

From my point of view...

A faculty perspective on life at graduate school

or

Here's the lowdown on how our department and groups work

or

What I think you should be doing and how

M. Scott Shell, 1/5/2016

Welcome to the research world

Science is hard. What makes grad school hard in particular is the high level of intellectual work and the pace of research. Not everyone can do it. You're part of a very select group, probably from the top 5% of your undergraduate class, but even so you will be challenged to reach substantially higher levels of accomplishment. A performance that you might have viewed as "excelling" in your undergraduate program will now only be "average" or possibly even "mediocre." That's what it takes to make it in scientific research; you must be a superstar. This can be intimidating, rightly so. But we believe that you can successfully take on the challenge, which is why you are here.

Your time in graduate school will be a defining point in your intellectual journey, and will shape your approach to problems for years to come. It is a very rich and rewarding time of your life. Still, it is easy to become stressed and perhaps bitter because of the immense challenges. One of the more difficult aspects of graduate life is the lack of regular, concrete metrics of success. Unlike your undergraduate career, it will not always be clear whether or not you're moving forward and making good progress. There are no assignments to gauge by, no regular A+ exam grades to feel proud of. Instead, you will spend great amounts of time trying new directions and failing, and then when you've had a positive result, convincing yourself (and your advisor, and the scientific world) that it is correct. Also unlike your undergraduate experience, your research success will be very independent and largely isolated from that of your peers. This can lead to high amounts of self-pressure and self-expectations.

It is easy to translate these stressors into bitterness, and then to displace that bitterness to your advisor, committee, the department, etc. This is especially likely if you lack a complete picture of the inner workings of academia, a sense of how research actually works, and a broader perspective on how a PhD fits into the scientific apparatus. With this document I hope to provide you with some calibration of these issues, and the general expectations that faculty have of graduate students, to give you a stronger context of the experience that will maximize your success and happiness in graduate school.

The math of proposals

First, however, let me tell you a little bit about my job. It is not always obvious what faculty do, because a great deal of it is administrative in nature and thus to some extent “behind the scenes.” Of all things that we do, raising money is one of the most important to a healthy research group.

Quite a bit of my time is spent writing proposals to obtain research grants. Each graduate student in the group requires me to raise ~\$70k per year, which covers stipend, tuition, benefits, and overhead as the university takes 51%. A typical 3-year grant from the NSF will cover about \$100k per year, which covers just about one graduate student when supplies, materials, and a half month of my summer salary are included.

If I have 5 graduate students in my group, I then need on average $5/3 = 1.7$ grants to be funded per year. If I’m really good, I will have a proposal success rate around 30%. (The typical rate is more like 20%, but let us consider the best case scenario.) That means that I need to write about 5-6 proposals per year to achieve the 1.7 needed to fund students. It helps tremendously, of course, if students apply to and secure fellowships on their own; such funding is also much more flexible.

Writing successful proposals takes substantial effort, around 25-40 hours each for a single-PI grant. The project narrative is typically 15 pages that are very dense in text and figures. I must also prepare: a summary suitable to the general public, a detailed line-item budget and budget justification, a specially formatted CV and publication record, mentoring plans for students, data archiving plans, and quite a bit of paperwork. For the latter, we have specialty staff that are experts in grant submission to help us.

Generally, the proposal plan must be flawless, and I am exacting in putting it together and revising its content. This is because a proposal will be scrutinized by a large panel of my peers and, sometimes, competitors. A single problem or misunderstood line could completely kill the appeal of the proposal when it is compared against competing proposals. Moreover, the review cycle is long and I only learn ~6 months after submission whether the funding has come through. If unsuccessful, I must resubmit anew and I am therefore set back a year in that funding opportunity. I don’t have a chance to revise and resubmit my proposal to be re-reviewed by the very same referees in the way that manuscripts submitted to journals are.

As you can imagine, I spend a tremendous amount of time on proposals. They are the support foundation of academic work. Graduate students tend to be shielded from these activities because we often work on proposals in the early mornings and late at night when we can be uninterrupted. But they are one of the most important and potentially time-consuming things that we do.

