What is quality in a PhD dissertation in mathematics education?

MOGENS NISS

The present paper discusses the issue of quality in PhD dissertations in mathematics education on the basis of the author's reflections, observations and experiences as a supervisor and as an assessor of PhD dissertations in several countries during the last three decades. Thus, the paper represents the personal stances and views of the author and does not claim to be written on behalf of any segment of the community of researchers in mathematics education. Two major components of quality in a PhD dissertation are being dealt with, quality of the underlying research, and quality of the dissertation as a reflective report of this research and its outcomes. Particular attention is being paid to the issue of what should cause a dissertation to be rejected. The paper emphasises, at the end, that because of the multiplicity of research paradigms and philosophies in research in mathematics education there is no royal road to quality in a PhD dissertation. Therefore, the student cannot avoid involving him- or herself in independent in-depth thinking.

During the last couple of decades the number of PhD degrees earned in mathematics education has increased markedly around the world. This is especially true of the Nordic countries, where national and regional PhD programmes and graduate schools have been established, in some cases with generous funding from various sources. The Nordic Graduate School (NoGSME), 2004–2009 has contributed considerably to promoting and furthering graduate education in the field (Grevholm, 2008). Along with the increasing number of PhD students, with the establishment and implementation of rather regulated programmes, in some places under close to "school like" conditions, and along with the involvement of growing pools of supervisors, sometimes with

Mogens Niss
Roskilde University

Niss, M. (2010). What is quality in a PhD dissertation in mathematics education? *Nordic Studies in Mathematics Education*, 15 (1), 5–23.

somewhat limited experience in that capacity, we have witnessed a growing demand for directions concerning desirable approaches and procedures (cf. Grevholm, Persson & Wall, 2005), as well as criteria for quality and relevance, concerning PhD studies and dissertations. ICME-11, (Mexico 2008) was the first ICME to include a discussion group (*Rethinking doctoral programmes*) devoted to doctoral studies. Moreover, in more and more countries there is a growing demand from politicians, authorities and institutions for comparability and harmonisation of studies within and across countries. Even though there certainly is, in practice, a fair amount of consensus in academia in general and amongst supervisors in particular about what counts as quality in a PhD dissertation, we can also find a lot of disagreement, or at least diversity, in the community of mathematics education researchers about this issue. Therefore, no one in the community can claim the right to speak *ex cathedra*, i.e. on a basis of settled and canonised rules and criteria and on behalf of a significant segment of the community. Needless to say, this is true of the present author as well. What I can offer here, therefore, are my own considerations and analyses, marked by my background, experiences, observations, and personal views. Of course, these have all been greatly influenced by exchanges and discussions with national and international colleagues during many decades. Some of the views expressed in this paper are likely to give rise to disagreements with colleagues, whereas others may turn out to be less controversial. In summary, I do not claim to speak for others than myself.

Structuring the discussion of quality

When attempting to answer the question which forms the title of this paper, *What is quality in a PhD dissertation in mathematics education?*, it is important to note that this question actually involves two different issues, namely (1) the issue of quality of the research underlying and preceding the dissertation, and (2) the issue of quality of the dissertation as a report of the research conducted and of its results. Obviously, there is a connection between the two components, although in actual practice the connection is not always as manifest as one would have expected. It is not so seldom to encounter what on closer inspection turns out to be an excellent piece of research, which, however, is rather poorly presented or argued for. The converse, an excellent report based on poorly conducted research, also occurs, except perhaps slightly less frequently.

This paper, then, is structured in four sections, dealing with quality of the underlying research, quality of the dissertation as a report, grounds for rejecting a dissertation, and final remarks.

Before moving on I would like to stress that everything said in this paper is specifically concerned with *mathematics education*. A dissertation which does not have substantive aspects of mathematics education in focus in all its major components might well be an excellent dissertation in other respects, but it cannot be of quality as a dissertation in mathematics education. In order to avoid tedious repetition I shall abstain from using the words "mathematics" or "mathematical" again and again. Should it turn out, therefore, that the text appears to make sense across disciplines and subjects, this is incidental. Whether or not the paper does in fact make sense in other domains, has to be an issue of concrete assessment with regard to these domains, an assessment which has to be undertaken by experts in them.

Quality of the underlying research

Quality of research for a PhD manifests itself in three places, in the research question(s) posed, in the research design and methods adopted, and in the research findings obtained.

