Grade Retention in Primary School has Robust and Long-term Positive Effects on Academic Performance and Psychosocial Well-being. How to Avoid Model Dependence and Selection-bias in Grade Retention Research.

**Wim Van den Broeck
Vrije Universiteit Brussel**

**Eva Staels
Vrije Universiteit Brussel**

Acknowledgement: The data in this research were made available and are the property of the 'Steunpunt Onderwijsonderzoek' (SONO), the authors of the SiBo database are Bieke De Fraine en Jan Van Damme, KU Leuven

Abstract

In this study, the effects of grade retention on school performance (mathematics) and psychosocial well-being are analyzed based on a large existing dataset that has already given rise to several publications. What is new in our study is that the data are analyzed for all years of grade retention and that different methods are used for which we thoroughly check the assumptions and if possible, test them. To improve analytical robustness, the effects of grade retention were investigated using different analysis techniques: latent growth models (LGM), cross-lagged panel models (CLPM), marginal structural models (MSM) and Sequential Conditional Mean Models (SCMM). Across the different methods, it appears that grade retention has a substantial and long-term positive effect on mathematics performance and well-being. The more the methods keep selection-bias and other forms of bias under control, the larger the estimates of the effects. Methods that make full use of longitudinal data are the least bias-sensitive. Finally, the size of the individual effects appears to depend exclusively on the intelligence of the pupil and on the ratio of grade repeaters in the school.

Keywords: grade retention, selection-bias, model-dependence, latent growth model, longitudinal design

GRADE RETENTION IN PRIMARY SCHOOL

The effects of grade retention on school performance and psychosocial well-being, according to recent research, are far from unequivocal. Worldwide, grade retention is the practice of having students who have great difficulty mastering the subject matter of the schoolyear repeat that same schoolyear. Advocates believe that in this way such pupils are given the opportunity to catch up with the other pupils so that they do not get completely stuck in the next schoolyears (e.g. Lorence, 2014). Opponents think that repeating a year mainly has negative effects because of the failure experience this entails and that it is better to promote these pupils to the next schoolyear, eventually with additional support (e.g. Jimerson, Woehr, Kaufman, & Anderson, 2004).

Most studies of the first generation found that grade retention generally has negative effects on school performance (cf. meta-analyses conducted by Holmes, 1989; Jimerson, 2001; Allen, Chen, Willson, & Hughes, 2009). Also, a few recent studies reported negative effects (e.g. studies based on the dataset that we will reanalyze in this article: Goos, Van Damme, Onghena, Petry, & de Bilde, 2013; Vandecandelaere, Vansteelandt, De Fraine, & Van Damme, 2016_a; Vandecandelaere, Vansteelandt, De Fraine, & Van Damme, 2016_b; see also: Hughes, Cao, West, Smith, & Cerda, 2017; Hughes, West, Kim, & Bauer, 2018)). However, Allen et al. (2009) concluded in their meta-analysis that the negative effects in many studies may have been the result of insufficient control of selection-bias. This implies that the negative effects may not have been due to grade retention per se, but rather to the fact that the group of repeaters was initially already weaker compared to the group of non-repeaters. The more selection-bias is brought under control, the less negative (and sometimes even positive) the effects turn out to be. Some more recent studies also reported positive effects of grade retention on performance motivation (Cham, Hughes, West, & Im, 2015) and on school performance (Marsh et al., 2017).

GRADE RETENTION IN PRIMARY SCHOOL

The conclusion of Allen et al. (2009) that the direction and magnitude of the estimated effects of grade retention depends on the extent to which confounding variables are under control and thus selection-bias can be avoided, leads to the research question that we want to answer in this article. Do research designs and data-analytical techniques that keep the confounding variables under control, if not perfectly, at least adequately, show positive effects of grade retention? Before reporting the results of a re-analysis of the data of one of the most comprehensive existing datasets in primary school students, we will discuss the main methodological issues necessary for a good understanding of the challenges faced by the researcher.

We successively discuss the problem of selection-bias, the implications of grade retention as a time-varying predictor, the assumption that the intervention in one individual is independent of the intervention in other individuals, i.e. the 'stable unit treatment value assumption' (SUTVA), and finally which comparison group is preferable, a group of same-aged peers or a group of grade classmates.

Selection-Bias as the Key Problem

An important problem in the study of the effects of grade retention is to find a group of non-repeaters that is as equal as possible to the group of retained pupils in all respects except that they are in fact not retained. In most studies into the effects of grade retention, the between subjects randomized experiment is taken as a reference paradigm. The researcher attempts to transform the initial situation in which natural selection determines the condition to which a participant ends up resemble as much as possible a real randomized experiment in which only chance determines the condition to which one belongs. The researcher tries to achieve this by measuring as many variables as possible on which the two groups, the retained

GRADE RETENTION IN PRIMARY SCHOOL

and the promoted children, could differ from each other¹. Then the two groups are somehow matched with each other on all these potentially confounding variables in the hope to simulate an ideal experiment in which the only essential difference between the groups is the treatment status. Formulated in terms of the 'potential outcomes' model (Rubin, 1974, 2005), which is often used as a frame of reference in causal modelling, the subjects of the two groups must be perfectly interchangeable (exchangeability assumption), given their scores on the covariates (conditional exchangeability). Only in that case are the scores of subjects in one group, as in a real experiment, representative of what would have happened if subjects from the other group had ended up in this group, and vice versa. In the non-experimental context of grade retention, it is precisely because several pupils have specific characteristics that do not occur in the other pupils that the school team advises them to repeat the year. Such differentiating characteristics can on the one hand relate to variables that the researcher does not know and therefore have not been measured, or they can relate to measured variables on which the two groups have little overlap. So, we distinguish two problems, namely the problem of *unmeasured confounders* and the problem of overlap (*positivity assumption*).

Measured and Unmeasured Confounders

If the school team makes use of characteristics that are unknown to the researcher or that were not included in the study, then selection-bias is inevitable, and the observed effects of retention will be less positive than they really are. It is therefore key to carefully include as many potential confounders as possible in the study. Often without empirical evidence researchers assume that no substantial unmeasured confounders remain. However, it is possible to determine the percentage of the explained variance of retaining vs. promoting based on all included covariates preceding retention. If that percentage is not high, then it may

¹ If these covariates are also related to the outcome variable, then they are confounding variables.

GRADE RETENTION IN PRIMARY SCHOOL

be assumed that not all confounders were included, and that selection-bias is likely. But also matching on known and measured variables is not always simple and can easily lead to selection-bias. We illustrate this with two published studies that also based their analyses on the dataset analyzed in this contribution.

A frequently used matching technique in research into the effects of grade retention is propensity score matching. For all subjects the probability of retention is calculated, e.g. through a logistic regression analysis, given a series of covariates considered as potential confounders. Next, the groups are matched on these propensity scores by constructing a subgroup of control subjects with similar propensity scores as the repeaters. After this matching, a final check is carried out on any remaining differences between the groups on the individual covariates. A variant of this technique, in which the propensity scores are categorized in ten equal strata, was used in the study by Goos et al. (2013). In a table that was not published in the article but in a separate online publication (Goos, Van Damme, Onghena, & Petry, 2011), it can be seen that the initial bias in the tenth stratum, which contains the highest propensity scores and in which 86% of all repeaters are located, is very large. The probability of retention in this stratum is much higher for the retained than for the non-retained pupils (logit = 0.36 vs. -1.25). Because there are only few retained children in the first nine strata, the average propensity score of all non-retained pupils is also substantially lower than that of the repeaters ($t = 33.66$, $df = 2423$, $p < 0.0001$). The effect size of this initial bias expressed in Cohen's d is 2.54. Thus, the imbalance between the two groups, even after this matching procedure, is enormous. Evidently, it is not surprising that negative effects were reported on both school performance and psychosocial functioning. It is striking that the negative effects found are much smaller than the size of the initial bias. In spite of this bias, the authors report that the intra-stratum balance was valid for both propensity scores (which is correct for the nine first strata sparsely populated by repeaters, but not for the tenth stratum

GRADE RETENTION IN PRIMARY SCHOOL

densely populated by repeaters) and for at least 95% of the 169 observed covariates. However, this 'covariate balance check' between the groups did not play a role in the analyses, because the strata were used and not the propensity scores themselves.

But also, a correct application of the 'covariate balance check' is not without problems and can cause bias. For example, in the study by Vandecandelaere et al (2016b), using another matching technique, i.e. 'inverse probability weighting', a covariate balance check was also performed. However, the balance appears to have succeeded only at the first timepoint, but not at the second and third timepoint. The initial imbalance without matching does decrease after matching, but clearly does not disappear. The applied rule of thumb that the standardized difference between the two groups on each covariate should not exceed .25 can easily lead to a situation where a large number or even most of the differences point in the same direction, and consequently to a systematic bias. This classical balance check (absence of significant differences² or differences smaller than .25 standard deviations), by considering each difference separately, does not consider capitalization on chance. Fortunately, matching methods have recently been developed that do not achieve the match via a propensity score, but directly match the covariates themselves with great precision between the two groups (e.g. Coarsened Exact Matching, Iacus, King, & Porro, 2012; and Entropy Balancing, Hainmueller, 2012).

Overlap in Distributions and the Importance of the Target Population

In order to make the groups comparable, it is essential that they overlap sufficiently in the distributions of the different covariates (*positivity assumption*). This implies that for each combination of the values of the covariates it must be possible to belong to one group

² This practice is based on a misinterpretation of the p-value. The absence of a significant difference does not imply that in the population the difference between the two averages is non-existent. For example, a p-value of .20 means that if there is no difference in the population, there is one chance in five that we see the observed difference or a larger difference in the sample. This probability does not give too much confidence in the null hypothesis.

GRADE RETENTION IN PRIMARY SCHOOL

(repeaters) or to the other (non-repeaters). In order to achieve a larger overlap, researchers often select the sample so that a minimum probability of being retained is guaranteed. For example, the study by Vandecandelaere et al (2016_b) used a cut-off propensity score of at least five percent chance of being retained, given the covariates. The consequence is that a significant proportion of the repeaters are excluded: for the first and second schoolyear this is about half of the retained children. One could assume that these children are mainly the 'better' ones because of their lower propensity score. However, the conclusions about the effects of retention are then no longer representative for the entire target population; this is also a form of selection-bias. Given that the logit of the propensity score of most of the repeaters is negative (i.e. they are more likely to promote than to repeat), this implies that some of the essential characteristics of being retained remain unnoticed ('unmeasured confounders').

It is remarkable that in many studies in which a selection is made of the group of repeaters, the conclusions are often formulated as if they referred to the non-selected, complete group of repeaters. The precise definition of the target population is essential in this type of research and should not depend on practical methodological considerations. The target population at least includes all pupils who were retained, and possibly also all pupils with a high risk of being retained, e.g. those who were advised by the school to repeat, but who in fact were promoted. Pupils with only a 1 in 20 chance of being retained, estimated based on the measured covariates, are not this target population. Exactly how the effects of retention can be studied for this complete well-defined target population, we discuss in the methods section.

Grade Retention is a Time-varying Predictor

Grade retention is not a one-off 'fixed' predictor because in each grade a pupil has a probability to repeat, and whether to repeat influences the probability to repeat in the

GRADE RETENTION IN PRIMARY SCHOOL

following grades. In other words, grade retention is a time-varying predictor.

Vandecandelaere et al. (2016_b) report that 31.9% of the non-repeaters in third kindergarten with at least a 5% chance of being retained, did in fact repeat in first grade. So, for this group, promoting leads to a higher probability of retention one year later, and vice versa, being retained leads to a smaller chance to be retained one year later. The latter can be the result of a direct effect of being retained in one grade on the probability to be retained in the next grade, e.g. because teachers want to avoid retention two times in a row, but it can also be the result of a positive effect of retention on the mathematical performance so that the need to be retained the following year is reduced. In any case, the fact that the probability of being retained in a later grade is higher among the promoted children who have been matched with children who have been retained in a previous year proves that in order to correctly estimate the effect of retention in a specific grade, the history of grade retention should also be taken into account.

In general, therefore, the more information about learning performance is included in the prediction model, e.g. about the history of retention, the better the match between the two groups. This implies that the quality of the matching will also increase as more preceding measurements of learning performance are available, i.e. as pupils are more advanced in their school careers.

