8

Study of teaching and learning activities in the lectures of a theoretical course

Anders Rønn-Nielsen

Department of Mathematical Sciences, University of Copenhagen

Introduction

This project concerns my teaching in the course Beting (conditioning and Markov properties) in Block 3 of the academic year 2012/2013. It was the second time, since I taught the course in the previous year as well. Although the students had been quite content with the teaching in 2011/2012, I was not entirely satisfied myself: During the teaching I was so occupied producing lecture notes and exercises that I was not really able to spend as much time on the actual teaching as I would have liked. The result was that the teaching ended up being mostly "old-fashioned lectures"' almost without student activities.

The students are mostly statistics and actuarial students, and they are typically on their 3rd or 4th year of study. This year the course had 9 participants (8 of them until the exam), but that number is likely to increase in the next years. The teaching lasts 7 weeks (1 block) and has 5 lectures and 4 exercise classes per week. The exam is a 24 hour written take-home-exam. The course is theoretical and furthermore has the purpose of giving the students some tools for later courses. Like most other mathematical courses the course is 'vertically' constructed such that almost all elements lie end to end with previous elements of the course. Hence it is not possible to divide the course into smaller independent parts. On the other hand testing the understanding of the entire course can be done by testing the understanding of the final elements.

The course description that corresponds to the course this year (I had to rewrite the description for next year already before I gave the course this year) was not entirely optimal. Basically the students are supposed to obtain two types of competences. They should be able to apply various techniques and methods, and they should obtain an understanding of the material deep enough to make them able to see new coherences. Generally I find that the intended learning outcomes (ILO's) of the course description agree with my own wishes for the students outcome of the course. Nevertheless, I believe that the ILO's should be more concrete in the quantification of the desired deep learning.

The version of the course that was given last year had in my opinion too much importance attached to proofs and too little to the true understanding of the material and also the ability to apply the theory to concrete examples. I certainly believe that proofs and technical arguments are important. Both in the general education of statisticians and in the deep understanding of the subject. Proof techniques simply form a craft that has to be learned. However, it is equally important to keep some focus on utility of the proved results.

It was my experience that 2 or 3 of the 6 students who followed the course last year did not really obtain a deep understanding. At the end of the course they were only able - in a rather shallow way - to use some of the simplest mechanical methods, and it was clear that they were not aware of why the methods works. This is also seen in the exam results from 2012, where 2 students were (alarmingly) close to failing the exam. See Table 8.1.

Goal for improving the course

My goal for the teaching this year has been to make sure that the students experienced deeper learning and gained familiarity with the subject. The hope was - not at least - that the level of the weakest students could be increased, such that they benefited more from the course than just having a shallow knowledge and mechanical calculation skills in the end.

I have not myself been teaching at the exercise classes, since a TA was attached to the course. Hence I have mostly made changes of the exercise classes on a more general level: I have tried to align the exercises to the lectures and the exam, and I have - together with the TA - made the program for the classes. My considerations concerning the exercise classes are collected in a later section.

I have not made any changes of the exam compared to last year. A 24-hour written exam can actually measure how well the ILO's of the course are satisfied. It is easy to find exercises that tests the students ability to apply the various desired techniques and skills to concrete examples. At a longer exam like this it is even meaningful to pose questions that tests the understanding of coherences in the material: Exercises can be constructed such that they can only be solved if the student is able to combine concepts in new ways. This will test for deeper understanding - possibly deeper than what is requested in the ILO's.

My main focus of this project has been on the lectures. The goal has been that they should contain more student involvement and more examples. These changes and improvements are described in the next section.

Improving the lectures

As mentioned the goal of the changed lectures was to increase the student activity. I also wanted to spend more time on examples (including simulation studies) since I believed that this could develop the students intuition for what the field is really about.

To find time for the new improvements at the lectures, I had to reduce something else in comparison with the lectures I gave last year. I chose to cut down on the amount of proofs shown at the lectures. The changes were not supposed to move all focus away from the theoretical part of the course: I still expected the students to gain insight in a lot of results and to become able to construct similar proofs at the exam. But I removed all proofs that I estimated were either without an interesting technique or did not add any intuition to the result.