Quality of the research question(s)

From my perspective, the issue of quality of the research undertaken takes its point of departure in the research question(s) posed. Academic traditions differ across disciplines and across countries and cultures, and in some disciplines – e.g. in some of the humanities – research questions do not play a prominent part at all, at least not explicitly. In other disciplines – e.g. oftentimes in mathematics – research questions occur only indirectly, as questions one may formulate only *post hoc*, i.e. once the results have been obtained. Although this is not the place to prove such a general claim, I do claim, however, that any activity that can be characterised as – or deserves the name – research does seek to answer research questions, even though these may not be formulated explicitly, especially not from the outset. So, as I see it, the first important component of quality of research is the quality of the research question(s) posed. Several aspects are in play here. For simplicity I shall refer to *the* research question, in singular, even though a dissertation may well, and often does, address more than one such question. The quality aspects discussed below are meant to pertain to each of the research questions posed.

The primary aspect of quality of a research question is its *clarity* and *precision*, both of which are prerequisites for a meaningful discussion of other properties of the question. Also, the question has to be *genuine*, in the sense that the answer is not already known, and *non-rhetorical* in

the sense that it does not have a trivial answer (in contrast to a rhetorical one, like "do teachers of mathematics have to know any mathematics in order to teach it?"). Then comes the scientific and scholarly *interest* of the question. Does the question deal with the very core of the area it concerns, or does it rather touch upon more marginal features? Does answering the question add to the body of knowledge and insight accumulated in the discipline? Closely related to "interest" is *significance*. A research question is significant if answers to it give rise to a wide array of consequences, either for mathematics education as a discipline or for the practice of teaching and learning of mathematics at large. Significance in the latter sense is sometimes referred to as *relevance*. Intellectual *depth* is to do with the numerosity of links that the question has to other questions. Questions with numerous links are likely to be hard questions to provide solid answers to. O*riginality* is another aspect of quality of a research question, since originality may open new perspectives on what is worth getting to know. Of course, originality of the question cannot be a "must". A dissertation succeeding in answering a research question of long standing is indeed likely to possess manifest quality even though the question is devoid of originality. Finally, for a question to be a research question, rather than just a question, it has to be *researchable* in the sense that scholarly or scientific approaches with a potential to provide answers to it have to be available.

Many reasonable questions in mathematics education are not research questions. For instance, "What can we do to improve to students' learning of mathematics in Norway?" is far too general, vague and imprecise to be researchable. Even though the context and the situation of the question may well be specified, sharpened and concretised (e.g. "What can we do to improve Peter's learning in grade 7 of linear first order equations with one unknown"), the very sub-phrase "what can we do" remains open, unfocused and vague, because there is no direction or end to what may be included in "what can we do". For instance, one potential answer that does not violate the nature of the question might be to offer Peter his favourite cake and hot cocoa when he works on equations, while an other answer might be to present a lot of real-world applications of the equations Peter is asked to solve. Therefore, no research approach could be designed that could answer the question as it stands. Similarly, I would not consider a question like "What can we say about the way Swedish students utilise ICT in their mathematics homework?" to be a research question, again because it is too open with no focus, direction or end. We can provide all sorts of completely different answers to this question, e.g. "Swedish students prefer calculators with silver coating when they do their homework" or "Swedish students tend to use ICT to replace mathematical reasoning".

Nevertheless, questions like the ones just mentioned may well inspire the posing of research questions proper, like, e.g. "Does the introduction of real-world applications further seventh grade students' ability to solve context-free linear first order numerical equations?".

Admittedly, it is not an easy task to give a precise criterion to single out research questions from other sorts of questions. However, research questions in any domain seem to arise from *two fundamental generic questions*:

- Is it/can it be *true that* ...?
- *Why is it* that ...?

Questions such as "Is it possible that ...?", "Does there exist ...?" and "If so, under what conditions/circumstances ...?" are derivatives of the first question, whereas "How can we explain that ...?" and "How can we characterise ...?" are variants of the second one. It goes without saying that even though research questions are – I submit – derivatives of one of these two generic questions, this does not imply that such formulations in themselves ensure that we are dealing with research questions. For instance, "Is it true that Finland got the best country score in PISA mathematics amongst the Nordic countries in 2006?" is not a genuine research question, because the answer is already known (to be "yes!"). Similarly, "why is it that students who do not obtain a pass mark in a mathematics exam flunk that exam?", is a rhetorical question.

In addition to research questions proper, research also makes use of what I term *auxiliary questions* like

- What are the *established facts* about ...?
- What are the available *data* concerning ...?
- What is the *standard explanation* of ...?

