

Architectural Inefficiencies and Educational Outcomes in STEM

Dan Sturtevant

Massachusetts Institute of Technology

Engineering Systems Division

dan.sturtevant@sloan.mit.edu

Abstract

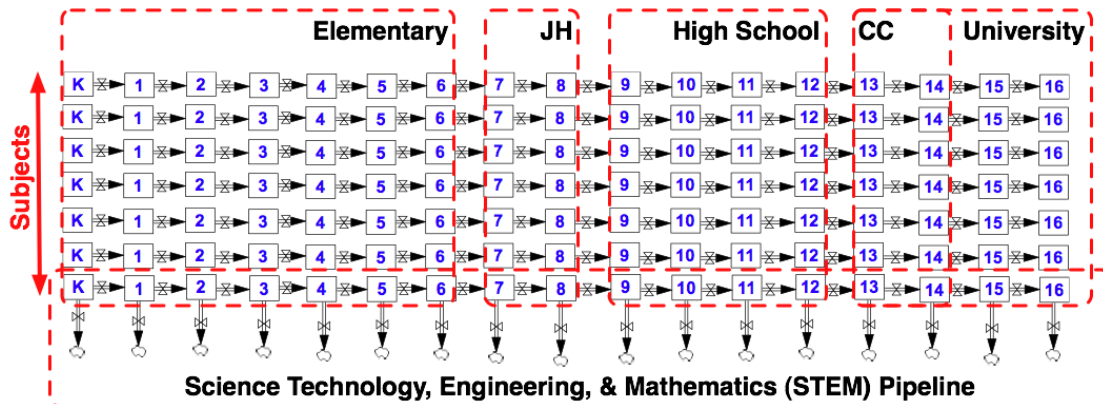
The modern education system is structured as a cellular ‘value stream’ in which students acquire new skills as they consume discrete courses and move incrementally through successive grade on a yearly clock. Within each classroom, these students are taught at a common regulated pace that is set by teachers based on the needs of the entire class. Inherent in common pace setting is an inescapable tradeoff between the amount of knowledge gained by those students who keep up and the number of students who fall behind. In ‘Disrupting Class: How Disruptive Innovation Will Change the Way the World Learns,’ (1) Christensen, Horn, and Johnson suggest that the ‘value stream’ model of education in which students move through cellular classrooms could move towards a more dynamic ‘network centric’ model within the next ten years. Under this new institutional regime, students could work at individualized paces and digest online material. Focus could be on mastery of material regardless of schedule. The teacher role would be that of a mentor and coach rather than the sole imparters of knowledge. This work presents a computational model designed to explore unavoidable losses inherent in the current ‘value stream’ structure under a variety of assumptions and education policy choices. System performance in these scenarios is compared against an ideal in which a network enabled paradigm eliminates the cellular pacing requirement. It is shown that seemingly benign structural rules embedded in the current system may have tremendous impact on level and character of STEM educational attainment across the population and that easing the pacing requirement could lead to improvements in educational attainment at all levels.

Introduction

While it is clear that the U.S. education system does many things well, it is currently thought to be failing in its mission to train sufficient numbers of future scientists, technologists, engineers, and mathematicians (STEM). (2) For instance, while the percentage of college graduates earning degrees in engineering approached 8% in 1985, today only 4.5% do so. (3) Furthermore, STEM disciplines are some of the least racially and ethnically diverse. For some reason, the U.S. education system seems incapable of helping people overcome initial socio-economic disadvantages in a way that allows them to eventually succeed in STEM disciplines. Many of those outside of STEM fields must also have quantitative and analytical competencies as well. Unfortunately, a significant fraction of citizens that graduate from high-school have attained only a 5th grade level of mathematical proficiency. Most people thinking about reasons for society’s failure to meet educational goals focus on individual teacher and student attributes such as ability, motivation, and socio-economic status. While these factors are surely important, the structure of the education system and the rules by which it operates may also play a key role. Identically endowed students might experience very different outcomes under a different set of institutional constraints.