What my job entails

It’s useful to describe how faculty are evaluated in routine promotions. At UCSB, we must clearly demonstrate outstanding performance in each of four separate categories: research, teaching, professional activities, and service. You know about research, so let’s talk about the others.

Teaching. Our normal department course load is three quarters per year. (Incidentally, for every nine quarters that we teach, we earn one quarter of sabbatical leave.) If I am teaching a course that I've taught before, it generally needs

- 5 hours a week for lecture and lecture prep
- 3-5 hours for office hours and meetings with students or TAs
- 2-6 hours for preparing assignments and exams
- 2 hours for grading (amortized over the entire quarter; usually comes in spurts)

That averages about 15 hours per week. On the other hand, new courses require significantly greater time commitment, say 10 hours more per week, in order to iron out lectures, assignments, online materials and notes, etc. I also work on classes the weeks before and after the official dates of the quarter in various activities (prep, meeting with students, etc).

Professional activities. I travel a fair amount to give (usually invited) talks at conferences and seminars in other departments. This is an important part of advancing the impact and visibility of our group research, and it also is a major component of my evaluation during promotions. For my colleagues in other fields, my presence at seminars and conferences helps them to calibrate the impact of my group research. I also contribute to the execution of conferences and meetings as a session chair and conference organizer from time to time. This too helps with visibility, but it also is an essential part of service to my field that keeps it healthy and vibrant – conferences do not happen on their own! Because of my travel, I will miss some of our regular meetings. Please do not assume this is a sign that you and your project are not important to me—quite the opposite!

Service. I contribute to many administrative tasks in the department, college, and university. The university is run and governed by faculty, not staff administrators, in that faculty are responsible for all major policy decisions. In the department, each faculty member sits on at least two administrative committees, such as undergraduate affairs, graduate affairs, graduate admissions, ABET accreditation, diversity, endowment, faculty recruiting, merits & promotions, public relations & alumni outreach, undergraduate lab & safety, and seminars. Post-tenure faculty also participate in administrative committees related to the college, various research centers, and the university.

There are many other activities that draw upon my time, but that do not fit into the three promotion categories above, because they are too routine:

Of course, I aim to meet with all of my graduate students weekly to discuss research progress and results. Really, these meetings are a placeholder in my schedule to make sure that the student has time with me, rather than for me to make a weekly assessment of the student's progress. These meetings typically last from 30-60 minutes.

I interact with collaborators and visitors through email, phone calls, and in person. I regularly meet with seminar speakers in our department and frequently with speakers in other departments that have common research overlap. I also plan research activities, proposals, grant reports, and papers with collaborators.

I spend significant time editing and giving feedback on documents written by students: progress reports, manuscripts, abstracts, thesis chapters, etc. When I travel, this is usually my go-to activity on the plane. More on manuscript drafts later...

For each grant I hold, I prepare yearly reports for the funding agency to show that we are making good progress and the grant funds are being spent responsibly. Generally such a report requires 10-15 hours of my time to put together a research update, list of products (publications, conference presentations, etc), and documentation relating to the training of students and/or postdocs. I usually ask the student on the grant to help prepare the research report part of this.

I often review manuscripts submitted to journals (15-20/year) and proposals submitted to funding agencies (5-10/year). Sometimes I am not only a reviewer of proposals, but a “panelist” that reviews about 9-12 proposals and then travels to DC for an in-person panel assessment that lasts two days. Reviewing is an essential part of the science biz that drives the credibility and quality of published research.

Incidentally, I will frequently ask you to help me review manuscripts. Don't let yourself view this as an annoyance or distraction from your research focus. Instead, it is a normal contribution to the process of science that also helps you formulate your own critical eye. Moreover, it is effort that will be reciprocated in your own manuscripts, which will be reviewed by 2-4 referees each; therefore, steady state happens when you contribute that number of reviews per each of your manuscripts.

How I work

For faculty, a work week averaging 60-70 hours is standard.

During the day, I am constantly switching gears to balance various priorities and deadlines associated with all of the activities above. Rarely can I work on or think about something for more than an hour or two. For that reason, I use mornings and evenings as times to address more substantial undertakings that require constant focus for multiple hours, like proposals.

