Evaluating mediator interventions for time-to-event outcomes: A causal framework for cancer disparities

Elisavet Syriopoulou

Department of Medical Epidemiology and Biostatistics, Karolinska Institutet

Email: elisavet.syriopoulou@ki.se

12th November 2025

About me

- BSc in Mathematics and MSc in Biostatistics
- Moved to the UK and worked at the University of Leicester
- Completed a PhD in Biostatistics at the University of Leicester in 2020
- In October 2020, I joined MEB as a postdoc and since 2023 I have been an Assistant Professor
- Research focuses on developing and applying statistical methods to cancer registry data, with an emphasis on understanding disparities in cancer survival outcomes

Motivation - Example, Sweden

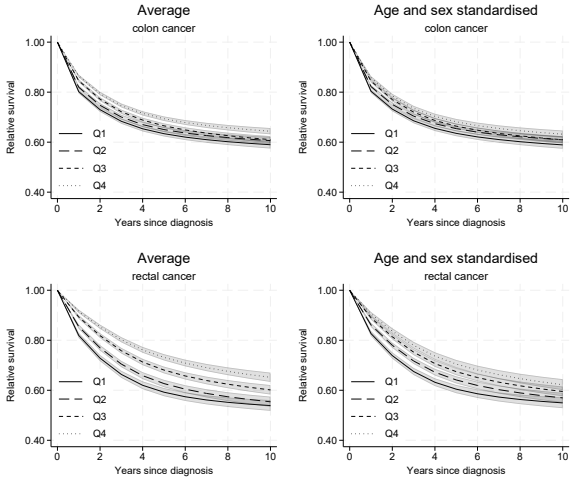

Syriopoulou E, Osterman E, Miething A, Nordenvall C, Andersson TM-L. Income disparities in loss in life expectancy after colon and rectal cancers: a Swedish register-based study. *J Epidemiol Community Health* 2024;78(6): 402–408.

Investigating other factors

Could differences in stage at diagnosis partly explain the survival differences between the least and most deprived groups?

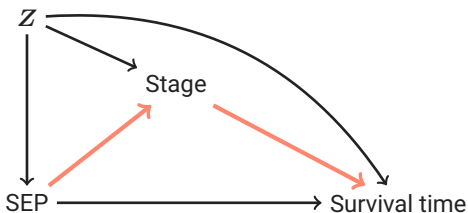

This is a mediation analysis question!

Mediation analysis

Mediation analysis allows to explore the role of a mediator on an observed association between an exposure - outcome of interest.

Mediation analysis

Mediation analysis allows to explore the role of a mediator on an observed association between an exposure - outcome of interest.

When working with cancer, we still need to deal with the complex mechanisms that contribute to cancer disparities:

- Cancer-related factors
- Other cause factors

Relative survival framework

Relative survival (*net* survival) allows comparisons of cancer survival between populations without distortions from differential background mortality.

We have adapted a formal causal framework to the settings of cancer registry-based epidemiology, extending mediation analysis methods to the relative survival framework[1]

- Main idea: using the relative survival framework allows to isolate cancer-related factors.

[1] Syriopoulou E, Rutherford MJ, Lambert PC. Understanding disparities in cancer prognosis: An extension of mediation analysis to the relative survival framework. *Biometrical Journal*. 2021; 63: 341–353.

Excess mortality rate and relative survival

Excess mortality rate

$$\begin{matrix}\text{excess}\\ \text{mortality}\end{matrix} = \begin{matrix}\text{all-cause}\\ \text{mortality}\end{matrix} - \begin{matrix}\text{expected}\\ \text{mortality}\end{matrix}$$

The survival analog of excess mortality is relative survival.

Relative survival

$$\text{relative survival} = \frac{\text{all-cause survival}}{\text{expected survival}}$$

$$R(t) = \frac{S(t)}{S^*(t)} \qquad \text{and} \qquad S(t) = S^*(t)R(t)$$

Notation

- X denotes the exposure of interest (here SEP), with $x = 1$ if exposed and $x = 0$ if unexposed
- M denotes the mediator of interest (here stage)
- $\boldsymbol{Z}$ denotes the confounding variables

Notation - potential outcomes

The framework used for formulating the effects of interest is that of potential outcomes: *the outcomes that would be observed if we intervene on X and M to set them on specific values*.

Notation - potential outcomes

The framework used for formulating the effects of interest is that of potential outcomes: *the outcomes that would be observed if we intervene on X and M to set them on specific values*.

- $R^x(t)$ is the value that $R(t)$ would have if we intervened on X and set it (possibly counter to fact) to the value x
- M^x is the value that M would take if we intervened on X and set it to x
- $R^{x,M^{x^*}}(t)$ is the the value that $R(t)$ would take if we intervened on X and set it to x and simultaneously intervened on M and set it to M^{x^*}, where x and x^* are not necessarily the same

Notation - potential outcomes

The framework used for formulating the effects of interest is that of potential outcomes: *the outcomes that would be observed if we intervene on X and M to set them on specific values*.

- $R^x(t)$ is the value that $R(t)$ would have if we intervened on X and set it (possibly counter to fact) to the value x
- M^x is the value that M would take if we intervened on X and set it to x
- $R^{x,M^{x^*}}(t)$ is the the value that $R(t)$ would take if we intervened on X and set it to x and simultaneously intervened on M and set it to M^{x^*}, where x and x^* are not necessarily the same

We will use those to define contrasts of marginal effects of the potential outcomes.

