# Thoughts on the roles of the student and of the supervisor in the different phases of a PhD

Tobias Galla
*Theoretical Physics, The University of Manchester, Manchester M13 9PL, UK**
(Dated: August 15, 2018)

I describe my personal view of the different phases of PhD studies in theoretical physics, and how I see the roles of the student and the supervisor in the different phases. I explain what I think PhD students can do to progress from being a 'student' to becoming a 'scientist'.

These notes are for PhD projects in theoretical physics, perhaps also applied mathematics. In these disciplines papers usually have a relatively small number of authors, often just two – the student and the supervisor. I don't feel qualified to have a view on PhDs in other areas. The notes are based on my own experience as a PhD student (some time ago), and from supervising and co-supervising PhD students. Perfectly sensible people may disagree with some of what I say, or with all of it. These notes would probably be very different had I written them five years ago, or if I were to write them again in five years from now.

## I. THE DIFFERENT TYPES OF PROJECTS

This refers to individual research projects, a paper perhaps, or one thesis chapter, but not to entire PhD theses.

Type A:
I will call a project type A if it has the following two characteristics:

1. The supervisor could complete the project on the time scale of one to two months, if he/she had time to work on this full time. By 'completing' I mean submission of a paper to a reputable peer reviewed journal.

2. The supervisor has a very good idea of the overall result from the outset. The story of the paper is clear from the beginning. Perhaps not every single figure, but the essential story.

Type B:
For these projects, the supervisor can be very confident (but not certain) that they can produce a paper within three months if they could work on this full time. The story is not clear from the outset, but the supervisor has possible ideas, and high confidence that 'a' publishable story will result.

Type C:
These projects are such that the supervisor cannot be sure whether a publishable paper will result.

Projects are not always uniquely type A, B or C. This is a continuous scale, a project can be between A and B, or between B and C.

* tobias.galla@manchester.ac.uk

Of course projects can change type as time goes by. All projects ending up in a publication turn into type A near the end. A project which is initially type C can turn into B, then finally A, and then a paper. Projects can also change from type A to type B, typically when a type-A project is absorbed in a more interesting type-B project which comes up along the way.

There is a further class of projects:

Type E (E='external'):
These are projects relying on collaborators outside the immediate research group of the supervisor, for example work with experimental colleagues, in another department, or at other institutions. This means that progress on type-E projects cannot be controlled by the supervisor or the student. Type-E projects have the properties of a type-C project, the outcome is uncertain even if the supervisor could work on this full time.

There is no type D in this document.

## II. THE PHASES OF A PHD

The following does not apply to every PhD, but overall I think this is roughly accurate. In some cases the phases may overlap, students move from one phase to the next gradually.

Phase I: Initial phase: orientation, getting into the research mode
You have no idea what you are doing (no offence); you do not have the experience to see your project in a wider research perspective. You are unable to place your project in relation to other things you will do in your PhD. You have no idea what the remaining 2–3 years of your PhD are going to look like. That's not surprising — it is the first time you are doing a PhD after all, and you have not seen one through from beginning to end.

At this point, you have no choice but to trust your supervisor that they have selected your project with the wider picture in mind. Your supervisor should tell you what that wider perspective is and how your initial project relates to plans for the later stages of your PhD. You may not fully understand this, but you (hopefully) develop some trust that your supervisor has this covered.

You understand from seeing students in higher years that doing a PhD is about writing papers. After some time you would like to work towards a paper. But — due to lack of experience — you have no idea how to do this. This is often the most critical phase of a PhD. It is stressful — you feel lost, not in control, that you are 'not good enough'. This is normal.

Phase II: Writing the first paper
At some point it is finally time to write the first paper. Your supervisor will hopefully start nudging you in this direction already in Phase I, and perhaps outline the goalposts from the beginning. You will be excited — finally your work gets written up, one step closer to the PhD. There is now a concrete goal: finishing off that paper, and getting it submitted. At the same time this is the phase in your PhD when you learn the most about the paper–writing process. How drafts go back between student and supervisor, or between collaborators. The supervisor makes major edits, asks you to re-write a section after you spent days and days, tears and sweat writing just that section. You were so proud of it, and now your supervisor says it has to be re-written from scratch. Frustrating, but this is how you learn. You realise

that the paper–writing process takes time and hard work, more than you expected probably. It can easily be a few months from the moment you create 'paper_v1.tex' and actual submission. How long exactly partly depends on the quality of the technical notes you have been writing, and on how careful and precise you have been in your work. And on your writing skills. I find that some students just know how to write — 80 or 90% of their first draft can stay unchanged. Others have to learn it (I was in that category). Of course the supervisor could just go and write the paper for you. But that's not the point. The point is that you learn how to write a paper so that it will be easier the next time.