Except during lunch, I almost never surf the internet or check up on news sites, etc. when I am working. That would only prolong my day.

My favorite thing to do is *think* – interpret results, generate new ideas, work out possible directions. Unfortunately, I don't have as much time as I would like to do this, but I tend to make sure I can allocate a few hours a week. Conferences and meetings tend to be excellent venues for thinking, and I frequently generate my best ideas at these. In fact, I maintain a long electronic document called “Research Ideas” in which I'm constantly workshopping new questions and potential projects. Upon revision many of these eventually turn into actual student projects or proposals. Others of them become blasé, apparent, or impractical, and get bumped to the bottom.

I enjoy staying on top of the newest literature and reading the latest-and-greatest published papers, to the extent that I can. I'd really like to spend more time here too, and there are occasions when I am able to allocate an evening for this purpose, especially if I am preparing a proposal. Fortunately,

conference presentations and interactions with colleagues during travel help fill in the gaps in my scanning of the literature.

I spend shockingly large amounts of time reading and writing emails. Most of these are short notes and so the time investment is really due to the sheer volume of email. Generally, I prefer not to have extended discussions over email with students. It takes much more time to type than talk, and direct two-way interaction will help me more quickly figure out what a student is talking about. Email is an inefficient mode of discussion, and is best reserved for logistics or resolving small issues. Thus, it's far better to email me to *schedule* a quick meeting.

You would be surprised by how many letters of recommendation I prepare and reference phone calls I provide. This is not a *major* activity per-se, but it does require a chunk of my time, especially in the Fall. I complete 10-20 undergraduate LORs per year, and many for graduate students for fellowships, jobs, etc. I generally love to do this for students (your success is my success!), but please, please, please... give me at least two weeks' notice and send me a recent CV every time. I want to make the best possible recommendation for you, and both time and organization enable that.

My expectations of you as a graduate student

Now that you know a little bit about my job, let me tell you how I see yours.

I expect you to be the lead on your project. Keep in mind that I am part of significantly more activities than those related to your project, including all of the projects of your groupmates and those of my collaborators. I function more as a kind of start-up CEO that establishes group priorities and directions (and raises money), while I see you as the project manager.

Therefore, I expect you to figure out many things on your own, especially the day-to-day details. You are smart, which is why you were admitted to this highly competitive program. I know that you can be resourceful and independent, even if you are new to research and don't quite feel that way yet. I am called an "advisor" and not a "boss" or "technical troubleshooter" for a good reason.

I work hard and I expect you to as well. I put in many hours on grant proposals to provide you with a paid education. I expect you to reciprocate with a high level of dedication and productivity on your project so that our collective efforts advance science and both of our careers. The "9 to 5" or "never on weekends" mentality does not work in graduate school. Grad students that start off coasting in undergraduate mode and not stepping up to put in their most energetic efforts will inevitably drag well into their sixth and seventh years. You (we) have to be able to keep up with the pace of science, which is blazingly fast. Working hard does not imply the lack of a social or family life by any means, and it is very important that you make time to nurture your personal life as it provides a healthy support system that will allow you to challenge yourself.

I also expect you to be efficient with your time and multitask between different responsibilities. I multitask as a matter of necessity. It is not usually efficient to drop everything for a week except for a single-minded focus on writing / reading / coding / etc. I tend to frown at that approach

because working on research problems in serial fashion will be harmful to your productivity (more on that later). Logistically, it also has a high probability of you getting “stuck,” bored, or slow and it wastes opportunities to do things in parallel. For example, when you are running of writing steam after 4-5 hours of working on a manuscript draft, you should switch to something else that you can approach with a fresh perspective, rather than painfully push through at a slower pace. I advise you to use a calendar to build a schedule for yourself and *actively* schedule your days. Take time each morning to assess your activities and then plan your day accordingly: one block for writing, one for data analysis, one for a new experiment, for example.

As a burgeoning young researcher, one of your most important activities is *reading*. Initially, this is key to bringing you up to speed with our field of research, through textbooks, classic papers, and review articles. Later on, this helps you stay up to date with ongoing research. Moreover, writing and publishing in peer-reviewed journals is the major mechanism by which new scientific information is documented, communicated, and evaluated, and you will only be able to develop good writing skills by becoming familiar with scientific idioms, writing styles, and audiences.

