Estimating adaptive treatment strategies for survival outcomes

Erica E M Moodie

Biostatistics, McGill University
erica.moodie@mcgill.ca

LSHTM Centre for Statistical Methodology

Road-map

1. Precision medicine in the statistical literature
2. Overview of regression-based methods of ATS estimation for continuous outcomes
3. Extending to the censored case
4. A case study: second- and third-line therapies for T2D using observational registry data

Evidence-based medicine: The statisticians' role

> [T]he medical statistician recognizes, and is familiar with the pros and cons of, that difficult question – should a fixed dose be given to all patients in a trial or should it be allowed to vary with the apparent needs of each patient as judged by the clinician?
>
> *Sir Austin Bradford Hill (1962)*

Evidence-based medicine: The statisticians' role

> Evidence-based medicine is the conscientious, explicit, and judicious use of the current best evidence in making decisions about the care of individual patients.
>
> *David Sackett (1992)*

Precision medicine

- **Precision:** Refers to the tailoring of medical treatment to the individual characteristics of each patient.

Precision medicine

- **Precision:** Refers to the tailoring of medical treatment to the individual characteristics of each patient.
  - Does not mean selection or creation of drugs that are unique to a patient.
  - Does mean the ability to *classify individuals* into subpopulations that differ in their response to a specific treatments – perhaps repeatedly over time.

Precision medicine

- **Precision:** Refers to the tailoring of medical treatment to the individual characteristics of each patient.
  - Does not mean selection or creation of drugs that are unique to a patient.
  - Does mean the ability to *classify individuals* into subpopulations that differ in their response to a specific treatments – perhaps repeatedly over time.

Precision medicine

- The context:
  - several treatments or doses available;
  - often more than one treatment decision (sequences);
  - (often) must account for clinician decisions about treatment allocation.

Precision medicine

- The context:
  - several treatments or doses available;
  - often more than one treatment decision (sequences);
  - (often) must account for clinician decisions about treatment allocation.

- Make use of information on patient characteristics such as
  - demographics, genetics, genomics;
  - physiologic or clinical measures;
  - medical history, etc.

  in order to determinate *which treatment* the patient should take *and when*.

When would we want treatment to be adaptive?

- Treatment tailoring is better because (when):
  - there is heterogeneity in patient response;
  - patient response may change over time;
  - response to treatment may inform future treatment choices;
  - patient compliance may be imperfect;
  - over-treating can lead to side-effects, treatment fatigue (poor compliance), and higher costs;
  - under-treating can lead to poorer patient outcomes.

When would we want treatment to be adaptive?

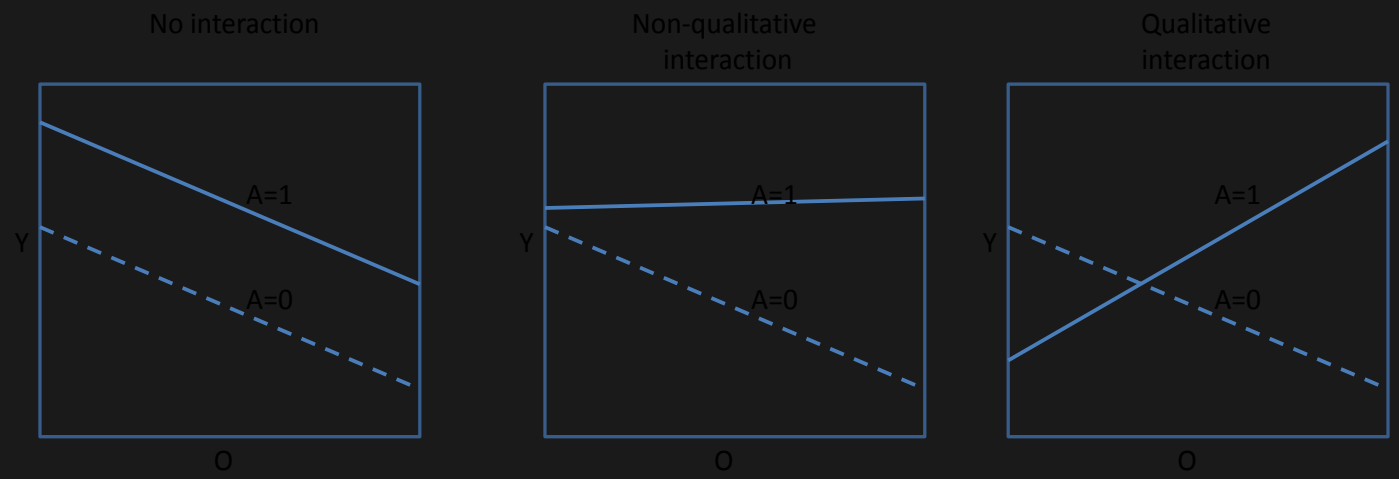

- Variables used to make treatment more targeted are called prescriptive or tailoring variables.

When would we want treatment to be adaptive?

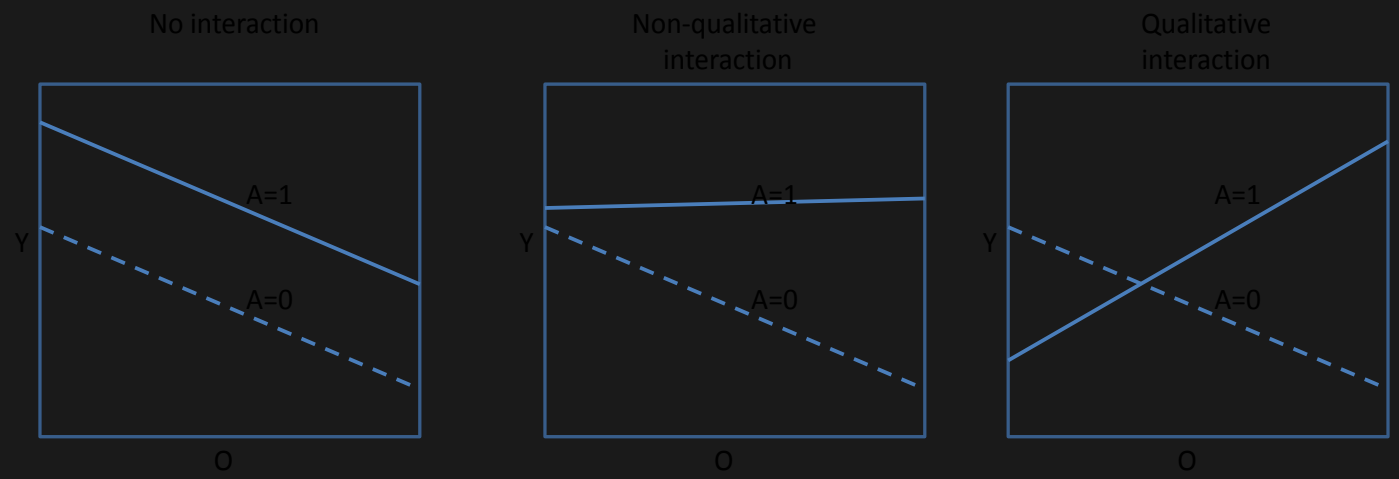

- Variables used to make treatment more targeted are called prescriptive or tailoring variables.
- Resulting treatment algorithms known as adaptive treatment strategy (ATSs), dynamic treatment regimes, adaptive interventions, or policies.

Analytic methods

- Considerable interest in estimation and inference for ATSs in statistics and CS over the $\sim$ 15 years.
- More common strategies include
  - Q-learning (sequential regression);
  - g-estimation;
  - dynamic weighted ordinary least squares (dWOLS);
  - weighted value search estimators, including weighted classifiers such as OWL and RWL, and (A)IPW.

The Single-stage Setting

Notation

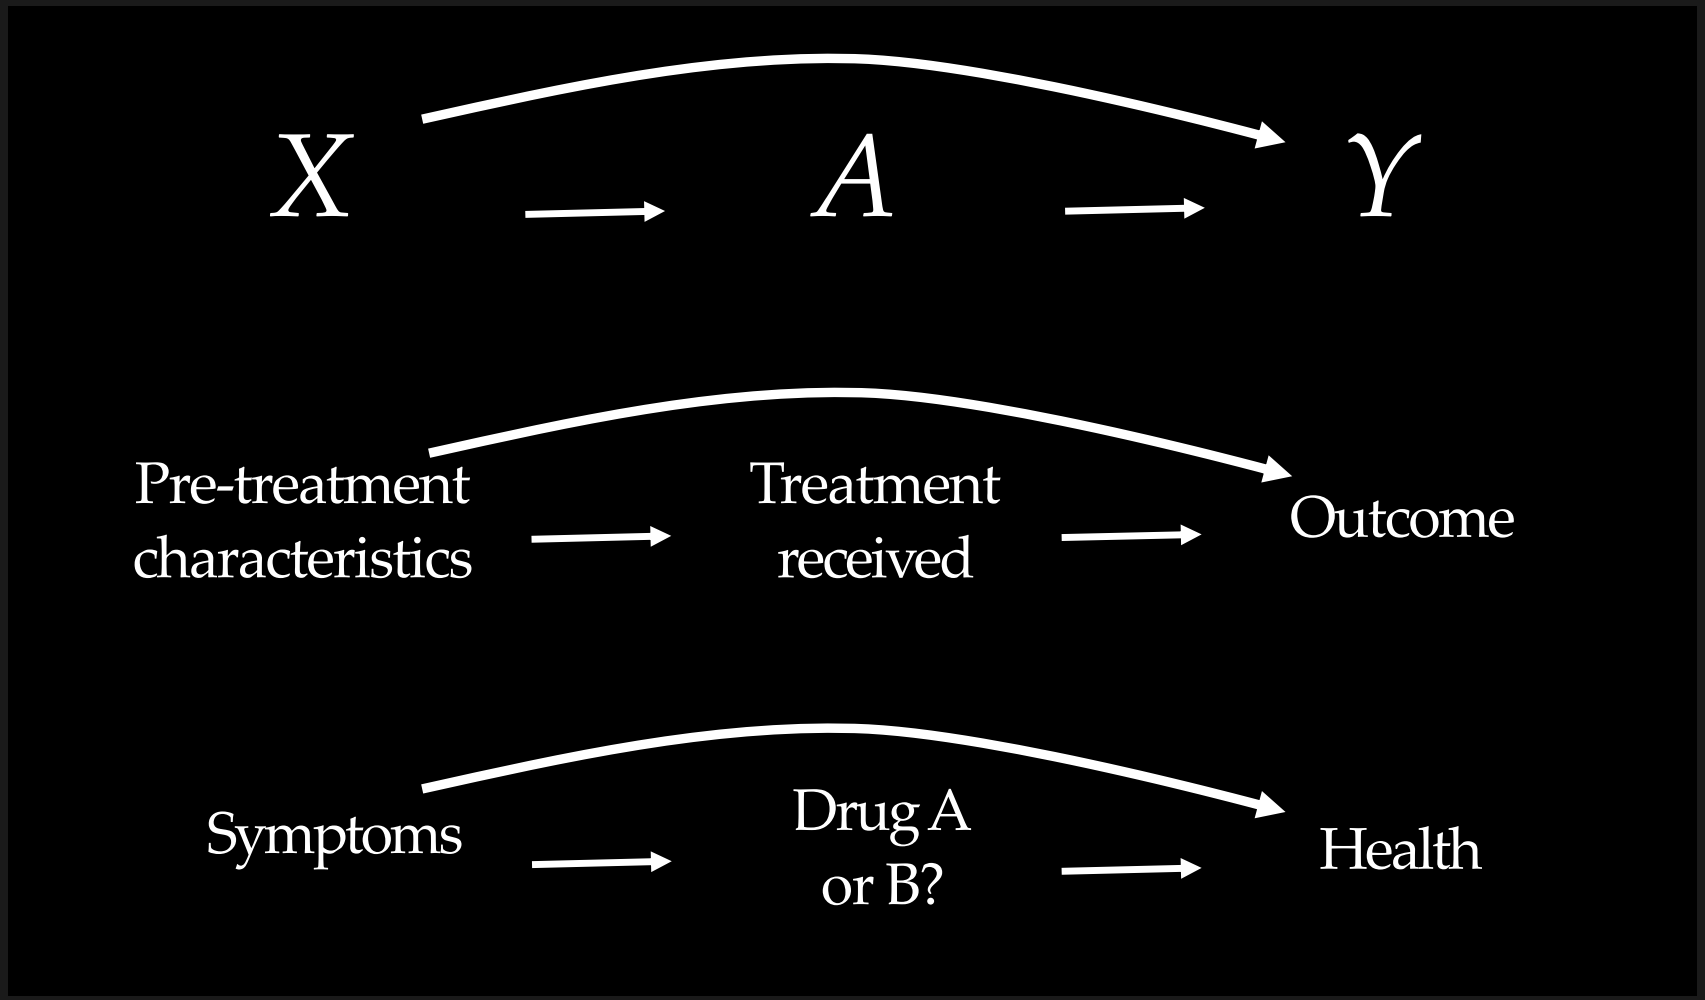

ATSs: how do we find treatment A^{opt} that maximizes Y?

