

Web Appendix to ‘Validating predictors of therapeutic success: A causal inference approach’

Ariel Alonso Abad, Wim Van der Elst, & Geert Molenberghs

April 13, 2015

Abstract

This document contains two main parts. In Appendix A, the data of the case study discussed in Alonso *et al.* (2015) are analyzed. All analyses are conducted with the new R package *EffectTreat*. In Appendix B (starting on page 13), some model checks are provided.

Appendix A. Individual causal treatment effects: Case study analysis using the R package *EffectTreat*

In personalized medicine one wants to determine, for a given patient and his outcome on a pretreatment predictor, which treatment will likely be more beneficial for him. The R library *EffectTreat* allows for the quantification of the predictive causal association, i.e., the association between the pretreatment predictor and the individual causal treatment effect. In this Appendix, the use of the library is illustrated based on the data of the case study discussed in Alonso *et al.* (2015). The data are introduced in Section 1 and analyzed in Sections 2-5.

1 The dataset: A clinical trial in opiate/heroin addiction

The data come from a randomized clinical trial in which the clinical utility of buprenorphine/naloxone (experimental treatment) was compared to clonidine (control treatment) for a short-term (13-day) opiate/heroin detoxification treatment. Before and after the treatment, patients were assessed for relapse, withdrawal symptoms, and treatment satisfaction.

Here, the potential pretreatment variable (S) is the Clinical Opiate Withdrawal Scale (COWS) score at screening. The COWS is an 11-item interviewer-administered questionnaire designed to provide a description of signs and symptoms of opiate withdrawal (e.g., sweating, runny nose, etc). A higher COWS score is indicative for more withdrawal symptoms. The number of days that heroin was used in the 30 days prior to the second follow-up (the second follow-up took place 3 months after the start of the treatment) was used as the true endpoint T .

Data were available for 335 patients, of whom $n = 106$ received the active control clonidine and $n = 229$ received the experimental treatment buprenorphine/naloxone. Study drop-out was substantial: T was observed for $n = 104$ patients and missing for 231 patients. Multiple imputation (MI) was used to handle the missing data in all the analyses and 5 imputed data sets were used. When a complete case analysis was conducted, the results were similar and the substantive conclusions were identical (data not shown). A detailed description of the results and the appropriate R code to carry out the complete case analysis can be obtained by contacting the authors (wim.vanderelst@gmail.com).

The data are not included in the *EffectTreat* library as they are not in the public domain. Nonetheless, the data can be downloaded (after registration) from the National Institute of Drug Abuse website (<http://datashare.nida.nih.gov/protocol/data>). In the following analyses the combined data of studies NIDA-CTN-0001 and NIDA-CTN-0002 were used.

2 Computing the Predictive Causal Association (PCA; ρ_ψ)

After installation of the *EffectTreat* package in R (`install.packages("EffectTreat")`), the package is loaded in memory:

```
library(EffectTreat)
```

The function `PCA.ContCont` (Predictive Causal Association in the Continuous Continuous case) implements a sensitivity analysis in which ρ_ψ (PCA) is estimated across a set of plausible values for the unidentified correlation $\rho_{T_0T_1}$ (details in Alonso *et al.*, 2015). This function requires the user to specify the following main arguments:

- `T0S=` and `T1S=`: the correlations between the pretreatment predictor S and the true endpoint in the control (ρ_{T_0S}) and experimental (ρ_{T_1S}) treatment groups.
- `T0T0=` and `T1T1=`: the variances of the true endpoint in the control ($\sigma_{T_0T_0}$) and experimental ($\sigma_{T_1T_1}$) treatment groups.
- `SS=`: the variance of the pretreatment predictor S (σ_{SS}).
- `T0T1=`: a vector (or scalar) that specifies plausible values for the unidentifiable correlation ($\rho_{T_0T_1}$) between the counterfactual outcomes T_0 and T_1 . Default `seq(-1, 1, by=.01)`, i.e., the values $-1, -0.99, -0.98, \dots, 1$.

In the opiate/heroin addiction study, $\hat{\rho}_{T_0S} = -0.3699$ ($p = 0.001$) and $\hat{\rho}_{T_1S} = -0.3367$ ($p = 0.001$). Notice that the negative correlations between S and T_0/T_1 indicate that patients who have higher $S =$ COWS scores (more withdrawal symptoms at screening) tend to use less heroin in the 30-day interval after the treatment in both treatment conditions. Importantly, the difference between $\hat{\rho}_{T_1S}$ and $\hat{\rho}_{T_0S}$ was not significant ($p = 0.75$). Further, $\hat{\sigma}_{T_0T_0} = 83.2872$, $\hat{\sigma}_{T_1T_1} = 95.6189$, and $\hat{\sigma}_{SS} = 356.2158$. The following command is used to conduct the analysis:

```
Results <- PCA.ContCont(T0S=-0.3699, T1S=-0.3367, T0T0=83.2872, T1T1=95.6189, SS=356.2158, T0T1=seq(-1, 1, by=.01))
```