Like any complex system, an education system has an *architecture*. An architecture is the underlying structure and set of relationships that (often tacitly) guide and constrain human action. A good architecture will allow a system to perform a primary function very effectively (in this case knowledge transfer to students) while also satisfying as many other stakeholder needs as possible. Under the current concept, education ‘happens’ when a teacher instructs a group of students within a class and then evaluates their work to ensure that knowledge was successfully transmitted. Over the course of an educational career, these classes join to form a ‘value stream’ in which students gain capabilities as they move through successive grades. The modern education system, which coalesced last century based on the principles of ‘Scientific Management,’ is decomposed into discrete subjects and grades-levels that determine which topics are to be covered in each classroom. Within the classroom, students are required to learn the same content as those around them at a regulated pace. The teaching profession has taken on a very cellular character that in many ways mirrors the segmentation of grades, subjects and classrooms. Because coordination costs are high and the workforce has a history of high turnover that negates investment in

cooperation, each teacher is generally expected to act as an autonomous modular unit and is treated as an interchangeable part (Weisberg, et al. 2009) (Lortie, 1975)



It is possible that the current problems afflicting U.S. STEM performance result in part from a misalignment between the current educational architecture and the demands of the ‘STEM pipeline’ that extends from Kindergarten through University. This is because the process of acquiring mathematics intensive knowledge may be disproportionately harmed by the current institutional design. Learning mathematics can be characterized as the serial acquisition of tightly interlinked knowledge with strong dependencies on prior work. (2) Failure to master arithmetic makes algebra impossible. Failure to master algebra renders calculus and probability unapproachable. For this reason, students who lose proficiency or interest at any point are almost never capable of recovering later in life. The education system can be thought of as a “leaky pipe” in which students who lose interest, proficiency, or confidence in STEM at any stage rarely get it back. Unlike the decision to pursue careers in law or medicine, which can be made during college, decision points that determine whether an individual can ever become a scientist or engineer occur as early as elementary school. Overall system performance may be determined by the weakest link in the value chain.

Under these circumstances, untended interfaces and expectation gaps within and between schools can leave students without prerequisite knowledge needed to continue and succeed. Furthermore, the demand that each class proceed at a regulated pace creates natural inefficiencies. Each teacher must set a start-point and a pace based on the starting-knowledge and abilities of the entering students. This demands an inescapable tradeoff between the amount of knowledge attained by the students who keep pace, and the number of students who fall behind. If a subject is taught more quickly, then more students fall behind. If a subject is taught more slowly, then more in the class do not reach their full learning potential. As variation increases, the system probably becomes less efficient and harder to manage.

In *‘Disrupting Class: How Disruptive Innovation Will Change the Way the World Learns,’* Christensen, Horn, and Johnson put forth the idea that computer based learning could play a key role in transforming the basic character of the education system within the next ten years. Their fundamental transformation is a change in the role of the teacher, from sole imparter of knowledge to mentor in the art of acquiring knowledge using internet enabled sources. (1) This change in the way value is delivered would enable the current organizational design, based on the ‘value stream,’ to give way to a ‘network centric’ paradigm. By decoupling the teacher somewhat from the curriculum and offloading content delivery to dynamically responsive internet sources, two important benefits accrue. First, the requirement of a standardized pace disappears. Secondly, the paradigm could allow students to fill in missing prerequisite knowledge in a dynamic “just in time” fashion. In the ideal, this ‘network centric’ organizational design would allow each student to perform as if (s)he were taught by a personal tutor who knew everything and could adapt to any learning style.