No one will ever “check on” your reading – this is your own responsibility and it takes self-discipline and curiosity. It is therefore critically important that you set aside consistent time each week for reading. I suggest that you subscribe to e-TOCs for relevant and major journals (I can make suggestions), and then scour them weekly for new papers. Beyond skimming abstracts, you should aim to completely read 100-200 papers per year (a mix of old and new). Of these, you will want to scrutinize about 30 of them, reworking derivations where appropriate and thinking deeply about results and implications. I expect you to quickly become an expert in the literature, moving beyond my own knowledge in your specific subproject. The most common candidacy exam blunder is a student that demonstrates a narrow knowledge of the literature, having focused only on group papers or a small handful that were directly suggested. Avoid this trap with proactive reading.

As a burgeoning researcher, you should have a notebook with you all of the time. If you can remember everything from the top of your head, you’re not thinking enough! A real research problem cannot simply be pondered, or it would have already been solved. Instead it’s hard, even for us smart folks. That’s what makes it good science. And that’s also what requires juggling tons of information, from meetings, conversations, papers, seminars, conferences, you name it. All of that is fodder for your notebook. Write down everything! I’m not going to keep track of things for you—remember, you’re the project manager—and in particular I will count on you to note and pursue the many ideas that we discuss in our meetings.

You should become an active member of the lab and greater research community. I expect you to be physically present in the lab for the vast majority of your research activities, and to establish a schedule for coming in during regular hours. It is tempting to work from home on theoretical or simulation projects, but your productivity will inevitably be compromised due to home-life distractions. Moreover, interactions with your lab mates are a critical part of the research process. You will learn as much from them as you will from me (perhaps more), and you will generate new ideas through detailed discussions and debates with them. That can be quite fun, actually.

You should attend seminars, including and especially those not in your immediate research area, as this is where creativity and cross-fertilization of ideas can happen. At the very least you should plan to attend every ChE seminar, but you can also add yourself to additional seminar mailing lists in Chemistry, Physics, Bioengineering, etc. "I'm too busy working on X to attend today's seminar" is a terrible habit that you must absolutely avoid. X can wait. Plan ahead around X by actively scheduling your calendar. If faculty can consistently attend seminars when they are in town, despite their frantic schedules, then so can you.

Seminars are actually a luxury. They give you a broad swath of ideas and areas in a single hour, and with great nuance as to the relative importance of ideas; that same information would require a significantly longer investment to glean from the literature. And yes, seminar speakers occasionally deliver boring talks. You have to make peace with that and realize those are par for the course, part of the odds of attending seminars. It only takes one phenomenal speaker in ten to transform your knowledge, inspire a new project, and more than compensate for the other nine lackluster ones.

I want you to immediately start working on your technical writing skills. The vast majority of first year graduate students are not experienced in writing technical publications and proposals, and so this is a key skill to start working on. There are plenty of resources to help you improve, online and otherwise, and I will gladly point you to them. I will also work to help you improve by putting many comments on your drafts; please take them to heart and learn from them. I am not simply being "picky," but truly want you to elevate your writing approach. I suggest that you also seek out comments from your peers on your writing drafts before sending them to me. Moreover, pay attention to the papers that are being published in journals in your area. Does your writing have the same level of depth, detail, and nuance? If not, let's figure out why and assess where you need to improve.

It generally should take a month to produce a manuscript draft once the results are clear and complete. A typical paper is about 7000 words. Let's say that you work on it for four hours a day for two weeks, which is what I would recommend. That's an average pace of 150 words per hour or *2.3 words per minute*. I know that you've probably written entire term papers the night before the due date as an undergrad, and so 2.3 words per minute is very doable. Therefore, if it is taking you many *weeks or multiple months* to produce a manuscript, let's reassess your writing approach to identify where you are wasting time. I can provide guidance.

How I interact with you

We will come up with many ideas in our meetings. I will also give you lots of suggestions, some great and some stupid, but by and large reasonable. The point of throwing these out is to effect the exploratory and creative approach necessary to getting to the bottom of the science. Therefore, you should take note of the ideas that we generate and then actually follow up on them – yes, all of them. It's hard to make progress on challenging research problems by keeping to a very narrow path.