Phase III: Production phase

The first paper is submitted. Time to work on the next, and the next. You have gone through the project cycle once now — from the beginnings to submission. You have some of the broader view that you didn't have in the earlier phases. You have acquired general research skills (Sec. VI), allowing you to be more efficient in the prosecution of your next project. You know how to do certain analyses or simulations. You have computer codes already that you can adapt. You may be working on two projects in parallel. Things are less frustrating now (but not entirely free of frustration), some confusion about the overall direction of your PhD remains. You may still feel a bit lost from time to time, but it is less intense than before. This phase is really the 'fun' phase of the PhD. You know you can produce. How much you produce depends on the number of hours you work, not on factors beyond your influence. In the earlier phases you may have felt that it didn't really matter whether you came into the office on any given day, it was all pointless anyway. But not now. You are in control. Work hard!

Phase IV: Creative phase

This phase usually overlaps with the production phase. Of course you have been creative in the earlier phases too. But it was creativity on a technical level — how to solve a particular equation, how to do a simulation, how to make it faster. All these things require creativity, and you have come up with solutions. Creativity in the sense of Phase IV means creativity at a different level — at the level of project design. You are now able to generate ideas for whole projects on your own, and to shape them. Perhaps you read a paper and have an idea how to improve or extend something in that paper. Or you listen to a talk at a conference and that starts a project. A key step in your education is reached when you are able to determine with good accuracy whether a particular idea is of type A, B or C, and how it fits into the bigger picture of your PhD. In fact type A now means that **you** are able to complete the project — perhaps with some input from the supervisor, but strictly speaking you are not really dependent on him or her any longer. If your supervisor disappeared from one day to the next, you could still complete.

Phase V: Wrapping and writing up

You have reached the point where two or three papers have been submitted or published. Perhaps even more than that — one or two further papers are almost finished. And few loose ends to tie up, but in essence your thesis is there, you know what the chapters will be, and what you need to do to complete them. You'll have to write an introduction too. And you may be looking for a job. This phase is more stressful again. The clock is ticking, your funding is running out. You work longer and longer hours. You want to finish.

## III. THE OVERALL GOAL: TRANSITION FROM STUDENT TO SCIENTIST

The goal is obviously to get a PhD. This means producing 2–3 papers on which the student *legitimately* is the first author, and making sure that the student is in a position to defend these in their 'viva' (=oral examination). The papers do not all need to be published, there is no formal requirement, at least not in the UK. But the general convention is that it is ok, and perhaps expected, to have two papers published and another submitted. Four or five papers, of course, is better. Producing papers is not the only goal, but it is a big part of getting a PhD. If you have three first–author papers published in reputable journals, it will be very hard for your examiners to fail you.

It helps to make the production phase and the creative phase as long and productive as possible. This does not mean that the initial phase should be as short as possible. A degree of initial struggle, feeling lost, and frustration is needed — you want *legitimate* first authorship on your papers. Of course, your supervisor could try to solve that initial problem for you. But that's not the point. The point is to get you to a stage where you can enter the production and creative phases. Entry into these phases is not a fixed point in time (end of first year, after 18 months, etc). You enter the production phase when you are 'ready'. It happens organically — no one will send you an email to tell you that the next phase has started. Students often go through some sort of phase transition, from being a slightly disoriented Phase-I student to a scientific collaborator and equal in Phases III and IV. Sometimes this transition happens quickly, over the course of a few months.

How long the initial phases last is a function of two things:

(a) the type of your initial project, and

(b) the student (you).

I am still trying to understand whether it is mostly (a) or mostly (b), and what the supervisor can do to make the 'phase transition' happen sooner. I do not know. Phase I tends to be shorter if the initial project is of type A. But then, if two different students were 'given' the same type-A project, they'd probably reach Phase II at different points in time. It is really a function of the project and the student. In my experience the technical ability of the student is often not the dominating factor. Instead it is more about how the student 'interacts' with their project, what approach they take to it and to research in general. See below ('Your jobs') for what you can do to progress as quickly as possible.

