

# Estimating Optimal Dynamic Treatment Strategies Under Resource Constraints Using Dynamic Marginal Structural Models

Ellen Caniglia (✉ [ellen.caniglia@nyulangone.org](mailto:ellen.caniglia@nyulangone.org))

New York University School of Medicine <https://orcid.org/0000-0002-7567-8428>

Eleanor Murray

Boston University School of Public Health

Miguel Hernan

Harvard T.H. Chan School of Public Health

Zach Shahn

IBM Research

---

## Research article

**Keywords:** dynamic strategy, HIV, marginal structural models, monitoring, optimal strategy, re- source constraints

**Posted Date:** December 6th, 2019

**DOI:** <https://doi.org/10.21203/rs.2.18323/v1>

**License:**  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

**Version of Record:** A version of this preprint was published at Statistics in Medicine on June 29th, 2021.  
See the published version at <https://doi.org/10.1002/sim.9107>.

# Estimating Optimal Dynamic Treatment Strategies Under Resource Constraints Using Dynamic Marginal Structural Models

Ellen Caniglia<sup>1,\*</sup>, Eleanor Murray<sup>2</sup>, Miguel Hernán<sup>3</sup>, Zach Shahn<sup>4</sup>

<sup>1</sup>Department of Population Health, New York University School of Medicine, New York, NY

<sup>2</sup>Department of Epidemiology, Boston University School of Public Health, Boston, MA

<sup>3</sup>Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, MA

<sup>4</sup>IBM Research, Yorktown Heights, New York, NY

\*Corresponding author

E-mail addresses for all authors:

[ellen.caniglia@nyulangone.org](mailto:ellen.caniglia@nyulangone.org)

[ejmurray@bu.edu](mailto:ejmurray@bu.edu)

[mhernan@hsph.harvard.edu](mailto:mhernan@hsph.harvard.edu)

[zachshahn@gmail.com](mailto:zachshahn@gmail.com)

## Abstract

**Background:** Existing methods for estimating the optimal treatment or monitoring strategy typically assume unlimited access to resources. However, when a health system has resource constraints, such as limited funds, access to medication, or monitoring capabilities, medical decisions must balance impacts on both individual and population health outcomes. That is, decisions should account for competition between individuals in resource usage. A simple solution is to estimate the (counterfactual) resource usage under the possible interventions and choose the optimal strategy for which resource usage is within acceptable limits.

**Methods:** We propose a method to identify the optimal dynamic intervention strategy that leads to the best expected health outcome accounting for a health system's resource constraints. We then apply this method to determine the optimal dynamic monitoring strategy for people living with HIV when resource limits on monitoring exist using observational data from the HIV-CAUSAL Collaboration. In our example, we identify the optimal CD4 cell count at which to decrease monitoring frequency from every 3-6 months to every 9-12 months among the subset of strategies that satisfy a hypothetical constraint on monitoring. We find the resource constrained dynamic strategy (RCDS) that minimizes the risk of virologic failure given a hypothetical resource constraint where individuals can only be monitored on average every 6 months.

**Results:** Given the hypothetical resource constraint, the optimal monitoring strategy is *'monitor CD4 cell count and HIV-RNA every 3-6 months when CD4 is below 320 cells/ $\mu$ l and every 9-12 months when CD4 is above 320 cells/ $\mu$ l.'*

**Conclusions:** We describe and apply a method to identify the optimal RCDS given a hypothetical resource constraint. Similar applications may be useful in settings with limited availability of or capacity for monitoring individuals living with HIV or other chronic diseases.

**Keywords:** dynamic strategy; HIV; marginal structural models; monitoring; optimal strategy; resource constraints.

## Background

Physicians repeatedly assess patients with chronic health conditions and make treatment decisions based on patient history at each assessment. Here, ‘treatment’ refers to any intervention, including monitoring through lab tests to inform future decisions. A ‘dynamic treatment strategy’ is a function mapping a patient’s treatment and covariate history up to the current visit to a treatment decision at that visit.

In resource limited settings where health system constraints prevent immediate initiation of treatment in all individuals, the optimal dynamic strategy is the strategy which, if implemented by all doctors, would lead to the best population health outcomes while ‘respecting’ the system’s resource constraints (in a sense we will make rigorous below). While randomized clinical trials of a wide range of dynamic strategies may be ideal for estimating this optimal strategy, they are usually financially or logistically infeasible. However, observational data can be used to estimate the optimal dynamic strategy under resource constraints. Note that, as would be the case with randomized trials, the optimal strategy here refers to the optimum from among a class of options assessed – the overall optimum may not be among these.