A summary of the results can be obtained by applying the `summary()` function to the fitted `Results` object:

```
summary(Results)

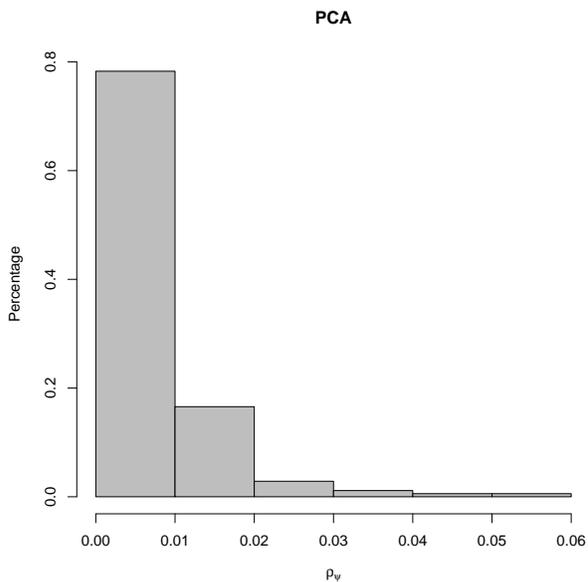
##
## Function call:
##
## PCA.ContCont(T0S = -0.3699, T1S = -0.3367, T0T0 = 83.2872, T1T1 = 95.6189,
##           SS = 356.2158, T0T1 = seq(-1, 1, by = 0.01))
##
##
## # Total number of matrices that can be formed by the specified vectors and/or scalars
## # of the correlations in the function call
## #~~~~~
##
## 201
##
## # Total number of positive definite matrices
## #~~~~~
##
## 175
```

```
##
##
## # Predictive causal association (PCA) results summary
## #~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~::~
##
## Mean (SD) PCA: 0.0088 (0.0065) [min: 0.0047; max: 0.0561]
## Mode PCA: 0.0057
##
## Quantiles of the PCA distribution:
##
##          5%          10%          20%          50%          80%
## 0.0048349119 0.0049663779 0.0052650549 0.0066422575 0.0103939852
##          90%          95%
## 0.0144561189 0.0198081950
```

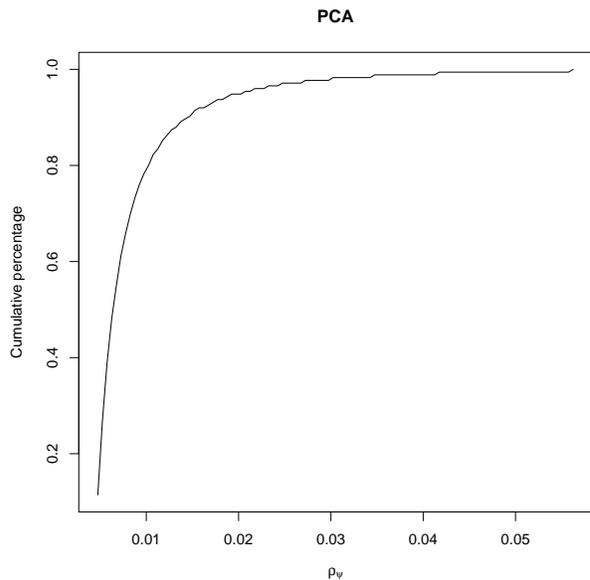
The output shows that out of the 201 matrices that can be formed, based on the specified vector for $\rho_{T_0T_1}$ and the identifiable variances and correlations, 175 were valid (positive definite) covariance matrices. The subsequent section in the output shows that the mean $\rho_\psi = 0.0088$, mode $\rho_\psi = 0.0057$, and median $\rho_\psi = 0.0066$. Further, 95% of the ρ_ψ values were ≤ 0.0198 and ρ_ψ was at most 0.0561. These results clearly show that in all ‘realities’ that are compatible with the observed data, ρ_ψ is rather low. It can thus be concluded that the accuracy by which a patient’s individual causal treatment effect on T ($\Delta_{Tj} = T_{1j} - T_{0j}$) can be predicted based on the COWS at screening is very poor.

The `plot()` function is a useful tool to further explore the frequency distribution of ρ_ψ . For example, plots of the relative frequencies (percentages) and cumulative frequencies can be requested by using the `Type="Percent"` and `Type="CumPerc"` arguments in the `plot()` call:

```
plot(Results, Type="Percent", breaks=7) #histogram with percentages
```



```
plot(Results, Type="CumPerc") #cumulative percentages
```



These plots confirm the earlier claim that ρ_ψ is rather low, irrespectively of the values assumed for the unidentifiable correlation $\rho_{T_0T_1}$.

