Monitoring vaccine safety using case series CUSUM charts

Patrick Musonda[1], Mounia N. Hocine[1],
Nick J. Andrews[2], C. Paddy Farrington[1]

[1]*Department of Mathematics and Statistics, The Open University, Milton Keynes, MK7 6AA, UK*
[2]*Statistics,Modelling and Bioinformatics Department, Centre for Infections, London, NW9 5EQ, UK*

Abstract

We adapt the self-controlled case series method for long-term surveillance of vaccine safety using cumulative sum (CUSUM) charts. The CUSUM surveillance method we propose is applicable for detecting associations that arise in a short pre-determined risk period following vaccination. The performance of the case series CUSUM is investigated through simulations. We illustrate the method using retrospective analyses of influenza vaccine and Bell's Palsy, and MMR vaccine and febrile convulsions.

Keywords: adverse event, self-controlled case series method, CUSUM, average run length, surveillance, vaccine.

Corresponding Author: Paddy Farrington, Department of Mathematics and Statistics, The Open University, Milton Keynes MK7 6AA, UK
Email: c.p.farrington@open.ac.uk

This research was supported with funding from the MRC, the EPSRC and the Fondation pour la Recherche Médicale.

1 Introduction

Ongoing surveillance of vaccine safety is important, for two reasons. First, it is desirable to identify hitherto undetected problems or new problems, for example resulting from changes in the manufacture of vaccines in routine use. Second, it is useful to build up a picture of a vaccine's safety profile over the long term, if only to accumulate evidence that the vaccine is safe.

Methods in common use for pharmacovigilance include voluntary notification systems such as the Yellow Card system in the UK, run by the Medicines and Healthcare products Regulatory Agency, and the Vaccine Adverse Event Reporting System (VAERS) in the United States [1]. However, because notification of adverse events is usually triggered by the suspicion that they are related to the vaccine, the interpretation of data from such systems is not straightforward [2-4], nor can they be used convincingly to provide evidence of safety.

In this paper we apply the self-controlled case series method [5] to obtain a cumulative sum (CUSUM) chart to monitor paediatric vaccine safety. This can be seen as a version of the sequential probability ratio test (SPRT), which we recently adapted for use with the self-controlled case series method [6]. The key difference is that whereas the SPRT is best used for monitoring a new vaccine, with the aim of deciding rapidly whether the vaccine is safe or otherwise with respect to events that cannot realistically be assessed in clinical trials, the CUSUM is appropriate for long-term surveillance of a vaccine for which the presumption is that it is safe.

The paper is arranged as follows. In section 2 we briefly give some background on the self-controlled case series method and CUSUM charts. In section 3 we describe the self-controlled case series CUSUM. We present some simulation results in Section 4, and applications in Section 5. We conclude with a discussion and recommendations in Section 6.

2 Background

2.1 The self-controlled case series method

The self-controlled case series method, or case series method for short, provides an alternative to more established cohort or case-control methods for investigating the association between a time-varying exposure and an outcome event [5 - 9]. The method has been widely used in pharmacoepidemiology, particularly in the study of vaccine safety. The main features of this method is that it uses cases only, implicitly adjusts for fixed confounders, and in some circumstances it has high efficiency relative to the cohort method [7]. A modelling guide, including details of software implementations, may be found in reference [9] and from the self-controlled case series website: http://statistics.open.ac.uk/sccs.

The basic idea of the method is as follows. Based on prior medical knowledge, or prior data, it is assumed that the incidence of the adverse event of interest is increased by a multiplicative factor ρ during a pre-specified time period after vaccination, known as the risk period. Outside the risk period, the vaccine has no effect on the incidence. Thus, ρ is the relative incidence associated with vaccination. Individuals

are observed over a pre-specified observation period, which is split up into successive intervals indexed by age group and level of vaccine-associated risk. The case series log-likelihood is derived by conditioning on the total number of events observed for each individual: it is this conditioning step that makes it possible to consider only cases, that is, individuals having experienced one or more events, vaccinated or otherwise. Conditioning also induces the self-matched property of the method whereby all fixed confounders, whether measured or not, are necessarily controlled in the analysis. The case series analysis provides estimates of the log relative incidence and of age effect parameters that determine how the incidence of the adverse event varies with age.

2.2 The cumulative sum (CUSUM) chart

The CUSUM procedure is a well-known sequential monitoring method, introduced by Page[10] and based on Wald's sequential probability ratio test [11]. The CUSUM is used to monitor an ongoing process governed by some parameter β. When all is well, $\beta = 0$ and the process is said to be in control. The basic idea of the CUSUM is to accumulate evidence against the null hypothesis $\beta = 0$, in such a way that evidence in favour of the null hypothesis is not allowed to build up. Thus, the CUSUM remains sensitive to changes in the process which might cause the value of β to increase. The accumulating evidence is represented by points on a time chart: once the trace crosses a pre-determined boundary, the process is deemed to be out of control and, if required, corrective measures are taken to bring it back under control, after which the monitoring sequence begins afresh.

We shall only consider one-sided CUSUMs, that is, charts with a single boundary. Barnard [12] developed the V-mask form of the CUSUM to allow detection of either an increase or a decrease of the parameter β of interest; further extensions are discussed in [13]. The initial use of the CUSUM was primarily in industrial applications, though more recently, CUSUMs have been used in a medical context [14-23], to monitor disease incidence and institutional performance. In the latter context there is particular interest in concurrent monitoring of large numbers of units, which poses particular problems of false discovery [24-26].

3 The case series CUSUM

We consider acute adverse events, potentially associated with vaccination during a pre-determined risk period of duration d. The incidence of the events may vary with age, and this is allowed for by assuming that the incidence is constant within pre-defined age groups. Data on cases and vaccination are collected sequentially at time intervals of length s, typically every 6 months or every year; these are the monitoring intervals. At the kth analysis starting from time zero, all cases occurring between times $(k-1) \times s$ and time $k \times s$ are identified, and their vaccination history is determined.

These cases are then used to obtain the case series loglikelihood ratio statistic Λ_k, for the test of H_1: $\beta = \log(\rho_A)$, where ρ_A is the pre-specified relative incidence to be detected ($\rho_A > 1$), against H_0: $\beta = 0$, corresponding to a relative incidence equal to 1. Two methods can be used to handle age effects. First, if a long series of historical data

are available, then age effects can be estimated accurately and these values substituted into the likelihood and treated as known, fixed values. This introduces bias if the age effects are estimated with error [27]. Alternatively, if the age effects are not known, they can be profiled out, and the loglikelihood is replaced by a profile loglikelihood in order to eliminate the age parameters at each monitoring interval. The procedure is described in detail in [6].

