Sudarshan Publications

Available online at https://s-epub.in/index.php/trp/index

Pharma Research

An International Peer Reviewed, Indexed Journal

10.62655/s-epub.2025.v17.i01.pp9-16

SIGNAL DETECTION METHODS IN PHARMACOVIGILANCE DATABASES: DISPROPORTIONALITY ANALYSIS EXPLAINED

HANNAH LINDQVIST[1]*, LUKAS BROCKHAUS[2], ISHITA D BHAT[3]

Affiliation

[1]Department of Pharmacy Practice, Shenandoah Valley College of Pharmacy, Harrisonburg, USA
[2]Department of Pharmacovigilance, Neckar Valley Institute of Pharmaceutical Sciences, Heilbronn, Germany
[3]Department of Clinical Pharmacology, Krishna Plateau College of Pharmacy, Khammam, India

*Corresponding author: Hannah Lindqvist, h.lindqvist@svcp.edu
ORCID iD: Hannah Lindqvist 0000-0001-0084-2089; Lukas Brockhaus 0000-0001-1228-773X; Ishita D Bhat 0000-0001-0111-4374

Article info

Received date:11-07-2024
Accepted date:19-11-2024
Published date:28-03-2025
ISSN;0975-8216

Keywords:

Pharmacovigilance, Signal detection, Disproportionality analysis, Spontaneous reporting, Data mining

Abstract

Spontaneous reports of suspected adverse drug reactions are the main source of early safety information after a medicine is marketed, and statistical screening of the databases that hold them is now routine. This review explains disproportionality analysis, the family of methods that compare how often an event is reported with a drug against how often it is reported with all other drugs. It describes the proportional reporting ratio, the reporting odds ratio and the Bayesian approaches used by the WHO and the US FDA, sets out the usual thresholds, and discusses biases such as under-reporting, notoriety and masking. The methods are useful for generating hypotheses, not for measuring risk. This is a review and presents no new analysis.

1. INTRODUCTION

Clinical trials are designed to show that a medicine works, not to find every harm it can cause. They enrol a few hundred or a few thousand people, usually exclude the very young, the old and those with several illnesses, and follow them for a limited time. Rare reactions, delayed reactions and reactions confined to particular groups therefore tend to come to light only after approval, when many thousands of patients are exposed in everyday practice [1, 2]. Safety monitoring after marketing has to rely on what clinicians, pharmacists and patients choose to report, and the problem is how to find the few meaningful patterns among the very large number of reports.

Early pharmacovigilance relied on expert review of individual case reports. As databases grew, this became impossible. The

WHO international database now holds many millions of reports from more than a hundred countries, and national and regional systems such as the US FDA Adverse Event Reporting System and the European EudraVigilance database are also very large [3, 4]. Statistical methods that screen all drug-event pairs and rank those that stand out were developed in response, and they are collectively called disproportionality analysis. They have become a standard part of the toolkit of regulators, companies and academic groups [5, 6].

This article explains how those methods work and where they go wrong. It begins with the idea of a signal and the nature of spontaneous reports, then sets out the frequentist and Bayesian methods in turn, with the thresholds usually applied. A later section deals with practical interpretation, including known biases, and the final section considers limitations. The article is a narrative review of published methodological and applied literature. It contains no new analysis of any database, and the figures that appear are illustrative of the logic of the methods, not results [7, 8, 9].

2. SPONTANEOUS REPORTING AND THE CONCEPT OF A SIGNAL

A spontaneous report is a description, sent voluntarily to a regulator, manufacturer or monitoring centre, of an adverse event that occurred in a patient taking a medicine and was suspected of being related to it. The report carries a few fields: age and sex, the suspected and concomitant drugs, the event terms, dates and outcome. A causal link is suspected, not established. Databases are coded with standard terminology for reactions, usually the Medical Dictionary for Regulatory Activities, and with a standard drug dictionary, which is what makes automated counting feasible in the first place [3, 4].

A signal, in pharmacovigilance, is information suggesting a new potentially causal association, or a new aspect of a known one, between a drug and an event, that is judged to be of sufficient likelihood to justify verification. The wording comes from the WHO, and reviews of decision support methods for post-marketing data stress that a signal is a hypothesis combining statistical association with clinical plausibility and the need for action [7, 10]. A signal is therefore not a proven hazard. It is a hypothesis flagged for assessment, and the work of evaluating it, by reviewing cases, looking at pharmacology and checking other data sources, comes after statistical screening, not before it.

