

# The Impact of Lessons-Learned Sessions on Effort Estimation and Uncertainty Assessments

M. JØRGENSEN, T. M. GRUSCHKE

*Abstract*--Inaccurate estimates of software development effort is a frequently reported cause of IT-project failures. We report results from a study that investigated the effect of introducing lessons learned sessions on estimation accuracy and the assessment of uncertainty. Twenty software professionals were randomly allocated to a Learning group or a Control group and instructed to estimate and complete the same five development tasks. Those in the Learning group, but not those in the Control group, were instructed to spend at least 30 minutes on identifying, analyzing, and summarizing their effort estimation and uncertainty assessment experience after completing each task. We found that the estimation accuracy and the realism of the uncertainty assessment were not better in the Learning group than in the Control group. A follow-up study with 83 software professionals was completed to better understand this lack of improvement from lessons-learned sessions. The follow-up study found that receiving feedback about other software professionals' estimation performance led to more realistic uncertainty assessments than receiving the same feedback of one's own estimates. Lessons-learned sessions, we argue, have to be carefully designed to avoid wasting resources on learning processes that stimulate rather than reduce learning biases related to assessment of own estimation performance.

*Index Terms*— cost estimation, process implementation and change, review and evaluation, software psychology

## I. MOTIVATION

A recent questionnaire-based survey of more than 1,000 IT professionals [1] reports that two out of the three most important causes of IT-project failure are perceived to be related to poor effort estimation. This perception is supported by empirical studies of software development effort overruns and over-confidence. The typical effort overrun seems to be about 30% [2] and 90% confidence (or “almost sure”) that the actual effort will be included in a minimum-maximum effort interval typically corresponds to a 60-70% interval hit rate [3]. Over-confidence in the accuracy of effort estimates typically means that the project plan is not reflecting the underlying uncertainties related to the use of effort. Clearly, both software vendors and clients would benefit from more accurate effort estimates and more realistic assessments of the uncertainty of these estimates. One possible way of achieving higher accuracy and more realistic uncertainty assessments is to develop and adopt better learning processes. One

candidate for improved learning is the use of individual lessons-learned sessions, i.e., structured reviews of own estimation experience. The learning effect of such lessons-learned sessions is the topic of this paper.

Judging by the number of scientific studies, it seems that software engineering researchers believe that estimation learning processes should focus on improving formal estimation models. This has been the main approach to learning and improvement among software researchers since at least the 1960s [4]. This strong focus on the improvement of formal estimation models can hardly be defended by its success. Even the use of unstructured and unsupported judgment-based effort estimation (“expert estimation”) seems, on average, to yield effort estimates that are just as accurate as those generated by the use of sophisticated formal estimation models, see our review in [5]. In addition, formal effort estimation models are not much used by the software industry [6]. The use of the Personal Software Process (PSP) [7] provides, in our opinion, an illustration of the limitations of an estimation learning process that is based mainly on the improvement of formal effort estimation models. Independent studies with control groups report a surprising lack of improvement in estimation accuracy when applying the PSP process [8, 9]. Note that, irrespective of these results, we believe that the PSP is based on many sound learning principles and may be useful for several purposes. In fact, the PSP and the lessons-learned sessions that we introduce in our study share many of the same learning principles, e.g., the tracking and feedback of estimates and actual effort, the tracking of estimation problems, and emphasis on reflections on how to avoid the problems in the future.

There has been much work on learning from experience in software organizations, e.g., work on experience databases [10-13] and work on project experience reviews [14, 15]. However, none of these studies report results on how the use of different types of lessons-learned processes improves judgment-based effort estimation and few of them compare the learning effect with on-the-job learning. The comparison with on-the-job learning is, we believe, essential if we are to avoid attributing to the introduced change in learning processes learning that would have happened without it. Therefore, in our study, we decided to compare the improvement in learning that followed lessons-learned sessions with that of a control group that represented on-the-job learning.

The study reported in [16] compares the accuracy of the estimation of software development effort of those that receive explicit feedback that compares the estimated and the actual effort (feedback on outcome), with that of a non-feedback group. The results from that study suggest that estimators that received feedback on outcome produced estimates that were no more accurate than those that did not receive such feedback. Studies on human judgment from other domains support this disappointing lack of improvement in estimation accuracy as a result of feedback on outcome; see, for example [17-19]. The unsatisfactory effect of feedback on outcome motivated us to emphasize lessons-learned sessions, i.e., a process that includes more analysis than examination of the outcome feedback alone. The structure we imposed on these lessons-learned sessions (see Section II) is based on the finding that information on how different events and variables are related to effort over-runs are reported to be required for improved judgment accuracy in several domains [17-20].

Lessons-learned sessions should be designed to solve problems pertaining to learning from experience that are typically part of on-the-job-learning, such as those reported in [21]:

- High time pressure. The next task should be started as soon as possible.
- Low reflection on mistakes.

- Lack of objective feedback.
- Lack of knowledge about how to summarize lessons learned.
- Perception of the task as unique and not relevant to learn from.

The lessons-learned sessions we introduce in Section II address at least three of these problems. They remove the high time pressure, force the software professionals to reflect on mistakes, and provide some relevant objective feedback. It is however possible that the two remaining problems, i.e., the lack of knowledge about how to summarize lessons learned and the perception of the task as unique and not relevant to learn from, are not addressed sufficiently. If these two problems are essential to achieve positive effect from lessons-learned sessions, if there is insufficient objective feedback, or if there are other learning problems that are not addressed by lessons-learned sessions, we may observe no or even negative learning from our lessons-learned sessions.

The remainder of this paper is organized as follows: Section II describes the design of the study, including the format of the lessons-learned sessions. Section III describes the results of the study. Section IV discusses limitations of the study and tries to explain the findings. That section includes a description of a follow-up study that was intended to shed more light on the findings. Section V concludes.

## II. DESIGN

### A. *Main Research Question*

The main research question addressed in this study is the following:

*To what extent do lessons-learned sessions (of the type implemented in this study) lead to improved accuracy of effort estimation and improved realism of uncertainty assessment in comparison to on-the-job learning?*

### B. *Structure of Lessons-Learned Sessions*

The lessons-learned sessions evaluated in this study contain the following elements, provided immediately after the completion of each development task:

- Feedback on the estimated effort, the actual effort, and the estimation error of the task just completed.
- Feedback on the mean estimation error of all preceding tasks.
- Feedback on the realism of the uncertainty assessments (the format of this feedback is described in Figure 1 and described in more detail in Section II.C).
- A required 30-minute learning session in which the following information should be provided by the software professional:
  - Perceived reasons for good/poor estimation performance.
  - Perceived realism of confidence in estimation accuracy.
  - Perceived reasons for realism/lack of realism of confidence in estimation accuracy.
  - Suggested lessons learned that are relevant for estimation and uncertainty assessment.

- Suggestions on how the lessons learned should affect the estimation and uncertainty assessment of future task(s).

**Figure 1: The Format of the Lessons Learned Session**

Task X		Comments
Estimated effort	Hour : Min 0 : 0	<b>1) Reasons for good/poor estimation performance on the last task.</b> If you considered your effort estimate on the just completed task to be inaccurate, provide reasons for the estimation error. If, on the other hand, you considered your effort estimate accurate, provide reasons for accurate/accurate effort estimate: Write comments here
Actual effort	Hour : Min :	
Estimation error		<b>2) Evaluation of the uncertainty of effort estimates (Interval 1, Interval 2, Interval 3)</b> <b>a) How realistic do you think your probabilities on Interval 1, 2 and 3 were on the task?</b> Write comments here
Your estimated likelihood of A inside the 90-120% interval (Interval 1): A actually inside 90-110% (yes/no)	[ , % ]	
Your estimated likelihood of A inside the 60-150% interval. (Interval 2): A actually inside 60-150% (yes/no)	[ , % ]	
Your estimated likelihood of A inside 50-200% interval. (Interval 3): A actually inside 50-200% (yes/no)	[ , % ]	<b>b) What were important reasons for realism/lack of realism?</b> Write comments here
Average estimation error:		<b>3) Learning from previous experience</b> <b>a) What have you learned, relevant for estimation and uncertainty assessment, from the</b> Write comments here
		<b>b) How should this impact the estimation and uncertainty assessment of the next task(s)?</b> Write comments here

Notice that our lessons-learned session raises the same type of questions as in common industrial project review meetings. The post-mortem reviews at Microsoft, for example, raise the questions: “What worked well in the last project, what went wrong, and what should the group do to improve in the next project?” [22]. An essential difference from many other types of project review is that we focus solely on individual learning in our study, while most other project reviews focus on individual learning, group-based learning, and experience sharing. Our focus on individual learning is similar to that of the PSP.