At the same time, argue with me. That is the only way that you can become the expert. I may also disagree or argue with you (politely, of course). Such interaction is not meant to frustrate or

belittle; it is about subjecting our research to the highest possible scrutiny and standards. It is also about ensuring that my understanding of things is correct.

I usually don't remember every detail of your project off of the top of my head. Remember that you know much more about what you're working on than I do. It is helpful to orient me before presenting results, both in terms of your general recent goals/activities as well as the details (i.e., what are the axes and what are you plotting on that graph?). Please also bring visual results to our meetings (printed or electronic) since that will help me quickly understand.

I like to put things in my schedule so that I can organize my day's work around them. That's not to say that I discourage walk-in chats – many of which are fun, and I especially like surprise news of exciting new results. But for more routine meetings, it's generally helpful to shoot me an email to schedule something, even if just the day before. Keep in mind that many parts of my schedule are booked months in advance, though I do try to be flexible with my graduate students.

One point of warning: if you try to schedule a meeting with multiple faculty (e.g., a committee meeting), you should start at least two months in advance, if not more. It will be very difficult to synchronize our schedules with just a few weeks' notice, and it will be hard to motivate folks to accommodate last minute meetings if you could have started earlier but didn't.

It takes me some time to give feedback on written material you've given me: papers, abstracts, thesis chapters, etc. I cannot get through them in a few days unless I've promised you I will and I've specifically allocated the time. Most often it will take at least a week, and during busy times, many weeks. I am not trying to neglect you and I will make sure that your work eventually gets published. But I have other responsibilities and often some with pressing deadlines, and papers require significant time to process carefully, which is very important to me. You can help by...

- ...writing well. I care about this a lot. If I start correcting spelling and grammar, I will probably put down your manuscript and work on something more pressing. The same goes if your writing is confusing or requires basic restructuring. In that case, I might ask to meet and go over the structure in person.
- ...giving me complete drafts. I don't want to read many different versions of your paper as you are finalizing it. I am happy to give general advice about how to write something up and structure a story, but I prefer to read a complete draft when I finally sit down to edit a manuscript, including figures, results, and references. This is for a very important reason: the more times I read through a manuscript, the more desensitized I become to it and lose a critical eye.
- ...sending me a single file primed to edit. I generally prefer Word documents because they are easy to edit. I can also do Latex but will likely put it into a Word document to track changes. If you send me something else as a PDF, please make sure that it is double spaced because I may edit it by hand.

By the way, yes that manuscript actually does have to be *perfect* before we submit it. I know that the editing process may seem tedious in that I make many changes, sometimes even revising my

earlier revisions. That is because papers become stamped in the annals of science forever. They are the one place where we work out the details in all of their glory for the rest of the scientific community to digest. Not only are mistakes and poor writing embarrassing, they are a disservice to the community, a dilution of the new knowledge we've discovered, and ineffective science. That's not how we operate. So every word must be carefully chosen, every appropriate reference included, every equation checked, every graph cleanly constructed and visually clear. We must be exacting, specific, and at the top of our game when we publish.

Logistical matters

Generally, it's *very* important to me that we spend grant funding wisely, partially because I dedicate so much time securing it but also because I am responsible to the funding agency for spending on it. Here are some broad guidelines. Be reasonable in spending on equipment, lab materials, travel, etc. Grant money is not "free" money to spend – it always costs the grant and thus future spending opportunities. More importantly, it costs federal funds and taxpayer money meaning that it is extraordinarily important that we spend reasonably. At any moment, NSF could decide to audit our spending and if they find anything that looks frivolous or unnecessary, it will compromise our ability to get future funding.

I want you to attend and present results at major meetings. If you want to travel to a meeting and I didn't bring it up, please discuss it with me. My usual policy is that I will send you to major meetings to present once you start publishing results.