There are numerous reasons why activating the students at lectures increase the learning. Among these are

- Even the most motivated student will have troubles staying concentrated for more than 15 minutes if the person only is supposed to sit and listen (see e.g. (Biggs & Tang 2007, p. 173)). Variation and small breaks are created by introducing various activities during the lecture.
- Learning is easier if the student can relate to the problem. This can be obtained if the student has had the opportunity to think about small problems by himself.
- The learning will be deeper if the student has been part of developing the material (or at least has a feeling of being part of the process)

- Student activities creates interaction between students and teacher, and thereby it is possible for the teacher to adjust the teaching according to the needs among the students.

Similar improvements of lectures using teaching and learning activities (TLA's) have been tried (with great success) in a series of previous KNUD projects (Schneider 2007, Xella 2008, Andersen 2010). I have found some inspiration in these projects.

I decided to use the following guidelines when planning my lectures:

- There should be at least one TLA (but possibly more) per 45 minutes of lecturing.
- I should try to ask as many questions as possible when lecturing.
- I should switch between blackboard and slides as often as possible (to create variation).
- I should give examples or show simulation studies whenever it was meaningful.

In the following subsections I will give examples of TLA's used in the lectures. They are divided into 4 different types depending on the context and the purpose of the TLA.

Activities developing intuition

When introducing new concepts it can be an advantage to think about how the students - possibly already from the beginning - obtain some intuition of the new theory or ideas. In figure 8.1 is seen 4 examples of TLA's at the lecture, where the purpose has been to enhance the intuitive understanding. Each of the 4 slides should not be seen as self-contained since notation and use of words may refer to previous slides presented at the relevant lectures.

Upper left slide: The aim of this activity is to introduce the concept of *conditional independence* before an actual definition is given. The students were - on an intuitive level - able to say that there is (probably) not independence between tomatoes and cucumbers, but *given* the knowledge that the dinner includes salad, then it makes sense to say that the events of tomatoes and cucumbers are independent. When the formal definition of conditional independence was introduced shortly after this, the students already had a skeleton of the concept in their heads: They were able to relate general events and sigma-algebras to tomatoes and salad.

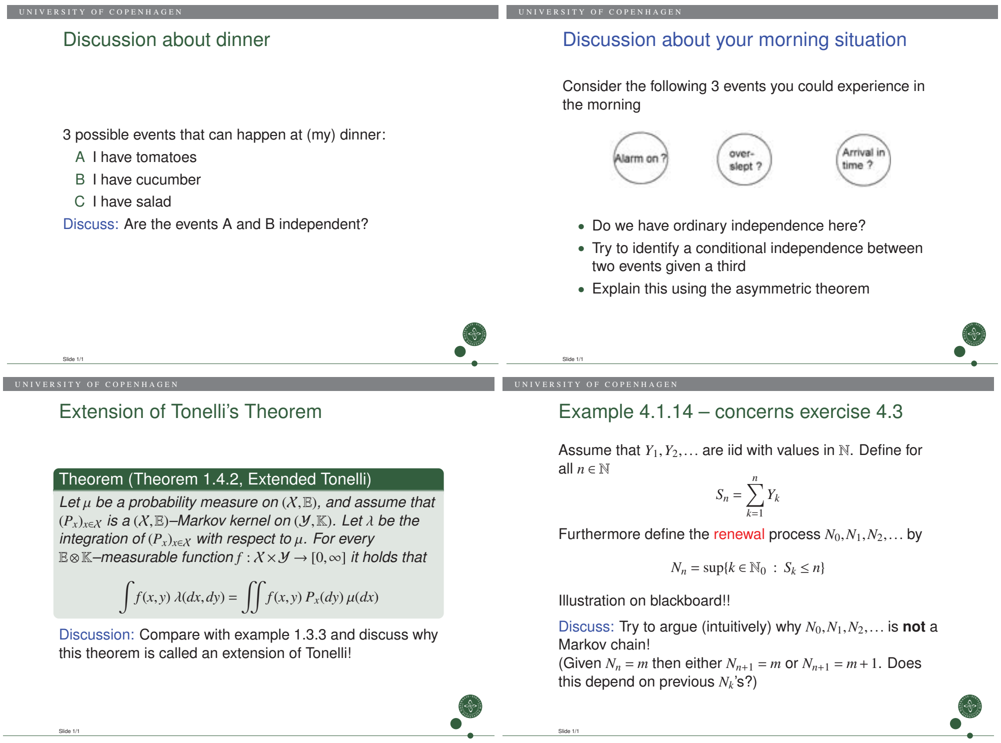

Fig. 8.1. Shows 4 examples of TLA's that helps to develop intuition.