The CUSUM value Z_k at step k is then defined as follows:

$$Z_k = \max\{Z_{k-1} + \Lambda_k, 0\}, \quad Z_0 = 0.$$

The CUSUM chart is obtained by plotting Z_k against k. Note that the case series CUSUM is group-sequential, that is, it is updated at the end of each monitoring interval, rather than as new cases arise. This is essential since the case series method is retrospective. Note also that the case series CUSUM is risk-adjusted: all fixed confounders are necessarily adjusted owing to the fact that a case series likelihood has been used, and age effects have been allowed for by one of the two methods described above. These methods of risk adjustment have the major advantage of retaining a simple recursive form for Z_k, unlike some other likelihood-based methods [13]. Also, under the null hypothesis, the increments Λ_k are independently and identically distributed.

The CUSUM differs from the sequential probability ratio test (SPRT) [6] because it has a holding barrier at zero rather than a lower absorbing barrier. A pre-determined upper boundary value h, known as the control limit, must also be specified. The process is deemed to be in state H_0, or 'in control', at step k if $Z_k < h$, and to have shifted to state H_1, and hence be 'out of control', if $Z_k \geq h$. A CUSUM that exceeds the control limit is said to have 'signalled'. This means that the chart has accumulated enough evidence to conclude that the value of β has increased. At this point, remedial action is taken, and the CUSUM is restarted from zero.

If left running long enough, a CUSUM chart will eventually signal even when it is in control. The performance of the CUSUM is traditionally measured using the average time to the first signal, known as the average run length or ARL. Ideally, the ARL should be short when the process is out of control from the start (this is the ARL under H_1, or ARL_1), and it should be long when the process is in control (ARL under H_0, or ARL_0). In the context of vaccine safety surveillance, a low value of ARL_1 is essential.

4 Simulations

The performance of the case series CUSUM depends on several parameters, including the underlying incidence of the adverse events, the age effect, the length of the risk period, the length of the monitoring interval, the choice of the design value of the relative incidence to be detected ρ_A, the choice of control limit h, the true relative incidence, and the annual incidence of the event λ. We investigate these dependences by simulation.

4.1 Design of the simulation study

We used the same design for the simulation as used for an earlier study of the SPRT [6]; see this paper for further details. Briefly, we consider a single vaccine and adverse events arising within one year of age, which for definiteness we chose to be the period 0 to 364 days of age. We assumed that the distribution of age at vaccination in the population was proportional to a gamma density with mean 120 days, shape parameter 1.2 and probability of remaining unvaccinated at 364 days equal to 0.1. The variable parameters in the simulations were as follows:

Design relative incidence ρ_A:	1.5, 2, 3, 5, 10
Monitoring interval length s:	0.25, 0.5, 1 year
Risk period length d:	1, 2, 4 weeks
Annual incidence of the event λ:	25, 50, 100 cases per year
Control limit h:	1, 2, 3
Age effect:	constant, increasing, decreasing
Age adjustment:	profiled out, effect known.

The age effect was obtained by assuming that the incidence was constant within thirteen 28-day intervals, varying by the factor 1, 1.2, or 1/1.2 in successive age groups for the three scenarios considered. The two age adjustment methods correspond to profiling out the (unknown) age effect as described above, and substituting known values of the age effect (as might be obtained from a long run of historical data). For out-of-control simulations, the true relative incidence was set equal to ρ_A. Two hundred simulations were run for each combination of the parameters investigated; this is sufficient to explore the properties of the system and its dependence on the different parameters, without being overly computationally demanding.

4.2 Results: baseline scenario

The baseline scenario corresponds to the following combination of parameters: 0.5 year monitoring interval, 2 week risk period, annual incidence 50 cases. The results for constant, increasing and decreasing age effects are shown in Tables 1, 2 and 3, respectively. As expected, both the average run length and the standard deviation of the run length (in and out of control) vary according to the age effect. Whatever the age effect, adjusting for it using profile likelihood methods has little impact on the average run length, compared to the situation in which the age effect is known exactly. Thus, this method of adjustment appears to be adequate.

The relationship between the average run length in control and the alternative hypothesis value of the relative incidence, ρ_A, is a complex one: unbounded when $\rho_A = 1$, ARL_0 decreases as ρ_A increases above 1, reaches a minimum, and thereafter increases to infinity. This is because the log likelihood ratio statistic Λ_k is a nonlinear function of ρ_A. On the other hand, the ratio of the average run lengths in and out of control, ARL_0/ARL_1, increases with ρ_A. As expected, ARL_0, ARL_1 and their ratio all increase as the control limit h increases. Note that the run lengths are all necessarily bounded below by the monitoring interval length.

4.3 Results: other scenarios

We investigated the performance of the system with events of annual incidence 25 and 100 cases. Figures 1 and 2 summarizes the results for 25 events expected per annum, with the age effect profiled out. As expected, average run lengths are longer than for the baseline scenario, and increase very rapidly with the control limit h. Conversely, for annual incidence 100 cases, average run lengths are shorter (data not shown).

Next, we varied the risk period using 1 week and 4 weeks, compared to 2 weeks for the baseline scenario (data not shown). The average run lengths and standard deviations decrease as the length of the risk period increases. This is as expected, since the information about the exposure effect is greatest when the expected number of events in and outside the risk period are similar [28].

We also studied the effect of varying the length s of the monitoring interval, using 3 months and 1 year, compared to 6 months for the baseline scenario (data not shown). The dependence of the ARL on the length of the monitoring interval is complex. On the one hand, reducing s reduces the lower bound on the ARL and hence allows shorter values to arise. On the other hand, the ratio of the risk period to the observation period increases, and the proportion of cases vaccinated decreases when s decreases, both of which reduce the information about the vaccine effect.

4.4 Surveillance of several vaccines

So far we have only considered surveillance of a single vaccine. In practice, we might wish to monitor several vaccines, and/or several adverse events. Note that the number of vaccines and events to be monitored is not expected to be extremely large, but might perhaps include up to 5 or 10 different combinations.

