

Article

How Should I Teach from This Month Onward? A State-Space Model That Helps Drive Whole Classes to Achieve End-of-Year National Standardized Test Learning Targets

Obed Ulloa  and Roberto Araya * 

Centro de Investigación Avanzada en Educación, Instituto de Educación, Universidad de Chile, Santiago 8320000, Chile

* Correspondence: roberto.araya.schulz@gmail.com; Tel.: +56-9-9599-0251

Abstract: Every month, teachers face the dilemma of what exercises their students should practice, and what their consequences are regarding long-term learning. Since teachers prefer to pose their own exercises, this generates a large number of questions, each one attempted by a small number of students. Thus, we could not use models based on big data, such as deep learning. Instead, we developed a simple-to-understand state-space model that predicts end-of-year national test scores. We used 2386 online fourth-grade mathematics questions designed by teachers, each attempted by some of the 500 students in 24 low socioeconomic schools. We found that the state-space model predictions improved month by month and that in most months, it outperformed linear regression models. Moreover, the state-space estimator provides for each month a direct mechanism to simulate different practice strategies and compute their impact on the end-of-year standardized national test. We built iso-impact curves based on two critical variables: the number of questions solved correctly in the first attempt and the total number of exercises attempted. This allows the teacher to visualize the trade-off between asking students to perform exercises more carefully or perform more exercises. To the best of our knowledge, this model is the first of its kind in education. It is a novel tool that supports teachers drive whole classes to achieve long-term learning targets.

Keywords: digital systems; educational systems; state-space models; optimal control; long-term learning prediction; learning analytics



Citation: Ulloa, O.; Araya, R. How Should I Teach from This Month Onward? A State-Space Model That Helps Drive Whole Classes to Achieve End-of-Year National Standardized Test Learning Targets. *Systems* **2022**, *10*, 167. <https://doi.org/10.3390/systems10050167>

Academic Editor: Jui-Che Tu

Received: 4 August 2022

Accepted: 19 September 2022

Published: 26 September 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Teachers have to constantly estimate the probable effect of their decisions on the long-term learning achieved by their students. For each learning objective of the national curriculum, teachers have to make the students practice enough exercises. A key decision they must make is between asking students to perform exercises more carefully or perform more exercises. Thus, every month, teachers face the dilemma of what and how many exercises their students should practice, and what their consequences are on long-term learning. For that, they need to have a good prediction mechanism. This dilemma presents several challenges.

First, the mechanism should predict their students' long-term learning. This is not straightforward because long-term learning is not always the same as short-term learning [1]. Long-term learning is the learning evaluated at the end of the year, while short-term learning is learning evaluated during the session. This is a well-known, unintuitive phenomenon, where fast progress measured on real time generates an illusion of mastery in the students [2,3]. This illusion makes it difficult for the teachers to make a good prediction of the long-term learning of each of their students. Second, we must consider that teachers prefer to design their own questions. This generates a huge diversity and number of questions, each answered by very few students. Therefore, the data gathered are not suitable for typical big-data mining, such as deep learning algorithms. For example, Ref. [4] found that

logistic regression leads to data sets of a moderate size, whereas deep knowledge tracing leads to data sets of a large size.

Third, from one month to the next, teachers need a simple update mechanism. For example, they could use the average score on the answers to questions posed on the current month and an accumulative score of the previous months. There is no need for the whole history of scores of all answers to previous questions. The mechanism must also combine in a simple way, hopefully additively, with information on other personal attributes of students, their class and their school. These are personal attributes, such as student beliefs, course attributes, such as classroom climate, and school attributes, such as its historical performance on national standardized tests.

Fourth, the mechanism should facilitate simulating different practice strategies for the coming months. These strategies, together with the assessment in the current month, should be enough to predict the effect on the long-term learning. Teachers normally use grade point average (GPA) for this purpose. GPAs are very simple to understand. They are just averages of partial grades obtained across the year. Teachers could welcome a better predictor of long-term learning if it is as simple and interpretable as GPAs. However, they do not consider all the exercises performed by the students. There is also no evidence of its predictive power. If we were to use nonlinear models for this prediction, such as gradient boosted machines, k-nearest neighbors or random forests, this goal would not be possible. Non-linear intertwined relationship between the estimated model parameters and the predicted score are difficult to interpret. Most machine learning models are designed with features and combinations of them that maximize prediction performance. There is no consideration to generate solutions that are meaningful to domain experts, such as teachers [5]. They are typically black boxes that are not transparent to users [6]. However, interpretability is critical for teachers to adopt the model, to simulate instructional strategies to follow in the coming months, and to use them to base their decisions.

Therefore, in order to improve the quality of education, teachers need to be able to predict for each student her long-term learning and how different strategies can impact this prediction. One way to do so is to know the current state of long-term learning of each student, and how different strategies could interact with each student's current state. Thus, based on an estimation of the actual state of long-term learning of each student and the effects of different strategies that actuate on the current state of each student, the teacher has to adjust her strategies of exercise practice in order to reach a specified long-term learning target.

2. Related Work

We review five references.

Ref. [7] developed linear predictors of students' scores of the end-of-year state test from dynamic testing metrics developed from monthly intelligent tutoring system data. They analyzed the data logs of 362 students, even though only 105 students had complete data in each of the months. They found that logs from an online tutoring system provide better end-of-year predictions than using paper and pencil benchmark tests. They found that R^2 adjusted is 0.637. However, all students attempted to solve a similar set of items. One of the challenges is that teachers prefer to have the flexibility to adapt the exercises to their own experience. The authors did not consider the situation where teachers select or design their own set of exercises. Ref. [7] also compared, month by month, their linear predictors without a pretest with baselines. They found an increasing performance as they gathered process information of more months.

On a sample of 23,000 learners in Grades 6–8 over three academic years, Ref. [8] analyzed the relative contribution of different types of learner data in statistical models that predict scores on summative assessments. They used six different categories of statistical models that predict summative scores. One of the best model categories turned out to be stepwise linear regression (SLR). In the best year, it achieved a R^2 of 0.734.

Ref. [9] analyzed 62 papers published between 2010 and 2020 that predicted mainly university students' academic outcomes and in computer sciences courses. They focused on the forms in which the learning outcomes are predicted, the predictive analytics models developed to forecast student learning, and the dominant factors impacting student outcomes. The authors concluded that the best performing predictive models were the hybrid random forest, feedforward 3-L neural network, and naive Bayes. However, there are a lack of explanations that go beyond predicting student performance.

Ref. [10] argued that multivariate linear regression models are more adequate to analyze which variables have a positive or negative impact on the predicted learning gain than deep learning nonlinear models. The authors claimed that these models do not have the best performance in terms of metrics such RMSE or R^2 , but they provide more understanding. We claim that state space models provide even better understandability.

Ref. [11] proposed a state-space model of skill formation in a home visiting program implemented in rural China. It is a dynamic learning model for multiple skills that include the sources of learning in the early life cycle. They used the model to estimate the impact of different interventions. An interesting feature of the model is that they extended the traditional measure of accuracy (% of correct answers) in two new variables: time to mastery and how stable the child's performance is after first mastery. The authors also included different rates of learning among students.

3. Theoretical Framework

An important part of the learning challenge of school mathematics is due to an evolutionary mismatch or evolutionary trap [12]. Our brain is the product of a long process of evolution that has taken place over millions of years. In this time, the mechanism of natural selection has evolved highly efficient mechanisms to acquire the knowledge and skills required for the life of hunter-gatherers. These are the learning mechanisms of imitation, play, storytelling, and teaching, which enable us to learn to walk, talk, gather food, hunt, cook, fight other tribes, and maintain a productive social life in our own small groups. These types of knowledge and skills are called biologically primary [13]. However, because of the enormous strides in cultural evolution over the past 5000 years, we now need students to learn an incredible volume of radically new concepts in just a few short years. Many of these concepts are unintuitive and seem to contradict our innate ideas. We need students to learn abstract concepts that are very different from those acquired by their ancient ancestors, ideas such as positional notation, negative numbers, fractions, and algebra. These are biologically secondary knowledge [13]. Given the learning challenges identified in this evolutionary framework, students need well-designed and tested lesson plans to guide them along proven learning paths.

A widely tested strategy in laboratories is that of deliberate practice. This is a strategy where the students occupy their time with performing exercises instead of passively listening, restudying or rereading. Thus, after reading a text, or studying it, or attending a lecture, students practice answering quizzes or tests. These quizzes and tests are low-stakes or no-stakes assessments. These formative assessments provide a reliable estimate of the state of knowledge reached by each student during the session. However, it is very important to distinguish between the knowledge reached just after doing learning activities and the definitive knowledge that the student will reveal in the long term [14]. They are not the same thing. There is a big difference between what the student demonstrates during or just after attending a lesson, studying or doing exercises, and what they reveal months or years later.

To estimate the long-term learning is a big challenge for the teacher since there are many examples in the literature of very good learning measured immediately after the intervention but declining in the long term. However, it is not only a phenomenon of decaying and forgetting. It is a much more complex dynamics. There are alternative interventions where the dynamics is the other way around [1]. These alternative interventions require more effort from the student and cause a slower progress. Students reach a weak

performance in the short term. However, they reach good performance in the long term. In these alternative teaching strategies, the long-term learning beats the long-term learning obtained with the first interventions. This reversal phenomenon is difficult to swallow. On each session, based on quizzes or formative assessments, the teacher can obtain an estimate of what each student knows at that moment. However, it is much more difficult to obtain a good prediction of what students will know at the end of the year. According to [1], during teaching, what we can observe and measure is performance, which is often an unreliable index of long-term learning. Ref. [1] make a critical distinction between performance, measured during acquisition, and learning, measured by long-term retention or transfer capacity. This is an unintuitive phenomenon, where fast progress generates an illusion of mastery in the students [2,3]. This illusion also hinders the teacher, and makes it difficult for her to make a good prediction of each of her students' long-term learning.

There is empirical evidence that the effect on long-term learning of practice testing is higher than using the time for studying, summarizing or rereading. This finding is known as the testing effect [15]. However, there are no implementation studies in schools with typical situations. These are situations where teachers design their own questions and there are very few answers per question. Additionally, teachers need clear guidelines on how to use the accumulated information. It is critical for the teacher to use that information to decide how to adjust teaching strategies for the months that follow. In this paper, we explore those challenges. More precisely, the research questions are the following:

Research Question 1: To what extent can teacher-designed questions (low- or no-stake quizzes) help predict students' long-term learning as measured by end-of-year standardized state tests?

Research Question 2: To what extent can a hidden Markovian state of the student's accumulated knowledge up to the current month be estimated so that, along with the probable deliberate practice on the next months, it is sufficient to predict with good accuracy the end-of-year standardized state tests?

Research Question 3: To what extent can a state-space model help teachers to visualize the trade-off between asking students to perform exercises more carefully or perform more exercises, and thus help drive whole classes to achieve long-term learning targets?

This hidden state would capture the knowledge stored in long-term memory as well as the ability to retrieve it and apply it properly.

The objective is to optimize long-term learning as measured in end-of-year state standardized tests. This is different to short-term learning, and particularly to the performance of the next question. One key restriction of the problem is that the questions of the formative assessments were proposed by the teacher. Therefore, they are questions that were not previously piloted. Moreover, for every question, there are very few students that have attempted it. This situation makes it very different from the ones where big data algorithms can be used.