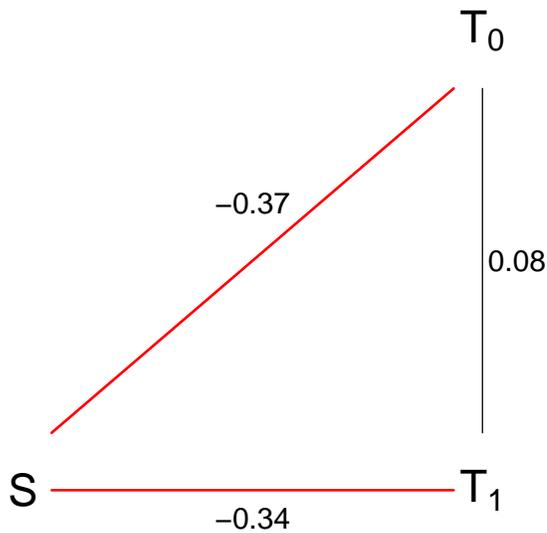
3 Relationship between $\rho_{T_0T_1}$ and PCA

In some applications one may want to explore in more detail some specific areas of the previous frequency distributions. For instance, one may want to evaluate which assumptions for the unidentified correlation $\rho_{T_0T_1}$ typically lead to a particular range of ρ_ψ values. In this context, the `CausalPCA.ContCont()` function (Causal diagram for the Predictive Causal Association in the Continuous Continuous case) can be a useful tool. The function provides a causal diagram depicting the median $\rho_{T_0T_1}$ value for a specified range of values of ρ_ψ . The following arguments are needed:

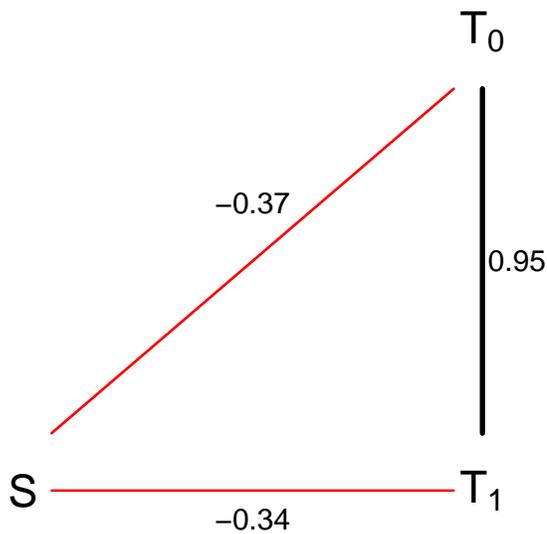
- `x=`: a fitted object of class `PCA.ContCont`.
- `Min=`, `Max=`: the minimum and maximum values for ρ_ψ that should be considered.

For example, the following code produces the causal diagrams for $0 < \rho_\psi < 0.02$ (an interval that contains the lowest 95% of the ρ_ψ values) and $0.02 < \rho_\psi < 0.06$ (the highest 5% of the ρ_ψ values):

```
CausalPCA.ContCont(x = Results, Min = 0, Max = 0.02)
```



```
CausalPCA.ContCont(x = Results, Min = 0.02, Max = 1)
```



The lines connecting the pairs (S, T_1) and (S, T_0) give the estimates for the identifiable correlations ρ_{T_1S} and ρ_{T_0S} . The lines are colored in red when the correlations are negative and thicker lines are used for stronger correlations. As noted earlier, the negative correlations indicate that patients with more (less) withdrawal symptoms take less (more) heroin in the 30 days interval after the treatment in both treatment conditions (T_0/T_1). The line connecting the pair (T_0, T_1) gives the median for the unidentified correlation between the counterfactuals that lead to the specified range of the ρ_ψ values.

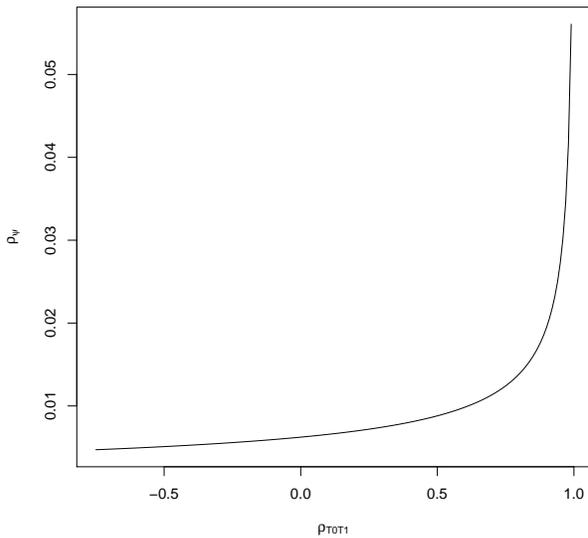
As it can be seen in the figures, a ρ_ψ value between 0 and 0.02 is typically associated with $\rho_{T_0T_1} = 0.08$. This suggests that T_0 , the patient's outcome under the control treatment (number of days that heroin is taken if the patient received the control treatment), carries essentially no information on T_1 , the patient's outcome under the experimental treatment (number of days that heroin is taken if the patient received the experimental treatment). On the other hand, a ρ_ψ between 0.02 and 0.06 is typically associated with $\rho_{T_0T_1} = 0.95$. This suggests that the patient's outcome under the control treatment carries a substantial

amount of information on his outcome under the experimental treatment.

