DOUBLE-LOOP FALSE LEARNING:

The micro-processes of learning under uncertainty

BY

KRISTIAN KREINER

Copenhagen Business School
Center for Management Studies of the Building Process
(Realdania Research)
Kilevej 14, DK-2000 Frederiksberg, Denmark
(kk.ioa@cbs.dk); www.clibyg.com

Paper submitted to EGOS 2006 – sub theme 35: The Social Complexity of Organizational Learning: Dynamics of micro-practices, processes and routines.

JUNI 2006

The research has been financed in part by The Foundation Realdania and the National Agency for Enterprise and Construction. The assistance of Lise Damkjær and Mette Løth Frederiksen is gratefully acknowledged.

DOUBLE-LOOP FALSE LEARNING:

The micro-processes of learning under uncertainty

1 THE BEGINNING

To build a building is a joint achievement of considerable complexity. Obviously, success depends on the disciplined effort of a very large number of individuals and organizations. Success also depends on a substantial amount of prior design, scheduling, and organizing. With all the contingencies known to impinge on the building process there can be nothing inevitable about success. We imagine that many things can go wrong, and that, in each and every project something always will go wrong. Potentially, many things that go wrong will unsettle the joint effort of the building project. Nonetheless, in spite of the often notoriously bumpy road to success very few projects fail to deliver a functional and durable building in the end. The joint effort is regularly challenged, occasionally unsettled temporarily, and very seldom destroyed altogether.

In this paper I analyse an example of a temporary collapse of the joint effort before the building was successfully completed. A serious accident brought the work to an unforeseen standstill. The accident happened in connection with the erection of prefabricated concrete panels – probably one of the most common and central tasks in the construction process. It is a puzzle that a project staffed with apparently experienced and skilled people can seriously fail in such a "trivial" task.

Such unexplainable (and some would say, inexcusable) snags dominate the public opinion about the construction industry. In spite of the observed reliability in terms of final output, the industry is notorious for its low stakeholder satisfaction. Outcomes may be reliable, but they are also

reliably late and above budget. In most countries, construction has a general reputation of low efficiency, low quality, and low trust. “There is something wrong with the construction industry,” say Ashworth & Hogg (2000) and continue, “The world is changing rapidly; according to the Latham Report the construction industry is not”. Thus, we find innumerable reports over at least the past 50 years, all calling for immediate action to improve the performance of the industry. Apparently, the many calls have had little impact. It may in fact seem that the problems have become graver and the scandals more spectacular, especially in connection with cost and time overruns (Flyvbjerg, Bruzelius, & Rothengatter, 2003).

If failures are opportunities for learning (Starbuck & Hedberg, 2001), why is it that the construction industry has not learned from its failures?[1] Of course, the literature is full of examples of individual firms that failed to learn to their own detriment. Lyytinen & Robey (1999) describe not only how organizations stick to “existing theories in use that were clearly in need of revision” (p. 89). They also show how organizations learn to live with their failures, “Organizations learn to live with inadequate performance and attribute negative outcomes to external causes rather than their own processes” (Lyytinen & Robey, 1999, p. 94). But how is it possible that a whole industry collectively avoids changing and continues to parade ever new generations of surviving failures? (Beckman, 1979) Why will random variation or rational competition not provide sufficient survival biases in the evolutionary process to “breed” an industry of higher efficiency?

Such questions are natural, but not necessarily precise. Logically, the lack of adaptation, change and improvement may reflect the lack of learning, or a systematic learning of a false lesson. While the general presumption exists that learning is only learning if it is valid, the literature does contain indications that learning, at least temporarily, may be wrong. One example

[1] In part this research is financed by the National Agency for Enterprise and Construction. The assignment was phrased precisely as the puzzle why the industry had failed to learn from its experience.

is the distinction between different types of learning, single-loop, double-loop and deutero-learning (Argyris & Schön, 1996). If single-loop learning is invalid, failure will ensue, which will mobilize processes of reflexivity in the form of double-loop or even deutero-learning (Berthoin Antal, Dierkes, Child, & Nonaka, 2001). Of course, feedback may be slow and ambiguous and for a while allow false or incomplete understanding (Weick, 2001) and inconsistency between individual beliefs and action to prevail (March & Olsen, 1988). "Major disruptions in business often result from events which are outside such terms of reference which are agreed, explicitly or tacitly, within the firm, or even with a whole industry," claims Brian Loasby (quoted from Mort, 2001). Thus, to give another example,

> *… For years the ferry boats left port with the doors open, in response to the assumption that customers favoured on-time departures more than safety precautions. Management had instilled in the organization a mind-set of getting the ship to sea at the earliest possible moment. Other factors overlooked include excessive numbers of passengers, over-loading, ignoring previous incidents of improper door closing and tolerance of frequent crew turnover"* (Blanco, Lewko, & Gillingham, 1996, *p. 8).*

But in all these cases, the implicit understanding is that eventually the truth will out. When The Herald of Free Enterprise went down, everybody could see the faulty sense in leaving the port with open doors. The experience of major disruptions will make people acutely aware of latent events previously neglected or ignored. Thus, we seem to share a strong belief in history as a process of increasing and improving knowledge in society. We also seem to share the belief that construction is somehow not collecting the fruits of such a history. The construction industry is constantly experiencing the major disruptions that forces people to become reflexive. Is it possible that such reflexivity is making things worse, not better? Is it possible that

e.g. double-loop learning does not correct false single-loop learning, but distorts it? If we accept the premise that the problems and disruptions repeat themselves over and over again, it is hard to find support in the literature for a belief that no learning takes place. Also, consider the obligatory inquiries into the causes of all the accidents and scandals; the multiple expert-reports on the misgivings of construction, and the innumerable experienced and skilled managers and workers having put their careers and lives on the line: It would be hard to accept that this is not a learning community. It would be hard to believe that the community would not experience to be continuously learning. But if they continue to learn, and nothing changes, then we are pushed toward the hypothesis that what they are learning is systematically wrong, not only temporarily, but fundamentally and permanently. Therefore, this is the hypothesis that will guide the subsequent case study.

