Supplementary Material for "Using Pilot Data for Power Analysis of Observational Studies for the Estimation of Dynamic Treatment Regimes"

A Technical Details

We proposed a test for empirical power (PWR) that rejects when $\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left\{\widehat{V}_n(\psi)-\frac{z_{1-\vartheta_2}\hat{\varsigma}_n(\psi)}{\sqrt{n}}\right\}\geq B_0$. Recall we chose ϑ_1 and ϑ_2 such that $\vartheta_1+\vartheta_2=\alpha$. The type 1 error rate is then given by

$$\begin{aligned}
&P\left[\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left\{\widehat{V}_n(\psi)-\frac{z_{1-\vartheta_2}\hat{\varsigma}_n(\psi)}{\sqrt{n}}\right\}\geq B_0\right]\\
&\quad\leq P\left[\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left\{\widehat{V}_n(\psi)-\frac{z_{1-\vartheta_2}\hat{\varsigma}_n(\psi)}{\sqrt{n}}\right\}\geq V(\boldsymbol{d}^{opt})\right]\\
&\quad\leq P\left\{\widehat{V}_n(\psi^*)-\frac{z_{1-\vartheta_2}\hat{\varsigma}_n(\psi^*)}{\sqrt{n}}\geq V(\boldsymbol{d}^{opt})\right\}+\vartheta_1+o(1)\\
&\quad= P\left[\frac{\sqrt{n}\{\widehat{V}_n(\psi^*)-V(\psi^*)\}}{\hat{\varsigma}_n(\psi^*)}\geq z_{1-\vartheta_2}\right]+\vartheta_1+o(1)\\
&\quad\leq \alpha+o(1).
\end{aligned}$$

Therefore, the proposed test is an α-level test for (PWR). The power for this test is given by

$$\begin{aligned}
&P\left[\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left\{\widehat{V}_n(\psi)-\frac{z_{1-\vartheta_2}\hat{\varsigma}_n(\psi)}{\sqrt{n}}\right\}\geq B_0\right]\\
&\quad= P\left\{\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left[\frac{\sqrt{n}\{\widehat{V}_n(\psi)-V(\psi)\}}{\hat{\varsigma}_n(\psi)}+\frac{\sqrt{n}\{V(\psi)-B_0\}}{\hat{\varsigma}_n(\psi)}\right]\geq z_{1-\vartheta_2}\right\}\\
&\quad\geq P\left\{\inf_{\psi\in\Psi_{n,1-\vartheta_1}}\left[\frac{\sqrt{n}\{\widehat{V}_n(\psi)-V(\psi)\}}{\hat{\varsigma}_n(\psi)}+\frac{\min\left[\sqrt{n}\{V(\psi)-B_0\},\sqrt{n}\eta\right]}{\hat{\varsigma}_n(\psi)}\right]\geq z_{1-\vartheta_2}\right\}.
\end{aligned}$$

B Confidence Interval for ψ

We used a simulation study to illustrate the coverage of the multistage m-out-of-n bootstrap proposed in Section 3.3. The same data generating model as the simulation study in Section 4 of the main paper was used, which was given by:

$$
\begin{aligned}
&X_1 \sim \mathrm{N}\{0,1\}, && P(A_1 = 1|\boldsymbol{H}_1 = \boldsymbol{h}_1) = \left\{1 + e^{-(\varpi_{1,0} + \varpi_{1,1} x_1)}\right\}^{-1} \\
&\tau_1 \sim \mathrm{N}(0,1), && X_2 = \mu_{20} + \mu_{21} X_1 + \tau_1, \\
&\boldsymbol{H}_2^\intercal = (X_1, A_1, X_2), && P(A_2 = 1|\boldsymbol{H}_2 = \boldsymbol{h}_2) = \left\{1 + e^{-(\varpi_{2,0} + \varpi_{2,1}^\intercal h_2)}\right\}^{-1}, \\
&\tau_2 \sim \mathrm{N}(0,1), && X_3 = \mu_{30} + \mu_{31} X_1 + \mu_{32} X_2 + \tau_2, \\
&\boldsymbol{H}_3^\intercal = (X_1, A_1, X_2, A_2, X_3), && P(A_3 = 1|\boldsymbol{H}_3 = \boldsymbol{h}_3) = \left\{1 + e^{-(\varpi_{3,0} + \varpi_{3,1}^\intercal h_3)}\right\}^{-1}, \\
&\boldsymbol{H}_{3,1}^\intercal = (1, X_1, X_2, X_3), && \boldsymbol{H}_{3,0}^\intercal = (1, X_1, A_1, A_1 X_1, X_2, A_2, A_2 X_1, A_2 X_2, X_3, X_1^2) \\
&\upsilon \sim \mathrm{N}(0,1), && Y = \boldsymbol{H}_{3,0}^\intercal \lambda_{3,0} + A_3 \boldsymbol{H}_{3,1}^\intercal \lambda_{3,1} + \upsilon.
\end{aligned}
$$

