PDS Pharmacoepidemiology & Drug Safety
Official Journal of the International Society for Pharmacoepidemiology
ISPE

Estimating time-varying drug adherence using electronic records: extending the Proportion of Days Covered (PDC) method

Journal:	*Pharmacoepidemiology and Drug Safety*
Manuscript ID:	PDS-15-0189.R1
Wiley - Manuscript type:	Original Research Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Bijlsma, Maarten; University of Groningen, Pharmacy Janssen, Fanny; University of Groningen, Population Research Centre Hak, Eelko; University of Groningen, PharmacoEpidemiology & PharmacoEconomics
Keywords:	adherence, methods, longitudinal, time dependence

Abstract:

Purpose: Accurate measurement of drug adherence is essential for valid risk-benefit assessments of pharmacologic interventions. To date, measures of drug adherence have almost exclusively been applied for a fixed-time interval, and without considering changes over time. However, patients with irregular dosing behavior commonly have a different prognosis than patients with stable dosing behavior.
Methods: We propose a method, based on the Proportion of Days Covered (PDC) method, to measure time-varying drug adherence and drug potency using electronic records. We use an irregularly dosing patient and a patient with stable adherence as examples. For these patients, we compare both a static PDC method with the time varying PDC method.
Results: We demonstrate that time varying PDC method better distinguishes an irregularly dosing patient from a stably dosing patient, and demonstrate how the static method can result in a biased estimate of drug adherence. Furthermore, the time varying PDC method may be better used to reduce certain types of confounding and misclassification of exposure.
Conclusions: The time varying PDC method may improve longitudinal and time-to-event studies that associate adherence with a clinical outcome, or (intervention) studies that seek to describe changes in adherence over time.

Estimating time-varying drug adherence using electronic records: extending the Proportion of Days Covered (PDC) method

Running head: Time varying proportion of days covered method

Maarten J. Bijlsma [a*], PhD Fellow

Fanny Janssen[b,c], Assistant Professor of Analytical Demography

Eelko Hak [a], Professor of Clinical Pharmacoepidemiology

[a] Unit PharmacoEpidemiology & PharmacoEconomics (PE^2), Department of Pharmacy, University of Groningen, Groningen, The Netherlands.

[b] Population Research Centre (PRC), Faculty of Spatial Sciences, University of Groningen, Groningen, The Netherlands.

[c] Netherlands Interdisciplinary Demographic Institute, the Hague, the Netherlands

* Corresponding author. E-mail address: maarten.bijlsma@rug.nl

University of Groningen, Department of Pharmacy, Unit PharmacoEpidemiology & PharmacoEconomics (PE2), A. Deusinglaan 1, PO BOX 9713 AV Groningen, the Netherlands.

Tel. +31 (0)50-363-7576; Fax: +31 (0)50-363-2772.

Key words: adherence, methods, longitudinal**,** time dependence

Key messages:

- To date, measures of drug adherence have almost exclusively been applied for a fixed time interval, and without considering changes over time. Yet time varying differences in drug adherence may have real effects on patient prognosis.
- We demonstrate a method to measure time varying drug adherence, which better distinguishes an irregularly dosing patient from a stably dosing patient, and which is less likely to produce biased estimates.
- The time varying PDC method may improve longitudinal and time-to-event studies that associate adherence with a clinical outcome, or (intervention) studies that seek to describe changes in adherence over time.

Conflict of interest statement: all authors report no conflict of interests.

Word count: 4405

Abstract

Purpose: Accurate measurement of drug adherence is essential for valid risk-benefit assessments of pharmacologic interventions. To date, measures of drug adherence have almost exclusively been applied for a fixed-time interval, and without considering changes over time. However, patients with irregular dosing behavior commonly have a different prognosis than patients with stable dosing behavior.

Methods: We propose a method, based on the Proportion of Days Covered (PDC) method, to measure time-varying drug adherence and drug dosage using electronic records. We use an irregularly dosing patient and a patient with stable adherence as examples. For these patients, we compare both a static PDC method with the time varying PDC method.

Results: We demonstrate that time varying PDC method better distinguishes an irregularly dosing patient from a stably dosing patient, and demonstrate how the static method can result in a biased estimate of drug adherence. Furthermore, the time varying PDC method may be better used to reduce certain types of confounding and misclassification of exposure.

Conclusions: The time varying PDC method may improve longitudinal and time-to-event studies that associate adherence with a clinical outcome, or (intervention) studies that seek to describe changes in adherence over time.

Introduction

Accurate measurement of drug adherence is essential for valid risk-benefit assessments of pharmacologic interventions [1-3]. Patient adherence has a direct influence on whether the patient receives the prescribed drug dose, or whether under- or overdosing of prescribed medication occurs. In clinical trials, due to strict protocols, higher levels of adherence are achieved than in observational study designs which may potentially lead to differences in drug efficacy or safety estimates between these designs [4-7]. Hence, accurate drug adherence measurements are a prerequisite for bridging the gap between biological efficacy estimates from experimental trials on the one hand and clinical effectiveness estimates from observational studies on the other hand.

To date, measures of drug adherence such as the Proportion of Days Covered (PDC) method have almost exclusively been applied for a fixed-time interval, and without considering changes over time. Such an application ignores the fact that adherence within patients may vary over time [8]. In a fixed time interval, a patient that receives the drug irregularly may have the same adherence estimate as a patient that steadily receives the drug in the same time interval, yet the real differences in dosing behavior may result in a totally different patient prognosis. In other words, using time-constant measures of drug adherence in a fixed time interval will bias the association between a clinical outcome and drug use. Furthermore, in studies with time-to-event analysis, the association between the drug use measured with an adherence measure taken over a fixed time interval and the clinical outcome may be biased because some patients may experience the event of interest long after the fixed time interval has passed. In all, time-constant drug adherence measures are disadvantageous both in studies assessing cumulative incidence ratios and incidence rate ratios.

There is a wide variety of methods to estimate adherence, each with their specific advantages and disadvantages [6, 9-13]. Methods that use electronic records (e.g. pharmacy

records), rather than patient reports or direct observation, have as their advantage that they are noninvasive and can often be used for large number of patients over a long time span. Given the fact that in Western countries chronic diseases are becoming more prevalent and both preventive and therapeutic drugs are used over a longer period of time and recorded in Big Data health care registries (e.g. [14, 15]), methods that use electronic records are indispensable. Of the methods designed for this purpose, the Proportion of Days Covered (PDC) method is most commonly applied (e.g. [16]).