1.1 The research question.

The central argument of the article revolves around the dilemma that while we seek to benefit from learning we also run the risk of learning something false. The more we try to learn from experience, the higher is probably the risk that we will learn something that will not benefit us. If we accept this proposition as an inherent dilemma we need not assume that people are necessarily stupid when they learn wrong things and continue to do so. Rather, we become obliged to consider a different type of questions, e.g. the conditions under which the learning take place, the methods and procedures used for learning, etc. My argument owes much to Busby (2001), when he observes that,

> *… it is the fact that inference can be systematically wrong that means experience may not correct behaviour and error may occur … Generally, the biases in people's inferences appeared to arise from the heuristics they had learned quite logically from experience – given their limited information processing capabilities, pressures on*

> *their time to make inferences, and the partly idiosyncratic nature of their experience" (pp. 305-306)*

Thus, learning in one respect may be logical and rational and lead to false learning in another respect. The invalidity of the second type of learning does not challenge the logics and rationality of the former type of learning. The tension remains in spite of the experienced disruptions and errors, because, in my terminology, no learning can eradicate the inherent learning dilemma.

The following analysis will attempt to answer two questions,

1. How, and with what implications, is the inherent learning dilemma enacted in construction projects?

2. Is there an alternative strategy for coping with the inherent learning dilemma that might reduce the risk of learning false lessons from repeated experience?

It will be clear at the end of the article that such questions have neither easy nor definitive answers.

1.2 The methodology

The case study to be analyzed below is part of a larger study of learning failures and construction snags (Kreiner, 2006). The study defined a snag as a situation that suddenly required additional work and effort.[2] We documented the preceding processes in detail, with a special focus on processes of learning and sensemaking. The documentation was based on written material as well as interviews with relevant actors in the projects. The actual course of events was in focus in the data collection. People's

[2] Notice that we define the snag situation not in terms of its alleged cause, but in terms of its implications, i.e. the required additional effort. Extra effort may be necessary for a number of legitimate and illegitimate reasons; and since such reasons were our object of inquiry we could not adhere to the rhetorical convention of defining snags as the unwarranted and blameworthy deviance from obligation, plan or norm.

analyses and interpretations of these events were considered relevant data inputs to our own analysis and interpretations, but not as templates for our own conclusions.

The data for the particular case study in this article were publicly available as a teaching case.[3] Our account of the course of events prior and subsequent to the accident we analyze has been published widely.[4] The fact that no one has questioned the account we take as validation of its accuracy. The analysis and interpretation of these data cannot be validated in the same manner. Since this analysis and interpretation are made explicit below they are offered for the reader to validate.

1.3 The plan

The article contains four main sections. The first section is an account of a series of dramatic events on a particular construction site. The second section is an extended analysis of these events, employing counterfactual analysis to excavate the significance of what actually took place. The third section draws implications from the analysis for our understanding of organization and joint work. And the fourth section draws implications for the organizational learning theory.

2 THE CASE

This case deals with an accident on a construction site during the erection of prefabricated concrete panels. The building under construction was commissioned by a major cultural institution and was designed by one of the leading architectural practices. The ambition was to create a new prestigious landmark and the insistence on success justified additional resources and provisions for quality assurance and accident prevention. It

[3] The data can be found on the following web page:
http://historier.entreprenoerskolen.cursum.net/client/CursumClientViewer.aspx

[4] It continues to be posted on the homepage of the National Agency for Enterprise and Construction, which is an institution under the Ministry of Economics in Denmark.

was agreed, for example, that drawings and instructions were to be checked and verified collectively, and that the prescription of unusual materials or methods of work was to be flagged and explicitly communicated. Thus, the will to success was strong and the conditions for achieving success seemed better than in most other cases.

The realities turned out quite differently. Having reached the third floor a whole section of concrete slabs suddenly falls to the ground. Miraculously, nobody is seriously hurt. But seriously hurt was the image of a flawlessly planned and managed project. The sense of shock and surprise fuelled an immediate inquiry into the cause of the accident. It was immediately clear that somebody had failed miserably and that the culprit had to be found and made accountable.

Eventually, the supplier of the concrete slaps was blamed for the accident. In court he was found guilty of failing to fulfil his obligations to inform the contractor and its site staff about the planned methods of erection.

2.1 Reconstructing the chain of events prior to the accident

Early it became clear that the construction workers had been using a wrong method of erection. The erection method has no implication for the integrity of the final construction. The method had to do with temporarily securing the concrete slabs and panels until the grouting and other finishing operations had been performed. “Merely two alternative methods of erection existed, and it was not obvious which of them was the correct one,” said the representative for the Danish Working Environment Authority. The workers had chosen to rest the concrete beams on small shelves on the columns and to pile panels on top. The intended method required the suspension of the slabs from the columns for which purpose rings were embedded in the slabs. The two methods would lead to the same outcome, but the methods differed in terms of the required actions and efforts of the construction workers.

The choice of the wrong erection method produced a major accident out of a small and normally inconsequential incident. The lower corner of the concrete panel being erected broke off. Thereby one side of the panel slid down below the shelf and in the process pushed the columns apart far enough to make all the panels tumble to the ground.

While playing an important role, the broken corner was not made the villain in the reconstruction of the chain of events. The wrong erection method was the focus of attention, and the manufacturer's omission to explicitly give instructions about the proper erection methods was determined as the ultimate cause of the accident. The logic is this: Had the manufacturer instructed the contractor, the assembly workers would have chosen the correct erection method and the missing corner would have been inconsequential. The construction and the joint effort would have remained intact. Since the manufacturers were responsible for giving proper instructions about the use of their products and materials, the absence of instructions in the present case was considered a punishable neglect. It left the workers without the necessary information to make the correct choice of erection method.

The punishment for failing to give proper instructions impressed on the parties the importance of paying heed to formal responsibilities and agreements. It also impressed an understanding of the nature and character of joint effort in construction (and other types of) work.

3 THE ANALYSIS

The analysis has two parts. The first part builds on the construction of a set of counterfactual courses of events. The explication of how things *might have* happened will help us recognize the significance of the ways they did in fact happen. The second part of the analysis (chapter 4) builds on the general lessons that can be abstracted from the specifics of the case. These general lessons will be related to theories of organizational learning.

3.1 Counterfactual courses of events

Three different junctions, as they are described in the above chain of events, are chosen for the exercise of imagining a different historical course of events. The junctions are,

1. The concrete slab with a broken corner;

2. The manufacturer's failure to submit instructions; and

3. The incompatibility between the practice of erection work and the designer's intentions.

What would have happened had not the corner broken off, had the manufacturer submitted the instructions, and had the workers chosen a method of erection that did not violate the intentions of the designer? Those are the questions over which we will ponder below. The intention is to learn the wider implications of the seemingly trivial choices made at each of these local junctions in the history of a construction project.