Part of the frustration in Phase I comes from the feeling that things seem to be progressing much more slowly than when you were an undergraduate. You are used to problem sheets, and you thought that you can solve any question if you only spend enough time on it. Example sheet 7 in Advanced Quantum Physics? No problem, give me six to eight hours and I'll get there. Suddenly things are different. Nearly all PhD students will have that experience in some form. Students who are used to getting full marks most of the time are especially prone to frustration — they don't know what it is like not being able to solve a physics problem. Is it normal that things go so slowly? Is it you? Are you 'too slow'? Or is it the project?

There is no way for the student to know. You do not yet know that projects can fail (especially types B and C), how they fail, and when it is time to abandon a project. If a project fails it is not necessarily your fault (or that of the supervisor). Things fail, this is research after all. What matters is how they fail.

The better the project was thought through at the beginning (by the supervisor), the lower the probability that it will fail. But that probability is never zero, not even for type-A projects. One can only develop a sense of this when one has seen tens of projects of the different types to completion (or failure). If it is your tenth project, and they are all slow or fail, then its likely to do with you. If you have only done one project and it is slow (or fails) it could be the luck of the draw, or it could be you. It is your supervisor's job to tell you, and what to do about it.

## IV. THE SUPERVISOR'S JOB

In Phase I:
In the beginning it is my job to come up with ideas for your first project. I normally try to suggest a few possible initial projects. This depends on the preferences you have expressed (what topic you want to work on) and the availability of suitable ideas for projects in that area. Ideas sort of come up as a random process (in conversations, walking in the street, in the shower, ....) — usually at a sufficient rate, but it is a little random. So good project ideas may not always be available in all areas at all times.

Anyway, at the beginning I will provide a few ideas for projects to work on. These will normally be of type A or B. It is partly up to you whether you'd like a type-A or type-B project at that point (or in fact types C or E). They all have advantages and disadvantages. Type-A projects can lead to a first paper quickly, this minimises frustration and gives you a feeling of satisfaction relatively early on. But often the results are not that exciting. Types B and C can lead to great results, genuinely new things, but they come with more frustration. Projects with external collaborators (type E) are particularly tricky, progress cannot be controlled 'in-house' by student and supervisor. But again, they often lead to the most exciting results. The best initial projects are probably of type A but with the potential of becoming type B, or projects of type C or E, with a type-A/B fallback option.

In Phase II:
Phase II is about writing the first paper. The supervisor now has two main jobs. The first is to look more closely at what you have done — his or her name will be on the paper after all, so they need to be convinced that what is in the paper is actually right. So far they may have followed your work 'from a distance' — they will know roughly what you have done, and they see the general picture. But they will not have checked every factor of two. Now is the time for your supervisor to make sure that they understand what you have done at a technical level, and to suggest corrections or modifications. He or she will ask you a lot of questions: 'What exactly did you do here?' 'How exactly is this equation derived?', 'Tell me exactly what parameters you have used for the simulations in Figure 3.' So far your supervisor may have not have cared about these details, but suddenly they do. Details matter. Having kept good notes and records is crucial so that you can answer these questions.

The second main job for the supervisor is to guide you through the paper–writing process. They have to make sure that you learn to write 'to publication standard'. Some supervisors will actively edit the paper, others just give feedback and expect you to implement it. Together with the supervisor you will need to decide what journal to submit to — the supervisor will normally have a better idea of what is suitable and realistic. Finally the supervisor needs to tell you how to write a covering letter, how to actually submit a paper, how to choose potential

referees, and of course how to respond to referee reports.

In Phase III:
In Phase III you have reached a degree of independence. This means that your supervisor's role is increasingly to take a step back, and to 'just let you do'. At this point your relation changes from supervisor–student to that of equal collaborators — the phase transition. Your supervisor will still be more experienced in running projects and writing papers, but by now you probably know more about the specific area of your project than the supervisor. You are now two scientists collaborating with each other. You'll still be doing most of the calculations, simulations and data analysis. Your supervisor will make suggestions, point out connections to other areas, propose projects etc. But over the course of Phase III (and hopefully IV) you complete the transition from being a 'student' to becoming a 'scientist'. The PhD is a licence to practice science, and you are getting closer.

Increasingly you and your supervisor should be thinking of your thesis as a whole. The first paper is submitted, work for the second in progress. That's half a thesis already. What are the other parts going to look like? What are projects three and four?