4. Materials and Methods

In this paper, we use the questions and answers in ConectaIdeas [16–21]. This is an online platform where students answer closed- and open-ended questions. On ConectaIdeas, teachers develop their own questions, designing them from scratch, or they select them from a library of questions designed previously by other teachers. Then teachers use those questions to build their own formative assessments composed by sets of 20 to 30 questions. Students answer them in laboratory sessions held once or twice a week, or at home.

The decision to work with fourth grade students is due to three facts. First, every year in Chile, there is a national standardized end-of-year test at that grade level. This is only true for the 4th, 8th and 10th grades. This fact is, therefore, very important for schools and teachers. Second, at a younger age, the return on investment is greater than with older students, as shown by the Heckman curve [22]; see, however, [23]. Third, in the fourth grade, students can autonomously login and answer on their devices.

We used data from 24 fourth-grade courses from 24 different schools with low socio-economic status in Santiago, Chile. Each school has two fourth-grade courses. In each school, an independent third party selected a course at random. We used data from all students in those selected courses. From those classes, there were a total of 256 boys and 244 girls that took the pretest and the national standardized end-of-year test. On average, we have 20.8 students per class. The average age at the beginning of the school year was 9.5 years with SD of 0.6 years. The courses belong to high vulnerability schools. The official vulnerability index of the schools is 0.28 SD above the national average. The historical performance of these schools in mathematics is very low. The historical score on the national standardized end-of-year test of the previous 3 years of the schools is 0.74 SD below the national average. All data correspond to the 2017 school year.

Unlike intelligent tutoring systems, teachers select a list of questions for the whole class and even write their own questions. Those students who manage to correctly answer 10 questions in a row in the session become candidates to be tutors of the session. From that list, the teacher selects 2 or 3 students as tutors. So, when a student asks for help, a tutor can help them. In those moments of the session when there is an option to select a peer or the teacher, students prefer a student tutor [24]. Additionally, the ConectaIdeas early warning system tracks the behavior of students' responses. If a student makes too many mistakes or is answering too slow, the system alerts the teacher to help the student or assign a tutor to help the student.

From a total of 2386 closed questions, only 60 questions were attempted by at least 90% of the students. Moreover, only 8.6% of the closed questions were attempted by at least one student of all of the participating classes. On the other hand, from a total of 1071 open-ended questions, only 8 questions were answered by more than 5% of the students. In this paper, we do not analyze the answers to open-ended questions. Although there are much fewer open questions than closed ones, the answers to open-ended questions also manage to make a contribution to learning and improve the prediction of the score in the national standardized end-of-year test [25]. This improvement is achieved even when the standardized national test at the end-of-the-year only has closed questions and they only are of the multiple-choice type.

Each month, the students carried out closed-ended exercises on the ConectaIdeas platform. The exercises are from the five strands of the fourth-grade Chilean mathematics curriculum:

1. Numbers and operations;
2. Patterns and algebra;
3. Geometry;
4. Measurement;
5. Data and probabilities.

In the school culture, teachers use periodic evaluations. In the vast majority of cases, these are monthly evaluations. Based on these evaluations, teachers make decisions to adjust the development of their lessons for the next month. Therefore, in this paper, we explore making long-term learning predictions based on the performance of the exercises done each month. Ref. [7] also made predictions based on monthly evaluations.

Table 1 shows the average number of closed-ended exercises performed each month and the corresponding standard deviations. The table shows that the students performed an average of 1089 exercises during the year. The month with the most exercises performed was October, just before the year-end national test. The strand with the most exercises is numbers and operations. A total of 65.2% of the exercises performed are from that strand.

Table 1. Number of closed-ended exercises performed by month for each strand, and monthly average.

Strand	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Average
Numbers	6	120	123	69	56	122	119	93	89
Patterns	0	2	4	18	7	1	14	27	9
Geometry	0	0	11	22	7	24	17	64	18
Measurement	0	0	5	35	11	17	11	28	13
Data	0	1	2	3	5	3	7	30	6
Total	6	124	146	147	86	168	168	243	136

In order to select the features to be used in the models, we calculated the correlation of 68 variables with the national standardized end-of-year test. From those, we chose 28 variables with which we trained linear models in a set of training courses to predict the end-of-year scores in the national standardized end-of-year test. We then selected the variables that achieved the best prediction in a set of test courses. These variables were the pre-test, a variable of students' beliefs (the degree of agreement on a scale of 1 to 5 with "Mathematics is easy for me"), one variable that indicates the climate of the class (the average of all the students of the class to the degree of agreement of each student on a scale of 1 to 5 with "My behavior is a problem for the math teacher"), and two performance variables on the formative assessments. These performance variables are the percentage of correct answers on the first attempt, and the difference between the number of questions answered correctly on the first attempt and the number of questions answered correctly not on the first attempt. Every time the student makes a mistake, the system notifies them of the error. Thus, in the long run, the student correctly answers the question, and thus can continue with the next question. We normalize each of these variables.

4.1. The Sequence of Linear Regression Models

We previously studied both linear regression models [16,21,26–28], and state-space models [27,28]. In this paper, we develop them in much more detail.

The first type of models is a sequence of linear regression models. For each month of the school year, we trained an optimal model. The model for each month is a linear regression for the month with an optimal forgetting rate (LRMOFR). This is a model that for each month finds an optimal combination of pre-test, one student and one classroom feature, and the two historic performance variables up to the current month on formative assessments. The historic performance is the addition of all the discounted average performance of each month up to the current month. The discounting is computed with an optimal forgetting rate. Thus, for each month, we have a prediction model. The model for a new month does not use the finding of the optimal parameters of the previous months. They are completely new optimization models. The only common optimal parameters are the forgetting or discount rate and the regression intercept.

The central problem in this type of models is to find the contribution to the prediction of each component of the model. It is then necessary to find in each month how much an increase of one standard deviation of the pretest contributes to the national standardized end-of-year test (SIMCE) test score, how much an increase of one standard deviation in the student's subjective appreciation that math is easy contributes to the SIMCE score, how much a one standard deviation increase in poor classroom climate contributes to the SIMCE score, and how much a one standard deviation increase in both historical performance features contributes to the SIMCE score. Here, the historical performance until each month is the weighted average of the monthly percentage of exercises solved correctly on the first attempt in the formative evaluations from the beginning until the month, and the difference between the number of exercises solved correctly in the first attempt and the number of exercises solved correctly not in the first attempt.

More precisely, for a student i from a school j , with information until month t , the linear regression for each month with the optimal forgetting rate model (LRMOFR) predicts the SIMCE score:

$$\hat{y}_{i,j,\tau} = \gamma_\tau p_i + \beta_0 + \beta_{1,\tau} \Theta_j^{(1)} + \beta_{2,\tau} \Theta_i^{(2)} + \frac{\sum_{t=1}^{\tau} \rho^{\tau-t} \left(\sum_{k=1}^5 \delta_{k,\tau} a_{i,j,k,t} + \eta_{k,\tau} d_{i,j,k,t} \right)}{\sum_{t=1}^{\tau} \rho^{\tau-t}} \tag{1}$$

- p_i : pretest of the student i ($p_i \in \mathbb{R}$, normalized).
- $\Theta_j^{(1)}$: class j average for “My behavior is a problem for the math teacher”.
- $\Theta_i^{(2)}$: student i response to “Math is easy for me”.
- $a_{i,j,k,t}$: percentage of exercises solved on the first attempt for a student i , on the k strand, at month t (normalized for each strand month).
- $d_{i,j,k,t}$: the difference between exercises solved on the first attempt and those that took more than one for a student i , on the k strand, at month t (normalized for each strand month).

On the model, the parameters are $\rho, \beta_0, \beta_{1,\tau}, \beta_{2,\tau}, \delta_{k,\tau}, \eta_{k,\tau}$ and γ_τ where $\tau \in \{1, \dots, 8\}$ (from March until October since the test was on November), $k \in \{1, \dots, 5\}$. This means that the ρ and β_0 parameters are shared by every \hat{y}_τ model and each has its specific $\beta_{1,\tau}, \beta_{2,\tau}, \delta_{k,\tau}, \eta_{k,\tau}$ and γ_τ estimates. Then, that makes it that each \hat{y}_τ LRMOFR has 15 parameters and for the 8 months overall, that is 106 estimates.

To find these values, we minimize the average RMSE across months:

$$(P) \quad \begin{aligned} \min_{\rho, \beta_0} \quad & \frac{1}{8} \sum_{\tau=1}^8 \sqrt{e_\tau(\rho, \beta_0)^2} \\ \text{s.t.} \quad & \rho \in [0, 1] \\ & \beta_0 \in [210, 275] \end{aligned}$$

$$(P_\tau) \quad \begin{aligned} e_\tau(\rho, \beta_0) = \min_{\theta} \quad & \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_{i,j,\tau})^2} \\ \text{s.t.} \quad & \theta = (\gamma_\tau, \beta_{1,\tau}, \beta_{2,\tau}, \delta_{1,\tau}, \dots, \delta_{5,\tau}, \eta_{1,\tau}, \dots, \eta_{5,\tau}) \\ & \hat{y}_{i,j,\tau} = \gamma_\tau p_i + \beta_0 + \beta_{1,\tau} \Theta_j^{(1)} + \beta_{2,\tau} \Theta_i^{(2)} \\ & \quad + \frac{\sum_{t=1}^{\tau} \rho^{\tau-t} \left(\sum_{k=1}^5 \delta_{k,\tau} a_{i,j,k,t} + \eta_{k,\tau} d_{i,j,k,t} \right)}{\sum_{t=1}^{\tau} \rho^{\tau-t}} \end{aligned}$$

and (P_τ) is the OLS problem for each linear regression, making it a 106-parameter model.

A good characteristic of these linear models is that each of their parameters is clearly understood and they have direct implications in the prediction. Their effects are independent of each other, and the total effect is just an addition of each separate effect. However, it is one model per month. In each month, there are 13 parameters. That means 104 parameters. If we add the 2 overall parameters, we have a total of 106 parameters. This is a huge number of parameters. It is very difficult to memorize, which can make analysis difficult for a teacher.

Ref. [28] analyzed other possibilities of models with linear regressions. The one presented in this paper has the best prediction error and, at the same time, is the one with the least number of parameters.

4.2. The State-Space Model

The second type of model is an accumulator model. In the first month, it is fed with three inputs: the contribution of the pretest, the contribution of the student's subjective appreciation that math is easy, and the contribution of poor classroom climate indicator. More precisely, as in the linear regression models, the contribution to the prediction of the national standardized end-of-year test of each of the three inputs is proportional to the difference in standard deviations of the input and its population average. Then on each of the next months, we add the contributions of the corresponding two monthly formative assessments' performance. More precisely, from month t to the next month $t + 1$, we have the following:

$$\begin{aligned}x(t+1) &= Ax(t) + Bu(t) + w(t) \\y(t) &= x(t) + v(t)\end{aligned}\tag{2}$$

where $t \in \{0, \dots, T\}$, $x(t), y(t) \in \mathbb{R}$, $u(t) \in \mathbb{R}^{13}$, $A \in \mathbb{R}$, $B \in \mathbb{R}^{13}$, $w(t) \sim \mathcal{N}(0, Q)$ and $v(t) \sim \mathcal{N}(0, R)$ are the noise with variances $Q, R \in \mathbb{R}$.

