University of Utah Mathematics Engineering Summary Report Spring '17

William H. Nesse

March 17, 2017

1 Introduction

In Fall of 2012 the University of Utah's Math Department introduced the Engineering Math (EM) course sequence, a series of four math courses covering calculus, differential equations, and partial differential equations that satisfies the curriculum requirements of all engineering majors in the College of Engineering (CoE), schematically shown in Figure 1. The four course EM sequence entails Math 1310 calculus 1, single variable differential and integral calculus; 1320 calculus 2, integration applications, sequences and series, and multivariable differential calculus of curves and surfaces; Math 2250, ordinary differential equations and linear algebra; and Math 3140, vector calculus and partial differential equations. The same topics are covered in the Traditional (Trad) five-course sequence, albeit in a slightly different order of topics, via Math 1210 calculus 1, 1220 calculus 2, 2210 calculus 3, 2250, 3150 partial differential equations. Math 2250 is shared between the two sequences.

The EM series differs in that it includes a TA-led 50-minute weekly recitation section in addition to the four 50-minute instructor lectures per week. The recitation section provides active-learning and group-work-based instruction with the aim to improve both basic math skills, but also to improve overall problem solving fluency on mathematics applications critical for engineering practice. The EM series also differs in the order of presentation: vector calculus is presented in Math 2210 in the traditional sequence, which precedes Math 2250. In the EM sequence, vector calculus is presented in Math 3140, after Math 2250 but before PDEs.

2 Enrollments and Student-Teacher Ratios

Figure 2 (top left) shows the enrollment figures by year for the years relevant for tracking. Math 1310 and 1320 show a clear increase over the years and appear to be approaching an asymptote of approximately 300 and 250 students, respectively. Math 2250, which is shared between the two course sequences taken most commonly by engineering majors naturally shows higher enrollment

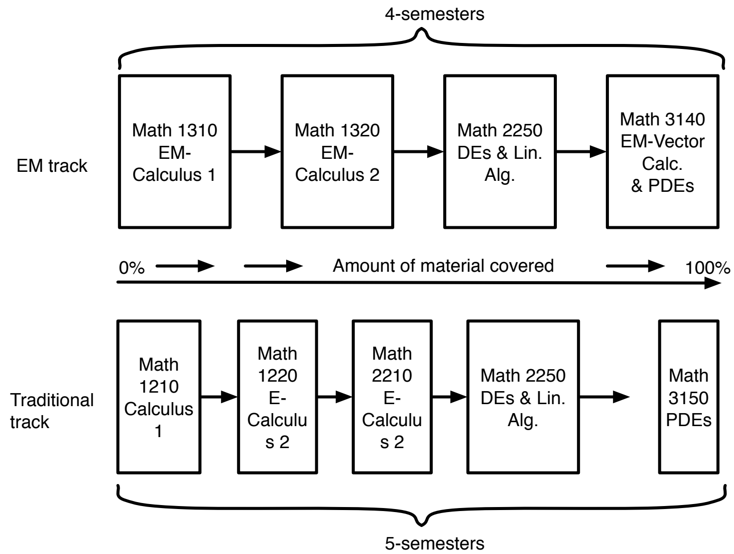

Figure 1: The 4-course EM and 5-course Traditional sequences satisfying the CoE math curriculum.

of around 750 students per year. The first two years of 2250 enrollments in this sample were omitted because the first EM cohorts were not yet fully saturated into Math 2250. Math 3140, the terminal capstone course in the EM sequence has approached about 150 students per year, about half the enrollment in the intake class. The drop is due to some students' majors do not require PDEs, so they elect to meet their requirements with another course (Math 2210), and attrition from the engineering or related science majors altogether.

Figure 2 (bottom left) shows the enrollment figures for the Traditional sequence over the same timeframe. Over the course of the EM sequence's lifespan the fraction of students taking calculus 1 and 2 via the EM sequence has asymptote to around 20% (Fig. 2 right). The student-teacher ratio in the EM courses is significantly lower than that of the Traditional course sequence, as illustrated in Figure 3

3 Achievement rates of EM versus Traditional sequences

The EM and Traditional sequences enjoy mostly comparable pass rates and earned grade distributions, but with notable differences as shown in Figure 4. Data on student achievement were obtained from students from 2012 up to and including Summer of 2016. Distributions were constructed by querying

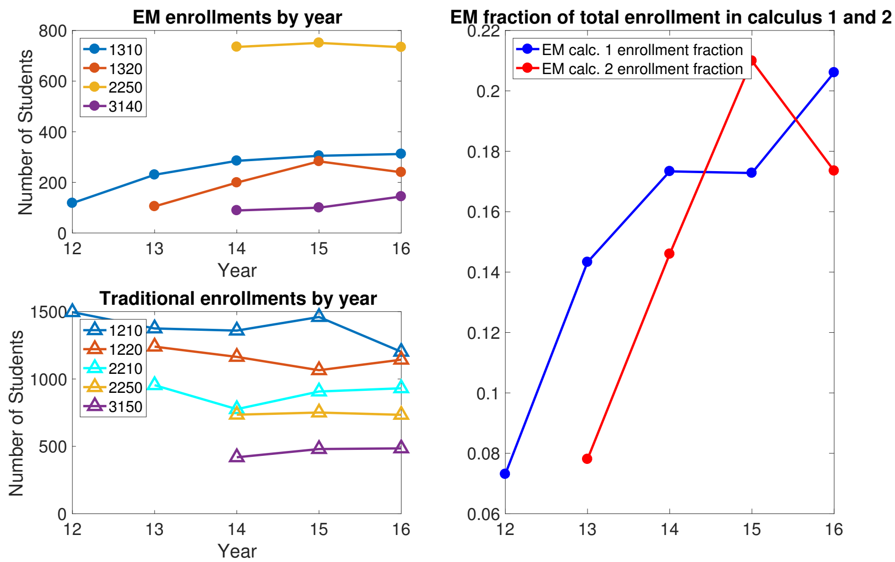

Figure 2: Enrollment numbers for the EM track (top left), the Traditional track (bottom left), and EM enrollment fraction of total calculus 1 and 2 enrollment (right) over the relevant years by course, since the EM program inception.

for populations that took all of the courses up to and including the course in question within the sequence. Cross-over students and students transferring with math credits from other institutions were excluded. For example, the Math 3150 population was formed from students who took 1210, 1220, 2210, and 2250 at the U of Utah within the aforementioned sample window—other 3150 students, of which many are out-of-institution transfers were not considered.