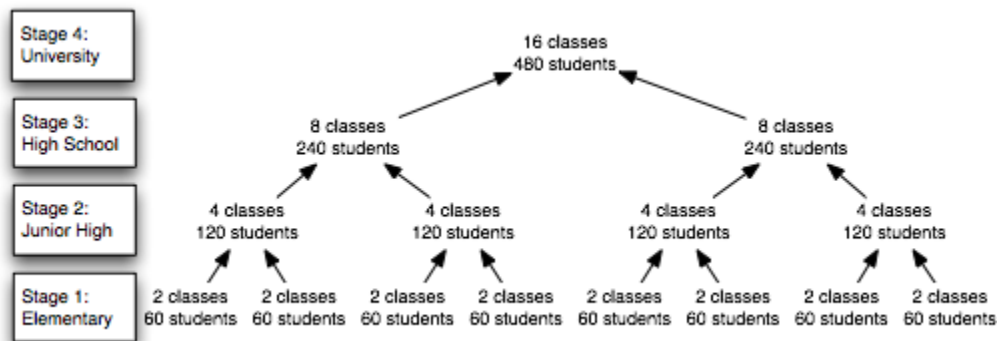
Simulation Modeling

Computer programs have been written to simulate a variety of social systems. (4) (5) (6) Such programs are models that embody explicit theories about how the world is structured. If a model can adequately capture the essence of some interesting phenomena, then it may lead to insights harder to discern by looking directly at the real world in all its complexity. Exploring the behavior of these simple models under a variety of conditions can help one build insight about the way a social system behaves and what policies may be effective when trying to improve its performance. Models can function as ‘flight simulators’ in which people can learn about the world by testing ideas in a risk free environment prior to costly implementation.

A simple simulation model was created to explore potential sources of loss in the current education system. Such sources include variance in prerequisite knowledge among students within a classroom, the effects of transitioning between different institutions (such as when multiple junior high schools feed a single high school), differences in individual student abilities, differences in teacher quality, and penalties imposed by the common pacing requirement. The impacts of policies related to the ‘tracking’ of students within a school based on perceived ability were explored. The impacts of policies for setting learning expectations and standards were also explored. Overall educational attainment under a variety of scenarios with different policies and constraints are compared and contrasted. The final simulation removes constraints associated with the common pacing requirement in an attempt to explore potential gains associated with a transition from a cellular to a network-centric learning paradigm.

Model Structure

The model simulated the mathematics attainment of 480 ‘students’ moving through sixteen years of education. These sixteen years are subdivided into four schools. For model simplicity, each student spends four years in each. The 480 students are initially separated into groups of 60 and allocated to 8 different elementary schools. At four year increments, students experience successive merges and eventually all meet in the same university. At the start of every year, students within a school are assigned to a new classroom containing one teacher and thirty pupils. At the elementary level, each school contains two classrooms per grade, while the university contains 16 classrooms per grade.



Children begin the simulation with a random amount of initial knowledge uniformly distributed between the values 0 and 1. Each year of schooling is intended to impart 1 unit of knowledge. Therefore, after 16 years, one may hope that each student will contain 16 additional units of knowledge.

Because mathematics knowledge is cumulative, knowledge must be gained in a serial fashion. Teachers must work with the knowledge level of the students in their classroom and attempt to advance them along the continuum. Each teacher must choose a starting point (SP) along this continuum and attempt to move students one unit beyond the SP by the end of the year. Multiple rules for setting this SP are explored within the simulations described below. Some rely on the perceived ability of students in the class while others rely on common standards.

Rules for setting Start Point (SP):

<i>Fixed</i>	$SP = gradeLevel - 1$, (initially 0) Common standards across the system dictate that students make 1 unit of progress each year and that we start at a point below the incoming abilities of all new students.
<i>Floating</i>	$SP = mean(students\ in\ class)$ Teacher chooses the mean student as the point at which instruction begins.
<i>HalfFixedFloat</i>	$SP = \frac{1}{2}(gradeLevel - 1) + \frac{1}{2}(mean(students\ in\ class))$ Teacher feels pressure to account for both factors.

Not all teachers and students are endowed with equal abilities. Some simulations test the impact of variation in teacher quality and student ability. Student ability (SA) is a modifier that affects each student's ability to learn every year. This value is different from the student's total amount of knowledge (SK) acquired. It is a modifier to the student's rate of gaining knowledge.

Rules for setting variation in student ability and teacher quality:

<i>NoIndividualVariation</i> TQ = teacher quality. SA = student ability.	TQ = 1 for all teachers. SA = 1 for all students.
<i>TeacherQualityVariation</i>	TQ is a random Normal variable with mean 1 and standard deviation 0.1. SA = 1 for all students.
<i>StudentAbilityVariation</i>	TQ = 1 for all teachers. SA is a random Normal variable with mean 1 and standard deviation 0.1.
<i>TeacherAndStudentVariation</i>	Both TQ and SA are random Normal variables with mean 1 and standard deviation 0.1.

In every year a student must attend a classroom in the school they belong to. Different policies for assigning students to individual classrooms within a school are tested.