We posited the following models for the treatment-free and blip functions:

$$
\begin{aligned}
&\gamma_1(\boldsymbol{h}_1, a_1; \psi_1) = a_1(\psi_{1,0} + \psi_{1,1} x_1), \\
&g_1(\boldsymbol{h}_1; \beta_1) = \beta_{1,0} + \beta_{1,1} x_1, \\
&\gamma_2(\boldsymbol{h}_2, a_2; \psi_2) = a_2(\psi_{2,0} + \psi_{2,1} x_1 + \psi_{2,2} x_2), \\
&g_2(\boldsymbol{h}_2; \beta_2) = \beta_{2,0} + \beta_{2,1} x_1 + \beta_{2,2} a_1 + \beta_{2,3} a_1 x_1 + \beta_{2,4} x_2, \\
&\gamma_3(\boldsymbol{h}_3, a_3; \psi_3) = a_3(\psi_{3,0} + \psi_{3,1} x_1 + \psi_{3,2} x_2 + \psi_{3,3} x_3), \\
&g_3(\boldsymbol{h}_3; \beta_3) = \beta_{3,0} + \beta_{3,1} x_1 + \beta_{3,2} a_1 + \beta_{3,3} a_1 x_1 + \beta_{3,4} x_2 + \beta_{3,5} a_2 + \beta_{3,6} a_2 x_2 + \beta_{3,7} x_3.
\end{aligned}
$$

Therefore, the blip models were correctly specified while the treatment-free models were all misspecified. We posited a correctly specified logistic regression model for the treatment model so that dynamic weighted ordinary least squares (dWOLS) produced consistent estimates of the blip parameters. The parameters of the data generating model were given

by:

$$\varpi_1 = (0.25, 1), \qquad \varpi_2 = (0.25, 1, -1, -1),$$
$$\varpi_3 = (0.25, 0.5, 0.5, -0.5, 1, -0.5),$$
$$\mu_2 = (0, 0.5), \qquad \mu_3 = (0, -0.5, 0.5),$$
$$\lambda_{3,0} = (1, 1, 0.5, -0.75, 0.5, -0.5, -0.5, 0.5, 0.5, 0.25), \quad \lambda_{3,1} = (0.25, 0.5, 0.5, -0.5).$$

In the three-stage setting, we have that ψ_1 and ψ_2 are both nonregular. Let $\psi_{\text{NR}} = (\psi_1, \psi_2)$ denote the vector of all nonregular parameters indexing the dynamic treatment regime. We used the proposed multistage m-out-of-n bootstrap to construct a 95% confidence region for ψ_{NR}. We also evaluated the nominal coverage of 95% confidence regions for ψ_1 and ψ_2 as well as the main effect of treatment in the first and second stage given by $\psi_{1,0}$ and $\psi_{2,0}$, respectively. For $\psi_{1,0}$ and $\psi_{2,0}$, we constructed confidence intervals. For the remaining parameters, which are multidimensional, we constructed joint confidence regions.