Identifying the best treatment regime

- If only one treatment decision: $\mathbb{E}[Y|X, A]$
- E.g., we might propose the following model

$$\mathbb{E}[Y|X, A; \psi, \beta] = \beta_0 + \beta_1 \mathsf{Sx} + A(\psi_0 + \psi_1 \mathsf{Sx})$$

Identifying the best treatment regime

- If only one treatment decision: $\mathbb{E}[Y|X, A]$
- E.g., we might propose the following model

$$\mathbb{E}[Y|X, A; \psi, \beta] = \beta_0 + \beta_1 \mathsf{Sx} + A(\psi_0 + \psi_1 \mathsf{Sx})$$

$$\underbrace{\mathbb{E}[Y|X, A; \psi, \beta]}_{\substack{\text{Expected outcome} \\ \text{(to be maximized)}}} = \overbrace{G(X; \beta)}^{\substack{\text{Impact of patient history} \\ \text{in the absence of treatment}}} + \underbrace{\gamma(X, A; \psi)}_{\substack{\text{Impact of treatment} \\ \text{on outcome}}}$$

Identifying the best treatment regime

- If only one treatment decision: $\mathbb{E}[Y|X, A]$
- E.g., we might propose the following model

$$\mathbb{E}[Y|X, A; \psi, \beta] = \beta_0 + \beta_1 \mathsf{Sx} + A(\psi_0 + \psi_1 \mathsf{Sx})$$

Impact of patient history in the absence of treatment

$$\underbrace{\mathbb{E}[Y|X, A; \psi, \beta]}_{\text{Expected outcome (to be maximized)}} = \overbrace{G(X; \beta)}^{\text{Impact of patient history in the absence of treatment}} + \underbrace{\gamma(X, A; \psi)}_{\text{Impact of treatment on outcome}}$$

- In practice, we specify a model for the contrast, e.g.:

$$\gamma(x, a; \psi) = a(\psi_0 + \psi_1 x_1 + \psi_2 x_2),$$

then if $a \in \{0, 1\}$, $a^{\text{opt}}(x) = \mathbb{I}[\psi_0 + \psi_1 x_1 + \psi_2 x_2 > 0]$.

Identifying the best treatment regime

- If only one treatment decision: $\mathbb{E}[Y|X, A]$
- E.g., we might propose the following model

$$\mathbb{E}[Y|X, A; \psi, \beta] = \beta_0 + \beta_1 Sx + A(\psi_0 + \psi_1 Sx)$$

Impact of patient history in the absence of treatment

$$\underbrace{\mathbb{E}[Y|X, A; \psi, \beta]}_{\text{Expected outcome (to be maximized)}} = \overbrace{G(X; \beta)} + \underbrace{\gamma(X, A; \psi)}_{\text{Impact of treatment on outcome}}$$

- In practice, we specify a model for the contrast, e.g.:

$$\gamma(x, a; \psi) = a(\psi_0 + \psi_1 x_1 + \psi_2 x_2),$$

then if $a \in \{0, 1\}$, $a^{\text{opt}}(x) = \mathbb{I}[\psi_0 + \psi_1 x_1 + \psi_2 x_2 > 0]$.

- Contrasts specified in this way $\rightarrow$ *linear* decision rules.
- All methods considered assume $\gamma(x, a; \psi)$ well-specified.

Q-learning

- Q-learning is a popular method in reinforcement/machine learning that has been shown, in a multi-stage setting, to be equivalent to a series of OLS regressions.
- In the single stage setting, it *is* simply linear regression.

$$\mathbb{E}[Y|X = x, A = a; \psi, \beta] = G(x; \beta) + \gamma(x, a; \psi)$$

- Two models to specify:
  1. Contrast model: $\gamma(x, a; \psi)$.
  2. Treatment-free model: $G(x; \beta)$.
- Use these models to estimate ψ (and β) via (e.g.) OLS.

Q-learning: editorial comments

- Trivial to implement in standard software;
- easy to explain to clinical collaborators;
- works nicely with continuous treatments (doses);
- extends easily to censored outcomes;
- many useful tools including standard variable selection methods (e.g. LRT) and residual diagnostics.

- Challenge/limitation: *not* robust to model mis-specification.

Dynamic Weighted OLS (dWOLS)

$$\mathbb{E}[Y|X = x, A = a; \psi, \beta] = G(x; \beta) + \gamma(x, a; \psi)$$

- Three models to specify:
  1. Contrast model: $\gamma(x, a; \psi)$.
  2. Treatment-free model: $G(x; \beta)$.
  3. Treatment model: $\mathbb{E}[A|X = x; \alpha]$.
- Estimate ψ via WOLS with weights satisfying

$$\pi(x)w(1, x; \alpha) = (1 - \pi(x))w(0, x; \alpha),$$

for $\pi(x) = \mathbb{E}[A|X = x; \alpha]$, e.g. $w = |A - \mathbb{E}[A|x; \hat{\alpha}]|$, inverse probability of treatment weighting, etc.

dWOLS: editorial comments

- Doubly robust;
- (quite) easy to explain to clinical collaborators;
- has been generalized to accommodate multiple treatments, and to continuous doses (GdWOLS);
- extended to handle censored outcomes using accelerated failure time models (dwSurv);
- many useful tools including residual diagnostics, model-validation based on double-robustness.

The Multi-stage Setting

Identifying the best treatment regime: multi-stage

Recall data set up in the ITR setting:

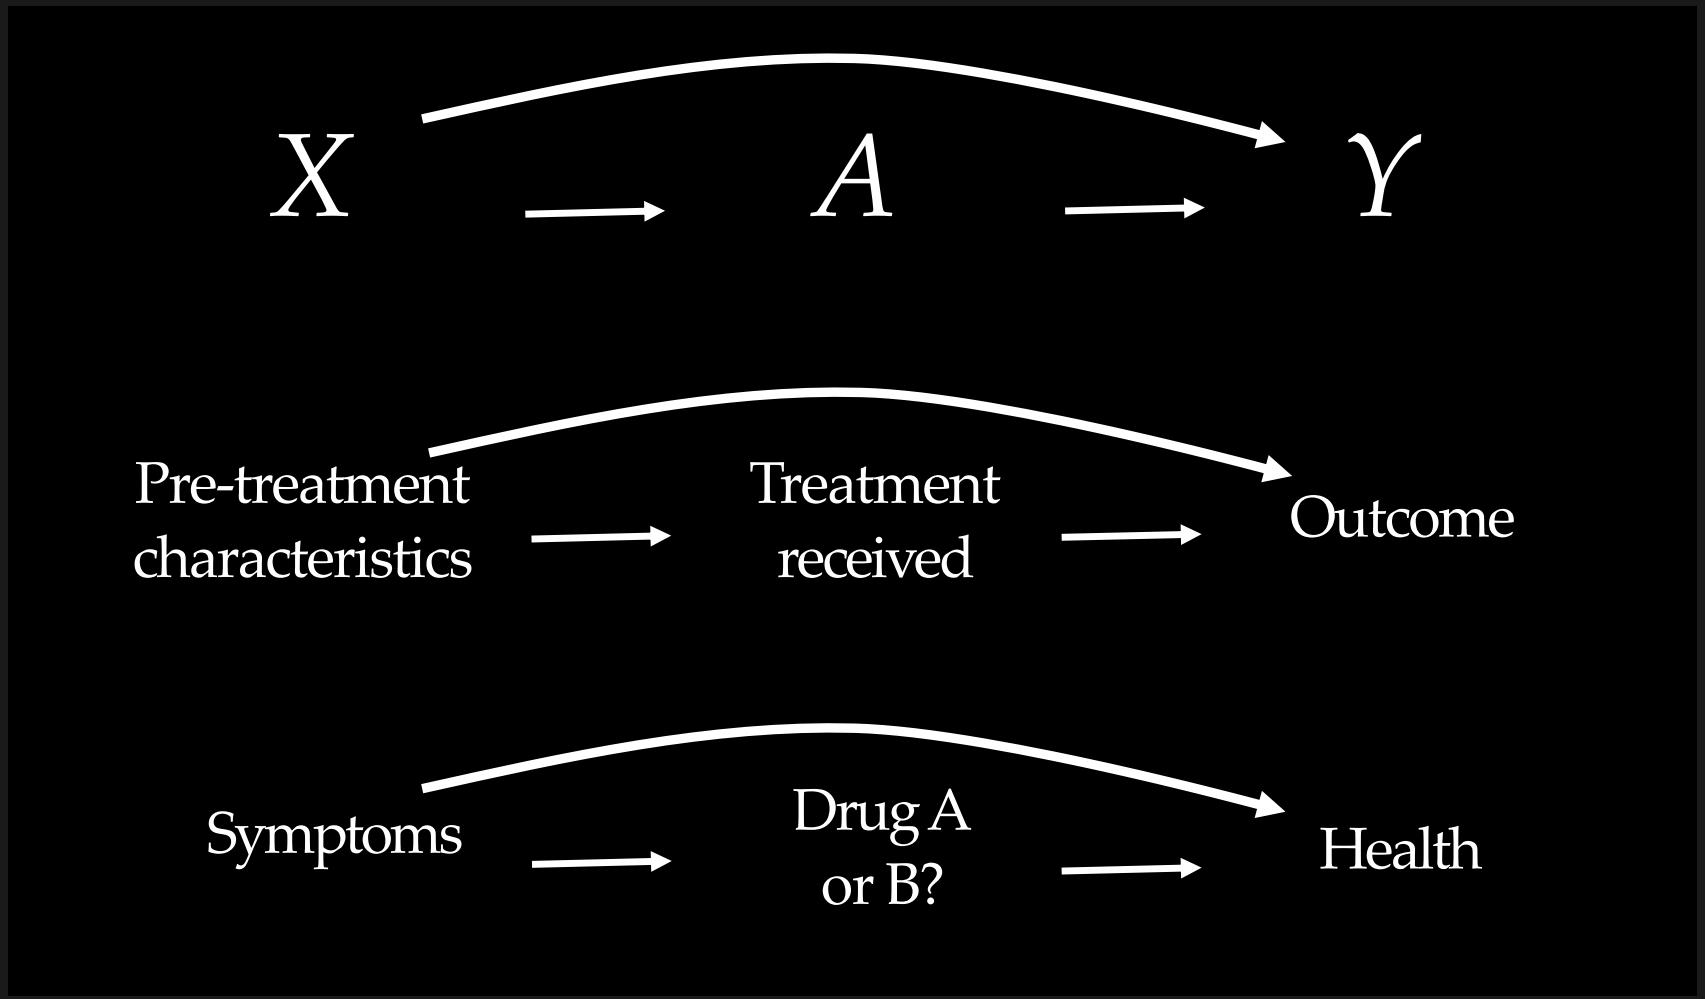

Identifying the best treatment regime: multi-stage

The multi-stage case is more complex:
Want a sequence $(a_1^{\text{opt}}, a_2^{\text{opt}}, a_3^{\text{opt}})$ that maximizes Y, but the choice of A_j affects future decisions.

