

EffectLite: User's Manual

A Program for the Uni- and Multivariate Analysis of
Unconditional, Conditional and Average Mean Differences
Between Groups

Rolf Steyer and Ivailo Partchev
University of Jena

March 2008

Version 3 for Mplus

A StatLite product
www.statlite.com



Contents

Preface	5
1 The scope of EffectLite	9
1.1 Analysis without covariates	9
1.1.1 Two groups	9
1.1.2 More than two groups	9
1.1.3 Null hypotheses	10
1.1.4 Optional assumptions on outcome (co)variances	10
1.2 Analysis with covariates	11
1.2.1 Two groups	11
1.2.2 Nonlinear effect functions	12
1.2.3 More than two groups	13
1.2.4 Average effects	14
1.2.5 Adjusted means	14
1.2.6 Average effects given a treatment condition	14
1.2.7 Null hypothesis: No average treatment effects	15
1.2.8 Null hypothesis: intercept function is constant	15
1.2.9 Null hypothesis: no interaction	16
1.2.10 Null hypothesis: No treatment effects	16
1.2.11 Null hypotheses about conditional effects	16
1.2.12 Several outcome variables	16
1.2.13 Sampling models	16
1.2.14 Optional assumptions about covariates	17
1.2.15 Residual variances and covariances	17
1.3 Latent covariates and latent outcome variables	17
1.3.1 Latent covariates	17
1.3.2 Two groups	18
1.3.3 More than two groups	19
1.3.4 Latent outcome variables	20
1.3.5 Optional assumptions about latent outcome variables	21
1.3.6 Method factors	21
1.4 Tests of unconfoundedness	21
1.4.1 Analysis without covariates	22
1.4.2 Analysis with covariates	23
1.5 Missing values	25
1.6 Inferential techniques in EffectLite	25

2 Detailed program description	29
2.1 Before you begin	29
2.2 A normal run of EffectLite	29
2.3 The Open Data screen	30
2.4 The Select Variables screen	31
2.5 The Analysis Options screen	31
2.6 The Measurement Model screen	33
2.7 The Main Results screen	33
2.8 The Conditional Effects screen	34
2.9 The Technical Output screen	34
2.10 The Wizard Stuff screen	35
2.11 The Mplus Locations screen	35
2.12 The Help screen	35
2.13 Types of data files you may use	36
2.14 The CSV data format	36
2.15 Models available in the current version	37
2.16 Types of variables	38
2.17 Output: Model without covariates	38
2.17.1 General information	38
2.17.2 Overall tests of significance	39
2.17.3 Detailed analysis of mean differences	40
2.17.4 Group means of the outcome variables	40
2.18 Output: Model with covariates	41
2.18.1 General information	41
2.18.2 Measurement model	41
2.18.3 Group means of the covariates	42
2.18.4 Group means of the outcome variable(s)	43
2.18.5 Overall tests of significance	43
2.18.6 Average treatment effects	44
2.18.7 Deviations of the adjusted means from their averages over groups	44
2.18.8 Average effects given a treatment condition	44
2.18.9 Effect functions	45
References	46

Preface

What is EffectLite?

EffectLite is a PC-program for the uni- and multivariate analysis of mean differences between groups on outcome (or response) variables in designs with and without covariates. If there is no covariate, then the unconditional mean differences are analyzed. If there is at least one covariate, the conditional and the average (or 'adjusted') mean differences are analyzed (see section 1.2.5).

EffectLite has been specially developed for the analysis of covariate-treatment-outcome designs with two or more treatment groups, one of which can be a control group. Covariates can be quantitative pretests and/or they can be qualitative blocking factors. There can be several dependent (outcome) variables, and a treatment variable with several levels, such as 'control', 'treatment 1', 'treatment 2'.

EffectLite has been developed by Rolf Steyer and Ivailo Partchev at the University of Jena, Germany. It acts as a pre- and post-processor for **Mplus** (Version 3 or later).¹ Although the use of **EffectLite** and the interpretation of its results do not require detailed knowledge of **Mplus**, the user does need to be acquainted with the theoretical background provided in this program description. A convenient introduction to **EffectLite** and its theoretical background is also available as a video online-course at www.metheval.uni-jena.de.

Why use EffectLite?

Outside the perfectly randomized experiment, unconditional mean differences with respect to an outcome variable Y can be very misleading if used for the evaluation of a treatment or intervention. It is well known that there can be a high positive (*unconditional*) *mean difference* in the outcome variable Y between an experimental ($X=1$) and a control group ($X=0$) in the total population, even though there are high negative (*conditional*) *mean differences* in all subpopulations,² such as 'males' ($Z=1$) and 'females' ($Z=0$), constituting the total population.

Phenomena like this have led to the development of statistical procedures such as the analysis of covariance, in which the conditional mean differences in Y between treatment and control are analyzed given the covariate Z . However, analysis of covariance does not allow for *interactions* between the covariate Z and the treatment variable X . Such an interaction means that mean differences (in the outcome variable Y) between the groups represented by X are different for different values (such as 'male' and 'female' or 'high score on pretest' or 'low score on pretest') of the (qualitative or quantitative) covariate Z . Oftentimes these interactions and the associated conditional mean differences are of considerable relevance for substantive science, because they describe which treatment is good or bad for which type of persons.

¹**Mplus** (Muthén & Muthén, 1998 - 2004) is a copyright protected commercial program.

²Simpson's paradox (see e.g. Steyer, 1992, 2003; Steyer, Partchev, Kröhne, Nagengast, & Fiege, 2007) is an example in case.

Aside from the conditional effects, one is often still interested in the average effect of the treatment as compared to the control. Examples like the Simpson paradox show that the unconditional mean difference does not represent the average effect if the treatment variable X and the covariate Z are correlated and Z has an effect on Y . Instead, the *average effect* or *average mean difference* has to be defined by the average of the conditional mean differences (in the outcome variable Y) between treatment and control, given the values of the covariate. These averages of the conditional mean differences are the ‘average effects’ analyzed in **EffectLite**. These average effects are also the differences between the ‘adjusted means’ which are also computed. If it can be assumed that the covariate-treatment regression is unbiased the average effects can also be interpreted as the average causal effects in the Neyman-Rubin sense (see Steyer et al., 2007).

Although there are many programs designed to deal with the issues mentioned above, **EffectLite** has the advantage of being more general than comparable programs because:

- (a) It does not assume homogeneity of variances (in the univariate case with a single outcome variable) or of covariance matrices (in the multivariate case with two or more outcome variables) of the outcome variables between treatment groups.
- (b) It allows analyzing mean differences between groups with respect to several manifest outcome variables, one or more *latent outcome variables*, and a mixture of the two kinds of outcome variables.
- (c) It allows analyzing conditional and average effects with respect to several manifest covariates or with respect to one or more *latent covariates*, and a mixture of the two kinds of covariates.
- (d) It estimates and tests average effects for non-orthogonal analysis of variance designs, provided that the covariates are specified as qualitative indicator variables. Other programs typically do not test the average effect at all,³ and they do not treat the covariates as stochastic regressors, which usually leads to invalid tests of the average effect.
- (e) **EffectLite** produces results which are easily interpretable in the analysis of conditional and of average effects (mean differences) between groups.
- (f) **EffectLite** estimates and tests the *average causal effects* (Rubin, 1974, 1978), provided that the covariates fulfil certain assumptions (see Steyer et al., 2007, for details).

New features in version 3

Version 3 of **EffectLite** introduces some new features:

- Cases with missing data can be deleted listwise (as before), or they can be kept; in the latter case, full information maximum likelihood (FIML) estimation in **Mplus** will be switched on automatically (cases with missing data on the treatment variable are always deleted).

³This includes all programs for nonorthogonal analysis of variance computing Type I, II, III and IV sums of squares.

- In models involving latent variables, there are two major improvements in defining the measurement model:
 - There is now a common measurement model for all latent variables instead of separate models for latent covariates and latent outcome variables. This makes it possible to define latent variables that are attached to both outcome variables and covariates, as in the case of *method factors*. **EffectLite** will treat such latent variables as covariates.
 - Loadings can be fixed to any numerical value (not just to 0 or 1), and equality constraints may be imposed on two or more loadings.
- Another, less obvious, new feature is that all chi-squared tests reported by **EffectLite** are now computed as Wald tests rather than likelihood ratio tests. The two kinds of tests are asymptotically equivalent, which means that they behave increasingly similar as sample sizes increase. Users will notice that the **Technical Output** screen now displays **Mplus** output for just one model (rather than four); this is because Wald tests can be computed from the unconstrained model alone.
- Starting with Version 3.1, there is a possibility to compute, examine, and save conditional treatment effects for each row of the data file.

Thanks and acknowledgements

We gratefully acknowledge substantive inputs by Andreas Klein, Safir Yousfi, Ulf Kröhne, Benjamin Nagengast, Karl Jöreskog, Bengt Muthén, Rüdiger Mutz, and Gerhard Mels.

EffectLite is written in the TCL/TK language (John Osterhout), using add-on modules by Jeffrey Hobbs, Edward C. Hume, Arjen Markus, Brent Welch, and Jean-Claude Wippler. Third-party data formats are translated with **DataLoad** by David Baird.

EffectLite is provided “as is” and without any express or implied warranties, including, without limitation, the implied warranties of merchantability and fitness for a particular purpose.

Jena, March 2008

Rolf Steyer and Ivailo Partchev

1 The scope of EffectLite

In this chapter, we outline the theoretical foundations of **EffectLite**. This is inevitably a very concise introduction, and the interested reader is referred to the much fuller discussion in Steyer et al., 2007. The focus here is on the available analyses, and on the types of statistical hypotheses tested by **EffectLite**. Explanations on how to actually use the program and how to read its output are postponed until chapter 2.

The first important distinction is between analyses without or with covariates. Covariates are variables that may correlate with the treatment variable, but are not affected (influenced) by the treatment variables. Often the covariate is a quantitative pretest, but it could also be a qualitative variable such as gender, or it could be a multivariate covariate consisting of a quantitative pretest and gender, for instance.

1.1 Analysis without covariates

In this section we will consider analyses without covariates. This is very close to what is known from uni- and multivariate analysis of variance with one between-group factor. The only extensions are: **EffectLite** also analyzes latent outcome variables and it does not presume homogeneity of variances (or of the variance-covariance matrices in the multivariate case) between treatment groups. In section 1.2 we will turn to corresponding analyses with covariates.

1.1.1 Two groups

Assume for a moment that there is a single numerical outcome or response variable Y , and a dichotomous variable X with values 0 (for the control or reference condition) and 1 (for the treatment condition). In this case, the regression $E(Y | X)$ can always be represented as a linear function of X :

$$E(Y | X) = \alpha_0 + \alpha_1 \cdot X, \quad (1.1)$$

with the *intercept*

$$\alpha_0 = E(Y | X=0) \quad (1.2)$$

and the *slope*

$$\alpha_1 = E(Y | X=1) - E(Y | X=0). \quad (1.3)$$

Hence, in the two-group case, the slope α_1 can be identified by the difference between the conditional expectations of Y between the two groups, and the intercept α_0 is the expectation $E(Y | X=0)$ of Y in the reference group coded by $X=0$.

1.1.2 More than two groups

If X has three values, 0 (for control or reference), 1 (for treatment 1) and 2 (for treatment 2), we need *two* indicator variables $I_{X=1}$ and $I_{X=2}$ indicating whether the observational unit is assigned to treatment 1 ($I_{X=1} = 1$ and $I_{X=2} = 0$) or to treatment 2 ($I_{X=1} = 0$ and $I_{X=2} = 1$).

Otherwise, $I_{X=1} = 0$ and $I_{X=2} = 0$, which means that the observational unit is assigned to the control (reference) condition coded with $X=0$. Hence, in the three-group case, the regression can be written:

$$E(Y | X) = \alpha_0 + \alpha_1 \cdot I_{X=1} + \alpha_2 \cdot I_{X=2}. \quad (1.4)$$

The intercept α_0 is now

$$\alpha_0 = E(Y | X) = 0, \quad (1.5)$$

and the two *slopes* are the two mean differences:

$$\alpha_1 = E(Y | X=1) - E(Y | X=0), \quad (1.6)$$

$$\alpha_2 = E(Y | X=2) - E(Y | X=0). \quad (1.7)$$

Hence, with α_1 we compare the expectation of Y under treatment 1 to its expectation under control, and with α_2 we compare the expectation of Y under treatment 2 to its expectation under control.