We fixed the sample size at $n = 500$ and the tuning parameter $\nu = 0.001$. We considered settings in which κ was fixed, and in which κ was chosen in a data-driven fashion. For the fixed κ approach, we considered three values, $\kappa \in \{0.01, 0.05, 0.1\}$, whereas for the data-adaptive approach, a modification of the double bootstrap methods of Chakraborty et al. [2013] was employed to choose κ from a fine grid of candidate values. In particular, this data-adaptive approach proceeded as follows:

1. Set a grid of candidate values for κ, e.g. $\kappa \in \{0.025, 0.05, \ldots, 1\}$.
2. Set κ to the smallest value in the grid of candidate values.
3. Draw B_1 bootstrap samples of size n and calculate $\psi_{k,n}^{(b_1)}$ for $k = 1, \ldots, K$, $b_1 = 1, \ldots, B_1$.
4. Calculate $\hat{\varrho}_K^{(b_1)} = \mathbb{P}_n \mathbb{I}\left[n\{\boldsymbol{H}_{K,\psi}^{(b_1)T} \hat{\psi}_K^{(b_1)}\}^2 \leq \tau_n\{\boldsymbol{H}_{K,\psi}^{(b_1)}\}\right]$ and $\hat{m}_K^{(b_1)} = n^{\frac{1+\kappa\{1-\hat{v}\varrho_K^{(b_1)}\}}{1+\kappa}}$ for $b_1 = 1, \ldots, B_1$.

5. Draw B_2 nested bootstrap samples of size $\hat{m}_K^{(b_1)}$ from each of the B_1 samples.

6. Use the B_2 samples to construct a confidence interval for $\boldsymbol{h}_{K-1,\psi}^{(b_1)T}\psi_{K-1,n}^{(b_1)}$ for $b_1 = 1, \ldots, B_1$.

7. Estimate $\hat{\varrho}_{K-1}^{(b_1)}$ by the proportion of individuals for which the confidence interval for $\boldsymbol{h}_{K-1,\psi}^{T}\psi_{K-1,n}^{(b_1)}$ contains 0 and calculate $\hat{m}_{K-1}^{(b_1)} = n^{\frac{1+\kappa\{1-\hat{\varrho}_{K-1}^{(b_1)}\}}{1+\kappa}}$ for $b_1 = 1, \ldots, B_1$.

8. Repeat steps 5-7 to calculate $\hat{\varrho}_k^{(b_1)}$ for $k = 1, \ldots, K$, $b_1 = 1, \ldots, B_1$.

9. Calculate $\hat{\varrho}^{(b_1)} = \max_k \hat{\varrho}_k^{(b_1)}$ and $\hat{m}^{(b_1)} = n^{\frac{1+\kappa\{1-\hat{\varrho}^{(b_1)}\}}{1+\kappa}}$ for $b_1 = 1, \ldots, B_1$.

10. Draw B_2 nested bootstrap samples of size $\hat{m}^{(b_1)}$ from each of the B_1 samples.

11. Calculate $\hat{\psi}_{k,n,\hat{m}^{(b_1)}}^{(b_1,b_2)}$ for $k = 1, \ldots, K$, $b_1 = 1, \ldots, B_1$, $b_2 = 1, \ldots, B_2$.

12. Calculate the $\epsilon_k/2 \times 100$ and $(1-\epsilon_k/2) \times 100$ percentiles of $\left[\sqrt{\hat{m}^{(b_1)}}\{\hat{\psi}_{k,n,\hat{m}^{(b_1)}}^{(b_1,b_2)} - \hat{\psi}_{k,n}^{(b_1)}\}, b_2 = 1, \ldots, B_2\right]$, which we will denote by $\hat{l}_k^{(b_1)}$ and $\hat{u}_k^{(b_1)}$ for $k = 1, \ldots, K$, $b_1 = 1, \ldots, B_1$.

13. Construct confidence intervals as $\left(\hat{\psi}_{k,n}^{(b_1)} - \hat{u}_k^{(b_1)}/\sqrt{\hat{m}^{(b_1)}}, \hat{\psi}_{k,n}^{(b_1)} + \hat{l}_k^{(b_1)}/\sqrt{\hat{m}^{(b_1)}}\right)$ for $k = 1, \ldots, K$, $b_1 = 1, \ldots, B_1$.

14. Estimate the coverage rate as

$$\frac{1}{B_1}\sum_{b_1=1}^{B_1}\left[\prod_{k=1}^{K-1}\mathbb{I}\left\{\hat{\psi}_{k,n}^{(b_1)} - \hat{u}_k^{(b_1)}/\sqrt{\hat{m}^{(b_1)}} \leq \hat{\psi}_{k,n} \leq \hat{\psi}_{k,n}^{(b_1)} + \hat{l}_k^{(b_1)}/\sqrt{\hat{m}^{(b_1)}}\right\}\right].$$

15. If the estimated coverage is greater than or equal to the desired coverage, then select the current value of κ. Otherwise, set κ to the next highest value in the grid of candidate values and repeat steps 3-15.

