Lessons from Simulations of Marginal Models for Dynamic Treatments with Survival Outcomes

Vanessa Didelez
School of Mathematics
University of Bristol

(joint work with various people)

ISCB — Utrecht, August 2015

Overview

- Motivation and target of inference.
- Dynamic marginal structural models and inverse probability weighting.
- Problems with simulating from DMSMs & examples.
- Conclusions.

Motivation and Target of Inference

Motivation

- Marginal structural models (MSMs), fitted by the method of inverse probability of treatment weighting (IPTW) have been developed to cope with time–dependent confounding.

- Their superiority to standard regression adjustment, especially with longitudinal data and survival analysis, has been emphasised.

- The parameters of MSMs always have a clear causal interpretation, e.g. like comparing dynamic treatment strategies in an RCT. In contrast, the parameters of e.g. proportional hazard models with time-dependent covariates do not necessarily have a causal interpretation.

- IPTW yields consistent estimates of marginal causal effects if the MSM is correctly specified and time-dependent confounding fully observed.

- It seems obvious that we would want to evaluate the use of MSMs and the performance of IPTW (e.g. for finite samples) on simulated data for which we know the 'truth',
  – e.g. to evaluate robustness towards violations of assumptions.

- Constructing a simulation algorithm can also be very instructive for understanding the model class used.

- **However**, especially for survival models, it is not straightforward to simulate from a given MSM. As we will see, the problem is essentially one of non–collapsibility.

Example: HIV Studies

Longitudinal HIV Studies (e.g. Cain et al., 2010, 2011)

Earlier research: is antiretroviral treatment helpful / by how much?

Current research: when should antiretroviral treatment start, e.g. as soon as CD4–count drops below 500 or below 300?

Future research: when, depending on patient's history, should we switch between different antiretroviral treatments?

Observational Data: so far mainly observational data available, doctors make treatment decisions as they think best.

Target of Inference

Ideally: target of inference is clear if we can

— formulate **decision problem** that will be informed

— describe **experiment** to measure the desired quantity explicitly

$\Rightarrow$ should guide the design, collection of data, assumptions, and analysis.

Here: ideally, randomise patients to different strategies

$\rightarrow$ "start treatment when CD4 drops below 600 / 500 ... / 200"

Aim: Compare survival chances under these different strategies.

$\Rightarrow$ e.g. 5–year or 10–year survival probabilities.

Data Situation

$A_1, A_2, \ldots$ "action" variables $\rightarrow$ can be 'manipulated'

$L_1, L_2 \ldots$ covariates $\rightarrow$ (obs.) background information (incl. 'at risk')

Y response variable, here (some aspect of) survival

all measured over time, L_t before A_t

$\bar{X}_t = (X_1, \ldots, X_t)$ past up to t $\qquad$ $\bar{X}^t = (X_t, X_{t+1} \ldots)$ future from t

(Continuous time also possible)

Strategies

Strategy $\mathbf{s} = (s_1, s_2 \ldots)$ set of functions assigning an action

$a_t = s_t(\bar{l}_t)$ to each history $(\bar{l}_t)$

Types of strategies

(Murphy, 2003)

static: fix $s_t \equiv a_t$ regardless of $\bar{l}_t$, $t = 1, 2 \ldots$, e.g.

"never treat" $\bar{\mathbf{s}} = 0$;

"start treatment at $t = 3$, stay on" $\bar{\mathbf{s}}_2 = 0$ and $\bar{\mathbf{s}}^3 = 1$.

dynamic: s_t indeed function of $\bar{l}_t$ for some $t = 1, 2 \ldots$, e.g.

"start treatment when L_t first drops below 350, never stop"

i.e. $s_t(\bar{l}_t) = 1$ if $\{\exists k \leq t$ s.t. $l_k < 350\}$ otherwise $= 0$.

Note: under dynamic strategy $\mathbf{s}$ we don't necessarily know actions in advance as they depend on the random histories.

Strategies Well Defined?

Strategies:

must be defined *in advance*, and functions of the past, as in actual randomised trial.

Counter-example:

"patients who die before they start treatment belong to control group."

Target of Inference (Formally)

Effect of a strategy: property / parameter of the **intervention distributions** (Pearl, 2000)

$$E(u(Y); do(\mathbf{s})), \qquad \mathbf{s} \in \mathcal{S}$$

where $do(\cdot)$ means we follow a given pre-specified treatment strategy, $u(\cdot)$ utility function. (Can also use potential responses $Y(\mathbf{s})$)

Marginal Structural Models (MSMs):

parameterises (aspect of) $E(u(Y); do(\mathbf{s}))$

Identifying Assumptions

Inference on MSMs from observational data relies on:

"No Unmeasured Confounding" (Robins, 1986; et al.)