It is also time to present your work at conferences, and to start making contact with researchers in other places working on things related to your project. It is your supervisor's job to point out who these groups are, and perhaps to put you in touch. In coffee breaks at conferences your supervisor can introduce you. Questions at poster sessions or after conference talks can be scary. It may look a little cut-throat sometimes. Your supervisor is there to protect you, and to interpret what you see. What may look aggressive to you in reality often is just normal scientific debate, and nothing to be scared of. At times these debates are heated — this is not personal, everyone is just very passionate about their work. You will meet scientists from other countries — they will have different personalities and cultures, and different ways of expressing criticism and sentiments. Your supervisor can help you interpret and 'renormalise' things.

Phase IV:
Your supervisor will (should) encourage you to be creative. This can mean a project based on an idea you have (not the supervisor), and that you then complete with the help of your supervisor. Ideally, you should also try write at least one paper during your PhD without the involvement of your supervisor. Something you and a postdoc start. Or perhaps a coffee–break conversation with someone at a conference that you follow up on. Something you come up with on your own, shape it from having the initial idea, to turning it into a viable project, carrying it out and writing it up. Careful though not just to modify something your supervisor once mentioned; by definition this has to be a genuinely separate idea coming from you, not from the supervisor in some direct or indirect way.

It is your supervisor's job to give you freedom and to encourage you to grow – *provided* they are sure that this does not put your PhD in jeopardy. Your supervisor has to keep the overall picture in mind. While they should give you freedom, it is also their duty to protect you and to make sure you do not take unreasonable risks — risks invisible to you due to your inexperience. If a supervisor lets a student do a project on their own, the supervisor has no control over it (by definition they are not involved). The supervisor may ask themselves the following questions: What if the student works on this for six months, but nothing comes out of it? Will that student still get a PhD? Has the student done enough other stuff? Doing a project on your own has huge potential benefits if things work out. But is the risk acceptable?

If the answer is yes, your supervisor should encourage you to work independently. If you have done enough 'other stuff' and your supervisor still insists that they need to be on all your papers, then something is wrong.

Phase V:
At this point the supervisor plays a relatively minor role. They can suggest examiners, they will (should) proof read your thesis and make comments. A good rule is to assume that your supervisor will read one draft of your thesis (maybe chapter by chapter) and return this to you with scribbles in the margins, or annotations in a PDF. But you can't expect them to go through draft after draft, or to edit the latex file. That's your job, not the supervisor's.

## V. YOUR JOB – PART I

Phase I
In this phase you have two main jobs:

(1) To get used to doing research, and

(2) Not to give up.

I describe (2) first. Phase I is the hardest phase of your PhD, due to the frustration described in Sec. II. Your supervisor will hopefully help to keep this to a reasonable level, with clear and specific instructions. But things are never going to be easy. Ideally your first project is relatively well defined and not one where 'finding the problem is part of the problem' (type C). However, occasionally no other projects are available. In this case your job is to read, to categorise what is known and to summarise it.

You also need to get used to doing research. You need to learn how to approach research problems. This entails writing notes, working with your supervisor, using existing resources (such as fellow students), etc. Sec. VI describes in more detail how you can learn these skills. It is important that you get used to the right approach to research in Phase I. You need the skills described in Sec. VI to progress to the later phases. It is part of the definition of a Phase-III student that they have mastered these skills.

Phase II:
In Phase II you are writing your first paper. This is hard and it takes time. Hopefully, you will have typed technical notes already, now it is time to morph these into a paper. You and your supervisor will probably sit down and plan the paper. Section 2 of the paper is going to be about this, in Section 3 like that; in Figure 1 we are going to show this, and so on. Your supervisor will then probably ask you to write a first draft, perhaps not the introduction and conclusions (these are the most difficult sections). He or she knows that you have probably never written a research paper. You may struggle a bit. That's fine. But within the possibilities you have do the best job you can.

Most of all, do this conscientiously and carefully. I have had students send me latex files that did not compile, or with text randomly typed in the middle of equations (clearly they had not even bothered to look at the PDF). Few things are more frustrating for a supervisor than a student who writes their first paper carelessly. It is your 'first baby' after all, so if you are not excited and want to get it right, then who will?

You will have read many papers by now, so you know what a good figure looks like, what different styles there are for legends, how people write figure captions, how existing literature is referenced. If you do not know how to do these things, look at papers you have read. How exactly do they cite journal papers and books? Any statement in your draft has to be backed up by either a reference, a calculation you did or a simulation. You can't just say things. Whenever you cite an existing paper or book, be prepared to explain why this reference and not another? Writing a paper is not about citing all references you have ever looked at; make sure the ones you cite are selected for a reason. If you need a reference for the so-and-so model, don't cite a random paper in which that model is mentioned. Find the paper where the model was first proposed.