This paper describes an extension of the static PDC method that enables the estimation of time-varying drug adherence and drug dosage using pharmacy prescription or dispensing records, it illustrates the method, and discusses its strengths and limitations. We provide an annotated syntax for the statistical programming language 'R' as online supplemental material [17].

Methods

Drug prescription or dispensing records

The extended PDC method is intended to be applied to data from drug prescription or dispensing records. Initially, data should be ordered such that each row represents a single drug prescription (or dispensed prescription). The information needed to apply the method is represented by variables (columns) in the dataset including a patient identification number (ID), date of dispensing, number of pills dispensed, and number of pills per day. Once prescriptions are chronologically ordered, a variable 'prescription number' can be added which is given value *k* for the *k*'th prescription (Figure 1).

Estimating time-varying drug adherence

First, we calculated the length of time in days for each interval ("Interval length" in Figure 1). To stabilize the adherence estimate, an interval is not the length in time between one prescription (*k*) and the next (*k* + *1*), but between each prescription *k* and the date of the second prescription afterwards (*k* + 2). Secondly, we calculated the expected number of days covered ("Total days" in Figure 1), by dividing the number of pills dispensed by the pills per day for each row, and summing these numbers for rows *k* and *k+1*. Then, "adherence" as a proportion (Figure 1) in each row was calculated by dividing the expected number of days covered by the length of time in the interval ("Total days" / "Interval length" in the figure). The adherence value may exceed 1 if the length of the interval is shorter than the expected number of days covered. This may occur if the patient is stockpiling the drugs (e.g. to go on holiday). In the case of stockpiling, we carried over the pills that are in excess of the expected number of days covered to the next interval until no interval has an "adherence" estimate above 1 (Figure 2). If stockpiling is not possible for the drug in question, for example if the

drug is not chemically stable for a long time, this estimation step should be skipped and intervals with adherence values above 1 should be set to 1.

When estimating "adherence" using intervals based on the length of time between prescriptions *k* and *k+2*, this leaves part of the information of the last two prescriptions unused because the length of time of the interval cannot be established; i.e. prescription *k+2* does not exist when *k* is the last or next to last prescription. This is not problematic if the final prescriptions take place outside the study period. In other cases, an end point of utilization of the drug can be established by assuming that the last observed adherence value will be continued in the final interval. The length of the last interval will then be the sum of the expected number of days covered by the last and second to last prescription, and divided by the last observed adherence value. This represents the length of time that a patient would be able to continue to use the drug if the last observed adherence is continued into the final interval.

Because intervals for adherence calculation are constructed between prescriptions *k* and *k+2*, two intervals will overlap at most time points. To any time point with such overlapping intervals, we assigned the adherence value from the first of these two intervals.

To calculate adherence over a longer time period (e.g. over 30-day periods), an average adherence over the desired time period can be computed after execution of the previous step.

Finally, patients may switch between drugs over time; if this is not detected, it will lead to erroneous estimates of drug adherence. A patient can be considered to have switched a drug if he or she receives a prescription for one drug, then later in time receives a prescription for a different drug in the same class as the first prescription and does not refill the old prescription [16]. Before calculating drug adherence, switchers should first be identified and rows of both

the old and the new drug can be ordered chronologically as in Figure 1. The calculating of time-varying drug adherence can then proceed as described in this paper.

Estimating time-varying drug dosage

Electronic pharmacy prescription or dispensing records commonly also contain information on dosage. Time-varying dosage was estimated using a simpler version of the drug adherence calculation. Drug dosage in an interval was calculated as a weighted average of milligrams (mg) dispensed, using number of days that should be covered by a prescription as weights. To calculate the weight of each row, we first divided the number of pills dispensed by the pills per day. Then, in each row, we multiplied the mg by the weight for that row. Then we summed the weighted mg in row *k* and *k+1*. Finally, to calculate the average mg in an interval, the sum is divided by the total days covered ("Total days"; see previous paragraph for its calculation) in that interval.

Example applications

Example 1: patient with irregular dosing behavior

The information in the first five columns of Figure 1 comes from a patient with patient number 003011. This patient had irregular dosing behavior: in the first 3.5 months, the patient visits the pharmacy about every 30 days to pick up enough pills for a month, then there is a gap in visits of about 3 months, then a short period with more frequent visits, and then once again a three month gap, and finally another set of frequent visits. For this individual, the length of the first interval was 72 days; starting on the 13th of January 2002 (date of first prescription) and ending on 26st of March 2002 (date of third prescription), thereby covering prescription numbers 1 and 2. Both of these prescriptions were dispensings of 30 pills of which 1 should be taken per day. Using the pills dispensed and the pills per day, we calculated the theoretical days covered by the drug in each interval. For the first interval, this was 30/1 + 30/1 = 60 days. To illustrate how this changes when the number of pills changes, in the sixth interval this was 60/2 + 75/2 = 67.5 days. Finally, adherence was calculated by dividing the theoretical days covered by the length of the interval. For the first interval this was 60/72 = 0.83 (rounded down). Stockpiling can be witnessed in the fifth interval; here the adherence value would exceed one (60/56 = 1.07). Therefore, the 4 excess days are carried forward to the sixth interval, which as a consequence receives the adherence value of (67.5 + 4)/117= 0.61, representing drug stockpiling (Figure 2). The last adherence value that can be calculated using this algorithm is the one corresponding to the third to last row. Using this adherence value (0.91) and placing it in the second to last row, we can calculate the length of the last interval, which is 60 / 0.91 = 66 days. The total length in days that we follow this patient is the difference between the first date (13th of January 2002) and the date of the start of the last interval (5th of February 2003), plus 66 days (length of last interval). This is 389 + 66 = 455 days. We then assigned adherence values to each individual day by using the adherence from

the interval that ends earliest after that day; therefore, all days from 13th of January 2002 to 26st of March 2002 were assigned an adherence of 0.83 (adherence of the first interval), the days from 27th of March 2002 to 1st of May were assigned an adherence of 0.82 (adherence of the second interval), etc.