"Stability" (Dawid & Didelez, 2010)

Here:

sufficient to observe L,

can then ignore U

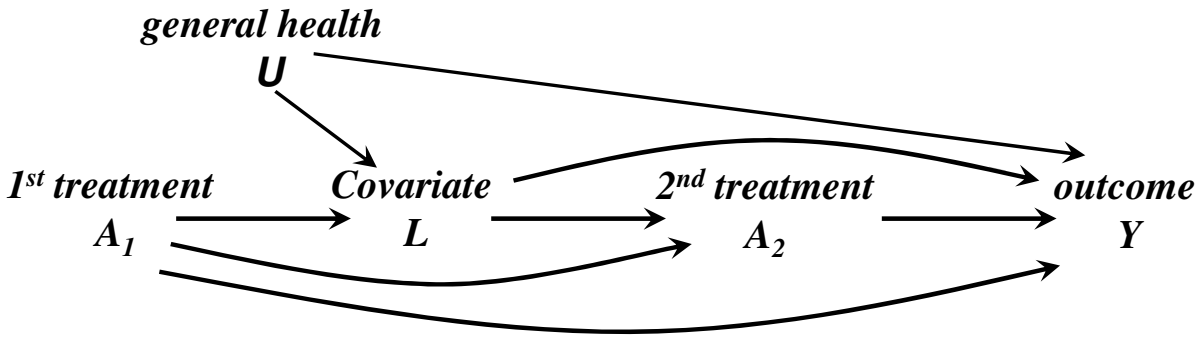

Dynamic Marginal Structural Models and Inverse Probability Weighting

G–Computation / Formula

(Robins, 1986)

Observational joint distribution:

$$P(Y, A_2, A_1, L, U) = P(Y|A_2, A_1, L, U)$$
$$P(A_2|A_1, L)P(L|A_1, U)P(A_1)P(U)$$

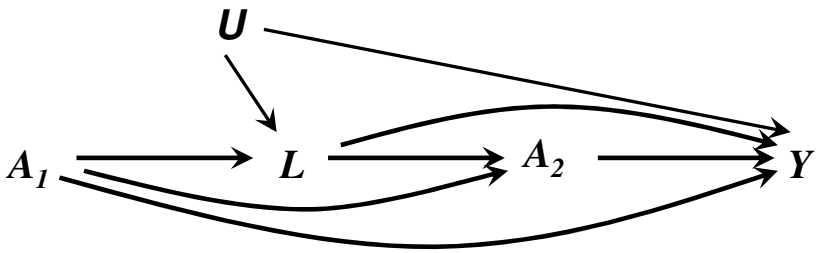

G–Formula: set $A_1 = a_1, A_2 = a_2$, marginalise the rest

$$P(Y; do(a_1, a_2)) = \sum_{L,U} P(Y|A_2 = a_2, A_1 = a_1, L, U)\ P(L|A_1 = a_1, U)P(U)$$
$$= \sum_L P(Y|A_2 = a_2, A_1 = a_1, L)$$
$$\times P(L|A_1 = a_1)$$

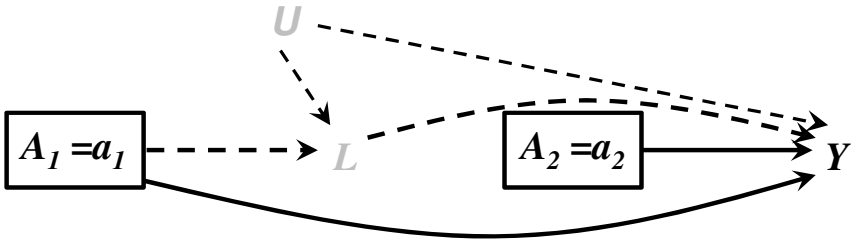

(Note: same as "extensive form analysis", Dawid & Didelez, 2010)

Estimation: estimate cond. distributions, use MC to predict effects.

Marginal Structural Models (MSMs)

(Robins et al., 2000)

MSM: parameterises aspect of $P(Y; do(\mathbf{s}))$,
e.g. marginal structural proportional hazard model.