- $x(t)$ is the hidden state that represents the long-term knowledge of a student.
- $y(t)$ are the observed measurement as a grade or score.
- $u(t)$ has historical variables at $t = 0$ and dynamic ones for $t > 0$.

With this structure, we can predict with information until month τ with the equations:

$$\begin{aligned}\hat{x}(\tau, t+1) &= (A - K)\hat{x}(\tau, t) + Bu(t) + Ky(t) & t \in \{0, \dots, \tau\} \\ \hat{x}(\tau, t+1) &= A\hat{x}(\tau, t) + B\hat{u}(\tau, t) & t \in \{\tau + 1, \dots, T\} \\ \hat{x}(\tau, T) &= \hat{x}(\tau, T) \\ \hat{y}(\tau, T) &= \hat{x}(\tau, T) \\ \hat{x}(0) &\text{ estimated}\end{aligned}$$

where $\hat{u}(\tau, t)$ is estimated repeating the last known value of $u(t)$ ($\hat{u}(\tau, t) = u(\tau), \forall t > \tau$) and K is Kalman gain matrix (from the respective steady state Riccati equation).

This is a much simpler model than the linear regression model. The main reasons are the following five.

First it is just one model, not a sequence of models. Therefore, it requires a much smaller number of parameters. The same parameters are used on the different months. We have a total of 17 parameters. This is a much smaller number than the 106 parameters in the LRMOFR models.

Second, the model is a simple accumulator. It is similar to a tank fed with water. From one moment to the next, we add water to the already accumulated water. In the accumulator model, instead of water, a hidden state accumulates 10 different contributions. These are two contributions for each one of the five strands of the curriculum. Additionally, the update of the hidden state is a very simple additive mechanism. In the next month, we have to just add the contributions of this new month. It is a Markovian model, which means that future states depend only on the current state, not on the events that occurred before it.

Third, to make predictions from the accumulator, we need to simulate teaching strategies in the next months with an estimated performance of the student. We can input different strategies and student performance, and compute the prediction. This is a powerful tool that gives great flexibility to the teacher. In the sequence of linear regression models, the process is more complex. There is no single place to input and simulate different teaching strategies. The model is fixed and makes predictions assuming the teacher follows the pattern of strategies used in 2017.

Fourth, the accumulator model is more robust than the sequence of linear regression models. The sequence of linear regression models is more prone to capture patterns that could have been produced in a particular month of 2017, for example, the impact

on some months of a teacher strike, social riots, or the educational effect of a big natural disaster, such as an earthquake. On the contrary, the accumulator model has a mechanism for updating from one month to the next, and therefore fits patterns of the whole year without overfitting to each month.

Fifth, the accumulator model provides a simple conceptual construct that is similar to a grade point average that all teachers use. It is a single number that represents the long-term learning of the student. It also provides an understandable and transparent updating mechanism. Moreover, it integrates personal and historic information of the student and her classroom with the student performance on the formative assessments. It converts all the information in a single currency, and provides a change mechanism for converting different educational, personal, and social information into that currency.

In both types of models, a detailed specification of the contribution of the formative assessments is critical. Given that students respond to the formative assessments in each session and that the teacher has their performance results for each month, she needs that information to conduct a what-if type of analysis. For example, to estimate the effect on the SIMCE predictor, if in a given month the student increases their performance by one standard deviation in the formative assessment of that month, clearly, the effect must be null in the months prior to the month in which the performance increases. However, in that and subsequent months, an effect should be provoked. In the language of dynamical systems, we are estimating the impulse response. That is, if in month s , the student experiences a shock or impulse of additional performance, then we estimate its effect $h(s, t)$ on the predictor of the SIMCE score that is computed in a month t .

5. Results

We tried 12 linear regression models and 12 state-space models. We tested four variants corresponding to the student's pretest: no pretest, with pretest, with third grade GPA instead of pretest, and with fourth grade GPA instead of pretest. On the other hand, we included three variants for performance month by month: without performance variables, for each of the 5 math strands of the curriculum, we included only the percentage of correct answers in the first attempt, and for each of the 5 math strands of the curriculum, we included the percentage of correct answers in the first attempt and the difference between the number of correct answers on the first attempt and the number of correct answers not on the first attempt.

In total, there are 12 linear regression models and 12 state-space models. We performed 200 cross validations, each time training in 18 classes and testing in 6 classes. The models with lowest RMSE were the linear regression and the space-state models with the following options: pretest and both monthly performance measures. In Table 2, we show, month by month, the percentage of times in which the linear regression model and the state-space model with those options had the lowest RMSE within the 24 models. The RMSE was measured in the testing classes.

Table 2. Percentage of times in which the model had the lowest RMSE within the 24 models in the testing classes.

Model	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Average
Linear regression	20.0	19.0	4.5	4.0	1.0	5.5	28.0	40.5	15.3
State-Space	19.5	24.0	61.0	62.5	69.0	69.0	35.5	32.0	46.6

Therefore, in what follows, we report only the two best models. These are the linear regression for each month with optimal forgetting rate (LRMOFR) and the state-space (SS) model with the following variables: pretest, the personal belief variable, the class climate variable, and the two monthly performance variables. These performance variables are the percentage of correct answers on the first attempt, and the difference between the

number of correct answers on the first attempt and the number of correct answers not on the first attempt.

5.1. Prediction Errors of Both Models

We found that in both models, their prediction error generally decreases as the school year progresses (Figure 1). Coincidentally, the fit between the prediction and the SIMCE results improve as the year progresses (Figure 2). This trend is expected since the models are receiving the results of the deliberate practice of the students. Every week, students carry out one or two sessions and in each of them, they answer about 24 questions that are part of the formative assessments. However, RMSE does not reach zero. This is partly because the error of the SIMCE test is between 0.26 and 0.35 standard deviations. Therefore, it is not possible for the models to go below that level of error.

It is interesting to observe that in August, both indicators worsen a little bit. Apparently, this is due to the two-week winter break in July. The models receive much less information. In July, students take half of the formative assessments of what they do on the other months.

Figure 1 shows that the RMSE of the state-space model is generally lower than that of the linear regression model, except for the first two and the last month.

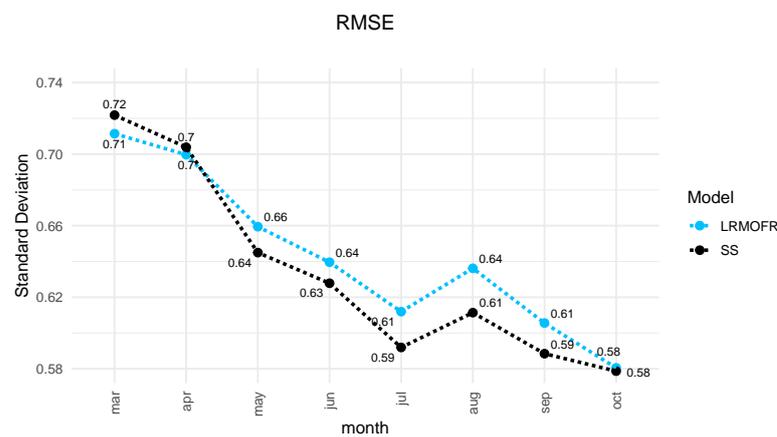


Figure 1. Root mean square error (RMSE) of the two models with best RMSE. The graph shows the RMSE in the test classes for each month of the school year. LRMOFR: linear regression for each month with an optimal forgetting rate model. SS: state-space model.

Figure 2 compares R^2 for the LRMOFR model with R^2 for the SS model. We see again that in most months, the state-state model has better R^2 .

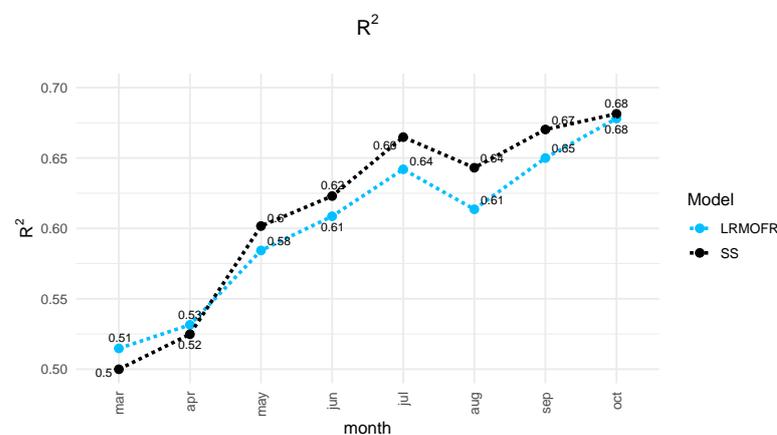


Figure 2. R^2 of the two best models for each month of the school year.

Since the state-space model has much fewer parameters, adjusted R^2 is even better for the state-space model (Figure 3).

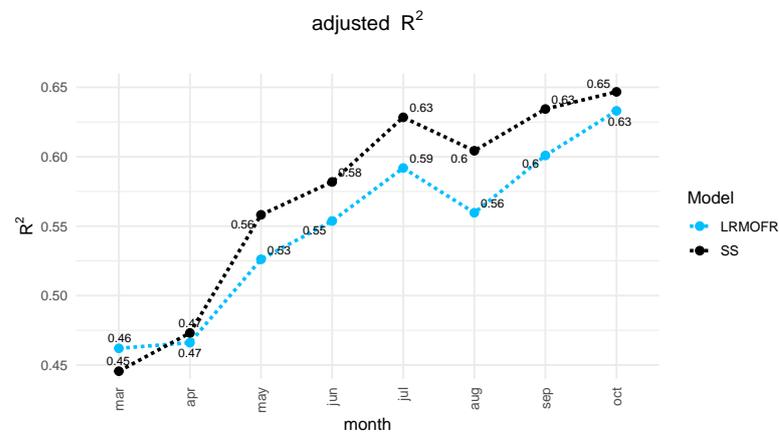


Figure 3. Adjusted R^2 of the two best models for each month of the school year.

The use of adjusted R^2 allows comparison with models developed by other authors in other educational systems, for example, with the models of Anozie et al. [7]. We changed the month in order to be able to compare with those of the Chilean school year. As shown in Figure 4, the adjusted R^2 of the state-space model is superior to the other models.

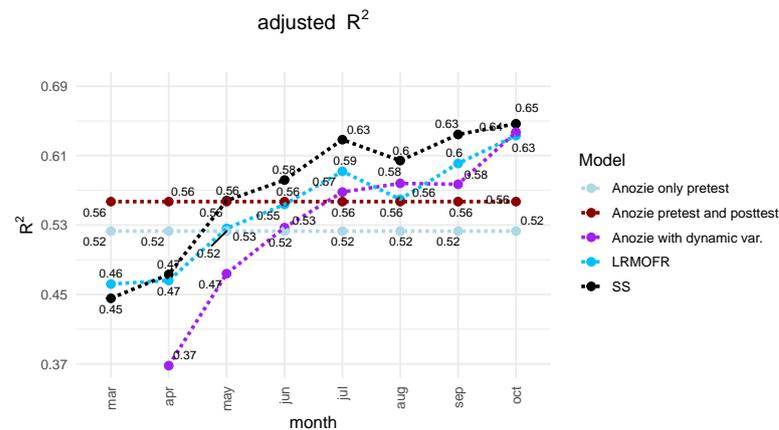


Figure 4. Adjusted R^2 of the two best models for each month of the school year.