The overall pass rates, around 71-72% for the first-semester EM and Trad calculus courses are not statistically different. However, the second-semester EM calculus course (Math 1320) exhibits significantly higher grade achievement, with a $\sim 6\%$ greater pass rate overall compared to the Traditional second-semester calculus. The EM sequence does not possess a third-semester calculus course. Students in the traditional and EM sequences both take Math 2250, so it can serve as an internal control. Both populations exhibit the same grade distributions and pass rates. The traditional and EM students then separate again and take distinct PDE courses: Math 3150 and 3140, respectively. The Math 3140 population enjoys a higher achievement- and pass rate than Math 3150 populations.

The higher or even achievement rate of EM students relative to Traditional the population cannot be interpreted to mean that EM students are necessarily learning any differently than their Traditional peers. It can only be concluded

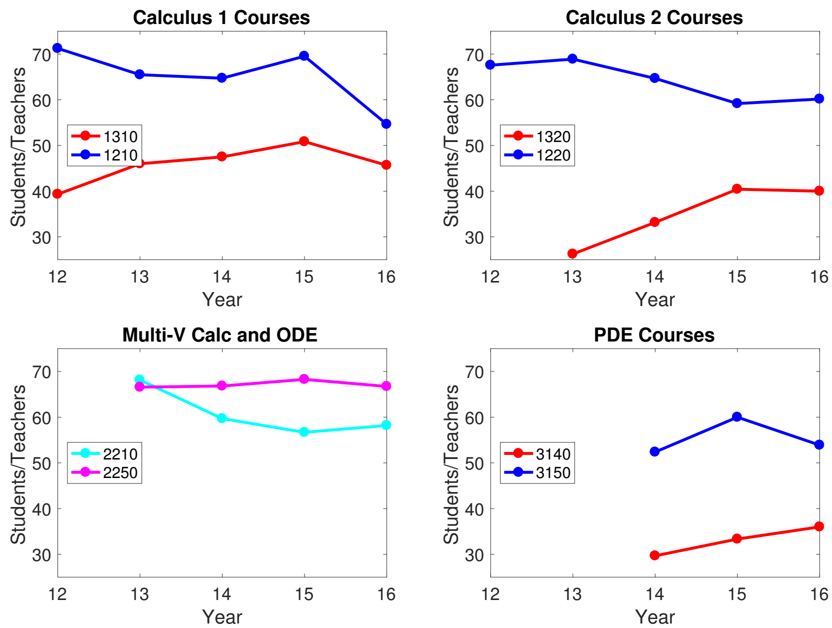

Figure 3: The student-teacher ratio: the number of students over the number of teachers for each year, broken out into all Traditional and EM courses.

that EM students who have passed through 1310-1320 are able to achieve the same average performance in 2250 as their Traditional peers who pass through 1210-1220-2210. Moreover, the Traditional sequence students are a very heterogeneous group, a majority of which are in the College of Engineering (CoE), but the sequence also serves College of Science (CoS) students, and students in many other majors on campus whose mathematics requirements vary greatly relative to CoE mathematics requirements. The EM sequence serves a more homogenous population of declared engineering majors with a more common math curriculum.

4 Time-to-complete EM and Traditional sequences

Figure 5 shows the distribution of time-to-complete each sequence, measured in semesters elapsed (including Summer semester) from the , broken down into the completion time for Math 2250 (left) and completion for Math 3140 or 3150 for EM and Trad, respectively. Students in the EM sequence exhibit significantly faster completion times. The EM sequence is four courses long and the traditional sequence is five courses, so it is to be expected that the time to complete

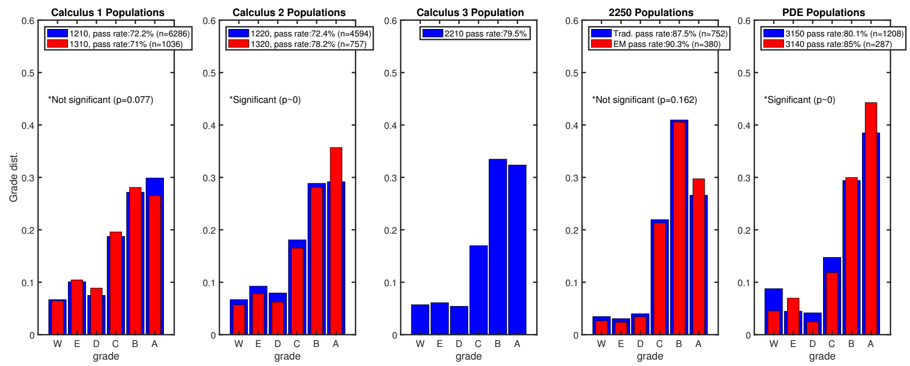

Figure 4: Grade distributions and pass rates by courses in the EM and Trad sequences.

be shorter for the EM sequence. The average relative time to complete through 2250 is 1.19 semesters, and the average relative time to complete 3140/3150 is 2.32 semesters. That is, while the EM sequence is one semester course less than the Traditional sequence, students complete the sequence over two semesters faster! The reasons for this speed-up, beyond the obvious, are not well understood, and probably involves greater pass rates of the EM sequence, and the greater convenience of scheduling four math courses within the freshman and sophomore Fall-Spring standard semesters; whereas enrollment in the fifth semester of the traditional sequence may often be delayed while taking intensive junior level engineering coursework (that involve PDEs, ironically).