Rationale

$$P(Y; do(a_1, a_2)) = \frac{\int \text{joint}}{P(A_2 = a_2 | A_1 = a_1, L) P(A_1 = a_1)}$$

$\Rightarrow$ fit an MSM to observational data but with weights

$\Rightarrow$ inverse probability of treatment weighting (IPTW)

$\Rightarrow$ can be seen as 'change of measure'

(weights = R-N derivative $P(\cdot; do(\mathbf{s}))/P(\cdot; obs)$).

Weights

Different ways to calculate weights:

— time-fixed outcome $\Rightarrow$ fixed weights

— survival outcome $\Rightarrow$ weights depend on 'at risk'

— static strategies: typically, each subject represents it's 'own' strategy

— dynamic strategies: subject contributes to more than one strategy

IPTW for Static Strategies (1)

Most popular: weights for subject i at time k (if still at risk)

$$W_{k,i} = \prod_{t=1}^{k} \frac{1}{P(A_t = a_{t,i} | \bar{L}_t = \bar{l}_{t,i}, \bar{A}_{t-1} = \bar{a}_{t-1,i})}$$

$\Rightarrow$ 'change of measure': in weighted data, each subject's treatment like randomised.

$\Rightarrow$ carry out any inference (on weighted data set) that would be valid under randomisation of treatment.

Note 1: problems with rare treatments $\Rightarrow$ stabilised weights.

Note 2: also for continuous time processes (Roysland, 2011)

IPTW for Static Strategies (2)

(Hernan et al., 2006)

Alternatively:
— identify strategies to be compared
— count every subject in every strategy that they comply with as long as possible, then artificially censor.

Example: "start treatment at time t"
$\Rightarrow$ every subject not on treatment by t can be included until t.

$\Rightarrow$ **Strategy specific weights** for individual i under strategy $\mathbf{s} = \mathbf{a}^*$:

$$W_{k,i}(\mathbf{s}) = \prod_{t=1}^{k} \frac{I\{a_{t,i} = a^*\}}{P(A_t = a_{t,i} | \bar{L}_t = \bar{l}_{t,i}, \bar{A}_{t-1} = \bar{a}_{t-1,i})}$$

(see also Gran et al., 2010)

IPTW for Dynamic Strategies

(Orellana et al., 2010)

$\Rightarrow$ **Strategy specific weights** for individual i under strategy $\mathbf{s}$:

$$W_{k,i}(\mathbf{s}) = \prod_{t=1}^{k} \frac{I\{a_{t,i} = s_t(\bar{l}_{t,i})\}}{P(A_t = a_{t,i} | \bar{L}_t = \bar{l}_{t,i}, \bar{A}_{t-1} = \bar{a}_{t-1,i})}$$

$\Rightarrow$ create 'replicants' for each strategy

$\Rightarrow$ weights + artificial censoring

$\Rightarrow$ uses all data

Example: patient starts treatment at $t = 4$

we recorded $L_1 = 500$, $L_2 = 450$, $L_3 = 400$, $L_4 = 350$ $\Rightarrow$

– agrees with "start when $L_t \leq 300$" at $t = 1, 2, 3$

– agrees with "start when $L_t \leq 400$" at $t = 1, 2$; etc.

IPTW = IPCW

(Hernan et al., 2006; Cain et al., 2010)

Artificial censoring:

can regard weights as 'inverse probability of censoring weights'

Rationale:

$P(\text{remaining uncensored wrt. } \mathbf{s}) = P(\text{action agrees with strategy } \mathbf{s})$.

$\Rightarrow$ IPCW prevents **selection bias** due to artificial censoring.

Note: 'natural censoring' $\rightarrow$ **separate** process.

Dynamic MSM — Example

Example: let $\mathbf{s}_x =$ "start treatment when CD4 drops below x"

Dynamic MSM: hazard function under strategy $\mathbf{s}_x$

$$\lambda(t; do(\mathbf{s}_x)) = \lambda_0(t)\varphi(x, t; \beta)$$

(Here, omitting baseline covariates from notation.)

Choose $\varphi(\cdot)$ sufficiently rich and flexible, e.g. constant hazard ratio implausible.

$\Rightarrow$ evaluate target of inference, e.g. 5-year survival.

Note: model in terms of strategy, not treatment sequence.

An Application

"When to start": Cain et al. (2010), French Hospital data on HIV.

Using IPW, artificial censoring, for dynamic strategy authors find

start when first	**Prob. 5-year survival**	95% CI
CD4 $<$ 500	0.95	[0.91,0.98]
CD4 $<$ 400	0.93	[0.90,0.97]
CD4 $<$ 300	0.94	[0.92,0.96]
CD4 $<$ 200	0.91	[0.89,0.94]

Note: analysis more plausible with 'grace period'.