3.1.1 The corner that broke

In the inquiry subsequent to the accident, apparently nobody paid much attention to the fact that one damaged concrete slab started a chain reaction. It is likely that it was not considered significant because such damages are the norm rather than the exception. The source of the problem can be traced back to the production in the factory, to the transportation from the factory and/or to the handling on the construction site. Normally, such small damages are easily repaired and compensated for in the assembly work. It is obviously true that had the concrete slabs been suspended and anchored the missing corner would have had no implications whatsoever. However, the small shelves on the columns on which the workers rested the slabs resulted in the low tolerance for even small variances in the shape of the panels. A specific slab, deviating from the norm, yet probably not outside the ordinary variation of things, was able to catapult the project into unappreciated fame.

With a bit of luck, however, it would not have happened. This is not a completely unrealistic thought, because after all it had not happened in the process of erecting the first three floors. Had it not happened, it is easily conceivable that the whole structure would have been erected successfully with the use of what we now consider an incorrect method. We would have had no reason to expect that accidents were latent, and we might easily have praised the workers for their competence and the management for achieving a high level of quality and safety.

We imagine that what was true before the accident might have remained true to the very end of the project. And what was true before the accident was that the erection method worked. Beams were rested on the shelves, and panels were rested on top of the beams. It worked, and they managed to progress with sufficient productivity to satisfy not only themselves but also the stakeholders around them. There was little reason to suspect that they were travelling the wrong road. And it is indeed difficult to claim that they did. They continued to do what they had learned would work well – and could have continued to work well, had not the accident happened.

It is important to understand that the workers did exactly as they had done all along. They acted consistently, and it was the consequences that showed unsuspected variance. The things that were established as true and correct in the previous process suddenly became untrue and incorrect. The sudden change is explained by the change in the criteria for truth and correctness. For the ongoing work at the construction site, the ultimate criterion was of a pragmatic (performative) kind: does it work, or doesn't it? According to this criterion, the chosen erection method worked, possibly due to the ingenuity and care of the gang. Concrete panels got erected, and the building took form, thus it worked! With the advent of the accident, apparently it didn't matter that it had worked. After the accident *the intentions* of the supplier (and his design engineer) became the criterion for assessing the correctness of the method. It was determined to be wrong because it did not comply with the manufacturer's design intentions. In the present case, we assume

that no conflict exists between the performativity and the design, in the sense that the gang would have been able to make also the suspension method work well. But pragmatics (performativity) and compliance with design premises are not the same type of criterion, and in principle they may easily be in conflict. While the period prior to the accident gave priority to the performative criterion, the accident suddenly gave priority to the design intentions as criterion. The accident also substituted 'knowledge by acquaintance' with 'knowledge by description' (Baron & Misovich, quoted from Weick (2005)).

> *…[K]nowledge by acquaintance […]is acquired through active exploration. Active exploration involves bottoms-up, stimulus-driven, on-line cognitive processing in order to take action. … Once people start working with names and concepts for the things that they see, they develop knowledge by description rather than knowledge by acquaintance, their cognitive processing is now schema-driven rather than stimulus-driven, and they go beyond the information given and elaborate their direct perceptions into types, categories, stereotypes, and schemas. … More formally, when social complexity increases, people shift from perceptually based knowing to categorically based knowing in the interest of coordination. The potential cost of doing so is greater intellectual and emotional distance from the details picked up by direct perception. … (p.163).´*

The direct experience of the workers, their knowledge of acquaintance with erecting the concrete panels, was of little use to all the other actors who were mobilized in the wake of the accident. In order to become involved and to negotiate the cause and to place the blame, they had to introduce concepts and notions like contract, designs, intentions, good managerial and

engineering practice, etc. The accident was analyzed in a completely different realm than the one in which it occurred – not by choice or manipulation, but by necessity. Otherwise, there would have been no role for all the new actors that were mobilized by the accident.

There was nothing inevitable about the breaking-off of the corner, and therefore there was nothing inevitable about the reinterpretation of the correctness of the assembly method. It was working well for a long period, and it was obviously true then. Things happened to change our interpretation, and we came to see it as the wrong method in view of the manufacturer's design intentions. While we may choose to claim that the method had been wrong all along,[5] realizing of course, that we would have continued to claim it was correct had not the corner broken off, we may also choose to say that methods and practices may be correct and valid sometimes, and wrong and invalid at other times. The latter options would direct us to pose the question how a method acquires (and loses) its quality of being correct and valid. We may try to identify and describe the underlying forces that make things acquire different qualities at different times and under different conditions. We imagine that the phase transition from being valid to being invalid happens abruptly (as in the case of the accident), and that such phase transitions may happen several times (although of course I have not studied such transitions back and forth in this specific case). These transitions happen when one underlying force suddenly gets to dominate all other forces in the situation.[6] In the present case, the sudden change in validity happened when the criterion of pragmatics was conquered by the criterion of compliance. The situation is depicted in figure 1.

--- Insert figure 1 ---

[5] This is a parallel to the claim that ice was water all along, the proof being that it melted.

[6] This argument is modeled on the phase transitions in physics, as described in (Ball, 2005)

In the figure I show the sudden transition in quality of the working method. For a long time, the method was clearly valid while the performativity was the dominating concern. Suddenly, the method becomes invalid while the dominating concern is the degree to which the method complies with the intentions of the designer and manufacturer of the concrete slabs. Performativity is not made irrelevant, of course; in the specific situation it is merely dominated by the compliance concern.

The figure also refers to other aspects that suddenly change in relative relevance and importance. Knowledge by acquaintance loses out to knowledge by description; and the site personnel have to give way to external experts and stakeholders. These aspects (and further aspects to be discussed below) are candidates for playing the roles as underlying forces to be modelled as constituting processes of organizational learning.

3.1.2 The missing instruction

The manufacturer's failure to submit instructions about the proper method of panel erection was determined as the ultimate cause of the accident. Such an omission may be hard to understand especially because the owner had required everybody to flag any unusual method or material. The omission gets a different meaning when it is viewed as a systemic phenomenon rather than as an individual breach. For that purpose, the manufacturer's justification is instructive. He claimed that the suspension method was the method *they* usually designed – thus, there was nothing unusual to report. It becomes clear to us that what is unusual, and what is usual, is a matter of prior experience, and it is of course very unlikely that the pools of experience of a distant concrete panel manufacturer and the local gang of construction workers overlap perfectly. The provision of flagging the unusual was meant to reduce the risk of misconceptions, but in fact it produced one. The lack of instructions helped sediment the expectation that there was nothing unusual about the assembly task and that the previous experience of the workers would safely apply. In fact, the prior experiences

of the workers did not apply. They were framed by the context to feel safe, when in fact they were in grave danger.