Luedtke and van der Laan [1] considered this problem for the case of a point exposure – that is, when there is only one time point at which a treatment decision is made. They considered resource constraints that place an upper limit  $\kappa$  on the expected proportion of treated patients, and defined the set of strategies which respected resource constraint  $\kappa$  as all strategies for which the expected proportion of treated patients in the population under that strategy is less than  $\kappa$ . The optimal point exposure strategy was then identified as the optimal strategy among those respecting the system constraint. However, optimal point exposure strategies are often of limited utility in clinical decision-making, especially in the context of chronic disease or long-term therapy. Such strategies cannot recommend, for example, “*Come back next month, and if the problem has progressed then begin treatment*”.

Here, we consider the optimal resource constrained dynamic strategy (RCDS) and present a method for estimating the optimal RCDS from a parameterized subclass of all strategies. For ex-

ample, suppose we restrict our attention to the class of monitoring strategies  $\{g_x : g_x(\bar{L}_{t-1}, \bar{N}_{t-1}) = \mathbf{1}\{L_{t-1} > x\}\}$  (where  $\mathbf{1}\{\cdot\}$  denotes the indicator function) that monitor  $(N_t)$  at time  $t$  if and only if the covariate  $L_{t-1}$  is greater than  $x$ . Such a class of strategies might be approximately appropriate, for example, for the decision of how often to monitor for anti-retroviral therapy (ART) failure or resistance in people living with HIV, where  $L_t$  represents CD4 count at time  $t$ .

In the absence of resource constraints, Orellana et al. [2] and Robins et al. [3] describe how to estimate the optimal strategy from a parameterized class of strategies using a dynamic Marginal Structural Model (dyn-MSM). A dyn-MSM models the expected counterfactual outcomes under strategies parameterized by  $x$  as a function  $b(x, \beta)$  of  $x$ . With an estimate  $\hat{\beta}$  of the dyn-MSM parameter  $\beta$ , estimating the optimal strategy simply reduces to finding  $x^{opt}$  maximizing  $b(x, \hat{\beta})$  (assuming larger values of  $b(x, \beta)$  are preferable). To accommodate resource constraints, we propose a fairly straightforward extension of this procedure that entails fitting two dyn-MSMs—one estimating the expected counterfactual clinical outcome  $b(x, \beta)$  and one estimating expected counterfactual treatment utilization  $h(x, \theta)$ . The optimal RCDS is then  $x_{rc}^{opt}$  maximizing  $b(x, \beta)$  subject to  $h(x, \theta) \leq \kappa$ .

We apply this approach to estimate the optimal RCDS for CD4 cell count and HIV-RNA monitoring in people living with HIV who have achieved viral suppression. CD4 cell count and HIV-RNA tests are used to monitor an individual's response to ART. More frequent monitoring has been shown to be associated with a lower risk of virologic failure (HIV RNA levels  $> 200$  copies/ml) at two years after viral suppression [4]. Guidelines recommend dynamic monitoring strategies in which virologically suppressed individuals on ART may be monitored less frequently once their CD4 cell count crosses above a certain threshold. However, the optimal point at which to decrease monitoring frequency is unclear even for high-income settings. In a health system with limited funds for monitoring, the CD4 cell count at which monitoring may be decreased must be chosen from a subset of strategies where the average number of tests does not exceed the available resources.

# Methods

## Notation

Let:

- $t \in \{0, \dots, K\}$  denote time, assumed discrete, with  $K$  the end of the study;
- $N_t$  be a variable indicating the whether an individual is monitored at time  $t$ ;
- $Y$  denote health outcome we aim to optimize;
- $L_t$  denote covariates at time  $t$  that may influence monitoring decisions at time  $t + 1$ ;
- $D$  denote the total number of monitoring tests received by an individual;
- $\bar{X}_t$  denote  $X_0, \dots, X_t$  and  $\underline{X}_t$  denote  $X_t, \dots, X_K$  for arbitrary time varying variable  $X$ .

We assume that we observe independent and identically distributed (*iid*) realizations of the random vector  $(\mathbf{L}_0, \dots, \mathbf{L}_K, N_0, \dots, N_K, Y)$ . We use capital letters to denote random variables and corresponding lower case letters to denote specific values that random variables might take.