Monitoring multiple vaccine/event combinations (referred to as units) increases the rate at which spurious signals, also called false discoveries, occur. The relevant average run length in control is no longer the ARL_0 for an individual vaccine, but the average time to the first signal from any of the units monitored, which we denote $SARL_0$ (system average run length). (The average run length out of control remains the same as before, assuming that at most one unit is out of control at any one time, a reasonable assumption in the present context.) We obtained the values of $SARL_0$ using the fact that, for unadjusted CUSUMs, the distribution of the run length in control is approximately geometric, the approximation improving as ARL_0 increases [29]. As the increments Λ_k for the profile likelihood case series CUSUM are identically and independently distributed, similar results apply here. Figure 3 illustrates this, using some of the simulations presented in Table 2. Goodness of fit tests yielded chi-squared values of 6.84 on 7 degrees of freedom ($p = 0.45$) for $\rho_A = 2$, $h = 2$, and 7.25 on 6 degrees of freedom ($p = 0.30$) for $\rho_A = 1.5$, $h = 3$.

The relationship between the system average run length (in years) for r vaccines in control, and the ARL_0 (in years) for each individual vaccine (assumed to be the same for simplicity), for monitoring intervals of duration s years is:

$$SARL_0 = \frac{s}{1-(1-s/ARL_0)^r}.$$

Table 4 shows the values of $SARL_0$ for r = 1, 2, 5 and 10, based on the data from Table 1. Different parameter combinations produce results with a similar pattern and so are not presented here. The more units are monitored, the shorter is the system average run length and the lower is the ratio $SARL_0/ARL_1$. The practical implications of these findings are discussed in Section 6.

5 Examples

In the following examples we present data in two contrasting situations. The control limit h = 3 is used throughout.

5.1 Bell's palsy and flu vaccine

Concern about the possible association between Bell's palsy, an acute facial paralysis affecting the 7th facial nerve, and some influenza vaccine formulations, was raised in Switzerland in October 2000, after the introduction of an inactivated nasal formulation of the influenza vaccine. In this case series study, the relative incidence within the 31-60 day post-vaccination risk period was estimated to be 35.6, 95% CI (14.1-89.8). A similar analysis was undertaken using data from the GPRD on 2263 episodes of Bell's palsy in the UK recorded from July 1st 1992 to 30th June 2005. The estimated relative incidence in the 3 months following parenteral inactivated influenza vaccine was not significant: RI = 0.92, 95% CI (0.78-1.08) [30].

We reanalysed these UK data using the case series CUSUM. We used a six month monitoring interval, 1-60 day risk period after any dose of influenza vaccine, with alternative hypothesis relative incidence $\rho_A = 1.5$ (Figure 4). In view of possible temporal confounding from the highly seasonal administration of influenza vaccine, the analysis was performed using a parametric case series model with 12 one month seasonal periods. Figure 4 shows that the CUSUM value remains less than 3 . Thus, in this retrospective analysis over 13 years, the CUSUM remains in control throughout.

5.2 MMR vaccine and convulsions

We reanalysed data from a study of convulsions after measles, mumps and rubella (MMR) vaccine in children in the UK [31]. Data were recorded in 418 children aged 12-24 months. This age group was chosen to cover the period in which most MMR vaccinations are administered. The estimated relative incidence of convulsions in the 6-11 day post-vaccination risk period using a case series analysis was 6.26, 95% CI (3.85-10.18) consistent with the known effect of the measles component of MMR vaccine.

We reanalyzed these data as if undertaking prospective monitoring using the case series CUSUM. We used a six month monitoring interval. Figure 5 shows the CUSUM graph for the alternative hypothesis relative incidence $\rho_A = 3$ against a null exposure effect. The control limit h = 3 is crossed in the first monitoring interval, indicating that the system is not in control. This is as expected: the association

between convulsions and live measles vaccine in the second week after vaccination is well-known. In Figure 5 the CUSUM has not been reset: the monotone increase over subsequent monitoring intervals provides further evidence that the CUSUM is not in control.

Since it is well-known that live measles vaccines cause convulsions with a relative incidence of the order of 2 – 4 , it makes sense to redefine the null hypothesis as corresponding to $\rho = 3$, say, and set the alternative hypothesis value $\rho_A = 5$, say, since interest resides not in detecting well-known associations but in finding evidence of changes for the worse. The resulting CUSUM (not reset) is shown in Figure 6. The chart provides some evidence that, after year 2, the in-control state $\rho = 3$ has shifted to $\rho = 5$. Also shown in Figure 6 is the CUSUM with $\rho = 4$ under the null and the alternative hypothesis value $\rho_A = 6$, which does not indicate a shift. This serves to illustrate an important feature with CUSUM charts, namely their sensitivity to the choice of hypotheses. Thus, if the true value of ρ when out of control is lower (respectively, greater) than ρ_A, then the average run length is greater (respectively, lower) than the ARL_1 for the alternative hypothesis $\rho = \rho_A$.

6 Discussion

The discussion is structured in two parts. In the first, we review the strengths and weaknesses of the methods described, and briefly discuss some alternatives. In the second, we present some recommendations for vaccine surveillance systems, based on the results obtained above.

6.1 Strengths, weaknesses and alternative approaches

A long-term monitoring system based on the self-controlled case series CUSUM has all the advantages and disadvantages of the case series method. Its main advantage is that it uses only cases, and hence, like the SPRT for surveillance of a new vaccine [6], can be applied to data from population and hospital-based reporting systems such as the General Practice Research Database and Hospital Episode Statistics database. A further advantage is that it controls for fixed confounders, and hence allows for more robust evaluation of potential associations than methods based on spontaneous adverse event reporting systems. Finally, as we have shown, control of age effects via profile likelihood methods works well; further work is warranted to study the properties of this scheme, in particular to obtain approximations to the average run lengths.

The weaknesses are those of the case series method [5,8], in particular the requirement that events should not influence subsequent exposures. A further limitation, specific to the use of the method in surveillance, is that the risk period should be short in relation to the monitoring interval. The method is thus wholly inappropriate for surveillance of reactions that may occur a long time after vaccination.

