Romain Neugebauer*, Malini Chandra, Antonio Paredes, David J. Graham,
Carolyn McCloskey and Alan S. Go

A Marginal Structural Modeling Approach with Super Learning for a Study on Oral Bisphosphonate Therapy and Atrial Fibrillation

Abstract: *Purpose:* Observational studies designed to investigate the safety of a drug in a postmarketing setting typically aim to examine rare and non-acute adverse effects in a population that is not restricted to particular patient subgroups for which the therapy, typically a drug, was originally approved. Large healthcare databases and, in particular, rich electronic medical record (EMR) databases, are well suited for the conduct of these safety studies since they can provide detailed longitudinal information on drug exposure, confounders, and outcomes for large and representative samples of patients that are considered for treatment in clinical settings. Analytic efforts for drawing valid causal inferences in such studies are faced with three challenges: (1) the formal definition of relevant effect measures addressing the safety question of interest; (2) the development of analytic protocols to estimate such effects based on causal methodologies that can properly address the problems of time-dependent confounding and selection bias due to informative censoring, and (3) the practical implementation of such protocols in a large clinical/medical database setting. In this article, we describe an effort to specifically address these challenges with marginal structural modeling based on inverse probability weighting with data reduction and super learning.

Methods: We describe the principles of, motivation for, and implementation of an analytical protocol applied in a safety study investigating possible effects of exposure to oral bisphosphonate therapy on the risk of non-elective hospitalization for atrial fibrillation or atrial flutter among older women based on EMR data from the Kaiser Permanente Northern California integrated health care delivery system. Adhering to guidelines brought forward by Hernan (Epidemiology 2011;22:290–1), we start by framing the safety research question as one that could be directly addressed by a sequence of ideal randomized experiments before describing the estimation approach that we implemented to emulate inference from such trials using observational data.

Results: This report underlines the important computation burden involved in the application of the current R implementation of super learning with large data sets. While computing time and memory requirements did not permit aggressive estimator selection with super learning, this analysis demonstrates the applicability of simplified versions of super learning based on select sets of candidate learners to avoid complete reliance on arbitrary selection of parametric models for confounding and selection bias adjustment. Results do not raise concern over the safety of one-year exposure to BP but may suggest residual bias possibly due to unmeasured confounders or insufficient parametric adjustment for observed confounders with the candidate learners selected.

Conclusions: Adjustment for time-dependent confounding and selection bias based on the ad hoc inverse probability weighting approach described in this report may provide a feasible alternative to extended Cox modeling or the point treatment analytic approaches (e.g. based on propensity score matching) that are often adopted in safety research with large data sets. Alternate algorithms are needed to permit the routine and more aggressive application of super learning with large data sets.

Keywords: marginal structural model, inverse probability weighting estimation, super learning, data reduction, comparative effectiveness and safety, atrial fibrillation, bisphosphonate therapy

***Corresponding author: Romain Neugebauer**, Division of Research, Kaiser Permanente Northern California, Oakland, CA, USA, E-mail: romain.s.neugebauer@kp.org

Malini Chandra, Division of Research, Kaiser Permanente Northern California, Oakland, CA, USA, E-mail: malini.chandra@kp.org
Antonio Paredes, Office of Biostatistics, Center for Drug Evaluation and Research, Food and Drug Administration, Silver Springs, MD, USA, E-mail: antonio.paredes@fda.hhs.gov

David J. Graham, Office of Surveillance and Epidemiology, Center for Drug Evaluation and Research, Food and Drug Administration, Silver Springs, MD, USA, E-mail: David.Graham1@fda.hhs.gov

Carolyn McCloskey, Office of Surveillance and Epidemiology, Center for Drug Evaluation and Research, Food and Drug Administration, Silver Springs, MD, USA, E-mail: Carolyn.McCloskey@fda.hhs.gov

Alan S. Go, Division of Research, Kaiser Permanente Northern California, Oakland, CA, USA; Departments of Epidemiology, Biostatistics and Medicine, University of California, San Francisco, CA, USA; Department of Health Research and Policy, Stanford University, Palo Alto, CA, USA, E-mail: alan.s.go@kp.org

1 Introduction

Osteoporosis and atrial fibrillation are two parallel epidemics that primarily affect the rapidly growing population of elderly persons in the US and worldwide [1, 2]. Osteoporotic fractures are associated with major morbidity and excess mortality which prompted the development of pharmacological approaches to lower this risk. Oral bisphosphonate drugs are one of the primary therapies for the primary and secondary prevention of osteoporosis and osteoporotic fractures. Atrial fibrillation also predominantly affects older persons (≈ 1 in 25 adults aged ≥ 60 years and ≈ 1 in 10 adults aged 80 years) and independently increases by fourfold to fivefold the risk of ischemic stroke [3]. Interestingly, recent data from randomized clinical trials [4, 5] suggest a possible excess risk of clinically significant atrial fibrillation events associated with oral bisphosphonate therapy. Subsequent observational studies have reported conflicting results [6–12]. However, existing data from clinical trials and observational studies have several important limitations including relatively low numbers of confirmed cases of atrial fibrillation, selected populations that are not necessarily representative of typical clinical settings, and other significant methodological limitations, including approaches to ascertaining atrial fibrillation and ability to control for confounding and selection bias. Given the widespread and increasing use of bisphosphonates for the prevention of osteoporotic fractures, delineation of any increased risk of clinically significant atrial fibrillation events has major public health implications, particularly in the elderly. To address these issues, we conducted a retrospective, cohort study of the effects of incident one-year exposure to oral bisphosphonate (BP) therapy on the risk of clinically significant atrial fibrillation (AF) events in patients with and/or without a prior history of AF based on Electronic Medical Record (EMR) data from a large, diverse community-based cohort of older women.

2 Method

2.1 Data

Study subjects included members of Kaiser Permanente Northern California (KPNC), a large, integrated health care delivery system that provides comprehensive medical care to >3.2 million members representing approximately one-third of insured adults in northern California. Health plan members are highly representative of the local surrounding and statewide population except for slightly lower representation of the extremes of age and income [13, 14]. We searched the entire adult membership of KPNC for patients meeting

the eligibility criteria described below and enrolled each in this study at the earliest date between 1 January 2002 and 31 December 2007 on which all criteria were met. These criteria identified 673,149 patients. All patients were followed up from study entry until the earliest of 31 December 2008 (administrative end of study), health plan disenrollment or death. Patients eligible to enter the study cohort were female aged >50 years who had continuous health plan membership of at least 12 months before entry and active pharmacy benefit which provides strong financial incentives to obtain medications from Kaiser Permanente pharmacies. We excluded subjects who had evidence of intravenous bisphosphonate therapy found in pharmacy databases before or during the follow-up period, oral BP therapy documented from pharmacy databases between 1995 and cohort entry, received etidronate before or during the follow-up period (given that it is a first-generation BP with cyclic prescription during the study period), or had diagnosed Paget's disease¹ of the bone before or during the follow-up period based on information from health plan hospital discharge and ambulatory diagnosis databases using International Classification of Diseases, Ninth Edition (ICD-9) code 731.0.

We assembled data on relevant exposure, outcomes, covariates, and right-censoring events for each study subject as described below. We focused on oral (nitrogen-containing) BPs (alendronate, risedronate, and ibandronate) exposure ascertained from information found in health plan outpatient pharmacy databases. We expanded upon previously described approaches [15–18] to facilitate estimation of time-varying current exposure and cumulative exposure to oral bisphosphonates during the study period. Specifically, we determined the use, frequency, and timing of oral BP therapy based on a systematic review of all possible combinations of physician prescription instructions, drug dose, drug frequency, quantity of pills dispensed, data of dispensing, and refill pattern for each individual agent (alendronate, risedronate, ibandronate) and oral formulation found in the cohort using information from filled/dispensed prescriptions in health plan pharmacy databases. We constructed a computerized algorithm with iterative refinements based on serial sampling of cohort members who underwent manual chart review to compare the accuracy of assigned drug exposure based on our algorithm compared with information in the electronic and paper medical record. As needed, individual electronic prescription data were corrected within the algorithm. After completion of the final longitudinal drug exposure algorithm and validation studies, we considered patients to have continuous drug exposure if there was a gap between consecutive prescriptions that was 60 days or less. For drug exposure gaps of >60 days, the patient was considered off BP therapy until the next filled prescription. In addition, if there was an overlap in consecutive drug prescriptions of 30 days or less, we added the number of overlapped days to the end of the next filled prescription. If the overlap was >30 days, we reset the end of the drug exposure to the end date of the second prescription. Deaths were identified from health plan administrative databases, proxy report, Social Security Administration vital status files, and California state death certificate information [19, 20]. Complete state death certificate information was only available through 2008 at the time of analysis. The primary outcome was a clinically relevant atrial fibrillation event defined as a non-elective hospitalization with a primary discharge diagnosis of atrial fibrillation (ICD-9 code 427.31) or atrial flutter (ICD-9 code 427.32). Among the 10,285 identified admissions in the cohort meeting this criterion, 7,918 (73.1%) had one or more electrocardiograms documenting AF and/or flutter identified from a health plan electrocardiography database. Among the remaining 2,367 admissions without a corresponding electronically available electrocardiogram documenting AF and/or atrial flutter, we conducted a manual review of corresponding medical records for a random sample of 408 non-elective admissions which showed that 378 (92.7%) had electrocardiographic evidence of AF and/or atrial flutter that was causing symptoms and/or clinical complications. Demographic characteristics (age, gender, and self-reported race/ethnicity) were identified from health plan databases. Information on annual household income and educational attainment was ascertained based on geocoding information from the 2000 US Census data and member's residential address information. Low annual income was

¹ Patients with Paget's disease were excluded because the administered BP dose is typically much higher and on an intermittent basis for treatment of a specific bone condition while this study focuses on BP exposure for primary or secondary prevention of osteoporotic fractures.

defined as living in a census block group with median annual household income $< \$35,000$; low education attainment was defined as living in a census block group where $\geq 25\%$ of residents have completed high school or less. We searched health plan automated databases (e.g. hospitalization, ambulatory visit, laboratory, and pharmacy) as well as regional cancer registry data to identify risk factors for the outcomes of interest and other potential confounders [3, 15, 21]. These include prior cardiovascular disease (coronary disease, aortic or mitral valvular heart disease, ischemic stroke or transient ischemic attack, heart failure, peripheral arterial disease), coronary revascularization, diabetes mellitus, diagnosed hypertension, dyslipidemia, hypothyroidism or hyperthyroidism, chronic lung disease, chronic liver disease, diagnosed dementia or psychiatric illness, and systemic cancer. Using data from dispensed prescriptions in pharmacy databases only (days supply in conjunction with dispense date), we characterized time-varying exposure to other prescription medications including other therapies for osteoporosis (e.g. hormone replacement therapy, raloxifene, calcitonin, parathyroid hormone), for atrial fibrillation (e.g. digoxin, antiarrhythmic medications), for other cardiovascular conditions (e.g. ACE inhibitor, angiotensin II receptor blocker, diuretic, β -blocker, calcium channel blocker, aldosterone receptor antagonist, other vasodilator, nitrate, digoxin, statin, other lipid-lowering therapy, warfarin, clopidogrel, ticlopidine, Aggrenox), and other relevant medications (e.g. aromatase inhibitors, systemic glucocorticoids, and thyroid hormone replacement therapy). The rules adopted for handling gaps between prescriptions and for overlapping prescriptions of these other medications were the same as the ones used for BP prescriptions.

The complete list of all baseline and time-varying covariates used in this study is described in Tables 1 and 2. Note that all covariates are binary indicators except for “age” which is a continuous variable. Hispanic ethnicity was defined as having Hispanic ethnicity noted in the health plan electronic database, so lack of documented evidence of Hispanic ethnicity was interpreted as “non-Hispanic.” For the two census variables, “hs.or.le” (low neighborhood education) and “med.inco” (low neighborhood income), if no data were available, we imputed these variables with the mode of the observed measurements, i.e. 0. We note that only 3.37% of patients had missing income status and educational attainment status. The comorbidity, hospitalization, and medication use variables were assigned a value of 1 if the patient’s EMR data indicated relevant information about these factors and otherwise they were assigned a value of 0.

2.2 Analytic approach

We formally describe a Marginal Structural Model (MSM) approach using the counterfactual statistical framework on which MSMs are based [22], but the approach is also framed informally by analogy with a series of ideal randomized experiments that we aim to emulate with observational data to address the research question of interest [23]. The presentation of the approach is decomposed in three steps. First, we introduce the approach by describing its principles if only data from the first year of follow-up were available to address the research question of interest. Second, we extend these principles by generalizing the approach to incorporate the complete follow-up data from each patient. Third, we adapt the approach for implementation with a reduced data set that is derived from the original data for the purpose of reducing computational burden.

2.2.1 Observed data structure

The observed data on any given patient in this study (referred to as the “BP/AF study” from now on) consist of the collected measurements on exposure, outcome, and confounding variables from study entry until the patient’s end of follow-up. In this analysis, time is measured by consecutive 30-day intervals after study entry. The time when the patient’s longitudinal data are right-censored is denoted by C and is defined by the

Table 1 Covariates used for confounding and selection bias adjustment in the BP/AF study (Part I).