Upper right slide: This discussion took place the day after the activity to the left. At this time an equivalent definition of conditional independence had been presented and the idea of the activity was that students tried to identify such a relation in an everyday example that was similar to the salad example from the day before, but where the new formulation method was more meaningful. This example was used repeatedly afterwards when further results were presented. Again, it was useful to have a very concrete example when results containing various events and sigma-algebras were to be explained.

Lower left slide: On this slide a theorem that appears rather complicated was presented - and it was to be proved afterwards. The theorem can be seen as a generalisation of a result that had been well-known to and also frequently used by the students for at least a year. The aim of the discussion was that the students deduced why this is true, and hopefully during this process got more familiar with both the notation and the content of the result. This helped a lot when we started going through the proof afterwards.

Lower right slide: This slide was presented at a lecture, where the concept of Markov chains was introduced. At this time the students had seen several examples of processes that actually *are* Markov chains, so here the idea is that the students get to think about how this "can go wrong". A consistent theoretical argument will be rather complicated to produce in a short time, but it is possible to give intuitive arguments that are instructive when understanding the so–called Markov property. Furthermore this example served as an introduction to an (notationally heavy) exercise they were supposed to work with at the following exercise class.

A possible challenge of this type of TLA's is that my intuition is not necessarily the type of intuitive arguments the students need to understand a given problem. I have experienced that some intuitive comparisons I asked the students to make, only made them more confused. They simply regarded some results in a completely different way. In these situations I had to spend some additional time on combining the different ways of thinking. Of course if the combining-process is successful it can be an advantage in the end that students have been exposed to several ways of interpreting the material.

Activities relating to known theory

For various reasons it can be profitable to use TLA's to remind the students of material they (in principle) already know about.

- An obvious reason could be that I get the opportunity to check whether they actually did understand what happened at the lectures, say yesterday.
- Another obvious reason is that it (almost) always helps the deep understanding of a subject if it is seen several times - possibly with varying applications and ways of formulation.
- It could also be useful to remind students about results (perhaps from previous courses) that are about to be applied in the lecture.

In Figure 8.2 is seen four examples of TLA's that relate to theory that is already known.

Upper left slide: The exercise in this slide was presented just before a rather complicated proof was started. The proof relies heavily on a particular calculation method that is used repeatedly. In principle this method - which is not very deep - is well-known to the students from previous courses, but I feared that using the method in this course (and thereby a changed framework) would confuse them. Hence I decided to prepare the students by giving them the exercise, where the method is naturally derived