An important consequence of including time-varying predictors is the threat of creating so-called "collider bias" (cf. Elwert, 2013; Elwert & Winship, 2014). When controlling for a variable (C) that is connected via two causal paths to two variables (A and B) that are initially uncorrelated, a spurious correlation between A and B emerges. In a longitudinal design in which a time-varying explanatory variable (treatment) and a time-varying covariate mutually influence each other, controlling for a confounder can create both 'over-control bias' and 'collider bias' if there is also an unmeasured confounder (see Daniel,

GRADE RETENTION IN PRIMARY SCHOOL

Cousens, De Stavola, Kenward, & Sterne, 2013). Vandecandelaere et al (2016_b) dedicated an excellent discussion to this problem and propose a solution in the form of "marginal structural models" with the aim to estimate patterns of the time-varying intervention variable or so-called "joint causal effects". This is in fact a collection of direct effects that can be estimated for each intervention pattern (e.g. never been retained, 0 0 0..., or only been retained on the second time point, 0 1 0 0 ...). In each of these patterns, the direct effect, i.e. not mediated by interventions on other time points, is examined. Comparing fixed intervention patterns is interesting, but what we want to know from a practical point of view is what the total effect is of grade retention at a specific moment. The total effect is the sum of direct and indirect (mediated) effects. How these total causal effects can be estimated without bias is explained in the methods section.

The 'Stable Unit Treatment Value Assumption' (SUTVA) and Indirect Effects

In the context of the 'potential outcomes' model, SUTVA is a basic premise on the stability of effects assuming that the potential outcome of an intervention in one individual is not influenced by potential outcomes of interventions in other individuals (Rubin, 1986; Morgan & Winship, 2015). Applied to this study: the effect of retaining an individual should not be influenced by the fact that other individuals may or may not be retained. This standard assumption is far from self-evident, even though it has seldom been addressed in existing research on retention. Pupils can influence each other directly, and it may also make a difference whether many pupils in the classroom are retained or none. There may also be an 'indirect' effect³ of using grade retention as an educational strategy, e.g. because pupils make an extra effort to avoid retention. The individual indirect effect of grade retention as a system is then the difference in outcome between a student who is not retained in a system where

³ This 'indirect effect' must be distinguished from the classical meaning of an indirect effect, as explained in the previous paragraph, i.e. as a mediated effect.

GRADE RETENTION IN PRIMARY SCHOOL

grade retention is applied, and an otherwise comparable student who is not retained in a system where grade retention is not applied (see also Morgan & Winship, 2015). The total effect of retention is then the difference in outcome between a student who is retained in a system where grade retention is applied, and a comparable student who is not retained in a system where no one is retained. If the literature mentions a zero effect of retention (cf. Allen et al., 2009), this only concerns the direct effect of retention, which is only a part of the total effect.

Age or Grade Comparisons?

In the study of grade retention, there has always been a discussion about what type of comparison, the comparison with peers or the comparison with classmates, is the most adequate to assess the effects of retention (see Lorence, 2006; Steiner, Park, & Kim, 2016). In the comparison with peers, those who are retained are compared with pupils of the same age who have moved on to the next year. The two groups are of course comparable in age, but not in terms of the subject matter they were taught. The comparison with classmates compares repeaters to either their younger classmates or their peers when they were in the same grade a year earlier. Here the two groups are equal in terms of the subject matter they were taught, but not in age.

Some researchers prefer peer comparison because it enables the intended counterfactual statement under the Neyman-Rubin 'potential outcomes' approach (Rubin, 2005), i.e. which school performance a repeater would have achieved if she/he had not been retained. This is also the preference of Goos et al. (2013), Vandecandelaere et al. (2016a) and Hong & Raudenbush (2005). Critics of this comparison method believe that this method is to the disadvantage of the repeaters, precisely because the pupils who move on have already been able to benefit from the subject matter of the next school year. The question is therefore whether this benefit concerns year-specific knowledge and skills, or a growth of the

GRADE RETENTION IN PRIMARY SCHOOL

underlying general mathematical knowledge and skills. In order to examine the effects of grade retention on the development of the underlying mathematical knowledge, it is necessary that the measurement measures the same construct over the schoolyears, i.e. 'measurement invariant'. Researchers therefore work with so-called vertically linked scores that are situated on a latent variable, which is comparable across schoolyears. In such a latent trait model, several items function as so-called "anchor items" that are repeated at the different measuring moments with the aim of constructing a common scale over the years. It has then to be examined whether the crucial assumption of measurement invariance across measurement moments is correct. A test is measurement invariant if the differences between groups or moments are only due to differences in the underlying latent variable and not to the fact that across groups or moments the items are solved in a different way.

To avoid the problem of measurement variance, some researchers opt for a grade comparison. However, the causal effect using such a comparison cannot be interpreted unambiguously. After all, the intervention of grade retention not only implies that the pupils repeated one year, but also that they are one year older than those in the comparison group. As a result, the effect of retention cannot unequivocally be attributed to retaining per se, but it can also be the result of being older (the biological or psychological age). This is called a double dose effect (Steiner et al., 2016). Although the causal effect is not unambiguous here, for practical reasons, the only question that matters is whether the intervention of retention (including the age effect) is more beneficial in the long term than being promoted. In principle, this question can be answered by means of a grade comparison.

Our position in this debate is that preference for one or the other method of comparison is largely an empirical question. If after testing measurement invariance across test moments it turns out that the tests are measurement invariant, then both age and grade

GRADE RETENTION IN PRIMARY SCHOOL

comparisons are possible and relevant. If measurement invariance cannot be demonstrated, then only grade comparison is viable.

Taking all these methodological issues into account, we are seeking answers to the following research questions: 1) If selection-bias is sufficiently under control, what is the effect of grade retention in each grade on mathematical performance and on psychosocial well-being up to and including the end of primary school? 2) Do the effects of retention vary according to pupil characteristics, such as socio-economic status, intelligence, etc.? 3) Is the effect of retention dependent on the number of repeaters in a classroom (indirect effect)?

Methods

Participants and Research Context

The reported research is based on the data of about 9000 Flemish pupils in more than 200 primary schools, collected in the context of the longitudinal SiBo project (now SONO) from 2002 to 2010/2011. The core of the SiBo project is based on a random stratified sample of 122 schools and about 6000 pupils who turned out to be representative of the Flemish school population (Verhaeghe, Maes, Gombeir, & Peeters, 2002). However, this core sample was extended to include schools with a high percentage of disadvantaged pupils, as well as a few method schools. 51% of students are boys, 49% girls. For 70% of the pupils the home language is Dutch, for 15% it is mixed and 15% has another home language. 77% of the pupils are of native origin, 10.5% of Moroccan or Turkish origin, 5% of other Western European origin, and 7.5% of other non-Western European origin. In the analyses presented, all pupils are followed from the last kindergarten class to the end of primary school (ages 5 to 12). In Flanders, standard tests are not used when deciding whether to repeat a year or to promote, but in practice the decision is made after consultation between the parents, the school team and the center for pupil guidance. The percentage of retained students in Flanders is relatively high compared to other European countries. However, this percentage has fallen

GRADE RETENTION IN PRIMARY SCHOOL

substantially in recent years, perhaps partly as a result of an intensive campaign against grade retention. In 2011, the Flemish government reported that 17.15% of all primary school pupils were at least one year behind, compared to 11.77% in 2018-2019. More detailed information about the participants and the Flemish educational context can be found in earlier publications using the same dataset (see Goos et al., 2013; Vandecandelaere et al., 2016a).

Measures

A large number of variables were measured at the level of the pupil, the teacher, and the school. We focus here exclusively on those variables that play a role in the analyses we will present. Each year, at the end of the schoolyear, mathematical skills were measured, as well as at the beginning of the third kindergarten class and first grade. Each test consisted of 50 to 80 items covering the following domains: number sense, number procedures, measurement, geometry, and applied math problem solving. The reliability coefficients (Cronbach's alpha) of these tests ranged between .88 and .94. The math scores at the different measurement moments were converted into a common latent scale score via a three-parameter IRT model with the aim of making the scores comparable over the different grades (Verhaeghe, 2010).

Psychosocial functioning of the pupils was examined annually by means of a questionnaire filled in by the teachers. Subscales were constructed for social relationships, cooperative behavior, attitude to work, self-confidence, aggressive behavior, anxious behavior, hyperactive behavior, anti-social behavior, independent behavior, well-being, and self-confidence. The reliability coefficients (Cronbach's alpha) of these scales varied between .83 and .94.

Information on potentially confounding variables was collected through official data, test material, and questionnaires completed by teachers or parents. The baseline covariates are month of birth, gender, socio-economic status, home language, ethnicity, and high-risk pupil

GRADE RETENTION IN PRIMARY SCHOOL

(as defined by the Flemish government). Also, intelligence, measured by Ravens' progressive matrices, was included as a fixed covariate. For more information see Goos et al., 2013.

Data-analysis and Results

Measurement Invariance

As a first step in the data analysis, we examine whether the scores on the crucial dependent variable, the math scores, are measurement invariant across grades. The question is whether the differences in the scores can be interpreted as differences in general mathematical skills, or whether the differences have rather to do with being able to solve specific items. Only if the score differences are caused by differences in general mathematical skills, age comparisons are possible. Figure 1 shows schematically what a growth model looks like over three measurement moments. η_0 and η_1 are the linear growth parameters (intercept and slope) of mathematical ability. Between every two time points there is at least one overlapping item (items 1 and 2 between T1 and T2, and item 4 between T2 and T3).

Before such a growth model can be applied to the data, it must first be examined whether the measurements are really measurement invariant over time, i.e. they measure the same underlying construct. In the procedure followed, measurement invariance is checked by sequentially testing several nested models, each time imposing additional restrictions (Meredith, 1993; Millsap, 2011). In a first step, an exploratory factor analysis is used to find a good fitting model per time moment. Then, for the overlapping items, it is successively tested whether there is configural invariance (the same factor structure applies to the different moments), metric or weak invariance (the items have the same factor loadings at the different moments), scalar or strong invariance (additionally, the intercepts or thresholds of the items

GRADE RETENTION IN PRIMARY SCHOOL

are kept equal), and strict invariance (additionally, the unique factor variances are kept equal). Each time, the fit of two successive models is compared, and if that fit is not statistically different, the most parsimonious model, which imposes more restrictions on the data, is retained. As is detailed in Appendix A, measurement invariance was not achieved for the transition between grades one and two, and grades three and four, and often there were not enough overlapping items to test measurement invariance.

These analyses on measurement invariance, and the partial absence of sufficient overlapping items taken in the complete sample, show that the hypothesis of measurement invariance for the mathematic scores cannot be maintained. This implies that these scores at the different measurement moments cannot simply be compared with each other and that the parametric form of a growth model, as shown in Figure 1, cannot be determined. In conclusion, a comparison of repeaters with their peers one grade further is not an option, precisely because such comparison only makes sense under the hypothesis of measurement invariance. A comparison with grade classmates is, however, possible. For practical purposes, the underlying unidimensional mathematics scale, as constructed by the authors of the original study, can be regarded as a monotonously increasing ordinal variable over time. Within each moment of measurement, the scale can be considered as a metric scale.

Data-analytical Strategy

In the context of observational studies, where the data generating process is unknown and the researcher tries to identify this process by specifying and fitting one or more models, the problem of model dependency is impending (King & Nielsen, 2019). Even small differences in the specification of models that fit the data approximately equally well can lead to substantially different conclusions about the underlying causal process. Model dependency gives the researcher discretion that opens the way for bias (King & Nielsen, 2019).

GRADE RETENTION IN PRIMARY SCHOOL

A typical way to make an observational design as similar as possible to an experimental study is to statistically control for confounding variables (C) by including them in a multiple regression analysis (see figure 2). In this case, the influence of grade retention (GR) on the math scores (M) is statistically corrected for the influence of the confounding variables (C). However, this method entails a series of choices that are not simply determined by the data; e.g. choice between a linear or a quadratic model, whether to add interaction terms, omitting non-significant predictors, etc.

Figure 2

Another way to eliminate the role of confounding variables is matching (Figure 2); in this study by equating the values of the confounding variables between the group of retained children and the group of promoted children. This method, of which there are numerous variants, in fact boils down to pruning the data of subjects for whom there is no match on the confounding variables. The advantage of this technique is that if the groups can be balanced, the problem of model dependence and the ensuing bias is solved. A crucial condition is that all potentially confounding variables must be known and measured, which is not always obvious.