The purpose of writing an initial draft of your paper is to be able to edit it later (you can't edit a blank page after all). Don't be disappointed if entire sections of the first or even the seventh draft get deleted in the end. That's normal. When your supervisor writes a paper, they will do it the same way. I have had students tell me 'But I spent a lot of time on this intro', after I had said it had to be re-written. Well, that's life when you are a writer – 'the first draft of anything is sh*t' (Hemingway). But the first, second and seventh drafts still need to be written, to be able to get to drafts eight, nine and seventeen.

If your supervisor is actively editing the paper, and you think something they write is wrong or inaccurate, then tell your supervisor. Don't assume things are right just because the supervisor wrote them. You share ownership of the paper, and all parties must stand by everything that is in the paper.

Paper writing usually proceeds from long length scales (sections) to smaller length scales (individuals words and commas). First get the overall structure agreed with the supervisor, then gradually work to finer and finer levels. Importantly, once the editing has reached a relatively fine level (for example, the scale of paragraphs) don't go back and make unilateral changes at longer scales. It is frustrating for a supervisor if a student undoes what the supervisor did. Loops of doing–undoing–redoing–re-undoing–... are not efficient. If you think you need to go back to longer scales, discuss with your co-authors first. If more senior authors changed something in what you wrote, and you don't like the change, don't just change it back. They probably did it for a reason. They may still be wrong of course. Best to discuss the matter directly and to make a joint decision.

Phase III:
Phase III is probably one of the most pleasant phases, you are now in the game, you know what you are doing. Things are exciting. On the other hand the clock is ticking. Keep going. Stay focused on the overall goal. Your first paper is submitted, it is time to think about papers two and three, and how to produce them efficiently. At the same time your job is to broaden your network and conferences and meetings. Be sociable, talk to people. This is particularly important if you want to continue in science after your PhD. Many postdoc positions are effectively won during coffee breaks.

Phase IV:
At this point you have gained some experience, have read many papers, and you are in a position to decide what is and is not viable. Once you have published/submitted two papers (usually with the supervisor), you can take risks, and come up with your own projects, ideally without your supervisor. When I had my second paper accepted I was told 'Now it is no longer a question *if* you get your PhD, but only *when*.'

If you want to do independent work, you should tell your supervisor. He or she can warn you if they think that your idea is hopeless. But bear in mind, you are now equals — if your supervisor tells you that an idea is no good they could be wrong (with hindsight I believe I have given incorrect advice on a few occasions). Use your own judgement — if you think your supervisor is wrong about your project idea, do it anyway. But also keep in mind that you alone are responsible for its success or failure. Developing the ability to gauge risk and to accept responsibility for projects is a big step towards being an independent scientist — one of the main goals of doing a PhD.

Phase V:
At this point your job is to finish the thesis and to submit it. Once you get to this stage, it will probably be relatively clear what you need to do. Keep in mind that this is *your* thesis and that you alone are responsible for it. You have written papers with your supervisor. Both names are on these papers, and you had joint responsibility for them. Your thesis only has one author.

## VI. YOUR JOB – PART II

**Top priority – Keep your supervisor engaged:**
This applies in all phases. Supervisors have several students to supervise, several projects to handle, many things to think about, emails to respond to, people to talk to. You probably won't see most of this 'other' activity of your supervisor, but what it means for you is this: it takes effort and energy on the part of the supervisor to actively make contact with you, to check how you are doing, and to go and see you and say: 'Should we have a meeting?' It is your job to push your project, not the supervisor's job. If you do not take charge of your project, interact with the supervisor, and make it easy for him or her, you'll simply be pushed to the background. Phase I in particular will take longer than necessary.

There are many ways of keeping your supervisor engaged:

(i) Send them notes and plots. Make sure these notes are self-contained. Don't assume that the supervisor knows what definitions you use. You have quantities $\alpha, x, \ldots$ and you know perfectly what they represent. But your supervisor talks to lots of people, and everyone has an $\alpha$ and an $x$. If you send figures with incomplete or no captions or without axes labels, it is easy for your supervisor to simply reply 'add labels and come back to me when you are ready'. If you want your supervisor to look at something, make it easy for him/her. Do not give them an easy way out (=an opportunity to return things to you unread). Print your notes and read through at least once before you send them to your supervisor. Your notes do not have to be long, but they should be professional. Messy notes make you look unprofessional, and your supervisor will not take them seriously ('Alright, this student is not yet mature enough, not yet ready for research.'). Needless to say that well-written and thoughtful notes also help your own understanding of the problem you are doing. It is not about writing notes because your supervisor said 'write some notes'. Think about what you want to achieve with them.