Although the data alone do not allow to discriminate between both scenarios (as ρ_{TOT1} is not identifiable), expert opinion may be used to evaluate the biological plausibility of these two settings. For example, given that both treatments are similar and T_0, T_1 are essentially repeated measures on the same patient, it may be argued that the plausibility of independent potential outcomes is biologically questionable.

The relationship between ρ_{TOT1} and ρ_ψ can be further studied using the `Effect.TOT1=TRUE` argument in the `plot()` call:

```
plot(Results, EffectTOT1=TRUE, PCA=FALSE)
```



As it can be seen, ρ_ψ is a monotonically increasing function of ρ_{TOT1} , but even when ρ_{TOT1} is close to 1 ρ_ψ is only about 0.06. A table containing all the combinations of ρ_{TOT1}, ρ_{TOS} and ρ_{T1S} that lead to valid (positive definite) covariance matrices and their corresponding PCA values can be obtained using the following command:

```
TableResults <- cbind(Results$Pos.Def, Results$PCA)[order(Results$PCA),]
```

For example, the combinations of ρ_{TOT1}, ρ_{TOS} and ρ_{T1S} that lead to the lowest and highest ρ_ψ values can be obtained in the following way:

```
head(TableResults) # lowest PCC values
##      TOT1      TOS      T1S Results$PCA
## 1 -0.75 -0.3699 -0.3367 0.0047133521
## 2 -0.74 -0.3699 -0.3367 0.0047268584
## 3 -0.73 -0.3699 -0.3367 0.0047404814
## 4 -0.72 -0.3699 -0.3367 0.0047542229
## 5 -0.71 -0.3699 -0.3367 0.0047680845
## 6 -0.70 -0.3699 -0.3367 0.0047820682

tail(TableResults) # highest PCC values
##      TOT1      TOS      T1S Results$PCA
## 170 0.94 -0.3699 -0.3367 0.024980880
```

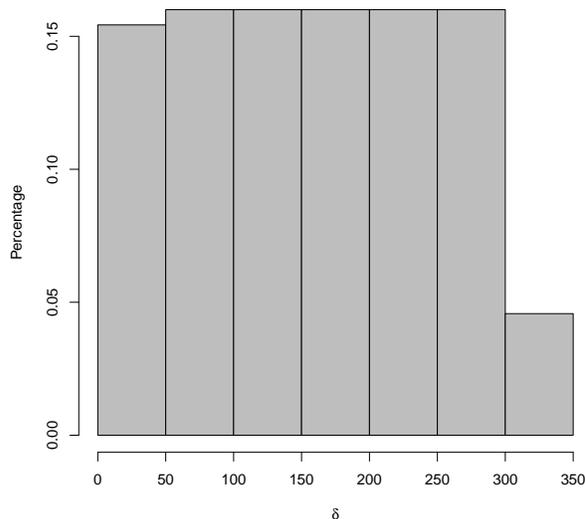
```
## 171 0.95 -0.3699 -0.3367 0.027261200
## 172 0.96 -0.3699 -0.3367 0.030307023
## 173 0.97 -0.3699 -0.3367 0.034672001
## 174 0.98 -0.3699 -0.3367 0.041703753
## 175 0.99 -0.3699 -0.3367 0.056067743
```

From the previous outcomes it can be learned that $\rho_\psi = 0.0047$ when $\rho_{TOT1} = -0.75$ and $\rho_\psi = 0.0561$ when $\rho_{TOT1} = 0.99$ (about 12 times higher).

4 Is there a good pretreatment predictor?

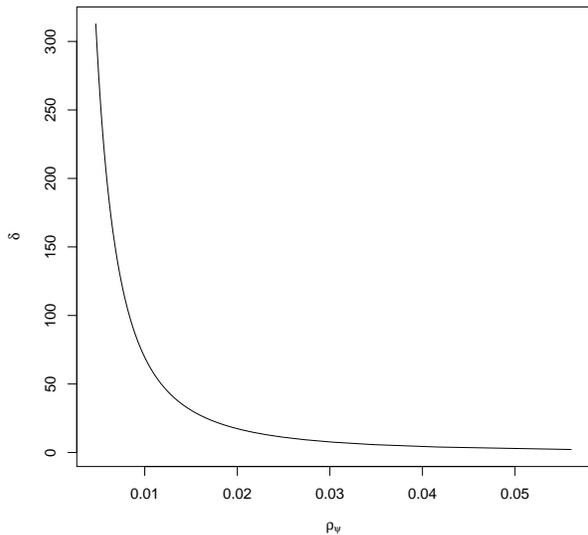
A plot that depicts all the prediction mean squared error (PMSE) values in the opiate/heroin study based on the fitted Results object can be obtained by using the Good.Pretreat=TRUE argument in the plot() function:

```
plot(Results, Good.Pretreat=TRUE, PCA=FALSE)
```