Variable name	Time varying	Known AF risk factor	Attribute description
ace	✓		Ace Inhibitor use
af.outc2	✓		History of hospitalization with primary discharge diagnosis (PDD) of AF
afib.hx	✓	✓	History of AF
age	✓	✓	Age
aggrenox	✓		Aggrenox use
aldoster	✓		Aldesterone use
alpha	✓		Alpha blocker use
ami.usa	✓		Hospitalization with PDD of acute myocardial infarction (AMI)/unstable angina
ami.usa.	✓	✓	History of AMI/unstable angina
arb	✓		Angiotensin receptor blocker use
aromatas	✓		Aromatase use
arrhyth	✓		Anti-arrhythmic use
b.blocke	✓		Beta-blocker use
c.blocke	✓		Calcium channel blocker use
calciton	✓		Calcitonin use
cancer.h	✓		History of cancer
chf	✓		Hospitalization with PDD of heart failure
chf.hx	✓	✓	History of heart failure
clopidog	✓		Clopidogrel use
dement.h	✓		History of dementia
depress.	✓		History of depression
digoxin	✓		Digoxin use
diuretic	✓		Diuretics use
dm.hx	✓	✓	History of diabetes
dyslipid	✓		History of dyslipidemia
fx	✓		Hospitalization with PDD of fracture
fx.hx	✓	✓	History of fracture
glucocor	✓		Glucocorticoid use
hispan			Hispanic ethnicity

shortest of (1) the time to the administrative end of study denoted by C^{admin} , (2) the time to disenrollment from the KPNC health plan denoted by C^{disen} , and (3) the time to death denoted by C^{death} , i.e. $C = \min(C^{\text{admin}}, C^{\text{disen}}, C^{\text{death}})$. The type of right-censoring event is encoded with a categorical variable denoted by Γ with possible values 1, 2, or 3 to represent end of follow-up by administrative end of study (i.e. $C = C^{\text{admin}}$), disenrollment from the health plan (i.e. $C = C^{\text{disen}}$), or death (i.e. $C = C^{\text{death}}$), respectively. At each time point $t = 0, \dots, C$, i.e. a 30-day interval poststudy entry, the patient's exposure to BP is represented by the binary variable $A_1(t)$, and the patient's right-censoring status is denoted by the indicator variable $A_2(t) = I(C \leq t)$. The combination $A(t) = (A_1(t), A_2(t))$ is referred to as the action at time t . At each time point $t = 0, \dots, C$, covariates (e.g. diagnosis of coronary disease or exposure to raloxifene) are denoted by the multidimensional variable $L(t)$ and defined as measurements that occur before $A(t)$ or are otherwise assumed not to be affected by the actions at time t or thereafter, $(A(t), \dots, A(C))$. For each time point $t > 0$, the outcome denoted by $Y(t)$ is defined as the indicator of failure (i.e. a clinically relevant AF event) at time $t - 1$ and is an element of the covariates at time t , $L(t)$. By convention, the outcome at baseline (an element of $L(0)$) is set to 0, i.e. $Y(0) = 0$. To simplify notation, we use overbars to denote covariate and action histories, e.g. a patient's action history through time t is denoted by $\bar{A}(t) = (A(0), \dots, A(t))$. We approached the observed data from the BP/AF study as realizations of $n = 673,149$ independent and identically distributed (i.i.d.) copies of $O = (\Gamma, C, \bar{L}(C), \bar{A}(C))$, where n denotes the sample size. The maximum follow-up time K for any

Table 2 Covariates used for confounding and selection bias adjustment in the BP/AF study (Part II).

Variable name	Time varying	Known AF risk factor	Attribute description
hrt	✓		Hormone replacement therapy use
hs.or.le			Lives in census block where $\geq 25\%$ of the residents have completed high school or less
htn.hx	✓	✓	History of hypertension
hyperthy	✓		History of hyperthyroidism
hypothy.	✓		History of hypothyroidism
liver.hx	✓		History of liver disease
lung.hx	✓		History of lung disease
med.inco			Lives in census block where median yearly household income is less than \$35,000
nitrate	✓		Nitrates use
non.lipi	✓		Non-statin lipid-lowering medications use
pad.hx	✓		History of peripheral artery disease
plavix	✓		Plavix use
race.AA		✓	African American
race.AD		✓	Mixed race
race.AP		✓	Asian/Pacific Islander
race.NA		✓	Native American
race.OT		✓	Other/unknown/missing race
race.WH		✓	White
raloxife	✓		Raloxifene use
revasc	✓		Cardiac revascularization procedure (PCI/CABG)
revasc.h	✓	✓	History of cardiac revascularization procedures (PCI/CABG)
statin	✓		Statin use
stroke	✓		Hospitalization with PDD of Stroke/transient ischemic attack (TIA)
stroke.h	✓		History of stroke/TIA
thyroid.	✓		Anti-hyperthyroid medication use
ticlopid	✓		Ticlopidine use
valve.hx	✓		History of heart valve disease
warfarin	✓		Warfarin use

given patient is $K = \max_{i=1, \dots, n} C_i - 1 = 84$ (i.e. 85 30-day intervals). In Section 2.2.5, we describe how EMR data were concretely mapped into this data structure O for each patient.

2.2.2 Simplified definition of the effect of interest

Informally, a causal effect may be defined by analogy with an ideal randomized experiment, i.e. a trial with perfect compliance and no right-censoring. In this study, the effects of incident one-year exposure to BP on the risk of AF could have ideally been studied by randomizing each patient of the study cohort to one of the following two treatment arms at study entry. In the first arm, patients would have remained unexposed to BP for 1 year, while in the second arm, patients would have been exposed to BP for the same time. The difference between the cumulative risk of AF events at 1 year in each arm is one effect measure that we aim to emulate with observational data in this study.

Formally in the counterfactual statistical framework [24], such causal effects are defined as contrasts between the distributions of potential outcomes in a conceptual experiment. We use notation consistent with that introduced in the previous section to describe the data from this experiment. At each time point $t = 0, \dots, K$ of this fictitious experiment, patient's covariates are observed under any given action intervention denoted by $\bar{a}(K) = (\bar{a}_1(K), \bar{a}_2(K))$ such that right-censoring events do not occur before and at time K ,

i.e. $\bar{a}_2(K) = 0$. These covariates, referred to as counterfactual covariates, are denoted by $L_{\bar{a}(K)}(t)$ and satisfy the time ordering assumption formalized by equality: $L_{\bar{a}(K)}(t) = L_{\bar{a}(t-1)}(t)$, i.e. actions after $t - 1$ cannot affect covariate levels at time t . The counterfactual covariates at t , $L_{\bar{a}(t-1)}(t)$, include the counterfactual outcome denoted by $Y_{\bar{a}(t-1)}(t)$. The collection of all counterfactual covariates collected over time during this ideal experiment are referred to as the full data denoted by X , i.e. $X = (\bar{L}_{\bar{a}(K)}(K+1))_{\bar{a}(K) \in \mathcal{A}(K)}$, where $\mathcal{A}(K)$ represents the collection of any possible action interventions such that right-censoring events do not occur before and at time K , i.e. $\mathcal{A}(K) = \{\bar{a}(K) = (\bar{a}_1(K), \bar{a}_2(K)) : \bar{a}_1(K) \in \{0, 1\}^K, \bar{a}_2(K) = 0\}$. The causal effect that was informally defined earlier by analogy with an ideal randomized trial can now be formally defined as a function of the distribution of these full data:

$$\psi_{0,11} = P(T_{\bar{a}^\bullet(11)} \leq 11) - P(T_{\bar{a}^\circ(11)} \leq 11), \quad [1]$$

where $T_{\bar{a}(11)}$ is implied by the full data and represents the counterfactual time to the first failure event under a given action intervention $\bar{a}(11)$ and $\bar{a}^\bullet(11)$ and $\bar{a}^\circ(11)$ represent the two action interventions of interest by which patients do not experience censoring events before and at time 11 (i.e. $\bar{a}_2^\bullet(11) = \bar{a}_2^\circ(11) = 0$) while experiencing continuous BP exposure (i.e. $\bar{a}_1^\bullet(11) = 1$) and no BP exposure (i.e. $\bar{a}_1^\circ(11) = 0$), respectively. More precisely, if the first failure occurs before time $t = 12$ then $0 \leq T_{\bar{a}(11)} \leq 11$ is defined by the time when the counterfactual outcome jumps to 1 for the first time, i.e. $(\bar{Y}_{\bar{a}(11)}(T_{\bar{a}(11)}) = 0, Y_{\bar{a}(11)}(T_{\bar{a}(11)} + 1) = 1)$. Otherwise, if failure does not occur during follow-up under action intervention $\bar{a}(11)$, i.e. $\bar{Y}_{\bar{a}(11)}(12) = 0$, the counterfactual failure time $T_{\bar{a}(11)}$ is defined as ∞ . The causal estimand $\psi_{0,11}$ can be viewed as one parameter of a nonparametric MSM for the distribution of the counterfactual survival times $T_{\bar{a}(11)}$.

2.2.3 Estimation approach

Figure 1 represents possible causal relationships between variables collected during the first two periods (on patients followed-up for at least two periods) in the BP/AF study using a causal diagram [25]. Standard estimation methods (e.g. extended Cox regression) to account for confounding involve modeling approaches

Figure 1 Causal diagram that represents possible causal relationships between a subset of the variables collected during the first two periods of the BP/AF study (on patients followed-up for at least two periods). T represents the failure time, i.e. time to the first AF event, and $I(T \leq 1) = 1 - I(Y(1) = Y(2) = 0)$ represents the indicator of failure occurrence within two periods from study entry. $A_1(0)$ and $A_1(1)$ represent exposure to BP during the first and second periods of follow-up, respectively. $L(0)$ represents baseline confounders of the effect of both early, $A_1(0)$, and late, $A_1(1)$, BP exposures on failure occurrence within two periods from study entry, $I(t \leq 1)$. $L(1)$ represents intermediate covariates that confound the effect of late BP exposure $A_1(1)$ on failure and that are also affected by the early BP exposure $A_1(0)$. Standard modeling approaches are not adequate to account for the time-dependent confounding that may result from covariates such as $L(1)$.

that compare the distributions of outcomes conditional on both the levels of confounders and exposure variables. It is well-established [26–28] that such methods do not properly handle time-dependent confounders affected by early therapy exposure such as covariate $L(1)$ (Figure 1) because (1) inclusion of such confounders in standard models may lead to bias, since they block some of the causal pathways under investigation; and (2) exclusion of such confounders from standard models may lead to bias due to the lack of confounding adjustment. In observational longitudinal studies that aim to compare the effectiveness of alternate treatment interventions, the existence of time-dependent confounders such as $L(1)$ is often expected since the indications for the treatment decisions under study are also typically risk factors for the clinical outcomes of interest (e.g. CD4 counts and antiretroviral therapy in HIV patients [29] or hemoglobin A1c and treatment intensification in diabetes patients [30]). Thus, standard analytic methods are often suspected to lead to biased inferences in such effectiveness studies. In longitudinal safety studies, however, the existence of time-dependent confounders such as $L(1)$ are often not as obvious since indications for the exposure under study such as history of fracture in the BP/AF study are not typically known risk factors for the clinical outcomes of interest. While time-dependent confounding by covariates such as $L(1)$ is not a priori expected in safety studies, bias from lack of or incorrect adjustment for unforeseen time-dependent confounders remains a concern. For example, it may be alleged that occurrence of major clinical events (e.g. hospitalization for AMI in the BP/AF study) that are risk factors for the outcome under study and also possibly the results of deleterious side effects from prior exposure to treatment will also often lead to the early interruption of that treatment. Such concerns motivate the adoption of an inverse probability of action weighted (IPAW) estimation approach [22, 31, 32] in longitudinal safety studies to not only properly handle possible observed time-dependent confounders on causal pathways of interest but also to adjust for possible time-dependent selection bias due to right-censoring [33].

Alternatively, one may attempt to alleviate the problem of time-dependent confounding through an intention-to-treat analytic approach where changes in exposure over time are ignored [34, 35]. Confounding adjustment by baseline covariates ($L(0)$) using standard estimation approaches (e.g. propensity score [36–39] matching) can then provide unbiased estimates of the effect of the *early* therapy exposure ($A(0)$). Such inferences are informative in randomized experiment in which changes in therapies are rare during the follow-up time or else in placebo-controlled trials conducted for evaluating efficacy since inference from an intention to treat (ITT) analysis is then conservative [40]. Placebo-controlled randomized clinical trials conducted for evaluating safety using an ITT approach only can be misleading and have been referred to as “randomized cynical trial” [40], since ITT results in such analyses are anticonservative and may thus conceal deleterious side effects. These considerations support our goal to emulate inference from a placebo-controlled with perfect compliance in the BP/AF study, instead of inference from an ITT analysis of a placebo-controlled trial with possible non-compliance.

Under assumptions detailed below, the causal estimand $\psi_{0,11}$ in the BP/AF study may be estimated consistently based on the following IPAW estimating function^[41]:

$$D_{0,11}(O \mid g, \psi) = \frac{\left(I(\bar{A}(\tilde{T}) = \bar{a}^\bullet(\tilde{T})) - I(\bar{A}(\tilde{T}) = \bar{a}^\circ(\tilde{T})) \right) Y(\tilde{T} + 1)}{\prod_{j=0}^{\tilde{T}} g(A(j) \mid \bar{L}(j), \bar{A}(j - 1))} - \psi, \quad [2]$$

where $I(\cdot)$ and g are the generic notation for indicator variables and conditional probabilities of actions, respectively, and where

$$\tilde{T} = \min(11, \{0 \leq t \leq \min(11, C - 1) : Y(t + 1) = 1\}, C) \quad [3]$$

represents the earliest time when a right-censoring or failure event occurs or when the action interventions of the hypothetical randomized experiment that we aim to emulate would end (i.e. time point 11). For conciseness, we adopt the following notation

$$\mathcal{F}(\bar{A}(\check{T})) = I(\bar{A}(\check{T}) = \bar{a}^\bullet(\check{T})) - I(\bar{A}(\check{T}) = \bar{a}^\circ(\check{T})) \quad [4]$$

to represent the difference between the two indicators of whether the patient's actual exposure to BP and right-censoring status (during the periods that would have been monitored in the ideal trial to be emulated) are concordant with the action interventions in one of the two arms of the ideal trial, i.e. (1) the indicator of both no right-censoring event and continuous exposure to BP up to and at time $\check{T}$, and (2) the indicator of both no right-censoring event and continuous non-exposure to BP up to and at time $\check{T}$.