Simulation Studies to Evaluate MSM/IPTW

How to Simulate from MSM?

Wanted: find data generating mechanism for $\{\bar{U}_t, \bar{L}_t, \bar{A}_t, Y\}$ s.t.

(1) when simulating randomised trial MSM assumptions satisfied

(2) when simulating obs. study exhibit time-dependent confounding.

Note: difference between (1) and (2) only in $P(A_t|$ past $)$.

Problem:

(1) *marginal* over $\bar{U}_t, \bar{L}_t$;

BUT (2) best implemented with *conditional* models.

DAGs and MSMs

Data generation via DAG — defined by conditional distributions.

$$P(Y, A_2, A_1, L, U) = P(Y|A_2, A_1, L, U)P(A_2|A_1, L)P(L|A_1, U)P(U)P(A_1)$$

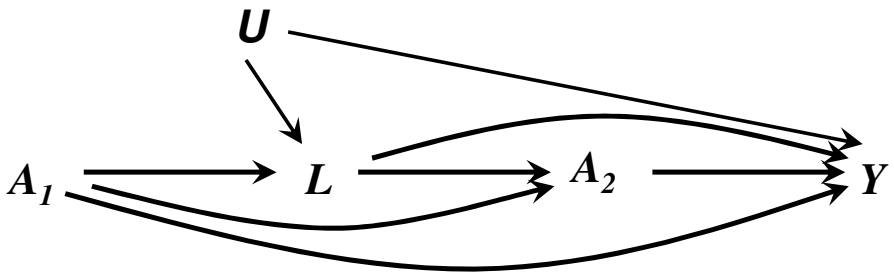

DAGs and MSMs

Data generation via DAG

$$P(Y, A_2, A_1, L, U) = P(Y|A_2, A_1, L, U)P(A_2|A_1, L)P(L|A_1, U)P(U)P(A_1)$$

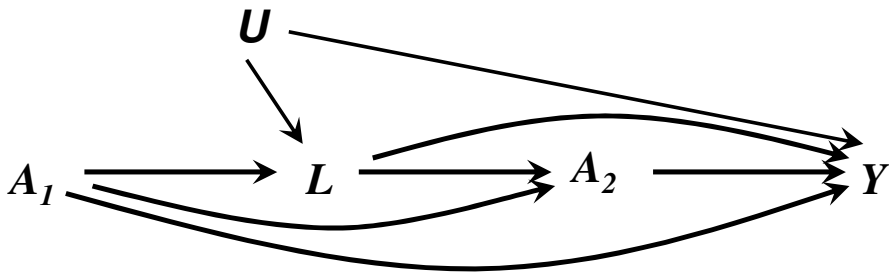

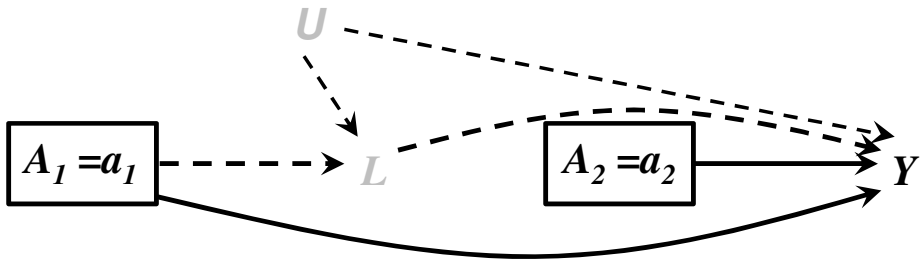

MSM — marginal model for the effect of intervention $do(a_1, a_2)$:

$$\underbrace{P(Y; do(a_1, a_2))}_{\text{MSM}} = \underbrace{\sum P(Y|A_2, A_1, L, U)P(L|A_1, U)P(U)I\{A_2 = a_2, A_1 = a_1\}}_{\text{G–formula}}$$

$\Rightarrow$ for desired MSM need appropriate choice of rhs factors!

$\Rightarrow$ not feasible for *non–collapsible* models (logistic, prop. hazard).

Simulating from MSMs

Static Strategies: with survival–type outcome:

Exact MSM Simulation:
Simulate from given MSM but without $L_t \longrightarrow Y$ effect:
Young et al. (2008, 2010) for continuous time;
Havercroft & Didelez (2011) for discrete time (pooled logistic)
(Bryan et al. (2004): similar, but no time depending confounding.)