In most settings PMSE lies between 50 and 300 and thus S_j predicts ΔT_j with a mean squared error that lies between about 7 and 17.5 days. Given that T = the number of days that heroin is used in a 30-day interval, the previous level of accuracy seems to be rather unsatisfactory. A figure that shows the relation between ρ_ψ and δ can be obtained using the following command:

```
plot(Results$GoodSurr$PCA, Results$GoodSurr$delta, xlab=expression(rho[psi]),
ylab=expression(delta), col=0)
lines(Results$GoodSurr$PCA, Results$GoodSurr$delta)
```



As expected, PMSE is a monotonically decreasing function of ρ_ψ . It may also be useful to examine the proportion of ‘realities’ for which a desired prediction accuracy is achieved. For example, the following code can be used to obtain the percentage of realities in which the PMSE ≤ 50 (corresponding with an average prediction error of about 7 days):

```
length(Results$GoodSurr$delta[Results$GoodSurr$delta<=50])/
length(Results$GoodSurr$delta)

## [1] 0.15428571
```

Thus the desired prediction accuracy is achieved in only about 15% of the realities compatible with the data.

The plausibility of finding a good pretreatment predictor The function `GoodPretreatContCont` (Good Pretreatment predictor in the Continuous Continuous setting) allows examining the plausibility of finding a good pretreatment predictor S . The function requires the following arguments:

- `T0T0=` and `T1T1=`: the variances of the true endpoint in the control (σ_{T0T0}) and experimental conditions (σ_{T1T1}).
- `Delta=`: the upper bound for the PMSE.
- `T0T1=`: a vector (or scalar) that specifies values for the unidentifiable correlation between the potential outcomes ρ_{T0T1} .

For example, suppose that one wants to examine the plausibility of finding a pretreatment predictor S that allows predicting ΔT_j with PMSE = 50 (corresponding with an average prediction error of about 7 days). The following code can be used for that purpose:

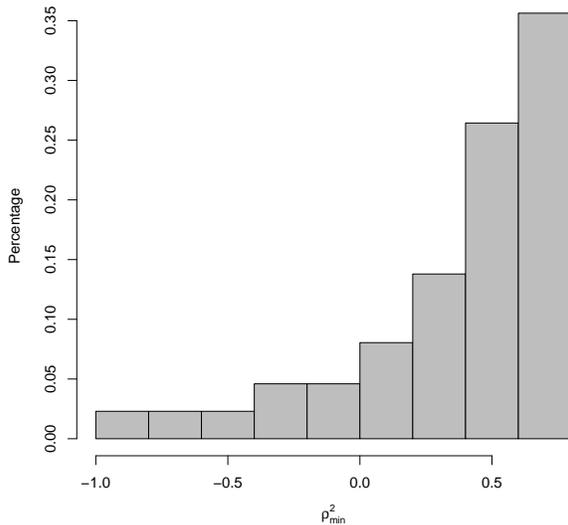
```
MinSurr <- GoodPretreatContCont(T0T0=83.2872, T1T1=95.6189, Delta=50,
T0T1=seq(from=0, to=1, by=.01))
```

The results can be examined by applying the `summary()` and `plot()` functions to the fitted `MinSurr` object:

```
summary(MinSurr)
```

```
##  
## Function call:  
##  
## GoodPretreatContCont(TOT0 = 83.2872, T1T1 = 95.6189, Delta = 50,  
##   TOT1 = seq(from = 0, to = 1, by = 0.01))  
##  
##  
##  
## # Rho2.Min results summary (Inf values are excluded)  
## #~~~~~  
##  
## Mean (SD) Rho^2_min: 0.3585 (0.3964) [min: -0.9675; max: 0.7205]  
##  
## Quantiles of the Rho2.Min distribution:  
##  
##      5%      10%      20%      50%      80%      90%  
## -0.51205111 -0.22719865  0.10877474  0.51056900  0.66262713  0.69428703  
##      95%  
##  0.70799085  
##  
## Note. Some Rho2.Min values were negative. This indicates that the PMSE is so large that  
## any Rho2.Min value suffices to achieve the desired prediction accuracy.
```

```
plot(MinSurr)
```



As it can be seen, about 80% of the ρ_{min}^2 values were above 0.1088. This indicates that a candidate S should produce a ρ_{ψ} of (at least) about $\sqrt{0.1088} = 0.3298$ to achieve the desired level of accuracy (PMSE=50) in the prediction of the individual causal effects on T .