The usefulness of spontaneous reporting is that it covers the entire marketed population and picks up rare, serious events with a high drug-attributable fraction, such as Stevens-Johnson syndrome, torsades de pointes and agranulocytosis. Case-by-case review remains the gold standard for these dramatic events. The difficulty arises with common events that also occur in untreated people, where a handful of reports cannot be distinguished from background without a comparison. Disproportionality analysis offers such a comparison within the database, using the reports for all other drugs as the reference [5, 6, 11].

Disproportionality has to be understood against the shortcomings of reporting itself. A comparison of a spontaneous reporting scheme with a sentinel system showed that under-reporting is widespread [2]. Reporting rates also change over a product's life, a pattern first described for nonsteroidal anti-inflammatory drugs and known as the Weber effect, in which reports rise in the first two years after launch and then decline [12]. No spontaneous report database gives incidence, because neither the number of patients exposed nor the number of events missed is known.

Data sources differ in size, population and access, and the choice affects what can be learned. VigiBase, maintained by the Uppsala Monitoring Centre, pools reports from national centres worldwide and therefore has wide geographical coverage, though variable completeness [4]. The FDA database is dominated by reports from the United States, and its public version has become a favourite of academic groups because it can be downloaded and analysed without special permission [13, 14]. EudraVigilance serves the European regulatory network [15]. Because reporting culture, legal obligations

and the share of reports from industry versus health professionals differ, the same drug-event pair can look quite different in each source.

3. FREQUENTIST METHODS

All disproportionality methods start from the same two-by-two table. For a given drug and a given event, reports are classified into four cells: reports with both the drug and the event (a), the drug with other events (b), other drugs with the event (c), and other drugs with other events (d). The expected number of drug-event reports can be calculated from the margins of the table on the assumption that the drug and the event are unrelated, and the question is whether the observed count exceeds it by more than chance would suggest [16, 17]. Everything else, from simple ratios to Bayesian shrinkage, is a different way of answering that question.

The proportional reporting ratio, developed at the UK Medicines Control Agency and discussed by Evans, compares the proportion of a drug's reports that concern the event, a/(a+b), with the proportion of all other reports that concern the event, c/(c+d) [16]. A PRR of 1 means no difference. The originators proposed that a signal be considered when the PRR is at least 2, the chi-squared statistic is at least 4 and there are at least three reports, and these criteria are still widely quoted. They were chosen pragmatically to cut the number of spurious signals to a manageable volume, not derived from a statistical theory of error rates [5, 16].

The reporting odds ratio, favoured in the Netherlands, is the ratio of the odds of the event among reports with the drug, a/b, to the odds among reports without it, c/d, which simplifies to ad/bc. A confidence interval is easily calculated from the standard error of the log odds ratio, and a signal is commonly declared when the lower limit of the 95 per cent interval exceeds one [17, 18]. Because it is an odds ratio, the ROR can be extended by logistic regression to adjust for age, sex and concomitant medicines, something the PRR cannot do in a simple way [8, 17]. That flexibility is a practical reason for the popularity of the method in academic studies.

The arithmetic is simple enough to do by hand. For the reporting odds ratio, the natural logarithm of ad/bc is approximately normally distributed with a standard error equal to the square root of the sum of the reciprocals of the four cell counts, 1/a + 1/b + 1/c + 1/d. The 95 per cent confidence interval is found by adding and subtracting 1.96 standard errors on the log scale and exponentiating [17, 18]. The formula makes the dependence on small cells obvious, since a count of one or two in cell a dominates the standard error and widens the interval. A reported ROR without its interval, and without the count in cell a, conveys very little.

PRR and ROR give very similar answers when the event is rare in the database, as it generally is, because the odds and the proportion converge. Discussion in the literature about which is better has been lively but has not produced a clear winner. Rothman and colleagues argued that the ROR had advantages, including a well-defined confidence interval and independence from the choice of reference group in certain situations, while others responded that in practice the differences were small and that both methods share the same weaknesses [18, 19]. Comparison across several databases found that the frequentist and Bayesian methods identified broadly similar sets of signals, with differences concentrated in pairs with few reports [20].

A weakness of both measures is instability for small counts. When a is one or two, the ratio can be very large by chance, and a confidence interval, though helpful, does not remove the problem entirely. Minimum case counts, such as the three reports in the PRR criteria, are one remedy and cost sensitivity for new and rare events. A second weakness is that the result depends on the dictionary level at which events are grouped, because analysing a preferred term, a group of related terms or a higher-level category leads to different tables and different ratios [5, 21]. Teams therefore agree in advance on the event definitions, often using standard queries created by medical experts.