Auxiliary questions serve to pave the way for answering a research question, by "charting the land" or providing an account of "the state of affairs". Put somewhat simplistically answering a research question aims at establishing new facts or insights which are not (yet) known to anyone, whereas answering an auxiliary question amounts to uncovering, combining or surveying known facts. That is what journalists do when they undertake the fact finding activities that they call "research". As examples of auxiliary questions we might take "How many high school teachers of mathematics in Iceland currently hold a university degree in mathematics at the master's level?" and "What were the official

mathematics syllabi at the upper secondary level in Denmark during the years 1903–2005?". Auxiliary questions can certainly play a crucial part along the road in providing answers to a research question. Also, it may well be a tedious and complicated, and in some cases impossible, task to answer an auxiliary question, but this does not make it a research question in my understanding of that term. But, once again, others may prefer to draw the demarcation line differently – or may even insist on seeing questions to form a continuum without any principal distinction between research questions and other kinds of questions.

Quality of the research design

With the term "research design" I refer to the overarching plan and lay-out that a researcher selects or creates and applies in order to guide and conduct the intended research. A research design can be theoretical or empirical or a combination of the two. Typically, a given research design involves a set of research methods for investigating particular phenomena and/or issues. These methods of investigation may well differ from one another in nature and scope. Some may be predominantly conceptual or theoretical, others may be empirical of qualitative orientation, while still others may be, say, statistical. So, for a dissertation there is only one research design, but probably several research methods. Choices of, discussions about, and reflections on research designs and research methods are covered by the term "methodology". Below, I shall use the term "research design" as a short-hand to also include methods and methodology.

It follows from my seeing the research question(s) as forming the pivot of the piece of research underlying a PhD dissertation that the research design and the ensuing methods of investigation to be adopted have to be derived from the research question. So, questions come first, design and methods come second. This is not to say that a research question in any automatic manner determines the research design, rather that the design has to be focused on answering the research question. Therefore, it is part of my view of quality of the research design that this has not been chosen or constructed prior to the posing of the questions. In other words a decision such as "I want to do empirical research making use of semi-structured interviews with teachers and classroom observations of their students' problem solving strategies, in the expectation and hope that I shall be able to see something interesting" is not a quality decision, even though choosing such a design may indeed prove to be a quality decision *provided* the research question posed calls for it and makes it sensible.

The main components of quality of the research design and the ensuing methods of investigation are the following.

The capacity of the design and methods to provide justifiable *answers to questions* in general (of whichever kind), and in particular justifiable answers to *those questions* that have been posed for the research endeavour at issue, is a primary quality parameter. Adopting a research design based on questionnaires to teachers is not likely to be the best way to identify students' views of the role of mathematics in society, whereas it may be part of an excellent design if the aim were to uncover the relationship between teachers' expectations about students' views and the actual views articulated by those students.

Since a good research question is likely to have arisen from a lengthy process, probably consisting of many steps, to identify, clarify, specify, sharpen and focus the researcher's initial generating research interest into something tractable, it is also an instance of quality if the research design is *compatible* with that generating interest. This is to say that the design should allow for answering the research question at the same level of specificity as that on which the question has been formulated. This is not the case, for instance, if the design adopted can only provide narrow and particular answers to special cases of the research question. As an example, imagine that the research question concerns students' ability to solve non-routine mathematical problems but the design adopted is to ask students to solve problems involving reduction of algebraic expressions only. That design might give rise to answers that have some bearing on the research question but at a much narrower level than intended. (However, asking students to do algebraic reductions might be an excellent idea if it were employed, along with other instruments, to uncover the relationship between students' manipulative skills and their problem solving competency.)

Research designs have different *scopes* and *ranges*. A research design that can only answer special aspects or parts of the research question has a narrower scope than a design that can answer the question in its entirety, as posed. Some research designs are created in order to answer only the particular research question posed, without having a wider range of applicability in mind. Such designs cannot have been tested in other contexts. This does not imply that they suffer from deficiencies, only that the burden of justifying their applicability in the current context has to be dealt with entirely in that context. Designs that can be used – or have been used – in a wider range of contexts to answering a variety of research questions are likely to have a more solid foundation, and to be of greater interest, than singular ones.

Investigational approaches that can give rise to clear, *strong* and complete answers to the research question are of a higher quality than approaches that can only provide relatively vague or tentative answers

to it. Designs that lend themselves to producing *generalisable* answers are better than ones that are limited to generating particular answers. Furthermore, the capacity of a research design to allow for *control* of the relationship between background factors and investigational variables is a quality parameter.