Partitioning the total survival difference

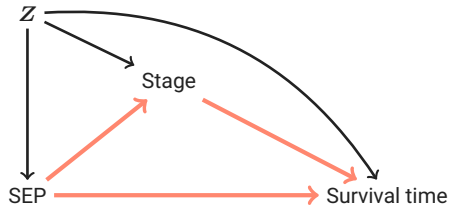

Partitioning the total survival difference

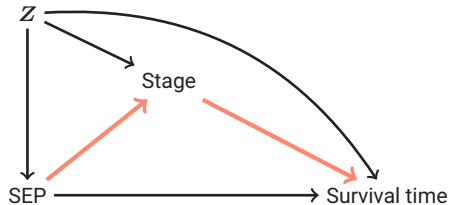

Natural indirect effect: quantifies how much of the observed difference is due to stage differences in the two groups

$$\text{NIE}(t) = R^{1,M^1}(t) - R^{1,M^0}(t) = E[R(t|X=1,\boldsymbol{Z},M^1)] - E[R(t|X=1,\boldsymbol{Z},M^0)]$$

Partitioning the total survival difference

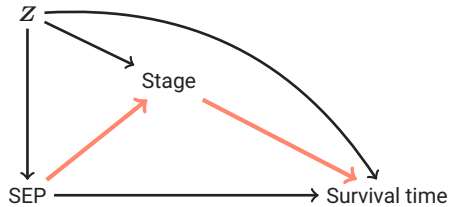

Natural indirect effect: quantifies how much of the observed difference is due to stage differences in the two groups

$$\text{NIE}(t) = R^{1,M^1}(t) - R^{1,M^0}(t) = E[R(t|X=1,\boldsymbol{Z},M^1)] - E[R(t|X=1,\boldsymbol{Z},M^0)]$$

This is a comparison of two hypothetical quantities:

- In the first, M is set to M^1 and in the second M is set to M^0, for everyone. In both, X is set to 1
- X influence $R(t)$ only through its influence on M
- Thus corresponds to an indirect effect through M

Partitioning the total survival difference

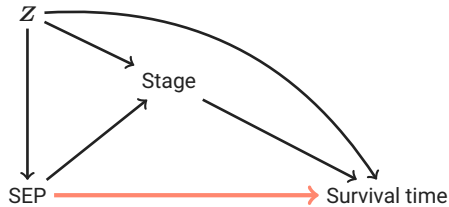

Natural indirect effect: quantifies how much of the observed difference is due to stage differences in the two groups

$$\text{NIE}(t) = R^{1,M^1}(t) - R^{1,M^0}(t) = E[R(t|X=1,\boldsymbol{Z},M^1)] - E[R(t|X=1,\boldsymbol{Z},M^0)]$$

Natural direct effect: quantifies the differences in relative survival that are *not* due to stage differences

$$\text{NDE}(t) = R^{1,M^0}(t) - R^{0,M^0}(t) = E[R(t|X=1,\boldsymbol{Z},M^0)] - E[R(t|X=0,\boldsymbol{Z},M^0)]$$

Comment

The definitions involve the following term:

$$R(t|X = 1, \boldsymbol{Z}, M^0)$$

i.e., the relative survival if setting the exposure to the level of the exposed and the mediator to the mediator value if unexposed!

An additional model is required for the mediator.

Assumptions for identification

To link the hypothetical quantities with the observed data, we need to assume no interference, consistency and conditional exchangeability.

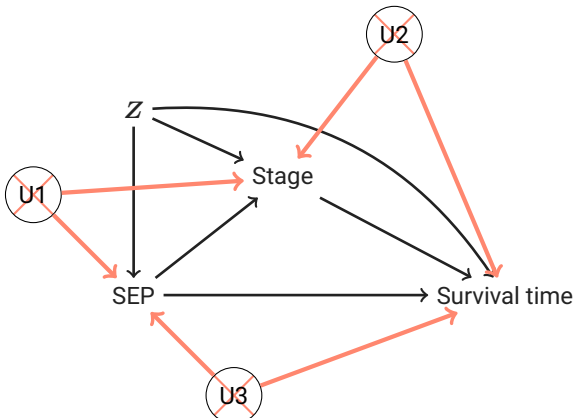

Assumptions for identification

We also need to assume that there are no intermediate confounders M.