For readers who want an intuition, consider a drug accounting for a modest part of all reports in a database. If the event under study

is mentioned in about one report in a hundred overall, but in one in ten of this drug's reports, the PRR is close to ten and the pair would be flagged. If the same ten per cent applied to a drug with only two such reports, the ratio would look similar, but the evidence would be almost worthless. The point of the thresholds on count and interval width is to separate these situations, which is why a ratio quoted without its count should not be trusted [16, 18].

4. BAYESIAN METHODS

Bayesian methods address the instability of simple ratios by shrinking extreme values toward the average, using what is known about the distribution of all drug-event pairs. The Bayesian Confidence Propagation Neural Network, used by the Uppsala Monitoring Centre for the WHO database, calculates an information component, which is the logarithm to base two of the ratio of the observed joint reporting probability to the probability expected if the drug and event were independent [22, 23]. A positive value indicates that the pair appears together more often than expected, and a credibility interval is derived from the posterior distribution.

The idea of shrinkage deserves a plain explanation. Suppose that in a large database most drug-event pairs hover around an observed-to-expected ratio of one, and that a few pairs with very small counts show ratios of ten or twenty. The Bayesian model treats these extreme values with suspicion, because with so little data they are more likely to reflect chance than a true tenfold excess, and pulls them toward the overall pattern. As the number of reports rises, the pull weakens and the estimate approaches the raw ratio [22, 23, 24]. This is why a pair that remains prominent after shrinkage is more convincing than one that is prominent only in the raw figures.

In the information component approach, a signal is usually flagged when the lower bound of the 95 per cent credibility interval, often written IC025, is above zero. Retrospective evaluation of the method on the WHO database found that it would have flagged a number of recognised associations at an early stage, and the subsequent extension to account for complex associations and improved accuracy for rare events refined the calculations [11, 23]. Because of shrinkage, a pair with two reports and a large observed-to-expected ratio produces a modest information component, and lacks the apparent drama of the corresponding PRR. That conservatism is both its strength and its cost.

The Multi-item Gamma Poisson Shrinker, developed for the US FDA by DuMouchel and colleagues, takes an empirical Bayes approach. It models the observed counts as Poisson variables whose relative reporting rates follow a mixture of two gamma distributions, the parameters of which are estimated from the full dataset [24]. The resulting empirical Bayes geometric mean, EBGM, is a shrunken estimate of the relative reporting rate, and its lower 5 per cent bound, EB05, is used as a conservative screen. A value of 2 for EB05 is a common choice, and the FDA has used the approach to screen its reports routinely [3, 24, 25].

Do Bayesian methods perform better? Evidence is mixed. Their principal advantage is fewer false positives for pairs with few reports, which reduces the workload for reviewers. Their main disadvantage is loss of sensitivity for new drugs and rare events, and greater complexity for a user who does not have software. Studies that compared methods against reference sets of known adverse reactions found small differences in overall performance, with the best choice depending on whether one values early detection or precision [20, 21, 26]. Many safety groups run more than one method and look at pairs flagged by several, a pragmatic way of combining their strengths.

5. APPLYING AND INTERPRETING THE RESULTS

Statistical screening is normally run on a schedule, with the whole database scanned for new or changed drug-event pairs. In Europe, the EudraVigilance guideline describes the use of proportional reporting ratio-type statistics at the level of drug and event, with thresholds that can be set to balance sensitivity and workload, and with output checked by pharmacovigilance medical assessors [15]. The UK, the Netherlands,

France and other countries operate their own systems. Pairs that pass the threshold are triaged on clinical grounds: is the event serious, is it already listed on the product label, is there a plausible mechanism, and do the individual case narratives support causality [9, 10]?

Published applied studies illustrate the range of use. Analyses of the publicly available FDA database have examined drug-induced torsades de pointes, cardiovascular events with various drug classes and the safety of newly marketed products, usually reporting ratios with confidence intervals for each drug of interest [14, 27]. Such studies are quick and cheap, which explains their number, and they sometimes provide the first hint of an association. They also bear a risk of overinterpretation, as the output of an algorithm is easy to present as a finding while the biases discussed below are left unmentioned [6, 13].

The findings of several kinds of method are drawn together in table 1, which sets out statistical measures, typical thresholds and biases, together with notes on practice. The number of entries in each group is displayed in fig. 1. The balance is informative. The statistical measures themselves are few and well defined, whereas the list of problems with them is longer, and reflects the fact that most of the practical difficulty in disproportionality analysis lies in interpretation and not in calculation. Any one ratio is best read as the start of an inquiry.