For each simulation setting, we performed 500 Monte Carlo repetitions. When constructing a confidence interval for only a subset of ψ_{NR} with an adaptive choice of κ, step 14 of the double

bootstrap algorithm is adjusted to use only the estimated coverage of the parameter(s) of interest.

Table 1 displays Monte Carlo estimates of the coverage of the proposed procedure, the average resample size, and the average area of the confidence regions for each of the simulation settings. The estimated coverage was slightly below the nominal coverage for $\psi_{1,0}$, ψ_1, and ψ_{NR} when $\kappa = 0.01$. As κ increased, the size of our confidence regions increased, leading to increased coverage. When $\kappa = 0.1$, the confidence regions were conservative, with estimated coverages all greater than 97.8%. The adaptive choice of κ using the double bootstrap led to estimated coverage that was close to the nominal coverage for all parameters, but at the cost of a significantly more computationally expensive procedure.

Table 1: Coverage of 95% confidence regions using the multistage m-out-of-n bootstrap. $\mathbb{E}(\hat{m})$ gives the average resample size and mean area gives the average area of the confidence regions produced for each parameter.

Parameter	κ	$\mathbb{E}(\hat{m})$	Coverage	Mean Area
$\psi_{1,0}$	0.1	339.60	99.2%	0.73
$\psi_{2,0}$	0.1	339.60	98.2%	0.64
ψ_1	0.1	339.60	97.8%	4.05
ψ_2	0.1	339.60	98.2%	5.08
ψ_{NR}	0.1	339.60	98.0%	9.52
$\psi_{1,0}$	0.05	407.77	95.6%	0.66
$\psi_{2,0}$	0.05	407.77	96.8%	0.58
ψ_1	0.05	407.77	95.4%	3.91
ψ_2	0.05	407.77	97.0%	4.94
ψ_{NR}	0.05	407.77	94.2%	9.22
$\psi_{1,0}$	0.01	479.36	94.8%	0.61
$\psi_{2,0}$	0.01	479.26	97.0%	0.54
ψ_1	0.01	479.26	92.8%	3.79
ψ_2	0.01	479.26	95.6%	4.79
ψ_{NR}	0.01	479.26	93.4%	8.91
$\psi_{1,0}$	Adaptive	433.45	95.8%	0.64
$\psi_{2,0}$	Adaptive	445.51	97.4%	0.55
ψ_1	Adaptive	399.46	95.0%	3.90
ψ_2	Adaptive	435.68	95.8%	4.87
ψ_{NR}	Adaptive	395.32	93.8%	9.23

C Simulation Results for Two-Stage Study

We examined the finite sample performance of our proposed power calculations and sample size procedure when applied to a two-stage dynamic treatment regime (DTR) using a simulation study. The data generating model for the two-stage study was given by

$$\begin{aligned}
&\boldsymbol{X}_1 \sim \mathrm{N}_3\{\mathbf{0}, \Omega_{AR1}(0.5)\}, && P(A_k = 1|\boldsymbol{H}_k = \boldsymbol{h}_k) = \left\{1 + e^{-(\varpi_{k,0} + \sum_{j=1}^{k} \varpi_{j,1}^{\intercal} x_j)}\right\}^{-1} \text{ for } k = 1, 2,\\
&\tau \sim \mathrm{N}(0, 1), && X_2 = \mu_0 + \mu_1^{\intercal}\boldsymbol{X}_1 + \tau,\\
&\boldsymbol{H}_{2,1} = (1, \boldsymbol{X}_1, X_2), && \boldsymbol{H}_{2,0} = (1, \boldsymbol{X}_1, A_1, A_1\boldsymbol{X}_1, X_2, X_{11}^2),\\
&\upsilon \sim \mathrm{N}(0, 1), && Y = \boldsymbol{H}_{2,0}^{\intercal}\lambda_{2,0} + A_2\boldsymbol{H}_{2,1}^{\intercal}\lambda_{2,1} + \upsilon,
\end{aligned}$$