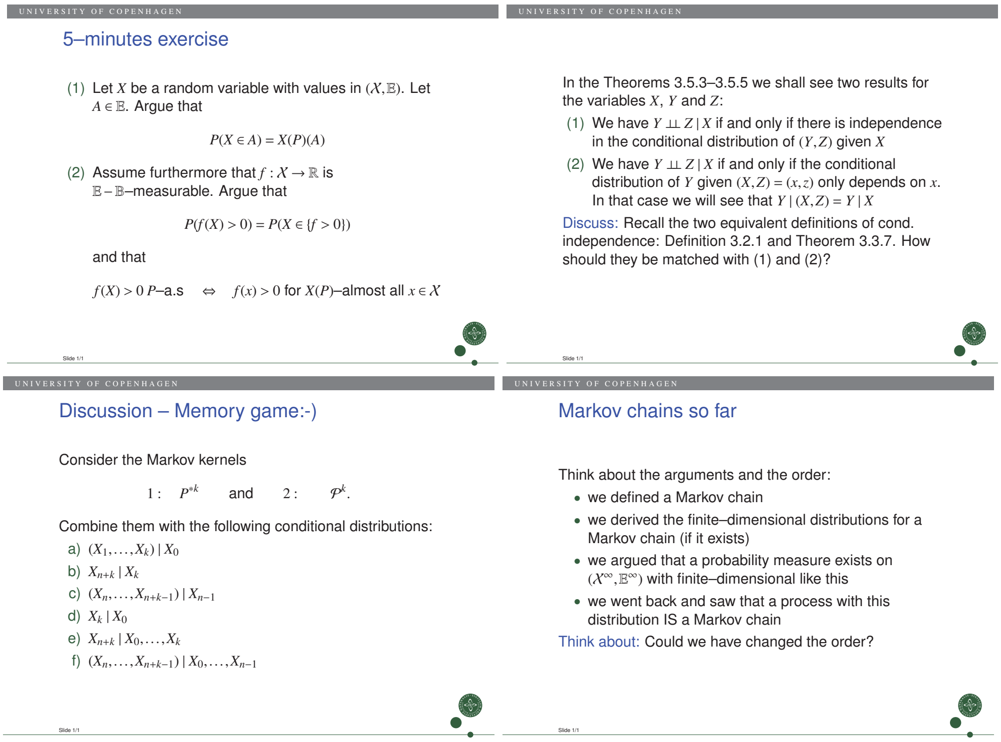

Fig. 8.2. Shows four examples of TLA's that help to understand known theory

in a simple and known framework. Most of them regarded it as extremely simple, but on the other hand it was really useful in the following proof: The students had no problems understanding the technical arguments.

Upper right slide: At a previous lecture the students had seen two equivalent ways of defining conditional independence. At this lecture we studied conditional independence in the special case when it is formulated for random variables. The two equivalent definitions of conditional independence are still meaningful but when formulated in the new framework it is not immediately clear how the two new definitions should be combined with the already known definitions. That is what the exercise in this slide is about, and the aim is that the students are forced to use known theory in new applications. This will hopefully give a deeper learning of both.

Lower left slide: This slide was used in the beginning of a lecture, where we were going to use various formulations of conditional distributions repeatedly. These formulations had been presented and used at the previous lectures, but I was not sure how well they were remembered and understood by the students. Hence I introduced this "memory game", where the students were supposed to compare different formulations of the same concepts. It actually worked really well: Almost all students gave input in the summarising part in the end, and they seemed more prepared than usual when we used the formulations afterwards.

Lower right slide: This slide was presented after the students had seen how Markov chains are constructed, and the purpose of the slide was that they should obtain a deeper understanding of how the 4 main parts of the construction fit together. The idea of the slide was to ask the students to reorder the 4 parts in a way such that the construction still makes sense (this is in fact possible). The reordering process forced the students to think about the content of each of the 4 parts and of how they interplay. Hence a correct answer of this exercise will be a valid indicator of deep learning.

Activities that gives involvement in a proof

Since Beting is a theoretical course and an important part of the ILO's of the course concerns proof techniques in the relevant field, some amount of proofs will necessarily be present at the lectures. In order to avoid situations, where I give one-man-shows that easily can last more than 30 minutes resulting in sleeping students with only a very superficial understanding of what happens, I have had to come up with ways to involve the students actively in the proofs.

Naturally, I try to ask small questions as often as possible, but typically this means that only one student can give an answer and that there is not much discussion going on between students. My solution has been to find more formalised ways of involvement. Four examples of ways to do this are seen in Figure 8.3.

Upper left slide: This is an example of a rather short proof, where the students can see all the steps but without any arguments. Previously the students have seen a similar - but not identical - proof of another result. In this exercise the students are asked to explain how the result is obtained. That makes them think about both the previous result and the similarities between the two proofs. My hope is that they are more prepared to give other similar proofs in the future (e.g. at the exam).

Upper right slide: This slide was presented in the middle of a rather long and possibly tedious proof shown on the blackboard. On this point in the proof a key step is about to be taken. The slide sums up what has been shown so far, and then asks the students to figure out how the important step can be shown. With this TLA the students became much more involved in a key part of the proof and furthermore got a break with time to think about what had happened so far.