5 Predicting ΔT based on S in an individual patient j

In practice, one is interested in the prediction of a patient's individual causal treatment effect (ΔT_j) given the patient's observed S_j . The function `Predict.Treat.ContCont` (Predict Treatment effect in the Continuous case) is useful in this context. It requires the following arguments:

- `x=`: a fitted object of class `PCA.ContCont`.
- `S=`: the observed pretreatment value S_j for a patient.
- `Beta=`: the expected causal treatment effect on T . Under SUTVA, $\beta = E(T_{1j} - T_{0j})$ can be estimated as $\beta = E(T_j | Z_j = 1) - E(T_j | Z_j = 0)$, i.e., the difference between the observed means of T_j in the experimental and control treatment groups, respectively.
- `SS=`: the variance of S .
- `mu_S=`: the mean of S .

In the heroin/opiate detoxification dataset, $\hat{\beta} = -0.9314$, $\hat{\sigma}_{SS} = 356.2158$, and $\hat{\mu}_S = 84.7994$. Notice that the negative β indicates that the average number of days that heroin is used post-treatment was slightly smaller in the experimental treatment group ($\hat{\mu}_E = 11.6437$) than in the control treatment group ($\hat{\mu}_C = 12.5751$) – albeit the difference was not significant, $p = 0.632$.

Suppose that a patient scores 60 on the COWS (i.e., a low level of opiate withdrawal symptoms). The following code can be used to predict ΔT_j for this patient:

```
Pred_S_60 <- Predict.Treat.ContCont(x=Results, S=60, Beta=-0.9314, SS=356.2158, mu_S=84.7994)
```

The results can be examined by applying the `summary()` function to the fitted object `Pred_S_60`:

```
summary(Pred_S_60)

##
## Function call:
##
## Predict.Treat.ContCont(x = Results, S = 60, Beta = -0.9314, SS = 356.2158,
##   mu_S = 84.7994)
##
## # Predicted (Mean) Delta_T_j | S_j
## #~-----
##
## -1.0409279
##
##
## # Variances and 95% support intervals of Delta_T_j | S_j for different values of rho_TOT1
## #~-----
##
##           rho_TOT1   Var Delta_T_j | S_j   95% supp. int. around Delta_T_j | S_j
##
## (min. value)   -0.750           312.760           [-35.702911; 33.621056]
## (max. value)    0.990            2.203           [-3.9502469; 1.8683912]
## (median value)  0.120            157.481          [-25.636834; 23.554979]
## (mean value)   0.120            157.481          [-25.636834; 23.554979]
##
##
##
```

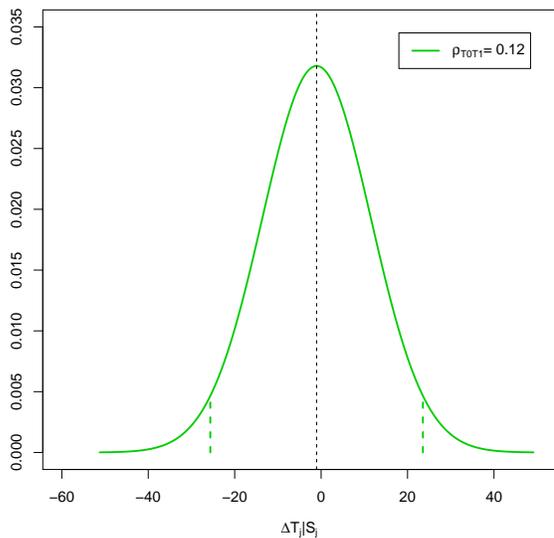
```
## # Proportion of 95% support intervals for Delta_T_j | S_j
## that include 0, are < 0, and are > 0
## #-----
##
## 0 included in support interval: 1      (obtained for rho_TOT1 values in range [-0.75; 0.99])
## Entire support interval below 0: 0
## Entire support interval above 0: 0
```

The expected $\Delta T_j|S_j = 60$ equals -1.0409 . Thus, the patient is expected to have about 1 more heroin-free day in the post-treatment interval with the experimental treatment than with the control treatment. Importantly, the expected $\Delta T_j|S_j = 60$ remains constant no matter what assumption regarding ρ_{TOT1} is made. However, the assumed ρ_{TOT1} does affect the *variance* of $\Delta T_j|S_j = 60$. This can be observed in the second part of the output, where variances and 95% support intervals for $\Delta T_j|S_j = 60$ are given for different ρ_{TOT1} values. For example, the 95% support interval around $\Delta T_j|S_j = 60$ is $[-35.7029; 33.6211]$ when it is assumed that $\rho_{TOT1} = -0.750$ (the minimum value of ρ_{TOT1} that is compatible with the observed data), whereas the 95% support interval around $\Delta T_j|S_j = 60$ is $[-3.9502; 1.8684]$ when it is assumed that $\rho_{TOT1} = 0.990$ (the maximum value of ρ_{TOT1}).