5.2. Optimal Values of the Parameters

For the regression model, we performed 200 cross validations for each month. The results are in Table 3. The parameters are quite stable.

Table 3. Estimated parameters for the linear regression for each month with optimal forgetting rate model obtained with 200 cross validations.

Parameter	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct
ρ	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)	0.71 (0.05)
γ_τ	33.16 (2.22)	23.78 (2.24)	18.77 (2.14)	17.17 (2.16)	14.56 (2.17)	17.38 (1.91)	15.87 (1.70)	15.34 (1.59)
β_0	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)	240.04 (3.45)
$\beta_{1,\tau}$	-11.80 (1.75)	-11.49 (2.04)	-11.38 (1.94)	-9.24 (1.73)	-9.26 (1.72)	-10.27 (1.76)	-8.96 (1.71)	-8.62 (1.65)
$\beta_{2,\tau}$	5.00 (1.12)	4.83 (1.05)	4.60 (0.95)	3.98 (0.95)	4.23 (0.83)	4.01 (0.93)	3.61 (0.95)	3.00 (0.89)
$\delta_{1,\tau}$	-1.00 (2.03)	1.73 (3.29)	7.93 (2.97)	5.57 (2.99)	6.75 (2.68)	7.22 (3.24)	12.21 (2.80)	12.01 (2.49)
$\delta_{2,\tau}$	-0.38 (1.73)	3.58 (1.63)	5.23 (1.71)	3.98 (2.25)	10.57 (1.85)	7.08 (2.08)	9.30 (1.90)	6.70 (2.23)
$\delta_{3,\tau}$	-5.69 (1.63)	-2.30 (2.12)	0.52 (2.06)	5.88 (2.31)	7.18 (2.71)	7.03 (2.69)	6.86 (2.74)	4.61 (2.40)
$\delta_{4,\tau}$	-2.08 (1.89)	-0.66 (2.12)	4.14 (2.24)	8.00 (2.67)	4.55 (2.90)	2.61 (2.66)	4.45 (2.35)	3.18 (2.94)
$\delta_{5,\tau}$	1.83 (1.64)	5.19 (1.92)	2.28 (2.51)	7.76 (1.48)	13.82 (2.35)	5.16 (2.54)	8.07 (2.56)	11.09 (2.30)
$\eta_{1,\tau}$	4.70 (1.69)	14.64 (2.61)	13.34 (2.22)	15.76 (2.33)	14.85 (2.07)	9.85 (2.59)	5.28 (2.54)	3.03 (2.19)
$\eta_{2,\tau}$	-	1.94 (1.21)	2.57 (0.99)	0.71 (1.52)	-0.82 (1.87)	1.82 (2.90)	2.13 (2.56)	7.64 (2.52)
$\eta_{3,\tau}$	4.10 (1.90)	1.20 (2.74)	1.86 (1.67)	-3.29 (3.02)	-6.90 (3.66)	-4.86 (3.17)	-3.59 (3.13)	0.50 (3.52)
$\eta_{4,\tau}$	-	1.21 (1.58)	2.59 (1.75)	1.59 (2.05)	3.49 (2.03)	10.44 (2.89)	5.24 (2.10)	3.55 (2.02)
$\eta_{5,\tau}$	-	-2.62 (1.83)	1.26 (2.46)	-2.72 (1.55)	-6.25 (2.42)	-0.79 (2.08)	1.09 (3.20)	1.23 (2.54)

For the state-space model, we also performed 200 cross validations, but not separated by month. The results are in Table 4. Again, the parameters are quite stable. We found that A is basically 1. That is, the state space is a pure accumulator.

Table 4. Estimated parameters for the state-space model obtained with 200 cross validations.

Parameter	Mean	SD
x_0	258.3901	3.4421
A	1.0010	0.0006
K	0.0835	0.0085
$Q^{\frac{1}{2}}$	0.8982	0.1080
$R^{\frac{1}{2}}$	10.4133	0.5388
B_1	0.3787	0.2285
B_2	0.5888	0.1422
B_3	0.0770	0.1739
B_4	0.1984	0.1855
B_5	0.6834	0.1631
B_6	1.1153	0.2337
B_7	0.3170	0.1014
B_8	-0.0979	0.1365
B_9	0.4526	0.1363
B_{10}	-0.0675	0.1544
B_{11}	24.6766	2.0893
B_{12}	-14.0852	2.0651
B_{13}	6.143	1.2942

A is 1, but K is 0.08. Therefore, $(A - K)$ is 0.9175. Thus, the predictor uses a forget rate that is close to 8% from one month to the next one. The standard deviation $R^{\frac{1}{2}}$ of the random forcing term $w(t)$ is low and equal to 0.9 SIMCE points. The standard deviation $Q^{\frac{1}{2}}$ of the measurement noise $v(t)$ is about 10 SIMCE points, which is in the order of that reported by the SIMCE year-end national statistical test.

For the accuracy indicators, the highest parameter B that translates accuracies in SIMCE points is for the data and probabilities strand. However, it does not stand out much from the others. For the other performance indicator, the difference between the number of questions answered correctly on the first attempt and the rest of the questions, the number and operatives strand has a B above the rest of the strands. It is a much larger B than the B 's of the rest of the strands. That is expected, given that the largest number of exercises performed belongs to the numbers and operations strand, and also to the fact that the year-end national standardized test has a lot more questions from that strand. There are some values of B that are negative. This may be because it is perhaps convenient to do

fewer exercises of those strands that have less weight in the end-of-year test, and then spend more time practicing exercises of the other strands.

5.3. Contribution of the Different Variables to the Prediction

If we expand the state-space expression, we can compare term by term with the LRMOFR model as shown in Table 5.

Table 5. Term-by-term comparison of each of the effects of the LRMOFR model with the SS model.

Effect	LRMOFR	SS
Intercept	β_0	$A^{T-\tau-1}(A-K)^\tau \hat{x}(0) + A^{T-\tau-1} \sum_{t=1}^{\tau} (A-K)^{\tau-t} Ky(t)$
Pretest	$\gamma \tau p_i$	$B_{11}(A^{T-\tau-1}(A-K)^\tau u_{11}(0))$
‘My behavior is a problem...’	$\beta_{1,\tau} \Theta_j^{(1)}$	$B_{12}(A^{T-\tau-1}(A-K)^\tau u_{12}(0))$
‘Math is easy for me’	$\beta_{2,\tau} \Theta_i^{(2)}$	$B_{13}(A^{T-\tau-1}(A-K)^\tau u_{12}(0))$
Accuracy	$\frac{\sum_{t=1}^{\tau} \sum_{k=1}^5 \rho^{\tau-t} \delta_{k,\tau} a_{i,j,k,t}}{\sum_{t=1}^{\tau} \rho^{\tau-t}}$	$\sum_{k=1}^5 B_k \left(A^{T-\tau-1} \sum_{t=1}^{\tau} (A-K)^{\tau-t} u_k(t) + \sum_{t=\tau+1}^{T-1} A^{T-t-1} \hat{u}_k(\tau,t) \right)$
Difference	$\frac{\sum_{t=1}^{\tau} \sum_{k=1}^5 \rho^{\tau-t} \eta_{k,\tau} d_{i,j,k,t}}{\sum_{t=1}^{\tau} \rho^{\tau-t}}$	$\sum_{k=6}^{10} B_k \left(A^{T-\tau-1} \sum_{t=1}^{\tau} (A-K)^{\tau-t} u_k(t) + \sum_{t=\tau+1}^{T-1} A^{T-t-1} \hat{u}_k(\tau,t) \right)$

We first analyze the intercept. It is the prediction made each month if we did not have any data on the student. In the case of the LRMOFR model, we impose the same value for all months. In the SS model, there is an initial estimate that falls but is corrected with the observations of each month. The expression for SS is

$$A^{T-\tau-1}(A-K)^\tau \hat{x}(0) + A^{T-\tau-1} \sum_{t=1}^{\tau} (A-K)^{\tau-t} Ky(t)$$

Since A is 1, this expression is

$$(A-K)^\tau \hat{x}(0) + \sum_{t=1}^{\tau} (A-K)^{\tau-t} Ky(t)$$

We see in Figure 5 that the values of the intercept of both models are similar and in the last months they coincide.

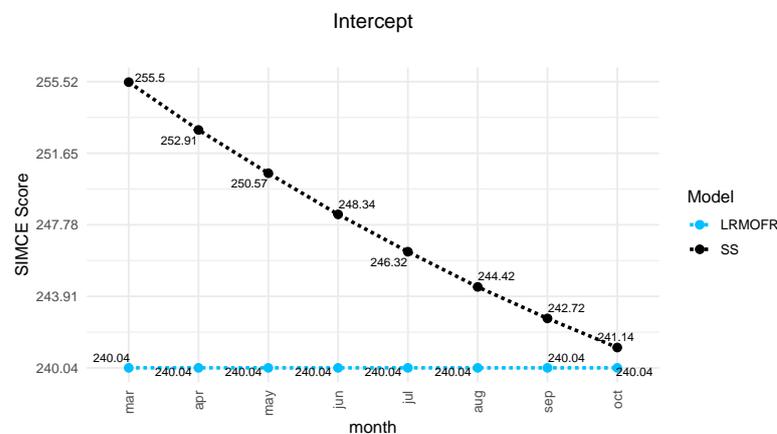


Figure 5. Intercept on the LRMOFR and SS models for each month of the school year.

Next, we analyze the effect of the pretest. We increase the pretest on standard deviation and compute its effect on the prediction. In both models, in the first months, the effect of

the pretest is much higher than in the later month of the school year (Figure 6). For the state-space model, the effect is monotonically and exponentially decreasing. In March, one standard deviation increase in the pre-test causes an increase of 0.54 standard deviations in the prediction of the SIMCE score, whereas in October it causes an increase of 0.28 standard deviations. This means that the effect decreases to one half. In the sequence of linear regressions, the effect is also decreasing but rapidly reaches a plateau in May. The graph shows the values of the parameter γ_τ presented in Table 3 and the values of the expression $(A - K)^\tau B_{11}$ with the value of the parameters presented in Table 4 but scaled to standard deviations of the SIMCE test score.

The decay of the effect of the pretest is natural and expected. Given that the platform constantly accumulates the performance of hundreds of exercises performed by the student, then as the times passes by, the pretest information is decreasingly important to predict the result of the end-of-year test. Only at the beginning of the year, the pretest is a highly valuable information. What these models add is the quantification of the decay. In the case of the state-space model, the decay is at a rate of 91.75% from one month to the next. This is a loss of more than 8% per month.

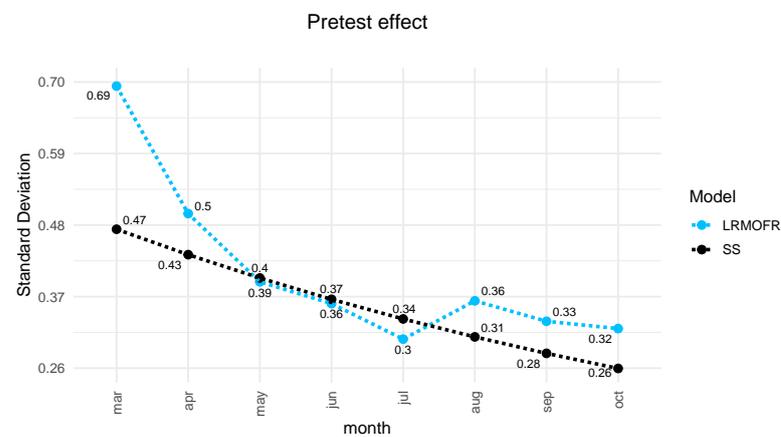


Figure 6. Pretest effect. This is the contribution on the SIMCE prediction of one standard deviation increase of the pre-test.