Rules for assigning students to classrooms:

<i>RandomAssignment</i>	Every year, students within a school at the same grade level are randomly assigned to classrooms within that school.
<i>TrackingAssignment</i>	Every year, students are sorted by their total knowledge (TK) and then assigned to classrooms with other similar students. Note that this assignment is <i>not</i> based on student ability (SA). This assignment scheme serves to reduce variability of prerequisite knowledge within each classroom.

A penalty may be imposed upon student performance in any year based on the difference between a student's prerequisite knowledge and the SP chosen by the teacher. A student with knowledge equal to the SP will incur no penalty. Students with knowledge greater than the SP will not gain a full unit of knowledge because they will not be required to. These students will fall closer to the class mean. Students with less knowledge than the SP will not gain a full unit of knowledge because they will be harmed by missing prerequisite knowledge. These students will fall further behind their class.

Rules for imposing learning penalty to student based on pacing requirement:

<i>PacingPenaltyOn</i> CSP = class start point. TK = student's total knowledge at the beginning of the year.	$PacingPenaltyModifier = \frac{e^{-\frac{1}{abs(CSP - TK) + 1}}}{e^{-1}}$ <p>This function peaks at 1 when the student is at the start point and decays exponentially as the distance increases. The function also has the property that $TK + PacingPenaltyModifier$ is monotonically increasing. This function was chosen because of some of its useful properties.</p>
<i>PacingPenaltyOff</i>	PacingPenaltyModifier = 1.

It should be noted that the pacing penalty rule conspires to harm learning of students more if variation in student total knowledge (TK) is higher. A teacher will be less effective and students will make less yearly progress if the teacher must contend with widely varying prerequisite knowledge among the classmates.

Finally, these rules interact to produce the main effect we hope to study in the model. Yearly progress of each student every year is calculated as:

$$StudentYearlyProgress = TeacherQuality * StudentAbility * PacingPenaltyModifier$$

Applying these rules every year over all 480 students produces distributions of the final total level of mathematics knowledge within our agent population.


Simulation Results

Simulation 1:


*SP = Floating
NoIndividualVariation
RandomAssignment
PacingPenaltyOn*

In this scenario, there is no individual variation in abilities and students are assigned to classrooms randomly. Only small random variation exists in initial conditions. Teachers set the yearly start point for classes based on the mean starting knowledge of their class. The pacing penalty is imposed.

The vertical axis represents knowledge units gained by students in every year of schooling. Individuals are arranged along the horizontal axis and sorted by their final TK. Note that most agents in this simulation approach 13 knowledge units, but a minority, perhaps 20%, fall off substantially. None reach their full potential of around 16.5. These identically endowed agents experience much different outcomes due entirely due to small random variations in starting position and class composition that are magnified over time.



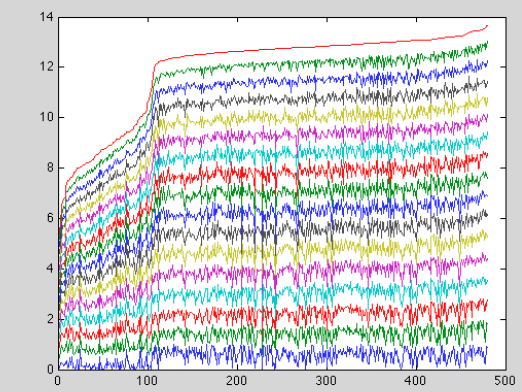
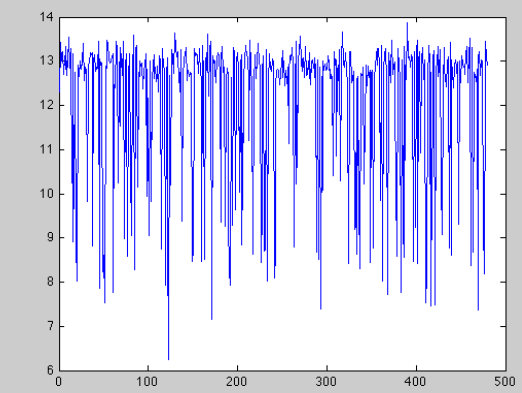
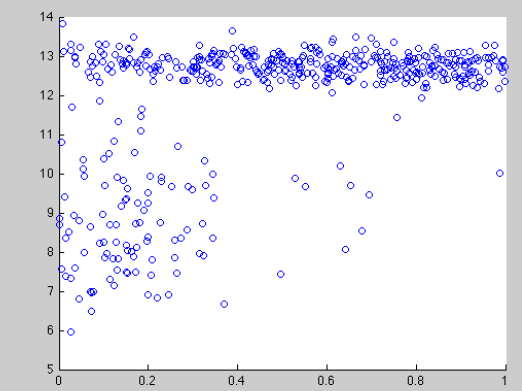
This plot shows the relationship between starting position (X axis) and final outcome (Y axis). Note that students with slightly poorer initial conditions experience substantial life-long penalties and that initial conditions are highly predictive of future success or failure in this scenario. This minority of students serves to pull down the mean, causing teachers to set expectations below the majority.