where $\Omega_{AR1}(0.5)$ was an autoregressive covariance matrix with $\{\Omega_{AR1}(0.5)\}_{ij} = 0.5^{|i-j|}$. We posited the following treatment-free and blip models:

$$\begin{aligned}
&\gamma_1(\boldsymbol{h}_1, a_1; \psi_1) = a_1(\psi_{1,0} + \psi_{1,1}^{\intercal}\boldsymbol{x}_1), && \gamma_2(\boldsymbol{h}_2, a_2; \psi_2) = a_2(\psi_{2,0} + \psi_{2,1}^{\intercal}\boldsymbol{x}_1 + \psi_{2,2}x_2),\\
&g_1(\boldsymbol{h}_1; \beta_1) = \beta_{1,0} + \beta_{1,1}^{\intercal}\boldsymbol{x}_1, && g_2(\boldsymbol{h}_2; \beta_2) = \beta_{2,0} + \beta_{2,1}^{\intercal}\boldsymbol{x}_1 + \beta_{2,2}a_1 + \beta_{2,3}^{\intercal}\boldsymbol{x}_1 a_1 + \beta_{2,4}x_2.
\end{aligned}$$

Therefore, the blip model at each stage was correctly specified, but both treatment-free models were misspecified. The model for the propensity score was given by a correctly specified logistic regression model so that dWOLS produced consistent estimates of the blip parameters. The parameter values used for this generating model are given by

$$\begin{aligned}
&\varpi_1 = (0.25, 1, -1, 1), && \varpi_2 = (0.25, 1, -1, -1, -1, -1),\\
&\lambda_{2,0} = (1, 1, 1, 1, 0.25, 0.5, -0.5, -0.5, 0.5, 0.25), && \lambda_{2,1} = (0.25, 0.5, 0.5, -0.5, 0.5),\\
&\mu = (0, 0.5, 0.5, 0.5).
\end{aligned}$$

The simulations for estimating the power for a given sample size for a two-stage study proceeded similarly to the simulations for the three-stage study (see main text). We let $\alpha = 0.05$ and $\eta = 1.5$, so we estimated the power for a 0.05 level test of $H_0 : V(\boldsymbol{d}^{\mathrm{opt}}) \leq B_0$.

We again evaluated the performance of the procedure when the effect size of tailoring was equal to $\eta = 1.5$ and examined how the results changed as the true effect size increased. We let $V(\boldsymbol{d}^{\text{opt}}) = B_0 + \eta + \Delta\eta$ and varied $\Delta \in \{0, 0.25, 0.5\}$. We let the size of the pilot study vary such that $n_0 \in \{200, 400\}$ and estimated the power for a set of different sample sizes given by $n \in \{250, 500, 750\}$.

Table 2 contains the mean, median, and standard deviation of the estimated power across 500 simulation repetitions for different combinations of pilot study sizes, sample sizes, and effect sizes. The mean of the estimated powers was close to the true power in almost all combinations of sample and effect sizes. As we saw for the three-stage simulations, when the true power is close to 1, the mean is slightly biased due to the distribution of the estimated power being truncated at 1. This bias decreased as the variability in the estimated powers decreased, which occurred when the true effect size increased and the size of the pilot study increased.

For the sample size calculations, we let $\alpha = 0.05$, $\phi = 0.1$, and $\eta = 1.5$ so that the first condition (PWR) was satisfied if we had a 0.05 level test of $H_0 : V(\boldsymbol{d}^{\text{opt}}) \leq B_0$ that had power 90% provided that $V(\boldsymbol{d}^{\text{opt}}) \geq B_0 + 1.5$. We examined the power as the effect size changed, similarly to the approach used in the three-stage simulations by letting $V(\boldsymbol{d}^{\text{opt}}) + \Delta = B_0 + 1.5$ and varying $\Delta \in \{0, 0.25, 0.5, 1\}$. We let $\epsilon = 0.5$ and $\zeta = 0.1$, so the second condition (OPT) was satisfied if $P\{V(\widehat{\boldsymbol{d}}_n) \geq V(\boldsymbol{d}^{\text{opt}}) - 0.5\} \geq 0.9$.