The generalization to $J + 1$ treatment conditions is obvious:

$$E(Y | X) = \alpha_0 + \alpha_1 \cdot I_{X=1} + \dots + \alpha_J \cdot I_{X=J}, \quad (1.8)$$

where each coefficient α_j , $j = 1, \dots, J$, compares the expectation of Y in treatment condition j to its expectation in the control condition $X=0$.

1.1.3 Null hypotheses

The null hypothesis of no effects (zero mean differences) between all groups simultaneously is then:

$$H_0 : \alpha_1 = \alpha_2 = \dots = \alpha_J = 0. \quad (1.9)$$

EffectLite also provides a significance test for each single effect α_j to be zero.

Note that you may also consider *several outcome variables* Y_i , $i = 1, \dots, I$, instead of just a single outcome variable Y . In such a case, **EffectLite** simultaneously tests the hypothesis (1.9) for *all* outcome variables.

As mentioned before, **EffectLite** does not assume homogeneity of variances (univariate case with a single outcome variable) or covariance matrices (multivariate case with two or more outcome variables) of the outcome variables between treatment groups. It also allows analyzing mean differences between groups with respect to one or more *latent outcome variables*, and with respect to a set of manifest *and* latent outcome variables. Last but not least, it allows analyzing conditional and average effects with respect to one or more *latent covariates*, and with respect to a set of the two kinds of covariates.

1.1.4 Optional assumptions on outcome (co)variances

In traditional univariate and multivariate analysis of variance, the variance (univariate case) and the covariance matrix (multivariate case) of the outcome variables are assumed to be invariant between groups. This assumption can also be made in **EffectLite**. However, the default is that the variances and covariances of the outcomes variables differ between groups. In fact, if we assume that the effects of the treatments differ between individuals, we expect heterogeneous variances and covariances of the outcome variables between treatment groups, both in the randomized experiment and in observational studies (see Steyer et al., 2007, for details).

1.2 Analysis with covariates

EffectLite also analyses models with covariates. A covariate, which will be denoted by Z , is a variable that may affect the treatment variable, e.g., by determining the treatment probabilities $P(X=j | Z=z)$. However, a covariate is defined such that it *can not be affected by the treatment variable* X . Some typical examples in case are

- (a) a manifest quantitative pretest,
- (b) a manifest qualitative variable characterizing the observational unit such as gender, race, or educational status,
- (c) a latent quantitative pretest (fallibly measured, e.g., by two parallel forms),
- (d) a second qualitative treatment variable that may be correlated with the focus treatment variable, but which is not affected by the focus treatment variable,
- (e) a squared covariate or a product $Z_1 \cdot Z_2$ of two covariates,
- (f) a combination of all types of covariates mentioned above.

Note, however, that qualitative covariates such as sex or race can be accommodated in the present version only by computing dummy variables and specifying them as new covariates, before **EffectLite** is used. Whatever you select as components of the covariate vector $Z \equiv (Z_1, \dots, Z_Q)$, **EffectLite** will work with a linear intercept function $g_0(Z)$ and linear effect functions $g_1(Z), \dots, g_J(Z)$ (see below). Also note that, if there are two treatment variables (or factors in the analysis of variance sense), one may define a new treatment variable, the values of which are the pairs of values of the original two treatment variables. This would allow for a simultaneous test of the effects of two treatment variables (factors).

1.2.1 Two groups

Again, assume for a while that there is a single continuous outcome variable Y , a dichotomous treatment variable X with values 0 (for control) and 1 (for treatment). Additionally, we now consider a (uni- or multivariate) covariate Z . In this case, the regression $E(Y | X, Z)$ can always be written:

$$E(Y | X, Z) = g_0(Z) + g_1(Z) \cdot X, \quad (1.10)$$

with (an unknown) *intercept function* $g_0(Z)$ and (an unknown) *slope or effect function* $g_1(Z)$.

If the covariate Z is also dichotomous with values 0 and 1, the functions $g_0(Z)$ and $g_1(Z)$ can always be parameterized as linear functions of Z , i.e., $g_0(Z) = \gamma_{00} + \gamma_{01}Z$ and $g_1(Z) = \gamma_{10} + \gamma_{11}Z$. Equation (1.10) may then also be written

$$E(Y | X, Z) = (\gamma_{00} + \gamma_{01}Z) + (\gamma_{10} + \gamma_{11}Z) \cdot X. \quad (1.11)$$

The first index of the regression coefficients γ refers to the intercept function $g_0(Z)$ and the effect function $g_1(Z)$, respectively. Considering the two conditional regressions of Y on X given $Z=z$ yields:

$$E_{Z=0}(Y | X) = \gamma_{00} + \gamma_{10}X \quad (1.12)$$

$$E_{Z=1}(Y | X) = (\gamma_{00} + \gamma_{01}) + (\gamma_{10} + \gamma_{11}) \cdot X. \quad (1.13)$$

Hence, both the intercept and the slope of the conditional regressions of Y on X given $Z=z$ depend on values z of the covariate.

Interpretation of the regression coefficients. The interpretation of the regression coefficients is simple: γ_{10} is the *conditional effect (slope)* of X if $Z = 0$, and $\gamma_{10} + \gamma_{11}$ is the *conditional effect (slope)* of X if $Z = 1$. Furthermore, γ_{00} is the *intercept of the conditional regression* of Y on X given $Z = 0$, and $\gamma_{00} + \gamma_{01}$ is the *intercept of the conditional regression* of Y on X given $Z = 1$. The parameter γ_{11} is called the *interaction parameter*: If $\gamma_{11} = 0$, then the conditional effects of X on Y given Z do *not* depend on the values z of Z , otherwise the conditional effects of X are modified by Z . Hence, we call Z a *modifier* and $\gamma_{10} + \gamma_{11}Z$ the *conditional effect (slope) function* or, simply, the *effect (slope) function*. A synonym for ‘modifier’ is *moderator* and the ‘effect function’ is sometimes called *moderator function* (see e.g. Saunders, 1956 or Baron & Kenny, 1986).

1.2.2 Nonlinear effect functions

The models described above can be easily generalized to include nonlinear effect functions. If Z is not dichotomous, the functions $g_0(Z) \equiv \gamma_{00} + \gamma_{01}Z$ and $g_1(Z) \equiv \gamma_{10} + \gamma_{11}Z$ in Equation (1.11) do not necessarily yield a saturated parameterization of the regression $E(Y | X, Z)$. In this case, the linear functions in Equation (1.11) may have to be replaced by more general functions, e.g., by polynomial functions of a higher degree, such as

$$E(Y | X, Z) = (\gamma_{00} + \gamma_{01}Z + \gamma_{02}Z^2) + (\gamma_{10} + \gamma_{11}Z + \gamma_{12}Z^2) \cdot X, \quad (1.14)$$

or, if Z is discrete, by a parameterization using indicator (dummy) variables for the values of the covariate Z , such as

$$\begin{aligned} E(Y | X, Z) &= (\beta_{00} + \beta_{01} \cdot I_{Z=1} + \beta_{02} \cdot I_{Z=2}) \\ &+ (\beta_{10} + \beta_{11} \cdot I_{Z=1} + \beta_{12} \cdot I_{Z=2}) \cdot X. \end{aligned} \quad (1.15)$$

In this equation we assume that Z has only the values 0, 1, and 2. We denote the parameters by β s in order to make clear that they are different from the γ s occurring in Equation (1.14).

Equations (1.14) and (1.15) represent saturated parameterizations of the regression $E(Y | X, Z)$ if X has only two and Z has only three different values. These values can also be qualitative, representing three educational backgrounds, for instance. This parameterization would be appropriate, e.g., in a 2×3 -factorial design with correlated factors, in which X is the treatment variable. In the present version, however, the user has to use indicator or dummy variables replacing the original qualitative covariate before using **EffectLite**.

EffectLite uses functions for $g_0(Z)$ and $g_1(Z)$ which are linear in the parameters, whatever the covariates provided by the user might be. If there are, e.g., two covariates Z_1, Z_2 , and the user provides also the product $Z_1 \cdot Z_2$ as a third covariate, a special case of Equation (1.10) is also

$$\begin{aligned} E(Y | X, Z) &= (\gamma_{00} + \gamma_{01}Z_1 + \gamma_{02}Z_2 + \gamma_{03}Z_1 \cdot Z_2) \\ &+ (\gamma_{10} + \gamma_{11}Z_1 + \gamma_{12}Z_2 + \gamma_{13}Z_1 \cdot Z_2) \cdot X, \end{aligned} \quad (1.16)$$

where $Z = (Z_1, Z_2)$. Version 3.0 of **EffectLite** does not produce these covariates by itself. Neither does it produce the indicator variables $I_{Z=1}$ and $I_{Z=2}$ occurring in Equation (1.15).

Summary. If X is a dichotomous variable with values 0 and 1, Equation (1.10) always holds with (an unknown) *conditional intercept function* $g_0(Z)$ and (an unknown) *conditional effect function* $g_1(Z)$, the values $g_0(z)$ and $g_1(z)$ of which are the conditional intercepts and conditional effects (slopes) in the conditional regression of Y on X given $Z=z$, respectively. *Interaction* refers to all constellations in which the effect function $g_1(Z)$ is not a constant, i.e., in which the (conditional) effects of X on Y depend on the values z of Z .

1.2.3 More than two groups

If (the treatment variable) X has $J + 1$ values, we generalize Equation (1.10) to:

$$E(Y | X, Z) = g_0(Z) + g_1(Z) \cdot I_{X=1} + \dots + g_J(Z) \cdot I_{X=J}. \quad (1.17)$$

In this equation, the variables $I_{X=1}, \dots, I_{X=J}$ are indicator (or dummy) variables for the different values of X indicating that $X=j$, for $j = 1, \dots, J$. If, for example, X is a qualitative variable with values 0 (for control), 1 (for treatment 1), and 2 (treatment 2), the variable $I_{X=1}$ could be the indicator (or dummy) variable with values 0 and 1 indicating (with 1) whether or not an observation is made under treatment 1 and $I_{X=2}$ could be the dummy variable with values 0 and 1 indicating (with 1) whether or not an observation is made under treatment 2.

In principle, the functions $g_0(Z), g_1(Z), \dots, g_J(Z)$ in Equation (1.17) can be any functions of Z . Oftentimes, they can be parameterized as polynomial functions or as indicator variables of the type just explained for the three-group case [see Equation (1.15)]. However, as mentioned before, **EffectLite** assumes linearity in the parameters. This means that the user has to provide the variables occurring in Equations (1.14) to (1.16), for instance. Note that each effect function $g_j(Z)$, $j = 1, \dots, J$, contains as its values $g_j(z)$ the $(Z=z)$ -conditional mean differences between group $X=j$ and the group represented by $X=0$ (chosen as a reference or control group), i.e.:

$$g_j(z) = E(Y | X=1, Z=z) - E(Y | X=0, Z=z). \quad (1.18)$$

In the special case in which there is only one single covariate Z and all effect functions $g_j(Z)$, $j = 1, \dots, J$, are linear, Equation (1.17) becomes

$$E(Y | X, Z) = (\gamma_{00} + \gamma_{01}Z) + (\gamma_{10} + \gamma_{11}Z) \cdot I_{X=1} + \dots + (\gamma_{J0} + \gamma_{J1}Z) \cdot I_{X=J}. \quad (1.19)$$

Considering the conditional regressions of Y on Z given $X=j$ in three groups then yields:

$$E_{X=0}(Y | Z) = \gamma_{00} + \gamma_{01} \cdot Z = \gamma_0^{(0)} + \gamma_1^{(0)} Z, \quad (1.20)$$

$$E_{X=1}(Y | Z) = (\gamma_{00} + \gamma_{10}) + (\gamma_{01} + \gamma_{11}) \cdot Z = \gamma_0^{(1)} + \gamma_1^{(1)} Z, \quad (1.21)$$

$$E_{X=2}(Y | Z) = (\gamma_{00} + \gamma_{20}) + (\gamma_{01} + \gamma_{21}) \cdot Z = \gamma_0^{(2)} + \gamma_1^{(2)} Z. \quad (1.22)$$

This shows that the regression coefficients occurring in Equation (1.19) can be computed by conducting a multi-group analysis of the linear regressions of Y on Z within each of the $J + 1$ groups, and then taking the differences between these group-specific linear regressions $\gamma_0^{(j)} + \gamma_1^{(j)} Z$ and $\gamma_0^{(0)} + \gamma_1^{(0)} Z$. This is exactly the way **EffectLite** proceeds in all its analyses, whether the covariates are manifest or latent.