During travel, I expect you to share an inexpensive hotel with a roommate, to find lowest available airline fare, and to incur reasonable parking/transportation costs. Please do not simply try to max out the per diem, but rather use it to mitigate any increased personal expense while traveling. My own policy is that I recharge dinners that are "work related" where I meet with colleagues, but I do not recharge personal meals during travel. I know that you are on a much tighter budget and food while traveling can exceed the costs of what you normally spend at home, but do be reasonable about spending on food and don't rush to the fanciest restaurant that your per diem will support.

When you submit reimbursements, know that you cannot simply take a per diem maximum on meals. You must submit itemized receipts that include the actual menu items ordered on them (not just a credit card receipt with the total). And this is very important: alcohol can *never* be charged to a grant, and must be subtracted from your meal and tax totals during travel reimbursement. When you bring me completed reimbursement requests to sign, make sure to attach all of the itemized receipts so that I can fulfill my obligation to the granting agency in compliance.

Please take a good 2-4 weeks of personal vacation each year. We who work like dogs and regularly pummel our brains need to have breaks to rejuvenate and to keep creative. I'll never tabulate the number of days that you take off if you are making good progress. Just make sure to discuss with me when you want to take your vacation so that we can figure out how to address any important matters/deadlines around that time (grant reports, papers, my seminars and talks, etc.). If you can, it's good to target times in between quarters and in the summers.

The most important thing that you do

You are most certainly not a technician or lab rat. I expect that you will constantly generate and evaluate new ideas. Early in your graduate studies, of course, it is important not to wander too far away from your start up project(s) until you have gained sufficient expertise and know the full range of subtleties of your field. But as you mature as a researcher into your third+ years, my goal is that you will find exciting new directions and results that I did not originally anticipate.

Evaluation of your ideas, on a systematic and regular basis, is essential here. The first part of evaluation is assessment of relevance and feasibility. Don't just rely on me to tell you whether or not your idea is good; there are many points that you can and should address on your own time as you are building your own critical eye.

- Has anyone ever done or thought about this before? How does the idea fit within recent results in the field? Is the idea consistent with established work? → Do a thorough literature search. Identify the major players in the field and look at their collective publication records.
- Is this an interesting and worthwhile research question? Will resolving it be publishable and make an impact on your field? → Compare to major unresolved challenges identified in the literature, in topics at conferences, or in initiatives at funding agencies.
- Does the idea make physical sense, at various scales (e.g., molecular, mesoscale, or macroscopic)? What are the expected orders of magnitude and trends? → Think about the problem and use your scientific and engineering intuition. Do back-of-the-envelope estimates or calculations if necessary.
- Are there subtleties or potential problems that will arise? → Workshop your ideas with labmates, students in other groups working on related problems, faculty around campus, seminars. Think about possible scenarios, outcomes, and back-up plans.
- What priority level should the idea have? → Assess the potential for success and impact in relation to your other activities and in the context of the effort required to make it happen.

The second part of evaluation is the actual research effort to critically assess your hypothesis(es). Don't start until your idea has passed the tests above to see if it is worth your time. If it has, then plan your research approach in detail, and anticipate backup strategies and controls. It's dangerous to simply "jump in" without formulating a set of definitive, quantitative questions because you can then be mired in the details and lose sight of the point.

Probably around 80-90% of your more creative and novel ideas will fail. That's normal, and I'm subject to similar statistics. If great ideas came easier, some other smart person around the world would have already thought of them. You might discover a flaw in your thinking, a subtlety that makes the problem too complicated/messy or too trivial, a damning negative result, or a detail that renders the idea uninteresting, impractical, or not impactful. But even with failure there is an

element of progress because you carve out boundaries that help define the problem and delineate the solution space. You've learned something and gained expertise when you fail.

And then there are the truly successful ideas, the new levels of understanding that provide a completely novel way of thinking about a problem. They don't come in swift flashes, but rather blossom out of consistent hard work and creative thinking. At a top program like ours, that's the game we're always playing, and that's what a PhD signifies. A passing PhD thesis might have just one of these revolutionary ideas, an excellent thesis two, and a really outstanding thesis will produce three or more (think: future faculty candidate). You'll be ready to graduate when you've not only been able to craft such knowledge, but when you're capable of a compelling and nuanced proposal for the next project or investigation, e.g., a proposal for the next graduate student that joins the group.