Furthermore, the provision of flagging the unusual reflects another reality about construction work. The provision was meant to conserve the limited capacity to process information, and to use the limited capacity on transferring the most important information. The strategy was rational, even if, in this particular case, the consequences were catastrophic. The logic of the provision indicates that the "correct" method of erection would not necessarily have been chosen, had the instructions in fact been submitted. Adding new pieces of information to a situation which is already overwhelmed with information will make it possible, but not probable that the new information will be picked up and processed mindfully. Unless the workers were acutely aware of the lack of instruction for the work to be done the fate of the additional instructions would be impossible to predict.

I am arguing that the instructions might easily have gone unnoticed had they been provided. It is impossible to determine the likelihood of oversight with any kind of precision, but we know that a very large part of the instructions reaching the construction site contains redundant information and are systematically ignored.[7] It would not be an unusual fate had the erection instruction been substituted with the workers' own judgments based on their prior training and experience.[8] I am also making a point of much wider significance. Had the workers in fact noticed the instructions, the "wrong" erection method might still have been chosen! The idea that deviance from instructions is all a matter of ignorance is probably wrong. In between the recognition of the instruction and the action it motivates falls the reading and the interpretation of the instruction and the context of its application (Feldman & Pentland, 2003). In practice few instructions are interpreted

[7] Consider for example the trivial example that each tin of paint reproduces instructions for proper use and handling. Every skilled painter will know how to use the materials and will never consult the instructions; unskilled painters will possibly need the instructions on the first tin, but not the instructions on all the subsequent ones.

literally, because when they are organizations fall apart.[9] Some kind of translation to make it relevant and applicable in the particular situation is to be expected. This is the middle ground between rigidly implementing the instructions and completely ignoring them. When taken seriously instructions are recognized, interpreted, evaluated and made pragmatic sense of in order to achieve a given task. It is completely possible that the workers would have questioned the validity of the instructions had they received and recognized them. They might have come to the conclusion that the designer was mistaken or ignorant about how the work is performed – that the instructions were based on knowledge by description rather than knowledge by acquaintance, and that their own experience and knowledge was superior. They might have weighted the pros and cons and still have come to the conclusion that the concrete panels should be rested on the shelves on the columns.

The above reflections point to the fact that the frame of analysis change with the advent of the accident. It is clear that in retrospect the workers are seen as choosing between two distinct alternatives (supporting or suspending the panels), and they made a bad choice because they didn't know better (i.e., they didn't know the design intentions of the manufacturer). The mental frame ascribed to the workers is one of choice and decision making. It is not clear that this was their frame of mind during the work prior to the accident. We do not know that they consciously considered erecting the panels by suspending them; it is more likely that they simply made sense of the task in front of them. In doing so they interpreted and made sense of whatever information and cue they picked up to define the task in terms of the support method. Rather than being deprived of information they struggled to make sense of the information they had. In Weick's terminology, they faced *confusion* rather than *uncertainty* (Weick, 1995). The cues came not only from the formal instructions, but from the totality of the perceived

[8] The responsibility for future misunderstandings and accidents would change place. In that case, the failure to seek information, not the failure to communicate it, would be penalized.

[9] This is one of the more effective conflict tactics of labor unions.

environment.[10] In the following section, I will analyse some of these environmental cues that apparently influenced the reading of the situation and framed the implications drawn.

--- Insert figure 2 ---

In figure 2 I summarize the forces that compete for primacy. The frame of mind changes abruptly from sense making to decision making. Suddenly the multiple physical and social cues present in the work situation become overshadowed by focus on the formal instructions (and especially the lack of such formal instructions). The ambiguity of all the cues (resulting in confusion) is replaced by the lack of information (resulting in uncertainty). Again, these competing aspects are candidates for playing the roles as underlying forces to be modelled as constituting processes of organizational learning.

3.1.3 The deviant practice

I made the point above that the workers faced too much information, not too little. Thus, the challenge was one of making sense, not decisions. They made sense of the task and the situation in a manner that subsequently was determined as being in conflict with the intentions of the designer. The inquiry established that the failure to communicate the design intentions was the ultimate cause of the accident. This interpretation revealed the underlying framework for understanding the process of joint work. Construction workers implement the decisions made by the designers, and without receiving full information about these intentions the work cannot be done correctly. The manufacturer failed to facilitate the work according to this frame of understanding. The accident was seen to reinforce the model, because it demonstrated the consequences of not providing the instructions, namely that people make mistakes and that the joint effort falls apart.

[10] The field of environmental psychology provides the theoretical foundations for understanding the ways in which information is extracted from physical cues on the scene of action. See *Ecological Psychology*, Special Issue. Vol. 8:2, 1996.

However, this interpretation neglects the fact that work had been going on for an extended period, indicating a well-functioning joint effort. I choose to define the problem not in terms of the communication of instructions, but in terms of the workers' ability to conduct work in the absence of instructions. How did they make sense of the situation? Why did they reach the conclusion that the concrete panels should be rested on the shelves on the column? And how might they have come to understand the task in compliance with the designer's intentions? Those are the questions I raise in this section.

Let me select two physical cues in the work environment that offer resources to the workers' sensemaking. The cues are contradictory and the choice of method of work hinges on the ways in which this contradiction is managed. The analysis rests on the premises already discussed. The instructions were missing and might have been ignored had they been provided. The fact that they did not request additional information (and by not doing so, unknowingly put their lives in danger) indicates that they did not perceive to be in need of additional information. They knew from previous experience the task of erection prefabricated concrete panels – that was their usual job. Probably, they had a strong preconception about the parameters of the task, and a strong competence in making things work under a variety of circumstances. These preconditions may explain that the workers (as well as the others around them on the construction site) approached the task with a disposition towards the traditional method of erection, the one in which you support the concrete panels from below instead of suspending them from above. In this case, such a disposition may have decided the relative prominence of two conflicting environmental cues. The first cue was the shelves on the columns. Such shelves extend an invitation to have something resting on them. The affordance seems obvious, but the designers did not intent them to be shelves and to offer such support (Gaver, 1996; Gibson, 1986; Latour, 2002). They played a minor role for the stability of the final structure. In the minds of the workers on the

site, the little excrescence on the column was mistaken for a shelf with a mission.