As shown in Appendix A, the IPAW estimating function [2] is unbiased at $\psi_{0,11}$ under the following identifiability assumptions [28, 42] which insure that the available observed data contain sufficient information to emulate inference from the ideal randomized trial that would address the research question of interest: (1) the sequential randomization assumption (SRA) formalizes the assumption of no unmeasured confounders in longitudinal studies: $X \perp A(t) \mid \bar{L}(t), \bar{A}(t-1)$ for $t = 0, \dots, K$. It cannot be tested with data alone. (2) The experimental treatment assignment (ETA) assumption (also known as positivity assumption) insures non-parametric identifiability of the causal effect of interest. In other words, satisfaction of the ETA assumption permits reliable estimation of the causal effect based on true information in the data instead of extrapolation from limited information based on parametric modeling assumptions [43]. Formally, this assumption insures that any patient may experience any of the interventions of interest at any point in time, regardless of her covariate history: $g(A(t) = a(t) \mid \bar{L}(t), \bar{A}(t-1) = \bar{a}(t-1)) > 0$ for $\bar{a}(K) \in \mathcal{A}(K)$ and $t = 0, \dots, K$.

The following solution of the estimating equations derived from the IPAW estimating function [2] defines an estimator of $\psi_{0,11}$:

$$\psi_{0,11,n} = \sum_{i=1}^n \frac{\mathcal{F}(\bar{A}_i(\check{T}_i))Y_i(\check{T}_i + 1)}{n \prod_{j=0}^{\check{T}_i} g(A_i(j) \mid \bar{L}_i(j), \bar{A}_i(j-1))}. \quad [5]$$

An estimator of the variance of this IPAW estimator can be derived from its influence curve [24, 44] and can be expressed with the following closed-form formula:

$$\frac{1}{n^2} \sum_{i=1}^n \left(\frac{\mathcal{F}(\bar{A}_i(\check{T}_i))Y_i(\check{T}_i + 1)}{\prod_{j=0}^{\check{T}_i} g(A_i(j) \mid \bar{L}_i(j), \bar{A}_i(j-1))} - \psi_{0,11,n} \right)^2. \quad [6]$$

With observational data, implementation of this IPAW estimator first requires estimation of the product of unknown conditional probabilities of actions g (referred to as action mechanism) and which can be decomposed into two products (referred to as the treatment and right-censoring mechanisms, respectively):

$$\underbrace{\prod_{j=0}^{\check{T}} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))}_{\text{action mechanism}} = \underbrace{\prod_{j=0}^{\check{T}} g(A_2(j) = 0 \mid \bar{L}(j), \bar{A}(j-1))}_{\text{right-censoring mechanism}} \times \underbrace{\prod_{j=0}^{\check{T}} g(A_1(j) \mid A_2(j) = 0, \bar{L}(j), \bar{A}(j-1))}_{\text{treatment mechanism}}.$$

The consistency of the IPAW estimator $\psi_{0,11,n}$ thus relies not only on the SRA and ETA assumption but also on consistent estimation of the action mechanism, i.e. proper adjustment for observed confounders. The inference drawn from formulas [5] and [6] is asymptotically conservative when the action mechanism is estimated [24]. Bootstrapping can be used instead to draw inference that incorporates variability from estimation of the action mechanism.

Maximum likelihood estimation based on a priori specified and more or less flexible parametric models have been used most commonly in practice to estimate the action mechanism. In few instances [45, 46],

data-adaptive estimation of the action mechanism has been implemented as an alternative to the more common reliance on a priori specification of parametric models which often do not reflect true subject-matter knowledge and may thus lead to incorrect inferences [47, 48]. Several machine learning algorithms (also known as learners) are potential candidates for estimating the different components of the action mechanism, e.g. random forests [49], least angle regression [50], logic regression [51], classification and regression trees [52], ridge regression [53], multivariate adaptive regression splines [54], Deletion/Substitution/Addition algorithm [55], high-dimensional propensity score algorithm [56], generalized boosted modeling [57], etc. Akin to the selection of a parametric model, the selection of a machine learning algorithm is rarely based on real subject-matter knowledge about the relative suitability of the different learners available, since “in practice it is generally impossible to know a priori which learner will perform best for a given prediction problem and data set” [58]. A poor choice in machine learning algorithms for estimation of the action mechanism may thus also result in biased inference. Recent theoretical work demonstrated that cross-validation may be used aggressively for estimator selection [59–62], i.e. to search through a large space of candidate estimators. From this work, a data-adaptive estimation approach referred to as super learning was developed [58]. Its application avoids the need for arbitrary selection of a candidate learner in a particular estimation problem. Instead, super learning combines predicted values from the various candidate learners considered through a weighted average. The selection of the optimal combination of the candidate learners is based on cross-validation to protect against over-fitting such that the resulting learner (called “super learner”) performs asymptotically as well or better than any of the candidate learners considered [58]. Super learning can thus be used for estimation of the action mechanism in this approach to avoid reliance on subjective specification of parametric models.

2.2.4 Extension of the approach over time

The analytic approach just described only makes use of data accumulated during the first 12 follow-up periods. Information in subsequent periods is thus left untapped while it could contribute to addressing the research question of interest. We thus extend the principles of the analytic approach presented earlier to make use of the complete follow-up data available from each patient to evaluate the safety of incident 1-year exposures to BP based on the 1-year cumulative risks of AF.

Informally, the extended approach over time aims to emulate inference from several ideal randomized trials. Each of these trials involves the same two one-year BP interventions described in the previous section but the start of each trial is staggered over time (Figure 2). More specifically, each trial starts at the end of one of the 30-day follow-up period of the AF/BP study and is based on the population of patients who remained unexposed to BP and active in the study (i.e. patients who did not experience a censoring event) through the end of that period. In each of these hypothetical ideal trials, the difference between the cumulative risk of AF events at 1 year in each arm is a measure of the effect of incident 1-year exposure to BP on the risk of AF. The approach described in this section aims to emulate an average of these effect measures by extending the analytic principles described earlier.

To formally describe the generalization of the analytic approach, we extend the overbar and counterfactual notation as follows. The history of a variable Z (action or covariate) between time t and t' is denoted by $\bar{Z}(t, t') = (Z(t), \dots, Z(t'))$ with $t \leq t'$. $\bar{Z}(t, t')$ is nil if $t' < t$. The counterfactual covariate at time t corresponding with an action history that is defined by the following sequence of observed actions and action interventions $(\bar{A}(j-1), \bar{a}(j, j'))$ with $0 \leq j \leq j' \leq K$, $t \leq j' + 1$ such that $\bar{A}_2(j-1) = 0$ and $\bar{a}_2(j, j') = 0$ is denoted by $L_{\bar{a}(j, j')}(t)$, i.e. $L_{\bar{a}(j, j')}(t) = L_{\bar{A}(j-1), \bar{a}(j, j')}(t)$, where $\bar{A}(-1)$ is nil. For patients whose follow-up data are not right-censored through time $t-1$ for $t = 0, \dots, K$ (i.e. for whom $\bar{A}_2(t-1) = 0$), the counterfactual time to first failure at or after t under intervention $\bar{a}(t, t+s)$ with any given $s \geq 0$ is denoted by $T_{\bar{a}(t, t+s)}^t$. If failure occurs at or after period t but before period $t+s+1$ under intervention $\bar{a}(t, t+s)$ then $t \leq T_{\bar{a}(t, t+s)}^t \leq t+s$ is defined by the time when the counterfactual outcome under intervention $\bar{a}(t, t+s)$ is 1 for the first time after t ,

Causal estimand $\psi_{0,11}$ defined by a trial that starts at period 0 and ends at period 11:

$$\left\{ \begin{array}{l} \psi_{0,11} = P(T_{\bar{a}^\bullet(0,11)}^0 \leq 11) - P(T_{\bar{a}^\circ(0,11)}^0 \leq 11) \\ D_{0,11}(O | g, \psi) = \frac{\mathcal{F}(\bar{A}(0, \tilde{T}^0))Y(\tilde{T}^0+1)}{\prod_{j=0}^{\tilde{T}^0} g(A(j) | \bar{L}(j), \bar{A}(j-1))} - \psi \end{array} \right.$$

Causal estimand $\psi_{1,12}$ defined by a trial that starts at period 1 and ends at period 12:

$$\left\{ \begin{array}{l} \psi_{1,12} = P(T_{\bar{a}^\bullet(1,12)}^1 \leq 12 | A(0)=0) - P(T_{\bar{a}^\circ(1,12)}^1 \leq 12 | A(0)=0) \\ D_{1,12}(O | g, \psi) = I(A(0)=0) \left(\frac{\mathcal{F}(\bar{A}(1, \tilde{T}^1))Y(\tilde{T}^1+1)}{\prod_{j=1}^{\tilde{T}^1} g(A(j) | \bar{L}(j), \bar{A}(j-1))} - \psi \right) \end{array} \right.$$

Causal estimand $\psi_{2,13}$ defined by a trial that starts at period 2 and ends at period 13:

$$\left\{ \begin{array}{l} \psi_{2,13} = P(T_{\bar{a}^\bullet(2,13)}^2 \leq 13 | \bar{A}(1)=0) - P(T_{\bar{a}^\circ(2,13)}^2 \leq 13 | \bar{A}(1)=0) \\ D_{2,13}(O | g, \psi) = I(\bar{A}(1)=0) \left(\frac{\mathcal{F}(\bar{A}(2, \tilde{T}^2))Y(\tilde{T}^2+1)}{\prod_{j=2}^{\tilde{T}^2} g(A(j) | \bar{L}(j), \bar{A}(j-1))} - \psi \right) \end{array} \right.$$

⋮ ⋮ ⋮

Causal estimand $\psi_{t-11,t}$ defined by a trial that starts at period $t-11$ and ends at period t :

$$\left\{ \begin{array}{l} \psi_{t-11,t} = P(T_{\bar{a}^\bullet(t-11,t)}^{t-11} \leq t | \bar{A}(t-12)=0) - \\ P(T_{\bar{a}^\circ(t-11,t)}^{t-11} \leq t | \bar{A}(t-12)=0) \\ D_{t-11,t}(O | g, \psi) = I(\bar{A}(t-12)=0) \left(\frac{\mathcal{F}(\bar{A}(t-11, \tilde{T}^{t-11}))Y(\tilde{T}^{t-11}+1)}{\prod_{j=t-11}^{\tilde{T}^{t-11}} g(A(j) | \bar{L}(j), \bar{A}(j-1))} - \psi \right) \end{array} \right.$$

Figure 2 Visual representation of the principles of the MSM approach in the BP/AF study. The approach aims to emulate inference from ideal randomized trials staggered over time. Each ellipse represents the intervention period for a given hypothetical trial. The green line represents the period during which patients from the original target population for the BP/AF study must remain untreated with BP and active in the study to form the target population of the trial. The red and blue lines represent the action interventions in the two arms of the trial (continuous 1- year exposure to BP is represented in red). The difference between the cumulative AF risks at 1- year in the two arms defines the causal effect measure of interest in each trial and is represented by the arched arrow. The formal definition of this effect using notation from the counterfactual statistical framework is given next to each diagram along with an unbiased IPAW estimating function for that effect.

i.e. $(\bar{Y}_{\bar{a}(t,t+s)}(t+1, T_{\bar{a}(t,t+s)}^t) = 0, Y_{\bar{a}(t,t+s)}(T_{\bar{a}(t,t+s)}^t + 1) = 1)$. Otherwise, if failure does not occur during the interval t to $t+s$ under intervention $\bar{a}(t, t+s)$ (i.e., $\bar{Y}_{\bar{a}(t,t+s)}(t+1, t+s+1) = 0$), the counterfactual failure time $T_{\bar{a}(t,t+s)}^t$ is defined as ∞ .

The causal estimand and its IPAW estimating function defined by formulas [1] and [2] in the previous section can now be expressed using the new notation above (with $t = 0$ and $s = 11$) as

$$\psi_{0,11} = P(T_{\bar{a}^\bullet(0,11)}^0 \leq 11) - P(T_{\bar{a}^\circ(0,11)}^0 \leq 11),$$

and

$$D_{0,11}(O | g, \psi) = \frac{\mathcal{F}(\bar{A}(0, \tilde{T}^0))Y(\tilde{T}^0+1)}{\prod_{j=0}^{\tilde{T}^0} g(A(j) | \bar{L}(j), \bar{A}(j-1))} - \psi, \quad [7]$$

where $\tilde{T}^0 = \tilde{T}$. This IPAW estimating functions can be used with observational data from the BP/AF study to emulate inference from a randomized trial that would start at time 0 and end at time 11 with two arms where patients are continuously exposed or continuously not exposed to BP for 12 30-day periods. The inference is generalizable to the target population of patients from which the patients in the BP/AF study cohort were drawn.

