

Self-controlled case series studies: just how rare does a rare non-recurrent outcome need to be?

Heather J. Whitaker*^{1,2}, Colin D. Steer³, and C. Paddy Farrington¹

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

² Faculty of Epidemiology and Population Health, London School of Hygiene and Tropical Medicine, London, UK.

³ Public Health Sciences, Bristol Medical School, University of Bristol, Bristol, UK.

Received Dec 2017, revised zzz, accepted zzz

The self-controlled case series method assumes that adverse outcomes arise according to a non-homogeneous Poisson process. This implies that it is applicable to independent recurrent outcomes. However, the self-controlled case series method may also be applied to unique, non-recurrent outcomes or first outcomes only, in the limit where these become rare. We investigate this rare outcome assumption when the self-controlled case series method is applied to non-recurrent outcomes. We study this requirement analytically and by simulation, and quantify what is meant by ‘rare’ in this context. In simulations we also apply the self-controlled risk interval design, a special case of the self-controlled case series design. To illustrate, we extract data on the incidence rate of some recurrent and non-recurrent outcomes within a defined study population to check whether outcomes are sufficiently rare for the rare outcome assumption to hold when applying the self-controlled case series method to first or unique outcomes.

The main findings are that the relative bias should be no more than 5% when the cumulative incidence over total time observed is less than 0.1 per individual. Inclusion of age (or calendar time) effects will further reduce bias. Designs that begin observation with exposure maximise bias, whereas little or no bias will be apparent when there is no time trend in the distribution of exposures, or when exposure is central within time observed.

Key words: Self controlled case series; Rare events; Poisson process

Supporting Information for this article is available from the author or on the WWW under <http://dx.doi.org/10.1022/bimj.XXXXXXX> (please delete if not applicable)

1 Introduction

The self-controlled case series (SCCS) method is an epidemiological study design for investigating the temporal association between a time-varying exposure and an adverse health outcome (Farrington, 1995; Whitaker *et al.*, 2006; Petersen *et al.*, 2016). Clearly defined dates must be available for both outcomes and windows of time that are hypothesized to be at increased risk due to exposure, thus the SCCS method is best suited to abrupt-onset outcomes and transient exposures, though it can also be used for progressive conditions and indefinite exposures in some circumstances. Its main area of application to date has been in the study of adverse outcomes following vaccination (Weldeselassie *et al.*, 2011), though it has also been applied more widely, for example to study the safety of prescription medications (Gault *et al.*, 2017; Nordmann *et al.*, 2012; Ryan *et al.*, 2012), and to investigate infections as triggers of cardiovascular outcomes.

The main advantages are that it requires only cases, those who have experienced the outcome of interest, and is self-controlled, so that any fixed confounders such as sex and ethnicity are automatically controlled for. However, the SCCS method makes some strong assumptions (Whitaker *et al.*, 2018). The SCCS

*Corresponding author: e-mail: heather.whitaker@open.ac.uk, Phone: +44-1908-654370

method is unusual amongst epidemiological study designs in that it uses time after an adverse outcome has occurred. The main limiting assumption is that outcomes must not influence subsequent exposure during time under observation in the study, nor must outcomes influence time under observation. Further assumptions are that time-varying covariates act multiplicatively on the baseline incidence and that outcomes arise according a non-homogeneous Poisson process or are non-recurrent and rare. It is this last assumption that we investigate further here, though other assumptions do bear some influence on our findings.

The SCCS method is derived from a model similar to a cohort study in which outcomes arise according to a non-homogeneous Poisson process, by conditioning on the total number of outcomes experienced by each individual. The Poisson assumption implies that the outcomes are potentially recurrent. However, the method also applies to non-recurrent outcomes, in the limit where these become rare (Farrington, 1995; Farrington and Whitaker, 2006). The aim of this paper is to quantify what is meant by ‘rare’ in this context, which we investigate both analytically in a simple scenario and by simulation.

In simulations we apply both the self-controlled case series (SCCS) and self-controlled risk interval (SCRI) designs. The self-controlled risk interval design is a special case of the SCCS design in which the total time under observation is cut down to a narrower interval defined in relation to exposure times. In Section 2 we present the key features of the SCCS method, including the SCCS likelihood and design choices, including the SCRI design. Section 3 contains some analytical results in simple scenarios to quantify the asymptotic bias involved when the SCCS method is applied to non-recurrent outcomes. In Section 4 we present simulations. Section 5 contains some examples of outcome incidence rates.

2 Self-controlled case series (SCCS) method and designs

In this section, we begin by introducing the likelihood for the SCCS model with categorical exposure and age effects, which we refer to as the ‘standard’ SCCS model. We then distinguish the self-controlled risk interval (SCRI) design as a special case within the broader class of SCCS designs. We outline a simple adaptation of SCCS designs to circumvent issues with outcomes changing the probability of exposure. We then end this section by briefly outlining how bias arises when non-recurrent outcomes are not sufficiently rare, and how this depends upon the distribution of exposures within the time under observation.

2.1 The standard SCCS likelihood

Suppose that individuals in a given cohort are followed up during observation periods $(a_i, b_i]$, for the occurrence of an adverse health outcome of interest. Each individual i included in a case series experiences $n_i \geq 1$ outcomes over this observation period, $i = 1, \dots, N$. In a standard SCCS design, exposure status is categorical and assigned to time windows of hypothesized excess risk due to exposure known as exposure risk windows, of which there may be several to capture dose, washout periods or varying risk with time since first exposure, denoted $k = 1, \dots, K$. All other time within an observation period, but outside an exposure risk window, constitute baseline or reference windows, denoted $k = 0$.

During their observation period, individual i 's outcome incidence rate is modified by age (or calendar time) group, exposure risk group, and fixed factors specific to him or her. In a standard SCCS design, age (or time / season) is divided into categories $j = 1, 2, \dots, J$. We assume that these influences are captured by the following multiplicative incidence model:

$$\lambda_{ijk} = \exp(\phi_i + \alpha_j + \beta_k),$$

where ϕ_i represents the combined effect of individual-specific factors, α_j is the effect of age in age group j (or calendar time, as required), assumed to be common to all individuals in the cohort, and β_k is the log relative incidence associated with exposure risk window k , the parameters of interest.