Plotting these adherence measures in a graph shows that the method adequately captures the irregular dosing behavior of the patient; the estimates of adherence fluctuate strongly over the time period (Figure 3). If we had instead made a time-constant PDC estimate of drug adherence over a 1 year time period, we would have counted the total days covered in the first year, noting that from the last dispensing in the first person-year only 8 pills can still be used in this year, we get (30/1 * 3 + 60/2 * 3 + 75/2 + 8)/365 = 0.62 (Figure 3). Compared to the information generated by the time-varying adherence method, this number provides very little information about the actual adherence behavior. Furthermore, the interval stops after 1 year, while the patient continues to receive the drug for about three additional months. Finally, because we also had information on other variables that were measured every 30 days for each patient, we chose to aggregate the adherence measurements to 30-day periods. The first two 30-day periods are in the period 13th of January 2002 to the 13th of March, 2002. Since the first interval does not end until the 25th of March, 2002, we can simply assign the adherence of the first interval (0.83) to the first two 30-day periods. The third 30-day period goes from the 14th of March, to the 12th of April. We therefore calculated adherence in this 30-day period as (12*0.83+18*0.82)/30 = 0.824. In the fourth period, it became (18*0.82+12*0.46)/30 = 0.676, etc. Using 30-day periods has a smoothing effect on the dynamic adherence measurements, but these still provide more detailed information than a time-constant measurement (Figure 3).

Example 2: patient with low intensity dosing behavior and with regular visits

Contrasting irregular dosing behavior, consider a patient that is not fully adherent but with a stable regularity of pharmacy visits (Figure 4). The patient started with a lower dose (prescribed 1 pill per day) for the first 30 days, and afterwards received a dose that should have lasted for 60 days each time, but the patient instead visited approximately every 90 days. Therefore, approximately (60/90 =) 0.66 adherent would be a correct estimate. For this patient, the estimates using the time-varying method are all close together and around 0.66 adherent; correctly showing a stable adherence over time. However, in this example, a time-constant adherence estimate over the 365 day period (1 year) would be biased upwards. The patient would be calculated as being covered for (30/1 + 60/1 + 3*120/2) = 270 days. Since 270/365 = 0.74, the patient was estimated to be more adherent than in the time-varying estimation. The reason for this is that the final interval that falls within the 365 day range (drugs dispensed in row 5, Figure 4) occurred on 2nd of March, 2007, and pills to last for (120/2 =) 60 days were dispensed on that date. Therefore, these pills could be said to have lasted until the 1st of May, 2007, while the 365 day period ends on 7th of May, 2007. Therefore, using the logic of time-constant PDC adherence calculation, all of these pills from the last dispensing could be used in this year. The problem with this is that the time-constant method does not take into account the timing between intervals; after the batch picked up on the 2nd of March, the next batch was picked up on the 14th of June; more than one month after the 365 static period ended. In other words, the time-constant method here assumes that the patient was highly adherent between the 2nd of March, and the 7th of May, but looking at the timing between intervals shows that this was unlikely to be true. This second example shows the usefulness of the time-varying method in calculating adherence in dynamically generated intervals.

Example 3: patient with time-varying dosage

In this example we show how information on dosage over time can be included in the proposed method. Figure 5 shows the adherence and dosage information of the irregularly dosing patient described in 'example 1' (rows 4 through 7). To calculate the average dosage in an interval, we use the number of pills dispensed as weights; in row 4 and 5, the number of pills dispensed is 60, and pills per day is 2, which means that the drugs in each row should cover 30 days (= 60/2). In row 4, the mg is 40, while in row 5 it is 60. Therefore, the average mg in the fourth interval is [(60/2)*40 + (60/2)*60)]/(60/2 + 60/2) = 50. Similarly, the average mg in the fifth interval is [(60/2)*60 + (60/2)*80)]/(60/2 + 60/2) = 70. In row 7, the number of pills dispensed changes from 60 to 75, so the pills in this row should last longer (as the pills per day remains at 2) than those in row 6. This means row 7 gets more weight in this interval than row 6 does. Therefore, in the sixth interval, the average mg is [(60/2)*80 + (75/2)*100)]/(60/2 + 75/2) = 91.1. Similar to the outcome from 'example 1', a time-constant estimate of drug dosage would produce only a single number, whereas the actually observed dosage changes strongly over time.

Discussion

In this paper we extend the existing PDC method by allowing it to estimate time-varying adherence and dosage. Compared to the time-constant PDC method, we demonstrated that time-varying adherence measures may lead to less biased associations between covariates and clinical outcomes.

Limitations of the time-varying PDC method

Like any method that uses electronic records, observation of drug utilization by the PDC method is indirect; therefore, its most important limitation is that it is unknown if patients actually take the drugs that are prescribed or dispensed. Nevertheless, these methods are considered a good alternative when direct observation of patient adherence is not feasible (e.g. when large sample sizes are desired) [18]. Note that the method can also be seen as a direct method of obtaining rather than taking the drug, which has also been shown to predict health outcomes. Furthermore, the method cannot determine adherence when less than 3 dispensings (or prescriptions) have been recorded, and therefore is not applicable to determine the effect of primary or early nonadherence on an outcome. We have here chosen to make intervals for adherence calculation on the basis of the timing of 3 dispensings; this choice is a bias-variance tradeoff; intervals based on more dispensings will vary less, but as interval size increases the measure will become more time-constant and thereby be subjected more to the biases that we have shown to be present in such a measure. The main limitation of the *time-varying* PDC method is likely the difficulty in estimating adherence for the final interval. We have suggested to continue with the previously observed adherence value, so that the length of the final interval can be determined. However, this assumption may not be realistic, depending on the setting. For example, the assumption is likely valid for patients with stable adherence behavior such as the patient from example 2, but may be less correct for patients

with erratic adherence behavior, such as the patient from example 1. Furthermore, for some drugs a tapering off period may be indicated by guidelines. In that instance, the period as indicated by the guideline could be substituted, granted that this does not directly interfere with the research objective. A more data-driven alternative would be to model the adherence trajectory based on some number of final observations for each individual patient and to extrapolate that pattern to estimate adherence in the final interval. A similar assumption is in place at the start of the interval, but is less apparent because we choose the date of dispensing as the start-date of being covered by the drug, while the true start date is unobserved.

In a sensitivity analysis using data from 100 real patients on statin therapy over 25 intervals (of 30 days each), we found that for 80 patients the time-constant measure was on average within 8% of the time-varying measure. However, for 20 patients, the difference was much larger; these patients were largely irregularly dosing patients (see appendix).