Unlike surveillance of a new vaccine, in which it is important rapidly to establish evidence of safety as well as identify problems should they occur, long-term surveillance of an established vaccine is based on the strong presumption that it is

safe. Thus, the CUSUM provides a more appropriate framework than the SPRT [6]. A further contrast with surveillance of new vaccines is that it might be desirable concurrently to monitor several vaccines or adverse events. As shown in Subsection 4.3, surveillance of several vaccines greatly reduces the average time to a false signal, as measured by the system average run length in control, $SARL_0$. This in turn will increase the false discovery rate (FDR), namely the expected proportion of signals that are genuine [32].

The problem of controlling the FDR when monitoring many units has been considered by several authors. Marshall et al [24] seek to choose h so as to control the FDR over a given time period. Others have applied the methods of [32] more directly to p-values, derived as inverse average run lengths [25] or from the null distribution of the CUSUM statistic [26].

In the setting considered in the present paper, obtaining such p-values is difficult and probably impractical, owing to the non-standard form of the case series log-likelihood. We do not envisage very large numbers of items to be monitored, so that the problem of false discovery, while remaining important, is not quite as overwhelming as in some other applications. Also, we expect the number of out-of control units to be very low, so that the power advantage of FDR-control methods over Bonferroni adjustments is limited [32]. Thus, we propose a more ad-hoc approach, informed by simulations of $SARL_0$ and ARL_1 and the ratio $SARL_0/ARL_1$. This is to base the choice of control limit h on a system average run length in control that produces a manageable number of signals, while retaining a low value of ARL_1, a key requirement of vaccine safety surveillance.

A further limitation of the CUSUM described here is that it does not provide an estimate of the relative incidence ρ. An alternative approach is to replace the profile likelihood under the alternative hypothesis $\rho = \rho_A$ by a full maximum likelihood, yielding an estimate of ρ at each monitoring interval. A further option is to use generalized likelihood ratio methods [13]. These would yield more robust estimates, though at additional computational cost. However, upon signalling, the interpretation of estimates obtained sequentially in these ways is not straightforward.

6.2 Recommendations

We recommend the use of a 6-month or 1-year monitoring interval, with a preference for the latter when monitoring very rare events and when obtaining data is costly or time-consuming. We found that reducing the monitoring interval below 6 months provided little gain in sensitivity except for very common events.

Age effects may be adjusted explicitly using profile likelihood methods. The simulations show that age effects do affect the results and hence should be controlled for, and that profiling them out incurs very little if any cost in timeliness of detection.

The choice of alternative hypothesis relative risk to be detected ρ_A is important, and also tricky, as shown by the examples in Section 5. The simulations suggest that, for events occurring at a rate of 50 per year, the ARL_1 for values of less than 3 can be in excess of 5 or even 10 years, depending on the age effect; such times to detection are far too long to be of any practical use. We therefore recommend that ρ_A is set no

lower than 3, unless the catchment area of the system can be expanded to ensure that events occur with an annual frequency of at least 100 cases.

We also recommend that a control limit $h \geq 3$ be used. The value $h = 3$ does not substantially affect the value of ARL_1 for ρ_A equal to 3 or more compared to $h = 1$ or 2, but greatly increases the value of ARL_0. With these choices, for adverse events occurring at the rate of 50 cases per year, the ARL_1 for $\rho_A = 3$ is about 2 years whereas ARL_0 is about 70 years.

The choices $h = 3$ and $\rho_A = 3$ may not yield an adequate surveillance mechanism for 5 or more vaccines (or vaccine and event combinations), in view of the increased frequency of false discoveries: with these values, the $SARL_0$ drops from 70.4 years to 14.3 years as the number of units monitored increases from 1 to 5. In this case, the control limit h may need to be increased. Since new safety problems with established vaccines are uncommon, it is unlikely that more than one unit will be out of control at any one time. In this case, a Bonferroni-type adjustment is adequate. The control limit h to be used should be determined by simulation, with values chosen to reflect the frequencies and age distributions of the events to be monitored.

The choice of the control limit h is to some extent determined by the action to be taken if the CUSUM signals. We envisage that, upon the CUSUM signalling, further investigations will be undertaken, for example to review changes in vaccine manufacture, to implement more focused surveillance, and to conduct other epidemiological studies, in particular to estimate the relative incidence.

References

[1]. Chen RT, Rastogi SC, and Mullen JR, The Vaccine Adverse Event Reporting System (VAERS). *Vaccine*, 1994. **12**: p. 542-550.

[2]. Evans SJW, Waller PC, and Davis S, Use of proportional reporting ratios (PRRs) for signal generation from spontaneous adverse drug reaction reports. *Pharmacoepidemiology and drug safety*, 2001. **10**: p. 483-486.

[3]. Rothman KJ, Lanes S, and Sacks ST, The reporting odds ratio and its advantages over the proportional reporting ratio. *Pharmacoepidemiology and drug safety*, 2004. **13**: p. 519-523.

[4]. Waller PC, Van Puijenbroek EP, Egberts ACG, and Evans SJW, The reporting odds ratio versus the proportional reporting ratio: 'deuce'. *Pharmacoepidemiology and drug safety*, 2004. **13**: p. 525-526.

[5]. Farrington CP, Relative incidence estimation from case series from vaccine evaluation. *Biometrics*, 1995. **51**: p. 228-235.

[6]. Hocine M, Musonda P, Farrington CP, and Andrews N, Sequential case series analysis for pharmacovigilance. *Tech. rept. available from http://statistics.open.ac.uk/sccs/JRSSA_submitted.pdf*.

[7]. Farrington CP, Nash J, and Miller E, Case series analysis of adverse reactions to vaccines: a comparative evaluation. *American Journal of Epidemiology*, 1996. **143**: p. 1165-1173 (Erratum 1998; **147**:93).

[8]. Farrington CP and Whitaker HJ, Semiparametric analysis of case series data. *Applied Statistics*, 2006. **55**: p. 553-594.