Approximate MSM Simulation:
Simulate from DAG, without U (no selection bias), with sufficiently rare event rate such that model is 'nearly collapsible': Xiao et al. (2010)

Exact (Static) MSM Simulation

Essential idea:

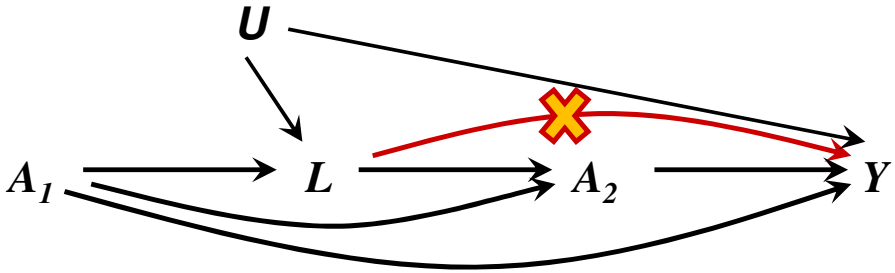

$$P(Y; do(a_1, a_2))$$
$$= \sum_U P(Y|A_2 = a_2, A_1 = a_1, U)P(U)$$

$\Rightarrow$ choose distribution of U, transform, so as to have desired distribution $P(\cdot; do(a_1, a_2))$. (inverse transform sampling.)

Note: cannot be used for dynamic strategies, as outcome must then depend on covariates used.

Approximate Simulation for Dynamic MSM

(Havercroft & Didelez, 2015?)

Wanted: dependence of survival on treatment and CD4 s.t. "starting at CD4 x" is optimal.

(Structural) Accelerated Failure Time Model:

$$P(Y; do(\bar{0})) =$$

$$P\left(\int_0^Y \exp\{\xi a_t g(L_t)\} dt \,;\, do(\bar{a})\right)$$

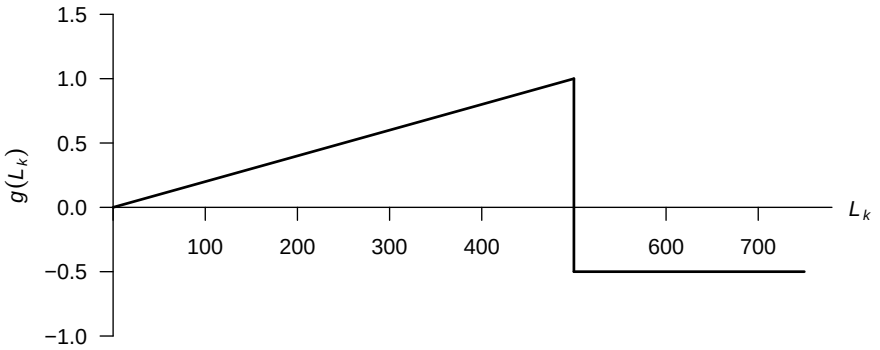

— at baseline $Y; do(\bar{0})$ Weibull;

— baseline CD4 depends on 'unobserved general health';

— CD4 downward trend, but one-off boost when treatment started;

— still all quite simplistic!

Simulate Randomised Trial

Truth: simulate N subjects for each $s_x =$ "start when CD4 $< x$."