Medication Possession Ratio

The (time-fixed) PDC method is similar to the Medication Possession Ratio (MPR) [19]. The PDC method results in estimates between 0 and 1, while the MPR can exceed 1. It should be possible to extend the MPR into a time-varying method, using a technique similar to the one presented in this paper. We have here chosen to extend the PDC method because this method was evaluated more positively [16].

Drug switching

We proposed a way to include drug switching by calculating adherence values of the old and new drug together, ordered chronologically. While this is technically feasible, this choice depends also on clinical sensibility; the new drug likely has other properties than the old, and may consequently have other effects on the outcome. In such a case, if those other properties

are relevant to the study at hand, it may be better to stop the adherence calculation for the old drug, with the day of switching as the stop date, and possibly calculate adherence for the new drug, starting from the moment of switching (if both drugs are included in a single analysis, e.g. identified through an indicator variable).

Misclassification of exposure

By using a time-varying adherence measure, the effect of interactions between drugs being used at the same time may be more realistically investigated than with a time-constant adherence measure. For example, using a time-constant measure during some fixed time interval, a patient may be 50% adherent to drug A and 40% adherent to drug B. Both drugs could have been used at low intensity throughout the whole time period, in which case they may have interacted. However, it is also possible that drug A was used intensively in the first half of the interval and drug B intensively during the second half of the interval; this means they would not have been used in the same time and therefore would not have interacted. If our time-varying method is applied to drug A and B separately, and adherence values then put on the same time axis, these two scenarios can be better distinguished from each other.

Longitudinal modelling

The time-varying PDC method is primarily intended for use in longitudinal analysis. In longitudinal analysis, time-varying adherence and dosage can be used either as outcomes or as explanatory variables. A major strength of following adherence within patients over time is that, depending on the study design, patients can act as their own control; the effect of changing adherence on some outcome can be measured within a patient. This design automatically controls for between-patient confounding factors. In a design where each patient has only one adherence value, the effect of adherence can only be assessed by doing a

between-patient comparison, e.g. comparing the outcomes of low adhering patients with those of high adhering patients. However, when using time-varying covariates, the possibility of time-varying confounding may arise and should therefore be considered [20]. Time-varying confounding can be guarded against by considering a causal diagram of the study and dealt with by using methods such as inverse probability weighting or the G-formula [21, 22].

When using time-varying covariates in general, including time-varying adherence, it may be wise to introduce time lag between the values of the covariate and the outcome. That is, the outcome at any point in time can be related to the adherence value that was observed a few days, weeks or even months earlier. Without time lag, the causal relations between variables may be reversed; for example, in a study of the effect of adherence on disease onset, we would expect the adherence value to affect disease onset, and not vice versa. However, patients with a worsening health condition may also become less adherent, thus the causal relations can become reversed. Implementing a time lag can prevent this from occurring. The exact size of the time lag is dependent on the study objective and drug in question; a time lag can be large if the drug is believed to have long term effects, but must be short if the drug primarily has short term effects.

For some drugs, it may be assumed that a patient's larger adherence history also plays a role. This could be represented by a variable that, at any time, contains the sum, mean, or some other mathematical transformation of observed adherence values of previous time points, depending on what is clinically sensible.

In longitudinal analysis, it is commonly a requirement that the timescale of the exposure (e.g. adherence) corresponds the timescale of the outcome. In our method, we have demonstrated how to calculate adherence in 30-day intervals, which is the shortest supply period for many chronic medications, and have noted that it can easily be changed to longer time periods. This makes the method especially suitable for measuring the associations

between adherence and chronic conditions, the aspects of which would also change in the scale of weeks or months. Other methods, for example those designed to use data from electronic monitoring of medication taking, should be employed for studying the associations with outcomes that vary on a daily or hourly basis.

Especially when building predictive models, using information from future observations should be limited, so as not to artificially increase the predictive power of a model. For this reason, when multiple intervals overlap, the adherence value that we assign to patients at any time point comes from the estimated adherence of the interval that will end most soon after that point. Finally, when the adherence variable is used as an outcome instead of as an explanatory variable, it should be noted that adherence observations within a patient that are close to each other in time are likely correlated. This should be taken into account by modelling some covariance structure for the data, such as an autoregressive covariance structure.

Adherence is often measured as a continuous variable, but then dichotomized [23]. For example, patients that are below 0.8 adherent may be categorized as non-adherent, whereas patients with 0.8 adherent or more are considered adherent. The time-varying adherence measure described in this paper may be similarly dichotomized, though we suggest that this is not needed; firstly, it is often unclear what the choice of the cutoff value should be based on, and secondly, by keeping the adherence as a continuous variable and using squared terms or splines (e.g. [24-26]), the response curve between adherence and some outcome may be described in greater detail. Knowing the response curve in detail can be useful because a desired outcome may already be achieved at lower levels of adherence. In such a situation, resources that would otherwise have been spent to achieve higher adherence levels in patients can be saved [3].

Similar patients with the same adherence over time may still have different clinical outcomes due to their dosage differing. In this paper we have also demonstrated how time-varying dosage can be estimated. We suggest that the dosage variable can be used in addition to the time-varying adherence variable. For this reason, our calculation of drug dosage did not take into account the total length of each interval; information on the total length of an interval was already used to estimate adherence. Depending on the research question, an interaction effect between drug adherence and dosage can be considered: by doing so, the model takes into account that the effect of drug dosage could differ between low adhering patients and highly adhering patients.

Conclusion

Accurate measurements of adherence are essential for the assessment of pharmacologic interventions. We have demonstrated that the extended proportion of days covered method better accounts for changes over time in drug utilization behavior, such as being better able to discern erratic dosing from continuous low intensity dosing behavior and the patient's regularity of visits to the pharmacy. This may improve longitudinal or time-to-event studies that associate adherence with another outcome, or (intervention) studies that seek to describe changes in adherence over time.

Conflict of interest statement

All authors report no conflict of interests.

Funding

This work was supported by means of an unrestricted personal grant by the Ubbo Emmius Programme of the University of Groningen to M. J. B.

Ethical approval

Ethical approval was not required to perform this study.

References

[1] Steiner JF. Rethinking adherence. Ann Intern Med. 2012;157(8):580-5.

[2] Anonymous. Patient compliance in therapeutic trials. Lancet 1991;337(8745):823–4.