[9]. Whitaker HJ, Farrington CP, Spiessens B, and Musonda P, Tutorial in Biostatistics: The self-controlled case series method. *Statistics in Medicine*, 2006. **25**: p. 1768-1797.
[10]. Page ES, Continuous inspection schemes. *Biometrika*, 1954. **41**: p. 100-115.
[11]. Wald A, Sequential analysis. 1947, New York: Wiley.
[12]. Barnard GA, Control charts and stochastic processes (with discussion). *Journal of the Royal Statistical Society, Series B*, 1959. **21**: p. 239-271.
[13]. Lai TL, Sequential changepoint detection in quality control and dynamical systems (with discussion). *Journal of the Royal Statistical Society, Series B*, 1995. **57**: p. 613-658.
[14]. Hutwagner LC, Maloney EK, Bean NH, Slutsker L, and Martin SM, Using Laboratory-Based Surveillance Data for Prevention: An Algorithm for Detecting *Samonella* Outbreaks. *Emerging Infectious Diseases*, 1997. **3**(3): p. 395-400.
[15]. Grigg OA, Farewell VT, and Speigelhalter DJ, Use of risk-adjusted CUSUM and RSPRT charts for monitoring in medical contexts. *Statistical Methods in Medical Research*, 2003. **12**: p. 147-170.
[16]. Steiner SH, Geyer PL, and Wesolowsky GO, Grouped Data-Sequential Probability Ratio Tests and Cumulative Sum Control Charts. *Technometrics*, 1996. **38**(3): p. 230-237.
[17]. Gallus G, Mandelli C, Marchi M, and Radeaelli G, On surveillance methods for congenital malformations. *Statistics in Medicine*, 1986. **5**: p. 565-571.
[18]. Nix AB, Rowlands RJ, and Kemp KW, Internal quality control in clinical chemistry: a teaching review. *Statistics in Medicine*, 1986. **6**: p. 425-440.
[19]. Tillett HE and Sepencer IL, Influenza surveillance in England and Wales using routine statistics. Development of 'cusum' graphs to compare 12 previous winters and to monitor the 1980/81 winter. *J.Hyg., Camb*, 1982. **88**: p. 83-94.
[20]. Williams SM, Parry BJ, and Schlup MM, Quality control: an application of the CUSUM. *British Medical Jounal*, 1992. **304**: p. 1359-1361.
[21]. DeLeval, Marc R, François K, Bull C, Brawn WB, and Speigelhalter DJ, Analysis of a cluster of surgical failures. *The Journal of Thoracic and Cardiovascular Surgery*, 1994. **104**: p. 914-924.
[22]. Steiner SH, Cook RJ, and Farewell VT, Monitoring paired binary surgical outcomes using cumulative sum charts. *Statistics in Medicine*, 1999. **18**: p. 69-86.
[23]. Ross G, Lampugnani L, and Marchi M, An approximate CUSUM procedure for surveillance of health events. *Statistics in Medicine*, 1999. **18**: p. 2111-2122.
[24]. Marshall C, Best N, Bottle A, and Aylin P, Statistical issues in the prospective monitoring of health outcomes across multiple units. *J R Statist Soc.A*, 2004. **167**(Part 3): p. 541-559.
[25]. Benjamini Y and Kling Y, A look at statistical process control through the p-values. *Tech. rept. RP-SOR-99-08. Tel Aviv University, Israel* , 2005.
[26] Grigg O and Spiegelhalter D, The null steady-state distribution of the CUSUM statistic. Tech. rept.
[27] Jensen WA, Jones-Farmer LA, Champ CW and Woodall WH. Effects of parameter estimation on control chart properties: a literature review. *Journal of Quality Technology*, 2006. **38**(4): p. 349-364.

[28] Musonda P, Mounia NH, Whitaker HJ, and Farrington CP, Self-controlled case series analyses: small sample performance. *Computational Statistics and Data Analysis*, 2007: p. doi:10.1016/j.csda.2007.06.016.
[29]. Brook D and Evans DA, An approach to the probability distribution of cusum run length. *Biometrika*, 1972. **3**: p. 539-549.
[30]. Stowe J, Andrews N, and Wise L and Miller E, Bell's palsy after parenteral inactivated influenza vaccine. *Human Vaccines*, 2006. **2**: p.110-112.
[31] Miller E, Andrews N, Stowe J, Grant A, Waight P and Taylor B, Risks of convulsion and aseptic meningitis following measles-mumps-rubella vaccination in the United Kingdom. *American Journal of Epidemiology*, 2007. **165**(6): p.704-709.
[32] Benjamini Y and Hochberg Y, Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. B*, 1995. **57**(1): p.289-300.

		Age effect profiled out		Age effect known			Age effect profiled out		Age effect known		Age effect profiled out	Age effect known
h	ρ_A	Run length in control (years)		Run length in control (years)			Run length out of control (years)		Run length out of control (years)		$\frac{ARL_0}{ARL_1}$	$\frac{ARL_0}{ARL_1}$
		ARL_0	Std. dev	ARL_0	Std. dev		ARL_1	Std. dev	ARL_1	Std. dev		
1	1.5	13.5	12.6	13.7	12.7		4.64	3.73	4.60	3.69	2.91	2.98
	2	8.71	8.65	8.73	8.69		2.05	1.53	2.05	1.54	4.25	4.26
	3	7.39	6.91	7.48	6.97		1.13	0.79	1.11	0.76	6.54	6.74
	5	9.42	8.80	9.61	9.00		0.73	0.38	0.73	0.40	12.9	13.2
	10	18.3	19.1	18.7	19.5		0.55	0.20	0.55	0.20	33.3	34.0
2	1.5	67.7	67.2	67.8	67.5		10.2	7.80	10.1	7.70	6.64	6.71
	2	32.7	30.7	32.8	30.7		3.82	2.77	3.79	2.77	8.56	8.66
	3	20.6	20.8	20.6	20.8		1.62	1.05	1.61	1.06	12.7	12.8
	5	23.4	23.2	23.5	23.3		0.86	0.51	0.85	0.51	27.2	27.6
	10	43.9	40.2	44.3	40.6		0.56	0.19	0.58	0.22	78.4	76.4
3	1.5	229.3	208.6	229.7	208.9		16.2	10.5	15.9	10.5	14.2	14.4
	2	109.3	108.2	109.6	108.3		5.55	3.90	5.52	3.93	19.7	19.9
	3	70.4	60.4	70.6	65.5		2.20	1.36	2.17	1.37	32.0	32.5
	5	68.3	75.8	68.5	76.1		1.05	0.57	1.03	0.57	65.0	66.5
	10	108.7	108.8	109.0	109.1		0.57	0.20	0.58	0.23	190.7	187.9

Table 1. Six months monitoring interval, two weeks risk period, 50 cases per year and constant age effect. $h =$ control limit, $\rho_A =$ relative incidence we wish to detect, $ARL_0 =$ average run length in control and $ARL_1 =$ average run length out of control.