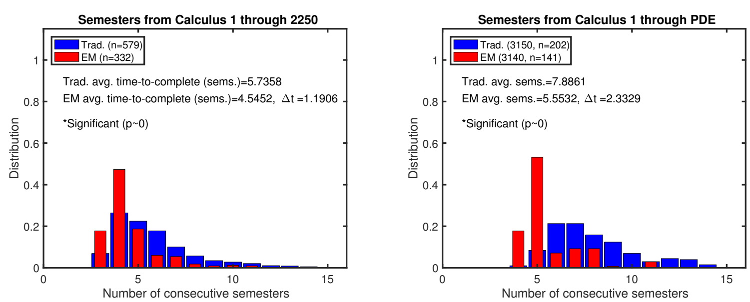

Figure 5: Distributions of time-to-complete the course sequence, measured in semesters since first course enrollment. Completion through Math 2250 and completion through math 3140 are shown in the left and right panels, respectively.

5 EM Attrition

Attrition from the EM program is overall fairly low. Attrition can occur for a variety of reasons: (1) students can elect to not take the next course in the series even with passing grade in the prerequisite course, or (2) the student can earn a non-passing grade (C- or lower) and not retake the course as many times as necessary in order to pass the course and then proceed to the next course. Persistence rates are defined for each course in the in the sequence as the fraction of students that persist to enroll in a stage relative to the total enrolled at the first class of the sequence (Fig. 6, 7 top). Attrition rates are the fraction lost between stages (e.g., the fraction not persisting from 1310 to 1320, 1320 to 2250, and so on; see Fig. 6, 7 bottom). While pass rates in any given EM course are between 71% and 90%, the overall attrition rates to the next class are between 11% and 15%. This means a significant fraction of failing students (C- or lower) do retake and pass the prerequisite course and persist to the next (Fig. 6). The attrition data are also segmented by gender, with 21% of incoming math 1310 declared engineering majors identifying as women, versus 79% as men. There is no clear gender-dependence on attrition at any juncture of the EM sequence. Note, that the data cannot indicate the reason for non-persistence, and it could be that the subsequent course is not required for a particular major (commonly 3140), or they could be electing to leave engineering or related STEM fields.

The Traditional sequence serves a heterogeneous group of students, a majority of which are in the College of Engineering (CoE), but it also serves College of Science (CoS) students, and many other majors on campus. Persistence rates are much lower and attrition rates are much higher in the first two courses of the Traditional sequence. There is a greater share of women (31%) in the Traditional sequence reflecting the greater diversity of the CoS as a whole. Moreover, there is a gender-dependent divergence in attrition rate in 1210-to-1220, and 1220-to-2210. The reason for this divergence is unclear but is a cause for concern. Many female-overrepresented CoS majors require less mathematics, so disentangling the social and achievement factors that drive attrition is not easy. Attrition in the later Traditional courses is near parity with EM courses, irrespective of gender.

6 Learning Outcomes Assessment Policy Development

The EM series is a highly articulated sequence of courses that serve STEM majors. These students will, at least in certain fields, use a great deal of mathematics in their professional practice. Nearly all of the mathematics topics/skills students learn in a given class of the sequence will be used at least once (if not continually) in topics presented in subsequent classes in the sequence and beyond. It is critical to establish minimum standards of demonstrable skills and knowledge that qualify for successful passage through the EM sequence.

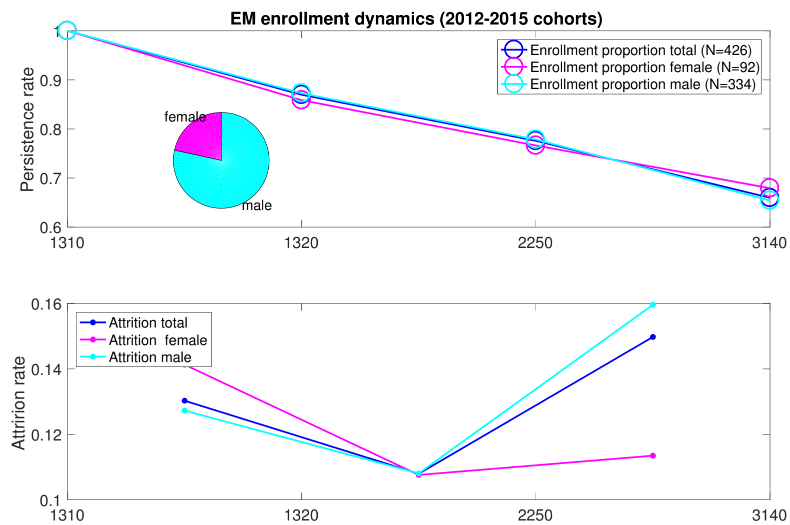

Figure 6: Course-to-course persistence rates from the EM sequence is defined as the fraction of students that continue to enroll in each stage within 1.5 years of original enrollment relative to size of the first course in the sequence (1310). Attrition rate is the fraction lost between stages. Data are taken from students in the first 1310 cohort in 2012 up to and including 1310 Spring 2015 enrollments. Data are broken out into total rates, and male and female rates.