The contradictory environmental cue was the rings embedded in the sides of the concrete slabs. They extended an invitation about hooking up to something and in doing so pointed towards the suspension as the method of erection. These rings might have challenged the interpretation that the slabs were to be supported from below, had they been seriously taken into account. However, the cues were discounted because, in the words of the workers themselves, the rings were mistaken for exposed iron rods. Therefore they had no relevance in relation to the erection work. With the prior disposition in direction of the support method, and with the immediate experience of being able to erect the panels by resting them on the shelves, the workers had plenty of confirmatory feedback in support of their understanding. Thus, we must assume that the searching of the scene for additional cues was soon discontinued. Had the disconfirming cues loomed higher, e.g. in terms of grave difficulties resting the beams on the small shelves, the search might have gone on longer .

This discussion has two important conclusions. First, the designer and manufacturer of the concrete slabs, who was blamed for failing to instruct the workers properly, failed, in the period prior to accident, to provide consistent physical cues to the workers. The problem was not the *missing information*, but the *misleading information* contained in the excrescence on the columns that was understood as shelves. Secondly, had the chosen method of work been in compliance with the intentions of the designer, we would probably falsely assume that the workers had diligently implemented the communicated instructions of the designer. We now know that this is not the only possible interpretation, and probably not the most likely interpretation. The physical cues that designers inadvertently sprayed over the work scene were decisive in the present case and might likely have influenced the search for and interpretation of the instructions, had they been more consciously designed and selected by the manufacturer. The

spontaneous correspondence between the sense made by the workers and the designer would indicate that the environmental cues pointed the sensemaking in the right direction. It is possible that the designers might become aware of the multiple media of instructing the work and would coordinate the information across media. And it is entirely possible that the design of the work method would adapt to the available range of physical cues. Thus, the correspondence between practice and intention might as well be a reflection of an adaptation of the intentions to the practice (and sensemaking), as the other way around.

--- Insert figure 3 ---

Figure 3 summarizes the competing aspects discussed here. The multiple and conflicting cues were rationalized to fit the dominant interpretation; but after the accident the rationalization took place at the more abstract level – at the level of the general model of organizing and management. It promotes an idealized picture of the joint effort in construction work and judges the legitimacy of what happened on the basis of its deviance from the idealization. The idealization assumes that information is not ambiguous, and that errors are avoidable if proper incentives are in place. Thus, it allows the missing instructions to be sanctioned negatively. However, the adoption of the support method of erection was caused not by missing information, but by misleading information extracted from the physical cues. In a sense, the ongoing work on the site is dominated by processes of single-loop learning, while the retrospective inquiry into the causes of the accident represents processes of double-loop learning by bringing the premises for learning from experience to our attention.

It is worth pointing out that the double-loop learning that took place in the wake of the accident did not invent a new frame of interpretation. It merely mobilized and reinforced an existing one which was different from the workers' frame of interpretation. It is still arguable that learning took place. The accident gave empirical support for the presupposition that joint action

requires shared knowledge and a correspondence between intentions and practice. When knowledge is not shared, and when practice is allowed to diverge from intentions, joint action is in jeopardy. Conversely, when things do in fact work we are able to infer that the stipulated conditions for joint work are in place. How could it work otherwise if agreement and consistency are the necessary preconditions for joint effort to work? Such learning is tautological, but to claim that it is wrong would require us to specify the criteria of truth, which we will refrain from doing here.

Before I close the present discussion, let me observe that the accident might have occasioned additional double-loop or deutero-learning by questioning the basic presupposition: the causal relation between the shared knowledge and the consistency between intentions and practice, on the one hand; and the effectiveness of the joint effort across the many actors in the construction project. The presupposed causality made it possible not only to predict that effectiveness will follow if preconditions were met, but also to infer that preconditions must have been met if the joint effort was observed to function well. The official inquiry into the accident chose to focus on the fact that practice deviated from the intentions, neglecting the fact that the joint effort had worked well so far. By ignoring an important data point the inquiry managed to learned nothing new or interesting. I have chosen to focus on the facts that the inquiry ignored. In the discussions ahead my premise will be the fact that the method of work actually worked quite well in spite of the lack of consistency with the intentions of the designer.

4 IMPLICATIONS: THE ORGANIZATIONAL CONDITIONS FOR LEARNING

I have analysed the ways in which methods of work change from being valid to becoming invalid. In the case study, an accident forced the redefinition of the method which had served the construction workers well for an extended period of time. The accident revealed to the people involved that it had been invalid all along. The accident also reinforced the prevailing

understanding of what makes things work in a construction project. This understanding presupposes the correspondence between intentions and practice. The accident clearly demonstrated what happens when intentions and practice diverge: the joint effort is ruined. The manufacturer was penalized for allowing intentions and practice to drift apart because he failed to explicitly communicate his intentions.

But if the chosen method of erection in the case study had been invalid all along, as it was determined in the subsequent inquiry, then we have a problem with the fact that it actually worked, not only randomly, but systematically and for an extended period of time – and it might have continued to work well, had not the corner on the infamous panel broken off. We might never have had the opportunity to determine that an invalid method had been used, and since it worked well we would have inferred that the preconditions for it to work were already met. The fact that we now know that the preconditions were not met cannot be used to claim that it didn't work. In retrospect the method was invalid in the sense of not complying with the designer's intentions, but it makes no sense to claim that the joint effort did not work before the accident. It worked well, in spite of the discrepancy between plans and practice. The method was not invalid all along when measured on the outcomes. The method changed from being valid to suddenly being invalid when the accident happened. The accident made it invalid in that moment and henceforward. It was invalid because it didn't work any longer – in addition to not being in compliance with the planned method. Before the accident, everybody was doing their respective jobs without interfering each other.[11] But the accident interfered with everybody's project. It enforced a change of plans on the workers, the site management, the owner and a large number of stakeholders and experts. The interruption of the "ordinary" practice of carrying out allocated plans signals a lack of order. It required adaptation and coordination across the contractual structures and the complex division of labour before the order

[11] The claim is only relative to the task of erecting the concrete panels. We have not studied the totality of tasks on the construction site.

was re-established and the individual efforts were once again adding up to a joint effort.

The above reflections make it clear that joint effort can be achieved both when the local projects of people involved in the construction task are compatible, and when incompatibilities are not confronted. The incompatibility of the gang's task and the panel designer's task was not confronted until the accident happened. Prior to the accident the local activities functioned well in the sense that panels were erected in a manner that did not create contingencies for anybody else. That they were erected in the wrong manner meant nothing in this connection. However, it is not implied that the panels could have been erected in any manner. There are many constraints on the method of work, not least the laws of gravity and the sequential interdependencies between the work activities. Had the gang not found a way of erecting the panels that worked sufficiently well under the totality of conditions faced, the inconsistency would have been revealed much earlier. Decoupling the work processes of designers and workers to allow multiple, conflicting understandings of reality was far from enough for saving the image of joint effort. It was decoupling *and* the sufficiently effective and coordinated practice that created no insights into the conflicting understandings – and no reason to suspect that they existed. The precondition of effectiveness and coordination made the decoupling inconsequential for the joint effort – until, of course, the accident happened.