1.2.4 Average effects

If X is dichotomous with values 0 and 1, Equation (1.10) always holds, and the values $g_1(z)$ of the effect function $g_1(Z)$ are the conditional effects of X on Y given $Z=z$. Hence, it is straightforward to define the *average effect* of X on Y as the expectation of the effect function, i.e. $E[g_1(Z)]$. Since the values $g_1(z)$ are the $(Z=z)$ -conditional effects of X on Y , the expectation $E[g_1(Z)]$ is the *average of the conditional effects* of X on Y . Note that this definition of an average effect is unique and meaningful even if X and Z are correlated and if there is an interaction between X and Z . Furthermore, note that this definition also applies to a multivariate vector Z consisting of several univariate variables Z_1, \dots, Z_Q , which might be quantitative or qualitative. In the latter case, these covariates may be transformed into indicator (or dummy) variables representing the values of a qualitative covariate such as “socio-economic status” or “education”, in such a way that the intercept and effect functions can be parameterized as functions which are linear in the parameters [see Equations (1.14) to (1.16)].

If (the treatment variable) X has $J+1$ values, there will be J average effects: $E[g_j(Z)]$, $j = 1, \dots, J$, where the functions $g_j(Z)$ are the effect functions occurring in Equation (1.17). Each of these average effects is the average of the $(Z=z)$ -conditional mean differences between group $X=j$ and the reference group $X=0$.

1.2.5 Adjusted means

The average effect $E[g_j(Z)]$ is also the difference between the adjusted means, i.e.:

$$E[g_j(Z)] = E[E_{X=j}(Y | Z)] - E[E_{X=0}(Y | Z)]. \quad (1.23)$$

This can be seen from Equation (1.17), considering the two cases $X=j$ and $X=0$, and then taking the expectation. Hence, if we are interested in the average effects of the treatment variable X , i.e., in the averages $E[g_j(Z)]$, $j = 1, \dots, J$, of the conditional effects, we are interested in the differences between the adjusted means $E[E_{X=j}(Y | Z)]$ and $E[E_{X=0}(Y | Z)]$. **EffectLite** estimates these adjusted means and their standards deviations for each treatment condition. If one of the sufficient conditions for unbiasedness is fulfilled, the average effects $E[g_j(Z)]$ are also the average causal effects of treatment j compared to treatment 0 (see Steyer et al., 2007, for details).

1.2.6 Average effects given a treatment condition

EffectLite also estimates the average effects given a treatment condition, i.e., it estimates the conditional expectations $E[g_j(Z) | X=i]$ for all pairs (j, i) of $j \in \{1, \dots, J\}$ and $i \in \{0, 1, \dots, J\}$. If there are just two values of the treatment variable X , this means that **EffectLite** also estimates the conditional expectations $E[g_1(Z) | X=0]$ and $E[g_1(Z) | X=1]$, i.e., the average effects of the non-treated and of the treated, respectively. These average effects will differ from the average effect $E[g_j(Z)]$ in the total population, if the treatment probabilities depend on the covariate Z , i.e., if $P(X=1 | Z) \neq P(X=1)$. In this case, we can evaluate the treatment assignment policy [described by $P(X=1 | Z)$] by comparing the two conditional expectations $E[g_1(Z) | X=0]$ and $E[g_1(Z) | X=1]$ against each other. If $E[g_1(Z) | X=0] > E[g_1(Z) | X=1]$, one might think about revising the treatment assignment policy, because this comparison would then indicate that the average effect of the non-treated would be greater *if they would be treated in the future*.

See Steyer et al. (2007) for more details and the necessary assumptions allowing for this causal inference.

1.2.7 Null hypothesis: No average treatment effects

Although, if there is interaction, it is certainly more informative to consider the effect function and its values, the conditional effects, researchers may often also be interested in whether or not the *average effect* of the treatment $X=j$ (on the outcome Y) compared to the control $X=0$, i.e., $E[g_j(Z)]$, differs from zero. The null hypothesis of no average effects (zero mean differences) between all J groups simultaneously is then:

$$H_0 : E[g_1(Z)] = E[g_2(Z)] = \dots = E[g_J(Z)] = 0. \quad (1.24)$$

If X is *dichotomous* with values 0 and 1, Z is a univariate covariate, and $g_1(Z)$ is a *linear effect function*, the null hypothesis of no average effect (zero average mean difference between the two groups) is:

$$H_0 : E[g_1(Z)] = E(\gamma_{10} + \gamma_{11} \cdot Z) = \gamma_{10} + \gamma_{11} \cdot E(Z) = 0. \quad (1.25)$$

[see Equations (1.11) and (1.24)].

For three treatment groups and a two-dimensional covariate $Z = (Z_1, Z_2)$, the model is

$$\begin{aligned} E(Y | X, Z) &= g_0(Z) + g_1(Z) \cdot I_{X=1} + g_2(Z) \cdot I_{X=2} \\ &= (\gamma_{00} + \gamma_{01}Z_1 + \gamma_{02}Z_2) + (\gamma_{10} + \gamma_{11}Z_1 + \gamma_{12}Z_2) \cdot I_{X=1} \\ &\quad + (\gamma_{20} + \gamma_{21}Z_1 + \gamma_{22}Z_2) \cdot I_{X=2}, \end{aligned} \quad (1.26)$$

and the null hypothesis of no average treatment effects is:

$$\begin{aligned} H_0 : \quad &E[g_1(Z)] = E[g_2(Z)] \\ &= E(\gamma_{10} + \gamma_{11}Z_1 + \gamma_{12}Z_2) = E(\gamma_{20} + \gamma_{21}Z_1 + \gamma_{22}Z_2) \\ &= \gamma_{10} + \gamma_{11}E(Z_1) + \gamma_{12}E(Z_2) = \gamma_{20} + \gamma_{21}E(Z_1) + \gamma_{22}E(Z_2) = 0. \end{aligned} \quad (1.27)$$

1.2.8 Null hypothesis: intercept function is constant

EffectLite also provides a test for the hypothesis that the intercept function $g_0(Z)$ is a constant, i.e.:

$$g_0(Z) = \gamma_{00}. \quad (1.28)$$

According to this hypothesis the expectation(s) of the outcome variable(s) *in the control (or reference) group* coded with $X=0$ do not depend on the covariate(s). If the covariate is a pre-test and the outcome variable the post-test, this hypothesis will usually not hold. EffectLite also estimates the function $g_0(Z)$. However, even if this hypothesis holds for a given covariate, this covariate should be considered as long as there is an interaction with the treatment variable.

1.2.9 Null hypothesis: no interaction

EffectLite also provides a test for the hypothesis that there is *no interaction between treatment and the covariates*. According to Equation (1.17) this means that the effect functions are constants, i.e.:

$$g_1(Z) = \gamma_{10}, \quad g_2(Z) = \gamma_{20}, \quad \dots, \quad g_J(Z) = \gamma_{J0}. \quad (1.29)$$

In any case, i.e., irrespective of whether or not there are interactions, the program also estimates the effect functions. This allows to test hypotheses about the individual coefficients in these effect functions.

1.2.10 Null hypothesis: No treatment effects

Another hypothesis tested in **EffectLite** is the hypothesis of no treatment effects. In terms of Equation (1.17) the hypothesis is:

$$H_0: g_1(Z) = g_2(Z) = \dots = g_J(Z) = 0. \quad (1.30)$$

This means that there neither conditional effects given Z nor average effects of the treatments on the outcome variable Y . If several outcome variables are considered, this hypothesis is tested simultaneously for all outcome variables.

1.2.11 Null hypotheses about conditional effects

Inspecting Equations (1.11) or (1.14) to (1.16) shows that the intercept function $g_0(Z)$ and the effect functions $g_j(Z)$ contain parameters to be estimated. Hence, the estimate of each of these parameters has a standard error and a t -value for the test of the null hypothesis that the parameter is zero. **EffectLite** provides these standard errors and the t -values.

As mentioned before, **EffectLite** does not assume homogeneity of variances (univariate case with a single outcome variable) or covariance matrices (multivariate case with two or more outcome variables) of the outcome variables between groups. This also applies to the estimates of the standard errors and to the t -values mentioned above.

1.2.12 Several outcome variables

Note that you may also consider several outcome variables Y_i , $i = 1, \dots, I$, instead of just a single outcome variable Y . In such a case, **EffectLite** provides overall significance tests of the hypotheses (1.24), (1.28), (1.29), and (1.30) for all outcome variables simultaneously.

1.2.13 Sampling models

The unconditional (w.r.t. the treatment groups) expectations of the covariates, $E(Z)$, appear in the formulae for the average effects and the adjusted means of the outcomes. Therefore, the sampling variability of $E(Z)$ must be considered in inferences (error estimation or hypothesis testing) about these quantities.

Since the underlying multi-group SEM-model does not include the $E(Z)$ as explicit parameters, our current approach is to construct these as weighted averages of the group-specific covariate expectations, $E(Z_j)$, with group sizes serving as weights. Now, group sizes can be

random quantities, as is usual in observational studies, or they can be fixed by design, as is more typical of experimental research. **EffectLite** allows a choice between these two possibilities. In the latter case, group sizes are still calculated from the data, but we pretend that they are known constants by fixing their variances and covariances to zero in subsequent inferences.

A similar choice is available for the covariate expectations, $E(Z)$. If there are good reasons to think of the Z as *fixed regressors*, one can treat the $E(Z)$ as known constants having no standard errors (and no correlation with other parameter estimates). In that case, **EffectLite** will also treat group sizes as known constants. The default is to treat both group sizes as covariate expectations as random variables subject to sampling variability.

1.2.14 Optional assumptions about covariates

In the perfectly randomized experiment, the distributions of all pre-treatment variables will be identical in the population. Since the covariates are assumed to be such pre-treatment variables, the means, the variances and the covariances of the covariates can, in the case of such a randomized experiment, be assumed to be invariant across groups, if the covariates are *manifest*. If the covariates are *latent* (see section 1.3.1), the variances of the measurement error variables can be assumed to be invariant across groups as well.

1.2.15 Residual variances and covariances

EffectLite also computes the group-specific *variance(s and covariances) of the residual(s)* of the regression of the outcome variable(s) on the covariate(s). These variance(s) (and covariances) of the residuals are also computed if the covariate(s) and/ or the outcomes variable(s) are latent. If the control-group-specific variance of the residual in the regression of the outcome variable on the covariate is not significantly different from zero, one might consider an *individual-causal-effect model* (see, e.g., Steyer, 2005). Such an *ICE-model cannot* be analyzed via **EffectLite**. It can be analyzed, however, using the traditional software for structural equations.

1.3 Latent covariates and latent outcome variables

1.3.1 Latent covariates

It was mentioned before that **EffectLite** can also handle models with one or more latent covariates. This means that we can also analyze conditional and average effects (group means) in the case when Z is replaced by a latent covariate ξ or (a vector of) several ξ s. Of course, latent covariates need a measurement model. In Version 3.0 of **EffectLite** we may use the model of essentially τ -equivalent or the model of τ -congeneric variables and we also allow for method factors as defined by Pohl, Steyer, & Kraus, 2007. These method factors, which are also considered to be covariates, are important for a valid measurement model if the manifest measures of the latent covariate(s) are pre-tests and the same measurement instruments are applied in assessing the post-tests, the manifest measures of the latent outcome variable(s).

For two manifest measures Z_1 and Z_2 and a *single latent covariate* ξ within each group $j = 0, 1, \dots, J$, the model of τ -congeneric variables is defined by the following assumptions:

$$Z_1 = \lambda_{10} + \lambda_{11} \cdot \xi + \varepsilon_{Z_1} = 0 + 1 \cdot \xi + \varepsilon_{Z_1}, \quad (1.31)$$

$$Z_2 = \lambda_{20} + \lambda_{21} \cdot \xi + \varepsilon_{Z_2}, \quad \lambda_{20}, \lambda_{21} \in \mathbb{R} \quad (1.32)$$

$$\text{Cov}(\xi, \varepsilon_{Z_i} | X=j) = \text{Cov}(\varepsilon_{Z_1}, \varepsilon_{Z_2} | X=j) = 0, \quad i = 1, 2 \quad j = 0, 1, \dots, J. \quad (1.33)$$

With intercept 0 and loading 1 in Equation (1.31) we fix the scale of the latent variable ξ (see Steyer, 2001 or Steyer & Eid, 2001, for a more detailed presentation of these models). The intercepts $\lambda_{10} = 0$ and λ_{20} as well as the loadings $\lambda_{11} = 1$ and λ_{21} in the group-specific measurement models are assumed to be invariant across treatment groups. This allows for mean differences of the latent covariates between treatment groups. The measurement error variances $\text{Var}(\varepsilon_{Z_1} | X=j)$ and $\text{Var}(\varepsilon_{Z_2} | X=j)$ can be assumed to be invariant across groups. The default is that they are different across groups.