The same is true of what is often called *triangulation* in the set of research methods adopted. The term "triangulation", which is borrowed from land surveying, needs to be defined as it is used by researchers in at least two different meanings. Some use it to indicate that different aspects of the research question are being investigated by different means. In contrast, I use the term to indicate that answers to the *same* question are being sought by means of different methods. If the application of different methods give rise to the same answers, more or less, or at least answers that are reasonably compatible for the context at issue, the entire research design has provided triangulation in the sense that the position of new a "vertex" (read: result) has been identified from "measurements" (read: inferences) that have been obtained from both endpoints of a base line segment (read: an established station of departure) in a consistent manner.

Quality of the research findings

A primary quality criterion concerning a research finding is that it is *genuine* and *non-trivial* (i.e. neither already established and well known, nor of a purely rhetorical nature). It is not as seldom as one would want it to be to encounter the "much ado about nothing" syndrome, where the findings are meagre and not at all surprising, and presented at the end of a huge conceptual or theoretical exposé which does not (appear to) have any real bearing on the findings. In the most problematic cases one may get the impression that what is labelled "research findings" is not much more than a permutation and confirmation of initial prejudices held by the author.

Another primary quality criterion is that the research finding is *falsifiable*, i.e. it must be possible to indicate circumstances that would render the finding false. In other words, the finding is not allowed to be immune, *a priori*, because of its form or its substance, to the claim that a contradictory conclusion is true instead. Next comes the requirement of *solidity* (or trustworthiness, as Schoenfeld (2007) prefers to call it) of the findings, i.e. the degree to which they are robust to various kinds of (*a posteriori*) objections and criticism, including the possibility of deriving alternative interpretations and claims from the evidence presented. It is also part of solidity that the findings can be defended if confronted with

contradictory findings from other studies. And it is part of the notion of solidity that a result obtained actually does follow from the methods employed to generate it. It is not unusual to see research findings that are not – and sometimes cannot be – justified by the methods adopted.

It is, of course, a feature of quality if the findings of a dissertation include *answers to the research questions* posed, either entirely or in part. Looking at actual practice, this is not a given thing as much as one would expect. It is not a rare event to encounter dissertations in which the findings presented are not really related to the research questions, but perhaps to some other questions that were not explicitly posed. It is even more usual to encounter findings that only cover parts of the research question. So, findings' connection to, and degree of coverage of, the research questions constitute a quality parameter in this context.

As much as it is an indication of insufficient quality if the findings display a lack of connections to, or a weak degree of coverage of, the research questions, it is, on the other hand, indeed a feature of quality if the *range* of the findings turn out to provide solid answers to a wider set of research questions, in addition to those posed. While "range" is to do with the set of research questions answered by way of the findings, *generalisability* is a matter of whether the findings can be claimed to hold for a wider spectrum of contexts than the ones within which the research has been carried out. If, for instance, a qualitative empirical investigation has been conducted on a small sample of study objects (e.g. students, teachers, classrooms, schools, tasks, textbooks, or curricula) it may happen that the findings have a potential for being generalised to other contexts as well, either immediately (because the questions have been answered in such a way that the answers from the outset go beyond the limitations of the context) or through replication of the research in new contexts. It is of particular interest, and a feature of quality, if the results obtained from a small scale study can be brought to hold for large scale populations as well. The opposite of generalisability is *over-particularity*, which occurs when the findings are so closely tied to and dependent on the specific context and circumstances investigated, that no transfer of results is likely to be possible beyond the context and circumstances in which they were generated. In such cases there is hardly anything we can learn from the findings, which is, indeed, a weakness as far as quality is concerned.

Here it may be worth pointing out that we should neither equate "small scale studies" and "*qualitative* studies", nor "large scale studies" and "*quantitative* studies". The fact that qualitative studies tend to be of a small scale and that large scale studies are quantitative more often than not, is not a matter of principle and does not imply that we cannot have small scale quantitative studies or large scale qualitative studies.

It is characteristic of research findings of quality that they give rise to a multitude of new research questions, or pave the way for the adoption of a successful research design in *further research* enterprises, sometimes even entire research programmes.

While the fundamental quality criterion of a research finding is that it adds to the scientific and scholarly edifice of knowledge and insight in our discipline, many mathematics educators (e.g. Lester & Lambdin, 1998) want to give high priority to *applicability* and *relevance* within the set of quality criteria for research findings. Research findings, which, in addition to being solid from a research point of view, can be *applied* to inform the design of curricula or materials, the practice of teaching of mathematics, the construction of assessment modes and instruments, and soforth, and findings which are relevant to teachers, policy makers, employers, parents, etc., may be seen as having a higher level of quality than findings that primarily are of internal interest to the discipline itself.