A treatment *strategy* is a rule or function that determines the value to which  $N_t$  will be set at each time  $t$  for a given observed history, i.e. a function  $g : (\bar{L}_{t-1}, \bar{N}_{t-1}) \rightarrow N_t$ . A strategy is said to be *static* if its recommendation for the present does not depend on past covariate and treatment values, and can therefore be specified from baseline. An example of a static strategy would be ‘monitor every 3 months’. A strategy is said to be *dynamic* if it does depend on past covariate and treatment values. An example of a dynamic strategy would be ‘monitor if time since last monitoring  $\geq 6$  months or if last observed viral load  $> 100$  copies/ml and time since last monitoring  $\geq 2$  months’. Most realistic, clinically-relevant strategies are dynamic.

We denote arbitrary strategies by  $g$  and we adopt the counterfactual framework of Robins [5] in which corresponding to each possible strategy  $g$  are counterfactual random variables  $Y(g)$ ,  $L_d(g)$ , and  $D(g)$  that would have been observed had strategy  $g$  been followed, possibly contrary to fact. Implicit in the notation for counterfactuals (e.g.  $Y(g)$ ) is the assumption that the treatment strategy followed by one patient does not influence any other patient. This implicit assumption is called the

‘No Interference Assumption’ by Rubin [6]. Note that in defining our resource constraint, we are optimizing based on the (counterfactual) average number of treatment or monitoring events per individual over a defined time period, but assuming no competition between individuals to access care under the optimal RCDS. We also make the additional assumptions [7]:

**Sequential Positivity:** If  $f(\bar{l}_t, \bar{n}_{t-1}) > 0$ ,  $f(N_t = n | \bar{L}_t = \bar{l}_t, \bar{N}_{t-1} = \bar{n}_{t-1}) > 0 \forall n, t$  (1)

**Consistency:** For any strategy  $g$ , if for a given individual  $N_t = g(\bar{L}_t, \bar{N}_{t-1})$  for each  $t$ , then  $Y = Y(g)$  and  $\bar{L}_K = \bar{L}_K(g)$  for that individual (2)

**Sequential Exchangeability:**  $Y(g) \perp N_t | \bar{L}_{t-1} = \bar{l}_t, \bar{N}_{t-1} = \bar{n}_{t-1} \forall t, g, \bar{n}, \bar{l}$  (3)

We define resource constraints as caps on the expected number of doses or monitoring events per patient over a defined time period. We say that strategy  $g$  respects resource constraint  $\kappa$  if

$$E[D(g)] \leq \kappa \tag{4}$$

We consider parameterized classes of strategies  $\{g_x\}$  and seek to estimate

$$x_\kappa^{opt} \equiv \operatorname{argmax}_x E[Y(g_x)] \text{ subject to } E[D(g_x)] \leq \kappa \tag{5}$$

### Review of Dynamic Marginal Structural Models (dyn-MSMs)

A dyn-MSM is a model for expected counterfactual outcomes under a class of strategies parameterized by  $x$  as a function of  $x$ , i.e.

$$E[Y(g_x)] = b(x, \beta) \tag{6}$$

To estimate  $\beta$ , first note that each individual might follow multiple strategies from the class  $\{g_x\}$ . Let  $\Lambda_i$  denote the number of strategies that individual  $i$  follows. Generate an artificial dataset with  $\Lambda_i$  contributions from each individual:  $(Y_i, x_{i1}), \dots, (Y_i, x_{i\Lambda_i})$ . Using the artificial dataset, we can

fit by weighted least squares the regression model

$$E[Y | x] = b(x; \gamma)$$

with weights

$$W(x) \equiv \prod_{k=0}^K \frac{f^*(x)}{f(N_k | \bar{L}_{k-1}, \bar{N}_{k-1})}$$

to obtain  $\hat{\gamma}$ , where  $f(\cdot|\cdot)$  denotes conditional density and  $f^*(x)$  is an arbitrary function of  $x$ . When treatment or monitoring probabilities are unknown, as they are in our application, a consistent estimator  $\hat{f}(N_k | \bar{L}_{k-1}, \bar{N}_{k-1})$  of  $f(N_k | \bar{L}_{k-1}, \bar{N}_{k-1})$  can be plugged into  $W(x)$ . Under sequential exchangeability (3) and consistency (2), [2] shows that the weighted regression parameter estimate  $\hat{\gamma}$  approaches the causal estimand  $\beta$  in the limit.