In the context of the discussion about the credibility of scientific findings in psychology and in the broader social sciences, it is of the utmost importance that the data and methods are transparent, that the analyses can be reproduced, that the results are robust across different data analytical methods, and that the results can be replicated in independent samples (LeBel, McCarthy, Earp, Elson, & Vanpaemel, 2018). The data of this study were

GRADE RETENTION IN PRIMARY SCHOOL

collected in the framework of the SiBo project⁴ (Verhaeghe et al., 2002) and can be consulted by independent researchers. The code files for the different analyses can be obtained on simple request from the first author. To answer the research questions, we use different analysis techniques: latent growth models (LGM), cross-lagged panel models (CLPM), marginal structural models (MSM) and Sequential Conditional Mean Models (SCMM). If the results of these analyses all point in the same direction, this contributes to the analytical robustness of the conclusions (LeBel et al., 2018).

Description of Growth Curves of Retained Children in Different Grades

Another way to avoid model dependence and promote analytical robustness is to make use of the richness of longitudinal data by plotting the data over time in a purely descriptive way and allowing the data to speak for itself as much as possible. For each individual, the mathematical ability, except for missing values or drop out, was measured at nine consecutive moments. A growth curve can then be determined for each child. The effect of an intervention becomes noticeable due to a sudden deviation in the growth curve. The methodological principle here is that of repeated measurements, in which the individual forms his own control and provides firm protection against the threat of stable unmeasured confounders. In order to detect systematic effects, it is of course advisable to plot the growth curves of different groups of retained individuals. We calculated the average math scores per measurement moment for each group of repeaters separately (see Table 1 and Figure 3). For example, for all pupils who were retained in third kindergarten class, we present their math scores for the first time they were in third kindergarten class, for first grade, second grade, etc. The math scores for the year in which the pupils were retained are not presented. In this way, we can describe the

⁴ SiBo is an acronym for 'Schoolloopbanen in het Basisonderwijs' ('School careers in primary education') and is now part of the SONO project (SONO stands for 'Steunpunt Onderwijsonderzoek' (Educational research support center)).

GRADE RETENTION IN PRIMARY SCHOOL

development of the math scores before and after retention of comparable pupils who were retained at least once.

Table 1

Figure 3

Without performing any statistical test, several important systematic observations can be reported. The average growth curve is monotonously increasing, but the rate of growth decreases with age. If the growth after retaining a grade is larger than what can be expected based on the average growth curve, this implies that the retained children are catching up. Even though strictly speaking the shape of the growth curve cannot be compared because of the lack of measurement invariance, it is still possible to determine the relative position of a group of retainers compared to all other students per measurement moment, and thus also a change in relative position. Figure 3 shows that retainers in the first to fifth grade score systematically very poorly in the year before they get retained compared to all other learners (under the first standard deviation), but each time in the grade after retention show a strong growth that is steeper than the average growth, so that they catapult themselves back into the group (within the first standard deviation of the total group of learners). For those who retained in third kindergarten class, the pattern is different. They are not necessarily the weakest pupils in the third kindergarten class. It is well known that the age of preschoolers in the context of the concept of 'school maturity' is often an argument for teachers to consider retention (Hong & Raudenbush, 2005). The positive effect of retention for this group manifests itself more slowly and apparently continues in the longer term. The lack of an immediate positive effect for this group in first grade is not surprising, because repeating third

GRADE RETENTION IN PRIMARY SCHOOL

kindergarten class does not yet involve direct instruction in mathematics that could be directly benefited from in the first year of primary school.

However, the data in Figure 3 on which these observations are based do not consist of groups of pupils who remain precisely the same over the measuring moments. As a result of early drop-out (censoring) or late intake, the composition of a group of retained students is always slightly different. To prevent selective drop-out from distorting these observations, we also calculated the average math scores for groups of repeaters who had both a measurement in third kindergarten class (beginning or end of third kindergarten class) and a measurement at the end of sixth grade. Intermittent missing values were allowed. Table 1 (second rows) shows these averages. Because the observed effects are the same, we do not present an additional figure. These data strongly suggest that the described effects are not substantially influenced by selective drop-out.

Carefully inspecting data for noticeable and systematic trends is an important and necessary first step in data analysis because they provide a normative framework for comparing model-based tests. In order to correctly estimate the magnitude of effects, statistical tests are of course a next step in data analysis.

Longitudinal models (Latent Growth Models and Cross-Lagged-Panel Models)

An analysis technique that is directly related to the idea of following-up the development of mathematical ability longitudinally, in function of an intervention (grade retention), is a longitudinal growth model with time-varying covariates (cf. Bollen & Curran, 2006). Precisely because grade retention is a time-varying predictor, and an individual may be retained one moment and not the next moment, each individual forms his own control. It is therefore possible to construct such a model without using a control group. This has the advantage compared to other techniques that all retained children can be included in the analysis. On the other hand, it is also possible to estimate such a model for a larger group, i.e.

GRADE RETENTION IN PRIMARY SCHOOL

including all pupils for whom it was advised to repeat. This group could be termed the target population, because it is this group for which we want to know what is best, to retain or not. Of the 9174 pupils for whom at least one measurement of math achievement was taken, 3310 (36%) were advised to retain at least once during their primary school period. Of these 3310 pupils, 2401 pupils (72.5% of the group with advice) were retained at least once. We present a latent growth model for the group of repeaters and one for the larger target population.

In a latent growth model, the consecutive math scores (third kindergarten to sixth grade) are regarded as indicators of the underlying growth curves in mathematical skills (see Figure 4). Two latent growth factors, the intercept (I) and the slope (S), on which individuals can differ, are estimated based on the indicators. The loadings of the intercept factor are kept constant at 1 because the influence of the intercept is the same for all indicators. The influence of the slope factor varies with time. Since we do not have an a priori hypothesis on the exact shape of the growth curve, except that it must be monotonously increasing, the slope factor loadings are freely estimated, except for the first two loadings (0 and 1) to determine the metric of the latent variable (see Meredith & Tisak, 1990).

Figure 4

The results of such a growth model can be interpreted in two different but equally valid ways (Bollen & Curran, 2006). An interpretation that is directly related to our research question is to consider the effects of the time-varying covariates (grade retention at different times) as the unique effects on the growth process, after the growth process has been estimated. Conversely, the growth process is also estimated as the net model regardless grade retention, i.e. if all retained children would never have been retained. For the model with only repeaters the estimation of the growth model therefore represents a counterfactual, for the

GRADE RETENTION IN PRIMARY SCHOOL

model with the target population the growth model is partly based on the observed math scores of pupils who at least once received advice to retain, but de facto never were retained.

The latent growth model for the group of retainers shows a reasonable fit with the data ($\chi^2=488.68$; $df=53$; $p=.000$; $RMSEA = .059$; $CFI = .929$).⁵ Table 2 and Figure 5 show that the effects of grade retention, in any grade, are positive and substantial (medium to large sized in terms of Cohen's d). Exactly as with the descriptive data, after grade repetition, the repeaters show an acceleration of growth after which the growth returns to normal. The effects of retention are greater as the effects are considered in later grades, especially in the fifth and sixth grades. In order to properly understand this increase in effects, it is necessary to correctly interpret the lower 'counterfactual' curve in Figure 5. At each point in time, this curve gives the estimation of the math scores of the pupils who were not retained until that moment. For example, in the second grade, the point on that curve gives the estimate of the math scores for students who did not repeat third kindergarten class or first grade, but who will be retained in a later grade. They score lower than students who were already retained. All points on this curve up to and including fifth grade are therefore not counterfactual because they are based on real data. Only the last point on the curve, for sixth grade, is counterfactual. This point represents the estimation of the model of the average math score in sixth grade that students would have achieved if they had never been retained. Except for those who were retained in third kindergarten class, pupils are initially less weak in mathematics the later they were retained (see Table 1 and Figure 3). Because the counterfactual growth curve for each time point is based on students who have not yet been

⁵ An alternative technique that also makes maximum use of the longitudinal character of the data, but avoids the question of whether the math scores can be considered as an underlying latent variable, is the cross-lagged panel model (CLPM; cf. Allison, Williams, & Moral-Benito, 2017). However, this model does use a latent variable, α , which explicitly embodies the idea that all unmeasured confounding variables are constant within each individual from time point to time point. Using this model, the estimates of the effects are very similar to those of the latent growth model.

GRADE RETENTION IN PRIMARY SCHOOL

retained, i.e. on the relatively better students, this implies that the estimates of the effect sizes of grade retention may underestimate the real effects. In other words, the counterfactual curve does not give an accurate picture of what would have happened to the math scores if students who were retained in the first, second or third grade had not been retained in these grades. The counterfactual growth curve of these students would probably have been lower.

Table 2

Figure 5

The latent growth model for the ‘target population’ shows a reasonable fit with the data ($\chi^2=554.83$; $df=53$; $p=.000$; $RMSEA = .054$; $CFI = .935$). The effects estimated for the growth model of the target population are somewhat smaller, especially for the last grades, but again significant and substantial (see Table 2). The reason why the estimated effects are somewhat smaller here can be explained by the fact that there is now for all ages a group that never was retained but was advised to do so. The estimated math scores for this relatively ‘better’ group (not counterfactual) are somewhat higher, especially in the last grades, so that the effects are smaller. For this population, the underestimation of the real effects of retention is therefore probably even greater than for the group of repeaters.

Effects on Psychosocial Well-being

In studies that report negative effects of grade retention on learning performance (e.g. Goos et al., 2013), these effects are often attributed to the negative influence that retention would have on pupils' psychosocial well-being. In this study, well-being was measured each schoolyear based on four items with a six-point scale. As the well-being scale was constructed by the authors of the original study (Verhaeghe et al., 2002) as the average of the four items,

GRADE RETENTION IN PRIMARY SCHOOL

we do not consider psychosocial well-being as a latent variable, and apply a cross-lagged panel model (see Figure 6). Such a model allows control for unmeasured time-invariant confounders, while mutual causation between the time-varying intervention and time-varying covariates is possible (Allison et al., 2017). This model showed an excellent fit with the data ($\chi^2=20.68$; $df=13$; $p=.0796$; $RMSEA = .016$; $CFI = .991$)⁶. Generally, the effects of grade retention on well-being appear to be positive (medium to large effects) and long-lasting over all grades (see Table 3).

Figure 6

Table3

Missing Values

In the models presented so far, the assumption is that the missing values of math scores are 'missing at random' (MAR) (Little & Rubin, 2002; Enders, 2010). This means that the missing values may depend on other included variables (e.g. previous math scores), but not on the values themselves of the variable under consideration. For example, if students drop out, precisely because their math score is very low, if it would have been measured, then it is 'missing not at random' (MNAR) ('nonignorable missingness'). If the missing values can be considered as MAR, then the growth models can be correctly estimated (without bias) using the Maximum Likelihood estimation procedure (Schafer & Graham, 2002). However, it is conceivable that the missingness, especially in case of permanent drop-out, may be the result of very poor math performance at that time. Therefore, we cannot simply assume the MAR assumption.

⁶ A latent growth model resulted in convergence problems and thus proved difficult to fit, perhaps due to the fact that for some timepoints there were too many missing values.

GRADE RETENTION IN PRIMARY SCHOOL

In order to meet this problem, different types of models were designed, selection models and 'pattern mixture' models (Enders, 2011; Muthén, Asparouhov, Hunter, & Leuchter, 2011). We have chosen the Diggle-Kenward selection model (Diggle & Kenward, 1994) which simultaneously estimates a substantive model (a latent growth model) and a set of regression equations that predict the probability of missingness. In this model, missingness is predicted based on the observed math scores and on the assumption that the repeated measurements are multivariate normally distributed (see Figure 7). In this model, R1 through R6 are discrete time survival indicators of missing values that indicate whether an outcome variable is missing or observed and the arrows in dotted line represent logistic regression equations (i.e. the probability of drop-out). In the model, 'missingness' is predicted based on the previous math score (observed) and on the current math score (by definition missing if $R = 1$). The logistic regression equations give an indication of the mechanism underlying missingness. If particularly the math scores of the previous measurement moment significantly influence the 'missingness', then this is according to a MAR mechanism, if the current math scores rather have a significant influence on the missingness, then this indicates an NMAR mechanism. Intermittent missing values are considered as MAR by default in this model.