It can be argued that reflecting on and learning from one’s own work is the basis of many learning processes. If individual learning is not present, there is frequently not much experience of value to share either. However, the focus on individual learning in this study also means that the potential improvement in estimation accuracy that may result from group-based learning and experience sharing in lessons-learned sessions is not evaluated.

### C. Measures

The measures used to evaluate estimation accuracy are the following:

$$\text{Magnitude of Relative Error} = \text{MRE} = |\text{actual effort} - \text{estimated effort}| / \text{actual effort}$$

$$\text{Relative Error} = \text{RE} = (\text{actual effort} - \text{estimated effort}) / \text{actual effort}$$

The median MRE is our main measure of estimation accuracy. The lower the median MRE, the better is the estimation accuracy. We use the median RE as our main measure of estimation bias. A positive median RE indicates a tendency towards over-optimism, a negative median RE a tendency towards over-pessimism, and a RE close to zero unbiased estimates. From previous experience, we expected that there would be a few MRE and RE-values (outliers) with a strong impact on the mean MRE and mean RE, i.e., that the mean MRE and RE would be misleading as the central value of a set of observations. That motivates our use of the more outlier-robust median MRE and RE.

The MRE and RE have their limitations and a lot of alternative accuracy measures have been proposed. For an overview and discussion, see [23]. We tested different variants of these accuracy measures on the data collected in this study, e.g., more symmetric accuracy measures, but found that they did not change the main results. We therefore decided to use the measures that are best known by the software engineering community, i.e., MRE and RE.

We measure the assessed uncertainty of an effort estimate by means of an effort prediction interval. An effort prediction interval is the combination of a stated confidence level (CL) and an effort minimum-maximum interval (EI). For example, a software developer may estimate that the most likely effort of a development task is 100 work-hours and that it is 80% probable that the actual effort will be between 50% (50 work-hours) and 200% (200 work-hours) of the estimated most likely effort. Then, the minimum-maximum interval [50; 200] work-hours is the 80% confidence, prediction interval of the developer's effort estimate of 100 work-hours. In order to ease the analysis in our study, we standardized on the following three effort minimum-maximum intervals:

$$EI-1 = [90\% * \text{estimated effort}; 110\% * \text{estimated effort}]$$

$$EI-2 = [60\% * \text{estimated effort}; 150\% * \text{estimated effort}]$$

$$EI-3 = [50\% * \text{estimated effort}; 200\% * \text{estimated effort}]$$

These three effort minimum-maximum intervals were selected to reflect what we believed were typical narrow, medium-wide, and wide minimum-maximum effort intervals. For each task, we asked the developers to assess how confident he was, i.e., how probable he thought it was that the actual effort would fall inside each of these three minimum-maximum intervals. If, as in the previous example, a software developer estimated that the most likely effort was 100 work-hours, he would be asked to provide the (subjective) probabilities that the actual effort would

fall inside the intervals [90; 110] work-hours (EI-1), [60; 150] work-hours (EI-2), and [50; 200] work-hours (EI-3). The developer may, for example, believe that it is 50% probable that the actual effort will be inside EI-1, 70% probable that it will be inside EI-2, and, 95% probable that it will be inside EI-3. The confidence level (CL) of EI-1 is consequently 50%, the CL of EI-2 is 70% and the CL of EI-3 is 95%.

The frequency of including the actual effort in an effort minimum-maximum interval is termed the “hit rate”. On the long run, i.e., when estimating and assessing the uncertainty of many tasks, a developer who has realistic uncertainty assessments will have a mean confidence level that is similar to his hit rate. If his mean confidence level is higher than his hit rate, we describe the developer’s uncertainty assessments as over-confident. If it is lower than the hit rate, we describe the developer’s uncertainty assessments as under-confident. In earlier studies, e.g., [3], we have documented a strong tendency towards over-confidence in the accuracy of the effort estimates.

We measure the level of over-confidence of an effort minimum-maximum interval, e.g., EI-1, and a set of tasks as follows:

$$\text{OverConfidence} = (\text{mean confidence level} - \text{“hit rate”}) / \text{“hit rate”}$$

An OverConfidence value close to zero indicates that the uncertainty assessments, on average, are realistic, a negative value that they are under-confident, and a positive value that they are over-confident. Assume, for example, that a developer has provided the confidence levels (subjective probabilities) 80%, 60%, 90%, 90%, and 80% related to the EI-1 effort interval, for the estimates of the tasks T1...T5, respectively. The mean confidence level of the developer’s estimates related to EI-1 and T1...T5 is then  $(80\%+60\%+90\%+90\%+80\%)/5 = 80\%$ . If the developer’s uncertainty assessments are realistic, we would expect a hit rate of 80%, i.e., that four out of the five estimates lay inside the EI-1 minimum-maximum effort intervals. If the EI-1 minimum-maximum effort intervals, for example, included only one out of the five actual effort values, i.e., if we observed a hit rate of only 20%, his level of over-confidence would be  $(80\% - 20\%)/20\% = 3.0$  and would suggest a high level of over-confidence in the accuracy of his estimates. Notice that our measure of over-confidence may not work well as an indicator of over-confidence for as small a set of tasks as in the above example. Consequently, we will only use this measure as an indicator of differences in levels of confidence between groups that include the uncertainty assessments of larger sets of tasks.

The defined measure of over-confidence only indicates the ability of developers to assess the mean level of estimation uncertainty of a set of tasks. Also relevant is the ability to distinguish between high and low uncertainty effort estimates in a set of tasks. For example, it is possible that the developers are strongly overconfident about the uncertainty of their effort estimates but are nevertheless able to distinguish between high and low uncertainty effort estimates. To assess the ability to distinguish between high and low uncertainty effort estimates, we apply the rank order correlation (Spearman rank order correlation) between confidence level and the estimation accuracy (MRE) for a minimum-maximum interval, e.g., EI-1, and a set of tasks:

*CorrConfAcc = rank order correlation between confidence level and MRE*

This measure uses the level of confidence as the uncertainty of an estimate as perceived by the software professional and the MRE-value as a substitute of the actual uncertainty of the effort estimate; i.e., the measure is based on the reasonable assumption that the higher the actual uncertainty of an effort estimate, the higher the typical estimation error will be. If a developer is able to rank his estimates in relation to the degree of uncertainty perfectly, we would expect a correlation of -1. Positive or low correlations would suggest that for the evaluated set of tasks, the developer is not good at separating effort estimates with high and low uncertainty.

#### *D. Research Hypotheses*

We designed the study to test the following four hypotheses:

**Hypothesis 1:** Lessons-learned sessions improve estimation accuracy in comparison with on-the-job learning (lower median MRE).

**Hypothesis 2:** Lessons-learned sessions reduce estimation bias in comparison with on-the-job learning (median RE closer to 0).

**Hypothesis 3:** Lessons-learned sessions improve the realism of uncertainty estimation in comparison with on-the-job learning (OverConfidence closer to 0).

**Hypothesis 4:** Lessons-learned sessions improve the ability to separate low and high uncertainty estimates in comparison with on-the-job learning (CorrConfAcc closer to -1).

#### *E. The Research Design*

##### *1) Subjects*

The subjects were recruited via a request for consultants that was sent to Norwegian consulting companies. The request specified a flexible range of time for which the consultants would be needed, along with the required education and expertise. Companies replied with *curricula vitae* of potential candidates and these were then screened to verify that they complied with the requirements. The subjects were required to at least have a Bachelor's

degree in informatics (or equivalent) and familiarity with the technology (UML, Struts, JSP, Java, HTML, the Eclipse IDE, and MySQL) of the system on which they were supposed to complete development tasks.