Broadly speaking, mathematical learning entails learning both concept knowledge and basic procedural skills, and *problem solving fluency* that combines concept and skills in a concerted way. All EM courses have the stated umbrella learning objective of problem solving fluency, which is operationalized on questions given on exams by the requirement that a student should be able to

(a) Identify problem objective and apply appropriate concepts.

(b) Select the appropriate method(s) that serve the objective.

(c) Perform the methods correctly.

(d) Interpret results and make appropriate conclusions concerning the problem objective.

The later requirement (d) can vary greatly depending on the problem type; it can be assessed when available, but is sometimes very obvious and follows directly from the execution of the methods, other times extremely challenging, and other times not defined at all.

For the past two years (Spring '15-Fall '16), a learning outcomes assessment has been regularly performed. The goal of these assessments was to provide

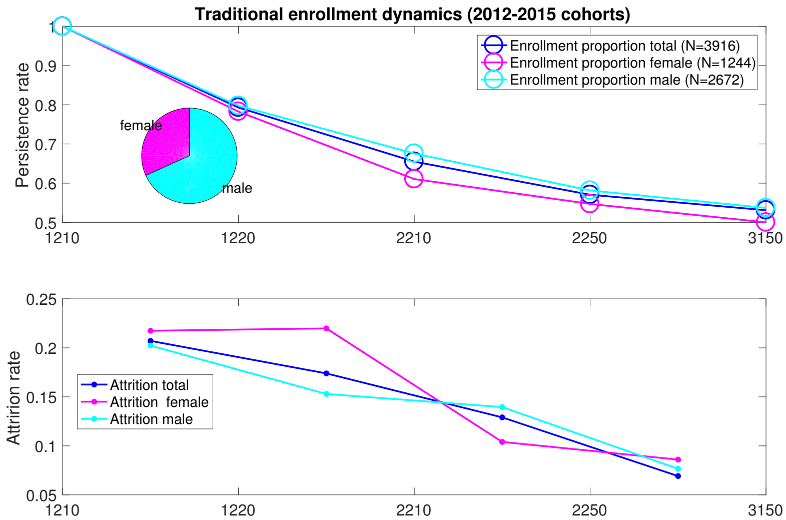

Figure 7: Course-to-course persistence rates from the Traditional sequence is defined as the fraction of students that continue to enroll in each stage within 1.5 years of original enrollment relative to size of the first course in the sequence (1210). Attrition rate is the fraction lost between stages. Data are taken from students in the first 1210 cohort in 2012 up to and including 1210 Spring 2015 enrollments. Data are broken out into total rates, and male and female rates.

a baseline measure of the EM teaching operation. At present, the results (see below), have been baselined sufficiently that we can now make recommendations for formal standards. The basis for standards are given by the following:

- A C grade or better in any EM or Traditional math course is required for graduation with an engineering degree and therefore represents a minimum standard of mathematical knowledge and skill sufficient for an engineering degree.
- Demonstration of problem solving fluency should result in a C or greater grade most of the time; a lack of problem solving fluency should result non passing grades most of the time.
- More specifically, a C grade or better should result for that a student should be able *at least* (a) identify problem objectives, and (b) select the appropriate method(s) that serve the objective.
- It is desirable that students be able to correctly perform the aforementioned selected methods and interpret results, but it is understandable that, within the confines of a timed exam, the execution is not performed

correctly all of the time. It is assumed that in an untimed environment conducive to problem solving, that the methods can be checked and rechecked until correct execution is guaranteed.

- A B grade or result most of the time when correctly selected methods are executed correctly, or only with minor errors; B grades should be very rare otherwise.
- The converse of grade selectivity is knowledge admittance: students with good grades should a high chance of demonstrating good problem solving fluency; whereas poor grades show a low chance of problem solving fluency.

7 Learning Outcomes Assessment of Fall '16 EM courses

A learning outcomes assessment (LOA) was performed on all 1310, 1320, 2250 EM courses in the Fall '16 semester. The assessment entailed the applying the following standards and procedures:

- Create three common final exam questions for each course (nine total questions) that assess a critical subset of the all stated learning objectives in each respective course.
- The exam questions must test problem solving fluency. Critically, the questions require the students to select a specific method that serves the problem objective. Correct application of the method then provides evidence of good mathematical concept learning.
- The following rubric was applied to score the questions:
  - 0: No evidence of correct identification of problem objective or methods selection.
  - 1: Weak evidence of correct identification of problem objective or methods selection.
  - 2: Clear evidence of correct identification of problem objective but serious errors in methods execution.
  - 3: Clear evidence of correct identification of problem objective and only minor errors in methods execution.
  - 4: Clear evidence of correct identification of problem objective and correct methods execution.
- Questions were scored by instructors and TAs of the courses, independent of the scores applied to the exam itself—students never see these 0-4 scores.
- Scores of 3-4 are considered a minimum requirement for mathematical competence.

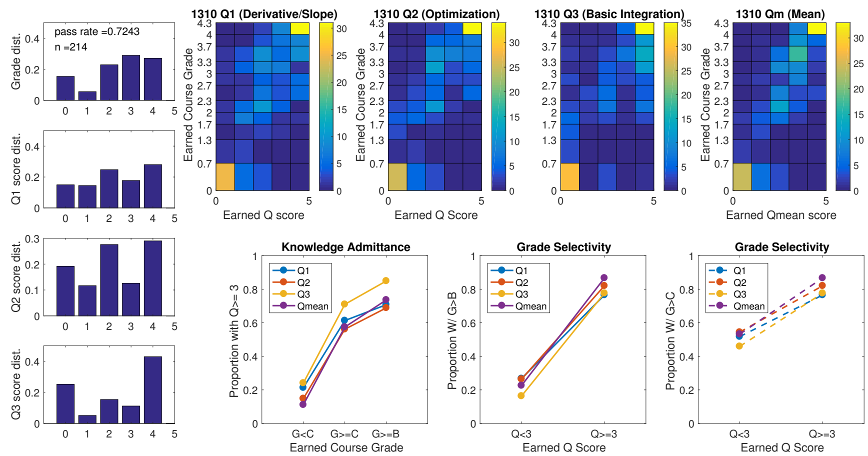

Figure 8: Math 1310 LOA results.