In the model of essentially τ -equivalent variables, both loadings λ_{11} and λ_{21} are fixed to 1. The generalization of both kinds of models to three or more manifest measures of a latent covariate is obvious. Usually, each manifest variable has only one single non-zero loading on one of the latent variables. However, the user may also specify measurement models with more than one non-zero loading of some of the manifest variables. This is necessary, e.g., in specifying method factors (see section 1.3.6). The only limit is the general identification problem in latent variable modeling.

If there are *several latent covariates*, we assume that each latent covariate has its own (two or more) manifest measures for which one of the two models described above holds. If we want to have a model consisting of a latent covariate and a manifest covariate, we assume that the manifest covariate measures its own 'latent covariate' (a 'dummy factor') with an intercept fixed to 0, a loading fixed to 1 and a zero variance of the measurement error variable.

1.3.2 Two groups

Figure 1.1 represents the path diagram of a simple model of this type with two manifest measures of the latent covariate ξ , a dichotomous group (treatment) variable X , and a single outcome variable Y . In version 3.0 of **EffectLite**, we allow for

- (a) several latent covariates,
- (b) linear intercept and effect functions of the latent covariates, and
- (c) measurement models with free loadings.

More than two groups are possible. The (distribution of the) group variable X may depend on the latent variable. However, although this dependency is allowed and taken into account in the model, it is not analyzed by **EffectLite**. Note that the analysis of conditional effect functions is also of interest even if X and ξ are independent, because they are a useful tool to study interaction, i.e., to study how the effect of X is modified by the latent covariates.

Figure 1.2 represents the same model as Figure 1.1. However, Figure 1.2 shows an extra path diagram for each of the two groups. These two path diagrams show that the slope of the

linear regression of the outcome Y on the latent covariate ξ may be different in each of the two (treatment) groups.¹




Figure 1.1: Path diagram of a conditional effects model with a latent covariate




Figure 1.2: A two-group presentation of the conditional effects model with a latent covariate

Furthermore, the difference between the two slope coefficients of the regression of Y on ξ in group $X=1$ and $X=0$ is γ_{11} , the interaction parameter occurring in the equation

$$E(Y | X, \xi) = (\gamma_{00} + \gamma_{01}\xi) + (\gamma_{10} + \gamma_{11}\xi) \cdot X. \quad (1.34)$$

This equation is the same as Equation (1.11) that was already considered in the case of a manifest covariate. The only difference is that the covariate is now denoted ξ instead of Z , and that, in specific empirical applications, Equation (1.34) can be wrong, because ξ is not dichotomous but assumed to be a latent continuous variable.

1.3.3 More than two groups

For the case of a latent covariate, Equation (1.17) may be rewritten

$$E(Y | X, \xi) = g_0(\xi) + g_1(\xi) \cdot I_{X=1} + \dots + g_J(\xi) \cdot I_{X=J}, \quad (1.35)$$

replacing the manifest covariate Z by the latent variable ξ . Assuming that all effect functions $g_j(\xi)$, $j = 1, \dots, J$ are linear, Equation (1.35) becomes

$$\begin{aligned} E(Y | X, \xi) &= (\gamma_{00} + \gamma_{01}\xi) + (\gamma_{10} + \gamma_{11}\xi) \cdot I_{X=1} + \dots \\ &+ (\gamma_{J0} + \gamma_{J1}\xi) \cdot I_{X=J}. \end{aligned} \quad (1.36)$$

Considering the conditional regressions of Y on ξ given $X = x$ in three groups yields:

$$E_{X=0}(Y | \xi) = \gamma_{00} + \gamma_{01} \cdot \xi, \quad (1.37)$$

$$E_{X=1}(Y | \xi) = (\gamma_{00} + \gamma_{10}) + (\gamma_{01} + \gamma_{11}) \cdot \xi, \quad (1.38)$$

$$E_{X=2}(Y | \xi) = (\gamma_{00} + \gamma_{20}) + (\gamma_{01} + \gamma_{21}) \cdot \xi. \quad (1.39)$$

¹Allowing for (a) different slopes in the groups and (b) the covariate to be latent, distinguishes this model from the traditional analysis of covariance model, which does not consider interaction between treatments and covariates.

This shows that the regression coefficients occurring in Equation (1.36) can be computed by performing a multi-group analysis, and then taking the differences between the group-specific linear regressions of the outcome variable Y on the covariate ξ . As noted before, this is exactly the way **EffectLite** proceeds in all its analyses, whether the covariates are latent, manifest, or not there at all.

1.3.4 Latent outcome variables

EffectLite handles latent outcome variables just in the same way as it handles latent covariates. This means that all analyses considered before, i.e., the analysis of conditional and average mean differences given manifest covariates or latent covariates, can be done also for the case in which the outcome variable(s) is (are) latent. The measurement model for the latent outcome variables is of the same type as the measurement model for the latent covariates. Hence, for two manifest measures Y_1 and Y_2 of a latent outcome variable η , this measurement model within each group j is defined by:

$$Y_1 = \kappa_{10} + \kappa_{11} \cdot \eta + \varepsilon_{Y_1} = 0 + 1 \cdot \eta + \varepsilon_{Y_1}, \quad (1.40)$$

$$Y_2 = \kappa_{20} + \kappa_{21} \cdot \eta + \varepsilon_{Y_2}, \quad \kappa_{20}, \kappa_{21} \in \mathbb{R}, \quad (1.41)$$

$$\text{Cov}(\eta, \varepsilon_{Y_i} | X=j) = \text{Cov}(\varepsilon_{Y_1}, \varepsilon_{Y_2} | X=j) = 0, \quad i = 1, 2, \quad j = 0, 1, \dots, J. \quad (1.42)$$

Hence, again the intercepts and the loadings (the slopes) in this measurement model are assumed to be invariant across groups. The variances of the measurement error variables $\text{Var}(\varepsilon_{Y_1} | X=j)$ and $\text{Var}(\varepsilon_{Y_2} | X=j)$ can be assumed to be equal or different between groups.

EffectLite also allows for three or more manifest measures of each latent outcome variable. Any combination of a manifest or a latent covariate with a manifest or latent outcome variables is possible. If both, covariates and outcome variables, are latent, it is assumed that the measurement error variables of the latent covariates and the latent outcomes are uncorrelated within treatment group j , i.e.:

$$\text{Cov}(\varepsilon_{Z_i}, \varepsilon_{Y_k} | X=j) = 0 \quad \text{for all pairs } (i, k), \quad i, k = 1, 2. \quad (1.43)$$

Note that this assumption is often violated in pretest-treatment-posttest designs. When a measurement instrument is applied repeatedly, it often happens that the associated measurement errors are correlated. This will lead to a bad fit of the model described above, and sometimes the estimation algorithm does not converge. As a simple remedy, one can specify the outcome variables as manifest rather than latent. A more sophisticated remedy is to introduce one or several method factors (see section 1.3.6).

As mentioned before, the measurement model is fixed between treatment groups in order to allow for between-group comparisons of the means of the latent variables. In pretest-treatment-posttest designs, we may want to fix the measurement also across time, between pre- and post-test measures. This is not absolutely necessary for effect estimation based on between-group comparisons. Nevertheless, it is a plausible assumption that may be introduced into the measurement model. In version 3.0 of **EffectLite** the user can specify such a model.

Equations (1.34) to (1.39) also apply to the case in which we have a *latent* outcome variable η instead of the manifest outcome variable Y . The path diagrams in Figures 1.1 and 1.2 are now replaced by Figures 1.3 and 1.4, respectively.




Figure 1.3: Path diagram of a conditional effects model with a latent covariate and a latent outcome variable

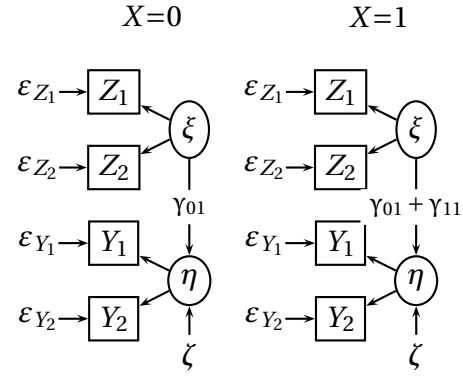


Figure 1.4: A two-group presentation of the conditional effects model with a latent covariate and a latent outcome variable. The latent variables ξ_1 and ξ_2 are allowed to correlate.

1.3.5 Optional assumptions about latent outcome variables

If the outcome variables are latent (see section 1.3.4), the variances of the measurement error variables can be assumed to be invariant across groups. This is an assumption that may or may not hold both, in the randomized experiment and in quasi-experiments.

1.3.6 Method factors

As mentioned before, the assumption of uncorrelated measurement errors is often violated in pretest-treatment-posttest designs. When a measurement instrument is applied repeatedly, it often happens that the associated measurement errors are correlated. This will lead to a bad fit of the model described in section 1.3.4. A remedy is to include a method factor, the values of which are the effects of using method 2 instead of method 1 (see Pohl et al., 2007, for more details). If two methods are designed to measure the same latent variable on the same scale, such a method factor is defined by the difference in the true-score variables of the two measurement methods. Assuming that this difference does not change over time leads to the measurement model depicted in Figure 1.5. This model is assumed to hold *within each treatment group* j . Only the effects of the latent covariates on the latent outcome variables, the residual variances, and the variances of the measurement error variables are allowed to differ between groups. The intercepts and loadings are assumed to be equal between groups.

1.4 Tests of unconfoundedness

EffectLite has been developed for the analysis of causal effects in experiments and quasi-experiments, i.e., in observational studies involving a treatment variable. However, it should be emphasized that a causal interpretation of the mean differences in the analyses without covariates, and of the average and conditional effects in the analyses with covariates is valid only if certain assumptions hold. In this case, we also say that the regressions $E(Y | X)$ and/or $E(Y | X, Z)$ are *causally unbiased* (see Steyer et al., 2007 for details of this and related concepts).

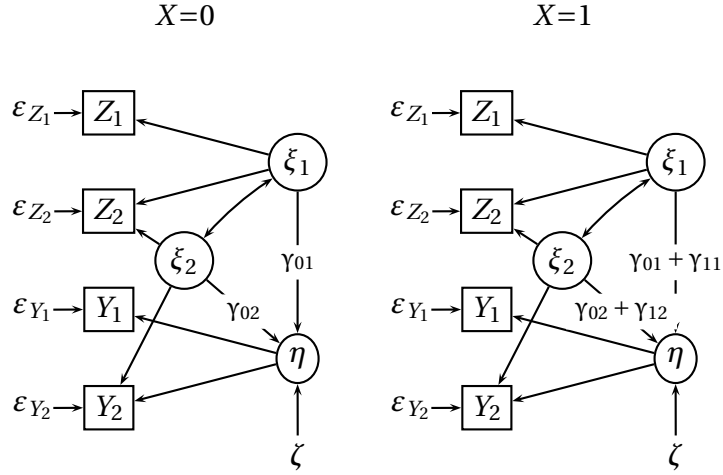


Figure 1.5: A two-group presentation of the conditional effects model with a latent covariate, η , a latent outcome variable, ξ_1 , and a method factor, ξ_2

Sufficient conditions for causal unbiasedness of the two regressions $E(Y|X)$ and $E(Y|X, Z)$ always refer to ‘potential confounders’. Potential confounders are (observed or unobserved) pre-treatment variables or other variables that cannot be affected by the treatment variable. Hence, potential confounders are covariates that are not specified as a covariate in **EffectLite**. Such a potential confounder will be denoted by W in the sequel.

1.4.1 Analysis without covariates

In analyses *without covariates*, a sufficient condition for the validity of a causal interpretation of the mean differences between treatment groups represented by X is that, for each potential confounder W (including also unobserved manifest and latent confounders), X and W are *independent*, i.e.:

$$P(X=j | W) = P(X=j), \quad \text{for each } j = 0, 1, \dots, J. \quad (1.44)$$

This condition will hold in the perfect randomized experiment, in which the probability of being in treatment j does not depend on the observational unit (individual), nor on any measure that is assessed (or that could be assessed) *before* the unit is assigned to one of the treatment conditions.