(ii) Make sure you go and see your supervisor on a regular basis. This does not have to be every week, but if you have not seen your supervisor for a week, send them an email to say what you are doing even if you don't think a meeting is necessary. Your supervisor

may not actively check on you, but they will appreciate if you explain why a meeting is not required.

(iii) When you go into a meeting, bring a printout of your notes with you, or the figures and calculations you'd like to discuss. Ideally send notes before the meeting. Your supervisor may not read them, but then they will feel guilty and spend more time talking to you. You can then talk through the notes in the meeting.

**Do what your supervisor says and more:**
Some students persistently resist suggestions. They simply do not do what the supervisor asks them to do. This could be a particular simulation, calculation or to work through a paper. It is very frustrating for the supervisor to have to ask the same student again and again to do a particular thing. The supervisor has had an idea, is excited, their fingers are itching to do it themselves, but they want to involve you. If there is no response you signal that you are not interested. The supervisor will not ask you again to work on new ideas. You lose out on opportunities to become involved in new projects. If you already have three projects to work on or are finishing a paper, then of course it's fine to say 'no' or 'yes, but in a few week's time'. But other than that, don't be that unresponsive student.

Just to avoid misunderstanding: you are not the supervisor's technician. A PhD is not a series of 10,000 plots that your supervisor asks you to do one after the other. If you think a particular suggestion does not make sense, you should say this to your supervisor, and discuss it. Voicing these concerns is part of your job, especially in the later parts of your PhD when you are more expert than your supervisor.

But in the beginning it is probably fair to assume that your supervisor makes suggestions for good reasons and with a plan in mind. This could be purely educational: they want you to do a certain thing, perhaps reproduce a figure from an existing paper, just for you to learn a particular method that you will need later. Or they have a hunch that a certain simulation could move your project forward. A hunch that requires experience, which you do not (yet) have.

You should follow your own ideas *in addition* to the suggestions from your supervisor, not instead of them. Productive meetings start like this: 'I did what we said last time, here are the plots, and it turns out that ... and so I also did this here ....'.

**Manage 'upwards':**
Understand what you can and cannot get from your supervisor. Some will read your notes in detail, others won't. Some will be available all the time, others won't (realistically this will probably vary at different times). Some will get involved in the calculations, perhaps even do simulations. Some will check your factors of two, others won't. Some will actively send you to conferences, others won't. Some will introduce you to people, others won't. Try to get a good idea of what your supervisor will realistically do, and what they won't provide. Then adapt your strategy accordingly to make sure you get different things from different sources (postdocs, fellow students, ...).

**Come into the office regularly:**
A PhD is not a 9-5 job, and no one checks attendance or records how many hours you work

and where (office, home, coffee shop ...). It is all up to you — what is expected is that you actually *want* to do a PhD. There are several good reasons for being in the office regularly:

- You benefit from interaction with other PhD students. You can ask them questions, discuss their projects and yours.
- A regular working schedule helps you to keep discipline, especially in the more frustrating phases of your PhD. It is easy to start drifting if you are not in regularly.
- Your supervisor may come to your office spontaneously and ask for a detail of a calculation or simulation. Maybe they just had an idea and want to talk to you. If you are not in, there is a barrier – your supervisor has to email you, or remember what they wanted to say the next time they see you. Not an issue usually, but if this happens on a regular basis it could turn into a problem. Make sure your supervisor knows when they can and cannot expect you in the office. Apply common sense. No need to tell them about each and every tutorial you give, or trip to the shop or doctor you make. But if you are away for a few days, let your supervisor know. If you regularly do 'something else' on Thursday afternoons, make sure your supervisor knows.

**Displacement activities:**
It is very easy to get trapped in displacement activities. This means that you are not doing your actual job, but instead you fill your day with seemingly purposeful activity. For example, you spend five hours marking tutorial work when one hour would be sufficient. It means fiddling endlessly with some tricky latex business when simple revtex would produce essentially the same outcome. Things like this. Everyone does it. Maybe writing these notes is a displacement activity — typing this here is easier than editing the paper I am supposed to work on. Anyway, a bit of it is fine, but keep it under control.