- Students with 3-4 scores on most assessment questions should yield C or better grades a majority of the time; sub 3-4 scores should yield non-passing grades most of the time.
- Students with 3-4 scores on most assessment questions should yield B or better grades most of the time.
- Students with C or worse grades should show sub 3-4 scores; C and greater and B and greater grades should have 3-4 scores a most of the time, and with high likelihood, respectively.

In the following we review the LOA results for Math 1310, 1320, 2250. In the Math 1310 section we develop and define key metrics that we will use in the other courses.

7.1 Math 1310 Results

The grade G and assessment question score Q results are shown in Figure 8. Students in all four sections of the Fall '16 semester were merged. The grade distribution is fairly broad (upper left) with modal grade of a B, and a pass rate (C or better grades; 'C-' is considered non-passing) of 72%, which is consistent with long-time averages shown in Figure 5. The assessment questions covered three core topics (the actual questions are listed in the Appendix):

1. Using the derivative of a function to determine the intervals and points where the graph of a function is increasing, decreasing, and staying the same.

2. Using the derivative to optimize a quantity under a constraint.

3. Using the integral of a function to find the area under a curve over a given interval.

The three questions did not direct students to use a particular method (e.g., compute a derivative or an integral, find roots, etc.) but students must deploy those methods to serve the objective and communicate the results. The distribution of scores on the three questions are listed in the left column. Compared to the grade distribution (top left), the borderline score of 2 appears with slightly greater frequency than is deemed acceptable by policy.

Grades and Q scores show a high degree of correlation, as shown in the 2D colormap histograms for the three questions and the question mean (right top row). Naturally students who flunk also commonly get zeros on the Q-scores; whereas 'A' grade earners commonly earn 4's on the q-scores. Intermediate scorers most commonly earn intermediate grades. This is natural. In fact, the overall grade entropy (using empirical distributions) is roughly 3.2 bits[1], whereas the grade entropy conditioned on the mean Q scores is 0.98 bits, resulting in a mutual information between grade and average Q score of 2.22 bits. This massive reduction in entropy indicates that with just three assessment questions 69% of the grade variation can be accounted for by the mean Q score.

The critical metrics that will be used to direct the fulfillment future learning outcomes and set standards are based on the likelihood of significant disagreements between Q-scores and grades G. Essentially, the learning outcomes policy mandates that a sufficiently small fraction of students appear in the upper left and lower right regions of the 2D histograms. Furthermore, only a small fraction of students should pass the course while earning sub 3-4 Q scores. The metrics are tuples of numbers defined as follows:

- Knowledge Admittance:

$$KA(X) = \begin{bmatrix} P(Q \geq 3 \& G < X)/P(G < X) \\ P(Q \geq 3 \& G \geq X)/P(G \geq X) \end{bmatrix},$$

  where $X = C, B$.

- Grade Selectivity:

$$GS(X) = \begin{bmatrix} P(Q < 3 \& G \geq X)/P(G \geq X) \\ P(Q \geq 3 \& G \geq X)/P(G \geq X) \end{bmatrix},$$

  where $X = C, B$.

The knowledge admittance and and grade selectivity are shown for the three questions and the mean question score in Figure 8 (lower right). It is most

[1]Entropy $-\sum_n p(n)\log_2(p(n))$ is interpreted as the average number of yes/no questions required to determine a student's grade. The mutual information is the number of yes/no questions that could be dispensed with if the Q score is known.

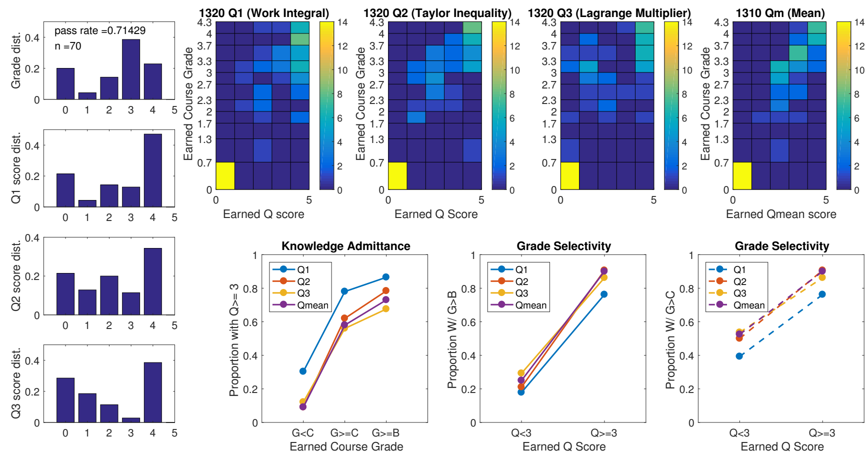

Figure 9: Math 1320 LOA results.

efficient to examine the Q-score mean because all three questions behave more-or-less similarly.

The knowledge admittance shows that among C grade earners less then 20% exhibited 'good' 3-4 score means. In past analyses this metric was higher leading to a hypothesis that many 1310 students enroll with calculus experience in high school and do fairly well on the questions, but a large fraction earn poor grades because they are not prepared for college level workloads. In this particular semester this hypothesis is not supported. Those earning passing C or B and above grades showed progressively higher 60% and 75% fractions scoring 3-4. I would like to see B and above grades have have above 80% earning 3-4 scores.