Unconfoundedness of the treatment regression. A less restrictive sufficient condition for the validity of a causal interpretation of the mean differences between treatment groups is that, *for each treatment condition* $j = 0, 1, \dots, J$, and each potential confounder W

(a) $P(X=j | W) = P(X=j)$

and/or

(b) $E_{X=j}(Y | W) = E_{X=j}(Y).$

This condition is called *unconfoundedness* of the treatment regression $E(Y | X)$.²

²Other, sufficient conditions for the validity of a causal interpretation of the mean differences between treatments groups are treated in Steyer et al. (2007).

Postulating (a) or (b) for a given potential confounder W implies:

$$E_{X=j}(Y) = E[E_{X=j}(Y | W)], \quad \text{for each value } j \text{ of } X. \quad (1.45)$$

This hypothesis [Equation (1.45)] is tested in version 3.0 of **EffectLite** if W is specified as the covariate. It holds if the means of the outcome variables in the treatment conditions do not differ significantly from their corresponding adjusted means.

If the analysis suggests that Equation (1.45) does not hold, and if there is no reason to assume that one of the other sufficient conditions for causal unbiasedness of the regression $E(Y | X)$ holds, then the potential confounder W should be included in the analysis as a covariate and causality tests should be conducted with another potential confounder, as outlined in the next section.

1.4.2 Analysis with covariates

As mentioned before, a covariate Z can be uni- or multivariate, manifest or latent. In analyses *with such a covariate*, a sufficient condition for the validity of a causal interpretation of the average effects and the conditional effects of the treatment variable X is that, for each potential confounder W , the treatment variable X and W are *conditionally independent given* Z , i.e.:

$$P(X=j | Z, W) = P(X=j | Z), \quad \text{for each } j = 0, 1, \dots, J. \quad (1.46)$$

This condition will hold in the perfect conditionally randomized experiment, in which the probability of being in treatment j may depend on Z , but on no other attribute of the observational unit (individual), nor on any other measure assessed before the unit is assigned to one of the treatment conditions.

Unconfoundedness of the covariate-treatment regression. A less restrictive sufficient condition for the validity of a causal interpretation of the conditional and average effects is that, for each treatment $j = 0, 1, \dots, J$, and each potential confounder W

$$(a) \quad P(X=j | Z, W) = P(X=j | Z)$$

and/or

$$(b) \quad E_{X=j}(Y | W, Z) = E_{X=j}(Y | Z).$$

This condition is called *Z-conditional unconfoundedness* of the covariate-treatment regression $E(Y | X, Z)$.³ Postulating (a) or (b) for a given potential confounder W implies:

$$E_{X=j}(Y | Z) = E[E_{X=j}(Y | Z, W) | Z], \quad \text{for each value } j \text{ of } X. \quad (1.47)$$

Although version 3.0 of **EffectLite** does not provide the test, let us outline how to test unconfoundedness of the covariate-treatment regression $E(Y | X, Z)$. Suppose for a moment that Z and W are unidimensional numerical variables. In this case, the hypothesis (1.47) that the

³Other, sufficient conditions for the validity of a causal interpretation of the conditional and average effects are treated in Steyer et al. (2007).

covariate-treatment regression is unconfounded can also be tested if we can make the assumptions

$$E_{X=j}(Y | Z, W) = \alpha_0^{(j)} + \alpha_1^{(j)} Z + \beta_1^{(j)} W, \quad \text{for each } j = 0, 1, \dots, J \quad (1.48)$$

and

$$E(W | Z) = \gamma_0 + \gamma_1 Z. \quad (1.49)$$

Under these assumptions, Equation (1.47) yields:

$$E_{X=j}(Y | Z) = \gamma_0^{(j)} + \gamma_1^{(j)} Z, \quad (1.50)$$

with

$$\gamma_0^{(j)} = \alpha_0^{(j)} + \beta_1^{(j)} \gamma_0 \quad \text{and} \quad \gamma_1^{(j)} = \alpha_1^{(j)} + \beta_1^{(j)} \gamma_1. \quad (1.51)$$

Hence, we can test the null hypothesis that Equations (1.51) hold, if we want to test unconfoundedness of the covariate-treatment regression $E(Y | X, Z)$.

This procedure can be generalized to several covariates Z_1, \dots, Z_Q and potential confounders W_1, \dots, W_R . In this case, unconfoundedness of the covariate-treatment regression [see Eq. (1.47)] can be tested if we can assume:

$$E_{X=j}(Y | Z, W) = \alpha_0^{(j)} + \sum_{q=1}^Q \alpha_q^{(j)} Z_q + \sum_{r=1}^R \beta_r^{(j)} W_r, \quad \text{for each } j = 0, 1, \dots, J, \quad (1.52)$$

and

$$E(W_r | Z) = \gamma_{r0} + \sum_{q=1}^Q \gamma_{rq} Z_q, \quad \text{for each } r = 1, \dots, R. \quad (1.53)$$

If the covariate-treatment regression is unconfounded and Equations (1.52) and (1.53) hold, then, for each treatment $j = 0, 1, \dots, J$,

$$E_{X=j}(Y | Z) = \gamma_0^{(j)} + \sum_{q=1}^Q \gamma_q^{(j)} Z_q, \quad (1.54)$$

with

$$\gamma_q^{(j)} = \alpha_q^{(j)} + \sum_{r=1}^R \beta_r^{(j)} \cdot \gamma_{rq}, \quad q = 0, 1, \dots, Q, \quad (1.55)$$

(for a proof see Steyer et al., 2007).

Hence, assuming Equations (1.52) and (1.53), we can test the null hypothesis that Equations (1.55) hold, if we want to test unconfoundedness of the covariate-treatment regression $E(Y | X, Z)$ [see Eq. (1.47)]. If the analysis suggests that Equation (1.47) has to be rejected and if none of the sufficient conditions for causal unbiasedness of the regression $E(Y | X, Z)$ holds for a given confounder W , then the confounder W should be included in the analysis as another covariate and tests of unconfoundedness should be conducted with another potential confounder W^* , if available, i.e., the procedure described above should be repeated for $E(Y | X, Z^*)$ with $Z^* \equiv (Z, W)$ replacing the original covariate Z and W^* replacing the original potential confounder W .

Note that we can be sure about causal unbiasedness of $E(Y | X, Z)$ only if the experimenter fixes the conditional treatment probabilities $P(X=j | Z)$. In all other cases, causal interpretations of conditional and average causal effects will always be preliminary and rest on assumptions that can only be falsified via causality tests, some of which have been described above. A *verification* of the sufficient conditions for the validity of those causal interpretations in observational studies *is not possible*.

1.5 Missing values

Version 3.0 of **EffectLite** will yield correct estimates of parameters and valid tests of null hypotheses even if there are missing values in the outcome variables. However, in this case we have to assume (a) that the covariates determining the probability that a value in one or more or even all of the outcome variables is missing are included as a covariate, and (b) that, within each treatment condition, the distribution of the covariates and outcome variables is *multivariate normal*. Assumption (b) assures that the full information maximum likelihood estimation (FIML) used in the background SEM program yields (i) unbiased estimates of the regressions of the outcome variables on the covariates within treatment conditions, and (ii) unbiased estimates of the expectations of the covariates in the treatment conditions. Without (i) and (ii) **EffectLite** may compute biased estimates. Since covariates only contain pretest information, assuming missing values only in the outcome variables is not too unrealistic, as long as systematic attrition only occurs after the covariates (often pre-tests) are assessed.

If assumptions (a) and (b) hold, then **EffectLite** will compute unbiased estimates of the expectations of the original outcome variable in the treatment groups ("raw means"), unbiased effect functions, unbiased estimates of the adjusted means, and unbiased average causal effects, provided that the FIML option is used (see section 2.3). Missing data as described above will only reduce statistical power as compared to a sample with no missing values. Note that multiple imputation (see, e.g., Little & Rubin, 2002) would *not* improve the quality of the estimates under these assumptions. Also note that, under assumptions (a) and (b), listwise deletion of cases with missing values may lead to seriously biased means of the outcome variable, biased average effects, and biased adjusted means.

Hence, in the case of missing values, it is essential to

- select 'FIML' on the **Open Data** screen, and
- include in the model all covariates that might play a role in predicting the probability that a value is missing on the outcome variable(s) or on the covariates.

To date it is still not clear if assumption (b) of multivariate normality can be replaced by a weaker assumption allowing also for non-normal covariates and whether or not we can allow for missingness completely at random in those covariates that do not determine missingness in the outcome variables. Bias can occur if there is missingness in a covariate that also determines missingness in one of the outcome variables.

1.6 Inferential techniques in EffectLite

In this section, we provide a brief and informal discussion of the statistical techniques used when testing statistical hypotheses and estimating the standard errors of the various statistics reported by **EffectLite**.

As mentioned above, **EffectLite** depends on the work of a standard program for structural equations models (SEM) such as **Mplus**. SEM programs typically implement some sort of maximum likelihood estimation (MLE), and report a vector of parameter estimates along with an estimate of their variance-covariance matrix. In our case, the SEM is a certain kind of multi-group regression model, possibly multivariate and/or involving latent variables. Once that the

SEM is estimated, **EffectLite** computes some new statistics, which are functions of the parameters in the original model and are relevant to causal analysis. **EffectLite** also estimates the standard errors of the new statistics and tests various statistical hypotheses. Each hypothesis can be expressed as a set of constraints on the parameters of the SEM.

Let us denote the vector of parameter estimates in a SEM with $\hat{\theta}$, and their (estimated) variance-covariance matrix as $\mathbf{V}(\hat{\theta})$. Suppose we have a new parameter, say ϑ , defined as a function of the original parameters, i.e., $\vartheta = f(\theta)$. From standard statistical theory, the function of the parameter estimates is the estimate of the function of parameters: $f(\hat{\theta}) = \hat{f}(\theta) = \hat{\vartheta}$. For the asymptotic distribution of $\hat{\vartheta}$, it can be shown that

$$\frac{\hat{\vartheta} - \vartheta}{\hat{\sigma}_{\hat{\vartheta}}} \rightsquigarrow N(0, 1). \quad (1.56)$$

The standard error, $\sigma_{\hat{\vartheta}}$, can be estimated as

$$\hat{\sigma}_{\hat{\vartheta}} = \sqrt{(\hat{\nabla} f)' \mathbf{V}(\hat{\theta}) (\hat{\nabla} f)}, \quad (1.57)$$

where $\hat{\nabla} f$ is the gradient (vector of partial derivatives with respect to $\theta_1, \theta_2, \dots, \theta_k$) of function f evaluated at $\theta = \hat{\theta}$. In statistics, Equations (1.56)–(1.57) are known as the *multiparameter delta method* (see, e.g., Wasserman, 2004).

In **EffectLite**, we are typically interested in sets of new parameters such as average effects or adjusted means. To estimate their standard errors, one simply collects the gradients for all new parameters as the column vectors of a matrix, \mathbf{J} , and computes the variance-covariance matrix of the new parameters as

$$\mathbf{V}(\hat{\theta}) = \mathbf{J}' \mathbf{V}(\hat{\theta}) \mathbf{J}. \quad (1.58)$$

The standard errors of the new parameters are then the square roots of the diagonal elements in the resulting matrix.

The many hypotheses tested by **EffectLite** are all expressed as sets of constraints imposed on the parameter estimates in the underlying SEM. There are many ways to construct statistical tests for such constraints, but the most popular ones are the *likelihood ratio test*, the *Wald test*, and the *Lagrange multiplier test*. Figure 1.6, adapted from Greene (1997), helps explain the basic idea behind each type of test.

Let $\log L(\theta)$ be the log likelihood function, which reaches its unconstrained maximum, $\log L$, at point $\hat{\theta}_{\text{MLE}}$, the maximum likelihood estimate. We also have a constrained estimate, $\hat{\theta}_{\text{R}}$, at which the log likelihood has a value of $\log L_{\text{R}}$. Because the constrained solution is sub-optimal, $\log L_{\text{R}}$ is smaller than the unconstrained maximum likelihood, $\log L$. The difference $\log L - \log L_{\text{R}}$ forms the basis for the likelihood ratio test of significance for the constraint. The test statistic is defined as $-2 \log(L_{\text{R}}/L)$; it has an asymptotic χ^2 distribution with as many degrees of freedom as the number of constraint functions imposed simultaneously on $\hat{\theta}$.

Just like the new parameters, the constraints defining a hypothesis to be tested are functions of θ — say, $c(\theta)$ for the simplest case of one constraint function on a single parameter θ . The value of this function is exactly zero at the constrained estimate. If the constrained and the unconstrained solutions are not too different, its absolute value at the unconstrained estimate should not be very large. This is the logic behind the Wald test.