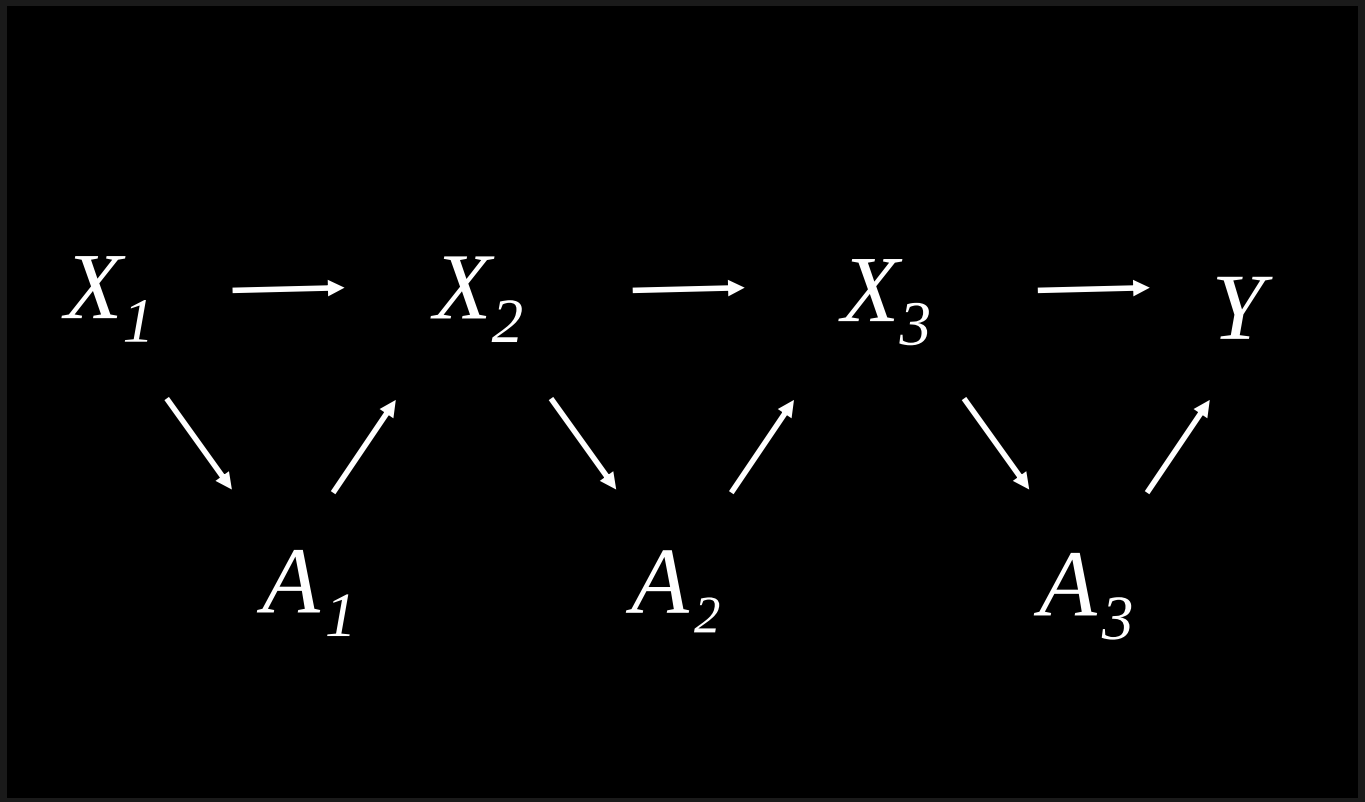

Identifying the best treatment regime: multi-stage

The multi-stage case is more complex:
Want a sequence $(a_1^{\text{opt}}, a_2^{\text{opt}}, a_3^{\text{opt}})$ that maximizes Y, but the choice of A_j affects future decisions.

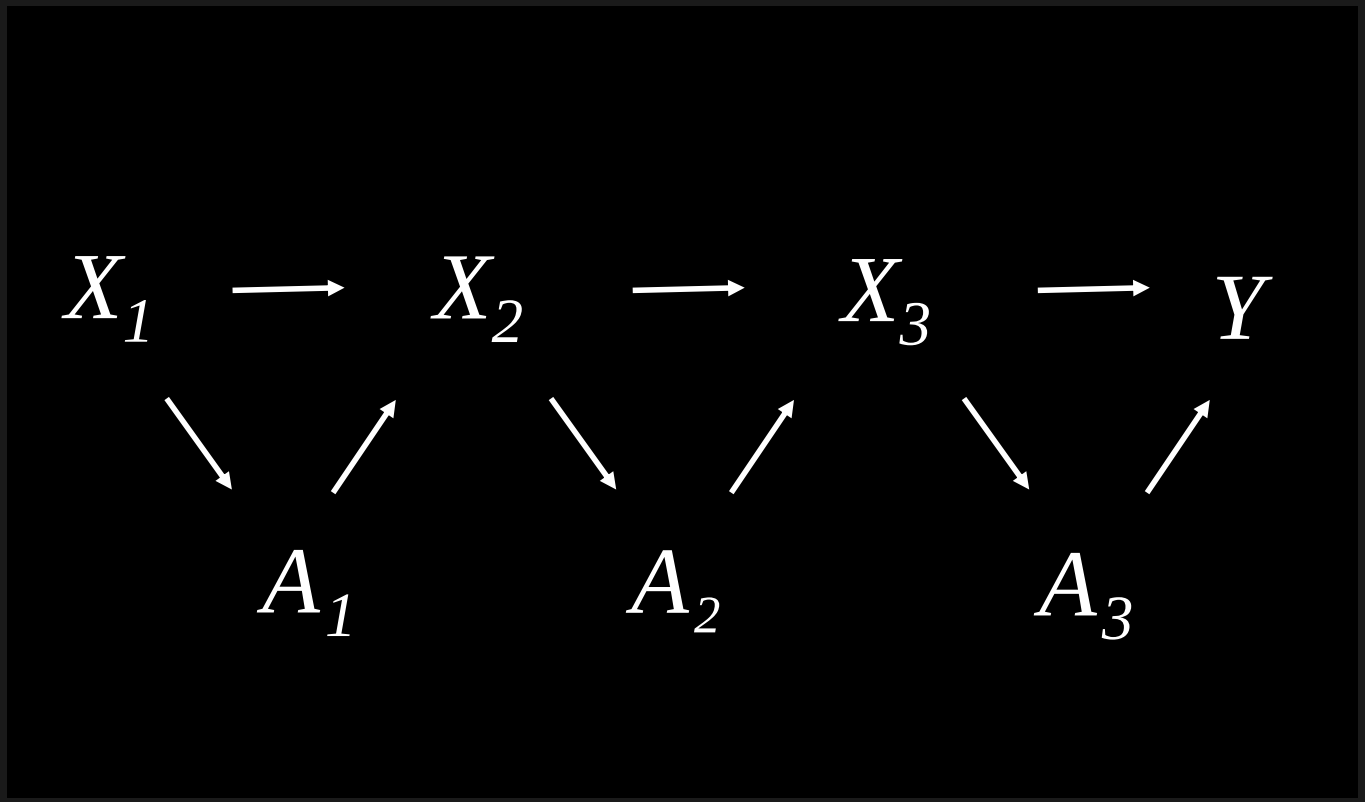

Recursive implementation reduces estimation to a series of one-stage problems.

Identifying the best treatment regime: multi-stage

Letting $H_3 = (X_1, A_1, X_2, A_2, X_3)$ reduces finding A_3^{opt} to a single-stage problem.

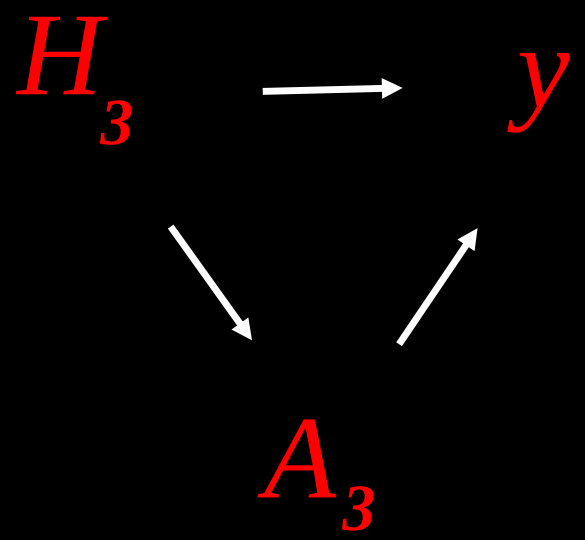

Identifying the best treatment regime: multi-stage

Writing $H_2 = (X_1, A_1, X_2)$ reduces finding A_2^{opt} to a single-stage problem, where $\tilde{Y}_2$ is taken to be the "best" possible stage 2 outcome for someone with H_2.

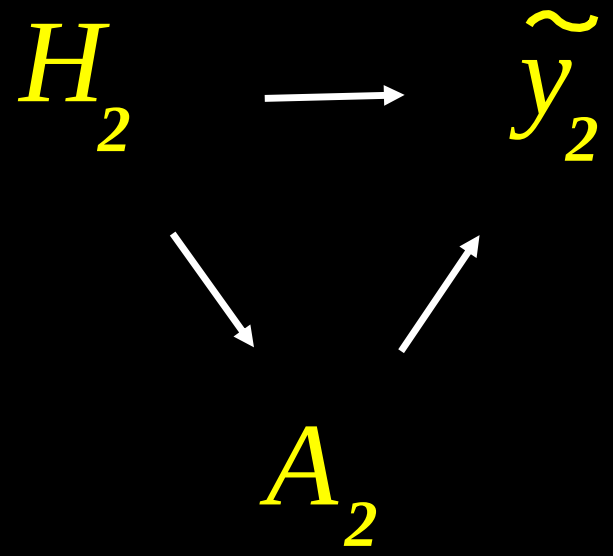

Identifying the best treatment regime: multi-stage

- The $\tilde{Y}$ term is key to the backwards induction, sequential regression approach.
- It is called a pseudo-outcome.
- Q-learning uses a different pseudo-outcome to dWOLS.
  - $\tilde{Y}$ in Q-learning is predicted at each stage using the assumed outcome model.
  - $\tilde{Y}$ in dWOLS (and G-estimation) equals the observed outcome for individuals who were treated optimally, otherwise, the observed outcome is "adjusted up" as predicted by the contrast function.

The Censored Outcome Setting

Setting (notation++)

Let's consider (at most) two stages:

- Treatment options are denoted a_k, and lie in the set $\mathcal{A}_k$
- Pre-treatment covariates x_k are measured at each stage
- In each stage, time is measured and denoted T_k
- Patients may be censored at any stage; let η_k denote whether a patient progressed to the k-th stage of intervention with $\eta_1 = 1$ for all.
- Total observation time is then $\sum_{k=1}^{2} \eta_k t_k = t_1 + \eta_2 t_2$
- Δ is an event indicator
- The totality of data for an individual is thus contained in $\{x_1, a_1, t_1, \eta_2, x_2, a_2, t_2, \Delta\}$ where x_2, a_2, t_2 are NA if $\eta_2 = 0$

dWSurv: dWOLS for censored outcomes

- dWSurv: similar robustness & ease of use as dWOLS.
- Relies on an AFT-based contrast function.
- Again, a weighted outcome regression, where weights depend *two* components: treatment (π) and censoring (ϖ) probabilities that depend on covariates, and must satisfy

$$\begin{aligned}[1-\varpi(0,\boldsymbol{x})][1-\pi(\boldsymbol{x})]w(0,0,\boldsymbol{x}) &= \varpi(0,\boldsymbol{x})[1-\pi(\boldsymbol{x})]w(0,1,\boldsymbol{x}) \\ &= [1-\varpi(1,\boldsymbol{x})]\pi(\boldsymbol{x})w(1,0,\boldsymbol{x}) = \varpi(1,\boldsymbol{x})\pi(\boldsymbol{x})w(1,1,\boldsymbol{x})\end{aligned}$$

where $\pi(\boldsymbol{x})$ is the treatment model, and $\varpi(a,\boldsymbol{x})$ a model for the probability of being censored.

dWSurv: dWOLS for censored outcomes

Steps:

1. Model treatment at each stage/decision point.
2. Model censoring (yes/no) at each stage/decision point.
3. Using log-time: weighted linear models for the event-time only in uncensored at a given stage.
4. Recursively estimate optimal rule at each stage, using pseudo-outcomes that 'extend' time according to the estimated contrast function for those not observed to have been optimally treated.

▸ (Details)

dWSurv: dWOLS for censored outcomes

- Easy to implement: requires only logistic regression or similar to model π and ϖ, and any standard survival software that can accommodate weighting.
- Double robustness requires correct model specification of either (i) the outcome model or (ii) *both* π and ϖ probability models.
- NB: dWSurv assumes that, given covariates, those who are and are not censored are exchangeable, and the timing of the censoring is irrelevant.
- In multi-stage settings requires *artificial censoring* if there is an administrative end of study (like G-estimation). $\Rightarrow$ In simulations, ignoring this has not led to notable bias.

Optimal T2D Treatment Strategies

Case study 2: Type 2 Diabetes

- Type 2 diabetes (T2D): body does not properly make or use insulin.
- Uncontrolled T2D: sugars build up in the blood rather than being used for energy.
- Complications: blindness, kidney failure, MI, stroke and amputation.

- Treatment begins with lifestyle changes, followed by metformin and additional drugs, with insulin typically the last option.