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 outcomes during the observation period. Case i might

experience $n_i \geq 1$ outcomes, so the total number of outcomes is $M = \sum_{i=1}^N n_i$. 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 up to b_i , and the number of outcomes n_i experienced within the observation period. The outcomes n_{ijk} for case i occur within age group j and exposure risk window k . The SCCS likelihood is then

$$L = \prod_{i=1}^n \prod_{jk} \left(\frac{e_{ijk} \exp(\alpha_j + \beta_k)}{\sum_{rs} e_{irs} \exp(\alpha_r + \beta_s)} \right)^{n_{ijk}}.$$

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 outcome incidence rate.

Here we focus on the standard model for which piecewise constant age and exposure effects are modelled. The likelihood can be generalised and alternative modelling approaches for both age and exposure effects are available (Ghebremichael-Weldeslassie *et al.*, 2014; Lee and Carlin, 2014; Farrington and Whitaker, 2006).

The SCCS method is also valid for non-recurrent outcomes (so $n_i = 1$), in the limit where the baseline incidence rate for individual i , $\exp(\phi_i) \rightarrow 0$. In practice this requires outcomes to be uncommon: this is the rare disease assumption. For further details of this derivation see Farrington (1995); Farrington and Whitaker (2006).

2.2 Standard SCCS and SCRI designs

Originally, the self-controlled case series design was conceptualized with an observation period defined by age and/or calendar time boundaries (Farrington, 1995). For vaccine safety studies these will be periods during which vaccines were in current use and at ages that they are usually administered within the population, such as the second year of life for mumps, measles, rubella (MMR) vaccine or an influenza season for seasonal influenza vaccines. Alternatively, observation periods may reflect the length of a database record. Observation periods can be long, and to account for the fact that the baseline hazard may change it is often important to include age effects in the model. Thus, a set up for an SCCS model with an age-defined observation period is illustrated in figure 1, panel (a)

An alternative is to ascertain exposure histories, and define observation periods in relation to exposure, thus cutting down observation time used and altogether dropping cases whose outcomes arose further in time from exposure. This approach is taken for the self-controlled risk interval (SCRI) design (Tse *et al.*, 2012; Baker *et al.*, 2015), a special case of an SCCS design. In a typical SCRI design, a single exposure risk window is defined, along with either one or two reference or control windows, before and after the exposure risk window or following the exposure risk window. Reference windows are not necessarily contiguous to exposure risk windows, as for example, allowance may need to be made for washout periods (whereas washout periods would be modelled in the original SCCS design, SCRI simply leaves gaps in observation time). This is illustrated in figure 1, panel (c).

Age may be dropped from the SCRI model under the assumption that age effects are constant over the relatively short observation period, though this may not always be reasonable and strategies for control of age effects in the SCRI model have been outlined by Li *et al.* (2015).

Much consideration has been given to the main limiting assumption in the family of SCCS designs that outcomes must not influence subsequent exposure (Whitaker *et al.*, 2006; Petersen *et al.*, 2016; Farrington *et al.*, 2009; Kuhnert *et al.*, 2011). Short-term influence is often accounted for by including a pre-exposure window in a SCCS modelling approach, or by including a gap in observation prior to exposure in the SCRI approach. Long-term or permanent influence is usually more complicated to account for (Farrington *et al.*, 2009), except when there can only be a single exposure within an observation period, where a simple strategy is to use only post-exposure time as illustrated in figure 1, panel (b) for SCCS and panel (d) for SCRI. For example, this design set up can be used when the outcome of interest is death, after

which exposure is impossible, post-death time is included until the planned end of observation (Hubbard *et al.*, 2005; Petersen *et al.*, 2016). That there should be only a single exposure within each observation period is more easily met using the SCRI approach which shortens observation time, for example there is often some minimum interval between vaccine doses. The SCRI approach using a single post-exposure reference window roughly matches the adapted SCCS method formulated to study the association between multi-dose vaccinations and death outlined by Kuhnert *et al.* (2011). Note that only a risk gradient can be estimated using a design that use post-exposure time only, such as those illustrated in figure 1 (b) and (d). It is necessary to assume that the risk returns to baseline in the post-exposure control period for the relative incidence to retain the same definition.

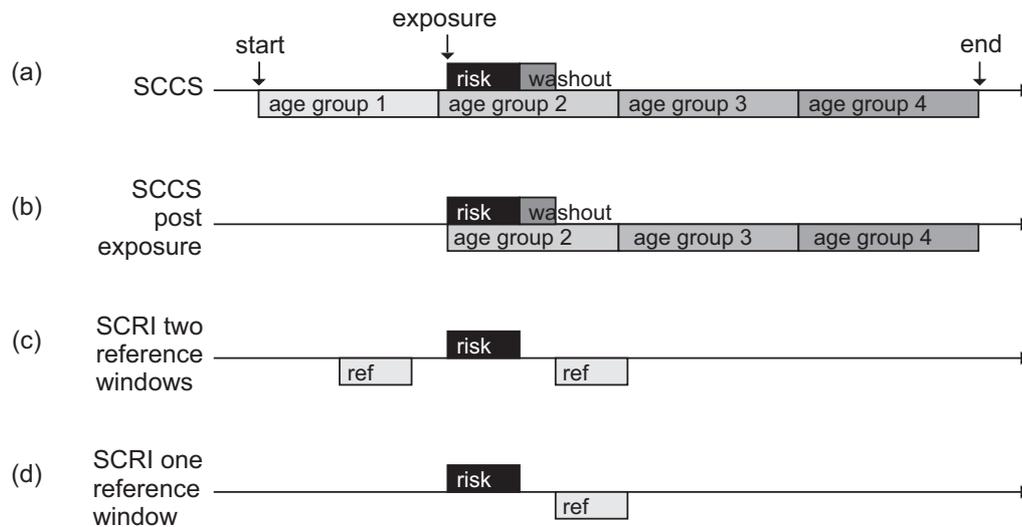


Figure 1 Example time lines for standard SCCS and SCRI designs with one exposure. The exposure risk window is labelled ‘risk’ and washout window is labelled ‘washout’. Panel (a) SCCS with observation period defined by age boundaries and four age groups. Panel (b) SCCS with observation period starting at exposure. Panel (c) SCRI with two reference windows (labelled ‘ref’) either side of the exposure risk window. Panel (d) SCRI with one reference window after the exposure risk window.

