10

One of Many: Building a Body of Work

One of the many ways that science is like pop music is the prevalence of one-hit wonders, youngsters who come out of nowhere, publish something striking, and then haunt fringe venues, like regional conferences and county fairs, for the next few decades. How can we avoid this? One of the many hard truths about academic publishing is that one paper, even if excellent, rarely makes a splash. It usually takes a line of work, a program of connected papers, a "network of enterprise" (Gruber, 1989) to get attention and change minds. Not every paper we write will be a hit, of course, but no one goes into science to publish only one notable thing. Like pop stars, researchers should seek to develop a substantial *oeuvre*, although I'm sure being huge in Finland has its comforts.

This chapter thus takes the long view of academic writing. How can we develop a program of research and bring attention to it? What isn't worth writing? And how will we get all this writing done?

One Is the Loneliest Number

Impact comes from a line of work, a series of linked papers on the same topic. Successful research programs often have a striking start, a first paper that is the best-known piece, but one rarely sees a rich orphan, an influential article that the author never followed up. My idealistic side wishes that quality would shine through, but the modern social sciences are too loud and crowded. Publishing a series of papers brings attention to your work through sheer mass: People are more likely to notice your ideas if they appear in several journals across several years. Beyond mere mass, a series of papers tells your readers that you're committed to your ideas, that you think they're important and worth years of your time. Some researchers hop capriciously from topic to topic, resulting in a vita with a quirky gallimaufry of publications. As a reader, I wonder why I should study an idea when its parent doesn't find it important enough to study more than once. And finally, a series of papers can reveal an idea's fecundity. By showing implications and extensions, moderators and boundaries, you show your readers what can be done with your idea, thus sparking their inspiration.

When thinking of research ideas, then, we should think expansively and plan ahead. The narrowest mind-set, and a common one among beginners, is simply trying to think of a study that could get published somewhere. As we argued in Chapter 1, strategic writers plan their research and think of it in terms of papers targeted to specific audiences and journals. Here we suggest thinking even more generally—instead of planning a lone paper, plan a series of studies, a network of linked papers that extend and elaborate your ideas. We're not trying to plan out our next 17 manuscripts—any number greater than one will suffice.

Trying to stretch your initial idea into a long program might take you farther than you thought, inspiring several new ideas for related papers—shockingly, creativity research has shown that deliberately trying to come up with creative ideas works (e.g., Christensen, Guilford, & Wilson, 1957). Or you might get the awkward experience of starting to talk but then finding yourself at a loss for words—some ideas are one-paper ideas, not cornerstones for a program. Either way, you've learned something essential: Your idea was more fertile or more barren than you thought. If you can think of a solid series, that's usually a sign that the idea is big enough to attract attention; if you can't, the idea might still be worth your time, but you can make a more realistic appraisal of whether it is worth the effort.

Ways to Build Impact

Be Discerning

Our lives are too hectic and brief to study every hypothesis that blithely wanders into our minds and trips over the extension cord duct taped to the carpet. Our bodies of work will have a bigger impact if we are discerning about our ideas. By committing only to our better ideas, we develop a better body of work and use our research time wisely. I write down most every idea I have, no matter how fringe, and they get placed into different mental piles.

- The first is the “must do” pile, the ideas I think are my best and connect directly to the problems I care most about. These are research ideas that are closest to my heart and that I think will be placed in the strongest journals. This pile gets written down in a document shared among the lab, so everyone can see and revise my ideas.
- The second is the “probably do” pile, the ideas that I’d like to do but lack the urgency or importance of the first pile. Often these are extensions or elaborations of “must do” ideas.
- The third is the “if time permits” pile, the ideas that seem neat but can’t compete for time. These usually only come to fruition if collaborators or students get excited about them and want to take the lead.
- The fourth pile is the “compost pile,” ideas that appeal to some fiendish or irrational streak of mine but are surely not worth studying. The main value of this pile is highlighting that not every idea is worth doing.

Beyond being discerning about our ideas for research, we should be discerning about what we seek to publish. Not all of our “must do” ideas pan out well:

File cabinets and flash drives are full of studies that ended badly, awkwardly, or mysteriously. People who seek mere publication will try to publish anything anywhere, and this is foolhardy. For one, the most insidious cost in life is opportunity cost. We can only do so much, so committing time to one project requires us to forsake others. And studies that don't end cleanly are usually hard to get published. As a result, a good paper takes less time and effort than a weak one, which will get kicked from journal to journal, undergoing extensive changes along the way, all for little gain in knowledge and impact. Efficiency alone suggests that we should euthanize our sickly projects rather than send them into the wild, where they'll be mauled by predators and ignored by potential mates.