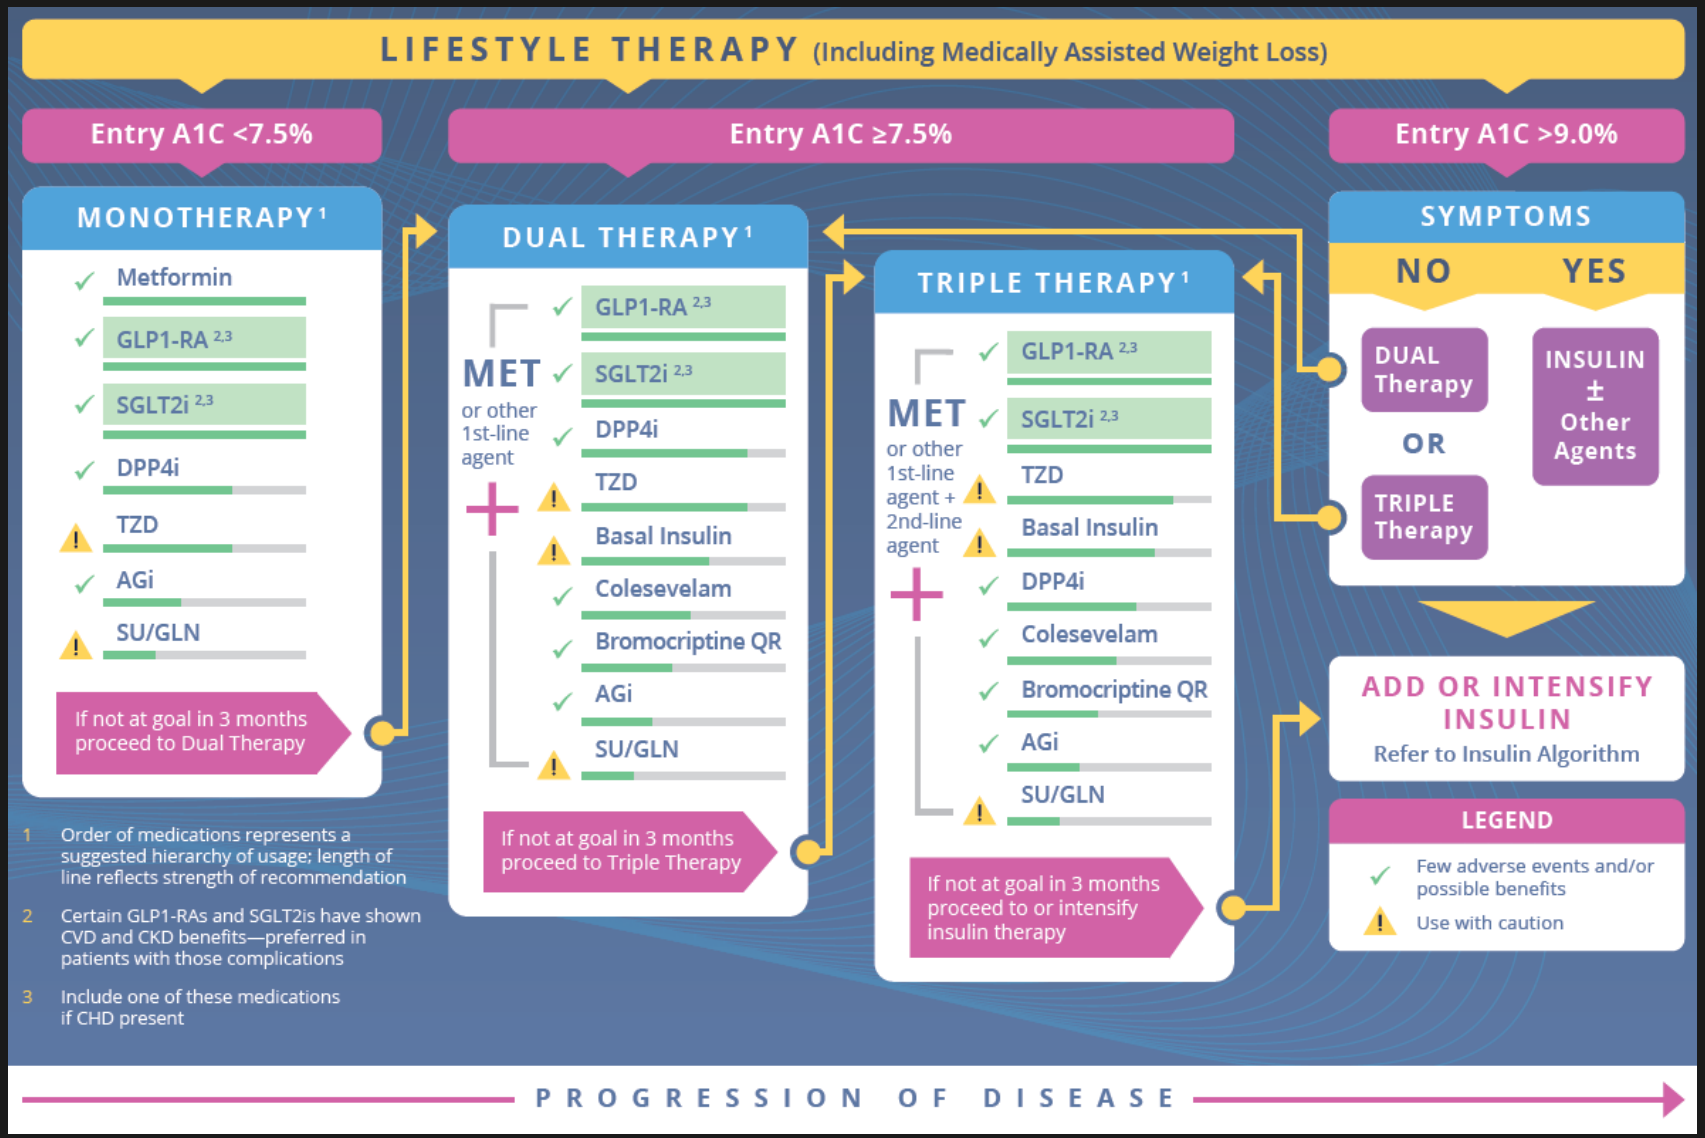
LIFESTYLE THERAPY (Including Medically Assisted Weight Loss)
Entry A1C <7.5%
Entry A1C ≥7.5%
Entry A1C >9.0%
MONOTHERAPY 1
Metformin
GLP1-RA 2,3
SGLT2i 2,3
DPP4i
TZD
AGi
SU/GLN
If not at goal in 3 months proceed to Dual Therapy
DUAL THERAPY 1
MET or other 1st-line agent +
GLP1-RA 2,3
SGLT2i 2,3
DPP4i
TZD
Basal Insulin
Colesevelam
Bromocriptine QR
AGi
SU/GLN
If not at goal in 3 months proceed to Triple Therapy
TRIPLE THERAPY 1
MET or other 1st-line agent + 2nd-line agent +
GLP1-RA 2,3
SGLT2i 2,3
TZD
Basal Insulin
DPP4i
Colesevelam
Bromocriptine QR
AGi
SU/GLN
If not at goal in 3 months proceed to or intensify insulin therapy
SYMPTOMS
NO
YES
DUAL Therapy
OR
TRIPLE Therapy
INSULIN ± Other Agents
ADD OR INTENSIFY INSULIN
Refer to Insulin Algorithm
LEGEND
Few adverse events and/or possible benefits
Use with caution
1 Order of medications represents a suggested hierarchy of usage; length of line reflects strength of recommendation
2 Certain GLP1-RAs and SGLT2is have shown CVD and CKD benefits—preferred in patients with those complications
3 Include one of these medications if CHD present
PROGRESSION OF DISEASE

Clinical Practice Research Datalink

- Primary care database in the United Kingdom
- Patient-level information on $\geq$ 13 million patients
  - prescriptions, demographics, diagnoses, laboratory data, etc.
- Hospital Episodes Statistics (HES) database from 1 April, 1997
  - hospital admission information (diagnoses, procedures)
  - Linkage possible for approx. 75% of the practices in the CPRD

Treatment of T2D

- **Wanted** to target a 2-stage regime for 2nd and 3rd line therapies to maximize the cardiovascular event-free survival:
  - For whom is sulfonylurea (sulfo) or dipeptidyl peptidase-4 inhibitors (DPP4) the best add-on treatment to metformin?
  - If on met+sulfo, should DPP4 or insulin be added? If on met+DPP4, should GLP-1 or insulin be added?

- A drug considered 'added' if a new prescription was recorded within 30 days after a prescription of metformin.
- 36,911 adults $\geq$40 from the CPRD, with a 1st Rx of metformin monotherapy 04/1997–03/2018, and $\geq$1 yr of follow-up prior to met.
- But...

Data ‘shortage’

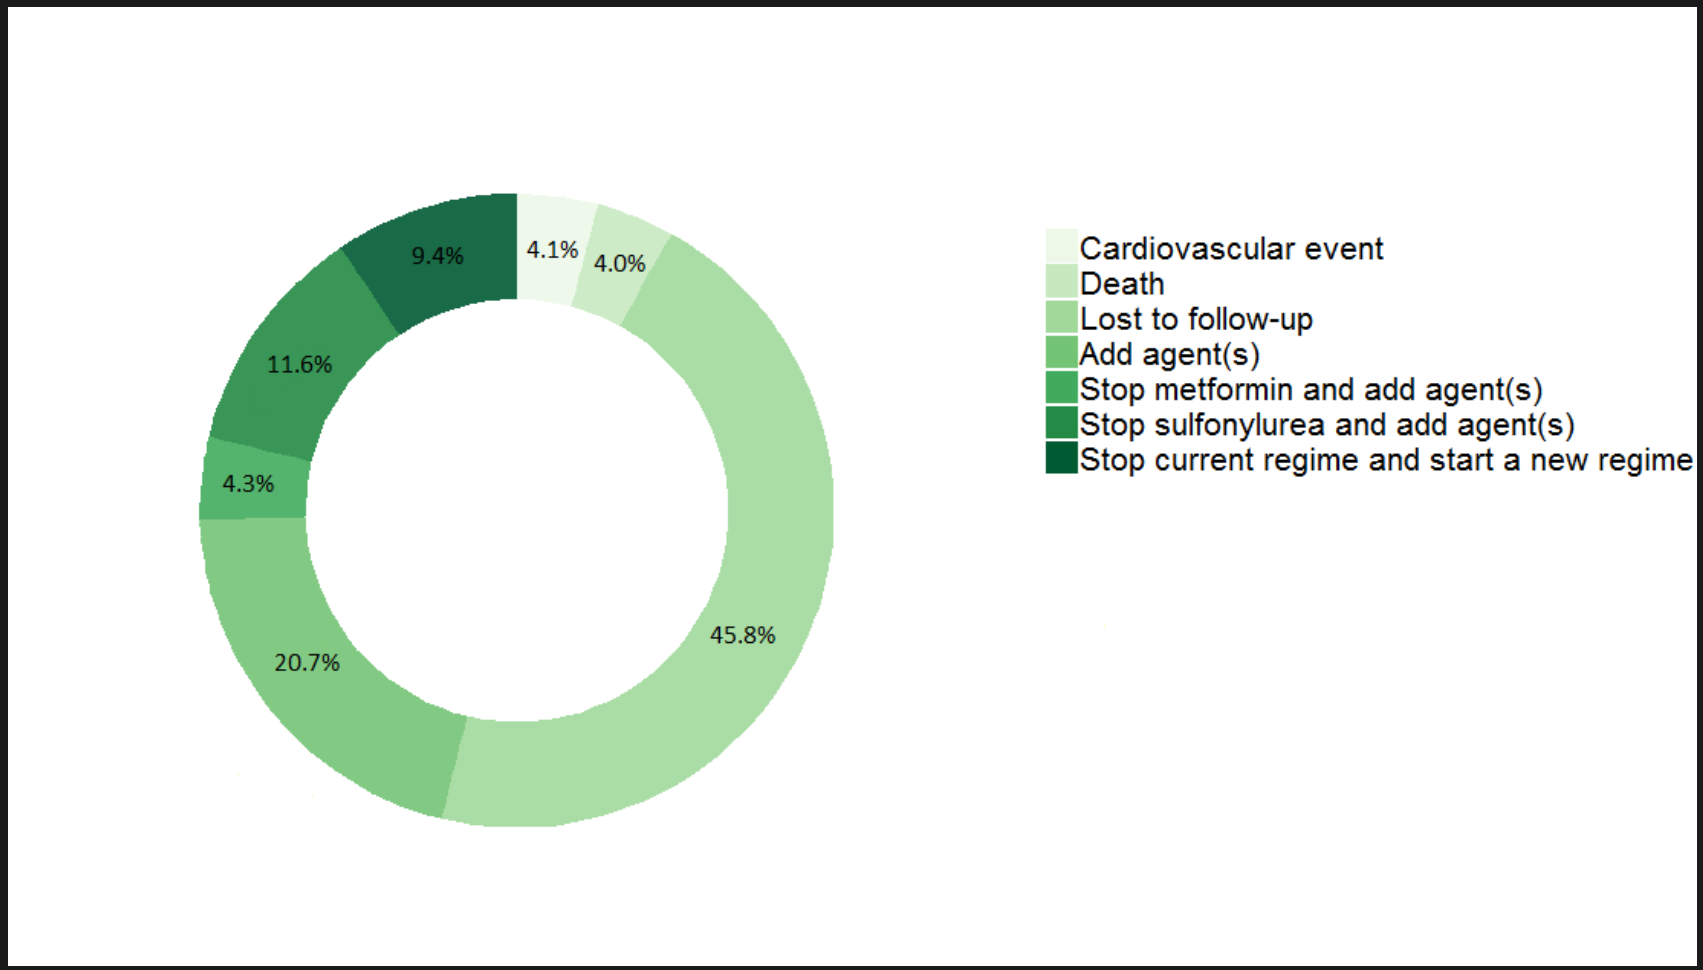

Data ‘shortage’