In total, 20 software professionals (all male) were selected. The software professionals were paid close to ordinary fees for their work and asked to treat the development work as ordinary consultancy work. The software professionals did not know that developers other than themselves were asked to complete the tasks.

## 2) *Tasks*

The experiment was conducted on the BESTweb system [4]. The BESTweb system is a web-based system developed in Java using the Struts framework. The system supports the research on software cost and effort estimation through a database front-end client that gives access to information about journal and conference papers on software effort estimation. An essential part of the system's functionality is the identification of relevant papers through the use of categories related to, for example, type of estimation methods. As an illustration of the size of the software, the BESTweb software consists of about 50 classes and 3000 lines of Java code. The BESTweb system can be accessed at [www.simula.no/BESTweb](http://www.simula.no/BESTweb).

All software professionals were instructed to estimate and complete the same five tasks in the same sequence. Below is a brief description of the development tasks:

- **Task 1:** Requires the addition of functionality to save a user's search query to persistent memory. In addition, the user's last search query must be displayed and re-executed automatically when they next log on to the system.
- **Tasks 2:** Requires that the system be extended to handle an additional piece of data from an input file (in XML format) used to update the publications in the BESTweb system. The system already partially handles the data: if it encounters the presence of the data in the file it warns the user that that data is not supported. The developer is asked to add support for this data by extending the domain model, the GUI, and the search functionality.
- **Task 3:** Requires the developer to add functionality to the system that extends the manner in which cost and effort estimation metadata associated with each publication are dealt with, specifically, the ability to add publication categories and corresponding codes.
- **Task 4:** Requires the developer to add caching logic to the system so that if statistics for all the publications in the system are requested, the cached results are used (so as to decrease the computational load on the system).
- **Task 5:** The developers are asked to add functionality such that the users could delete existing publication codes from the system.

## 3) *Experiment Process*

The software professionals were allocated randomly, by the toss of a coin, to the Control or Learning group. The groups were the same size, 10 software professionals in each group. Following their allocation to one of the groups, the software professionals followed these steps:

1. They received a description of the work process to be followed.

3. They completed an initial questionnaire that captured professional background and experience.
4. They received a task specification, starting with Task 1.
5. They estimated the most likely effort needed to complete the specified task and the assessment of the perceived uncertainty of that effort estimate, i.e., the subjective probabilities (confidence levels) of including the actual effort in the minimum-maximum intervals EI-1, EI-2 and EI-3.
6. They described, in brief, the strategy for the estimation and the assessment of uncertainty that they used.
7. They performed the task (design, programming, testing, and documentation).
8. Upon completion of the task, it was sent in for acceptance testing. The system was then tested using a pre-defined system acceptance test plan:
  - a. If the test fails, the software professional is told the problem and asked to fix it and to submit the solution again.
  - b. If the test passes and the software professional belongs to the Learning group, he receives the feedback on estimation accuracy and spends about 30 minutes on lessons learned. Then, he receives the next task, i.e., repeats the process from Step 4.
  - c. If the test passes and the software professional belongs to the Control group, he receives the next task, i.e., repeats the process from Step 4.
9. Upon completion of all five tasks, the developers were debriefed about the experiment and provided tape-recorded responses to questions about estimation strategies and learning processes used.

### III. RESULTS

#### A. *Descriptive Characteristics*

A random allocation of as few as 20 subjects into two groups of the same size does not guarantee that the groups are similar with respect to estimation-relevant background and skill. Fortunately, we observed no essential differences, as illustrated in Table 1.

**Table 1: Comparison of the Subjects' Backgrounds**

<b>Variable</b>	<b>Group</b>	<b>Mean</b>	<b>Lower Quartile</b>	<b>Median</b>	<b>Upper Quartile</b>
Experience total (years)	Control	7,2	5,0	6,0	8,0
	Learning	7,0	4,0	5,5	10,3
Degree (1= Bachelors, 2= Masters)	Control	1,6	1,0	2,0	2,0
	Learning	1,9	2,0	2,0	2,0
Average grade in computer science courses	Control	2,0	1,7	2,0	2,4
	Learning	2,0	1,9	2,1	2,5
Self-assessed estimation competence (1=much worse than average, 5=much better than average)	Control	3,2	3,0	3,0	3,0
	Learning	2,9	3,0	3,0	3,0
Self-assessed average estimation error on previous tasks	Control	28%	19%	20%	43%
	Learning	25%	10%	23%	35%
Self-assessed programming skill (1=much worse than average, 5=much better than average)	Control	3,3	3,0	3,0	4,0
	Learning	3,2	3,0	3,0	4,0

Table 2 extends the analysis of potential differences between the groups by comparing the performance on Task 1, i.e., the performance before the first exposure to the lessons-learned session.

**Table 2: Comparison of Effort, Quality, Estimation Accuracy and Bias of Task 1**

Variable	Group	Mean	Lower Quartile	Median	Upper Quartile
Effort (work-hours)	Control	5,8	4,1	5,8	7,4
	Learning	6,0	4,1	4,8	9,1
Programming quality (mean number of corrections)	Control	1,4	0,8	1,0	2,3
	Learning	1,5	0,0	1,0	3,0
Estimation accuracy (MRE)	Control	1,4	0,4	0,6	1,9
	Learning	0,8	0,3	0,7	1,2
Estimation bias (RE)	Control	0,11	-0,3	0,2	0,4
	Learning	0,05	-0,4	0,2	0,4

It seems to be safe to claim, on the basis of Table 2, that both the effort and quality indicators of Task 1 were similar, which indicates that there were no large differences in programming skill. If anything, the lower median effort of the Learning group on Task 1 may suggest that those in that group were slightly more efficient programmers. The difference is, however, not large and statistically non-significant. The higher mean MRE of the Control group is due mainly to one outlier and the small difference in median values is probably more representative for the difference in estimation accuracy between the two groups.

A comparison of the uncertainty assessment related values for IE-1, IE-2 and IE-3 for Task 1, see Table 3, shows that there were no large difference here, either.

**Table 3: Comparison of Uncertainty Assessment Values of Task 1**

Variable	Group	Mean Confidence Level	Mean Hit Rate	OverConfidence
Over-confidence IE-1	Control	52%	20%	1,6
	Learning	57%	20%	1,9
Over-confidence IE-2	Control	77%	50%	0,5
	Learning	80%	50%	0,6
Over-confidence IE-3	Control	96%	90%	0,7
	Learning	96%	90%	0,7

An informal analysis of the described strategies for estimating and for assessing uncertainty (Step 6 of the experiment process) used for Task 1 showed no essential differences. All of the developers relied on expert judgment supported by a separation of the task into sub-activities, i.e., a bottom-up strategy. The processes used to assess uncertainty were poorly described and difficult to evaluate. This may indicate that the software professionals

had no explicit strategy for uncertainty assessments, i.e., that the process was based more or less on expert judgement.