Generalizing from the case study we may draw the conclusion that decoupling is a manner in which local task performance is given the license to establish effective work practices in relation to the totalities of conditions – not only in relation to the formal and prescribed conditions. The claim can be made that joint action requires a balancing of different – and often opposing – concerns. The case study suggested many such concerns, and practice was seen to rest on the horns of performativity and conformity with the design intentions. *Implicitly, the successful balancing of the multiple concerns in practice was a necessity for joint effort to be sustained.*

Allowing any one of the concerns to dominate the other would terminate the order of things and undermine the joint effort. The decoupling of operations was less significant for its implied avoidance of confronting the conflicting understandings. Remember that in many ways the work remained highly coupled in spite of the decoupling of operations. The manufacturer and the construction workers were still coupled, e.g. in the sense that construction workers erected the panels received from the manufacturer, and that the final structure was closely predetermined both in shape, quality and time. Until the accident, the workers were able to meet the output expectations with the deviant method of work. In this sense, the decoupling was not just a way of maintaining the illusion of the joint effort by not confronting the discrepancies earlier. On the contrary, the loosely coupled processes were signs that the joint effort worked well. The accident showed beyond doubt that when things didn't work well, and people were not able to pursue their local programs effectively, the system returned to being tightly coupled instantly.

One important generalization remains, however. The case study portrayed the accident as the occasion for the joint effort to cease. I do not want to suggest that joint effort is only threatened by accidents. Many more occasions exist for that to happen, the primary reason being that there exists much interdependence across the highly elaborate division of labour in construction. The well-known "work break-down structure" portrays the hierarchical relationships between local tasks that form the basis for the design of the organizational structure. But the break-down of the total task into small, easily distributed subtasks requires that the task achievements need to be reassembled at some point in time. There exists a "work add-up structure", similar to the work break-down structure, but extended in time. Tasks can be performed locally without much coordination and interference, except for the fact that the outcomes need to be compatible with the outcomes from other lines of work. The moments when outcomes are added up are important moments of truth for the joint effort. The door that arrives at the construction site better fit the hole made in the wall – not exactly,

because the fitting can easily adjust for some types of variation, but sufficiently close so the skills of the workers putting in the door are relevant and sufficient for the task.

--- Insert figure 4 ---

Figure 4 gives an image of the work add-up structure. For construction projects this structure is immensely complex. Furthermore, the physical nature of the multitude of outcomes being added up to produce a complete building puts additional constraints on the adding process. To return to the above example, a door supplied must fit the whole in the wall. Whether it fits is still to a certain degree dependent on social definitions, and the efforts on making it fit (or seeing it as fitting) is not independent of the skills and motivation of the people involved. Yet, even good skills and high motivation meet physical constraints of a kind that fitting university courses into an educational program does not.[12] The multiple points at which things need to add up put enormous demands on the coordination of activities. This is true no matter how the coordination is achieved. If local processes of work did not find a way of adapting ahead of time to the requirements of the adding-up points, the joint effort would fall apart immediately. The case study showed that the compliance with the intentions of the designers was not the only way of achieving such coordination. On the contrary, it showed the power of local sensemaking in finding ways of balancing multiple demands of a complex work process. Accidentally, it also showed the inherent risks in sensemaking. Strictly implementing the intentions of the designers and coordinating activities by planning has its inherent risks as well, especially if local conditions vary much. These risks were not exposed in the present case, while the risks of making sense locally were. Presuming that such risks did not exist was naïve, but such naivety helped reinforce the idealized image of an efficient organization as being completely planned

[12] The distinction may not be defensible across all philosophical schools. But in a pragmatic sense, in construction there is less interpretive flexibility in deciding whether an oval window fit a square frame, than there is in deciding the consistency of educational programs.

and centrally controlled, the proof being the accidents and errors that were the result of people making sense of more cues than the official instructions. We now know that local sensemaking was not absent when things worked well, and that the accident did not reflect the risk of this practice of sensemaking, but reflected the risk of any kind of learning. It would not have disappeared, had the workers been prevented from reading the environmental cues for instructions. The risk of making false sense of these cues might have been reduced, had the misleading cues inadvertently put there by the designers and planners been part of their "terms of reference". They were not, and were therefore (to return to Loasby's point quoted above) the source of major disruption. After the inquiry and the official learning from the failure, they remain outside the terms of reference and therefore also remain a continued source of disruption.

5 THE CONCLUSION

> **"***… organizational members need to have a highly sophisticated kind of situation awareness if they are to anticipate the ways that error can arise from this [situational] variety. The organizational learning that is required is not a particular response to a particular situation, but an appreciation of the extent and manner of variation in situations"* (Busby, 2001, p. 305).

I return to Busby to draw the conclusions from my analysis. The sophistication that he requires for errors to be prevented (and in my case, for the joint effort to remain joint) translates into the necessity of balancing multiple concerns and aspects of the situation. In a sense, when reading the multiple cues from the immediate physical context the construction workers must bear in mind that somebody designed the concrete panels and made assumptions about the manner in which they would be erected. It is perfectly imaginable that the workers got more and more entrenched in their

own understanding of the situation and that they ended up disregarding the many subtleties of the task (e.g. the too small shelves on the columns). Thus, their understanding and awareness may have grown less sophisticated, not more sophisticated, with the accumulating experience of success and accomplishment. The accident introduced a new imbalance of equally unsophisticated character, but now at the other extreme of things. The inquiry installed a primacy of the intentions of the designers. If taken literally, the reflexivity sparked by the accident made things worse and people less knowledgeable. But if taken less literally, we may come to recognize a meta-message in the conclusions of the inquiry, not directed at the manufacturer of the concrete panels, but directed at the workers. The meta-message reminds them that the intentions of the designers are important and should be taken seriously. I.e., they should be taken into account, even if they should not be literally translated into action. Thus, an unsophisticated message may help balance an unsophisticated practice by lending a voice to the neglected aspects of the task. It revitalizes the learning dilemma and restores a certain amount of doubt that facilitates a more heedful reading of the situation (Weick, 2005).