Survival functions and Survival probabilities

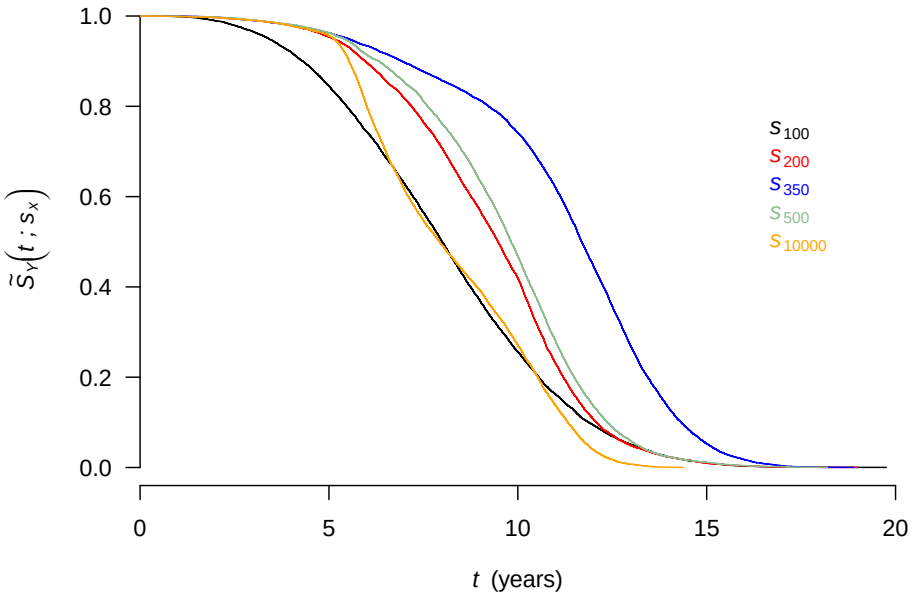

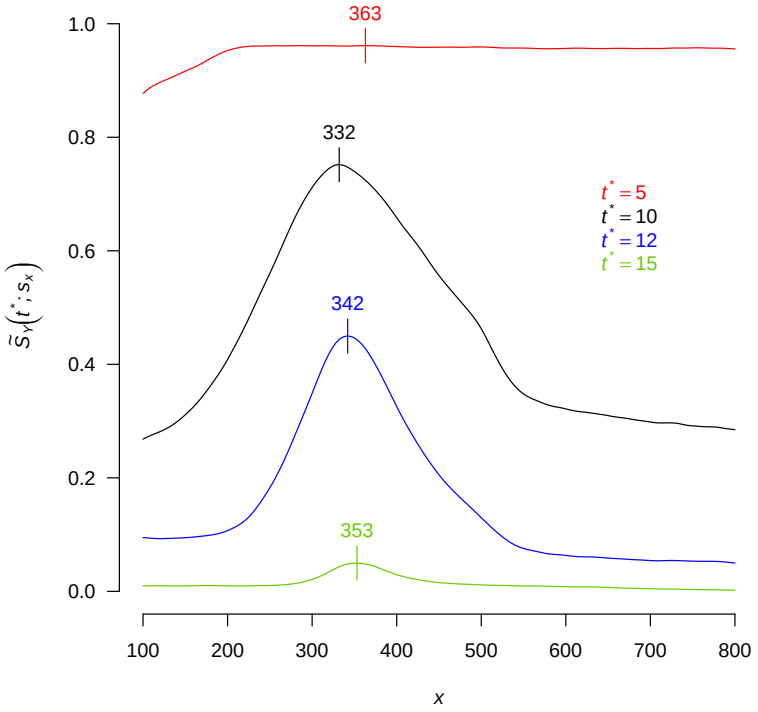

Note: model for Y alone does not determine optimal x!

Flexible Dynamic MSM?

Now: simulate observational data, $P(A_t|L_t, A_{t-1})$.

Apply: time-dependent MSM hazard model, restricted cubic spline in x (5, 8, 9 knots) — find max. survival prob.