Several biases deserve a name. Notoriety bias arises when publicity about a safety concern boosts reporting of that event for the drug, inflating the ratio without any change in true risk; one study showed that disproportionality measures rose after safety alerts in France [28]. Masking, or competition bias, occurs when a drug with a very large number of reports of an event lowers the apparent disproportionality of other drugs for the same event, hiding true signals [29]. Protopathic bias, confounding by indication and co-reporting of drugs taken together also distort the picture. The sertindole experience showed how reporting patterns peculiar to one drug can distort a PRR, and why the composition of the reference group matters [30].

A related issue is that disproportionality compares a drug with other drugs, not with no drug. If the events of interest are linked to the underlying disease, then a drug used for that disease will show an association whatever it does. Antidiabetic drugs and diabetic complications, or antipsychotics and sudden death in psychiatric patients, are classic illustrations. Analysts try to reduce this by restricting the comparator to drugs with the same indication, a step that is simple in principle but demands a good understanding of the clinical area [21, 30]. Without it, a signal may reflect the illness more than the medicine.

Validation of the results is the step that distinguishes a responsible analysis from a speculative one. The pair should be checked against the product information and the literature, the cases reviewed for plausible timing, dechallenge, rechallenge and alternative explanations, and the pattern compared across other databases where possible [9, 10]. Retrospective studies on label changes show that statistical signals frequently precede regulatory action, which supports their value as an early warning, though they also generate many pairs that never become validated [31, 32]. Reports should state which databases, time periods, event definitions and comparators were used, so that others can reproduce the analysis.

Table 1: Disproportionality measures, usual thresholds and common biases in published pharmacovigilance literature

Item	Key point reported	Ref.
Proportional reporting ratio PRR	Ratio of proportions; usual screen is PRR of 2 or more, chi-squared 4 or more, three or more reports	[5, 16]
Reporting odds ratio ROR	Odds ratio ad/bc; signal when lower 95% confidence limit exceeds one; can be adjusted by regression	[17, 18]
PRR versus ROR comparison	Near-identical results for rare events; differences are small in practice	[19, 20]
Information component BCPNN	Log ratio of observed to expected; IC025 above zero used as screen in WHO database	[22, 23]

Item	Key point reported	Ref.
Empirical Bayes MGPS	Shrunken relative reporting rate; lower bound EB05 used as conservative screen	[3, 24]
Method comparison reference sets	Broadly similar signals across methods; differences concentrated in low counts	[20, 26]
Under-reporting and Weber effect reporting rates	Most reactions go unreported; reporting varies with time on market	[2, 12]
Notoriety bias publicity	Safety alerts raise disproportionality without a change in risk	[28]
Masking competition	Heavily reported drugs can hide signals for others	[29, 30]
Regulatory use screening aid	Statistical output triaged by medical assessors in EudraVigilance	[15]
Complementary data health records	Record databases supply denominators and support spontaneous report signals	[31, 33]

6. LIMITATIONS AND FUTURE DIRECTIONS

The most important limitation is conceptual. A disproportionality ratio measures how a pair stands out within a database, not how likely a patient is to suffer the reaction. It cannot be converted into relative risk, since the denominators are unknown, and statements such as increased risk of an event on the strength of an ROR alone go beyond what the data permit [6]. Reports of high ratios for newly launched drugs may reflect stimulated reporting, and low ratios for older ones may reflect fatigue and reduced attention [12, 28]. Readers of published analyses should look for caution on exactly this point.

Data quality also matters. Duplicate reports, missing information, inconsistent drug names and multiple reporters for the same case all affect counts. The suspected and concomitant drug fields are not reliable indicators of causality, because the choice of which drug to name suspect is made by the reporter. Reports are more numerous from countries with active reporting cultures, which tilts global databases toward some populations [2, 4]. Different threshold conventions, event groupings and software make results hard to compare, and the lack of a widely accepted reference standard makes performance evaluation unavoidably dependent on the choice of reference set [21, 26].