What are signs that something shouldn't be published at all? Here are a few: the methods have severe flaws; the results lack a take-home message; the findings seem unlikely to replicate; and the message feels tortured, often because the data were collected for another purpose and don't dovetail with the new argument. The alternative is to try to "put a good face" on the data, to try to conceal the project's weaknesses and hope reviewers won't notice. I can assure you that they always notice. But even if they don't, a crueler fate awaits when readers notice, and the writer will diminish in their esteem. Remember that the only face on your work is your own.

Good writers should take pride in what they withhold—someone who publishes everything has no standards. As the culinary dictum goes, "Don't let your mistakes leave the kitchen."

Write Review Articles

Once you have a small program of research in print, you should think about synthesizing it into a review article. Review papers get freakish attention. Anyone getting started in a field, planning readings for a grad seminar, or developing lectures about an unfamiliar topic will turn first to review papers. This has always been true, but it is probably becoming more so: There's too much primary literature to keep up with, so readers need the distilled version.

The phrase *review article* is inapt because the goal isn't to merely review what others have done. Like empirical articles, review articles must make a point. The best reviews advocate for a position, review the literature on both sides, address nuances that aren't easily worked into an empirical paper, and point to where the literature should go. Baumeister and Leary (1997) provided some excellent advice on writing review articles. Some of their reviews have been cited thousands of times, so they know of what they speak.

If they have written any, most researchers will find that their review articles are their most cited works. I'm thus always surprised at how few people consider writing review articles. One reason is surely the scope of the task: A long review article can be herculean. Another, I suspect, is that people have already said most of what they would want to say in edited book chapters, which eat up as much time as review articles but have less impact.

When planning a new area of research, you can use a hypothetical review article as a heuristic for coming up with ideas for the research program. If you were to write a review article, what kinds of studies would need to be done? What problems would need to be tackled, and with what methods? What new ideas should be infused into the literature? Think about it, do those studies, and then write the review article.

Collaborate

Collaborating expands the range of things you can do and plugs you into a network of productive peers. Chapter 3 had much to say about collaboration—including a caution about choosing collaborators wisely that bears repeating—so here we'll simply reinforce the many virtues of working together with your peers. Strong teams can pool time, resources, and expertise, so they can execute projects that are hard for one person. Over the years, you'll find that working with other experts provides a constructive sense of humility, a discerning appraisal of what you do well and what others do better. And finally, one of the collaborators, if the team is big enough, will probably bring bagels.

Organize a Community

Most areas of research are small areas—no problem is too technical or obscure to evade the keen eye of science. Even when the audience is big, the papers are usually generated by a ragtag band of merry researchers. By organizing this band into activities and institutions that attract attention and promote the group's work, you can promote your own work. Some obvious examples include proposing a special issue for a journal or sessions at conferences. In some cases you can make it easier for people to share their work, such as by creating an email list or social-media page that serves the dual goals of connecting interested researchers and allowing them to waste time on the Internet in the guise of science. And if you have time on your hands and a knack for administrative tasks—two facts you must never let your dean know—you can organize a preconference, a free-standing conference, or even a new scholarly society devoted to your area of research.

Encourage Disagreement

Once your papers get out there, you'll start to see them cited and discussed in other papers. The first times you see your work cited are eerie—it finally sinks in that people do read these papers and that we can't change what we committed to print. And eventually you'll see unflattering portrayals of your work. It might be gentle, such as pointing out a minor limitation or oversight, but it might be ghastly, such as mocking everything but your running head.

Disagreement is good. Your critics, be they gentle or ghastly, often have the most interest in your research:

They're the ones who are following your program, reading your papers, and conducting research inspired by it. Our pettier natures are tempted to thwart our critics, but you should encourage them. They don't bite, although they might nibble. These people will be good colleagues and collaborators.

Seek External Funding for Your Research

When you have research grants, it's hard to answer a self-evident question like "Why should I seek external funding?" It's like asking cat owners why they like cats—you have to have one to really get it. Viewed pragmatically, grant writing gets done for three reasons. First, people write grants to avoid penury and homelessness. If you have a soft-money job or work in a department that requires external funding for tenure, you know what I mean. Getting grants is less about scientific pride and more about preventing the kids from wearing burlap clothes. If you plan to go into such a job after grad school, don't give away your ramen noodle stash to your office mates when you graduate. Second, people write grants to fund a good idea. For obvious reasons—employing staff, paying participants, and buying equipment—some research won't happen without external funding. Getting a grant thus lets you do research that is expansive and exciting. And finally, people write grants to develop their ideas for a research program. This shouldn't be the main reason for writing a grant proposal, but it's a nice consolation prize if the proposal isn't funded. Consistent with the writing-to-learn approach (Zinsser, 1988), writing about 3 to 5 years' worth of research is a great way to find out how good your ideas really are.