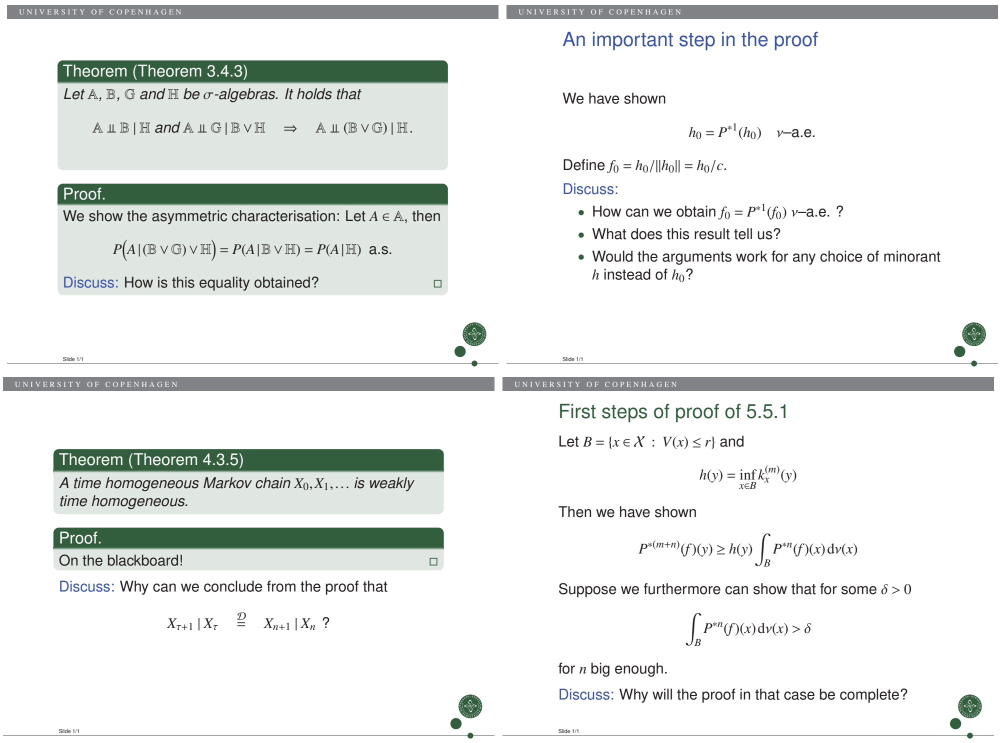

Fig. 8.3. Activities that gives involvementin a proof

Lower left slide: This slide was presented to the students just before I started to give the proof (on the blackboard) of the theorem seen on the slide. The proof lasted about 10 minutes, and the students got the task to observe my proof and then figure out why something more than the stated results actually could be derived from the proof. This exercise worked very well. Giving the students this TLA forced them to pay special attention to what was going on and furthermore it yielded some excitement because they had to look for a hidden result.

Lower right slide: The idea of the TLA on this slide is quite similar to the upper right slide. The students had seen me showing a series of sub-results on the blackboard, and now they got the exercise to combine these sub-result to the result of the theorem. Again this forced them to take part in a key step of the proof.

Activities on the form of exercises

Although the course has separate exercise classes it is very reasonable that some of the TLA's have the form of small exercises. This way the students have the opportunity to apply the theory to concrete examples immediately after it has been presented. Similar to the intuition-based TLA's the exercises help the students obtaining a deeper understanding of the material: When they know about possible applications it becomes easier to relate to the content of a given theorem.

In Figure 8.4 a collection of TLA's on the form of exercises is displayed.