It is difficult to discuss the issue of quality of research findings without taking into account that there are very different kinds of research, and hence of research findings. Even though the far majority of papers and, indeed, PhD dissertations, are of an empirical nature, research in mathematics education is not, and should not be, limited to being empirical.

Some research identifies new, *important issues* or *problématiques*, mostly generated by systematic reflection on observations and experiences gained from looking at practice or from introspection. Other research results in proposing new *distinctions* or *concepts*, in some cases even *theoretical* constructs, based on in-depth conceptual analyses of situations and phenomena. Empirical research may consist in *organising and reporting experiences* stemming from *ad hoc* observations and impressions. It may also consist in *systematic observation*, based on pre-established protocols, or in designing and carrying out *controlled experiments*, which in their classical form involve comparisons between experiment and control groups. C*ase studies* are usually conducted in order to generate interpretations and models that can help us make sense of complex, multivariate webs of contexts, situations, phenomena, agents, background factors, and suchlike, in the hope that eventually we shall be able to create *interpretive theoretical constructs* – in a very few instances even entire *theories* – that are significant and applicable beyond the cases for which they were first created and developed. The final type of outcome of research is *design* and *constructions* (of, say, curricula, teaching approaches, assessment modes or instruments), where it is part of the research task to show that these designs and constructions possess certain desirable properties.

Clearly, the quality components listed above have very different roles and weights with regard to so many different kinds of research and research outcomes. So, concrete discussions about the quality of a given research result have to take these differences into account. For instance, it doesn't make sense to apply quality criteria that are particularly focused on conducting large scale quantitative empirical studies, based on statistical methodology, to in-depth qualitative studies of a couple of singular cases. Nor does it make sense to transfer quality criteria specifically pertaining to empirical studies to a conceptual analysis of different proof structures in tertiary mathematics.

Quality of the dissertation as a report of the work done

In this part of our deliberations, we distinguish between two different aspects of the dissertation as a reflective account of the research work done:

- the *scientific/scholarly* quality of the exposition contained in the dissertation, and
- the *communicative quality* of the dissertation.

Even though these two aspects are, of course, closely linked, as the communication at issue deals with academic matters, they are not identical, since the former aspect focuses on substance whereas the latter rather focuses on form. We shall deal with them one by one.

Scientific/scholarly quality of the exposition

A good dissertation places the work that has been done *in a context*. By outlining what other researchers have done before, in the area under investigation, with particular regard to the *problématiques*, approaches and results already established, and by identifying and characterising the focus of the present PhD project and presenting its outcomes relatively to others' accomplishments.

It gives a *clear and exhausive account of the research* done for the dissertation, ranging from the initial *problématique*, the research questions, the conceptual and theoretical framework(s) employed, over the research design and the methods adopted, through to the findings obtained. In so doing, it demonstrates a high degree of *balance* and *harmony* amongst these dimensions, so that the dissertation is neither heavy at the front end (putting forward a huge machinery at the beginning without counterbalancing it by the research design or the extent and significance of

findings at the end) nor at the rear end (presenting piles of outcomes at the end which are not counterbalanced by a satisfactory setting of the stage in the beginning, or by sufficient conceptual or methodological deliberations in the middle).

A good dissertation *argues exhaustively and sufficiently* for the results obtained and offers a *thorough discussion* of their solidity (reliability and validity), scopes, limitations and significance, as well as of their strengths, also vis-à-vis possible alternative interpretations of the evidence on which they rely. In so doing the dissertation adopts a clear distinction between *normative* and *descriptive* questions, statements and results.

A high degree of demonstrated *methodological and philosophical reflexivity* about the nature of the research work accomplished and presented, and about what has not been done, is one of the characteristics of a quality dissertation, as is the identification of concrete research trajectories in continuation of the dissertation.

A good dissertation handles what we might call *"non-trivial trivialities"* in an appropriate manner, i.e. it clarifies explicitly what the author has borrowed from others and what is original with him- or herself, the grounds on which the literature review has been carried out (what criteria have been employed for selecting literature, how – and how systematically – has the actual selection procedure been conducted) and how referencing has been tackled.