The principles of this conceptual ideal trial can now be extended to define another hypothetical ideal trial that would start at time $t-11$ and end at time t for any given $t = 11, \dots, K+11$. The target population for this trial may be defined as the population of patients in the original target population above with the added restriction that patients should not experience BP exposure nor a censoring event before time $t-11$ (i.e. $I(\bar{A}_i(t-12)=0) = 1$ for any patient i in this new target population). Note that while patients with documented previous exposure to BP are excluded from the target population, patients with a history of AF are included in the target population. The inference from this new trial is also based on the same two action interventions (continuous exposure to BP versus continuous absence of exposure to BP with no censoring event before and at time t denoted by $\bar{a}^*(t-11, t)$ and $\bar{a}^\circ(t-11, t)$, respectively). In this new conceptual trial, the causal estimand of interest representing the difference between the cumulative risks of failure in the two arms is defined as:

$$\psi_{t-11,t} = P(T_{\bar{a}^*}^{t-11}(t-11,t) \leq t \mid \bar{A}(t-12)=0) - P(T_{\bar{a}^\circ}^{t-11}(t-11,t) \leq t \mid \bar{A}(t-12)=0).$$

Inference from such a trial with the available observational data from the BP/AF study may be emulated under the SRA and ETA assumption with IPAW estimation based on the following IPAW estimating function:

$$D_{t-11,t}(O \mid g, \psi) = I(\bar{A}(t-12)=0) \left(\frac{\mathcal{F}(\bar{A}(t-11, \tilde{T}^{t-11}))Y(\tilde{T}^{t-11}+1)}{\prod_{j=t-11}^{\tilde{T}^{t-11}} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))} - \psi \right), \quad [8]$$

where $\tilde{T}^{t-11}$ is defined as the earliest time when a right-censoring or failure event occurs at or after time $t-11$ (start of the trial) or when the action interventions of the trial would end (i.e. time point t):

$$\tilde{T}^{t-11} = \min(t, \{t-11 \leq j \leq \min(t, C-1) : Y(j+1)=1\}, C).$$

Note that definition [8] reduces to definition [7] when $t=11$.

Instead of estimating each effect $\psi_{t-11,t}$ by solving the IPAW estimating equations associated with $D_{t-11,t}(O \mid g, \psi)$ separately for each $t = 11, \dots, K^*$ (for a given $K^* \leq K+11$), we aimed to estimate a single average effect measure ψ by solving the estimating equations associated with

$$D(O \mid g, \psi) = \sum_{t=11}^{K^*} D_{t-11,t}(O \mid g, \psi). \quad [9]$$

The solution is the IPAW estimator denoted by ψ_n of an average measure of the effects of interest $\psi_{t-11,t}$ for $t = 11, \dots, K^*$:

$$\psi_n = \sum_{i=1}^n \sum_{t=11}^{K^*} \frac{I(\bar{A}_i(t-12)=0) \mathcal{F}(\bar{A}_i(t-11, \tilde{T}_i^{t-11})) Y_i(\tilde{T}_i^{t-11}+1)}{\left[\sum_{i'=1}^n \sum_{t'=11}^{K^*} I(\bar{A}_{i'}(t'-12)=0) \right] \prod_{j=t-11}^{\tilde{T}_i^{t-11}} g(A_i(j) \mid \bar{L}_i(j), \bar{A}_i(j-1))}. \quad [10]$$

As shown in Appendix B, the average effect measure denoted by ψ that is targeted by this IPAW estimator ψ_n is given below and can be alternatively defined by a working [63] semi-parametric [64] history-restricted MSM (HRMSM) [65, 66] which, if correct, corresponds with the assumption that the risk differences $\psi_{t-11,t}$ are constant over time:

$$\psi = \sum_{t=11}^{K^*} \frac{g(\bar{A}(t-12)=0)}{\sum_{t'=11}^{K^*} g(\bar{A}(t'-12)=0)} \psi_{t-11,t}. \quad [11]$$

An estimator of the variance of the IPAW estimator ψ_n if the conditional probabilities of actions g (action mechanism) were known can be derived from its influence curve and can be expressed by the following closed-form formula:

$$\sum_{i=1}^n \left[\sum_{t=11}^{K^*} \frac{I(\bar{A}_i(t-12)=0)}{\sum_{i'=1}^n \sum_{t'=11}^{K^*} I(\bar{A}_{i'}(t'-12)=0)} \left(\frac{\mathcal{F}(\bar{A}_i(t-11, \bar{T}_i^{t-11})) Y_i(\bar{T}_i^{t-11}+1)}{\bar{T}_i^{t-11} \prod_{j=t-11} g(A_i(j) | \bar{L}_i(j), \bar{A}_i(j-1))} - \psi_n \right) \right]^2. \quad [12]$$

Under the SRA and ETA assumption, the inference for causal estimand ψ from formulas [10] and [12] is asymptotically conservative if the action mechanism is estimated consistently. To avoid reliance on a priori specified parametric models for the action mechanism, super learning may be implemented to select estimators for the different components of the action mechanism as described later. Simulations and applications in previous work [58, 67] have demonstrated the important gains achieved with the super learning approach in prediction problems. To our knowledge, this methodology has however not been implemented in any analyses that approach the scale of observation numbers ($\approx 700,000$ patients followed-up for up to 85 30-day periods) involved in the BP/AF study. Even though we expected a heavy computational burden (computing times and memory requirements) for running the R implementation of the super learner (version 1.1-18) [68] for estimating the action mechanism in the BP/AF study, the reliance on proper confounding adjustment for drawing causal inferences motivated the exploration of the applicability of super learning in a safety study of this scale. We later discuss practical obstacles to the application of super learning in this analysis and practical solutions that permitted preliminary application of its principles.

2.2.5 Implementation challenges and practical solutions

To balance computational burden and misclassification of the time-varying actions and covariates, we chose the 30-day interval as the unit of time in this analysis (see Sections 2.2.1 and 2.2.2). To approximate the observed data structure O resulting from this choice, we first needed to map raw data from the KPNC EMR collected on a continuous time scale to the coarsened time scale chosen for the analysis. The following practical approach was implemented to uphold the time-ordering assumption between measurements of the actions and covariates, i.e. $\bar{L}_{\bar{a}(K)}(t) = \bar{L}_{\bar{a}(t-1)}(t)$ which encodes that covariate $L(t)$ represents attribute measurements that occur before the action $A(t)$ or that are otherwise assumed not to be affected by the actions at time t or thereafter, $\bar{A}(t, C)$:

- Each patient's follow-up time measured in days from study entry was discretized into consecutive 30-day intervals denoted by $t = 0, \dots, C$. In particular, $t = 0$ represents the first $\tau = 30$ days of a patient's follow-up: $[0, \tau]$ and C represent the last 30-day interval when a censoring event occurs. Note that unlike t for $t = 0, \dots, C-1$ which represents a monitoring interval of 30 day $[t\tau, (t+1)\tau]$, the follow-up time $t = C$ typically represents a shorter monitoring interval since the actual censoring event often falls within the 30-day interval $[C\tau, (C+1)\tau]$.
- The indicator of right-censoring $A_2(t)$ is 0 for $t = 0, \dots, C-1$ and $A_2(C) = 1$.
- BP exposure $A_1(t)$ is set to 0 for all time intervals t except for the interval (if any) when BP is first initiated and for all subsequent intervals where BP is used for at least 15 days.
- For any attribute of the covariate $L(t)$ which does not represent a given drug exposure (e.g. diuretic or statins):
 - For all t representing intervals $[t\tau, (t+1)\tau]$ that do not contain the day when BP exposure is first initiated or first discontinued (i.e. when the value for $A_1(t)$ is 1 for the first time or 0 for the first time after a sequence of 1's, respectively), the value for each covariate attribute of $L(t)$ is set to the last measurement of that attribute (if any) 1) at $t-1$ (i.e. the last measurement obtained during interval $[(t-1)\tau, t\tau]$) if $t > 0$ and 2) if $t = 0$, within the 6 years preceding study entry if the attribute is "cancer.h" and since 1996 for all other attributes.

- For t representing the interval $[t\tau, (t+1)\tau[$ when BP exposure is first initiated or first discontinued, the value for each covariate attribute of $L(t)$ is set to the last measurement of that attribute (if any) 1) at $t - 1$ or t but always prior to the actual day when BP was initiated/discontinued if $t > 0$ and 2) if $t = 0$, during the first follow-up interval but always prior to the actual day when BP was initiated and otherwise within the 6 years preceding study entry if the attribute is “cancer.h” and since 1996 for all other attributes.
- For any attribute of the covariates $L(t)$ which represents exposure to a given drug (other than BP) at interval t , its measurement was set to 0 except if the drug was used at least 15 days of that interval. This convention relies on the assumption that BP exposure during any given interval should not affect exposure to other drugs during the same interval.

To mitigate computational burden, we implemented a data reduction step based on sampling from the original cohort data. Given the relatively rare onset of treatment with BP during follow-up in the original cohort, i.e.

$$\sum_{i=1}^n \underbrace{(1 - I(\bar{A}_{1,i}(C_i - 1) = 0))}_{\text{“indicator of BP use”}} \ll \sum_{i=1}^n \underbrace{I(\bar{A}_{1,i}(C_i - 1) = 0)}_{\text{“indicator of no BP use”}},$$

we differentially sampled a patient’s follow-up data based on whether she was ever treated with BP during follow-up (i.e. $I(\bar{A}_1(C - 1) = 0)$) to insure sufficient information on the treated patients in the resulting reduced data. This complex sampling scheme can be viewed as an artificial censoring event that is superimposed on the original data structure O collected on each patient in the BP/AF study cohort. The resulting reduced data from such complex sampling can be viewed as n i.i.d. copies of $O' = (C, I(\bar{A}_1(C - 1) = 0), \Delta, \Delta O)$, where Δ is the indicator of a patient’s data being sampled. The probability of a patient’s follow-up data being sampled conditional on her original data is known by design and only based on her “ever/never” exposure status, i.e. $g'(\Delta = 1 | O) = g'(\Delta = 1 | C, I(\bar{A}_1(C - 1) = 0))$, where g' denotes conditional probabilities of a patient’s data being sampled. We chose the following sampling probabilities by which follow-up data from all patients who ever used BP during follow-up in the original cohort ($\approx 60,000$ patients) are retained in the reduced data along with follow-up data from an approximately equal number of patients who never used BP during follow-up in the original cohort: $g'(\Delta = 1 | C, I(\bar{A}_1(C - 1) = 0) = 0) = 1$ and $g'(\Delta = 1 | C, I(\bar{A}_1(C - 1) = 0) = 1) = \frac{\sum_{i=1}^n (1 - I(\bar{A}_{1,i}(C_i - 1) = 0))}{\sum_{i=1}^n I(\bar{A}_{1,i}(C_i - 1) = 0)}$. The resulting reduced data consist of approximately 8.3 million person-time observations from about 120,000 patients monitored for up to 85 30-day periods each. As shown in Appendix C, the IPAW estimating function for ψ based on the complete data O given by formula [9] can be mapped into the following unbiased estimating function for ψ based on the reduced data O' :

$$D(O' | g, g', \psi) = \frac{\Delta D(O | g, \psi)}{g'(\Delta = 1 | C, I(\bar{A}_1(C - 1) = 0))}. \quad [13]$$

The resulting reduced-data IPAW estimator of ψ that solves the estimating equations associated with the previous estimating function $D(O' | g, \psi)$ is

$$\psi'_n = \sum_{i=1}^n \sum_{t=11}^{K^*} \frac{\Delta_i I(\bar{A}_i(t - 12) = 0) \mathcal{F}(\bar{A}_i(t - 11, \tilde{T}_i^{t-11})) Y_i(\tilde{T}_i^{t-11} + 1)}{\Lambda g'(\Delta_i = 1 | C_i, I(\bar{A}_{1,i}(C_i - 1) = 0)) \prod_{j=t-11}^{\tilde{T}_i^{t-11}} g(A_i(j) | \bar{L}_i(j), \bar{A}_i(j - 1))}, \quad [14]$$

with $\Lambda = \sum_{i=1}^n \sum_{t=11}^{K^*} \frac{\Delta_i I(\bar{A}_i(t - 12) = 0)}{g'(\Delta_i = 1 | C_i, I(\bar{A}_{1,i}(C_i - 1) = 0))}$. An estimator of the variance of this reduced-data IPAW estimator ψ_n if the action mechanism g were known can be derived from its influence curve and can be expressed by the following closed-form formula:

$$\sum_{i=1}^n \left[\sum_{t=11}^{K^*} \frac{\Delta_i I(\bar{A}_i(t-12)=0)}{g'(\Delta_i=1 | C_i, I(\bar{A}_{1,i}(C_i-1)=0))\Lambda} \left(\frac{\mathcal{F}(\bar{A}_i(t-11, \tilde{T}_i^{t-11}))Y_i(\tilde{T}_i^{t-11}+1)}{\tilde{T}_i^{t-11}} - \psi'_n \right) \right]^2. \quad [15]$$

$$\prod_{j=t-11} g(A_i(j) | \bar{L}_i(j), \bar{A}_i(j-1))$$

Akin to the estimation function for ψ , estimating functions for the different components of the action mechanism may be mapped into estimating functions based on the reduced data only. In practice, estimation of the action mechanism with the reduced data may thus be implemented using standard estimators (e.g. based on a priori specified parametric models or super learning) where the contribution of each observation is weighted by the inverse probability of sampling $\frac{\Delta}{g'(\Delta=1|C, I(\bar{A}(C-1)=0))}$ to account for sampling bias. Under the SRA and ETA assumption, the inference for causal estimand ψ from formulas [14] and [15] is conservative if the action mechanism is estimated consistently. While bootstrapping can be used instead to draw inference that incorporates variability from estimation of the action mechanism, its implementation was not practical in this study, since it would have required to iterate super learning estimation of each component of the action mechanism on each bootstrap sample.