Figure 7

The estimated model shows that drop-out at the end of third kindergarten class and at the end of first grade is strongly determined by the weaker math skills at those times points (NMAR). At the end of third kindergarten, the estimated odds of dropping out for children who score one unit lower on the current math scale are 1.88 times the odds for children who score one unit higher on the current math scale ($SE = .053$, $p = .000$) and 1.86 times the odds ($SE = .090$, $p = .000$) at the end of first grade. For the later grades a MAR mechanism

GRADE RETENTION IN PRIMARY SCHOOL

underlies missingness. The odds of dropping out for children who score one unit lower on the previous math scale are 1.09 to 1.41 times the odds for children who score one unit higher on the previous math scale (SE's between .022 and .036, all p 's = .000). This implies that the estimate of the counterfactual growth curve in the model shown in Figure 5 (lower curve) may still be strongly overestimated because the weakest pupils already disappeared from the data set at the beginning of the data collection (end of third kindergarten class or end of first grade). If they had remained in the dataset, the estimated performance of the children who never were retained (bottom curve in Figure 5) might have been even lower. The estimated effects (see Table 4) confirm this suspicion. Especially the estimated effects for the last measurement (the only real counterfactual data point) are almost double compared to those of the standard latent growth model (Figure 4). Most of the weakest pupils probably leave regular education as early as after third kindergarten class or after first grade and end up in special education. It is therefore perhaps unrealistic to consider these pupils as belonging to the target population for which grade retention is a real option.

Table 4

Do children's characteristics moderate the effects of grade retention?

A next step in the analysis of the growth models concerns the question whether the effect size of grade retention on the development of math achievement depends on specific pupil characteristics. We therefore extend the latent growth model in Figure 4 by including a latent random variable (slope of time-varying covariates, STVC, see Figure 8). This variable allows the magnitude of the effect of grade retention to be different for each individual, and we can also examine whether the magnitude of the effects of grade retention depends on specific learner characteristics by estimating the influences of different time-invariant covariates on STVC. Not unexpectedly, this estimated model shows that a later month of birth,

GRADE RETENTION IN PRIMARY SCHOOL

a lower SES, a different home language, a different origin, being a high-risk student, and a lower IQ has a significant negative effect on the intercept (see Table 5). The impact of these factors on the slope of the growth curve is less clear to interpret, because there is a strong negative correlation between intercept and slope ($r = -.73$), which means that a student who starts math development well has less margin to grow than a student who starts weaker. However, there is only one variable related to the effect size of grade retention, i.e. IQ, the higher the IQ, the more positive the effects of retention. The other factors mentioned do not moderate the effect.

Figure 8

Table 5

Artifact Explanations of Positive Retention Effects

Before continuing with other analysis techniques, we discuss two related alternative explanations of the observed positive effects of grade retention referring to a possible artifact. The first is the regression effect (cf. Shepard, 2002; Lorence, 2014). That the positive effects of retention cannot be explained by referring to the regression effect can easily be deduced from Table 1 and Figure 3. The predicted math score in the schoolyear after retention, based on the set of preceding math scores (e.g. in multiple regression analysis) if these pupils would not have been retained, is much lower than the observed score. Contrary to the general trend of an ever-decreasing rise in the curve, retained students deviate exactly from this general trend in the year following retention as they show a strong increase.

According to another artifact explanation, which is less easy to refute, the increase in the math scores after repeating a grade is mainly due to the fact that repeaters in the year

GRADE RETENTION IN PRIMARY SCHOOL

before they repeat a grade score lower than they are really capable of because they know in advance that they will have to repeat the grade. In other words: the expectation to repeat leads to underperformance of their math potential. Whether this artifact explanation makes sense can be verified by testing a logical prediction of this explanation. For example, if students who will repeat second grade, have a lower math score at the end of second grade due to the expectation to repeat, then these students are predicted to score higher on a previous measurement (at the end of first grade) compared to students who achieved the same math score in the second grade as the retained students, but who were in fact not retained.

Appendix B explains in detail how this prediction can be tested via matching. These analyses show that this prediction was not upheld. We can conclude from this that the low score of the retainers prior to the repeat year is not a consequence of the expectation to get retained, but vice versa, grade retention is at least partly a consequence of the low math score.

Marginal Structural Models (MSM)

In contrast to the latent growth model, which uses successive measurements to estimate the influence of grade retention at different points in time on an underlying growth model, marginal structural models are based on the idea of a sequentially randomized experiment, trying to simulate the exchangeability of an intervention group and a control group as closely as possible. Successive measurements of mathematic achievement prior to an intervention (retention) are considered here as confounding variables that need to be controlled. The motive to work with marginal structural models is based on the criticism on standard methods, i.e. regression analyses, that by controlling for a time-varying confounder they lead to both overcontrol bias and 'collider stratification bias'. However, both forms of bias are only problematic for the standard methods if one is interested in so-called 'joint effects' and in direct effects (see introduction). Total effects can be correctly estimated with standard methods (Keogh, Daniel, Vanderweele, & Vansteelandt., 2017). In marginal

GRADE RETENTION IN PRIMARY SCHOOL

structural models (Robins, Hernan, & Brumback, 2000; Robins & Hernan, 2009), by assigning weights to the participants, a pseudo-population is created in such a way that the treatment status is statistically independent of all preceding fixed and time-varying covariates at any point in time, exactly as in a randomized experiment. The weights are calculated using so-called 'Inverse Probability of Treatment Weighting' (IPTW) in which the inverse of the propensity score (the probability to be retained or promoted given several covariates) is taken. Individuals that are exceptional in the original dataset, or vice versa, are numerous, given their scores on a series of covariates, are equally common in the pseudo-population after weighting, thus removing any link between the covariates and the treatment status. All kinds of causal (or 'structural') contrasts, e.g. to examine direct, 'joint', or total effects, can then be defined in the pseudo-population by calculating the 'marginal' effects across the strata of the different covariates.

However, there are also disadvantages to this method (Zhou & Wodtke, 2020). A first problem is that determining the weights depends on formulating a model (e.g. a logistic regression equation) in which choices have to be made, and research has shown that such models are sensitive to misspecifications (Howe, Cole, Chmiel, & Muñoz, 2011; Kang & Schafer, 2007; Lefebvre, Delaney, & Platt, 2008). A second problem is that even with so-called 'stabilized weights'⁷, which are less sensitive to extreme weights (in the case of very small propensity scores), there may still be sample specific bias and a lack of precision in the estimates (Cole & Hernan, 2008; Wang, Petersen, Bangsberg, & van der Laan, 2006). Finally, marginal structural models can only be used if the 'positivity' assumption is plausible (see introduction). This means that for each combination of the values of the covariates both

⁷ With stabilized weights, the denominator is not equal to one, as with non-stabilized weights, but equal to the so-called 'baseline probability' to be retained or to be promoted, which is given the treatment history and the fixed covariates.

GRADE RETENTION IN PRIMARY SCHOOL

retained and promoted children must occur. For participants for whom this overlap does not exist, this method cannot make any statements.

In order to deal with these problems, we have replaced the IPTW procedure that uses propensity scores by the in Appendix B discussed and recently developed 'entropy balancing' method (Hainmueller, 2012). As a result, the first two problems, model dependence and sample specific bias due to extreme weights, are effectively addressed. Using this method, the two groups, repeaters and non-repeaters, can be directly matched on the original covariates at any point in time. This avoids cumbersome 'covariate balancing checking' procedures, which often do not lead to a satisfactory match on the covariates.

For each time point, the groups of repeaters and non-repeaters are matched for the first three central moments on all baseline covariates (month of birth, gender, SES, home language, ethnic origin, high risk pupil, IQ, language and numeracy skills in 3rd kindergarten class), on all language, reading, spelling and math tests prior to retention, the history of retention, all variables of psychosocial functioning taken in the year before retention, and all interaction terms between the baseline covariates. If the algorithm used (Hainmueller & Xu, 2013) did not succeed in determining the weights in such a way that a match was found for all included variables (between 50 and 100 variables depending on the time point), a match was attempted for the first two central moments (mean and variance), if this also failed, then only the mean was matched. In most cases, at least a match was achieved for the first two central moments. To meet the problem of missing values, a multiple imputation procedure with 25 imputations was first performed in SPSS v.26. This procedure considers the missing values as Missing at Random (MAR). Next, we performed a weighted regression analysis with as dependent variable the math score of the year in which the effect of retention is investigated, and as independent variable the group (repeaters or not). Table 6 gives an overview for the different grades of the effect sizes pooled across the 25 imputations using Rubin's rules (Rubin, 1987).

 Table 6

This overview of the total effects shows a different picture than the latent growth models. The effects of being held back in third kindergarten class and first grade are usually negative (often not statistically significant), while the effects of being retained from second grade onwards are positive (for half of the effects significant). The magnitude of these positive effects is about half as large compared to the estimates of the latent growth models (see Table 2). However, before drawing the substantive conclusion that retention has rather negative effects at the beginning of primary education and then moderately positive effects, two important methodological observations must be made.

A first problem is that in most cases these effects pertain to only part of the repeaters. Only for the effects of retention in the third kindergarten class could all repeaters be matched with non-repeaters. Table 7 shows that for more than one third of primary school pupils, this method does not allow a conclusion to be made about the effects of retention.

 Table 7

A second problem with this method is that it is likely that the matching for the pupils who were retained relatively early is much less accurate and can therefore easily lead to an underestimation of the real effects of retention. The more preceding data is available, the more accurate the match will be. This is again essentially the problem of 'unmeasured confounders', where repeaters have characteristics that are not or insufficiently covered by the measured covariates. If all characteristics of being retained are well covered by the measured covariates, the explained variance of being retained or not should be very high. Table 7 shows that the explained variances in treatment status for each grade, given all previous covariates,

GRADE RETENTION IN PRIMARY SCHOOL

are only between .30 and .47. This implies that it is very likely that there are still unmeasured confounders that largely determine the treatment status. If early repeaters are matched with non-repeaters based on a limited number of previous data, there is a high probability that the repeaters are really weaker than the non-repeaters and will later return to their real population average. One of the most important time-varying confounders are the preceding math scores. Table 7 shows that the more prior math scores are available, the better the math scores can be predicted. This means that math scores at the beginning of primary school can be predicted less reliably and that weak pupils will usually be overestimated. In contrast to latent growth models, marginal structural models only use the previous data and not the entire growth trajectory.

Sequential Conditional Mean Models

Another method that is based on the idea of a sequentially randomized experiment is the 'Sequential Conditional Mean Model' (SCMM). In a recent publication Keogh et al. (2018) describe this method as a way to estimate total effects in longitudinal studies with time-varying covariates and interventions, without time-varying confounders inducing 'overcontrol bias' or 'collider stratification bias'. In this technique based on regression analysis (see Figure 2), the expected value of the outcome variable (mathematical ability) at a given point in time, given preceding covariates and interventions, is estimated by including all preceding interventions and covariates as predictors in the regression equation, in addition to the effect of the specific intervention. For example, to examine the total effect of repeating fourth grade on the math scores in sixth grade, a regression equation is written with as dependent variable the math scores of sixth grade, and as independent variables retaining in fourth grade (or not) and all previous retention possibilities, all math scores up to and including fourth grade, all time invariant covariates, all psychosocial variables measured in fourth grade, all interaction terms between the time invariant covariates, and all relevant interaction terms of treatment

GRADE RETENTION IN PRIMARY SCHOOL

status and math scores (in this example: GRK3 x MK3, GR1 x MK3, GR1 x M1, GR2 x MK3, GR2 x M1, GR2 x M2, ..., GR4 x M4).

An advantage of this technique is that, just as with latent growth models, all participants are involved in the analysis. The disadvantage, however, is that in the absence of overlap between the groups for the distribution of the covariates, participants from one group are compared with statistically extrapolated but non-existing participants from the other group. Whether these extrapolations cause bias or not depends on the correctness of the regression model (model dependence). The effects are estimated by using robust standard errors in MPlus. The target population is all students who were retained at least once and all students who were advised to retain at least once (see also latent growth models). Table 8 shows the total effects, again considering the missing values as MAR.