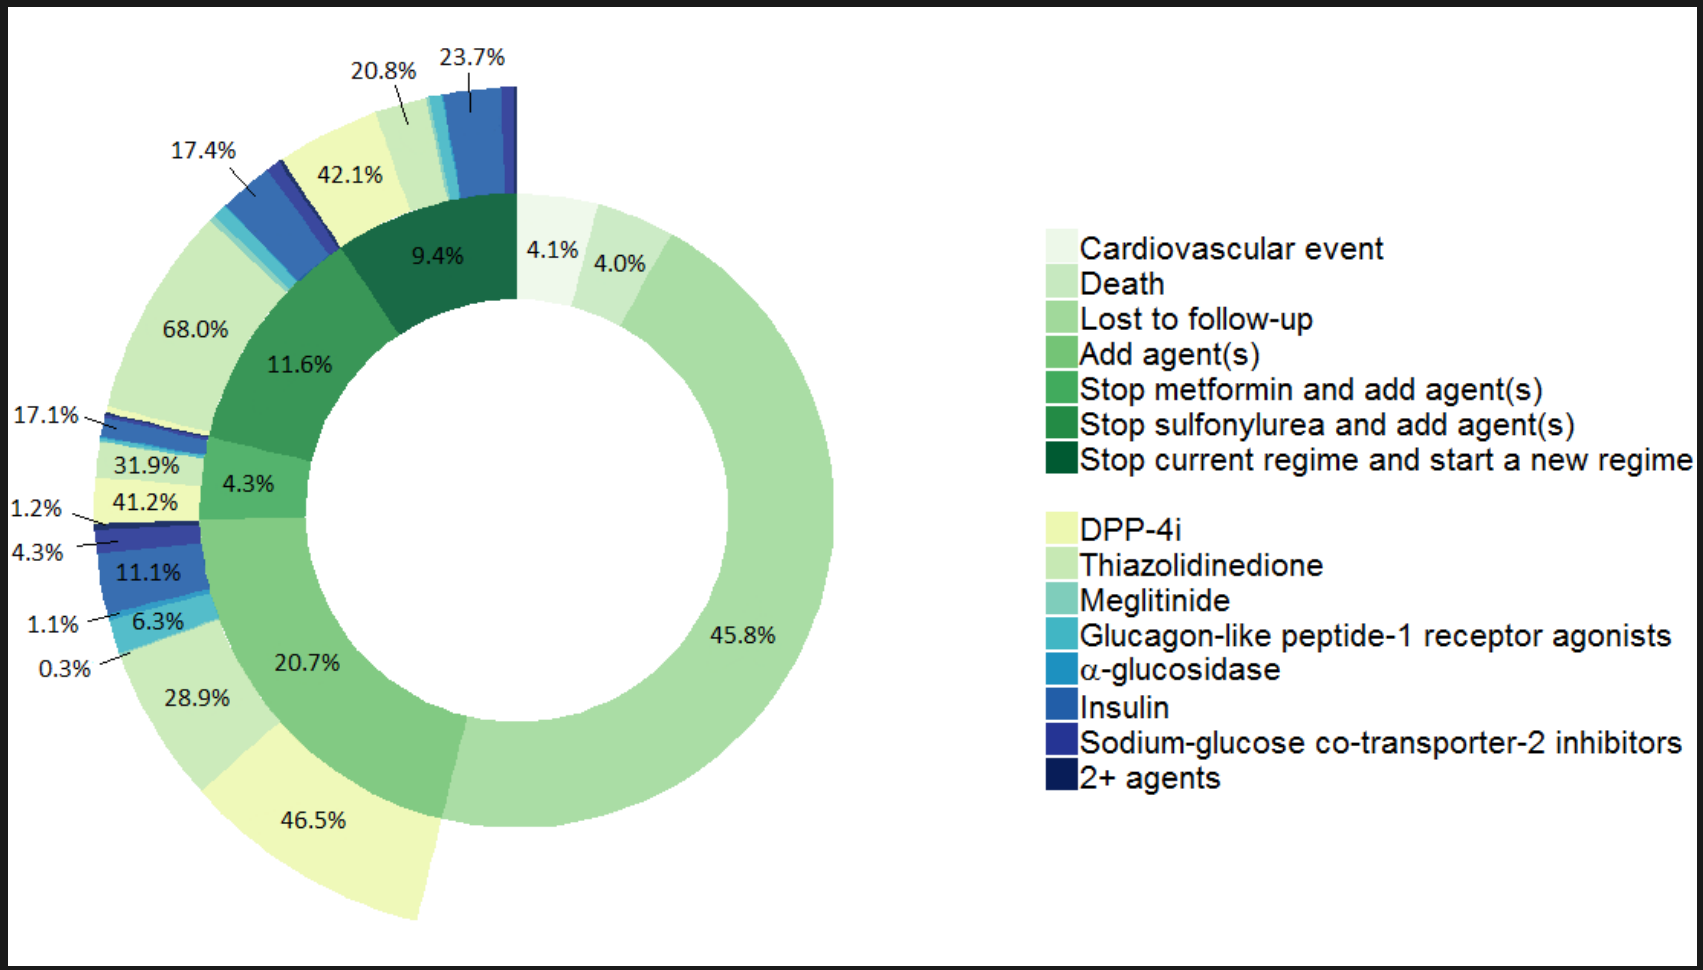

Data 'shortage'

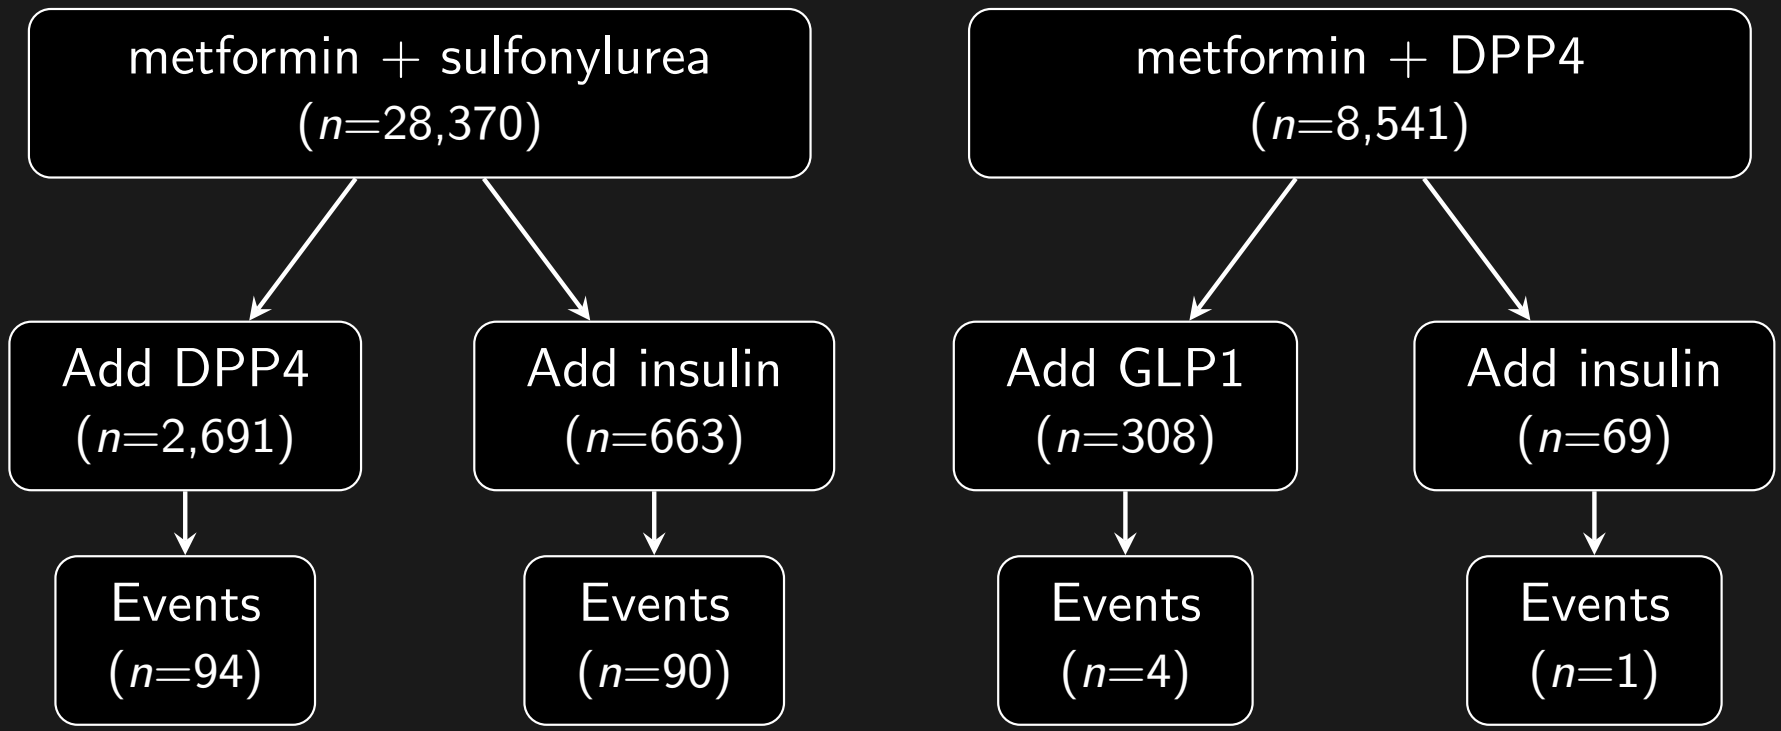

- Impossible to estimate a two-stage optimal ATS with so few events at the end of the treatment sequence.

Instead, Part I: A single decision rule

Research question:

- Can we find an individualized decision rule that chooses between adding sulfonylurea or DPP4 to metformin when metformin alone does not control the symptoms?
- Again, aim is to maximizes the time until a cardiovascular event or death.

Study cohort

- Patients aged ≥ 40 with a first-ever prescription of metformin, excluding other known indications for metformin