2.3 Bias related to non-recurrent outcomes

It is assumed that outcomes arise according to a non-homogeneous Poisson process. Under any standard SCCS model (including SCRI), the Poisson rate is assumed to be constant within each window of time defined by age and exposure categories. Outcomes are counts within each time window. Naturally, such a model allows outcomes to be recurrent, but outcomes must arise in time independently of one another. Such independence is frequently not present in practice, for example the timing of a second myocardial infarction will often not be independent of the timing of the first myocardial infarction. A simple work-around is to include only first outcomes in a study. In considering how bias arises it is conceptually easier to think of a non-recurrent outcome as the first of a potentially recurrent outcome, even if recurrence is impossible. The shape of the distribution of first outcomes over age or time will be shifted toward lower ages or earlier dates in the observation period than the distribution of all outcomes. In other words, if second and subsequent outcomes are omitted, an increasing deficit of outcomes is created as time progresses. If this shift in temporal distribution of common outcomes is not (or insufficiently) taken into account, estimates of the exposure-related relative incidence may become biased. The extent of the bias will depend heavily upon the distribution of exposures throughout observation periods. Where exposures tend to come toward the

beginning of observation, exposures will coincide more often with first outcomes and the estimates of the exposure-related relative incidence $\exp(\beta_k)$ may become biased upward. Conversely, where exposures tend to come toward the end of observation, estimates of $\exp(\beta_k)$ may become biased toward 0. Little or no bias should result when the probability of exposure is roughly constant throughout observation periods, or when there is an equal distribution of reference time either side of an exposure risk window. Thus, any SCCS design set-up where observation starts with a single exposure represents an extreme where the potential for upward bias in $\exp(\beta_k)$ is maximised (i.e. for no other SCCS design set up can bias be greater). Whereas little or no bias should be present for an SCRI design with two reference windows of equal length either side of the exposure risk window.

3 Analytic evaluation for a simple scenario

The rare outcome assumption cannot be investigated from case series data alone: it requires external information. Such information is usually not difficult to obtain, especially since the precise disease frequency is not required. However, it helps to know what is meant by ‘rare’ in this context. To this end, we undertake some calculations in a special, but extreme, scenario.

Assume that the outcome hazard λ is constant and that outcome times T are exponentially distributed, $T \sim M(\lambda)$. Suppose that the observation period is $(0, b]$ and that there is a single exposure risk window $(c, c + d]$ of length d with $0 \leq c \leq b - d$. We investigate the asymptotic (large sample) bias in the relative incidence ρ associated with this exposure risk window, in an SCCS analysis that makes no allowance for age. Including age effects would be expected to reduce the bias in ρ . Note that investigating large sample bias allows us to distinguish systematic bias relating to unique outcomes from bias relating to small samples (Musonda *et al.*, 2008).

Let Λ denote the cumulative incidence over the entire observation period $(0, b]$, including the effect of exposure. Thus,

$$\Lambda = \rho\lambda d + \lambda(b - d).$$

The probability that an outcome occurs in the exposure risk window $(c, c + d]$, conditional on it occurring in $(0, b]$, is

$$P_1 = P(c < T \leq c + d | T \leq b) = \frac{\exp(-\lambda c)(1 - \exp(-\rho\lambda d))}{1 - \exp(-\Lambda)}.$$

Similarly, the conditional probability that an outcome does not occur in the exposure risk window is

$$P_0 = P(T \leq c \text{ or } T \geq c + d | T \leq b) = \frac{1 - \exp(-\lambda(b - d) - \rho\lambda d) - \exp(-\lambda c)(1 - \exp(-\rho\lambda d))}{1 - \exp(-\Lambda)}.$$

The likelihood for an SCCS model with common observation period $(0, b]$, common exposure risk window $(c, c + d]$ and no age effects is binomial. If N is the total number of cases (and therefore of unique outcomes), N_1 is the number of outcomes in the exposure risk window, and N_0 the number of outcomes in the reference windows, then the maximum likelihood estimator of ρ is

$$\hat{\rho} = \frac{N_1}{N_0} \times \frac{b - d}{d}.$$

Asymptotically as $N \rightarrow \infty$,

$$\hat{\rho} \rightarrow \bar{\rho} = \frac{P_1}{P_0} \times \frac{b - d}{d} \quad \text{in probability.}$$

We consider the limit in which λ , and hence Λ , are small. Using standard Taylor approximations,

$$\begin{aligned}\bar{\rho} &= \frac{\exp(-\lambda c)(1 - \exp(-\rho\lambda d))}{1 - \exp(-\lambda(b-d) - \rho\lambda d) - \exp(-\lambda c)(1 - \exp(-\rho\lambda d))} \times \frac{b-d}{d} \\ &\simeq \frac{(1 - \lambda c)[\rho\lambda d - \frac{1}{2}(\rho\lambda d)^2]}{\lambda(b-d) + \rho\lambda d - \frac{1}{2}[\lambda(b-d) + \rho\lambda d]^2 - (1 - \lambda c)[\rho\lambda d - \frac{1}{2}(\rho\lambda d)^2]} \times \frac{b-d}{d} \\ &\simeq \rho(1 - \lambda c)\left(1 - \frac{1}{2}\rho\lambda d\right)\left[1 + \frac{1}{2}\lambda(b-d) + \rho\lambda d - \rho\lambda\frac{dc}{b-d}\right] \\ &\simeq \rho\left[1 + \frac{1}{2}\Lambda\left(1 - 2\frac{c}{b-d}\right)\right],\end{aligned}$$