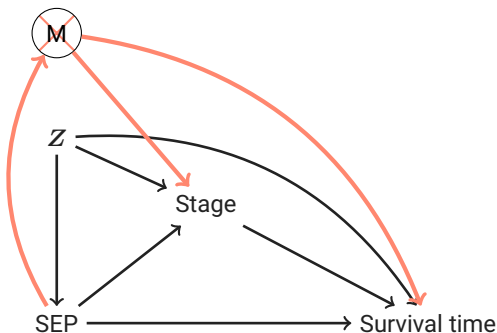

Estimation - regression standardisation

$$\widehat{NIE}(t) = \frac{1}{N}\sum_{i=1}^{N}\sum_{m}\hat{R}(t|X=1,\boldsymbol{Z},M=m)\hat{P}(M=m|X=1,\boldsymbol{Z}) - \frac{1}{N}\sum_{i=1}^{N}\sum_{m}\hat{R}(t|X=1,\boldsymbol{Z},M=m)\hat{P}(M=m|X=0,\boldsymbol{Z})$$

$$\widehat{NDE}(t) = \frac{1}{N}\sum_{i=1}^{N}\sum_{m}\hat{R}(t|X=1,\boldsymbol{Z},M=m)\hat{P}(M=m|X=0,\boldsymbol{Z}) - \frac{1}{N}\sum_{i=1}^{N}\sum_{m}\hat{R}(t|X=0,\boldsymbol{Z},M=m)\hat{P}(M=m|X=0,\boldsymbol{Z})$$

Identification algorithm

1. Fit a relative survival model including X, M, $\boldsymbol{Z}$.
2. Fit model for the mediator including X, $\mathbf{Z}$.
3. For each individual in the study population obtain estimates for $\widehat{P}(M = m|X = x, \boldsymbol{Z} = \boldsymbol{z_i})$, at each $X = x$.
4. Obtain estimates of standardised $\widehat{R}(t|X = x, \boldsymbol{Z} = \boldsymbol{z_i}, M = m)$ at $X = x$, using the predictions of Step 3 as weights. Form contrasts to obtain the $\widehat{NDE_{RS}}$ and $\widehat{NIE_{RS}}$.

Identification algorithm

1. Fit a relative survival model including X, M, $\boldsymbol{Z}$.
2. Fit model for the mediator including X, $\mathbf{Z}$.
3. For each individual in the study population obtain estimates for $\widehat{P}(M = m|X = x, \boldsymbol{Z} = \boldsymbol{z_i})$, at each $X = x$.
4. Obtain estimates of standardised $\widehat{R}(t|X = x, \boldsymbol{Z} = \boldsymbol{z_i}, M = m)$ at $X = x$, using the predictions of Step 3 as weights. Form contrasts to obtain the $\widehat{NDE_{RS}}$ and $\widehat{NIE_{RS}}$.
5. Repeat from Step 3, k times, while performing parametric bootstrap for the parameter estimates for both models.
6. Calculate 95% confidence intervals either by taking the 2.5% and 97.5% percentiles of the estimates across the bootstrapped samples or by using their standard deviation.

Fitting flexible parametric survival models

- Various modelling approaches are available for estimating relative survival
- This paper focuses on flexible parametric survival models (Royston–Parmar), which use restricted cubic splines to model the baseline log cumulative excess hazard [2]
- Non-proportional excess hazards and interactions can easily be incorporated in the model
- We model continuous survival functions rather than discrete time-point outcomes

[2] Nelson CP, Lambert PC, Squire IB, Jones DR. Flexible parametric models for relative survival, with application in coronary heart disease. Stat Med 2007.

Estimation in Stata

You can find an example of this on GitHub:
`https://github.com/syriop-elisa/mediation-example-stpm3`

Example - simulated colon cancer data

I will use an example of simulated colon cancer data that is available with the methodological paper.

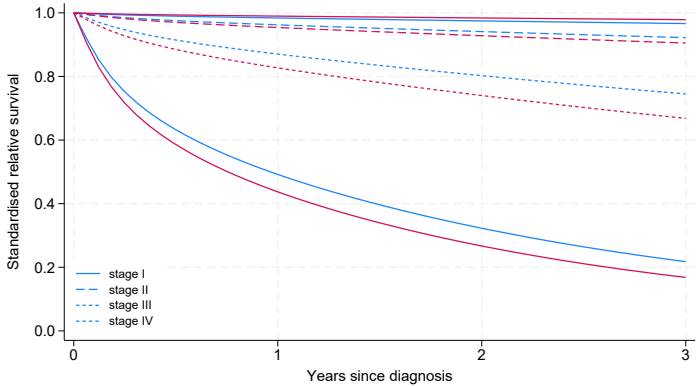

Example - partitioning the total causal effect

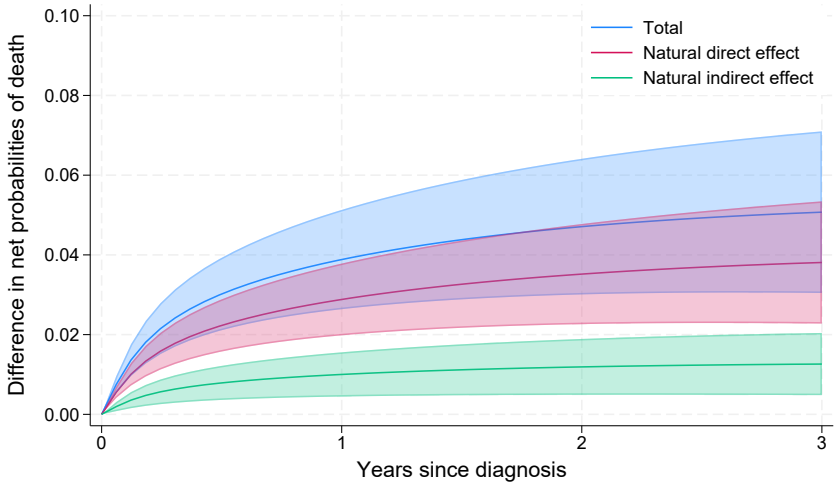

The effect among the exposed

- We can also define contrasts within specific subsets such as the NDE and NIE among the exposed (e.g. among the most deprived) which is identifiable by

$$NIE_{exposed} = E\left[R(t|X=1, \boldsymbol{Z_{sub}}, M^1)\right] - E\left[R(t|X=1, \boldsymbol{Z_{sub}}, M^0)\right]$$