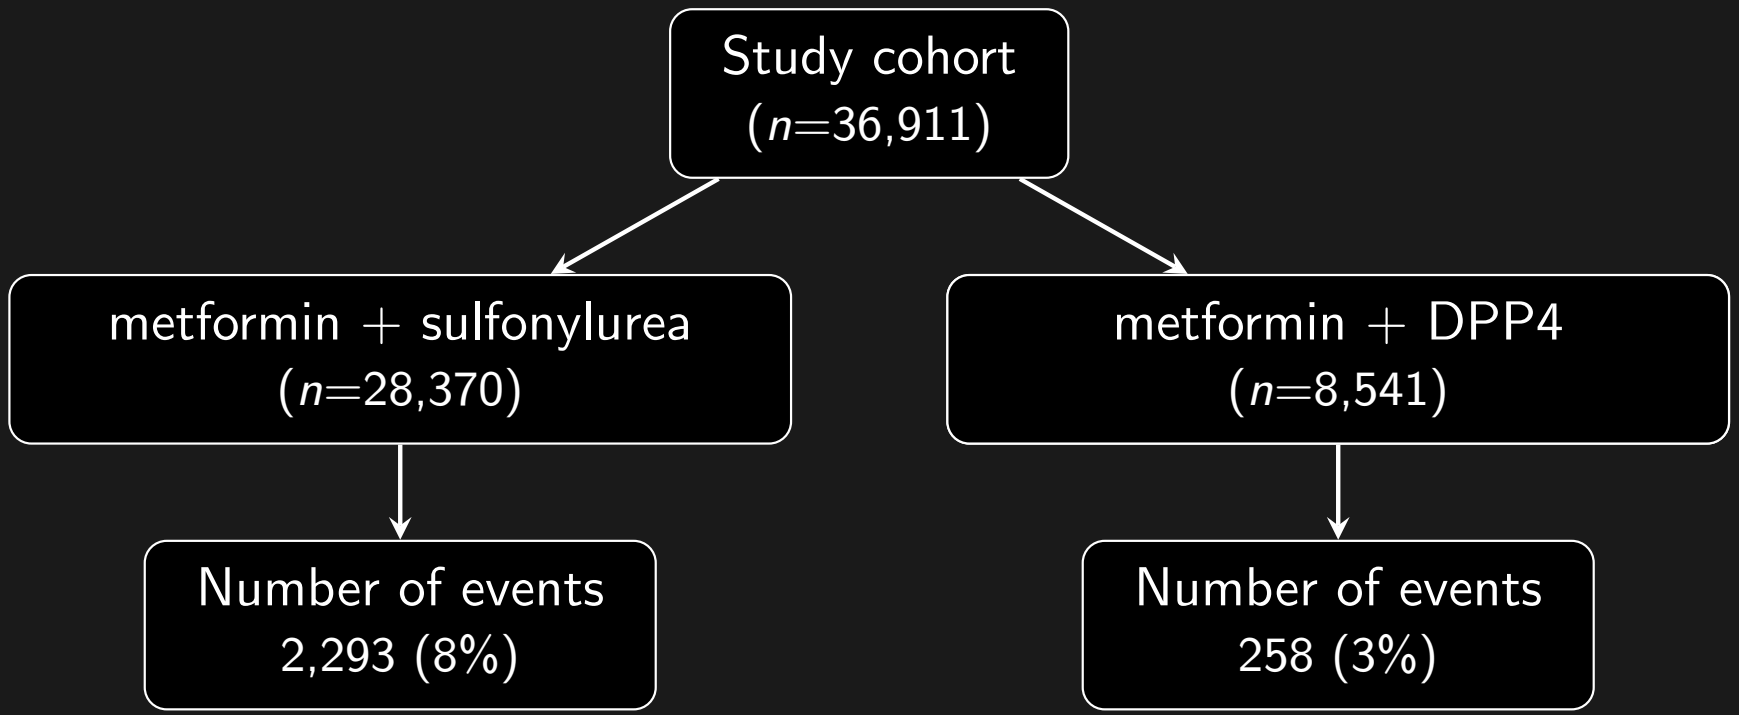

Model fitting

Tailoring variables

- glycemic control based on HbA1c (good, borderline, bad)
- history of severe hypoglycemia
- BMI

Individualized treatment rule

Recommend DPP4 if

$$\mathbb{I}\left(\psi_0 + \psi_1 \times \text{borderline} + \psi_2 \times \text{bad} + \psi_3 \times \text{hypoglycemia} + \psi_4 \times \text{BMI} > 0\right)$$

and sulfonylurea otherwise.

Estimated treatment rule

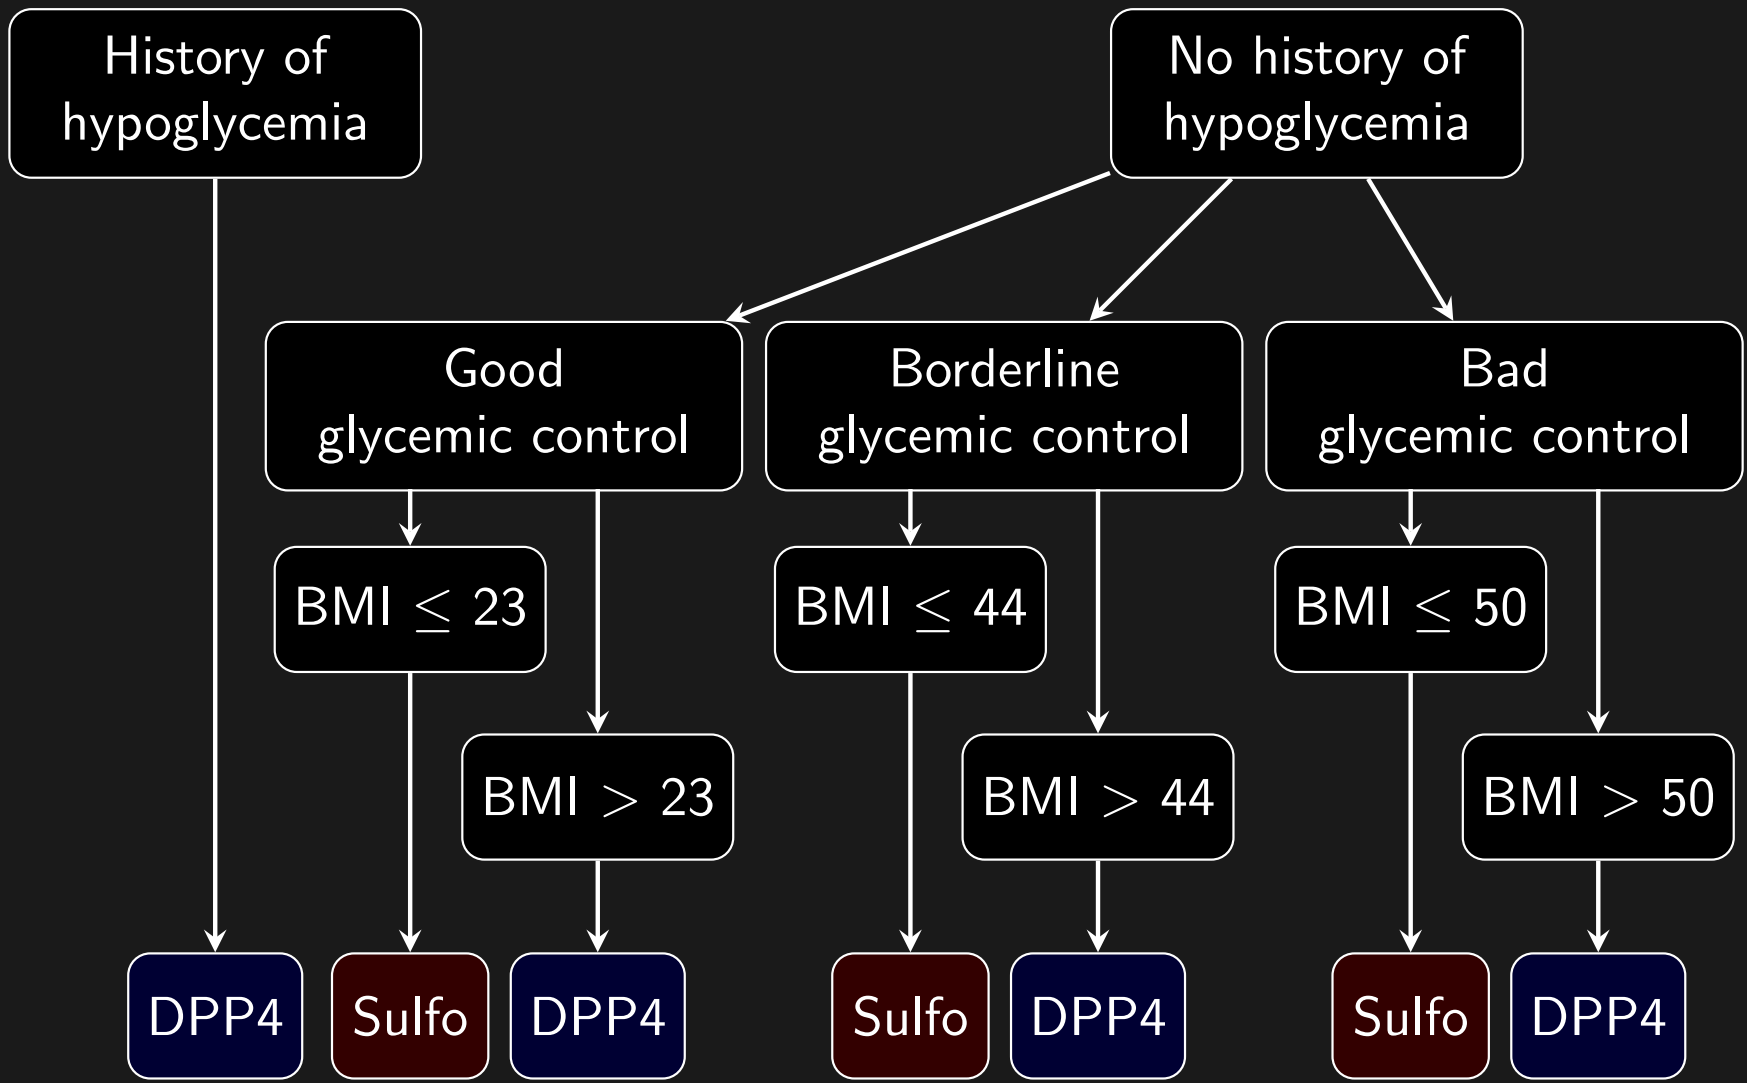

Instead, Part II: Two-stage treatment strategies

Research question:
Can we find a sequence of two adaptive decision rules:

1) when metformin fails, add/switch to sulfonylurea or *any other* drug class?

2.1) if met-sulfo dual therapy fails, add/switch to DPP4 or *any other* drug class?

2.2) if met-other drug class dual therapy fails, add/switch to sulfo or *any other* drug class?