How can I feel confident about how I am doing?

Understandably, the nature of exploring uncertain problems makes it very challenging – and frustrating – to calibrate your own progress throughout graduate school. You may invest substantial effort over a period of days or even months only to find a negative result or a seeming dead end. Even as faculty, it is sometimes difficult to assess your progress on short time scales because of the inherent unknowns of research. That being said, we know that a solid approach will eventually lead you to major publishable results in the long term. We have a good sense of what this is and thus generally know when you are on the right track, and can give advice and feedback. Here are some common questions that we would consider to assess progress, and that you should ask of yourself regularly.

Are you putting in a sound investment of time? It is not about quality *versus* quantity of time because our scientific field already operates and advances on the basis of a high quality. So, we must put in both quality *and* quantity to keep up. A good rule of thumb is that 50 hours of quality effort a week is a minimum that will allow you to “get by,” unless you happen to be astoundingly brilliant. 60 hours a week is more typical of a successful scientist. These numbers stem from the pace of science, not from advisor task-mastering. It's often quoted that it takes about 10,000 total hours to become an expert in anything. Lop off your first and last six months in graduate school, and the remaining four years already amount to 50 hours per week of solid study. Then, add in low-quality time in email, logistics, random off-topic chats, etc. Granted, some weeks will be light and easy – everything works out, you've been very productive, able to tick off many activities and accomplishments from your to-do list. Go enjoy the weekend! Others, though, may require a major dedication in order to resolve a critical issue in a timely manner and avoid (gasp) scooping. Go enjoy a nice coffee with your labmates to refresh your energy!

Are you doing enough early in your PhD? Towards the first or middle part of your second year, we want you to be “up to speed” with the major methods, approaches, results, ideas, and questions in your immediate project area. This is for a very good reason: we want you to spend the majority of your subsequent time working on relevant, new, interesting problems and making good research progress, rather than rediscovering what's been done or repeating earlier capabilities. In scientific

research, you will never be given credit for the latter. This is why it is so important that you have exacting standards for your knowledge of the literature, and read and think prodigiously, in that first 1.0-1.5 years of research. If you are at the point where you can quote papers, rattle off ~10 of the major other groups in your area (and their recent results), and confidently explain known findings in regular discussions with us, then you are at the right level.

Are you spending too much time on “low-bar” activities? It is very easy to work on tasks that resolve details important to your project but that have an overall small or incremental impact in terms of research insight. We all like the feeling that comes with crossing off such activities from our to-do lists. Examples include: understanding/implementing an existing model, debugging code, learning an established experimental technique, organizing your data, creating nice reports of your results, etc. All of these need to and must happen, but you should be cognizant that they are not taking up more than 50% of your time. If they are, you need to step up your game and likely need to reprioritize or invest more time per week until you iron out the inefficiencies of the easy stuff. If you look back and you have spent most of the last 2-3 months in such activities and not making strides in research results, then this indicates the need to change strategies.

Are you approaching your project from multiple angles in parallel, and constantly (re)prioritizing your efforts? Because many novel ideas will not work out, a serial work style in which you evaluate one idea after the other will impair your scientific productivity. Instead, you should have a handful of different approaches and “leads” that you are multitasking simultaneously. Keep a list in your notebook of possible directions to try. When you get stuck on one front, move to another. As new results come in, rebalance your time distribution between them, on a daily basis.

Are you working at a deep level of thought? This is probably the most intimidating aspect of research because it will push you to your limit, and it is at the core of what differentiates a top from an average scientist. You must do a serious, thorough evaluation of ideas and results that come your way, not once a month or quarter, but regularly each day. A good rule to live by is: *never accept anything at face value!* You must constantly scrutinize everything. Some examples:

- When generating a new idea, how does it fit into the entire relevant literature? Have you explored all possible quick or back-of-the-envelope estimates and calculations that can be done to evaluate it? If chatting with a colleague, can you defend it off of the top of your head, pulling or evaluating details as needed? What are all the ways the idea could be incorrect, or rendered problematic due to complications? How many multiple, independent ways can it be tested?
- When results come in, do they have the expected direction? Do they violate any underlying principles (e.g., from thermodynamics or transport)? Would they be surprising to your research community? Do they conform to any known model or scaling law? Do they change in the expected way with conditions like temperature or concentration? What is actually being measured by the instrument or in the simulation? What complications are you possibly ignoring (e.g., impurities, small system size, etc.)? Is the level of error expected?