- For the estimation, we restrict standardisation to a specific subset of the population

Moving to a real-world setting

- So far, we only talked about a net setting that requires elimination of competing events and looked at
  - differences in relative survival
- Differences can also be quantified in a real-word setting where no elimination of competing events is required
  - difference in all-cause survival
  - avoidable deaths
- To do so, we need to incorporate the expected mortality rates, $S^*(t)$, in the contrast of interest [Recall that $S(t) = S^*(t)R(t)$]
  - There are many different ways to do it

Direct & indirect effects - all cause setting

- Use the observed distribution of the exposure for $S^*(t)$:

$$NIE_{AC} = E\left[\boldsymbol{S^*(t|X, Z'_{sub})}R(t|X=1, \boldsymbol{Z_{sub}}, M^1)\right] - E\left[\boldsymbol{S^*(t|X, Z'_{sub})}R(t|X=1, \boldsymbol{Z_{sub}}, M^0)\right]$$

Differences in all-cause survival can only be due to the cancer of interest (not due to other-cause differences).

- The difference in all-cause survival if we could intervene and shift the stage distribution of the lowest SEP group to that of the highest SEP group, while keeping their background mortality unchanged?

A conceptually similar measure has recently been proposed for the standard survival setting, so called separable effects[3]

[3] Stensrud et al. J Am Stat Assoc 2022.

Moving to more realistic settings

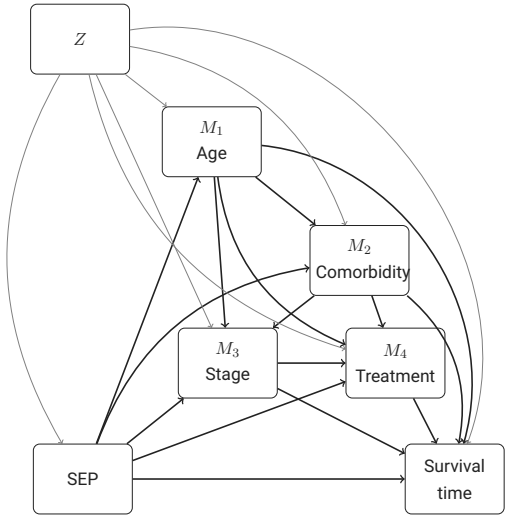

Moving to more realistic settings

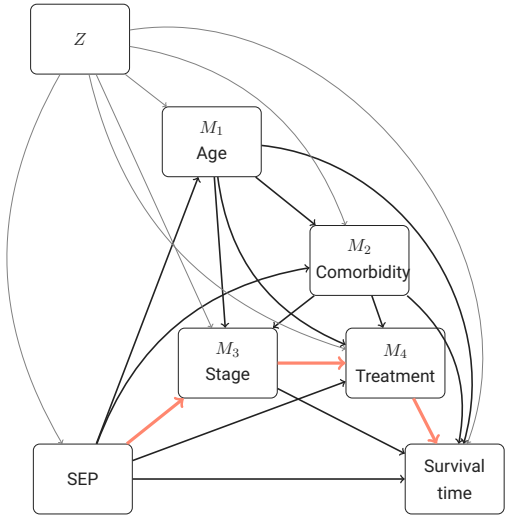

Interventional effects

Randomised interventional analogues of the NDE and the NIE have been proposed as alternatives that can be identified under weaker assumptions than those required for natural (in)direct effects [4]

- Recall that NDE involves fixing the mediator at the counterfactual value it would take under the reference exposure level ($x = 0$)
- In contrast, interventional direct effects set the mediator for each individual to a random draw from the distribution of the mediator under exposure level $X = x$, given covariate values $\boldsymbol{Z}$

[4] Vansteelandt and Daniel. Epidemiology 2017.

Interventional effects

More recently, another paper introduced interventional effects that map target trials and focus on evaluating shifts in mediator distributions[5]

- First, select policy-relevant questions related to mediator-shifting interventions
- Then, define and simulate mediating effects corresponding to such target trial
- This approach shifts focus away from merely discovering mechanistic pathways to the more tangible goal of evaluating mediator interventions
- Relevant when there are no-well defined interventions
- In the context of my example, one could argue that no well-defined mediator interventions exist, as the available data do not permit direct assessment of such interventions
- The next step is to extend this approach within the relative survival framework

[5]Moreno-Betancur et al. Stat Methods Med Res 2021.

Notation

- Let B_k denote the intervention targeting mediator M_k, with $k = 1, 2, \ldots, K$ (in our example it is only $k = 1, \ldots, 4$)
- $B_k = 1$ if the intervention is received and $B_k = 0$ if it does not
- Let M_k^x be the value that M_k would take if we intervened on X and set it to x

Question 1 - one policy premise

If targeting only one mediator, which intervention would provide the “biggest gain”, in terms of reducing disparities between exposure groups?