The final part of the output provides an overview of the proportion of support intervals for $\Delta T_j|S_j = 60$ that included 0 (no difference between treatments expected for the patient), that lay entirely below 0 (experimental treatment more beneficial), and that lay entirely above 0 (control treatment more beneficial). As it can be seen, the 95% support interval of $\Delta T_j|S_j = 60$ included 0 in all cases.

The results can be graphically displayed by applying the `plot()` function to the fitted object `Pred_S_60`. By default, the distribution of $\Delta T_j|S_j$ is shown for the median ρ_{TOT1} value, i.e., for $\rho_{TOT1} = 0.120$. The following command can be used to obtain the plot:

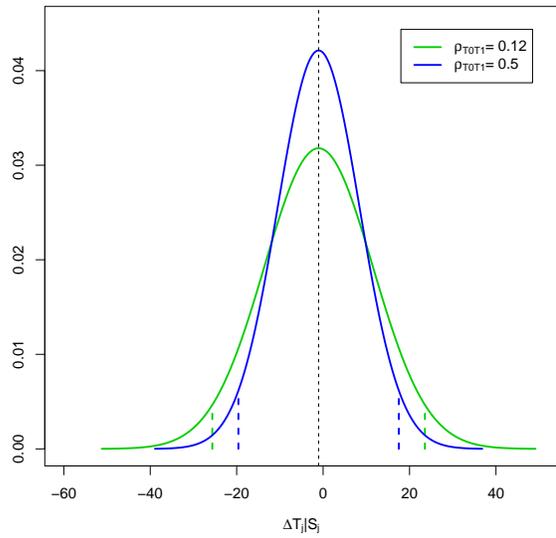
```
plot(Pred_S_60, xlim=c(-60, 50), ylim=c(0, .035))
```



The vertical black dashed line is the expected $\Delta T_j|S_j = 60$ value, and the dashed green lines depict the 95% support interval. In line with the earlier results, the 95% support interval for $\Delta T_j|S_j = 60$ assuming $\rho_{TOT1} = 0.120$ contains 0 and thus no difference between both treatments is expected for the patient.

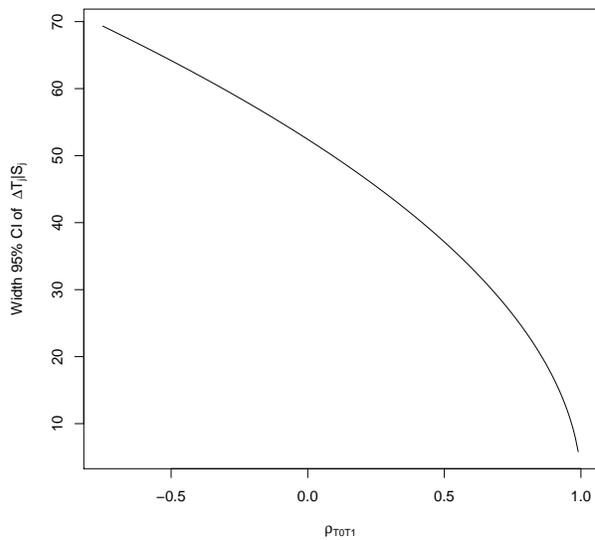
It is also possible to request the 95% support interval for a particular value of ρ_{TOT1} by using the `Specific.TOT1=` argument in the `plot()` call. For example, the 95% support interval around $\Delta T_j|S_j$ assuming $\rho_{TOT1} = 0.5$ can be requested using the following command:

```
plot(Pred_S_60, Specific.TOT1 = 0.5, xlim=c(-60, 50), ylim=c(0, .045))
```



As expected, the width of the 95% support interval decreases when ρ_{TOT1} increases. A plot that shows the relation between ρ_{TOT1} and the width of the 95% support interval can be obtained with the following command:

```
plot(x=Pred_S_60$TOT1, y= (sqrt(Pred_S_60$Var_Delta.T_S)*1.96)*2, type="l",
xlab=expression(rho[TOT1]), ylab=expression(paste("Width 95% CI of ",
Delta, T[j], "|", S[j])))
```