Instead, Part II

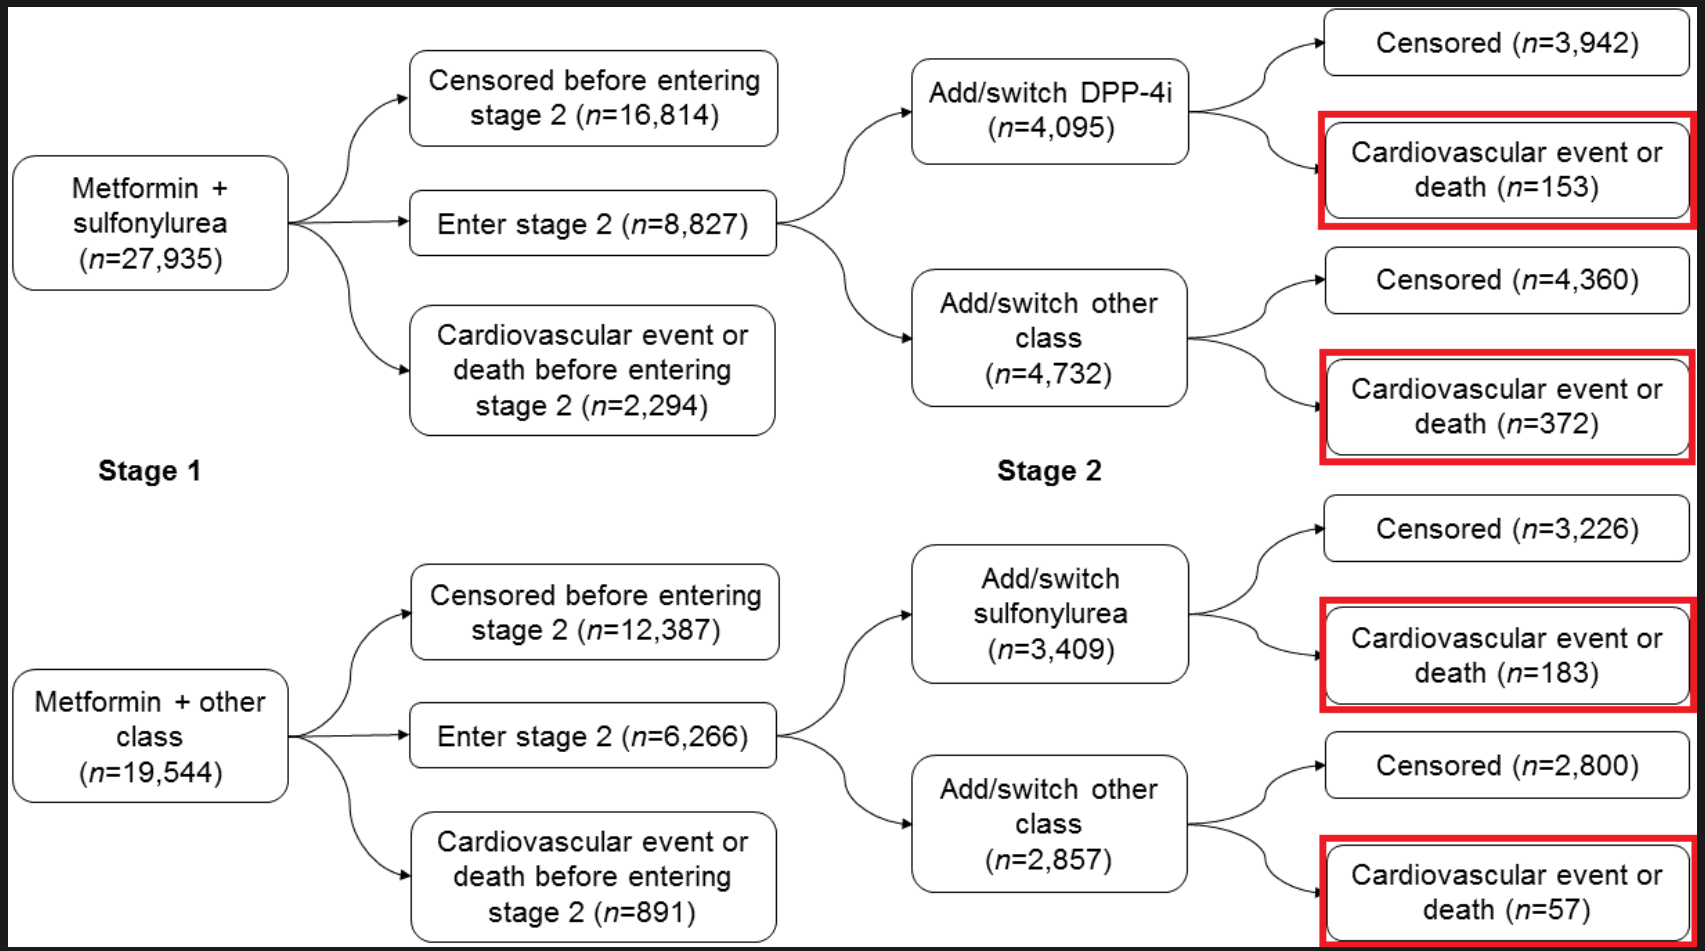

Stage 1 estimated rule

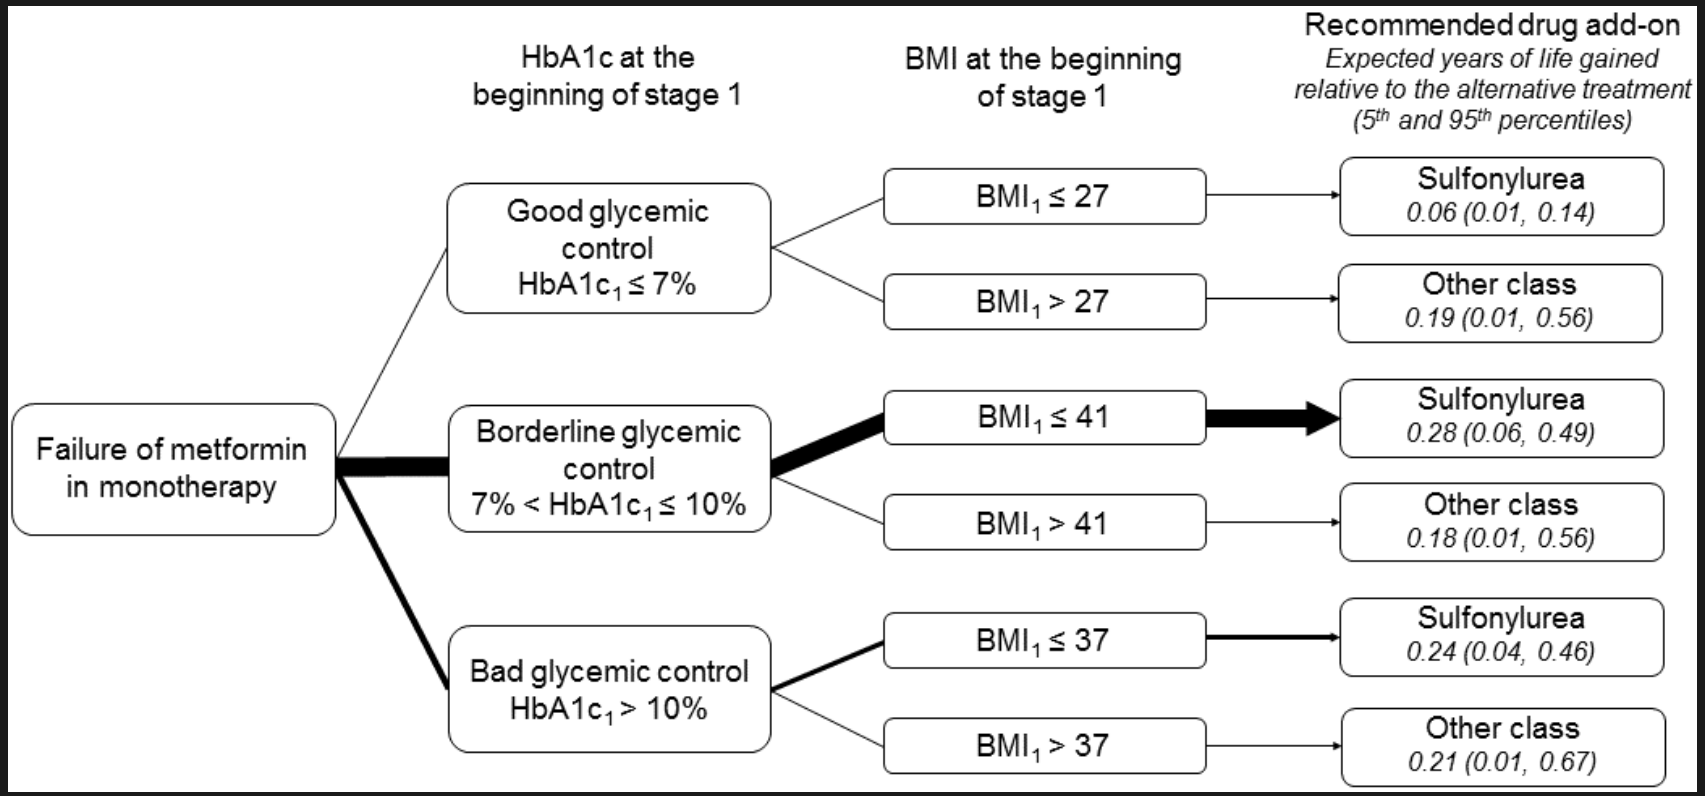

Right-most box shows average expected years of life gained assuming that, if needed, optimal estimated stage 2 treatment is taken thereafter. The width of the lines is proportional to the number of patients following each path.

Stage 2 estimated rule (Stage 1: met-sulfo)

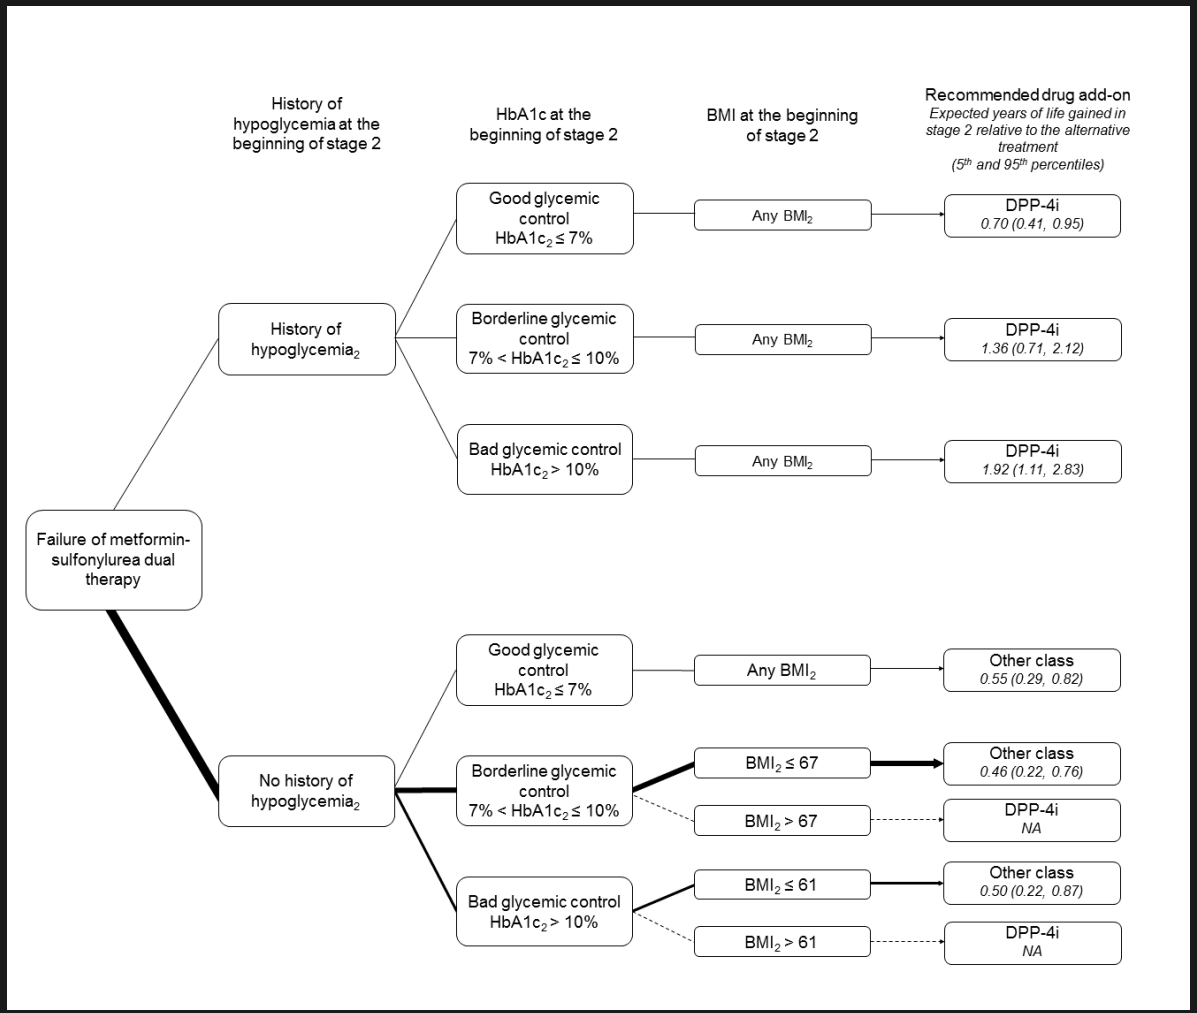

Right-most box shows average expected years of life gained. The width of the lines is proportional to the number of patients following each path. Dashed lines indicate paths with no data support.

Stage 2 estimated rule (Stage 1: met-other)

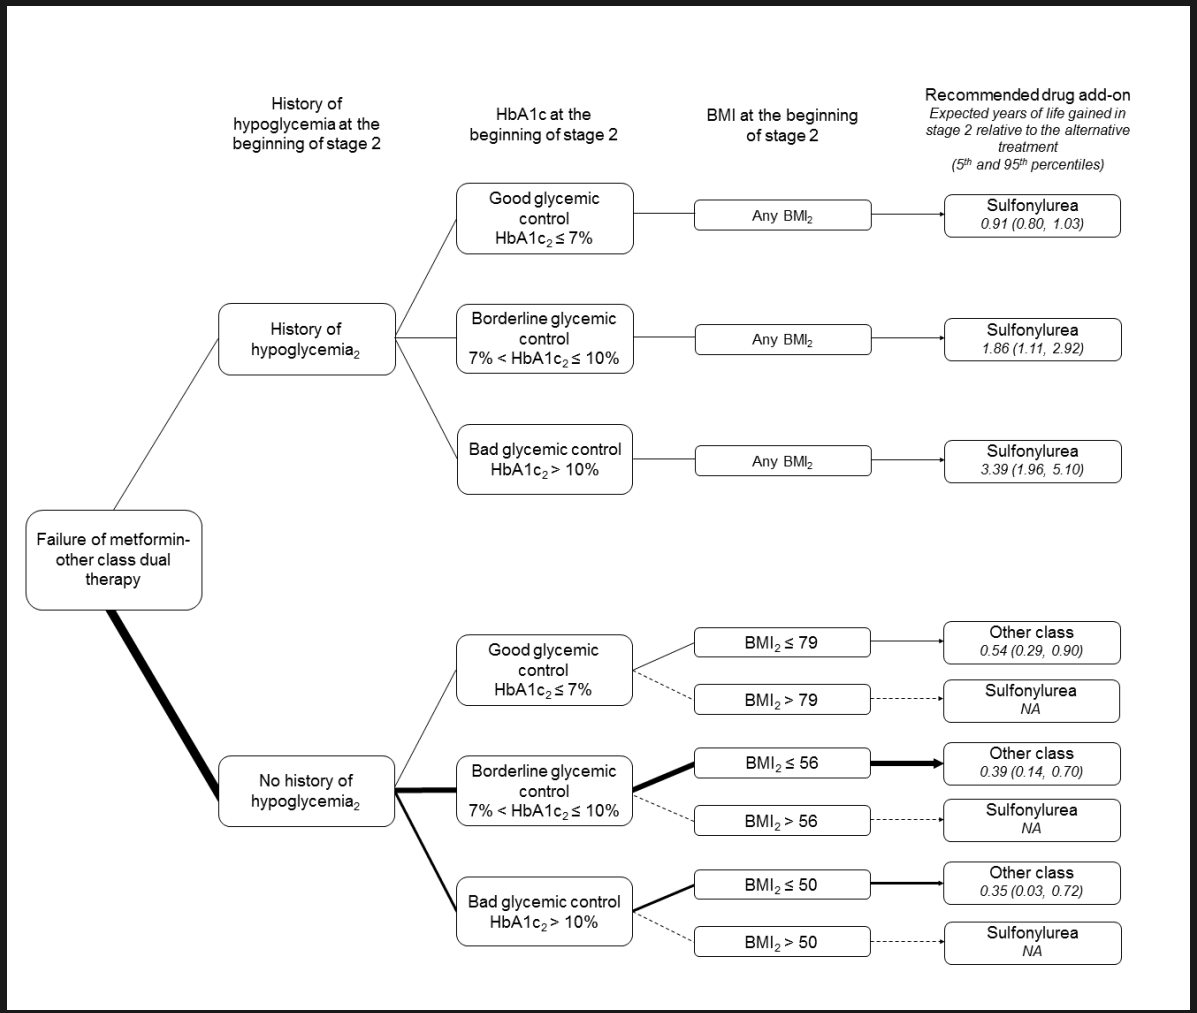

Right-most box shows average expected years of life gained. The width of the lines is proportional to the number of patients following each path. Dashed lines indicate paths with no data support.