[3] Arnet I, Abraham I, Messerli M, Hersberger KE. A method for calculating adherence to polypharmacy from dispensing data records. Int J Clin Pharm. 2014;36(1):192-201.

[4] Servick K. 'Nonadherence': a bitter pill for drug trials. Science. 2014; 346(6207):288-9.

[5] Andrade SE, Walker AM, Gottlieb LK, Hollenberg NK, Testa MA, Saperia GM, Platt R. Discontinuation of antihyperlipidemic drugs – do rates reported in clinical trials reflect rates in primary care settings? N Engl J Med. 1995; 332: 1125-1131.

[6] Farmer KC. Methods for measuring and monitoring medication regimen adherence in clinical trials and clinical practice. Clin Ther. 1999;21(6):1074-90.

[7] Potter LS. Oral contraceptive compliance and its role in the effectiveness of the method. In: Cramer JA, Spilker B, eds. Patient Compliance in Medical Practice and Clinical Trials. New York: Raven Press; 1991:195-207.

[8] Franklin JM, Shrank WH, Pakes J, Sanfélix-Gimeno G, Matlin OS, Brennan TA, Choudhry NK. Group-based trajectory models: a new approach to classifying and predicting long-term medication adherence. Med Care. 2013;51(9):789-96.

[9] Ho PM, Bryson CL, Rumsfeld JS. Medication adherence. Circulation. 2009;119:3028–3035.

[10] Jeffery RA, Navarro T, Wilczynski NL, Iserman EC, Keepanasseril A, Sivaramalingam B, Agoritsas T, Haynes RB. Adherence measurement and patient recruitment methods are poor in intervention trials to improve patient adherence. J Clin Epidemiol. 2014;67(10):1076-82.

[11] Hansen RA, Kim MM, Song L, Tu W, Wu J, Murray MD. Comparison of methods to assess medication adherence and classify nonadherence. Ann Pharmacother. 2009 Mar;43(3):413-22.

[12] Clifford S, Perez-Nieves M, Skalicky AM, Reaney M, Coyne KS. A systematic literature review of methodologies used to assess medication adherence in patients with diabetes. Curr Med Res Opin. 2014;30(6):1071-85.

[13] Vrijens B, Goetghebeur E. Comparing compliance patterns between randomized treatments. Control Clin Trials. 1997; 18(3):187-203.

[14] Clinical Practice Research Datalink. Internet: http://www.cprd.com/intro.asp Last visited 15 december 2014.

[15] The Mondriaan Project. Internet: http://www.tipharma.com/pharmaceutical-research-projects/completed-projects/mondriaan-project.html Last visited 15 december 2014.

[16] Martin BC, Wiley-Exley EK, Richards S, Domino ME, Carey TS, Sleath BL. Contrasting measures of adherence with simple drug use, medication switching, and therapeutic duplication. Ann Pharmacother. 2009;43(1):36-44.

[17] R Core Team. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. Internet: http://www.R-project.org/.

[18] Steiner JF, Prochazka AV. The assessment of refill compliance using pharmacy records: methods, validity, and applications. J Clin Epidemiol. 1997;50(1):105-16.

[19] Fairman K, Motheral B. Evaluating medication adherence: which measure is right for your program? Journal of Managed Care Pharmacy. 2000;6: 499-506.

[20] Platt RW, Schisterman EF, Cole SR. Time-modified Confounding. Am. J. Epidemiol. 2009; 170 (6): 687-694.

[21] Daniel RM, Cousens S, De Stavola B, Kenward M, Sterne J. Methods for dealing with time-dependent confounding. Stat Med. 2013;32(9):1584-618.

[22] Young JG, Cain LE, Robins JM, O'Reilly EJ, Hernán MA. Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula Stat Biosci. 2011;3(1):119-143.

[23] Karve S, Cleves MA, Helm M, Hudson TJ, West DS, Martin BC. Good and poor adherence: optimal cut-point for adherence measures using administrative claims data. Curr Med Res Opin. 2009 25(9):2303-10.

[24] Helms HJ, Benda N, Zinserling J, Kneib T, Friede T. Spline-based procedures for dose-finding studies with active control. Stat Med. 2015; 34(2):232-48.

[25] Gurrin LC, Scurrah KJ, Hazelton ML. Tutorial in biostatistics: spline smoothing with linear mixed models. Stat Med. 2005;24(21):3361-81.

[26] Tilling K, Macdonald-Wallis C, Lawlor DA, Hughes RA, Howe LD. Modelling childhood growth using fractional polynomials and linear splines. Ann Nutr Metab. 2014;65(2-3):129-38.

Figure legends

Figure 1. Electronic records of a patient with irregular dosing behavior. Rows of pharmacy dispensing records showing patient ID, date of dispensing, number of pills dispensed and pills per day. Interval length, total days and adherence are added later; they are intentionally left blank for the 11th row because that row belongs to a new patient.

Figure 2. Incorporating drug stockpiling.

Figure 3. Comparisons of dynamic versus static PDC estimates of drug adherence from a patient with irregular dosing behavior. Each interval is represented by a horizontal line and labeled by # and its number. Interval #8 continues beyond the displayed range.

Figure 4. Electronic records of a patient with low adherence and a stable visit pattern.

Figure 5. Rows of pharmacy dispensing records including dosage information.

	ID	*Prescription number*	Date (dd-mm-yy)	Pills dispensed	Pills per day	*Interval length*	*Total days*	*Adherence*
1	003011	1	13-01-02	30	1	72	60	0.83
2	003011	2	17-02-02	30	1	73	60	0.82
3	003011	3	26-03-02	30	1	131	60	0.46
4	003011	4	01-05-02	60	2	132	60	0.45
5	003011	5	04-08-02	60	2	56	60	1.07?
6	003011	6	10-09-02	60	2	117	67.5	0.58
7	003011	7	29-09-02	75	2	129	67.5	0.52
8	003011	8	05-01-03	60	2	66	60	0.91
9	003011	9	05-02-03	60	2	?	60	?
10	003011	10	12-03-03	60	2	?	?	?
11	088610	1	27-01-94	30	1			

		Interval length	*Total days*	*Adherence*
4		132	60	0.45
5		56	60 - 4 = 56	1.00
6		117	67.5 + 4 = 71.5	0.61
7		129	67.5	0.52