At the classroom level, the greatest effect was one that is a proxy for the classroom climate. It is the average of the responses of all the students in the class to the degree of agreement in a scale from 0 to 5 with the statement that “My behavior is a problem for the teacher”. At the beginning of the school year, in March, an increase of one standard deviation in this response generates a decrease of 0.27 standard deviations in the prediction of the SIMCE score of each student in the course (Figure 7). By the end of the year, this negative contribution decreases by approximately half. As Figure 7 shows, the decay of this classroom climate effect in the state-space model is exponential, while in the linear regression model, it is more complex. Moreover, it basically reaches a plateau in June. In the state space, again the decay from one month to the next of the classroom climate is 91.75%. This is the same as the decay of the pretest. This means that in each variable, the contribution to the end-of-year test drops by 8% from one month to the next. Figure 7 shows the values of the parameter $\beta_{1,\tau}$ presented in Table 3 and the values of the expression $(A - K)^\tau B_{12}$ with the value of the parameters presented in Table 4, but scaled to standard deviations of the SIMCE test score.

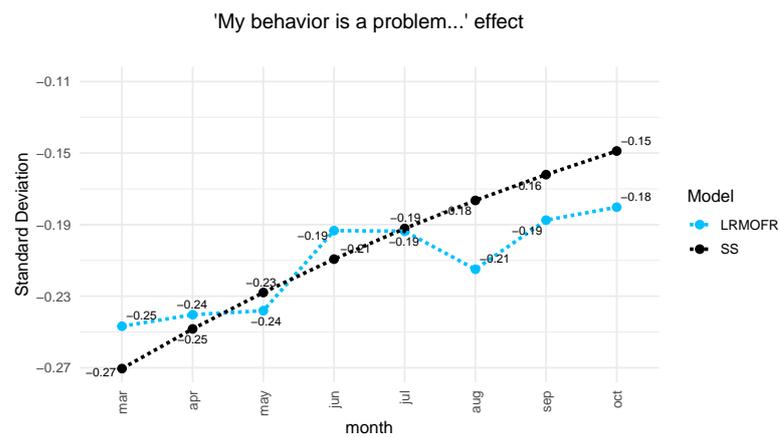


Figure 7. Effect of the class average of the degree of agreement with “My behavior is a problem for the teacher”. This is the contribution to the SIMCE prediction of an increase of one standard deviation of the student belief.

The personal variable that contributes the most to the prediction of the SIMCE score is the belief that mathematics is easy for me (Figure 8). This is the degree of agreement on a scale of 1 to 5 with “Mathematics is easy for me”. At the beginning of the school year, in March, a student’s one standard deviation increase in that belief causes a 0.12 standard deviation increase in the predicted SIMCE score. However, as the school year progresses, this contribution decreases. At the end of the year, the contribution of this belief reaches half of what it had in March. As Figure 8 shows, the decay of this effect in the state-space model is exponential, while in the linear regression model, it is more complex. In the state space, the decay of this belief is again a 91.75% from one month to the next, the same as the decay of the pretest and of the class climate. Figure 8 shows the values of the parameter $\beta_{2,\tau}$ presented in Table 3 and the values of the expression $(A - K)^T B_{13}$ with the value if the parameters presented in Table 4 but scaled to standard deviations of the SIMCE test score.

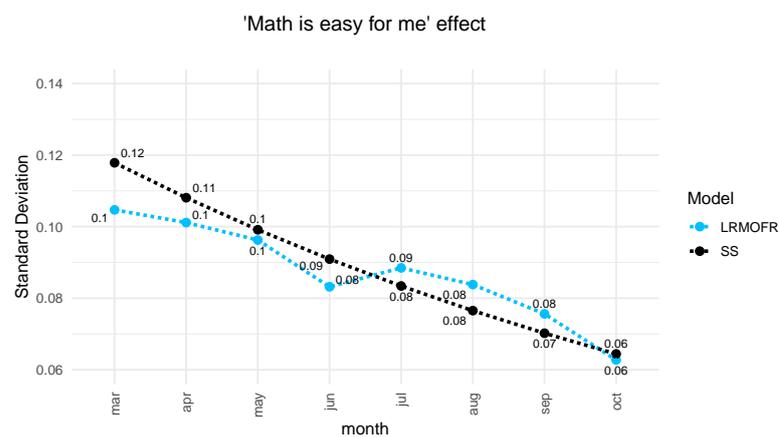


Figure 8. Effect of the “Math is easy for me” subjective belief. This is the contribution to the SIMCE prediction of an increase of one standard deviation of the student belief.

We now compare the effects of the deliberate practice of exercises from each of the five strands of the national curriculum. In each month, we have two performance metrics for the five strands. On the one hand, we have the accuracy. This is the rate of exercises answered correctly on the first attempt. On the other hand, we have the difference between the number of exercises answered correctly on the first attempt and the number of exercises answered correctly on other attempts.

For the LRMOFR model, Table 3 shows the effects on the prediction of the end-of-year test of a one standard deviation increase in the accuracies. These effects are measured by

the parameters δ for each curriculum strand and for each month. In general, the effects are positive. The first month of the school year is March. However, this is a month with very little activity since the implementation started in the fourth week of that month. In general, the numbers and operations strand is the one with the greatest effect. Students perform most of the exercises corresponding to this strand. The percentage of exercises that the students performed that belong to this strand is 58.78%. On the other hand, this strand is the one with the most items in the end-of-year SIMCE test. The second strand with the greatest effect is patterns and algebra, which has the greatest effect in July. Then there is the measurement strand, which is the one with the most effect in June. In July and August, the geometry strand has a great effect. Only in October, just before the SIMCE test, does the data and probability strand have the higher effect. This pattern is typical in fourth grade classes in Chile. The data and probability strand is taught at the end of the year.

On the other hand, an increase of one standard deviation in the difference between the number of questions answered correctly on the first attempt and the number of questions answered correctly on other attempts generates far more effect on the strand of numbers and operations. This large effect is shown in Table 3 in the η parameters. The only different month is the final month, October, where the patterns and algebra strand has the highest effect.

Reviewing Table 4 and comparing the effect values between the LRMOFR and SS models, we see that the effects of the LRMOFR model decay at the rate $\rho = 0.71$. This decay is much faster than the one in the SS model, whose decay rate is $(A - K) = 0.9175$. On the other hand, in the SS model, the effect of the difference in the number of exercises answered correctly on the first attempt and the rest of the exercises is much greater in the numbers and operations strand than in the rest of the strands. However, additionally, that effect is much greater than that of accuracy. Since increasing accuracy can require much more effort, the large effect of this difference suggests that in many cases, it may be worthwhile to increase the number of exercises. For example, if we have an accuracy of 0.7 with 10 exercises, then the difference is $7 - 3 = 4$. If we go up to 20 exercises while maintaining the accuracy, then the difference is $14 - 6 = 8$. Since B6 is more than three times B1, and since the variables are normalized, then the performance in terms of the difference between questions answered correctly in the first attempt and the number of those not answered correctly in the first attempt has three times the effect of the accuracy. That is, increasing one standard deviation in the difference has three times the effect of raising the accuracy by one standard deviation.

We show in Figures 9 and 10 only the case of the numbers strand, which is by far the strand with the highest percentage of exercises performed. The percentage is 65% of the exercises performed in the year.

Figure 9 shows the effect of a student who is one standard deviation above the mean in accuracy each month. In the SS model, we assume that in the future, the student will perform one standard deviation above the mean.

Figure 10 shows the effect of a student who in each month is one standard deviation above the mean in the difference between the number of correct answers in the first attempt and the number of correct answers in another attempt, but assuming that the accuracy is equal to the mean. This means that the student has performed many more exercises, but with accuracy equal to the average of all the students. For example, if the average accuracy is 0.7, then if they performed 7 good exercises on the first attempt and 3 correct exercises on other attempts, there is an accuracy of 0.7 and a difference of 4, whereas if they performed 70 good exercises on the first attempt and 30 correct ones on other attempts, then there is also an accuracy of 0.7 but the difference is 40.

In the SS model, we assume that in the future, the student will continue with a performance of one standard deviation above the mean.

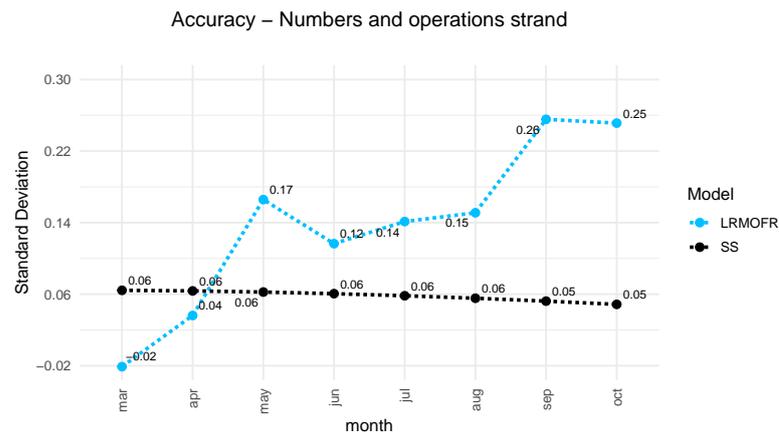


Figure 9. Effect of an accuracy one standard deviation above the mean on every month. This is the contribution to the SIMCE prediction of an increase of one standard deviation of the student accuracy.

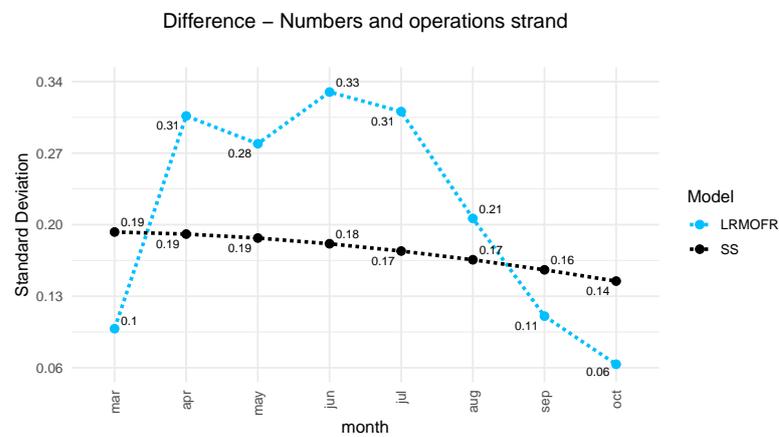


Figure 10. Effect of a difference one standard deviation above the mean on every month. This is the contribution to the SIMCE prediction of the difference between number of exercises solved correctly in the first attempt and the number solved correctly on other attempts.

5.4. The Optimal Control Problem: Ask Students to Do Exercises More Carefully or Do More Exercises

In each session, the teacher must decide how many exercises to place and what level of performance will be acceptable. If students make a lot of mistakes, then the teacher may need to stop the practice and explain the core concepts in more detail. The teacher can also assign student monitors to explain and give help to peers as well as asking the students to perform the exercises more carefully and seek help. There may be students responding very quickly and even randomly.