to first order in Λ . Thus, the relative bias is of order $\Lambda/2$ if $c = 0$, $-\Lambda/2$ if $c = b - d$, and 0 if $c = (b - d)/2$. To put this in context, recall subsection 2.3; if outcome events arise according to a non-homogeneous Poisson process and unique outcomes are thought of as the first event only, an excess of outcomes will occur in exposure risk windows that fall at the beginning of the observation period ($c = 0$) resulting in upward bias, which is of order $\Lambda/2$. Vice versa, there will be a dearth of outcomes in exposure risk windows that fall at the end of the observation period ($c = b - d$) resulting in downward bias, which is of order $-\Lambda/2$. In general, to first order,

$$\left|\frac{\bar{\rho} - \rho}{\rho}\right| \leq \frac{1}{2}\Lambda.$$

Thus the relative bias in this setting is at most $\frac{1}{2}\Lambda$ in absolute value.

For example, if the cumulative incidence over the observation period is about 0.1, the relative bias is of the order of 5% or less. The additional inclusion of age effects will correct this bias to some degree, by allowing for the shift in age effects produced by the absence of second or subsequent outcome events.

4 Simulations

Simulation choices using a single exposure risk period were made to demonstrate minimal, intermediate and maximal bias when applying SCCS and SCRI models to the first of multiple outcomes.

4.1 Method

A study period was fixed at $(0, 100]$ days for 1000 individuals $i = 1, \dots, 1000$. The start of the exposure risk period c_i for individual was simulated within the period $(0, 90]$ days and had fixed length 10 days, so exposure risk periods were $(c_i, c_i + 10]$. Exposure status is denoted $k = 1$ inside the exposure risk window and $k = 0$ otherwise. Two scenarios for the distribution of exposure starts were explored, uniformly distributed throughout the observation period (to produce no bias using a standard SCCS approach) and linear decreasing over the observation period (to produce some bias using a standard SCCS approach). SCRI reference windows were fixed at $(\max(0, c_i - 10), c_i]$ (before exposure) and $(c_i + 10, \min(c_i + 20, 100)]$ (after exposure), within the study period boundaries. Four age groups $j = 1, 2, 3, 4$ of length 25 days were fixed. For each individual, nine boundaries at 0, $(\max(0, c_i - 10), c_i]$, $(c_i, c_i + 10]$, $(c_i + 10, \min(c_i + 20, 100)]$, 25, 50, 75 and 100 were ordered to form eight segments, indexed ijk , of length $e_{ijk} \geq 0$. Segments are illustrated in figure 2. Scenarios were simulated for three true exposure-related relative incidences $\exp(\beta_1)$, either $\exp(\beta_1) = 1, 2, 5$ (always with $\exp(\beta_0) = 1$). No effect of age was simulated, the true age-related relative incidence for multiple outcomes $\exp(\alpha_j) = 1$, $j = 1, 2, 3, 4$. Let Λ denote the cumulative outcome rate over the entire observation period, excluding the effect of exposure (note that this differs from section 3 in which Λ included the effect of exposure). Scenarios were simulated with six choices of outcome rate $\Lambda = 0.02, 0.05, 0.1, 0.2, 0.5, 1$. An outcome hazard was determined for each segment $\lambda_{ijk} = e_{ijk} \times \Lambda/100 \times \exp(\beta_k)$ and an overall outcome count simulated $n_{ijk} \sim \text{Poisson}(\lambda_{ijk})$. If

$\sum_i n_{ijk} = 0$ for an individual, outcomes were re-simulated for that individual; this was repeated until all 1000 individuals had least one outcome $\sum_i n_{ijk} \geq 1$. The segment during which the first outcome occurred was identified.

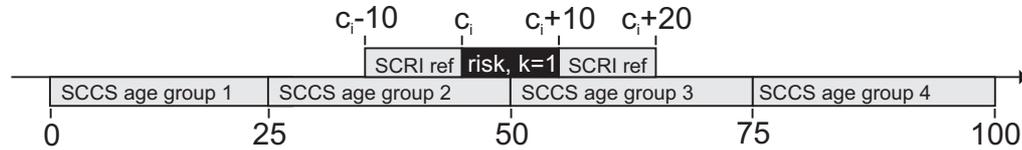


Figure 2 Segments for an example simulated individual. The exposure risk window $k = 1$ is represented by the black window, outside this window $k = 0$.

Segments were then restricted and/or combined to fit six models to the simulated data on each of the recurrent outcome counts and first outcomes only (12 models total). The six models were: SCCS over the full study period $(0, 100]$ both with no age effect and with 4 age groups, SCCS starting from exposure $(c_i, 100]$ both with no age effect and with 4 age groups, SCRI with two reference windows either side of the exposure risk window $(c_i - 10, c_i + 20]$ and SCRI with one reference window after the exposure risk window $(c_i, c_i + 20]$. Only SCCS models over the full study period included all 1000 simulated cases. Restrictions on the observation periods for other models resulted in the exclusion of cases outside the given boundaries.

Simulations were carried out using R, and models were fitted by conditional Poisson regression using the package `gnm` (Turner and Firth, 2015). Simulations were replicated 5000 times. A relative bias was calculated as

$$\text{relative bias} = \frac{|\exp(\text{mean } \hat{\beta}) - \exp(\text{true } \beta)|}{\exp(\text{true } \hat{\beta})}.$$

4.2 Results

Simulation results (mean $\hat{\beta}$) for the six models applied to first outcomes only and linear decreasing exposure start times over the observation period are shown in Table 1. Results for the same model applied to all outcomes is available in the supporting information, all displayed very little bias. Table 2 contains mean relative biases for the same scenario, along with an approximate cumulative incidence for each model. When reference periods are included both before and after exposure, SCCS with no age effect becomes increasingly biased as the baseline cumulative outcome rate increases. Bias in the SCCS model is greatly reduced by the inclusion of age effects. As expected, the SCRI model with two reference windows either side of the exposure risk window showed very little bias.