Simulation 2:

SP = Floating
TeacherAndStudentVariation
RandomAssignment
PacingPenaltyOn

In scenario 2, the only change is to add variability to both student abilities and teacher quality.

<p>This plot shows a picture similar to results from simulation 1.</p>	
<p>The final positions of each of the 480 individuals. Each individual is arranged along the X axis by their starting position in year 1. Classmates in year 1 are grouped in increments of 30. Note that some classes (shown towards the middle) have no failing students. This is represented by white bands that reach the top of the curve. This most likely occurred because these early classrooms experienced low variance and/or high average starting TK by pure chance.</p>	
<p>Initial starting TK still affects final TK, but is no longer a perfect predictor. Some individuals with excellent starting positions still fell behind due to peculiarities along their path through the pipeline.</p>	

This plot shows individual student ability (X axis) versus final TK (Y axis). Note the fact that many students with below average abilities still manage to end up in the dark band in the proficient region while others with stronger abilities end well below.

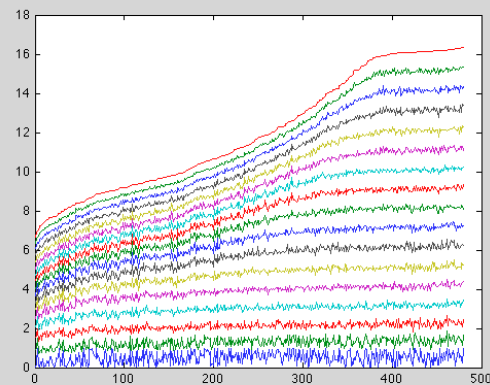


Simulation 3:


*SP = Fixed
TeacherAndStudentVariation
RandomAssignment
PacingPenaltyOn*

Simulation 3 is identical to simulation 2 with the exception that teachers now follow mandated curriculum without any regard for the composition of the class. These teachers start each year at grade level with the expectation that the class is ready and will make 1 unit of progress.


Note that the performance among the most proficient is higher (at around 16 knowledge units) than in previous simulations. The performance of students in the middle of the distribution is much lower however.



Initial starting conditions now have very little impact on final outcome.



Individual student ability is now a very strong predictor of final total knowledge.




Simulation 4:


*SP = HalfFixedHalfFloat
TeacherAndStudentVariation
TrackedAssignment
PacingPenaltyOn*

In simulation four, teachers set their starting position at a compromise point between the demands of class composition and globally set standards. In addition, students are now tracked into different classrooms within a school by their total knowledge at the start of each grade.


Tracking appears to have added some discrete jumps along TK curves.




Note that starting position has no great impact on outcomes.



Student abilities are strongly related to outcomes.



This graph shows only final TK and is arranged by original position. Note also that the composition of classes in early years can also have an impact on ultimate outcome in some cases. A few original classes contain no individuals with less than 12 knowledge units while some have many. These differences are likely due to the variance of student starting knowledge within those early classrooms.




Simulation 5:


*SP = Float
TeacherAndStudentVariation
TrackedAssignment
PacingPenaltyOn*

The following simulation is identical to simulation 4 with the exception that teachers now set expectations entirely by composition of their class without regard to global standards.


Peak performance is much higher than in previous simulations, approaching 18 knowledge units. Drop-off is more gradual. Discontinuous steps emerge between ability levels among different classes.




Initial individual starting conditions have little bearing on final outcomes.



Individual abilities have a stronger impact on final outcomes. Note also that the curvature of this plot appears to be concave down.



Note that initial class and school conditions have a very strong impact on final outcomes. Membership in a particular elementary school (or even junior high) strongly impacts final position.