That is, the teacher must decide whether to give more importance to the percentage of correct answers in the first attempt or to the total number of exercises attempted by the students. For example, the teacher has to decide between getting the class to perform 30 exercises per student with 70% of them solved correctly at the first attempt, or to perform only 20 exercises but with 80% them solved correctly at the first attempt. The dilemma is: Which strategy causes better long-term learning? Or, more precisely, which strategy of exercise practicing causes a prediction of a highest score at the end-of-the-year national test?

Consider that we are in the month of τ , that is, with complete information on the students up to that month. For the following month s , the teacher is planning to carry out exercises of strand k . Therefore, if a student achieves in that month the performance given

by $\hat{u}_k(\tau, s)$ and $\hat{u}_{k+5}(\tau, s)$, then the predictor of the SS model year-end test result will add the effect $\Delta_{k,\tau,s}$

$$\Delta_{k,\tau,s} = A^{T-s-1}(B_k\hat{u}_k(\tau, s) + B_{k+5}\hat{u}_{k+5}(\tau, s))$$

In terms of the accuracy a , the number of exercises solved correctly at the first attempt N , and number of exercises solved correctly in other attempts M , the effect is

$$\Delta_\tau(N_{k,s}, M_{k,s}) = A^{T-s-1} \left[B_k \left(\frac{N_{k,s} - M_{k,s} - \mu_{k,\tau}^a}{\sigma_{k,\tau}^a} \right) + B_{k+5} \left(\frac{N_{k,s} - M_{k,s} - \mu_{k,\tau}^d}{\sigma_{k,\tau}^d} \right) \right]$$

with mean and deviation $\mu_{k,t}^a, \sigma_{k,t}^a$ for $a_{k,t}$, and $\mu_{k,t}^d, \sigma_{k,t}^d$ the mean and deviation for $d_{k,t}$, $\forall t \in \{1, \dots, T-1\}$.

Rewriting in terms of accuracy and total number of exercises L , the effect on the predictor of an accuracy a , and a number of exercises L to be performed in a month s after the current month τ is given by

$$\Delta_\tau(a_{k,s}, L_{k,s}) = A^{T-s-1} \left[B_k \left(\frac{a_{k,s} - \mu_{k,\tau}^a}{\sigma_{k,\tau}^a} \right) + B_{k+5} \left(\frac{L_{k,s}(2a_{k,s} - 1) - \mu_{k,\tau}^d}{\sigma_{k,\tau}^d} \right) \right]$$

Then for a level δ , the pairs of accuracies and number of exercises L for a strand k with the effect δ on month s posterior to τ , the last month with student data, is given by

$$a_{k,s} = \frac{\delta A^{s-T+1} + \frac{B_k \mu_{k,\tau}^a}{\sigma_{k,\tau}^a} + \frac{B_{k+5}(L_{k,s} + \mu_{k,\tau}^d)}{\sigma_{k,\tau}^d}}{\frac{B_k}{\sigma_{k,\tau}^a} + \frac{2L_{k,s} B_{k+5}}{\sigma_{k,\tau}^d}}$$

This expression describes a curve in the $a - L$ plane. Figure 11 shows the iso-effect curves as a function of the accuracy a , and the number of exercises L . Since A is basically 1, the iso-effect curves do not depend on the month s . However, they do depend on τ , through the means and standard deviations of the accuracy and the difference.

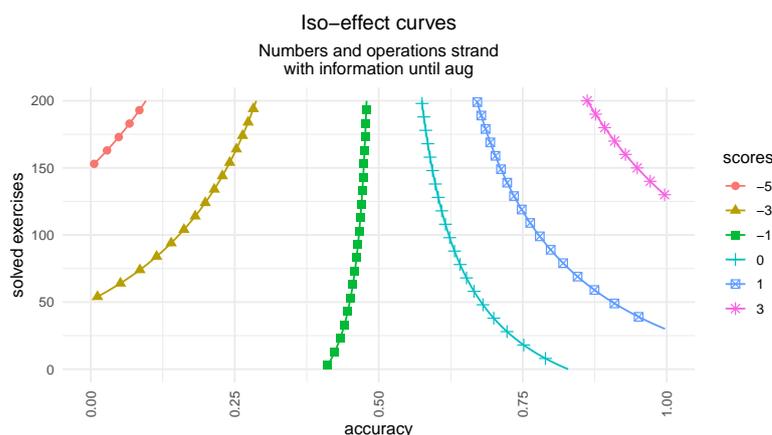


Figure 11. Iso-effect curves for the number and operations strand, in September, where the last month with student data is August.

The gradient of the effect with respect to a and L is

$$\nabla \Delta_\tau(a_{k,s}, L_{k,s}) = \begin{pmatrix} \frac{\partial \Delta_\tau}{\partial a}(a_{k,s}, L_{k,s}) \\ \frac{\partial \Delta_\tau}{\partial L}(a_{k,s}, L_{k,s}) \end{pmatrix} = \begin{pmatrix} A^{T-s-1} \left[\frac{B_k}{\sigma_{k,\tau}^a} + \frac{2L_{k,s} B_{k+5}}{\sigma_{k,\tau}^d} \right] \\ A^{T-s-1} \left[\frac{B_{k+5}(2a_{k,s} - 1)}{\sigma_{k,\tau}^d} \right] \end{pmatrix}$$

Figure 12 shows the vector field of the gradients superposed to the iso-effect curves for the number and operations strand for the month September when we have information up to August.

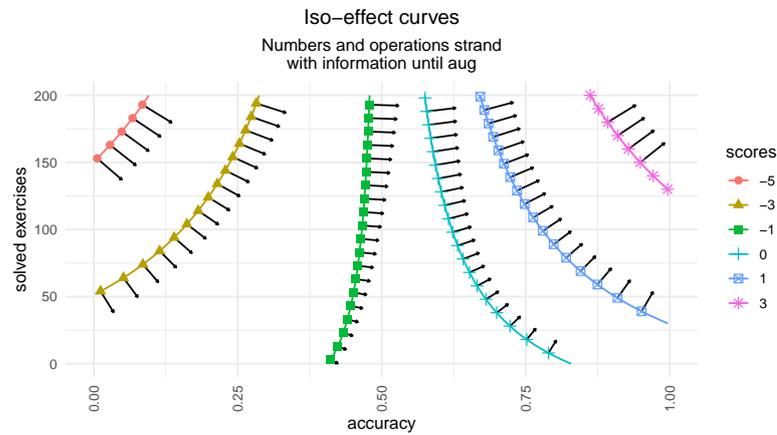


Figure 12. Effect gradients and iso-effect curves for the Number and Operations strand, in September, where the last month with student data is August.

From Figure 12, we see that if the accuracy is below 50%, then it is convenient to reduce the number of exercises to be done but to do them more carefully to improve the accuracy. For example, if we have been doing 130 exercises with an accuracy of 30%, then in order to increase the result of the year-end test, it is better to lower the number of exercises and improve the accuracy. This is due to the fact that in the coordinate $(0.30, 130)$ of the plane $a - L$, the gradient points to having higher accuracy a , but to do fewer exercises, L .

Let us consider that performing L exercises with an accuracy a means an average effort $e(a, L)$ for the students of the class. It is reasonable, within a range, that by doubling the number of exercises to be performed in a session, the effort doubles. However, the dependence of effort on accuracy is more complex. Let us assume the following dependency:

$$e = \frac{cL}{(1.1 - a)}$$

Then, for a given strand of the curriculum, the iso-effort curves are those of Figure 13. By considering the iso-effect curves or the vector field of gradients, for each strand of the curriculum, for a given level of effort, we can find the optimal combination of accuracy and number of exercises to perform in the session.

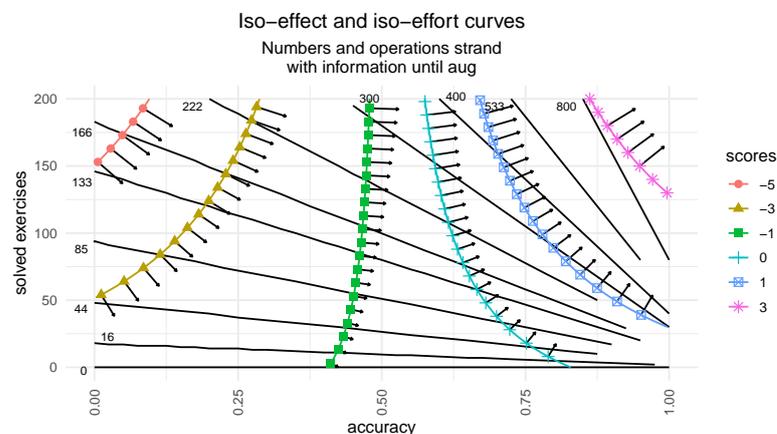


Figure 13. Iso-effort, effect gradients, and iso-effect curves for the number and operations strand in September, where the last month with student data is August.

In Figure 13, we can see that for an effort less than or equal to 166, in the numbers and operation strand, the best combination is to plan to perform 40 exercises and with an accuracy of 95%. At that point on the $a - L$ plane, the effect gradient is perpendicular to the iso-effort curve.

Figure 14 shows that in the data and probability strand, the gradient always points to lower accuracy. One possible explanation is that this may be because it is more effective to spend that effort on other strands.

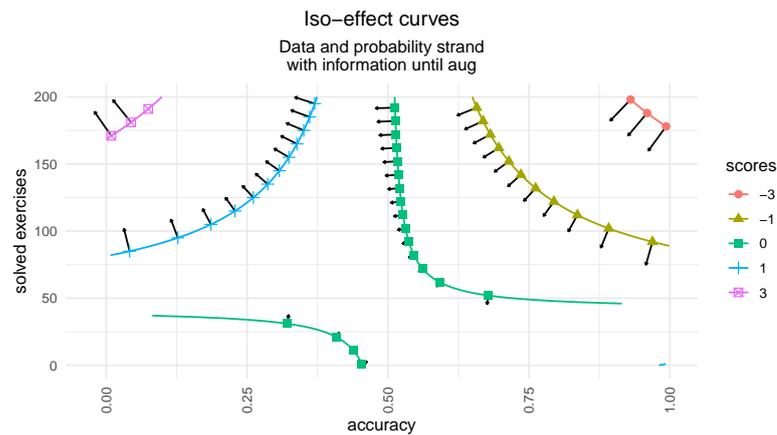


Figure 14. Iso-effort curves and effect gradients for the data and probability strand in September, where the last month with student data is August.

6. Discussion

Improving the quality of education depends on the quality of transmission of pedagogical practices. Like any cultural phenomenon, improvement depends on access to reliable performance indicators. Good indicators of teaching practices are those that achieve good long-term learning in students. It is therefore necessary for each teacher to be able to estimate and understand the effect of teaching practices on long-term learning. The teacher can then use these estimates to make teaching decisions for every lesson. Therefore, teachers need to be able to predict the effect of applying, in the remaining months, the different options of teaching strategies. In this paper, we consider for each month and curriculum strand two options: how many exercises and with what level of accuracy.

To solve this problem, teachers face several challenges.

First, given that they prefer to pose their own exercises, this generates a large number of questions, each one attempted by a small number of students. Thus, we could not use models based on big data, such as deep learning.