## B. Estimation Accuracy and Bias

### 1) Hypothesis 1

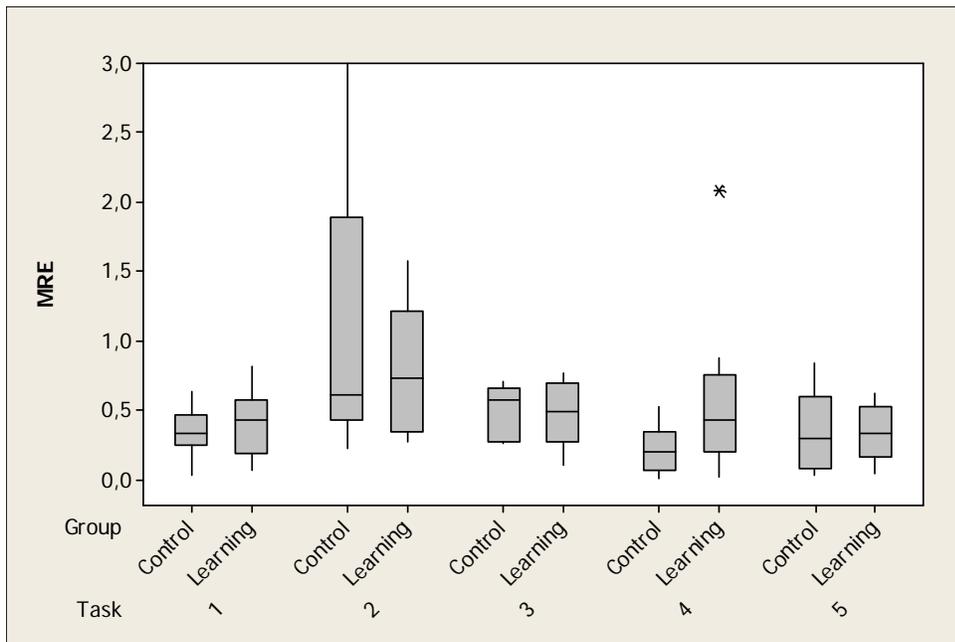
Our first hypothesis states that the Learning group will improve the accuracy of their estimates more than the Control group, which will be indicated by the median MRE of the estimates being lower for Tasks 2 to 5 (the tasks following the first lessons learned session). Table 4 and Figure 2 show that this was hardly the case. The median MRE of the Learning and the Control group across Tasks 2 to 5 are about the same (0,50 vs 0,49). A Kruskal-Wallis test on difference in median MRE of the two groups for Tasks 2 to 5 gives  $p=0,88$ .

Table 4 and Figure 2 also show that the MRE decreases over time for both groups. In the absence of a proper control group, as is the case in typical field settings and poorly designed experiments, the observed improvement in MRE may falsely have been credited to the lessons-learned sessions. This supports our emphasis on introducing proper control groups in effect studies of this type. Notice also that a control group, as we use it here, enables much stronger cause-effect analyses than the introduction of a so-called baseline typically recommended for field studies, see for example [7]. While baseline data are typically collected before the process change take place, our control group assumes a random allocation of process change and the parallel use of old and new processes. This means that any improvement that is measured in comparison to the baseline data must be able to isolate the effect of the process change from all other changes that took place in the relevant period. In our experience, this is a very difficult analysis.

As can be seen in Figure 2 the *variation* of MRE is not systematically lower in the Learning group. This result differs from the findings reported in [9], where the use of the Personal Software Process on students did lead to lower variance in estimation error (although not lower mean estimation error). This difference in results from those reported in [9] may have been caused by the use of formal effort estimation models when applying the Personal Software Process, as opposed to the use of judgment-based effort estimation processes that was the case in our study. As pointed out in our review comparing models and expert judgment (see [5]), estimation models sometimes reduce the number of very large estimation errors and consequently the variance in estimation error. An increased level of consistency is also expected from the use of formal estimation models.

**Table 4: MRE (Tasks 2-5)**

Task	Group	Lower Quartile	Median	Upper Quartile
2	Control	0,4	0,6	1,9
	Learning	0,3	0,7	1,2
3	Control	0,3	0,6	0,6
	Learning	0,3	0,5	0,7
4	Control	0,06	0,2	0,3
	Learning	0,2	0,4	0,8
5	Control	0,08	0,3	0,6
	Learning	0,2	0,3	0,5

**Figure 2: Boxplot of MRE (All Tasks Included)**

## 2) Hypothesis 2

Our second hypothesis states that the Learning group's estimates would be less biased, which would be indicated by the RE of the estimates provided by those in the Learning group being closer to zero than those in the Control group for Tasks 2 to 5. Table 5 and Figure 3 show that this was not the case. The median RE of the Learning and the Control group turned out to be almost the same (-0,01 vs 0,00), which shows that the median effort estimate was unbiased in both groups. Unbiased effort estimates are not untypical when the tasks to be performed are small and

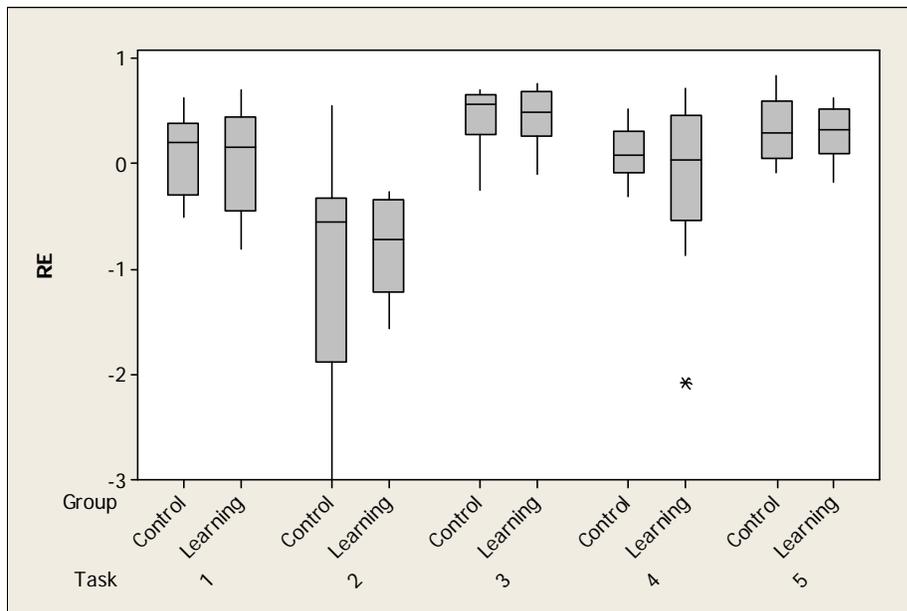
payment is made on a per work-hour basis [6], as was the case in our study. A Kruskal-Wallis test on difference in median RE of the two groups for Tasks 2 to 5 gives  $p=1,0$ .

Notice the zigzag pattern of the RE in Figure 3. This pattern suggests that over-optimism on one task (e.g., Tasks 1) is easily followed by over-pessimism on the next task (e.g., Tasks 2). If this suggestion is correct, it may be that there is a tendency to over-react to the experience of the task immediately before the one to be estimated.

**Table 5: RE (Tasks 2-5)**

Task	Group	Lower Quartile	Median	Upper Quartile
2	Control	-1,9	-0,6	-0,3
	Learning	-1,2	-0,7	-0,4
3	Control	0,3	0,6	0,6
	Learning	0,3	0,5	0,7
4	Control	-0,08	0,08	0,3
	Learning	-0,5	0,04	0,5
5	Control	0,05	0,3	0,6
	Learning	0,1	0,3	0,5

**Figure 3: Boxplot of RE (All Tasks Included)**



### C. Uncertainty Assessment

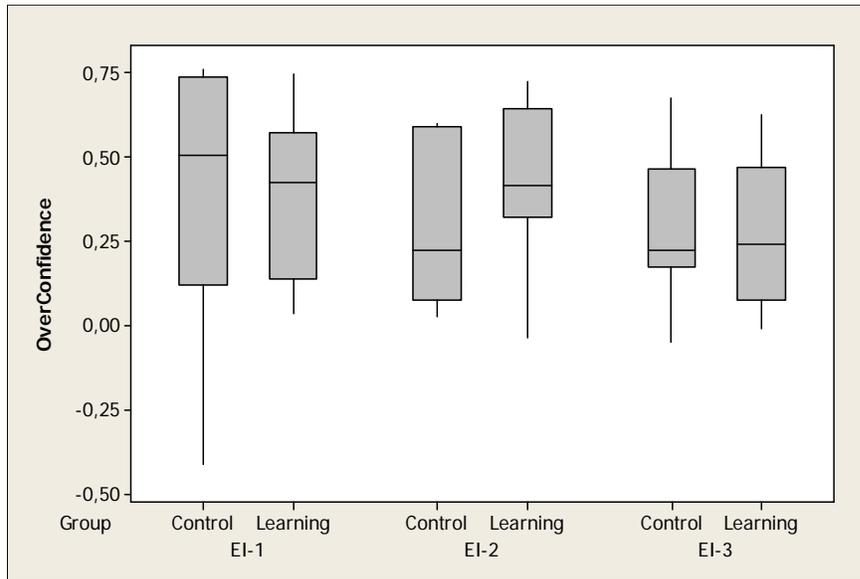
#### 1) Hypothesis 3

Our third hypothesis states that the Learning group would improve the realism in their uncertainty assessment, which would be indicated by the OverConfidence of those in the Learning group being closer to zero than that of the Control group. We considered the number of observations to be too low for meaningful use of the OverConfidence measure on individual tasks; see the discussion in Section II.C. That being so, the values presented are based on the distribution of OverConfidence of the developers on the combined set of Tasks 2 to 5.