It helps to realise that the purpose of your day in the office is not to get to the end of the day in some way that you can justify to yourself and then to go home with a clear conscience ('I was very busy marking tutorial work'). Your main job is to progress your PhD. At the end of each day ask yourself what you did that day to advance your project. Ask yourself at the end of each week: what did I do this week? If an outsider looked at this, would they agree that I did a reasonable amount of work *for the PhD* that week?

**Exchanging files on Dropbox:**
Starting in Phases I and II you will exchange notes and paper drafts with your supervisor. A shared directory on Dropbox is a very convenient tool.

Keep files organised and with clear version numbers in the name of the file. For example, 'topic_notes_31July_TG.tex', where the initials at the end indicate who last edited the file. During the paper editing process, you are likely pass the manuscript back and forth between you and your supervisor. Whenever you receive a version from your supervisor, create a copy and give it a new name (new version). Best to always work in a 'live' shared folder. When you have completed your edits email your supervisor. Tell them 'Version X is now on Dropbox for you to read/edit'. Unless you tell your supervisor they won't know that it is their turn.

Make sure your Dropbox folder is synced. People work at odd times, late at night, or early in the morning. Perhaps your supervisor suddenly has an hour or two on a Saturday afternoon to look at something. It is frustrating if they can't find the latest version of your draft manuscript. This is made worse if the response to the email asking for the file is slow, see the next item.

**Read your email:**
Make sure you look at your work email on a regular basis during working hours. Your supervisor may email you at 10 and suggest a meeting for 2. They may ask for a particular file or plot. Of course no one expects continuous online presence during the night or at the weekend. But your supervisor's frustration levels will rise if responses persistently come with a delay of two working days. You should be the main driver of your project(s).

**Read the arXiv regularly:**
Other people may be working on projects related to yours. In most cases you and your supervisor will not know about this. That's why it is important that you regularly look at the relevant sections of the preprint server arXiv, ideally each day. Depending on your research area also look at other servers (INSPIRE, bioRxiv ,...). It is useful to subscribe to the preprint servers' email lists so that you receive daily messages with new abstracts. Scan these for anything that is broadly relevant for your project(s). Then investigate this selection more closely. Best to make this a routine. Spend the first 15-30 minutes of the day looking at new preprints. If you find something that's very close to your own project, alert your supervisor.

Apart from the obvious benefits of being up to date, reading new things, building up expertise and experience, the idea is to avoid the following scenario: you work on a project for nine months or a year. Your supervisor follows this 'from a distance', but they do not necessarily look at new literature in detail. When it is time to write the paper, you and your supervisor find that there has been a paper on the exact same thing on the arXiv six months ago. By now that preprint from six month ago is published in a journal. Your own paper dies here and you can't put this piece of work in your thesis any longer, at least not in the way you had planned to. Someone else did it before you, so it's not new anymore. Back to the drawing board.

This is an extreme case of course – it is rare, and you can protect yourself against it by following the literature. Project ideas can usually be adapted in the early to medium stages if a new preprint comes out. Even if that other preprint comes close to what you wanted to do and it all looks very scary that morning, projects can normally be 'rescued' — alright, these people did it, but there is this interesting other question here, so let's go in that direction. In fact a preprint by someone else along the way can sometimes make your own work better. But you can only react if you and your supervisor know immediately. You don't want to discover that other work from six months ago at the very end, after you have spent weeks polishing your own paper.

Your supervisor may also be checking the arXiv, but perhaps they have less time and do not look every day. At any rate, you can't be sure if they do or don't. Assume that it is your responsibility to follow the literature, and to let your supervisor know about anything important.

All this is very difficult in Phase I when you do not yet know what is and is not relevant. But try anyway. It gets easier with time. After a while you'll know who the key players are in your area, so you can watch out for certain authors. With a bit of experience you'll find that you do not need to look at a paper for very long to tell whether it is worth reading in more detail. But you need to build up that experience first. It is good to get in the habit of checking the arXiv from Phase I on.

**Summarise literature:**
The previous item was about looking at *new* literature as it comes out. Another aspect is to be careful about studying *existing* literature from the past. In Phase I you'll spend a lot of time doing that. Your supervisor will probably point you in the direction of a few papers and authors. In the beginning you will not understand much of these papers. That's normal. Papers are written by experts for experts, and you are still on your way to becoming one. Keep in mind that papers are not read from beginning to end, line by line. At least not in the initial reading. Instead try to understand the main message of the paper – the key results. Look out for the main equations and for something you recognise. What are the main figures?