Table 8

In general, the effects of grade retention estimated by SCMM lie between those of the latent growth models and those of the marginal structural models. They are generally positive, smaller than in the latent growth models, but larger than in the marginal structural models.

Indirect Effect of Grade Retention

The indirect effect of grade retention, as defined in the introduction discussing SUTVA, could not be verified in our dataset because there was no school in which none of the pupils were retained. To repeat a grade was therefore always one of the possibilities. The percentage of retained students per school varied greatly from 2% to 72% of the students who were retained at least once over all grades. This allowed us to investigate the effect of the frequency of retention in schools on the strength of the direct effect of retention (the individual effect of being retained or not). To this end, the latent growth model from Figure 8

GRADE RETENTION IN PRIMARY SCHOOL

with a latent random variable (STVC) was adapted by including the percentage of repeaters per school as a time invariant predictor in the model. We included the linear as well as the quadratic effect of the percentage of retained children in the model. The percentage of repeaters appeared to have a significant influence on the magnitude of the effects of being retained. Both the linear effect was substantial and significant ($\beta = .056$, $SE = .018$, $p = .002$), and the quadratic effect ($\beta = -0.36$, $SE = .015$, $p = .013$). This means that the positive effect of retention increases with the ratio of repeaters in the school, but that this increase strongly flattens with high percentages of repeaters. This implies that the classic assumption of non-influence of treatment status between students (SUTVA) is not valid. The observed effects indicate mutual influence between pupils but are mute about the indirect influence of grade retention as a system.

Conclusions and Discussion

Given the contradictory results in existing research on the effects of grade retention on school performance and psychosocial functioning, the aim of this article was to use new analyses and methods that excluded selection-bias as much as possible to find out, in one of the largest existing datasets, in which direction the effects point and how large they really are. The main conclusion is that the effects on mathematic achievement are positive, also in the longer term. It appears that the effects on well-being are also positive. This means that the reason that is usually given to explain any negative effects of grade repetition on school performance is not substantiated. After all, it is often assumed that the awareness of being retained has such a negative effect on the self-concept and the motivation of the student that school performance is more likely to suffer (Alexander, Entwistle, & Dauber, 2003; Kretschmann, Vock, Lüdtke, Jansen, & Gronostaj, 2019). However, on the basis of research into social comparison theories (Huguet et al., 2009) and the extensively studied and robust 'Big-Fish-Little-Pond-Effect' (BFLPE) (Marsh, 1987; Marsh et al., 2014), one may expect that

GRADE RETENTION IN PRIMARY SCHOOL

a pupil who, by repeating a grade, performs better in comparison with the new class group will experience a more favorable self-concept and associated well-being (Marsh, 2016). The BFLPE is based on a contrast effect in which the comparison with the direct reference group (here the class) is more salient and has more impact than the comparison with the previous class group (the awareness of being retained). The finding that the positive effect of retention increases as there are proportionally more repeaters in the school can also be understood from comparative processes. If you are not alone as a repeater in the classroom, this can also be experienced as support for your own position. If the putative negative psychosocial effect of retention is at least compensated by the BFLPE, then the positive effect of retention can mainly be attributed to repeating the subject matter in the year of retention. An additional finding in an analysis in which the strength of the effect was modelled as a random variable is that the effect size is not dependent on the month of birth, gender, socio-economic status, home language, ethnic group, risk status, but only on the intelligence of the pupil. The higher the intelligence of the student who is retained, the stronger the effect of retention. Intelligence is, as always, an excellent predictor of school learning (cf. Deary, Strand, Smith, & Fernandes, 2007).

However, the reported positive effects in this study contradict the conclusion from previous studies and publications based on the same dataset that the overall effects of retention are negative. We showed that the reason why the study by Goos et al. (2013) came to different conclusions is due to a strong bias between the groups of repeaters and non-repeaters. The other studies carried out on the same dataset (Vandecandelaere, Schmitt, Vanlaar, De Fraine, & Van Damme, 2015; Vandecandelaere et al., 2016a; Vandecandelaere et al., 2016b), just like the study by Goos et al., only examined the effects of grade retention in the beginning of primary school (third kindergarten until second grade). Given the methods used in these studies, propensity score matching and marginal structural models, the

GRADE RETENTION IN PRIMARY SCHOOL

conclusions of these authors are not very different from the results we presented using marginal structural models. With this method, we also observed rather small negative effects (not always significant) of retention in third kindergarten class and first grade.

What is new in our study, however, is that we examined the effects of grade retention for all grades, and that we used multiple methods to counteract model dependence, while also investigating to what extent the specific assumptions associated with each model were met. In order to reach non-biased conclusions and to avoid so-called 'statisticism', which is the naïve idea that statistics are a complete and sufficient basis for scientific methodology (Duncan, 1984), it is crucial that the conditions, assumptions and limitations of each statistical technique are thoroughly studied in the context of the specific research question.

In the discussion of the different analysis techniques, we paid extensive attention to some crucial assumptions, i.e. 'no unmeasured confounders' and the positivity assumption, and to the choice of the type of comparison group. However, these three methodological aspects are strongly intertwined. The overarching question here is to what extent the two groups, repeaters and non-repeaters, are comparable or can be made comparable. The choice in many existing studies for an age comparison is in fact based on the fundamental assumption that the two groups are comparable. However, this assumption must be vindicated, otherwise no valid conclusions can be drawn from the analysis techniques. When discussing the results, we have narrowed the question of comparability of the groups to the question of the comparability of the measurement, the measurement invariance. Our analyses showed that at least for the mathematics tests used, there is no measurement invariance between two consecutive grades. This means that not the same mathematical ability is measured between one and the other year. In our analyses, therefore, only grade class comparisons were used, though theoretically more ambiguous, i.e. as the result of repeating the academic year and getting a year older.

GRADE RETENTION IN PRIMARY SCHOOL

Proponents of an age comparison, however, could still argue that this incomparability of measurement does not invalidate the simple fact that the promoted children are already further advanced in their mathematical development than the repeaters. However, this ignores the central problem, i.e. are the two groups really comparable, or in other words, if the repeaters had been promoted, could they have taken lessons in a normal way in the next year? This question cannot be bypassed.

For techniques such as MSMs and SCMMs, this question translates into specific assumptions that need to be verified in order to allow valid conclusions to be drawn. Both in matching and regression techniques, an important assumption is that there are no 'unmeasured confounders'. The finding that the probability to be retained could not be predicted very well based on the very numerous covariates is a strong indication that not all confounders were measured. Because prior math scores should also be understood as confounders, the finding that the math scores themselves are less well predicted by preceding math scores as the measurement falls at an earlier point in time is also relevant. This implies two things: 1) one of the central assumptions of MSMs and SCMMs has not been met and hence the results based on these models are negatively biased when estimating the effects of retention, 2) the bias is largest at early points in time.

A second central assumption that MSMs and SCMMs make is the 'positivity' assumption, which is the requirement that the groups of repeaters and non-repeaters overlap for all measured covariates on which the groups are to be matched. Complete overlap was only possible for the repeaters of third kindergarten class, while for the other grades only 62 to 68% of the repeaters could be matched. The reason why several repeaters could not be matched is obviously because these repeaters are the weakest pupils. The 'counterfactual' question as to what would have happened to these children if they had been promoted, cannot be answered with these methods, but is crucial. However, with only partial overlap, the

GRADE RETENTION IN PRIMARY SCHOOL

'Average Treatment Effect on the Treated' (ATT) will be underestimated, in the first case (MSM) because an important subgroup is removed from the analysis, in the second case (SCMM) because data from non-existing subjects are extrapolated statistically on the basis of existing subjects.

Unlike MSMs and SCMMs, which are based on the sequentially randomized experiment, preceding math scores in a latent growth model (LGM) and in a cross-lagged panel model (CLPM) are not confounding variables that need to be controlled, but these successive scores are modelled directly to determine the extent to which time-varying interventions affect them. This has several advantages. Because in LGMs and CLPMs each subject forms its own control over time, both the measured and unmeasured stable characteristics of each individual also remain constant over time. As a result, the problem of 'unmeasured confounders' is largely solved. The problem of overlap also disappears because all the retained participants are included in the analysis. The biases that were associated with both assumptions in the previous models also disappear as a result. However, like any model, LGMs and CLPMs make some assumptions.

The most important assumption in a LGM concerns the interpretation of the lower curve (Figure 5) as a counterfactual curve, i.e. the curve of the population of repeaters if they would never have been retained. However, as explained, the data points on this curve up to and including the fifth grade are not counterfactual but based on the scores of students who had not repeated until that time. Because pupils who were retained earlier are usually weaker than those who were retained later, this implies that the bottom curve overestimates performance, especially of pupils who were retained early. The crucial question "how would pupils who repeated the first, second, third, etc. grade have performed if they were not retained" does not appear to be estimated in any method without bias. The LGM does, however, make it possible to determine a lower limit of the real effects of grade retention. In

GRADE RETENTION IN PRIMARY SCHOOL

short, all the presented methods underestimate the effects of retention, but the latent growth models the least. The fundamental reason why the 'true' effects of retention are so difficult to estimate is, of course, related with the simple fact that the groups of repeaters and non-repeaters are so different that 'exchangeability' is a fiction in this context. The counterfactual question as to what would have happened to the weakest pupils or early repeaters if they had not been retained can then only be answered by imagining as realistically as possible what would have happened to them. It is reasonable to assume that it would be precisely these pupils who would have get stuck first and get retained in a later year or would have been referred to special education. Some will see this as an argument for questioning more fundamentally the education system with grade classes. But even then, the educational question will have to be asked as to how to deal with learning delays that will be inevitable also in a system without grade classes. The conclusion from the meta-analysis by Allen et al. (2009) that stronger research designs show fewer negative and even positive effects is confirmed and even reinforced by this study: if the bias is eliminated even more effectively, there are clearly large positive effects.

References

- Alexander, K. L., Entwisle, D. R., & Dauber, S. L. (2003). *On the success of failure: A reassessment of the effects of retention in the primary school grades*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511500091>
- Allen, C.S., Chen, Q., Willson, V.L., & Hughes, J.N. (2009). Quality of research design moderates effects of grade retention on achievement: A meta-analytic, multilevel analysis. *Educational Evaluation and Policy Analysis*, 31(4), 480-499. <https://doi.org/10.3102/0162373709352239>
- Allison, P.D., Williams, R., & Moral-Benito, E. (2017). Maximum likelihood for cross-lagged panel models with fixed effects. *Socius: Sociological Research for a Dynamic World*, 3, 1-17. <https://doi.org/10.1177/2378023117710578>
- Bollen, K. A., & Curran, P. J. (2006). *Latent curve models: A structural equation perspective*. John Wiley & Sons. <https://doi.org/10.1002/0471746096>
- Cham, H., Hughes, J. N., West, S. G., & Im, M. H. (2015). Effect of retention in elementary grades on grade 9 motivation for educational attainment. *Journal of School Psychology*, 53(1), 7-24. <https://doi.org/10.1016/j.jsp.2014.10.001>
- Chen, F. F. (2007). Sensitivity of goodness of fit indexes to lack of measurement invariance. *Structural Equation Modeling: A Multidisciplinary Journal*, 14(3), 464-504. <https://doi.org/10.1080/10705510701301834>
- Cole, S. R., & Hernán, M. A. (2008). Constructing inverse probability weights for marginal structural models. *American Journal of Epidemiology*, 168(6), 656-664. <https://doi.org/10.1093/aje/kwn164>
- Daniel, R.M., Cousens, S.N., De Stavola, B.L., Kenward, M.G., & Sterne, J.A.C. (2013). Methods for dealing with time-dependent confounding. *Statistics in Medicine*, 32(9), 1584-1618. <https://doi.org/10.1002/sim.5686>
- Deary, I. J., Strand, S., Smith, P., & Fernandes, C. (2007). Intelligence and educational achievement. *Intelligence*, 35(1), 13-21. <https://doi.org/10.1016/j.intell.2006.02.001>
- Diggle, P., & Kenward, M. G. (1994). Informative drop-out in longitudinal data analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 43(1), 49-73. <https://doi.org/10.2307/2986113>
- Duncan, O. D. (1984). *Notes on social measurement: Historical and critical*. Russell Sage Foundation.
- Elwert, F. (2013). Graphical causal models. In *Handbook of causal analysis for social research* (pp. 245-273). Springer, Dordrecht. https://doi.org/10.1007/978-94-007-6094-3_13

GRADE RETENTION IN PRIMARY SCHOOL

Elwert, F., & Winship, C. (2014). Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40, 31-53. <https://doi.org/10.1146/annurev-soc-071913-043455>

Enders, C. K. (2010). *Applied missing data analysis*. Guilford press.