In our case, there is usually a whole vector of parameters, θ , and most of the hypotheses that we want to test involve more than one constraint function on θ . Let $\mathbf{c}(\hat{\theta})$ be a vector containing




Figure 1.6: An illustration of the likelihood ratio (LR), test, the Wald test, and the Lagrange multiplier (LM) test; see text for explanation

the values of all constraint functions related to a hypothesis, at the unconstrained solution. The Wald test for the hypothesis $H_0 : \mathbf{c}(\hat{\theta}) = \mathbf{0}$ can then be computed as

$$W = \mathbf{c}(\hat{\theta})' (\hat{\mathbf{V}}[\mathbf{c}(\hat{\theta})])^{-1} \mathbf{c}(\hat{\theta}). \quad (1.59)$$

Because the constraints are functions of the original SEM parameters, their variance-covariance matrix, $\mathbf{V}[\mathbf{c}(\hat{\theta})]$, can be estimated by the delta method [Equation (1.58)]. The Wald test is asymptotically equivalent to the likelihood ratio test: in large samples, it has the same χ^2 distribution with as many degrees of freedom as the number of constraint functions imposed on $\hat{\theta}$.

The first derivative of the log likelihood, $\partial \log L(\theta) / \partial \theta$, is zero at the unconstrained MLE. For any constrained, i.e., suboptimal solution, the derivative deviates from zero, and can therefore be used to test whether the constraint brings about a sizeable departure from the optimum. This is the logic behind the Lagrange multiplier test, which is asymptotically equivalent to the likelihood ratio test and the Wald test.

Hence, the Lagrange multiplier test requires the constrained estimates and tests for the improvement in fit associated with *lifting* one or more constraints, while the Wald test requires the unconstrained estimates and tests for the loss of fit associated with *imposing* one or more constraints. Finally, the likelihood ratio test requires both the constrained and the unconstrained estimates, and is hence more expensive to calculate.

In earlier versions of **EffectLite**, we used likelihood ratio tests, which have now all been replaced with Wald tests. This has two important advantages:

- it is now sufficient to estimate just the unconstrained model. The earlier approach required the estimation of many constrained models, one for each hypothesis tested. This was a lot of work and ran a higher risk that estimation might fail for at least one of the models;
- some of the parameters used in defining the constraints are not a part of the underlying SEM, and had to be obtained in fairly artificial ways. We still have to compute them, but the only extraneous information needed are the proportions of the treatment groups in the sample. We simply augment the variance-covariance matrix reported by **Mplus** with the variances and covariances of these proportions; when group sizes are fixed by design, we augment with a zero matrix instead.

2 Detailed program description

2.1 Before you begin

Because it depends on **Mplus**, **EffectLite** must know if, and where, **Mplus** is installed. The first time it is run, **EffectLite** will remind you to complete the installation by providing this vital information. Basically, you have two options: search for the **Mplus .exe** file, or use **EffectLite** without **Mplus**; in the latter case, you can generate and examine **Mplus** command language for all models in **EffectLite**, but you cannot fit the models and obtain additional output.

To search for **Mplus**, navigate the file system until you reach the folder where **Mplus** is installed. Double-click on the .exe file, e.g. **mplus.exe** (NOT **Mpluswin.exe**!).

In English-speaking countries, the standard location for the **Mplus .exe** file is

```
C:\Program Files\Mplus\Mplus.exe
```

For the demo version (freely available at www.statmodel.com and sufficient for small-scale problems), this will change to

```
C:\Program Files\Mplus Demo\Mpdemo3.exe
```

In other countries, **Program Files** is replaced with a nationalized version — for instance, it is **Programme** in Germany.

Normally, the search for **Mplus** is only necessary before you run **EffectLite** for the first time. You will have to repeat it if you log on as a different user, or if **EffectLite** cannot find **Mplus** where it is supposed to be. You can repeat it at any time by going, via the **EffectLite** menu, to the **Mplus Location** screen.

2.2 A normal run of EffectLite

EffectLite is organized in a series of screens. Most screens have a **Continue** button. Set the options on each screen as necessary and click on **Continue** — this takes you to the next screen, or it may display a warning message if you forgot to set some vital option.

The menu above the main window provides access to all screens, but some of the entries may be dimmed depending on the stage of analysis. You can use the menu to go back to a previous screen and change some analysis options, open a new data set, and so on. However, *only clicking the Continue button* will remember the new options and eventually run **Mplus** to produce a new analysis. If you change some options on a screen and then use the menu to go *forward* to some of the results screens, all you will see are the results of your previous analysis.

- A normal run of **EffectLite** begins with the **Open Data** screen. Open a data set and click on **Continue** to go to the **Select Variables** screen.
- Select the variables for the analysis and click on **Continue** to open the **Main Options** screen.

- Fill in the options for the analysis and click the **Continue** button.
- If the covariate and / or the outcome is latent, **EffectLite** will open additional screens to specify the measurement model.
- Finally, **EffectLite** will open the **Main Results** screen and display the results of the analysis.

2.3 The Open Data screen

EffectLite can open data files in a variety of formats, such as **SPSS**, **Systat**, **Statistica**, **Stata**, **SAS**, **Excel**, **Matlab**, **S+**, and others. The file type is recognized automatically. For more information, see section 2.13.

To open a data file, click on the **Select file** button and navigate to the directory where the data file is saved. Double-click on the data file, or click on the data file and then on **Open**. The name of the file will be displayed to the right of the **Select file** button.

Select file	c:/mydata/myfile.csv
--------------------	----------------------



If **EffectLite** can open the file, it will display it in the window below. This is not a data editor, just a facility to make sure that the right data file was loaded, and identify possible problems. Use your favourite statistical program to prepare and edit data files.

EffectLite translates the data file to the comma-separated-values (CSV) format (for more information, see section 2.14). This new file is saved under the same name and in the same directory as your original data file, but with the extension `.csv`. Later on, you can open this file instead of the original one, or you can delete it if you wish. It is not the data file that will be passed on to **Mplus**: after you have selected the variables and other options for the analysis, **EffectLite** will produce a separate data file for each non-empty treatment group. These files will contain only the selected variables, and will be saved in the same directory with the extension `.dat`.

Missing data — <input type="radio"/> listwise deletion <input checked="" type="radio"/> FIML

Right under the data window, there is a device to choose how missing data will be treated. The option **listwise deletion** will remove all cases for which any of the variables selected on the **Select Variables** screen has a missing value. The option **FIML** will keep all cases (except those where the treatment variable is missing), code missing values appropriately, and switch on **Mplus**'s Full Information Maximum Likelihood (FIML) estimation option. Section 1.5 describes in more detail when and how to use this option to avoid biases due to missing data.

2.4 The Select Variables screen

On the **Select Variables** screen, you select the variables for the analysis. Initially, all variables in the data set appear in the left-hand window. To select the outcome variable(s), click on one or more variables, and then on the  button to the left of the **Outcome variable(s)** window. The selected variables will move to the **Outcome variable(s)** window. If you selected the wrong variable, select it in the **Outcome variable(s)** window, and click on the  button to its left.

Proceed in a similar fashion to select the treatment variable. Note that there can be only one treatment variable in the analysis. If you select more than one variable, there are two possibilities:

- If the **Group (Treatment) variable** window is still empty, the first variable of your selection will be moved to it.
- If the **Group (Treatment) variable** window already contains a variable, nothing will happen.

Optionally select at least one covariate if you want to do an analysis with covariates. The **Covariate(s)** window works in the same way as the **Outcome variable(s)** window.

The **Select Variables** screen is accessible only if a data file has been opened. If that is not the case, you will be returned to the **Open Data** screen.

2.5 The Analysis Options screen

The **Select control group** window on the left displays the category codes of the independent variables and their observed frequencies. Click on any row of this display to select the control (reference) group. The row for the selected group will be highlighted.

Measurement model

☒ all manifest

☐ some latent

The **Measurement model** window is accessible when more than one variable have been selected for the analysis. Choose the **all manifest** to treat all selected variables as manifest, or the **some latent** option if your model involves at least one latent variable. The measurement model will be specified on a separate screen.

Sampling model

☒ fully stochastic

☐ group sizes fixed

☐ group sizes and $E(Z)$ fixed

The **Sampling model** window contains options regarding the sampling model. Group sizes can be treated either as random variables (the default), or as known constants. When group

sizes are random variables (as is typical of observational studies), their variances and covariances will be considered when computing Wald tests and the standard errors of average effects and adjusted means — this is the default setting. In experimental studies, group sizes are typically fixed by design, and can be treated as known constants; in that case, they will still be calculated from the data, but their sampling variances and covariances will be assumed to be zero. In addition, the unconditional expectations of the covariates, $E(Z)$, can also be treated as known constants; this option is not available when no covariates have been selected for the analysis.

The controls on the right hand side may be used to set some equality constraints between groups: the available options depend on the kind of model being tested. In addition, some options may be unavailable (dimmed) for particular models.

Models without covariates

Assuming group invariance of: —

☐ Covariate means
☐ Outcome (co)variances
☐ Covariate measurement errors
☐ Outcome measurement errors

Models with covariates

Assuming group invariance of: —

☐ Covariate means
☐ Covariate (co)variances
☐ Covariate measurement errors
☐ Outcome measurement errors

When the model involves covariates and the data come from a randomized experiment, it may be interesting to set the covariate means, covariate variances and covariances, and – if the covariate(s) are latent – covariate measurement errors equal among groups.

Finally, there is a possibility to ask for shorter or longer output. In the present version, the additional output shown by the **long** option includes:

- in models without covariates: the covariance/correlation matrix of the outcome variables;
- in models with covariates: the covariance/correlation matrix of the residuals of the outcomes;
- in models with latent outcomes and/or latent covariates: standard errors of the parameters in the measurement model(s).

The **EffectLite** output group contains a checkbox labeled **Conditional effects**. Turn this on to calculate conditional treatment effects for each row in the data file. The effects and their standard errors can be examined on the **Conditional effects** screen, and they can be saved in a CSV file along with all variables used in the analysis. Conditional effects are not available in models involving latent variables — in that case, the checkbox will be dimmed and the **Conditional effects** screen will not be accessible.

When you have set all analysis options, click the **Continue** button. **EffectLite** will generate the setup language for **Mplus**, run **Mplus**, and display the **Main Results** screen. If you have forgotten to select the control group, you will be taken back to the **Analysis Options** screen.

2.6 The Measurement Model screen

Unlike some earlier versions of **EffectLite**, there is now a single **Measurement Model** screen for all outcome variables and covariates. It is accessible only when the **some latent** option has been selected on the **Main Options** screen.

The **Measurement Model** screen contains a button to add a latent variable (factor), a button to delete a factor, and a device to assign manifest variables to factors and constrain some of the factor loadings.

All manifest variables selected as outcomes (or, respectively, covariates) are displayed in the first column; the outcome variables come first, followed by the covariates (if any). Latent variables will be named F1, F2... They can be defined in any arbitrary order, and **EffectLite** will recognize them automatically as latent outcomes or latent covariates. Latent variables attached to both outcomes and covariates (such as method factors) will be treated as latent covariates (for a definition of method factors, see Pohl et al., 2007).

Loadings are shown in the next columns. Click on any button to change its state. The possible states for factor loadings are 1, *, and 0, meaning fixed to 1, freely estimated, or fixed to zero.

When a button is blue and displays an asterisk, its contents can be edited. The button selected for editing will have a grey background and will display a cursor. Remove the asterisk with the backspace key, and type in either a number (using a decimal dot rather than comma), or a letter. Typing a number will fix a loading to the corresponding value, and letters are useful for defining equality constraints among two or more loadings: all loadings sharing the same letter will be set equal by **Mplus**.

Before it continues with the analysis, **EffectLite** will check that:


- each manifest variable is attached to a factor;
- each factor has at least one manifest measure; and
- at least one factor loading for each factor is set to 1.

To mix latent and manifest outcomes/covariates, create a dummy factor for each variable that must enter the analysis as manifest. The dummy factor has a loading of 1 for the manifest variable in question, and loadings of 0 for all other manifest variables. The associated error variable term will automatically be fixed to 0.

The **Continue** button will run **Mplus**, compute some additional output, and open up the **Main Results** screen.

2.7 The Main Results screen

EffectLite will typically prepare the setup language for a **Mplus** model, run it, interpret the output, compute some additional parameters, and display the **Main Results** screen. If **Mplus** estimation has converged without fatal errors, you will see the **EffectLite** output on a light blue background. You can scroll up and down, and even type in your own comments. For detailed information on the output, see section 2.18 and section 2.17.



Click the **Save** button if you wish to save the output to a file. You will be prompted for a file name. If you check **Include Mplus output**, the detailed output shown in the window on the **Technical Output** screen will be appended to your output file.