There are two special, inter-related points I would like to pay attention to here. One very often sees dissertations full of *second- or third-hand accounts and quotations*, i.e. primary authors' work is cited or even "quoted" from secondary authors' writings. This may be justified in cases when it is not possible to get access to an original text, either because it is not easily retrievable in libraries or on the internet, or because it has been written in a remote language. However when the original text is more or less readily accessible it is unreasonable that the author of a dissertation has not made the effort to get hold of the original literature and to present or quote it first hand. This is particularly important when it comes to outlining others' positions or views where second-hand accounts – which in the dissertation then become third-hand accounts – often give an unbalanced and distorted representation of the original position. For instance, I have experienced myself in a number of cases that some of my writings have been cited (not quoted!) second or third hand in ways that make it difficult for me to recognise myself in the accounts. A related point is that one often encounters – not only in dissertations but in papers by established researchers as well – that distinctions, concepts, terms, findings etc. are attributed to an author who has applied them as if that author had also invented them, even though they have, as a matter of fact, been created by someone else. It may be true that priority disputes

do not constitute the most important issue in academic work, but it is the obligation of any author to make sure that (s)he is not inadvertently giving credit to the wrong predecessor. If you do not protect the intellectual property and integrity of others, you may eventually experience that others do not protect yours.

Communicative quality

A good dissertation has a *transparent structure* and a *clear and omnipresent "read thread"* that explicates the nature and role of each individual section/chapter within the entire dissertation, such that it becomes clear to the reader throughout the text what (s)he is reading about right now and, especially, why. Typically, such a dissertation contains summaries at key places that serve the purpose of taking stock of what has been accomplished so far, and to stipulate what the next item on the agenda is about. On the other hand, it is also important to avoid excessive repetition. In order to strike a balance between taking too many things for granted and being unhelpfully pedantic or repetitive, the author of a good dissertation demonstrates *empathy* with regard to the backgrounds, prerequisites and needs of the indended readership.

It is crucial for the communicative quality of a dissertation that *sharp definitions* of all key concepts are provided. This is important even with concepts that have been borrowed from other researchers, both because it is so easy to introduce unintended distortions and transformations of established concepts by just using them uncritically, and because it may be worthwhile, and sometimes necessary, to settle the definition of concepts that display some vagueness.

In academic literature in general and in philosophy of science in particular, a number of so-called "razors" can be found, including what is usually referred to as "Occam's razor" (although one cannot find the standard version of it in William of Occam's own writings): When choosing between competing hypotheses or theories, all capable of answering a given question or explaining a given phenomenon, we should select the one which posits the smallest set of assumptions and entities. I would like to propose to consider yet another razor, which we might call *the exposition razor*: "When writing an academic text, strive to express yourself in the simplest possible way, without compromising conceptual clarity and exhaustive argumentation." A good dissertation employs this razor, guides the reader through the text, and when specialised or technical language becomes necessary makes a big effort to explain things clearly and concisely. Jargon and pompous "impressionator lingo" is avoided in a good dissertation, as is unspecific and unnecessary name dropping.

In addition to these major points, a good dissertation also tackles yet another set of *"non-trivial trivialities"* in a satisfactory manner Thus it provides a vivid, engaged and engaging exposition, and avoids excessive repetition. Sentences are formed under respect for formal and material logic. Presentations of others' investigations and theories are balanced and to the point, and are never longer than necessary. Quotations are never used to replace the author's own efforts, or as mental crutches – "by quoting the famous Dr. X I must be on safe ground" – but as objects of investigation or pieces of evidence only. Figures, tables and diagrammes are only used when relevant for the exposition, and when this is the case, their origin and content are explained and their role in the text is clarified.

There is one issue which continues to generate debate amongst researchers in mathematics education, namely: *Should a dissertation be written as a monograph or as a pack of (published) papers supplemented with an introduction and a cover?* From my perspective, this question does not have a definitive answer, as there are merits and disadvantages to both options. Moreover, personal and pragmatic circumstances and taste are likely to have a say as well. Suffice it here to briefly outline the most important *pros* and *cons* of the two options.

A good *monograph* dissertation can provide a thorough, substantial and coherent exposition of a complex and multi-faceted piece of research, where attention and respect is being paid to in-depth conceptual development and to detailed argumentation and discussion. It can situate different aspects of the research conducted, including different sub-investigations, within one comprehensive edifice. A monograph makes it possible for the author to indulge in deliberations and analyses that will not be allowed in journal papers. Moreover, it offers a once-in-a-lifetime opportunity for the candidate to produce a substantial piece of work which will be scrutinised in its entirety and in detail by expert assessors. There are, however, also drawbacks of this approach. One is that in most cases only a few people are going to read a lengthy monograph, which implies that the outcomes will only be known by a wider readership if subsequently transformed into published papers. Also, as refereed papers rather than monographs constitute the predominant form of academic publication of our time, a candidate who has written a monograph and published no papers has fewer authentic experiences with the rules of the game than does one who has chosen to write a set of singular papers. A monograph has to be under continual "surveillance" and kept alive for quite some time, and if major changes are made to one component of the dissertation, ensuing changes to other parts of the text are likely to be needed as well. This constitutes a challenge to most doctoral students.