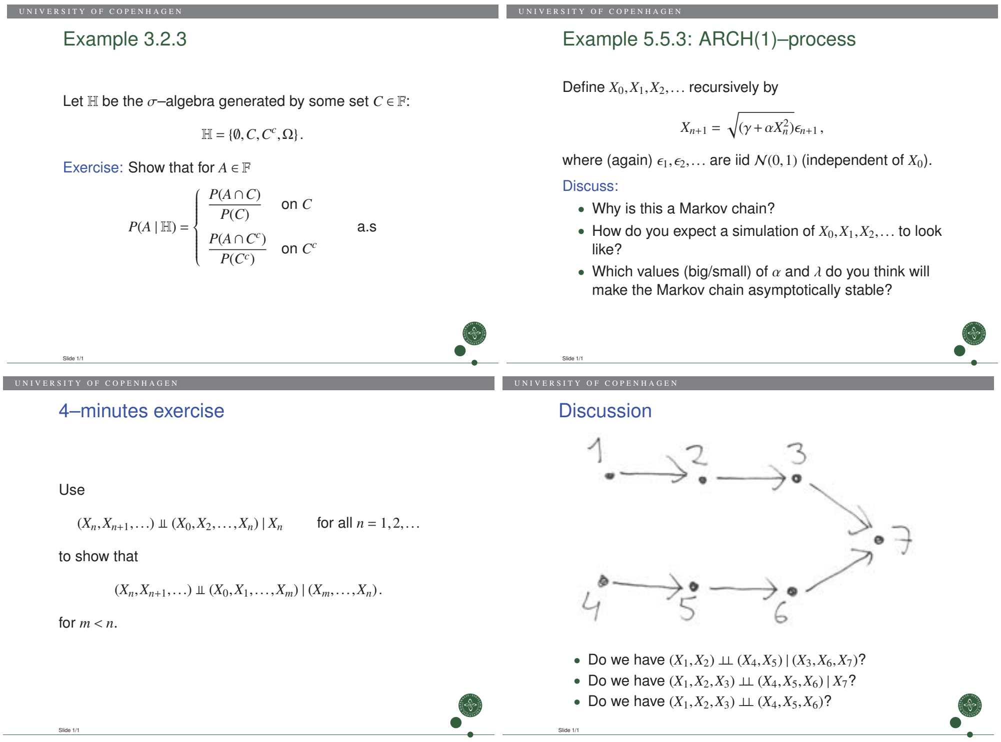

Fig. 8.4. Activities on the form of exercises.

Upper left slide: This is a "standard" exercise, where the students were asked to use a new (it was presented to them 5 minutes before) definition to establish a general property in a simple example.

Upper right slide: This is also a standard exercise. At this time the students had seen several results concerning Markov chains, and in this exercise they had the chance to apply these results to a concrete (but not trivial) example.

Lower left slide: Before this exercise the students had seen several proofs where a certain argument technique was used. I expect them to be able to use this technique at the exam. In the exercise they are asked to prove a similar result using the technique.

Lower right slide: This is a classical voting exercise with 3 questions each with the possible answers *yes* and *no*. The questions are not at all very complicated but you need to understand some basic concepts in order to be able to answer correctly. With this type of exercise it is possible to get answers from the entire class by asking: "How many say yes?".

A problem that is often encountered when exercise-based TLA's are used, is that the exercise question may be either too simple or too complicated. If the questions are too difficult the students may become insecure and perhaps also less willing to give an answer the next time they are presented to a similar exercise. In such situations it can be helpful that the teacher admits that the exercise in fact *is* difficult.

If on the other hand an exercise appears to be too easy to solve, the students can actually be reluctant to answer as well: In such cases they are afraid that they have misunderstood what the question is about. Here a solution could be that the teacher admits that the question *is* an easy (or even silly) question.

General comments on TLA's

Almost all TLA's had a form where I presented a slide with a question or an exercise. I made it a tradition that all such questions and exercises had a blue headline. Typically I explained the exercise and then gave the students some minutes (I told them how many) to discuss the solution. In this period I walked around ready to help and answer questions. When the time was up or all student clearly were ready, I asked for answers. In some situations - but not often - this developed into a longer discussion.

The students were not particularly willing to discuss the exercises with each other. Only a few of them said something in the time period for discussion. The rest worked on the problems by themselves - but it was clear that they were actually working. I think that the reason for this behaviour among the students was that most of them did not really know each other before this course, and furthermore group discussions are not very common in courses at the mathematics department. It will probably take some time until they get used to this type of activities at lectures. I tried to convince them to take part in discussions, but it is my personal opinion that students should be allowed to stay out of the discussions if they choose to; the best I can do is to make sure that there is a positive atmosphere in the classroom, where the students are not afraid of saying something wrong or even stupid.

It was my experience that a larger part of the students were willing to answer questions, when the answer could be expressed in a single word (e.g. yes/no questions or a multiple choice question). On the other hand, this type of questions (and answers) are less informative about what the students are thinking: Has the student developed a useful intuition, or is he in fact a little confused? I decided to use a combination of the two question types.

Exercise classes

In courses at The Department of Mathematical sciences there has been a long (but not necessarily reasonable) tradition for how the exercise classes should progress: The students are supposed to have prepared solutions for all exercises, and by turns they go to the blackboard and present the solution to the other students. If no student is willing to go to the blackboard (and that happens quite often), then the TA will explain the solution.

This system has some obvious drawbacks. Students who have not prepared the exercises from home only experience another lecture where a lot of theory is presented. They gain no deep learning from observing how an exercise could be solved if they have not worked with the exercise themselves. The students who have solved the exercise from home are not very interested in seeing the solution repeated on the blackboard. Only the student at the blackboard and the students with minor problems in solving the exercise, will gain anything from this system.