Table 6 and Figure 4 illustrate that the Learning group did not make more realistic uncertainty assessments than the Control group. Both groups were strongly over-confident. A Kruskal-Wallis test of difference in median OverConfidence values for IE-1, IE-2 and IE-3 gives the p-values 0,6, 0,1 and 0,9, respectively. In other words, the only difference in over-confidence between the groups with a low p-value, i.e., IE-2, is in favour of the Control group. If anything, the lessons-learned sessions seem to have resulted in an increase, rather than reduction, in the level of the participants' over-confidence in the accuracy of their estimates.

**Table 6: Over-confidence (Tasks 2 to 5)**

Variable	Group	Mean	Lower Quartile	Median	Upper Quartile
OverConfidence EI-1	Control	0,4	0,12	0,5	0,7
	Learning	0,4	0,14	0,4	0,6
OverConfidence EI-2	Control	0,3	0,08	0,2	0,6
	Learning	0,4	0,3	0,4	0,6
OverConfidence EI-3	Control	0,3	0,2	0,2	0,5
	Learning	0,3	0,08	0,2	0,5

**Figure 4: Boxplots of OverConfidence (Tasks 2 to 5)**

## 2) Hypothesis 4

Our fourth hypothesis states that the Learning group would be better able to separate low and high uncertainty estimates than the Control group. We measure this ability through the Spearman rank order correlation for each developer and each type of minimum-maximum interval, i.e., three correlations per developer. The first relative uncertainty evaluation is on Task 2, where the confidence in the accuracy of Task 2 is compared with the confidence in the accuracy of Task 1. Given that the first relative uncertainty assessment is on Task 2, but involves Task 1, we include all five tasks in the correlation analysis. The presented median correlation of each group's minimum-maximum interval, e.g., the EI-1 interval, is consequently based on 10 individual correlation coefficients.

Descriptive statistics for the rank order correlations are displayed in Table 7. We ranked the MRE and confidence values so that the highest values get the highest ranks. This means, as described earlier, that we should expect strong negative correlations if the developers are good at separating high and low uncertainty tasks, i.e., when a high confidence in the accuracy of an estimate correlates with a low estimation error (MRE). A perfect correlation between confidence and estimation error gives the value -1.

**Table 7: Rank Order Correlation between Confidence and MRE**

<b>Correlation between MRE and ...</b>	<b>Group</b>	<b>Mean</b>	<b>Lower Quartile</b>	<b>Median</b>	<b>Upper Quartile</b>
... confidence in EI-1	Control	-0,3	-0,7	-0,3	-0,1
	Learning	0,12	-0,3	0,4	0,5
... confidence in EI-2	Control	0,2	-0,3	0,3	0,7
	Learning	-0,04	-0,3	0,09	0,3
... confidence in EI-3	Control	0,09	-0,4	-0,05	0,7
	Learning	0,16	-0,3	0,18	0,7

An analysis of the median values suggests that the Learning group are somewhat worse than the Control group for EI-1 and EI-3, but better for EI-2. However, the general conclusion is that the software professionals, regardless of group, are in general not very good at separating high and low uncertainty effort estimates. If anything, the Control group participants were slightly better. This lack of ability to separate high and low uncertainty effort estimates is not only a result of low ability, but also a result of the similarity of the tasks. If we had introduced less homogenous tasks, it is likely that the ability to separate high and low uncertainty tasks would improve substantially; see, for example, [3].

#### IV. DISCUSSION

Our results show that the introduced lessons-learned session did not improve the learning in comparison with on-the-job learning. While our study may be the first on lessons-learned-based improvement of judgment-based effort estimation and uncertainty assessment, there are several other studies that report a similar lack of positive effects from lessons-learned sessions in other software project contexts, e.g., [24-26], and other project domain contexts, e.g., [27]. Shortcomings of software professionals' analyses of reasons for errors of estimation, such as the tendency to neglect indirect and contributing reasons and the tendency towards a biased attribution of causes (such as the tendency to claim that success factors are controlled by oneself, while failures are due to external events [28]), may contribute to this lack of learning from lessons learned sessions.

Our results and those of previous studies on the use of lessons-learned sessions should, of course, not be used to claim that lessons learned sessions will never result in improved estimation accuracy or more realistic assessment of uncertainty. There are many ways of designing lessons-learned sessions. It may well be that imposing other structures, other types of feedback, better training in advance of the lessons-learned sessions, or other measures, will yield different results. To understand how to improve lessons-learned sessions and other learning processes, it may be useful to look at possible reasons for the observed lack of improved learning from the lessons-learned session in our study. After all, to give up on learning from previous experience of estimation is not an attractive option.

To structure this discussion, we will analyse the following three types of possible reasons for the observed lack of improved learning:

- Limitation of the chosen study design. (Section IV.A)

- Limitations of the lessons learned process studied. (Section IV.B)
- Observed problems related to learning from experience of the estimation of software development effort and the assessment of uncertainty. (Section IV.C)

We will apply the qualitative data that we collected as part of our study to support this analysis. We collected this data (see description in Section II) immediately after each estimate (where all the software professionals described their estimation processes) and as part of the lessons-learned sessions (where the software professionals in the Learning group summarized their reasons for estimation errors and what they had learned). In addition, we will describe and apply the results from a follow-up experiment. In this experiment, other software professionals were asked to assess the uncertainty of the estimates provided in the main study. The purpose of the experiment was to examine whether there is a difference in realism between assessing the accuracy of one's own and estimates and those of others, which would enable us to determine whether some of the observed problems with learning were due to the learning from one's own estimation performance as opposed to that of others.

#### A. *Limitations of the Study Design*

We assess that the major limitations of our study design are as follows:

- *Individual learning*: The software developers worked and learned individually, not as part of a development team. For example, two of the developers remarked in the debriefing session that in an ordinary work context they would have asked for some support by a colleague on at least one of the problems. Instead, they had to spend time to figure out the problem by themselves. This means that our results may mainly be representative for individual learning and working, not so much for team learning. However, as argued earlier in this paper, individual learning is an essential component of group learning and consequently of relevance, too.
- *Pseudo-realistic context*: The software professionals were paid close to their ordinary fees, instructed to behave as close to normal as possible, and did not know that several others were completing the same tasks. However, the extensive logging (data about themselves, estimation strategy used, detailed time logging, etc.), the lessons-learned sessions (for the Learning group), and the fact that we as researchers were the clients meant that the situation was hardly perceived to be the same as in ordinary software development work. The software professionals expressed different opinions about the effect of the unusual process elements in the debriefing session. Our main impression from observing the work and analysing the debriefing interviews is that, in spite of the unusual process elements, the estimation and programming work had the same challenges as ordinary programming work and hence that the differences from ordinary work situations had only small, if any, effects on the learning processes.
- *Representativeness*: The high cost of hiring software professionals placed restrictions on the number of tasks that we could assign and the number of subjects that we could use. This means that we should be careful when extrapolating the results to tasks and situations that are different from those we have studied.

Notice that many of the unusual process elements of the study may have led to a stronger focus on learning than in typical field situations. Thus, we find it difficult to see how individual, lessons-learned session-based learning would lead to improvements in typical field situations similar to the one we studied if it did not succeed here. Of

course, this belief assumes that the software professionals were motivated to learn. We find that this assumption is likely to be true, judging from our observation of their behaviour and the documents that were produced in the lessons-learned sessions. As an illustration, the software professionals wrote, on average, about 1,000 words each when describing lessons learned (about 200 words per lessons-learned session). The amount of lessons-learned work was impressive. It was also of good quality (more on this in the next section). Consequently, despite the limitations of the study, we interpret the findings of the study as describing a situation in which those in the Learning group really tried to use the learning sessions to learn, but nevertheless did not improve effort estimation accuracy more than the Control group, and were just as overconfident in the accuracy of their effort estimates.