		Age effect profiled out		Age effect known			Age effect profiled out		Age effect known		Age effect profiled out	Age effect profiled out
h	ρ_A	Run length in control (years)		Run length in control (years)			Run length out of control (years)		Run length out of control (years)		$\frac{ARL_0}{ARL_1}$	$\frac{ARL_0}{ARL_1}$
		ARL_0	Std. dev	ARL_0	Std. dev		ARL_1	Std. dev	ARL_1	Std. dev		
1	1.5	21.4	21.0	21.3	21.0		7.24	5.99	7.14	5.96	2.96	2.98
	2	11.3	10.6	11.2	10.5		3.05	2.36	2.98	2.36	3.71	3.76
	3	9.52	8.60	9.58	8.70		1.42	1.01	1.40	1.02	6.70	6.84
	5	8.22	7.19	8.42	7.40		0.86	0.53	0.86	0.55	9.56	9.79
	10	11.0	8.90	11.2	9.07		0.58	0.21	0.58	0.19	19.0	19.3
2	1.5	97.6	90.4	97.7	90.6		16.5	12.5	16.2	12.5	5.92	6.03
	2	47.6	42.8	47.6	42.9		5.98	4.28	5.88	4.25	7.96	8.10
	3	28.5	29.9	28.7	30.1		2.29	1.60	2.26	1.60	12.4	12.7
	5	24.8	24.0	24.9	24.1		1.10	0.71	1.10	0.74	22.5	22.6
	10	31.3	28.0	31.6	28.2		0.64	0.29	0.65	0.29	48.9	48.6
3	1.5	314.6	328.7	314.3	328.4		24.6	16.0	24.1	16.0	12.8	13.0
	2	152.4	157.2	152.5	157.5		8.79	6.08	8.68	5.98	17.3	17.6
	3	77.4	75.3	77.7	75.6		3.59	2.56	3.52	2.57	21.6	22.1
	5	68.9	72.4	69.1	72.8		1.44	0.86	1.41	0.88	47.8	49.0
	10	86.5	88.3	86.9	88.7		0.71	0.36	0.70	0.34	121.8	124.1

Table 2. Six months monitoring interval, two weeks risk period, 50 cases per year and increasing age effect. $h =$ control limit, $\rho_A =$ relative incidence we wish to detect, $ARL_0 =$ average run length in control and $ARL_1 =$ average run length out of control.

		Age effect profiled out		Age effect known			Age effect profiled out		Age effect known		Age effect profiled out	Age effect profiled out
h	ρ_A	Run length in control (years)		Run length in control (years)			Run length out of control (years)		Run length out of control (years)		$\frac{ARL_0}{ARL_1}$	$\frac{ARL_0}{ARL_1}$
		ARL_0	Std. dev	ARL_0	Std. dev		ARL_1	Std. dev	ARL_1	Std. dev		
1	1.5	13.7	12.8	13.7	12.7		3.99	3.57	3.91	3.51	3.43	3.50
	2	7.50	7.05	7.52	7.04		1.58	1.14	1.55	1.08	4.75	4.85
	3	7.09	7.00	7.18	7.06		0.88	0.58	0.88	0.60	8.06	8.16
	5	10.4	11.5	10.5	11.5		0.66	0.32	0.68	0.39	15.8	15.4
	10	26.8	25.8	27.1	26.1		0.51	0.08	0.52	0.09	52.5	52.1
2	1.5	53.2	50.2	53.1	50.1		7.96	6.05	7.92	6.14	6.68	6.71
	2	33.6	31.6	33.7	31.6		3.31	2.86	3.29	2.85	10.2	10.2
	3	22.2	20.9	22.3	20.9		1.35	0.86	1.33	0.91	16.4	16.8
	5	27.3	24.8	27.4	25.0		0.72	0.32	0.74	0.35	37.9	37.0
	10	62.5	57.6	62.9	58.0		0.54	0.14	0.54	0.14	115.7	116.5
3	1.5	164.1	174.5	164.3	174.6		11.4	6.82	11.3	6.61	14.4	14.5
	2	98.7	89.5	98.7	89.4		4.38	3.06	4.36	3.03	22.5	22.6
	3	61.5	61.5	61.6	61.6		1.69	1.00	1.67	1.00	36.4	36.9
	5	66.6	60.6	66.8	60.7		0.86	0.46	0.86	0.46	77.4	77.7
	10	130.6	132.1	130.9	132.4		0.56	0.16	0.56	0.16	233.2	233.8

Table 3. Six months monitoring interval, two weeks risk period, 50 cases per year and decreasing age effect. $h =$ control limit, $\rho_A =$ relative incidence we wish to detect, $ARL_0 =$ average run length in control and $ARL_1 =$ average run length out of control.

	e^{β_1}	$SARL_0$				ARL_1	$SARL_0 / ARL_1$			
h		r = 1	r = 2	r = 5	r = 10		r = 1	r = 2	r = 5	r = 10
1	1.5	13.5	6.88	2.91	1.59	4.64	2.91	1.48	0.63	0.34
	2	8.71	4.48	1.95	1.12	2.05	4.25	2.19	0.95	0.55
	3	7.39	3.82	1.69	0.99	1.13	6.54	3.38	1.50	0.88
	5	9.42	4.84	2.09	1.19	0.73	12.9	6.63	2.87	1.63
	10	18.3	9.28	3.87	2.07	0.55	33.3	16.9	7.03	3.76
2	1.5	67.7	34.0	13.7	7.00	10.2	6.64	3.33	1.35	0.69
	2	32.7	16.5	6.74	3.50	3.82	8.56	4.31	1.77	0.92
	3	20.6	10.4	4.32	2.30	1.62	12.7	6.44	2.67	1.42
	5	23.4	11.8	4.88	2.57	0.86	27.2	13.8	5.68	2.99
	10	43.9	22.1	8.98	4.62	0.56	78.4	39.4	16.04	8.25
3	1.5	229.3	114.8	46.1	23.2	16.2	14.2	7.09	2.84	1.43
	2	109.3	54.8	22.1	11.2	5.55	19.7	9.87	3.98	2.01
	3	70.4	35.3	14.3	7.27	2.20	32.0	16.1	6.49	3.30
	5	68.3	34.3	13.9	7.06	1.05	65.0	32.6	13.2	6.72
	10	108.7	54.5	21.9	11.1	0.57	190.7	95.6	38.5	19.47

Table 4. System average run length in control ($SARL_0$) for 1, 2, 5 and 10 units, average run length out of control (ARL_1), and ratio, for constant age effect profiled out, 0.5 year monitoring interval, two weeks risk period, and 50 cases per year. All run lengths are in years.