Estimation of the action mechanism needed as a prerequisite to IPAW estimation in this study was based on further decomposition of both the treatment and right-censoring mechanisms:

$$\underbrace{g(A_1(j) | A_2(j)=0, \bar{L}(j), \bar{A}(j-1))}_{\text{treatment mechanism at time } j} = \left[(1 - I(\bar{A}_1(j-1)=0) - I(A_1(j-1)=1)) \times \right. \\ \left. g(A_1(j) | A_2(j)=0, \bar{L}(j), \bar{A}(j-1), A_1(j-1)=0, \bar{A}_1(j-1) \neq 0) + \right. \\ \left. I(\bar{A}_1(j-1)=0) \underbrace{g(A_1(j) | A_2(j)=0, \bar{L}(j), \bar{A}(j-1), \bar{A}_1(j-1)=0)}_{\text{“BP initiation mechanism” at time } j} + \right. \\ \left. I(A_1(j-1)=1) \underbrace{g(A_1(j) | A_2(j)=0, \bar{L}(j), \bar{A}(j-1), A_1(j-1)=1)}_{\text{“BP continuation mechanism” at time } j} \right],$$

$$\underbrace{g(A_2(j)=0 | \bar{L}(j), \bar{A}(j-1), \bar{A}_2(j-1)=0)}_{\text{right-censoring mechanism at time } j} = \underbrace{g(I(C=C^{\text{disen}})=0 | \bar{L}(j), \bar{A}(j-1), \bar{A}_2(j-1)=0))}_{\text{“Disenrollment mechanism” at time } j} \times \\ \underbrace{g(I(C=C^{\text{death}})=0 | \bar{L}(j), \bar{A}(j-1), \bar{A}_2(j-1)=0, I(C=C^{\text{disen}})=0)}_{\text{“Death mechanism” at time } j} \times \\ \underbrace{g(I(C=C^{\text{admin}})=0 | \bar{L}(j), \bar{A}(j-1), \bar{A}_2(j-1)=0, I(C=C^{\text{disen}})=0, I(C=C^{\text{death}})=0)}_{\text{“End of study mechanism” at time } j}.$$

Along with the three components of the right-censoring mechanism, only the BP initiation and BP continuation mechanisms need to be estimated for IPAW estimation of the causal estimand ψ . A separate estimator based on pooled data over time j was derived by super learning not only for each of these two components of the treatment mechanism but also for each of the three components of the right-censoring mechanism. Our choice of candidate learners for building these five super learners was mainly based on practical considerations, i.e. the computational burden in terms of both computing speed and memory requirements associated with each learner available in version 1.1-18 of the SuperLearner package in R. The following three types of learners were selected: main term logistic regression,² k-nearest neighbor classification [69], and polychotomous regression [70]. We refer to these estimators as glm, knn, and polyclass learners, respectively, from now on.

² Each term in the working logistic model corresponds to an explanatory variable. In particular, no interaction terms nor terms that involve a transformation of an explanatory variable (e.g., polynomial terms) are included in the working logistic model.

For constructing each of the five super learners, six candidate glm learners defined by an increasing number of explanatory variables were considered. More specifically, these glm learners involved the first 13, 20, 30, 40, 50 variables, and all variables, respectively, from an ordered list of explanatory variables. The same list of explanatory variables was considered for building the two super learners predicting BP exposure at time j , $A_1(j)$. It consists of the 30-day interval number j (named “period”) and all covariates (Tables 1 and 2) associated with interval j , $L(j)$, excluding the outcome $Y(j)$. Note that all explanatory variables are binary variables except for “age” and “period”. The same list of explanatory variables was considered for building the other three super learners and consists of the previous list of covariates but also includes BP exposure at time $j - 1$, $A_1(j - 1)$. For constructing each of the five super learners, the list of relevant explanatory variables was ordered based on not only a combination of subject-matter knowledge but also the magnitude and significance of the crude association between each relevant explanatory variable and the corresponding dependent variable as measured by the odds ratio (OR) and p -value from a univariate logistic regression accounting for repeated measurements.³ More specifically, the ordering of covariates for predicting a given dependent variable was derived based on the following approach: the 14 variables that are established risk factors for AF events (Tables 1 and 2) appear first in the list. They are ordered by decreasing magnitude of association with the dependent variable (as measured by a decreasing distance of the OR from 1). The remaining variables that are significantly associated with the dependent variable at the 0.05 level are ordered by decreasing magnitude of association with the dependent variable and appeared next in the list just before all other variables ordered by decreasing significance of association with the dependent variable (as measured by an increasing p value). Finally, one of the dummy variables for “race” (an established risk factor for AF) is removed from the ordered list due to collinearity with the other dummy variables. Note that the default screening routines available as part of version 1.1-18 of the SuperLearner package in R do not permit the screening of explanatory variables to automatically generate the six glm learners described earlier. Alternate, customized screening routines can however be easily developed and integrated to the super learner algorithm implemented in R.

Due to the heavy computational burden, the polyclass learner considered for constructing each super learner was based on only the 14 variables that are established risk factors for the outcome. We used a customized routine that implements the polyclass learner based on BIC as the model selection criterion to improve computing speed over the default `SL.polymer` routine that implements the polyclass learner based on cross-validation when the `binomial` link function is specified.

Given the inability of the default knn learner to incorporate inverse probability of sampling weights, we only considered⁴ the knn learner with $k = 10$ as a candidate learner when estimating the BP continuation mechanism since the sampling weights for all relevant observation is then 1 by design (i.e. weighting is obsolete). Due to the heavy computational burden, the candidate knn learner was only based on the 14 variables that are established risk factors for the outcome.

2.3 Results

The nomenclature for all baseline and time-varying covariates used for potential confounding and selection bias adjustment in this analysis is described in Tables 1 and 2. To illustrate the protocol implemented for ordering the lists of explanatory variables involved in the construction of each of the five super learners, we

³ For each explanatory variable, a point estimate and p -value associated with the slope of a working logistic regression of the dependent variable on the explanatory variable was obtained with generalized estimating equations using the independence structure for the matrix of variance-covariance. The exponential of the point estimates (OR) was used as a measure of the magnitude of the crude association between the two variables.

⁴ To permit completion without failure of the knn routine in the `class` R package (version 7.3-2) on which the default knn learner is based, a customized knn learner was developed and used as part of the SuperLearner package with the knn argument `use.all` set to “FALSE” and after changing the `MAX_TIES` variable in the C file “class.c” to “10,000”.

Table 3 Ranking of explanatory variables for predicting BP initiation (Part I).

Variable name	OR	p-value	Known AF risk factor
fx.hx	4.12	0	✓
afib.hx	1.73	5.21e-84	✓
race.AP	1.52	1.04e-117	✓
race.AA	0.52	4.94e-105	✓
chf.hx	1.39	6.65e-13	✓
race.AD	1.33	2.81e-28	✓
ami.usa.	1.32	2.5e-17	✓
race.OT	0.69	1.78e-77	✓
revasc.h	1.27	3.6e-09	✓
htn.hx	1.27	1.84e-88	✓
dm.hx	0.83	3.61e-26	✓
age	1.05	0	✓
race.NA	0.96	6.63e-01	✓
race.WH	1.03	2.13e-02	✓
fx	17.59	0	
calciton	9.67	2.7e-183	
glucocor	5.64	0	
aromatas	2.9	1.21e-111	
chf	2.14	3.76e-14	
stroke	2.08	4.87e-11	
arrhythmi	1.84	4.24e-22	
raloxife	1.82	5.33e-32	
dement.h	1.77	2.12e-95	
aggrenox	1.76	3.91e-12	
warfarin	1.76	9.09e-67	
valve.hx	1.73	2.04e-64	
af.outc2	1.61	7.31e-15	
thyroid.	1.61	4.18e-04	
ami.usa	1.59	1.09e-06	

provide the ordering of the explanatory variables that define the six glm learners considered for predicting BP initiation in Tables 3 and 4. The five weighted combination of candidate learners selected as the five super learners for estimating the IPAW weights are given in Tables 5–9. The average cross-validate risk associated with each learner for estimating each of the five components of the action mechanism is given in Tables 10–14. The resulting IPAW estimate for the effect ψ (defined by formula [11] with $K^* = 95$) based on the reduced data is $\psi'_n = -8.776e-04$ (derived with formula [14]) with 95% confidence interval $[-1.128e-03; -6.271e-04]$ (the standard error derived with formula [15] is $1.278e-04$ and corresponding p -value is $6.646e-12$). The 99th percentile of the inverse probability of action weights for the observations contributing to the IPAW estimate (i.e. the 99th percentile of the estimates of the weights

$\left(\frac{1}{\prod_{j=t-11}^{\bar{T}_i^t-11} g(A_i(j)|\bar{L}_i(j), \bar{A}_i(j-1))} \right)_{i: \Delta_i I(\bar{A}_i(t-12)=0) \mathcal{F}(\bar{A}_i(t-11), \bar{T}_i^t-11)) \neq 0}$ is 5.41. Truncation of the IPAW weights [71] was

implemented to evaluate the extent to which these results could be driven by possibly extreme weighting of less than 1% of the observations. Given the large imbalance between the number of observations for patients continuously treated with BP versus the number of control observations in the reduced data (Tables 15–17), we determined the observation with IPAW weights deemed extreme by identifying observations with

stabilized [72, 73] IPAW weights greater than 50. Stabilized weights were defined as $\frac{g(\bar{A}_i(t-11, \bar{T}_i^t-11)|\bar{A}_i(t-12)=0)}{\prod_{j=t-11}^{\bar{T}_i^t-11} g(A_i(j)|\bar{L}_i(j), \bar{A}_i(j-1))}$.

The numerator of the weights for the observations contributing to the IPAW estimate (i.e. observations

Table 4 Ranking of explanatory variables for predicting BP initiation (Part II).

Variable name	OR	p-value	Known AF risk factor
clopidog	1.53	6.59e-20	
stroke.h	1.51	6.3e-22	
plavix	1.5	1.91e-17	
cancer.h	1.47	6.9e-69	
pad.hx	1.47	1.2e-13	
nitrate	1.46	2.79e-18	
aldoster	1.43	1.4e-08	
statin	1.4	1.9e-169	
digoxin	1.4	1.22e-17	
hrt	0.61	1.98e-202	
c.blocke	1.38	3.12e-83	
hyperthy	1.35	4.11e-38	
liver.hx	1.27	1.78e-08	
lung.hx	1.27	8.88e-65	
dyslipid	1.27	7.89e-90	
b.blocke	1.25	1.83e-65	
arb	1.21	5.09e-15	
hypothy.	1.2	1.09e-33	
ace	1.15	1.22e-22	
alpha	1.14	6.01e-04	
depress.	1.14	1.95e-20	
diuretic	1.11	1.32e-17	
med.inco	0.9	5.69e-08	
hs.or.le	0.91	3.51e-13	
period	1.01	0	
non.lipi	1.06	1.73e-01	
revasc	0.82	2.29e-01	
ticlopid	1.37	2.42e-01	
hispan	0.98	3.28e-01	

Table 5 Weights applied to each candidate learner to form the Super Learner for predicting BP initiation based on 6,853,387 person-time observations.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0	0	0	0	0	0.63	0.37

Table 6 Weights applied to each learner to form the Super Learner for predicting BP continuation based on 1,411,585 person-time observations.

glm13	glm20	glm30	glm40	glm50	glmAll	knn	poly
0	0	0	0	0.038	0.89	0.0024	0.065

Table 7 Weights applied to each learner to form the Super Learner for predicting right-censoring due to disenrollment from the health plan based on 8,264,972 person-time observations.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0	0	0	0	0	0.95	0.053

Table 8 Weights applied to each learner to form the Super Learner for predicting right-censoring due to death based on 8,240,633 person-time observations.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0	0.25	0	0	0	0.32	0.43

Table 9 Weights applied to each learner to form the Super Learner for predicting right-censoring due to administrative end of study based on 8,230,239 person-time observations.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0	0	0	0	0	1	0

Table 10 Five fold cross-validated risks associated with the candidate learners considered for predicting BP initiation.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0.01452	0.01451	0.01451	0.0145	0.0145	0.0145	0.01451

Table 11 Five fold cross-validated risks associated with the candidate learners considered for predicting BP continuation.

glm13	glm20	glm30	glm40	glm50	glmAll	knn	poly
0.0489	0.04885	0.04876	0.04873	0.04869	0.04869	0.05387	0.04889

Table 12 Five fold cross-validated risks associated with the candidate learners considered for predicting right-censoring due to disenrollment from the health plan.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0.01983	0.01982	0.01982	0.01982	0.01982	0.0198	0.01983

Table 13 Five fold cross-validated risks associated with the candidate learners considered for predicting right-censoring due to death.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0.006088	0.006082	0.006112	0.006112	0.006111	0.006116	0.006084

Table 14 Five fold cross-validated risks associated with the candidate learners considered for predicting right-censoring due to administrative and of study.

glm13	glm20	glm30	glm40	glm50	glmAll	poly13
0.05461	0.05455	0.05291	0.05234	0.05218	0.04311	0.0546

Table 15 Part I – Distribution of observations in the reduced data informing the risk estimates in each arm of the staggered hypothetical randomized trials emulated by the MSM approach of the BP/AF study.