Furthermore, a monograph dissertation is subjected to a 0–1 decision, the outcome of which is in principle uncertain until it has been made. In other words, the work stands and falls in its entirety, which may introduce a more massive stress factor in the completion of a monograph dissertation than in the completion of a papers-based one. Finally, since a monograph dissertation is subject to fewer constraints than peer reviewed papers, such a dissertation may run the risk of becoming longer, more talkative and unfocused than a papers-based dissertation.

Naturally, the advantages and the drawbacks of a *papers-based dissertation* are to some extent dual reflections of the drawbacks and the advantages of a monograph dissertation. To be precise, for a dissertation to be papers-based, it is not enough that a number of papers have been written in the usual format of scholarly papers, they also have to have actually been published. In view of the fact that one may encounter dissertations containing a number of "papers" that have been submitted but not published, and even draft papers, this caveat is not as trivial as it may look. The obvious advantage of a papers-based dissertation is that the work gets parcelled out, completed and accepted along the road, so that the writing of the introduction and the cover – the "cloak" ("kappa"), as the Swedes call it – is likely to be much less strenuous and marked by uncertainty than is the case with a monograph. Also the candidate, when defending his or her work, has already gained experiences in publishing academic papers, and is thus on the road to becoming a member of the academic republic in mathematics education. However, drawbacks exist as well. The most important one is that each of the published papers has to be singular, to stand alone, to be short and concise, and to be stripped of all deeper reflections and discussions. Since a dissertation will usually (have to) contain a number of papers that are linked together, each paper has to spend some space on setting the same stage, presenting the same background, and possibly refer to findings in one or more of the other papers. In view of the space limitations imposed on papers by most journal and anthology editors, a fair amount of space in each paper then has to be spent on repeating things from other papers. This may tend to make the papers, and hence the dissertation, more superficial and less reflective than desirable. Also, a papers-based dissertation may well appear rather heterogenous, and sometimes even incoherent, in particular if there is a considerable time span between the first and the last paper in the dissertation, and the doctoral student has undergone a marked development during the process.

In summary, it is not clear which approach is the better of the two. A monograph may resemble a classical symphony, a papers-based dissertation a classical suite, and as we know, both genres have given rise

to master-pieces in the past. This musical analogy may be carried a little further. If the material of the dissertation forms an integrated and coherent whole, where all components are part of the same organism, like in a symphony by Beethoven, a monograph may be the better option. If, in contrast, the material lends itself to partitioning in rather clear components without sizeable overlap, like in a suite by Bach, the papers-based option may prevail. The role of the supervisor in advising his or her PhD student as to which approach to adopt is a crucial one, as some of the drawbacks of each approach can be counteracted and remedied by the supervision provided. Furthermore, it is actually possible to adopt both approaches, at least partly, even though this is, of course, a bit more demanding than choosing between the two options. My own preference – but other supervisors may settle the account differently – is to encourage my PhD students to do both, aim at writing a monograph while writing and publishing papers for conferences, journals etc. along the road.

What could/should lead to rejection of a dissertation?

In the previous sections we have primarily been looking at what gives a dissertation quality. It may seem as if there is no end to the requirements that have to be fulfilled for quality to emerge, which may discourage some graduate students. Perhaps it might be worthwhile to also consider the dual question, as posed above. Answers to that question may appear to be more specific to a graduate student who is thinking about how (s) he is going to fare once the dissertation has been submitted. Once again, the answers I am going to offer are my own. Colleagues in supervision or assessment of PhD dissertations may take different stances.

First of all, a dissertation should be rejected if the *research conducted* is unsatisfactory in one or several of the following respects:

- The *problématique* is too general or diffuse, the research questions are not genuine or mainly rhetorical (answers are known from the outset and require no further investigation), or they are not researchable.
- The *research design* adopted (or invoked) and the corresponding methods of investigation are not applicable to the research question posed – or they have been misunderstood or misused, for instance because the conditions for their use have not been observed.
- Only *trivial results* have been obtained.
- Results are *insufficiently* (or not at all) *justfied or argued for.*

If there are crucial deficiencies in these respects, even a highly polished exposition and a carefully crafted text should not suffice to prevent rejection of the dissertation.