Second, the questions of the national end-of-year tests are unknown to the teachers and there is no public information available on the results of students in each of those questions. For this reason, predicting the results on those tests is much more complex than in other tests. Strategies to predict students' performance in a question based on the data of the performance of many similar students in similar questions, such as the algorithms used by matrix or tensor factorization or deep learning networks, cannot be directly applied.

Third, another difficulty with predicting long-term learning is that it is not the same as short-term learning. It is not only a phenomenon of decaying and forgetting. It is much more complex. According to [1], during teaching, what we can observe and measure is performance, which is often an unreliable index of long-term learning. There are strategies that lead to less learning than others in the short term, but generate more learning than the rest over in the long term.

With the information from the responses of 500 students from 24 courses who completed an average of 1089 closed-question exercises on the ConectaIdeas platform during the year 2017, we built two types of models: linear regressions and state-space models. We chose the best option of variables for each of them. We selected those variables analyzing

information from 18 courses and we calculated the prediction errors in the remaining 6 courses. We carried out this cross-validation process 200 times. With these results, we can now answer affirmatively the three research questions.

Research Question 1: To what extent can teacher-designed questions (low- or no-stake quizzes) help predict students' long-term learning as measured by end-of-year standardized state tests?

Despite having questions designed by teachers and each of them answered by few students, the linear regression and state-space models achieve predictions with errors comparable to those of the literature [7,29]. We found that with both types of models, we were able to improve the month-to-month predictions. Additionally, the state-space model is in general better than the linear regression model. This better performance is very important because it achieves it with much fewer parameters. The state-space model has only 17 parameters to fit, whereas the linear regression has 106 parameters. That is, the SS model has only 16% of the parameters of the linear regression model. This means that it must be more robust and generalizable to other schools and years.

Research Question 2: To what extent can a hidden Markovian state of the student's accumulated knowledge up to the current month be estimated so that, along with the probable deliberate practice on the next months, it is sufficient to predict with good accuracy the end-of-year standardized state tests?

The SS model is a simple accumulator that adds monthly performance. From month to month, the predictor adds the score for the new month to the score for previous months. Additionally, the mechanism of the SS model adds the score that it estimates for the following months according to the strategy that the teacher plans to follow. In the first month, the mechanism includes a base score, and the contribution of three sources of information: a pretest, the beliefs of each student, and estimation of the climate of the class. The predictor translates these three contributions into three scores, and adds them. Thus, the teacher has the initial prediction in view and, with the performance of each month, an additive combination to predict long-term learning.

Both models estimate the impact of the pretest and how its impact decreases as the year goes by. This is expected since, as the ConectaIdeas platform accumulates information on student performance, the initial information is less relevant. Likewise, they estimate the impact of students' beliefs about mathematics. This is very relevant for the teacher, as it gives a clue to determine the importance of this emotional factor and how much effort should be put into dealing with it. Both models also estimate the impact of the class climate on the prediction of each student's end-of-year test. Again, dimensioning this effect is very relevant for the teacher. It guides to quantify how much effort should be dedicated to improving the climate of the class. In both cases, these effects also decline over the course of the year. In the case of the state-space model, the three effects decay at the same rate. Both models also estimate the effect of the practice of exercises in each of the five strands of the curriculum. The teacher can thus compare the effect of performance on each strand on the prediction of each student's results on the end-of-year national test. This helps to decide how much intensity to devote to each strand of the curriculum.

A central element of the models is the information for each month. It has two performance elements: the accuracy, which is the percentage of questions answered correctly on the first try; and the difference between the number of questions answered correctly on the first attempt and the rest of the questions. They are very meaningful and very common elements in the daily practice of the teachers. The predictor mechanism just adds those contributions.

A particular characteristic of the SS model is that it includes the effect of strategies on future months. If we use the typical pattern of the student shown in previous months, the predictor uses that pattern and achieves a good prediction in the national end-of-year test. However, the teacher can also try different options for the following month.

Research Question 3: To what extent can a state-space model help teachers to visualize the trade-off between asking students to perform exercises more carefully or perform more exercises, and thus help drive whole classes to achieve long-term learning targets?

The state-space model provides a very transparent and simple prediction mechanism. First, it has very few parameters. They are just what is needed. As there are 5 strands in the Chilean mathematics curriculum, and in each one we have 2 performance metrics, this already brings together 10 parameters. They are parameters that translate both performances in the five strands into the corresponding contribution to the test score. The rest are that the parameters translate into the test score the contribution of the pretest, the beliefs about ease in mathematics, and the climate of the class. Finally, there is the rate of forgetting from one month to another.

Second, the whole mechanism is an additive one. The score of each contribution is added with the rest.

Third, the peculiarity of the state space model is that from one month to another, the mechanism updates the prediction in a Markovian way. It is enough to have the predictor of the last month, not the whole history of information. This is a huge advantage, as last month's predictor sums it all up. It is not necessary to go to gather information from previous months.

Fourth, for each month and for each strand, the mechanism is a tool that helps the teacher to see the effect of her decisions. She can visualize ahead the effect of choosing how many exercises to request and what minimal accuracy to request. In a 2D graph, the teacher can see the iso-impact curves of these two decisions.

Fifth, the teacher can also visualize the arrows of the gradients. They are an intuitive visualization that indicates how to move to achieve more impact in the national test at the end-of-the-year. This allows the teacher to design her optimal strategy for driving the entire class toward the goal of greatest long-term learning. It is not a black box recommending what to do. It is a tool that allows the teacher to simulate and visualize the effect of different strategies.

Sixth, the teacher can also superimpose the iso-effort curves in order to find an optimal solution. That is, she can seek the combination of the number of exercises and accuracy in order to achieve the greatest long-term learning while maintaining a limited and predefined level of effort. That is, not only the state-space model gives the teacher an understanding of the factors involved and how they contribute to the result of the end-of-year summative test. It also gives her a tool to simulate the effect of various practice strategies. It is a very natural and intuitive way to visualize the trade-off between asking students to perform exercises more carefully or perform more exercises. In summary, in this paper, we built a state-space model with control variables that allow the teacher to visualize an optimal control strategy in order to conduct sessions that achieve the best long-term learning.

Although there are applications of optimal control to student learning strategies [30–33], these are problems formulated with abstract situations. They do not include empirical data on classes with students, models that fit those data, and end-of-year national standardized tests that independently measure long-term learning. To the best of our knowledge, this is the first time that the problem faced by each teacher is presented as an optimal control problem, with whole year empirical data of hundreds of students in several classrooms, and where a very practical solution is proposed. Moreover, we developed a graphical tool, intuitive and easily understandable, that helps teachers visualize the effect of strategies. With this state-space formulation and with the iso-impact and iso-effort curves, at each session, the teacher can review what students have achieved in previous sessions, and decide how to drive the entire class from then on.

There are several aspects to investigate and develop in the future. One is to include an interaction between strands of the curriculum. In a previous study, we empirically identified the effect of some topics on others [16]. We plan to include this interaction in the state-space model. Another aspect is to incorporate the effect of written answers to open questions [21,25]. We already did this for a linear regression model. We plan to include

this information in the state-space model. There is also the effect of different strategies that the teacher can choose for deliberate practice. For example, the effect of spacing and interleaving [34,35]. Spacing is not doing all the exercises at once. For example, 6 h of deliberate practice in a type of exercise is not the same as doing it in spaces of 1 h for six different days. On the other hand, in interleaving, in each session, students solve a random mixture of different types of problems. Interleaving and spacing are cases of variability in practice. Experimental evidence has shown that variability generates more long-term learning and generalization, and in various domains and tasks [36]. Another important topic is to include the effect of conducting peer reviews of the written response to open-ended questions. Finally, there is the effect of peer tutors [24]. We plan to include these strategies in future state-space models.

7. Conclusions

This work connects with previous literature in several respects. On the one hand, it implements in several schools the strategy of systematic deliberate practice. This strategy has been studied mainly in research laboratories for the study of human cognition and learning [1–3]. Instead, in this paper, we study the typical situations found in schools. It is a situation where instructors choose or create their own questions to adjust better to the realities of their classes. Furthermore, we adhere to the tight restrictions of low SES schools. On the other hand, in this paper, we connect with the long-term learning prediction literature [7–9]. Moreover, we connect with the month-by-month predictions proposed by [7]. However, we do so with a state-space model that has the advantage of few parameters and greater interpretability. Additionally, we connect with the literature on optimal control in learning [30–33], but we do it contextualizing in elementary education and with empirical data from hundreds of students who answered a large number of questions during the year. We also do it by proposing a graph that gives the teacher a direct visualization of the impact of two basic strategies: for each strand of the curriculum, how many more exercises to include and what percentage of correct answers students should try to achieve in the following month.

7.1. Theoretical Implications

In the theoretical framework of the evolutionary educational psychology, the challenge of teaching school mathematics is due to an evolutionary mismatch or evolutionary trap [12]. The evolved cognitive systems and inferential biases that define the natural knowledge required for hunting–gathering life are not sufficient for academic learning [37]. Our brain is not adapted to learning these types of radically new concepts. It does not cost us very much to learn to walk or speak. We do it even without a teacher. However, learning to read and write provokes a huge change that greatly affects different brain areas [38]. Similarly, learning the mathematical concepts of the curriculum requires a lot of effort and deliberate and constant practice, day by day, for many years.

However, implementing a deliberate practice strategy is not easy. There are several challenges. On the one hand, the teacher has a huge motivational problem to solve. It is easy to say that you have to make students practice, but actually making all students do it systematically is not that simple. The teacher has to deal with the motivation and beliefs of each student, as well as with the climate of the entire class. The teacher must also take into account the prior knowledge of each student and the progress they are making. The teacher must continually decide between how much to exercise, what level of difficulty and what percentage of correct answers to achieve. There is a curriculum to cover and a national end-of-year test target. All students must perform well. This is an optimal control problem, where the teacher must guide dozens of autonomous agents, each one with their own strengths and motivations. The teacher has to lead a complex swarm of agents in order to reach a predefined target. Every student must achieve the optimal long-term learning, given their initial condition and the condition of the class.

To reach this goal, the teacher needs to be able to predict the effect of deliberate practice on the long-term learning of each student. However, the teacher also needs to have a simple-to-understand tool that allows to simulate deliberate practice alternatives for the next month and know the impact they will have on long-term learning.

In summary, we found that this type of system is feasible to implement. With the ConectaIdeas platform, students carry out intensive deliberate practices that allow the teacher to predict their long-term learning each month. In addition, the platform can provide the teacher with tools to guide on how to conduct the practice of the whole class in the next month.

7.2. Managerial Implications

The results found in this paper give some ideas to introduce management components in online platforms that support the teacher to lead the whole class to the learning target. We learned three practical ideas.

First, it is possible to build a module that each month tells the teacher the prediction of long-term learning for each student. We found that the prediction error is reasonable and decreases month by month.

Second, the module can indicate the weight of socio-emotional factors, such as students' beliefs about mathematics. It can also include the effect that the climate of the class can have for the prediction of the long-term learning of each student. This will allow the teacher to take personalized strategies to handle socio-affective states of each student. It also supports the teacher to design strategies directed to improve the learning environment of the whole class. The teacher will know how important these factors are for her course and then decide on a plan of action. She will be able to compare the effect of these variables with respect to others, such as initial performance and the effect of practice.

Third, the module can help the teacher visualize how to plan a deliberate practice schedule for the coming month. It will help her to determine, for each student and each strand of the curriculum, the effect of planning a given number of exercises and trying to achieve a percentage of correct answers. These are basic indicators for action in the following month.