### Estimating Optimal Treatment Strategies With Resource Constraints

To estimate the optimal RCDS, we simply estimate the parameters of two dyn-MSMs:

$$E[Y(g_x)] = b(x; \beta) \tag{7}$$

$$E[D(g_x)] = h(x; \theta) \tag{8}$$

and then estimate the  $x$  indexing the optimal strategy as

$$\begin{aligned} \hat{x}_\kappa^{opt} &\equiv \operatorname{argmax}_x b(x; \hat{\beta}) \\ &\text{subject to } h(x; \hat{\theta}) \leq \kappa \end{aligned} \tag{9}$$

Standard errors for  $\hat{\beta}$ ,  $\hat{\theta}$ , and certain derived quantities such as  $E[Y(g_x)]$  and  $E[D(g_x)]$  for any  $x$  can be computed by bootstrap or analytically using formulas in [2]. We note that confidence intervals for  $\hat{x}_\kappa^{opt}$  and  $E[Y(g_{\hat{x}_\kappa^{opt}})]$  are not generally available in the resource constrained setting through standard methods. This is due to complications arising from an estimator obtained from a constrained optimization where the constraint itself contains a plugin estimator.

### Application to Monitoring of HIV Patients

We apply the method described above to estimate the optimal RCDS for monitoring CD4 cell count and HIV-RNA in people living with HIV using data from the HIV-CAUSAL collaboration. The HIV-CAUSAL collaboration combines data from prospective cohorts of people living with HIV enrolled in universal health care systems in Brazil, Canada, France, Greece, Netherlands, Spain, Switzerland, UK, and USA.

We have previously reported on the optimal dynamic monitoring strategy in this cohort and showed that decreasing monitoring when CD4 cell count  $>200$  cells/ $\mu$ l compared to  $>500$  cells/ $\mu$ l does not worsen short-term clinical and immunologic outcomes in virologically suppressed individuals living with HIV but may increase the risk of virological failure [4]. We now extend these results to identify the optimal RCDS under constraints on the average number of monitoring events over a two-year period. However, because the majority of the HIV-CAUSAL data comes from high-income countries, we apply an artificial resource constraint selected to demonstrate the methodologic approach.

First we briefly describe the eligibility criteria and monitoring strategies under consideration. We then describe the estimation of the optimal RCDS.

*Eligibility criteria:* Previously antiretroviral therapy naive HIV-positive individuals who initiate antiretroviral therapy in 2000 or later and achieve confirmed virologic suppression (2 consecutive HIV-RNA  $\leq 200$  copies/ml) within 12 months of initiating therapy are eligible for inclusion in the study. Individuals must meet the following additional eligibility criteria at baseline (date of confirmed virologic suppression): 18 years of age or older, CD4 cell count measurement within the previous 3 months, no history of an AIDS-defining illness, and no pregnancy (when information was available).

*Monitoring strategies:* We consider 31 dynamic monitoring strategies, based loosely on current clinical guidelines. Under each strategy, CD4 cell count and HIV-RNA are monitored every 3-6 months when CD4 is below the strategy's threshold and every 9-12 months when CD4 is above the threshold. Each strategy corresponds to a CD4 threshold ranging from 200-500 cells/ $\mu$ l in increments of 10 cells/ $\mu$ l. All of the monitoring strategies further require individuals to be monitored

once every 3-6 months when HIV-RNA >200 copies/ml or after diagnosis of an AIDS-defining illness, and that CD4 cell count and HIV-RNA be monitored concurrently.

*Follow-up period and outcome:* Individuals are followed from baseline until death, pregnancy (if known), loss to follow-up, or the administrative end of follow-up. The outcome of interest is virologic failure (HIV-RNA >200 copies/ml) at 24 months of follow-up.

*Statistical methods:* We compare the 31 monitoring strategies using the replication and censoring approach. Briefly, we create an expanded dataset by making 31 exact replicates of each individual (1 per strategy). If and when an individual's data are no longer consistent with a given strategy, we artificially censor the corresponding replicate at that time. We compute inverse probability weights to adjust for the potential selection bias induced by the artificial censoring.

We then fit an inverse-probability weighted Poisson regression model to estimate the risk ratio of virologic failure at 24 months of follow-up among those with measurements at  $24 \pm 2$  months. The model includes a flexible functional form of the strategy variable (restricted cubic splines) and the baseline covariates: sex, CD4 cell count (<200, 200-349, 350-499,  $\geq 500$  cells/L), years since HIV diagnosis (<1, 1 to 4,  $\geq 5$  years, unknown), race (white, black, other or unknown), geographic origin (N. America/W. Europe, Sub-Saharan Africa, other, unknown), acquisition group (heterosexual, homosexual or bisexual, injection drug use, other or unknown), calendar year (restricted cubic splines with 3 knots at 2001, 2007, and 2011), age (restricted cubic splines with 3 knots at 25, 39, and 60), cohort, and months from cART initiation to virologic suppression (2-4, 5-8,  $\geq 9$ ). Under the assumptions described above, the parameters of the regression model consistently estimate the parameters of a dynamic marginal structural model. The model's estimated parameters are used to estimate the standardized risk of virologic failure at 24-months for each monitoring strategy.