Period	# eligible	# treated	# control
11	124,904	393	106,772
12	123,846	402	104,751
13	122,308	396	103,062
14	120,987	418	101,256
15	119,528	456	99,446
16	118,222	456	97,734
17	116,763	588	95,952
18	115,036	683	94,252
19	113,192	666	92,637
20	111,434	592	91,007
21	109,839	619	89,250
22	108,167	499	87,586
23	106,755	482	86,031
24	104,734	492	84,153
25	103,045	561	82,648
26	101,234	548	81,170
27	99,422	516	79,774
28	97,709	519	78,311
29	95,925	492	76,903
30	94,227	512	75,538
31	92,614	519	74,180
32	90,987	543	72,837
33	89,223	548	71,584
34	87,557	440	70,307
35	86,000	356	69,079
36	84,124	470	67,590
37	82,621	471	66,444
38	81,146	466	65,308
39	79,750	472	64,042

Notes: For any given period t : # eligible = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0)$, # treated = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^*(t-11, \tilde{T}_i^{t-11}))$, and # control = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^o(t-11, \tilde{T}_i^{t-11}))$.

i such that $\Delta_i I(\bar{A}_i(t-12) = 0) \mathcal{F}(\bar{A}_i(t-11, \tilde{T}_i^{t-11})) \neq 0$ were estimated separately for each $t = 11, \dots, K^*$ as follows to account for the complex sampling that lead to the reduced data:

$$g_n(\bar{a}_i^*(t-11, \tilde{T}_i^{t-11}) \mid \bar{A}_i(t-12) = 0) = \frac{\sum_{i=1}^n \frac{\Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^*(t-11, \tilde{T}_i^{t-11}))}{g'(\Delta_i = 1 \mid O)}}{\sum_{i=1}^n \frac{\Delta_i I(\bar{A}_i(t-12) = 0)}{g'(\Delta_i = 1 \mid O)}},$$

for $\bar{a}_i^*(t-11, \tilde{T}_i^{t-11}) \in \{\bar{a}_i^*(t-11, \tilde{T}_i^{t-11}), \bar{a}_i^o(t-11, \tilde{T}_i^{t-11})\}$. The unstabilized weights greater than 50 for observations contributing to the IPAW estimate were truncated at for $50/g_n(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) \mid \bar{A}_i(t-12) = 0)$. Of the 4,342,902 observations contributing to the IPAW estimate, only 898 observations had their weights truncated based on the truncation scheme above: none are from patients who experience failure (i.e., $Y(\tilde{T}_i^{t-11} + 1) = 0$), all are observations associated with time points t greater than 66, and 896 are from patients with continuous BP exposure (i.e., $\bar{a}_i^*(t-11, \tilde{T}_i^{t-11})$). The truncated IPAW estimate resulting from the weight truncation scheme above are thus identical to the untruncated IPAW estimate above. Results also remained unchanged when using a truncation level of 20 instead of 50. In addition, an analysis which is restricted to observations collected

Table 16 Part II – Distribution of observations in the reduced data informing the risk estimates in each arm of the staggered hypothetical randomized trials emulated by the MSM approach of the BP/AF study.

Period	# eligible	# treated	# control
40	78,286	439	62,868
41	76,880	397	61,813
42	75,515	436	60,678
43	74,160	438	59,616
44	72,814	405	58,601
45	71,565	362	57,521
46	70,288	379	56,385
47	69,060	269	55,420
48	67,567	316	54,145
49	66,418	323	53,117
50	65,281	372	52,083
51	64,018	321	51,046
52	62,843	303	50,055
53	61,788	298	49,163
54	60,654	266	48,231
55	59,592	262	47,381
56	58,577	296	46,535
57	57,497	341	45,739
58	56,363	253	44,936
59	55,398	189	44,098
60	54,128	248	42,784
61	53,100	286	41,900
62	52,067	228	41,020
63	51,033	262	40,143
64	50,038	228	39,356
65	49,145	213	38,583
66	48,213	212	37,824
67	47,364	227	37,038
68	46,520	211	36,258

Notes: For any given period t : # eligible = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0)$, # treated = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^*(t-11, \tilde{T}_i^{t-11}))$, and # control = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^*(t-11, \tilde{T}_i^{t-11}))$.

before period 67 (i.e. an analysis targeting the effect measure ψ defined by formula [11] with $K^* = 66$) lead to the following result: $\psi'_n = -7.716e-04$ $[-1.054e-03; -4.891e-04]$. The combination of results above did not raise concern over violation of the ETA assumption and thus identifiability of the effect measure ψ of interest.

3 Discussion

In this report, we described the motivation for, principles of, and implementation of a history restricted marginal structural modeling approach based on IPAW estimation and super learning to evaluate the safety of 1-year incident exposure to BP with respect to 1-year cumulative risks of hospitalization for AF using longitudinal observational EMR data from a large cohort of patients.

Following data structuring with SAS software⁵, all computations were implemented in R [74] version 2.13.0 on a computer server operating under RedHat RHEL 6.1 with four Physical CPU Sockets (Intel (R) Xeon

⁵ Version 9.1.3 of the SAS System for Unix. Copyright © 2002-2003 SAS Institute Inc. SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc., Cary, NC, USA.

Table 17 Part III – Distribution of observations in the reduced data informing the risk estimates in each arm of the staggered hypothetical randomized trials emulated by the MSM approach of the BP/AF study.

Period	# eligible	# treated	# control
69	45,723	192	35,475
70	44,921	229	34,699
71	44,087	171	33,975
72	42,772	214	33,367
73	41,892	196	32,057
74	41,010	203	31,268
75	40,133	179	30,477
76	39,346	189	29,704
77	38,570	187	28,970
78	37,807	167	28,199
79	37,018	201	27,497
80	36,237	207	26,775
81	35,454	191	26,023
82	34,676	180	25,328
83	33,955	152	24,612
84	33,351	144	23,903
85	32,041	0	47
86	31,255	1	43
87	30,465	0	34
88	29,692	0	32
89	28,959	0	26
90	28,187	0	23
91	27,486	0	21
92	26,764	0	20
93	26,011	0	17
94	25,316	0	9
95	24,601	0	2
96	23,892	0	0

Notes: For any given period t : # eligible = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0)$, # treated = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^*(t-11, \tilde{T}_i^{t-11}))$, and # control = $\sum_{i=1}^n \Delta_i I(\bar{A}_i(t-12) = 0) I(\bar{A}_i(t-11, \tilde{T}_i^{t-11}) = \bar{a}_i^c(t-11, \tilde{T}_i^{t-11}))$.

(R) CPU X7560 @ 2.27 GHz with 16 Cores) and 128 GB physical memory. The cumulative time elapsed to sequentially derive the five super learners needed for IPAW estimation amounts to about 8 days (≈ 196 hours) of continuous computing time with version 1.1-18 of the SuperLearner package [68]. In addition, up to 118 GB of physical memory was required for deriving these five super learners. While the computing time for super learning is long relative to standard analytic approaches, it remains reasonable due to the choice of candidate learners considered in the analysis. To illustrate the unpracticality of the current implementation of super learning with a different selection of candidate learners for this analysis, we report our experience of approximately 35 days of computing time for super learning aimed at predicting BP initiation when the list of candidate learners was expanded to include the knn learner based on the 14 established risk factors for AF. Through judicious selection of candidate learners, one can thus substantially decrease the computing time for super learning at the cost of a more limited estimator search and thus possible residual confounding or selection bias. To our knowledge, this study provides the first practical evaluation of super learning with such a large dataset. The amount of data involved in the BP/AF study is expected to become a common feature of many future Comparative Effectiveness Research (CER) studies based on EMR data from healthcare databases. This experience highlights the important computation burden associated with applications of the current implementation of super learning in such large CER studies. It motivates alternate

implementations of super learning that aim to improve not only the computing speed but also the ability to apply this machine learning approach in more typical computing environments with limited physical memory.

Unlike typical IPAW estimation approaches [29, 30, 75] to adjust for time-dependent confounding and selection bias with survival outcomes, the IPAW estimation approach in this report does not rely on the specification of a parametric MSM for the hazards of failure over time. IPAW estimates based on such models are derived using data from all person-time observations with action histories consistent with at least one of the action interventions of interest. Only a subset of such person-time observations are used to estimate the cumulative risks of interest in the IPAW approach described in this report, e.g. person-time observations from patients who are censored before the end of each hypothetical trials do not contribute to the effect estimate regardless of whether or not these patients were following continuously (or partially) one of the treatment strategies of interest before the right censoring event. The IPAW estimation approach in this report is thus expected to be less efficient in general compared to the IPAW estimation approaches that have been applied most frequently in survival analysis.

In addition, the IPAW estimator implemented in this study does not have the “boundedness property” [48] because the resulting risk difference estimate is not guaranteed to be in the $[-1, 1]$ interval. For improved efficiency, an alternate IPAW estimator with the boundedness property could be derived based on separate “bounded Horvitz-Thompson” estimators [48] of the two weighted average risks of interest. Implementation of such a bounded IPAW estimator would be particularly desirable in problems with smaller sample sizes or in which violation of the positivity assumption is a concern.

The results of this analysis do not raise concern for the safety of 1-year exposure to BP and, instead, suggest a very small protective effect of continuous 1-year exposure to BP with respect to 1-year cumulative risks of hospitalization for AF. Given our premise that exposure to BP does not protect against AF events, we are concerned with residual biases due to violation of the assumption of no unmeasured confounders and/or the assumption of consistent estimation of the action mechanism. Super learning was specifically adopted in this analysis to avoid biases caused by violation of the second assumption due to arbitrary reliance on incorrect parametric models for the different components of the action mechanism [47]. For pragmatic reasons, however, we only considered a restricted set of candidate learners in this analysis. Newer versions [68] of the SuperLearner package in R now include alternate algorithms that make use of parallel computations for implementing super learning. It will be of interest to evaluate the computational burden of these newer algorithms in future analytic efforts with datasets of comparable dimensions to investigate whether a more aggressive estimator search with super learning based on a richer set of candidate learners is feasible. For further protection against bias due to model misspecification, double robust estimators [24, 44, 76] have been developed to provide two chances for the analyst to obtain correct inferences in observational studies. In addition, these estimators can be more efficient than IPAW estimators. While double robust estimators have been applied in point treatment problems [77], applications in problems with time-varying treatments have been lagging [41, 78]. This may largely be explained by the complexity in implementing the proposed algorithms in problems with long follow-up and high-dimensional, time-dependent covariates such as the one tackled in this report. Recent algorithmic results for implementation of a double robust targeted minimum loss-based estimation (TMLE) approach [79] may mitigate such pragmatic challenges. It will be of interest to evaluate the feasibility of this TMLE approach for estimating the types of effect measures evaluated in this study.

Appendix

A Proof 1

We first show that the following estimating function $D^*(O | \gamma, g)$ is unbiased for estimating $\gamma^* = P(T_{\bar{a}^*(11)} \leq 11)$ defined for any given action intervention $\bar{a}^*(11)$ such that $\bar{a}_2^*(11) = 0$:

$$D^*(O \mid \gamma, g) = \frac{I(\bar{A}(\check{T}) = \bar{a}^*(\check{T}))}{\prod_{j=0}^{\check{T}} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))} Y(\check{T} + 1) - \gamma.$$

The consistency assumption, i.e. for each $t \leq C$ we have $\bar{L}(t) = \bar{L}_{\bar{a}(K)}(t)$ for all $\bar{a}(K) \in \mathcal{A}(K)$ such that $\bar{a}(t) = \bar{A}(t)$, and the time-ordering assumption, i.e. $L_{\bar{a}(K)}(t) = L_{\bar{a}(t-1)}(t)$, imply equality $I(\bar{A}(\check{T}) = \bar{a}^*(\check{T}))Y(\check{T} + 1) = I(\bar{A}(\check{T}) = \bar{a}^*(\check{T}))I(T_{\bar{a}^*(11)} \leq 11)$. From this equality and equality $I(\bar{A}(\check{T}) = \bar{a}^*(\check{T})) = I(\bar{A}_1(\check{T}) = \bar{a}_1^*(\check{T}))I(\bar{A}_2(\check{T}) = 0)$, we have:

$$D^*(O \mid \gamma, g) = \frac{I(\bar{A}_1(\check{T}) = \bar{a}_1^*(\check{T}))I(\bar{A}_2(\check{T}) = 0)}{\prod_{j=0}^{\check{T}} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))} I(T_{\bar{a}^*(11)} \leq 11) - \gamma.$$

The SRA implies the equality $g(\bar{A}(\check{T}) \mid X) = \prod_{j=0}^{\check{T}} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))$ and we thus have:

$$D^*(O \mid \gamma, g) = \frac{I(\bar{A}_1(\check{T}) = \bar{a}_1^*(\check{T}))I(\bar{A}_2(\check{T}) = 0)}{g(\bar{A}_1(\check{T}), \bar{A}_2(\check{T}) = 0 \mid X)} I(T_{\bar{a}^*(11)} \leq 11) - \gamma.$$