1. Intervention B_k would be applied independently of the other mediators
2. Intervention B_k would shift the distribution of mediator M_k to what it would be in the unexposed given $\boldsymbol{Z}$. This is equivalent to setting M_k to a random draw from the distribution it would have *under no exposure* given $\boldsymbol{Z}$.
3. Intervention B_k would sever the dependence on average between M_k and the other mediators, so that the joint distribution of the other mediators is held at what it would be *under exposure* given $\boldsymbol{Z}$ (worst-case scenario)

Question 1 - one policy premise

If targeting only one mediator, which intervention would provide the “biggest gain”, in terms of reducing disparities between exposure groups?

1. Intervention B_k would be applied independently of the other mediators
2. Intervention B_k would shift the distribution of mediator M_k to what it would be in the unexposed given $\boldsymbol{Z}$. This is equivalent to setting M_k to a random draw from the distribution it would have *under no exposure* given $\boldsymbol{Z}$.

Relaxing assumption 3 and replacing with:

a. The order of the mediators is $M_1, \ldots M_K$
b. Under intervention B_k, the joint distribution of the causally antecedent mediators of M_k is unaffected, remaining at what it would be under exposure
c. Under intervention B_k, the *conditional* joint distribution of the causally descendent mediators of M_k given $M_1, \ldots M_k$ and $\boldsymbol{Z}$ is what it would be under exposure

Random draws

The previous assumptions mean that the hypothetical intervention B_k would set the mediators to a random draw from the following joint distribution

$$P(\boldsymbol{M}^{\mathbf{1}}_{\overline{\boldsymbol{k-1}}} = \boldsymbol{m}_{\overline{\boldsymbol{k-1}}} \mid \boldsymbol{Z}) \times P(M^0_k = m_k \mid \boldsymbol{Z})$$
$$\times \; P(\boldsymbol{M}^{\mathbf{1}}_{\underline{\boldsymbol{k+1}}} = \boldsymbol{m}_{\underline{\boldsymbol{k+1}}} \mid \boldsymbol{Z}, \boldsymbol{M}^{\mathbf{1}}_{\overline{\boldsymbol{k-1}}} = \boldsymbol{m}_{\overline{\boldsymbol{k-1}}}, M^1_k = m_k)$$

Here, we assume that the hypothetical intervention would have no impact on the interdependence between the descendent mediators, given previous ones, with these associations remaining as they would be under exposure at a population level (given confounders).

Interventional indirect effect via M_k

The interventional indirect effect via M_k is defined as the difference between the outcome expectation in the exposed group and the arm in which the M_k distribution is shifted by B_k

$$\text{IIE}_k = p_{exp} - p_k$$

If interest is on relative survival

Under identifiability assumptions, this can be identified by:

$$E\left[\sum_{\boldsymbol{m}} R(t|X=1,\boldsymbol{Z},\boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=1,\boldsymbol{Z})\right]$$
$$- E\Bigg[\sum_{\boldsymbol{m}} R(t|X=1,\boldsymbol{Z},\boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}_{\overline{\boldsymbol{k-1}}}=\boldsymbol{m}_{\overline{\boldsymbol{k-1}}}|X=1,\boldsymbol{Z})$$
$$\times P(M_k=m_k|X=0,\boldsymbol{Z}) \times P(\boldsymbol{M}_{\underline{\boldsymbol{k+1}}}=\boldsymbol{m}_{\underline{\boldsymbol{k+1}}}|X=1,\boldsymbol{Z},\boldsymbol{M}_{\overline{\boldsymbol{k-1}}}=\boldsymbol{m}_{\overline{\boldsymbol{k-1}}},M_k=m_k)\Bigg]$$

Estimation can be performed using the Monte Carlo simulation-based g-computation approach that relies on the factorization of joint mediator distributions as the sequential product of conditional distributions.

Estimation of p_k

1. Fit regression models to estimate the require distributions.
   For instance, for p_1 i.e. this would be $P(M_1 = m_1|X, \boldsymbol{Z})$,
   $P(M_2 = m_2|X, M_1, \boldsymbol{Z})$, $P(M_3 = m_3|X, M_1, M_2, \boldsymbol{Z})$,
   $P(M_4 = m_4|X, M_1, M_2, M_3, \boldsymbol{Z})$ and $R(t|X, M_1, M_2, M_3, M_4, \boldsymbol{Z})$.
2. For each individual, sequentially draw mediator values from the fitted distributions, given their covariate vector $\boldsymbol{Z}$, the relevant exposure value and draws of previous mediators as required.
3. Using the fitted outcome model, for each individual predict the outcome given the value of their covariate vector $\boldsymbol{Z}$, the relevant exposure value and the mediator draws.
4. Repeat steps 1 to 3 a large number of times and average the outcome predictions across the whole sample and these repetitions to obtain an estimate of the required term.