Next, we fit an inverse-probability weighted log-linear regression model to estimate the mean number of measurements at 24 months of follow-up. As above, the model includes a flexible functional form of the strategy variable and the baseline covariates. The predicted values are used to estimate the standardized mean number of measurements at 24-months for each monitoring strategy.

After estimating the counterfactual risk of virologic failure at 24 months and counterfactual

mean number of measurements at 24 months, we rank the strategies by the counterfactual mean number of measurements at 24 months. We then restrict our consideration to the strategies that satisfy the resource constraint. For our example, we consider a hypothetical constraint allowing an average of one CD4 cell count and one HIV-RNA test per person every 6 months, for an average of 4 measurements per person over 24 months. Under this constraint, only the strategies that lead to an average of 4 measurements per person over the 24-months of follow-up will be considered. Importantly, any given individual may have more than 4 measurements over 24 months, but the population average must be 4 measurements or fewer.

Finally, we find the optimal strategy for minimizing the risk of virologic failure among the strategies that satisfy the constraint.

## Results

Figure 1 shows the estimates obtained from the dyn-MSMs for virologic failure at 24 months and for mean number of monitoring events over 24 months across the range of CD4 cell count thresholds considered. In this example, the estimated risk of virologic failure is monotonically increasing and the estimated mean number of measurements is monotonically decreasing as the CD4 cell count threshold increases, so the optimal RCDS can be identified graphically as the lowest CD4 cell count threshold for which the mean number of monitoring events over 24 months is at or below the resource constraint,  $\kappa$ . Table 1 gives the same information with 95% confidence intervals obtained via 500 bootstrap samples.

In our example, we consider the case of  $\kappa = 4$  and identify the optimal threshold for switching monitoring frequency as 320 cells/ $\mu\text{l}$ . The optimal RCDS is then ‘monitor CD4 cell count and HIV-RNA every 3-6 months when CD4 is below 320 cells/ $\mu\text{l}$  and every 9-12 months when CD4 is above 320 cells/ $\mu\text{l}$ ’.

## Discussion

Dynamic treatment strategies are often a better representation of real-world clinical decision-making processes than static or point intervention strategies. However, resource utilization of dynamic strategies is difficult to assess, since the number of individuals requiring intervention over time under a given strategy cannot be straightforwardly determined at baseline. When a health system faces resource constraints that prohibit implementing the true optimal dynamic treatment strategy, the optimal RCDS is instead required.

Here we propose a method to identify the optimal RCDS within a parameterized class of strategies of interest by estimating the counterfactual resource usage. We apply this method to estimate the optimal RCDS for monitoring frequency among individuals living with HIV who achieve virologic suppression. This application may be useful in resource limited settings with limited availability of or capacity for monitoring individuals living with HIV or other chronic diseases.

Our choice of  $\kappa = 4$  was somewhat arbitrary. If we had instead chosen  $\kappa = 3$ , we would have found that none of the strategies under consideration would satisfy this resource constraint. Interestingly, if we had chosen  $\kappa = 4.7$ , we would have identified the optimal threshold for switching monitoring frequency as 410 cells/ $\mu\text{l}$ , even though all the strategies in the range from 200-450 cells/ $\mu\text{l}$  would have satisfied the resource constraint (however, the confidence intervals around our estimates are quite wide). A more conservative approach when selecting the optimal threshold for switching may be to restrict our consideration to the strategies for which the lower and upper bound of the 95% confidence interval are less than or equal to the resource constraint. In our example, we would then select 300 cells/ $\mu\text{l}$  as the optimal threshold. Finally, we considered a discrete number of strategies corresponding to CD4 thresholds that are divisible by 10, which may be beneficial from a clinical perspective. Future work could consider continuous CD4 cell count or multidimensional strategies (e.g., strategies based on both CD4 cell count and HIV RNA thresholds).