If **Mplus** cannot estimate the model, the main output window will be pink. If **Mplus** is not installed or could not be found in the presumed location, the main output window will be yellow.

A blue window is generally good news, but it does not guarantee that everything was fine with the model. There may be non-fatal errors or warnings. You may want to switch to the **Technical Output** screen to check for any of these.

2.8 The Conditional Effects screen

The **Conditional Effects** screen displays, in a scrollable table, estimates of the conditional treatment effects and their standard errors for each row of the data file. All variables used in the model: the group variable, the outcome variable(s), and the covariate(s), are included as well.

Conditional effects are computed just like average treatment effects; the only difference is that covariate expectations are replaced by fixed covariate values. The columns containing the conditional effects are labeled Y1_G1, Y1_G2... where Y1 means the first outcome variable, and G1 refers to group 1, the first treatment group (the control group is always group 0). The labels for standard errors have SE instead of Y. While this system may not be very convenient, it ensures simple and short variable names that will not lead to trouble with other programs. The numbers for the outcome variables correspond to their ordering in the table. The numbers for the treatment groups correspond to the ordering shown under **Groups**, by order of analysis near the beginning of the **EffectLite** output.

Starting with v.3.1.1, the table can be sorted on any column — just place the cursor anywhere in the column and control-click the left mouse button (i.e., press and hold the Control button and click). Alt-click in any column containing conditional effects to sort the table on the ratio of the conditional effects to their standard errors; this is quietly ignored if the cursor is not in a column containing conditional effects.

Click on the **Save CSV file** button to save the contents of the table into a file. Conditional effects are only available for models with covariates but without latent variables.

2.9 The Technical Output screen

The chi-squared tests and other statistics reported by **EffectLite** are computed from the parameter estimates of an underlying **Mplus** model. The **Technical Output** screen displays the original **Mplus** output for the model. This contains much additional information that may be confusing for the novice but helpful in troubleshooting or interesting for those who wish to understand the technical aspects of the models.

To help you identify possible problems, we have adopted the following convention:

- Mplus command language that is not supposed to run or could not be run at all is displayed on a *yellow* background;
- The output for a model that ran with a fatal error is displayed on a *pink* background;
- When Mplus has produced *at least* chi-squared statistics for a model, the output is displayed on a *blue* background. This is not a guarantee that everything is fine with the model: watch out for warning messages in the output.

You can save all technical output together with the main results from the analysis – see section 2.7.

2.10 The Wizard Stuff screen

The models in EffectLite are fairly complicated, and estimation may fail for particular data sets. Sometimes the problem may be remedied by minor changes to some Mplus settings.

Examples of such changes are introducing user-provided start values for some parameters, or changing the convergence criterion for the iterative algorithm.

The Wizard Stuff screen, which is only accessible over the menu, gives you the possibility to introduce such small changes, and rerun the models. The layout is similar to the Technical Output screen; however, the window has a pale green background and shows the Mplus command language for the model. Edit the input as necessary and click the **Save file** button to save the file. When you have made all the necessary changes to the input files, click **Run Mplus** to re-estimate the model.

It is very important to avoid any changes to the input files that would modify the model itself — for instance, introduce new parameters or constrain existing ones. Doing so would confuse the post-processing of Mplus output by EffectLite and produce misleading results or unpredictable errors.

2.11 The Mplus Locations screen

This screen displays the location of the executable files for Mplus, or Mplus not available when EffectLite is configured to work without Mplus. It also allows you to change the configuration for the location of Mplus. This might be useful, for instance, if you install a new version of Mplus.

Full instructions on how to search for the Mplus executable are provided in section 2.1.

2.12 The Help screen

The Help screen displays the help files for EffectLite. Home displays the help contents, Back and Forward page through the help screens shown last, and About shows some general information on the program.

You can search the Help contents on any word. Add * to the word to do a complete search through all help files.

2.13 Types of data files you may use

Thanks to David Baird's remarkable utility **DataLoad**, **EffectLite** can open data files in a vast variety of formats: Excel 2-5, 95, 97, 2000, XP; Lotus WK1; dBase 2-5; Quattro WQ1, WB*, QPW; Paradox 3-5; comma, tab, or space delimited text files; GenStat GSH; S+ PC/Unix/Text; SPSS for Windows; Minitab 8-13, Portable; SAS PC 6.03-12; SAS Transport; Stata 4-7; Systat; Epi-Info; Statistica 5-6; Gauss; MatLab; and others. We have not tried **EffectLite** with all of them, but SPSS files (.sav) and SYSTAT files (.syd; .sys) have worked fine.

The type of the data file is recognized automatically. When you open a data file, it is first translated to the CSV data format: the resulting CSV file is saved in the same folder as the original data file, and it will have the same name but the extension .csv. **EffectLite** actually works with the CSV file: if you wish to re-use the same data, you can save a couple of seconds by opening the CSV file instead of your original file.

Very large or complex SPSS files may cause problems. One remedy is to keep the files as small and simple as possible. Another possibility is to save the data from within SPSS to a different format such as SAS Transport or SYSTAT .syd, and then open that in **EffectLite**. However, we do not recommend creating text files with SPSS – they may contain decimal commas and other symbols that can confuse Mplus.

2.14 The CSV data format

All data formats (e.g. SPSS data files) are translated to a Comma Separated Values (CSV) file before they are loaded into **EffectLite**. CSV files are very simple: you could prepare them with a simple text editor such as Notepad.

Because the CVS format can differ slightly between programs, here is the specification of the version exported by **DataLoad** and expected by **EffectLite**:

- The first line contains the variable names, separated by commas.
- Each subsequent line has the data for one case, separated by commas.
- The values of all non-numeric variables are enclosed with simple quotes, e.g. 'value'. They may contain blanks and commas. Quotes are also allowed and must be represented as two simple quotes: ' '. Long strings are truncated to a length of 12 characters.
- The decimal sign is a period .
- Missing values are represented with * . Values of non-numeric variables consisting entirely of white space or equal to the empty string are also considered missing.
- Variable names need not be enclosed in quotes. **DataLoad** will automatically change variable names to a minimal standard, replacing potentially confusing characters.

Here is a very short example of a CSV file; note the embedded quote in Dewar's, and the missing value in the last case:

```
treatment,age,sex,pulse
'control grou',23,0,66
'Johnnie Walk',26,1,49
'Dewar''s',44,0,55
'Johnnie Walk',*,1,70
```

2.15 Models available in the current version

EffectLite can handle two kinds of statistical analysis: with or without covariates. The more general case of an **analysis with covariates** (see section 1.2) involves:

- one categorical group variable,
- one or more outcome variable(s),
- one or more covariate(s).

In principle, the outcome variables and the covariate variables may be quantitative, ordinal, or nominal; however, the current version of **EffectLite** works only with quantitative outcome variables. A simple trick to use qualitative covariates is to represent them by dichotomous indicator (dummy) variables.

Given a sufficient number of indicators, both outcomes and covariates can be latent, or even a mixture of latent and manifest variables. In that case, the *measurement model* is specified on a special screen (see section 2.6). A new feature in **EffectLite** is that a latent variable may be attached both to outcomes and covariates; such latent variables are treated as latent covariates, and may be used to represent method factors.

Another new feature is the ability to impose equality constraints between two or more factor loadings. An obvious application is measurement invariance for latent variables that have been measured repeatedly over time.

The unconditional (w.r.t. treatments) expectations of the covariates, $E(Z)$, appear in the formulae for the average treatment effects and the adjusted outcome means. Since $E(Z)$ do not appear as explicit parameters of the **Mplus** model, our current approach is to construct them as weighted averages of the treatment-specific expectations, with group sizes serving as weights. The user has the option to consider group sizes either as random, which is a realistic assumption in an observational study, or as fixed and known constants, which is more appropriate for a perfectly controlled experiment. **EffectLite** will always compute group sizes from the data, but in the random case it will also estimate their sampling variances and covariances and consider them in Wald tests and standard error calculations. A similar choice is available for the expectations $E(Z)$: they can also be treated as known constants, which makes sense in research designs with fixed regressors.

In the present version of **EffectLite**, all measures of a latent outcome or a latent covariate must be numeric, and are treated as quantitative variables. Future versions might add options for ordinal measures. A further constraint is that each manifest variable can have a non-zero loading on just one latent variable (see the model of congeneric variables in Steyer, 2001).

A mixture of manifest and latent outcomes and/or covariates can be achieved with dummy factors. A dummy factor is a latent variable measured by just one manifest variable, with a factor loading of 1 and a measurement error variance of 0. To create a dummy factor, set up the

factor on the appropriate **Measurement Model** screen, and set the loading to 1; the measurement error will be fixed to zero automatically if a factor has only one manifest variable.

The current version of **EffectLite** can only handle models with a linear intercept function and a linear slope (effect) function. However, the user can easily prepare the data in ways that allow for a wider choice of models. For instance,

- product terms like $Z_1 \cdot Z_2$ and powers like $Z_2 \cdot Z_2$ may be computed and specified as a new covariate to allow for models that are linear in the parameters but nonlinear in the covariates;
- nominal covariates such as sex can be accommodated by pre-computing indicator variables and specifying them as new covariates. This results in nonorthogonal ANOVA and MANOVA models with and without covariates.

EffectLite can also handle **models without covariates** (see section 1.1). These can be used as alternatives to ANOVA and MANOVA designs. However, they are more general than ordinary ANOVA and MANOVA models, since **EffectLite** does not presume homogeneity of variances or covariance matrices between groups. Furthermore, you have the choice between group comparisons of the means of the manifest outcome variables or group comparison of the means of one or more common underlying latent variable(s).

2.16 Types of variables

EffectLite works with three kinds of variables:

- a single nominal treatment variable, which can have two or more categories. The categories may be coded as integers (1, 2), characters (A, B, C), or strings (control, treatment).
- at least one outcome variable, which should be quantitative. Ordinal variables coded with integers are accepted but treated as quantitative; this may change in future versions. Non-numeric variables will cause an error in the present version.
- optionally, one or more covariates. All covariates should be numeric, and will be treated as quantitative. Ordinal variables coded with integers are accepted but treated as quantitative; this may change in future versions. Nominal covariates can be represented by dummy variables; these must be prepared before the data file is opened in **EffectLite**.

Data files may contain missing values. When the data file is translated to the CSV data format, all missing values are represented with a `*`. The user has the choice between listwise deletion of cases having missing values on any variable selected for the analysis, or full information maximum likelihood estimation (FIML).

2.17 Output: Model without covariates

2.17.1 General information

The output for a model without covariates starts with some general information:

```
EffectLite v.3.0 for Mplus: A program for the
analysis of conditional and average effects
(c) Rolf Steyer and Ivailo Partchev, 2007
Internet: www.statlite.com
Email:    Rolf.Steyer@uni-jena.de
          Ivailo.Partchev@uni-jena.de
```

Thursday, June 21 2007, 16:18

The file from which the data was read, and the options for missing data treatment and for the sampling model:

```
Data from file D:/current/EffectLite/simul.csv
Missing data treatment is listwise deletion
Sampling model is fully stochastic
```

A list of the levels of the group variable and their counts (frequencies):

```
Group variable: x
Groups and their frequencies:
      Group  Frequency
        a         348
        b         306
        c         344
with Group a as control group
```

The groups are then shown again, reordered as they will be seen and processed by Mplus. The control group always comes first. This table is useful when trying to understand the details shown on the Technical output screen.

```
Groups, by order of analysis:
      Group  Rel.Freq.
        0         a    0.3487
        1         b    0.3066
        2         c    0.3447
```

A list of the outcome variables:

```
OUTCOME(S): y1 y2, treated as manifest
```

2.17.2 Overall tests of significance

The overall test of significance for the average effect, now computed as a Wald test. This is a test that, for each outcome variable, the means are equal in all treatment groups. The variances need not be equal.

```
* Simultaneous tests for all treatment groups and all dependent variables *

=====
Hypothesis                Chi-sq      DF      Prob
-----
No treatment effect       64.071      4    0.0000
=====
```

2.17.3 Detailed analysis of mean differences

The detailed analysis shows, for each outcome variable, the mean difference of each group with respect to the control group.

```
*** Detailed analysis of the effects ***
```

The following results are for the first outcome variable – in this case, `y1`:

```
Results for outcome variable 1: y1
```

The first comparison is between the second group (`b`) and the control group (`a`):

```
Group b - Control group a
```

`Effect` is the difference between the two means, `Std.error` is the standard error of this difference, and `Effect/SE` is the ratio of the effect to its standard error.

```
Effect          0.455
Std.error       0.186
Effect/SE       2.444
Effect size     0.201
```

Effect size is estimated by dividing the effect by the standard deviation of the outcome variable in the control group.