### *B. Limitations of the Learning Process*

To examine what actually took place in the lessons-learned session we examined the provided:

- Reasons for good/poor estimation performance and uncertainty assessment. (IV.B.1)
- Lessons learned relevant for estimation and uncertainty assessment and how the lessons learned may affect the estimation and uncertainty assessment of future tasks (IV.B.2)

#### *1) Reasons for Accuracy/Inaccurate Estimates and Realistic/Unrealistic Uncertainty Assessments*

The reasons described by the developers were as expected from previous studies on this topic; see [28] for an overview of typically provided reasons. The most common responses related to reasons for inaccurate estimates in the current study were:

- Spent too little time on effort estimation and uncertainty assessment work
- Too little knowledge about the problem or the technology
- Unexpected events or problems
- Forgotten/strongly underestimated activities (such as documentation)
- Fewer and less severe problems than expected (led to over-estimation)
- More complex task than expected
- Task specification unclear or misunderstood the specification
- Error corrections needed
- Poor impact analysis of consequences of code changes
- Design errors made
- Incorrect assumptions about the code

Typical responses related to reasons for accurate estimates were:

- Good impact analysis of consequences of code changes
- Task similar to the one previously completed (on the same system)
- Good understanding
- Luck

- Simple task and no problems

When we asked for reasons for realistic/unrealistic uncertainty assessments, we expected to be provided with reasons that were directly related to the uncertainty, such as insufficient risk analysis and lack of learning from the actual uncertainty of the previous effort estimates. Instead of these types of reason, the reasons for realistic/unrealistic uncertainty assessments provided by the developers were, almost without exception, related to the accuracy/inaccuracy of the effort estimate. This suggests that they believed that the dominant means to improve the realism of the uncertainty assessment was to improve estimation accuracy. The alternative, which would have been to change their level of confidence in the effort estimates, was hardly considered. An analysis of the confidence levels supports this suggestion. The analysis indicated that the software professionals changed their levels of confidence only a little in response to feedback about the accuracy of their previous estimate.

## 2) *Lessons Learned*

The main estimation accuracy related lessons learned provided by those in the Learning group were the following:

- Spend more time on estimation work (including a more detailed impact analysis).
- Add more time for unknown events.
- Add more time for debugging, testing and error correction.
- Add more time for documentation.
- Assess more carefully the ripple effects of a change.
- Add more time when the task involves changes many different places in the code (“distributed” task)
- Read the requirement specification and system documentation more carefully and check the assumptions made.
- Pay special attention to the parts where the knowledge and experience is low.
- Trust one’s instincts more.
- Expect unexpected problems.
- Control the use of effort better (do not do more than that which is required by the task specification).

All these lessons learned were reasonable and have the potential to improve estimation accuracy. In that sense, our lessons-learned session process was a success in relation to the estimation of the most likely effort. The same cannot be said in relation to the lessons learned regarding the assessment of uncertainty. In spite of a direct instruction to summarize the lessons learned that were related to uncertainty assessment, we found, as noted in Section IV.B.1, no lessons learned that could be interpreted as practical changes in the process of assessing uncertainty.

We find it especially strange that the software professionals did not update their levels of confidence in response to feedback strongly suggesting over-confidence. Consider the following example. One software developer had assessed the probability of including the actual effort in the minimum-maximum interval IE-1 ([90% of estimate; 110% of estimate]) to be about 50% of all the first three tasks. The developer then received feedback that told him

that none of the IE-1 intervals had included the actual effort. A proper lesson learned would then be for him to admit that he had a tendency toward overconfidence and that he had to be less confident that he would include the actual effort in IE-1 of future estimates. Instead, we observed that he typically responded with no or only a minor small adjustment and a strong emphasis on what he needed to do to achieve more accurate effort estimates. We discuss possible reasons for this unwillingness to update confidence levels in spite of accurate, relevant, and timely data on estimation error in [29, 30]. Possible reasons include: i) the software professionals do not have proper, probabilistic mental models to enable proper learning and, ii) there is a conflict between the developers' self-images of being predictable and skilled developers and providing realistic uncertainty assessments that suggest the opposite (the so-called "cognitive dissonance" effect.) We discuss these two reasons in more detail in Section IV.C.

Given the intuitively meaningful lessons learned that were provided by the developers in relation to the estimation of most likely effort, why did the accuracy of their estimates not improve more than that of the estimates of those in the Control group? Possible reasons for this include: (i) the software developers did not translate the lessons learned into practical actions, (ii) the Control group developers made similar progress in learning without the lessons-learned session, or (iii) there were negative consequences of the lessons learned that cancelled out the positive ones.

It is not easy to assess the importance of these possible reasons in isolation. The estimation strategy descriptions provided by the software professionals did not contain as much valuable information about the practical use of the lessons learned in the estimation work as we had hoped for. Nevertheless, it was clear that at least a few of the software professionals actually applied the lessons learned. As an illustration, when a developer stated that he needed to spend more time on the effort estimation work, we observed that he typically did so.

We counted the number of references to previous experience in the software professionals' description of the processes that they used for estimating and assessing uncertainty (Step 6 in the experiment process) and found no large differences in the number and/or types of references. This similarity in the amount of use of experience, together with the observed similarity in the improvement in the estimates of the Control and the Learning group, do provide some support for reason (ii) above. Those in the Control group seemed to learn quite a lot from experience without spending time on lessons-learned sessions.

We also found some evidence in support of reason (iii), namely that the software professionals in the Learning group may have over-reacted more strongly to the experience of the previous task. While the median absolute difference in RE between two tasks for those in the Control group was 0,6, the corresponding value of those in the Learning group was 0,8. A high difference in median RE is, for example, a result of going from strong under-estimation to a strong over-estimation of effort. This observation provides (weak) support of that lessons learned may lead, not only to positive effects, but also to stronger over-reactions to previous experience. More studies are needed to examine whether, and if so the extent to which, this is a typical effect of lessons-learned sessions or an effect that is caused by the particular context in our study. As reported in Section I, however, our study is not the only one to observe possible negative effects of lessons-learned sessions.

### C. Learning Problems

Learning from experience is documented to be difficult, both in software development [31] and other contexts. Hammond [32, p. 278] summarizes the situation: *“Yet in nearly every study of experts carried out within the judgment and decision-making approach, experience has been shown to be unrelated to the empirical accuracy of expert judgments”*.

Reason for the learning problems includes the following:

- The “hindsight bias”, e.g., the tendency to interpret a cause-effect relationship as more obvious after it has happened than before [33, 34].
- The tendency to confirm rules and disregard conflicting evidence, as illustrated in studies on human judgment [35, 36].
- The tendency to apply a “deterministic” instead of a “probabilistic” learning model. This ability to think in probability-based terms can, according to [37] hardly be derived from experience alone, but must be taught. Hammond [32] suggests that the ability to understand relationships in terms of probabilities instead of purely deterministic connections is important for correct learning in situations in which the level of uncertainty is high. When assessing the uncertainty of effort estimates, probabilistic thinking and learning models are essential. The observed problems with proper learning from feedback on uncertainty assessment in our study may be caused by a lack of training in the use of proper probabilistic learning models.
- The high amount of complexly interconnected reasons for high or low estimation accuracy. In practice, we may need to include both system dynamics and game theory to understand the network of reasons for high or low estimation accuracy.

In many situations in which human judgment is used that have high uncertainty and unstable task relations, there are indications that even feedback of high quality, e.g., high-quality task-relation-oriented feedback, is not sufficient for learning [20, 38]. The underlying reason is that it is frequently difficult to transfer experience from one context to another. For this reason, it is important to recognize when there is nothing to learn from experience, as reported in the software estimation studies [31, 39].