Author Contributions: Conceptualization, R.A.; methodology, R.A.; software, O.U.; validation, R.A. and O.U.; formal analysis, R.A. and O.U.; investigation, R.A. and O.U.; resources, R.A.; data curation, R.A. and O.U.; writing—original draft preparation, R.A.; writing—review and editing, R.A.; visualization, R.A. and O.U.; supervision, R.A.; project administration, R.A.; funding acquisition, R.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by ANID/PIA/Basal Funds for Centers of Excellence FB0003 and by UCE, Ministry of Education of Chile.

Data Availability Statement: Not applicable.

Acknowledgments: Support from ANID/PIA/Basal Funds for Centers of Excellence FB0003 and UCE, Ministry of Education of Chile, is gratefully acknowledged. We also thank the International Development Research Center (IDRC) and the Interamerican Development Bank (IDB) for supporting the RCT project where the data were obtained, Julián Cristia from IDB, Abelino Jiménez and Raúl Gormaz from Universidad de Chile for many suggestions for analysis.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

GPA	Grade Point Average
LRMOFR	Linear Regression for the Month with an Optimal Forgetting Rate
OLS	Ordinary Least Squares
RMSE	Root Mean Squared Error
SD	Standard Deviation
SIMCE	Sistema de Medición de la Calidad de la Educación (Measurement System of the Quality of Education)
SS	State Space

References

- Soderstrom, N.; Bjork, R. Learning Versus Performance: An Integrative Review. *Perspect. Psychol. Sci.* **2015**, *10*, 176–199. <https://doi.org/10.1177/1745691615569000>.
- Bjork, R.A.; Bjork, E.L. Desirable difficulties in theory and practice. *J. Appl. Res. Mem. Cogn.* **2020**, *9*, 475–479. <https://doi.org/10.1016/j.jarmac.2020.09.003>.
- Bjork, E.; Bjork, R. Making Things Hard on Yourself, But in a Good Way: Creating Desirable Difficulties to Enhance Learning. In *FABBS Foundation, Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*; Worth Publishers: London, UK, 2016; pp. 56–64.
- Gervet, T.; Koedinger, K.; Schneider, J.; Mitchell, T. When is Deep Learning the Best Approach to Knowledge Tracing? *J. Educ. Data Min.* **2020**, *12*, 31–54. <https://doi.org/10.5281/zenodo.4143614>.
- Zytek, A.; Arnaldo, I.; Liu, D.; Berti-Equille, K.; Veeramachaneni, K. The Need for Interpretable Features: Motivation and Taxonomy. *arXiv* **2022**, arXiv:2202.11748.
- Bagaloplan, A.; Zhang, H.; Hamidieh, K.; Hartvigsen, T.; Rudzicz, F.; Ghassemi, M. The Road to Explainability is Paved with Bias: Measuring the Fairness of Explanations. In Proceedings of the FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Korea, 21–24 June 2022. <https://doi.org/10.1145/3531146.3533179>.
- Anozie, N.; Junker, B. Predicting end-of-year accountability assessment scores from monthly student records in an online tutoring system. In Proceedings of the AAAI Workshop on Educational Data Mining, Boston, MA, USA, 16–17 July 2006.
- Zheng, G.; Fancsali, S.; Ritter, S.; Berman, S. Using Instruction-Embedded Formative Assessment to Predict State Summative Test Scores and Achievement Levels in Mathematics. *J. Learn. Anal.* **2019**, *6*, 153–174. <https://doi.org/10.18608/jla.2019.62.11>.
- Namoun, A.; Alshantiti, A. Predicting Student Performance Using Data Mining and Learning Analytics Techniques: A Systematic Literature Review. *Appl. Sci.* **2021**, *11*, 237. <https://doi.org/10.3390/app11010237>.
- Ruipérez-Valiente, J.A.; Muñoz-Merino, P.J.; Delgado, C. Improving the prediction of learning outcomes in educational platforms including higher level interaction indicators. *Expert Syst.* **2018**, *35*, e12298. <https://doi.org/10.1111/exsy.12298>.
- Heckman, J.; Zhou, J. *Interactions as Investments: The Microdynamics and Measurement of Early Childhood Learning*; Working paper; Center for the Economics of Human Development and Department of Economics, University of Chicago: Chicago, IL, USA, 2021.
- Lloyd, E.; Wilson, D.; Sober, E. Evolutionary Mismatch and What to Do About It: A Basic Tutorial. *Evol. Appl.* **2011**, *2–4*.
- Geary, D. Educating the evolved mind: Conceptual foundations for an evolutionary educational psychology. In *Psychological Perspectives on Contemporary Educational Issues*; Carlson, J.S., Levin, J.R., Eds.; Information Age Publishing: Charlotte, NC, USA, 2007.
- Gezer, T.; Wang, C.; Polly, A.; Martind, C.; Pugalee, D.; Lambert, R. The Relationship between Formative Assessment and Summative Assessment in Primary Grade Mathematics Classrooms. *International Electronic J. Elem. Educ.* **2021**, *13*, 673–685.
- Yang, C.; Luo, L.; Vadillo, M.; Yu, R.; Shanks, D. Testing (Quizzing) Boosts Classroom Learning: A Systematic and Meta-Analytic Review. *Psychol. Bull.* **2020**, *147*, 399–435. <https://doi.org/10.1037/bul0000309>.
- Araya, R.; Van der Molen, J. Causal Dependence among Contents Emerges from the Collective Online Learning of Students, In Proceedings of the International Conference on Computational Collective Intelligence, Craiova, Romania, 11–13 September 2013; Springer: Berlin, Germany, 2013; Volume 8083, pp. 641–650. https://doi.org/10.1007/978-3-642-40495-5_64.
- Araya, R.; Van der Molen, J. Impact of a blended ICT adoption model on Chilean vulnerable schools correlates with amount of on online practice. In Proceedings of the Workshops at the 16th International Conference on Artificial Intelligence in Education AIED 2013, Memphis, TN, USA, 9–13 July 2013.
- Araya, R.; Gormaz, R.; Bahamondez, M.; Aguirre, C.; Calfucura, P.; Jaure, P.; Laborda, C. ICT supported learning rises math achievement in low socio economic status schools. In Proceedings of the 10th European Conference on Technology Enhanced Learning, Toledo, Spain, 15–18 September 2015; Springer: Berlin, Germany, 2015; Lecture Notes in Computer Science, Volume 9307, pp. 383–388. https://doi.org/10.1007/978-3-319-24258-3_28.
- Araya, R. Integrating Classes from Different Schools Using Intelligent Teacher Support Systems. In Proceedings of the 1st International Conference on Intelligent Human Systems Integration (IHSI 2018), Dubai, United Arab Emirates, 7–9 January 2018; Springer: Cham, Switzerland, 2018; Volume 722, pp. 294–300. https://doi.org/10.1007/978-3-319-73888-8_46.
- Araya, R. Teacher Training, Mentoring or Performance Support Systems? In Proceedings of the International Conference on Applied Human Factors and Ergonomics 2018, Orlando, FL, USA, 21–25 July 2018; Springer: Cham, Switzerland, 2019; Volume 785, pp. 306–315. https://doi.org/10.1007/978-3-319-93882-0_30.

21. Araya, R.; Diaz, K. Implementing Government Elementary Math Exercises Online: Positive Effects Found in RCT under Social Turmoil in Chile. *Educ. Sci.* **2020**, *10*, 244. <https://doi.org/10.3390/educsci10090244>.
22. Heckman, J.J. Skill formation and the economics of investing in disadvantaged children. *Science* **2006**, *312*, 1900–1902. <https://doi.org/10.1126/science.1128898>.
23. Rea, D.; Burton, T. Does an Empirical Heckman Curve Exist? 2018. Available online: https://www.wgtn.ac.nz/__data/assets/pdf_file/0005/1716953/WP18-03-does-an-empirical-Heckman-curve-exist.pdf (accessed on 1 August 2022).
24. Araya, R.; Gormaz, R. Revealed Preferences of Fourth Graders When Requesting Face-to-Face Help While Doing Math Exercises Online. *Educ. Sci.* **2021**, *11*, 429. <https://doi.org/10.3390/educsci11080429>.
25. Urrutia, F.; Araya, R. Do written responses to open-ended questions on fourth-grade online formative assessments in mathematics help predict scores on end-of-year standardized tests? 2022, *submitted*.
26. Araya, R.; Arias Ortiz, E.; Botta, N.; Cristia, J. Does Gamification in Education Work? Experimental Evidence from Chile. In *IDB Working Paper Series N° IDB-WP-982 Inter-American Development Bank*; IDB: Washington, DC, USA, 2019. <https://doi.org/10.18235/0001777>.
27. Van der Molen, J. *Minería de Datos Educativas: Modelos de Predicción del Desempeño Escolar en Alumnos de Enseñanza Básica*; Universidad de Chile: Santiago, Chile, 2013.
28. Ulloa, O. Estimación de Desempeño en Evaluación Sumativa, Con Base en Evaluaciones Formativas Usando Modelos Espacio Estado. Master's Thesis, University of Chile, Santiago, Chile, 2021.
29. Pardos, Z.; Baker, R.; San Pedro, M.; Gowda, S.; Gowda, S. Affective States and State Tests: Investigating How Affect and Engagement during the School Year Predict End-of-Year Learning Outcomes. *J. Learn. Anal.* **2014**, *1*, 107–128. <https://doi.org/10.18608/jla.2014.11.6>.
30. Chen, Y.H. A revisit to the student learning problem. *Optim. Control. Appl. Methods* **2016**, *12*, 263–272.
31. Lewis, D. Modeling student engagement using optimal control theory. *J. Geom. Mech.* **2022**, *14*, 131–150.
32. Raggett, G.F.; Hempson, P.W.; Jukes, K.A. A Student-Related Optimal Control Problem. *Bull. Inst. Math. Its Appl.* **1981**, *17*, 133–136.
33. Teklu, S.W.; Terefe, B.B. Mathematical modeling analysis on the dynamics of university students animosity towards mathematics with optimal control theory. *Sci. Rep.* **2022**, *12*, 11578. <https://doi.org/10.1038/s41598-022-15376-3>.
34. Rohrer, D.; Dedrick, R.; Stershic, S. Interleaved Practice Improves Mathematics Learning. *J. Educ. Psychol.* **2015**, *107*, 900–908. <https://doi.org/10.1037/edu0000001>.
35. Rohrer, D.; Dedrick, R.; Hartwig, M. The Scarcity of Interleaved Practice in Mathematics Textbooks. *Educ. Psychol. Rev.* **2020**, *32*, 873–883. <https://doi.org/10.1007/s10648-020-09516-2>.
36. Raviv, L.; Lupyan, G.; Green, S.C. How variability shapes learning and generalization. *Trends Cogn. Sci.* **2022**, *26*, 462–483. <https://doi.org/10.1016/j.tics.2022.03.007>.
37. Geary, D.C. Principles of evolutionary educational psychology. *Learn. Individ. Differ. Evol. Educ. Psychol.* **2002**, *12*, 317–345. [https://doi.org/10.1016/S1041-6080\(02\)00046-8](https://doi.org/10.1016/S1041-6080(02)00046-8).
38. Henrich, J. *The Weirdest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous*; Farrar, Straus and Giroux: New York, NY, USA, 2020.