Estimation of p_k

- Steps 1 to 4 are a way of estimating a weighted average, where the weights reflect the required population-level intervention on the joint distribution of the mediators
- The estimated terms are then added or subtracted to obtain the required contrast i.e. effect estimate
- The nonparametric bootstrap can be used to obtain standard errors, confidence intervals and p-values
- To reduce the risk of misspecification bias, use rich parametric models, including various interaction terms and higher-order terms (for continuous variables)

Example (simulated data) - one policy premise

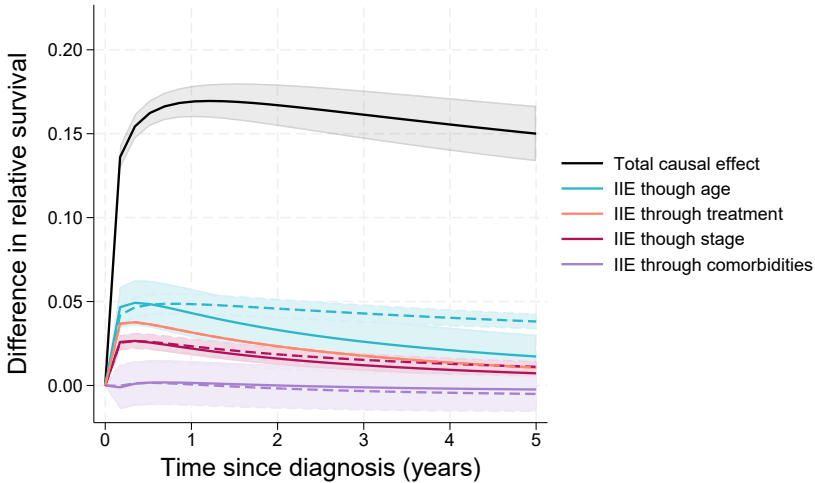

Question 2 - joint intervention

What would be the remaining disparities between exposure groups if it was possible to jointly target all the mediators?

- Let B_{all} denote an intervention targeting all the mediators
- Assuming that the hypothetical intervention B_{all} shifts the joint distribution of the mediators to be as in the unexposed (including the mediators interdependence) given $\boldsymbol{Z}$ we set the mediators to a random draw from the joint distribution

$$P(\boldsymbol{M}_{\cdot}^{\mathbf{0}} = \boldsymbol{m}|\boldsymbol{Z})$$

- Large remaining disparities after this intervention would suggest that there is a need to investigate additional intermediate processes

Interventional direct effect not via any mediator

Let p_{all} denote the outcome expectation in the arm shifting the joint distribution of all mediators. The interventional direct effect not via any mediator

$$\text{IDE} = p_{all} - p_{ctr}$$

is identifiable as

$$E\left[\sum_{\boldsymbol{m}} R(t|X=1, \boldsymbol{Z}, \boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=0, \boldsymbol{Z})\right]$$
$$- E\left[\sum_{\boldsymbol{m}} R(t|X=0, \boldsymbol{Z}, \boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=0, \boldsymbol{Z})\right]$$

Other related estimands

The effect of the joint intervention, defined in terms of the reduction achieved

$$\text{IIE}_{all} = p_{exp} - p_{all}$$

which is identifiable under assumptions by

$$E\left[\sum_{\boldsymbol{m}} R(t|X=1, \boldsymbol{Z}, \boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=1, \boldsymbol{Z})\right]$$
$$- E\left[\sum_{\boldsymbol{m}} R(t|X=1, \boldsymbol{Z}, \boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=0, \boldsymbol{Z})\right]$$

Similar estimands can also be defined when interest in on intervening on a subset of mediators, e.g. $M_{\underline{k}}$ but not on $M_{\overline{k-1}}$.

Example - joint intervention

Joint intervention in M_1, M_2, M_3, M_4

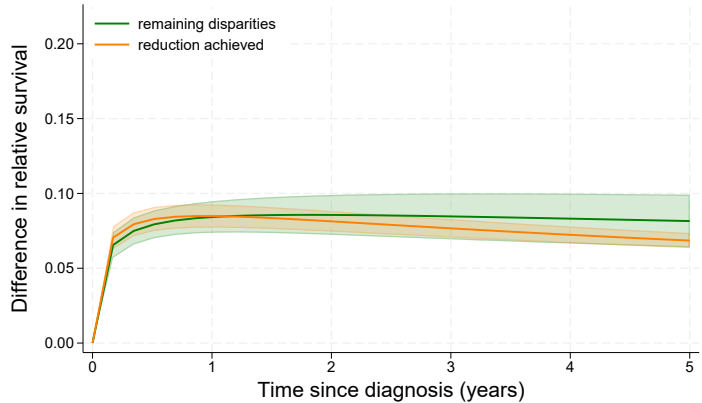

Example - joint intervention

Joint intervention in M_2, M_3, M_4

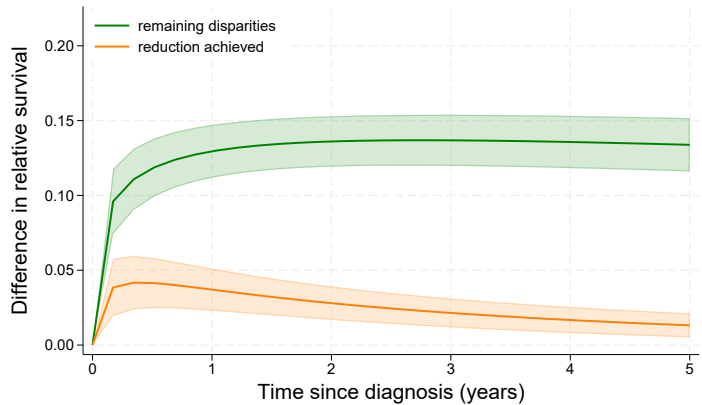

Question 3 - sequential policies

What would be the benefit of sequential policies, applying the separate mediator interventions under Question 1 (worst-case scenario) sequentially?