The grade selectivity shows performance in-line with the previously stated 80-20 rule: Only 20% of poor scorers earned B and above grades while 80% of good scorers earned a B and above grade. Passing grades show a different story: 50% of poor scorers earned a passing grade. Ideally I would like to see closer to 30% earn passing grades.

7.2 Math 1320 and 2250 Results

The results in Math 1320 and 2250 (Figs. 9) are broadly similar to 1310. In 1320, the assessment questions entailed

1. Using an integral to compute the work (energy) done on an object.

2. Using a Taylor polynomial to approximate a function and use Taylor's inequality to measure the approximation error.

3. Lagrange multiplier method for finding maxima/minima.

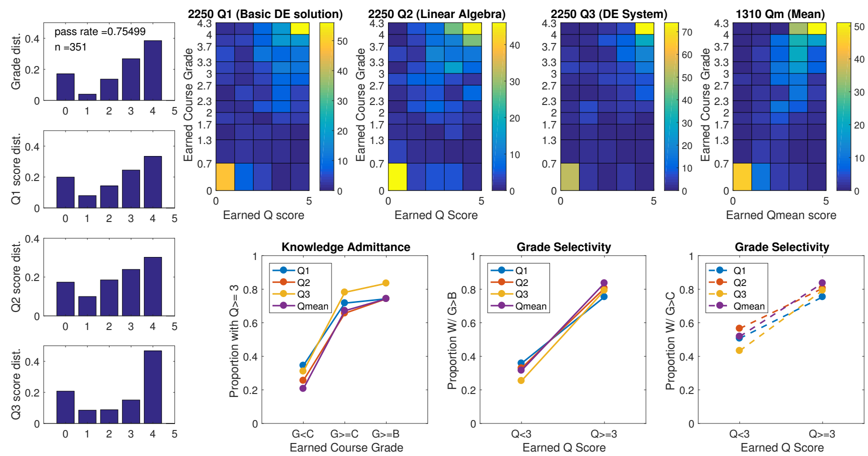

Figure 10: Math 2250 LOA results.

In 2250, the assessment questions entailed

1. Solving a DE $y' = f(y)$ with initial condition via one of the standard methods.

2. A linear algebra question entailing finding a basis for a subspace.

3. Solving a system of ODEs $\vec{x}' = A\vec{x}$ using matrix diagonalization.

As in the 1310 results, the key metric of interest is the proportion of low Q-scorers earning passing or good grades. While few ($\sim 20-30\%$) of low, sub 3-4 scorers earn B-grades, the 50-60% of sub 3-4 scorers earn C and above grades. It would be ideal if less sub 3-4 scorers earned C grades—closer to 30%.

8 Learning objective achievement and success passing through the EM sequence

In the previous section we established that B and better grades do ensure a reasonable chance that learning objectives are achieved; however, C-level passing grades do not ensure a reasonable level of learning objective attainment at the time of the grade assignment. The EM sequence relies on ensuring a reasonable level of learning attainment in order to succeed in the subsequent courses. So, an obvious question is what performance is observed in later classes conditioned earning C, B, and A grades in prior courses. Figure 11 shows the conditional grade distributions $P(X|Y)$ for the population of students who have persisted to a given course within the sequence (students must not enroll in the next course n+1 in the sequence within at least 1.5 years to be considered in the sample of those stopping at point n in the sequence). The top row shows the reference distributions for entire population of each course in the sequence, and the subsequent rows show, respectively, populations that made 3140, stopped at 2250, stopped at 1320, and stopped at 1310. There is a small fraction of students who switch to the traditional sequence midway through the EM sequence that is not considered.

As stated before, we are interested in specifically what fraction of C-grade earners (excluding C- grades) in early classes persist to the terminal classes 2250 and 3140. Roughly 68% of the EM sequence persists through 3140 (see Fig. 11, and Fig. 6). Of those that persist 9.5%, 11.5%, 19% earned a C grade in 1310, 1320, and 2250, respectively. C grade earners make up 17%, 14.5%, and 16%, of the respective total 1310, 1320, and 2250 populations. This persisting population is less than 68% of the total (there are other populations that show up in the total that are excluded from the EM sample). So, roughly $(100(0.095\times 0.68)/0.17 = 38\%$ of C-grade earners in 1310 make it to 3140. Similarly, roughly 53% and 80% of C grade earners in 1320 and 2250, respectively, make it to 3140. Surprisingly, those making it to 3140 via the EM sequence enjoy a much higher pass rate (94%) relative to the population as a whole (85%)—some students outside the EM program will take 3140, including some physics and math majors and students who need to take 3150 but opt for 3140 for either scheduling or other reasons.

Those that do not persist to 3140 show progressively weaker grades in the preceding courses to which they stop at. Those that stop at 1320 and 1310 are particularly telling: the model preceding grades are Cs. Taken together, its clear that borderline-but-passing C grade earners in early courses (1310, 1320) of the EM sequence are likely to not persist through their math requirements. Students either discover that they are not up of the task and opt or fail out, or they make changes to their learning strategies and succeed. However, having roughly 80% of C grade earners in 2250 moving on to 3140 appears to be a potential risk. Differential equations (2250) is the most important course in the sequence and the heart of engineering mathematics.

To understand how grades in previous courses determine performance in subsequent courses, forward conditional distributions were generated for the population of students who make it to 3140. Figure 12 shows grades in 1320, 2250, and 3140 (columns) conditioned on 1310 grades C, B, and A on the rows from top to bottom, with row proportions given on the ordinates. There is significant grade diffusion in subsequent courses. A C grade in 1310 yielded no As in 1320 and 2250, but by 3140 there was a broad distribution of outcomes centered B grades. Clearly what is not learned well early on can be remediated to some extent, but again, the most likely grade in 2250 for for a 1310 C earner is a C. B and higher grades show ever greater assurance of 3140 passage.