Instead, Part II: Strengths and limitations

- By
  - considering together drug add-ons and switches, and
  - merging drug classes for comparison

  the sample size and the number of events at the end of the sequence increase.

- However,
  - makes assumptions about mechanism of action of the different drugs,
  - impact on clinical practice is modest.

Concluding remarks

- ATSs have enormous potential to improve clinical practice in an evidence-based way.
- Non-experimental data will remain a key source of information, but many challenges arise:
  - All relevant confounders?
  - Reasons for changes in treatment?
  - Reasons for a measurement at all (frequency/timing)?
  - Sample size may still be small for 'rarities'!
- Many approaches to estimation available, however ease of use remains a significant challenge for many.

Concluding remarks

- dWOLS/dwSurv quite accessible and interpretable method of analyzing data, but many challenges remain.

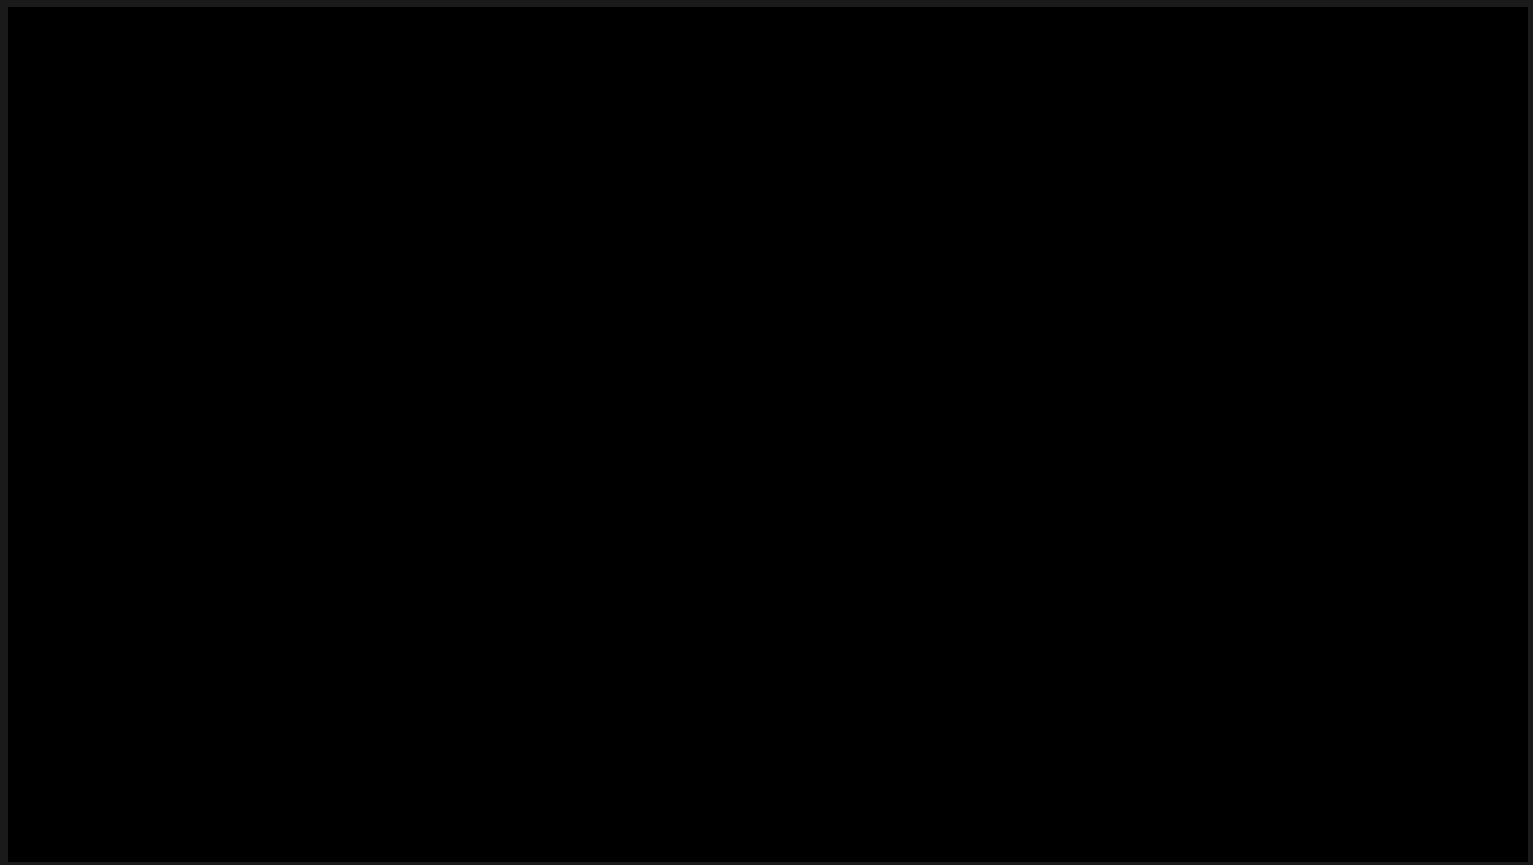

Concluding remarks

- dWOLS/dwSurv quite accessible and interpretable method of analyzing data, but many challenges remain.

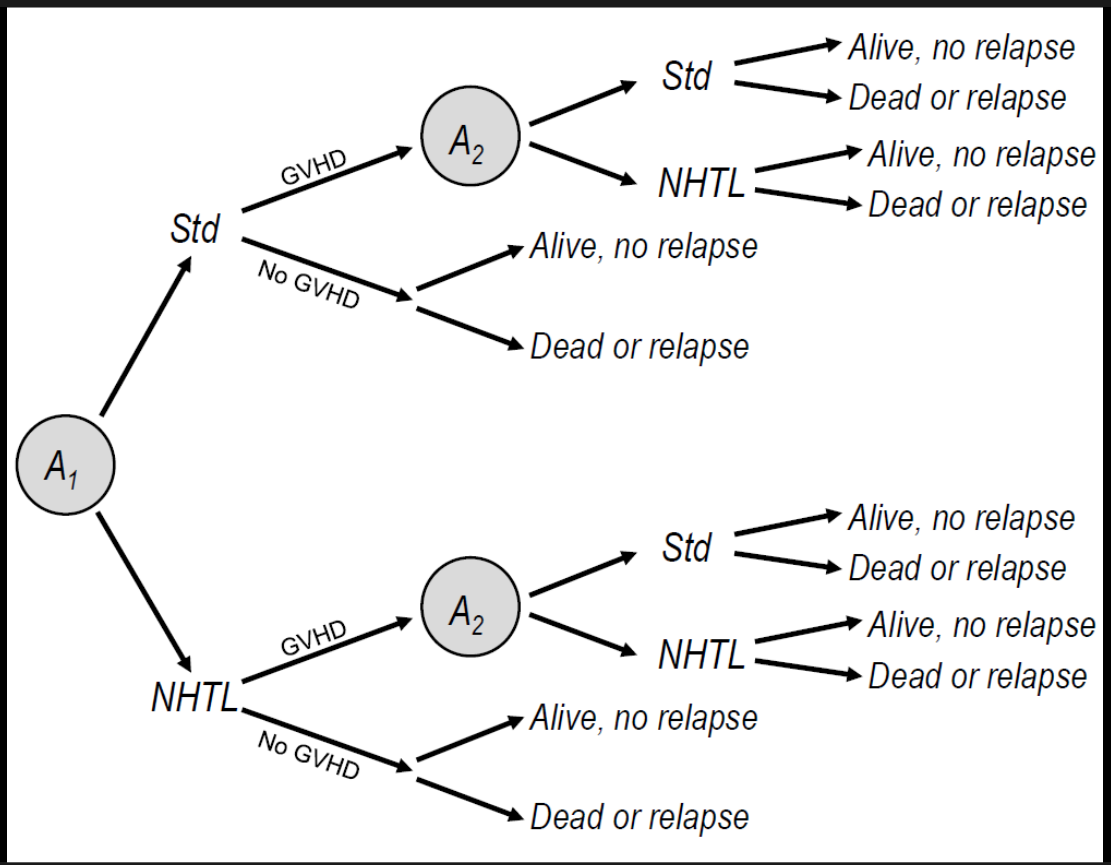

Concluding remarks

- dWOLS/dwSurv quite accessible and interpretable method of analyzing data, but many challenges remain.

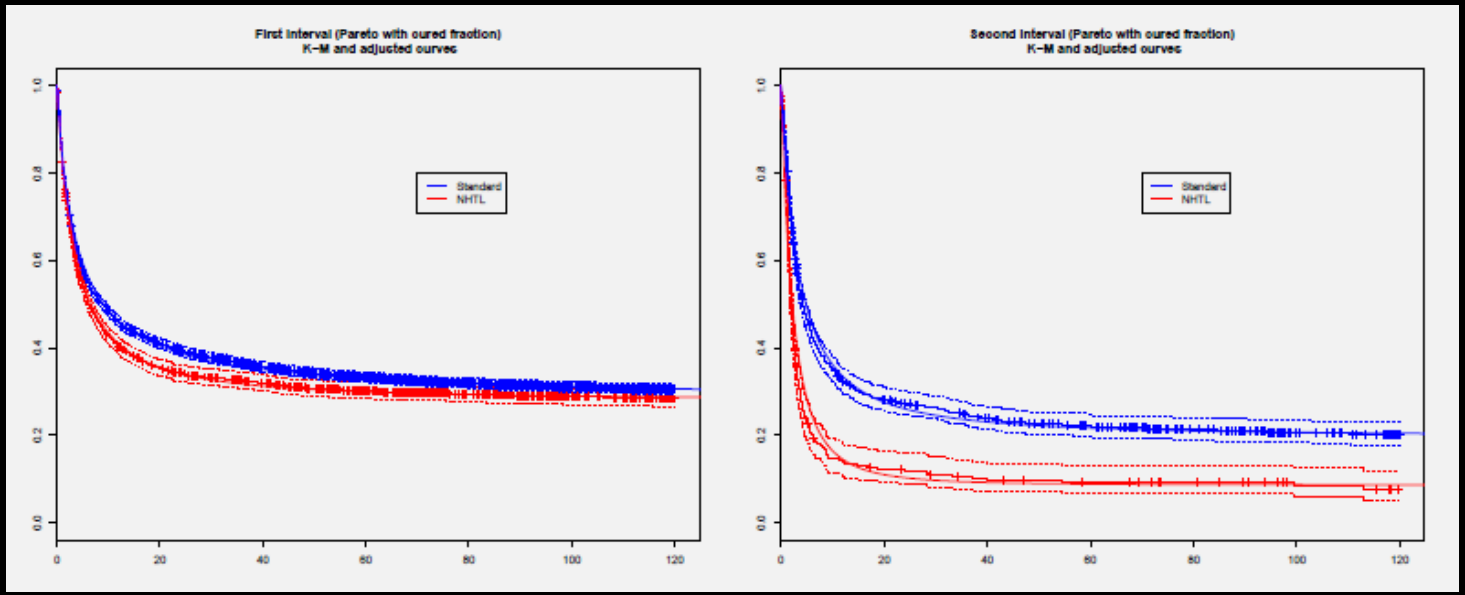

Thanks


CAN-AIM
CAN-AIM
NSERC
CRSNG
Fonds de recherche
Nature et
technologies
Québec
NSERC
CRSNG
CIHR IRSC
Canadian Institutes of Health Research
Instituts de recherche en santé du Canada
Fonds de recherche
Santé
Québec
HEALTHY BRAINS
HEALTHY LIVES