I find the following approach very useful: whenever you start a new project and want to summarise the existing literature, write a 3-sentence *biased* summary for each relevant paper. By 'biased' I mean 'from the viewpoint of your own project'. Perhaps only one subsection of the paper is really relevant. Then summarise this part of the paper, and leave out the remaining 80%. Identify what in the paper is really important *for your project*.

Summaries could read like this 'Paper by X, Y and Z. Title, journal reference. Were the first to introduce the so-and-so model of phenomenon x. Based on simulations. As parameter A is varied, find a transition between a phase in which this happens, and another phase where that happens.' That's good enough. Or perhaps: 'Paper by A, B and C. Analysed the model of F and G (see paper so-and-so). Confirmed earlier finding that this causes that, but with different method. Unlike F and and G they measure $\gamma$, not $\delta$. Simulations more precise, were the first to integrate the so-and-so-equation up to times of ten-to-the-seven.' Or maybe 'Paper by U. On a completely unrelated topic. But their equation (17) looks similar to Eq. (4) in my notes from August 15, 2018. In subsection III.C they use the bla-bla method. May be able to use this for our problem. But not sure, could be problematic because of ... .' You get the idea.

Writing summaries like this is useful for several reasons:

(i) You'll know where to look things up if you need them. You'll probably revisit these papers again and again. Having a list of concise summaries makes this process easier as you don't have to start Googling around each time (I have been there — 'I think there was a paper about X, but I can't quite remember by whom. Was this Smith, or Muller? Around 2013 I think, but not sure. Phys. Rev. E, or maybe not ...?')

(ii) When the time comes, these summaries will make it easy to compile the list of references for your own papers, and to know what paper to cite where in your manuscript.

(iii) Perhaps most importantly, making these summaries trains you in quickly determining what is and is not important in a paper. You learn how to identify the key messages of a paper, and to summarise them concisely. As you progress this skill becomes more and more important.

**Go to seminars and talks:**
This one of the most important part of your education. You will come out of many many seminars thinking that this was a waste of time. But ultimately you build up a body of 'things that you have heard about', and you will benefit from this later in Phases III and IV, and in your viva when you'll be asked about things in the broader area around your project(s).

Even talks on topics which are seemingly unrelated to your project can turn out important (within reason – if you work on complex systems there is no need to go to each and every seminar in cosmology, and vice versa). But the default assumption should be that everything in the general area of your PhD is interesting (e.g. 'statistical physics', 'particle physics'). Not going to seminars because 'my time is better spent working on my project' is a little short sighted. It reveals that you are not yet in full research mode. Going to seminars is a sign of maturity.

**Scientific debates:**
Scientific debates or discussions can be heated, and sometimes confrontational. This includes discussions with your supervisor. They may make blunt statements, sometimes intentionally so to get a point across. Always keep in mind that this is not personal. It is all about the science. The best collaborations are ones in which collaborator A says to collaborator B 'Your approach to this is crazy. Already line 3 in your calculation is wrong.' – B then replies equally, and in the end they write a great paper together. Towards the end of your PhD discussions with your supervisor may well have that flavour. In Phases III and IV you should certainly 'answer back', you are as likely to be right as your supervisor (perhaps you are even *more* likely to be right). Good students tell their supervisor: 'Look supervisor, you are wrong about this and I am right. And here is why ...'. Then they present the evidence.
In Phases I and II things are a little more tricky, chances are that your supervisor is right, and you are not. Tread with care.

**If something bothers you, tell your supervisor:**
One of the worst scenarios is one in which frustration builds up on the part of the student, but the supervisor is not aware. Supervisors may simply miss the signs. They may misinterpret what you say. That is why it is important that you let your supervisor know (in unambiguous terms) if you feel frustrated or if something bothers you. If you feel that your project isn't going anyhere, if you lose motivation, if you think that you are slow, that it is all a bit pointless etc. This is particularly important in Phase I.

## VII. FINAL NOTE

I hope these notes do not put you off. It may all sound very scary. It is, but then it isn't. The purpose of these notes is not to criticise anyone or to tell anyone off. It is to help you grow. Writing these notes I learned a lot about the process of doing a PhD myself, and what I can do better as a supervisor.

Always keep in mind that many students have completed their PhD before you. Your supervisor accepted you — this means that they have confidence that you can finish yours. They believe in you. They want to help you, and they are there to guide you through to the end. I hope you'll have fun along the way, despite the inevitable frustration at different points.

Tobias Galla – August 15, 2018
Comments and criticism welcome.
(tobias.galla@manchester.ac.uk)