In reality, determining the number of CD4 cell count and HIV-RNA measurements a setting is willing to allocate depends on a complex assessment of the costs and health benefits of monitoring. In our illustrative application, we imposed a constraint on the number of tests, which we imagined

was derived from a hypothetical corresponding cost constraint. In other applications, it might be useful to directly bound cost instead. For example, since HIV-RNA tests cost more than CD4 tests, an optimal strategy satisfying a total cost constraint may be a more flexible joint strategy that allows CD4 cell count and HIV-RNA to be monitored with different frequencies. Future studies should also assess other health outcomes such as quality-adjusted life years associated with various monitoring strategies. Finally, even in the absence of a single hard resource constraint, examining outcomes of optimal strategies over a range of hypothetical cost constraints could allow for computation of incremental cost-effectiveness ratios, which could be useful for key stakeholders and decision-makers.

## **Conclusions**

We propose a method to identify the optimal RCDS given a hypothetical resource constraint and apply this method to estimate the optimal RCDS for monitoring frequency among individuals living with HIV who achieve virologic suppression. Similar applications may be useful in settings with limited availability of or capacity for monitoring individuals living with HIV or other chronic diseases.

## **Abbreviations**

RCDS: resource constrained dynamic strategy

dyn-MSM: dynamic marginal structural model

## **Declarations**

Ethics approval and consent to participate: not applicable

Consent for publication: not applicable

Availability of data and materials: The raw data analysed during the current study are not publicly available due to laws and regulations in each country from which data is received. Data from

each participating study can be requested from the sponsoring institution in accordance with the applicable laws or regulations in each country.

Competing interests: The authors declare that they have no competing interests.

Funding: This work was partially supported by NIH grants R37 AI102634 and T32 AI007433. The funders had no role in the design of the study, data collection, analysis, interpretation of data or writing of the manuscript.

Authors' contributions: All authors contributed to the conception and design of this study. EC analyzed the data with assistance from EM and ZS. EC, EM, and ZS drafted the manuscript. All authors reviewed and approved the final manuscript.

Acknowledgements: We thank Andrew Phillips, Linda Wittkop, Giota Touloumi, and Hansjakob Furrer for useful comments on an earlier draft of this paper and James Robins for helpful discussions. Preprint available on Arxiv.org under the standard Arxiv.org licensing agreement [8].

Table 1: Risk of Virologic Failure and cumulative number of measurements at 24 months for CD4 threshold

CD4 Threshold* (cells/ $\mu$ l)	Virologic Failure		Cumulative # Measurements	
	Risk (%)	95% CI	Expected Value	95% CI
500	6.91	(3.99, 9.82)	4.94	(4.82, 5.06)
490	6.85	(4.09, 9.61)	4.89	(4.78, 5.00)
480	6.80	(4.17, 9.43)	4.84	(4.73, 4.94)
470	6.74	(4.21, 9.27)	4.78	(4.69, 4.88)
460	6.69	(4.23, 9.15)	4.73	(4.64, 4.83)
450	6.64	(4.22, 9.05)	4.68	(4.59, 4.78)
440	6.59	(4.18, 8.99)	4.63	(4.54, 4.73)
430	6.54	(4.13, 8.95)	4.58	(4.49, 4.68)
420	6.52	(4.08, 8.95)	4.53	(4.43, 4.63)
410	6.51	(4.03, 8.99)	4.48	(4.38, 4.58)
400	6.54	(4.02, 9.05)	4.43	(4.32, 4.53)
390	6.60	(4.05, 9.15)	4.37	(4.27, 4.48)
380	6.71	(4.14, 9.28)	4.31	(4.21, 4.42)
370	6.87	(4.26, 9.48)	4.25	(4.15, 4.36)
360	7.08	(4.41, 9.76)	4.19	(4.08, 4.30)
350	7.33	(4.55, 10.12)	4.13	(4.02, 4.24)
340	7.62	(4.67, 10.57)	4.07	(3.95, 4.18)
330	7.93	(4.79, 11.08)	4.01	(3.89, 4.13)
<b>320</b>	<b>8.27</b>	<b>(4.91, 11.62)</b>	<b>3.96</b>	<b>(3.84, 4.08)</b>
310	8.61	(5.09, 12.14)	3.92	(3.79, 4.04)
300	8.97	(5.31, 12.61)	3.88	(3.75, 4.00)
290	9.33	(5.59, 13.06)	3.84	(3.71, 3.97)
280	9.70	(5.91, 13.49)	3.81	(3.68, 3.93)
270	10.08	(6.22, 13.94)	3.78	(3.65, 3.91)
260	10.48	(6.52, 14.43)	3.75	(3.62, 3.88)
250	10.88	(6.78, 14.99)	3.73	(3.59, 3.86)
240	11.31	(6.98, 15.64)	3.70	(3.55, 3.85)
230	11.75	(7.12, 16.38)	3.67	(3.52, 3.83)
220	12.21	(7.21, 17.72)	3.65	(3.48, 3.82)
210	12.68	(7.25, 18.12)	3.62	(3.44, 3.80)
200	13.18	(7.23, 19.13)	3.60	(3.41, 3.79)

\*The CD4 Threshold corresponds to the CD4 cell count at which monitoring frequency changes from once every 2-7 months (if CD4 cell count is below the threshold) to once every 8-13 months (if CD4 cell count is above the threshold). Each strategy also includes monitoring once every 2-7 months when HIV-RNA > 200 copies/ml or after diagnosis of an AIDS-defining illness.