As expected, bias was greater for the models with observation starting from exposure. The relative bias should be approximately $\frac{1}{2}\Lambda$ for post-exposure reference only models without age effects, and this represents the maximum possible relative bias. From Table 2 it can be seen that simulation results are roughly, albeit not exactly, in line with this. When no age effects were fitted in the SCCS model, bias became very large when the cumulative outcome rate rose above 0.1. The simulation results demonstrate that inclusion of age effects reduces bias considerably for the SCCS model, though where the baseline outcome rate within an age group remains above 0.1, relative bias remains greater than 5%. The SCRI model, with observation length of only 20 days, displays greater bias than the SCCS model with age groups of length 25-days, though bias is considerably less than the SCCS model with no age effects. However, the proportion of exposure risk time over the observation period is greater for the SCRI model, which pushes up the overall cumulative outcome rate over the 20-day period relative to the average cumulative outcome rate within a 25-day SCCS age group.

Table 1 Simulation results, 5000 replications. Exposure start day linear decreasing over observation period. Analyses of first outcomes only. Base rate is the baseline cumulative outcome rate over 100 days. True β is the true natural logarithm of the relative incidence associated with the exposure risk window. Results are the means of the estimated log relative incidences $\hat{\beta}$. *Starred results show $> 5\%$ relative bias. Means of s.e. ($\hat{\beta}$) are also given.

base rate	true β	reference before and after exposure			post-exposure reference only		
		SCCS	SCCS + age	SCRI, 2 ref	SCCS	SCCS + age	SCRI, 1 ref
		mean $\hat{\beta}$ (mean s.e. $\hat{\beta}$)					
0.02	0.000	0.000(0.106)	-0.004(0.107)	0.000(0.124)	0.004(0.110)	-0.003(0.135)	0.005(0.142)
0.05	0.000	0.005(0.105)	-0.003(0.107)	-0.002(0.124)	0.015(0.110)	0.001(0.135)	0.005(0.142)
0.1	0.000	0.012(0.105)	-0.004(0.107)	-0.002(0.123)	0.031(0.109)	0.002(0.134)	0.007(0.141)
0.2	0.000	0.028(0.104)	-0.004(0.106)	-0.001(0.123)	*0.069(0.109)	0.008(0.134)	0.021(0.141)
0.5	0.000	*0.071(0.103)	-0.005(0.104)	0.001(0.120)	*0.175(0.107)	0.023(0.133)	*0.050(0.139)
1	0.000	*0.137(0.100)	-0.007(0.102)	0.002(0.117)	*0.353(0.106)	*0.053(0.133)	*0.102(0.137)
0.02	0.693	0.694(0.082)	0.691(0.084)	0.693(0.107)	0.699(0.087)	0.695(0.117)	0.700(0.129)
0.05	0.693	0.700(0.082)	0.692(0.084)	0.696(0.106)	0.713(0.087)	0.699(0.117)	0.707(0.129)
0.1	0.693	0.710(0.082)	0.693(0.084)	0.697(0.106)	*0.735(0.087)	0.706(0.117)	0.714(0.129)
0.2	0.693	0.725(0.081)	0.691(0.083)	0.694(0.105)	*0.773(0.087)	0.713(0.117)	0.726(0.129)
0.5	0.693	*0.772(0.080)	0.691(0.082)	0.698(0.104)	*0.897(0.087)	*0.746(0.118)	*0.772(0.128)
1	0.693	*0.840(0.078)	0.688(0.080)	0.699(0.101)	*1.106(0.086)	*0.803(0.119)	*0.849(0.128)
0.02	1.609	1.615(0.066)	1.611(0.069)	1.614(0.101)	1.622(0.073)	1.617(0.111)	1.625(0.131)
0.05	1.609	1.620(0.066)	1.610(0.069)	1.614(0.101)	1.638(0.073)	1.627(0.111)	1.633(0.131)
0.1	1.609	1.633(0.066)	1.613(0.069)	1.616(0.101)	1.668(0.074)	1.636(0.111)	1.647(0.131)
0.2	1.609	1.651(0.066)	1.610(0.069)	1.614(0.100)	*1.723(0.074)	1.662(0.112)	1.677(0.132)
0.5	1.609	*1.706(0.065)	1.609(0.068)	1.614(0.098)	*1.899(0.076)	*1.738(0.115)	*1.770(0.135)
1	1.609	*1.762(0.065)	1.597(0.067)	1.598(0.096)	*2.187(0.079)	*1.867(0.123)	*1.926(0.142)

The mean standard errors (s.e.) in the table reflect the number of cases and observation time included in each analysis. Age effects will be estimated with much less precision for the SCCS analyses starting observation with exposure, which affects the precision of the exposure estimates.

Simulation results with a uniform distribution of first exposure over the 100 day study period are given in the supporting information. Results are similar, except that for the full SCCS model with reference before and after exposure with no age effect, results are very close to unbiased, as expected.

4.3 Evaluation

For SCCS models, it appears to be key that the cumulative outcome rate over the observation period is less than 0.1 to ensure bias is less than 5%, and that inclusion of age effects clearly reduces bias further. The same rule of thumb appears to apply for the SCRI model with 1 reference, that the cumulative outcome rate over all windows should be less than 0.1 to ensure relative bias below 5%. Bias is considerably less for any bi-directional design choice, which should also apply for standard SCCS models where there are multiple intermittent exposures.