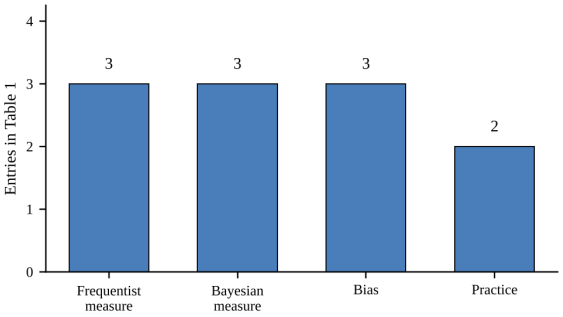


Fig. 1: Number of entries in table 1 by group

Future directions include better reference sets for evaluation, more routine adjustment for confounders, and the integration of spontaneous reports with other sources, including social media and electronic records, with careful attention to their own biases. Machine learning approaches are being explored for case triage and signal prioritisation, though their benefit over simple disproportionality has not been consistently shown. For now, the practical advice for those using the methods is modest: know the database, state the method and the thresholds, pair statistical output with clinical review, and describe the result as a hypothesis [6, 9, 21].

7. CONCLUSION

Disproportionality analysis converts a very large collection of spontaneous reports into a ranked list of drug-event pairs that deserve attention. The frequentist measures are simple and transparent, the Bayesian ones are more conservative for small counts, and in practice the two families identify largely overlapping signals [16, 20, 22, 24]. Their value is in focusing scarce clinical review time, and many label changes have followed pairs that were first picked up statistically [31, 32].

Their limits are just as clear. The ratios depend on reporting behaviour, the comparator and the dictionary level, and they cannot be read as risks [2, 28, 29]. A sensible user states the method, thresholds and data, checks plausible confounding by indication, and treats every flagged pair as a question for further investigation using cases, mechanism and other data sources [6, 9, 15]. Used in that way, disproportionality analysis is a

worthwhile part of modern drug safety surveillance.

8. ABBREVIATIONS

BCPNN-Bayesian confidence propagation neural network, CIOMS-Council for International Organizations of Medical Sciences, EBGM-Empirical Bayes geometric mean, FAERS-FDA Adverse Event Reporting System, IC-Information component, MedDRA-Medical Dictionary for Regulatory Activities, MGPS-Multi-item Gamma Poisson Shrinker, PRR-Proportional reporting ratio, ROR-Reporting odds ratio, WHO-World Health Organization

9. FUNDING

Nil

10. AUTHORS CONTRIBUTIONS

All the authors have contributed equally.

11. CONFLICT OF INTERESTS

Declared none

12. REFERENCES

1. Lasser KE, Allen PD, Woolhandler SJ, Himmelstein DU, Wolfe SM, Bor DH. Timing of new black box warnings and withdrawals for prescription medications. JAMA 2002;287:2215–20.
2. Bäckström M, Mjörndal T, Dahlqvist R. Under-reporting of serious adverse drug reactions in Sweden. Pharmacoepidemiol Drug Saf 2004;13:483–7.
3. Szarfman A, Machado SG, O'Neill RT. Use of screening algorithms and computer systems to efficiently signal higher-than-expected combinations of drugs and events in the US FDA's spontaneous reports database. Drug Saf 2002;25:381–92.
4. Lindquist M. VigiBase, the WHO Global ICSR Database System: basic facts. Drug Inf J 2008;42:409–19.
5. Hauben M, Madigan D, Gerrits CM, Walsh L, Van Puijenbroek EP. The role of data mining in pharmacovigilance. Expert Opin Drug Saf 2005;4:929–48.
6. Montastruc JL, Sommet A, Bagheri H, Lapeyre-Mestre M. Benefits and strengths of the disproportionality analysis for identification of adverse drug reactions in a pharmacovigilance database. Br J Clin Pharmacol 2011;72:905–8.
7. Hauben M, Bate A. Decision support methods for the detection of adverse events in post-marketing data. Drug Discov Today 2009;14:343–57.
8. Bate A, Evans SJW. Quantitative signal detection using spontaneous ADR reporting. Pharmacoepidemiol Drug Saf 2009;18:427–36.
9. European Medicines Agency. Guideline on good pharmacovigilance practices (GVP) Module IX: signal management (Rev 1). London: EMA; 2017.
10. Meyboom RHB, Lindquist M, Egberts ACG, Edwards IR. Signal selection and follow-up in pharmacovigilance. Drug Saf 2002;25:459–65.
11. Lindquist M, Ståhl M, Bate A, Edwards IR, Meyboom RHB. A retrospective evaluation of a data mining approach to aid finding new adverse drug reaction signals in the WHO international database. Drug Saf 2000;23:533–42.
12. Weber JCP. Epidemiology of adverse reactions to nonsteroidal antiinflammatory drugs. Adv Inflamm Res 1984;6:1–7.
13. Poluzzi E, Raschi E, Piccinni C, De Ponti F. Data mining techniques in pharmacovigilance: analysis of the publicly accessible FDA Adverse Event Reporting System (AERS). In: Karahoca A, editor. Data mining applications in engineering and medicine. Rijeka: InTech; 2012. p. 267–301.
14. Sakaeda T, Tamon A, Kadoyama K, Okuno Y. Data mining of the public version of the FDA Adverse Event Reporting System. Int J Med Sci 2013;10:796–803.
15. European Medicines Agency. Guideline on the use of statistical signal detection methods in the EudraVigilance data analysis system. London: EMA; 2008.
16. Evans SJW. Pharmacovigilance: a science or fielding emergencies? Stat Med 2000;19:3199–209.
17. Egberts ACG, Meyboom RHB, van Puijenbroek EP. Use of measures of disproportionality in pharmacovigilance. Drug Saf 2002;25:453–458.
18. Rothman KJ, Lanes S, Sacks ST. The reporting odds ratio and its advantages over the proportional reporting ratio. Pharmacoepidemiol Drug Saf 2004;13:519–23.
19. Waller P, van Puijenbroek E, Egberts A, Evans S. The reporting odds ratio versus the proportional reporting ratio: 'deuce'. Pharmacoepidemiol Drug Saf 2004;13:525–6.