Given the equality $I(\bar{A}_2(\check{T}) = 0) = I(C > \check{T})$ and given that the value for $\check{T}$ is completely determined by both the censoring time C and the full data X , we have

$$E(D^*(O \mid \gamma, g) \mid C, X) = \sum_{\bar{a}_1(\check{T}): g(\bar{a}_1(\check{T}) \mid C, X) > 0} \frac{I(\bar{a}_1(\check{T}) = \bar{a}_1^*(\check{T}))I(C > \check{T})}{g(\bar{a}_1(\check{T}), \bar{A}_2(\check{T}) = 0 \mid X)} I(T_{\bar{a}^*(11)} \leq 11) g(\bar{a}_1(\check{T}) \mid C, X) - \gamma.$$

Under the SRA and positivity assumption,⁶ this equality becomes

$$E(D^*(O \mid \gamma, g) \mid C, X) = \frac{I(C > \check{T})}{g(\bar{a}_1^*(\check{T}), \bar{A}_2(\check{T}) = 0 \mid X)} I(T_{\bar{a}^*(11)} \leq 11) g(\bar{a}_1^*(\check{T}) \mid C, X) - \gamma.$$

By evaluating the conditional expectation of each side of the equality given the full data X , we have:

$$E(D^*(O \mid \gamma, g) \mid X) = \sum_{c: g(c \mid X) > 0} \frac{I(c > \check{T})}{g(\bar{a}_1^*(\check{T}), \bar{A}_2(\check{T}) = 0 \mid X)} I(T_{\bar{a}^*(11)} \leq 11) g(\bar{a}_1^*(\check{T}), c \mid X) - \gamma.$$

Under the SRA and positivity assumption⁶ and since $\check{T}$ is only a function of X when $C > \check{T}$, this equality becomes

$$\begin{aligned} E(D^*(O \mid \gamma, g) \mid X) &= \frac{I(T_{\bar{a}^*(11)} \leq 11)}{g(\bar{a}_1^*(\check{T}), \bar{A}_2(\check{T}) = 0 \mid X)} \sum_{c: c > \check{T}} g(\bar{a}_1^*(\check{T}), c \mid X) - \gamma \\ &= I(T_{\bar{a}^*(11)} \leq 11) - \gamma \end{aligned}$$

and we thus have

$$E(D^*(O \mid \gamma^*, g)) = 0. \quad [16]$$

Using the notation introduced above, the causal estimand $\psi_{0,11}$ defined by formula (1) can be expressed as $\psi_{0,11} = \gamma^\bullet - \gamma^\circ$ and we have

⁶ The following weaker positivity assumption is sufficient: $g(A(t) \mid \bar{L}(t), \bar{A}(t-1) = \bar{a}(t-1)) > 0$ for $t = 0, \dots, 11$ and $\bar{a}(11) \in \{\bar{a}^\circ(11), \bar{a}^\bullet(11)\}$.

$$\begin{aligned} E(D_{0,11}(O \mid g, \psi_{0,11})) &= E(D^\bullet(O \mid g, \gamma^\bullet) - D^\circ(O \mid g, \gamma^\circ)) \\ &= 0 \end{aligned}$$

where the second equality follows from result [16] applied separately to the estimation of the cumulative risks $\gamma^\bullet$ and γ° .

B Proof 2

We first show that the following estimating function $D^*(O \mid \gamma, g)$ is unbiased for estimating

$\gamma^* = \sum_{t=11}^{K^*} \frac{g(\bar{A}(t-12)=0)}{\sum_{t'=11}^{K^*} g(\bar{A}(t'-12)=0)} P(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t \mid \bar{A}(t-12)=0)$ defined for any given action intervention $\bar{a}^*(t-11, t)$ such that $\bar{a}_2^*(t-11, t) = 0$:

$$D^*(O \mid \gamma, g) = \sum_{t=11}^{K^*} I(\bar{A}(t-12)=0) \left(\frac{I(\bar{A}(t-11, \tilde{T}^{t-11}) = \bar{a}^*(t-11, \tilde{T}^{t-11}))}{\tilde{T}^{t-11}} Y(\tilde{T}^{t-11} + 1) - \gamma \right) \prod_{j=t-11}^{t-1} g(A(j) \mid \bar{L}(j), \bar{A}(j-1)).$$

From equality $I(\bar{A}(t-11, \tilde{T}^{t-11}) = \bar{a}^*(t-11, \tilde{T}^{t-11})) Y(\tilde{T}^{t-11} + 1) = I(\bar{A}(t-11, \tilde{T}^{t-11}) = \bar{a}^*(t-11, \tilde{T}^{t-11})) I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t)$ implied by the consistency and time ordering assumptions and the equality $I(\bar{A}(t-11, \tilde{T}^{t-11}) = \bar{a}^*(t-11, \tilde{T}^{t-11})) = I(\bar{A}_1(t-11, \tilde{T}^{t-11}) = \bar{a}_1^*(t-11, \tilde{T}^{t-11})) I(\bar{A}_2(t-11, \tilde{T}^{t-11}) = 0)$, we have:

$$\begin{aligned} D^*(O \mid \gamma, g) &= \sum_{t=11}^{K^*} I(\bar{A}(t-12)=0) \left(\frac{I(\bar{A}_1(t-11, \tilde{T}^{t-11}) = \bar{a}_1^*(t-11, \tilde{T}^{t-11})) I(\bar{A}_2(t-11, \tilde{T}^{t-11}) = 0)}{\tilde{T}^{t-11}} \times \right. \\ &\quad \left. \prod_{j=t-11}^{t-1} g(A(j) \mid \bar{L}(j), \bar{A}(j-1)) I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t) - \gamma \right). \end{aligned}$$

We define the so-called t -specific full data [66] as $X^{t-11} = (\bar{A}(t-12), \bar{L}(t-11), \bar{L}_{\bar{a}(t-11,t)}(t-10, t+1))_{\bar{a}(t-11,t): \bar{a}(K) \in \mathcal{A}(K)}$ for $t = 11, \dots, K^*$. The SRA implies the equality $g(\bar{A}(t-11, \tilde{T}^{t-11}) \mid X^{t-11}) = \prod_{j=t-11}^{t-1} g(A(j) \mid \bar{L}(j), \bar{A}(j-1))$ and we thus have $D^*(O \mid \gamma, g) = \sum_{t=11}^{K^*} D^{*,t}(O \mid \gamma, g)$ with:

$$\begin{aligned} D^{*,t}(O \mid \gamma, g) &= I(\bar{A}(t-12)=0) \left(\frac{I(\bar{A}_1(t-11, \tilde{T}^{t-11}) = \bar{a}_1^*(t-11, \tilde{T}^{t-11})) I(\bar{A}_2(t-11, \tilde{T}^{t-11}) = 0)}{g(\bar{A}_1(t-11, \tilde{T}^{t-11}), \bar{A}_2(t-11, \tilde{T}^{t-11}) = 0 \mid X^{t-11})} \times \right. \\ &\quad \left. I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t) - \gamma \right). \end{aligned}$$

Given the equality $I(\bar{A}_2(\tilde{T}^{t-11}) = 0) = I(C > \tilde{T}^{t-11})$ and given that the value for $\tilde{T}^{t-11}$ is completely determined by both the censoring time C and the full data X^{t-11} , we have

$$\begin{aligned} E(D^{*,t}(O \mid \gamma, g) \mid C, X^{t-11}) &= \sum_{\bar{a}_1(t-11, \tilde{T}^{t-11}): g(\bar{a}_1(t-11, \tilde{T}^{t-11}) \mid C, X^{t-11}) > 0} \left(I(\bar{A}(t-12)=0) I(\bar{a}_1(t-11, \tilde{T}^{t-11}) = \bar{a}_1^*(t-11, \tilde{T}^{t-11})) \times \right. \\ &\quad \left. \frac{I(C > \tilde{T}^{t-11}) I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t)}{g(\bar{a}_1(t-11, \tilde{T}^{t-11}), \bar{A}_2(t-11, \tilde{T}^{t-11}) = 0 \mid X^{t-11})} g(\bar{a}_1(t-11, \tilde{T}^{t-11}) \mid C, X^{t-11}) \right) \\ &\quad - \gamma I(\bar{A}(t-12)=0). \end{aligned}$$

Under the SRA and positivity assumption,⁷ this equality becomes

$$E(D^{*,t}(O \mid \gamma, g) \mid C, X^{t-11}) = \left(\frac{I(\bar{A}(t-12)=0)I(C > \tilde{T}^{t-11})I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t)}{g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}), \bar{A}_2(t-11, \tilde{T}^{t-11})=0 \mid X^{t-11})} \times \right. \\ \left. g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}) \mid C, X^{t-11}) \right) - \gamma I(\bar{A}(t-12)=0)$$

By evaluating the conditional expectation of each side of the equality given the t -specific full data X^{t-11} , we have:

$$E(D^{*,t}(O \mid \gamma, g) \mid X^{t-11}) = \sum_{c: g(c \mid X^{t-11}) > 0} \left(\frac{I(\bar{A}(t-12)=0)I(c > \tilde{T}^{t-11})I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t)}{g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}), \bar{A}_2(t-11, \tilde{T}^{t-11})=0 \mid X^{t-11})} \times \right. \\ \left. g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}), c \mid X^{t-11}) \right) - \gamma I(\bar{A}(t-12)=0)$$

Under the SRA and positivity assumption* and since $\tilde{T}^{t-11}$ is only a function of X^{t-11} when $C > \tilde{T}^{t-11}$, this equality becomes

$$E(D^{*,t}(O \mid \gamma, g) \mid X^{t-11}) = \left(\frac{I(\bar{A}(t-12)=0)I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t)}{g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}), \bar{A}_2(t-11, \tilde{T}^{t-11})=0 \mid X^{t-11})} \times \right. \\ \left. \sum_{c: c > \tilde{T}^{t-11}} g(\bar{a}_1^*(t-11, \tilde{T}^{t-11}), c \mid X^{t-11}) \right) - \gamma I(\bar{A}(t-12)=0). \\ = I(\bar{A}(t-12)=0)I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t) - \gamma I(\bar{A}(t-12)=0).$$

and we thus have

$$\begin{aligned} E(D^*(O \mid \gamma^*, g)) &= \sum_{t=11}^{K^*} E \left(I(\bar{A}(t-12)=0)I(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t) - \gamma^* I(\bar{A}(t-12)=0) \right) \\ &= \sum_{t=11}^{K^*} g(\bar{A}(t-12)=0)P(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t \mid \bar{A}(t-12)=0) - \gamma^* \sum_{t=11}^{K^*} g(\bar{A}(t-12)=0) \\ &= 0. \end{aligned} \quad [17]$$

Using the notation introduced above, the causal estimand ψ defined by formula (11) can be expressed as $\psi = \gamma^\bullet - \gamma^\circ$ and we have

$$\begin{aligned} E(D(O \mid g, \psi)) &= E(D^\bullet(O \mid g, \gamma^\bullet) - D^\circ(O \mid g, \gamma^\circ)) \\ &= 0, \end{aligned}$$

where the second equality follows from result [17] applied separately to the estimation of the cumulative risks $\gamma^\bullet$ and γ° . Thus, the estimator ψ_n that is solution of the estimating equations associated with the estimating function defined by formula [9] is a consistent and asymptotically linear estimator of the target parameter ψ which can also be defined as:

$$\psi = \operatorname{argmin}_{\mu} \sum_{t=11}^{K^*} g(\bar{A}(t-12)=0) \left(P(T_{\bar{a}^\bullet(t-11,t)}^t \leq t \mid \bar{A}(t-12)=0) - P(T_{\bar{a}^\circ(t-11,t)}^{t-11} \leq t \mid \bar{A}(t-12)=0) - \mu \right)^2,$$

With this representation, the target parameter ψ can be viewed as defined by the following pooled (over time) working semi-parametric history-restricted MSM for $P(T_{\bar{a}^\bullet(t-11,t)}^t \leq t \mid \bar{A}(t-12)=0)$:

⁷ The following weaker positivity assumption is sufficient: for $t = 11, \dots, K^*$ $g(A(j) = a(j) \mid \bar{L}(j), \bar{A}(t-12)=0, \bar{A}(t-11, j-1) = \bar{a}(t-11, j-1)) > 0$ for $j = t-11, \dots, t$ and $\bar{a}(t-11, t) \in \{\bar{a}^\circ(t-11, t), \bar{a}^\bullet(t-11, t)\}$

$P(T_{\bar{a}^*(t-11,t)}^{t-11} \leq t \mid \bar{A}(t-12) = 0) + \psi$. Note that the assumption that this working HRMSM is correct corresponds with the assumption $\psi = \psi_{t-11,t}$ for $t = 11, \dots, K^*$ (i.e., constant risk differences $\psi_{t-11,t}$ over time).