You should seek external funding, even if your department doesn't require it. Grant writing connects you into a broader world, allows you to execute bigger kinds of research, and can be weirdly gratifying even when the proposals go unfunded. I'm not snobbish about amounts and mechanisms and sponsors. There are only two kinds of grants: Those you get and those you don't. A $450 grant from a small foundation that lets you buy some necessary equipment gets you farther than a $450,000 proposal that a federal agency declined.

Writing to Avoid

We never have enough time to write, so we need to be selective. Here are some kinds of writing worth avoiding.

Chapters for Edited Books

In a quirky and compelling case study, Dorothy Bishop (2012) analyzed the citations of her publications. She sorted her papers into three categories: edited book chapters, empirical papers in journals, and conceptual and review articles in journals. Her conclusion? "Quite simply, if you write a chapter for an edited book, you might as well write the paper and then bury it in a hole in the ground." The journal articles—empirical or conceptual—dwarfed the book chapters, which had attracted few citations. This pattern is true for my work—some of my book chapters are buried deep enough to require spelunking equipment—and, I suspect, true for nearly everyone who works in a field that privileges traditional peer-reviewed journals over books, conference proceedings, and open-access outlets.

Why do book chapters get buried? One can consider many reasons, but I think Bishop (2012) nailed it:

> Accessibility is the problem. However good your chapter is, if readers don't have access to the book, they won't find it. In the past, there was at least a faint hope that they may happen upon the book in a library, but these days, most of us don't bother with any articles that we can't download from the Internet.

Journal articles are easily accessed online from anywhere in the world. Some publishers have placed some edited books into online databases, but reading the typical book chapter still requires a trip across campus to the library, where physical books sit crammed together like dogs in an animal shelter, each hoping for a short stint outside.

On the basis of impact alone, then, we should choose our invitations to contribute chapters wisely. But chapters have some other strikes against them. In most of the social, educational, and health sciences, book chapters count for less than articles when one's beans are counted. Book chapters also have deadlines, those banes of busy professors everywhere, and the deadlines are always when time is tight and enthusiasm is low. Many readers view chapters as dumping grounds for orphaned data and rewarmed ideas—not unfairly, I think—and thus expect to find little of value. Finally, chapters tend to be far longer than a typical manuscript, so their impact payoff is poor relative to the time spent writing them.

Given this bleak view of book chapters, why write them at all? Exhibit 10.1 shows some guidelines for when a chapter might be worth writing. If a chapter is about your own work, it will be easier, faster, and more interesting to write, and you can speculate and integrate in ways that are intellectually gratifying for you and your seven readers. If you're pretenure and the project looks interesting, go ahead and write the chapter. It connects you to other researchers and shows that your work is attracting attention. But if time is tight and you have better things to write, think thrice. I decline most of the invitations I receive. A simple

EXHIBIT 10.1. Some Reasons for Writing a Book Chapter

- The book seems prestigious, and it would be nice to be a part of the project.
- The editor is a friend or someone to whom you owe a favor.
- The chapter is a writing opportunity for your graduate students, who can take the lead under your mentorship (also known as "You guys write that and get back to me").
- You have a lot of time for writing and a thin backlog, so the chapter isn't crowding out more important projects.
- The chapter is about your research and ideas, which you know well.
- It pays a surprising amount of money.

e-mail that thanks the editor for thinking of you but notes that your writing backlog is too deep to take on another project will suffice. Either way, always accept or decline quickly so the editors can keep moving on their own project.

Encyclopedias, Book Reviews, and Ephemera

The more seasoned among us remember a time when interested people would adopt a curious visage, walk to a shelf, and look something up in an encyclopedia or dictionary. Those books were a great way to avoid feeling left out when people were abuzz about the hottest trends in science and technology, like watches powered by small batteries instead of mainsprings, and they had a good run. Publishers still publish encyclopedias and dictionaries focused on professional topics, and they need experts to write entries. Libraries still buy these volumes, ensuring wide availability if not a wide readership. I suspect that all but a few of these books fall off the face of the earth. It's a shame—these books are great, and we would like our students to look things up using legitimate texts with entries composed by experts, not from the Internet, the home of ignorance in all caps—but it is what it is. There's no great harm in writing a short entry for a dictionary or encyclopedia, apart from more anxiety from your ever-looming backlog of more important projects, so use your time well.