On the basis of our previous experience, reported in [40], we believed that one essential reason for the sustained strong tendency towards overconfident uncertainty assessments could be related to the difference between assessing the uncertainty of one’s own and others’ effort estimates. When assessing one’s own estimate of one’s own work there may for example be, as pointed out earlier, learning problems related to the cognitive dissonance effect. This effect is, for example, present when the software professional tries to avoid a conflict between his image of himself as a skilled and predictable developer and historical data that suggests that the opposite is the case. A possible way to avoid this conflict is, for each new task, to assume either that one has learned much from previous experience or that the feedback about previous performance was, in some way, not relevant for the future. To test whether the problem of uncertainty assessment learning was related to this effect, we conducted a follow-up experiment at a seminar on effort estimation with software professionals as participants. If the assessment of other software professionals estimates, applying the same feedback and information, improves the learning, this may be used to

improve the lessons-learned sessions. Alternatively, if the learning problems remain, this supports the belief that the lack of proper probabilistic learning models is the main obstacle to uncertainty assessment learning, and, that better training in probability and statistics are needed.

### 1) *The Follow-Up Experiment*

*Participants:* Eighty-three software professionals who attended an estimation seminar at Simula Research Laboratory. These software professionals were similar to those in the main experiment with respect to length of experience and background.

*Hypothesis:* Assessment of the uncertainty of other developers' effort estimates based on feedback about the estimation error of previous estimates will be more realistic than the assessment of the uncertainty of effort estimates of one's own work. (The realism of the uncertainty assessment will, as before, be indicated by the OverConfidence measure. The hypothesis implies that the participants in the follow-up experiment achieve an OverConfidence closer to zero than that of those in the main experiment.)

*Process:* Each of the 83 software professionals:

1. Was randomly allocated to the estimates and feedback related to one of the 20 developers who participated in the main experiment.
2. Read the description of what he was supposed to do, i.e., that he was to evaluate the uncertainty of effort estimates produced by a software developer on real software maintenance tasks.
3. Read the description of the BESTweb-system. (A short version of the description given to the developers in the main experiment).
4. Read the instruction on how to assess the uncertainty of the effort estimates. This instruction was the same as the one in the main experiment.
5. Received a database with information about the estimation and uncertainty assessment of Tasks 1-3 as completed by the allocated developer in the main experiment. This database contained the task descriptions, the estimates of most likely effort, the actual use of effort values, and the estimation errors of all the three tasks.
6. Received a description of Task 4 and the allocated developer's effort estimate of that task.
7. Assessed the probability that the actual effort of the allocated developer would fall within IE-1 ([90%;110%] of the estimate), IE-2 ([60%;150%] of the estimate), and IE-3 ([50%;200%] of the estimate) for Task 4.
8. Received the allocated developer's actual use of effort of Task 4.
9. Received an updated version of the database, now including information about the first four tasks.
10. Received a description of Task 5 and the allocated developer's effort estimate of that task.
11. Assessed the probability that the actual effort would fall within IE-1 ([90%;110%] of the estimate), IE-2 ([60%;150%] of the estimate), and IE-3 ([50%;200%] of the estimate) for Task 5.

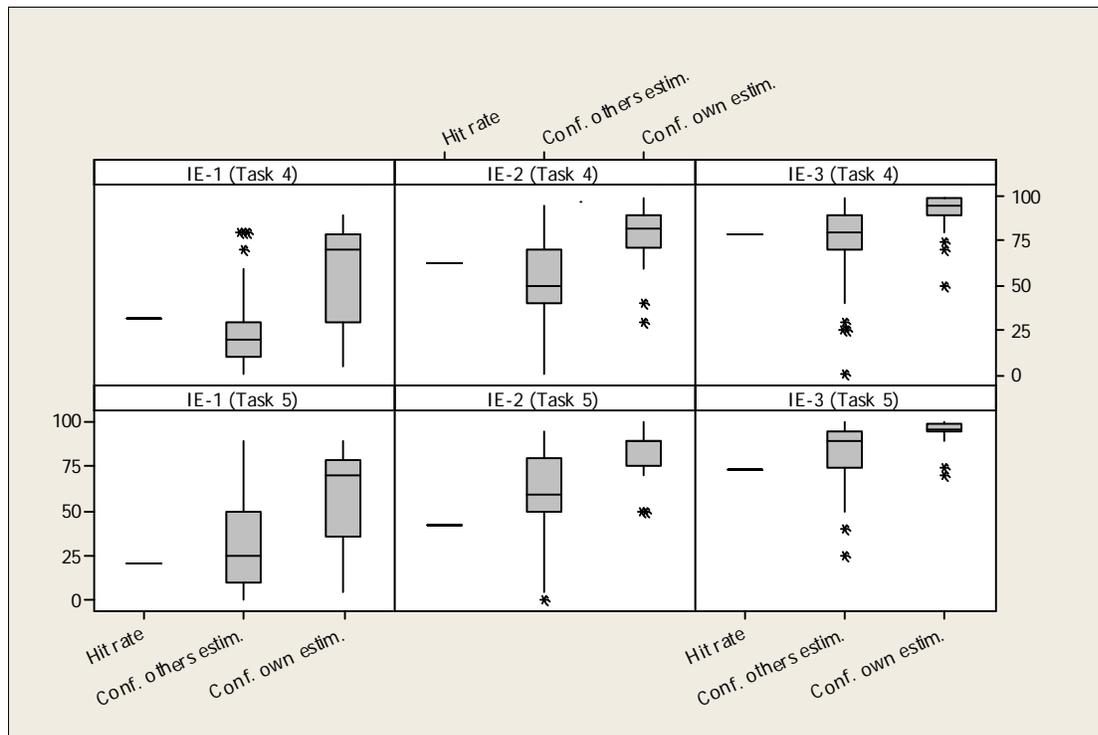
Table 8 compares the mean OverConfidence of the developers' assessment of the uncertainty of others (data from the follow-up experiment) and their own (data from the main experiment) estimates. As can be seen in Table 8, the uncertainty assessments of other developers' effort estimates, when supported by the same feedback of previous performance, are much more realistic (OverConfidence closer to zero) and less biased (both negative and positive OverConfidence values). While the uncertainty assessments in the main experiment, which were based on the developers' assessment of the accuracy of their own estimates, were strongly biased towards overconfidence, the assessments in the follow-up experiment, which were based on other developers' estimates, were more realistic and even tended towards underconfidence for one of the tasks (Task 4).

**Table 8: Mean OverConfidence of Others (Follow-up Experiment Data) vs Own (Main Experiment Data) Estimates**

Task	Assessment Type	IE-1	IE-2	IE-3
4	Others	-0,14	-0,21	-0,10
	Own	0,79	0,25	0,15
5	Others	0,71	0,39	0,09
	Own	1,75	0,98	0,28

Figure 5 displays the distributions of confidence levels together with the Hit rate of both Tasks 4 and 5 and all three uncertainty intervals. As before, the closer the confidence level to the Hit rate, the more realistic are the confidence levels. Figure 5 suggests that not only does the realism increase, but the variance decreases, as well.

**Figure 5: Boxplot of IE-1, IE-2 and IE-3 Confidence Levels and Hit Rates of Tasks 4 and 5**



The results from the follow-up experiment suggest that the uncertainty assessment learning problems observed in the main experiment is not so much related to lack of proper probabilistic mental models, but more to psychological biases related to evaluation of own estimation accuracy performance.

## V. CONCLUSIONS

Lessons-learned-based processes, such as post-mortem analyses and project reviews, are frequently suggested for the purposes of improving processes [14, 41, 42]. Our results suggest that the type of individual lessons-learned processes examined in this study may have no, or even a negative, effect on the accuracy of estimates and the realism of assessments of the uncertainty of effort estimates in comparison with on-the-job learning due to, for example, learning biases related to the assessment of own estimation performance. Without better lessons-learned processes, there may be a substantial risk of wasting a lot of resources with no, or even a negative, effect in comparison with pure on-the-job-learning. We recommend that there be a shift from the current tendency to propose lessons-learned processes that assume that it is sufficient to ask of the type “what can you learn from this” towards evidence-based learning processes that take into account, for example, the learning biases related to interpretation of own estimation performance.