C Proof 3

We show that the reduced-data estimating function defined by formula [13] is unbiased for estimating the target parameter ψ defined by formula [11]:

$$\begin{aligned} E(D(O' \mid g, g', \psi)) &= E\left(E\left[\frac{\Delta D(O \mid g, \psi)}{g'(\Delta = 1 \mid C, I(\bar{A}_1(C-1) = 0))} \mid O\right]\right) \\ &= E\left(\sum_{\delta: g'(\delta \mid O) > 0} \frac{\delta D(O \mid g, \psi)}{g'(\Delta = 1 \mid C, I(\bar{A}_1(C-1) = 0))} g'(\delta \mid O)\right) \\ &= E\left(\frac{D(O \mid g, \psi)}{g'(\Delta = 1 \mid C, I(\bar{A}_1(C-1) = 0))} g'(\Delta = 1 \mid O)\right) \\ &\quad \text{because } g'(\Delta = 1 \mid O) > 0 \text{ by design of the data} \\ &\quad \text{reduction scheme adopted (section 2.2.5)} \\ &= E(D(O \mid g, \psi)) \text{ because} \\ &\quad g'(\Delta = 1 \mid O) = g'(\Delta = 1 \mid C, I(\bar{A}_1(C-1) = 0)) \text{ by} \\ &\quad \text{design of the data reduction scheme adopted} \\ &= 0 \text{ as shown in Appendix B.} \end{aligned}$$

Acknowledgments: This study was supported by a research contract (HHSF223200510008C) from the U.S. Food and Drug Administration. The contents of this publication are solely the responsibility of the authors and do not represent the official views of the Food and Drug Administration. The authors thank Eric C. Polley and Mark J. van der Laan for helpful discussions.

References

1. Go AS. The epidemiology of atrial fibrillation in elderly persons: the tip of the iceberg. *Am j Geriatr Cardiol* 2005;14:56–61.
2. Rejnmark L, Vestergaard P, Mosekilde L. Fracture risk in patients treated with amiodarone or digoxin for cardiac arrhythmias: a nationwide case-control study. *Osteoporos Int* 2007;18:409–17.
3. Go AS, Hylek EM, Phillips KA, et al. Prevalence of diagnosed atrial fibrillation in adults: national implications for rhythm management and stroke prevention: the Anticoagulation and Risk Factors in Atrial Fibrillation (ATRIA) Study. *JAMA* 2001;285:2370–25.
4. Black DM, Delmas PD, Eastell R, et al. Once-yearly zoledronic acid for treatment of postmenopausal osteoporosis. *New Engl J Med* 2007;356:1809–22.
5. Cummings SR, Schwartz AV, Black DM. Alendronate and atrial fibrillation. *New Engl J Med* 2007;356:1895–6.
6. Vestergaard P, Schwartz K, Pinholt EM, Rejnmark L, Mosekilde L. Risk of atrial fibrillation associated with use of bisphosphonates and other drugs against osteoporosis: a cohort study. *Calcif Tissue Int* 2010;86:335–42.
7. Abrahamsen B, Eiken P, Brixen K. Atrial fibrillation in fracture patients treated with oral bisphosphonates. *J Intern Med* 2009;265:581–92.
8. Bunch TJ, Anderson JL, May HT, et al. Relation of bisphosphonate therapies and risk of developing atrial fibrillation. *Am J Cardiol* 2009;103:824–8.
9. Grosso A, Douglas I, Hingorani A, MacAllister R, Smeeth L. Oral bisphosphonates and risk of atrial fibrillation and flutter in women: a self-controlled case-series safety analysis. *PLoS ONE*. 2009;4:e4720.
10. Heckbert SR, Li G, Cummings SR, Smith NL, Psaty BM. Use of alendronate and risk of incident atrial fibrillation in women. *Arch Intern Med* 2008;168:826–831.
11. Huang WF, Tsai YW, Wen YW, Hsiao FY, Kuo KN, Tsai CR. Osteoporosis treatment and atrial fibrillation: alendronate versus raloxifene. *Menopause* 2010;17:57–63.

12. Sorensen HT, Christensen S, Mehner F, et al. Use of bisphosphonates among women and risk of atrial fibrillation and flutter: population based case-control study. *BMJ* 2008;336:813–6.
13. Krieger N. Overcoming the absence of socioeconomic data in medical records: validation and application of a census-based methodology. *Am j Public Health* 1992;82:703–10.
14. Gordon NP. Characteristics of Adult Health Plan Members in Kaiser Permanente's Northern California Region, as Estimated from the 2008 Member Health Survey. Oakland CA: Division of Research, Kaiser Permanente Medical Care Program, 2008.
15. Lo JC, Pressman AR, Omar MA, Ettinger B. Persistence with weekly alendronate therapy among postmenopausal women. *Osteoporos Int* 2006;17:922–8.
16. Go A S, Lee WY, Yang J, Lo JC, Gurwitz JH. Statin therapy and risks for death and hospitalization in chronic heart failure. *JAMA* 2006;296:2105–11.
17. Go AS, Yang J, Ackerson LM, et al. Hemoglobin level, chronic kidney disease, and the risks of death and hospitalization in adults with chronic heart failure: the Anemia in Chronic Heart Failure: outcomes and Resource Utilization (ANCHOR) Study. *Circulation* 2006;113:2713–23.
18. Go AS, Hylek EM, Chang Y, et al. Anticoagulation therapy for stroke prevention in atrial fibrillation: how well do randomized trials translate into clinical practice? *JAMA* 2003;290:2685–92.
19. Newman TB, Brown AN. Use of commercial record linkage software and vital statistics to identify patient deaths. *J Am Med Inform Assoc* 1997;4:233–7.
20. Arellano MG, Petersen GR, Petitti DB, Smith RE. The California Automated Mortality Linkage System (CAMLIS). *Am j Public Health* 1984;74:1324–30.
21. Go AS, Chertow G M, Fan D, McCulloch CE, Hsu CY. Chronic kidney disease and the risks of death, cardiovascular events, and hospitalization. *N Engl J Med* 2004;351:1296–305.
22. Robins JM. (1998) Marginal structural models. In 1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science, pp. 1–10. Reproduced courtesy of the American Statistical Association. <http://www.biostat.harvard.edu/robins/msm-web.pdf>
23. Hernan MA. With great data comes great responsibility: publishing comparative effectiveness research in epidemiology. *Epidemiology* 2011;22:290–1.
24. van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. New York: Springer, 2003.
25. Hernan MA, Hernandez-Diaz S, Werler MM, Mitchell AA. Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology *Am J Epidemiol* 2002;155:176–84.
26. Rosenbaum PR. The consequence of adjustment for a concomitant variable that has been affected by the treatment. *J Roy Stat Soc A, General* 1984;147:656–66.
27. Robins JM. A new approach to causal inference in mortality studies with a sustained exposure period; application to the healthy worker survivor effect. *Math Model* 1986;7:1393–1512.
28. Robins JM. Association, causation and marginal structural models. *Synthese* 1999;121:151–179.
29. Hernan MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* 2000;11:561–70.
30. Neugebauer R, Fireman B, Roy JA, O'Connor PJ, Selby JV. Dynamic marginal structural modeling to evaluate the comparative effectiveness of more or less aggressive treatment intensification strategies in adults with type 2 diabetes. *Pharmacoepidemiol Drug Saf* 2012;21(2):99–113.
31. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;11:550–60.
32. Bryan J, Yu Z, van der Laan MJ. Analysis of longitudinal marginal structural models. *Biostatistics* 2004;5:361–80.
33. Hernan MA, Hernandez-Diaz S, Robins JM. A structural approach to selection bias. *Epidemiology* 2004;15:615–25.
34. Toh S, Garcia Rodriguez LA, Hernan MA. Confounding adjustment via a semi-automated high-dimensional propensity score algorithm: an application to electronic medical records. *Pharmacoepidemiol Drug Saf* 2011;20:849–57.
35. Rassen JA, Schneeweiss S. Using high-dimensional propensity scores to automate confounding control in a distributed medical product safety surveillance system. *Pharmacoepidemiol Drug Saf* 2012;21 (Suppl 1):41–49.
36. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika* 1983;70:41–55.
37. Jixian W, Peter TD. Propensity score methods in drug safety studies: practice, strengths and limitations. *Pharmacoepidemiology Drug Saf* 2001;10:341–44.
38. Glynn RJ, Schneeweiss S, Sturmer T. Indications for propensity scores and review of their use in pharmacoepidemiology. *Basic Clin Pharmacol Toxicol* 2006;98:253–9.
39. Perkins SM, Tu W, Underhill MG, Zhou X, Murray MD. The use of propensity scores in pharmacoepidemiologic research. *Pharmacoepidemiol Drug Saf* 2000;9:93–101.
40. Hernan MA, Hernandez-Diaz S. Beyond the intention-to-treat in comparative effectiveness research. *Clin Trials*, 2012; 9(1):48–55.
41. Neugebauer R, Silverberg MJ, van der Laan MJ. Observational study and individualized antiretroviral therapy initiation rules for reducing cancer incidence in HIV-infected patients: New York: Springer, 2011;436–456.

42. Young JG, Hernan MA, Picciotto S, Robins JM. Relation between three classes of structural models for the effect of a time-varying exposure on survival. *Lifetime Data Anal* 2010;16:71–84.
43. Tan Z. Comment: understanding OR, PS and DR. *Stat Sci*. 2007;22:560–8.
44. Tsiatis AA. *Semiparametric theory and missing data*. New York: Springer, 2006.
45. Petersen ML, Wang Y, van der Laan MJ, Guzman D, Riley E, Bangsberg DR. Pillbox organizers are associated with improved adherence to HIV antiretroviral therapy and viral suppression: a marginal structural model analysis. *Clin Infect Dis* 2007;45:908–15.
46. Lippman SA, Shade SB, Hubbard AE. Inverse probability weighting in sexually transmitted infection/human immunodeficiency virus prevention research: methods for evaluating social and community interventions. *Sex Transm Dis* 2010;37:512–8.
47. Kang JD Y, Schafer JL. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stati Sci* 2007;22:523–39.
48. Robins J, Sued M, Lei-Gomez Q, Rotnitzky A. Comment: performance of double-robust estimators when “Inverse Probability” weights are highly variable. *Stati Sci* 2007;22:544–59.
49. Breiman L. Random forests. *Machine Learning* 2001;45:5–32.
50. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Ann Stati* 2004;32:407–99.
51. Ruczinski I, Kooperberg C, LeBlanc M. Logic regression. *J Comput Graph Stati* 2003;12:475–511.
52. Breiman L, Friedman JH, Olshen RA, Stone CJ. *Classification and Regression Trees*. Monterey: Wadsworth & Brooks/Cole, 1984.
53. Hoerl AE, Kennard RW. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* 1970;12:55–67.
54. Friedman JH. Multivariate adaptive regression splines. *Ann Stati* 1991;19:1–141.
55. Sinisi SE, van der Laan MJ. Deletion/substitution/addition algorithm in learning with applications in genomics. *Stat Appl Genet Mol Biol* 2004;3:Article 18.
56. Schneeweiss S, Rassen JA, Glynn RJ, Avorn J, Mogun H, Brookhart MA. High-dimensional propensity score adjustment in studies of treatment effects using health care claims data. *Epidemiology* 2009;20:512–22.
57. Ridgeway G, McCaffrey DF. Comment: demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007;22:540–3.
58. van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stati Appli Genet Mol Biol* 2007;6: Article 25.
59. Dudoit S, van der Laan MJ. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Stati Methodol* 2005;2:131–54.
60. van der Laan MJ, Dudoit S, Keles S. Asymptotic optimality of likelihood-based cross-validation. *Stati Appli Genet Mol Biol* 2004;3: Article 4.
61. van der Vaart AW, Dudoit S, van der Laan MJ. Oracle inequalities for multi-fold cross-validation. *Stat Decisions* 2006;24:351–71.
62. van der Laan MJ, Dudoit S, van der Vaart AW. The cross-validated adaptive epsilon-net estimator. *Stat Decisions* 2006;24:373–95.
63. Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. *J Stat Planning Inference* 2007;137:419–34.
64. van der Laan MJ. Statistical inference for variable importance. *Int J Biostat* 2006;2: 1008.
65. van der Laan MJ, Petersen ML, Joffe MM. History-adjusted marginal structural models and statically-optimal dynamic treatment regimens. *Int J Biostat* 2005;1:10–20.
66. Neugebauer R, van der Laan M J, Joffe M, Tager IB. Causal inference in longitudinal studies with history-restricted marginal structural models. *Elect J Stat* 2007;1:119–57.
67. Sinisi SE, Polley E C, Petersen M L, Rhee S Y, van der Laan M J. Super learning: an application to the prediction of HIV-1 drug resistance. *Stat Appl Genet Mol Biol* 2007;6:Article7.
68. Polley EC. SuperLearner R package; <https://github.com/ecpolley/SuperLearner>, 2011.
69. Ripley BD. *Pattern Recognition and Neural Networks*. Cambridge: Cambridge University Press, 1996.
70. Kooperberg C, Bose S, Stone C J. Polychotomous regression. *J Am Stat Assoc* 1997;92:117–27.
71. Cole SR, Hernan AM. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol* 2008;168:656–64.
72. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;11:550–60.
73. Hernan MA, Brumback BA, Robins JM. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Stat Med* 2002;21:1689–1709.
74. R Development Core Team . *R: a language and environment for statistical computing*. Austria: R Foundation for Statistical Computing Vienna, 2011. ISBN 3-900051-07-0.
75. Danaei G, Garcia Rodríguez LA, Cantero OF, Logan R, Hernan MA. Observational data for comparative effectiveness research: an emulation of randomised trials of statins and primary prevention of coronary heart disease. *Stat Methods Med Res* 2013;22(1):70–96.

76. van der Laan MJ., Rose S. Targeted learning: causal inference for observational and experimental data. New York: Springer, 2011.
77. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics* 2005;61:962–73.
78. Yu Z, van der Laan M. Double robust estimation in longitudinal marginal structural models. *Jo Stati Planning Inference* 2006;136:1061–89.
79. van der Laan MJ, Gruber S. Targeted minimum loss based estimation of an intervention specific mean outcome. tech. Rep. 290, Division of Biostatistics, UC Berkeley, 2011.