Book reviews are another quirky kind of publication. I enjoy writing them—they force me to read and reflect about a book I'd like to read—but they reach a limited audience. The author is always in that audience, even if your review appears in a regional newsletter printed with a purple ditto machine, so think twice before ridiculing the book if you hate it. Book reviews also take surprisingly long to do, so don't do one unless you'd really like to read the book.

Our final category is ephemeral writing, the catch-all for publications that aren't archived or cataloged: your blog, if you have one; guest posts for someone else's blog; and essays for newsletters, ranging from humble department house organs to periodicals that reach thousands of peers in a professional society. Such publications aren't peer reviewed, unless you count the flamboyant histrionics in a blog's comments section as a kind of review and your peers are rogues and desperados, but they can reach surprisingly large audiences. Some of my best writing is in this category, and I suspect that more people have read my guest blog posts and newsletter essays than most of my journal articles. Nevertheless, ephemeral writing, like anything usually requiring the Internet, can be a tremendous time sink. Blogging in particular can be an insidious form of procrastination—writing something fun and ostensibly productive to avoid writing something important—so don't neglect your journal articles.

How to Write It All

If you're freaking out, wondering how you'll ever write all that stuff when the paper you just finished was so hard, freak not. You can do it. Most people have bad habits and mind-sets that hold them back, like believing that they need to wait for big blocks of time, inspiration, a feeling of readiness, a whole day at home, an uncluttered desk, and other fictions that let us procrastinate with a clean conscience. A bit of behavior change and a solid routine are usually enough to get a lot of papers written.

Motivational aspects of productive writing are a book unto themselves. One of those books is *How to Write a Lot* (Silvia, 2007), but there are many more that address time management, procrastination, and good habits (e.g., Boice, 1990; Goodson, 2013; Lambert, 2013). My perspective is that we have much less control over our time than we think. Most of us believe, against all reason and experience, that there's time in the week to be found for writing. But, of course, our time is quickly set upon by the usual brigands: teaching, service, fires to put out and start, and bushels upon bushels of e-mail.

If the workweek is largely a maelstrom of chaos, then I think it is fruitless to adopt intricate time-management systems or to set temporal goals farther out than 1 week. I have seen people set 6 or 8 weeks' worth of goals for their paper—"I'll spend two weeks on the Intro, one on the Method, two on the Results . . . "—but such plans usually come to grief, dashed on the shoals of grading, web browsing, and providing vital service on the Associate Vice Provost for Parking's Utilization Committee.

We should accept that much is out of our hands and then control what we can—our own behavior. We can choose a time for writing, sit down and write during that time, and then stop when that time is over. Scheduling writing is how people who publish a lot write a lot. Making a writing schedule guarantees time to write and shelters your writing from the inanity and chaos of the workweek. And after a couple weeks, writing at that place at that time becomes a sturdy habit, and writing is no longer something you choose, hope, or want to do—it's just another habitual reflex, like brushing your teeth, winding your watch, and grousing about the kids these days.

Try it—start with 4 to 6 hours a week for writing. Four hours is enough to write most of what you'd like to write and more than most people spend writing. You'll be surprised how much you get written.

Wrapping Everything Up

As professors, teachers, and mentors, we know all too well that giving people fish feeds them only for a day. But teaching them to fish—ideally with PowerPoint slides, short essay tests, and discursive lectures on history and theory—feeds them for a semester, after which they sell their ichthyology textbook back. This book, with its combination of trust-me-on-this and here's-why, has tried to give and to teach. Life is too short, and the publication process too long, to learn the hard way—it's easier to learn from other people's mistakes, especially when they are embarrassing and hilarious. But most of writing's many decisions require people to develop their own informed perspective, so some things you'll learn only from practice and rejection.

This book developed a few themes: We should write for impact and influence, not merely for publication; we should respect the opportunity cost of writing and be selective in what we pursue; we should view writing as a craft, a skill to honor, not as a mere step in the research process; we should be reflective and plan for writing's many decisions, ranging from picking journals to selecting references; and we should sweat the small stuff until it breaks down and confesses.

Psychologists are a hardened lot—it comes from dealing with students who say, "I majored in psych because all my friends say I'm a good listener" yet don't listen when told the assigned readings—so this journey through the netherworld of academic writing and publishing will end with a taciturn goodbye, a flinty nod instead of sassy kisses on both cheeks. These chapters distilled most of what I've seen and heard from my years in the peer-reviewed trenches. Now you should get to writing and dig some trenches of your own—pack a big thermos and give me a flinty nod if our trenches should cross.