**Acknowledgement:** We thank all James Dzidek, Hans Christian Benestad, Stein Grimstad, Erik Arisholm, Dag Sjøberg and other staff at Simula Research Laboratory for their contributions to this study.

**Reference:**

1. CompTIA, *Survey: Poor communication causes most IT project failures. Inadequate resource planning, unrealistic deadlines also cited in CompTIA study in Computerworld.* 2007.
2. Moløkken, K. and M. Jørgensen. *A review of software surveys on software effort estimation.* Proceedings of *International Symposium on Empirical Software Engineering.* 2003. Rome, Italy: Simula Res. Lab. Lysaker Norway: p. 223-230.
3. Jørgensen, M., K.H. Teigen, and K. Moløkken, *Better sure than safe? Over-confidence in judgement based software development effort prediction intervals.* Journal of Systems and Software, 2004. **70**(1-2): p. 79-93.
4. Jørgensen, M. and M. Shepperd, *A systematic review of software cost estimation studies.* IEEE Transactions on software engineering, 2007. **33**(1): p. 33-53.
5. Jørgensen, M., *Estimation of Software Development Work Effort: Evidence on Expert Judgment and Formal Models.* International Journal of Forecasting, 2007. **23**(3): p. 449-462.
6. Jørgensen, M., *A review of studies on expert estimation of software development effort.* Journal of Systems and Software, 2004. **70**(1-2): p. 37-60.
7. Humphrey, W.S., *Introduction to the Personal Software Process.* 1996: Addison-Wesley.
8. Abrahamsson, P. and K.H. Kautz, *Personal software process: Classroom experiences from Finland.* Software Quality (ECSQ 2002), 2002. **2349**: p. 175-185.
9. Prechelt, L. and B. Unger, *An experiment measuring the effects of personal software process (PSP) training.* IEEE Transactions on Software Engineering, 2000. **27**(5): p. 465-472.
10. Basili, V., H. Caldiera, and D. Rombach, *The Experience Factory,* in *Encyclopedia of Software Engineering,* J.J. Marciniak, Editor. 1994, Wiley. p. 469-476.
11. Engelkamp, S., S. Hartkopf, and P. Brössler. *Project Experience Database: A Report Based on First Practical Experience.* Proceedings of *PROFES.* 2000. Oulu, Finland: Springer-Verlag: p. 204-215.
12. Houdek, F., S. K, and W. E. *Establishing Experience Factories at Daimler-Benz An Experience Report.* Proceedings of *International Conference on Software Engineering.* 1998. Kyoto, Japan: p. 443-447.
13. Jørgensen, M., D.I.K. Sjøberg, and R. Conradi. *Reuse of software development experience at Telenor Telecom Software.* Proceedings of *European Software Process Improvement Conference (EuroSPI'98).* 1998. Gothenburg, Sweden: p. 10.19-10.31.
14. Birk, A., T. Dingsøy, and T. Stålhane, *Postmortem: never leave a project without it.* IEEE Software, 2002. **May/June**: p. 43-45.
15. Dingsøy, T., *Postmortem reviews: purpose and approaches in software engineering.* Information and Software Technology, 2005. **47**: p. 293-303.
16. Ohlsson, M.C., C. Wohlin, and B. Regnell, *A project effort estimation study.* Information and Software Technology, 1998. **40**(14): p. 831-839.
17. Balzer, W.K., M.E. Doherty, and R. O'Connor, *Effects of cognitive feedback on performance.* Psychological Bulletin, 1989. **106**(3): p. 410-433.

18. Benson, P.G., *The effects of feedback and training on the performance of probability forecasters*. International Journal of Forecasting, 1992. **8**(4): p. 559-573.
19. Stone, E., R and R.B. Opel, *Training to improve calibration and discrimination: The effects of performance and environmental feedback*. Organizational Behaviour and Human Decision Processes, 2000. **83**(2): p. 282-309.
20. Schmitt, N., B.W. Coyle, and L. King, *Feedback and task predictability as determinants of performance in multiple cue probability learning tasks*. Organizational Behaviour and Human decision processes., 1976. **16**(2): p. 388-402.
21. Schindler, M. and M.J. Eppler, *Harvesting project knowledge: a review of project learning methods and success factors*. International journal of project management, 2003. **21**: p. 219-228.
22. Cusomano, M. and R. Selby, *Microsoft secrets—How the world's most powerful software company creates technology, shapes markets, and manages people*. 1995, New York: The Free Press.
23. Jørgensen, M. *A Critique of How We Measure and Interpret the Accuracy of Software Development Effort Estimation*. Proceedings of *1st International Workshop on Software Productivity Analysis and Cost Estimation*. 2007. Nagoya, Japan: Information Processing Society of Japan: p. 15-22.
24. Pan, G., S.L. Pan, and M. Newman, *Information systems project post-mortem: Insights from an attribution perspective*. Journal of the american society for information science and technology, 2007. **58**(14): p. 2255-2268.
25. Lyytinen, K. and D. Robey, *Learning failure in information systems development*. Information systems journal, 1999. **9**(2): p. 85-101.
26. Wastell, D., *Learning dysfunctions in information systems development: Overcoming the social defenses with transitional objects*. MIS Quarterly, 1999. **23**(4): p. 581-600.
27. Urban, M. and A. Witt, *Self-serving biases in group member attributions of success and failures*. The journal of social psychology, 1990. **130**(3): p. 417-419.
28. Jørgensen, M. and K. Moløkken-Østvold, *Reasons for software effort estimation error: impact of respondent role, information collection approach, and data analysis method*. IEEE Transactions on Software Engineering, 2004. **30**(12): p. 993-1007.
29. Jørgensen, M. and K. Moløkken-Østvold. *Eliminating Over-Confidence in Software Development Effort Estimates*. Proceedings of *Conference on Product Focused Software Process Improvement*. 2004. Japan: Springer-Verlag: p. 174-184.
30. Jorgensen, M., K.H. Teigen, and K. Molokken, *Better sure than safe? Over-confidence in judgement based software development effort prediction intervals*. Journal of Systems and Software, 2004. **70**(1-2): p. 79-93.
31. Jørgensen, M. and D. Sjøberg. *The importance of not learning from experience*. Proceedings of *European Software Process Improvement 2000 (EuroSPI2000)*. 2000. Copenhagen: p. 2.2-2.8.
32. Hammond, K.R., *Human judgement and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. 1996, New York: Oxford University Press.

33. Fischhoff, B., *Hindsight <> foresight: The effect of outcome knowledge on judgement under uncertainty*. Journal of Experimental Psychology: Human Perception and Performance, 1975. **1**: p. 288-299.
34. Stahlberg, D., et al., *We knew it all along: Hindsight bias in groups*. Organizational Behaviour and Human Decision Processes, 1995. **63**(1): p. 46-58.
35. Camerer, C.F. and E.J. Johnson, *The process-performance paradox in expert judgment: How can experts know so much and predict so badly?*, in *Towards a general theory of expertise*, K.A. Ericsson and J. Smith, Editors. 1991, Cambridge University Press. p. 195-217.
36. Sanbonmatsu, D.M., A.A. Sharon, and E. Biggs, *Overestimating causality: Attributional effects of confirmatory processing*. Journal of Personality and Social Psychology, 1993. **65**(5): p. 892-903.
37. Brehmer, B., *In one word: Not from experience*. Acta Psychologica, 1980. **45**: p. 223-241.
38. Bolger, F. and G. Wright, *Assessing the quality of expert judgment: Issues and analysis*. Decision support systems, 1994. **11**(1): p. 1-24.
39. Shanteau, J., *Competence in experts: The role of task characteristics*. Organizational Behaviour and Human Decision Processes, 1992. **53**(2): p. 252-266.
40. Jørgensen, M., *Realism in assessment of effort estimation uncertainty: it matters how you ask*. Software Engineering, IEEE Transactions on, 2004. **30**(4): p. 209-217.
41. Cannon, M. and A. Edmundson, *Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve*. Long range planning, 2005. **38**: p. 299-319.
42. Collier, B., T. Demarco, and P. Fearey, *A defined process for project post-mortem review*. IEEE Software, 1996. **13**(4): p. 65-71.