These simulations proceeded in the same manner as the three-stage simulations (see main text). We let the size of the pilot study vary such that $n_0 \in \{200, 400\}$. The power was then estimated on a grid of candidate sample sizes using 500 bootstrap repetitions for each sample size. The confidence region for ψ was constructed via the m-out-of-n bootstrap with $\kappa = 0.2$. Using an adaptive choice of κ is best in practice, but is too computationally intensive for each repetition in a simulation study. To estimate the smallest sample size that achieved the desired power, we again regressed the estimated power on the tested sample sizes and used

Table 2: Estimated power from the proposed power calculations using a pilot study of size n_0 for varying sample sizes n. We assume the effect size under the alternative hypothesis is given by $\eta = 1.5$. Δ denotes the difference between the true value of the optimal regime and B_0 which is given by $\eta(1+\Delta)$, so $\Delta = 0$ corresponds to the true effect size being equal to η. The remaining columns contain the mean, median, and standard deviation of the estimated powers across 500 simulated pilot studies as well as the true power which is calculated via simulation. We do not estimate the power for $n = 250$ when the pilot is of size $n_0 = 400$, since the full study we are using the pilot to estimate the power for would be larger than the pilot.

Δ	n	n_0	True PWR	Mean PWR	Med PWR	SD PWR
0	250	200	0.55	0.57	0.62	0.30
0.25	250	200	0.96	0.86	0.92	0.17
0.5	250	200	1.00	0.95	0.96	0.05
0	500	200	0.85	0.76	0.91	0.30
0.25	500	200	1.00	0.97	1.00	0.08
0.5	500	200	1.00	1.00	1.00	0.00
0	750	200	0.97	0.87	0.99	0.23
0.25	750	200	1.00	0.99	1.00	0.04
0.5	750	200	1.00	1.00	1.00	0.00
0	500	400	0.84	0.80	0.89	0.23
0.25	500	400	1.00	0.99	1.00	0.03
0.5	500	400	1.00	1.00	1.00	0.00
0	750	400	0.96	0.90	0.98	0.18
0.25	750	400	1.00	1.00	1.00	0.02
0.5	750	400	1.00	1.00	1.00	0.00

the fitted model to solve for the sample size that resulted in the targeted power.

Tables 3 and 4 contain the results for sizing for (PWR) and (OPT), respectively. The results were very similar to those for the three-stage simulation study. When $\Delta = 0$ and $n_0 = 200$, the true value of the optimal regime equaled the comparison mean, B_0, plus the effect size of interest, and the procedure was slightly underpowered with an estimated power of 78.03%. Increasing the size of the pilot study to $n_0 = 400$ increased the power to 85.00%. As Δ increased, the power increased to 100% as expected and the variance in the estimated sample size decreased. Table 4 shows that the estimated concentration when

Table 3: Empirical power (PWR) using the projection-based sample size procedure at a nominal level of 90 using a pilot study of size $n_0 = 200$ and $n_0 = 400$. Δ denotes the difference between the true value of the optimal regime and B_0 which is given by $\eta(1 + \Delta)$. $P(\hat{n} = \infty)$ represents the proportion of pilot studies for which $\hat{n}(\mathcal{D}_{n_0}) = \infty$. The remaining columns give the mean, median, quartiles, and standard deviation of the estimated sample sizes across the 500 simulation repetitions.

Δ	n_0	$\mathbb{E}(\hat{n})$	Q1($\hat{n}$)	Med($\hat{n}$)	Q3($\hat{n}$)	SD($\hat{n}$)	$P(\hat{n} = \infty)$	PWR
0	200	610.03	274.00	441.00	712.00	486.54	0.03	78.03
0.25	200	276.31	212.00	236.50	279.00	133.94	0.00	96.20
0.5	200	220.45	201.00	217.00	234.00	34.47	0.00	99.80
1	200	216.76	198.00	216.00	233.00	25.83	0.00	100.00
0	400	651.44	376.75	532.00	750.00	441.22	0.00	85.00
0.25	400	255.86	210.00	224.00	255.00	105.85	0.00	94.08
0.5	400	210.88	199.00	210.00	220.00	20.84	0.00	100.00
1	400	211.15	200.00	210.00	222.00	16.82	0.00	100.00

sizing for condition (OPT) was 100% for all values of ϵ and n_0. When ϵ or n_0 increased, the mean and variance of the estimated sample size decreased.