Captions for figures

Figure 1. Average run lengths in control by control limit h (see text for details). Age effect constant: ----, increasing: ____, decreasing:…… .

Figure 2. Average run lengths out of control by control limit h (see text for details). Age effect constant: ----, increasing: ____, decreasing:…… .

Figure 3. Run length distributions based on 200 simulations (bar charts) and superimposed expected values from the geometric distribution (lozenges). Data from Table 2, increasing age effect profiled out. (a) $\rho_A = 1.5$, h = 3; (b) $\rho_A = 2$, h = 2.

Figure 4. Bell's palsy and flu vaccine. CUSUM graph for 1-60 day risk period and $\rho_A=1.5$.

Figure 5. Convulsions and MMR vaccine. CUSUM for 6-11 day risk period, $\rho_A = 3$ and h=3. The pairs (no. events in risk period, total no. events) are given for each monitoring interval.

Figure 6. Convulsions and MMR vaccine. CUSUM for 6-11 day risk period and h=3. ------ $\rho=3$ vs $\rho_A=5$, _____ $\rho=4$ vs $\rho_A=6$.

Vaccine
Musonda et al
Figure 1

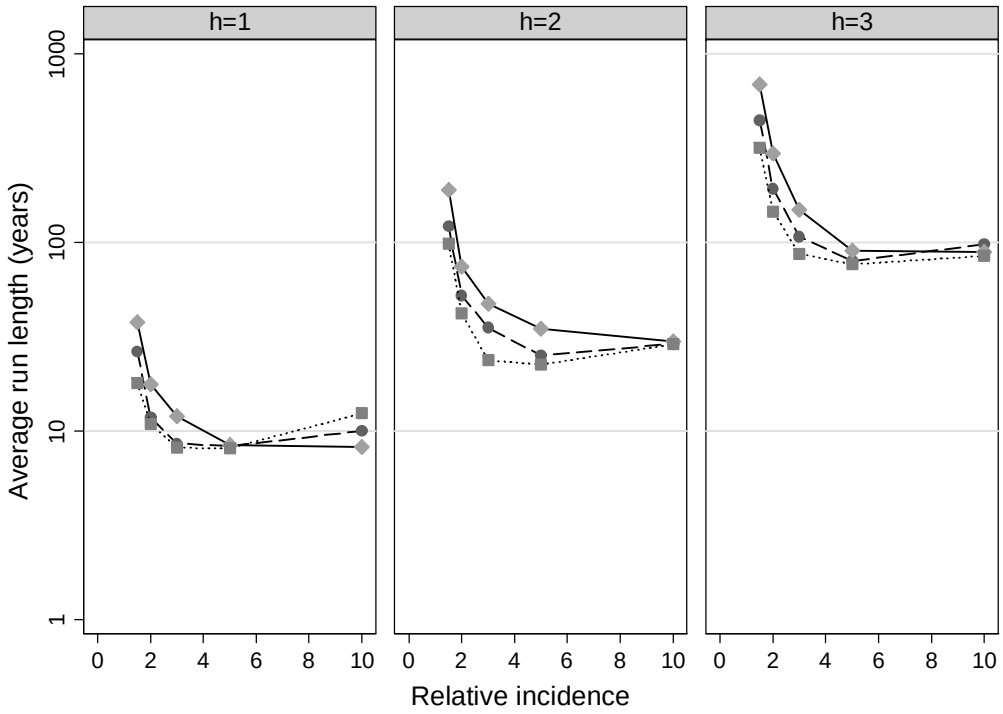

Vaccine
Musonda et al
Figure 2

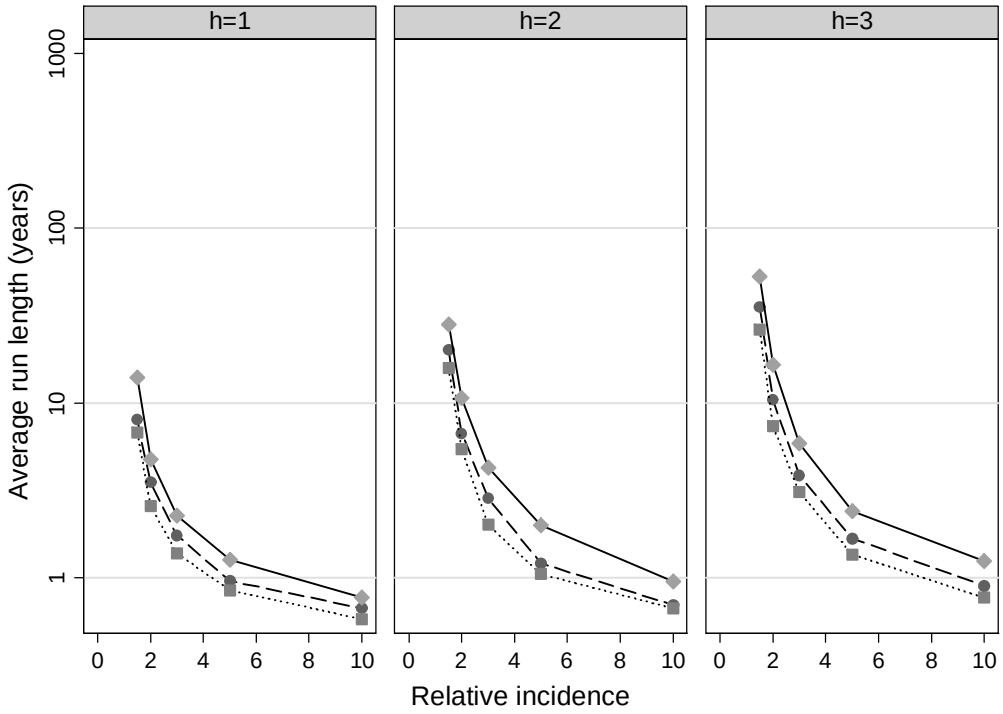

Vaccine
Musonda et al
Figure 3

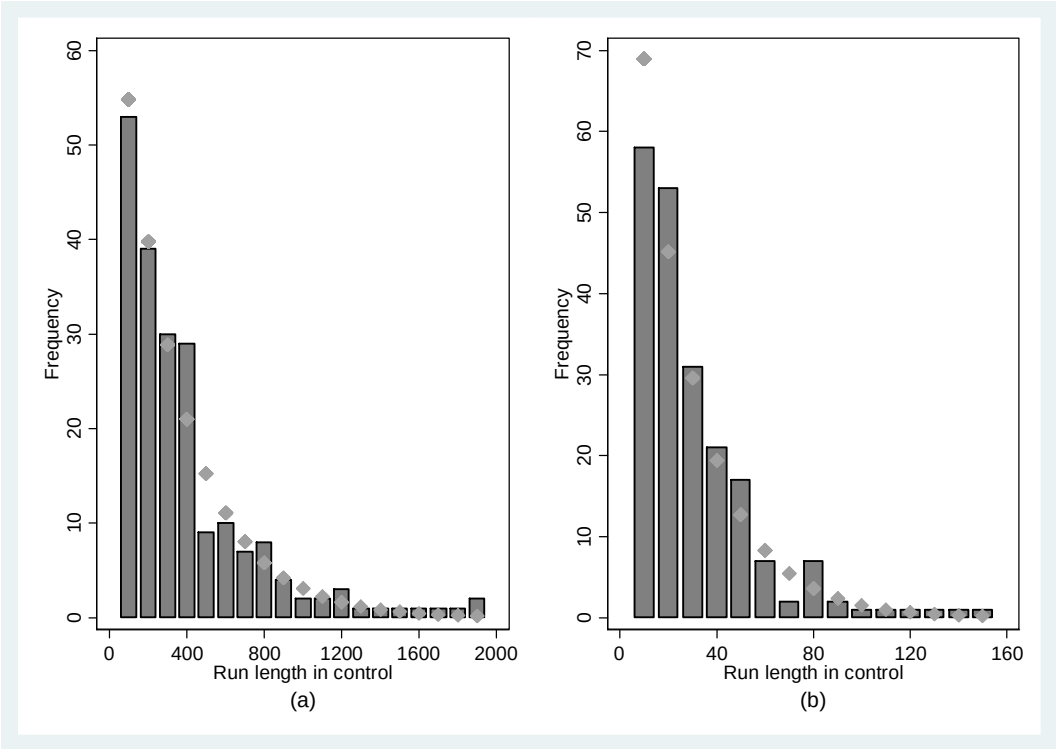

Vaccine
Musonda et al
Figure 4

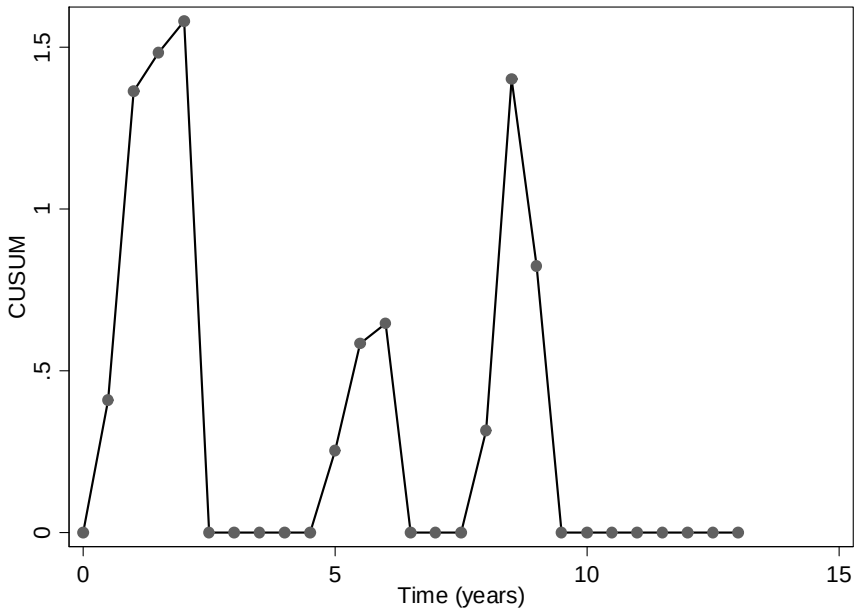

Vaccine
Musonda et al
Figure 5

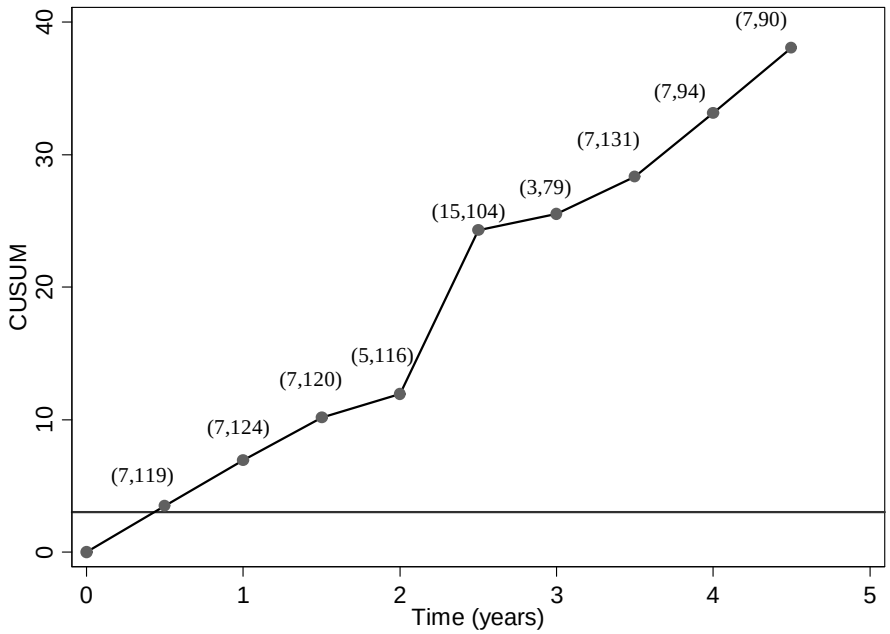

Vaccine
Musonda et al
Figure 6