Secondly, a dissertation should be rejected if the text as a report is unsatisfactory in one or several of the following respects:

- It is *unclear what the research contributions of the author are*, or how contributions relate to other researchers' findings.
- The dissertation displays *problems of intellectual property* and ownership, in particular concerning the author's share of this.
- The exposition *displays a mismatch* amongst some of the following: the *problématique*, the conceptual framework created or employed, the research design and the methods applied, the research findings, the basis on which conclusions have been drawn.
- There are serious deficiencies with regard to *explicit methodological or epistemological reflection* in the dissertation.
- The typical reader is likely to be left in such confusion about the structure, content and language of the dissertation that *comprehension is jeopardised*.

If there are crucial deficiencies in these respects, even a well conducted piece of underlying research should not suffice to prevent rejection of the dissertation.

In my view, these grounds for rejection hold even if the dissertation is based on published, peer-reviewed papers, although in that case the probability of rejection diminishes, at least in practice if not in principle.

Final remarks

This paper is a pretty personal one as it reflects my own experiences and positions. It places a marked emphasis on the role of the research question as the pivotal component of research in mathematics education in general and in PhD dissertations in particular. There is not universal agreement on attributing such a crucial role to the research question, let alone on insisting on a very sharp interpretation of what may count as a research question. For example, many researchers would not make the distinction between a research question proper and an auxiliary question that I am pleading for here. So let me underline, once again, that I am not claiming to speak on behalf of established quarters in mathematics

education research. Therefore, in my own supervision of research students I make a point of explicitly stating that there are other positions than mine in the community, and that, therefore, they should listen to my arguments rather than to my conclusions.

It is characteristic of *mathematics education as a field of research* (see Niss, 2007b) that it is relatively weak in terms of established paradigms, and with respect to relying on well-founded and universally accepted methods of inference. The fact that the concept, nature and role of theory are all extremely blurred is a particularly serious problem (see Niss, 2007a). Much of what is referred to as "theory" does not, in my opinion live up to basic requirements to such a notion. Concepts of quality and relevance are under constant debate, which is a healthy indication of self-reflection and critique.

Against that background there is no established royal road to guaranteed quality or relevance in mathematics education research. In that sense every PhD project is unique, which puts high demands on graduate students' independence and originality. Every PhD student must "think from scratch", even if (s)he is supervised by one of the most experienced, renowned and respected supervisors.

Acknowledgements

The author is grateful to the reviewers for careful reading of the manuscript and for excellent suggestions which have helped improve the paper.

References

Grevholm, B. (2008). Nordic doctoral programmes in mathematics education. In R. Reys & J. Dossey (Eds.), *U.S. doctorates in mathematics education: developing stewards of the discipline* (pp. 189–194). Providence: American Mathematical Society.

Grevholm, B., Persson, L.-E. & Wall, P. (2005). A dynamic model for education of doctoral students and guidance of supervisors in research groups. *Educational Studies in Mathematics*, 60 (2), 173–197.

Lester, F. & Lambdin, S. (1998). The ship of Theseus and other metaphors for thinking about what we value in mathematics education research. In A. Sierpinska & J. Kilpatrick (Eds.), *Mathematics education as a research domain: a search for identity* (pp. 415–426). Dordrecht: Kluwer Academic Publishers.

Niss, M. (2007a). The concept and role of theory in mathematics education. In C. Bergsten, B. Grevholm, H. Måsøval & F. Rønning (Eds.), *Relating practice and research in mathematics education. Proceedings of Norma05, fourth Nordic conference on mathematics education* (pp. 97–110). Trondheim: Tapir Academic Press.

Niss, M. (2007b). Reflections on the state of and trends in research on mathematics teaching and learning. From here to Utopia. In F. K. Lester (Ed.), *Second handbook of research on mathematics teaching and learning* (pp. 1293–1312). Charlotte, NC: Information Age Publishing.

Schoenfeld, A. (2007). Method. In F. K. Lester (Ed.), *Second handbook of research on mathematics teaching and learning* (pp. 69–107). Charlotte, NC: Information Age Publishing.

Mogens Niss

Mogens Niss is a professor of mathematics and mathematics education at Roskilde university, Denmark, where he has been working since 1972, the year the university was founded. He was a member of the Executive committee of ICMI, 1987–1991, the last eight years as the Secretary general of the commission. He is currently the chair of the ICMI Awards committee, a member of the Education committee of the European mathematical society, and a member of the board of National centre for mathematics education (NCM) in Sweden. His research interests and many publications focus on mathematics education, in particular mathematical competencies, mathematical modelling, the nature of mathematics education as a research discipline, and the justification problem of mathematics education in society.

mn@ruc.dk