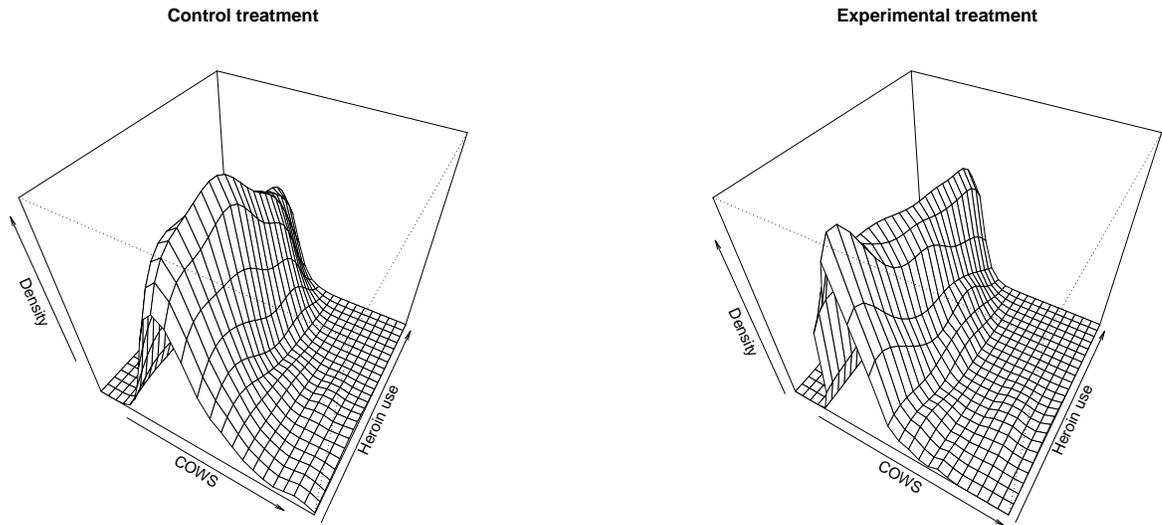
The plot shows that if one wants to be 95% confident that the true $\Delta T_j | S_j = 60$ deviates at most 10 days (in positive or negative direction) from the expected $\Delta T_j | S_j$ (width of the 95% support interval ≤ 20), then $\rho_{TOT1} \geq 0.85$.

Appendix B. Distributions of the endpoints

In this Appendix the assumptions underlying the causal inference model introduced in Alonso *et al.* (2015) are evaluated using the identifiable marginals. Notice that although the entire model is basically unidentifiable, the bivariate normality of (S_0, T_0) and (S_1, T_1) and univariate normality of the corresponding marginals can be assessed with the data at hand. The results were similar for all five multiply imputed datasets. For succinctness, attention is restricted here to only one of the multiply imputed datasets.

Bivariate density plots for (S_0, T_0) and (S_1, T_1) are shown in Figure 1. Histograms S and T in the control and experimental treatment conditions are shown in Figure 2.

As it can be seen, the normality assumption seems reasonable for COWS but is violated for T in both treatment conditions. Likewise, bivariate normality is questionable.



[H]

Figure 1: Bivariate density plots of $S = \text{COWS}$ and $T = \text{heroin use}$ in the control (left) and experimental (right) treatment groups.

Obviously, the findings of the previous sections are predicated on the assumption that the potential outcomes are continuous and normally distributed. However, many of these results will still be valid for non-normal potential outcomes, although their interpretation will change if the normality assumption is questioned. For instance, the correlation between the individual causal treatment effect and the pretreatment predictor could still be quantified using ρ_ψ , but these expressions could not be interpreted any longer as the predictive causal association. Indeed, even though ρ_ψ will still be a valid measure of causal correlation, the equivalence between association and correlation will be broken if the normality assumption is dropped. One may also consider transforming the response variable in order to make the distributional assumptions more plausible, but this may affect the interpretation of the results. In general, departures from normality will have just a mild effect on the estimation of the parameters of interest but, as already pointed out, the interpretation of the results will become more restricted and limited.

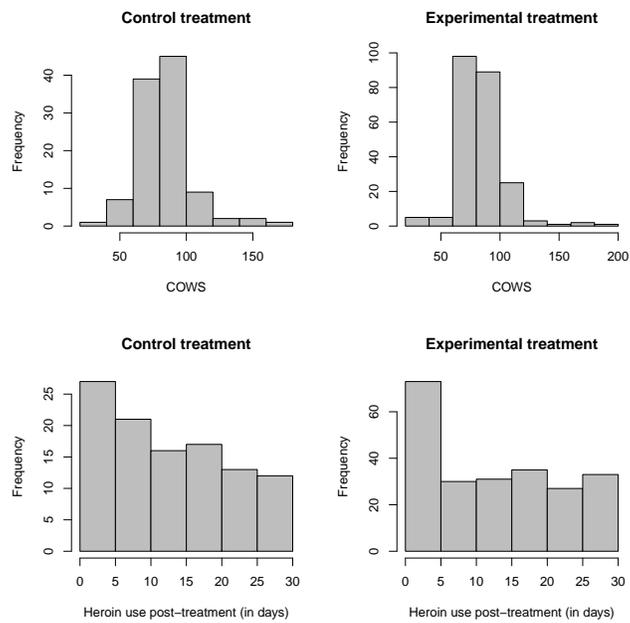


Figure 2: Histograms of $S = \text{COWS}$ and $T = \text{heroin use}$ in the control (left) and experimental (right) treatment groups.

References

Alonso, A., Van der Elst, W., and Molenberghs, G. (2015). Validating predictors of therapeutic success: A causal inference approach. *Statistical Modelling*.