The need for a revitalization of the learning dilemma suggests the potential falsity of double-loop learning. Argyris and Schön are not completely oblivious of the risk of learning errors, but they trust that proper procedures will eradicate such errors,

> *The second-order errors that arise in processes of organizational inquiry, such as a failure to question existing practices, allow … first-order errors [e.g. outcomes below or different from expectations] to arise and persist. Double-loop learning in organizational inquiry consists in the questioning, information-gathering, and reflection that get at second-order errors. When it is successful, it results in change toward values for inquiry*

> *that yields valid and actionable learning about second-order error* (Argyris & Schön, 1996, p. 28).

But Blackman, Connelly & Henderson (2004) make it clear that double-loop learning in the Argyris & Schön version has little reliability.

> *"Double loop learning is not an ongoing process of knowledge creation and testing as the process only commences when there is a mismatch, or problem, between expectations derived from world 3 [in Popperian terms, the world of objective knowledge independent of a knowing subject] and experiences in world 2 [socially constructed knowledge]. … Double loop learning is therefore unreliable in Popperian terms as it is not activated in all circumstances where knowledge could be falsified but only where the knowledge in world 2 recognises a difference between current ideas and those formalized in world 3" (p.22-23).*

The sense of double-loop learning that comes with the inquiry after the accident, and the mobilization and application of world 3 knowledge (the principles of best practice in construction) to produce and reinforce instructional procedures between the actors in the project, reveals itself as a false sense of learning from experience. It is false because it is unreliable, and it is unreliable because it was never tested against the situations in which things actually work. Learning from the accidental collapse of joint work is unreliable when the learned lessons are not tested against situations where joint work is effective. The loosely coupled structures of construction projects make it difficult to take lack of problems (as recognized in world 2) as proof of the validity of the principles of world 3 knowledge. Applying knowledge critically, i.e., to base action on knowledge while doubting and testing its validity, could seem to be a wiser strategy in view of the

fundamental learning dilemma and the not very conducive organizational conditions for learning.

Finally, to return to the second question posed in the beginning of the article, if the learning dilemma is the premise, are there any way of reducing the risk of false learning, except the already stated strategy of taking the lessons of double-loop learning seriously but not literally? The issue is not to eliminate the risk altogether, but to reduce the combined risk of missing important lessons and drawing mistaken lessons from experience. The issue revolves around the relationship between the designed solutions and prescribed methods of work, on the one hand, and the sense making in particular contexts, on the other hand. Or expressed in slightly different terms by Surowiecki (2005),

> *[The dilemma of information is] finding the right balance between the two imperatives: making individual knowledge globally and collectively useful (as we know it can be), while still allowing it to remain resolutely specific and local. (Surowiecki, 2005, p.72)*

The inquiry into the accident would seem to suggest that the designs should prevail, while too much local sense making contributed to the errors. In idealized conceptions of organizations, local sense making (and the human variability it introduces) is considered superfluous by design, planning and control. In practice, few would dare to pursue such a strategy – and the literature is replete with arguments why it would be an impossible strategy (Polanyi, 1966; Polanyi & Prosch, 1975). The alternative extreme that local sense making was given primacy would soon destroy the organization by making the adding-up of the various pieces of work highly uncertain. Somehow the strategy must imply that the instructions and the sense making come to supplement each other, or even support each other in practice. If, for example, the instructions are designed in such a way that the local sense making would point in the "correct" directions; or if the physical and social

cues were so edited that the sense making would point in the direction of the designed solutions, it would be more likely that things would add up and the joint effort would remain joint. It would reduce the risk of mistakes and accidents, but would not eliminate the risk completely. Still, whatever strategy we adopt, the instructions are dependent upon the sense made of them in practice, and the sense made of them in practice remains dependent on very many different factors besides the instructions. But the primary impression one gets from the case study is the bias towards learning too quickly and too easily from experience. If such a bias exists we also know that the risk of learning something false is probably high. In the interest of lowering this risk, perhaps the strategy should be to make it less easy to learn. Inquiries should be conducted not only when things go wrong, but also when things go right. That might give new impetus to a more valid double-loop learning. The rationalization of conflicting cues might be made harder for the workers and thereby introduce a critical stance among the workers on the construction site. Painting the rings in the side of the concrete panels, or making them in a material not easily mistaken for the iron rods, would make it harder to discard them as irrelevant to the sensemaking process. That would slow down the readiness to conclude the lessons of experience without serious scrutiny. It would implant a seed of doubt – which would be rational to have in the view of the fundamental dilemma of learning.

The article has shown that work can be joint work without necessarily being performed as planned and anticipated. It has shown that experiential learning can accumulate and make people's awareness of the situation *less sophisticated*. The reflexivity in the wake of accidents and other errors may deplete the community of such false learning by introducing an alternative, equally false and unsophisticated awareness. By neglecting the importance and necessity of sense making, and by presupposing a causal model in a community characterized by loose couplings, the difficulties and multi-dimensionality of learning are underestimated. This breeds a situation awareness that is not a highly sophisticated one. And since the double-loop

learning in construction reduces the sophistication of the awareness, errors are bound to be repeated. The failures in construction are failures in its learning processes. It is not the failure to learn, but a disposition to learn things that are false, that is the reason why there is something wrong with the construction industry.

6 REFERENCES

Argyris, C., & Schön, D.A. 1996. ***Organizational learning II. Theory, methods, and practice***. Reading, MA: Addison-Wesley.

Ashworth, Allan, & Hogg, Keith. 2000. ***Added value in design and construction***. New York: Longman

Ball, P. 2005. ***Critical mass. How one thing leads to another***. London: Arrow Books.

Beckman, S. 1979. Immunization - in defense of status quo. In B. Persson (Ed.), ***Surviving failures. patterns and cases of project mismanagement***. Stockholm: Almqvist & Wiksell.

Berthoin Antal, A., Dierkes, M., Child, J., & Nonaka, I. 2001. Organizational learning and knowledge: Reflections on the dynamics of the field and challenges for the future. In M. Dierkes, A. Berthoin Antal, J. Child, & I. Nonaka (Ed.), ***Handbook of organizational learning & knowledge***: 921-939. Oxford: Oxford University Press.

Blackman, D., Connelly, J. & Henderson, S. 2004. Does double loop learning create reliable knowledge? ***The Learning Organization,*** 11 (1): 11-27.

Blanco, José, Lewko, John H., & Gillingham, David. 1996. Fallible decisions in management: Learning from errors. ***Disaster Prevention and Management***, 5(2): 5-11.

Busby, J. S. 2001. Adaptation failures that cause error in technical organizations ***Human Systems Management***, 20: 301-312.

Feldman, M.S. & Pentland, B.T. 2003. Reconceptualizing Organizational Routines as a Source of Flexibility and Change. ***Administrative Science Quarterly***, 48: 94-118.