Enders, C. K. (2011). Missing not at random models for latent growth curve analyses. *Psychological Methods*, 16(1), 1-16. <https://doi.org/10.1037/a0022640>

Goos, M., Van Damme, J., Onghena, P., & Petry, K. (2011). *First-grade retention: Effects on children's academic and psychosocial growth throughout primary education*. SSL-rapport SSL/OD1/2010.33. Steunpunt SSL.

Goos, M., Van Damme, J., Onghena, P., Petry, K., & de Bilde, J. (2013). First-grade retention in the Flemish educational context: Effects on children's academic growth, psychosocial growth, and school career throughout primary education. *Journal of School Psychology*, 51(3), 323-347. <https://doi.org/10.1016/j.jsp.2013.03.002>

Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1), 25-46. <https://doi.org/10.1093/pan/mpr025>

Hainmueller, J., & Xu, Y. (2013). Ebalance: A Stata package for entropy balancing. *Journal of Statistical Software*, 54(7). <https://doi.org/10.18637/jss.v054.i07>

Holmes, C. T. (1989). Grade level retention effects: A meta-analysis of research studies. In L.A. Shepard & M.L. Smith (Eds.) *Flunking Grades: Research and Policies on Retention* (pp. 16-33). London: The Falmer Press.

Hong, G., & Raudenbush, S. W. (2005). Effects of kindergarten retention policy on children's cognitive growth in reading and mathematics. *Educational Evaluation and Policy Analysis*, 27(3), 205-224. <https://doi.org/10.3102/01623737027003205>

Howe, C.J., Cole, S.R., Chmiel, J.S., & Muñoz, A. (2011). Limitation of inverse probability-of-censoring weights in estimating survival in the presence of strong selection bias. *American Journal of Epidemiology*, 173(5), 569-577. <https://doi.org/10.1093/aje/kwq385>

Hu, L. T., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, 6(1), 1-55. <https://doi.org/10.1080/10705519909540118>

Hughes, J. N., Cao, Q., West, S. G., Smith, P. A., & Cerda, C. (2017). Effect of retention in elementary grades on dropping out of school early. *Journal of School Psychology*, 65, 11-27. <https://doi.org/10.1016/j.jsp.2017.06.003>

GRADE RETENTION IN PRIMARY SCHOOL

Hughes, J. N., West, S. G., Kim, H., & Bauer, S. S. (2018). Effect of early grade retention on school completion: A prospective study. *Journal of Educational Psychology*, 110(7), 974. <https://doi.org/10.1037/edu0000243>

Huguet, P., Dumas, F., Marsh, H., Régner, I., Wheeler, L., Suls, J., Seaton, M., & Nezlek, J. (2009). Clarifying the role of social comparison in the big-fish–little-pond effect (BFLPE): An integrative study. *Journal of Personality and Social Psychology*, 97(1), 156. <https://doi.org/10.1037/a0015558>

Iacus, S. M., King, G., & Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1), 1-24. <https://doi.org/10.1093/pan/mpr013>

Jimerson, S.R. (2001). Meta-analysis of grade retention research: Implications for practice in the 21st century. *School Psychology Review*, 30(3), 420-437.

Jimerson, S., Woehr, S., Kaufman, A., & Anderson, G. (2004). Grade retention and promotion: Important information for educators. In A.S. Canter, S.A. Carroll, L. Paige, & I. Romero (Eds.), *Helping children at home and school: Handouts from your school psychologist* (2nd ed., Section 3, pp. 61–64). Bethesda, MD: National Association of School Psychologists

Kang, J. D., & Schafer, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, 22(4), 523-539. <https://doi.org/10.1214/07-STS227>

Keogh, R.H., Daniel, R.M., Vanderweele, T.J., & Vansteelandt, S. (2018). Analysis of longitudinal studies with repeated outcome measures: Adjusting for time-dependent confounding using conventional methods. *American Journal of Epidemiology*, 187(5), 1085-1092. <https://doi.org/10.1093/aje/kwx311>

King, G., & Nielsen, R. (2019). Why propensity scores should not be used for matching. *Political Analysis*, 27(4), 435-454. <https://doi.org/10.1017/pan.2019.11>

Kretschmann, J., Vock, M., Lüdtke, O., Jansen, M., & Gronostaj, A. (2019). Effects of grade retention on students' motivation: A longitudinal study over 3 years of secondary school. *Journal of Educational Psychology*, 111(8), 1432-1446. <https://doi.org/10.1037/edu0000353>

LeBel, E. P., McCarthy, R. J., Earp, B. D., Elson, M., & Vanpaemel, W. (2018). A unified framework to quantify the credibility of scientific findings. *Advances in Methods and Practices in Psychological Science*, 1(3), 389-402. <https://doi.org/10.1177/2515245918787489>

GRADE RETENTION IN PRIMARY SCHOOL

- Lefebvre, G., Delaney, J.A.C., & Platt, R.W. (2008). Impact of mis-specification of the treatment model on estimates from a marginal structural model. *Statistics in Medicine*, 27(18), 3629–3642. <https://doi.org/10.1002/sim.3200>
- Little, R. J., & Rubin, D. B. (2002). Statistical analysis with missing data. John Wiley & Sons. *New York*. <https://doi.org/10.1002/9781119013563>
- Liu, Y., Millsap, R. E., West, S. G., Tein, J. Y., Tanaka, R., & Grimm, K. J. (2017). Testing measurement invariance in longitudinal data with ordered-categorical measures. *Psychological Methods*, 22(3), 486-506. <https://doi.org/10.1037/met0000075>
- Lorence, J. (2006). Retention and academic achievement research revisited from a United States perspective. *International Education Journal*, 7(5), 731-777.
- Lorence, J. (2014). Third-grade retention and reading achievement in Texas: A nine year panel study. *Social Science Research*, 48, 1-19. <https://doi.org/10.1016/j.ssresearch.2014.05.001>
- Marsh, H. W. (1987). The big-fish-little-pond effect on academic self-concept. *Journal of Educational Psychology*, 79(3), 280-295. <https://doi.org/10.1037/0022-0663.79.3.280>
- Marsh, H. W. (2016). Cross-cultural generalizability of year in school effects: Negative effects of acceleration and positive effects of retention on academic self-concept. *Journal of Educational Psychology*, 108(2), 256-273. <https://doi.org/10.1037/edu0000059>
- Marsh, H. W., Abduljabbar, A. S., Parker, P. D., Morin, A. J., Abdelfattah, F., & Nagengast, B. (2014). The Big-Fish-Little-Pond effect in mathematics: A cross-cultural comparison of US and Saudi Arabian TIMSS responses. *Journal of Cross-Cultural Psychology*, 45(5), 777-804. <https://doi.org/10.1177/0022022113519858>
- Marsh, H. W., Lüdtke, O., Muthén, B., Asparouhov, T., Morin, A. J., Trautwein, U., & Nagengast, B. (2010). A new look at the big five factor structure through exploratory structural equation modeling. *Psychological Assessment*, 22(3), 471. <https://doi.org/10.1037/a0019227>
- Marsh, H. W., Pekrun, R., Parker, P. D., Murayama, K., Guo, J., Dicke, T., & Lichtenfeld, S. (2017). Long-term positive effects of repeating a year in school: Six-year longitudinal study of self-beliefs, anxiety, social relations, school grades, and test scores. *Journal of Educational Psychology*, 109(3), 425. <https://doi.org/10.1037/edu0000144>
- Meredith, W. (1993). Measurement invariance, factor analysis and factorial invariance. *Psychometrika*, 58(4), 525-543. <https://doi.org/10.1007/BF02294825>
- Meredith, W., & Tisak, J. (1990). Latent curve analysis. *Psychometrika*, 55(1), 107-122. <https://doi.org/10.1007/BF02294746>

GRADE RETENTION IN PRIMARY SCHOOL

- Millsap, R.E. (2011). Statistical approaches to measurement invariance. *Routledge, Arizona State University*. <https://doi.org/10.4324/9780203821961>
- Millsap, R. E., & Yun-Tein, J. (2004). Assessing factorial invariance in ordered-categorical measures. *Multivariate Behavioral Research*, 39(3), 479-515. https://doi.org/10.1207/S15327906MBR3903_4
- Muthén, B., Asparouhov, T., Hunter, A. M., & Leuchter, A. F. (2011). Growth modeling with nonignorable dropout: alternative analyses of the STAR* D antidepressant trial. *Psychological Methods*, 16(1), 17-33. <https://doi.org/10.1037/a0022634>
- Muthén, L., & Muthén, B. (2019). *Mplus User's Guide* (1998–2019). Muthén & Muthén: Los Angeles, CA.
- Morgan, S. L., & Winship, C. (2015). *Counterfactuals and causal inference. Methods and principles for social research*. (2nd Ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9781107587991>
- Robins, J. M., Hernán, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* 11(5), 550–560. <https://doi.org/10.1097/00001648-200009000-00011>
- Robins, J.M., & Hernán, M.A . (2009). Estimation of the causal effects of time-varying exposures. In G. Fitzmaurice, M. Davidian, G. Verbeke, & G. Molenberghs (Eds.), *Longitudinal data analysis*. (pp. 553-599). New York: Chapman and Hall/CRC Press. <https://doi.org/10.1201/9781420011579.ch23>
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688. <https://doi.org/10.1037/h0037350>
- Rubin, D. B. (1986). “Which ifs have causal answers (Comment on ‘Statistics and causal inference’ by Paul W. Holland).” *Journal of the American Statistical Association*, 81, 961-962. <https://doi.org/10.2307/2289065>
- Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. Hoboken, NJ: Wiley. <https://doi.org/10.1002/9780470316696>
- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469), 322-331. <https://doi.org/10.1198/016214504000001880>
- Schafer, J.L. & Graham, J.W. (2002). Missing data: Our view of the state of the art. *Psychological Methods*, 7(2), 147-177. <https://doi.org/10.1037/1082-989X.7.2.147>
- Shepard, L.A., 2002. Comment by Lorrie A. Shepard. In: Ravitch, D. (Ed.), *Brookings Papers on Education Policy 2002*, pp. 56–63. <https://doi.org/10.1353/pep.2002.0021>

GRADE RETENTION IN PRIMARY SCHOOL

- Steiner, P. M., Park, S., & Kim, Y. (2016). Identifying causal estimands for time-varying treatments measured with time-varying (age or grade-based) instruments. *Multivariate Behavioral Research*, 51(6), 865-8780. <https://doi.org/10.1080/00273171.2016.1205470>
- Tjur, T. (2009). Coefficients of determination in logistic regression models—A new proposal: The coefficient of discrimination. *The American Statistician*, 63(4), 366-372. <https://doi.org/10.1198/tast.2009.08210>
- Vandecandelaere, M., Schmitt, E., Vanlaar, G., De Fraine, B., & Van Damme, J. (2015). Effects of kindergarten retention for at-risk children's mathematics development. *Research Papers in Education*, 30(3), 305-326. <https://doi.org/10.1080/02671522.2014.919523>
- Vandecandelaere, M., Vansteelandt, S., De Fraine, B., & Van Damme, J. (2016a). The effects of early grade retention: Effect modification by prior achievement and age. *Journal of School Psychology*, 54, 77-93. <https://doi.org/10.1016/j.jsp.2015.10.004>
- Vandecandelaere, M., Vansteelandt, S., De Fraine, B., & Van Damme, J. (2016b). Time-varying treatments in observational studies: Marginal structural models of the effects of early grade retention on math achievement. *Multivariate Behavioral Research*, 51(6), 843-864. <https://doi.org/10.1080/00273171.2016.1155146>
- Verhaeghe, J., Maes, F., Gombeir, D., & Peeters, E. (2002). Longitudinaal onderzoek in het basisonderwijs: steekproeftrekking (Longitudinal research in primary education: sampling). *LOA-rapport*, 5.
- Verhaeghe, J. P. (2010). Technische fiche voor IRT-analyse wiskunde basisonderwijs [Technical report IRT analysis math primary education]. *Leuven, Belgium: Steunpunt Studietoelagen en Schoolloopbanen*.
- Wang, Y., Petersen, M. L., Bangsberg, D., & van der Laan, M. J. (2006). Diagnosing bias in the inverse probability of treatment weighted estimator resulting from violation of experimental treatment assignment. *U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 211*.
- Zhou, X., & Wodtke, G. T. (2020). Residual balancing: A method of constructing weights for marginal structural models. *Political Analysis*, published online. <https://doi.org/10.1017/pan.2020.2>

APPENDIX A: TESTS ON MEASUREMENT INVARIANCE

If the items are continuous variables, then at least scalar or strong invariance is required in order to unambiguously attribute the differences in the mean item scores between measurement moments to differences in the mean factor scores (Chen, 2007; Millsap, 2011). However, if the items are categorical variables, then strict or unique factor invariance is required (Liu et al., 2017). If the categorical variables are dichotomous, as is the case here, additional restrictions must be imposed in order to identify the model, i.e. the thresholds between response categories must be the same for the different measurement moments (Millsap & Tein, 2004). This implies that in the case of dichotomous items, the step of scalar or strong invariance cannot be tested and is thus obsolete.