The SCRI model with 2 equal length reference periods either side of the exposure risk window minimises bias due to common unique events by design and can sometimes be convenient in terms of data collection. Potential disadvantages of SCRI over SCCS to be mindful of when designing a study are that reference windows need to be carefully chosen, age or season effects if present will be more difficult to allow for, and reduced power, which not only depends on case sample size (as cases whose outcome occurred

Table 2 Relative bias calculated from simulation results, 5000 replications. Exposure start day linear decreasing over observation period. Analyses of first outcomes only. The column labelled ‘base rate’ gives the baseline cumulative outcome rate over 100 days, not taking the exposure effect into account. Columns labelled ‘approx Λ ’ give an approximate cumulative incidence, this has been calculated for each of full SCCS, SCRI with two reference windows, SCCS post exposure only and SCRI with one post exposure reference only. Columns labelled ‘relative bias’ contain the mean relative bias defined at the end of section 4.1.

base rate	true β	reference before and after exposure					post-exposure reference only				
		SCCS			SCRI (2 reference)		SCCS			SCRI (1 reference)	
		no age	age	relative bias	approx Λ	relative bias	no age	age	relative bias	approx Λ	relative bias
0.02	0.000	0.020	0.000	0.004	0.006	0.000	0.014	0.004	0.003	0.004	0.005
0.05	0.000	0.050	0.005	0.003	0.015	0.002	0.035	0.016	0.001	0.010	0.005
0.1	0.000	0.100	0.012	0.004	0.030	0.002	0.070	0.032	0.002	0.020	0.007
0.2	0.000	0.200	0.028	0.004	0.060	0.001	0.140	0.072	0.008	0.040	0.021
0.5	0.000	0.500	0.074	0.005	0.150	0.001	0.350	0.191	0.023	0.100	0.051
1	0.000	1.000	0.147	0.007	0.300	0.002	0.700	0.424	0.055	0.200	0.108
0.02	0.693	0.022	0.001	0.002	0.008	0.000	0.016	0.006	0.002	0.006	0.007
0.05	0.693	0.055	0.007	0.001	0.020	0.003	0.040	0.020	0.006	0.015	0.013
0.1	0.693	0.110	0.017	0.000	0.040	0.004	0.080	0.043	0.013	0.030	0.021
0.2	0.693	0.220	0.032	0.002	0.080	0.001	0.160	0.083	0.020	0.060	0.033
0.5	0.693	0.550	0.082	0.002	0.200	0.005	0.400	0.227	0.054	0.150	0.082
1	0.693	1.100	0.158	0.005	0.400	0.006	0.800	0.511	0.116	0.300	0.169
0.02	1.609	0.028	0.005	0.001	0.014	0.004	0.022	0.013	0.007	0.012	0.016
0.05	1.609	0.070	0.010	0.001	0.035	0.005	0.055	0.029	0.017	0.030	0.024
0.1	1.609	0.140	0.024	0.003	0.070	0.007	0.110	0.060	0.027	0.060	0.039
0.2	1.609	0.280	0.043	0.001	0.140	0.005	0.220	0.120	0.054	0.120	0.070
0.5	1.609	0.700	0.102	0.000	0.350	0.005	0.550	0.336	0.137	0.300	0.174
1	1.609	1.400	0.165	0.012	0.700	0.011	1.100	0.781	0.294	0.600	0.372

far in time from exposure are dropped), but also on the ratio of the length of exposure risk windows to reference windows (with optimal efficiency when the ratio of the total length of all exposure risk windows to the total length of all reference windows is small). Further, use of defined reference windows in relation to exposure rather than use of all available baseline time will in practice subtly change interpretation of results; relative incidences gained from a full SCCS analysis with baseline time both before and after exposure should target the relative risk of a cohort study, whereas relative incidences gained from SCRI models and SCCS including only reference time after exposure can be thought of as a risk gradient.

5 Examples

The information required to assess the potential for bias relating to unique events are incidence rates in the study population. Unfortunately, this cannot be quantified from a case series alone. Sometimes information on a full cohort will be available, but where it is not, information on incidence rates is readily available in the literature and a rough idea of cumulative incidence should be sufficient to assess the potential for bias relating to unique events.

5.1 Febrile convulsions at ages 0-24 months

Febrile convulsions after childhood vaccinations have been studied using SCCS (Weldeslassie *et al.*, 2011; Hanf *et al.*, 2013; Huang *et al.*, 2010). Most studies have included recurrent convulsions, such as Hanf *et al.* (2013) in a study with mumps, measles, rubella (MMR) vaccination as the exposure; re-admissions within 72 hours were counted as the same episode. Observation periods span the period during which vaccinations are normally given. For example, MMR is often given during the second year of life, so Hanf *et al.* (2013) included children between 240 and 730 days of age. Diphtheria-tetanus-acellular pertussis (DTaP) is given during the first two years of life, so Huang *et al.* (2010) included children aged 6 weeks to 23 months. There is a strong age trend in febrile convulsions, with incidence peaking during the second year of life, so studies include age in the SCCS model.

A British cohort study followed 13,135 children from birth to age 5 (Verity *et al.*, 1985). Excluding 13 children with a prior condition and 97 with missing information, it was estimated that 197 children experienced a febrile convulsion before the age of 2 years, of which 82 experienced more than one convulsion at some point before the age of 5 years. Exact numbers of total febrile convulsions by age are not given, but assuming that all recurrences occurred before age 2 (so as to over-estimate incidence), the cumulative incidence of first and recurrent convulsions is approximately 0.034 for the first two years of life. This is under 0.05, febrile convulsions are sufficiently rare such that little or no bias should be present in studies of first events only (most studies include multiple events).

5.2 Myocardial infarction

SCCS has been used to study the association between prescription medicines and first myocardial infarction (Gault *et al.*, 2017). For example, Wong *et al.* (2016) studied cardiovascular outcomes associated with use of clarithromycin. The study period spanned 10 years from 2003 to 2012, and included age bands of length 1 year. Thus, observation periods can be relatively long.

ARIC surveillance (USA) provide data on the number of coronary events per year per 1000 persons in the population, by age, race and sex (Benjamin *et al.*, 2017). Events are defined as definite or probable myocardial infarctions (new or recurrent) and definite coronary heart disease deaths. The average incidence of events is highest for black males aged 75-84 years, at approximately 19 per year per 1000 persons for the period 2005-2013. The cumulative incidence over, say, 10 years, for an individual in this group would be 0.19. Depending on design choices, two or more age bands (of length 5 years or less) may be needed in an SCCS study to ensure relative bias is less than 5%. However, populations studied are likely to be mixed in terms of age, sex and race, and overall cumulative incidence will likely be lower. For example, average incidence of events at ages 55-64 is reported to be between 2 and 8 per year per 1000 persons depending on race and sex, thus cumulative incidence over 10 years for this age group is clearly less than 0.1.