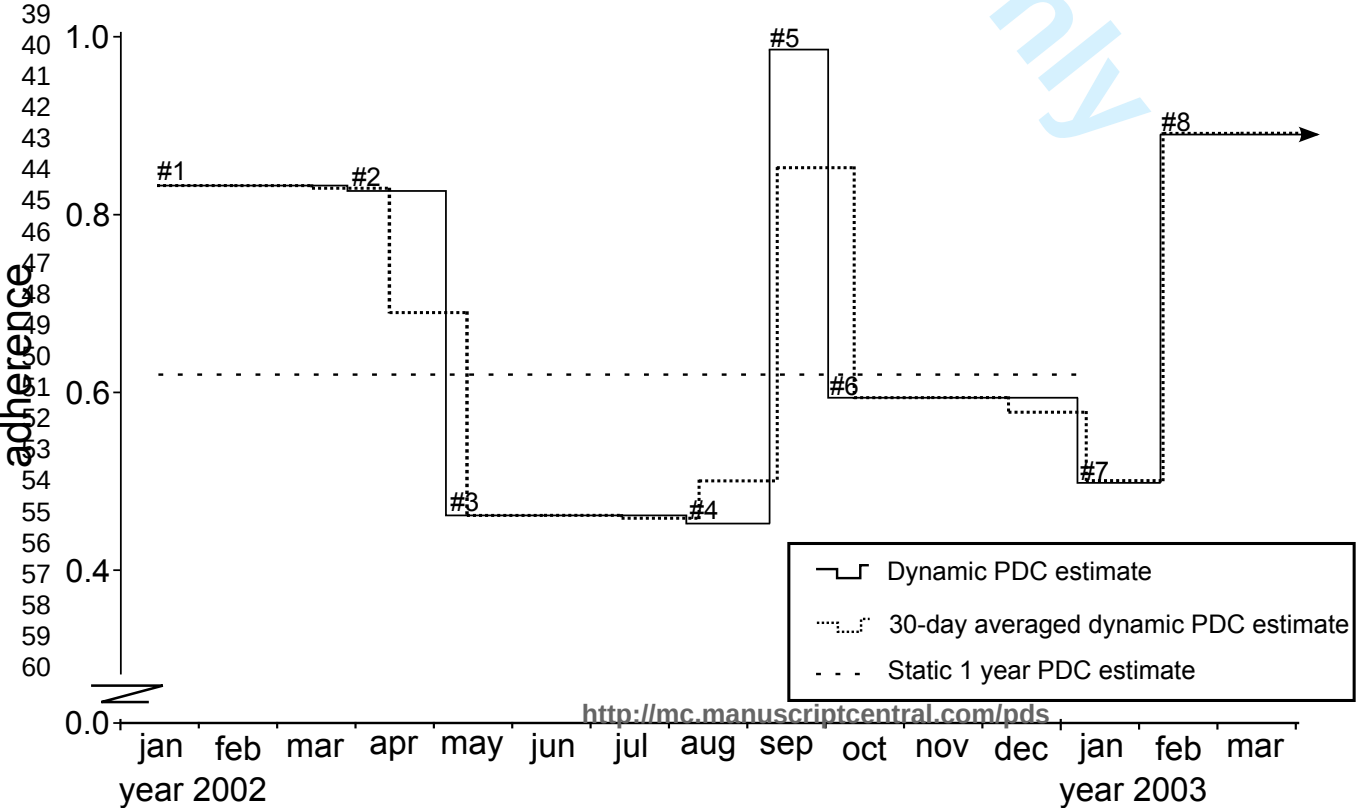

	ID	*Prescription number*	Date (dd-mm-yy)	Pills dispensed	Pills per day	*Interval length*	*Total days*	*Adherence*
225	004312	1	07-05-06	30	1	141	90	0.64
226	004312	2	16-06-06	60	1	179	120	0.67
227	004312	3	25-09-06	120	2	183	120	0.66
228	004312	4	12-12-06	120	2	185	120	0.65
229	004312	5	02-03-07	120	2	185	120	0.65
230	004312	6	14-06-07	120	2			

	ID	Prescription number	Date (dd-mm-yy)	Pills dispensed	Pills per day	Total days	Adherence	mg	Average mg
4	003011	4	01-05-02	60	2	60	0.45	40	50
5	003011	5	04-08-02	60	2	60	1.00	60	70
6	003011	6	10-09-02	60	2	67.5	0.61	80	91.1
7	003011	7	29-09-02	75	2	67.5	0.52	100	etc.

279x361mm (300 x 300 DPI)

Time varying PDC method applied to empirical statin user data

We selected 100 starters of statin therapy (ATC code C10AA) at random from a drug dispensing database (iadb.nl). We applied our method to the dispensing data of these patients, and then calculated the average adherence in the first year, as may be done in a time-constant PDC measure. We then compared the time-dependent adherence (measured in 30-day intervals as explained in the paper, and hence 25 time points for which we have time-varying adherence values) to this time-constant measure in the first two years of follow-up. We summed the absolute difference for each patient (which we shall here refer to as the total absolute residual), in order to identify how much the time-constant measure departs from a time-varying measure. eFigure 1 shows the histogram of these residuals.