Together with the TA I chose to introduce some changes:

- We still expected that the students tried to solve most of the exercises at home
- At each exercise class (an exercise class lasts 2 hours), the first hour was spent on students (in groups or individually) trying to solve the remaining exercises - with help from the TA.
- The second hour was spent on the traditional presentation of exercises by the students or the TA. Here only exercises chosen by either the students or the TA would be presented - there was not time to present all exercises.
- To obtain a proper institutionalisation with respect to the remaining exercises, written solutions of all exercises were handed out after the class.

Although this was a slight change of the didactic contract the students received the changes well.

Another change of the exercise classes compared with the previous year was that I in numerous situations tried to use exercises from the exercise classes as examples at lectures. This helped to institutionalise the student work, and furthermore the students seemed more confident with examples they already knew.

Results

Course evaluation

The results from the student evaluations can be seen in Appendix A. The general impression from the evaluations is that the students liked the course and that they gained a lot from the teaching. The workload was adequate and also the difficulty of the course has been reasonable. More importantly for this project, the students liked the balance between proofs and examples and the amount of TLA's. The answer to the final question "*At lectures I achieve a better understanding of a subject if there are student activities (as opposed to the subject being presented without student activities)*" is a bit interesting. Only 16% of them agreed, while the rest of them were neutral or "didn't know". I asked them afterwards if this meant that the TLA's were unnecessary, but that was definitely not the case: They did not want less or changed TLA's. Some of them said - with a certain amount of logic - that it was impossible for them to know whether they would have obtained the same amount of understanding without the TLA's; not having TLA's had simply not been an option (they are mathematicians after all...).

Exam

To the best of my knowledge the exam this year had the same level of difficulty as the exam last year. However, the results had improved remarkably: 6 out of 8 students obtained the best grade 12 by handing in excellent solutions showing an almost perfect understanding of the subject. The two remaining students handed in solutions on a much higher level than the two weakest students last year. As a statistician I am reluctant to conclude too much in a comparison of two groups with 6 and 8 observations - it may very well be that the student population this year simply was more able.

However, it is my impression that the changed teaching method actually had an effect on the exam result.

Grade	-3	00	02	4	7	10	12	average
2012	0	0	2	0	1	1	2	7.5
2013	0	0	0	1	1	0	6	10.4

Table 8.1. Exam results from last year (2012) and this year (2013)

Conclusion

In this project I have studied how to improve deep learning in the course Beting. I have realised that although the course is theoretical and contains a lot of proofs it is in fact possible and meaningful to make student activities in the teaching. My experience has been that it is essential to vary the teaching as much as possible. Not at least when it comes to the use of TLA's, where different types of TLA's are useful in very different situations.

I believe that the changes I introduced during the course have been welcomed by the students and have helped them obtaining deep learning. This was far from being contradicted by the exam result - which was remarkably good.

Appendix A – Student evaluations

Question 1: I experienced a good correspondence between the teaching and the course objectives

Strongly agree	83,3%
Agree	16,7%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 2: I think that the practical execution of the course was successful (facilities, equipment, information dissemination etc.)

Strongly agree	83,3%
Agree	16,7%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 3: I experience a good coherence between the various course elements (lectures, practical work, etc.)

Strongly agree	83,3%
Agree	16,7%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 4: I experience the course as relevant to my personal educational objectives

Strongly agree	66,7%
Agree	33,3%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 5: In cases where I needed feedback on my work (presentations, assignments, papers, reports) I was able to adequately get such feedback from the teachers

Strongly agree	50%
Agree	33,3%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	16,7%

Question 6: For me, the teaching material is adequate for this course.

Strongly agree	66,7%
Agree	33,3%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 7: Compared to my background knowledge I experience that the academic level of the course is

Far too low	0%
Low	16,7%
Adequate	83,3%
High	0%
Far too high	0%
Don't know	0%

Question 8: I experience the work load of the course as

Much too low	0%
Somewhat low	16,7%
Adequate	66,7%
Somewhat high	16,7%
Much too high	0%
Don't know	0%

Question 9: In this course, for me the average work load per week was (including classes, preparation, written assignments etc.)