Simulation 6:


*TeacherAndStudentVariation
RandomAssignment
PacingPenaltyOff*

The final simulation explores the impact of removing the pacing penalty altogether. The purpose of doing so is to provide a theoretical upper limit for comparison with other scenarios.


The outcome of this simulation shows variation resulting only from differences in student ability and teacher quality. Student abilities play the dominant role. Architectural constraints have been eliminated.



Outcomes are largely determined by student ability. Variation is introduced by randomness in student starting conditions and teacher quality.



The following plot shows the distributions of final total knowledge for each of the six scenarios:



The rules explored resulted in very different outcomes at a societal level. Some demonstrate a tradeoff between higher attainment at the upper end and in the middle of the distribution.

Discussion

Much of education research focuses on the relationship between individual attributes (such as student ability, motivation, teacher quality, and socioeconomic status) on educational outcomes. Model results presented here suggest that the architecture of the education system and seemingly benign policy choices made within that institutional regime could play a large role in determining the distribution of social outcomes as well. Increased focus on institutional design and the way structure leads to behavior might provide insights of significant value.

The results presented in these simulations should not be interpreted to endorse any particular policy in the real world. Many important factors were not represented in this simple model, and those that were represented were tuned to maximize conceptual clarity rather than real-world fidelity. It should be noted, however, that the distributions shown of total knowledge at intermediate and final points are qualitatively similar to those observed in the real world. National NAEP test score distributions indicate that students at the 10th, 25th, and 50th percentiles in mathematics ability in the eighth grade have abilities comparable to fourth grade students at the 50th, 75th, and 90th percentiles respectively. (7) This very wide disparity most likely gets even wider between eighth and twelfth grade.

This modeling work suggests that the common pacing requirement within the cellular ‘value chain’ system may be leading to significant losses relative to an ideal scenario in which students are pushed at individually maximal rates. Further work using more realistic assumptions should be done to explore the potential economic, social, and human-capital value that could accrue from a technology enabled ‘network centric’ transformation focused on individualization of performance expectations, emphasis on mastery independent of schedule, and the dynamic delivery of prerequisite information in a ‘just-in-time’ fashion.

Finally, simulation modeling could be productively employed within the education domain. Qualitative education literature contains descriptions of a host of interesting problems in which disequilibrium conditions, non-linear dynamics, balancing and reinforcing feedback loops, momentum and delays, and path-dependence play important roles. Systems characterized by such complex causality are often not well understood or analyzed using traditional statistical or econometric methods. Using system dynamics (6), agent based (4), or network modeling (8) methodologies to represent structure and reproduce behavior described in the education literature could help bring clarity to some of education’s thornier issues.

References

1. **Christensen, Clayton, Horn, Michael B. and Johnson, Curtis W.** *Disrupting Class: How Disruptive Innovation Will Change The Way The World Learns.* New York : McGraw Hill, 2008.
2. **NAS.** *Rising Above The Gathering Storm: Energizing and Employing America for a Brighter Economic Future.* Washington, D.C. : The National Academies Press, 2007.
3. **National Science Foundation WebCASPAR System.** WebCASPAR Integrated Science and Engineering Resources Data System. *National Science Foundation Web site.* [Online] <http://webcaspar.nsf.gov/TableBuilder>.
4. **Epstein, Joshua M.** *Generative Social Science: Studies in Agent Based Computational Modeling.* Princeton : Princeton University Press, 2007.
5. **Forrester, Jay W.** *Principles of Systems.* Waltham, MA : Pegasus Communications, 1971.
6. **Sterman, John D.** *Business Dynamics: Systems Thinking and Modeling for a Complex World.* Boston : McGraw Hill, 2000.
7. **National Center of Education Statistics.** NAEP Data Explorer. [Online] <http://nces.ed.gov/nationsreportcard/naepdata/>.
8. **Watts, Duncan J.** *Six Degrees: The Science of a Networked World.* New York : W. W. Norton & Company, 2002.
9. **Lortie, D. C.** *Schoolteacher: A sociological study.* Chicago : University of Chicago Press, 1975.
10. **Weisberg, D., et al.** *The Widget Effect: Our national failure to acknowledge and act on differences in teacher effectiveness.* s.l. : The New Teacher Project., 2009.