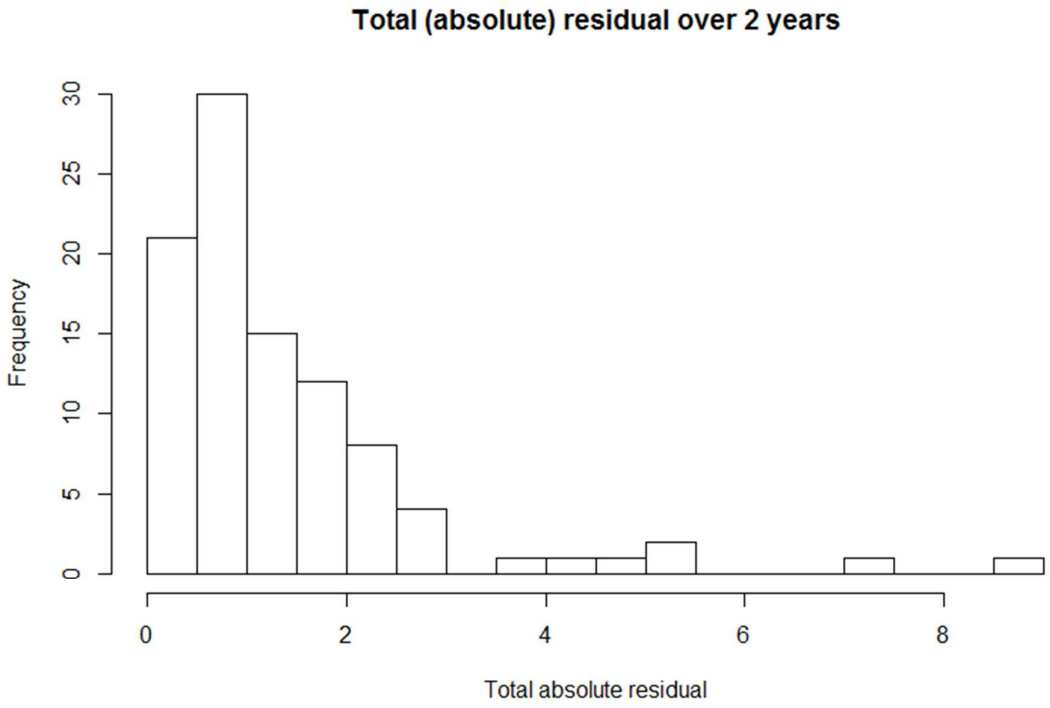

eFigure 1. Total absolute residual of the 100 statin users.

We see that for 80 of the 100 patients, the time-constant measure and the time-varying measure are fairly close; at most a total residual of 2, meaning that on average the time-constant measure was at a distance of about 8 percentage point too high or too low (2/25 = 0.08). This is to be expected, as other studies have shown that time-constant adherence is a decent predictor of various outcomes, which wouldn't be true if the measure did not do a good job of catching a large part of adherence behaviour.

However, for 20% of patients, the distance was much larger. For some of these patients, we show below their adherence trajectories.

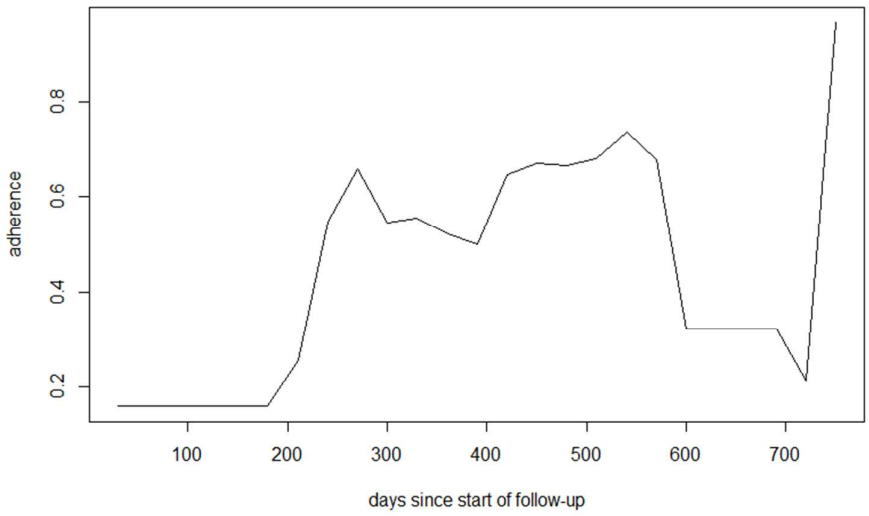

eFigure 2. Adherence value of patient # 8013483

This patient (eFigure 2) initially has a low adherence value, and at the end of the first year of follow-up has an increase in adherence. It then remains stable, then drops again, and at the very end of follow-up a strong improvement appears to occur. Clearly, such behavior cannot be captured in a time-constant adherence measure.

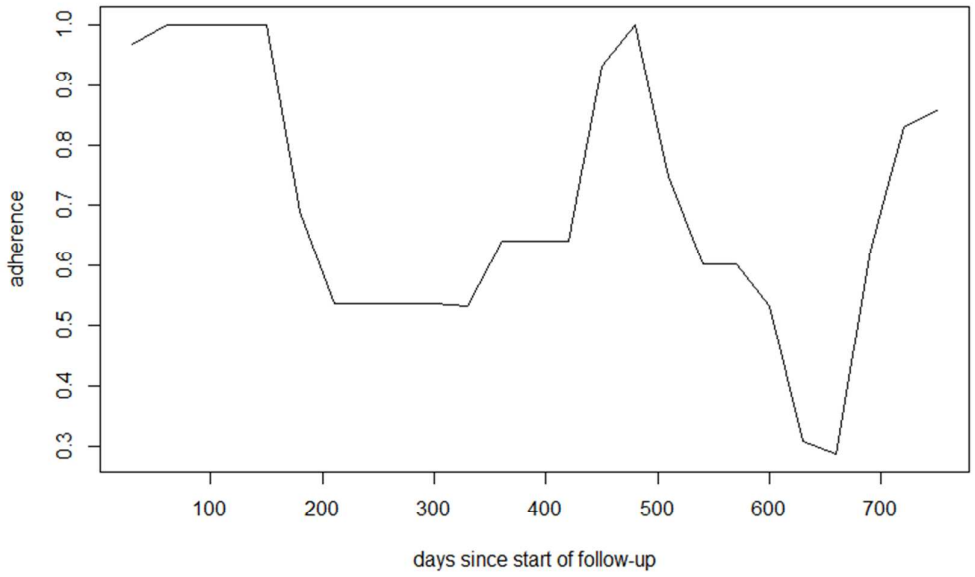

eFigure 3. Adherence of patient 29034318

The patient in eFigure 3 is similar to the patient from eFigure 2, but starts out adherent and then has a drop in adherence.