- A hypothetical intervention $B_{\{k\}}$ would set the mediators to a random draw from the following joint distribution

$$P(M_1^0 = m_k|\boldsymbol{Z}) \times \cdots \times P(M_k^0 = m_k|\boldsymbol{Z}) \times P(\boldsymbol{M}_{\underline{\boldsymbol{k+1}}}^{\boldsymbol{1}} = \boldsymbol{m}_{\underline{\boldsymbol{k+1}}}|\boldsymbol{Z})$$

- The interventional indirect effect of the k-th intervention in the sequence

$$\text{IIE}_{\{k\}} = p_{\{k-1\}} - p_{\{k\}}$$

Example - Sequential policy decomposition

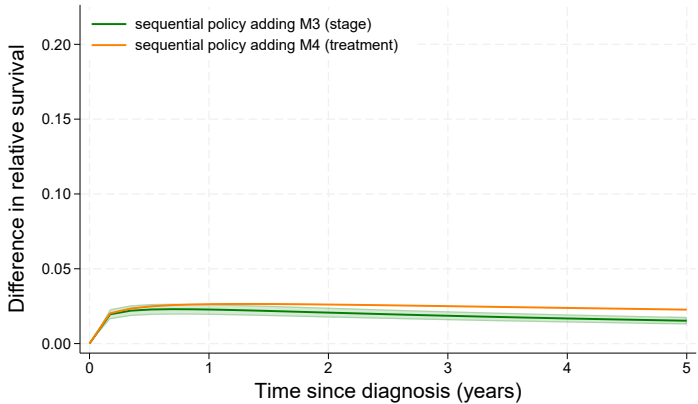

Example - sequential policies

It is also possible to form different contrasts and have a decomposition of the total effect

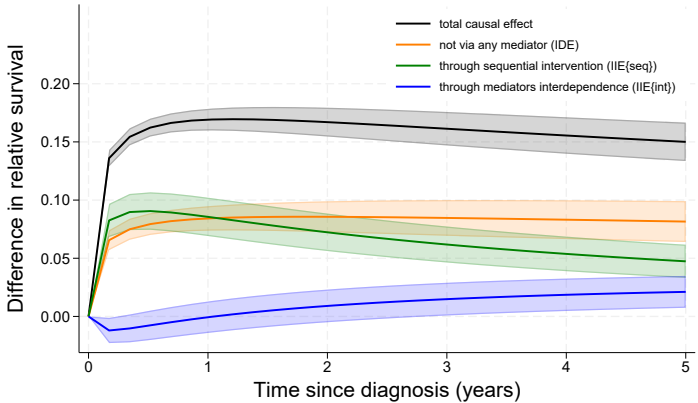

Moving to a real-world setting

- To form contrasts of all-cause survival we need to incorporate the expected survival probabilities $S^*(t)$
- There are many ways to do this, for example the IIE_k can be defined by multiplying both terms with $S^*(t|X = 1)$

Example - real world setting

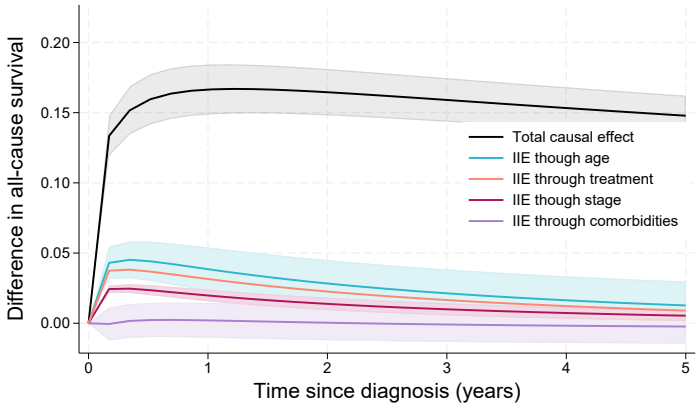

Assumptions

- When focusing on all-cause survival differences, we assume that the intervention considered has no impact on other cause mortality rates
- The two competing events (cancer death and other-cause death) are assumed to be conditionally independent
- Any observed effect of the intervention on survival is attributed solely to changes in cancer-specific mortality
- This distinction is motivated by the idea that certain interventions may have separable effects across competing causes; for example, increased engagement in cancer screening is expected to impact cancer mortality directly, while only indirectly affecting the probability of other-cause death without altering its underlying hazard
- However, this assumption may be difficult to justify in some contexts (e.g. interventions on comorbidities)

Other estimands

- All the effects discussed so far can also be defined within subsets of the whole population
- For instance, the IIE_k among the exposed that can be estimated by standardising only to patients of the exposed group
- Such estimands can be useful for assessing the potential impact of interventions using the covariate distribution of the actual population we plan to intervene on
- Moreover, effects in a real-world setting can also be visualised in terms of the avoidable/postponables deaths under hypothetical interventions
- The avoidable deaths is a time-specific measure as eventually all deaths will be realised