Figure 13 shows forward distributions conditioned on 1320 grades. C grade earners in 1320, which covers the most critical and difficult concepts in calculus, shows a more severe effect on future performance—only 66% pass rate in 3140 versus 85% overall. B and A earners in 1320 are all but assured passage in 3140

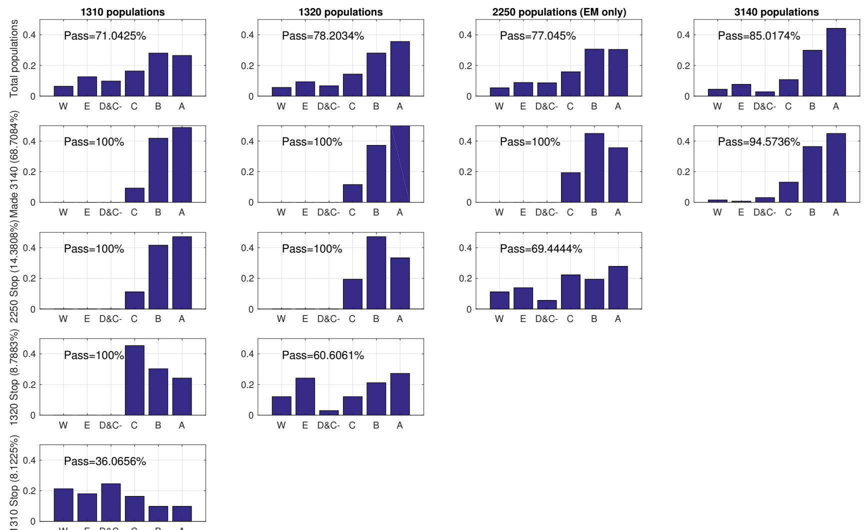

Figure 11: Conditional grade distributions: Top row shows grade distributions for the total population over the last five years. The second row from the top: grade distributions in the EM sequence that were successful in passing to 3140. Third row: students who stopped at 2250, did not elect enroll in 3140 within a year or more time. Fourth row: students stopping at 1320. Fifth: students stopping at 1310.

(provided they also pass 2250).

Figure 14 shows forward distributions conditioned on 2250 grades. C grade earners in 2250 only weakly affects their success rate in 3140. Interestingly the 1320 C grade (combined with 2250 passage) had a more significant effect on 3140 performance than a C grade in 2250 alone. Either 2250 does a very good job at selecting 3140 enrollees, or 1320, which contains the most conceptually challenging material, forms lynchpin knowledge that is a greater predictor of future success, or 3140 standards should be examined more closely—3140 is a course who's curriculum is part calculus and part partial differential equations. Its curriculum has been evolving so formal assessment has not yet been performed, but it will soon.

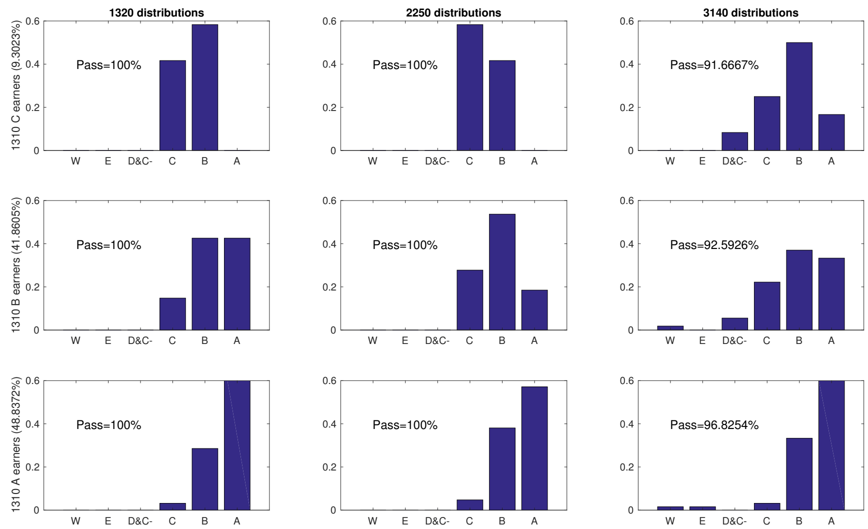

Figure 12: Forward grade distributions conditioned on 1310 grade: Within the population of 3140 enrollees, the top, middle, and bottom row shows the grade distributions for passing C, B and A grade earners in 1310, respectively. The percentages of the population in each row are given on the left ordinate and the pass rates through each class are within each panel.

9 Preliminary Conclusions of the LOA and Grade Achievement

It is clear that greater effort should be made to ensure C grades meet minimum learning outcome standards. At the beginning EM courses, it appears learning outcomes performance tied to passing C grades is partly mitigated by time: students can make up for poor performance early on by cementing those skills as they progress—most that earn Cs early on do leave the program and those that remain often improve, albeit weakly. However, 2250 (and 3140 eventually) should be the primary focus of increased standards learning outcomes. Strategies to enact such a program change, while maintaining reasonable student experience and fair pass rates will be a topic of future discussion.