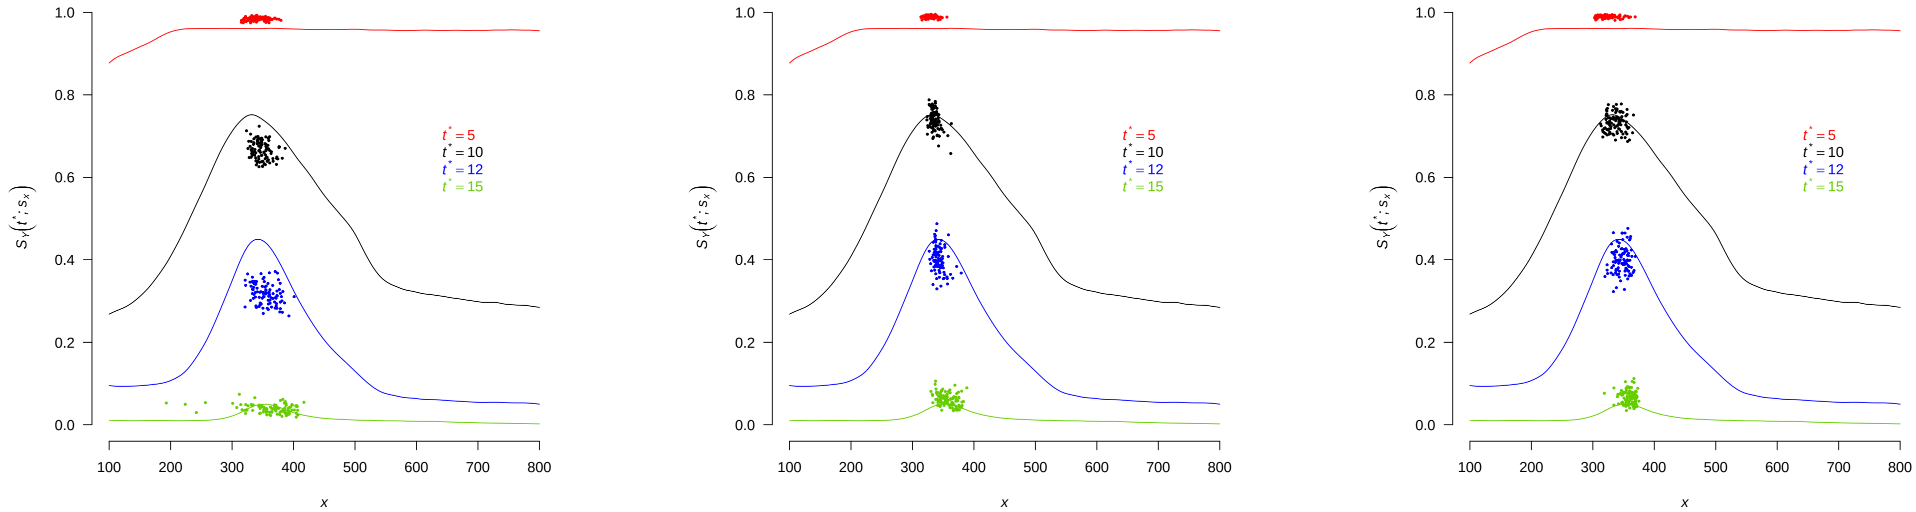

Simulation of a Misinterpretation

Kitahata et.al. (2009) "when to start" analysis: compare two groups

(1) "early" start within 6 months while still $350 <$CD4< 500;

(2) "deferred" start treatment within 6 months after CD4< 350.

$\Rightarrow$ do not correspond to strategies.

Also: groups constraint to be disjoint and 'not (1)' is included in (2), e.g. patients who die before treatment.

$\Rightarrow$ induces selection bias.

Simulation of a Misinterpretation

Simulation:

1) true curves — as in RCT

2) correct d-IPTW analysis

3) Kitahata's two groups.

⇒ Difference b/w
the two groups is
wrongly exaggerated.

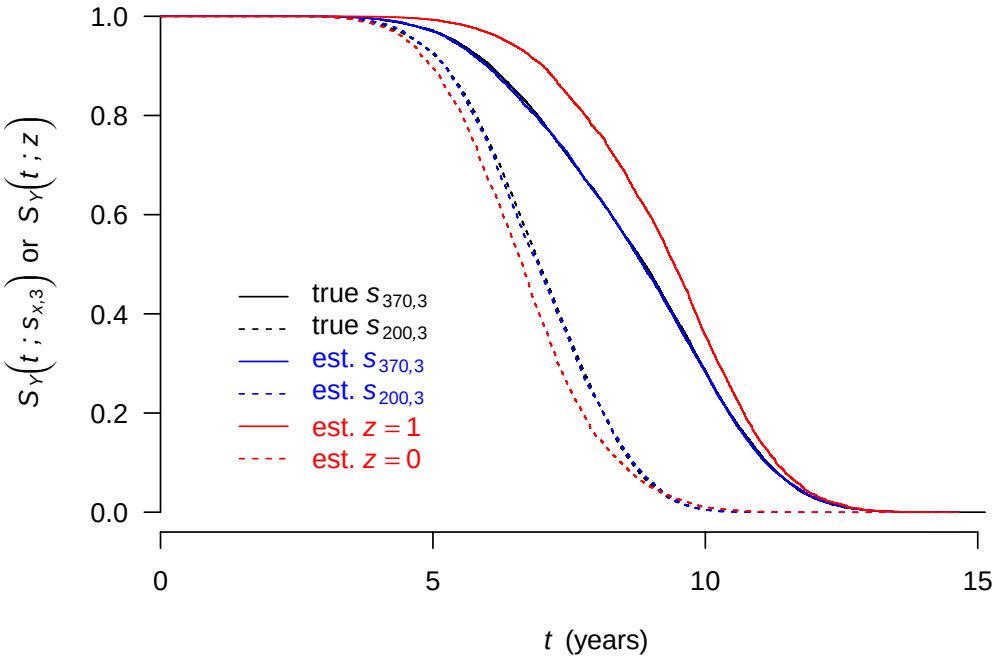

Conclusions and Outlook

- Dynamic strategies need to be well-defined.
- MSMs can be used for evaluating dynamic strategies by artificial censoring; can find optimal strategy out of a limited set of strategies...
- ... relatively simple to implement, easy to interpret
- ... but not designed to provide insight into mechanisms
- ... and sufficiently but not too flexible model choice open problem.
- Simulations: data generating process $\leftrightarrow$ model assumptions.