- When you read a new published paper, what is new and novel? Is the analysis sound and free of questionable assumptions? What earlier papers does it support/disagree with? How does the paper fit into the historical work of the lead author? Are there variations or ideas that should have been explored by the authors but weren't? Are the conclusions supported by the data, or could alternate hypotheses be relevant? Could you explain and teach the paper to someone else? Could you recapitulate the main methodological details yourself? Could you rederive the main results?

This all sounds a little intimidating

It should! But part of the reason is that you've never been pushed to your absolute maximum capability before. You were never in the lower half of your undergraduate class and struggling to keep up. Now you are in a super-elite group that is actually capable of creating new knowledge, and the requirements (those of the proverbial bar) are much higher. It is as if you graduated from your school track team to become an elite Olympic sprinter. You're primed for the experience, but it will require a level of intense training and dedication beyond what you've already experienced. Welcome to our Olympic team.

You should know that it is equally challenging for us. We are trying to run world-class research groups that push the boundaries of knowledge. That's no easy task. Admittedly, we faculty can be irritable at times. It could be because we've had four hours of sleep for days in a row, e.g., on proposals, administrative activities, or the pre-tenure push. Or it could be because we've put much effort into a recent proposal that was rejected and thus suddenly have to work hard and quickly to ensure continued group funding. Alternatively, we could simply have an exhaustingly busy day of back-to-back meetings (visitors, admin, students, class, etc.), and we just can't focus well. We try to diffuse resultant "bitchiness," but sometimes it comes through. You should not take it personally unless we have legitimate feelings that you could be more productive, i.e., if we are working much harder than you.

In contrast, your graduate research is lean on administrative activities, and in turn your experience while challenging should be a largely creative endeavor. It has the potential to be consistently rewarding and exciting if you are playing your cards right. Unlike me, you can spend the vast majority of your time thinking, exploring, learning, creating. That is the absolute best part of research, the *crème de la crème*. You may work yourself hard, but you will not spend the majority of your hours in logistics. Instead, you have the opportunity to immerse yourself in discovery, dynamic debates, a hunt for clues... all of the things that make learning exhilarating. You'll never again have the same situation after graduate school (and postdoc, if you take one up).

Part of playing your cards right and avoiding cynicism is simply realizing that research is both intrinsically hard and independent. You may sometimes feel isolated as a graduate student, both in terms of feedback and your research efforts. Those kinds of feelings indicate your new presence in an incredibly elite group. Remind yourself that you do have the inner abilities to step up to the plate and succeed fabulously, if you work to nurture them. We believe so!

What keeps us in this game is the opportunity to learn constantly and to push the boundaries of knowledge. So few people have jobs as creative, independent, and potentially impactful as ours. As an undergraduate, mastering challenging exams or projects gave you a sense of great pride and accomplishment. That wonderful, addictive feeling of learning probably led you to graduate school. Now imagine that you've not only learned something but you've carved it into being as entirely new science. It did not exist before your efforts. It is truth that will bear your mark forever. That sense of triumph, success, and permanence is a reward many times stronger than any A+ exam grade, and it keeps us all working hard. That is why we are in this business.

Final notes

Feeling a bit pressured and thus "stressed" in research in graduate school is a normal and potentially good thing. You should think of it as a mechanism pushing you towards greatness. Suddenly one day you will realize that you are an elite athlete with incredible skills. We hope that this will make you want to take on even greater challenges in the next stage of your career. In some ways, that is the most important measure of our success as educators.

We are very concerned with your success. Lots of money and time is being invested in it. We do not want you to fail! Keep in mind that you are part of an extremely elite group: a graduate student at a top department in a top university. Take maximum advantage of this opportunity while you are here. These five years will be among the most exciting and creative times in your life, and will launch and help define your entire career. But they will go by fast. Don't let the experience happen to you passively or languidly.