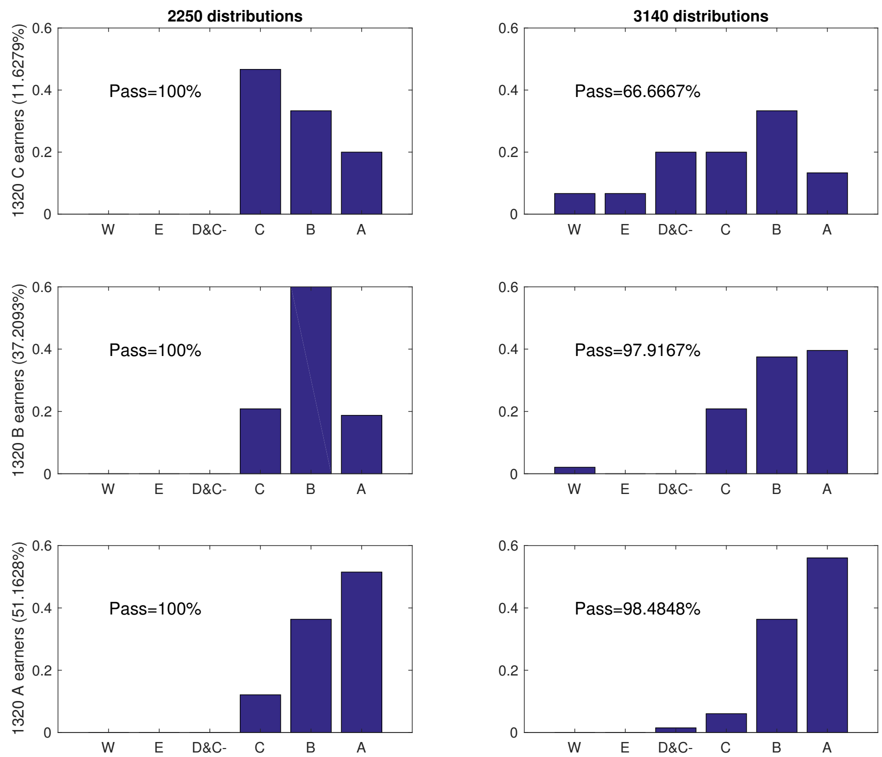

Figure 13: Forward grade distributions conditioned on 1320 grade: Within the population of 3140 enrollees, the top, middle, and bottom row shows the grade distributions for passing C, B and A grade earners in 1320, respectively. The percentages of the population in each row are given on the left ordinate and the pass rates through each class are within each panel.

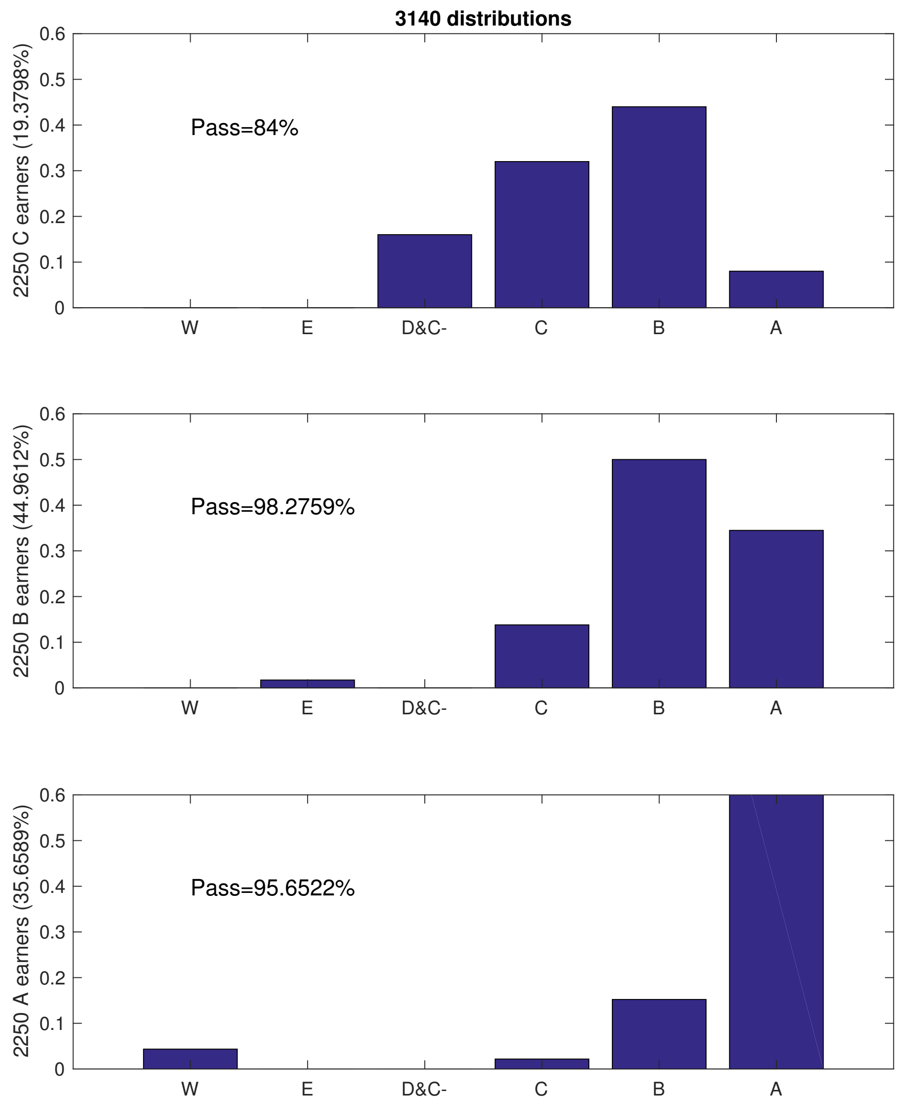

Figure 14: Forward grade distributions conditioned on 2250 grade: Within the population of 3140 enrollees, the top, middle, and bottom row shows the grade distributions for passing C, B and A grade earners in 2250, respectively. The percentages of the population in each row are given on the left ordinate and the pass rates through each class are within each panel.