A similar comparison is made between the third group (`c`) and the control group:

```
Group c - Control group a
...
```

and then follow results for the second outcome variable, `y2`:

```
Results for outcome variable 2: y2
...
```

2.17.4 Group means of the outcome variables

```
*** Group means of the outcome variable(s) ***

      Group      Outcome      Mean      Std.dev.      Std.error
      a         y1        -0.402        2.258        0.121
      a         y2        -0.367        2.306        0.124
      ...
```

There will be some additional output if the `EffectLite` output option on the Main Options screen is set to `long`: for details, see section 2.5.

Note that the interpretation of the results as average causal effects is based on the additional assumption of unbiasedness of the treatment regression $E(Y | X)$.

2.18 Output: Model with covariates

2.18.1 General information

The output for a model with covariates starts with some general information:

```
EffectLite v.3.0 for Mplus: A program for the
analysis of conditional and average effects
(c) Rolf Steyer and Ivailo Partchev, 2007
Internet: www.statlite.com
Email:    Rolf.Steyer@uni-jena.de
          Ivailo.Partchev@uni-jena.de
```

```
Thursday, June 21 2007, 16:18
```

The file from which the data was read, and the options for missing data treatment and for the sampling model:

```
Data from file D:/current/EffectLite/simul.csv
Missing data treatment is listwise deletion
Sampling model is fully stochastic
```

A list of the levels of the group variable and their counts (frequencies):

```
Group variable: x
Groups and their frequencies:
  Group   Frequency
    a         348
    b         306
    c         344
with Group a as control group
```

The groups are then shown again, reordered as they will be seen and processed by Mplus. The control group always comes first.

```
Groups, by order of analysis:
      Group   Rel.Freq.
0         a      0.3487
1         b      0.3066
2         c      0.3447
```

A list of the outcome variables:

```
OUTCOME(S):    latent, measured by y1 y2
```

and a list of the covariates:

```
COVARIATE(S):  latent, measured by z1 z2 z3 z4
```

2.18.2 Measurement model

Similar to the manual to the previous version, we happen to treat the two outcome variables as manifest, and the four covariate variables as measures of one latent covariate. However, because we no longer specify the measurement model separately for the outcomes and the covariates, we have to use the 'dummy factor trick' to accommodate the manifest variables. Each dummy factor has only one indicator (manifest variable) loading on it, the loading is 1, and the error variance is zero, as it is easily seen from the output information on the measurement model:

```

*** Measurement model ***

y1 =          0          +1.000 * F1

y2 =          0          +1.000 * F2

z1 =          0          +1.000 * F3      + Error(z1)

z2 =          0.017      +0.966 * F3      + Error(z2)

z3 =         -0.004      +1.002 * F3      + Error(z3)

z4 =         -0.008      +1.021 * F3      + Error(z4)

Error variances

Group      y1      y2      z1      z2      z3      z4
a          ---      ---      0.240  0.232  0.229  0.259
b          ---      ---      0.205  0.245  0.287  0.248
c          ---      ---      0.239  0.254  0.257  0.267

```

Covariances between latent variables are shown below the main diagonal, and correlations above the main diagonal:

```

Covariance / correlation matrix of the latent variables

Group a

      F1      F2      F3
F1    5.101    0.952    0.301
F2    4.959    5.319    0.319
F3    0.655    0.708    0.928

... (similar output for groups b and c)

```

2.18.3 Group means of the covariates

EffectLite reports the group means of the covariates, their standard errors, and the standard deviations:

```

*** Group means of the covariate(s) ***

Covariate  Group      Mean  Std.error  Std.dev.
F3         a        -0.502   0.055     0.963
F3         b         0.039   0.057     0.938
F3         c         0.492   0.057     0.983

```

There is also a simultaneous test that covariate means do not differ across treatment groups:

```

Simultaneous test that covariate means are equal between treatment groups:
Chi-sq    DF    Prob
164.617    2    0.0000

```

If the covariate means are equal across all groups, we have a situation similar to a perfectly randomized experiment, and the average treatment effects (after adjusting for the covariates) will remain close to the unconditional mean differences. This is certainly not the case in the example above: the group means are markedly different and the overall test is highly significant, so it makes sense to proceed further.

2.18.4 Group means of the outcome variable(s)

The next section of the output concentrates on the means of the outcome variable(s). **EffectLite** shows the raw mean, the adjusted mean, and the standard deviation of each outcome variable and each group, along with the standard errors for the two kinds of means:

```
*** Group means of the outcome variable(s) ***

Group Outcome Raw mean SE(Raw mean) Std.dev. Adj.mean SE(Adj.mean)

  a      F1    -0.402      0.121      2.258   -0.043      0.134
  a      F2    -0.367      0.124      2.306    0.022      0.136

... (similar output for groups b and c)
```

Much attention is given to comparing each raw mean to its corresponding adjusted mean:

```
Differences between raw and adjusted means

Group Outcome Difference SE(Difference) Difference/SE

  a      F1      0.359      0.071      5.075
  a      F2      0.388      0.073      5.332

... (similar output for groups b and c)

Simultaneous test that all raw means are equal to the corresponding
adjusted means:
  Chi-sq   DF   Prob
  103.235   6   0.0000
```

When the overall test that all adjusted means are equal to the corresponding raw means rejects the null hypothesis, this is further indication that it makes sense to consider an analysis involving covariates. On the contrary, a non-significant test implies that an analysis without covariates could have been sufficient. Even in that case, including covariates may be valuable in increasing the power of significance tests or in the study of possible interactions between covariates and the treatment variable. A non-significant test does *not* necessarily mean that the covariate-adjusted average effects shown below are causally unbiased: that kind of test is not yet available in the current version of **EffectLite**.

2.18.5 Overall tests of significance

The overall tests of significance are arranged somewhat similar to an ANOVA table. Each test is now computed as a Wald test:

```
* Simultaneous tests for all treatment groups and all dependent variables *

=====
Hypothesis                               Chi-sq   DF   Prob
-----
No average treatment effect: E(g1) = ... = E(g2) = 0      3.076    4  0.5453
No covariate effect in control group: g0 = constant       36.365    2  0.0000
No treatment*covariate interaction: g1,...,g2 = constant  55.370    4  0.0000
No treatment effects: g1,...,g2 = 0                     64.748    8  0.0000
=====
```

The null hypotheses are discussed in detail in sections 1.2.7 (no average treatment effect), 1.2.8 (no covariate effect), 1.2.9 (no treatment \times covariate interaction), and 1.2.10 (no treatment effects).

2.18.6 Average treatment effects

The next section of the output shows results for the average effects of the treatment. Averaging is over the distribution of the covariates.

```
| *** Detailed analysis of the average effects ***
```

Results are shown for each outcome variable...

```
| Results for outcome variable 1: y1
```

...and for each group compared to the control (or reference) group:

```
| Group b - Control group a
|
| Effect E(g1)                0.045
| Std.error                  0.178
| Effect/Std.error            0.252
| Effect size                  0.020
|
| Group c - Control group a
|
| ...
| Results for outcome variable 2: y2
|
| ...
```

Effect size is computed by dividing the average effect estimate by the standard deviation of the outcome variable in the control group.

2.18.7 Deviations of the adjusted means from their averages over groups

A new feature in Version 3 are the deviations of the adjusted means of each outcome variable from their (unweighted) average over all groups. These may be more appropriate than the average treatment effects of section 2.18.6 in situations where there is no prescribed control group against which to measure effects.

```
| Deviations of adjusted means from their average over groups
|
| Outcome      Group      Deviation  SE(Deviation)  Deviation/SE
| F1           a          -0.094      0.108          -0.874
| F1           b          -0.049      0.099          -0.500
| F1           c           0.144      0.106           1.355
| ... (similar output for outcome variable F2)
```

2.18.8 Average effects given a treatment condition

The next section of the output shows the *average effects given a treatment condition*. These effects are perhaps easiest to understand when there are only two groups, treatment and control. In that case, the average effect given the control group is the *average effect for the non-treated*, while the average effect given the treatment group is the *average effect for the treated*. For more details, see section 1.2.6.

Because the number of these effects increases rapidly with the number of groups, we have set an upper limit (currently 40) for the total number of effects that can be shown. Hence, average effects given a treatment condition will not be shown if the number of outcome variables exceeds the following limits:

Number of treatment groups	Maximum number of outcome variables
2	20
3	5
4	3
5	2
6	1
More than 6	No effects shown

```
*** Average effects given a treatment condition ***
```

```
Results for outcome variable 1: y1
```

```
Group b - Control group a
```

Effect given treatment:	0	1	2
Effect $E(g_1 X=i)$	-0.381	0.072	0.451
Std.error	0.184	0.182	0.218
Effect/Std.error	-2.076	0.398	2.072
Effect size	-0.169	0.032	0.200

```
Group c - Control group a
```

```
...
```

2.18.9 Effect functions

The next section of the output shows results useful for the analysis of the effect functions:

```
*** Detailed analysis of the conditional effects ***
```

Again, results are presented for each outcome variable in turn:

```
Results for outcome variable 1: y1
```

The *intercept function* is the regression of the outcome variable on the covariate(s) in the control group. This is the g_0 function discussed above and in Steyer et al. (2007). The output displays the point estimate for each regression coefficient, its standard error, and the ratio of the estimate to its standard error:

	Intercept	F3
Coefficient	-0.048	0.706
Std.error	0.132	0.125
Coeff./SE	-0.361	5.664

For each group other than the control group, there is an *effect function* — these are the g_1 function, g_2 function, etc., discussed above and in Steyer et al. (2007):

```
Effect function g_1: Group b - Control group a
```

	Intercept	F3
Coefficient	0.039	0.838
Std.error	0.176	0.179
Coeff./SE	0.223	4.672

```
Effect function g_2: Group c - Control group a
```

```
...
```

```
Results for outcome variable 2: y2
```

```
...
```

There will be some additional output if the **EffectLite** output option on the Main Options screen is set to **long**: for details, see section 2.5.

Note that the interpretation of the results as average causal effects and conditional causal effects is based on certain assumptions (see Steyer et al., 2007 for details). The present version of **EffectLite** does not test whether this is the case. For background information on causality tests, see e.g. Steyer, Gabler, Davier, and Nachtigall (2000).

References

- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic and statistical considerations. *Journal of Personality and Social Psychology*, 51, 1173–1182.
- Greene, W. H. (1997). *Econometric analysis* (3 ed.). Prentice-Hall.
- Little, R. J. A., & Rubin, D. B. (2002). *Statistical analysis with missing data* (2 ed.). Hoboken, NJ: Wiley.
- Muthén, L. K., & Muthén, B. O. (1998 - 2004). *Mplus. statistical analysis with latent variables. user's guide*. Los Angeles, CA: Muthén & Muthén.
- Pohl, S., Steyer, R., & Kraus, K. (2007). Modelling method effects as individual causal effects. (manuscript submitted for publication)
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66, 688–701.
- Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3), 581–590.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *Annals of Statistics*, 6, 34–58.
- Saunders, D. R. (1956). Moderator variables in prediction. *Educational and Psychological Measurement*, 16, 209–222.
- Steyer, R. (1992). *Theorie kausaler Regressionsmodelle [theory of causal regression models]*. Stuttgart: Gustav Fischer Verlag.
- Steyer, R. (2001). Classical test theory. In C. Ragin & T. Cook (Eds.), *International Encyclopedia of the Social and Behavioural Sciences. Logic of inquiry and research design* (pp. 481–520). Pergamon, Oxford.
- Steyer, R. (2003). *Wahrscheinlichkeit und Regression [probability theory and regression]*. Berlin: Springer.
- Steyer, R. (2005). Analyzing individual and average causal effects via structural equation models. *Methodology*, 1, 39–54.
- Steyer, R., & Eid, M. (2001). *Messen und Testen*. Berlin: Springer.
- Steyer, R., Gabler, S., Davier, A. A. von, & Nachtigall, C. (2000). Causal regression models II: Unconfoundedness and causal unbiasedness. *Methods of Psychological Research Online*, 5, 55–86.
- Steyer, R., Partchev, I., Kröhne, U., Nagengast, B., & Fiege, C. (2007). *Causal effects in between-group experiments and quasi-experiments: Theory*. (Available at www.causal-effects.de)
- Wasserman, L. (2004). *All of statistics: A concise course in statistical inference*. Springer.