The number of overlapping math items between successive measurements varied greatly: 9 overlapping items between the beginning and the end of the third kindergarten class, 18 overlapping items between the end of the third kindergarten class and the beginning of first grade, no overlapping items between the beginning and the end of first grade, 13 overlapping items between the end of first grade and the end of second grade, 2 overlapping items between the end of second grade and the end of third grade, 5 overlapping items between the end of third grade and the end of fourth grade, 2 overlapping items between the end of fourth grade and the end of fifth grade, and 1 overlapping item between the end of fifth grade and the end of sixth grade. Due to the sometimes very small number of overlapping items and in one case even the absence of overlapping items, measurement invariance for all tests taken is difficult or impossible to determine. We therefore only present the results for transitions between schoolyears (not within a year), for which there were at least five overlapping items: for the transition between the end of the third kindergarten class and the beginning of first grade (18 items), the transition between the end of first grade and the end of second grade (13 items), and the transition between the end of third grade and the end of fourth grade (5 items).

GRADE RETENTION IN PRIMARY SCHOOL

In order to deal with the problem of scarcity or lack of overlapping items, the authors of the original study organized additional test runs with a separate and limited group of students. For example, in order to compensate for the missing link between the beginning and the end of first grade, an extra math test was taken in the middle of first grade with about 80 pupils from special education. Because the characteristics of measurement invariance are sample specific, this attempt to create an indirect link does not offer a solution for testing measurement invariance. The reason for taking an extra test in the middle of first grade was precisely the concern that the items that are suitable at the beginning of first grade would be far too easy at the end of first grade. In such circumstances, measurement invariance may be an unattainable goal.

Table A1

In Table A1 we present the results for the tests on measurement invariance of math scores for the mentioned transitions. For the transition between the end of the third kindergarten class and the beginning of first grade, we see that a 1-factor model has a good fit for each of the time points, and also for a model in which the two measurement moments are taken together with one factor per measurement moment (configural invariance model). The imposition of more restrictions in the following models does not lead to a weakening of the fit. Over these two measurement moments, we can therefore consider the math scores as measurement invariant and thus as indicators of one and the same underlying math factor. Given the short time lapse and the absence of instruction between the end of the third kindergarten class and the beginning of first grade, this result is in line with expectations. For the transition between the end of first and the end of second grade, we see that a 1-factor model has a bad fit with the data for both times separately and together. Based on an

GRADE RETENTION IN PRIMARY SCHOOL

exploratory factor analysis, a four-factor model was tested in which each item loads only in one factor. The four factors correspond to a large extent with the sub-domains of mathematics mentioned in the measures section. It is striking that imposing the restriction of equal factor loadings in the metric invariance model immediately leads to a strong deterioration of the fit. The imposition of equal unique factor variances causes the fit to drop even further. The math scores are therefore not measurement invariant at all between these two measurement moments. For the transition between the end of third and the end of fourth grade, we see that a 1-factor model for the five overlapping items does provide an excellent fit for each time point separately and a reasonably good fit if the two measurement moments are taken together in the configural model. Imposing restrictions that are required for measuring invariance also leads to a substantial deterioration of the fit.

APPENDIX B: TEST ON ARTIFACT EXPLANATION

In order to match the retained students as precisely as possible with students from the same class who were promoted, we use a recently developed matching technique, namely entropy balancing (Hainmueller, 2012). This technique makes it possible to match two groups at the same time on different variables for the first three central moments of the respective distributions (mean, variance and skewness). The corresponding program (Hainmueller & Xu, 2013) automatically calculates the weights for all participants so that the match can be realized. Because this program has no built-in procedure to deal with the problem of missing values, we first performed a multiple imputation procedure with 10 imputations in SPSS v.26.⁸ This procedure considers the missing values as Missing at Random (MAR).

We matched each group of repeaters (per grade) with the group of non-repeaters on the math score administered the first time they are in the grade in which they will be retained, as well as on the following covariates: month of birth, gender, SES, home language, ethnic origin, high risk student, IQ, and the language proficiency tests taken up to that point. As an example, for second grade, for the first imputed dataset we give in Table B1 and B2 the mean, the variance, and the skew for all covariates, both before and after matching. The matching after application of the weights is obviously almost perfect. Next, we performed a weighted regression analysis with as dependent variable the math score at the end of first grade and as independent variable the group (repeaters or not). The regression coefficient for group was negative ($-.48$, $SE = .35$, $p = .173$). Contrary to the expectation of the artifact explanation, at the end of first grade the repeaters did not score higher, but lower (albeit not statistically

⁸ The dataset on which the imputations were computed via 'Fully Conditional Specification' (FCS), also known as 'Multivariate Imputation by Chained Equations' (MICE), was first purged for participants with less than two math scores, and for participants for whom the information on retention was missing up to and including the fifth year of study. All interaction terms of the following fixed covariates were added to the dataset: month of birth, ethnicity, gender, SES, home language, high risk student, language proficiency K3 and mathematics proficiency K3.

GRADE RETENTION IN PRIMARY SCHOOL

significant) than the at the end of second grade matched non-repeaters. Table B3 gives an overview of the effect sizes pooled across the ten imputations based on Rubin's well-known rules (Rubin, 1987) for the various grades. Table B3 shows that no group of repeaters scores higher on the math score (M_{X-1}) prior to the schoolyear in which they were retained than the non-repeaters who achieved the same math score at time T_X . In the first, second, and third grade, the effect is even negative (almost significant), which implies that even after matching on one math score, the repeaters are a weaker group.

Table B1

Table B2

Table B3

GRADE RETENTION IN PRIMARY SCHOOL

Table 1***Mean Math Scores for Grade Retention (GR) Groups (With or Without Censoring or Later Intake)***

	All pupils	GR 3K	GR G1	GR G2	GR G3	GR G4	GR G5
K3	51.00 (5862)	43.09 (589)	41.53 (518)	42.98 (253)	43.49 (154)	44.04 (96)	45.72 (78)
	53.16 (3823)	46.06 (194)	42.33 (157)	44.73 (121)	42.87 (64)	46.54 (42)	46.66 (44)
G1	67.78 (4943)	61.38 (398)	57.99 (508)	59.87 (205)	60.08 (135)	61.78 (70)	63.08 (73)
	69.47 (3678)	64.14 (171)	59.82 (159)	60.94 (110)	60.47 (61)	63.15 (35)	63.97 (42)
G2	79.86 (5935)	74.19 (320)	74.51 (708)	70.89 (268)	71.52 (166)	72.09 (88)	74.47 (92)
	81.73 (3543)	77.01 (145)	75.94 (104)	71.00 (101)	71.43 (62)	74.17 (30)	76.36 (38)
G3	88.28 (5668)	83.77 (269)	82.16 (602)	83.08 (167)	77.60 (159)	78.34 (88)	81.22 (93)
	90.54 (3367)	85.64 (150)	83.68 (110)	83.65 (80)	77.30 (61)	79.35 (30)	81.04 (36)
G4	95.36 (5650)	89.72 (287)	89.03 (524)	89.04 (167)	89.46 (86)	84.15 (88)	88.00 (92)
	97.77 (3269)	92.75 (143)	91.64 (101)	90.12 (82)	90.31 (44)	84.79 (28)	89.11 (33)
G5	101.10 (5525)	97.11 (208)	93.41 (455)	94.23 (163)	93.11 (92)	93.32 (60)	90.58 (113)
	103.40 (3255)	98.44 (146)	94.96 (110)	96.14 (84)	95.16 (49)	95.44 (24)	94.13 (35)
G6	105.91 (5682)	102.88 (216)	99.54 (396)	99.66 (160)	99.64 (77)	100.70 (58)	101.81 (73)
	106.95 (3920)	102.83 (204)	99.95 (166)	99.78 (129)	99.52 (67)	100.26 (43)	103.11 (45)

Note: all average math scores are the math scores for the first year pupils followed that year, not the year of retention. The sample sizes are given in brackets. For each schoolyear, the first row presents the mean math scores for the sample with drop-out or late intake, and the second row for the sample without drop-out or late intake.

GRADE RETENTION IN PRIMARY SCHOOL

Table 2

Direct Effects of Grade Retention on Math Scores (M) for Group of Repeaters (R) and Target Population (TP)

		M-G1b	M-G1e	M-G2e	M-G3e	M-G4e	M-G5e	M-G6e
GR-K3	R	0.84* (.41)	1.87* (.81)	3.51** (1.15)	4.89** (1.40)	4.59** (1.64)	7.16** (1.77)	9.96** (1.84)
	TP	0.69 (.41)	1.82* (.77)	3.20** (1.10)	4.32** (1.35)	4.01** (1.56)	5.55** (1.72)	6.71** (1.82)
GR-G1	R			5.53** (.56)	5.16** (.70)	4.37** (.82)	4.52** (.91)	7.30** (1.03)
	TP			5.68** (.51)	4.79** (.63)	3.96** (.72)	3.05** (.80)	3.96** (.87)
GR-G2	R				4.83** (.61)	3.76** (.69)	4.95** (.76)	8.25** (.89)
	TP				4.83** (.59)	3.67** (.63)	3.94** (.68)	5.28** (.74)
GR-G3	R					4.13** (.70)	4.70** (.71)	8.43** (.90)
	TP					4.19** (.67)	4.03** (.67)	5.73** (.77)
GR-G4	R						5.24** (.74)	8.80** (.88)
	TP						4.72** (.72)	6.48** (.80)
GR-G5	R							7.75** (.82)
	TP							5.37** (.67)

Note: Standard errors between brackets.

*p < .05, **p < .01

GRADE RETENTION IN PRIMARY SCHOOL

Table 3

Direct Effects of Grade Retention on Psycho-social Well-being (WB) for Group of Repeaters in a Cross-Lagged Panel Model

	WB-G1	WB-G2	WB-G3	WB-G4	WB-G5	WB-G6
GR-K3	.99* (.48)	1.05* (.49)	1.07* (.49)	1.03* (.48)	.92 (.48)	1.21* (.50)
GR-G1		.33** (.07)	.40** (.08)	.48** (.09)	.27** (.09)	.69** (.11)
GR-G2			.20* (.09)	.29** (.10)	-.00 (.10)	.25* (.12)
GR-G3				.26* (.12)	-.02 (.12)	.30* (.14)
GR-G4					.11 (.90)	-.04 (.17)
GR-G5						.39* (.16)

Note: Standard errors between brackets.

*p < .05, **p < .01

GRADE RETENTION IN PRIMARY SCHOOL

Table 4

Direct Effects of Retention on Math Scores (M) for Group of Repeaters in Diggle-Kenward Model

	M-G1b	M-G1e	M-G2e	M-G3e	M-G4e	M-G5e	M-G6e
GR-K3	1.16** (.45)	1.37 (.89)	4.14** (1.27)	5.53** (1.55)	5.59** (1.80)	9.92** (1.94)	14.19** (1.92)
GR-G1			6.13** (.56)	5.96** (.73)	5.99** (.84)	7.66** (1.01)	13.03** (1.06)
GR-G2				5.25** (.65)	4.46** (.70)	6.99** (.85)	12.87** (.95)
GR-G3					5.75** (.71)	7.05** (.84)	13.83** (.97)
GR-G4						7.84** (.90)	13.20** (.95)
GR-G5							12.47** (.95)

Note: Standard errors between brackets.