Flyvbjerg, B., Bruzelius, N., & Rothengatter, W. 2003. ***Megeprojects and risk. An anatomy of ambition***. Cambridge: Cambridge University Press.

Gaver, W. W. 1996. Affordances for interaction: The social is material for design. ***Ecological Psycology***, 8(2): 111-129.

Gibson, J.J. 1986. ***The ecological approach to visual perception***. Hillsdale, N.J.: Lawrence Erlbaum Associates.

Kreiner, K. 2006. ***Læringssvigt i byggeriet. En empirisk analyse af nogle gode grunde til dårlige resultater***.

Latour, B. 2002. Morality and technology. the end of the means. ***Theory, Culture & Society***, 19(5/6): 247-260.

Lyytinen, Lasse, & Robey, Daniel. 1999. Learning failure in information systmes development. ***Information Systems Journal***, 9: 85-101.

March, J.G., & Olsen, J.P. 1988. The uncertainty of the past: Organizational learning under ambiguity. In ***Decisions and organizations***: 335-358. Oxford: Basil Blackwell Ltd.

Mort, Joe. 2001. Nature, value and pursuit of reliable corporate knowledge. ***Journal of Knowledge Management***, 5(3): 222-230.

Polanyi, M. 1966. Tacit knowledge. In L. Prusak (Ed.), ***Knowledge in organizations***: 135-146. Boston: Butterworth-Heinemann.

Polanyi, M., & Prosch, H. 1975. ***Meaning***. Chicago and London: The University of Chicago Press.

Starbuck, W.H., & Hedberg, B. 2001. How organizations learn from success and failure. In M. Dierkes, A. Berthoin Antal, J. Child, & I. Nonaka (Ed.), ***Handbook of organizational learning & knowledge***: 327-350. Oxford: Oxford University Press.

Surowiecki, J. 2005. ***The wisdom of crowds. why the many are smarter than the few***. London: Abacus.

Weick, K.E. 1995. ***Sensemaking in organizations***. Thousand Oaks: Sage Publications.

Weick, K.E. 2001. Technology as equivoque: Sensemaking in new technologies. In ***Making sense of the organization***: 148-175. Oxford: Blackwell.

Weick, K.E. 2005. Making sense of blurred images: Mindful organziaing in mission STS-107. In W.H. Starbuck & M. Farjoun (Ed.), ***Organization at the limit. lessons from the columbia disaster***: 159-177. Oxford: Blackwell Publishing.

Figure 1

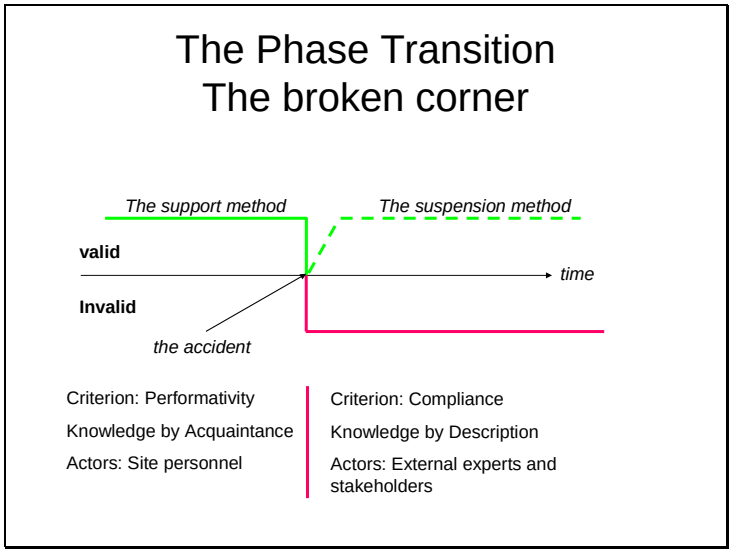

Figure 2

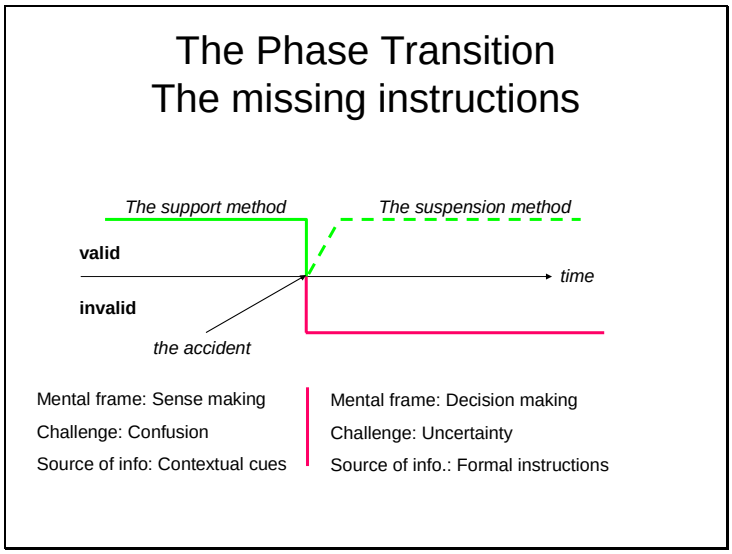

Figure 3

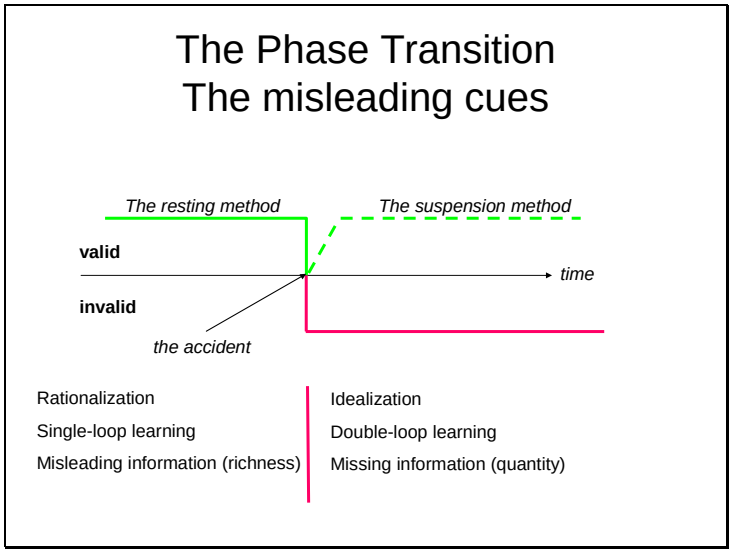

Figure 4