  Much scope for more realistic simulations.

References

Arjas, E. and Saarela, O. (2010). Optimal dynamic regimes: Presenting a case for predictive inference. Intern. J. Biostat. 6. `http://tinyurl.com/33dfssf`

Cain, LE.; Robins, JM.; Lanoy, E; Logan, R; Costagliola, D; and Hernan, MA. (2010). When to Start Treatment? A Systematic Approach to the Comparison of Dynamic Regimes Using Observational Data. Intern. J. Biostat. 6, 2, Article 18. DOI: 10.2202/1557-4679.1212

Cain LE, Logan R, Robins JM, Sterne JA, Sabin C, Bansi L, Justice A, Goulet J, van Sighem A, de Wolf F, Bucher HC, von Wyl V, Esteve A, Casabona J, del Amo J, Moreno S, Seng R, Meyer L, Perez-Hoyos S, Muga R, Lodi S, Lanoy E, Costagliola D, Hernan MA. (2011). When to initiate combined antiretroviral therapy to reduce mortality and AIDS-defining illness in HIV-infected persons in developed countries: an observational study. Ann. Intern. Med. 154(8), 509-15.

Dawid, AP., Didelez, V. (2008). Identifying optimal sequential decisions, Proceedings of the 24th Annual Conference on Uncertainty in Artifical Intelligence, 113-120.

Dawid, AP., Didelez, V. (2010). Identifying the consequences of dynamic treatment strategies: A decision theoretic overview, Statistics Surveys, 4, 184-231.

References

Gran, J. M.; Roysland, K.; Wolbers, M.; Didelez, V; Sterne, J.A.C.; Ledergerber, B.; Furrer, H.; von Wyl, V.; Aalen, O.O. (2010). A sequential Cox approach for estimating the causal effect of treatment in the presence of time-dependent confounding applied to data from the Swiss HIV Cohort Study, Statistics in Medicine, DOI: 10.1002/sim.4048

Havercroft, W., Didelez, V. (2012). Simulating from marginal structural models with time-dependent confounding. *Submitted*

Henderson, R., Ansel, P. and Alshibani, D. (2010). Regret-regression for optimal dynamic treatment regimes. Biometrics, 66, 4, 1192-1201.

Hernan, Miguel A., Emilie Lanoy, Dominique Costagliola and James M. Robins (2006). Comparison of Dynamic Treatment Regimes via Inverse Probability Weighting. Basic & Clinical Pharmacology & Toxicology, 98, 237242.

Murphy, S.A. (2003). Optimal dynamic treatment regimes (with Discussion). Journal of the Royal Statistical Society, Series B 65 331-366.

Orellana, L.; Rotnitzky, A.; and Robins, JM. (2010). Dynamic Regime Marginal Structural Mean Models for Estimation of Optimal Dynamic Treatment Regimes, Part I: Main Content. Intern. J. Biostat. 6, 2, Article 8. DOI: 10.2202/1557-4679.1200

References

Pearl, J. and Robins, J. (1995). Probabilistic evaluation of sequential plans from causal models with hidden variables. Proceedings 11th UAI Conference (eds. Besnard and Hanks) 444-453. Morgan Kaufmann Publishers, San Francisco.

Robins, J.M. (1997). Causal inference from complex longitudinal data. In Latent Variable Modeling and Applications to Causality, (M. Berkane, ed.). Lecture Notes in Statistics 120 69117. Springer-Verlag, New York.

Robins JM, Hernan MA, Brumback B. (2000). Marginal structural models and causal inference in epidemiology. Epidemiology 11, 550-560.

Rosthoj, S., Fullwood, C., Henderson, R. and Stewart, S. (2006). Estimation of optimal dynamic anticoagulation regimes from observational data: A regret-based approach. Statistics in Medicine, 25, 41974215.

Sterne, J.A.C., Hernan, M.A., Ledergerber B., Tilling, K., Weber, R., Sendi, P. and Rickenbach, M., Robins, J., Egger, M. and the Swiss HIV Cohort Study. (2005) Long-term effectiveness of potent antiretroviral therapy in preventing AIDS and death: the Swiss HIV Cohort Study. Lancet, 366, 378-84.