The monitoring strategies falling in the grey area meet the restriction that CD4 cell count and HIV-RNA may only be monitored every six months. Among the monitoring strategies that meet the restriction, the 320 threshold strategy is the optimal strategy.

## References

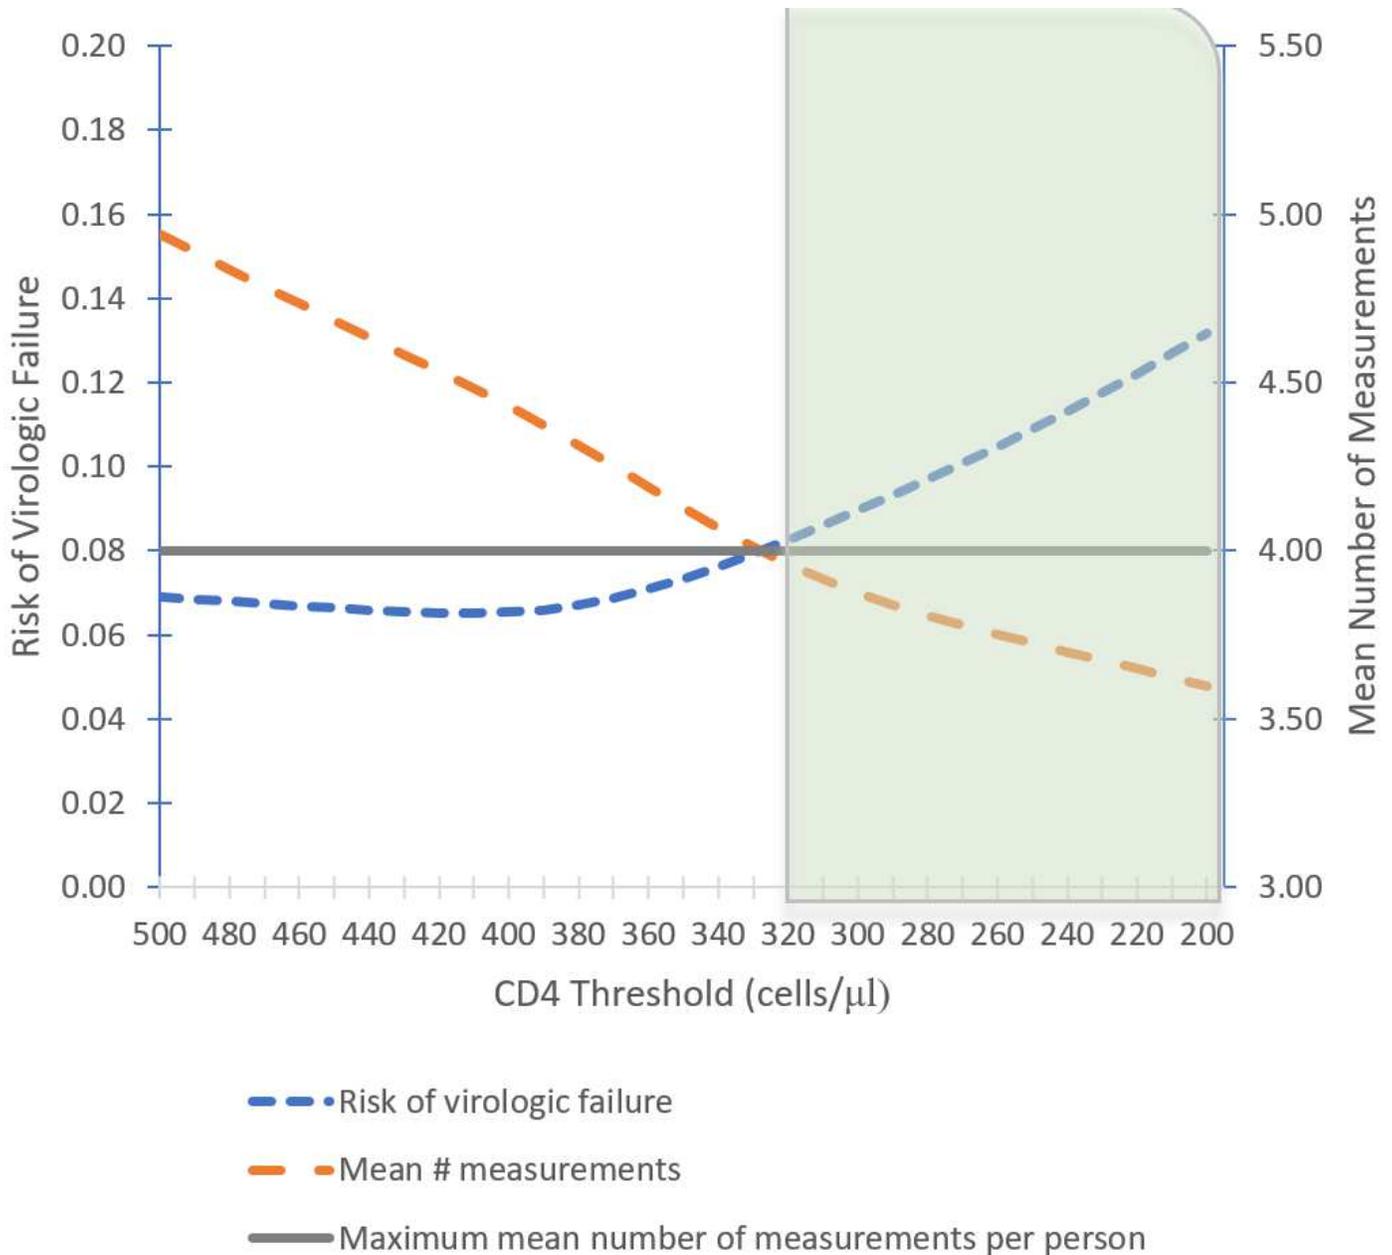
- [1] Luedtke A.R., van der Laan M. J. (2016). Optimal Individualized Treatments in Resource-Limited Settings *International Journal of Biostatistics* **12(1)**: 283-303.
- [2] Orellana L., Rotnitzky A.G., Robins J.M. (2006). Generalized Marginal Structural Models for Estimating Optimal Treatment Regimes *Technical Report*, Department of Biostatistics, Harvard School of Public Health.
- [3] Robins J.M, Orellana L., Rotnitzky A.G. (2008). Estimation and extrapolation of optimal treatment and testing strategies *Statistics in Medicine* **27(23)**:4678-4721.
- [4] Caniglia E.C., Cain L.E., Sabin C.A., Robins J.M., et al. (2017). Comparison of Dynamic Monitoring Strategies Based on CD4 Cell Counts in Virally Suppressed, HIV-Positive Individuals on Combination Antiretroviral Therapy in High-Income Countries: A Prospective, Observational Study *The Lancet HIV* **4(6)**: e251-e259.
- [5] Robins J.M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period — Application to the healthy worker survivor effect. *Mathematical Modelling* **7**: 1393-1512.