*p < .05, **p < .01

GRADE RETENTION IN PRIMARY SCHOOL

Table 5***Effects of Time Invariant Predictors on Intercept, Slope and STVC***

	Intercept			Slope			STVC		
	Estimate	SE	p	Estimate	SE	p	Estimate	SE	p
MB	-0.111	.050	.027	0.045	.018	.011	-0.021	.057	.719
SEX	0.034	.329	.918	0.005	.069	.937	-0.093	.065	.152
SES	0.038	.009	.000	-0.005	.003	.037	0.000	.008	.966
HL	0.038	.011	.000	-0.009	.003	.002	0.014	.009	.109
ETN	0.019	.007	.010	-0.005	.002	.014	-0.003	.007	.684
HR	-0.050	.008	.000	0.012	.003	.000	-0.009	.008	.260
IQ	0.008	.004	.052	-0.002	.001	.132	0.013	.004	.003

Note: STVC is the individually varying slope of time-varying covariate (the individual effect of grade retention); MB is month of birth, HL is home language, HR is high risk student.

GRADE RETENTION IN PRIMARY SCHOOL

Table 6

Total Effects of Grade Retention on Math Achievement Using Marginal Structural Models

	M-G1	M-G2	M-G3	M-G4	M-G5	M-G6
GR-K3	-0.50 (.53)	-1.28* (.60)	-0.55 (.63)	-1.32* (.62)	-1.03 (.72)	-0.37 (.77)
GR-G1		1.20 (.74)	-0.73 (.94)	-1.10 (.83)	-2.80** (.91)	-1.69 (.90)
GR-G2			2.80** (.93)	1.80* (.91)	0.93 (1.06)	0.80 (.98)
GR-G3				3.07** (1.03)	3.09** (1.15)	2.03 (1.19)
GR-G4					2.14 (1.16)	1.89 (1.20)
GR-G5						3.82** (1.01)

Note: Standard errors between brackets; effects calculated over 25 imputations.

*p < .05, **p < .01

GRADE RETENTION IN PRIMARY SCHOOL

Table 7***Degree of Overlap, and Explained Variance of Grade Retention and Math Scores***

	% matched	R ² grade retention	Math scores predicted on previous math scores	R ² math scores
GR-K3	100%	.30	M-G1: M-K3	.48
GR-G1	64.6%	.47	M-G2: M-K3, M-G1	.58
GR-G2	66.2%	.40	M-G3: M-K3, M-G1, M-G2	.63
GR-G3	68.2%	.37	M-G4: M-K3, M-G1, M-G2, M-G3	.72
GR-G4	62.7%	.33	M-G5: M-K3, M-G1, M-G2, M-G3, M-G4	.76
GR-G5	62.0%	.30	M-G6: M-K3, M-G1, M-G2, M-G3, M-G4, M-G5	.76

Note: % matched is percentage repeaters that could be matched with non-repeaters via entropy balancing. R² grade retention is Tjur's 'coefficient of discrimination' (Tjur, 2009), it is the probability that repeaters are predicted to repeat minus the probability that non-repeaters are predicted to repeat, given the covariates; R² math scores is the explained variance of the math scores for each grade, predicted on the math scores of the previous grades.

GRADE RETENTION IN PRIMARY SCHOOL

Table 8

Total Effects of Grade Retention on Math Achievement Using Sequential Conditional Mean Models

	M-G1	M-G2	M-G3	M-G4	M-G5	M-G6
GR-K3	0.50 (.44)	0.30 (.56)	1.91* (.76)	-0.42 (.72)	1.90* (.81)	1.42* (.68)
GR-G1		2.39** (.45)	1.12 (.58)	0.17 (.57)	-1.63** (.60)	-1.17* (.55)
GR-G2			4.78** (.72)	2.51** (.81)	2.80** (.86)	1.67* (.75)
GR-G3				4.19** (.87)	3.80** (.95)	1.97* (.95)
GR-G4					1.94 (1.46)	3.11** (1.16)
GR-G5						4.37** (1.21)

Note: Standard errors between brackets; effects calculated over 25 imputations.

*p < .05, **p < .01

GRADE RETENTION IN PRIMARY SCHOOL

Table A1***Test of Measurement Invariance for Math Scores on Three Transition Moments***

	χ^2/df	p	CFI	RMSEA	$\Delta\chi^2/\text{df}$	p
Transition end third Kindergarten class (K3e) – begin first grade (G1b) (1-factor model)						
CFA K3e	2051.5/135	0	.965	.049		
CFA G1b	2446.13/135	0	.965	.054		
Configural Invariance	5331.32/575	0	.969	.034		
Metric Invariance	5447.65/592	0	.968	.034	174.26/17	0
Unique Factor Invariance	5182.59/609	0	.970	.033	302.28/17	0
Transition end first grade (G1e) – end second grade (G2e) (1-factor model)						
CFA G1e	5716.81/65	0	.892	.121		
CFA G2e	2670.47/65	0	.798	.083		
Configural Invariance	9990.41/285	0	.860	.071		
Metric Invariance	9597.97/297	0	.866	.068	13.45/12	.34
Unique Factor Invariance	9832.47/309	0	.863	.067	558.33/12	0
Transition end first grade (G1e) – end second grade (G2e) (4-factor model)						
CFA G1e	448.68/59	0	.993	.033		
CFA G2e	221.83/59	0	.987	.022		
Configural Invariance	989.84/258	0	.989	.020		
Metric Invariance	13987.61	0	.802	.087	6024.09/10	0
Unique Factor Invariance	15521.63/277	0	.780	.090	1363.65/9	0
Transition end third grade (G3e) – end fourth grade (G4e) (1-factor model)						
CFA G3e	9.76/5	.082	.998	.013		
CFA G4e	2.49/5	.778	1.00	.000		
Configural Invariance	583.86/29	0	.952	.053		
Metric Invariance	894.43/33	0	.926	.062	268.92/4	0
Unique Factor Invariance	819.89/37	0	.932	.056	34.48	0

GRADE RETENTION IN PRIMARY SCHOOL

Note: values for CFI > .95 and for RMSEA < .05 indicate a good fit (Hu & Bentler, 1999; Marsh et al., 2010). For the invariance models the Theta parametrization in MPlus was used (Muthen & Muthen, 1998-2019).

GRADE RETENTION IN PRIMARY SCHOOL

Table B1*Descriptive Measures for the Covariates in the First Imputed Dataset, Before Matching in 2nd Grade*

Treated units: 568		total of weights: 568				
Control units: 6654		total of weights: 568				
	Treated			Control		
	Mean	Variance	Skewness	Mean	Variance	Skewness
Math G2	73.85	63.03	.259	80.32	61.32	.0001438
Birth Month	7.278	12.08	-.1822	6.645	11.55	-.062
Sex	.5211	.25	-.08458	.4941	.25	.02345
SES	-.4984	.4654	-.07135	-.06442	.7345	.09419
Home language	2.111	.7196	-.2127	2.529	.5519	-1.205
Ethnicity	2.644	1.566	-.1863	3.311	1.216	-1.234
High Risk	.6567	.2258	-.66	.3371	.2235	.6892
IQ	-.5564	.8813	-.7727	.06608	.7421	-.9897
Language K3	3.539	.8842	-.6635	3.824	.8697	-.5577
Language K3b	37.16	91.51	.02742	43.97	108.1	-.395
Language K3e	46.68	73.18	-.2907	52.22	76	-.2618
Language G1b	48.91	51.23	-.3618	53.34	56.89	-.0453

GRADE RETENTION IN PRIMARY SCHOOL

Table B2*Descriptive Measures for the Covariates in the First Imputed Dataset, After Matching in 2nd Grade*

	Treat			Control		
	Mean	Variance	Skewness	Mean	Variance	Skewness
Math G2	73.85	63.03	.259	73.84	63.02	.263
Birth Month	7.278	12.08	-.1822	7.277	12.08	-.1814
Sex	.5211	.25	-.08458	.521	.2496	-.0842
SES	-.4984	.4654	-.07135	-.4984	.4654	-.07161
Home Language	2.111	.7196	-.2127	2.111	.7194	-.2125
Ethnicity	2.644	1.566	-.1863	2.644	1.566	-.1854
High Risk	.6567	.2258	-.66	.6562	.2256	-.6578
IQ	-.5564	.8813	-.7727	-.5563	.8812	-.773
Language K3	3.539	.8842	-.6635	3.538	.8841	-.662
Language K3b	37.16	91.51	.02742	37.16	91.49	.02904
Language K3e	46.68	73.18	-.2907	46.68	73.17	-.2884
Language L1b	48.91	51.23	-.3618	48.9	51.22	-.359

GRADE RETENTION IN PRIMARY SCHOOL

Table B3***Effects of grade retention (GR) at T_x on M_{x-1} after matching on M_x pooled over ten imputations***

<i>Effect</i>	<i>Effect size</i>	<i>Within Imputation Variance</i>	<i>Between Imputation variance</i>	<i>Standard Error</i>	<i>df</i>	<i>t</i>	<i>P</i>
GR-K3	.2403	.2644	.0297	.5451	597.07	.441	.670
GR-G1	-1.0539	.0808	.1881	.5365	10.37	-1.965	.039
GR-G2	-.8199	.1233	.1043	.4878	19.39	-1.681	.055
GR-G3	-1.3497	.1651	.2234	.6410	13.06	-2.106	.028
GR-G4	-.5686	.3269	.1149	.6733	69.15	-.845	.201
GR-G5	.0643	.3307	.1977	.7404	29.80	.087	.534

Figure 1***Schematic Growth Model of Mathematical Ability***

Note: η_0 and η_1 are the linear growth parameters (intercept and slope) of mathematical ability. v_1 and v_2 are anchor items between T1 and T2, v_4 is an anchor item between T2 and T3.

Figure 2

Two Ways to Control for Confounding Variables.

Note: The causal effect of grade retention (GR) on math achievement (M) can be identified and estimated after controlling for all confounding variables (C), either by equating (matching) the group of repeaters and non-repeaters on all confounding variables, or by statistically adjusting the effect of GR on M for the influence of C.

GRADE RETENTION IN PRIMARY SCHOOL

Figure 3*Growth curves for retention groups per grade*

Figure 4***Latent Growth Model for Math Scores.***

Note: The factor loading for all math scores in the intercept latent variable (*I*) is equal to 1.

The factor loading of the math score at the beginning of first grade (*M1b*) in the slope latent variable (*S*) is fixed to zero, and the loading of the math score at the end of first grade (*M1e*) in *S* is fixed to 1. All the other loadings in *S* are freely estimated (*).

GRADE RETENTION IN PRIMARY SCHOOL

Figure 5*Growth Curves Based on Latent Growth Model for Grade Repeaters*

Figure 6.***Cross-Lagged Panel Model for Grade Retention Effects on Well-being***

Note. All GR-variables are allowed to freely correlate with each other, as well as with WB3K, which is treated as an exogenous variable. The latent variable alpha, which captures all unmeasured time invariant confounders, is also allowed to correlate with all exogenous variables. WB1 to WB6, all have their loading in Alpha fixed to one. Also, all GR-variables are allowed to correlate with the error terms (not shown) of all dependent variables (WB) at prior time points (dotted double-edged arrows), which makes the GR-variables predetermined.

GRADE RETENTION IN PRIMARY SCHOOL

Figure 7*Diggle-Kenward Model for Grade Retention Effects on Math Scores*

Note: $R1$ to $R6$ are survival indicators, and the dashed lines towards these indicators represent logistic regression equations. Otherwise, same as in Figure 4.

GRADE RETENTION IN PRIMARY SCHOOL

Figure 8***Latent Growth Model for Math Scores with Latent Random Variable STVC and Time******Invariant Predictors***

Note: MB is month of birth, HL is home language, and HR is 'high risk' learner. To keep the figure readable, only the influences of STVC are shown on the effects of grade retention in third kindergarten class (dotted lines). The model also estimates all influences of STVC on the effects at the other times of retention. For the sake of clarity, only the effects of MB on Intercept (I) and Slope (S) are shown, all other effects of the time-invariant predictors on I and S are also estimated in the model. Otherwise same as in Figure 4.