5.3 Death within population of opioid drug abusers

SCCS methods have recently been applied to all cause mortality amongst drug abusers being treated in UK primary care (Steer *et al.*, 2018).

In total, 11,033 patients were identified who had received methadone or buprenorphine treatment for opioid drug abuse between 1998 and 2014. For this cohort, the average time of follow-up was 2.76 years with 587 deaths being observed. This suggested an overall incidence rate per patient of 0.053. The analyses adjusted for age (roughly using age bands of length 10 years, which were constant within many individuals with shorter observation periods), calendar year and an index of comorbidity based upon 17 chronic illnesses (which could vary with time).

Much care was taken to reduce bias in analyses by use of extended SCCS methods for death outcomes (Farrington *et al.*, 2009; Kuhnert *et al.*, 2011) with additional modification of the treatment end during the planned follow-up period. However, these extended methods use post-exposure time only to estimate exposure-related effects, which should maximise systematic bias relating to unique outcomes.

Within this cohort of drug abusers, it was observed that death is more common within certain population subgroups. For example, for age, the incidence rate rose to 0.176 for those patients aged 50+ years while, for comorbidity, the incidence rate was 0.439 for the highest category.

6 Final remarks

We have offered some guidance about how rare a rare unique adverse health outcome is required to be in order that relative bias is less than about 5%, or that relative bias should be to the order of no more than $\frac{1}{2}\Lambda$, where Λ is the cumulative incidence over the entire observation period. We acknowledge that our choice of 5% here is arbitrary, and that what is acceptable within a particular epidemiological study setting may vary. However, it may be useful to acknowledge the potential for bias in SCCS studies of non-rare unique outcomes, particularly when results are borderline statistically significant. Evaluation of the potential order of bias can be approximate, using information external to a study if necessary. In our examples, we saw that incidence rates vary between subgroups based on age, ethnicity and comorbidity. We acknowledge that it is a limitation of this work that differing incidence between subgroups has not been evaluated, although we believe that a rough assessment can be based upon average incidence over observation periods.

The direction of bias due to common unique outcomes is fairly predictable, with upward bias created when exposures tend to fall toward the beginning of observation and downward bias when exposures tend to fall toward the end of observation periods. Regardless of how common unique outcomes are, little bias will arise when exposures are randomly scattered throughout, or fall toward the middle of observation periods.

Our findings highlight the importance of including age (and/or calendar time) effects in SCCS studies. This is particularly important for non-rare unique outcomes to allow for the shift in age effects produced by the absence of second or subsequent outcome events that would be present according to the assumption that events arise according to a non-homogeneous Poisson process. When unique outcomes are not rare, fine control of age effects is preferable. This can be achieved by including sufficiently narrow age categories, or by modelling age effects using splines, fractional polynomials or the semi-parametric SCCS model (Ghebremichael-Weldeslassie *et al.*, 2014; Lee and Carlin, 2014; Farrington and Whitaker, 2006). Note that sample size plays no role in bias related to common unique events, but increasing sample size can help estimates age effects more reliably. Age is frequently ignored in SCRI studies under the assumption that incidence changes very little over the time observed. Study of common unique outcomes coupled with the use of only one reference window may give reason to take age effects into account, and Li *et al.* (2015) outline some strategies for controlling for age in SCRI studies, by either including unexposed cases or by using external information on outcome incidence rates.

There are several potential sources of bias that researchers should be mindful of when conducting SCCS or SCRI studies, including time-varying confounders, small sample estimation bias (Musonda *et al.*, 2008), systematic bias resulting from outcomes that prohibit or precipitate subsequent exposure (Farrington *et al.*, 2009; Kuhnert *et al.*, 2011), systematic bias resulting from outcomes that censor subsequent observation (Farrington *et al.*, 2011), and bias resulting from outcomes that reasonably do not arise according to a non-homogeneous Poisson process (Simpson, 2013), which includes common unique outcomes. Outcomes that potentially censor the observation period, such as myocardial infarction or stroke that carry high mortality risk, have not been mentioned in the present paper, but similar to bias relating to unique common outcomes, the magnitude and direction of bias depends on the distribution of exposure risk windows over observation periods. Simple sensitivity analyses can be carried out to check whether results are robust to such bias (Whitaker *et al.*, 2018). In our experience, bias resulting from outcomes that permanently prohibit or precipitate subsequent exposure is the most problematic and where such an issue is present, use of an SCCS design that circumvents it is necessary. However, SCCS models that begin observation from exposure maximise bias relating to non-rare unique outcomes, and this includes the modified SCCS models outlined in Farrington *et al.* (2009) and Kuhnert *et al.* (2011) for outcomes that prohibit subsequent exposure.

Outcomes that can be studied using such SCCS models include death (Petersen *et al.*, 2016), and we have demonstrated that death can be relatively common within certain population subgroups. To reiterate, use of these models that use observation time from exposure forwards is essential to reduce bias relating to the assumption that outcomes must not influence subsequent exposure, and any bias relating to common unique non-recurrent outcomes is likely to be small in comparison with this.

Acknowledgements HJW and CPF were partially supported by an MRC Methodology Grant MR/L009005/1. CDS was funded by a NIHR Health Services and Delivery Research grant (PI Prof Matt Hickman).

Conflict of Interest

The authors have declared no conflict of interest.