- [6] Rubin D.B. (1978). Bayesian Inference for Causal Effects: The Role of Randomization *The Annals of Statistics* **1**: 34-58.
- [7] Robins J.M., Hernán M.A. (2008). Estimation of the causal effects of time-varying exposures. In *Longitudinal Data Analysis*, G. Fitzmaurice, M. Davidian, G. Verbeke, and G. Molenberghs (eds), 533-599. New York: Chapman and Hall/CRC Press.
- [8] Caniglia E.C., Murray E.J., Hernán M.A., Shahn Z (2019). Note on Estimating Optimal Dynamic Treatment Strategies Under Resource Constraints Using Dynamic Marginal Structural Models. Preprint at <https://arxiv.org/abs/1903.06488>

## List of Figures

Figure 1 Title: Risk of virologic failure and mean number of measurements per person at 24 months of follow-up by CD4 threshold strategy.

Figure 1 Legend: The grey line represents one potential resource constraint – a cap on the average per person number of measurements over 24 months of follow-up of 4. Strategies in the green shaded area meet this restriction, and the CD4 threshold 320 strategy is the optimal RCDS.

# Figures



**Figure 1**

Risk of virologic failure and mean number of measurements per person at 24 months of follow-up by CD4 threshold strategy. Figure 1 Legend: The grey line represents one potential resource constraint – a cap on the average per person number of measurements over 24 months of follow-up of 4. Strategies in the green shaded area meet this restriction, and the CD4 threshold 320 strategy is the optimal RCDS.