D Pseudocode for Simulation Study

A single repetition of the simulation study when sizing for condition (PWR) proceeds as follows:

1. Generate a pilot study, $\mathcal{D}_{n_0}$, of size n_0.
2. Use bootstrap oversampling of the pilot data to estimate the power for a grid of potential sample sizes.
3. Regress the estimated power on the tested sample sizes using only the observations that resulted in estimated power within a small range (5%) of the targeted power of 90%.
4. Use the regression model to estimate the sample size, $\hat{n}(\mathcal{D}_{n_0})$, that achieves 90% power.

Table 4: Empirical concentration (OPT) using the projection-based sample size procedure at a nominal level of 90 using a pilot study of size $n_0 = 200$ and $n_0 = 400$. We test whether the true value of the estimated regime is within ϵ of the true value of the true optimal regime. The remaining columns give the mean, median, quartiles, and standard deviation of the estimated sample sizes across the 500 simulation repetitions.

ϵ	n_0	$\mathbb{E}(\hat{n})$	Q1($\hat{n}$)	Med($\hat{n}$)	Q3($\hat{n}$)	SD($\hat{n}$)	OPT
0.3	200	966.48	852.00	950.50	1066.75	161.64	100.00
0.5	200	532.02	485.00	523.50	573.25	68.83	100.00
0.7	200	375.92	343.00	371.50	401.00	44.93	100.00
0.3	400	936.70	865.00	931.00	1000.00	108.95	100.00
0.5	400	509.71	479.00	500.00	535.00	46.80	100.00
0.7	400	361.62	340.75	361.00	382.25	30.62	100.00

5. Simulate a study of size $\hat{n}(\mathcal{D}_{n_0})$.
6. Use the simulated study to estimate an optimal treatment regime and test the null hypothesis, $H_0 : V(\boldsymbol{d}^{\text{opt}}) \leq B_0$.

When sizing for condition (OPT), a single repetition is given by:

1. Generate a pilot study, $\mathcal{D}_{n_0}$, of size n_0.
2. Use bootstrap oversampling of the pilot data to estimate $\mathfrak{Q}^{(b)}_{n_0,n,1-\vartheta_2,1-\vartheta_1}$ for a grid of potential sample sizes.
3. Regress $\mathfrak{Q}^{(b)}_{n_0,n,1-\vartheta_2,1-\vartheta_1}$ on the tested sample sizes.
4. Use the regression model to estimate the smallest n such that $\mathfrak{Q}^{(b)}_{n_0,n,1-\vartheta_2,1-\vartheta_1} \leq \epsilon$ and denote it $\hat{n}(\mathcal{D}_{n_0})$.
5. Simulate a study of size $\hat{n}(\mathcal{D}_{n_0})$.
6. Use the simulated study to estimate an optimal treatment regime we denote by $\widehat{\boldsymbol{d}}_n$.
7. Calculate the value of the estimated optimal regime and evaluate whether $V(\widehat{\boldsymbol{d}}_n) \geq V(\boldsymbol{d}^{\text{opt}}) - \epsilon$.

E Details of Electronic Health Record Data

For the electronic health record data of Kaiser Permanente Washington members, the antidepressants in the selective serotonin reuptake inhibitor (SSRI) class were citalopram, escitalopram, fluoxetine, fluvoxamine, paroxetine, sertraline, and vilazodone. The alternative classes of antidepressants were norepinephrine and dopamine reuptake inhibitors (NDRI), tricyclic antidepressants (TCA), serotonin and norepinephrine reuptake inhibitors (SNRI), tetracyclic antidepressants (TeCA), and monoamine oxidase inhibitors (MAOI). The drugs prescribed that fit into the alternative classes of antidepressants were bupropion (NDRI), desipramine (TCA), desvenlafaxine (SNRI), duloxetine (SNRI), imipramine (TCA), mirtazapine (TeCA), nortriptyline (TCA), selegiline (MAOI), and venlafaxine (SNRI).

References

Bibhas Chakraborty, Eric B Laber, and Yingqi Zhao. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics*, 69(3):714–723, 2013.