References

- Baker M.A., Lieu T.A., Li L., Hua W., Qiang Y., Tse Kawai A., Fireman B.H., Martin D.B. and Nguyen M.D. (2015). A vaccine safety study design selection framework for the psotlicensure rapid immunization safety monitoring program. *American Journal of Epidemiology* **181**, 608-618.
- Benjamin E.J. et al (2017). Heart Disease and Stroke Statistics 2017 Update: A Report From the American Heart Association. *Circulation* **135**, e146-e603. (Figures were taken from Chart 20-4 e516).
- Farrington C.P. (1995). Relative incidence estimation from case series for vaccine safety evaluation. *Biometrics* **51**, 228-235.
- Farrington C.P., Whitaker H.J. (2006). Semiparametric analysis of case series data (with Discussion). *Journal of the Royal Statistical Society Series C, Applied Statistics* **55**, 553-580.
- Farrington C.P., Whitaker H.J., Hocine M.N. (2009). Case series analysis for censored, perturbed or curtailed post-event exposures. *Biostatistics* **10**, 3-16.
- Farrington C.P., Anaya-Izquierdo K., Whitaker H.J., Hocine M.N., Douglas I. and Smeeth L. (2011). Self-controlled case series analysis with event-dependent observation periods. *Journal of the American Statistical Association*, **106**, 417-426.
- Gault N., Castaeda-Sanabria J., De Rycke Y., Guillo S., Foulon S., Tubach F. (2017). Self-controlled designs in pharmacoepidemiology involving electronic healthcare databases: A systematic review. *BMC Medical Research Methodology* **17**.
- Ghebremichael-Weldeslassie Y., Whitaker H. and Farrington C.P. (2014). Self-controlled case series method with smooth age effect. *Statistics in Medicine* **33**, 639-649.
- Hanf M., Quantin C., Farrington P., Benzenine E., Hocine N.M., Velten M., Tubert-Bitter P. and Escolano S. (2013). Validation of the French national health insurance information system as a tool in vaccine safety assessment: application to febrile convulsions after pediatric measles/mumps/rubella immunization. *Vaccine* **31**(49), 5856-5862
- Huang W.T., Gargiullo P.M., Broder K.R., Weintraub E.S., Iskander J.K., Klein N.P., Baggs J.M. and the Vaccine Safety Datalink Team (2010). Lack of association between acellular pertussis vaccine and seizures in early childhood. *Pediatrics* **126**(2), 263-9.
- Hubbard R., Lewis S., West J., Smith C., Godfrey C., Smeeth L., et al (2005). Bupropion and the risk of sudden death: A self-controlled case-series analysis using The Health Improvement Network. *Thorax* **60**, 848-850.
- Kuhnert R., Hecker H., Poethko-Muller C. et al. (2011). A modified self-controlled case series method to examine association between multidose vaccinations and death. *Statistics in Medicine* **30**, 666-677.
- Lee K.J. and Carlin J.B. (2014). Fractional polynomial adjustment for time-varying covariates in a self-controlled case series analysis. *Statistics in Medicine* **33**, 105-116.
- Li L., Kulldorf M., Russek-Cohen E., Tse Kawai A., Hua W. (2015). Quantifying the impact of time-varying baseline risk adjustment in the self-controlled risk interval design. *Pharmacoepidemiology and Drug Safety* **24**, 1304-1312.
- Musonda P., Hocine M.N., Whitaker H.J. and Farrington C.P. (2008). Self-controlled case series method: small sample performance. *Computational Statistics and Data Analysis* **52**, 1942-1957.

- Nordmann S., Biard L., Revaud P., Esposito-Farese M. and Tubach F. (2012). Case-only designs in Pharmacoepidemiology: A Systematic Review. *Plos One* **7**, e49444.
- Petersen I., Douglas I. and Whitaker H. (2016). Self-controlled case series methods - an alternative to standard epidemiological study designs. *British Medical Journal* **354**, i4515.
- Ryan P.B., Madigan D., Stang P.E., Marc Overhage J., Racoosin J.A., Hartzema A.G. (2012). Empirical assessment of methods for risk identification in healthcare data: Results from the experiments of the Observational Medical Outcomes Partnership. *Statistics in Medicine* **31**, 4401-15.
- Simpson S.E. (2013). A positive event dependence model for self-controlled case series with applications in postmarketing surveillance. *Biometrics*, **69**, 128-136.
- Steer C.D., Macleod J., Tilling K., Lim A.G., Marsden J, Millar T., Strang J., Telfer M., Whitaker H., Vickerman P. and Hickman M. Chapter 8. In: *The impact of opiate substitution treatment on mortality in the UK: a natural experiment using the Clinical Practice Research Datalink. Health Services and Delivery Research*, **6**, to appear.
- Tse A., Fu Tseng H., Greene S.K., Vellozzi C., Lee G.M. (2012). Signal detection and evaluation for risk of febrile seizures in children following trivalent inactivated influenza vaccine in the Vaccine Safety Datalink Project, 2010-2011. *Vaccine* **30**, 2024-2031.
- Turner H. and Firth D. (2015). Generalized nonlinear models in R: An overview of the gnm package. R package version 1.0-8. <http://CRAN.R-project.org/package=gnm>
- Verity C.M., Butler N.R. and Golding J. (1985). Febrile convulsions in a national cohort followed up from birth. I - Prevalence and recurrence in the first five years of life. *British Medical Journal* **290**, 1307-1310.
- Weldeselassie Y.G., Whitaker H.J., Farrington C.P. (2011). Use of the self-controlled case-series method in vaccine safety studies: review and recommendations for best practice. *Epidemiology and Infection* **139**, 1805-1817.
- Whitaker H.J., Farrington C.P., Spiessens B., Musonda P. (2006). Tutorial in biostatistics: The self-controlled case series method. *Statistics in Medicine* **25**, 1768-1797.
- Whitaker H.J., Ghebremichael-Weldeselassie Y., Douglas I.J., Smeeth L. and Farrington C.P. (2018). Investigating the assumptions of the self-controlled case series method. *Statistics in Medicine* **37**, 643-658.
- Wong A.Y., Root A., Douglas I.J., Chui C.S., Chan E.W., Ghebremichael-Weldeselassie Y., Siu C.W., Smeeth L., Wong I.C. (2016). Cardiovascular outcomes associated with use of clarithromycin: population based study. *British Medical Journal* **352**, h6926.