaker HJ. Semiparametric analysis of case series data (with Discussion). *J R Stat Soc Ser C Appl Stat* 2006;55:553–580.
- [4] Petersen I, Douglas I and Whitaker H. Self-controlled case series methods - an alternative to standard epidemiological study designs. *BMJ* 2016;354:i4515.
- [5] Weldelessie YG, Whitaker HJ, Farrington CP. Use of the self-controlled case-series method in vaccine safety studies: review and recommendations for best practice. *Epidemiol Infect* 2011;139:1805–1817.
- [6] Nordmann S, Biard L, Revaud P, Esposito-Farese M and Tubach F. Case-only designs in Pharmacoepidemiology: A Systematic Review. *Plos One* 2012;7:e49444.
- [7] Maclure M, Fireman B, Nelson JC, Hua W, Shoaibi A, Paredes A and Madigan D. When should case-only designs be used for safety monitoring of medical products? *Pharmacoepidemiol Drug Saf* 2012;21:50–61.
- [8] Farrington P, Whitaker H. Mortality and the self-controlled case series method: Letter to the Editor. *Pharmacoepidemiol Drug Saf* 2012;21:8.
- [9] Maclure M. The case-crossover design: A method for studying transient effects on the risk of acute events. *Am J Epidemiol* 1991;133:144–153.
- [10] Kalbfleisch JD and Prentice RL. *The Statistical Analysis of Failure Time Data*, Second Edition. Hoboken, NJ, USA: John Wiley and Sons; 2002.
- [11] Robins J. A new approach to causal inference in mortality studies with a sustained exposure period – Application to control of the healthy worker survivor effect. *Mathematical Modelling* 1986;7:1393–1512.
- [12] Farrington CP, Whitaker HJ, Hocine MN. Case series analysis for censored, perturbed or curtailed post-event exposures. *Biostatistics* 2009;10:3–16.
- [13] Kuhnert R, Hecker H, Poethko-Müller C. A modified self-controlled case series method to examine association between multidose vaccinations and death. *Stat Med* 2011;30:666–677.
- [14] Farrington CP, Anaya-Izquierdo K, Whitaker HJ, Hocine MN, Douglas I, and Smeeth L. Self-controlled case series analysis with event-dependent observation periods. *J Am Stat Assoc* 2011;106:417–426.
- [15] Lee KJ and Carlin JB. Fractional polynomial adjustment for time-varying covariates in a self-controlled case series analysis. *Stat Med* 2014;33:105–116.
- [16] Ghebremichael-Weldelessie Y, Whitaker HJ, Farrington CP. Self controlled case series method with smooth age effect. *Stat Med* 2014;33(4):639–649.
- [17] Ghebremichael-Weldelessie Y, Whitaker HJ, Farrington CP. Flexible modelling of vaccine effect in self-controlled case series models. *Biom J* 2016;58(3):607–622.
- [18] Ghebremichael-Weldelessie Y, Whitaker HJ, Farrington CP. Spline-based self-controlled case series method. *Stat Med* 2017;36(19):3022–3038.

- [19] Miller E, Waight P, Farrington P, Andrews N, Stowe J, Taylor B. Idiopathic thrombocytopenic purpura and MMR vaccine. *Arch Dis Child* 2001;84:227–229.
- [20] Murphy TV, Gargiullo PM, Massoudi MS, et al. Intussusception among infants given an oral rotavirus vaccine. *N Engl J Med* 2001;344:564–572.
- [21] Douglas IJ, Smeeth L. Exposure to antipsychotics and risk of stroke: Self controlled case series study. *BMJ* 2008;337(7670):616–618.
- [22] McCullagh P and Nelder JA. *Generalized Linear Models*, 2nd edition. London: Chapman and Hall; 1989.