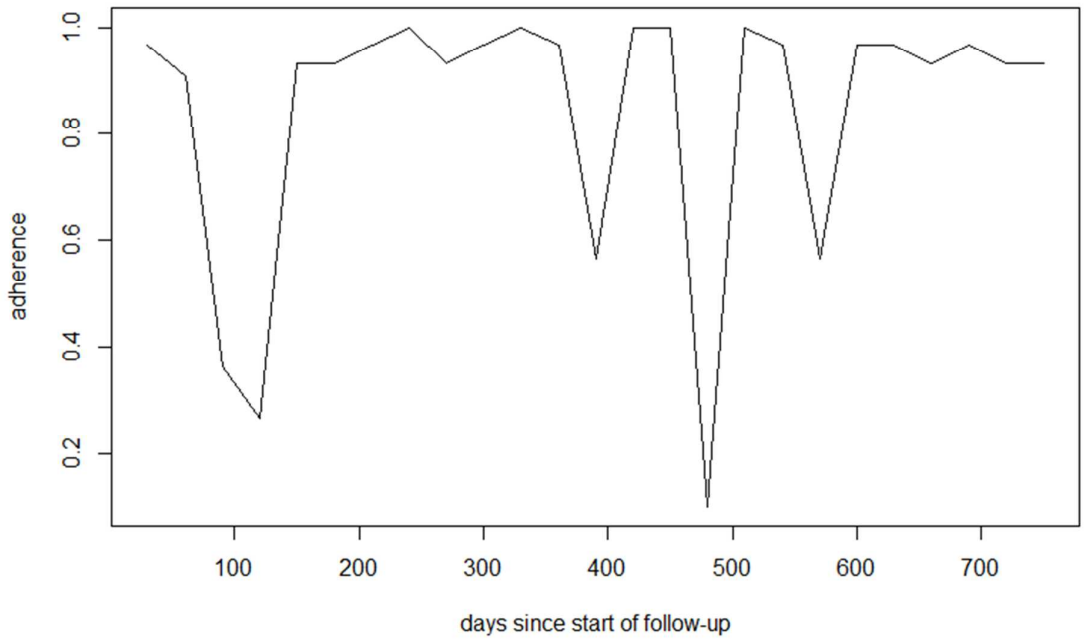

eFigure 4. Adherence of patient 17009932

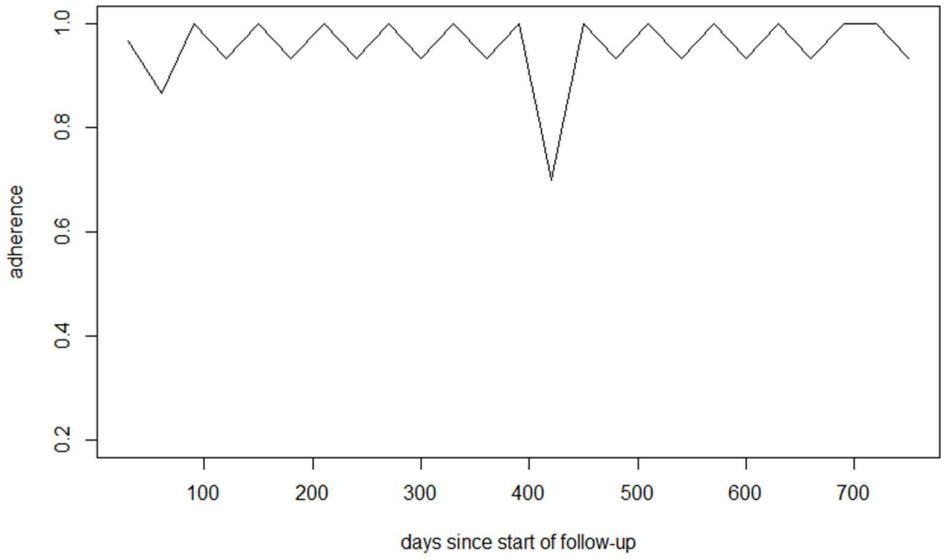

eFigure 5. Adherence of patient 25033087

eFigure 4 shows a patient who appears to be dosing irregularly, as the dosage sees a number of large spikes. We see a similar pattern in eFigure 5, except that here it occurs on a much smaller scale (y-axis scales correspond to show this). For these patients, due to the regularity of the oscillating adherence patterns, while these patients may truly be irregular dosers, the oscillating pattern gives an indication that perhaps a 60-day averages should be used instead of a 30-day averages; in that case, the patient in Figure 5 would have a fairly smooth adherence over time, with a dip at around day 400.

We also applied the method to a much larger sample of ca. 50,000 statin users (from initiation of statin therapy onwards) and for a much longer time period. eFigure 6 shows the mean adherence over time in that group of patients. This figure is for reviewers only, as it will be used in a study of the effect of time-varying adherence on cardiovascular mortality. The figure shows that in the first 3 years of followup, on average there was a decline is adherence to statin therapy, and hence a single measure over the first year would not be a correct summary measure of overall adherence. After these 3 years, adherence in our sample stabilized.

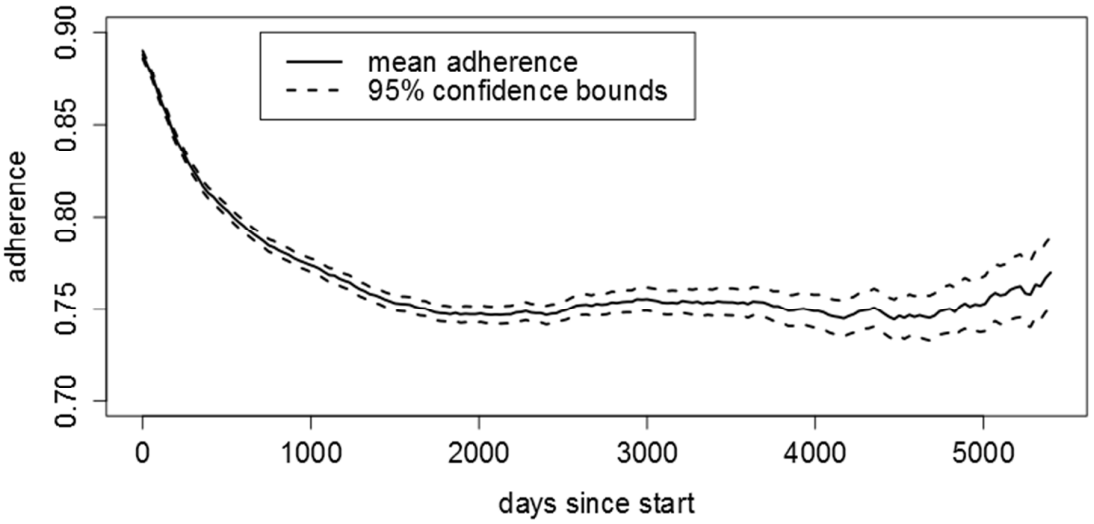

eFigure 6. Mean adherence over time in a cohort of ca. 50,000 statin users.