Under 10 hours	0%
10-15 hours	16,7%
15-20 hours	50%
20-25 hours	16,7%
25-30 hours	16,7%
Over 30 hours	0%

Question 10: If you have further suggestions for improving the course - or other comments and or elaborations on your answers above (please refer to question number)

- The lecture notes works great and is nicely structured. There are, however, a lot of hints in the exercises. The hints are generally good, but I usually read them by mistake before I get to try to solve the exercise by myself. Maybe you could collect the hints in a seperate section, like in the VidSand-book.
- 1) I would have liked the textbook to be posted in one piece at the beginning of the course instead of chapter-by-chapter at the end of each week. However, this was not really a big problem.
  2) As the textbook is well-written/easy to read and the exercise problems are so slow and elaborate that even I could solve them, I think the idea of walking through curriculum during lectures could be skipped altogether.
  Maybe the lectures in the future could be biased more towards showing where, and how the theory has "real-world"-applications. By this I don't intend to suggest that the lecturer should be merely increasing the number of "examples".
  For instance, one could focus more on problems in other branches of mathematics, where knowing about conditioning and markov properties, was crucial for solving a given problem (if any exists). Or really diving full scale into the acceptance-rejection algorithm, simulated annealing, regenerating markov processes or maybe even going through some professors reasearch?
  I realize, one has to strike a balance, as some students prefer sticking to the textbook.... hmm..

Question 11: Please comment on how much you have benefitet from the lectures in this course while learning the subject at hand

- The lectures are great, with nice examples and regular keeping track on how the whole intuition is (e.g. when it comes to conditional independe/moments/distribution and such). The 1-min.-discussions during the lectures are a good idea, however usually there aren't really any discussion at all. Its a matter of us (the students) giving it a try and not being afraid of saying something wrong and a matter of you trying to push us to start discussing. The last thing is definitely the most difficult part, because there is no obvious way to do it. However, I think what you are doing works, e.g. asking some of the students individually how they would go about the problems.
- Really good lectures, Anders makes the very theoretical definitions and theorems understandable through examples and "word-explanations".
- On a scale 1-10, I would say 7-8, depending on coffee. To the best of my understanding, Anders Rønn did a good job.
- Very much. The balance between motivation and proofs is very good. Also the activating questions work fine. I am impressed. Just keep up the good work!

Question 12: Please comment on how much you have benefitet from the problem sessions in this course while learning the subject at hand

- The problem sessions were generally very good. The TA was always ready to give hints and presented the correct solutions in a nice way. Also you put a lot of effort in to typping in the correct solutions, which is really great.
- Sometimes when you present the solutions you skip parts, e.g. simple symbol-manipulations using extended tonelli etc., it is nice to know that one can find the formal way to go about the problems in the typped-in solutions. Maybe the first time you present a solution involving extended tonelli, you should emphasize that this strategy is used a lot, but in the future you will not give all the details. In that way it will be easier for me to recall that particular strategy at a later point.
- Good exersice classes. I find that questions are always welcome in class and the written solutions are really helpfull and understandable.
- The instructor seemed competent.
- A lot. It is important to be able to check whether you are on the right track. There is no chance in hell you can solve all the exercises during class (at least I can't), but even when I have solved the majority of the exercises, I believe I have benefited from the sessions.

Question 13: The lectures include both theoretical proofs and examples: I think that

More time should be spent on theoretical proofs	0%
More time should be spent on examples	0%
The balance between theoretical proofs and examples in adequate	100%
Don't know	0%

Question 14: I experience good possibilities to ask questions during lectures.

Strongly agree	50%
Agree	50%
Neutral	0%
Disagree	0%
Strongly disagree	0%
Don't know	0%

Question 15: The lectures include some student activities (questions, discussions, small exercises, ect). I think that

More time should be spent on student activities	0%
Less time should be spent on student activities	0%
The amount of time spent on student activities is adequate	100%
Don't know	0%

Question 16: At lectures i achieve a better understanding of a subject if there are student activities (as opposed to the subject being presented without student activities).

Strongly agree	0%
Agree	16,7%
Neutral	50%
Disagree	0%
Strongly disagree	0%
Don't know	33,3%

All contributions to this volume can be found at:

http://www.ind.ku.dk/publikationer/up_projekter/2013-6/

The bibliography can be found at:

http://www.ind.ku.dk/publikationer/up_projekter/kapitler/2013_vol6_bibliography.pdf/