20. Candore G, Juhlin K, Manlik K, Thakrar B, Quarcoo N, Seabroke S, *et al.* Comparison of statistical signal detection methods within and across spontaneous reporting databases. Drug Saf 2015;38:577–87.
21. Harpaz R, DuMouchel W, Shah NH, Madigan D, Ryan P, Friedman C. Novel data-mining methodologies for adverse drug event discovery and analysis. Clin Pharmacol Ther 2012;91:1010–21.
22. Bate A. Bayesian confidence propagation neural network. Drug Saf 2007;30:623–625.
23. Norén GN, Bate A, Orre R, Edwards IR. Extending the methods used to screen the WHO drug safety database towards analysis of complex associations and improved accuracy for rare events. Stat Med 2006;25:3740–57.
24. DuMouchel W. Bayesian data mining in large frequency tables, with an application to the FDA spontaneous reporting system. Am Stat 1999;53:177–90.
25. Almenoff JS, Pattishall EN, Gibbs TG, DuMouchel W, Evans SJW, Yuen N. Novel statistical tools for monitoring the safety of marketed drugs. Clin Pharmacol Ther 2007;82:157–66.
26. Harpaz R, DuMouchel W, LePendu P, Bauer-Mehren A, Ryan P, Shah NH. Performance of pharmacovigilance signal-detection algorithms for the FDA adverse event reporting system. Clin Pharmacol Ther 2013;93:539–46.
27. Poluzzi E, Raschi E, Moretti U, De Ponti F. Drug-induced torsades de pointes: data mining of the public version of the FDA Adverse Event Reporting System (AERS). Pharmacoepidemiol Drug Saf 2009;18:512–8.
28. Pariente A, Gregoire F, Fourrier-Reglat A, Haramburu F, Moore N. Impact of safety alerts on measures of disproportionality in spontaneous reporting databases: the notoriety bias. Drug Saf 2007;30:645–55.
29. Pariente A, Avillach P, Salvo F, Thiessard F, Miremont-Salamé G, Fourrier-Réglat A, *et al.* Effect of competition bias in safety signal generation: analysis of a research database of spontaneous reports in France. Drug Saf 2012;35:855–64.
30. Moore N, Hall G, Sturkenboom M, Mann R, Lagnaoui R, Begaud B. Biases affecting the proportional reporting ratio (PRR) in spontaneous reports pharmacovigilance databases: the example of sertindole. Pharmacoepidemiol Drug Saf 2003;12:271–81.
31. Hauben M, Reich L. Safety related drug-labelling changes: findings from two data mining algorithms. Drug Saf 2004;27:735–44.
32. Bate A, Lindquist M, Edwards IR. The application of knowledge discovery in databases to post-marketing drug safety: example of the WHO database. Fundam Clin Pharmacol 2008;22:127–40.
33. Trifirò G, Pariente A, Coloma PM, Kors JA, Polimeni G, Miremont-Salamé G, *et al.* Data mining on electronic health record databases for signal detection in pharmacovigilance: which events to monitor? Pharmacoepidemiol Drug Saf 2009;18:1176–84.