Avoidable deaths

The avoidable deaths after an intervention on mediator M_k that keeps the joint distribution of all other mediator as it would be under exposure (worst-case scenario) is identifiable by:

$$N^* \times \left(1 - E\left[\sum_{\boldsymbol{m}} S^*(t|X=1, \boldsymbol{Z}'_{sub}) \times R(t|X=1, \boldsymbol{Z}_{sub}, \boldsymbol{M}=\boldsymbol{m}) \times P(\boldsymbol{M}=\boldsymbol{m}|X=1, \boldsymbol{Z}_{sub})\right]\right)$$
$$- N^* \times \left(1 - E\left[\sum_{\boldsymbol{m}} S^*(t|X=1, \boldsymbol{Z}'_{sub}) \times R(t|X=1, \boldsymbol{Z}_{sub}, \boldsymbol{M}=\boldsymbol{m})\right.\right.$$
$$\left.\left. \times P(M_k = m_k|X=0, \boldsymbol{Z}_{sub}) \times P(\boldsymbol{M_{(-k)}} = \boldsymbol{m_{(-k)}}|X=1, \boldsymbol{Z}_{sub})\right]\right)$$

Example - avoidable deaths per 10,000 patients

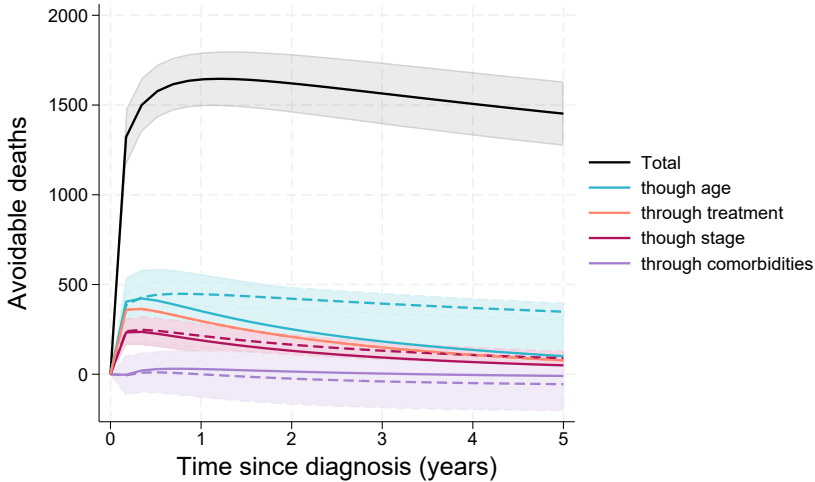

Example - avoidable deaths after sequential policies

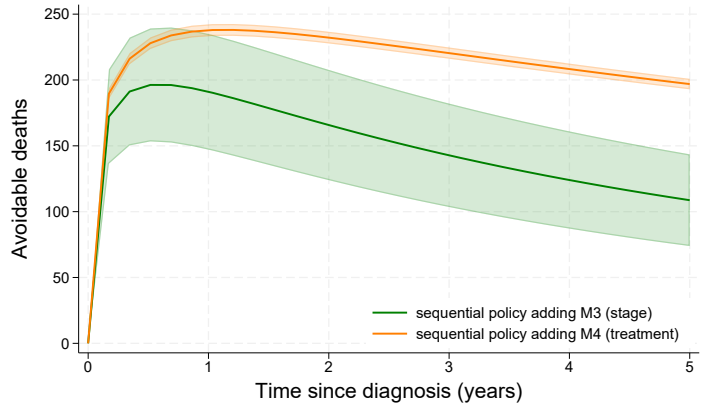

Mediation analysis with categorical mediators

Is staging I-IV adequate or does it result in substantial within-group heterogeneity in survival outcomes?

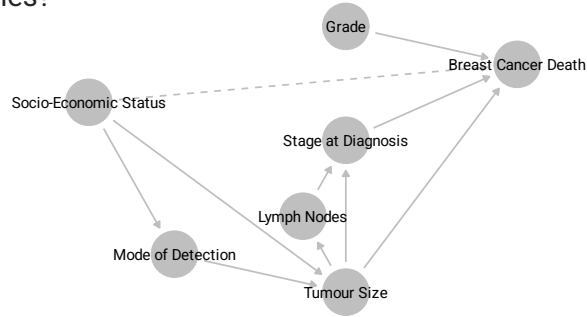

We performed a microsimulation[6] using a breast cancer natural history model developed in Swedish settings to assess this and found substantial bias in the estimated mediated effects.

[6] Gasparini A, Humphreys K, Booth S, Rutherford M, Syriopoulou E. Mediation Analysis with Imperfectly Defined Mediators: a Microsimulation Experiment with Breast Cancer Survival, Socio-Economic Status, and Stage at Diagnosis (under review)

Final thoughts

- Mediation analysis offers a valuable framework for investigating cancer disparities
- Using the relative survival framework allows to isolate cancer-related differences
- Interventional effects that evaluate shifts in mediator distributions are particularly relevant to our context
- However, the utility of such analyses depends on careful consideration of underlying assumptions and methodological challenges
- For instance, considering age at diagnosis as a mediator fails to account for the potential introduction of lead-time through interventions that influence the timing of diagnosis